



NETGEAR 8800

Chassis Switch CLI Manual

Software Version 12.4

350 East Plumeria Drive
San Jose, CA 95134
USA

March 2011
202-10802-01
v1.0

© 2011 NETGEAR, Inc. All rights reserved.

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form or by any means without the written permission of NETGEAR, Inc.

Technical Support

Thank you for choosing NETGEAR. To register your product, get the latest product updates, or get support online, visit us at <http://support.netgear.com>.

Phone (US and Canada only): 1-888-NETGEAR

Phone (Other Countries): See Support information card.

Trademarks

NETGEAR, the NETGEAR logo, ReadyNAS, ProSafe, Smart Wizard, Auto Uplink, X-RAID2, and NeoTV are trademarks or registered trademarks of NETGEAR, Inc. Microsoft, Windows, Windows NT, and Vista are registered trademarks of Microsoft Corporation. Other brand and product names are registered trademarks or trademarks of their respective holders.

Statement of Conditions

To improve internal design, operational function, and/or reliability, NETGEAR reserves the right to make changes to the products described in this document without notice. NETGEAR does not assume any liability that may occur due to the use, or application of, the product(s) or circuit layout(s) described herein.

Revision History

Publication Part Number	Version	Publish Date	Comments
202-10802-01	v1.0	March 2011	First publication

Contents

Chapter 1 Command Reference Overview

Introduction	6
Audience	6
Structure of this Guide	7
Understanding the Command Syntax	7
Access Levels	7
Syntax Symbols	8
Syntax Helper	8
Object Names	9
Command Shortcuts	10
Port Numbering	10
Numerical Ranges	10
Line-Editing Keys	11
Command History	12

Chapter 2 Commands for Accessing the Switch

Chapter 3 Commands for Managing the Switch

SNMP	57
Telnet	58
TFTP	58
System Redundancy with Dual Management Modules Installed	58
Power Supply Management	58
Simple Network Time Protocol	59

Chapter 4 Commands for Managing the NETGEAR 8800 Software

Chapter 5 Commands for Configuring Slots and Ports on a Switch

Chapter 6 Commands for Configuring LLDP

Chapter 7 PoE Commands

Summary of PoE Software Features	300
--	-----

Chapter 8 Commands for Status Monitoring and Statistics

Event Management System	330
sFlow Statistics	331
RMON	331

Chapter 9 VLAN Commands**Chapter 10 FDB Commands****Chapter 11 Commands for Virtual Routers****Chapter 12 Policy Manager Commands****Chapter 13 ACL Commands****Chapter 14 QoS Commands****Chapter 15 Security Commands**

SSH	554
SSL	554
User Authentication	554
Denial of Service	555

Chapter 16 Network Login Commands**Chapter 17 STP Commands**

STP	737
RSTP	737
MSTP	738
Spanning Tree Domains	738
Member VLANs	738
Carrier VLAN	738
Protected VLAN	739
STPD Modes	739
Encapsulation Modes	740
STP Rules and Restrictions	741

Chapter 18 VRRP Commands**Chapter 19 IP Unicast Commands****Chapter 20 IPv6 Unicast Commands****Chapter 21 RIP Commands**

Chapter 22 RIPng Commands

Chapter 23 OSPF Commands

[OSPF Edge Mode](#)995

Chapter 24 OSPFv3 Commands

[OSPF Edge Mode](#)1037

Chapter 25 BGP Commands

Chapter 26 IP Multicast Commands

Chapter 27 IPv6 Multicast Commands

Chapter 28 MSDP Commands

Chapter 29 vMAN (PBN) Commands

Appendix A Configuration and Image Commands

Appendix B Troubleshooting Commands

[Event Management System](#)1345

Command List

Command Reference Overview

1

Introduction

This guide provides details of the command syntax for all NETGEAR 8800 Chassis Switch commands as of Software Version 12.4.

The guide does not provide feature descriptions, explanations of the technologies, or configuration examples. For information about the various features and technologies supported by NETGEAR switches, see the *NETGEAR 8800 User Manual*.

This chapter includes the following sections:

- [Audience](#) on page 6
- [Structure of this Guide](#) on page 7
- [Understanding the Command Syntax](#) on page 7
- [Port Numbering](#) on page 10
- [Line-Editing Keys](#) on page 11
- [Command History](#) on page 12

Audience

This guide is intended for use by network administrators who are responsible for installing and setting up network equipment. It assumes a basic working knowledge of the following:

- Local area networks (LANs)
- Ethernet concepts
- Ethernet switching and bridging concepts
- Routing concepts
- Internet Protocol (IP) concepts
- Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Intermediate System-Intermediate System (IS-IS) concepts
- Border Gateway Protocol (BGP-4) concepts
- IP Multicast concepts

- Protocol Independent Multicast (PIM) concepts
- Simple Network Management Protocol (SNMP)

Structure of this Guide

This guide documents each NETGEAR 8800 OS command. Related commands are grouped together and organized into chapters based on their most common usage. The chapters reflect the organization of the *NETGEAR 8800 User Manual*. If a specific command is relevant to a wide variety of functions and could be included in a number of different chapters, we have attempted to place the command in the most logical chapter. Within each chapter, commands appear in alphabetical order. You can use the Index of Commands to locate specific commands if they do not appear where you expect to find them.

For each command, the following information is provided:

- **Command Syntax**—The actual syntax of the command. The syntax conventions (the use of braces, for example) are defined in the section [Understanding the Command Syntax](#) on page 7.
- **Description**—A brief one sentence summary of what the command does.
- **Syntax Description**—The definition of any keywords and options used in the command.
- **Default**—The defaults, if any, for this command. The default can be the default action of the command if optional arguments are not provided, or it can be the default state of the switch (such as for an enable/disable command).
- **Usage Guidelines**—Information to help you use the command. This may include prerequisites, prohibitions, and related commands, as well as other information.
- **Example**—Examples of the command usage, including output, if relevant.

Understanding the Command Syntax

This section covers the following topics:

- [Access Levels](#) on page 7
- [Syntax Symbols](#) on page 8
- [Syntax Helper](#) on page 8
- [Object Names](#) on page 9
- [Command Shortcuts](#) on page 10

Access Levels

When entering a command at the prompt, ensure that you have the appropriate privilege level. Most configuration commands require you to have the administrator privilege level.

Syntax Symbols

You may see a variety of symbols shown as part of the command syntax. These symbols explain how to enter the command, but you do not type them as part of the command itself.

Table 1 summarizes the command syntax symbols.

Note: NETGEAR 8800 software does not support the ampersand (&), left angle bracket (<), or right angle bracket (>), because they are reserved characters with special meaning in XML.

Table 1. Command Syntax Symbols

Symbol	Description
angle brackets < >	Enclose a variable or value. You must specify the variable or value. For example, in the syntax <pre>configure vlan <vlan_name> ipaddress <ip_address></pre> you must supply a VLAN name for <vlan_name> and an address for <ip_address> when entering the command. Do not type the angle brackets and do not include spaces within angle brackets.
square brackets []	Enclose a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax <pre>use image [primary secondary]</pre> you must specify either the primary or secondary image when entering the command. Do not type the square brackets.
vertical bar	Separates mutually exclusive items in a list, one of which must be entered. For example, in the syntax <pre>configure snmp community [readonly readwrite] <alphanumeric_string></pre> you must specify either the read or write community string in the command. Do not type the vertical bar.
braces { }	Enclose an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax <pre>reboot {time <month> <day> <year> <hour> <min> <sec>} {cancel} {msm <slot_id>} {slot <slot-number> node-address <node-address> stack-topology {as-standby} }</pre> you can specify either a particular date and time combination, or the keyword <code>cancel</code> to cancel a previously scheduled reboot. (In this command, if you do not specify an argument, the command will prompt asking if you want to reboot the switch now.) Do not type the braces.

Syntax Helper

The CLI has a built-in syntax helper. If you are unsure of the complete syntax for a particular command, enter as much of the command as possible and press TAB. The syntax helper

provides a list of options for the remainder of the command, and places the cursor at the end of the command you have entered so far, ready for the next option.

If the command is one where the next option is a named component, such as a VLAN, access profile, or route map, the syntax helper also lists any currently configured names that might be used as the next option. In situations where this list might be very long, the syntax helper lists only one line of names, followed by an ellipses (...) to indicate that there are more names than can be displayed.

Some values (such as the <node-address>) are lengthy, but limited in number. The NETGEAR 8800 places these values into a “namespace.” This allows command completion on these values.

The syntax helper also provides assistance if you have entered an incorrect command.

Abbreviated Syntax

Abbreviated syntax is the shortest unambiguous allowable abbreviation of a command or parameter. Typically, this is the first three letters of the command. If you do not enter enough letters to allow the switch to determine which command you mean, the syntax helper provides a list of the options based on the portion of the command you have entered.

Note: When using abbreviated syntax, you must enter enough characters to make the command unambiguous and distinguishable to the switch.

Object Names

All named components within a category of the switch configuration, such as VLAN, must be given a unique object name. Object names must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but they cannot contain spaces. The maximum allowed length for a name is 32 characters.

Object names can be reused across categories (for example, STPD and VLAN names). If the software encounters any ambiguity in the components within your command, it generates a message requesting that you clarify the object you specified.

Note: If you use the same name across categories, NETGEAR recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Reserved Keywords

Keywords such as `vlan`, `stp`, and other 2nd level keywords, are determined to be reserved keywords and cannot be used as object names. This restriction applies to the specific word (`vlan`) only, while expanded versions (`vlan2`) can be used.

A complete list of the reserved keywords for NETGEAR 8800 12.4 and later software is displayed in Table 8 of the *NETGEAR 8800 User Manual*. Any keyword that is not on this list can be used as an object name.

Command Shortcuts

Components are typically named using the `create` command. When you enter a command to configure a named component, you do not need to use the keyword of the component. For example, to create a VLAN, enter a VLAN name:

```
create vlan engineering
```

Once you have created the VLAN with a unique name, you can then eliminate the keyword `vlan` from all other commands that require the name to be entered (unless you used the same name for another category, such as STPD). For example, instead of entering the command:

```
configure vlan engineering delete port 1:3,4:6
```

you could enter the following shortcut:

```
configure engineering delete port 1:3,4:6
```

Port Numbering

Commands that require you to enter one or more port numbers use the parameter `<port_list>` in the syntax.

Note: The keyword `all` acts on all possible ports; it continues on all ports even if one port in the sequence fails.

Numerical Ranges

On the NETGEAR 8800, the port number is a combination of the slot number and the port number. The nomenclature for the port number is as follows:

```
slot:port
```

For example, if an I/O module that has a total of four ports is installed in slot 2 of the chassis, the following ports are valid:

- 2:1

- 2:2
- 2:3
- 2:4

You can also use wildcard combinations (*) to specify multiple modular slot and port combinations. The following wildcard combinations are allowed:

- slot:*—Specifies all ports on a particular I/O module.
- slot:x-slot:y—Specifies a contiguous series of ports on a particular I/O module.
- slot:x-y—Specifies a contiguous series of ports on a particular I/O module.
- slota:x-slotb:y—Specifies a contiguous series of ports that begin on one I/O module or node and end on another node.

Line-Editing Keys

Table 2 describes the line-editing keys available using the CLI.

Table 2. Line-Editing Keys

Key(s)	Description
Left arrow or [Ctrl] + B	Moves the cursor one character to the left.
Right arrow or [Ctrl] + F	Moves the cursor one character to the right.
[Ctrl] + H or Backspace	Deletes character to left of cursor and shifts remainder of line to left.
Delete or [Ctrl] + D	Deletes character under cursor and shifts remainder of line to left.
[Ctrl] + K	Deletes characters from under cursor to end of line.
Insert	Toggles on and off. When toggled on, inserts text and shifts previous text to right.
[Ctrl] + A	Moves cursor to first character in line.
[Ctrl] + E	Moves cursor to last character in line.
[Ctrl] + L	Clears screen and moves cursor to beginning of line.
[Ctrl] + P or Up Arrow	Displays previous command in command history buffer and places cursor at end of command.
[Ctrl] + N or Down Arrow	Displays next command in command history buffer and places cursor at end of command.
[Ctrl] + U	Clears all characters typed from cursor to beginning of line.
[Ctrl] + W	Deletes previous word.
[Ctrl] + C	Interrupts the current CLI command execution.

Command History

The NETGEAR 8800 saves the commands you enter. You can display a list of these commands by using the following command:

```
history
```

If you use a command more than once, consecutively, the history will list only the first instance.

Commands for Accessing the Switch

2

This chapter describes commands used for:

- Accessing and configuring the switch including how to set up user accounts, passwords, date and time settings, and software licenses
- Managing passwords
- Configuring the Domain Name Service (DNS) client
- Checking basic switch connectivity
- Enabling and displaying licenses
- Returning the switch to safe defaults mode

NETGEAR 8800 supports the following two levels of management:

- User
- Administrator

A user-level account has viewing access to all manageable parameters, with the exception of:

- User account database
- SNMP community strings

A user-level account can change the password assigned to the account name and use the `ping` command to test device reachability.

An administrator-level account can view and change all switch parameters. It can also add and delete users and change the password associated with any account name. The administrator can disconnect a management session that has been established by way of a Telnet connection. If this happens, the user logged on by way of the Telnet connection is notified that the session has been terminated.

The DNS client in NETGEAR 8800 augments certain commands to accept either IP addresses or host names. For example, DNS can be used during a Telnet session when you are accessing a device or when using the `ping` command to check the connectivity of a device.

The switch offers the following commands for checking basic connectivity:

- `ping`
- `traceroute`

The `ping` command enables you to send Internet Control Message Protocol (ICMP) echo messages to a remote IP device. The `traceroute` command enables you to trace the routed path between the switch and a destination endstation.

This chapter describes commands for enabling and displaying software, security, and feature pack licenses.

clear account logout

```
clear account [all | <name>] logout
```

Description

This command re-enables an account that has been locked out (disabled) for exceeding the permitted number failed login attempts, which was configured by using the `configure account password-policy logout-on-login-failures` command.

Syntax Description

all	Specifies all users.
name	Specifies an account name.

Usage Guidelines

This command applies to sessions at the console port of the switch as well as all other sessions. You can re-enable both user and administrative accounts, once they have been disabled for exceeding the three failed login attempts.

Note: The failsafe accounts are never locked out.

This command clears only the locked-out (or disabled) condition of the account. The action of locking out accounts following the failed login attempts *remains* until you turn it off by issuing the `configure account [all | <name>] password-policy logout-on-login failures off` command.

Example

The following command re-enables the account `finance`, which had been locked out (disabled) for exceeding 3 consecutive failed login attempts:

```
clear account finance logout
```

clear license-info

```
clear license-info
```

Description

This command, which should be used only in conjunction with a representative from NETGEAR, clears the licensing information from the switch.

Syntax Description

This command has no variables or parameters.

Default

N/A.

Usage Guidelines

Note: Use this command only under the guidance of an NETGEAR representative.

This command clears licensing information from the switch. When you issue this command, the system requests a confirmation. If you answer yes, the system sends a Warning message to the log.

Example

The following command removes licensing information from the switch:

```
clear license-info
```

clear session

```
clear session [history | <sessId> | all]
```

Description

Terminates a Telnet and/or SSH2 sessions from the switch.

Syntax Description

Default

N/A.

Usage Guidelines

An administrator-level account can disconnect a management session that has been established by way of a Telnet connection. You can determine the session number of the session you want to terminate by using the `show session` command. The `show session` output displays information about current Telnet and/or SSH2 sessions including:

- The session number
- The login date and time
- The user name
- The type of Telnet session
- Authentication information

Depending on the software version running on your switch, additional session information may be displayed. The session number is the first number displayed in the `show session` output.

When invoked to clear the session history, the command clears the information about all the previous sessions that were logged. The information about the active sessions remains intact.

Example

The following command terminates session 4 from the system:

```
clear session 4
```

configure account

```
configure account [all | <name>]
```

Description

Configures a password for the specified account, either user account or administrative account.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.

Default

N/A.

Usage Guidelines

You must create a user or administrative account before you can configure that account with a password. Use the `create account` command to create a user account.

The system prompts you to specify a password after you enter this command. You must enter a password for this command; passwords cannot be null and cannot include the following characters: "<", ">", and "?".

Note: Once you issue this command, you cannot have a null password. However, if you want to have a null password (that is, no password on the specified account), use the `create account` command.

Passwords can have a minimum of 0 character and can have a maximum of 32 characters. Both passwords and user names are case-sensitive.

Note: If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

You must have administrator privileges to change passwords for accounts other than your own.

Example

The following command defines a new password *green* for the account *marketing*:

```
configure account marketing
```

The switch responds with a password prompt:

```
password: green
```

Your keystrokes will not be echoed as you enter the new password. After you enter the password, the switch will then prompt you to reenter it.

```
Reenter password: green
```

Assuming you enter it successfully a second time, the password is now changed.

configure account encrypted

```
configure account [all | <name>] encrypted <e-password>
```

Description

Encrypts the password that is entered in plain text for the specified account, either user account or administrative account.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.
e-password	Enter in plain text the string you for an encrypted password. See Usage Guidelines for more information.

Default

N/A.

Usage Guidelines

You must create a user or administrative account before you can configure that account with a password. Use the `create account account` command to create a user account.

When you use this command, the following password that you specify in plain text is entered and displayed by the switch in an encrypted format. Administrators should enter the password in plain text. The encrypted password is then used by the switch once it encrypts the plain text password. The encrypted command should be used by the switch only to show, store, and load a system-generated encrypted password in configuration; this applies with the following commands: `save configuration`, `show configuration`, and `use configuration`.

Note: Once you issue this command, you cannot have a null password. However, if you want to have a null password (that is, no password on the specified account), use the `create account` command.

Passwords can have a minimum of 0 character and can have a maximum of 32 characters. Both passwords and user names are case-sensitive.

Note: If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

You must have administrator privileges to change passwords for accounts other than your own.

Example

The following command encrypts the password *red* for the account *marketing*:

```
configure account marketing encrypted red
```

configure account password-policy char-validation

```
configure account [all | <name>] password-policy char-validation [none | all-char-groups]
```

Description

Requires that the user include an upper-case letter, a lower-case letter, a digit, and a symbol in the password.

Syntax Description

all	Specifies all users (and future users).
name	Specifies an account name.
none	Resets password to accept all formats.
all-char-groups	Specifies that the password must contain at least two characters from each of the four groups.

Note: The password minimum length will be 8 characters if you specify this option.

Default

N/A.

Usage Guidelines

This feature is disabled by default.

Once you issue this command, each password must include at least *two* characters of each of the following four types:

- Upper-case A-Z
- Lower-case a-z
- 0-9
- !, @, #, \$, %, ^, *, (,)

The minimum number of characters for these specifically formatted passwords is 8 characters and the maximum is 32 characters.

Use the `none` option to reset the password to accept all formats.

Example

The following command requires all users to use this specified format for all passwords:

```
configure account all password-policy char-validation all-char-groups
```

configure account password-policy history

```
configure account [all | <name>] password-policy history [<num_passwords> | none]
```

Description

Configures the switch to verify the specified number of previous passwords for the account. The user is prevented from changing the password on a user or administrative account to any of these previously saved passwords.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.
num_passwords	Specifies the number of previous passwords the system verifies for each account. The range is 1 to 10 passwords.
none	Resets the system to not remember any previous passwords.

Default

N/A.

Usage Guidelines

Use this command to instruct the system to verify new passwords against a list of all previously used passwords, once an account successfully changes a password. The limit is the number of previous passwords that the system checks against in the record to verify the new password.

If this parameter is configured, the system returns an error message if a user attempts to change the password to one that is saved by the system (up to the configured limit) for that account; this applies to both user and administrative accounts. This also applies to a configured password on the default admin account on the switch.

The limit of previous passwords that the system checks for previous use is configurable from 1 to 10. Using the `none` option disables previous password tracking and returns the system to the default state of no record of previous passwords.

Example

The following command instructs the system to verify that the new password has not been used as a password in the previous 5 passwords for the account engineering:

```
configure account engineering password-policy history 5
```

configure account password-policy lockout-on-login-failures

```
configure account [all | <name>] password-policy lockout-on-login-failures [on | off]
```

Description

Disables an account after the user has 3 consecutive failed login attempts.

Syntax Description

all	Specifies all users (and future users).
name	Specifies an account name.

on	Specifies an account name.
off	Resets the password to never lockout the user.

Default

N/A.

Usage Guidelines

If you are not working on SSH, you can configure the number of failed logins that trigger lockout, using the `configure cli max-failed-logins <num-of-logins>` command.

This command applies to sessions at the console port of the switch as well as all other sessions and to user-level and administrator-level accounts. This command locks out the user after 3 consecutive failed login attempts; the user's account must be specifically re-enabled by an administrator.

Using the `off` option resets the account to allow innumerable consecutive failed login attempts, which is the system default. The system default is that 3 failed consecutive login attempts terminate the particular session, but the user may launch another session; there is no lockout feature by default.

Note: The failsafe accounts are never locked out, no matter how many consecutive failed login attempts.

Example

The following command enables the account finance for lockout. After 3 consecutive failed login attempts, the account is subsequently locked out:

```
configure account finance password-policy lockout-on-login-failures on
```

configure account password-policy max-age

```
configure account [all | <name>] password-policy max-age [<num_days> | none]
```

Description

Configures a time limit for the passwords for specified accounts. The passwords for the default admin account and the failsafe account do not age out.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.

num_days	Specifies the length of time that a password can be used. The range is 1 to 365 days.
none	Resets the password to never expire.

Default

N/A.

Usage Guidelines

The passwords for the default admin account and the failsafe account never expire.

The time limit is specified in days, from 1 to 365 days. Existing sessions are not closed when the time limit expires; it will not open the next time the user attempts to log in.

When a user logs into an account with an expired password, the system first verifies that the entered password had been valid prior to expiring and then prompts the user to change the password.

Note: This is the sole time that a user with a user-level (opposed to an administrator-level) account can make any changes to the user-level account.

Using the `none` option prevents the password for the specified account from ever expiring (it resets the password to the system default of no time limit).

Example

The following command sets a 3-month time limit for the password for the account marketing:

```
configure account marketing password-policy max-age 90
```

configure account password-policy min-length

```
configure account [all | <name>] password-policy min-length [<num_characters> | none]
```

Description

Requires a minimum number of characters for passwords.

Syntax Description

all	Specifies all accounts (and future users).
name	Specifies an account name.

num_characters	Specifies the minimum number of characters required for the password. The range is 1 to 32 characters. Note: If you configure the <code>configure account password-policy char-validation</code> parameter, the minimum length is 8 characters.
none	Resets password to accept a minimum of 0 characters. Note: If you configure the <code>configure account encrypted</code> parameter, the minimum length is 8 characters.

Default

N/A.

Usage Guidelines

Use this command to configure a minimum length restriction for all passwords for specified accounts. This command affects the minimum allowed length for the *next* password; the current password is unaffected.

The minimum password length is configurable from 1 to 32 characters. Using the `none` option disables the requirement of minimum password length and returns the system to the default state (password minimum is 0 by default).

Note: If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

Example

The following command requires a minimum of 8 letters for the password for the account management:

```
configure account management password-policy min-length 8
```

configure banner

```
configure banner {acknowledge}
```

Description

Configures the banner string that is displayed at the beginning of each login prompt of each session.

Syntax Description

acknowledge	Specifies that the system return the user-defined message after the banner is displayed. The user must then press a key (any key) to accept before the login displays. Certain systems require this configuration (for example, the U.S. Department of Defense).
-------------	--

Default

N/A.

Usage Guidelines

Press [Return] at the beginning of a line to terminate the command and apply the banner. To clear the banner, press [Return] at the beginning of the first line. You can enter up to 24 rows of 79-column text that is displayed before the login prompt of each session. To disable the acknowledgement feature, use the `configure banner` command omitting the `acknowledge` parameter.

Note: The system does not wait for a keypress when you use SSH for access; this only applies to the serial console login sessions and telnet sessions.

Example

The following command adds a banner, *Welcome to the switch*, before the login prompt:

```
configure banner [Return]
Welcome to the switch
```

configure cli max-sessions

```
configure cli max-sessions <num-of-sessions>
```

Description

Limits number of simultaneous CLI sessions on the switch.

Syntax Description

num-of-sessions	Specifies the maximum number of concurrent sessions permitted. The range is 1 to 16.
-----------------	--

Default

The default is eight sessions.

Usage Guidelines

The value must be greater than 0; the range is 1 to 16.

Example

The following command limits the number of simultaneous CLI sessions to ten:

```
configure cli max-sessions 10
```

configure cli max-failed-logins

```
configure cli max-failed-logins <num-of-logins>
```

Description

Establishes the maximum number of failed logins permitted before the session is terminated.

Syntax Description

num-of-logins	Specifies the maximum number of failed logins permitted; the range is 1 to 10.
---------------	--

Default

The default is three logins.

Usage Guidelines

The value must be greater than 0; the range is 1 to 10.

Example

The following command sets the maximum number of failed logins to five:

```
configure cli max-failed-logins 5
```

configure dns-client add

```
configure dns-client add [domain-suffix <domain_name> | name-server <ip_address> {vr  
<vr_name>}]
```

Description

Adds a domain suffix to the domain suffix list or a name server to the available server list for the DNS client.

Syntax Description

domain-suffix	Specifies adding a domain suffix.
domain_name	Specifies a domain name.

name-server	Specifies adding a name server.
ip_address	Specifies an IP address for the name server.
vr	Specifies use of a virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
vr_name	Specifies a virtual router.

Default

N/A.

Usage Guidelines

The domain suffix list can include up to six items. If the use of all previous names fails to resolve a name, the most recently added entry on the domain suffix list will be the last name used during name resolution. This command will not overwrite any exiting entries. If a null string is used as the last suffix in the list, and all other lookups fail, the name resolver will attempt to look up the name with no suffix.

Up to eight DNS name servers can be configured. The default value for the virtual router used by the DNS client option is VR-Default.

Examples

The following command configures a domain name and adds it to the domain suffix list:

```
configure dns-client add domain-suffix xyz_inc.com
```

The following command specifies that the switch use the DNS server 10.1.2.1:

```
configure dns-client add name-server 10.1.2.1
```

The following command specifies that the switch use the virtual router Management:

```
configure dns-client add name-server 10.1.2.1 vr "VR-Mgmt"
```

configure dns-client default-domain

```
configure dns-client default-domain <domain_name>
```

Description

Configures the domain that the DNS client uses if a fully qualified domain name is not entered.

Syntax Description

domain_name	Specifies a default domain name.
-------------	----------------------------------

Default

N/A.

Usage Guidelines

The default domain name will be used to create a fully qualified host name when a domain name is not specified. For example, if the default domain name is set to “food.com” then when a command like “ping dog” is entered, the ping will actually be executed as “ping dog.food.com”.

Example

The following command configures the default domain name for the server:

```
configure dns-client default-domain xyz_inc.com
```

configure dns-client delete

```
configure dns-client delete [domain-suffix <domain_name> | name-server <ip_address> {vr <vr_name>}]
```

Description

Deletes a domain suffix from the domain suffix list or a name server from the available server list for the DNS client.

Syntax Description

domain-suffix	Specifies deleting a domain suffix.
domain_name	Specifies a domain name.
name-server	Specifies deleting a name server.
ip_address	Specifies an IP address for the name server.
vr	Specifies deleting a virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
vr_name	Specifies a virtual router.

Default

N/A.

Usage Guidelines

Specifying a domain suffix removes an entry from the domain suffix list. If the deleted item was not the last entry in the list, all items that had been added later are moved up in the list. If no entries in the list match the domain name specified, an error message will be displayed.

The default value for the virtual router used by the DNS client option is VR-Default.

Examples

The following command deletes a domain name from the domain suffix list:

```
configure dns-client delete domain-suffix xyz_inc.com
```

The following command removes a DNS server from the list:

```
configure dns-client delete name-server 10.1.2.1
```

configure failsafe-account

```
configure failsafe-account {[deny | permit] [all | control | serial | ssh  
{vr <vr-name>} | telnet {vr <vr-name>}]}
```

Description

Configures a name and password for the failsafe account, or restricts access to specified connection types.

Syntax Description

deny	Prohibits failsafe account usage over the specified connection type(s).
permit	Allows a failsafe account to be used over the specified connection type(s).
all	Specifies all connection types.
control	Specifies internal access between nodes in a NETGEAR 8800 or between MSMs/MMs in a chassis.
serial	Specifies access over the switch console port.
ssh	Specifies access using SSH on specified or all virtual routers.
telnet	Specifies access using Telnet on specified or all virtual routers.

Default

The failsafe account is always configured. The default connection types over which failsafe account access is permitted are the same as if “permit all” is configured.

Usage Guidelines

The failsafe account is the account of last resort to access your switch.

If you use the command with no parameters, you are prompted for the failsafe account name and prompted twice to specify the password for the account. The password does not appear on the display at any time. You are not required to know the current failsafe account and password in order to change it.

If you use the command with the permit or deny parameter, the permitted connection types are altered as specified.

The failsafe account or permitted connection types are immediately saved to NVRAM on all MSMs/MMs or active nodes.

Note: The information that you use to configure the failsafe account cannot be recovered by NETGEAR. Technical support cannot retrieve passwords or account names for this account. Protect this information carefully.

Once you enter the failsafe account name, you are prompted to enter the password. Once you successfully log in to the failsafe account, you are logged in to an admin-level account.

Example

The following command changes the failsafe account: username to `blue5green` and the password to `red5yellow`.

```
XCM8806.1 # configure failsafe-account
enter failsafe user name: blue5green
enter failsafe password:
enter password again:
XCM8806.2
```

The following example restricts usage of the failsafe account to the series console port and to access between MSMs.

```
XCM8810.1 # configure failsafe-account deny all
XCM8810.2 # configure failsafe-account permit serial
XCM8810.3 # configure failsafe-account permit control
XCM8810.4 #
```

configure idletimeout

```
configure idletimeout <minutes>
```

Description

Configures the time-out for idle console, SSH2, and Telnet sessions.

Syntax Description

minutes	Specifies the time-out interval, in minutes. Range is 1 to 240 (1 minute to 4 hours).
---------	---

Default

The default time-out is 20 minutes.

Usage Guidelines

This command configures the length of time the switch will wait before disconnecting idle console, SSH2, or Telnet sessions. The `idletimeout` feature must be enabled for this command to have an effect (the `idletimeout` feature is enabled by default).

Example

The following command sets the time-out for idle login and console sessions to 10 minutes:

```
configure idletimeout 10
```

configure safe-default-script

```
configure safe-default-script
```

Description

Allows you to change management access to your device and to enhance security.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command runs an interactive script that prompts you to choose to enable or disable SNMP, Telnet, and enabled ports. Refer to the “Safe Defaults Setup Method” section in the *NETGEAR 8800 User Manual* for complete information on the safe default mode.

Once you issue this command, the system presents you with the following interactive script:

```
Telnet is enabled by default. Telnet is unencrypted and has been the target of
security exploits in the past.
```

```
Would you like to disable Telnet? [y/N]:
```

```
SNMP access is enabled by default. SNMP uses no encryption, SNMPv3 can be
configured to eliminate this problem.
```

```
Would you like to disable SNMP? [y/N]:
```

```
All ports are enabled by default. In some secure applications, it maybe more
```

desirable for the ports to be turned off.

Would you like unconfigured ports to be turned off by default? [y/N]:

Changing the default failsafe account username and password is highly recommended. If you choose to do so, please remember the username and password as this information cannot be recovered by NETGEAR.

Would you like to change the failsafe account username and password now? [y/N]:

Would you like to permit failsafe account access via the management port? [y/N]:

Since you have chosen less secure management methods, please remember to increase the security of your network by taking the following actions:

- * change your admin password
- * change your failsafe account username and password
- * change your SNMP public and private strings
- * consider using SNMPv3 to secure network management traffic

Example

The following command reruns the interactive script to configure management access:

```
configure safe-default-script
```

configure time

```
configure time <month> <day> <year> <hour> <min> <sec>
```

Description

Configures the system date and time.

Syntax Description

month	Specifies the month. The range is 1-12.
day	Specifies the day of the month. The range is 1-31.
year	Specifies the year in the YYYY format. The range is 2003 to 2036.
hour	Specifies the hour of the day. The range is 0 (midnight) to 23 (11 pm).
min	Specifies the minute. The range is 0-59.
sec	Specifies the second. The range is 0-59.

Default

N/A.

Usage Guidelines

The format for the system date and time is as follows:

```
mm dd yyyy hh mm ss
```

The time uses a 24-hour clock format. You cannot set the year earlier than 2003 or past 2036. You have the choice of inputting the entire time/date string. If you provide one item at a time and press TAB, the screen prompts you for the next item. Press <cr> to complete the input.

Example

The following command configures a system date of February 15, 2002 and a system time of 8:42 AM and 55 seconds:

```
configure time 02 15 2002 08 42 55
```

configure timezone

```
configure timezone {name <tz_name>} <GMT_offset>
{autodst {name <dst_timezone_ID>} {<dst_offset>}}
{begins [every <floatingday> | on <absoluteday>] {at <time_of_day>}}
{ends [every <floatingday> | on <absoluteday>] {at <time_of_day>}}}
| noautodst}
```

Description

Configures the Greenwich Mean Time (GMT) offset and Daylight Saving Time (DST) preference.

Syntax Description

tz_name	Specifies an optional name for this timezone specification. May be up to six characters in length. The default is an empty string.
GMT_offset	Specifies a Greenwich Mean Time (GMT) offset, in + or - minutes.
autodst	Enables automatic Daylight Saving Time.
dst-timezone-ID	Specifies an optional name for this DST specification. May be up to six characters in length. The default is an empty string.
dst_offset	Specifies an offset from standard time, in minutes. Value is in the range of 1 to 60. Default is 60 minutes.

floatingday	<p>Specifies the day, week, and month of the year to begin or end DST each year. Format is:</p> <p><week> <day> <month> where:</p> <ul style="list-style-type: none"> • <week> is specified as [first second third fourth last] or 1-5. • <day> is specified as [sunday monday tuesday wednesday thursday friday saturday] or 1-7 (where 1 is Sunday). • <month> is specified as [january february march april may june july august september october november december] or 1-12. <p>Default for beginning is second sunday march; default for ending is first sunday november.</p>
absoluteday	<p>Specifies a specific day of a specific year on which to begin or end DST. Format is:</p> <p><month> <day> <year> where:</p> <ul style="list-style-type: none"> • <month> is specified as 1-12. • <day> is specified as 1-31. • <year> is specified as 2003-2035. <p>The year must be the same for the begin and end dates.</p>
time_of_day	<p>Specifies the time of day to begin or end Daylight Saving Time. May be specified as an hour (0-23) or as hour:minutes. Default is 2:00.</p>
noautodst	<p>Disables automatic Daylight Saving Time.</p>

Default

`Autodst`, beginning every second Sunday in March, and ending every first Sunday in November.

Usage Guidelines

Network Time Protocol (NTP) server updates are distributed using GMT time. To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographic location.

The `gmt_offset` is specified in +/- minutes from the GMT time.

Automatic DST changes can be enabled or disabled. The default configuration, where DST begins on the second Sunday in March at 2:00 AM and ends the first Sunday in November at 2:00 AM, applies to most of North America (beginning in 2007), and can be configured with the following syntax:

```
configure timezone <gmt_offset> autodst.
```

The starting and ending date and time for DST may be specified, as these vary in time zones around the world.

- Use the `every` keyword to specify a year-after-year repeating set of dates (for example, the last Sunday in March every year)
- Use the `on` keyword to specify a non-repeating, specific date for the specified year. If you use this option, you will need to specify the command again every year.
- The `begins` specification defaults to `every second sunday march`.

- The `ends` specification defaults to `every first sunday november`.
- The `ends` date may occur earlier in the year than the `begins` date. This will be the case for countries in the Southern Hemisphere.
- If you specify only the starting or ending time (not both) the one you leave unspecified will be reset to its default.
- The `time_of_day` specification defaults to `2:00`.
- The timezone IDs are optional. They are used only in the display of timezone configuration information in the `show switch` command.

To disable automatic DST changes, re-specify the GMT offset using the `noautodst` option:
`configure timezone <gmt_offst> noautodst.`

NTP updates are distributed using GMT time. To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographical location. [Table 3](#) describes the GMT offsets.

Table 3. Greenwich Mean Time offsets

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+0:00	+0	GMT - Greenwich Mean UT or UTC - Universal (Coordinated) WET - Western European	London, England; Dublin, Ireland; Edinburgh, Scotland; Lisbon, Portugal; Reykjavik, Iceland; Casablanca, Morocco
-1:00	-60	WAT - West Africa	Cape Verde Islands
-2:00	-120	AT - Azores	Azores
-3:00	-180		Brasilia, Brazil; Buenos Aires, Argentina; Georgetown, Guyana;
-4:00	-240	AST - Atlantic Standard	Caracas; La Paz
-5:00	-300	EST - Eastern Standard	Bogota, Columbia; Lima, Peru; New York, NY, Trevor City, MI USA
-6:00	-360	CST - Central Standard	Mexico City, Mexico
-7:00	-420	MST - Mountain Standard	Saskatchewan, Canada
-8:00	-480	PST - Pacific Standard	Los Angeles, CA, Cupertino, CA, Seattle, WA USA
-9:00	-540	YST - Yukon Standard	
-10:00	-600	AHST - Alaska-Hawaii Standard CAT - Central Alaska HST - Hawaii Standard	
-11:00	-660	NT - Nome	
-12:00	-720	IDLW - International Date Line West	

Table 3. Greenwich Mean Time offsets (Continued)

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+1:00	+60	CET - Central European FWT - French Winter MET - Middle European MEWT - Middle European Winter SWT - Swedish Winter	Paris, France; Berlin, Germany; Amsterdam, The Netherlands; Brussels, Belgium; Vienna, Austria; Madrid, Spain; Rome, Italy; Bern, Switzerland; Stockholm, Sweden; Oslo, Norway
+2:00	+120	EET - Eastern European, Russia Zone 1	Athens, Greece; Helsinki, Finland; Istanbul, Turkey; Jerusalem, Israel; Harare, Zimbabwe
+3:00	+180	BT - Baghdad, Russia Zone 2	Kuwait; Nairobi, Kenya; Riyadh, Saudi Arabia; Moscow, Russia; Tehran, Iran
+4:00	+240	ZP4 - Russia Zone 3	Abu Dhabi, UAE; Muscat; Tblisi; Volgograd; Kabul
+5:00	+300	ZP5 - Russia Zone 4	
+5:30	+330	IST – India Standard Time	New Delhi, Pune, Allahabad, India
+6:00	+360	ZP6 - Russia Zone 5	
+7:00	+420	WAST - West Australian Standard	
+8:00	+480	CCT - China Coast, Russia Zone 7	
+9:00	+540	JST - Japan Standard, Russia Zone 8	
+10:00	+600	EAST - East Australian Standard GST - Guam Standard Russia Zone 9	
+11:00	+660		
+12:00	+720	IDLE - International Date Line East NZST - New Zealand Standard NZT - New Zealand	Wellington, New Zealand; Fiji, Marshall Islands

For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

Example

The following command configures GMT offset for Mexico City, Mexico and disables automatic DST:

```
configure timezone -360 noautodst
```

The following four commands are equivalent, and configure the GMT offset and automatic DST adjustment for the US Eastern timezone, with an optional timezone ID of EST:

```
configure timezone name EST -300 autodst name EDT 60 begins every second sunday march at 2:00
ends every first sunday november at 2:00
```

```
configure timezone name EST -300 autodst name EDT 60 begins every 1 1 4 at 2:00 ends every 5
1 10 at 2:00
```

```
configure timezone name EST -300 autodst name EDT
```

```
configure timezone -300 autodst
```

The following command configures the GMT offset and automatic DST adjustment for the Middle European timezone, with the optional timezone ID of MET:

```
configure timezone name MET 60 autodst name MDT begins every last sunday march at 1 ends every
last sunday october at 1
```

The following command configures the GMT offset and automatic DST adjustment for New Zealand. The ending date must be configured each year because it occurs on the first Sunday on or after March 5:

```
configure timezone name NZST 720 autodst name NZDT 60 begins every first sunday october at 2
ends on 3/16/2002 at 2
```

create account

```
create account [admin | user] <account-name> {encrypted <password>}
```

Description

Creates a new user account.

Syntax Description

admin	Specifies an access level for account type admin.
user	Specifies an access level for account type user.
account-name	Specifies a new user account name. See Usage Guidelines for more information.
encrypted	Specifies the encrypted option.
password	Specifies a user password. See Usage Guidelines for more information.

Default

By default, the switch is configured with two accounts with the access levels shown in [Table 4](#).

Table 4. User account levels

Account Name	Access Level
admin	This user can access and change all manageable parameters. The admin account cannot be deleted.
user	<p>This user can view (but not change) all manageable parameters, with the following exceptions:</p> <ul style="list-style-type: none"> • This user cannot view the user account database. • This user cannot view the SNMP community strings. • This user cannot view SSL settings. <p>This user has access to the <code>ping</code> command.</p>

You can use the default names (*admin* and *user*), or you can create new names and passwords for the accounts. Default accounts do not have passwords assigned to them. For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

Usage Guidelines

The switch can have a total of 16 user accounts. The system must have one administrator account.

When you use the `encrypted` keyword, the following password that you specify in plain text is entered and displayed by the switch in an encrypted format. Administrators should not use the encrypted option and should enter the password in plain text. The encrypted option is used by the switch after encrypting the plain text password. The encrypted option should be used by the switch only to show, store, and load a system-generated encrypted password in configuration; this applies with the following commands: `save configuration`, `show configuration`, and `use configuration`.

The system prompts you to specify a password after you enter this command and to reenter the password. If you do not want a password associated with the specified account, press *Enter twice*.

You must have administrator privileges to change passwords for accounts other than your own. User names and passwords are case-sensitive. User account names must have a minimum of 1 character and can have a maximum of 32 characters. Passwords must have a minimum of 0 characters and can have a maximum of 32 characters.

Note: If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

Example

The following command creates a new account named John2 with administrator privileges:

```
create account admin John2
```

delete account

```
delete account <name>
```

Description

Deletes a specified user account.

Syntax Description

name	Specifies a user account name.
------	--------------------------------

Default

N/A.

Usage Guidelines

Use the `show accounts` command to determine which account you want to delete from the system. The show accounts output displays the following information in a tabular format:

- The user name
- Access information associated with each user
- User login information
- Session information

Depending on the software version running on your switch and the type of switch you have, additional account information may be displayed.

You must have administrator privileges to delete a user account. The system must have one administrator account; the command will fail if an attempt is made to delete the last administrator account on the system.

To ensure security, change the password on the default account, but do not delete it. The changed password will remain intact through configuration uploads and downloads.

If you must delete the default account, first create another administrator-level account.

Example

The following command deletes account John2:

```
delete account John2
```

disable cli space-completion

```
disable cli space-completion
```

Description

Disables the NETGEAR 8800 feature that completes a command automatically with the spacebar. If you disable this feature, you can still use the TAB key for auto-completion.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables using the spacebar to automatically complete a command:

```
disable cli space-completion
```

disable clipaging

```
disable clipaging
```

Description

Disables pausing at the end of each show screen.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment. Most `show` command output will pause when the display reaches the end of a page. This command disables the pause mechanism and allows the display to print continuously to the screen.

CLI paging is only active on a per-shell session basis. In other words, when you enable or disable CLI paging from within the current configuration, it only affects that session. For new or existing sessions, paging is enabled by default. This setting cannot be saved.

To view the status of CLI paging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for CLI paging.

Example

The following command disables clipaging and allows you to print continuously to the screen:

```
disable clipaging
```

disable idletimeout

```
disable idletimeout
```

Description

Disables the timer that disconnects idle sessions from the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled. Timeout 20 minutes.

Usage Guidelines

When idle time-outs are disabled, console sessions remain open until the switch is rebooted or until you logoff. Telnet sessions remain open until you close the Telnet client.

If you have an SSH2 session and disable the idle timer, the SSH2 connection times out after 61 minutes of inactivity.

To view the status of idle time-outs on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for idle time-outs.

Example

The following command disables the timer that disconnects all sessions to the switch:

```
disable idletimeout
```

enable cli space-completion

```
enable cli space-completion
```

Description

Enables the NETGEAR 8800 feature that completes a command automatically with the spacebar. You can also use the TAB key for auto-completion.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables using the spacebar to automatically complete a command:

```
enable cli space-completion
```

enable clipaging

```
enable clipaging
```

Description

Enables the pause mechanism and does not allow the display to print continuously to the screen.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment. Most `show` command output will pause when the display reaches the end of a page.

To view the status of CLI paging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for CLI paging.

If CLI paging is enabled and you use the `show tech` command to diagnose system technical problems, the CLI paging feature is disabled.

CLI paging is only active on a per-shell session basis. In other words, when you enable or disable CLI paging from within the current configuration, it only affects that session. For new or existing sessions, paging is enabled by default. This setting cannot be saved.

Example

The following command enables clipaging and does not allow the display to print continuously to the screen:

```
enable clipaging
```

enable idletimeout

```
enable idletimeout
```

Description

Enables a timer that disconnects Telnet, SSH2, and console sessions after a period of inactivity (20 minutes is default).

Syntax Description

This command has no arguments or variables.

Default

Enabled. Timeout 20 minutes.

Usage Guidelines

You can use this command to ensure that a Telnet, Secure Shell (SSH2), or console session is disconnected if it has been idle for the required length of time. This ensures that there are no hanging connections.

To change the period of inactivity that triggers the timeout for a Telnet, SSH2, or console session, use the `configure timezone` command.

To view the status of idle timeouts on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for idle timeouts. You can configure the length of the timeout interval.

Example

The following command enables a timer that disconnects any Telnet, SSH2, and console sessions after 20 minutes of inactivity:

```
enable idletimeout
```

enable license software

```
enable license {software} <key>
```

Description

Enables software license or feature pack that allows you to use advanced features.

Syntax Description

key	Specifies your hexadecimal license key in format xxxx-xxxx-xxxx-xxxx-xxxx.
-----	--

Default

N/A

Usage Guidelines

The software license levels that apply to NETGEAR 8800 software are described in Appendix A of the *NETGEAR 8800 User Manual*.

To obtain a software license, specify the `key` in the format `xxxx-xxxx-xxxx-xxxx-xxxx`.

You obtain the software license key (or feature pack key) either by ordering it from the factory or by obtaining a license voucher from your NETGEAR supplier. You can obtain a regular software license or a trial software license, which allows you use of the license for either 30, 60 or 90 days; you cannot downgrade software licenses.

The voucher contains all the necessary information on the software license, whether regular or trial, and number of days for trial software license.

After you enable the software license or feature pack by entering the software key, the system returns a message that you either successfully or unsuccessfully set the license.

Once you enable the software license (or if you do not use the correct key, attempt to downgrade the license, or already installed the software license) you see one of the following messages:

```
Enabled license successfully.
```

```
Error: Unable to set license using supplied key.
```

```
Error: Unable to set license - downgrade of licenses is not supported.
```

```
Error: Unable to set license - license is already enabled.
```

```
Error: Unable to set license - trial license already enabled.
```

If you enable a trial license, the system generates a daily message showing the number of days until expiry.

Once installed (or enabled), the software license goes with the switch chassis itself (not with the MSM/MM module). The software license information is stored in EEPROM; the information persists through reboots, software upgrades, power outages, and reconfigurations.

If you attempt to execute a command and you do not either have the required software license or have reached the limits defined by the current software license level, the system returns one of the following messages:

```
Error: This command cannot be executed at the current license level.
```

```
Error: You have reached the maximum limit for this feature at this license level.
```

If you attempt to execute a command and you do not have the required feature pack, the system also returns a message.

To protect against attacks to install maliciously created license keys, the system has an exponential delay of each failed attempt to install a license.

To view the type of software license you are currently running on the switch, use the `show licenses` command. The license key number is not displayed, but the type of software

license is displayed in the `show licenses` output. This command can be run on any node in a NETGEAR 8800, regardless of its node role (Master, Standby, or Backup).

Example

The following command enables a software license on the switch:

```
enable license 2d5e-0e84-e87d-c3fe-bfff
```

enable license file

```
enable license file <filename>
```

Description

Enables the text file that applies software licenses and feature packs licenses to more than one switch at a time.

Syntax Description

filename	Specifies the filename that you download onto the switch using TFTP; the file extension is .xlic.
----------	---

Default

N/A

Usage Guidelines

You download the license file to the switch using TFTP or SCP. The file name extension for this file is <xlic>; for example, you may see a file named systemlic.xlic.

Using this file, you enable the software and feature pack licenses for more than one switch simultaneously. The file can contain licenses for some or all of the NETGEAR switches that the customer owns. During upload, only those license keys destined for the specific switch are used to attempt enabling the licenses. The license file is a text file that has the switch serial number, software license type, and license key; it is removed from the switch after the licenses are enabled.

After you enable the license file, the system returns one or more of the following messages:

```
Enabled license successfully.
```

```
Error: Unable to set license <license_name> using supplied key.
```

```
Error: Unable to set license <license_name> - downgrade of licenses is not supported.
```

```
Error: Unable to set license <license_name> - license is already enabled.
```

```
Error: Unable to set license <license_name> - trial license already enabled.
```

To protect against attacks to install maliciously created license keys, the system has an exponential delay of each failed attempt to install a license.

Example

The following command enables a license file on the specified NETGEAR switches:

```
enable license file santaclara.xlic
```

history

```
history
```

Description

Displays a list of all the commands entered on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

NETGEAR 8800 saves the commands you entered on the switch. Use the `history` command to display a list of these commands.

Example

The following command displays all the commands entered on the switch:

```
history
```

If you use a command more than once consecutively, the history will list only the first instance.

ping

```
ping {count <count> {start-size <start-size>} | continuous {start-size <start-size>} |  
{start-size <start-size> {end-size <end-size>}}} {udp} {dont-fragment} {ttl <ttl>} {tos  
<tos>} {interval <interval>} {vr <vrid>} {ipv4 <host> | ipv6 <host>} {from} {with  
record-route}
```

Description

Enables you to send User Datagram Protocol (UDP) or Internet Control Message Protocol (ICMP) echo messages or to a remote IP device.

Syntax Description

count	Specifies the number of ping requests to send.
-------	--

start-size	Specifies the size, in bytes, of the packet to be sent, or the starting size if incremental packets are to be sent.
continuous	Specifies that UDP or ICMP echo messages to be sent continuously. This option can be interrupted by pressing [Ctrl] + C.
end-size	Specifies an end size for packets to be sent.
udp	Specifies that the ping request should use UDP instead of ICMP.
dont-fragment	Sets the IP to not fragment the bit.
ttl	Sets the TTL value.
tos	Sets the TOS value.
interval	Sets the time interval between sending out ping requests.
vr	Specifies the virtual route to use for sending out the echo message. If not specified, VR-Default is used.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
ipv4	Specifies IPv4 transport.
ipv6	Specifies IPv6 transport.
	Note: If you are contacting an IPv6 link local address, you must specify the VLAN you are sending the message from: <code>ping <ipv6> <link-local address> %<vlan_name> <host></code> .
host	Specifies a host name or IP address (either v4 or v6).
from	Uses the specified source address. If not specified, the address of the transmitting interface is used.
with record-route	Sets the traceroute information.

Default

N/A.

Usage Guidelines

The `ping` command is used to test for connectivity to a specific host.

You use the `ipv6` variable to ping an IPv6 host by generating an ICMPv6 echo request message and sending the message to the specified address. If you are contacting an IPv6 link local address, you must specify the VLAN you sending the message from, as shown in the following example (you must include the % sign): `ping <ipv6> <link-local address> %<vlan_name> <host>`.

The `ping` command is available for both the user and administrator privilege level.

Example

The following command enables continuous ICMP echo messages to be sent to a remote host:

```
ping continuous 123.45.67.8
```

reboot

```
reboot {time <month> <day> <year> <hour> <min> <sec>} {cancel} {msm <slot_id>} {slot
<slot-number> | node-address <node-address> | stack-topology {as-standby} }
```

Description

Reboots the switch or the module in the specified slot at a specified date and time.

Syntax Description

time	Specifies a reboot date in mm dd yyyy format and reboot time in hh mm ss format.
cancel	Cancels a previously scheduled reboot.
msm	Specifies rebooting the MSM module.
slot_id	Specifies the slot--A or B--for an MSM module.
slot-number	Specifies the slot number currently being used by the active stack node that is to be rebooted
node-address	Specifies the MAC address of the node to be rebooted
stack-topology	Specifies that the entire NETGEAR 8800 is to be rebooted whether or not nodes are active
as-standby	Specifies that all stack nodes that are to be rebooted are to operate as if configured to not be master-capable

Default

N/A.

Usage Guidelines

If you do not specify a reboot time, the switch will reboot immediately following the command, and any previously scheduled reboots are cancelled. Prior to rebooting, the switch returns the following message:

```
Do you want to save configuration changes to primary and reboot?
(y - save and reboot, n - reboot without save, <cr> - cancel command)
```

To cancel a previously scheduled reboot, use the `cancel` option.

The modules that can be rebooted are management switch fabric modules (MSM)/management modules (MM).

On the NETGEAR 8800 series switches, if your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a sharp object into the “A” and “R” holes on the MSM and applying slight pressure. Refer to the hardware documentation for information on the MSM.

The `reboot MSM` option on the 8800 series switches affects the entire module.

Example

The following command reboots the switch at 8:00 AM on April 15, 2005:

```
reboot time 04 15 2005 08 00 00
```

show accounts

```
show accounts
```

Description

Displays user account information for all users on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You need to create a user account using the `create account` command before you can display user account information.

To view the accounts that have been created, you must have administrator privileges.

The `show accounts` command displays the following information in a tabular format:

- **User Name**—The name of the user. This list displays all of the users who have access to the switch.
- **Access**—This may be listed as R/W for read/write or RO for read only.
- **Login OK**—The number of logins that are okay.
- **Failed**—The number of failed logins.
- **Accounts locked out**—Account configured to be locked out after 3 consecutive failed login attempts (using the `configure account password-policy lockout-on-login-failures` command).

Note: This command does not show the failsafe account.

Example

The following command displays user account information on the switch:

```
show accounts pppuser
```

Output from this command looks similar to the following:

User Name	Access	LoginOK	Failed
admin	R/W	3	1
user	RO	0	0
dbackman	R/W	0	0
ron*	RO	0	0
nocteam	RO	0	0

(*) - Account locked

show accounts password-policy

```
show accounts password-policy
```

Description

Displays password policy information for all users on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

To view the password management information, you must have administrator privileges.

The `show accounts password-policy` command displays the following information in a tabular format:

- Global password management parameters applied to new accounts upon creation:
 - Maximum age—The maximum number of days for the passwords to remain valid.
 - History limit—The number of previous password that the switch scans prior to validating a new password.
 - Minimum length—The minimum number of characters in passwords.
 - Character validation—The passwords must be in the specific format required by the `configure account password-policy char-validation` command.
 - Lockout on login failures—If enabled, the system locks out users after 3 failed login attempts.

- Accounts locked out—Number of accounts locked out.
- **User Name**—The name of the user. This list displays all of the users who have access to the switch.
- **Password Expiry Date**—Date the password for this account expires; may be blank.
- **Password Max. age**—The number of days originally allowed to passwords on this account; may show None.
- **Password Min. length**—The minimum number of characters required for passwords on this account; may show None.
- **Password History Limit**—The number of previous passwords the system scans to disallow duplication on this account; may show None.

Example

The following command displays the password management parameters configured for each account on the switch:

```
show accounts password-policy
```

Output from this command looks similar to the following:

```
-----
Accounts global configuration(applied to new accounts on creation)
-----
Password Max. age           : None
Password History limit      : None
Password Min. length        : None
Password Character Validation : Disabled
Accts. lockout on login failures: Disabled
Accounts locked out         : No
-----

      User Name      Password  Password Password Password Flags
                   Expiry    Max. age Min. len History
                   Date
-----
              admin          None    None    None    ---
              user           None    None    None    ---
              test Apr-17-2005   12     32     9     C--
-----

Flags: (C) Password character validation enabled, (L) Account locked out
       (l) Account lockout on login failures enabled
```

show banner

```
show banner
```

Description

Displays the user-configured banner string.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed before the login prompt.

Example

The following command displays the switch banner:

```
show banner
```

Output from this command varies depending on your configuration; the following is one example:

```

NETGEAR 8800 Switch
#####
  Unauthorized Access is strictly prohibited.
  Violators will be prosecuted
#####

```

show dns-client

```
show dns-client
```

Description

Displays the DNS configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the DNS configuration:

```
show dns-client
```

Output from this command looks similar to the following:

```
Number of domain suffixes: 2
Domain Suffix 1:          njudah.local
Domain Suffix 2:          dbackman.com
Number of name servers: 2
Name Server 1:  172.17.1.104
Name Server 2:  172.17.1.123
```

show failsafe-account

```
show failsafe-account
```

Description

Displays whether the user configured a username and password for the failsafe account or shows the configured connection type access restrictions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the failsafe account configuration.

The command shows the access permissions and whether or not the user configured a username and password. It does not show the configured username or password.

Example

The following command displays the failsafe account configuration.

```
show failsafe-account
```

Output from this command looks similar to the following when a failsafe account username and password have been configured with all connections types permitted for failsafe account access:

```
BD-8810.7 # show failsafe-account
User-Specified Failsafe Account Username and Password are in effect for these connection
types:
- Serial Console
- Control Fabric (inter-node)
- Mgmt VR Telnet
- Mgmt VR SSH
- User VR Telnet
- User VR SSH
BD-8810.8 #
```

show licenses`show licenses`**Description**

Displays current software license level and feature packs enabled on your switches.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The command displays information on the software license level and feature packs enabled on the switch, including the trial license and days left to expiry.

Note: Refer to the specific chapter that discusses each feature of the *NETGEAR 8800 User Manual* to determine if a license is required for some functionality. If not noted, all functionality is available, and license is not required.

Example

The following command displays the license level configuration:

`show licenses`

Output from this command looks similar to the following:

```
XCM8806.2 # show license
Enabled License Level:
    NETGEAR AdvancedCore
Enabled Feature Packs:
    None
XCM8806.3 #
```

show switch`show switch {detail}`**Description**

Displays the current switch information.

This command displays the Master and Backup node information if executed on the Master, and displays the current node and the Master node information if executed on any other node.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show switch` command displays:

- sysName, sysLocation, sysContact
- MAC address
- System type
- System health check
- Recovery mode
- Watchdog state
- Current date, time, system boot time, and time zone configuration
- Any scheduled reboot information
- System up time
- Master and Backup information
- Current state (available only on stand-alone switches)
 - OPERATIONAL
 - OPERATIONAL (OverHeat)
 - FAILED
- Software image information (primary/secondary image and version)
- Configuration information (primary/secondary configuration and version)

This information may be useful for your technical support representative if you have a problem.

Depending on the software version running on your switch, additional or different switch information may be displayed.

On a stack the following additional information will be available:

- System Type
- System UpTime
- Details of Master and Backup, or current node and Master

Example

The following command displays current switch information:

```
show switch
```

Output from this command looks similar to the following:

```

SysName:          BD-8810Rack3
SysLocation:
SysContact:
System MAC:       00:04:96:1D:00:C0
System Type:      BD-8810

SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled

Current Time:     Fri Feb 13 02:25:24 1925
Timezone:         [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:        Wed Feb 11 21:39:56 1925
Boot Count:       159
Next Reboot:      None scheduled
System UpTime:    1 day 4 hours 45 minutes 28 seconds

Slot:             MSM-A *                MSM-B
                  -----                -----
Current State:    MASTER                  BACKUP (In Sync)

Image Selected:   secondary                secondary
Image Booted:     primary                  primary
Primary ver:      12.0.0.4                 12.0.0.4
Secondary ver:    12.0.0.4                 12.0.0.4

Config Selected:  primary.cfg              primary.cfg
Config Booted:    primary.cfg              primary.cfg

primary.cfg       Created by NETGEAR 8800 version 11.6.0.30
                  574246 bytes saved on Wed Jul 30 19:39:55 1924

```

The `show switch detail` command displays the same information shown above.

traceroute

```

traceroute {vr <vrid>} {ipv4 <host>} {ipv6 <host>} {ttl <number>} {from <from>} {[port
<port>] | icmp}

```

Description

Enables you to trace the routed path between the switch and a destination endstation.

Syntax Description

vr	Specifies a virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
vrid	Specifies which virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
ipv4	Specifies IPv4 transport.
ipv6	Specifies IPv6 transport.
host	Specifies the host of the destination endstation.
ttl <number>	Configures the switch to trace up to the time-to-live number of the switch.
from <from>	Uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used.
port <port>	Specifies the UDP port number.
icmp	Configures the switch to send ICMP echo messages to trace the routed path between the switch and a destination endstation.

Default

N/A.

Usage Guidelines

Use this command to trace the routed path between the switch and a destination endstation. Each router along the path is displayed.

Example

The following command enables the traceroute function to a destination of 123.45.67.8:

```
traceroute 123.45.67.8
```

The following is sample output that displays when the traceroute fails:

```
traceroute to 10.209.10.37, 30 hops max
 1  0.0.0.0                                * !u          * !u          * !u

--- Packet Response/Error Flags ---
(*) No response, (!N) ICMP network unreachable, (!H) ICMP host unreachable,
(!P) ICMP protocol unreachable, (!F) ICMP fragmentation needed,
(!S) ICMP source route failed, (!u) Transmit error, network unreachable,
(!f) Transmit error, fragmentation needed, (!t) General transmit error
```


Commands for Managing the Switch

3

This chapter describes commands for:

- Configuring Simple Network Management Protocol (SNMP) parameters on the switch
- Managing the switch using Telnet
- Transferring files using the Trivial File Transfer Protocol (TFTP)
- Configuring system redundancy
- Displaying power management statistics on the switch
- Configuring Simple Network Time Protocol (SNTP) parameters on the switch

SNMP

Any network manager running the Simple Network Management Protocol (SNMP) can manage the switch, if the Management Information Base (MIB) is installed correctly on the management station. Each network manager provides its own user interface to the management facilities.

The following SNMP parameters can be configured on the switch:

- Authorized trap receivers—An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMP traps to all trap receivers. Entries in this list can be created, modified, and deleted using the RMON2 trapDestTable MIB table, as described in RFC 2021, and the SNMPv3 tables.
- Authorized managers—An authorized manager can be either a single network management station, or a range of addresses (for example, a complete subnet) specified by a prefix and a mask.
- Community strings—The community strings allow a simple method of authentication between the switch and the remote network manager. The default read-only community string is *public*. The default read-write community string is *private*. The community strings for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps.
- System contact (optional)—The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the switch.
- System name (optional)—The system name enables you to enter a name that you have assigned to this switch. The default name is the model name of the switch (for example, BD-1.2).

- System location (optional)—Using the system location field, you can find the location of the switch.

Note: If you specify volatile storage when configuring SNMP parameters, that configuration is not saved across a switch reboot.

Telnet

Telnet allows you to access the switch remotely using TCP/IP through one of the switch ports or a workstation with a Telnet facility. If you access the switch via Telnet, you will use the command line interface (CLI) to manage the switch and modify switch configurations.

TFTP

NETGEAR 8800 supports the Trivial File Transfer Protocol (TFTP) based on RFC 1350. TFTP is a method used to transfer files from one network device to another. The NETGEAR 8800 TFTP client is a command line application used to contact an external TFTP server on the network. For example, the NETGEAR 8800 uses TFTP to download software image files, switch configuration files, and access control lists (ACLs) from a server on the network to the switch.

System Redundancy with Dual Management Modules Installed

If you install two MSMs/MMs, one assumes the role of primary and the other assumes the role of backup. The primary MSM/MM provides all of the switch management functions including bringing up and programming the I/O modules, running the bridging and routing protocols, and configuring the switch. The primary also keeps synchronized with the backup MSM/MM in case the backup MSM/MM needs to take over the management functions if the primary MSM/MM fails.

Power Supply Management

On the NETGEAR 8800, the 8800 OS monitors and manages power consumption on the switch by periodically checking the power supply units (PSUs) and testing them for failures. To determine the health of the PSU, the 8800 OS checks the voltage, current, and temperature of the PSU.

The power management capability of the NETGEAR 8800 OS:

- Protects the system from overload conditions.

- Monitors all installed PSUs, even installed PSUs that are disabled.
- Enables and disables PSUs as required .
- Powers up or down I/O modules based on available power and required power resources.
- Logs power resource changes, including power budget, total available power, redundancy, and so on.
- Detects and isolates faulty PSUs.

Simple Network Time Protocol

The NETGEAR 8800 supports the client portion of the Simple Network Time Protocol (SNTP) Version 3 based on RFC1769. SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. When enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean time (GMT) offset and the use of Daylight Saving Time.

configure node priority

```
configure node slot <slot_id> priority <node_pri>
```

Description

Configures the priority of the node.

Syntax Description

slot_id	Specifies the slot of the node. A is for the MSM/MM installed in slot A. B is for the MSM/MM installed in slot B.
node_pri	Specifies the priority of the node. The default 0 gives MSM-A a higher priority over MSM-B. The range is 1 to 100; 0 means you have not configured a node priority.

Default

Default node priority is 0.

Usage Guidelines

Use this command to configure the priority of the node. The lower the number, the higher the priority.

The node priority is part of the selection criteria for the primary node. The following list describes the parameters used to determine the primary node:

- Node state—The node state must be STANDBY to participate in leader election and to be selected primary. If the node is in the INIT, DOWN, or FAIL states, the node will not participate in leader election.
- Configuration priority—This is a user assigned priority. The configured priority is compared only after the node meets the minimum thresholds in each category for it to be healthy. Required processes and devices must not fail.
- Software health—This represents the percent of processes available.
- Health of secondary hardware components—This represents the health of switch components, such as the power supplies, fans, and so forth.
- Slot ID—The MSM/MM slot where the node is installed (MSM-A or MSM-B).

If you do not configure any priorities, MSM-A has a higher priority than MSM-B.

Example

The following command configures a priority of 2 for MSM-B:

```
configure node slot B priority 2
```

configure power supply

```
configure power supply <ps_num> {auto | on}
```

Description

Configures a power supply for either automatic power management, or forced on, regardless of the impact to the total available system power.

Syntax Description

ps_num	Specifies the slot number of the installed power supply unit (PSU) to which this command applies.
auto	Specifies that the NETGEAR 8800 determine the enabled or disabled state of the PSU to maximize total system power. This is the default.
on	Specifies that the PSU be enabled even if the NETGEAR 8800 determines it should be disabled. This action may reduce the total available system power and may result in one or more I/O modules powering down.

Default

The default setting is auto; the NETGEAR 8800 either enables or disables the PSU in order to maximize total system power.

Usage Guidelines

If a switch has PSUs with a mix of both 220V AC and 110V AC inputs, the NETGEAR 8800 maximizes system power by automatically taking one of two possible actions:

- If all PSUs are enabled then all PSUs must be budgeted at 110V AC to prevent overload of PSUs with 110V AC inputs.

OR

- If the PSUs with 110V AC inputs are disabled, then the PSUs with 220V AC inputs can be budgeted with a higher output per PSU.

The NETGEAR 8800 computes the total available power using both methods and automatically uses the PSU configuration that provides the greatest amount of power to the switch. [Table 5](#) lists combinations where the NETGEAR 8800 maximizes system power by disabling the PSUs with 110V AC inputs.

Table 5. PSU Combinations Where 110V PSUs Are Disabled

Number of PSUs with 220V AC Inputs	Number of PSUs with 110V AC Inputs
2	1
3	1
3	2
4	1
4	2
5	1

For all other combinations of 220V AC and 110V AC PSUs, the NETGEAR 8800 maximizes system power by enabling all PSUs and budgeting each PSU at 110V AC.

In addition to the PSU, you can specify the following options:

- `auto`—Specifies that the NETGEAR 8800 determine the enabled or disabled state of the PSU to maximize total system power. This is the default.
- `on`—Specifies that the PSU be enabled even if the NETGEAR 8800 determines it should be disabled. This action may reduce the total available system power and may result in one or more I/O modules powering down.

You can override automatic power supply management to enable a PSU with 110V AC inputs that the NETGEAR 8800 disables if the need arises, such as for a planned maintenance of 220V AC circuits. If the combination of AC inputs represents one of those listed in [Table 5](#), you can turn on a disabled PSU using the `configure power supply <ps_num> on` command.

Note: If you override automatic power supply management, you may reduce the available power and cause one or more I/O modules to power down.

To resume using automatic power supply management on a PSU, use the `configure power supply <ps_num> auto` command. The setting for each PSU is stored as part of the switch configuration.

To display power supply status and power budget information use the `show power` and `show power budget` commands.

Example

The following command configures the PSU in slot 1 to be forced on when either 110V AC or 220V AC power input is present, overriding automatic power management:

```
configure power supply 1 on
```

The switch displays the following message:

```
In a mixed environment of 110V and 220V AC inputs, power management may
automatically disable 110V supplies to maximize the system power budget.
By specifying 'on', you wish to override power management and enable the
specified power supply. This may cause the system power budget to decrease
and one or more I/O cards may be powered off as a result.
```

```
Are you sure you want to continue? (y/n)
```

Enter `y` to continue.

configure snmp access-profile

```
configure snmp access-profile [<profile_name> | none] {readonly | readwrite}
```

Description

Configures SNMP to use an ACL policy for access control.

Syntax Description

profile_name	Configures SNMP to use an ACL policy.
none	Cancels a previously configured ACL policy.
readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.

Default

SNMP access is enabled by default, with no ACL policies.

Usage Guidelines

You must be logged in as administrator to configure SNMP parameters.

You can restrict SNMP access by using an ACL and implementing an ACL policy. You create an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for SNMP. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

Use the `none` option to remove a previously configured ACL.

In the ACL policy file for SNMP, the `source-address` field is the only supported match condition. Any other match conditions are ignored.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see the chapters entitled “Policy Manager” and “ACLs” in the *NETGEAR 8800 User Manual*.

If you attempt to implement a policy that does not exist, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists. To confirm the existence of the policies, use the `ls` command. If the policy does not exist, create the ACL policy file.

Viewing SNMP Information

To display the current management configuration, including SNMP access related information, whether SNMP access is enabled or disabled, and whether any ACL policies are configured for SNMP, use the following command:

```
show management
```

Example

This example assumes that you already created an ACL to apply to SNMP.

The following command applies the ACL `MyAccessProfile_2` to SNMP:

```
configure snmp access-profile MyAccessProfile_2
```

configure snmp add community

```
configure snmp add community [readonly | readwrite] <alphanumeric_string>
```

Description

Adds an SNMP read or read/write community string.

Syntax Description

<code>readonly</code>	Specifies read-only access to the system.
<code>readwrite</code>	Specifies read and write access to the system.

alphanumeric_string	Specifies an SNMP community string name. See “Usage Guidelines” for more information.
---------------------	---

Default

The default read-only community string is *public*. The default read/write community string is *private*.

Usage Guidelines

Community strings provide a simple method of authentication between a switch and a remote network manager. Read community strings provide read-only access to the switch. The default read-only community string is *public*. Read-write community strings provide read and write access to the switch. The default read/write community string is *private*. Sixteen read-only and sixteen read/write community strings can be configured on the switch, including the defaults.

An authorized trap receiver must be configured to use the correct community strings on the switch for the trap receiver to receive switch-generated traps. In some cases, it may be useful to allow multiple community strings so that all switches and trap receivers are not forced to use identical community strings. The `configure snmp add community` command allows you to add multiple community strings in addition to the default community string.

An SNMP community string can contain up to 32 characters.

NETGEAR recommends that you change the defaults of the community strings. To change the value of the default read/write and read-only community strings, use the `configure snmp delete community` command.

Example

The following command adds a read/write community string with the value *netgear*:

```
configure snmp add community readwrite netgear
```

configure snmp add trapreceiver

```
configure snmp add trapreceiver [<ip_address> | <ipv6_address>] community [[hex
<hex_community_name>] | <community_name>] {port <port_number>} {from [<src_ip_address> |
<src_ipv6_address>]} {vr <vr_name>} {mode <trap_mode>}
```

Description

Adds the IP address of a trap receiver to the trap receiver list and specifies which SNMPv1/v2c traps are to be sent.

Syntax Description

ip_address	Specifies an SNMP trap receiver IPv4 address.
ipv6_address	Specifies an SNMP trap receiver IPv6 address

hex_community_name	Specifies that the trap receiver is to be supplied as a colon separated string of hex octets.
community_name	Specifies the community string of the trap receiver to be supplied in ASCII format.
port_number	Specifies a UDP port to which the trap should be sent. Default is 162.
src_ip_address	Specifies the IPv4 address of a VLAN to be used as the source address for the trap.
src_ipv6_address	Specifies the IPv6 address of a VLAN to be used as the source address for the trap.
vr_name	Specifies the name of the virtual router.
trap_mode	Specifies the mode of the traps: <ul style="list-style-type: none"> • <i>enhanced</i>—Contains extra varbinds at the end. • <i>standard</i>—Does not contain extra varbinds.

Default

Trap receivers are in enhanced mode by default, and the version is SNMPv2c by default.

Usage Guidelines

The IP address can be unicast, multicast, or broadcast.

An authorized trap receiver can be one or more network management stations on your network. Authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. The switch sends SNMP traps to all trap receivers configured to receive the specific trap group.

To view the SNMP trap receivers configured on the switch, use the `show management` command. The `show management` command displays information about the switch including the destination and community of the SNMP trap receivers configured on the switch.

Example

The following command adds the IP address 10.101.0.100 as a trap receiver with community string *purple*:

```
configure snmp add trapreceiver 10.101.0.100 community purple
```

The following command adds the IP address 10.101.0.105 as a trap receiver with community string *green*, using port 3003:

```
configure snmp add trapreceiver 10.101.0.105 community green port 3003
```

The following command adds the IP address 10.101.0.105 as a trap receiver with community string *blue*, and IP address 10.101.0.25 as the source:

```
configure snmp add trapreceiver 10.101.0.105 community blue from 10.101.0.25
```

configure snmp delete community

```
configure snmp delete community [readonly | readwrite] [all | <alphanumeric_string>]
```

Description

Deletes an SNMP read or read/write community string.

Syntax Description

readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.
all	Specifies all of the SNMP community strings.
alphanumeric_string	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

The default read-only community string is *public*. The default read/write community string is *private*.

Usage Guidelines

You must have at least one community string for SNMP access. If you delete all of the community strings on your system, you will no longer have SNMP access, even if you have SNMP enabled.

The community strings allow a simple method of authentication between the switch and the remote network manager. There are two types of community strings on the switch. Read community strings provide read-only access to the switch. The default read-only community string is *public*. read/write community strings provide read and write access to the switch. The default read/write community string is *private*. Sixteen read-only and sixteen read-write community strings can be configured on the switch, including the defaults. The community string for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. SNMP community strings can contain up to 32 characters.

For increased security, NETGEAR recommends that you change the defaults of the read/write and read-only community strings.

Use the `configure snmp add` commands to configure an authorized SNMP management station.

Example

The following command deletes a read/write community string named *netgear*:

```
configure snmp delete community readwrite netgear
```

configure snmp delete trapreceiver

```
configure snmp delete trapreceiver [[<ip_address> | <ipv6_address>] {<port_number>} | all]
```

Description

Deletes a specified trap receiver or all authorized trap receivers.

Syntax Description

ip_address	Specifies an SNMP trap receiver IPv4 address.
ipv6_address	Specifies an SNMP trap receiver IPv6 address.
port_number	Specifies the port associated with the receiver.
all	Specifies all SNMP trap receiver IP addresses.

Default

The default port number is 162.

Usage Guidelines

Use this command to delete a trap receiver of the specified IPv4 or IPv6 address, or all authorized trap receivers.

This command deletes only the first SNMPv1/v2c trap receiver whose IP address and port number match the specified value.

Example

The following command deletes the trap receiver 10.101.0.100 from the trap receiver list:

```
configure snmp delete trapreceiver 10.101.0.100
```

The following command deletes entries in the trap receiver list for 10.101.0.100, port 9990:

```
configure snmp delete trapreceiver 10.101.0.100 9990
```

Any entries for this IP address with a different community string will not be affected.

configure snmp sysContact

```
configure snmp syscontact <sysContact>
```

Description

Configures the name of the system contact.

Syntax Description

sysContact	An alphanumeric string that specifies a system contact name.
------------	--

Default

N/A.

Usage Guidelines

The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the switch. A maximum of 255 characters is allowed.

To view the name of the system contact listed on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the name of the system contact.

Example

The following command defines FredJ as the system contact:

```
configure snmp syscontact fredj
```

The following output from the `show switch` command displays *FredJ* as the system contact:

```
SysName:      engineeringlab
SysLocation:   englab
SysContact:    FredJ
```

configure snmp sysLocation

```
configure snmp syslocation <sysLocation>
```

Description

Configures the location of the switch.

Syntax Description

sysLocation	An alphanumeric string that specifies the switch location.
-------------	--

Default

N/A.

Usage Guidelines

Use this command to indicate the location of the switch. A maximum of 255 characters is allowed.

To view the location of the switch on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the location of the switch.

Example

The following command configures a switch location name on the system:

```
configure snmp syslocation englab
```

The following output from the `show switch` command displays *englab* as the location of the switch:

```
SysName:          engineeringlab
SysLocation:      englab
SysContact:       FredJ
```

configure snmp sysName

```
configure snmp sysname <sysName>
```

Description

Configures the name of the switch.

Syntax Description

sysName	An alphanumeric string that specifies a device name.
---------	--

Default

The default `sysname` is the model name of the device (for example, `XCM8806`).

Usage Guidelines

You can use this command to change the name of the switch. A maximum of 32 characters is allowed. The `sysname` appears in the switch prompt.

To view the name of the system listed on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the name of the system.

Example

The following command names the switch:

```
configure snmp sysname engineeringlab
```

The following output from the `show switch` command displays *engineeringlab* as the name of the switch:

```
SysName:          engineeringlab
SysLocation:      englab
SysContact:       FredJ
```

configure snmpv3 add access

```
configure snmpv3 add access [[hex <hex_group_name>] | <group_name>] {sec-model [snmpv1 |
snmpv2c | usm]} {sec-level [noauth | authnopriv | priv]} {read-view [[hex
<hex_read_view_name>] | <read_view_name>]} {write-view [[hex <hex_write_view_name>]] |
<write_view_name>]} {notify-view [[hex <hex_notify_view_name>] | <notify_view_name>]}
{volatile}
```

Description

Creates (and modifies) a group and its access rights.

Syntax Description

hex_group_name	Specifies the group name to add or modify. The value is to be supplied as a colon separated string of hex octets.
group_name	Specifies the group name to add or modify. The value is to be supplied in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.
read-view	Specifies the read view name: <ul style="list-style-type: none"> • hex_read_view_name—Specifies a hex value supplied as a colon separated string of hex octets • read_view_name—Specifies an ASCII value
write-view	Specifies the write view name: <ul style="list-style-type: none"> • hex_write_view_name—Specifies a hex value supplied as a colon separated string of hex octets • write_view_name—Specifies an ASCII value
notify-view	Specifies the notify view name: <ul style="list-style-type: none"> • hex_notify_view_name—Specifies a hex value supplied as a colon separated string of hex octets • notify_view_name—Specifies an ASCII value
volatile	Specifies volatile storage.

Default

The default values are:

- `sec-model`—USM
- `sec-level`—noauth
- `read view name`—defaultUserView
- `write view name`—""
- `notify view name`—defaultNotifyView
- `non-volatile storage`

Usage Guidelines

Use this command to configure access rights for a group. All access groups are created with a unique default context, "", as that is the only supported context.

Use more than one character when creating unique community strings and access group names.

A number of default groups are already defined. These groups are: *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*.

- The default groups defined are *v1v2c_ro* for security name *v1v2c_ro*, *v1v2c_rw* for security name *v1v2c_rw*, *admin* for security name *admin*, and *initial* for security names *initial*, *initialmd5*, *initialsha*, *initialmd5Priv* and *initialshaPriv*.
- The default access defined are *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*, and *v1v2cNotifyGroup*.

Example

In the following command, access for the group *defaultROGroup* is created with all the default values: security model `usm`, security level `noauth`, read view *defaultUserView*, no write view, notify view *defaultNotifyView*, and storage nonvolatile.

```
configure snmpv3 add access defaultROGroup
```

In the following command, access for the group *defaultROGroup* is created with the values: security model `USM`, security level `authnopriv`, read view *defaultAdminView*, write view *defaultAdminView*, notify view *defaultAdminView*, and storage nonvolatile.

```
configure snmpv3 add access defaultROGroup sec-model usm sec-level authnopriv read-view
defaultAdminView write-view defaultAdminView notify-view defaultAdminView
```

configure snmpv3 add community

```
configure snmpv3 add community [[hex <hex_community_index>] | <community_index>] name [[hex
<hex_community_name>] | <community_name>] user [[hex <hex_user_name>] | <user_name>] {tag
[[hex <hex_transport_tag>] | <transport_tag>]} {volatile}
```

Description

Adds an SNMPv3 community entry.

Syntax Description

hex_community_index	Specifies the row index in the snmpCommunity table as a hex value supplied as a colon separated string of hex octets.
community_index	Specifies the row index in the snmpCommunity Table as an ASCII value.
hex_community_name	Specifies the community name as a hex value supplied as a colon separated string of hex octets
community_name	Specifies the community name as an ASCII value.
hex_user_name	Specifies the USM user name as a hex value supplied as a colon separated string of hex octets.
user_name	Specifies the USM user name as an ASCII value.
tag	Specifies the tag used to locate transport endpoints in SnmpTargetAddrTable. When this community entry is used to authenticate v1/v2c messages, this tag is used to verify the authenticity of the remote entity. <ul style="list-style-type: none"> • hex_transport_tag—Specifies a hex value supplied as a colon separated string of hex octets • transport_tag—Specifies an ASCII value
volatile	Specifies volatile storage.

Default

N/A.

Usage Guidelines

Use this command to create or modify an SMMPv3 community in the community MIB.

Example

The following command creates an entry with the community index *comm_index*, community name *comm_public*, and user (security) name *v1v2c_user*:

```
configure snmpv3 add community comm_index name comm_public user v1v2c_user
```

The following command creates an entry with the community index (hex) of *12:0E*, community name (hex) of *EA:12:CD:CF:AB:11:3C*, user (security) name *v1v2c_user*, using transport tag *34872* and *volatile* storage:

```
configure snmpv3 add community hex 12:0E name hex EA:12:CD:CF:AB:11:3C user v1v2c_user tag 34872 volatile
```

configure snmpv3 add filter

```
configure snmpv3 add filter [[hex <hex_profile_name>] | <profile_name>] subtree
<object_identifier> {/<subtree_mask>} type [included | excluded] {volatile}
```


Description

Adds a filter to a filter profile.

Syntax Description

hex_profile_name	Specifies the filter profile that the current filter is added to. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile that the current filter is added to in ASCII format.
object identifier	Specifies a MIB subtree.
subtree_mask	Specifies a hex octet string used to mask the subtree. For example, f7a indicates 1.1.1.1.0.1.1.1.1.0.1.0.
included	Specifies that the MIB subtree defined by <object identifier>/<mask> is to be included.
excluded	Specifies that the MIB subtree defined by <object identifier>/<mask> is to be excluded.
volatile	Specifies volatile storage.

Default

The default values are:

- mask value—empty string (all 1s)
- type—`included`
- storage—`non-volatile`

Usage Guidelines

Use this command to create a filter entry in the `snmpNotifyFilterTable`. Each filter includes or excludes a portion of the MIB. Multiple filter entries comprise a filter profile that can eventually be associated with a target address. Other commands are used to associate a filter profile with a parameter name, and the parameter name with a target address.

This command can be used multiple times to configure the exact filter profile desired.

Example

The following command adds a filter to the filter profile *prof1* that includes the MIB subtree *1.3.6.1.4.1/f0*:

```
configure snmpv3 add filter prof1 subtree 1.3.6.1.4.1/f0 type included
```

configure snmpv3 add filter-profile

```
configure snmpv3 add filter-profile [[hex <hex_profile_name>] | <profile_name>] param [[hex <hex_param_name>]] | <param_name>] {volatile}
```

Description

Associates a filter profile with a parameter name.

Syntax Description

hex_profile_name	Specifies the filter profile name. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile name in ASCII format.
hex_param_name	Specifies a parameter name to associate with the filter profile. The value to follow is to be supplied as a colon separated string of hex octets.
param_name	Specifies a parameter name to associate with the filter profile in ASCII format.
volatile	Specifies volatile storage.

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to add an entry to the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. The parameter name is associated with target addresses, and the filter profile is associated with a series of filters, so, in effect, you are associating a series of filters with a target address.

Example

The following command associates the filter profile *prof1* with the parameter name *P1*:

```
configure snmpv3 add filter-profile prof1 param P1
```

configure snmpv3 add group user

```
configure snmpv3 add group [[hex <hex_group_name>] | <group_name>] user [[hex  
<hex_user_name>] | <user_name>] {sec-model [snmpv1| snmpv2c | usm]} {volatile}
```

Description

Adds a user name (security name) to a group.

Syntax Description

hex_group_name	Specifies the group name to add or modify. The value is to be supplied as a colon separated string of hex octets.
group_name	Specifies the group name to add or modify in ASCII format.

hex_user_name	Specifies the user name to add or modify. The value to follow is to be supplies as a colon separated string of hex octets.
user_name	Specifies the user name to add or modify in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
volatile	Specifies volatile storage.

Default

The default values are:

- sec-model—USM
- non-volatile storage

Usage Guidelines

Use this command to associate a user name with a group.

As per the SNMPv3 RFC, a security name is model independent while a username is model dependent. For simplicity, both are assumed to be same here. User names and security names are handled the same. In other words, if a user is created with the user name *username*, the security name value is the same, *username*.

Every group is uniquely identified by a security name and security model. So the same security name can be associated to a group name but with different security models.

Example

The following command associates the user *userV1* to the group *defaultRoGroup* with SNMPv1 security:

```
configure snmpv3 add group defaultRoGroup user userV1 sec-model snmpv1
```

The following command associates the user *userv3* with security model *usm* and storage type *volatile* to the access group *defaultRoGroup*:

```
configure snmpv3 add group defaultRoGroup user userV3 volatile
```

configure snmpv3 add mib-view

```
configure snmpv3 add mib-view [[hex <hex_view_name>] | <view_name>] subtree
<object_identifier> {/ <subtree_mask>} {type [included | excluded]} {volatile}
```

Description

Adds (and modifies) a MIB view.

Syntax Description

hex_view_name	Specifies the MIB view name to add or modify. The value is to be supplies as a colon separated string of hex octets.
view_name	Specifies the MIB view name to add or modify in ASCII format.
object_identifier	Specifies a MIB subtree.
subtree_mask	Specifies a hex octet string used to mask the subtree. For example, f7a indicates 1.1.1.1.0.1.1.1.0.1.0.
included	Specifies that the MIB subtree defined by <subtree>/<mask> is to be included.
excluded	Specifies that the MIB subtree defined by <subtree>/<mask> is to be excluded.
volatile	Specifies volatile storage.

Default

The default `mask` value is an empty string (all 1s). The other default values are `included` and `non-volatile`.

Usage Guidelines

Use this command to create a MIB view into a subtree of the MIB. If the view already exists, this command modifies the view to additionally include or exclude the specified subtree.

In addition to the created MIB views, there are three default views. They are: *defaultUserView*, *defaultAdminView*, and *defaultNotifyView*.

Example

The following command creates the MIB view *allMIB* with the subtree *1.3* included as non-volatile:

```
configure snmpv3 add mib-view allMIB subtree 1.3
```

The following command creates the view *netgearMib* with the subtree *1.3.6.1.4.1.1916* included as non-volatile:

```
configure snmpv3 add mib-view netgearMib subtree 1.3.6.1.4.1.1916
```

The following command creates a view *vrrpTrapNewMaster* which excludes VRRP notification *.1* and the entry is volatile:

```
configure snmpv3 add mib-view vrrpTrapNewMaster 1.3.6.1.2.1.68.0.1/ff8 type excluded volatile
```

configure snmpv3 add notify

```
configure snmpv3 add notify [[hex <hex_notify_name>] | <notify_name>] tag [[hex <hex_tag>] | <tag>] {volatile}
```

Description

Adds an entry to the snmpNotifyTable.

Syntax Description

hex_notify_name	Specifies the notify name to add. The value is to be supplied as a colon separated string of hex octets.
notify_name	Specifies the notify name to add in ASCII format.
hex_tag	Specifies a string identifier for the notifications to be sent to the target. The value is supplied as a colon separated string of octets.
tag	Specifies a string identifier for the notifications to be sent to the target in ASCII format.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to add an entry to the snmpNotifyTable. When a notification is to be sent, this table is examined. For the target addresses that have been associated with the tags present in the table, notifications are sent based on the filters also associated with the target addresses.

Example

The following command sends notifications to addresses associated with the tag *type1*:

```
configure snmpv3 add notify N1 tag type1
```

configure snmpv3 add target-addr

```
configure snmpv3 add target-addr [[hex <hex_addr_name>] | <addr_name>] param [[hex
<hex_param_name>] | <param_name>] ipaddress [ [ <ip_address> | <ip_and_tmask> ] | [
<ipv6_address> | <ipv6_and_tmask> ]] {transport-port <port_number>} {from [<src_ip_address> |
<src_ipv6_address>]} {vr <vr_name>} {tag-list <tag_list>} {volatile}
```

Description

Adds and configures an SNMPv3 target address and associates filtering, security, and notifications with that address.

Syntax Description

hex_addr_name	Specifies a string identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
addr_name	Specifies a string identifier for the target address in ASCII format.
hex_param_name	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
param_name	Specifies the parameter name associated with the target in ASCII format.
ip_address	Specifies an SNMPv3 target IPv4 address.
ip_and_tmask	Specifies the IPv4 address and hexadecimal mask in form A.B.C.D/NN...
ipv6_address	Specifies an SNMPv3 target IPv6 address.
ipv6_and_tmask	Specifies an IPv6 address and hexadecimal mask in form A:B:C:D:E:F:G:H/NN...
port_number	Specifies a UDP port. Default is 162.
src_ip_address	Specifies the IPv4 address of a VLAN to be used as the source address for the trap.
src_ipv6_address	Specifies the IPv6 address of a VLAN to be used as the source address for the trap.
vr_name	Specifies the name of the virtual router.
tag-list	Specifies a list of comma separated string identifiers for the notifications to be sent to the target.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- transport-port—port 162
- non-volatile storage

If you do not specify `tag-list` the single tag *defaultNotify*, a pre-defined value in the `snmpNotifyTable`, is used.

Usage Guidelines

Use this command to create an entry in the SNMPv3 `snmpTargetAddressTable`. The `param` parameter associates the target address with an entry in the `snmpTargetParamsTable`, which specifies security and storage parameters for messages to the target address, and an entry in the `snmpNotifyFilterProfileTable`, which specifies filter profiles to use for notifications to the target address. The filter profiles are associated with the filters in the `snmpNotifyFilterTable`.

The list of tag-lists must match one or more of the tags in the snmpNotifyTable for the trap to be sent out.

Example

The following command specifies a target address of *10.203.0.22* with the name *A1*, and associates it with the security parameters and target address parameter *P1*:

```
configure snmpv3 add target-addr A1 param P1 ipaddress 10.203.0.22
```

The following command specifies a target address of *10.203.0.22* with the name *A1*, and associates it with the security parameters and target address parameter *P1*, and the notification tags *type1* and *type2*:

```
configure snmpv3 add target-addr A1 param P1 ipaddress 10.203.0.22 from 10.203.0.23 tag-list type1,type2
```

configure snmpv3 add target-params

```
configure snmpv3 add target-params [[hex <hex_param_name>] | <param_name>] user [[hex <hex_user_name>] | <user_name>] mp-model [snmpv1 | snmpv2c | snmpv3] sec-model [snmpv1 | snmpv2c | usm] {sec-level [noauth | authnopriv | priv]} {volatile}
```

Description

Adds and configures SNMPv3 target parameters.

Syntax Description

hex_param_name	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
param_name	Specifies the parameter name associated with the target in ASCII format.
hex_user_name	Specifies a user name. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies a user name in ASCII format.
mp-model	Specifies a message processing model; choose from SNMPv1, SNMPv2, or SNMPv3.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.

volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.
----------	--

Default

The default values are:

- sec-level—noauth
- non-volatile storage

Usage Guidelines

Use this command to create an entry in the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

To associate a target address with a parameter name, see the command `configure snmpv3 add target-addr`.

Example

The following command specifies a target parameters entry named *P1*, a user name of *guest*, message processing and security model of SNMPv2c, and a security level of no authentication:

```
configure snmpv3 add target-params P1 user guest mp-model snmpv2c sec-model snmpv2c sec-level noauth
```

configure snmpv3 add user

```
configure snmpv3 add user [[hex <hex_user_name>] | <user_name>] {authentication [md5 | sha]
[hex <hex_auth_password> | <auth_password>]} {privacy {des | 3des | aes {128 | 192 | 256}}}
[[hex <hex_priv_password>] | <priv_password>]} {volatile}
```

Description

Adds (and modifies) an SNMPv3 user.

Syntax Description

hex_user_name	Specifies the user name to add or modify. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to add or modify in ASCII format.
MD5	Specifies MD5 authentication.
SHA	Specifies SHA authentication.
authentication	Specifies the authentication password or hex string to use for generating the authentication key for this user.

privacy	Specifies the privacy password or hex string to use for generating the privacy key for this user.
des	Specifies the use of the 56-bit DES algorithm for encryption. This is the default.
3des	Specifies the use of the 168-bit 3DES algorithm for encryption.
aes	Specifies the use of the AES algorithm for encryption.
128	Specifies the use of the 128-bit AES algorithm for encryption.
192	Specifies the use of the 192-bit AES algorithm for encryption.
256	Specifies the use of the 256-bit AES algorithm for encryption.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- authentication—no authentication
- privacy—no privacy
- non-volatile storage

Usage Guidelines

Use this command to create or modify an SNMPv3 user configuration.

The default user names are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*. The initial password for *admin* is *password*. For the other default users, the initial password is the user name.

If hex is specified, supply a 16 octet hex string for MD5, or a 20 octet hex string for SHA.

You must specify authentication if you want to specify privacy. There is no support for privacy without authentication.

Note: 3DES, AES 192, and AES 256 bit encryptions are proprietary implementations and may not work with some SNMP managers.

Example

The following command configures the user *guest* on the local SNMP Engine with security level `noauth` (no authentication and no privacy):

```
configure snmpv3 add user guest
```

The following command configures the user *authMD5* to use `MD5` authentication with the password *palertyu*:

```
configure snmpv3 add user authMD5 authentication md5 palertyu
```

The following command configures the user *authShapriv* to use *SHA* authentication with the hex key shown below, the privacy password *palertyu*, and *volatile* storage:

```
configure snmpv3 add user authShapriv authentication sha hex
01:03:04:05:01:05:02:ff:ef:cd:12:99:34:23:ed:ad:ff:ea:cb:11 privacy palertyu volatile
```

configure snmpv3 add user clone-from

```
configure snmpv3 add user [[hex <hex_user_name>] | <user_name>] clone-from [[hex
<hex_user_name>] | <user_name>]
```

Description

Creates a new user by cloning from an existing SNMPv3 user.

Syntax Description

hex_user_name	Specifies the user name to add or to clone from. The value is to be supplies as a colon separated string of hex octets.
user_name	Specifies the user name to add or to clone from in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to create a new user by cloning an existing one. After you have successfully cloned the new user, you can modify its parameters using the following command:

```
configure snmpv3 add user [[hex <hex_user_name>] | <user_name>] {authentication [md5 |
sha] [hex <hex_auth_password> | <auth_password>]} {privacy {des | 3des | aes {128 | 192
| 256}} [[hex <hex_priv_password>] | <priv_password>]} {volatile}
```

Users cloned from the default users will have the storage type of non-volatile. The default names are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*.

Example

The following command creates a user *cloneMD5* with same properties as the default user *initialmd5*. All authorization and privacy keys will initially be the same as with the default user *initialmd5*.

```
configure snmpv3 add user cloneMD5 clone-from initialmd5
```

configure snmpv3 delete access

```
configure snmpv3 delete access [all-non-defaults | {[hex <hex_group_name>] | <group_name>]
{sec-model [snmpv1 | snmpv2c | usm] sec-level [noauth | authnopriv | priv]]}]
```

Description

Deletes access rights for a group.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) security groups are to be deleted.
hex_group_name	Specifies the group name to be deleted. The value is to be supplies as a colon separated string of hex octets.
group_name	Specifies the group name to be deleted in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.

Default

The default values are:

- sec-model—USM
- sec-level—noauth

Usage Guidelines

Use this command to remove access rights for a group. Use the `all-non-defaults` keyword to delete all the security groups, except for the default groups. The default groups are: *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*.

Deleting an access will not implicitly remove the related group to user association from the VACMSecurityToGroupTable. To remove the association, use the following command:

```
configure snmpv3 delete group {[[hex <hex_group_name>] | <group_name>]} user
[all-non-defaults | {[hex <hex_user_name>] | <user_name>}] {sec-model
[snmpv1|snmpv2c|usm]}]}
```

Example

The following command deletes all entries with the group name *userGroup*:

```
configure snmpv3 delete access userGroup
```

The following command deletes the group *userGroup* with the security model `snmpv1` and security level of authentication and no privacy (`authnopriv`):

```
configure snmpv3 delete access userGroup sec-model snmpv1 sec-level authnopriv
```

configure snmpv3 delete community

```
configure snmpv3 delete community [all-non-defaults | {[hex <hex_community_index>] | <community_index>} | {name [[hex <hex_community_name>] | <community_name>}]
```

Description

Deletes an SNMPv3 community entry.

Syntax Description

all-non-defaults	Specifies that all non-default community entries are to be removed.
hex_community_index	Specifies the row index in the <code>snmpCommunityTable</code> . The value is to be supplied as a colon separated string of hex octets.
community_index	Specifies the row index in the <code>snmpCommunityTable</code> in ASCII format.
hex_community_name	Specifies the community name. The value is to be supplied as a colon separated string of hex octets.
community_name	Specifies the community name in ASCII format.

Default

The default entries are *public* and *private*.

Usage Guidelines

Use this command to delete an SMMPv3 community in the community MIB.

Example

The following command deletes an entry with the community index *comm_index*:

```
configure snmpv3 delete community comm_index
```

The following command creates an entry with the community name (hex) of *EA:12:CD:CF:AB:11:3C*:

```
configure snmpv3 delete community name hex EA:12:CD:CF:AB:11:3C
```

configure snmpv3 delete filter

```
configure snmpv3 delete filter [all | [[hex <hex_profile_name>] | <profile_name>] {subtree <object_identifier>}]
```

Description

Deletes a filter from a filter profile.

Syntax Description

all	Specifies all filters.
hex_profile_name	Specifies the filter profile of the filter to delete. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile of the filter to delete in ASCII format.
object_identifier	Specifies the MIB subtree of the filter to delete.

Default

N/A.

Usage Guidelines

Use this command to delete a filter entry from the snmpNotifyFilterTable. Specify **all** to remove all entries. Specify a profile name to delete all entries for that profile name. Specify a profile name and a subtree to delete just those entries for that filter profile and subtree.

Example

The following command deletes the filters from the filter profile *prof1* that reference the MIB subtree *1.3.6.1.4.1*:

```
configure snmpv3 delete filter prof1 subtree 1.3.6.1.4.1
```

configure snmpv3 delete filter-profile

```
configure snmpv3 delete filter-profile [all | [[hex <hex_profile_name>] | <profile_name>]
{param [[hex <hex_param_name>] | <param_name>}]
```

Description

Removes the association of a filter profile with a parameter name.

Syntax Description

all	Specifies all filter profiles.
hex_profile_name	Specifies the filter profile name to delete. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile name to delete in ASCII format.
hex_param_name	Specifies to delete the filter profile with the specified profile name and parameter name. The value is to be supplied as a colon separated string of hex octets.

param_name	Specifies to delete the filter profile with the specified profile name and parameter name in ASCII format.
------------	--

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to delete entries from the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. Specify `all` to remove all entries. Specify a profile name to delete all entries for that profile name. Specify a profile name and a parameter name to delete just those entries for that filter profile and parameter name.

Example

The following command deletes the filter profile *prof1* with the parameter name *P1*:

```
configure snmpv3 delete filter-profile prof1 param P1
```

configure snmpv3 delete group user

```
configure snmpv3 delete group {[[hex <hex_group_name>] | <group_name>]} user
[all-non-defaults | {[[hex <hex_user_name>] | <user_name>] {sec-model [snmpv1|snmpv2c|usm]}]}
```

Description

Deletes a user name (security name) from a group.

Syntax Description

hex_group_name	Specifies the group name to delete or modify. The value is to be supplied as a colon separated string of hex octets.
group_name	Specifies the group name to delete or modify in ASCII format.
all-non-defaults	Specifies that all non-default (non-permanent) users are to be deleted from the group.
hex_user_name	Specifies the user name to delete or modify. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to delete or modify in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).

Default

The default value for sec-model is USM.

Usage Guidelines

Use this command to remove the associate of a user name with a group.

As per the SNMPv3 RFC, a security name is model independent while a username is model dependent. For simplicity, both are assumed to be same here. User names and security names are handled the same. In other words, if a user is created with the user name *username*, the security name value is the same, *username*.

Every group is uniquely identified by a security name and security model. So the same security name can be associated to a group name but with different security models.

The default groups are: *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*.

The default users are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*.

Example

The following command deletes the user *guest* from the group *UserGroup* for the security model *snmpv2c*:

```
configure snmpv3 delete group UserGroup user guest sec-model snmpv2c
```

The following command deletes the user *guest* from the group *userGroup* with the security model *USM*:

```
configure snmpv3 delete group userGroup user guest
```

configure snmpv3 delete mib-view

```
configure snmpv3 delete mib-view [all-non-defaults | {[hex <hex_view_name>] | <view_name>}]
{subtree <object_identifier>}]
```

Description

Deletes a MIB view.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) MIB views are to be deleted.
hex_view_name	Specifies the MIB view to delete. The value is to be supplied as a colon separated string of hex octets.
view_name	Specifies the MIB view name to delete in ASCII format.
object_identifier	Specifies a MIB subtree.

Default

N/A.

Usage Guidelines

Use this command to delete a MIB view. Views which are being used by security groups cannot be deleted. Use the `all-non-defaults` keyword to delete all the MIB views (not being used by security groups) except for the default views. The default views are: *defaultUserView*, *defaultAdminView*, and *defaultNotifyView*.

Use the `configure snmpv3 add mib-view` command to remove a MIB view from its security group, by specifying a different view.

Example

The following command deletes all views (only the permanent views will not be deleted):

```
configure snmpv3 delete mib-view all-non-defaults
```

The following command deletes all subtrees with the view name *AdminView*:

```
configure snmpv3 delete mib-view AdminView
```

The following command deletes the view *AdminView* with subtree 1.3.6.1.2.1.2

```
configure snmpv3 delete mib-view AdminView subtree 1.3.6.1.2.1.2
```

configure snmpv3 delete notify

```
configure snmpv3 delete notify [{[hex <hex_notify_name>] | <notify_name>}] |  
all-non-defaults]
```

Description

Deletes an entry from the snmpNotifyTable.

Syntax Description

hex_notify_name	Specifies the notify name to add. The value is to be supplied as a colon separated string of hex octets.
notify_name	Specifies the notify name to add in ASCII format.
all-non-defaults	Specifies that all non-default (non-permanent) notifications are to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the snmpNotifyTable. When a notification is to be sent, this table is examined. For the target addresses that have been associated with the tags present in the table, notifications will be sent, based on the filters also associated with the target addresses.

Example

The following command removes the *N1* entry from the table:

```
configure snmpv3 delete notify N1
```

configure snmpv3 delete target-addr

```
configure snmpv3 delete target-addr [[[hex <hex_addr_name>] | <addr_name>]] | all]
```

Description

Deletes SNMPv3 target addresses.

Syntax Description

hex_addr_name	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
addr_name	Specifies a string identifier for the target address.
all	Specifies all target addresses.

Default

N/A.

Usage Guidelines

Use this command to delete an entry in the SNMPv3 snmpTargetAddressTable.

Example

The following command deletes target address named *A1*:

```
configure snmpv3 delete target-addr A1
```

configure snmpv3 delete target-params

```
configure snmpv3 delete target-params [[[hex <hex_param_name>] | <param_name>]] | all]
```

Description

Deletes SNMPv3 target parameters.

Syntax Description

hex_param_name	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
param_name	Specifies the parameter name associated with the target in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to delete an entry in the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

Example

The following command deletes a target parameters entry named *P1*:

```
configure snmpv3 delete target-params P1
```

configure snmpv3 delete user

```
configure snmpv3 delete user [all-non-defaults | [[hex <hex_user_name>] | <user_name>]]
```

Description

Deletes an existing SNMPv3 user.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) users are to be deleted.
hex_user_name	Specifies the user name to delete. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to delete.

Default

N/A.

Usage Guidelines

Use this command to delete an existing user.

Use the `all-non-defaults` keyword to delete all users, except for the default users. The default user names are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*.

Deleting a user will not implicitly remove the related group to user association from the VACMSecurityToGroupTable. To remove the association, use the following command:

```
configure snmpv3 delete group {[[hex <hex_group_name>] | <group_name>]} user
[all-non-defaults | {[hex <hex_user_name>] | <user_name>}] {sec-model
[snmpv1|snmpv2c|usm]]}]
```

Example

The following command deletes all non-default users:

```
configure snmpv3 delete user all-non-defaults
```

The following command deletes the user *guest*:

```
configure snmpv3 delete user guest
```

configure snmpv3 engine-boots

```
configure snmpv3 engine-boots <(1-2147483647)>
```

Description

Configures the SNMPv3 Engine Boots value.

Syntax Description

(1-2147483647)	Specifies the value of engine boots.
----------------	--------------------------------------

Default

N/A.

Usage Guidelines

Use this command if the Engine Boots value needs to be explicitly configured. Engine Boots and Engine Time will be reset to zero if the Engine ID is changed. Engine Boots can be set to any desired value but will latch on its maximum, 2147483647.

Example

The following command configures Engine Boots to 4096:

```
configure snmpv3 engine-boots 4096
```

configure snmpv3 engine-id

```
configure snmpv3 engine-id <hex_engine_id>
```

Description

Configures the SNMPv3 `snmpEngineID`.

Syntax Description

<code>hex_engine_id</code>	Specifies the colon delimited hex octet that serves as part of the <code>snmpEngineID</code> (5-32 octets).
----------------------------	---

Default

The default `snmpEngineID` is the device MAC address.

Usage Guidelines

Use this command if the `snmpEngineID` needs to be explicitly configured. The first four octets of the ID are fixed to 80:00:11:AE, which represents the NETGEAR Vendor ID. Once the `snmpEngineID` is changed, default users will be reverted back to their original passwords/keys, while non-default users will be reset to the security level of no authorization, no privacy.

In a chassis, the `snmpEngineID` will be generated using the MAC address of the MSM/MM with which the switch boots first. For MSM/MM hitless failover, the same `snmpEngineID` will be propagated to both of the MSMs/MMs.

Example

The following command configures the `snmpEngineID` to be 80:00:11:AE:00:0a:1c:3e:11:

```
configure snmpv3 engine-id 00:0a:1c:3e:11
```

configure sntp-client

```
configure sntp-client [primary | secondary] <host-name-or-ip> {vr <vr_name>}
```

Description

Configures an NTP server for the switch to obtain time information.

Syntax Description

<code>primary</code>	Specifies a primary server name.
<code>secondary</code>	Specifies a secondary server name.
<code>host-name-or-ip</code>	Specifies a host name or IPv4 address or IPv6 address.
<code>vr</code>	Specifies use of a virtual router.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the *NETGEAR 8800 User Manual*.

vr_name	Specifies the name of a virtual router.
---------	---

Default

N/A.

Usage Guidelines

Queries are first sent to the primary server. If the primary server does not respond within 1 second, or if it is not synchronized, the switch queries the second server. If the switch cannot obtain the time, it restarts the query process. Otherwise, the switch waits for the `sntp-client update interval` before querying again.

Example

The following command configures a primary NTP server:

```
configure sntp-client primary 10.1.2.2
```

The following command configures the primary NTP server to use the management virtual router *VR-Mgmt*:

```
configure sntp-client primary 10.1.2.2 vr VR-Mgmt
```

configure sntp-client update-interval

```
configure sntp-client update-interval <update-interval>
```

Description

Configures the interval between polls for time information from SNTP servers.

Syntax Description

update-interval	Specifies an interval in seconds.
-----------------	-----------------------------------

Default

64 seconds.

Usage Guidelines

None.

Example

The following command configures the interval timer:

```
configure sntp-client update-interval 30
```

configure telnet access-profile

```
configure telnet access-profile [<access_profile> | none]
```

Description

Configures Telnet to use an ACL policy for access control.

Syntax Description

<code>access_profile</code>	Specifies an ACL policy.
<code>none</code>	Cancels a previously configured ACL policy.

Default

Telnet is enabled with no ACL policies and uses TCP port 23.

Usage Guidelines

You must be logged in as administrator to configure Telnet parameters.

You can restrict Telnet access by using an ACL and implementing an ACL policy. You create an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for the Telnet port. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

Use the `none` option to remove a previously configured ACL.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see the chapters entitled “Policy Manager” and “ACLs” in the *NETGEAR 8800 User Manual*.

In the ACL policy file for telnet, the “source-address” field is the only supported match condition. Any other match conditions are ignored.

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `ls` command. If the policy does not exist, create the ACL policy file.

Viewing Telnet Information

To display the status of Telnet, including the current TCP port, the virtual router used to establish a Telnet session, and whether ACLs are controlling Telnet access, use the following command:

```
show management
```

Example

This example assumes that you already created an ACL to apply to Telnet.

The following command applies the ACL MyAccessProfile_2 to Telnet:

```
configure telnet access-profile MyAccessProfile_2
```

configure telnet port

```
configure telnet port [<portno> | default]
```

Description

Configures the TCP port used by Telnet for communication.

Syntax Description

portno	Specifies a TCP port number. The default is 23. The range is 1 through 65535. The following TCP port numbers are reserved and cannot be used for Telnet connections: 22, 80, and 1023.
default	Specifies the default Telnet TCP port number. The default is 23.

Default

The switch listens for Telnet connections on Port 23.

Usage Guidelines

You must be logged in as administrator to configure the Telnet port.

The `portno` range is 1 through 65535. The following TCP port numbers are reserved and cannot be used for Telnet connections: 22, 80, and 1023. If you attempt to configure a reserved port, the switch displays an error message similar to the following:

```
configure telnet port 22
Error: port number is a reserved port
```

If this occurs, select a port number that is not a reserved port.

The switch accepts IPv6 connections.

Example

The following command changes the port used for Telnet to port 85:

```
configure telnet port 85
```

The following command returns the port used for Telnet to the default port of 23:

```
configure telnet port default
```

configure telnet vr

```
configure telnet vr [all | default | <vr_name>]
```

Description

Configures the virtual router used on the switch for listening for Telnet connections.

Syntax Description

all	Specifies to use all virtual routers for Telnet connections.
default	Specifies to use the default virtual router for Telnet connections. The default router is VR-Mgmt.
vr_name	Specifies the name of the virtual router to use for Telnet connections.
<p>Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i>.</p>	

Default

The default is `all`.

Usage Guidelines

You must be logged in as administrator to configure the virtual router.

The switch accepts IPv6 connections.

If you specify `all`, the switch listens on all of the available virtual routers for Telnet connections.

The `vr_name` specifies the name of the virtual router to use for Telnet connections.

If you specify a virtual router name that does not exist, the switch displays an error message similar to the following:

```
configure telnet vr vr-ttt
                        ^
%% Invalid input detected at '^' marker.
```

Example

The following command configures the switch to listen for and receive Telnet requests on all virtual routers:


```
configure telnet vr all
```

create snmp trap

```
create snmp trap severity <severity> event <EventName> <msg>
```

Description

Creates and sends an SNMP trap containing the information defined in the command.

Syntax Description

severity	Specifies one of the eight severity levels defined in the NETGEAR 8800 software. Enter one of the following values: <code>critical</code> , <code>error</code> , <code>warning</code> , <code>notice</code> , <code>info</code> , <code>debug-summary</code> , <code>debug-verbose</code> , <code>debug-data</code> .
EventName	Specifies the event name. Enter a name using alphanumeric characters.
msg	Specifies a message. Enter the message using alphanumeric characters.

Default

N/A.

Usage Guidelines

None.

Example

The following example sends a trap of severity `info` for event `AAA` with the message `user XYZ logged in`:

```
create snmp trap severity info event AAA "user XYZ logged in"
```

disable dhcp vlan

```
disable dhcp vlan [<vlan_name> | all]
```

Description

Disables the generation and processing of DHCP packets on a VLAN to obtain an IP address for the VLAN from a DHCP server.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs

Default

Disabled for all VLANs.

Usage Guidelines

None.

Example

The following command disables the generation and processing of DHCP packets on a VLAN named *accounting*:

```
disable dhcp vlan accounting
```

disable snmp access

```
disable snmp access {snmp-v1v2c | snmpv3}
```

Description

Selectively disables SNMP on the switch.

Syntax Description

snmp-v1v2c	Specifies SNMPv1/v2c access only.
snmpv3	Specifies SNMPv3 access only.

Default

Enabled.

Usage Guidelines

Disabling SNMP access does not affect the SNMP configuration (for example, community strings). However, if you disable SNMP access, you will be unable to access the switch using SNMP.

This command allows you to disable either all SNMP access, v1/v2c access only, or v3 access only.

To allow access, use the following command:

```
enable snmp access {snmp-v1v2c | snmpv3}
```

Example

The following command disables all SNMP access on the switch:

```
disable snmp access
```

disable snmp access vr

```
disable snmp access vr [<vr_name> | all]
```

Description

Selectively disables SNMP access on virtual routers.

Syntax Description

vr_name	Specifies the virtual router name.
all	Specifies all virtual routers.

Default

Enabled on all virtual routers.

Usage Guidelines

Use this command to disable SNMP access on any or all virtual routers.

When SNMP access is disabled on a virtual router, the incoming SNMP request is dropped and an EMS message is logged.

To enable SNMP access on virtual routers use the `enable snmp access vr` command.

To display the SNMP configuration and statistics on a specified virtual router, use the `show snmp vr_name` command.

Example

The following command disables SNMP access on the virtual router *vr-finance*:

```
disable snmp access vr vr-finance
```

disable snmp community

```
disable snmp community <alphanumeric-community-string>
```

Description

Disables SNMP community strings on the switch.

Syntax Description

alphanumeric-community-string	Specifies the SNMP community string name.
-------------------------------	---

Default

N/A

Usage Guidelines

This command allows the administrator to disable an snmp community. It sets the `rowStatus` of the community to `NotInService`. When disabled, SNMP access to the switch using the designated community is not allowed.

Example

The following command disables the community string named *netgear*.

```
disable snmp community netgear
```

disable snmp traps

```
disable snmp traps
```

Description

Prevents SNMP traps from being sent from the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command does not clear the SNMP trap receivers that have been configured. The command prevents SNMP traps from being sent from the switch even if trap receivers are configured.

To view if SNMP traps are being sent from the switch, use the `show management` command. The `show management` command displays information about the switch including the enabled/disabled state of SNMP traps being sent.

Example

The following command prevents SNMP traps from being sent from the switch to the trap receivers:

```
disable snmp traps
```

disable snmpv3

```
disable snmpv3 [default-group | default-user]
```

Description

Selectively disables SNMPv3 default-group or default-user access on the switch.

Syntax Description

default-group	Specifies SNMPv3 default-group.
default-user	Specifies SNMPv3 default-user.

Default

Enabled

Usage Guidelines

This command is used to disable SNMPv3 default-group or default-user access.

Disabling SNMPv3 default-group access removes access to default-users and user-created users who are part of the default-group. The user-created authenticated SNMPv3 users (who are part of a user-created group) are able to access the switch. By disabling default-users access, the end-user is not able to access the switch/MIBs using SNMPv3 default-user.

The default groups are: *admin*, *initial*, *v1v2c_ro*, *v1v2c_rw*.

The default users are: *admin*, *initial*, *initialmd5*, *initialsha*, *initialmd5Priv*, *initialshaPriv*.

Example

The following command disables the default group on the switch:

```
disable snmp default-group
```

disable snmp-client

```
disable snmp-client
```

Description

Disables the SNTP client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. After the SNTP client has been enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean Time (GMT) offset and the use of Daylight Savings Time (DST).

Example

The following command disables the SNTP client:

```
disable sntp-client
```

disable telnet

```
disable telnet
```

Description

Disables external Telnet services on the system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You must be logged in as an administrator to enable or disable Telnet.

Note: Telnet sessions between MSMs/MMs are not affected by this command.

Example

With administrator privilege, the following command disables external Telnet services on the switch:

```
disable telnet
```

disable watchdog

```
disable watchdog
```

Description

Disables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer monitors the health of the switch hardware and software events. For example, the watchdog timer reboots the switch if the system cannot reset the watchdog timer. This can be caused by a long CPU processing loop, any unhandled exception, or a hardware problem with the communication channel to the watchdog. In most cases, if the watchdog timer expires, the switch captures the current CPU status and posts it to the console and the system log. In some cases, if the problem is so severe that the switch is unable to perform any action, the switch reboots without logging any system status information prior to reboot.

This command takes affect immediately.

The watchdog settings are saved in the configuration file.

To display the watchdog state of your system, use the `show switch` command.

Example

The following command disables the watchdog timer:

```
disable watchdog
```

enable dhcp vlan

```
enable dhcp vlan [<vlan_name> | all]
```

Description

Enables the generation and processing of DHCP packets on a VLAN to obtain an IP address for the VLAN from a DHCP server.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
<code>all</code>	Specifies all VLANs.

Default

Disabled for all VLANs.

Usage Guidelines

None.

Example

The following command enables the generation and processing of DHCP packets on a VLAN named *accounting*:

```
enable dhcp vlan accounting
```

enable snmp access

```
enable snmp access {snmp-v1v2c | snmpv3}
```

Description

Selectively enables SNMP access on the switch.

Syntax Description

snmp-v1v2c	Specifies SNMPv1/v2c access only.
snmpv3	Specifies SNMPv3 access only.

Default

Enabled.

Usage Guidelines

To have access to the SNMP agent residing in the switch, at least one VLAN must have an IP address assigned to it.

Any network manager running SNMP can manage the switch for v1/v2c/v3, provided the MIB is installed correctly on the management station. Each network manager provides its own user interface to the management facilities.

For SNMPv3, additional security keys are used to control access, so an SNMPv3 manager is required for this type of access.

This command allows you to enable either all SNMP access, no SNMP access, v1/v2c access only, or v3 access only.

To prevent any SNMP access, use the following command:

```
disable snmp access {snmp-v1v2c | snmpv3}
```

The 8800 OS introduced the concept of safe defaults mode. Safe defaults mode runs an interactive script that allows you to enable or disable SNMP, Telnet, and switch ports. When you set up your switch for the first time, you must connect to the console port to access the switch. After logging in to the switch, you enter safe defaults mode. Although SNMP, Telnet, and switch ports are enabled by default, the script prompts you to confirm those settings.

If you choose to keep the default setting for SNMP—the default setting is enabled—the switch returns the following interactive script:

Since you have chosen less secure management methods, please remember to increase the security of your network by taking the following actions:

- * change your admin password
- * change your SNMP public and private strings
- * consider using SNMPv3 to secure network management traffic

In addition, you can return to safe defaults mode by issuing the following command:

```
configure safe-default-script
```

If you return to safe defaults mode, you must answer the questions presented during the interactive script.

For more detailed information about safe defaults mode, see the section “Safe Defaults Setup Method” in the *NETGEAR 8800 User Manual*.

Example

The following command enables all SNMP access for the switch:

```
enable snmp access
```

enable snmp access vr

```
enable snmp access vr [<vr_name> | all]
```

Description

Selectively enables SNMP access on virtual routers.

Syntax Description

vr_name	Specifies the virtual router name.
all	Specifies all virtual routers.

Default

Enabled on all virtual routers.

Usage Guidelines

Use this command to enable SNMP access on any or all virtual routers.

To disable SNMP access on virtual routers, use the `disable snmp access vr` command.

To display the SNMP configuration and statistics on a specified virtual router, use the `show snmp vr_name` command.

Example

The following command enables SNMP access on the virtual router *vr-finance*:

```
enable snmp access vr vr-finance
```

enable snmp community

```
enable snmp community <alphanumeric-community-string>
```

Description

Enables SNMP community strings.

Syntax Description

alphanumeric-community-string	Specifies the SNMP community string name.
-------------------------------	---

Default

N/A

Usage Guidelines

This command allows the administrator to enable an snmp community that has been disabled. It sets the `rowStatus` of the community to `Active`.

Example

The following command enables the community string named *netgear*:

```
enable snmp community netgear
```

enable snmp traps

```
enable snmp traps
```

Description

Turns on SNMP trap support.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMP traps to all trap receivers.

To view if SNMP traps are being sent from the switch, use the `show management` command. The `show management` command displays information about the switch including the enabled/disabled state of SNMP traps being sent.

Example

The following command enables SNMP trap support on the switch:

```
enable snmp traps
```

enable snmpv3

```
enable snmpv3 [default-group | default-user]
```

Description

Selectively enables SNMPv3 default-group or default-user access on the switch.

Syntax Description

default-group	Specifies SNMPv3 default-group.
default-user	Specifies SNMPv3 default-user.

Default

Enabled

Usage Guidelines

This command is used to enable SNMPv3 default-group or default-user access.

Enabling SNMPv3 default-group access activates the access to an SNMPv3 default-group and the user- created SNMPv3-user part of default-group. Enabling the SNMPv3 default-user access allows an end user to access the MIBs using SNMPv3 default-user. This command throws an error if the SNMPv3 access is disabled on the switch.

The default groups are: *admin, initial, v1v2c_ro, v1v2c_rw*.

The default users are: *admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv*.

Example

The following command enables the default users on the switch:

```
enable snmp default-user
```

enable snmp-client

```
enable snmp-client
```

Description

Enables the SNTP client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. After the SNTP client has been enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean Time (GMT) offset and the use of Daylight Savings Time (DST).

Example

The following command enables the SNTP client:

```
enable sntp-client
```

enable telnet

```
enable telnet
```

Description

Enables external Telnet services on the system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You must be logged in as an administrator to enable or disable Telnet.

The 8800 OS introduces the concept of safe defaults mode. Safe defaults mode runs an interactive script that allows you to enable or disable SNMP, Telnet, and switch ports. When you set up your switch for the first time, you must connect to the console port to access the switch. After logging in to the switch, you enter safe defaults mode. Although SNMP, Telnet, and switch ports are enabled by default, the script prompts you to confirm those settings.

If you choose to keep the default setting for Telnet—the default setting is enabled—the switch returns the following interactive script:

Since you have chosen less secure management methods, please remember to increase the security of your network by taking the following actions:

- * change your admin password
- * change your SNMP public and private strings
- * consider using SNMPv3 to secure network management traffic

In addition, you can return to safe defaults mode by issuing the following command:

```
configure safe-default-script
```

If you return to safe defaults mode, you must answer the questions presented during the interactive script.

For more detailed information about safe defaults mode, see the section “Safe Defaults Setup Method” in the *NETGEAR 8800 User Manual*.

Example

With administrator privilege, the following command enables Telnet services on the switch:

```
enable telnet
```

enable watchdog

```
enable watchdog
```

Description

Enables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer monitors the health of the switch hardware and software events. For example, the watchdog timer reboots the switch if the system cannot reset the watchdog timer. This is caused by a long CPU processing loop, any unhandled exception, or a hardware problem with the communication channel to the watchdog. In most cases, if the watchdog timer expires, the switch captures the current CPU status and posts it to the console and the system log. In some cases, if the problem is so severe that the switch is unable to perform any action, the switch reboots without logging any system status information prior to reboot.

This command takes affect immediately.

The watchdog settings are saved in the configuration file.

To display the watchdog state of your system, use the `show switch` command.

Example

The following command enables the watchdog timer:

```
enable watchdog
```

exit

```
exit
```

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
exit
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

logout

```
logout
```

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
logout
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

quit

```
quit
```

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter `y` if you want to save your changes. Enter `n` if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
quit
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter **y** if you want to save your changes. Enter **n** if you do not want to save your changes.

show checkpoint-data

```
show checkpoint-data {<process>}
```

Description

Displays the status of one or more processes being copied from the primary MSM/MM to the backup MSM/MM.

Syntax Description

process	Specifies the name of the processes being copied.
---------	---

Default

N/A.

Usage Guidelines

This command displays, in percentages, the amount of internal state copying completed by each process and the traffic statistics between the process on both the primary and the backup MSMs/MMs.

This command is also helpful in debugging synchronization problems that occur at run-time. To check the status of synchronizing the MSMs/MMs, use the `show switch` command.

Depending on the software version running on your switch and the type of switch you have, additional or different checkpoint status information may be displayed.

Example

The following command displays the checkpointing status and the traffic statics of all of the processes between the primary and the backup MSM:

```
show checkpoint-data
```

The following is sample output from this command:

Process	Tx	Rx	Errors	Sent	Total	% Chkpt	Debug-info
devmgr	3812	1731	0	3	3	100% ON OK	1 (00008853)

dirser	0	0	0	0	0	0%	ON	OK	1 (000008D3)
ems	5	0	0	0	0	100%	ON	OK	1 (000008D3)
nodemgr	0	0	0	0	0	0%	ON	OK	1 (000008D3)
snmpSubagent	0	0	0	0	0	0%	ON	OK	1 (000018D3)
snmpMaster	0	0	0	0	0	0%	ON	OK	1 (000008D3)
cli	0	0	0	0	0	0%	ON	OK	1 (000018D3)
cfgmgr	82	82	0	1	1	100%	ON	OK	1 (000018D3)
elrp	0	0	0	0	0	0%	ON	OK	1 (000008D3)
vlan	1047	1	0	0	0	100%	ON	OK	1 (000008D3)
aaa	0	0	0	0	0	0%	ON	OK	1 (000008D3)
fdb	957	2	0	0	0	100%	ON	OK	1 (000008D3)
msgsrv	0	0	0	0	0	100%	ON	OK	1 (000008D3)
stp	1	0	0	0	0	0%	ON	OK	1 (000008D3)
polMgr	0	0	0	0	0	0%	ON	OK	1 (000008D3)
mcmgr	2	2	0	0	0	100%	ON	OK	1 (000008D3)
acl	0	0	0	0	0	100%	ON	OK	1 (000008D3)
netLogin	0	0	0	0	0	0%	ON	OK	1 (000008D3)
ospf	0	0	0	0	0	0%	ON	OK	1 (000008D3)
netTools	1	0	0	0	0	100%	ON	OK	1 (000008D3)
telnetd	0	0	0	0	0	0%	ON	OK	1 (000008D3)
rtmgr	4	4	0	0	0	100%	ON	OK	1 (000008D3)
vrrp	378	0	0	0	0	0%	ON	OK	1 (000008D3)
tftpd	0	0	0	0	0	0%	ON	OK	1 (000008D3)
thttpd	0	0	0	0	0	0%	ON	OK	1 (000008D3)
rip	0	0	0	0	0	0%	ON	OK	1 (000008D3)
dosprotect	0	0	0	0	0	0%	ON	OK	1 (000008D3)
epm	0	0	0	0	0	0%	ON	OK	1 (000008D3)
hal	0	0	0	0	0	0%	ON	OK	1 (000008D3)
bgp	0	0	0	0	0	0%	ON	OK	1 (000008D3)
pim	0	0	0	0	0	0%	ON	OK	1 (000008D3)
etmon	185	185	0	0	0	100%	ON	OK	1 (000008D3)

To view the output for a specific process, use the `process` option. The following command displays detailed information for the STP process:

```
show checkpoint-data stp
```

The following is sample output from this command:

Process	Tx	Rx	Errors	Sent	Total	%	Chkpt	Debug-info
stp	1	0	0	0	0	0%	ON OK	1 (000008D3)

show dhcp-client state

```
show dhcp-client state
```

Description

Displays the current DHCP/BOOTP client state for each vlan.

Syntax Description

This command has no arguments or variables.

Default

Displays the client state for all existing VLANs.

Usage Guidelines

None.

Example

The following command displays the DHCP/BOOTP status for all VLANs:

```
show dhcp-client state
```

Depending on your configurations, output from this command is similar to the following:

Client VLAN	Protocol	Server	Current State
Default	BOOTP	10.1.1.2.3	Received IP address configured on vlan
accounting	DHCP	10.2.3.4	DHCP state; Requesting
Mgmt	None	0.0.0.0	

A total of 3 vlan(s) were displayed

show management

```
show management
```

Description

Displays the SNMP and CLI settings configured on the switch and the SNMP statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines:

The following management output is displayed:

- Enable/disable state for Telnet, and SNMP access
- Login statistics
 - Enable/disable state for idle timeouts
 - Maximum number of CLI sessions

- SNMP community strings
- SNMP trap receiver list
- SNMP trap receiver source IP address
- SNMP statistics counter
- SSH access states of enabled, disabled, and module not loaded
- CLI configuration logging
- SNMP access states of v1, v2c disabled and v3 enabled

If all three types of SNMP access are enabled or disabled, SNMP access is displayed as either Enabled or Disabled.

- Enable/disable state for RMON
- Access-profile usage configured via Access Control Lists (ACLs) for additional Telnet and SSH2 security
- CLI scripting settings
 - Enable/disable state
 - Error message setting
 - Persistence mode
- Dropped SNMP packet counter.

Example

The following command displays configured SNMP settings on an 8800 switch:

```
show management
```

The following is sample output from this command:

```
CLI idle timeout           : Enabled (20 minutes)
CLI max number of login attempts : 3
CLI max number of sessions  : 8
CLI paging                 : Enabled (this session only)
CLI space-completion       : Disabled (this session only)
CLI configuration logging   : Disabled
CLI scripting              : Disabled (this session only)
CLI scripting error mode    : Ignore-Error (this session only)
CLI persistent mode        : Persistent (this session only)
Telnet access              : Enabled (tcp port 23 vr all)
                           : Access Profile : not set
SSH Access                 : ssh module not loaded.
Web access                 : Disabled (tcp port 80)
Total Read Only Communities : 1
Total Read Write Communities : 1
RMON                       : Disabled
SNMP access                : Enabled
                           : Access Profile Name : not set
SNMP Traps                 : Enabled
```

```
SNMP v1/v2c TrapReceivers      :
  Destination      Source IP Address  Flags
  10.120.91.89 /10550              2E

Flags:  Version: 1=v1 2=v2c
        Mode: S=Standard E=Enhanced

SNMP stats:      InPkts 582      OutPkts  588      Errors 0      AuthErrors 0
                  Gets   0        GetNexts 582      Sets   0        Drops   12294
SNMP traps:      Sent    6        AuthTraps Enabled
```

show node

```
show node {detail}
```

Description

Displays the status of the nodes in the system as well as the general health of the system.

Syntax Description

detail	Displays the information on a per-node basis rather than in a tabular format.
--------	---

Default

N/A.

Usage Guidelines

Use this command to display the current status of the nodes and the health of the system. The information displayed shows the node configurations (such as node priority) and the system and hardware health computations. You can use this information to determine which node will be elected primary in case of a failover.

Table 6 lists the node statistic information collected by the switch.

Table 6. Node States

Node State	Description
BACKUP	In the backup state, this node becomes the primary node if the primary fails or enters the DOWN state. The backup node also receives the checkpoint state data from the primary.
DOWN	In the down state, the node is not available to participate in leader election. The node enters this state during any user action, other than a failure, that makes the node unavailable for management. Examples of user actions are: <ul style="list-style-type: none"> • Upgrading the software • Rebooting the system using the <code>reboot</code> command • Initiating an MSM/MM failover using the <code>run msm-failover</code> command • Synchronizing the MSM's/MM's software and configuration in non-volatile storage using the <code>synchronize</code> command
FAIL	In the fail state, the node has failed and needs to be restarted or repaired. The node reaches this state if the system has a hardware or software failure.
INIT	In the initial state, the node is being initialized. A node stays in this state when it is coming up and remains in this state until it has been fully initialized. Being fully initialized means that all of the hardware has been initialized correctly and there are no diagnostic faults.
MASTER	In the primary state, the node is responsible for all switch management functions.
STANDBY	In the standby state, leader election occurs—the primary and backup nodes are elected. The priority of the node is only significant in the standby state.

Example

The following command displays the status of the node, the priority of the node, and the general health of the system:

```
show node
```

The following is sample output from this command:

```
Node   State      Priority  SwHealth  HwHealth
-----
MSM-A  MASTER      0        49        7
MSM-B  BACKUP      0        49        7
```

If you specify the `detail` option, the same information is displayed on a per node basis rather than in a tabular format.

```
Node MSM-A information:
```

```
Node State:    MASTER
Node Priority:  0
Sw Health:     49
Hw Health:     7
```

```
Node MSM-B information:
```

```
Node State:    BACKUP
Node Priority:  0
```

```
Sw Health:      49
Hw Health:      7
```

show odometers

```
show odometers
```

Description

Displays a counter for each component of a switch that shows how long it has been functioning since it was manufactured.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output from this command displays how long individual components in the switch have been functioning since it was manufactured. This odometer counter is kept in the EEPROM of each monitored component. This means that even if you plug in the component into a different chassis, the odometer counter is available in the new switch chassis.

Monitored Components

On the 8800, the odometer monitors the following components:

- Chassis
- MSMs/MMs
- I/O modules
- Power controllers

Recorded Statistics

The following odometer statistics are collected by the switch:

- Service Days—The amount of days that the component has been running
- First Recorded Start Date—The date that the component was powered-up and began running

Depending on the software version running on your switch, the modules installed in your switch, and the type of switch you have, additional or different odometer information may be displayed.

Example

The following command displays how long each component of a switch has been functioning since its manufacture date:

```
show odometers
```

The following is sample output from the NETGEAR 8800 series switch:

Field Replaceable Units	Service Days	First Recorded Start Date
-----	-----	-----
Chassis : BD-8810	209	Dec-07-2004
Slot-1 : G48T	208	Dec-07-2004
Slot-2 : 10G4X	219	Nov-02-2004
Slot-3 : G48T	228	Oct-26-2004
Slot-4 : G24X	226	Oct-19-2004
Slot-5 : G8X	139	Dec-07-2004
Slot-6 :		
Slot-7 : 10G4X	160	Dec-16-2004
Slot-8 : 10G4X	133	Dec-14-2004
Slot-9 : G48P	111	Nov-04-2004
Slot-10 :		
MSM-A : MSM-G8X	137	Dec-07-2004
MSM-B :		
PSUCTRL-1 :	209	Dec-07-2004
PSUCTRL-2 :	208	Dec-07-2004

show power

```
show power {<ps_num>} {detail}
```

Description

Displays the current status of the installed power supplies.

Command Syntax

ps_num	Specifies the slot number of the installed power supply.
detail	The detail option is reserved for future use.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the power supplies.

This status information may be useful for your technical support representative if you have a network problem.

The switch collects the following power supply information:

- **State**—Indicates the current state of the power supply. Options are:
 - **Empty**—There is no power supply installed.
 - **Power Failed**—The power supply has failed.
 - **Powered Off**—The power supply is off.
 - **Powered On**—The power supply is on and working normally.

Located next to the “State” of the power supply, the following information provides more detailed status information. Options are:

- **Disabled for net power gain**—Indicates that the power supply is disabled in order to maximize the total available system power
- **Configured ON**—Indicates that the user requested to enable a disabled power supply regardless of the affect on the total available system power
- **Configured ON when present**—Indicates that the power supply slot is currently empty, but the user requested to enable the power supply regardless of the affect on the total available system power
- **Unsupported**—Indicates that a 600/900 W AC PSU is inserted in a chassis other than the XCM8806 and XCM8810.
- **PartInfo**—Provides information about the power supply. Depending on your switch, options include:
 - **Serial number**—A collection of numbers and letters, that make up the serial number of the power supply.
 - **Part number**—A collection of numbers and letters that make up the part number of the power supply.
- **Revision**—Displays the revision number of the power supply.
- **Odometer**—Specifies how long the power supply has been operating.
- **Temperature**—Specifies, in Celsius, the current temperature of the power supply.
- **Input**—Specifies the input voltage and the current requirements of the power supply and whether the input is AC or DC.
- **Output 1 and Output 2**—Specifies the output voltage and the current supplied by the power supply. The values are only displayed if known for the platform.

Example

The following command displays the status of the power supply installed in slot 1:

```
show power 1
```

The following is sample output from this command:

```
PowerSupply 1 information:
State:           Powered On
```



```

PartInfo:      PS 2336 5003J-00479 4300-00137
Revision:      2.0
Odometer:      90 days 5 hours
Temperature:   29.0 deg C
Fan 1:         6473 RPM
Fan 2:         6233 RPM
Input:         230.00 V AC
Output 1:      48.50 V,  7.25 A   (48V/1104W Max)
Output 2:      12.44 V,  0.62 A   (12V/48W Max)

```

If power management needs to disable a power supply to maximize the total available power, you see `Disabled for net power gain` next to the state of the power supply, as shown in the sample truncated output:

```

PowerSupply 1 information:
State:          Powered Off (Disabled for net power gain)
PartInfo:       PS 2336 0413J-00732 4300-00137
...

```

If you choose to always enable a power supply, regardless of the affect on the total available power, you see `Configured ON` next to the state of the power supply, as shown in the sample truncated output:

```

PowerSupply 1 information:
State:          Powered On  (Configured ON)
PartInfo:       PS 2336 0413J-00732 4300-00137

```

If you install the 600/900 W AC PSU in a chassis other than a NETGEAR 8806, you see `unsupported` next to the state of the power supply, as shown in this sample truncated output:

```

PowerSupply 3 information:
State:          Unsupported
PartInfo:       PS 2431 0622J-00013 4300-00161

```

show power budget

```
show power budget
```

Description

Displays the power status and the amount of available and required power.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the amount of power available on the switch.

This status information may be useful if the `show slot` command displays a state of Powered OFF for any I/O module, for monitoring power, or for power planning purposes.

The first table of the `show power budget` command displays:

- Slot number of the power supply.
- Current state of the power supply. Options are:
 - Empty—There is no power supply installed.
 - Power Failed—The power supply has failed.
 - Power Off—The power supply is off.
 - Power On—The power supply is on.
- Watts and voltage amounts of the power supply.
- Redundant power information. Redundant power is the amount of power available if power to one PSU is lost. If a switch has PSUs with a mix of both 220V AC and 110V AC inputs, the amount of redundant power shown is based on the worst-case assumption that power to a PSU with 220V AC input is lost.

The second table of the `show power budget` command displays:

- Slot number and name of the component installed in the slot. Options include:
 - I/O modules
 - MSMs/MMs
 - Fan trays
- Current state of the module. Options include, among others:
 - Empty: There is no component installed.
 - Operational: The component is installed and operational.
 - Present: The component is installed but not operational.
 - Down: The module is installed, but the administrator has taken the module offline.
 - Power ON: There is sufficient system power to power up the module.
 - Powered OFF: There is insufficient system power to keep the module up and running, or there is a mismatch between the module configured for the slot and the actual module installed in the slot.
 - Booting: The module has completed downloading the software image and is now booting.
 - Initializing: The module is initializing.
- Watts and voltage amounts of the modules.
- Power Surplus or Power Shortfall.
 - If the amount of available power meets or exceeds the required port, the excess is displayed as the Power Surplus.

- If the available power is insufficient to meet the required power, the deficit is displayed as Power Shortfall.
- Redundant power information. If the amount of redundant power meets or exceeds the required power, the system has (N+1) power.
 - Yes—The system has redundant (N+1) power.
 - No—The system does not have redundant (N+1) power.

The information contained in this display is for planning purposes since the system operates without redundant power as long as a power surplus is shown. However, if power is lost to a single PSU when the system is not redundant, I/O modules are powered down. Sefer to the section “Understanding Power Supply Management” in Chapter 2 of the *NETGEAR 8800 User Manual*.

Depending on the software version running on your switch, the modules installed in your switch, and the type of switch you have, additional or different power information may be displayed.

Example

The following command displays the distribution of power and the available power on the switch:

```
show power budget
```

The following is sample output of this command from a NETGEAR 8800 series switch:

```
PS   State                               48V
-----
 1   Powered On                         624.00
 2   Powered On                         624.00
 3   Empty
 4   Empty
 5   Empty
 6   Empty
-----
Power Available:                        1248.00
Redundant (N+1) Power Available:       648.00

Slots   Type           State      Watts
-----
Slot-1                               Empty
Slot-2   GM-20T        Operational  149.00
Slot-5   GM-20T        Operational  149.00
Slot-6                               Empty
MSM-A    MSM-5         Operational  185.00
MSM-B                               Empty      185.00
FanTray                               Operational  45.00
-----
Power Required:                        713.00
```

```
Power Allocated:                713.00
Power Surplus:                  535.00
Redundant Power Supply(s) Present?: NO
```

show power controller

```
show power controller {<num>}
```

Description

Displays the current status of the installed power supply controllers.

Command Syntax

num	Specifies the slot number of the installed power supply controller.
-----	---

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the power supply controllers. Power controllers collect data about the installed power supplies and report the results to the MSM/MM.

This status information may be useful for your technical support representative if you have a network problem.

The switch collects the following power supply controller information:

- **State**—Indicates the current state of the power supply controller. Options are:
 - **Empty**: There is no power supply controller installed.
 - **Operational**: The power supply controller is installed and operational.
 - **Present**: The power supply controller is installed.
- **PartInfo**—Provides information about the power supply controller including the:
 - **Slot number** where the power supply controller is installed.
 - **Serial number**, a collection of numbers and letters, that make up the serial number of the power supply controller.
 - **Part number**, a collection of numbers and letters that make up the part number of the power supply controller.
- **Revision**—Displays the revision number of the power supply controller.
- **FailureCode**—Specifies the failure code of the power supply controller.
- **Odometer**—Specifies the date and how long the power supply controller has been operating.

- Temperature—Specifies, in Celsius, the current temperature of the power supply controller.
- Status—Specifies the status of the power supply controller.

Example

The following command displays the status of the installed power supply controllers:

```
show power controller
```

The following is sample output from this command:

PSUCTRL-1 information:

```
State:           Operational
PartInfo:        PSUCTRL-1 04334-00021 450117-00-01
Revision:        1.0
FailureCode:     0
Odometer:        337 days 7 hours   since Nov-30-2004
Temperature:     32.14 deg C
Status:          PSU CTRL Mode:    Master
```

PSUCTRL-2 information:

```
State:           Empty
```

If you have two power supply controllers installed, the switch displays output about both of the power supply controllers:

PSUCTRL-1 information:

```
State:           Operational
PartInfo:        PSUCTRL-1 04334-00021 450117-00-01
Revision:        1.0
FailureCode:     0
Odometer:        17 days 5 hours 30 minutes since Oct-19-2004
Temperature:     35.1 deg C
Status:          PSU CTRL Mode:    Master
```

PSUCTRL-2 information:

```
State:           Operational
PartInfo:        PSUCTRL-2 04334-00068 450117-00-01
Revision:        1.0
FailureCode:     0
Odometer:        4 days 13 hours   since Sep-21-2004
Temperature:     33.56 deg C
Status:          PSU CTRL Mode:    Backup
```

show session

```
show session {{detail}} {<sessID>}} {history}
```

Description

Displays the currently active Telnet and console sessions communicating with the switch.

Syntax Description

detail	Specifies more detailed session information.
sessID	Specifies a session ID number.
history	Displays a list of all sessions.

Default

N/A.

Usage Guidelines

The `show session` command displays the username and IP address of the incoming Telnet session, whether a console session is currently active, and the login time. Each session is numbered.

The switch accepts IPv6 connections. If the incoming session is from an IPv6 address, the `show session` output indicates IPv6.

You can specify the following options to alter the session output:

- `detail`—The output for all current sessions is displayed in a list format.
- `sessID`—The output for the specified session is displayed in a list format.
- `history`—Displays a list of current and previous sessions, including the user, type of session, location, and start and end time of the session.

The `show session` command fields are defined in [Table 7](#).

Table 7. Show Command Field Definitions

Field	Definition
#	Indicates session number.
Login Time	Indicates login time of session.
User	Indicates the user logged in for each session.
Type	Indicates the type of session, for example: console, telnet, http, https.
Auth	Indicates how the user is logged in.
CLI Auth	Indicates the type of authentication (RADIUS and TACACS) if enabled.
Location	Indicates the location (IP address) from which the user logged in. The output also indicates if the location is an IPv6 address.

Example

The following command displays the active sessions on the switch:

```
show session
```

The following is sample output from this command:

#	Login Time	User	Type	Auth	Auth Location
1	Thu Apr 28 20:16:56 2005	admin	console	local	dis serial
*2	Thu Apr 28 23:36:20 2005	admin	ssh2	local	dis 3001::20d:88ff:fec5:ad40
3	Fri Apr 29 11:14:27 2005	admin	telnet	local	dis 10.255.44.55

The following command displays a list of current and previous sessions on the switch:

```
show session history
```

The following is sample output from this command:

```

Session History:
admin                console    serial                Mon Jun 21 09:19:
00 2004  Mon Jun 21 10:00:16 2004
admin                console    serial                Tue Jun 22 07:28:
11 2004  Tue Jun 22 11:46:48 2004
admin                console    serial                Wed Jun 23 10:05:
44 2004  Wed Jun 23 14:11:47 2004
admin                console    serial                Thu Jun 24 07:07:
25 2004  Thu Jun 24 07:08:55 2004
admin                console    serial                Thu Jun 24 13:30:
07 2004  Active

```

show snmp

```
show snmp [get | get-next] <object_identifier>
```

Description

Displays the contents of an SNMP MIB object.

Syntax Description

object_identifier	Specifies the object identifier for an SNMP MIB object.
-------------------	---

Default

N/A.

Usage Guidelines

Use the `get` option to establish an index into the SNMP MIB. After the `get` option is executed, you can use the `get next` option to step through the MIB objects.

Example

The following gets the contents of SNMP object 1.3.6.1.2.1.1.5.0:

```
show snmp get 1.3.6.1.2.1.1.5.0
system.5.0 = BD-12804
```

show snmp vr_name

```
show snmp {vr} <vr_name>
```

Description

Displays the SNMP configuration and statistics on a virtual router.

Syntax Description

vr_name	Specifies the virtual router.
---------	-------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the SNMP configuration and statistics on a virtual router.

Example

The following command displays configuration and statistics for the virtual router *VR-Default*:

```
show snmp vr VR-Default
```

Following is sample output for the command:

```
SNMP access                : Disabled
SNMP Traps                  : Enabled
SNMP v1/v2c TrapReceivers  :
    Destination            Source IP Address    Flags
    10.120.91.89 /162      2E
Flags:  Version: 1=v1 2=v2c
        Mode: S=Standard E=Enhanced

SNMP stats:  InPkts 300    OutPkts  300    Errors 0    AuthErrors 0
              Gets   0     GetNexts  300    Sets   0     Drops    0
SNMP traps:  Sent    0     AuthTraps Enabled
```


show snmpv3 access

```
show snmpv3 access {[[hex <hex_group_name>] | <group_name>]}
```

Description

Displays SNMPv3 access rights.

Syntax Description

hex	Specifies that the value to follow is to be supplied as a colon separated string of hex octets.
group_name	Specifies the name of the group to display.

Default

N/A.

Usage Guidelines

The `show snmpv3 access` command displays the access rights of a group. If you do not specify a group name, the command will display details for all the groups.

This command displays the SNMPv3 vacmAccessTable entries.

Example

The following command displays all the access details:

```
show snmpv3 access
```

The following is sample output from this command:

```
X450a-24t.5 # show snmpv3 access
```

```
Group Name      : admin
Context Prefix  :
Security Model   : USM
Security Level   : Authentication Privacy
Context Match    : Exact
Read View        : defaultAdminView
Write View       : defaultAdminView
Notify View      : defaultNotifyView
Storage Type     : NonVolatile
Row Status       : Active

Group Name      : initial
Context Prefix   :
Security Model    : USM
Security Level    : No-Authentication No-Privacy
```

```

Context Match      : Exact
Read View          : defaultUserView
Write View         :
Notify View        : defaultNotifyView
Storage Type       : NonVolatile
Row Status         : Active

```

```

Group Name         : initial
Context Prefix     :
Security Model      : USM
Security Level      : Authentication No-Privacy
Context Match      : Exact
Read View          : defaultUserView
Write View         : defaultUserView
Notify View        : defaultNotifyView
Storage Type       : NonVolatile
Row Status         : Active

```

```

Group Name         : vlv2c_ro
Context Prefix     :
Security Model      : snmpv1
Security Level      : No-Authentication No-Privacy
Context Match      : Exact
Read View          : defaultUserView
Write View         :
Notify View        : defaultNotifyView
Storage Type       : NonVolatile
Row Status         : Active

```

```

Group Name         : vlv2c_ro
Context Prefix     :
Security Model      : snmpv2c
Security Level      : No-Authentication No-Privacy
Context Match      : Exact
Read View          : defaultUserView
Write View         :
Notify View        : defaultNotifyView
Storage Type       : NonVolatile
Row Status         : Active

```

```

Group Name         : vlv2c_rw
Context Prefix     :
Security Model      : snmpv1
Security Level      : No-Authentication No-Privacy
Context Match      : Exact
Read View          : defaultUserView

```

```

Write View      : defaultUserView
Notify View     : defaultNotifyView
Storage Type    : NonVolatile
Row Status      : Active

```

```

Group Name      : vlv2c_rw
Context Prefix  :
Security Model   : snmpv2c
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View       : defaultUserView
Notify View      : defaultNotifyView
Storage Type     : NonVolatile
Row Status      : Active

```

```

Group Name      : vlv2cNotifyGroup
Context Prefix  :
Security Model   : snmpv1
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       :
Write View       :
Notify View      : defaultNotifyView
Storage Type     : NonVolatile
Row Status      : Active

```

```

Group Name      : vlv2cNotifyGroup
Context Prefix  :
Security Model   : snmpv2c
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       :
Write View       :
Notify View      : defaultNotifyView
Storage Type     : NonVolatile
Row Status      : Active

```

Total num. of entries in vacmAccessTable : 9

The following command displays the access rights for the group *group1*:

```
show snmpv3 access group1
```

show snmpv3 community

```
show snmpv3 community
```

Description

Displays information about SNMP community strings.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

This command displays information about and status of the SNMP community on the switch. This information is available to Administrator Accounts.

Example

The following command displays the community:

```
show snmpv3 community
```

The following is sample output from this command.

```
X450a-24t.4 # show snmpv3 community
```

```

Community Index   : private
Community Name    : private
Security Name     : vlv2c_rw
Context EngineID  : 80:00:07:7c:03:00:04:96:27:b6:7b
Context Name      :
Transport Tag     :
Storage Type      : NonVolatile
Row Status        : Active

```

```

Community Index   : public
Community Name    : public
Security Name     : vlv2c_ro
Context EngineID  : 80:00:07:7c:03:00:04:96:27:b6:7b
Context Name      :
Transport Tag     :
Storage Type      : NonVolatile
Row Status        : Active

```

```
Total num. of entries in snmpCommunityTable : 2
```

show snmpv3 context

```
show snmpv3 context
```

Description

Displays information about the SNMPv3 contexts on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines:

This command displays the entries in the View-based Access Control Model (VACM) context table (VACMContextTable).

Example

The following command displays information about the SNMPv3 contexts on the switch:

```
show snmpv3 context
```

The following is sample output from this command:

```
VACM Context Name :
Note : This Version Supports one global context ("")
```

show snmpv3 counters

```
show snmpv3 counters
```

Description

Displays SNMPv3 counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show snmpv3 counters` command displays the following SNMPv3 counters:

- snmpUnknownSecurityModels
- snmpInvalidMessages
- snmpUnknownPDUHandlers
- usmStatsUnsupportedSecLevels

- usmStatsNotInTimeWindows
- usmStatsUnknownUserNames
- usmStatsUnknownEngineIDs
- usmStatsWrongDigests
- usmStatsDecryptionErrors

Issuing the command `clear counters` resets all counters to zero.

Example

The following command displays all the SNMPv3 counters.

```
show snmpv3 counters
```

The following is sample output from this command:

```
snmpUnknownSecurityModels      : 0
snmpInvalidMessages            : 0
snmpUnknownPDUHandlers         : 0
usmStatsUnsupportedSecLevels   : 0
usmStatsNotInTimeWindows       : 0
usmStatsUnknownUserNames       : 0
usmStatsUnknownEngineIDs       : 0
usmStatsWrongDigests           : 0
usmStatsDecryptionErrors       : 0
```

show snmpv3 engine-info

```
show snmpv3 engine-info
```

Description

Displays information about the SNMPv3 engine on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines:

The following show engine-info output is displayed:

- Engine-ID—Either the ID auto generated from MAC address of switch, or the ID manually configured.
- Engine Boots—Number of times the agent has been rebooted.
- Engine Time—Time since agent last rebooted, in centiseconds.

- Max. Message Size—Maximum SNMP Message size supported by the Engine (8192).

Example

The following command displays information about the SNMPv3 engine on the switch:

```
show snmpv3 engine-info
```

The following is sample output from this command:

```
SNMP Engine-ID       : 80:0:11:AE:3:0:30:48:41:ed:97 'H'
SNMP Engine Boots    : 1
SNMP Engine Time     : 866896
SNMP Max. Message Size : 8192
```

show snmpv3 filter

```
show snmpv3 filter {[[hex <hex_profile_name>] | <profile_name>] {{subtree}
<object_identifier>}}
```

Description

Displays the filters that belong a filter profile.

Syntax Description

hex_profile_name	Specifies the filter profile to display. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile to display in ASCII format.
object_identifier	Specifies a MIB subtree.

Default

N/A.

Usage Guidelines

Use this command to display entries from the snmpNotifyFilterTable. If you specify a profile name and subtree, you will display only the entries with that profile name and subtree. If you specify only the profile name, you will display all entries for that profile name. If you do not specify a profile name, then all the entries are displayed.

Example

The following command displays the part of filter profile *prof1* that includes the MIB subtree *1.3.6.1.4.1*:

```
show snmpv3 filter prof1 subtree 1.3.6.1.4.1
```

The following is sample output from this command:

```
Profile Name       : prof1
```

```

Subtree      : 1.3.6.1.4.1
Mask         :
Type         : Included
Storage Type : NonVolatile
Row Status   : Active

```

show snmpv3 filter-profile

```
show snmpv3 filter-profile {[[hex <hex_profile_name>] | <profile_name>]} {param [[hex
<hex_param_name>] | <param_name>]}
```

Description

Displays the association between parameter names and filter profiles.

Syntax Description

hex_profile_name	Specifies the filter profile name. The value is to be supplied as a colon separated string of hex octets.
profile_name	Specifies the filter profile name in ASCII format.
hex_param_name	Specifies the parameter name. The values is to be supplied as a colon separated string of hex octets.
param_name	Specifies the parameter name in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. The parameter name is associated with target addresses, and the filter profile is associated with a series of filters, so, in effect, you are associating a series of filters with a target address.

Example

The following command displays the entry with filter profile *prof1* with the parameter name *P1*:

```
show snmpv3 filter-profile prof1 param P1
```

The following is sample output of this command:

```

Filter Profile Params Name : p1
Name                       : prof1
Storage Type               : NonVolatile
Row Status                 : Active

```


show snmpv3 group

```
show snmpv3 group {[[hex <hex_group_name>] | <group_name>] {user [[hex <hex_user_name>] | <user_name>]]}}
```

Description

Displays the user name (security name) and security model association with a group name.

Syntax Description

hex_group_name	Specifies the group name to display. The value is to be supplied as a colon separated string of hex octets.
group_name	Specifies the group name to display. The value is to be supplied in ASCII format.
hex_user_name	Specifies the user name to display. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

The `show snmpv3 group` command displays the details of a group with the given group name. If you do not specify a group name, the command will display details for all the groups.

This command displays the SNMPv3 vacmSecurityToGroupTable.

Example

The following command displays information about all groups for every security model and user name:

```
show snmpv3 group
```

The following is sample output from this command:

```
X450a-24t.9 # sh snmpv3 group
```

```
Group Name      : v1v2c_ro
Security Name   : v1v2c_ro
Security Model  : snmpv1
Storage Type    : NonVolatile
Row Status      : Active
```

```
Group Name      : v1v2c_rw
Security Name   : v1v2c_rw
```

```

Security Model : snmpv1
Storage Type   : NonVolatile
Row Status     : Active

```

```

Group Name     : v1v2c_ro
Security Name   : v1v2c_ro
Security Model  : snmpv2c
Storage Type    : NonVolatile
Row Status     : Active

```

```

Group Name     : v1v2c_rw
Security Name   : v1v2c_rw
Security Model  : snmpv2c
Storage Type    : NonVolatile
Row Status     : Active

```

```

Group Name     : admin
Security Name   : admin
Security Model  : USM
Storage Type    : NonVolatile
Row Status     : Active

```

```

Group Name     : initial
Security Name   : initial
Security Model  : USM
Storage Type    : NonVolatile
Row Status     : Active

```

```

Group Name     : initial
Security Name   : initialmd5
Security Model  : USM
Storage Type    : NonVolatile
Row Status     : Active

```

```

Group Name     : initial
Security Name   : initialsha
Security Model  : USM
Storage Type    : NonVolatile
Row Status     : Active

```

```

Group Name     : initial
Security Name   : initialmd5Priv
Security Model  : USM
Storage Type    : NonVolatile
Row Status     : Active

```

```

Group Name       : initial
Security Name    : initialshaPriv
Security Model   : USM
Storage Type     : NonVolatile
Row Status       : Active

```

Total num. of entries in vacmSecurityToGroupTable : 10

The following command shows information about the group *testgroup* and user name *testuser*:

```
show snmpv3 group testgroup user testuser
```

The following is sample output from this command:

```

Group Name       : testgroup
Security Name    : testuser
Security Model   : USM
Storage Type     : NonVolatile
Row Status       : Active

```

show snmpv3 mib-view

```
show snmpv3 mib-view {[[hex <hex_view_name>] | <view_name>] {subtree <object_identifier>}}
```

Description

Displays a MIB view.

Syntax Description

hex_view_name	Specifies the name of the MIB view to display. The value is to be supplied as a colon separated string of hex octets.
view_name	Specifies the name of the MIB view to display. The value is to be supplied in ASCII format.
object_identifier	Specifies the object identifier of the view to display.

Default

N/A.

Usage Guidelines

The `show snmpv3 mib-view` command displays a MIB view. If you do not specify a view name, the command will display details for all the MIB views. If a subtree is not specified, then all subtrees belonging to the view name will be displayed.

This command displays the SNMPv3 vacmViewTreeFamilyTable.

Example

The following command displays all the view details:

```
show snmpv3 mib-view
```

The following is sample output from this command:

```
X450a-24t.10 # sh snmpv3 mib-view
```

```
View Name      : defaultUserView
MIB Subtree    : 1
Mask           :
View Type      : Included
Storage Type   : NonVolatile
Row Status     : Active
```

```
View Name      : defaultUserView
MIB Subtree    : 1.3.6.1.6.3.16
Mask           :
View Type      : Excluded
Storage Type   : NonVolatile
Row Status     : Active
```

```
View Name      : defaultUserView
MIB Subtree    : 1.3.6.1.6.3.18
Mask           :
View Type      : Excluded
Storage Type   : NonVolatile
Row Status     : Active
```

```
View Name      : defaultUserView
MIB Subtree    : 1.3.6.1.6.3.15.1.2.2.1.4
Mask           :
View Type      : Excluded
Storage Type   : NonVolatile
Row Status     : Active
```

```
View Name      : defaultUserView
MIB Subtree    : 1.3.6.1.6.3.15.1.2.2.1.6
Mask           :
View Type      : Excluded
Storage Type   : NonVolatile
Row Status     : Active
```

```
View Name      : defaultUserView
MIB Subtree    : 1.3.6.1.6.3.15.1.2.2.1.9
Mask           :
```

```
View Type      : Excluded
Storage Type   : NonVolatile
Row Status     : Active
```

```
View Name      : defaultAdminView
MIB Subtree    : 1
Mask           :
View Type      : Included
Storage Type   : NonVolatile
Row Status     : Active
```

```
View Name      : defaultNotifyView
MIB Subtree    : 1
Mask           :
View Type      : Included
Storage Type   : NonVolatile
Row Status     : Active
```

Total num. of entries in vacmViewTreeFamilyTable : 8

The following command displays a view with the view name *Roview* and subtree 1.3.6.1.2.1.1:

```
show snmpv3 mib-view Roview subtree 1.3.6.1.2.1.1
```

show snmpv3 notify

```
show snmpv3 notify {[hex <hex_notify_name>] | <notify_name>]}
```

Description

Displays the notifications that are set. This command displays the snmpNotifyTable.

Syntax Description

hex_notify_name	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
notify_name	Specifies the parameter name associated with the target. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display entries from the SNMPv3 snmpNotifyTable. This table lists the notify tags that the agent will use to send notifications (traps).

If no notify name is specified, all the entries are displayed.

Example

The following command displays the notify table entry for *N1*:

```
show snmpv3 notify N1
```

The following is sample output from this command:

```
Notify Name      : N1
Tag              : type1
Type             : Trap
Storage Type     : NonVolatile
Row Status       : Active
```

show snmpv3 target-addr

```
show snmpv3 target-addr {[[hex <hex_addr_name>] | <addr_name>]}
```

Description

Displays information about SNMPv3 target addresses.

Syntax Description

hex_addr_name	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
addr_name	Specifies a string identifier for the target address.

Default

N/A.

Usage Guidelines

Use this command to display entries in the SNMPv3 snmpTargetAddressTable. If no target address is specified, the entries for all the target addresses will be displayed.

To view the source IP address, use the [show management](#) command.

Example

The following command displays the entry for the target address named *A1*:

```
show snmpv3 target-addr A1
```

The following is sample output from this command:

```
Target Addr Name : A1
TDomain          : 1.3.6.1.6.1.1
TAddress         : 10.201.31.234, 162
```

```

TMask          :
Timeout        : 1500
Retry Count    : 0
Tag List       : defaultNotify
Params         : vlv2cNotifyParam1
Storage Type   : NonVolatile
Row Status     : Active
Storage Type   : NonVolatile
Row Status     : Active

```

show snmpv3 target-params

```
show snmpv3 target-params {[[hex <hex_target_params>] | <target_params>]}
```

Description

Displays the information about the options associated with the parameter name.

Syntax Description

hex_target_params	Specifies the parameter to display. The value is to be supplied as a colon separated string of hex octets.
target_params	Specifies the parameter name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display entries from the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

If no parameter name is specified, all the entries are displayed.

Example

The following command displays the target parameter entry named *P1*:

```
show snmpv3 target-params P1
```

The following is sample output from this command:

```

Target Params Name    : p1
MP Model              : snmpv2c
Security Model        : snmpv2c
User Name             : testuser

```

```
Security Level      : No-Authentication No-Privacy
Storage Type       : NonVolatile
Row Status         : Active
```

show snmpv3 user

```
show snmpv3 user {[[hex <hex_user_name>] | <user_name>]}
```

Description

Displays detailed information about the user.

Syntax Description

hex_user_name	Specifies the user name to display. The value is to be supplied as a colon separated string of hex octets.
user_name	Specifies the user name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

The `show snmpv3 user` command displays the details of a user. If you do not specify a user name, the command will display details for all the users. The authentication and privacy passwords and keys will not be displayed.

The user entries in SNMPv3 are stored in the USMUserTable, so the entries are indexed by EngineID and user name.

Example

The following command lists all user entries:

```
show snmpv3 user
```

The following is sample output from this command:

```
X450a-24t.11 # sh snmpv3 user
```

```
Engine-ID      : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name      : admin
Security Name   : admin
Authentication  : HMAC-MD5
Privacy        : DES
Storage Type    : NonVolatile
Row Status      : Active
```



```

Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initial
Security Name    : initial
Authentication   : No-Authentication
Privacy         : No-Privacy
Storage Type     : NonVolatile
Row Status      : Active

```

```

Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initialmd5
Security Name    : initialmd5
Authentication   : HMAC-MD5
Privacy         : No-Privacy
Storage Type     : NonVolatile
Row Status      : Active

```

```

Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initialsha
Security Name    : initialsha
Authentication   : HMAC-SHA
Privacy         : No-Privacy
Storage Type     : NonVolatile
Row Status      : Active

```

```

Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initialmd5Priv
Security Name    : initialmd5Priv
Authentication   : HMAC-MD5
Privacy         : DES
Storage Type     : NonVolatile
Row Status      : Active

```

```

Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initialshaPriv
Security Name    : initialshaPriv
Authentication   : HMAC-SHA
Privacy         : DES
Storage Type     : NonVolatile
Row Status      : Active

```

Total num. of entries in usmUserTable : 6

The following command lists details for the specified user, *testuser*.

```
show snmpv3 user testuser
```

show sntp-client`show sntp-client`**Description**

Displays the DNS configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays configuration and statistics information of SNTP client.

Example

The following command displays the SNTP configuration:

```
show sntp-client
```

The following is sample output from this command:

```

SNTP client is enabled
SNTP time is valid
Primary server: 172.17.1.104
Secondary server: 172.17.1.104
Query interval: 64
Last valid SNTP update: From server 172.17.1.104, on Wed Oct 30 22:46:03 2003
SNTPC Statistics:
  Packets transmitted:
    to primary server:          1
    to secondary server:        0
  Packets received with valid time:
    from Primary server:        1
    from Secondary server:       0
    from Broadcast server:       0
  Packets received without valid time:
    from Primary server:         0
    from Secondary server:        0
    from Broadcast server:        0
  Replies not received to requests:
    from Primary server:         0
    from Secondary server:        0

```

telnet

```
telnet {vr <vr_name>} [<host_name> | <remote_ip>] {<port>}
```

Description

Allows you to Telnet from the current command-line interface session to another host.

Syntax Description

vr	Specifies use of a virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
vr_name	Specifies the name of the virtual router.
host_name	Specifies the name of the host.
remote_ip	Specifies the IP address of the host.
port	Specifies a TCP port number. The default is port 23.

Default

- Telnet—enabled
- Virtual router—Uses all virtual routers on the switch for outgoing Telnet requests
- Port—23

Usage Guidelines

Only VT100 emulation is supported.

Before you can start an outgoing Telnet session, you need to configure the switch IP parameters. To open a Telnet connection, you must specify the host IP address or the host name of the device you want to connect to. Check the user manual supplied with the Telnet facility if you are unsure of how to do this. Although the switch accepts IPv6 connections, you can only Telnet from the switch to another device with an IPv4 address.

You must configure DNS in order to use the `host_name` option.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)

- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Virtual Router Requirements

The `vr_name` option specifies the name of the virtual router. The valid virtual router names at system boot-up are *VR-Mgmt*, *VR-Control*, and *VR-Default*; however, you can Telnet only on *VR-Mgmt* and *VR-Default*. For more information about virtual routers, see the section “Virtual Routers” in the *NETGEAR 8800 User Manual*.

Example

The following command starts a Telnet client communication to the host at IP address 123.45.67.8:

```
telnet 123.45.67.8
```

The following command starts a Telnet client communication with a host named *sales*:

```
telnet sales
```

telnet msm

```
telnet msm [a | b]
```

Description

Allows you to Telnet to either the primary or the backup MSM regardless of which console port you are connected to.

Syntax Description

a	Specifies the MSM installed in slot A.
b	Specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

Use this command to access either the primary or the backup MSM regardless of which console port you are connected to. For example, if MSM A is the primary MSM and you are connected to MSM A via its console port, you can access the backup MSM installed in slot B by issuing the `telnet msm b` command.

Example

The following example makes the following assumptions:

- The MSM installed in slot A is the primary
- The MSM installed in slot B is the backup
- You have a console connection to MSM B

The following command accesses the primary MSM installed in slot A from the backup MSM installed in slot B:

```
My8800.6 # telnet msm b
```

```
Entering character mode
Escape character is '^]'.
```

```
telnet session telnet0 on /dev/ptyb0
```

```
login: admin
password:
```

```
NETGEAR 8800
```

```
Copyright (C) 2000-2007 NETGEAR. All rights reserved.
```

```
Protected by US Patent Nos: 6,678,248; 6,104,700; 6,766,482; 6,618,388; 6,034,957; 6,859,438;
6,912,592; 6,954,436; 6,977,891; 6,980,550; 6,981,174; 7,003,705; 7,012,082; 7,046,665;
7,126,923; 7,142,509; 7,149,217; 7,152,124; 7,154,861.
```

```
=====
```

```
You are connected to a Backup node. Only a limited command set is supported.
You may use "telnet msm A" to connect to the Master node to access
the full set of commands.
```

```
Press the <tab> or '?' key at any time for completions.
Remember to save your configuration changes.
```

```
My8800.1 >
```

tftp

```
tftp [<host-name> | <ip-address>] {-v <vr_name>} [-g | -p] [{-l [internal-memory
<local-file-internal> | memorycard <local-file-memcard> | <local-file>} {-r <remote-file>} |
{-r <remote-file>} {-l [internal-memory <local-file-internal> | memorycard
<local-file-memcard> | <local-file>}}]
```

Description

Allows you to TFTP from the current command line interface session to a TFTP server.

Syntax Description

host-name	Specifies the name of the remote host.
ip-address	Specifies the IP address of the TFTP server.
vr_name	Specifies the name of the virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
-g	Gets the specified file from the TFTP server and copies it to the local host.
-p	Puts the specified file from the local host and copies it to the TFTP server.
internal-memory	Specifies the internal memory card.
local-file-internal	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash card.
local-file-memcard	Specifies the name of the file on the external compact flash card.
local-file	Specifies the name of the file (configuration file, policy file) on the local host.
remote-file	Specifies the name of the file on the remote host.

Default

If you do not specify a virtual router, *VR-Mgmt* is used.

Usage Guidelines

NetASCII and mail file type formats are not supported.

TFTP Server Requirements

NETGEAR recommends using a TFTP server that supports blocksize negotiation (as described in RFC 2348, *TFTP Blocksize Option*), to enable faster file downloads and larger file downloads. If the TFTP server does not support blocksize negotiation, the file size is limited to 32 MB. Older TFTP servers that do not support blocksize negotiation have additional implementation limits that may decrease the maximum file size to only 16 MB, which may be too small to install NETGEAR 8800 images.

If your TFTP server does not support blocksize negotiation, the switch displays a message similar to the following when you attempt a get (-g) or put (-p) operation:

```
Note: The blocksize option is not supported by the remote TFTP server.
      Without this option, the maximum file transfer size is limited to 32MB.
      Some older TFTP servers may be limited to 16MB file.
```

Using TFTP

Use TFTP to download a previously saved configuration file or policy file from the TFTP server to the switch. When you download a file, this command does not automatically apply it to the switch. You must specify that the downloaded file be applied to the switch. For example, if you download a configuration file, issue the `use configuration` command to apply the saved configuration on the next reboot. You must use the `reboot` command to activate the new configuration. If you download a policy file, use the `refresh policy` command to reprocess the text file and update the policy database.

You also use TFTP to upload a saved configuration file or policy file from the switch to the TFTP server.

If your download from the TFTP server to the switch is successful, the switch displays a message similar to the following:

```
Downloading megtest2.cfg to switch... done!
```

If your upload from the switch to the TFTP server is successful, the switch displays a message similar to the following:

```
Uploading megtest1.cfg to TFTPHost ... done!
```

Up to eight active TFTP sessions can run on the switch concurrently.

You must configure DNS in order to use the `host_name` option.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Local and Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)

- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a local or remote file, remember the requirements listed above.

Virtual Router Requirements

The `vr_name` option specifies the name of the virtual router. The valid virtual router names at system boot-up are *VR-Mgmt*, *VR-Control*, and *VR-Default*; however, you can TFTP only on *VR-Mgmt* and *VR-Default*. On the NETGEAR 8800 switch, you can also create and configure your own virtual routers. For more information about virtual routers, see the section “Virtual Routers” in the *NETGEAR 8800 User Manual*.

Internal Memory and Core Dump Files

Core dump files have a `.gz` file extension. The filename format is:

`core.<process-name.pid>.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process. If you save core dump files to an external memory card, the filename also includes the affected MSM/MM: MSM-A or MSM-B.

If you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to transfer those files from the internal memory card to a TFTP server. You can also transfer core dump information to and from an external compact flash card.

If the switch has not saved any debug files, you cannot transfer other files to or from the internal memory. For example if you attempt to transfer a configuration file from the switch to the internal memory, the switch displays a message similar to the following:

```
Error: tftp transfer to internal-memory not allowed.
```

For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

For more detailed information about core dump files, see the troubleshooting appendix in the *NETGEAR 8800 User Manual*.

If you specify the `memorycard` option, you can copy and transfer files to and from the external memory card using TFTP.

Other Useful Commands

To upgrade the image, use the `download image` command. This command utilizes TFTP to transfer the software image file from your TFTP server to the switch. For more information about this command, see [download image](#) on page 1308.

Example

The following command downloads the configuration file named *XOS1.cfg* from the TFTP server with an IP address of 10.123.45.67:

```
tftp 10.123.45.67 -v "VR-Default" -g -r XOS1.cfg
```

The following command uploads the configuration file named *XOS2.cfg* to the TFTP server with an IP address of 10.123.45.67:

```
tftp 10.123.45.67 -v "VR-Default" -p -r XOS2.cfg
```

The following command retrieves and transfers files from an external memory card:

```
tftp 10.1.2.3. -g -l memorycard test.pol -r august23.pol
```

tftp get

```
tftp get [<host-name> | <ip-address>] {-vr <vr_name>} [{[internal-memory  
<local-file-internal> | memorycard <local-file-memcard> | <local_file>] {<remote_file>} |  
{<remote_file>} {[internal-memory <local-file-internal> | memorycard <local-file-memcard> |  
<local_file>}}] {force-overwrite}
```

Description

Allows you to use TFTP from the current command line interface session to copy the file from a TFTP server and copy it to a local host, including the switch, internal memory card, or external compact flash card.

Syntax Description

host-name	Specifies the name of the remote host.
ip-address	Specifies the IP address of the TFTP server.
vr_name	Specifies the name of the virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
internal-memory	Specifies the internal memory card.
local-file-internal	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash card.
local-file-memcard	Specifies the name of the file on the external compact flash card.
local_file	Specifies the name of the file (configuration file, policy file) on the local host.
remote_file	Specifies the name of the file on the remote host.
force-overwrite	Specifies the switch to automatically overwrite an existing file.

Default

If you do not specify a virtual router, *VR-Mgmt* is used; if you transfer a file with a name that already exists on the system, the switch prompts you to overwrite the existing file.

Usage Guidelines

NetASCII and mail file type formats are not supported.

By default, the switch prompts you to overwrite an existing file. For example, if you have a file named `test.cfg` on the switch and download a file named `test.cfg` from a TFTP server, the switch displays a message similar to the following:

```
test.cfg already exists, do you want to overwrite it? (y/n)
```

Enter `y` to download the file and overwrite the existing file. Enter `n` to cancel this action.

If you successfully download the file, the switch displays a message similar to the following:

```
Downloading test.cfg to switch... done!
```

If you cancel this action, the switch displays a message similar to the following:

```
Tftp download aborted.
```

If you specify the `force-overwrite` parameter, the switch automatically overwrites an existing file. For example, if you have a file named `test.cfg` on the switch and download a file named `test.cfg` from a TFTP server, the switch automatically overrides the existing file. If you successfully download the file, the switch displays a message similar to the following:

```
Downloading test.cfg to switch... done!
```

This command was introduced to simplify using TFTP to transfer configuration, policy, and if configured, core dump files from the switch to the TFTP server. You can continue to use the original `tftp` command.

For more information about TFTP, including:

- TFTP server requirements
- How to use TFTP
- Host name and remote IP address character restrictions
- Local and remote filename character restrictions
- Virtual router requirements
- Internal memory and core dump files
- Other useful commands

See the `tftp` command [on page 149](#).

Example

The following command retrieves and transfers the file `test.pol` from a TFTP server with an IP address of `10.1.2.3` and renames the file `august23.pol` when transferred to an external memory card installed the switch:

```
tftp get 10.1.2.3 vr "VR-Mgmt" test.pol memory-card august23.pol
```

The following command retrieves the configuration file named meg-upload.cfg from a TFTP server with an IP address of 10.10.10.10:

```
tftp get 10.10.10.10 vr "VR-Mgmt" meg_upload.cfg
```

tftp put

```
tftp put [<host-name> | <ip-address>] {-vr <vr_name>} [{[internal-memory  
<local-file-internal> | memorycard <local-file-memcard> | <local_file>} {<remote_file>} |  
{<remote_file>} {[internal-memory <local-file-internal> | memorycard <local-file-memcard> |  
<local_file>}]}
```

Description

Allows you to use TFTP from the current command line interface session to copy the file from the local host, including the switch, internal memory card, or external compact flash card and put it on a TFTP server.

Syntax Description

host-name	Specifies the name of the remote host.
ip-address	Specifies the IP address of the TFTP server.
vr_name	Specifies the name of the virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A of the <i>NETGEAR 8800 User Manual</i> .
internal-memory	Specifies the internal memory card.
local-file-internal	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash card.
local-file-memcard	Specifies the name of the file on the external compact flash card.
local_file	Specifies the name of the file (configuration file, policy file) on the local host.
remote_file	Specifies the name of the file on the remote host.

Default

If you do not specify a virtual router, *VR-Mgmt* is used.

Usage Guidelines

NetASCII and mail file type formats are not supported.

This command was introduced to simplify using TFTP to transfer configuration, policy, and if configured, core dump files from the switch to the TFTP server. You can continue to use the original TFTP command.

For more information about TFTP, including:

- TFTP server requirements
- How to use TFTP
- Host name and remote IP address character restrictions
- Local and remote filename character restrictions
- Virtual router requirements
- Internal memory and core dump files
- Other useful commands

See the [tftp](#) command [on page 149](#).

Example

The following command transfers a saved, not currently used configuration file named *XOS1.cfg* from the switch to the TFTP server:

```
tftp put 10.123.45.67 vr "VR-Mgmt" XOS1.cfg
```

Commands for Managing the NETGEAR 8800 Software

4

This chapter describes commands for:

- Working with the configuration and policy files used by the switch
- Starting, stopping, and displaying information about processes on the switch
- Viewing system memory resources
- Monitoring CPU utilization

Note: For information about downloading and upgrading a new software image, saving configuration changes, and upgrading the BootROM, see [Appendix A, “Configuration and Image Commands.”](#)

Like any advanced operating system, NETGEAR 8800 OS gives you the tools to manage your switch and create your network configurations. The following enhancements and functionality are included in the switch operating system:

- File system administration—You can move, copy, and delete files from the switch. The file system structure allows you to keep, save, rename, and maintain multiple copies of configuration files on the switch. In addition, you can manage other entities of the switch such as policies and access control lists (ACLs).
- Configuration file management—You can oversee and manage multiple configuration files on your switch. In addition, you can upload, download, modify, and name configuration files used by the switch.
- Process control—You can stop and start processes, restart failed processes, and update the software for a specific process or set of processes.
- Memory protection—With memory protection, the NETGEAR 8800 protects each process from every other process in the system. If one process experiences a memory fault, that process cannot affect the memory space of another process.
- CPU monitoring—You can monitor CPU utilization for Management Switch Fabric Modules (MSMs)/Management Modules (MMs) and the individual processes running on the switch. Monitoring the workload of the CPU allows you to troubleshoot and identify suspect processes.

Note: Filenames are case-sensitive.

clear cpu-monitoring

```
clear cpu-monitoring {process <name>} {slot <slotid>}
```

Description

Clears, resets the CPU utilization history and statistics stored in the switch.

Syntax Description

name	Specifies the name of the process.
slotid	Specifies the slot number of the MSM/MM module: <ul style="list-style-type: none"> • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

When you do not specify any keywords, this command clears the CPU utilization history for the entire switch, including processes, and resets the statistics to zero (0). This command also clears the CPU utilization history of the installed MSMs/MMs.

When you specify `process`, the switch clears and resets the CPU utilization history for the specified process.

When you specify `slot`, the switch clears and resets the CPU utilization history for the specified MSM/MM.

Example

The following command resets the CPU history and resets the statistics to 0 for the TFTP process running on the MSM/MM installed in slot A:

```
clear cpu-monitoring process tftpd slot A
```

cp

```
cp [internal-memory <old-name-internal> internal-memory <new-name-internal> | internal-memory <old-name-internal> memorycard <new-name-memorycard> | memorycard <old-name-memorycard> memorycard <new-name-memorycard> | memorycard <old-name-memorycard> <new-name> | <old-name> memorycard <new-name-memorycard> | <old-name> <new-name>]
```

Description

Copies an existing configuration, policy, or if configured, core dump file stored in the system.

Syntax Description

internal-memory	Specifies the internal memory card.
old-name-internal	Specifies the name of the core dump file located on the internal memory card that you want to copy.
new-name-internal	Specifies the name of the newly copied core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash memory card.
old-name-memorycard	Specifies the name of the file located on the external compact flash memory card that you want to copy. Depending on your switch configuration, you can have configuration, policy, or core dump files stored in this card.
new-name-memorycard	Specifies the name of the newly copied file located on the external compact flash memory card.
old-name	Specifies the name of the configuration or policy file that you want to copy.
new-name	Specifies the name of the newly copied configuration or policy file.

Default

N/A.

Usage Guidelines

Use this command to make a copy of an existing file before you alter or edit the file. By making a copy, you can easily go back to the original file if needed.

When you copy a configuration or policy file, remember the following:

- XML-formatted configuration files have a .cfg file extension. The switch only runs .cfg files.
- ASCII-formatted configuration files have a .xsf file extension. For more information, see Appendix B in the *NETGEAR 8800 User Manual*.
- Policy files have a .pol file extension.
- Core dump files have a .gz file extension. See “[Internal Memory and Core Dump Files](#)” below.

When you copy a configuration or policy file from the system, make sure you specify the appropriate file extension. For example, when you want to copy a policy file, specify the filename and *.pol*.

When you copy a file on the switch, the switch displays a message similar to the following:

```
Copy config test.cfg to config test1.cfg on switch? (y/n)
```

Enter `y` to copy the file. Enter `n` to cancel this process and not copy the file.

When you enter `y`, the switch copies the file with the new name and keeps a backup of the original file with the original name. After the switch copies the file, use the `ls` command to display a complete list of files. In this example, the switch displays the original file named `test.cfg` and the copied file named `test_rev2.cfg`.

The following is sample output from the `ls` command:

```
...
-rw-r--r--    1 root    root      100980 Sep 23 09:16 test.cfg
-rw-r--r--    1 root    root      100980 Oct 13 08:47 test_rev2.cfg
...
```

When you enter `n`, the switch displays a message similar to the following:

```
Copy cancelled.
```

Case-sensitive Filenames

Filenames are case-sensitive. In this example, you have a configuration file named `Test.cfg`. If you attempt to copy the file with the incorrect case, for example `test.cfg`, the switch displays a message similar to the following:

```
Error: cp: /config/test.cfg: No such file or directory
```

Since the switch is unable to locate `test.cfg`, the file is not copied.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Internal Memory and Core Dump Files

Core dump files have a `.gz` file extension. The filename format is:

`core.<process-name.pid>.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process. If you save core dump files to an external memory card, the filename also includes the affected MSM/MM: MSM-A or MSM-B.

By making a copy of a core dump file, you can easily compare new debug information with the old file if needed.

When you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option and associated internal-memory name options to copy an existing core dump file. If your switch has an external compact flash memory card installed, you can copy the core dump file to that card.

For information about configuring and sending core dump information to the internal memory card, see the [configure debug core-dumps](#) and [save debug tracefiles memorycard](#) commands.

For more detailed information about core dump files, see Appendix D in the *NETGEAR 8800 User Manual*.

This command also replicates the action from the primary MSM/MM to the backup MSM/MM. For example, when you copy a file on the primary MSM, the same file is copied to the backup MSM/MM.

For the `memorycard` option, the source and/or destination is the memorycard. You must mount the memory card for this operation to succeed. The `cp` command copies a file from the switch to the external memory card or a file already on the card. If you copy a file from the switch to the external memory card, and the new filename is identical to the source file, you do not need to re-enter the filename.

When you send core dump information to the external memory card, specify the `memorycard` option and associated memorycard name options to copy an existing core dump file.

Example

The following command makes a copy of a configuration file named *test.cfg* and gives the copied file a new name of *test_rev2.cfg*:

```
cp test.cfg test_rev2.cfg
```

The following command makes a copy of a configuration file named *primary.cfg* from the switch to an external memory card with the same name, *primary.cfg*:

```
cp primary.cfg memorycard
```

The above command performs the same action as entering the following command:

```
cp primary.cfg memorycard primary.cfg
```

disable cpu-monitoring

```
disable cpu-monitoring
```

Description

Disables CPU monitoring on the switch.

Command Syntax

This command has no arguments or variables.

Default

CPU monitoring is enabled and occurs every 5 seconds.

Usage Guidelines

Use this command to disable CPU monitoring on the switch.

This command does not clear the monitoring interval. Therefore, if you altered the CPU monitoring interval, this command does not return the CPU monitoring interval to 5 seconds. To return to the default frequency level, use the `enable cpu-monitoring {interval <seconds>} {threshold <percent>}` and specify 5 for the interval.

Example

The following command disables CPU monitoring on the switch:

```
disable cpu-monitoring
```

disable xml-mode

```
disable xml-mode
```

Description

Disables XML configuration mode on the switch.

Command Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable the XML configuration mode on the switch. XML configuration mode is not supported for end users.

See the command:

```
enable xml-mode
```

Example

The following command disables XML configuration mode on the switch:

```
disable xml-mode
```

enable cpu-monitoring

```
enable cpu-monitoring {interval <seconds>} {threshold <percent>}
```

Description

Enables CPU monitoring on the switch.

Command Syntax

seconds	Specifies the monitoring interval, in seconds. The default is 5 seconds, and the range is 5 to 60 seconds.
threshold	Specifies the CPU threshold value. CPU usage is measured in percentages. The default is 90%, and the range is 0% to 100%.

Default

CPU monitoring is enabled and occurs every 5 seconds. The default CPU threshold value is 90%.

Usage Guidelines

CPU monitoring allows you to monitor the CPU utilization and history for all of the processes running on the switch. By viewing this history on a regular basis, you can see trends emerging and identify processes with peak utilization. Monitoring the workload of the CPU allows you to troubleshoot and identify suspect processes before they become a problem.

To specify the frequency of CPU monitoring, use the `interval` keyword. NETGEAR recommends the default setting for most network environments.

CPU usage is measured in percentages. By default, the CPU threshold value is 90%. When CPU utilization of a process exceeds 90% of the regular operating basis, the switch logs an error message specifying the process name and the current CPU utilization for the process. To modify the CPU threshold level, use the `threshold` keyword. The range is 0% to 100%.

Example

The following command enables CPU monitoring every 30 seconds:

```
enable cpu-monitoring interval 30
```

enable xml-mode

```
enable xml-mode
```

Description

Enables XML configuration mode on the switch.

Command Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the XML configuration mode on the switch, however XML configuration mode is not supported for end users, and NETGEAR strongly cautions you not to enable this mode. Use this command only under the direction of NETGEAR.

If you inadvertently issue this command, the switch prompt will be changed by adding the text (xml) to the front of the prompt. If you see this mode indicator, please disable XML configuration mode by using the following command:

```
disable xml-mode
```

Example

The following command enables XML configuration mode on the switch:

```
enable xml-mode
```

ls

```
ls {[internal-memory | memorycard]} {<file-name>}
```

Description

Lists all configuration, policy, and if configured, core dump files in the system.

Syntax Description

internal-memory	Lists the core dump (debug) files that are present and saved in the internal memory card.
memorycard	Lists all of the files on the removable external compact flash memory card.
file-name	Lists all the files that match the wildcard.

Default

N/A.

Usage Guidelines

When you use issue this command without any options, the output displays all of the configuration and policy files stored on the switch.

When you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to display the core dump files stored on the internal memory card. For more information, see [Core Dump Files](#) on page 165.

When you specify the `memorycard` option, the output displays all of the files stored on the external compact flash memory card, including core dump files if so configured. For more information, see [Core Dump Files](#) on page 165.

When you specify the `<file-name>` option, the output displays all of the files that fit the wildcard criteria.

Understanding the Output

Output from this command includes the following:

- The first column displays the file permission using the following ten place holders:
 - The first place holder displays `-` for a file.
 - The next three place holders display `r` for read access and `w` for write access permission for the file owner.
 - The following three place holders display `r` for read access permission for members of the file owner's group.
 - The last three place holders display `r` for read access for every user that is not a member of the file owner's group.
- The second column displays how many links the file has to other files or directories.
- The third column displays the file owner.
- The remaining columns display the file size, date and time the file was last modified, and the file name.

Core Dump Files

Core dump files have a `.gz` file extension. The filename format is:

`core.<process-name.pid>.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process. If you save core dump files to an external memory card, the filename also includes the affected MSM/MM: MSM-A or MSM-B.

When the switch has not saved any debug files, no files are displayed. For information about configuring and sending core dump information to the internal memory card or the external memory card, see the [configure debug core-dumps](#) and [save debug tracefiles memorycard](#) commands.

For more detailed information about core dump files, see Appendix D in the *NETGEAR 8800 User Manual*.

Example

The following command displays a list of all current configuration and policy files in the system:

```
ls
```

The following is sample output from this command:

```
total 424
-rw-r--r--  1 root    root          50 Jul 30 14:19 hugh.pol
```

```
-rw-r--r--    1 root    root          94256 Jul 23 14:26 hughtest.cfg
-rw-r--r--    1 root    root        100980 Sep 23 09:16 megtest.cfg
-rw-r--r--    1 root    root          35 Jun 29 06:42 newpolicy.pol
-rw-r--r--    1 root    root        100980 Sep 23 09:17 primary.cfg
-rw-r--r--    1 root    root          94256 Jun 30 17:10 roytest.cfg
```

The following command displays a list of all current configuration and policy files in an external memory card:

```
ls memorycard
```

The following is sample output from this command:

```
-rwxr-xr-x    1 root    0          15401865 Mar 30 00:03 NG8800-12.4.3.5-1-4.xos
-rwxr-xr-x    1 root    0           10 Mar 31 09:41 test-1.pol
-rwxr-xr-x    1 root    0           10 Apr  4 09:15 test.pol
-rwxr-xr-x    1 root    0           10 Mar 31 09:41 test_1.pol
-rwxr-xr-x    1 root    0        223599 Mar 31 10:02 v11_1_3.cfg
```

The following command displays a list of all configuration and policy files with a filename beginning with the letter “a.”

```
(debug) BD-12804.1 # ls a*
```

Following is sample output from this command:

```
-rw-r--r--    1 root    0           2062 Jan  6 09:11 abc
-rw-rw-rw-    1 root    0           1922 Jan  7 02:19 abc.xsf
```

```
lk-blocks      Used Available Use%
    16384         496    15888    3%
```

The following command displays a list of all .tgz files

```
(debug) BD-12804.24 # ls internal-memory *.tgz
```

Following is sample output from this command:

```
-rwxr-xr-x    1 root    0           79076 Jan  6 09:47 old_traces.tgz
```

```
lk-blocks      Used Available Use%
    49038        110    48928    0%
```

mv

```
mv [internal-memory <old-name-internal> internal-memory <new-name-internal> | internal-memory
<old-name-internal> memorycard <new-name-memorycard> | memorycard <old-name-memorycard>
memorycard <new-name-memorycard> | memorycard <new-name-memorycard> <new-name> | <old-name>
memorycard <new-name-memorycard> | <old-name> <new-name>]
```

Description

Moves or renames an existing configuration, policy, or if configured, core dump file in the system.

Syntax Description

internal-memory	Specifies the internal memory card.
old-name-internal	Specifies the current name of the core dump file located on the internal memory card.
new-name-internal	Specifies the new name of the core dump file located on the internal memory card.
memorycard	Specifies the removable external compact flash card.
old-name-memorycard	Specifies the current name of the file located on the external compact flash memory card. Depending on your switch configuration, you can have configuration, policy, or cord dump files stored in this card.
new-name-memorycard	Specifies the new name of the file located on the external compact flash memory card.
old-name	Specifies the current name of the configuration or policy file on the system.
new-name	Specifies the new name of the configuration or policy file on the system.

Default

N/A.

Usage Guidelines

When you rename a file with a given extension, remember the following:

- XML-formatted configuration files have the .cfg file extension. The switch only runs .cfg files.
- ASCII-formatted configuration files have the .xsf file extensions. See Appendix B in the *NETGEAR 8800 User Manual* for more information.
- Policy files have the .pol file extension.
- Core dump files have the .gz file extension. See [Internal Memory and Core Dump Files](#) on page 168 for more information.

Make sure the renamed file uses the same file extension as the original file. If you change the file extensions, the file may be unrecognized by the system. For example, if you have an existing configuration file named *test.cfg*, the new filename must include the .cfg file extension.

You cannot rename an active configuration file (the configuration currently selected to boot the switch). To verify the configuration that you are currently using, issue the `show switch {detail}` command. If you attempt to rename the active configuration file, the switch displays a message similar to the following:

```
Error: Cannot rename current selected active configuration file.
```

When you rename a file, the switch displays a message similar to the following:

```
Rename config test.cfg to config megtest.cfg on switch? (y/n)
```

Enter y to rename the file on your system. Enter n to cancel this process and keep the existing filename.

Case-sensitive Filenames

Filenames are case-sensitive. In this example, you have a configuration file named Test.cfg. If you attempt to rename the file with the incorrect case, for example test.cfg, the switch displays a message similar to the following:

```
Error: mv: unable to rename `/config/test.cfg': No such file or directory
```

Since the switch is unable to locate test.cfg, the file is not renamed.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local or remote file, remember the requirements listed above.

Internal Memory and Core Dump Files

Core dump files have a .gz file extension. The filename format is:

`core.<process-name.pid>.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process. If you save core dump files to an external memory card, the filename also includes the affected MSM/MM: MSM-A or MSM-B.

When you configure the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to rename an existing core dump file. If your switch has an external compact flash memory card installed, you can move and rename the core dump file to that card.

For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

This command also replicates the action from the primary MSM/MM to the backup MSM/MM. For example, when you rename a file on the primary MSM/MM, the same file on the backup MSM/MM is renamed.

For the `memorycard` option, this command moves files between the external memory card and the switch. If you use the `memorycard` option for both the `old-name` and the `new-name`, this command just renames a file on the external memory card.

For information about core dump files, see the previous section “[Internal Memory and Core Dump Files](#).”

Example

The following command renames the configuration file named *Testb91.cfg* to *Activeb91.cfg*:

```
mv Testb91.cfg Activeb91.cfg
```

If the switch has an external memory card installed, the following command moves the configuration file named *test1.cfg* from the switch to the external memory card:

```
mv test1.cfg memorycard test1.cfg
```

If you do not change the name of the configuration file, you can also use the following command to move the configuration file *test1.cfg* from the switch to the external memory card:

```
mv test1.cfg memorycard
```

If the switch has an external memory card installed, the following command moves the policy file named *bgp.pol* from the memorycard to the switch:

```
mv memorycard bgp.pol bgp.pol
```

restart process

```
restart process [class <cname> | <name> {msm <slot>}]
```

Description

Terminates and restarts the specified process during a software upgrade on the switch.

Syntax Description

cname	Specifies the name of the process to restart. With this parameter, you can terminate and restart all instances of the process associated with a specific routing protocol on all VRs. You can restart the OSPF routing protocol and associated processes.
-------	--

name	Specifies the name of the process to terminate and restart. You can use this command with the following processes: <ul style="list-style-type: none"> • bgp • exsshd • lldp • netLogin • netTools • ospf • snmpSubagent • snmpMaster • telnetd • tftpd • vrrp • xmld
slot	Specifies the MSM/MM where the process should be terminated and restarted. A specifies the MSM/MM installed in slot A, and B specifies the MSM/MM installed in slot B.

Default

N/A.

Usage Guidelines

Use this command to terminate and restart a process during a software upgrade on the switch. You have the following options:

- `cname`—Specifies that the software terminates and restarts all instances of the process associated with a specific routing protocol on all VRs.
- `name`—Specifies the name of the process.

Depending on the software version running on your switch and the type of switch you have, you can terminate and restart different or additional processes. To see which processes you can restart during a software upgrade, enter `restart process` followed by TAB. The switch displays a list of available processes.

You can also use the `restart process` command when upgrading a software modular package. For more information, see the section “Upgrading a Modular Software Package” in Appendix B of the *NETGEAR 8800 User Manual*.

Example

The following command stops and restarts the process `tftpd` during a software upgrade:

```
restart process tftpd
```

The following command stops and restarts all instances of the OSPF routing protocol for all VRs during a software upgrade:

```
restart process class ospf
```

rm

```
rm {internal-memory | memorycard} <file-name>
```

Description

Removes/deletes an existing configuration, policy, or if configured, core dump file from the system.

Syntax Description

internal-memory	Specifies the internal memory card.
memorycard	Specifies the removable external compact flash card.
file-name	Specifies the name of the configuration, policy file, or if configured, the core dump file.

Default

N/A.

Usage Guidelines

After you remove a configuration or policy file from the system, that file is unavailable to the system. For information about core dump files, see [Internal Memory Card and Core Dump Files](#) on page 172.

You cannot remove an active configuration file (the configuration currently selected to boot the switch). To verify the configuration that you are currently using, issue the `show switch {detail}` command. If you attempt to remove the active configuration file, the switch displays a message similar to the following:

```
Error: Cannot remove current selected active configuration file.
```

When you delete a file from the switch, a message similar to the following appears:

```
Remove testpolicy.pol from switch? (y/n)
```

Enter `y` to remove the file from your system. Enter `n` to cancel the process and keep the file on your system.

Case-sensitive Filenames

Filenames are case-sensitive. In this example, you have a configuration file named `Test.cfg`. If you attempt to remove a file with the incorrect case, for example `test.cfg`, the system is unable to remove the file. The switch does not display an error message; however, the `ls` command continues to display the file `Test.cfg`. To remove the file, make sure you use the appropriate case.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local or remote file, remember the requirements listed above.

Internal Memory Card and Core Dump Files

When you delete a core dump file from the system, that file is unavailable.

When you configure the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to remove/delete the specified core dump file.

For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

You can use the `*` wildcard to delete core dump files from the internal memory card.

If you configure the switch to write core dump files to the internal memory card and attempt to download a new software image, you might have insufficient space to complete the image download. When this occurs, you must decide whether to continue the software download or move or delete the core dump files from the internal memory. For example, if your switch has an external memory card installed with space available, transfer the files to the external memory card. Transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

This command also replicates the action from the primary MSM/MM to the backup MSM/MM. For example, when you delete a file on the primary MSM/MM, the same file on the backup MSM/MM is deleted.

For the `memorycard` option, this command removes/deletes an existing file on the card, including core dump files if configured. See the section “[Internal Memory Card and Core Dump Files](#)” for information about core dump files.

You can use the `*` wildcard to delete all of a particular file type from the external memory card; currently running and in use files are not deleted.

Example

The following command removes the configuration file named *Activeb91.cfg* from the system:

```
rm Activeb91.cfg
```

The following command removes all of the core dump files stored on the internal memory card:

```
rm internal-memory *
```

If your switch has an external memory card installed, the following command removes the policy file named `test.pol` from the external memory card:

```
rm memorycard test.pol
```

If your switch has an external memory card installed, the following command removes all of the configuration files from the external memory card:

```
rm memorycard *.cfg
```

show cpu-monitoring

```
show cpu-monitoring {process <name>} {slot <slotid>}
```

Description

Displays the CPU utilization history of one or more processes.

Command Syntax

name	Specifies the name of the process.
slotid	Specifies the slot number of the MSM/MM module: <ul style="list-style-type: none"> • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

By default, CPU monitoring is enabled and occurs every 20 seconds. The default CPU threshold value is 60%.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different CPU and process information might be displayed.

When you issue the command without any parameters, the switch displays CPU utilization history for all of the processes running on the MSMs/MMs installed in your system.

Reading the Output

The `show cpu-monitoring` command is helpful for understanding the behavior of a process over an extended period of time. The following information appears in a tabular format:

- **Card**—The location (MSM A or MSM B).
- **Process**—The name of the process.
- **Range of time (5 seconds, 10 seconds, and so forth)**—The CPU utilization history of the process or the system. The CPU utilization history goes back only 1 hour.
- **Total User/System CPU Usage**—The amount of time recorded in seconds that the process spends occupying CPU resources. The values are cumulative meaning that the values are displayed as long as the system is running. You can use this information for debugging purposes to see where the process spends the most amount of time: user context or system context.

Example

The following command displays CPU utilization on the switch:

```
show cpu-monitoring
```

The following is sample truncated output from an 8800 switch:

CPU Utilization Statistics - Monitored every 5 seconds

Card	Process	5	10	30	1	5	30	1	Max	Total	
		secs	secs	secs	min	mins	mins	hour		User/System	
		util	util	util	util	util	util	util	util	CPU Usage	
		(%)	(%)	(%)	(%)	(%)	(%)	(%)	(%)	(secs)	
MSM-A	System	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.9		
MSM-B	System	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		
MSM-A	GNSS_cpuif	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_ctrlif	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_esmi	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_fabric	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_mac_10g	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_pbusmux	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_pktengine	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_pktif	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	GNSS_switch	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MSM-A	aaa	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.4	0.82	0.56
MSM-A	acl	0.0	0.0	0.0	0.0	0.0	0.0	0.0	7.5	0.37	0.33
MSM-A	bgp	0.0	0.0	0.0	0.0	0.0	0.0	0.0	5.2	0.27	0.42
MSM-A	cfgmgr	0.0	0.9	0.3	3.7	1.2	1.2	1.3	27.3	7.70	7.84
MSM-A	cli	0.0	0.0	0.0	48.3	9.6	2.5	2.1	48.3	0.51	0.37
MSM-A	devmgr	0.0	0.0	0.0	0.9	0.3	0.2	0.2	17.1	2.22	2.50

```

MSM-A  dirser      0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  9.5  0.0      0.0
MSM-A  dosprotect  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  3.8  0.20    0.26
MSM-A  ems         0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  12.2  1.1     1.16
MSM-A  epm         0.0  0.0  0.0  0.9  0.1  0.2  0.2  4.7  2.6     4.18
MSM-A  etmon       0.9  0.4  0.6  1.2  1.1  1.0  1.0  23.3  21.84   7.24
...

```

show heartbeat process

```
show heartbeat process {<name>}
```

Description

Displays the health of the NETGEAR 8800 processes.

Command Syntax

name	Specifies the name of the process.
------	------------------------------------

Default

N/A.

Usage Guidelines

The software monitors all of the XOS processes running on the switch. This process monitor creates and terminates XOS processes on demand (for example, when you log in or log out of the switch) and restarts processes if an abnormal termination occurs (for example, if your system crashes). The process monitor also ensures that only version-compatible processes and processes with proper licenses are started.

The `show heartbeat process` command is a resource for providing background system health information because you can view the health of the processes on the switch.

Use this command to monitor the health of the NETGEAR 8800 processes. The switch uses two algorithms to collect process health information: polling and reporting. Both polling and reporting measure the heartbeat of the process. Polling occurs when a HELLO message is sent and a HELLO_ACK message is received. The two counts are the same. Reporting occurs when a HELLO_ACK message is sent only. Therefore, no HELLO messages are sent and the HELLO count remains at zero.

The `show heartbeat process` command displays the following information in a tabular format:

- Card—The name of the module where the process is running.
- Process Name—The name of the process.
- Hello—The number of hello messages sent to the process.
- HelloAck—The number of hello acknowledgement messages received by the process manager.

- **Last Heartbeat Time**—The timestamp of the last health check received by the process manager. (Unknown specifies kernel modules and they do not participate in heartbeat monitoring.)

This status information may be useful for your technical support representative if you have a network problem.

You may find it useful to capture the process information under normal operating conditions to establish a baseline. By having a baseline, if you experience a problem, you and your technical support representative can more easily identify the problem.

Example

To display the health of all processes on your system, use the following command:

```
show heartbeat process
```

The following is sample output:

Card	Process Name	Hello	HelloAck	Last Heartbeat Time
MSM-A	aaa	0	180324	Wed Dec 10 15:06:04 2003
MSM-A	acl	36069	36069	Wed Dec 10 15:05:57 2003
MSM-A	bgp	0	180348	Wed Dec 10 15:06:05 2003
MSM-A	cfgmgr	72139	72139	Wed Dec 10 15:06:02 2003
MSM-A	cli	60116	60116	Wed Dec 10 15:06:03 2003
MSM-A	devmgr	0	180339	Wed Dec 10 15:06:03 2003
MSM-A	dirser	0	180324	Wed Dec 10 15:06:03 2003
MSM-A	ems	45087	45087	Wed Dec 10 15:06:03 2003
MSM-A	epm	0	0	Unknown
MSM-A	exacl	0	0	Unknown
....				

To display the health of the STP process on your system, use the following command:

```
show heartbeat process stp
```

The following is sample output:

Card	Process Name	Hello	HelloAck	Last Heartbeat Time
MSM-A	stp	34921	34921	Wed Dec 10 11:54:37 2003

show memory

```
show memory {slot [slotid | a | b]}
```

Description

Displays the current system memory information.

Syntax Description

slot a	Specifies the MSM module installed in slot A.
slot b	Specifies the MSM module installed in slot B.
slotid	Specifies slot number for the node in a stack. The value can be from 1 to 8.

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. When you keep simple daily records, you see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different memory information might be displayed.

You can also use the `show memory process <name> {slot <slotid>}` command to view the system memory and the memory used by the individual processes.

When you issue the command without any parameters, the switch displays information about all of the MSMs/MMs installed in your system.

Reading the Output

The `show memory` command displays the following information in a tabular format:

- System memory information (both total and free).
- Current memory used by the individual processes.

The current memory statistics for the individual process also includes the following:

- The module (MSM A or MSM B) and the slot number of the MSM.
- The name of the process.

In general, the `free` memory count for an MSM/MM decreases when one or more running processes experiences an increase in memory usage.

If you observe a continuous decrease in the `free` memory over an extended period of time, and you have not altered your switch configuration, please contact NETGEAR Technical Support.

Example

The following command displays current system memory information for the MSM installed in slot A of the switch:

```
show memory slot a
```

The following is sample output from this command:

System Memory Information

```
MSM-A    Total DRAM (KB): 524288
MSM-A    System      (KB): 45912
MSM-A    User        (KB): 102264
MSM-A    Free        (KB): 376112
```

Memory Utilization Statistics

Card Slot	Process Name	Memory (KB)
MSM-A 9	aaa	7772
MSM-A 9	acl	6716
MSM-A 9	bgp	16708
MSM-A 9	cfgmgr	3484
MSM-A 9	cli	33964
MSM-A 9	devmgr	3656
MSM-A 9	ems	5832
MSM-A 9	epm	8084
MSM-A 9	etmon	11356
MSM-A 9	exacl	13
MSM-A 9	exosmc	22
MSM-A 9	exosq	29
MSM-A 9	exsflow	8
MSM-A 9	exsnoop	15
MSM-A 9	exvlan	252
MSM-A 9	fdb	8760
MSM-A 9	hal	22624
MSM-A 9	mcmgr	13128
MSM-A 9	msgsrv	2972
MSM-A 9	netLogin	4564
MSM-A 9	netTools	4696
MSM-A 9	nettx	56
MSM-A 9	nodemgr	5388
MSM-A 9	ospf	12476
MSM-A 9	pim	10012
MSM-A 9	polMgr	3272
MSM-A 9	rip	10392

```

MSM-A  9    rtmgr          9748
MSM-A  9    snmpMaster     6400
MSM-A  9    snmpSubagent   8104
MSM-A  9    stp            6896
MSM-A  9    telnetd        3236
MSM-A  9    tftpd          3080
MSM-A  9    vlan           5816
MSM-A  9    vrrp           6584

```

The following command displays current system memory information for a stack, where slot 1 is the master and slot 6 is the backup:

```
Slot-1 stack.3 # show memory
```

System Memory Information

```

-----
Slot-1    Total DRAM (KB): 262144
Slot-1    System      (KB): 25476
Slot-1    User        (KB): 132256
Slot-1    Free        (KB): 104412
Slot-6    Total DRAM (KB): 262144
Slot-6    System      (KB): 25476
Slot-6    User        (KB): 122820
Slot-6    Free        (KB): 113848

```

Memory Utilization Statistics

```

-----
Card Slot Process Name      Memory (KB)
-----
Slot-1 1   aaa              2548
Slot-1 1   acl              2960
Slot-1 1   bgp              0
Slot-1 1   brm              2428
Slot-1 1   cfgmgr           3256
Slot-1 1   cli              16932
Slot-1 1   devmgr           2708
Slot-1 1   dirser           1916
Slot-1 1   dosprotect       1972
Slot-1 1   elsm             2592
Slot-1 1   ems              2764
Slot-1 1   epm              3092
Slot-1 1   etmon            16264
...
Slot-6 6   aaa              2440
Slot-6 6   acl              2872
Slot-6 6   bgp              0
Slot-6 6   brm              2396

```

```
Slot-6 6  cfmgr          2776
Slot-6 6  cli            16292
Slot-6 6  devmgr         2672
Slot-6 6  dirser         1836
Slot-6 6  dosprotect     1944
Slot-6 6  elsm           2564
Slot-6 6  ems            2744
Slot-6 6  epm            2976
Slot-6 6  etmon          10068
...
```

show memory process

`show memory process <name> {slot <slotid>}`

Description

Displays the current system memory and that of the specified process.

Command Syntax

name	Specifies the name of the process.
slotid	Specifies the slot number of the MSM/MM module: <ul style="list-style-type: none"> • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B. Specifies the slot number of the node in the stack topology. The value can be from 1 to 8.

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. When you keep simple daily records, you see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different memory information might be displayed.

You can also use the `show memory {slot [slotid | a | b]}` command to view the system memory and the memory used by the individual processes, even for all processes on all MSMs/MMs installed in the switch.

Reading the Output

The `show memory process` command displays the following information in a tabular format:

- System memory information (both total and free).
- Current memory used by the individual processes.

The current memory statistics for the individual process also includes the following:

- The module (MSM A or MSM B) and the slot number of the MSM/MM.
- The name of the process.

Example

The following command displays system memory and VRRP memory usage:

```
show memory process vrrp
```

The following is sample output:

```
System Memory Information
```

```
-----
```

```
MSM-A    Total (KB): 512508 KB
```

```
MSM-A    Free  (KB): 395796 KB
```

```
Memory Utilization Statistics
```

```
-----
```

```
Card Slot Process Name      Memory (KB)
```

```
-----
```

```
MSM-A  9    vrrp              6596
```

show process

```
show process {<name>} {detail} {description} {slot <slotid>}
```

Description

Displays the status of the NETGEAR 8800 processes.

Command Syntax

name	Specifies the name of the process.
detail	Specifies more detailed process information.
description	Describes the name of all of the processes or the specified process running on the switch.
slotid	Specifies the slot number of the MSM/MM module: <ul style="list-style-type: none"> • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

The NETGEAR 8800 process manager monitors all processes. The process manager also ensures that only version-compatible processes are started.

Using this command without the optional keywords displays summary process information. When you specify the `slot` keyword, summary information is displayed for that particular slot only.

The `show process` and `show process slot <slotid>` commands display the following information in a tabular format:

- Card—The name of the module where the process is running.
- Process Name—The name of the process.
- Version—The version number of the process. Options are:
 - Version number—A series of numbers that identify the version number of the process. This is helpful to ensure that you have version-compatible processes and if you experience a problem.
 - Not Started—The process has not been started. This can be caused by not having the appropriate license or for not starting the process.
- Restart—The number of times the process has been restarted. This number increments by one each time a process stops and restarts.
- State—The current state of the process. Options are:
 - No License—The process requires a license level that you do not have. For example, you have not upgraded to that license, or the license is not available for your platform.
 - Ready—The process is running.
 - Stopped—The process has been stopped.
- Start Time—The current start time of the process. Options are:
 - Day/Month/Date/Time/Year—The date and time the process began. When a process terminates and restarts, the start time is also updated.
 - Not Started—The process has not been started. This can be caused by not having the appropriate license or for not starting the process.

When you specify the `detail` keyword, more specific and detailed process information is displayed. The `show process detail` and `show process slot <slotid> detail` commands display the following information in a multi-tabular format:

- Detailed process information
- Memory usage configurations
- Recovery policies
- Process statistics
- Resource usage

This status information may be useful for your technical support representative if you have a network problem.

Depending on the software version running on your switch or your switch model, additional or different process information might be displayed.

You may find it useful to capture the process information under normal operating conditions to establish a baseline. By having a baseline, if you experience a problem, you and your technical support representative can more easily identify the problem.

Example

To display the processes on your system, use the following command:

```
show process
```

The following is sample output:

Card	Process Name	Version	Restart	State	Start Time
MSM-A	aaa	3.0.0.2	0	Ready	Sat Dec 6 10:54:24 2003
MSM-A	acl	3.0.0.2	0	Ready	Sat Dec 6 10:54:25 2003
MSM-A	bgp	3.0.0.2	0	Ready	Sat Dec 6 10:54:24 2003
MSM-A	cfgmgr	3.0.0.20	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	cli	3.0.0.21	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	devmgr	3.0.0.2	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	dirser	3.0.0.2	0	Ready	Sat Dec 6 10:54:21 2003
MSM-A	ems	3.0.0.2	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	epm	3.0.0.2	0	Ready	Sat Dec 6 10:54:21 2003
MSM-A	exacl	3.0.0.2	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	exosmc	3.0.0.2	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	exosq	3.0.0.2	0	Ready	Sat Dec 6 10:54:22 2003
MSM-A	exsnoop	3.0.0.2	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	exvlan	3.0.0.2	0	Ready	Sat Dec 6 10:54:22 2003
MSM-A	fdb	3.0.0.2	0	Ready	Sat Dec 6 10:54:24 2003
....					

The following example specifies the process `aaa` along with the `detail` keyword:

```
show process aaa detail
```

The following is sample output from this command:

Name	PID	Path	Type	Link	Date	Build By	Peer
aaa	284	./aaa	App	Thu Dec 4 13:23:07 PST 2003	release-manager	2	
3							
Virtual Router(s):							

Configuration:							
Start Priority	SchedPolicy	Stack	TTY	CoreSize	Heartbeat	StartSeq	

```

1      0      0      0      0      0      1      1
Memory Usage Configuration:
Memory(KB) Zones: Green Yellow Orange Red
-----
0      0      0      0      0

Recovery policies
-----
failover-reboot
-----

Statistics:
ConnexionLost  Timeout  Start  Restart  Kill  Register  Signal  Hello  Hello Ack
-----
0      0      0      0      0      1      0      0      173199

Memory Zone  Green  Yellow  Orange  Red
-----
Green      0      0      0      0
-----

Commands:
Start      Stop      Resume      Shutdown      Kill
-----
0      0      0      0      0
-----

Resource Usage:
UserTime SysTime  PageReclaim PageFault Up Since      Up Date  Up Time
-----
2.160000 0.560000    546      966  Sat Dec  6 10:54:24 2003 00/00/04 00:14:02
-----

Thread Name      Pid      Tid      Delay  Timeout Count
-----
tacThread        0      2051     10      0
radiusThread     0      1026     10      1
main             0      1024      2      1
-----

```

The following example describes the name of all of the processes running on the switch:

```
show process description
```

The following is sample output from this command:

```

Process Name      Description
-----
aaa               Authentication, Authorization, and Accounting Server
acl               Access Control List Manager
bgp               Border Gateway Protocol

```


brm	Bandwidth Resource Manager
cfgmgr	Configuration Manager
cli	Cli Manager
devmgr	Device Manager
dirser	Directory Services
dosprotect	Protection against Denial of Service attacks application
elsm	NETGEAR Link State Monitor
ems	Event Management System Server
epm	NETGEAR Process Manager
etmon	Traffic monitoring and sampling utility
exacl	Access Control List Module
exdhcpsnoop	DHCP snooping module
exdos	Detection of potential Denial of Service attacks module
exfib	Routing interface to manage missing routes in ASIC
exosipv6	IPv6 Custom Interface Module
exosmc	Multicast Forwarding Module
exosnvram	Interface to non-volatile RAM
exosq	EXOS Queue Module
exsflow	Sflow interface to gather sflow samples
exsnoop	IGMP/MLD Snooping Module
exvlan	Layer 2 configuration module
fdb	Forwarding Data Base Manager
hal	Hardware Abstraction Layer
ipSecurity	IP Security
isis	Intermediate System to Intermediate System Routing Protocol
lACP	Link Aggregation Control Protocol
lldp	802.1AB; Station and Media Access Control Connectivity Discover
mcmgr	Multicast Cache Manager
msdp	Multicast Source Discovery Protocol
msgsrv	Message Server
netLogin	Network Login includes MAC, Web-Based and 802.1X authentication
netTools	Network Toolset includes ping/tracert/bootprelay/dhcp/dns/sntp
nettx	Layer 2 forwarding engine module
nodemgr	Fault Tolerance Manager
ospf	Open Shortest Path First Routing Protocol
ospfv3	Open Shortest Path First Routing Protocol for IPv6
pim	Protocol Independent Multicast
poe	Power Over Ethernet Manager
polMgr	Policy Manager
rip	Routing Information Protocol
ripng	Routing Information Protocol for IPv6
rtmgr	Route Table Manager
snmpMaster	Simple Network Management Protocol - Master agent
snmpSubagent	Simple Network Management Protocol - Subagent
stp	Spanning Tree Protocol
telnetd	Telnet server
tftpd	Tftp server

thttpd	Web Server
upm	Universal Port Manager
vlan	VLAN Manager - L2 Switching application
vrrp	Virtual Router Redundancy Protocol (RFC 3768)
xmld	XML server

start process

```
start process <name> {msm <slot>}
```

Description

Starts the specified process on the switch. (Used to restart a process after it has been terminated.)

Syntax Description

name	Specifies the name of the process to start. You can start the following processes: <ul style="list-style-type: none"> • bgp • exsshd • lldp • netLogin • netTools • ospf • snmpMaster • snmpSubagent • telnetd • thttpd • tftpd • vrrp • xmld
slot	Specifies the MSM/MM where the process should be started. A specifies the MSM installed in slot A, and B specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

Use this command after you have stopped a process and you want to restart it. To stop a process, use the [terminate process](#) command.

You are unable to start a process that is already running. If you try to start a currently running process, an error message similar to the following appears:

```
Error: Process telnetd already exists!
```

Depending on the software version running on your switch and the type of switch you have, you can restart different or additional processes. To see which processes you can restart, enter `start process` followed by TAB. The switch displays a list of available processes.

To display the status of NETGEAR 8800 processes on the switch, including how many times a process has been restarted, use the `show process {<name>} {detail} {description} {slot <slotid>}` command.

You can also use the `start process` command when upgrading a software modular package. For more information, see the section “Upgrading a Modular Software Package” in Appendix B of the *NETGEAR 8800 User Manual*.

Note: After you stop a process, do not change the configuration on the switch until you start the process again. A new process loads the configuration that was saved prior to stopping the process. Changes made between a process termination and a process start are lost. Else, error messages can result when you start the new process.

Example

The following restarts the process *tftpd*:

```
start process tftpd
```

terminate process

```
terminate process <name> [forceful | graceful] {msm <slot>}
```

Description

Terminates the specified process on the switch.

Syntax Description

name	Specifies the name of the process to terminate. You can terminate the following processes: <ul style="list-style-type: none"> • bgp • exsshd • lldp • netLogin • netTools • ospf • snmpMaster • snmpSubagent • telnetd • thttpd • tftpd • vrrp • xmld
forceful	Specifies a forceful termination.
graceful	Specifies a graceful termination.
slot	For a modular chassis, specifies the MSM/MM where the process should be terminated. A specifies the MSM installed in slot A, and B specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

If recommended by NETGEAR Technical Support personnel, you can stop a running process.

The `forceful` option quickly terminates a process on demand. Unlike the `graceful` option, the process is immediately shutdown without any of the normal process cleanup. The status of the operation is displayed on the console. After a successful forceful termination of a process, a message similar to the following appears:

```
Forceful termination success for snmpMaster
```

The `graceful` option terminates the process by allowing it to close all opened connections, notify peers on the network, and other types of process cleanup. After this phase, the process is finally terminated. After a successful graceful termination of a process, a message similar to the following appears:

```
Successful graceful termination for snmpSubagent
```

Note: Do not terminate a process that was installed since the last reboot unless you have saved your configuration. If you have installed a software module and you terminate the newly installed process without saving your configuration, your module may not be loaded when you attempt to restart the process with the `start process` command.

To preserve a process's configuration during a terminate and (re)start cycle, save your switch configuration before terminating the process. Do not save the configuration or change the configuration during the process terminate and re(start) cycle. If you save the configuration after terminating a process, and before the process (re)starts, the configuration for that process is lost.

You can also use the `terminate process` command when upgrading a software modular package. For more information, see the section "Upgrading a Modular Software Package" in Appendix B of the *NETGEAR 8800 User Manual*.

Example

The following initiates a graceful termination of the process *tftpd*:

```
terminate process tftpd graceful
```

Commands for Configuring Slots and Ports on a Switch

5

This chapter describes commands related to:

- Enabling, disabling, and configuring individual ports
- Configuring port speed (Fast Ethernet ports only) and half- or full-duplex mode
- Creating link aggregation groups on multiple ports
- Displaying port statistics
- Configuring mirroring
- Configuring software-controlled redundant ports and Smart Redundancy

By default, all ports on the switch are enabled. After you configure the ports to your specific needs, you can select which ports are enabled or disabled.

Fast Ethernet ports can connect to either 10BASE-T or 100BASE-T networks. By default, the ports autonegotiate (automatically determine) the port speed. You can also configure each port for a particular speed (either 10 Mbps or 100 Mbps). In general Gigabit Ethernet ports with fiber interfaces are statically set, and their speed cannot be modified.

The switch comes configured to use autonegotiation to determine the port speed and duplex setting for each port. You can manually configure the duplex setting and the speed of 10/100 Mbps ports, and you can manually configure the duplex setting on gigabit Ethernet ports.

All ports on the switch (except gigabit Ethernet ports) can be configured for half-duplex or full-duplex operation. The ports are configured to autonegotiate the duplex setting, but you can manually configure the duplex setting for your specific needs.

Flow control is supported only on gigabit Ethernet ports. It is enabled or disabled as part of autonegotiation. If autonegotiation is set to off, flow control is disabled. When autonegotiation is turned on, flow control is enabled. (See the *NETGEAR 8800 User Manual* for more detailed information on flow control on NETGEAR devices.)

Link aggregation, or load sharing, with NETGEAR switches allows you to increase bandwidth and resilience between switches by using a group of ports to carry traffic in parallel between switches. The sharing algorithm allows the switch to use multiple ports as a single logical port. For example, VLANs see the link aggregation group (LAG) as a single logical port. The algorithm also guarantees packet sequencing between clients.

NETGEAR 8800 software supports two broad categories of load sharing, or link aggregation: static load sharing and dynamic load sharing.

If a port in a link aggregation group fails, traffic is redistributed to the remaining ports in the LAG. If the failed port becomes active again, traffic is redistributed to include that port.

You can view port status on the switch using the `show ports` commands. These commands, when used with specific keywords and parameters, allow you to view various issues such as collision statistics, link speed, flow control, and packet size. These port information displays show real-time statistics, or you can configure the display to show a snapshot of real-time statistics.

You can configure WAN PHY OAM on those interfaces that connect 10G Ethernet ports to the SONET/SDH network.

Commands that require you to enter one or more port numbers use the parameter `<port_list>` in the syntax. On the 8800, a `<port_list>` can be a list of slots and ports. For a detailed explanation of port specification, see [Port Numbering](#) in [Chapter 1, "Command Reference Overview."](#)

clear counters ports

```
clear counters ports
```

Description

Clears the counters associated with the ports.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Note: If you use the `clear counters` command with no keyword, the system clears the counters for all applications.

This command clears the counters for the ports, including the following:

- Statistics
- Transmit errors
- Receive errors
- Collisions
- Packets

Example

The following command clears the counters on all ports:

```
clear counters ports
```

clear lacp counters

```
clear lacp counters
```

Description

Clears the counters associated with Link Aggregations Control Protocol (LACP).

Syntax Description

This command has no parameters or variables.

Default

N/A

Usage Guidelines

This command clears the following counters for LACP; it sets these counters back to 0 for every LACP port on the device:

- LACP PDUs dropped on non_LACP ports
- Stats
 - Rx - Accepted
 - Rx - Dropped due to error in verifying PDU
 - Rx - Dropped due to LACP not being up on this port
 - Rx - Dropped due to matching own MAC
 - Tx - Sent Successfully
 - Tx - Transmit error

Example

The following command clears the LACP counters on all ports:

```
clear lacp counters
```

clear slot

```
clear slot <slot>
```

Description

Clears a slot of a previously assigned module type.

Syntax Description

slot	Specifies the slot number.
------	----------------------------

Default

N/A.

Usage Guidelines

All configuration information related to the slot and the ports on the module is erased. If a module is present when you issue this command, the module is reset to default settings.

If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state (where the inserted module does not match the configured slot), and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type. Use the `enable mirroring to port tagged` command to configure the slot.

Example

The following command clears slot 2 of a previously assigned module type:

```
clear slot 2
```

The following command clears slot 4 of a previously assigned module type in a stack:

```
clear slot 4
```

configure ip-mtu vlan

```
configure ip-mtu <mtu> vlan <vlan_name>
```

Description

Sets the maximum transmission unit (MTU) for the VLAN.

Syntax Description

mtu	Specifies the IP maximum transmission unit (MTU) value. Range is from 1500 to 9194.
vlan_name	Specifies a VLAN name.

Default

The default IP MTU size is 1500.

Usage Guidelines

The 8800 switches support IP fragmentation and path MTU discovery.

Use this command to enable jumbo frame support or for IP fragmentation with jumbo frames. Jumbo frames are Ethernet frames that are larger than 1522 bytes, including 4 bytes used for CRC. Both endstations involved in the transfer must be capable of supporting jumbo frames. The switch does not perform IP fragmentation or participate in MTU negotiation on behalf of devices that do not support jumbo frames.

When enabling jumbo frames and setting the MTU size for the VLAN, keep in mind that some network interface cards (NICs) have a configured maximum MTU size that does not include the additional 4 bytes of CRC included in a jumbo frame configuration. Ensure that the NIC maximum MTU is at or below the maximum MTU size configured on the switch. Frames that are larger than the MTU size configured on the switch are dropped at the ingress port.

If you use IP fragmentation with jumbo frames and you want to set the MTU size greater than 1500, all ports in the VLAN must have jumbo frames enabled.

Example

The following command sets the MTU size to 2000 for VLAN *sales*:

```
configure ip-mtu 2000 vlan sales
```

configure jumbo-frame-size

```
configure jumbo-frame-size <framesize>
```

Description

Sets the maximum jumbo frame size for the switch.

Syntax Description

framesize	Specifies a maximum transmission unit (MTU) size for a jumbo frame. The range is 1523 to 9216; the default is 9216.
-----------	---

Default

Jumbo frames are disabled by default. The default size setting is 9216.

Usage Guidelines

Jumbo frames are used between endstations that support larger frame sizes for more efficient transfers of bulk data. Both endstations involved in the transfer must be capable of supporting jumbo frames.

The `framesize` keyword describes the maximum jumbo frame size “on the wire,” and includes 4 bytes of cyclic redundancy check (CRC) plus another 4 bytes if 802.1Q tagging is being used.

To enable jumbo frame support, you must configure the maximum transmission unit (MTU) size of a jumbo frame that will be allowed by the switch.

Note: NETGEAR recommends that you set the MTU size so that fragmentation does *not* occur.

Some network interface cards (NICs) have a configured maximum MTU size that does not include the additional 4 bytes of CRC. Ensure that the NIC maximum MTU size is at or below the maximum MTU size configured on the switch. Frames that are larger than the MTU size configured on the switch are dropped at the ingress port.

Example

The following command configures the jumbo frame size to 5500:

```
configure jumbo-frame-size 5500
```

configure lacp member-port priority

```
configure lacp member-port <port> priority <port_priority>
```

Description

Configures the member port of an LACP to ensure the order that ports are added to the aggregator. The lower value you configure for the port's priority, the higher priority that port has to be added to the aggregator.

Syntax Description

port	Specifies the LACP member port that you are specifying the priority for.
port_priority	Specifies the priority you are applying to this member port to be assigned to the LACP aggregator. The range is from 1 to 65535; the default is 0. The lower configured value has higher priority to be added to the aggregator.

Default

The default priority is 0.

Usage Guidelines

The port must be added to the LAG prior to configuring it for LACP. The default value is 0, or highest priority.

You can configure the port priority to ensure the order in which LAG ports join the aggregator. If you do not configure this parameter, the lowest numbered ports in the LAG are the first to be added to the aggregator; if there are additional ports configured for that LAG, they are put in standby mode.

Use this command to override the default behavior and ensure the order in which LAG ports are selected. Also, if more than one port is *configured* with the same priority, the lowest numbered port joins the aggregator.

Example

The following command sets the port priority for the LAG port 5:1 to be 55 (which will probably put that port in standby initially):

```
configure lacp member-port 5:1 priority 55
```

configure mirror add ports anomaly

```
configure mirror add ports <port list> anomaly
```

Description

Mirrors detected anomaly traffic to the mirror port.

Syntax Description

port list	Specifies the list of ports.
-----------	------------------------------

Default

N/A.

Usage Guidelines

The command mirrors detected anomaly traffic to the mirror port. You must enable a mirror port and enable protocol anomaly protection on the slot that has the port to be monitored before using this command. After configuration, only detected anomaly traffic from these ports are dropped or mirrored to the mirror port, and legitimate traffic is not affected.

This command takes effect after enabling anomaly-protection.

configure mirroring add

```
configure mirroring add [vlan <name> {port <port>}] port <port> {vlan <name>}] {ingress | egress | ingress-and-egress}
```

Description

Adds a particular mirroring filter definition on the switch.

Syntax Description

vlan	Specifies a VLAN.
name	Specifies a VLAN name.
port	Specifies a port or slot and port.
port	Specifies particular ports or slots and ports.
ingress	Specifies packets be mirrored as they are received on a port.

egress	Specifies packets be mirrored as they are sent from a port.
ingress-and-egress	Specifies all forwarded packets be mirrored. This is the default setting on the NETGEAR 8800 series switches for port-based mirroring.

Default

N/A.

Usage Guidelines

You must enable port-mirroring using the `enable mirroring to port` command before you can configure the mirroring filter definitions.

Port mirroring configures the switch to copy all traffic associated with one or more ports to a monitor port on the switch. The switch uses a traffic filter that copies a group of traffic to the monitor port.

Up to 16 mirroring filters and one monitor port can be configured on the switch. Frames that contain errors are not mirrored.

Guidelines for configuring mirroring

This section summarizes the guidelines for configuring mirroring:

- When you disable mirroring, all the filters are unconfigured.
- You cannot mirror the monitor port.
- The mirroring configuration is removed when you:
 - Delete a VLAN (for all VLAN-based filters).
 - Delete a port from a VLAN (for all VLAN-, port-based filters).
 - Unconfigure a slot (for all port-based filters on that slot).
- Any mirrored port can also be enabled for load sharing (or link aggregation); however, each individual port of the load-sharing group must be explicitly configured for mirroring.
- The mirroring filters are not confined to a single module; they can have ports that span multiple modules.
- You cannot use the management port at all in mirroring configurations.
- You cannot run ELSM and mirroring on the same port. If you attempt to enable mirroring on a port that is already enabled for ELSM, the switch returns a message similar to the following:

```
Error: Port mirroring cannot be enabled on an ELSM enabled port.
```

The traffic filter can be defined based on one of the following criteria:

- **Physical port**—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port(s). You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.
 - Egress—Mirrors traffic sent from the port.

- Ingress and egress—Mirrors traffic either received at the port or sent from the port.
(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).
- **VLAN**—All data to a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- **Virtual port**—All data specific to a VLAN on a specific port is copied to the monitor port.
- EXOS supports up to 16 mirror filters where each filter can be a port, a VLAN, or a port + VLAN.
- EXOS supports up to 16 monitor ports for one-to-many mirroring.
- Only traffic *ingressing* a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- When routing between VLANs, ingress mirrored traffic is presented to the monitor port as *modified* for routing. This is the default behavior and the behavior when you use the command, `configure mirroring mode standard`. When you use the command, `configure mirroring mode enhanced`, ingress traffic is mirrored as it is received (on the wire).
- When using standard mode mirroring, a packet which matches both an ingress mirroring filter and an egress mirroring filter can only be ingress mirrored. The behavior depends on the location of the ingress port, egress port and monitor port within the switch as well as the type of module on which the packet ingresses. The behavior also varies depending on the configuration of daisy chain or ring mode stacking. When using enhanced mode mirroring, two packets are mirrored when a packet encounters both an ingress and egress mirroring filter.
- When traffic is modified by hardware on egress, egress mirrored packets may not be transmitted out of the monitor port as they egressed the port containing the egress mirroring filter. For example, an egress mirrored packet that undergoes VLAN translation is mirrored with the untranslated VLAN ID. In addition, IP multicast packets which are egress mirrored contain the source MAC address and VLAN ID of the unmodified packet.
- You cannot include the monitor port for a NETGEAR 8800 series switch in a load-sharing group.
- Tagged and untagged traffic is mirrored slightly differently depending on the module that the mirrored port and the monitor port are on:
 - With a monitor port or ports on an 8800 switch, the mirrored packet is tagged *only* if the ingress packet is tagged (regardless of what module the ingressing port is on). If the packet arrived at the ingress port as untagged, the packet egress the monitor port(s) as untagged.
 - With the 8800 series switches, you may see a packet mirrored twice. This occurs only if both the ingress mirrored port and the monitor port or ports are on the same one-half of the module *and* the egress mirrored port is either on the other one-half of that module or on another module.
 - On NETGEAR 8800 series switches, when traffic is modified by hardware on egress, egress mirrored packets may not be transmitted out of the monitor port as they egressed the port containing the egress mirroring filter. For example, an egress mirrored packet that undergoes VLAN translation is mirrored with the untranslated VLAN ID. In addition,

IP multicast packets which are egress mirrored contain the source MAC address and VLAN ID of the unmodified packet.

- Enhanced mirroring mode must be configured if you are going to configure a remote mirroring tag. Enhanced mirroring mode is configured using the following command:

```
configure mirroring mode enhanced
```

- The configuration of `remote-tag` does not require the creation of a VLAN with the same tag; on these platforms the existence of a VLAN with the same tag as a configured `remote-tag` is prevented. This combination is allowed so that an intermediate remote mirroring switch can configure remote mirroring using the same remote mirroring tag as other source switches in the network. Make sure that VLANs meant to carry normal user traffic are not configured with a tag used for remote mirroring.
- When a VLAN is created with `remote-tag`, that tag is locked and a normal VLAN cannot have that tag. The tag is unique across the switch. Similarly if you try to create a `remote-tag` VLAN where `remote-tag` already exists in a normal VLAN as a VLAN tag, you cannot use that tag and the VLAN creation fails.

Example

The following example sends all traffic coming into a NETGEAR 8800 series switch on slot 3, port 2 to the mirror port:

```
configure mirroring add port 3:2 ingress
```

configure mirroring delete

```
configure mirroring delete [all | port <port> {vlan <name>} | vlan <name> {port <port>}]
```

Description

Deletes a particular mirroring filter definition on the switch.

Syntax Description

all	Specifies all mirroring filter definitions.
port	Specifies a port or a slot and port.
port	Specifies particular ports or slots and ports.
vlan	Specifies a VLAN.
name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

On the switch, <port_list> must be a slot and port in the form <slot>:<port>. For a detailed explanation of port specification, see [Port Numbering](#) in Chapter 1, “Command Reference Overview.”

Example

The following example deletes the mirroring filter on an 8800 series switch defined for slot 7, port 1:

```
configure mirroring delete ports 7:1
```

configure mirroring mode

```
configure mirroring mode [enhanced | standard]
```

Description

Configures the mirroring mode which affects mirroring behavior globally in the system.

Syntax Description

enhanced	Specifies the mirroring mode that provides enhanced mirroring operation.
standard	Specifies the standard mirroring mode that is required when the mirroring configuration involves ports or VLANs on 8800 series modules.

Default

Standard mode is the default.

Usage Guidelines

When the mirroring configuration involves only ports or VLANs on 8800 series switches, enhanced mode is recommended since it provides enhanced behavior. (For more information, see Chapter 5 in the *NETGEAR 8800 User Manual*.)

Example

The following example configures a system to use enhanced mirroring mode:

```
configure mirroring mode enhanced
```

configure ports auto off

```
configure ports <port_list> auto off speed <speed> duplex [half | full]
```


Description

Manually configures port speed and duplex setting configuration on one or more ports on a switch.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
speed	Specifies the port speed as either 10, 100, 1000 (1 Gigabit), or 10000 (10 Gigabit) Mbps ports.
duplex [half]	Specifies half duplex; transmitting and receiving data one direction at a time.
duplex [full]	Specifies full duplex; transmitting and receiving data at the same time.

Default

Auto on for 1G ports.

Usage Guidelines

You can manually configure the duplex setting and the speed on 10/100 and 10/100/1000 Mbps and fiber SFP gigabit Ethernet ports.

In general, SFP gigabit Ethernet ports are statically set to 1 Gbps, and their speed cannot be modified. However, there are GBICs supported by NETGEAR that can have a configured speed:

- 100 FX GBICs, which must have their speed configured to 100 Mbps
- 100FX/1000LX GBICs, which can be configured at either speed
- SFP+ optics, must have their speed configured to 10G auto off

In certain interoperability situations, it is necessary to turn autonegotiation off on a fiber gigabit Ethernet port. Even though a gigabit Ethernet port runs only at full duplex and gigabit speeds, the command that turns off autonegotiation must still include the duplex setting.

Gigabit Ethernet ports support flow control only when autonegotiation is turned on. When autonegotiation is turned off, flow control is not supported. (See the *NETGEAR 8800 User Manual* for more detailed information on flow control on NETGEAR devices.)

Example

The following example turns autonegotiation off for slot 2, port 1 at full duplex:

```
configure ports 2:1 auto off speed 100 duplex full
```

The following example turns autonegotiation off for port 2 with copper medium and a port speed of 100 Mbps at full duplex:

```
configure ports 2 medium copper auto off speed 100 duplex full
```

configure ports auto on

```
configure ports <port_list> auto on [{speed <speed>} {duplex [half | full]}] | [{duplex [half | full]} {speed <speed>}]
```

Description

Enables autonegotiation for the particular port type.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
speed	Specifies the port speed as either 10, 100, 1000 (1 Gigabit), or 10000 (10 Gigabit) Mbps ports.
duplex [half]	Specifies half duplex; transmitting and receiving data one direction at a time.
duplex [full]	Specifies full duplex; transmitting and receiving data at the same time.

Default

Auto on for 1 Gbps ports.

Auto off for 10 Gbps ports.

Usage Guidelines

The type of ports enabled for autonegotiation are 802.3u for 10/100 Mbps ports or 802.3z for gigabit Ethernet ports.

Flow control on gigabit Ethernet ports is enabled or disabled as part of autonegotiation. If autonegotiation is set to off, flow control is disabled. When autonegotiation is turned on, flow control is enabled. (See the *NETGEAR 8800 User Manual* for more detailed information on flow control on NETGEAR devices.)

Example

The following command configures the switch to autonegotiate for slot 1, ports 2 and 4:

```
configure ports 1:2, 1:4 auto on
```

The following command configures the switch to autonegotiate for port 2, with copper medium at a port speed of 100 Mbps at full duplex:

```
configure ports 2 medium copper auto on speed 100 duplex full
```

configure ports auto-polarity

```
configure ports [<port_list> | all] auto-polarity [off | on]
```

Description

Configures the autopolarity detection feature on the specified Ethernet ports.

Syntax Description

port_list	Specifies one or more ports on the switch.
all	Specifies all of the ports on the switch.
off	Disables the autopolarity detection feature on the specified ports.
on	Enables the autopolarity detection feature on the specified ports.

Default

Enabled.

Usage Guidelines

This feature applies to only the 10/100/1000 BASE-T ports on the switch.

Use the `all` keyword to enable or disable the autopolarity detection feature on all of the Ethernet ports on 8800 series switches.

When autopolarity is disabled on one or more Ethernet ports, you can verify that status by using the command:

```
show ports information detail
```

Example

The following command disables the autopolarity detection feature on ports 5 to 7 on the NETGEAR 8800 switch:

```
configure ports 5-7 auto-polarity off
```

configure ports display-string

```
configure ports <port_list> display-string <string>
```

Description

Configures a user-defined string for a port or group of ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
string	Specifies a user-defined display string.

Default

N/A.

Usage Guidelines

The display string can be up to 15 characters. Display strings do not need to be unique for each port—you can assign the same string to multiple ports. For example, you could give all the ports that connected to a particular department a common display string.

The string is displayed in certain commands such as the `show ports information` command.

Note: Do not use a port number as a display string. For example, do not assign the display string “2” to port 2.

Example

The following command configures the user-defined string *corporate* for port 1 on a stand-alone switch:

```
configure ports 1 display-string corporate
```

The following command configures the user-defined string *corporate* for ports 3, 4, and 5 on slot 1:

```
configure ports 1:3-5 display-string corporate
```

configure ports redundant

```
configure ports <primaryPort> redundant <secondaryPort> {link [on | off]}
```

Description

Configures a software-controlled redundant port.

Syntax Description

primaryPort	Specifies one primary port or slot and port.
redundantPort <secondaryPort>	Specifies one or redundant port or slot and port.
link	Specifies state of link: <ul style="list-style-type: none"> • on—Specifies keeping the redundant port active, but block traffic • off—Specifies forcing the link down on the redundant port

Note: The default value is off.

Default

N/A.

Usage Guidelines

The first port specifies the primary port. The second port specifies the redundant port.

A software-controlled redundant port is configured to back up a specified primary port; both ports are on the same device. The redundant port tracks the link state of the associated primary port, and if the link on the primary port fails, the redundant port establishes a link and becomes active. You can back up a specified Ethernet port with a redundant, dedicated Ethernet port.

You configure the redundant link to be always physically up but logically blocked or to be always physically down. The default is off, or the redundant link is down.

The following criteria must be considered when configuring a software-controlled redundant port:

- You can configure only one redundant port for each primary port.
- You cannot have any Layer 2 protocols configured on any of the VLANs that are present on the ports. (You will see an error message if you attempt to configure software redundant ports on ports with VLANs running Layer 2 protocols.)
- The primary and redundant port must have identical VLAN memberships.
- The master port is the only port of a load-sharing group that can be configured as either a primary or redundant port. (The entire trunk must go down before the software-controlled redundant port takes effect.)
- Only one side of the link should be configured as redundant.

Example

The following command configures a software-controlled redundant port:

```
configure ports 1:3 redundant 2:3
```

configure sharing add ports

```
configure sharing <port> add ports <port_list>
```

Description

Adds ports to a load-sharing, or link aggregation, group. By using link aggregation, you use multiple ports as a single logical port. Link aggregation also provides redundancy because traffic is redistributed to the remaining ports in the link aggregation group (LAG) if one port in the group goes down.

Syntax Description

port	Specifies the logical port for a load-sharing group or link aggregation group (LAG). This number also functions as the LAG Group ID.
port_list	Specifies one or more ports or slots and ports to be grouped in the LAG.

Default

N/A.

Usage Guidelines

Use this command to dynamically add ports to a load-sharing group, or link aggregation group (LAG).

Note: You must create a LAG (or load-sharing group) before you can configure the LAG. To create a LAG, see `enable sharing <port> grouping <port_list> {algorithm [address-based {L2 | L3 | L3_L4 | custom}]}` `{lacp | health-check}`.

vMAN ports can belong to LAGs. If any port in the LAG is enabled for vMAN, all ports in the group are automatically enabled to handle jumbo size frames. Also, vMAN is automatically enabled on all ports of the untagged LAG.

To verify your configuration, use the `show ports sharing` command.

Note: All ports that are designated for the LAG must be removed from all VLANs prior to configuring the LAG.

The following guidelines apply to link aggregation on the NETGEAR 8800 series switch:

- A static LAG can include a maximum of 8 ports.
- An LACP LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- A Health Check LAG can include a maximum of 8 ports.
- Any broadcast, multicast, or unknown unicast packet is transmitted on a single port in the LAG.

Note: You cannot configure port-based load sharing algorithm on the 8800 series switch; you configure only address-based load-sharing algorithms.

- The available address-based parameters on the 8800 series switch are L2 for Layer 2 and L3 for Layer 3. If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting.

Example

The following example adds port 3:13 to the LAG with the logical port 3:9 on the switch:

```
configure sharing 3:9 add port 3:13
```

configure sharing address-based custom

```
configure sharing address-based custom [ipv4 [L3-and-L4 | source-only | destination-only |
source-and-destination] | hash-algorithm [xor | crc-16]]
```

Description

On NETGEAR 8800 series switches, this command configures the part of the packet examined by the switch when selecting the egress port for transmitting link aggregation, or load-sharing, data.

Syntax Description

ipv4	Specifies that the user configuration applies to IPv4 traffic.
L3-and-L4	Indicates that the switch should examine the IP source and destination address and the TCP or UDP source and destination port number.
source-only	Indicates that the switch should examine the IP source address only.
destination-only	Indicates that the switch should examine the IP destination address only.
source-and-destination	Indicates that the switch should examine the IP source and destination address.
xor	Use exclusive-OR for load sharing hash computation.
crc-16	Use CRC-16 for load sharing hash computation.

Default

Algorithm: L3-and-L4

Hash algorithm: xor

Usage Guidelines

This command specifies the part of the packet header that the switch examines to select the egress port for address-based load-sharing trunks. The address-based load-sharing setting is global and applies to all load-sharing trunks, or LAGs, that are address-based and configured with a custom algorithm. You change this setting by issuing the command again with a different option.

The addressing information examined is based on the packet protocol as follows:

- IPv4 packets—Uses the source and destination IPv4 addresses and Layer 4 port numbers as specified with this command.

- IPv6 packets—Uses the source and destination IPv6 addresses and Layer 4 port numbers.
- MPLS packets—Uses the top, second, and reserved labels and the source and destination IP addresses.
- Non-IP Layer 2—Uses the VLAN ID, the source and destination MAC addresses, and the ethertype.

The `xor` hash algorithm guarantees that the same egress port is selected for traffic distribution based on a pair of IP addresses, Layer 4 ports, or both, regardless of which is the source and which is the destination.

For IP-in-IP and GRE tunneled packets, the switch examines the *inner header* to determine the egress port.

To verify your configuration, use the `show ports sharing` command.

Example

The following example configures the switch to examine the source IP address:

```
configure sharing address-based custom ipv4 source-only
```

configure sharing delete ports

```
configure sharing <port> delete ports <port_list>
```

Description

Deletes ports from a link aggregation, or load-sharing, group.

Syntax Description

<code>port</code>	Specifies the logical port for a load-sharing group or a link aggregation group (LAG). This number also functions as the LAG Group ID.
<code>port_list</code>	Specifies one or more ports or slots and ports to be grouped in the LAG.

Default

N/A.

Usage Guidelines

Use this command to dynamically delete ports from a load-sharing group, or link aggregation group (LAG). This command applies to static and dynamic link aggregation.

Example

The following example deletes port 3:12 from the LAG with the logical port, or LAG Group ID, 3:9:


```
configure sharing 3:9 delete port 3:12
```

configure sharing health-check member-port add tcp-tracking

```
configure sharing health-check member-port <port> add tcp-tracking <IP Address> {tcp-port
<TCP Port> frequency <sec> misses <count>}
```

Description

Configures monitoring for each member port of a health check LAG.

Syntax Description

port	Specifies the member port.
IP Address	Specifies the IP address to monitor.
TCP Port	Specifies the TCP port to watch. The default is port 80.
sec	Specifies the frequency in seconds at which tracking takes place. The default is 10 seconds.
count	Specifies the number of misses before a connection loss is reported. The default is 3 misses.

Default

N/A.

Usage Guidelines

To configure a health check LAG, you first create a health check type of LAG using the [enable sharing grouping](#) command. Then use this command to configure the monitoring for each member port. You can configure each member port to track a particular IP address, but only one IP address per member port.

To display the monitoring configuration for a health check LAG, use the [show sharing health-check](#) command.

To display the link aggregation configured on a switch, use the [show ports sharing](#) command.

Example

The following commands configure four different member ports:

```
# configure sharing health-check member-port 10 add track-tcp 10.1.1.1 tcp-port 23
# configure sharing health-check member-port 11 add track-tcp 10.1.1.2 tcp-port 23
# configure sharing health-check member-port 12 add track-tcp 10.1.1.3
# configure sharing health-check member-port 13 add track-tcp 10.1.1.4
```

When the TCP port, seconds, or counts are not specified, they default to the values described in the Syntax Description.

configure sharing health-check member-port delete tcp-tracking

```
configure sharing health-check member-port <port> delete tcp-tracking <IP Address> {tcp-port <TCP Port>}
```

Description

Unconfigures monitoring for each member port of a health check LAG.

Syntax Description

port	Specifies the member port.
IP Address	Specifies the IP address.
TCP Port	Specifies the TCP port.

Default

N/A.

Usage Guidelines

Use this command to remove the monitoring configuration on the ports of a health check link aggregation group. Each port must be unconfigured separately, specifying the IP address and TCP port.

Example

The following command removes the configuration setting on port 12 that monitors IP address 10.1.1.3:

```
# configure sharing health-check member-port 12 delete track-tcp 10.1.1.3
```

configure sharing health-check member-port tcp-tracking

```
configure sharing health-check member-port <port> [disable | enable] tcp-tracking
```

Description

Enables or disables configured monitoring on a member port of a health check LAG.

Syntax Description

port	Specifies the member port.
------	----------------------------

Default

N/A.

Usage Guidelines

This disables/enables monitoring on a particular member port. When monitoring is disabled, the member port is added back to the LAG if it has not already been added. This allows a member port to be added back to LAG even though connectivity to the host is down.

Example

The following command disables port 12:

```
configure sharing health-check member-port 12 disable tcp-tracking
```

configure sharing lacp activity-mode

```
configure sharing <port> lacp activity-mode [active | passive]
```

Description

Configures the whether the switch sends LACPDUs periodically (active) or only in response to LACPDUs sent from the partner on the link (passive).

Syntax Description

port	Specifies the master logical port for the LAG you are setting the activity mode for.
active	Enter this value to have the switch periodically sent LACPDUs for this LAG.
passive	Enter this value to have the switch only respond to LACPDUs for this LAG.

Default

Active.

Usage Guidelines

You must enable sharing and create the LAG prior to assigning this LACP activity mode.

Note: One side of the link must be in active mode in order to pass traffic. If you configure your side in the passive mode, ensure that the partner link is in LACP active mode.

To verify the LACP activity mode, use the `show lacp lag <group-id> detail` command.

If you attempt to enter a port number that is different that a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```

Example

The following command changes the activity mode to passive for the specified LAG group ID:

```
configure sharing 5:1 lacp activity-mode passive
```

configure sharing lacp defaulted-state-action

```
configure sharing <port> lacp defaulted-state-action [add | delete]
```

Description

Configures whether a defaulted LAG port is removed from the aggregator.

Syntax Description

port	Specifies the master logical port for the LAG you are setting the default action for.
add	Enter this value to have the switch add defaulted ports to the aggregator for this LAG.
delete	Enter this value to have the switch delete defaulted ports from the aggregator for this LAG.

Default

Delete.

Usage Guidelines

You must enable sharing and create the LAG prior to configuring this LACP parameter.

You can configure whether you want a defaulted LAG port removed from the aggregator or added back into the aggregator. If you configure the LAG to remove ports that move into the default state, those ports are removed from the aggregator and the port state is set to unselected.

If you configure the LAG to add the defaulted port into the aggregator, the system takes inventory of the number of ports currently in the aggregator:

- If there are fewer ports in the aggregator than the maximum number allowed, the system adds the defaulted port to the aggregator (port set to selected and collecting-distributing).
- If the aggregator has the maximum ports, the system adds the defaulted port to the standby list (port set to standby).

Note: If the defaulted port is assigned to standby, that port automatically has a lower priority than any other port in the LAG (including those already in standby).

To verify the LACP default action, use the `show lacp lag <group-id> detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```

Note: To force the LACP trunk to behave like a static sharing trunk, use this command to add ports to the aggregator.

Example

The following command deletes defaulted ports from the aggregator for the specified LAG group ID:

```
configure sharing 5:1 lacp defaulted-state-action delete
```

configure sharing lacp system-priority

```
configure sharing <port> lacp system-priority <priority>
```

Description

Configures the system priority used by LACP for each LAG to establish the station on which end assumes priority in determining those LAG ports moved to the collecting/distributing state of the protocol. That end of the LAG with the lowest system priority is the one that assumes control of the determination. This is optional; if you do not configure this parameter, LACP uses system MAC values to determine priority. If you choose to configure this parameter, enter a value between 1 and 65535.

Syntax Description

port	Specifies the master logical port for the LAG you are setting the priority for.
priority	Enter the value you want for the priority of the system for the LACP. The range is 1 to 65535; there is no default.

Default

N/A.

Usage Guidelines

The LACP uses the system MAC values to assign priority to one of the systems, and that system then determines which LAG ports move into the collecting/distributing state and exchange traffic. That end of the LAG with the lowest system priority is the one that assumes control of the determination. If you wish to override the default LACP system priority for a specific LAG, use this command to assign that LAG a specific LACP priority. Enter a value between 1 and 65535.

You must enable sharing and create the LAG prior to assigning this LACP priority.

To verify the LACP system priority, use the `show lacp` command.

To change the system priority you previously assigned to a specific LAG, issue the `configure sharing lacp system-priority` using the new priority you want. To remove the assigned system priority entirely and use the LACP priorities, issue the `configure sharing lacp system-priority` using a value of 0.

Example

The following command assigns LAG 10 an LACP system priority of 3:

```
configure sharing 10 lacp system-priority 3
```

configure sharing lacp timeout

```
configure sharing <port> lacp timeout [long | short]
```

Description

Configures the timeout used by each LAG to stop transmitting once LACPDUs are no longer received from the partner link. You can configure this timeout value to be either 90 seconds, long, or 3 seconds, short.

Syntax Description

port	Specifies the master logical port for the LAG you are setting the timeout value for.
long	Enter this value to use 90 seconds as the timeout value.
short	Enter this value to use 3 seconds as the timeout value.

Default

Long.

Usage Guidelines

You must enable sharing and create the LAG prior to assigning this LACP timeout value.

To verify the LACP timeout value, use the `show lacp lag <group-id> detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```

Example

The following command changes the timeout value for the specified LAG group ID to short:

```
configure sharing 5:1 lacp timeout short
```

configure slot module

```
configure slot <slot> module <module_type>
```

Description

Configures a slot for a particular I/O module card.

On a stack, this command configures a slot for a particular type of node.

Syntax Description

slot	Specifies the slot number.
module_type	Specifies the type of module or node for which the slot should be configured. The list of modules you can enter will vary depending on the type of switch and version of the NETGEAR 8800 you are running. Certain modules are supported only with specific releases.

Default

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated.

Usage Guidelines

The command displays different module parameters depending on the type of switch you are configuring and the version of NETGEAR 8800 running on the switch.

You can also preconfigure the slot before inserting the module card. This allows you to begin configuring the module and ports before installing the card in the chassis.

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated. If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state, and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type.

Upon powering up the chassis, or when an I/O module is hot-swapped, the NETGEAR 8800 automatically determines the system power budget and protects the switch from any potential overpower configurations. If power is available, the NETGEAR 8800 powers on and initializes the module. When the NETGEAR 8800 detects that a module will cause an overpower condition, the module remains powered down, and is not initialized. An entry is made to the system log indicating the condition.

On a stack, the module type must be a switch that supports NETGEAR 8800.

Example

The following command configures slot 2 for a 10/100/1000, 48-port, copper module:

```
configure slot 2 module XCM8848T
```

configure slot restart-limit

```
configure slot <slot-number> restart-limit <num_restarts>
```

Description

Configures the number of times a slot can be restarted on a failure before it is shut down.

Syntax Description

slot-number	Specifies the slot number
num_restarts	Specifies the number of times the slot can be restarted. The range is from 0 to 10,000.

Default

The default is 5.

Usage Guidelines

This command allows you to configure the number of times a slot can be restarted on a failure before it is shut down. If the number of failures exceeds the restart-limit, the module goes into a “Failed” state. If that occurs, use the `disable slot` and `enable slot` commands to restart the module.

Example

The following command configures slot 2 on the switch to be restarted up to 3 times upon a failure:

```
configure slot 2 restart-limit 3
```

disable flow-control rx-pause ports

```
disable flow-control rx-pause ports [<port_list> | all]
```

Description

Disables the processing of received pause flow control messages.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
-----------	---

Default

Enabled

Usage Guidelines

With autonegotiation enabled, the NETGEAR 8800 series switches advertise the ability to support pause frames. This includes receiving and reacting to (stopping transmission) pause frames.

Use this command to disable the processing of IEEE 802.3x pause flow control messages received from the remote partner. Disabling rx-pause processing avoids dropping packets in the switch and allows for better overall network performance in some scenarios where protocols such as TCP handle the retransmission of dropped packets by the remote partner.

To disable RX flow-control, TX flow-control must first be disabled. Refer to the [disable flow-control tx-pause ports](#) command. If you attempt to disable RX flow-control with TX flow-control enabled, an error message is displayed.

Example

The following command disables the rx flow-control feature on ports 5 through 7 on the NETGEAR 8800 switch:

```
disable flow-control rx-pause ports 5-7
```

disable flow-control tx-pause ports

```
disable flow-control tx-pause ports [<port_list> | all]
```

Description

Disables the transmission of pause frames.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
-----------	---

Default

Disabled

Usage Guidelines

Use this command to stop the transmission of flow control pause frames and revert to the default.

Example

The following command disables the tx flow-control feature on ports 5 through 7 on a NETGEAR 8800:

```
disable flow-control tx-pause ports 5-7
```

disable jumbo-frame ports

```
disable jumbo-frame ports [all | <port_list>]
```

Description

Disables jumbo frame support on a port.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

- You can enable or disable jumbo frames for the entire module or switch globally only.
- You can enable and disable jumbo frames on individual ports.

Example

The following command disables jumbo frame support on slot 1, port 2 on a NETGEAR 8800 switch:

```
disable jumbo-frame ports 1:2
```

The following command disables jumbo frame support on a NETGEAR 8800 switch:

```
disable jumbo-frame ports all
```

disable learning port

```
disable learning {drop-packets | forward-packets} port [<port_list> | all]
```

Description

Disables MAC address learning on one or more ports for security purposes.

Syntax Description

port	Specifies the port.
port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports and slots.
drop-packets	Specifies that packets with unknown source MAC addresses be dropped.

forward-packets	Specifies that packets with unknown source MAC addresses be forwarded.
-----------------	--

Default

Enabled.

Usage Guidelines

Use this command in a secure environment where access is granted via permanent forwarding databases (FDBs) per port.

Example

The following command disables MAC address learning on port 4:3:

```
disable learning ports 4:3
```

disable mirroring

```
disable mirroring
```

Description

Disables port mirroring.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use the `disable mirroring` command to stop all configured copied mirroring traffic. Use this command to unconfigure all the filters on the system.

Example

The following command disables port mirroring:

```
disable mirroring
```

disable port

```
disable port [<port_list> | all]
```

Description

Disables one or more ports on the switch.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command for security, administration, and troubleshooting purposes.

When a port is disabled, the link is brought down.

Example

The following command disables ports 3, 5, and 12 through 15 on a stand-alone switch:

```
disable ports 3,5,12-15
```

The following command disables slot 1, ports 3, 5, and 12 through 15:

```
disable port 1:3,1:5,1:12-1:15
```

disable sharing

```
disable sharing <port>
```

Description

Disables a load-sharing group of ports, also known as a link aggregation group (LAG).

Syntax Description

port	Specifies the logical port of a load-sharing group or link aggregation group (LAG). Specifies a port or a combination of the slot and port number.
------	--

Default

Disabled.

Usage Guidelines

When sharing is disabled, the logical port retains all configuration including VLAN membership. All other member ports are removed from all VLANs to prevent loops and their configuration is reset to default values.

Example

The following command disables sharing on master logical port 9 in slot 3, which contains ports 9 through 12:

```
disable sharing 3:9
```

disable slot

```
disable slot <slot> {offline}
```

Description

Disables slot and leaves that module in a power down state.

Syntax Description

slot	Specifies the slot to be disabled.
offline	Specifies that the slot be disabled offline.
<p>Note: This variable is supported only on the NETGEAR 8800 series switches; that is, those switches that support offline diagnostics.</p>	

Default

Enabled.

Usage Guidelines

This command allows the user to disable a slot. When the user types this command, the I/O card in that particular slot number is brought down, and the slot is powered down. The LEDs on the card go OFF.

A disabled slot can be re-enabled using the `enable slot` command. When the slot is re-enabled, the software on the I/O module is updated to match the software on the primary MSM/MM.

The `show slot` command, if invoked after the user disables the slot, shows this slot state as “Power Off/Disabled.”

If there is no I/O card present in a slot when the user disables the slot, the slot still goes to the “Disable” state. If a card is inserted in a slot that has been disabled, the card does not come up and stays in the “Power Off/Disabled” state until the slot is enabled by using the `enable slot` command. below.

If you do not save the configuration before you do a switch reboot, the slot will be re-enabled upon reboot. If you save the configuration after disabling a slot, the slot will remain disabled after a reboot.

On Power over Ethernet (PoE) modules, disabling a slot also disables any inline power that is flowing to that slot.

This command applies only to the data, or I/O ports on slots holding an MSM. The slots holding an MSM on the NETGEAR 8810 switch are 5 and possibly 6; the slots holding an MSM on the NETGEAR 8806 switch are 3 and possibly 4. Use the `offline` parameter to run the diagnostics offline.

Example

The following command disables slot 5 on the switch:

```
disable slot 5
```

disable smartredundancy

```
disable smartredundancy <port_list>
```

Description

Disables the Smart Redundancy feature.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
-----------	---

Default

Enabled.

Usage Guidelines

The Smart Redundancy feature works in concert with the software-controlled redundant feature. When Smart Redundancy is disabled, the switch attempts only to reset the primary port to active if the redundant port fails. That is, if you disable Smart Redundancy, the traffic does not automatically return to the primary port once it becomes active again; the traffic continues to flow through the redundant port even after the primary port comes up again.

Example

The following command disables the Smart Redundancy feature on ports 1:1 to 1:4:

```
disable smartredundancy 1:1-4
```

disable snmp traps port-up-down ports

```
disable snmp traps port-up-down ports [<port_list> | all]
```

Description

Disables port up/down trap reception for specified ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command to stop receiving SNMP trap messages when a port transitions between being up and down.

Example

The following command stops ports 3, 5, and 12 through 15 on a stand-alone switch from receiving SNMP trap messages when the port goes up/down:

```
disable snmp traps port-up-down ports 3,5,12-15
```

enable flow-control rx-pause ports

```
enable flow-control rx-pause ports [<port_list> | all]
```

Description

Enables the switch to process received pause frames.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
-----------	---

Default

Enabled

Usage Guidelines

Use this command to configure the switch to return to the default behavior of processing received pause frames.

Example

The following command enables the tx flow-control feature on ports 5 through 7 on a NETGEAR 8800:

```
enable flow-control rx-pause ports 5-7
```

enable flow-control tx-pause ports

```
enable flow-control tx-pause ports [<port_list> | all]
```

Description

Enables the switch to transmit pause frames.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
-----------	---

Default

Disabled

Usage Guidelines

With autonegotiation enabled, NETGEAR 8800 series switches advertise the ability to support pause frames. This includes receiving, reacting to (stopping transmission), and transmitting pause frames. However, the switch does not actually transmit pause frames unless it is configured to do so.

IEEE 802.3x flow control provides the ability to configure different modes in the default behaviors. Use this command to configure the switch to transmit link-layer pause frames when congestion is detected.

To enable TX flow-control, RX flow-control must first be enabled. Refer to the [enable flow-control rx-pause ports](#) command. If you attempt to enable TX flow-control with RX flow-control disabled, an error message is displayed.

Example

The following command enables the tx flow-control feature on ports 5 through 7 on a NETGEAR 8800:

```
enable flow-control tx-pause ports 5-7
```

enable jumbo-frame ports

```
enable jumbo-frame ports [all | <port_list>]
```

Description

Enables support on the physical ports that will carry jumbo frames.

Syntax Description

all	Specifies ports.
-----	------------------

port_list	Specifies one or more slots and ports.
-----------	--

Default

Disabled.

Usage Guidelines

Increases performance to back-end servers or allows for vMAN 802.1Q encapsulations.

You can configure the maximum size of a jumbo frame if you want to use a different size than the default value of 9216. Use the `configure jumbo-frame-size` command to configure the size.

This setting is preserved across reboots.

You can enable and disable jumbo frames on individual ports.

Example

The following command enables jumbo frame support on slot 3, port 5 on a NETGEAR 8800 switch:

```
enable jumbo-frame ports 3:5
```

The following command enables jumbo frame support on a NETGEAR 8800 switch:

```
enable jumbo-frame ports all
```

enable learning port

```
enable learning port [all | <port_list>]
```

Description

Enables MAC address learning on one or more ports.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

N/A.

Example

The following command enables MAC address learning on slot 1, ports 7 and 8:

```
enable learning ports 1:7-8
```

enable mirroring to port

```
enable mirroring to [port <port> | port-list <port-list> loopback-port <port> ] {remote-tag <vlan tag>}
```

Description

Dedicates a port on the switch to be the mirror output port, or the monitor port.

Syntax Description

port	Specifies the mirror output port.
port-list	Specifies the list of ports where traffic is to be mirrored.
loopback-port	Specifies an otherwise unused port required when mirroring to a port-list. The loopback-port is not available for switching user data traffic.
port	Specifies a single loopback port that is used internally to provide this feature.
remote-tag	Specifies the value of the VLAN ID used by the mirrored packets when egressing the monitor port.

Default

Disabled.

Usage Guidelines

Port mirroring configures the switch to copy all traffic associated with one or more ports, VLANs or virtual ports. A virtual port is a combination of a VLAN and a port. The monitor port(s) can be connected to a network analyzer or RMON probe for packet analysis. The switch uses a traffic filter that copies a group of traffic to the monitor port.

Up to 16 mirroring filters and one monitor port can be configured on the switch. After a port has been specified as a monitor port, it cannot be used for any other function. Frames that contain errors are not mirrored.

You cannot run ELSM and mirroring on the same port. If you attempt to enable mirroring on a port that is already enabled for ELSM, the switch returns a message similar to the following:

```
Error: Port mirroring cannot be enabled on an ELSM enabled port.
```

The traffic filter on NETGEAR 8800 series switches can be defined based on one of the following criteria:

- **Physical port**—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port. You can specify which traffic the port mirrors:

- Ingress—Mirrors traffic received at the port.
- Egress—Mirrors traffic sent from the port.
- Ingress and egress—Mirrors all traffic forwarded by the port.

(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).

- **VLAN**—All data to a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- **Virtual port**—All data specific to a VLAN on a specific port is copied to the monitor port.
- Only 8 VLANs can be mirrored on a given physical port.
- Only traffic *ingressing* a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- When routing between VLANs, ingress mirrored traffic is presented to the monitor port as *modified* for routing. This is the default behavior and the behavior when you use the command, `configure mirroring mode standard`. When you use the command, `configure mirroring mode enhanced`, ingress traffic is mirrored as it is received (on the wire).
- In standard mode (see `configure mirroring mode` command), even if you select ingress and egress traffic, the packet is mirrored only the *first* time it matches a mirror filter and is not mirrored on subsequent configured filters. In enhanced mode, packets which match both an ingress filter and an egress filter will result in two packets egressing the monitor port or ports.
- You cannot include the monitor port for the NETGEAR 8800 series switch in a load-sharing group.
- You can run mirroring and sFlow on the same device when you are running NETGEAR 8800.
- With a monitor port on a NETGEAR 8800 original-series module, *all* traffic egressing the monitor port is tagged (regardless of what module the ingressing port is on). Even if some untagged ports send mirrored traffic to the monitor port, that traffic also egresses the monitor port tagged with the internal VLAN ID.
- When you are using standard mode mirroring on an 8800, a packet that matches both an ingress mirroring filter and an egress mirroring filter may only be ingress mirrored. The behavior depends on the location of the ingress port, egress port and monitor port within the switch as well as the type of switch on which the packet ingresses. When using enhanced mode mirroring, two packets are mirrored when a packet encounters both an ingress and egress mirroring filter. r one-half of that module or on another module.

Enhanced mirroring mode must be configured if you are going to configure a remote mirroring tag. Enhanced mirroring mode is configured using the following command:

```
configure mirroring mode enhanced
```

Note: This parameter is used for the remote port mirroring feature only.

Example

The following example selects slot 3, port 4 as the mirror, or monitor, port on the NETGEAR 8800 switch:

```
enable mirroring to port 3:4
```

The following example selects slot 1, port 3 as the tagged mirror, or monitor, port on the NETGEAR 8800 switch:

```
enable mirroring to port 1:3 tagged
```

enable port

```
enable port [<port_list> | all]
```

Description

Enables a port.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

All ports are enabled.

Usage Guidelines

Use this command to enable the port(s) if you disabled the port(s) for security, administration, or troubleshooting purposes.

Example

The following command enables ports 3, 5, and 12 through 15 on the stand-alone switch:

```
enable ports 3,5,12-15
```

The following command enables slot 1, ports 3, 5, and 12 through 15:

```
enable port 1:3, 1:5, 1:12-1:15
```

enable sharing grouping

```
enable sharing <port> grouping <port_list> {algorithm [address-based {L2 | L3 | L3_L4 | custom}]} {lacp | health-check}
```

Description

Enables the switch to configure port link aggregation, or load sharing. By using link aggregation, you use multiple ports as a single logical port. Link aggregation also provides redundancy because traffic is redistributed to the remaining ports in the LAG if one port in the group goes down. LACP allows the system to dynamically configure the LAGs.

Syntax Description

port	Specifies the master logical port for a load-sharing group or link aggregation group (LAG).
port_list	Specifies one or more ports or slots and ports to be grouped to the logical port.
address-based	Specifies link aggregation by address-based algorithm.
L2	Specifies address-based link aggregation by Layer 2. This is the default value.
L3	Specifies address-based link aggregation by Layer 3.
L3_L4	Specifies address-based link aggregation by Layer 3 IP plus Layer 4 port.
custom	<p>Selects the custom link aggregation algorithm configured with the following command:</p> <pre>configure sharing address-based custom [ipv4 [L3-and-L4 source-only destination-only source-and-destination] hash-algorithm [xor crc-16]].</pre> <p>The custom option applies to all LAGs on the switch.</p>
lACP	Specifies dynamic link aggregation, or load sharing, using the LACP.
health-check	Specifies a health check type of link aggregation group.

Default

Disabled.

Usage Guidelines

Link aggregation, or load sharing, allows you to increase bandwidth and availability between switches by using a group of ports to carry traffic in parallel between switches. The aggregation algorithm allows the switch to use multiple ports as a single logical port. For example, VLANs see the link aggregation group (LAG) as a single logical port. Groups can span multiple modules.

Note: All ports that are designated for the LAG must be removed from all VLANs prior to configuring the LAG.

You can enable and configure dynamic link aggregation, using LACP or health-check link aggregation. Static link aggregation is the default link aggregation method.

Note: Always verify the LACP configuration by issuing the `show ports sharing` command; look for the ports listed as being in the aggregator.

If a port in a LAG fails, traffic is redistributed to the remaining ports in the LAG. If the failed port becomes active again, traffic is redistributed to include that port.

Link aggregation must be enabled on both ends of the link, or a network loop will result.

Note: See *NETGEAR 8800 User Manual* for information on the interaction of port-based ACLs and LAGs of ports.

LAGs are defined according to the following rules:

- Although you can *reference* only the logical port of a LAG to a Spanning Tree Domain (STPD), *all* the ports of a load-sharing group actually belong to the specified STPD.
- When using link aggregation, you should always reference the logical port of the LAG when configuring or viewing VLANs. VLANs configured to use other ports in the LAG will have those ports deleted from the VLAN when link aggregation becomes enabled.

Link aggregation, or load-sharing, algorithms allow you to select the distribution technique used by the LAG to determine the output port selection. Algorithm selection is not intended for use in predictive traffic engineering.

- **Port-based**—Uses the ingress port to determine which physical port in the LAG is used to forward traffic out of the switch.
- **Address-based**—Uses addressing information to determine which physical port in the LAG to use for forwarding traffic out of the switch. Refer to `configure sharing address-based custom` for more information on using addressing information.

The following guidelines apply to link aggregation on the NETGEAR 8800 series switch:

- A static LAG can include a maximum of 8 ports.
- An LACP LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- A Health Check LAG can include a maximum of 8 ports.
- The available address-based parameters on the NETGEAR 8800 series switch are L2 for Layer 2 and L3 for Layer 3.

If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting. The switch can use IPv6 addresses.

- Broadcast, multicast, or unknown unicast packets are transmitted differently depending on the device you are using:

- On the 8800 original-series modules, these packets are transmitted on a single port of a LAG.
- On the 8800, these packets are distributed across all members of a LAG. The distribution of these packets depends on the type of the traffic. Broadcast, L2 multicast and unknown unicast traffic distribution is based on the source and destination MAC addresses. IP multicast traffic distribution is based on the source and destination IP addresses. This behavior is not configurable.
- The `custom` keyword is supported only on NETGEAR 8800 switches. If the `custom` keyword is specified on a NETGEAR 8800 switch that includes a mix of 8800 series modules, the individual modules use algorithms as follows:
 - The XCM8848T, XCM8824F, and XCM8808X I/O modules forward unicast traffic using the L3 algorithm.
 - All other modules forward unicast traffic using the L3_L4 algorithm.
 - All modules forward non-unicast traffic (broadcast, multicast, and unknown unicast packets) using a separate internal hash algorithm.

Example

The following example defines a static link aggregation group (LAG) on a switch that contains ports 9 through 12 on slot 3, ports 7 through 10 on slot 5, and uses the first port on slot 3 as the logical port 9:

```
enable sharing 3:9 grouping 3:9-3:12, 5:7-5:10
```

In this example, logical port 3:9 represents physical ports 3:9 through 3:12 and 5:7 through 5:10.

The following example defines a dynamic LAG on a stand-alone switch containing ports 10 through 15, with port 10 being the logical port:

```
enable sharing 10 grouping 10-15 lacp
```

The following example selects the custom option on a NETGEAR 8800 switch:

```
XCM8810.1 # enable sharing 2:1 grouping 2:1-2 algorithm address-based custom
```

The following example defines a health check LAG containing ports 10 through 13 with port 10 as the master logical port and specifies address-based link aggregation by Layer 3 IP plus Layer 4 port:

```
enable sharing 10 grouping 10,11,12,13 algorithm address L3_L4 health-check
```

To configure a health-check LAG, refer to the `configure sharing health-check member-port add tcp-tracking` command.

enable slot

```
enable slot <slot>
```

Description

Enables slots.

Syntax Description

slot	Specifies the slot to be enabled.
------	-----------------------------------

Default

Enabled.

Usage Guidelines

This command allows the user to enable a slot that has been previously disabled using the `disable slot` command.

Note: On the NETGEAR 8800 series switches, this command applies only to the data or I/O ports on slots holding an MSM.

When the user enters the enable command, the disabled I/O card in the specified slot is brought up, and the slot is made operational, if possible, or goes to the appropriate state as determined by the card state machine. The LEDs on the card are brought ON as usual. When the slot is enabled, the software on the I/O module is updated to match the software on the primary MSM/MM.

After the user enables the slot, the `show slot` command shows the state as “Operational” or will display the appropriate state if the card could not be brought up successfully. Note that there is no card state named “Enable” and the card goes to the appropriate states as determined by the card state machine when the `enable slot` command is invoked.

Only slots that have their state as “disabled” can be enabled using this command. If this command is used on slots that are in states other than “disabled,” the card state machine takes no action on these slots.

To enable inline power to a slot, the slot must be enabled as well as inline power for that slot. Use the `enable inline-power` command to enable inline power.

Note: If your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be enabled; the slot will not function in data-only mode without enough power for inline power.

Example

The following command enables slot 5 on the switch:

```
enable slot 5
```


enable smartredundancy

```
enable smartredundancy <port_list>
```

Description

Enables the Smart Redundancy feature on the primary port.

Syntax Description

portlist	Specifies one or more ports or slots and ports.
----------	---

Default

Enabled.

Usage Guidelines

You must configure the software-controlled redundant port using the [configure ports redundant](#) command prior to enabling Smart Redundancy.

The Smart Redundancy feature works in concert with the software-controlled redundant port feature. With Smart Redundancy enabled on the switch, when the primary port becomes active the switch redirects all traffic to the primary port and blocks the redundant port again. (If you disable Smart Redundancy, the primary port is blocked because traffic is now flowing through the redundant, port.)

Example

The following command enables the Smart Redundancy feature on slot 1, port 4:

```
enable smartredundancy 1:4
```

enable snmp traps port-up-down ports

```
enable snmp traps port-up-down ports [<port_list> | all]
```

Description

Enables port up/down trap reception for specified ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command to begin receiving SNMP trap messages when a port transitions between being up and down.

Example

The following command enables ports 3, 5, and 12 through 15 on a stand-alone switch to receive SNMP trap messages when the port goes up/down:

```
enable snmp traps port-up-down ports 3,5,12-15
```

restart ports

```
restart ports [all | <port_list>]
```

Description

Resets autonegotiation for one or more ports by resetting the physical link.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command resets autonegotiation on slot 1, port 4:

```
restart ports 1:4
```

run failover

```
run failover {force}
```

Description

Causes a user-specified node failover.

Syntax Description

force	Force failover to occur.
-------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to cause the primary MSM/MM to failover to the backup MSM/MM, or the Master node to failover to the Backup node.

Before you initiate failover, use the `show switch {detail}` command to confirm that the nodes are in sync and have identical software and switch configurations. If the output shows MASTER and BACKUP (InSync), the two MSMs/MMs or nodes are in sync.

If the MSM/MM's software and configuration are not in sync, use the `synchronize` command to get the two MSMs/MMs or nodes in sync. This command ensures that the backup has the same software in flash as the master.

Example

The following command causes a failover:

```
run failover
```

run msm-failover

```
run msm-failover {force}
```

Description

Causes a user-specified node failover.

Syntax Description

force	Force failover to occur.
-------	--------------------------

Default

N/A.

Usage Guidelines

This command is being replaced with the `run failover` command. For usage guidelines, see the description for the `run failover` command.

Example

The following command causes a user-specified MSM failover:

```
run msm-failover
```

show lacp

```
show lacp
```

Description

Displays LACP, or dynamic link aggregation, settings on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following information about the LACP LAGs configured on the switch:

- Up or Down
- Enabled or disabled (not configurable)
- System MAC
 - MAC address for the system, which is used for LACP priority in the absence of a specifically configured priority.
- LACP PDUs dropped on non-LACP ports
- LAG
 - Identifies the particular LAG. This number comes from logical port assigned to the LAG and is the LAG group ID.
- Actor Sys-Pri
 - Shows the system priority for that LAG.
 - If this number is lower than the number displayed for the Partner Sys-Pri, the system you are working on is the controlling partner in the LAG.
- Actor Key
 - Automatically generated LACP key.
- Partner MAC
 - Identifies the MAC address for the system connecting to the LAG on the remote end.
- Partner Sys-Pri
 - Shows the system priority for that LAG on the remote end.

- If this number is lower than the number displayed for the Actor Sys-Pri, the system at the remote end is the controlling partner in the LAG.
- Partner Key
 - LACP key automatically generated by the system to which this aggregator is connected.
 - If this number is lower than the number displayed for the Actor Key, the partner system is the controlling partner in the LAG.
- Agg Count
 - Identifies the number of ports added to the aggregator for that LAG.

Example

The following command displays the LACP LAGs on the switch:

```
show lacp
```

The following is sample output from this command:

```
LACP Up                : Yes
LACP Enabled           : Yes
System MAC              : 00:04:96:10:33:60
LACP PDUs dropped on non-LACP ports : 0
```

Lag	Actor Sys-Pri	Actor Key	Partner MAC	Partner Sys-Pri	Partner Key	Agg Count
2:1	90	0x07d1	00:01:30:f9:9c:30	601	0x1391	2
4:5	100	0x0fa5	00:01:30:f9:9c:30	321	0x1f47	16
4:9	677	0x0fa9	00:01:30:f9:9c:30	87	0x0fa9	8

show lacp counters

```
show lacp counters
```

Description

Displays all LACP, or dynamic link aggregation, counters for all member ports in the system.

Syntax Description

This command has no parameters or variables.

Default

N/A.

Usage Guidelines

This command displays the following information for all link aggregation groups (LAGs):

- LACP PDUs dropped on non-LACP ports
- LACP bulk checkpointed messages sent
- LACP bulk checkpointed messages received
- LACP PDUs checkpointed sent
- LACP PDUs checkpointed received
- LAG group ID
- Member port
- Packets received
- Packets dropped from PDU error
- Packets dropped because LACP is not enabled on this port
- Packets dropped because sender's system MAC address matches that of receiver
- Packets successfully transmitted
- Packets with errors during transmission

Example

The following command displays LACP counters:

```
show lacp counters
```

The following is sample output from this command:

```
LACP PDUs dropped on non-LACP ports : 519392
LACP Bulk checkpointed msgs sent    : 1
LACP Bulk checkpointed msgs rcv     : 0
LACP PDUs checkpointed sent        : 575616
LACP PDUs checkpointed rcv         : 0
```

Lag Group	Member Port	Rx Ok	Rx Drop PDU Err	Rx Drop Not Up	Rx Drop Same MAC	Tx Sent	Tx Ok	Tx Err
-----------	-------------	-------	-----------------	----------------	------------------	---------	-------	--------

1:1	1:1	2169	0	0	0	2170	0	
	1:2	2169	0	0	0	2170	0	
	1:3	2169	0	0	0	2170	0	
	1:4	2169	0	0	0	2170	0	
	1:5	2169	0	0	0	2170	0	
	1:6	2169	0	0	0	2170	0	
	1:7	2169	0	0	0	2170	0	
	1:8	2168	0	0	0	2169	0	

show lacp lag

```
show lacp lag <group-id> {detail}
```

Description

Displays LACP, or dynamic link aggregation, settings for the specified LAG.

Syntax Description

group-id	Specifies the LAG group ID you want to display. This is the number of the port you configured as the logical port of the LAG.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

This command displays the following information about the specified LACP LAG:

- LAG
 - Identifies the particular LAG. This number comes from logical port assigned to the LAG and is the LAG group ID.
- Actor Sys-Pri
 - Shows the system priority for that LAG.
 - If this number is lower than the number displayed for the Partner Sys-Pri, the system you are working on is the controlling partner in the LAG.
- Actor Key
 - Automatically generated LACP key.
- Partner MAC
 - Identifies the MAC address for the system connecting to the LAG on the remote end.
- Partner Sys-Pri
 - Shows the system priority for that LAG on the remote end.
 - If this number is lower than the number displayed for the Actor Sys-Pri, the system at the remote end is the controlling partner in the LAG.
- Partner Key
 - LACP key automatically generated by the system to which this aggregator is connected.
 - If this number is lower than the number displayed for the Actor Key, the partner system is the controlling partner in the LAG.
- Agg Count
 - Identifies the number of ports added to the aggregator for that LAG.
- Member port
- Port priority

- Rx State—Receiving state of the port
 - Idle
 - Initialized
 - Current—Receiving LACP PDUs
 - Expired
 - Defaulted
- Sel Logic—Selection state of the port
 - Selected—Ports with a matching admin key on the remote end.
 - Unselected—Ports that failed to meet with a matching admin key on the remote end.
 - Standby—Ports that exceed the number of ports that can be active in the LAG simultaneously. These ports can be moved into selected mode if one of the currently selected ports in the LAG goes down.
- Mux State—Ability to transmit and collect data of the port
 - Waiting—Selected port that is waiting for LACP to determine if it can join the aggregator.
 - Attached—Ports ready to be added to the aggregator.
 - Collecting-Dist—Ports that are added to the aggregator and are transferring data.
 - Detached—Ports that cannot be added to the aggregator.
- Actor Flag—Mux state of the port
 - A—Activity
 - T—Timeout
 - G—Aggregation
 - S—Synchronization
 - C—Collecting
 - D—Distributing
 - F—Defaulted
 - E—Expired
- Partner Port
 - The operational value of the port number assigned to this link by partner.
- Up—Yes or no
- Enabled—Yes or no
- Unack count
- Wait-for-count
- Current timeout
- Activity mode
- Defaulted action
- Receive state

- Transmit state
- Selected count—Number of selected ports in the LAG
- Standby count—Number of standby ports in the LAG
- LAG Id flag
 - S—Displays information on controlling partner of LAG.
 - T—Displays information on controlled partner of LAG.

Example

The following command displays information on the specified LACP LAG:

```
show lacp lag 4:9
```

The following is sample output from this command:

Lag	Actor	Actor	Partner	Partner	Partner	Agg
Sys-Pri	Key	MAC	Sys-Pri	Key	Count	
4:9	2110	0x0fa9	00:04:96:10:33:60	2110	0x0fa9	16

Port list:

Member	Port	Rx	Sel	Mux	Actor	Partner
Port	Priority	State	Logic	State	Flags	Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009
4:10	301	Current	Selected	Collect-Dist	A-GSCD--	4010
4:11	302	Current	Standby	Detached	A-G-----	4011
4:12	303	Current	Standby	Detached	A-G-----	4012
4:29	200	Current	Selected	Collect-Dist	A-GSCD--	4029
4:30	0	Current	Selected	Collect-Dist	A-GSCD--	4030
4:31	202	Current	Selected	Collect-Dist	A-GSCD--	4031
4:32	203	Current	Selected	Collect-Dist	A-GSCD--	4032
8:7	101	Current	Selected	Collect-Dist	A-GSCD--	8013
8:8	10	Current	Selected	Collect-Dist	A-GSCD--	8014
8:9	9	Current	Selected	Collect-Dist	A-GSCD--	8015
8:10	8	Current	Selected	Collect-Dist	A-GSCD--	8016
8:11	7	Current	Selected	Collect-Dist	A-GSCD--	8017
8:12	6	Current	Selected	Collect-Dist	A-GSCD--	8018
8:13	5	Current	Selected	Collect-Dist	A-GSCD--	8019
8:14	3	Current	Selected	Collect-Dist	A-GSCD--	8020
8:15	0	Current	Selected	Collect-Dist	A-GSCD--	8043
8:16	3	Current	Selected	Collect-Dist	A-GSCD--	8044
8:17	2	Idle	Unselected	Detached	-----	0
8:18	37	Idle	Unselected	Detached	-----	0
8:19	36	Idle	Unselected	Detached	-----	0
8:20	35	Idle	Unselected	Detached	-----	0

```
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
              C-Collecting, D-Distributing, F-Defaulted, E-Expired
```

The following command displays detailed information on the specified LACP LAG:

```
show lacp lag 4:9 detail
```

The following is sample output from this command:

```

Lag          Actor    Actor    Partner          Partner  Partner  Agg
              Sys-Pri  Key      MAC              Sys-Pri  Key      Count
-----
4:9          2110     0x0fa9   00:04:96:10:33:60  2110     0x0fa9   16

Up           : Yes
Enabled      : Yes
Unack count  : 0
Wait-for-count : 0
Current timeout : Long
Activity mode : Active
Defaulted Action : Delete
Receive state : Enabled
Transmit state : Enabled
Selected count : 16
Standby count : 2
LAG Id flag   : Yes
S.pri:2110, S.id:00:01:30:f9:9c:30, K:0x0fa9
T.pri:2110, T.id:00:04:96:10:33:60, L:0x0fa9

```

Port list:

Member Port	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009
4:10	301	Current	Selected	Collect-Dist	A-GSCD--	4010
4:11	302	Current	Standby	Detached	A-G----	4011
4:12	303	Current	Standby	Detached	A-G----	4012
4:29	200	Current	Selected	Collect-Dist	A-GSCD--	4029
4:30	0	Current	Selected	Collect-Dist	A-GSCD--	4030
4:31	202	Current	Selected	Collect-Dist	A-GSCD--	4031
4:32	203	Current	Selected	Collect-Dist	A-GSCD--	4032
8:7	101	Current	Selected	Collect-Dist	A-GSCD--	8013
8:8	10	Current	Selected	Collect-Dist	A-GSCD--	8014
8:9	9	Current	Selected	Collect-Dist	A-GSCD--	8015
8:10	8	Current	Selected	Collect-Dist	A-GSCD--	8016

```

8:11      7      Current      Selected      Collect-Dist      A-GSCD--      8017
8:12      6      Current      Selected      Collect-Dist      A-GSCD--      8018
8:13      5      Current      Selected      Collect-Dist      A-GSCD--      8019
8:14      3      Current      Selected      Collect-Dist      A-GSCD--      8020
8:15      0      Current      Selected      Collect-Dist      A-GSCD--      8043
8:16      3      Current      Selected      Collect-Dist      A-GSCD--      8044
8:17      2      Idle        Unselected      Detached          -----      0
8:18      37     Idle        Unselected      Detached          -----      0
8:19      36     Idle        Unselected      Detached          -----      0
8:20      35     Idle        Unselected      Detached          -----      0
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
              C-Collecting, D-Distributing, F-Defaulted, E-Expired

```

show lacp member-port

```
show lacp member-port <port> {detail}
```

Description

Displays LACP, or dynamic link aggregation, settings for the specified port that is a member of any LAG.

Syntax Description

port	Specifies the port number.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

This command displays the following information about the specified port:

- Member Port
- Port Priority
- Rx State—Receiving state of the port
 - Idle
 - Initialized
 - Current—Receiving LACP PDUs
 - Expired
 - Defaulted
- Sel Logic—Selection state of the port

- Selected—Ports with a matching admin key on the remote end.
- Unselected—Ports that failed to meet with a matching admin key on the remote end.
- Standby—Ports that exceed the number of ports that can be active in the LAG simultaneously. These ports can be moved into selected mode if one of the currently selected ports in the LAG goes down.
- Mux State—Ability to transmit and collect data of the port
 - Waiting—Selected port that is waiting for LACP to determine if it can join the aggregator.
 - Attached—Ports ready to be added to the aggregator.
 - Collecting-Dist—Ports that are added to the aggregator and are transferring data.
 - Detached—Ports that cannot be added to the aggregator.
- Actor Flag
 - A—Activity
 - T—Timeout
 - G—Aggregation
 - S—Synchronization
 - C—Collecting
 - D—Distributing
 - F—Defaulted
 - E—Expired
- Partner Port
 - The operational value of the port number assigned to this link by partner.
- Up or Down—LACP protocol running or not on specified port
- Enabled or disabled (not configurable)
- Link State—Link state on this port up or down
- Actor Churn—True or false
- Partner Churn—True or false
- Ready_N—Ready to be added to aggregator.
- Wait pending
- Ack pending
- LAG Id
 - S—Displays information on controlling partner of LAG.
 - T—Displays information on controlled partner of LAG.
- Stats
 - Rx - Accepted
 - Rx - Dropped due to error in verifying PDU
 - Rx - Dropped due to LACP not being up on this port

- Rx - Dropped due to matching own MAC
- Tx - Sent Successfully
- Tx - Transmit error

Example

The following command displays LACP information on the specified port:

```
show lacp member-port 4:9
```

The following is sample output from this command:

Member Port	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009

=====

Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
C-Collecting, D-Distributing, F-Defaulted, E-Expired

The following command displays detailed LACP information on the specified port:

```
show lacp member-port 4:9 detail
```

The following is sample output from this command:

Member Port	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009

Up : Yes

Enabled : Yes

Link State : Up

Actor Churn : False

Partner Churn : False

Ready_N : Yes

Wait pending : No

Ack pending : No

LAG Id:

S.pri:2110, S.id:00:01:30:f9:9c:30, K:0x0fa9, P.pri:300 , P.num:4009

T.pri:2110, T.id:00:04:96:10:33:60, L:0x0fa9, Q.pri:300 , Q.num:4009

Stats:

Rx - Accepted : 2174

Rx - Dropped due to error in verifying PDU : 0

Rx - Dropped due to LACP not being up on this port : 0

Rx - Dropped due to matching own MAC : 0

Tx - Sent successfully : 2175

Tx - Transmit error : 0

=====

Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
C-Collecting, D-Distributing, F-Defaulted, E-Expired

show mirroring

show mirroring

Description

Displays the port-mirroring configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You must enable mirroring on the switch prior to configuring mirroring, and you must configure mirroring to display mirroring statistics. Use the `enable mirroring to port` command to enable mirroring and the `configure mirroring add` command to configure mirroring.

You can use this command to display mirroring statistics and determine if mirroring is enabled or disabled on the switch.

Example

The following command displays switch mirroring statistics:

```
show mirroring
```

Following is sample output from this command for a NETGEAR 8810 switch that is configured for port-based mirroring for single monitor ports:

```
Mirror port: 3:15 is up
Number of Mirroring filters: 3
Mirror Port configuration:
    Port number 3:12 in  all vlans ingress only
    Port number 5:4 in  all vlans egress only
    Port number 8:30 in  all vlans
```

show ports

```
show ports {<port_list>} {no-refresh}
```

Description

Display port summary statistics.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of the data.

Default

N/A.

Usage Guidelines

Use this command to display the port number, display string, and some of the port states in tabular form.

The VLAN name is displayed only if that port contains a single VLAN. If the port contains more than one VLAN, then the number of the VLANs are displayed.

Example

The following command displays on slot 2-3 on port 1 and slot 12 on port 10:

```
show ports 1:2-3,10:12
```

Following is sample output from this command:

```
show ports 1:2-3,10:12
Port Summary Monitor                               Thu Feb 14 14:19:50 2008
Port  Display          VLAN Name          Port  Link  Speed  Duplex
#    String            (or # VLANs)      State State Actual Actual
=====
1:2   2nd-Floor-Lab      Lab-Backbone      E     A     1000   FULL
1:3                               Building2         E     A D
10:12 AllBackboneLANs (34)         E     R                               FULL
=====
Port State: D-Disabled, E-Enabled
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback,
            D-ELSM enabled but not up
            U->page up  D->page down  ESC->exit
```

show ports anomaly

```
show ports <port list> anomaly {no-refresh}
```

Description

Display statistics of anomaly violation events in real time.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

N/A.

Usage Guidelines

If you do not specify a port number or range of ports, statistics are displayed for all ports. To clear the counters, use the `clear counters ports` command. The default display is a constantly refreshing real-time display. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

This command takes effect after enabling anomaly-protection.

Example

The following command displays real-time anomaly statistics on slot 2, all ports:

```
show ports 2:* anomaly
```

Following is sample output from this command:

```
Port Statistics Thu Nov  9 22:44:31 2006
Port   Link      Rx Pkt ===== Anomaly Violation =====
      State      Count          L3 Count          L4 Count          ICMP Count          Frag Count
=====
2:1     A         191585           1              2              0              0
2:2     R           0              0              0              0              0
2:3     R           0              0              0              0              0
2:4     R           0              0              0              0              0
2:5     R           0              0              0              0              0
2:6     R           0              0              0              0              0
2:7     R           0              0              0              0              0
2:8     R           0              0              0              0              0
2:9     R           0              0              0              0              0
2:10    R           0              0              0              0              0
2:11    R           0              0              0              0              0
2:12    A         178024           0              0              0              0
2:13    A         196956           0              0              0              0
2:14    R           0              0              0              0              0
2:15    R           0              0              0              0              0
2:16    R           0              0              0              0              0
2:17    R           0              0              0              0              0
=====
```


Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
 0->Clear Counters U->page up D->page down ESC->exit

show ports collisions

show ports {mgmt | <port_list>} collisions {no-refresh}

Description

Displays real-time collision statistics.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number or range of ports, collision statistics are displayed for all ports. To clear the counters, use the `clear counters ports` command. The default display is a constantly refreshing real-time display. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

This status information may be useful for your technical support representative if you have a network problem.

Example

The following command displays real-time collision statistics on slot 1, ports 1 and 2:

```
show ports 1:1-2 collisions
```

Following is sample output from this command:

```
Port Collision Monitor
Port          Link          Collision Histogram
              State    1    2    3    4    5    6    7    8    9   10   11   12   13   14   15   16
=====
   1:1        A         0    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
   1:2        R         0    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
=====
Link State: A-Active R-Ready, NP-Port not present, L-Loopback
```

The numbers 1 to 16 represent the number of collisions encountered prior to successfully transmitting the packet; this is applicable only for half-duplex links.

show ports configuration

```
show ports {mgmt | <port_list>} configuration {no-refresh}
```

Description

Displays port configuration statistics, in real time or snapshot.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number or range of ports, configuration statistics are displayed for all ports. If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

This status information may be useful for your technical support representative if you have a network problem.

This command displays port configuration, which includes:

- Virtual router
- Port state
- Link state
- Autonegotiation information
- Link speed
- Duplex mode
- Flow control
- Load sharing information
- Link media information

Note: On 10 Gbps ports, the Media Primary column displays NONE when no module is installed, and SR, LR, or ER depending on the module installed when there is one present.

Example

The following command displays the port configuration for all ports:

```
show ports configuration
Port Configuration Monitor                               Fri Apr 13 10:22:29 2007
Port      Virtual      Port Link Auto  Speed      Duplex      Flow  Load  Media
          router       State State Neg  Cfg Actual  Cfg Actual Cntrl Master Pri Red
=====
1         VR-Default    E      R    ON  AUTO      AUTO                        NONE UTP
2         VR-Default    E      R    ON  AUTO      AUTO                        NONE UTP
3         VR-Default    E      R    ON  AUTO      AUTO                        NONE UTP
4         VR-Default    E      R    ON  AUTO      AUTO                        NONE UTP
5         VR-Default    E      R    ON  AUTO      AUTO                        NONE
6         VR-Default    E      R    ON  AUTO      AUTO                        NONE
7         VR-Default    E      R    OFF 100      FULL                        SX
8         VR-Default    E      R    ON  AUTO      AUTO                        NONE
9         VR-Default    E      R    ON  AUTO      AUTO                        NONE
10        VR-Default    E      R    ON  AUTO      AUTO                        NONE
11        VR-Default    E      R    ON  AUTO      AUTO                        NONE
12        VR-Default    E      R    ON  AUTO      AUTO                        NONE
13        VR-Default    E      R    ON  AUTO      AUTO                        NONE
14        VR-Default    E      R    ON  AUTO      AUTO                        NONE
15        VR-Default    E      R    ON  AUTO      AUTO                        NONE
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled, Media: !-Unsupported Optic Module
Media Red: * - use "show port info detail" for redundant media type
0->Clear Counters  U->page up  D->page down  ESC->exit
```

The following command displays the port configuration statistics for slot 2, port 2:

```
show ports 2:2 configuration
```

Following is sample output from this command:

```
Port Configuration
Port      Virtual      Port Link Auto  Speed      Duplex      Flow  Load  Media
          router       State State Neg  Cfg Actual  Cfg Actual Cntrl Master Pri Red
=====
2:2       VR-Default    E      R    ON  AUTO      AUTO                        UTP
=====
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
Port State: D-Disabled E-Enabled, Media: !-Unsupported Optic Module
0->Clear Counters  U->page up  D->page down  ESC->exit
```

show ports information

```
show ports {mgmt | <port_list>} information {detail}
```

Description

Displays detailed system-related information.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports of slots and ports.
detail	Specifies detailed port information.

Default

N/A.

Usage Guidelines

This command displays information, including the following:

- Port number
- Port configuration
 - Virtual router
 - Type of port
 - Admin state
 - Link state and speed
 - Link counter
 - VLAN configuration
 - STP configuration
 - Trunking, or load sharing
 - ELSM (disabled; or if enabled, the ELSM link state is shown as well)
 - Load balancing
 - Learning
 - Egress flooding
 - Jumbo frames
 - Link port up/down traps
 - QoS profiles
 - vMAN status
 - Smart Redundancy status
 - SRP status
 - Additional platform-specific information

If you do not specify a port number or range of ports, detailed system-related information is displayed for all ports. The data is displayed in a table format.

This status information may be useful for your technical support representative if you have a network problem.

The `detail` parameter is used to provide more specific port information. The data is called out with written explanations versus displayed in a table format.

Note: The keyword `detail` displays slightly different information depending on the platform and configuration you are working with.

The link filter counter displayed with the `detail` keyword is calculated at the middle layer on receiving an event. The link filter up indicates the number of link transitions from down to up at the middle layer filter.

Example

The following command displays port system-related information on a NETGEAR 8810 switch:

```
show port 1:1 info
```

Following is sample output from this command:

```
* XCM8806.1 # show port 1:1 info
Port   Flags           Link           OAM   Link Num Num   Num   Jumbo QOS      Load
              State           UPS   STP VLAN Proto Size  profile Master
=====
1:1     Em-----fMB---x ready   -/-   0    1    1    1   9216  none
=====
> indicates Port Display Name truncated past 8 characters
Flags : a - Load Sharing Algorithm address-based, D - Port Disabled,
        E - Port Enabled,
        g - Egress TOS Enabled, j - Jumbo Frame Enabled,
        l - Load Sharing Enabled, m - MACLearning Enabled,
        n - Ingress TOS Enabled, o - Dot1p Replacement Enabled,
        P - Software redundant port(Primary),
        R - Software redundant port(Redundant),
        q - Background QOS Monitoring Enabled,
        s - diffserv Replacement Enabled,
        v - Vman Enabled, f - Unicast Flooding Enabled,
        M - Multicast Flooding Enabled, B - Broadcast Flooding Enabled
        O - Ethernet OAM Enabled
        w - MACLearning Disabled with Forwarding
        b - Rx and Tx Flow Control Enabled, x - Rx Flow Control Enabled
```

The following command displays detailed port system-related information on the NETGEAR 8800 switch:

```
show ports 3:1 information detail
```

Following is sample output from this command:

```
Port:      3:1
  Virtual-router: VR-Default
  Type:        UTP
  Random Early drop:      Unsupported
  Admin state:   Enabled with auto-speed sensing (100M Advertised), auto-duplex
(half-duplex Advertised)
  ELSM Link State:   Up
  Link State:      Active, 1 Gbps, full-duplex
  Link Counter: Up      1 time(s)
  VLAN cfg:
      Name: Default, Internal Tag = 1 (MAC-Based), MAC-limit = No-limit

  STP cfg:
      s0(disable), Tag=(none), Mode=802.1D, State=FORWARDING

  Protocol:
      Name: Default      Protocol: ANY      Match all protocols.
  Trunking:      Load sharing is not enabled.
  ELSM:          Enabled
  Learning:      Enabled
  Unicast Flooding:      Enabled
  Multicast Flooding:    Enabled
  Broadcast Flooding:    Enabled
  Jumbo: Enabled, MTU= 9194
  Flow Control: Rx-Pause: Disabled Tx-Pause: Disabled
  Link up/down SNMP trap filter setting: Enabled
  Egress Port Rate:      128 Kbps, Max Burst Size: 200 Kb
  Broadcast Rate:        No-limit
  Multicast Rate:        No-limit
  Unknown Dest Mac Rate: No-limit
  QoS Profile:   QP3 configured by user
  Ingress Rate Shaping :      Unsupported
  Ingress IPTOS Examination:  Disabled
  Ingress 802.1p Examination: Enabled
  Ingress 802.1p Inner Exam:  Disabled
  Egress IPTOS Replacement:   Disabled
  Egress 802.1p Replacement:  Disabled
  NetLogin:      Enabled
  NetLogin authentication mode: MAC based
  NetLogin port mode:      MAC based VLANs
  Smart redundancy:      Enabled
```

```
Software redundant port:      Disabled
autopolarity:                Enabled
```

show ports packet

```
show ports {mgmt | <port_list>} packet {no-refresh}
```

Description

Displays a snapshot or real-time histogram of packet statistics.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number or range of ports, the system displays information for all ports; if you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command. To clear the counters, use the `clear counters ports` command.

This status information may be useful for your technical support representative if you have a network problem.

The following packet statistics are displayed:

- Port number
- Link state
- Packet size

Example

The following command displays packet statistics for slot 1, port 1, slot 2, port 1, and slot 5, ports 1 through 8:

```
show ports 1:1, 2:1, 5:1-5:8 packet
```

Following is sample output from this command:

```
Port   Link           Packet Sizes
      State  0-64   65-127  128-255  256-511  512-1023  1024-1518  Jumbo
=====
1:1   A         0       0       0       0       0       0       0
```

2:1	R	0	0	0	0	0	0	0
5:1	R	0	0	0	0	0	0	0
5:2	R	0	0	0	0	0	0	0
5:3	R	0	0	0	0	0	0	0
5:4	R	0	0	0	0	0	0	0
5:5	R	0	0	0	0	0	0	0
5:6	R	0	0	0	0	0	0	0
5:7	R	0	0	0	0	0	0	0
5:8	R	0	0	0	0	0	0	0

=====

Link State: A-Active, R-Ready, NP-Port not present, L-Loopback

show ports redundant

show ports redundant

Description

Displays detailed information about redundant ports.

Syntax

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information on software-controlled redundant ports on the switch:

show ports redundant

Following is sample output from this command:

```
Primary: *1:1          Redundant: 3:1, Link on/off option: OFF
Flags: (*)Active, (!) Disabled, (g) Load Share Group
```

show ports sharing

show ports sharing

Description

Displays port load-sharing groups, or link aggregation groups (LAGs).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Output from this command displays the following information:

- **Config Master**—The port that is configured as the master logical port of the link aggregation group (LAG). This number is also the LAG group ID.
- **Current Master**—In LACP, this is the port that is currently the LAG group ID, or master logical port for the LAG.
- **Agg Control**—This is the aggregation control for the specified LAG; it can be either static, LACP or health-check. In LACP, it is the aggregation control for the specified LAG.
- **Ld Share Algorithm**—The algorithm used for the link aggregation. The available link aggregation algorithms vary among platforms; see the *NETGEAR 8800 User Manual* for more information.
- **Ld Share Group**—The specific ports that belong to each LAG, or the port numbers in the trunk. A port can belong to only one LAG, either static or dynamic.
- **Agg Mbr**—In LACP, this shows whether the port has been added to the aggregator or not; it will be either Y for yes or - for no.
- **Link State**—This is the current status of the link
- **Link Up transitions**—Number of times the link has cycled through being up, then down, then up.

Example

The following is an example display for an 8800 switch that uses a custom load sharing algorithm

```
BD-8810.8 # show port sharing
```

```
Load Sharing Monitor
```

Config Master	Current Master	Agg Control	Ld Share Algorithm	Ld Share Group	Agg Mbr	Link State	Link Up Transitions
2:1	2:1	Static	L2	2:1	Y	A	1
			L2	2:2	Y	A	1
3:1	3:1	Static	L3_L4	3:1	Y	A	1
			L3_L4	3:2	Y	A	1
4:1	4:1	Static	custom	4:1	Y	A	1
			custom	4:2	Y	A	1

```
Link State: A-Active, D-Disabled, R-Ready, NP-Port not present, L-Loopback
```

```
Load Sharing Algorithm: (L2) Layer 2 address based, (L3) Layer 3 address based
```

(L3_L4) Layer 3 address and Layer 4 port based
 (custom) User-selected address-based configuration

Custom Algorithm Configuration: ipv4 source-only, xor

Note - Layer 4 ports are not used for distribution for traffic ingressing

MSM-G8X I/O ports and ports on G48T, G48P, G24X, and 10G4X modules.

- The 'custom' algorithm is not used for traffic ingressing on current slot 1, 2, 3, 5 and 10. Refer to XOS Command Reference.

Number of load sharing trunks: 3

show port transceiver information

show port <port-list> transceiver information

Description

Displays basic information about the optical transceiver.

Syntax Description

port-list	Specifies the port number(s).
-----------	-------------------------------

Default

N/A.

Usage Guidelines

Digital Diagnostic Monitoring Interface (DDMI) provides critical system information about 10G XFP optical modules. Use this command to monitor the condition of the XFP modules.

If you try to execute this command on one of the ports in the port list that is non-compliant with DDMI, the following error message is displayed and the command does not go through:

Port 3:1 This command is not supported on this port. All ports and transceiver of the ports requested in the command need to support DDMI.

If you try to execute this command on one of the ports in the port list on which the transceiver is non-compliant with DDMI, the following error message is displayed:

Port 3:1 This media/transceiver does not support enhanced digital diagnostic monitoring interface (DDMI). All ports and transceiver of the ports requested in the command need to support DDMI.

For more detailed information, use the `show port transceiver information detail` command.

Example

The following display shows output for the command `show port 1:1-2 transceiver information`:

```
BD-8810.2 # sh port 1:1-2 transceiver information
```

```

Port          Temp      TxPower  RxPower  TxBiasCurrent  Voltage-Aux1  Voltage-Aux2
      (Celcius)    (dBm)    (dBm)    (mA)           (Volts)       (Volts)
=====
1:1          30.60     -25.20   -18.70    0.40           5.09          5.07
1:2          30.60     -25.20   -18.70    0.40           5.09          N/A
=====
N/A indicates that the parameter is not applicable
to the optics connected to the port

```

show port transceiver information detail

show port <port-list> transceiver information detail

Description

Displays detailed information about the optical transceiver.

Syntax Description

port-list	Specifies the port number(s).
-----------	-------------------------------

Default

N/A.

Usage Guidelines

Digital Diagnostic Monitoring Interface (DDMI) provides critical system information about 10G XFP optical modules. Use this command to monitor the condition of the XFP modules.

If you try to execute this command on one of the ports in the port list that is non-compliant with DDMI, the following error message is displayed and the command does not go through:

```
Port 3:1 This command is not supported on this port. All ports and transceiver of the ports
requested in the command need to support DDMI.
```

If you try to execute this command on one of the ports in the port list on which the transceiver is non-compliant with DDMI, the following error message is displayed:

```
Port 3:1 This media/transceiver does not support enhanced digital diagnostic monitoring
interface (DDMI). All ports and transceiver of the ports requested in the command need to
support DDMI.
```

Example

The following display shows output for the command `show port 1:1-2 transceiver information detail`:

```
BD-8810.2 # sh port 1:1 transceiver information detail
```

```
Port : 1:1
```

Media Type : XFP_LR
 Part Number : 1234567890
 Serial Number : A12345B78

Temp (Celsius) : 30.60
 Low Warn Threshold : 20.60 High Warn Threshold : 45.60
 Low Alarm Threshold : 10.60 High Alarm Threshold : 50.60
 Status : Normal

Tx Power (dBm) : -25.20
 Low Warn Threshold : -35.20 High Warn Threshold : 15.20
 Low Alarm Threshold : -40.20 High Alarm Threshold : 25.20
 Status : Normal

Rx Power (dBm) : -18.70
 Low Warn Threshold : -35.20 High Warn Threshold : 15.20
 Low Alarm Threshold : -40.20 High Alarm Threshold : 25.20
 Status : Normal

Tx Bias Current (mA) : 0.40
 Low Warn Threshold : -35.20 High Warn Threshold : 15.20
 Low Alarm Threshold : -40.20 High Alarm Threshold : 25.20
 Status : Normal

Voltage AUX-1 (Volts) : 5.09
 Low Warn Threshold : 5.01 High Warn Threshold : 6.30
 Low Alarm Threshold : 5.00 High Alarm Threshold : 6.50
 Status : Normal

Voltage AUX-2 (Volts) : 5.07
 Low Warn Threshold : 5.01 High Warn Threshold : 6.30
 Low Alarm Threshold : 5.00 High Alarm Threshold : 6.50
 Status : Normal

Port : 1:2

Media Type : XFP_LR
 Part Number : 1234567890
 Serial Number : A12345B78

Temp (Celsius) : 30.60
 Low Warn Threshold : 20.60 High Warn Threshold : 45.60
 Low Alarm Threshold : 10.60 High Alarm Threshold : 50.60
 Status : Normal

```

Tx Power (dBm)           : -25.20
  Low Warn Threshold    : -35.20    High Warn Threshold   : 15.20
  Low Alarm Threshold   : -40.20    High Alarm Threshold  : 25.20
  Status : Normal

```

```

Rx Power (dBm)           : -18.70
  Low Warn Threshold    : -35.20    High Warn Threshold   : 15.20
  Low Alarm Threshold   : -40.20    High Alarm Threshold  : 25.20
  Status : Normal

```

```

Tx Bias Current (mA)     : 0.40
  Low Warn Threshold    : -35.20    High Warn Threshold   : 15.20
  Low Alarm Threshold   : -40.20    High Alarm Threshold  : 25.20
  Status : Normal

```

```

Voltage AUX-1 (Volts)    : 5.09
  Low Warn Threshold    : 5.01      High Warn Threshold   : 6.30
  Low Alarm Threshold   : 5.00      High Alarm Threshold  : 6.50
  Status : Normal

```

```

Voltage AUX-2 (Volts)    : N/A
  Low Warn Threshold    : N/A       High Warn Threshold   : N/A
  Low Alarm Threshold   : N/A       High Alarm Threshold  : N/A
  Status : N/A

```

show ports utilization

```

show ports {mgmt | <port_list> | stack-ports <stacking-port-list>} utilization {bandwidth |
bytes | packets}

```

Description

Displays real-time port utilization information. The total utilization displays as real-time information, constantly refreshing. and the parameter displays show a snapshot of the activity on the port when you issue the command.

Syntax Description

mgmt	Specifies the management port.
port_list	Specifies one or more ports or slots and ports.
stacking-port-list	Specifies one or more stacking slots and ports.
bandwidth	Specifies port utilization as percentage of bandwidth.
bytes	Specifies port utilization in bytes per second.
packets	Specifies port utilization in packets per second.

Default

N/A.

Usage Guidelines

The software continuously monitors port utilization and calculates bandwidth as a function of each port's maximum link capacity.

The total utilization display presents real-time statistics. Use the <spacebar> to toggle the real-time displayed information for packets, bytes, and bandwidth in that order. When you use a parameter (packets, bytes, or bandwidth) with the command, the display for the specified type shows a snapshot per port when you issued the command. When the `show ports utilization` command is run with the `bandwidth`, `bytes`, or `packets` options, the command may need to be repeated a few times in order for the NETGEAR 8800 software to gather enough statistics to calculate appropriate values.

If you do not specify a port number or range of ports, port utilization information is displayed for all ports.

This status information may be useful for your technical support representative if you have a network problem.

Example

The following command displays utilization statistics for port 1 on a stand-alone switch:

```
show ports 1 utilization
```

The following command displays utilization statistics for slot 3, port 1:

```
show ports 3:1 utilization
```

The following example shows sample output from the `show ports utilization packets` command:

```
Link Utilization Averages                               Mon Oct  6 22:38:25 2008
Port    Link    Rx          Peak Rx          Tx          Peak Tx
      State pkts/sec    pkts/sec    pkts/sec    pkts/sec
=====
1:1     A         47           191             0             0
```

```

1:2      A          0          0          0          0
2:1      R          0          0          0          0
2:2      R          0          0          0          0
3:1      R          0          0          0          0
3:2      R          0          0          0          0
4:1      R          0          0          0          0
4:2      R          0          0          0          0
5:1      R          0          0          0          0
5:2      R          0          0          0          0
6:1      R          0          0          0          0
6:2      R          0          0          0          0
7:1      R          0          0          0          0
7:2      R          0          0          0          0

```

```

=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up D->page down ESC->exit

```

Note: Use the <spacebar> to toggle this real-time display for all ports from packets to bytes to bandwidth, in that order.

The following example shows sample output from the `show ports utilization bytes` command:

```

Link Utilization Averages                               Mon Oct  6 22:39:22 2008
Port      Link      Rx          Peak Rx          Tx          Peak Tx
          State    bytes/sec    bytes/sec    bytes/sec    bytes/sec
=====
1:1      A          0          0          0          63
1:2      A          0          63         63         63
2:1      R          0          0          0          0
2:2      R          0          0          0          0
3:1      R          0          0          0          0
3:2      R          0          0          0          0
4:1      R          0          0          0          0
4:2      R          0          0          0          0
5:1      R          0          0          0          0
5:2      R          0          0          0          0
6:1      R          0          0          0          0
6:2      R          0          0          0          0
7:1      R          0          0          0          0
7:2      R          0          0          0          0

```

```

=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback

```

Spacebar->toggle screen U->page up D->page down ESC->exit

The following example shows sample output of the `show ports utilization bandwidth` command:

```
Link Utilization Averages                               Mon Oct  6 22:39:46 2008
Port      Link      Link  Rx      Peak Rx      Tx      Peak Tx
          State    Speed % bandwidth % bandwidth % bandwidth % bandwidth
=====
1:1       A        100    0.00    0.03    0.00    0.00
1:2       A        100    0.00    0.00    0.00    0.00
2:1       R         0    0.00    0.00    0.00    0.00
2:2       R         0    0.00    0.00    0.00    0.00
3:1       R         0    0.00    0.00    0.00    0.00
3:2       R         0    0.00    0.00    0.00    0.00
4:1       R         0    0.00    0.00    0.00    0.00
4:2       R         0    0.00    0.00    0.00    0.00
5:1       R         0    0.00    0.00    0.00    0.00
5:2       R         0    0.00    0.00    0.00    0.00
6:1       R         0    0.00    0.00    0.00    0.00
6:2       R         0    0.00    0.00    0.00    0.00
7:1       R         0    0.00    0.00    0.00    0.00
7:2       R         0    0.00    0.00    0.00    0.00
=====

> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up D->page down ESC->exit
```

show sharing health-check

`show sharing health-check`

Description

Displays the configured health check LAGs on a switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the health-check LAGs that have been configured on the switch.

Example

The following is sample output from this command:

```
(debug) BD-8810.1 # show sharing health-check
```

Group	Member Port	Agg Mbr	Admin State	Track IP Addr	Track TCP Port	Miss Freq	State	Dn	Up
2:8	2:1*	Y	En	30.1.1.1	23	3	3 Up	0	1
	2:2*	Y	En	30.1.1.2	23	3	3 Up	0	1
	2:3*	Y	En	30.1.1.3	23	3	3 Up	0	1
	2:8*	-	En	30.1.1.8	80	3	10 Down	0	0
	2:11*	Y	-	-	-	-	-	-	-
	2:12*	-	En	44.1.3.2	80	3	4 Down	0	0
	2:16	-	En	30.1.1.16	80	3	10 Dis	0	0
2:20	2:20*	Y	En	192.1.1.1	80	10	3 Up	0	1
	2:21*	Y	En	192.1.1.2	80	10	3 Up	0	1

Member Port Flags: (*)Active, (!) Disabled

show slot

```
show slot {<slot> {detail} | detail }
```

Description

Displays the slot-specific information.

Syntax Description

slot	Specifies a slot on the switch.
detail	Specifies detailed port information.

Default

N/A.

Usage Guidelines

The `show slot` command displays the following information:

- The slot number
- The type of module installed in the slot
- The type of module configured for the slot
- The state of the module, whether the power is down, if the module is operational, if a diagnostic being run, if there is a mismatch between the slot configuration and the module in the slot

- The number of ports on the module
- The current number of times the module has been restarted after a failure and the configured restart-limit.

Note: You may see slightly different information displayed depending on the platform and configuration you are using.

If you do not specify a slot number, information for all slots is displayed.

The display also includes a notice of insufficient power, should that arise.

The `show slot` command displays the following states, among others:

- Empty (This is also displayed if you have a module in the chassis that is unsupported by the current software you are running.)
- Down
- Power ON
- Powered OFF
- Booting
- Initializing
- VLAN sync
- FDB sync
- ACL sync
- RT sync
- Operational

The following example displays module information for all slots:

Slots	Type	Configured	State	Ports	Flags

Slot-1			Empty	0	
Slot-2	XCM8824F	XCM8824F	Operational	24	M S
Slot-3			Empty	0	
Slot-4			Empty	0	
Slot-5	XCM8808X	XCM8808X	Operational	8	M S
Slot-6			Empty	0	
Slot-7		XCM8848T	Empty	48	
Slot-8	XCM8848T		Operational	48	M S
Slot-9	XCM8808X	XCM8808X	Powered OFF	8	SI
Slot-10			Empty	0	
MSM-A	XCM88S1		Operational	0	S
MSM-B			Empty	0	

Flags : M - Backplane link to Master MSM is Active

B - Backplane link to Backup MSM is also Active
 D - Slot Disabled, S - Slot Secured
 I - Insufficient Power (refer to "show power budget")

The following example displays module information for a specified slot on a NETGEAR 8810 switch:

```
XCM8810.3 # show slot 2
```

Slot-2 information:

```
State:           Operational
Download %:      100
Flags:           MB
Restart count:   0 (limit 5)
Serial number:   800114-00-04 04364-00013
Hw Module Type:  xcm8848T
SW Version:      12.1.0.56
SW Build:        v1210b56
Configured Type: G48P
Ports available: 48
Recovery Mode:   Reset
```

Flags : M - Backplane link to Master is Active
 B - Backplane link to Backup is also Active
 D - Slot Disabled, S - Slot Secured
 I - Insufficient Power (refer to "show power budget")

unconfigure ports display string

```
unconfigure ports <port_list> display-string
```

Description

Clears the user-defined display string from one or more ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
-----------	---

Default

N/A.

Usage Guidelines

This command removes the display string that you configured using the `configure ports display-string` command.

Example

The following command clears the user-defined display string from slot 2, port 4:

```
unconfigure ports 2:4 display-string
```

unconfigure ports redundant

```
unconfigure ports <port_list> redundant
```

Description

Clears a previously configured software-controlled redundant port.

Syntax Description

port_list	This refers to the primary port of the redundant pair and specifies one or more ports or slots and ports.
-----------	---

Default

N/A.

Usage Guidelines

The list of port numbers or the port display string specifies the primary port(s).

Example

The following command unconfigures a software-controlled redundant port:

```
unconfigure ports 2:3 redundant
```

Commands for Configuring LLDP

6

This chapter describes commands for doing the following:

- Configuring LLDP
- Managing LLDP
- Displaying LLDP information

For an introduction to LLDP, see the *NETGEAR 8800 User Manual*.

configure lldp med fast-start repeat-count

```
configure lldp med fast-start repeat-count <count>
```

Description

The fast-start feature is automatically enabled when you enable the LLDP MED capabilities TLV. This command configures how many times, from 1 to 10, the switch sends out an LLDP MED packet with an interval of 1 second.

Syntax Description

count	Specifies the number of times the switch transmits LLDP MED TLVs each second (once it detects a neighbor transmitting LLDP MED TLVs). The range is 1 to 10.
-------	---

Default

3.

Usage Guidelines

When the switch detects a MED-capable device, this count determines how many times the switch sends a LLDP MED TLVs with an interval of 1 second. The fast-start feature enables the MED-capable device to quickly learn information; this command changes the value from the default 3. The fast-start feature is automatically enabled when you enable the LLDP MED capabilities TLV.

Note: After you configure the LLDP MED capability TLV, the fast-start feature automatically runs. To configure the LLDP MED capability TLV, use the `configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities` command.

Example

The following command configures fast learning on the switch to a value of 2:

```
configure lldp med fast-start repeat-count 2
```

configure lldp ports management-address

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] management-address
```

Description

Configures the LLDP port to advertise or not to advertise management address information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

You can add only one management address TLV per LLDPDU and the information must be the IP address configured on the management VLAN. If no IP address is assigned to the management VLAN, the system sends the system MAC address. LLDP does not send out IPv6 addresses in this field.

Example

The following command advertises the management address information for port 1:5:

```
configure lldp ports 1:5 advertise management-address
```

configure lldp ports port-description

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] port-description
```

Description

Configures the LLDP port to advertise or not advertise port description information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

N/A.

Example

The following command configures port 1:7 to not advertise the port description information to neighbors:

```
configure lldp ports 1:7 no-advertise port-description
```

configure lldp ports system-capabilities

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] system-capabilities
```

Description

Configures the LLDP port to advertise or not to advertise its system capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When at least one VLAN exists with more than two ports, bridging is sent to enabled.

When at least one VLAN on the switch has IP forwarding enabled, the system automatically sets the router bit.

Example

The following command configures all ports to advertise system capability information to neighbors:

```
configure lldp ports all advertise system-capabilities
```

configure lldp ports system-description

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] system-description
```

Description

Configures the LLDP port to advertise or not to advertise its system description to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

Advertise.

Usage Guidelines

Although not mandatory according to the standard, this TLV is included in the LLDPDU by default when you enable LLDP.

When enabled, the system sends the following image (from the `show version` command) in the system description TLV:

```
NETGEAR 8800 version 11.2.0.12 v1120b12 by release-manager
on Fri Mar 18 16:01:08 PST 2005
```


Example

The following command configures port 1:4 through port 1:8 to not advertise the system description information to neighbors:

```
configure lldp ports 1:4 - 1:8 no-advertise system-description
```

configure lldp ports system-name

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] system-name
```

Description

Configures the LLDP port to advertise or not to advertise its system name to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

N/A.

Example

The following command configures port 1:6 to advertise the system name to neighbors:

```
configure lldp ports 1:4 - 1:8 advertise system-name
```

configure lldp ports vendor-specific dot1 port-vlan-ID

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot1  
port-vlan-ID
```

Description

Configures the LLDP port to advertise or not advertise port vlan ID information to its neighbors. This allows a VLAN bridge port to advertise the port VLAN identifier that is associated with untagged or priority-tagged frames.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

The port VLAN ID TLV allows the port to transmit the VLAN ID associated with untagged VLANs. There can be only one port VLAN ID in each LLDPDU.

If no untagged VLANs are configured on the specified port, the TLV is not added to the LLDPDU, even if you configured this to advertise.

Example

The following command configures all ports to advertise port vlan ID information to neighbors:

```
configure lldp ports all advertise vendor-specific dot1 port-vlan-ID
```

configure lldp ports vendor-specific dot1 port-protocol-vlan-ID

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot1  
port-protocol-vlan-ID {vlan [all | <vlan_name>]}
```

Description

Configures the LLDP port to advertise or not advertise port VLAN information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
all	Specifies all VLANs on the port.
vlan_name	Specifies the VLAN on the port that you want to advertise.

Default

No advertise.

Usage Guidelines

When configured to advertise, the switch inserts a port and protocol VLAN ID TLV for each VLAN configured on the ports. The port and protocol VLAN ID TLV allows the port to advertise if it supports protocol and/or tagged VLANs, along with the associated tagged values. A separate TLV is sent for each VLAN that you want to advertise.

By default, once you configure this TLV, the system sends all protocol-based VLANs on the port. However, the LLDPDU cannot exceed 1500 bytes, so you should configure the port to advertise only the *specified* VLANs.

Note: The total LLDPDU size is 1500 bytes; any TLVs after that limit are dropped.

This TLV does not send information on the type of protocol that the VLAN has enabled; it just says whether the port is enabled or disabled for protocol-based VLANs. As NETGEAR devices are always capable of supporting protocol-based VLANs, once you configure this TLV, the system *always* advertises support these VLANs.

Example

The following command configures all ports to advertise port and protocol VLAN information to neighbors for all VLANs on all ports:

```
configure lldp ports all advertise vendor-specific dot1 port-protocol-vlan-id
```

configure lldp ports vendor-specific dot1 vlan-name

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot1  
vlan-name {vlan [all | <vlan_name>]}
```

Description

Configures the LLDP port to advertise or not advertise VLAN name information to its neighbors. Use this TLV to advertise information for the tagged VLANs you want to specify on the port. This allows an IEEE 802.1Q-compatible 802 LAN station to advertise the assigned name of any VLAN with which it is configured.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.

no-advertise	Specifies not to send the information to neighbors.
vlan	Specifies all VLANs on the port.
vlan_name	Specifies the VLAN on the port that you want to advertise.

Default

No advertise.

Usage Guidelines

The VLAN name TLV sends the VLAN name and the tag used; it associates a name to a tag for the specified VLAN. This allows an IEEE 802.1Q-compatible 802 LAN station to advertise the assigned name of any VLAN with which it is configured.

You can enable this TLV for tagged and untagged VLANs. When you enable this TLV for tagged VLANs, the TLV advertises the IEEE 802.1Q tag for that VLAN. (For untagged VLANs, the internal tag is advertised.) You can specify exactly *which* VLANs to advertise.

When configured to advertise, the switch inserts a VLAN name TLV for every VLAN configured on the ports. By default, once you configure this TLV, the system sends all VLAN names on the port. However, each VLAN name can require up to 32 bytes and the LLDPDU cannot exceed 1500 bytes, so you should configure the port to advertise only the *specified* VLANs, using the keyword `vlan_name`.

Note: The total LLDPDU size is 1500 bytes; any TLVs after that limit are dropped.

Example

The following command configures all ports to not advertise VLAN name information to neighbors:

```
configure lldp ports all no-advertise vendor-specific dot1 vlan-name
```

configure lldp ports vendor-specific dot3 link-aggregation

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot3
link-aggregation
```

Description

Configures the LLDP port to advertise or not advertise link-aggregation capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, this TLV is added to each LLDP port LLDPDU indicating the link-aggregation capabilities, status, and value of the master port of the load-sharing group.

Example

The following command configures port 1:12 to not advertise link-aggregation capabilities to neighbors:

```
configure lldp ports 1:12 no-advertise vendor-specific dot3 link-aggregation
```

configure lldp ports vendor-specific dot3 mac-phy

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot3
mac-phy
```

Description

Configures the LLDP port to advertise or not advertise MAC and physical layer capabilities to its neighbors. The capabilities include duplex and bit rate.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, the system add information about the speed capabilities, as well as autonegotiation support and status, of the LLDP port.

Example

The following command configures all ports to advertise MAC/PHY capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific dot3 mac-phy
```

configure lldp ports vendor-specific dot3 max-frame-size

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot3
max-frame-size
```

Description

Configures the LLDP port to advertise or not advertise its maximum frame size to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When jumbo frames are not enabled on the specified port, the TLV reports a value of 1518 once you configure it to advertise. If jumbo frames are enabled, the TLV inserts the configured value for the jumbo frames.

Example

The following command configures ports 1:12 and 1:13 to advertise the maximum frame size to neighbors:

```
configure lldp ports 1:12 - 1:13 advertise vendor-specific dot3 max-frame-size
```

configure lldp ports vendor-specific dot3 power-via-mdi

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific dot3
power-via-mdi
```

Description

Configures the LLDP port to advertise or not advertise Power over Ethernet (PoE) capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, the system includes this TLV. NETGEAR recommends enabling this TLV only on PoE-capable ports.

The following information is transmitted for LLDP ports with this TLV:

- Support PoE or not
- Port class
 - Power sourcing equipment (PSE)
 - Powered device (PD)
- Power pairs used to supply power
 - Signal
 - Spare
- Power status
- Support pairs control or not
- Power class
 - Class0
 - Class1
 - Class2
 - Class2
 - Class3
 - Class4

Note: For more information on advertising power support, see the `configure lldp ports vendor-specific med power-via-mdi` command.

Example

The following command configures all ports to advertise power capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific dot3 power-via-mdi
```

configure lldp ports vendor-specific med capabilities

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities
```

Description

Configures the LLDP port to advertise or not advertise MED capabilities. This TLV must be enabled before any of the other MED TLVs can be enabled. Also, this TLV must be set to no-advertise after all other MED TLVs are set to no-advertise.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

This command enables the LLDP media endpoint discovery (MED) capabilities TLV, which allows LLDP-MED network connectivity devices to definitively determine that particular endpoints support LLDP MED, and if so, to discover which LLDP MED TLVs the particular endpoint devices are capable of supporting and to which specific device class the device belongs to.

This TLV must be enabled before any of the other MED TLVs can be enabled; and this TLV must be set to no-advertise after all other MED TLVs are set to no-advertise.

As with all the LLDP MED TLVs, the switch sends this TLV *only* after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

Note: Network connectivity devices wait to detect LLDP MED TLVs from endpoints before they send out LLDP MED TLVs; so L2 network connectivity devices do not exchange LLDP MED messages.

The following information is included in the LLDP MED capabilities TLV when it is transmitted:

- The supported LLDP MED TLVs—For NETGEAR 8800 devices, these are capabilities, network policy, location, and extended power (extended power only advertised only on PoE-capable ports).
- The MED device type—For NETGEAR 8800 devices, this is advertised as a network connectivity device (set to 4).

Example

The following command configures all ports to advertise MED capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific med capabilities
```

configure lldp ports vendor-specific med location-identification

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med
location-identification [coordinate-based <hex_value> | civic-based <hex_value> | ecs-elin
<elin>]
```

Description

Configures the LLDP port to advertise or not advertise MED location information. You configure up to 3 different location identifiers.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
advertise	Specifies to send the information to neighbors.
coordinate-based	Specifies using the coordinate-based location identifier. This value is exactly 16 bytes long; see RFC 3825 for details.
hex_value	Enter a hexadecimal value with each byte separated by a colon. Or, you can obtain this value from a network management application. Note: This parameter is not used when the no-advertise parameter is configured.
civic-based	Specifies using the civic-based location identifier. This value must have a minimum length of 6 bytes; see RFC3825 for details.

ecs-elin	Specifies using the ecs location identifier. (Emergency Call Service, as defined in the TIA-TSB-146.)
elin	Enter a numerical string; the range is 10 to 25 characters. Or, you can obtain this value from a network management application. (See the TIA-TSB-146 standard for a definition of these numbers; also, the network management application must be able to handle the LLDP MED MIB.)
	Note: This parameter is not used when the no-advertise parameter is configured.

Default

No advertise.

Usage Guidelines

You might need to use a specific format for your specific VoIP implementation; see the VoIP manufacturer's manual for details.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV *only* after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

Example

The following command configures all ports to advertise MED location information to neighbors using the ECS format:

```
configure lldp ports all advertise vendor-specific med location-identification ecs-elin
423233455676
```

configure lldp ports vendor-specific med policy application

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med policy
application [voice | voice-signaling | guest-voice | guest-voice-signaling | softphone-voice |
video-conferencing | streaming-video | video-signaling] vlan <vlan_name> dscp <dscp_value>
{priority-tagged}
```

Description

Configures the LLDP port to advertise or not advertise MED network policy TLVs. This TLV advertises VLAN configuration and associated Layer 2 and Layer 3 attributes that apply for a set of specific applications on that port. You can advertise up to 8 TLVs, each for a specific application, per port/VLAN. Each application type can exist only *once* per port. This TLV tells the endpoint the specific VLAN to use for the specific application, along with its unique priority.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
advertise	Specifies to send the information to neighbors.
voice	Specifies voice application on specified port/VLAN(s).
voice-signaling	Specifies voice signaling application on specified port/VLAN(s).
guest-voice	Specifies guest voice application on specified port/VLAN(s).
guest-voice-signaling	Specifies guest voice signaling application on specified port/VLAN(s).
softphone-voice	Specifies soft phone voice application on specified port/VLAN(s).
video-conferencing	Specifies videoconferencing application on specified port/VLAN(s).
streaming-video	Specifies streaming video application on specified port/VLAN(s).
video-signaling	Specifies video signaling application on specified port/VLAN(s).
vlan_name	Specifies the VLAN the specified application is using.
	Note: This parameter does not apply when the no-advertise parameter is configured.
dscp_value	Specifies the DSCP value for the specified application. This is a 6-bit value from 0 to 63.
	Note: This parameter does not apply when the no-advertise parameter is configured.
priority-tagged	Use this if you want priority tagging, and the VLAN is configured as untagged on the port. (The endpoint sends out frames for the specified application with a tag of 0.)
	Note: This parameter does not apply when the no-advertise parameter is configured.

Default

No advertise.

Usage Guidelines

This command enables the LLDP MED network policy TLV, which allows network connectivity devices and endpoint devices to advertise VLAN configuration and associated Layer 2 and Layer 3 attributes that apply for a set of specific application on that port. This TLV can be enabled on a per port/VLAN basis. Each application type can exist only once on a port.

You can enable the transmission of a TLV policy for each application. A maximum of 8 TLVs can be enabled, and each can have a unique DSCP value and/or priority tagging.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV *only* after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

The following information is transmitted for LLDP ports with this TLV:

- Application type

Used as configured.

- Unknown policy flag

Set to 0.

- Tagged flag

Set to tagged for tagged VLANs; set to untagged for untagged VLANs. By default, set to 0.

- VLAN ID

Copied from the VLAN. However, if you configure the `priority-tagged` parameter, this value is set to 0.

- Layer 2 priority

Copied from the VLAN priority.

- DSCP value

Uses the value configured in the `dscp` parameter.

Note: See the documentation provided by the manufacturer of connected devices regarding values.

Example

The following command configures all ports to advertise videoconferencing on the VLAN video with a DSCP of 7 to neighbors:

```
configure lldp ports all advertise vendor-specific med policy application video-conferencing
vlan video dscp 7
```

configure lldp ports vendor-specific med power-via-mdi

```
configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med
power-via-mdi
```

Description

Configures the LLDP port to advertise or not advertise MED power requirement details. This TLV can only be enabled on a PoE-capable port and is used for advanced power management between the MED network connectivity and endpoint devices.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When enabled, this LLDP MED TLV advertises fine-grained power requirement details about PoE settings and support. This TLV can be enabled only on a PoE-capable port; the switch returns an error message if this TLV is configured for a non-PoE-capable port.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | <port_list>] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV *only* after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

Note: For additional information on power support, see the `configure lldp ports vendor-specific dot3 power-via-mdi` command.

The following information is transmitted for LLDP MED PoE-capable ports with this TLV:

- Power type

Set to PSE.

- Power source

Set to primary power source.

- Power priority

Taken from PoE port configuration.

- Power value

Taken from PoE port configuration.

Example

The following command configures all ports to advertise MED power information to neighbors:

```
configure lldp ports all advertise vendor-specific med power-via-mdi
```

configure lldp reinitialize-delay

```
configure lldp reinitialize-delay <seconds>
```

Description

Configures the delay before the receive state machine is reinstalled once the LLDP transmit mode has been disabled.

Syntax Description

seconds	Specifies the delay that applies to the reinitialization attempt. The range is 1 to 10 seconds.
---------	---

Default

2 seconds.

Usage Guidelines

N/A.

Example

The following command configures a reinitialization delay of 10 seconds:

```
configure lldp reinitialize-delay 10
```

configure lldp snmp-notification-interval

```
configure lldp snmp-notification-interval <seconds>
```

Description

Configures the allowed interval at which Simple Network Management Protocol (SNMP) notifications are sent.

Syntax Description

seconds	Specifies the interval at which LLDP SNMP notifications are sent. The range is 5 to 3600 seconds.
---------	---

Default

5 seconds.

Usage Guidelines

This is a global timer. If one port sends a notification, no notifications for other ports go out for the configured interval.

Example

The following command configures an interval of 60 seconds for LLDP SNMP notifications:

```
configure lldp snmp-notification-interval 60
```

configure lldp transmit-delay

```
configure lldp transmit-delay [ auto | <seconds>]
```

Description

Configures the delay time between successive frame transmissions initiated by a value change or status change in any of the LLDP local systems Management Information Base (MIB). The `auto` option uses a formula ($0.25 * \text{transmit-interval}$) to calculate the number of seconds.

Syntax Description

<code>auto</code>	Uses the formula ($0.25 * \text{transmit-interval}$) to calculate the seconds.
<code>seconds</code>	Specifies the interval at which LLDP notifications are sent. The range is 1 to 8291.

Default

2 seconds.

Usage Guidelines

This is the timer between triggered updates.

Example

The following command configures the delay between LLDP frame transmissions for triggered updates to be automatically calculated:

```
configure lldp transmit-delay auto
```

configure lldp transmit-hold

```
configure lldp transmit-hold <hold>
```

Description

Calculates the actual time-to-live (TTL) value used in the LLDPDU messages. The formula is `transmit-interval * transmit-hold`; by default the TTL value is (30*4) 120 seconds.

Syntax Description

hold	Used to calculate the TTL value; the range is 2 to 10.
------	--

Default

4.

Usage Guidelines

N/A.

Example

The following command configures the transmit-hold value (which is used to calculate the TTL of the LLDP packets) to 5:

```
configure lldp transmit-hold 5
```

configure lldp transmit-interval

```
configure lldp transmit-interval <seconds>
```

Description

Configures the periodic transmittal interval for LLDPDUs.

Syntax Description

seconds	Specifies the time between LLDPDU transmissions. The range is 5 to 32768.
---------	---

Default

30 seconds.

Usage Guidelines

N/A.

Example

The following command configures a transmittal interval of 20 seconds for LLDPDUs.

```
configure lldp transmit-interval 20
```


disable lldp ports

```
disable lldp ports [all | <port_list>] {receive-only | transmit-only}
```

Description

Disables LLDP transmit mode, receive mode, or transmit and receive mode on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
receive-only	Specifies that only the receive mode for LLDP is disabled.
transmit-only	Specifies that only the transmit mode for LLDP is disabled.

Default

Disabled.

Usage Guidelines

If you do not specify an option, both LLDP modes (transmit and receive) are disabled.

Example

The following example disables the LLDP receive mode on ports 1:2 to 1:6.

```
disable lldp ports 1:2-1:6 receive-only
```

disable snmp traps lldp

```
disable snmp traps lldp {ports [all | <port_list>]}
```

Description

Disables the sending of LLDP-specific SNMP traps on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system stops sending LLDP traps from all ports on the switch.

Example

The following example disables sending LLDP SNMP traps on all switch ports:

```
disable snmp traps lldp ports all
```

disable snmp traps lldp-med

```
disable snmp traps lldp-med {ports [all | <port_list>]}
```

Description

Disables the sending of LLDP MED-specific SNMP traps on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system stops sending LLDP MED traps from all ports on the switch.

Example

The following example disables sending LLDP MED SNMP traps on all switch ports:

```
disable snmp traps lldp-med ports all
```

enable lldp ports

```
enable lldp ports [all | <port_list>] {receive-only | transmit-only}
```

Description

Enables LLDP transmit mode, receive mode, or transmit and receive mode. If the transmit-only or receive-only option is not specified, both transmit and receive modes are enabled.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
receive-only	Specifies that the port only receives LLDP messages.
transmit-only	Specifies that the port only transmits LLDP messages.

Default

Disabled.

Usage Guidelines

If you do not specify an option, the port is enabled to both transmit and receive LLDP messages.

Once the port is enabled for LLDP in one mode and you issue another `enable lldp ports` command for another mode, that second mode replaces the original mode. For example, you might originally enable several ports to only receive LLDP messages and then want those ports to both receive and transmit LLDP messages. In that case, you issue the `enable lldp ports` command with no variables (and the receive-and-transmit mode replaces the receive-only mode).

To verify the port setting for LLDP, use the `show lldp {port [all | <port_list>]} {detailed}` command.

Example

The following example enables LLDP transmit and receive mode on port 1:4.

```
enable lldp port 1:4
```

enable snmp traps lldp

```
enable snmp traps lldp {ports [all | <port_list>]}
```

Description

Enables the transmission of LLDP SNMP trap notifications.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

Note: To enable SNMP traps for LLDP MED TLVs, you must issue a separate command; use the `enable snmp traps lldp-med {ports [all | <port_list>]}`.

If you do not specify any ports, the system sends LLDP traps for all ports.

Example

The following command enables LLDP SNMP traps for all ports:

```
enable snmp traps lldp ports all
```

enable snmp traps lldp-med

```
enable snmp traps lldp-med {ports [all | <port_list>]}
```

Description

Enables the transmission of LLDP SNMP trap notifications related to LLDP MED extension TLVs.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system sends LLDP-MED traps for all ports.

Example

The following command enables LLDP-MED SNMP traps for all ports:

```
enable snmp traps lldp-med ports all
```

show lldp

```
show lldp {port [all | <port_list>]} {detailed}
```

Description

Displays LLDP configuration information for the specified port or ports. Use the `detailed` keyword to display the configured VLANs on the port and the enabled VLAN-specific TLVs.

Syntax Description

<code>all</code>	Specifies all ports on the switch.
<code>port_list</code>	Specifies one or more ports or slots and ports.
<code>detailed</code>	Shows information on the configured VLANs on the port.

Default

N/A.

Usage Guidelines

Use the `detailed` variable to display information regarding configured VLANs on the ports and any enabled VLAN-specific TLVs.

Example

The following example displays LLDP configuration information for the switch:

```
# show lldp
```

```

LLDP transmit interval      : 30 seconds
LLDP transmit hold multiplier : 4  (used TTL = 120 seconds)
LLDP transmit delay        : 2 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay     : 2 seconds
LLDP-MED fast start repeat count : 4

```

LLDP Port Configuration:

Port	Rx Mode	Tx Mode	SNMP Notification	Optional enabled transmit TLVs
				LLDP 802.1 802.3 MED AvEx
=====				
1:1	Enabled	Enabled	--	--D-- --- ---- CLP- ----
1:2	Enabled	Enabled	L-	--D-- --- ---- C-P- ----
7:1	Enabled	Enabled	LM	--D-- --- ---- CLP- ----
=====				

Notification: (L) lldpRemTablesChange, (M) lldpXMedTopologyChangeDetected

```

LLDP Flags : (P) Port Description, (N) System Name, (D) System Description
              (C) System Capabilities, (M) Mgmt Address
802.1 Flags : (P) Port VLAN ID, (p) Port & Protocol VLAN ID, (N) VLAN Name
802.3 Flags : (M) MAC/PHY Configuration/Status, (P) Power via MDI
              (L) Link Aggregation, (F) Frame Size
MED Flags   : (C) MED Capabilities, (P) Network Policy,
              (L) Location Identification, (p) Extended Power-via-MDI
AvEx Flags  : (P) PoE Conservation Request, (C) Call Server, (F) File Server
              (Q) 802.1Q Framing

```

The following example includes detailed information on the LLDP configuration for port 1:1:

```
# show lldp port 1:1 detailed
```

```

LLDP transmit interval      : 30 seconds
LLDP transmit hold multiplier : 4   (used TTL = 120 seconds)
LLDP transmit delay        : 2 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay     : 2 seconds
LLDP-MED fast start repeat count : 4

```

LLDP Port Configuration:

Port	Rx Mode	Tx Mode	SNMP Notification	Optional enabled transmit TLVs
				LLDP 802.1 802.3 MED AvEx
1:1	Enabled	Enabled	--	--D-- --- ---- CLP- ----
VLAN: Default				----- --- ---- ----
VLAN: voice				----- --- ---- ----
AvEx Call-Server: IP Address(es)=10.0.0.20, 10.0.0.21				
AvEx File-Server: IP Address(es)=10.0.0.20, 10.0.0.21, 10.0.0.22				
AvEx 802.1Q Framing: Mode=tagged				
MED LCI: Location Format=ECS ELIN based				
1234567890				
MED Policy: Application=voice				
VLAN=voice, DSCP=40				

```

=====
Notification: (L) lldpRemTablesChange, (M) lldpXMedTopologyChangeDetected
LLDP Flags : (P) Port Description, (N) System Name, (D) System Description
              (C) System Capabilities, (M) Mgmt Address
802.1 Flags : (P) Port VLAN ID, (p) Port & Protocol VLAN ID, (N) VLAN Name
802.3 Flags : (M) MAC/PHY Configuration/Status, (P) Power via MDI
              (L) Link Aggregation, (F) Frame Size
MED Flags   : (C) MED Capabilities, (P) Network Policy,
              (L) Location Identification, (p) Extended Power-via-MDI
AvEx Flags  : (P) PoE Conservation Request, (C) Call Server, (F) File Server

```

(Q) 802.1Q Framing

show lldp neighbors`show lldp {port [all | <port_list>]} neighbors {detailed}`**Description**

Displays the information related to the LLDP neighbors detected on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.
detailed	Shows detailed information on the neighbors.

Default

N/A.

Usage Guidelines

You must use the `detailed` parameter to display detailed information about the received LLDP TLVs.

Example

The following example displays LLDP neighbor information for all switch ports:

```
# show lldp port all neighbors
```

Port	Neighbor Chassis ID	Neighbor Port ID	TTL	Age
1:2	00:04:96:26:A4:70	1:1	120	7
2:6	(5.1)10.201.41.146	00:04:0D:EC:EA:5C	120	3
2:7	(5.1)10.201.41.147	00:04:0D:ED:41:9B	120	3
2:10	00:01:30:F9:9E:80	8:10	120	15

NOTE: The Chassis ID and/or Port ID might be truncated to fit the screen.

The following command lists detailed LLDP neighbor information for all switch ports:

```
# show lldp all neighbors detailed
```

```
-----
LLDP Port 1:2 detected 1 neighbor
  Neighbor: 00:04:96:26:A4:70/1:1, age 12 seconds
    - Chassis ID type: MAC address (4)
```

```

Chassis ID      : 00:04:96:26:A4:70
- Port ID type: ifName (5)
  Port ID       : "1:1"
- Time To Live: 120 seconds
- System Description: "NETGEAR 8800 version 12.0.0.6 v1200b6 by release-ma\
                      nager on Mon Mar 19 00:37:59 PDT 2007"

```

```

-----
LLDP Port 2:6 detected 1 neighbor

```

```

Neighbor: (5.1)10.201.41.146/00:04:0D:EC:EA:5C, age 8 seconds
- Chassis ID type: Network address (5); Address type: IPv4 (1)
  Chassis ID      : 10.201.41.146
- Port ID type: MAC address (3)
  Port ID       : 00:04:0D:EC:EA:5C
- Time To Live: 120 seconds
- System Name: "AVAECEA5C"
- System Capabilities : "Bridge, Telephone"
  Enabled Capabilities: "Bridge, Telephone"
- Management Address Subtype: IPv4 (1)
  Management Address      : 10.201.41.146
  Interface Number Subtype : System Port Number (3)
  Interface Number       : 1
  Object ID String       : "1.3.6.1.4.1.6889.1.69.2.3"
- IEEE802.3 MAC/PHY Configuration/Status
  Auto-negotiation      : Supported, Enabled (0x03)
  Operational MAU Type   : 100BaseTXFD (16)
- MED Capabilities: "MED Capabilities, Network Policy, Inventory"
  MED Device Type : Endpoint Class III (3)
- MED Network Policy
  Application Type  : Voice (1)
  Policy Flags      : Known Policy, Tagged (0x1)
  VLAN ID          : 0
  L2 Priority       : 6
  DSCP Value       : 46
- MED Hardware Revision: "9650D01A"
- MED Firmware Revision: "hb96xxual_20r30s.bin"
- MED Software Revision: "ha96xxual_20r30s.bin"
- MED Serial Number: "06N537900335"
- MED Manufacturer Name: "Avaya"
- MED Model Name: "9650"

```

```

-----
LLDP Port 2:7 detected 1 neighbor

```

```

Neighbor: (5.1)10.201.41.147/00:04:0D:ED:41:9B, age 8 seconds
- Chassis ID type: Network address (5); Address type: IPv4 (1)
  Chassis ID      : 10.201.41.147
- Port ID type: MAC address (3)
  Port ID       : 00:04:0D:ED:41:9B

```


- Time To Live: 120 seconds
- System Name: "AVAED419B"
- System Capabilities : "Telephone"
Enabled Capabilities: "Telephone"
- Management Address Subtype: IPv4 (1)
Management Address : 10.201.41.147
Interface Number Subtype : System Port Number (3)
Interface Number : 1
Object ID String : "1.3.6.1.4.1.6889.1.69.2.5"
- IEEE802.3 MAC/PHY Configuration/Status
Auto-negotiation : Supported, Enabled (0x03)
Operational MAU Type : 100BaseTXFD (16)
- MED Capabilities: "MED Capabilities, Network Policy, Inventory"
MED Device Type : Endpoint Class III (3)
- MED Network Policy
Application Type : Voice (1)
Policy Flags : Known Policy, Tagged (0x1)
VLAN ID : 0
L2 Priority : 6
DSCP Value : 46
- MED Hardware Revision: "9610D01A"
- MED Firmware Revision: "hb96xxual_20r30s.bin"
- MED Software Revision: "ha96xxual_20r30s.bin"
- MED Serial Number: "06N538825133"
- MED Manufacturer Name: "Avaya"
- MED Model Name: "9610"

LLDP Port 2:10 detected 1 neighbor

Neighbor: 00:01:30:F9:9E:80/8:10, age 20 seconds

- Chassis ID type: MAC address (4)
Chassis ID : 00:01:30:F9:9E:80
- Port ID type: ifName (5)
Port ID : "8:10"
- Time To Live: 120 seconds
- System Description: "NETGEAR 8800 version 12.0.0.6 v1200b6 by release-ma\
nager on Mon Mar 19 00:43:19 PDT 2007"

show lldp statistics

show lldp {port [all | <port_list>]} statistics

Description

Displays statistical counters related to the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The following counters are presented with the standard command (taken from the IEEE 802.1ab MIB definition):

- Last table change time: Last time an entry in the LLDP database was added, changed or deleted.
- Number of table inserts: The number of times the complete set of information advertised by a particular neighbor has been inserted into tables.
- Number of table deletes: The number of times the complete set of information advertised by a particular neighbor has been deleted from tables.
- Number of table drops: The number of times the complete set of information advertised by a particular neighbor could not be stored in memory because of insufficient resources.
- Number of table age outs: The number of times the complete set of information advertised by a particular neighbor has been deleted from tables because the information timeliness interval has expired.
- Tx Total: The number of LLDP frames transmitted by this switch on the indicated port.
- Tx Total Length Exceeded: The number of LLDP frames sent out on this port that could not hold all the information configured because the total frame length would exceed the maximum LDDPDU size of 1500 bytes.
- Rx Total: The number of valid LLDP frames received by this switch on the indicated port, while this LLDP agent is enabled.
- Rx Discarded: The number of LLDP frames received by this switch on the indicated port, and then discarded for any reason.
- Rx Errors: The number of invalid LLDP frames received by this switch on the indicated port, while this LLDP agent is enabled.
- TLVs Discarded: The number of LLDP TLVs discarded for any reason by this switch on the indicated port.
- TLVs Unrecognized: The number of LLDP TLVs received on the given port that are not recognized by the switch.

Example

The following example lists statistical counters for all ports on the switch:

```
# show lldp port all statistics
```

```

Last table change time   : Fri Dec 17 10:42:33 2004
Number of Table Inserts  : 3
Number of Table Deletes  : 0
Number of Table Drops    : 0
Number of Table Age Outs : 0

```

Port	Tx Total	Tx Length Exceeded	Rx Total	Rx Discarded	Rx Errors	TLVs Discarded	TLVs Unrecogn.
1:1	189	0	5654	0	0	0	0
2:2	188	0	565	0	0	0	0

unconfigure lldp

```
unconfigure lldp {ports [all | <port_list>]}
```

Description

Leaves LLDP enabled and configured; restores the LLDP timer default values.

Syntax Description

all	Specifies all ports on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

When you issue the global `unconfigure lldp`, only the LLDP timers are reset to default values. All the configured TLVs remain on the ports remain, and LLDP remains enabled.

When you use the keyword `ports`, the TLVs for each port are returned to the five default TLVs. LLDP remains enabled.

Example

The following command restores LLDP factory default TLVs for ports 1:4 to 1:8:

```
unconfigure lldp ports 1:4 - 1:8
```

Power over Ethernet (PoE) is an effective method of supplying 48 VDC power to certain types of powered devices (PDs) through Category 5 or Category 3 twisted pair Ethernet cables. PDs include wireless access points, IP telephones, laptop computers, web cameras, and other devices. With PoE, a single Ethernet cable supplies power and the data connection, reducing costs associated with separate power cabling and supply. PoE for NETGEAR 8800 includes a method of detection to assure that power is delivered to devices that meet the IEEE 802.3af specification for PoE, as well as to many legacy devices.

Summary of PoE Software Features

The NETGEAR 8800 PoE devices support the following PoE software features:

- Configuration and control of the power distribution for PoE at the system, slot, and port levels
- Real-time discovery and classification of 802.3af-compliant PDs and many legacy (non-standard) devices
- Monitor and control of PoE fault conditions
- Support for configuring and monitoring PoE status at the system, slot, and port levels
- LED control for indicating the port's PoE inline power state
- Management of an over-subscribed power budget

For more information about configuring and managing PoE, see the *NETGEAR 8800 User Manual*.

clear inline-power stats ports

```
clear inline-power stats ports [all | <port_list>]
```

Description

Clears the inline statistics for the selected port to zero.

Syntax Description

all	Specifies all ports.
-----	----------------------

port_list	Specifies one or more ports or slots and ports.
-----------	---

Default

N/A.

Usage Guidelines

Use this command to clear all the information displayed by the `show inline-power stats ports <port_list>` command.

Example

The following command clears the inline statistics for ports 1-8 on slot 3:

```
clear inline-power stats ports 3:1-3:8
```

The following command displays cleared inline power configuration information for ports 1-8 in slot 3:

```
show inline-power stats ports 3:1-3:8
```

Following is sample output from this command:

STATISTICS COUNTERS							
Port	State	Class	Absent	InvSig	Denied	OverCurrent	Short
3:1	delivering	class3	0	0	0	0	0
3:2	delivering	class3	0	0	0	0	0
3:3	searching	class0	0	0	0	0	0
3:4	searching	class0	0	0	0	0	0
3:5	searching	class0	0	0	0	0	0
3:6	searching	class0	0	0	0	0	0
3:7	searching	class0	0	0	0	0	0
3:8	searching	class0	0	0	0	0	0

configure inline-power budget

```
configure inline-power budget <num_watts> {slot <slot>}
```

Description

Sets the reserved power on the switch or specified slot to the specified watts.

Syntax Description

num_watts	Specifies the number of watts to reserve for specified switch or slot for inline power. Enter an integer. The minimum value is 37, or 0 if the slot is disabled; the maximum is 768; and the default value is 50.
slot	Specifies a slot. The slot must be configured to hold a PoE module.

Default

50 W.

Usage Guidelines

This command sets the budgeted power reserved for all PDs connected to the switch or specified slot in Watts. None of the power budget on a specified slot can be used to power other slots or PDs on other slots.

If you specify a slot that is not configured to hold a PoE module, the system returns the following error message:

```
Error: Slot 2 is not capable of inline-power.
```

You can modify the power budget without disabling the switch or slot.

If the power consumption of the PDs on the switch or a specified slot exceeds this configured power budget, the system disconnects the lowest priority ports. (Refer to [configure inline-power priority ports](#) for information on configuring this parameter.)

If you attempt to configure this power budget for a value that the system cannot safely provide, the system returns an error message. To display inline power settings, use the command `show inline-power`; to display the power for the entire switch, use the command `show power budget`.

Note: You must disable inline power for the switch or the specified slot using the `disable inline-power slot` command prior to setting the budget to 0.

To reduce the chances of ports fluctuating between powered and non-powered states, newly inserted PDs are not powered when the actual delivered power for the module is within approximately 19 W of the configured inline power budget for that switch or slot. However, actual aggregate power can be delivered up to the configured inline power budget for the switch or slot (for example, when delivered power from ports increases or when the configured inline power budget for the switch or slot is reduced).

Example

The following command sets the power for slot 4 to 150 W on NETGEAR 8800 switches:

```
configure inline-power budget 150 slot 4
```

configure inline-power disconnect-precedence

```
configure inline-power disconnect-precedence [deny-port | lowest-priority]
```

Description

Configures the disconnect precedence priority for the switch when a new PD is detected and the measured inline power for that switch or specified slot is within 19 W of the switch's or slot's PoE power budget.

Syntax Description

deny-port	Specifies power be denied to PD requesting power, regardless of priority.
lowest-priority	Specifies power be withdrawn from lowest-priority port(s) when next PD requesting power connects.

Default

Deny-port.

Usage Guidelines

You configure this parameter for the switch; you cannot configure this per slot or per port.

If the power supplied to the PDs on a switch or specified slot exceeds the power that was budgeted for that switch or specified slot, the system disconnects power to one or more ports to prevent power overload. Refer to [configure inline-power budget](#) for information on configuring and modifying the power budgeted for each switch or specified slot.

You configure the switch to either deny power to the next PD that requests power on that switch or slot, regardless of the priority, or to disconnect those PDs on ports with lower priorities until there is enough power for the new PD. If you select this last argument and you did not configure port priorities or if several ports have the same priority, the switch withdraws power (or disconnects) those ports with the *highest* port number (s). Refer to [configure inline-power priority ports](#) for information on configuring the PoE priority for the ports.

The default value is deny-port. So, if you do not change the default value and the switch's or slot's power is exceeded, the next PD requesting power will not be connected.

When the setting is lowest priority, the switch continues dropping ports with the lowest configured PoE port priorities, or the highest port number in the case of equal PoE port priorities, until there is enough power for the requesting PD.

Example

The following command sets the switch to withdraw power from the lowest-priority port(s):

```
configure inline-power disconnect-precedence lowest-priority
```

configure inline-power label ports

```
configure inline-power label <string> ports <port_list>
```

Description

Lets you create your own label for a specified PoE port or group of PoE ports.

Syntax Description

string	Specifies a name up to 15 characters in length to identify the specified power port(s).
port_list	Specifies one or more ports or slots and ports.

Default

No label.

Usage Guidelines

Use the `show inline-power configuration ports` command, as shown in the following example, to display inline power configuration information, including the label (if any) for each port:

```
show inline-power configuration port 3:1-10
```

Following is sample output from this command on a NETGEAR 8800:

Port	Config	Operator Limit	Priority	Label
3:1	Enabled	16000 mW	Low	finance
3:2	Enabled	15000 mW	Low	finance
3:3	Enabled	15000 mW	Low	
3:4	Enabled	15000 mW	Low	
3:5	Enabled	15000 mW	Low	
3:6	Enabled	15000 mW	Low	marketing
3:7	Enabled	15000 mW	Low	marketing
3:8	Enabled	15000 mW	Low	marketing
3:9	Enabled	15000 mW	Low	
3:10	Enabled	15000 mW	Low	

Example

The following command assigns the name “alpha-test_1” to port 1 on slot 4:

```
config inline-power label alpha-test_1 ports 4:1
```

configure inline-power operator-limit ports

```
configure inline-power operator-limit <milliwatts> ports [all |<port_list>]
```

Description

Sets the power limit allowed for PDs connected to the specified ports.

Syntax Description

milliwatts	An integer specifying the maximum allowed power in milliwatts; the range is 3000 to 16800 mW.
	Note: If you attempt to enter a higher value, the switch returns an error message.
port_list	Specifies one or more ports or slots and ports.

Default

15400 mW.

Usage Guidelines

This command sets the power limit that a PD can draw on the specified ports. Range is 3000 to 16800 mW; the default value is 15400 mW.

If the measured power for a specified port exceeds the port's operator limit, the power is withdrawn from that port and the port moves into a fault state.

If you try to set an operator-limit outside the accepted range, the system returns the following error message:

```
Error: Invalid operator-limit value. Must be in the range of 3000-16800 mW
```

Example

The following command sets the limit for legacy PDs on ports 3 – 6 of slot 5 to 10000 mW:

```
configure inline-power operator-limit 10000 ports 5:3-5:6
```

configure inline-power priority ports

```
configure inline-power priority [critical | high | low] ports <port_list>
```

Description

Sets the PoE priority on the specified ports.

Syntax Description

critical high low	Sets the PoE priority for the specified ports.
port_list	Specifies one or more ports or slots and ports.

Default

Low.

Usage Guidelines

The system allocates power to those ports with the highest priorities first. This command can also be used in conjunction with the `configure inline-power disconnect-precedence` command. If you configure the disconnect precedence as lowest priority, then newly detected PDs will be powered if that port has higher priority than the existing powered ports.

If there are multiple ports at the same priority level (either configured or by default) and one of the ports must have power withdrawn because of excessive power demands, those ports with the lower port number are powered first. The higher port numbers have power withdrawn first in the case of equal PoE port priorities.

Example

The following command assigns a critical PoE priority on ports 4 – 6 on slot 3:

```
configure inline-power priority critical ports 3:4-3:6
```

configure inline-power usage-threshold

```
configure inline-power usage-threshold <threshold>
```

Description

Sets the inline power usage SNMP event threshold.

Syntax Description

threshold	Specifies the percentage of budgeted power used on any PoE module or stand-alone switch that causes the system to send an SNMP event and create a log message. The range 1 to 99; the default value is 70.
-----------	--

Default

70.

Usage Guidelines

This command sets the threshold for generating an SNMP event and an Event Management System (EMS) message. This threshold is reached when the measured power for a PoE module compared to the budgeted power for that slot exceeds a certain value. On stand-alone switches, this threshold applies to the total power available to the entire switch. The configured threshold value initiates the event and message once that percentage of the budgeted power is being used.

On the NETGEAR 8800, the PoE threshold applies only to the percentage *per slot* of measured to budgeted power use; it does not apply systemwide.

The system generates an additional SNMP event and EMS message once the power usage falls *below* the threshold again; once the condition clears.

Example

The following command sets the inline power usage alarm threshold at 75%:

```
configure inline-power usage-threshold 75
```

disable inline-power

```
disable inline-power
```

Description

Shuts down PoE power currently provided on all ports on all slots.

Syntax Description

This command has no arguments or variables

Default

Enable.

Usage Guidelines

You can control whether inline power is provided to the system by using the `disable inline-power` command and the `enable inline-power` command. Using the `disable inline-power` command shuts down inline power currently provided on the entire switch or to specified ports and slots. Disabling inline power to a switch, port, or slot immediately removes power to any connected PDs. By default, inline power provided to all ports is enabled.

Note: Disabling inline power using the `disable inline-power` command does *not* affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

On the 8800, disabling inline power does not allow PoE power reserved for slots to be allocated to other slots that may be needing more power to become operational. However, when you issue the command `disable slot` on a slot holding a PoE module, the inline power is also disabled; that slot is totally offline.

Note: Inline power cannot be delivered to connected PDs unless the NETGEAR 8800 chassis and module are powered on.

Example

The following command shuts down inline power currently provided to all ports and all slots:

```
disable inline-power
```

disable inline-power legacy

```
disable inline-power legacy
```

Description

Disables the non-standard (or capacitance) power detection mechanism for the switch.

Syntax Description

This command has no arguments or variables

Default

Disable.

Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used *only* if this parameter is enabled and after an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.

The reason legacy detection is configurable is that it is possible for a normal (non-PoE) device to have a capacitance signature that causes the device to be detected as a legacy PoE device and have power delivered to it, potentially causing damage to the device.

Example

The following command disables capacitance detection of PDs on the switch:

```
disable inline-power legacy
```

disable inline-power legacy slot

```
disable inline-power legacy slot <slot>
```

Description

Disables the non-standard (or capacitance) power detection mechanism for the specified slot.

Syntax Description

slot	Disables non-standard power detection for specified slot.
------	---

Default

Disable.

Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch or specified slot. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used *only* if this parameter is enabled and after an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.

The reason legacy detection is configurable is that it is possible for a normal (non-PoE) device to have a capacitance signature that causes the device to be detected as a legacy PoE device and have power delivered to it, potentially causing damage to the device.

On a stack if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command disables capacitance detection of PDs on slot 3 of the NETGEAR 8800:

```
disable inline-power legacy slot 3
```

disable inline-power ports

```
disable inline-power ports [all | <port_list>]
```

Description

Shuts down PoE power currently provided to all ports or to specified ports.

Syntax Description

all	Disables inline power to all ports on the switch.
port_list	Disables inline power to the specified ports.

Default

Enable.

Usage Guidelines

Disabling inline power to ports immediately removes power to any connected PDs. By default, the capability to provide inline power to all ports is enabled.

Note: Disabling inline power using the `disable inline-power` command does *not* affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

Disabling inline power to a port providing power to a PD immediately removes power to the PD.

Note: On the NETGEAR 8800, PoE power removed from ports using this command can be used by other ports on the same module.

Example

The following command shuts down inline power currently provided to ports 4 and 5 on slot 3 on the NETGEAR 8800:

```
disable inline-power ports 3:4-5
```

disable inline-power slot

```
disable inline-power slot <slot>
```

Description

Shuts down PoE power currently provided to the specified slot.

Syntax Description

slot	Disables inline power to specified slot.
------	--

Default

Enable.

Usage Guidelines

Disabling inline power to a slot immediately removes power to any connected PDs. By default, the capability to provide inline power to a slot is enabled.

Disabling a slot using this command does not change the power budgeted to a specified slot using the `configure inline-power budget` command; nor can that power be used by PDs connected to any other slot.

Note: You can set the reserved power budget to 0 for a slot if, and only if, you first issue this command.

On a stack if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command removes power to all PDs on slot 3:

```
disable inline-power slot 3
```

enable inline-power

```
enable inline-power
```

Description

Enables PoE power to all ports on all slots.

Syntax Description

This command has no arguments or variables.

Default

Enable.

Usage Guidelines

You can control whether inline power is provided to the system by using the `disable inline-power` command and the `enable inline-power` command. By default, inline power provided to all ports is enabled.

Enabling inline power starts the PoE detection process used to discover, classify, and power remote PDs.

Note: If your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the `disable inline-power` command does *not* affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

However, when you issue the command `disable slot` for the switch on a slot holding a PoE module, the inline power is also disabled; that slot is totally offline.

Note: Inline power cannot be delivered to connected PDs unless the NETGEAR 8800 chassis and module are powered on.

Example

The following command enables inline power currently provided to all ports and all slots:

```
enable inline-power
```

enable inline-power legacy

```
enable inline-power legacy
```

Description

Enables the non-standard (or capacitance) power detection mechanism for the switch.

Syntax Description

This command has no arguments or variables

Default

Disable.

Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used *only* if this parameter is enabled and *after* an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.



CAUTION:

A normal (non-PoE) device may have a capacitance signature that causes the device to be detected as a legacy PoE device (and have power supplied), potentially causing damage to the device.

Example

The following command enables capacitance detection of PDs on the switch:


```
enable inline-power legacy
```

enable inline-power legacy slot

```
enable inline-power legacy slot <slot>
```

Description

Enables non-standard (or capacitance) power detection mechanism for the specified slot on the switch.

Syntax Description

slot	Enables non-standard power detection for specified slot.
------	--

Default

Disable.

Usage Guidelines

This command enables the non-standard power-detection mechanism on the specified slot. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used *only* if this parameter is enabled and *after* an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.



CAUTION:

A normal (non-PoE) device may have a capacitance signature that causes the device to be detected as a legacy PoE device (and have power supplied), potentially causing damage to the device.

On stack, if you do not specify a slot number, the command operates on all active nodes. The command operates *only* on nodes in the active topology.

Example

The following command enables capacitance detection of PDs on slot 3 on the switch:

```
enable inline-power legacy slot 3
```

enable inline-power ports

```
enable inline-power ports [all | <port_list>]
```

Description

Enables PoE power currently provided to all ports or to specified ports.

Syntax Description

all	Enables inline power to all ports on the switch.
port_list	Enables inline power to the specified ports.

Default

Enable.

Usage Guidelines

Disabling inline power to a port immediately removes power to any connected PD. By default, inline power provided to all ports is enabled.

To deliver inline power to ports with connected PDs, you must also reserve power for the slot with the PDs using the `configure inline-power budget` command. If you do not have enough reserved power for the port, that port moves into a Denied state.

Note: If your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the `disable inline-power` command does *not* affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

Example

The following command enables inline power to ports 4 and 5 on slot 3 on the switch:

```
enable inline-power ports 3:4-5
```

enable inline-power slot

```
enable inline-power slot <slot>
```

Description

Enables PoE power to the specified slot on the switch.

Syntax Description

slot	Enables inline power to specified slot.
------	---

Default

Enable.

Usage Guidelines

Disabling inline power to a slot immediately removes power to any connected PDs. By default, inline power provided to all slots is enabled.

To deliver inline power to slots, you must reserve power for that slot using the `configure inline-power budget` command. By default, each PoE module has 50 W of power reserved for inline power.

Note: If your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the `disable inline-power` command does *not* affect the data traffic traversing the slot. And, disabling the slot using the `disable slot` command does not affect the inline power supplied to the slot.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command makes inline power available to slot 3:

```
enable inline-power slot 3
```

reset inline-power ports

```
reset inline-power ports <port_list>
```

Description

Power cycles the specified ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports for which power is to be reset.
-----------	--

Default

N/A.

Usage Guidelines

This command power cycles the specified ports. Ports are immediately disabled and then re-enabled, allowing remote PDs to be power-cycled.

This command affects only inline power; it does not affect network connectivity for the port(s).

Example

The following command resets power for port 4 on slot 3 on the switch:

```
reset inline-power ports 3:4
```

show inline-power

```
show inline-power
```

Description

Displays inline power status information for the specified PoE switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output varies depending on the PoE device you are using.

- Inline power status—The status of inline power. The status conditions are:
 - Enabled
 - Disabled
- Power usage threshold
- Disconnect precedence
- Firmware status—The operational status of the slot. The status conditions are:
 - Operational
 - Not operational
 - Disabled
 - Subsystem failure
- Measured power—The amount of power, in watts, that currently being used by the switch.

- Legacy—The status of the legacy mode, which allows detection of many non-standard PDs.

Note: For additional information on inline power parameters, refer to the `show power budget` command.

Example

The following command displays inline power status for the switch:

```
show inline-power
```

```
(Demo) XCM8806.2 # show inline-power
Inline Power System Information
Configured : Enabled
System Power Surplus : 2473 Watts available for budgeting
Redundant Power Surplus : 1438 Watts available for budgeting to maintain N+1
Power Usage Threshold : 70 percent (per slot)
Disconnect Precedence : deny-port
Budgeted Measured
Slot Inline-Power Firmware Status Power (Watts) Power (Watts) Legacy
6 Enabled Operational 50 W 0 W Disabled
```

show inline-power configuration ports

```
show inline-power configuration ports <port_list>
```

Description

Displays inline power configuration information for the specified ports.

Syntax Description

<code>port_list</code>	Specifies one or more ports.
------------------------	------------------------------

Default

N/A.

Usage Guidelines

The output displays the following inline power configuration information for the specified ports:

- Config—Indicates whether the port is enabled to provide inline power:
 - Enabled: The port can provide inline power.

- Disabled: The port cannot provide inline power.
- Operator Limit—Displays the configured limit, in milliwatts, for inline power on the port.
- Label—Displays a text string, if any, associated with the port.

The following also displays for this command on modular PoE devices:

- Priority—Displays inline power priority of the port, which is used when the disconnect precedence is set to lowest priority:
 - Low
 - High
 - Critical

Example

The following command displays inline power configuration information for ports 1 to 10 in slot 3 on the switch:

```
show inline-power configuration port 3:1-10
```

Following is sample output from this command:

Port	Config	Operator Limit	Priority	Label
3:1	Enabled	15000 mW	Low	
3:2	Enabled	15000 mW	Low	
3:3	Enabled	15000 mW	Low	
3:4	Enabled	15000 mW	Low	
3:5	Enabled	15000 mW	Low	
3:6	Enabled	15000 mW	Low	
3:7	Enabled	15000 mW	Low	
3:8	Enabled	15000 mW	Low	
3:9	Enabled	15000 mW	Low	
3:10	Enabled	15000 mW	Low	

show inline-power info ports

```
show inline-power info {detail} ports <port_list>
```

Description

Displays inline power information for the specified ports.

Syntax Description

port_list	Specifies one or more ports.
-----------	------------------------------

Default

N/A.

Usage Guidelines

Note: Ports in the `denied` or `faulted` state periodically display the `searching` state as the hardware retests the PD state.

You can use this command to generate a summary report or a detailed report.

Summary output displays the following inline power information for the specified ports:

- **State**—Displays the port power state:
 - Disabled
 - Searching
 - Delivering
 - Faulted
 - Disconnected
 - Other
 - Denied
- **PD's power class**—Displays the class type of the connected PD:
 - “----”: disabled or searching
 - “class0”: class 0 device
 - “class1”: class 1 device
 - “class2”: class 2 device
 - “class3”: class 3 device
 - “class4”: class 4 device
- **Volts**—Displays the measured voltage. A value from 0 to 2 is valid for ports that are in a searching state.
- **Curr**—Displays the measured current, in milliamperes, drawn by the PD.
- **Power**—Displays the measured power, in watts, supplied to the PD.
- **Fault**—Displays the fault value:
 - None
 - UV/OV fault
 - UV/OV spike
 - Over current
 - Overload
 - Undefined
 - Underload
 - HW fault
 - Discovery resistance fail

- Operator limit violation
- Disconnect
- Discovery resistance, A2D failure
- Classify, A2D failure
- Sample, A2D failure
- Device fault, A2D failure
- Force on error

The detail command lists all inline power information for the selected ports. Detail output displays the following information:

- Configured Admin State—Displays the port's configured state; Enabled or Disabled.
- Inline Power State—Displays the port power state.
- MIB Detect Status—Displays the port state as reported by SNMP; valid values are as follows:
 - disabled
 - searching
 - delivering
 - fault
 - test
 - otherFault
 - denyLowPriority
- Label—Displays the port's configured label.
- Operator Limit—Displays the port's configured operator limit value.
- PD Class—Displays the class type of connected PD:
- Max Allowed Power—Displays the amount of maximum allowed power for a device of this class.
- Measured Power—Displays the measured power, in watts, supplied to the PD.
- Line Voltage—Displays the measured voltage. A value from 0 to 2 is valid for ports in a searching state.
- Current—Displays the measured current, in milliamperes, drawn by the PD.
- Fault Status—Displays the fault value.
- Detailed Status

The following information displays only with modular PoE devices:

- Priority—Displays the port's configured PoE priority value, as follows:
 - Critical
 - High
 - Low

Example

The following command displays summary inline power information for ports 1 to 3 on slot 3 on the switch:

```
show inline-power info ports 3:1-3
```

Following is sample output from this command:

Port	State	Class	Volts	Curr (mA)	Power (Watts)	Fault
3:1	delivering	class3	48.3	192	9.300	None
3:2	delivering	class3	48.3	192	9.300	None
3:3	searching	-----	0.0	0	0.0	None

The following command displays detail inline power information for port 1 on slot 3:

```
show inline-power info detail port 3:1
```

Following is sample output from this command:

```
Port 3:1
```

```
Configured Admin State: enabled
Inline Power State      : delivering
MIB Detect Status      : delivering
Label                  :
Operator Limit         : 16800 milliwatts
PD Class               : class3
Max Allowed Power      : 15.400 W
Measured Power         : 9.400 W
Line Voltage           : 48.3 Volts
Current                : 193 mA
Fault Status           : None
Detailed Status        :
```

show inline-power slot

```
show inline-power slot <slot>
```

Description

Displays inline power information for the specified slot on the switch.

Syntax Description

slot	Specifies the slot.
------	---------------------

Default

N/A.

Usage Guidelines

The output indicates the following inline power status for each system:

- Configured power
 - Enabled
 - Disabled
- System power surplus
- Redundant power surplus
- Power usage threshold
- Disconnect precedence
- Legacy—The status of the legacy mode, which allows detection of many non-standard PDs.

The output indicates the following inline power status information for each slot:

- Inline power status—The status of inline power. The status conditions are:
 - Enabled
 - Disabled
- Firmware status—The operational status of the slot. The status conditions are:
 - Operational
 - Not operational
 - Disabled
 - Subsystem failure
 - Card not present
 - Slot disabled
- Budgeted power—The amount of power, in watts, that is available to the slot.
- Measured power—The amount of power, in watts, that currently being used by the slot.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command displays inline power information for slot 3 on the switch:

```
show inline-power slot 3
```

Following is sample output from this command:

```

      Inline Power System Information
Configured           : Enabled
System Power Surplus : 1500 Watts available for budgeting
Redundant Power Surplus : 465 Watts available for budgeting to maintain N+1
Power Usage Threshold : 70 percent (per slot)
Disconnect Precedence : lowest-priority
```

```
Legacy Mode           : Disabled
```

			Budgeted	Measured
Slot	Inline-Power	Firmware Status	Power (Watts)	Power (Watts)
3	Enabled	Operational	50 W	9 W
4	Enabled	Card Not Present	(50 W)	n/a
7	Enabled	Operational	50 W	0 W

Note: A budget value in parentheses is not allocated from the system power

show inline-power stats

```
show inline-power stats
```

Description

Displays inline power statistics for the specified switch.

Syntax Description

There are no variables or parameters for this command.

Default

N/A.

Usage Guidelines

Use this command to produce a report that shows the firmware status and version plus how many ports are currently faulted, powered, and waiting for power for the switch. Unlike the values displayed with the `show inline-power stats ports` command, these values are current readings, not cumulative counters.

Example

The following command displays inline power statistics information for the NETGEAR 8800 switch:

```
show inline-power stats
```

Following is sample output from this command:

```
Inline-Power Slot Statistics
Firmware status           : Operational
Firmware revision         : 292b1

Total ports powered       : 7
Total ports awaiting power : 17
Total ports faulted       : 0
Total ports disabled      : 0
```

show inline-power stats ports

```
show inline-power stats ports <port_list>
```

Description

Displays inline power statistics for the specified ports.

Syntax Description

port_list	Specifies one or more slots and ports.
-----------	--

Default

N/A.

Usage Guidelines

The output displays the following inline power statistics for the specified ports:

- State—Displays the port power state:
 - Disabled
 - Searching
 - Delivering
 - Faulted
 - Disconnected
 - Other
 - Denied
- PD's power class—Displays the class type of the connected PD:
 - “----”: disabled or searching
 - “class0”: class 0 device
 - “class1”: class 1 device
 - “class2”: class 2 device
 - “class3”: class 3 device
 - “class4”: class 4 device
- Absent—Displays the number of times the port was disconnected.
- InvSig—Displays the number of times the port had an invalid signature.
- Denied—Displays the number of times the port was denied.
- Over-current—Displays the number of times the port entered an overcurrent state.
- Short—Displays the number of times the port entered undercurrent state.

Example

The following command displays inline power configuration information for ports 1 to 10 in slot 3 on the switch:

```
show inline-power stats ports 3:1-10
```

Following is sample output from this command:

STATISTICS COUNTERS							
Port	State	Class	Absent	InvSig	Denied	OverCurrent	Short
3:1	delivering	class3	0	0	0	18	0
3:2	delivering	class3	0	0	0	0	0
3:3	searching	class0	0	0	0	0	0
3:4	searching	class0	0	0	0	0	0
3:5	searching	class0	0	0	0	0	0
3:6	searching	class0	0	0	0	0	0
3:7	searching	class0	0	0	0	0	0
3:8	searching	class0	0	0	0	0	0
3:9	searching	class0	0	0	0	0	0
3:10	searching	class0	0	0	0	0	0

show inline-power stats slot

```
show inline-power stats slot <slot>
```

Description

Displays inline power statistics for the specified slot on the switch.

Syntax Description

slot	Specifies the slot.
------	---------------------

Default

N/A.

Usage Guidelines

Use this command to produce a report that shows the firmware status and version plus how many ports are currently faulted, powered, and waiting for power for the selected slots. Unlike the values displayed with the `show inline-power stats ports` command, these values (displayed with the `show inline-power stats slot` command) are current readings; not cumulative counters.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates *only* on nodes in the active topology.

Example

The following command displays inline power statistics information for slot 3 on the switch:

```
show inline-power stats slot 3
```

Following is sample output from this command:

```
Inline-Power Slot Statistics
```

```

Slot: 3
Firmware status      : Operational
Firmware revision    : 292b1

Total ports powered   : 7
Total ports awaiting power : 41
Total ports faulted   : 0
Total ports disabled  : 0

```

unconfigure inline-power budget slot

```
unconfigure inline-power budget slot <slot>
```

Description

Unconfigures the inline reserved power on the 8800 on the specified slot and returns the power budget on that slot to the default value of 50 W.

Syntax Description

slot	Specifies the slot.
------	---------------------

Default

50 W.

Usage Guidelines

This command unconfigures any previously configured power budget for the specified slot and resets the budgeted power reserved for all PDs connected to this slot to 50 W. The rest of the previously configured power budget on this slot cannot be used to power other slots or PDs on other slots (unless you explicitly reconfigure the power budget for other slots).

If you specify a slot that does not have a PoE module, the system returns the following error message:

```
Error: Slot 2 is not capable of inline-power.
```

Example

The following command resets the power for slot 4 to 50 W:

```
unconfigure inline-power budget slot 4
```

unconfigure inline-power disconnect-precedence

```
unconfigure inline-power disconnect-precedence
```

Description

On a NETGEAR 8800 switch, unconfigures the disconnect precedence setting and returns the switch to the default disconnect precedence value of deny port.

Syntax Description

This command has no arguments or variables.

Default

Deny-port.

Usage Guidelines

You configure this parameter for the entire switch; you cannot configure this per slot or per port.

Unconfigures the PoE disconnect precedence previously set for the NETGEAR 8800 switch and returns the disconnect precedence to the default value of deny port. Deny port denies power to the next PD that requests inline power from the slot when the inline power budget for the switch or slot is reached, regardless of the inline power port priority.

Example

The following command resets the switch to the PoE disconnect precedence value, which is deny port:

```
unconfigure inline-power disconnect-precedence
```

unconfigure inline-power operator-limit ports

```
unconfigure inline-power operator-limit ports [all |<port_list>]
```

Description

Unconfigures the PoE operator limit setting and resets the power limit allowed for PDs connected to the specified ports to the default value of 15400 mW.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more slots and ports.

Default

15400 mW.

Usage Guidelines

This command unconfigures any previously configured operator limit for the specified ports. It resets the maximum power that any PD can draw to 15400 mW.

Example

The following command resets the limit on ports 3 to 6 of slot 5 on the switch to the default value of 15400 mW:

```
unconfigure inline-power operator-limit ports 5:3-5:6
```

unconfigure inline-power priority ports

```
unconfigure inline-power priority ports [all | <port_list>]
```

Description

On NETGEAR 8800 switches, unconfigures the PoE priority on the specified ports and returns the ports to the default PoE port priority value of low.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more ports or slots and ports.

Default

Low.

Usage Guidelines

Use this to reset the PoE port priority on specified ports on the NETGEAR 8800 switch to the default value of low.

If there are multiple ports on the NETGEAR 8800 switch at the same priority level (either configured or by default), and one of the ports must have power withdrawn because of excessive power demands, those ports with the lower port number are powered first. The higher port numbers have power withdrawn first in the case of equal PoE port priorities.

Example

The following command resets the PoE priority on ports 4 – 6 on slot 3 to low:

```
unconfigure inline-power priority ports 3:4-3:6
```

unconfigure inline-power usage-threshold

```
unconfigure inline-power usage-threshold
```


Description

Unconfigures the inline power usage alarm threshold and returns threshold to the default value of 70%.

Syntax Description

This command has no arguments or variables.

Default

70.

Usage Guidelines

This command unconfigures the PoE usage threshold setting for initiating SNMP event and EMS messages and returns the switch's inline power usage threshold for to 70%. The system initiates an event and message once that percentage of the budgeted power is being used.

The system generates an additional SNMP event and EMS message once the power usage falls *below* the threshold again; once the condition clears.

Example

The following command resets the inline power usage alarm threshold to 70%:

```
unconfigure inline-power usage-threshold
```

Commands for Status Monitoring and Statistics

8

This chapter describes commands for:

- Configuring and managing the Event Management System/Logging
- Configuring and monitoring system health and statistics
- Enabling and disabling the collection of remote monitoring (RMON) statistics on the switch
- Enabling, disabling, and configuring sFlow® statistics collection

Event Management System

When an event occurs on a switch, the Event Management System (EMS) allows you to send messages generated by these events to a specified log target. You can send messages to the memory buffer, NVRAM, the console display, the current session, to a syslog host, or to the other Management Switch Fabric Module (MSM) or Management Module (MM). The log messages contain configuration and fault information pertaining to the device. You can format the log messages to contain various items of information, but typically a message consists of:

- Timestamp—The timestamp records when the event occurred.
- Severity level:
 - Critical—A desired switch function is inoperable. The switch may need to be reset.
 - Error—A problem is interfering with normal operation.
 - Warning—An abnormal condition exists that may lead to a function failure.
 - Notice—A normal but significant condition has been detected; the system is functioning as expected.
 - Info—Actions and events that are consistent with expected behavior.
 - Debug-Summary, Debug-Verbose, and Debug-Data—Information that is useful when performing detailed trouble shooting procedures.

By default, log entries that are assigned a critical, error, or warning level are considered static entries and remain in the NVRAM log target after a switch reboot.

- Component—The component refers to the specific functional area to which the error refers.

- **Message**—The message contains the log information with text that is specific to the problem.

The switch maintains a configurable number of messages in its internal (memory-buffer) log (1000 by default). You can display a snapshot of the log at any time. In addition to viewing a snapshot of the log, you can configure the system to maintain a running real-time display of log messages on the console display or telnet session. In addition to maintaining an internal log, the switch supports remote logging by way of the UNIX syslog host facility.

EMS supports IPv6 as a parameter for filtering events.

sFlow Statistics

sFlow® is a technology for monitoring traffic in data networks containing switches and routers. It relies on statistical sampling of packets from high-speed networks, plus periodic gathering of the statistics. A User Datagram Protocol (UDP) datagram format is defined to send the information to an external entity for analysis. sFlow consists of a (Management Information Base) MIB and a specification of the packet format for forwarding information to a remote agent. Details of sFlow specifications can be found in RFC 3176 and at the following website:

<http://www.sflow.org>

NETGEAR 8800 allows you to collect sFlow statistics on a per port basis. An agent, residing locally on the switch, sends data to a collector that resides on another machine. You configure the local agent, the address of the remote collector, and the ports of interest for sFlow statistics gathering. You can also modify default values for how frequently on average a sample is taken, how often the data is sent to the collector, and the maximum load allowed on the CPU before throttling the statistics gathering.

For information about software licensing, including how to obtain and upgrade your license, see Appendix A in the *NETGEAR 8800 User Manual*.

RMON

RMON is the common abbreviation for the Remote Monitoring Management Information Base (MIB) system defined by the Internet Engineering Task Force (IETF) documents RFC 1757 and RFC 2021, which allows you to monitor LANs remotely.

Using the RMON capabilities of the switch allows network administrators to improve system efficiency and reduce the load on the network.

The IETF defines nine groups of Ethernet RMON statistics. The switch supports the following four of these groups, as defined in RFC 1757:

- Statistics
- History
- Alarms
- Events

The switch also supports the following parameters for configuring the RMON probe and the trap destination table, as defined in RFC 2021:

- probeCapabilities
- probeSoftwareRev
- probeHardwareRev
- probeDateTime
- probeResetControl
- trapDestTable

clear counters

```
clear counters
```

Description

Clears all switch statistics and port counters, including port packet statistics, bridging statistics, IP statistics, and log event counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You should view the switch statistics and port counters before you clear them. Use the `show ports` command to view port statistics. Use the `show log counters` command to show event statistics.

The CLI also provides a number of options that you can specify with the `clear counters` command. If you specify an option, the switch only clears the statistics for that option. For example, if you want to clear, reset only the STP statistics and counters, use the `clear counters stp` command. Please refer to the specific chapter in this guide for more detailed information about those commands.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counters, you can see fresh statistics for the time period you are monitoring.

Example

The following command clears all switch statistics and port counters:

```
clear counters
```

clear log

```
clear log {error-led | static | messages [memory-buffer | nvram]}
```

Description

Clears the log messages in memory and NVRAM, and clears the ERR LED on the MSM/MM.

Syntax Description

error-led	Clears the ERR LED on the MSM/MM.
static	Specifies that the messages in the NVRAM and memory-buffer targets are cleared, and the ERR LED on the MSM/MM is cleared.
memory-buffer	Clears entries from the memory buffer.
nvram	Clears entries from NVRAM.

Default

N/A.

Usage Guidelines

The switch log tracks configuration and fault information pertaining to the device.

By default, log entries that are sent to the NVRAM remain in the log after a switch reboot. The `clear log` and `clear log messages memory-buffer` commands remove entries in the memory buffer target; the `clear log static` and `clear log messages nvram` commands remove messages from the NVRAM target. In addition, the `clear log static` command will also clear the memory buffer target.

There are three ways to clear the ERR LED: clear the log, reboot the switch, or use the `clear log error-led` command. To clear the ERR LED without rebooting the switch or clearing the log messages, use the `clear log error-led` command.

Example

The following command clears all log messages, from the NVRAM:

```
clear log static
```

clear log counters

```
clear log counters [<event-condition> | [all | <event-component>] {severity <severity> {only}}]
```

Description

Clears the incident counters for events.

Syntax Description

event-condition	Specifies the event condition counter to clear.
all	Specifies that all events counters are to be cleared.
event-component	Specifies that all the event counters associated with a particular component should be cleared.
severity	Specifies the minimum severity level of event counters to clear (if the keyword only is omitted).
only	Specifies that only event counters of the specified severity level are to be cleared.

Default

If severity is not specified, then the event counters of any severity are cleared in the specified component.

Usage Guidelines

This command sets the incident counters to zero for each event specified. To display event counters, use the following command:

```
show log counters
```

See the command [show log on page 387](#) for more information about severity levels.

To get a listing of the event conditions in the system, use the following command:

```
show log events {detail}
```

To get a listing of the components present in the system, use the following command:

```
show log components
```

Execution of these commands on a backup or standby node results in the clearing of that node's information only. Execution of these commands on the master node results in the clearing of information on all nodes in the system.

Example

The following command clears the event counters for event conditions of severity error or greater in the component *BGP*:

```
clear log counters "BGP" severity error
```

clear sys-recovery-level

```
clear sys-recovery-level
```

Description

If configured and the switch detects a hardware fault and enters the shutdown state, this command clears the shutdown state and renders the switch, I/O, or MSM/MM module(s) operational.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If you configure the switch or one or more modules to shutdown upon detecting a hardware fault, and the switch or module enters the shutdown state, you must explicitly clear the shutdown state and reset the switch or the affected modules for the switch to become operational.

To clear the shutdown state, use the following command:

```
clear sys-recovery-level
```

The switch prompts you to confirm this action. The following is a sample confirmation message:

```
Are you sure you want to clear sys-recovery-level? (y/n)
```

Enter **y** to confirm this action and clear the shutdown state. Enter **n** or press [Enter] to cancel this action.

On the NETGEAR 8800, after using the `clear sys-recovery-level` command, you must reset each affected module.

If you configured only a few I/O modules to shutdown, reset each affected I/O module as follows:

1. Disable the slot using the `disable slot <slot>` command.
2. Re-enable the slot using the `enable slot <slot>` command.

Note: You must complete this procedure for each module that enters the shutdown state.

If you configured all I/O modules or one or more MSMs/MMs to shut down, use the `reboot` command to reboot the switch and reset all affected modules.

After you clear the shutdown state and reset the affected module, each port is brought offline and then back online before the module and the entire system is operational.

Example

The following command clears the shutdown state:

```
clear sys-recovery-level
```

configure log display

```
configure log display <severity> {only}
```

Description

Configures the real-time log-level message to display.

Syntax Description

severity	Specifies a message severity. Severities include critical, error, warning, notice, info, debug-summary, debug-verbose, and debug-data.
only	Specifies only log messages of the specified severity level.

Default

If not specified, messages of all severities are displayed on the console display.

Usage Guidelines

You must enable the log display before messages are displayed on the log display. Use the `enable log display` command to enable the log display. This allows you to configure the system to maintain a running real-time display of log messages on the console.

Severity filters the log to display messages with the selected severity or higher (more critical). Severities include critical, error, warning, info, notice, debug-summary, debug-verbose, and debug-data.

You can also control log data to different targets. The command equivalent to `configure log display` is the following:

```
configure log target console-display severity <severity>
```

To display the current configuration of the log display, use the following command:

```
show log configuration target console-display
```

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Example

The following command configures the system log to maintain a running real-time display of log messages of critical severity or higher:

```
configure log display critical
```


The following command configures the system log to maintain a running real-time display of only log messages of critical severity:

```
configure log display critical only
```

configure log filter events

```
configure log filter <name> [add | delete] {exclude} events [<event-condition> | [all | <event-component>] {severity <severity> {only}}]
```

Description

Configures a log filter to add or delete detailed feature messages based on a specified set of events.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

name	Specifies the filter to configure.
add	Add the specified events to the filter
delete	Remove the specified events from the filter
exclude	Events matching the specified events will be excluded
event-condition	Specifies an individual event.
all	Specifies all components and subcomponents.
event-component	Specifies all the events associated with a particular component.
severity	Specifies the minimum severity level of events (if the keyword only is omitted).
only	Specifies only events of the specified severity level.

Default

If the `exclude` keyword is not used, the events will be included by the filter. If `severity` is not specified, then the filter will use the component default severity threshold (see the note on [page 338](#) when `delete` or `exclude` is specified).

Usage Guidelines

This command controls the incidents that pass a filter by adding, or deleting, a specified set of events. If you want to configure a filter to include or exclude incidents based on event parameter values (for example, MAC address or BGP Neighbor) see the command `configure log filter events match` on [page 340](#).

When the `add` keyword is used, the specified event name is added to the beginning of the filter item list maintained for this filter. The new filter item either includes the events specified, or if the `exclude` keyword is present, excludes the events specified.

The `delete` keyword is used to remove events from the filter item list that were previously added using the `add` command. All filter items currently in the filter item list that are identical to, or a subset of, the set of events specified in the `delete` command will be removed.

Event Filtering Process

From a logical standpoint, the filter associated with each enabled log target is examined to determine whether a message should be logged to that particular target. The determination is made for a given filter by comparing the incident with the most recently configured filter item first. If the incident matches this filter item, the incident is either included or excluded, depending on whether the `exclude` keyword was used. Subsequent filter items on the list are compared if necessary. If the list of filter items has been exhausted with no match, the incident is excluded.

Events, Components, and Subcomponents

As mentioned, a single event can be included or excluded by specifying the event's name. Multiple events can be added or removed by specifying a NETGEAR 8800 component name plus an optional severity. Some components, such as *BGP*, contain subcomponents, such as *Keepalive*, which is specified as *BGP.Keepalive*. Either components or subcomponents can be specified. The keyword `all` in place of a component name can be used to indicate all NETGEAR 8800 components.

Severity Levels

When an individual event name is specified following the `events` keyword, no severity value is needed since each event has pre-assigned severity. When a component, subcomponent, or the `all` keyword is specified following the `events` keyword, a severity value is optional. If no severity is specified, the severity used for each applicable subcomponent is obtained from the pre-assigned severity threshold levels for those subcomponents. For example, if *STP* were specified as the component, and no severity is specified for the add of an include item, then only messages with severity of `error` and greater would be passed, since the threshold severity for the *STP* component is `error`. If *STP.InBPDU* were specified as the component, and no severity is specified, then only messages with severity of `warning` and greater would be passed, since the threshold severity for the *STP.InBPDU* subcomponent is `warning`. Use the `show log components` command to see this information.

The severity keyword `all` can be used as a convenience when `delete` or `exclude` is specified. The use of `delete` (or `exclude`) with severity `all` deletes (or excludes) previously added events of the same component of all severity values.

Note: If no severity is specified when `delete` or `exclude` is specified, severity `all` is used

If the `only` keyword is present following the severity value, then only the events in the specified component at that exact severity are included. Without the `only` keyword, events in the specified component at that severity or more urgent are included. For example, using the option `severity warning` implies critical, error, or warning events, whereas the option `severity warning only` implies warning events only. Severity `all only` is not a valid choice.

Any EMS events with severity `debug-summary`, `debug-verbose`, or `debug-data` will not be logged unless debug mode is enabled. See the command [enable log debug-mode](#) on [page 1350](#).

Filter Optimization

Each time a `configure log filter` command is issued for a given filter name, the events specified are compared against the current configuration of the filter to try to logically simplify the configuration.

For example, if the command:

```
configure log filter bgpFilter1 add events bgp.keepalive severity error only
```

were to be followed by the command:

```
configure log filter bgpFilter1 add events bgp severity info
```

the filter item in the first command is automatically deleted since all events in the *BGP.Keepalive* subcomponent at severity `error` would be also included as part of the second command, making the first command redundant.

More Information

See the command [show log](#) on [page 387](#) for more information about severity levels.

To get a listing of the components present in the system, use the following command:

```
show log components
```

To get a listing of event condition definitions, use the following command:

```
show log events
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

Example

The following command adds all STP component events at severity `info` to the filter *mySTPFilter*:

```
configure log filter myStpFilter add events stp severity info
```

The following command adds the *STP.OutBPDU* subcomponent, at the pre-defined severity level for that component, to the filter *myStpFilter*:

```
configure log filter myStpFilter add events stp.outbpdu
```

The following command excludes one particular event, *STP.InBPDU.Drop*, from the filter:

```
configure log filter myStpFilter add exclude events stp.inbpdu.drop
```

configure log filter events match

```
configure log filter <name> [add | delete] {exclude} events [<event-condition> | [all |  
<event-component>] {severity <severity> {only}}] [match | strict-match] <type> <value>
```

Description

Configures a log filter to add or delete detailed feature messages based on a specified set of events and match parameter values.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

name	Specifies the filter to configure.
add	Add the specified events to the filter.
delete	Remove the specified events from the filter.
exclude	Events matching the filter will be excluded.
event-condition	Specifies the event condition.
all	Specifies all events.
event-component	Specifies all the events associated with a particular component.
severity	Specifies the minimum severity level of events (if the keyword only is omitted).
only	Specifies only events of the specified severity level.
match	Specifies events whose parameter values match the <type> <value> pair.
strict-match	Specifies events whose parameter values match the <type> <value> pair, and possess all the parameters specified.
type	Specifies the type of parameter to match. For more information about types and values see Types and Values on page 341.
value	Specifies the value of the parameter to match. For more information about types and values see Types and Values on page 341.

Default

If the `exclude` keyword is not used, the events will be included by the filter. If `severity` is not specified, then the filter will use the component default severity threshold (see the note on [page 338](#) when `delete` or `exclude` is specified).

Usage Guidelines

This command controls the incidents that pass a filter by adding, or deleting, a specified set of events that match a list of `<type> <value>` pairs. This command is an extension of the command `configure log filter events`, and adds the ability to filter incidents based on matching specified event parameter values to the event.

See the `configure log filter events` command [on page 337](#) for more information on specifying and using filters, on event conditions and components, and on the details of the filtering process. The discussion here is about the concepts of matching `<type> <value>` pairs to more narrowly define filters.

Types and Values

Each event in NETGEAR 8800 is defined with a message format and zero or more parameter types. The `show log events` command [on page 401](#) can be used to display event definitions (the event text and parameter types). The syntax for the parameter types (represented by `<type>` in the command syntax above) is:

```
[address-family [ipv4-multicast | ipv4-unicast | ipv6-multicast | ipv6-unicast]
| bgp-neighbor <ip address>
| bgp-routerid <ip address>
| {destination | source} [ipaddress <ip address> | L4-port | mac-address ]
| {egress | ingress} [slot <slot number> | ports <portlist>]
| ipaddress <ip address>
| L4-port <L4-port>
| mac-address <mac_address>
| netmask <netmask>
| number <number>
| port <portlist>
| process <process name>
| slot <slotid>
| string <exact string to be matched>
| vlan <vlan name>
| vlan tag <vlan tag>]
```

You can specify the `ipaddress` type as IPv4 or IPv6, depending on the IP version. The following examples show how to configure IPv4 addresses and IPv6 addresses:

- IPv4 address

To configure an IP address, with a mask of 32 assumed, use the following command:

```
configure log filter myFilter add events all match ipaddress 12.0.0.1
```

To configure a range of IP addresses with a mask of 8, use the following command:

```
configure log filter myFilter add events all match ipaddress 12.0.0.0/8
```

- IPv6 address

To configure an IPv6 address, with a mask of 128 assumed, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::1
```

To configure a range of IPv6 addresses with a mask of 16, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::/16
```

- IPv6 scoped address

IPv6 scoped addresses consist of an IPv6 address and a VLAN. The following examples identify a link local IPv6 address.

To configure a scoped IPv6 address, with a mask of 128 assumed, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::1%Default
```

To configure a range of scoped IPv6 addresses with a mask of 16, use the following command:

```
configure log filter myFilter add events all match ipaddress
3ffe::/16%Default
```

To configure a scoped IPv6 address with any VLAN, use the following command:

```
configure log filter myFilter add events all match ipaddress 3ffe::/16%*
```

To configure any scoped IPv6 address with a specific VLAN, use the following command:

```
configure log filter myFilter add events all match ipaddress ::/0%Default
```

Note: In the previous example, if you specify the VLAN name, it must be a full match; wild cards are not allowed.

The <value> depends on the parameter type specified. As an example, an event may contain a physical port number, a source MAC address, and a destination MAC address. To allow only those incidents with a specific source MAC address, use the following in the command:

```
configure log filter myFilter add events aaa.radius.requestInit security notice match source
mac-address 00:01:30:23:C1:00
configure log filter myFilter add events bridge severity notice match source mac-address
00:01:30:23:C1:00
```

The string type is used to match a specific string value of an event parameter, such as a user name. The exact string is matched with the given parameter and no regular expression is supported.

Match Versus Strict-Match

The `match` and `strict-match` keywords control the filter behavior for incidents whose event definition does not contain all the parameters specified in a `configure log filter events match` command. This is best explained with an example. Suppose an event in the XYZ component, named `XYZ.event5`, contains a physical port number, a source MAC address, but no destination MAC address. If you configure a filter to match a source MAC address and a destination MAC address, `XYZ.event5` will match the filter when the source MAC address matches regardless of the destination MAC address, since the event contains no destination

MAC address. If you specify the `strict-match` keyword, then the filter will never match, since `XYZ.event5` does not contain the destination MAC address.

In other words, if the `match` keyword is specified, an incident will pass a filter so long as all parameter values in the incident match those in the match criteria, but all parameter types in the match criteria need not be present in the event definition.

More Information

See the command [show log](#) on page 387 for more information about severity levels.

To get a listing of the components present in the system, use the following command:

```
show log components
```

To get a listing of event condition definitions, use the following command:

```
show log events
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

Example

By default, all log targets are associated with the built-in filter, *DefaultFilter*. Therefore, the most straightforward way to send additional messages to a log target is to modify *DefaultFilter*. In the following example, the command modifies the built-in filter to allow incidents in the *STP* component, and all subcomponents of *STP*, of severity critical, error, warning, notice and info. For any of these events containing a physical port number as a match parameter, limit the incidents to only those occurring on physical ports 3, 4 and 5 on slot 1, and all ports on slot 2:

```
configure log filter DefaultFilter add events stp severity info match ports 1:3-1:5, 2:*
```

If desired, issue the `unconfigure log DefaultFilter` command to restore the *DefaultFilter* back to its original configuration.

configure log target filter

```
configure log target [console | memory-buffer | primary-msm | primary-node | backup-msm |
backup-node | nvram | session | syslog [all | <ipaddress> | <ipPort> {vr <vr_name>} [local0
... local7]]] filter <filter-name> {severity <severity> {only}}
```

Description

Associates a filter to a target.

In a stack, this command is applicable only to Master and Backup nodes. This command is not applicable to standby nodes.

Syntax Description

target	Specifies the device to send the log entries.
--------	---

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
primary-msm	Specifies the primary MSM.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.
backup-node	Specifies the backup node in a stack.
nvrn	Specifies the switch NVRAM.
session	Specifies the current session (including console display).
syslog	Specifies a syslog remote server.
all	Specifies all of the syslog remote servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
<p>Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i>.</p>	
local0 ... local7	Specifies the local syslog facility.
filter-name	Specifies the filter to associate with the target.
severity	Specifies the minimum severity level to send (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be sent.

Default

If severity is not specified, the severity level for the target is left unchanged. If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command associates the specified filter and severity with the specified target. A filter limits messages sent to a target.

Although each target can be configured with its own filter, by default, all targets are associated with the built-in filter, *DefaultFilter*. Each target can also be configured with its own severity level. This provides the ability to associate multiple targets with the same filter, while having a configurable severity level for each target.

A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified. By default, the memory buffer, NVRAM, primary MSM/MM, and backup MSM/MM targets are enabled. For other targets, use

the command `enable log target` on page 380. **Table 8** describes the default characteristics of each type of target.

Table 8. Default target log characteristics

Target	Enabled	Severity Level
console display	no	info
memory buffer	yes	debug-data
NVRAM	yes	warning
primary MSM/MM	yes	warning
backup MSM/MM	yes	warning
session	no	info
syslog	no	debug-data

The built-in filter, *DefaultFilter*, and a severity level of `info` are used for each new telnet session. These values may be overridden on a per-session basis using the `configure log target filter` command and specify the target as `session`. Use the following form of the command for per-session configuration changes:

```
configure log target session filter <filter name> {severity <severity> {only}}
```

Configuration changes to the current session target are in effect only for the duration of the session, and are not saved in FLASH memory. The `session` option can also be used on the console display, if the changes are desired to be temporary. If changes to the console-display are to be permanent (saved to FLASH memory), use the following form of the command:

```
configure log target console filter <filter name> {severity <severity> {only}}
```

If the condition for the `backup-msm` target is met by a message generated on the primary, the event is sent to the backup MSM/MM. When the backup MSM/MM receives the event, it will see if any of the local targets (`nvr`, `memory`, or `console`) are matched. If so it gets processed. The `session` and `syslog` targets are disabled on the backup MSM/MM, as they are handled on the primary. If the condition for the `primary-msm` target is met by a message generated on the backup, the event is sent to the primary MSM.

Note that the `backup-msm` target is only active on the primary MSM/MM, and the `primary-msm` target is only active on the backup MSM/MM.

Example

The following command sends log messages to the previously syslog host at 10.31.8.25, port 8993, and facility `local3`, that pass the filter *myFilter* and are of severity `warning` and above:

```
configure log target syslog 10.31.8.25:8993 local3 filter myFilter severity warning
```

The following command sends log messages to the current session, that pass the filter *myFilter* and are of severity `warning` and above:

```
configure log target session filter myFilter severity warning
```

configure log target format

```
configure log target [console | memory-buffer | nvram | session | syslog [all | <ipaddress> |
<ipPort>] {vr <vr_name>} {local0 ... local7}]]
format [timestamp [seconds | hundredths | none]
| date [dd-mm-yyyy | dd-Mmm-yyyy | mm-dd-yyyy | Mmm-dd | yyyy-mm-dd | none] | severity
| event-name [component | condition | none | subcomponent]
| host-name
| priority
| process-name
| process-slot
| source-line
```

Description

Configures the formats of the displayed message, on a per-target basis.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .	
local0 ... local7	Specifies the local syslog facility.
timestamp	Specifies a timestamp formatted to display seconds, hundredths, or none.
date	Specifies a date formatted as specified, or none.
severity	Specifies whether to include the severity.
event-name	Specifies how detailed the event description will be. Choose from none, component, subcomponent, or condition.
host-name	Specifies whether to include the syslog host name.
priority	Specifies whether to include the priority.
process-name	Specifies whether to include the internal process name.

process-slot	Specifies which slot number the message was generated.
source-line	Specifies whether to include the source file name and line number.

Default

The following defaults apply to console display, memory buffer, NVRAM, and session targets:

- timestamp—hundredths
- date—mm-dd-yyyy
- severity—on
- event-name—condition
- host-name—off
- priority—off
- process-name—off
- process-slot—off
- source-line—off

The following defaults apply to syslog targets (per RFC 3164):

- timestamp—seconds
- date—mmm-dd
- severity—on
- event-name—none
- host-name—off
- priority—on
- process-name—off
- process-slot—off
- source-line—off

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command configures the format of the items that make up log messages. You can choose to include or exclude items and set the format for those items, but you cannot vary the order in which the items are assembled.

When applied to the targets `console` or `session`, the format specified is used for the messages sent to the console display or telnet session. Configuration changes to the `session` target, be it either a telnet or console display target session, are in effect only for the duration of the session, and are not saved in FLASH.

When this command is applied to the target `memory-buffer`, the format specified is used in subsequent `show log` and `upload log` commands. The format configured for the internal

memory buffer can be overridden by specifying a format on the `show log` and `upload log` commands.

When this command is applied to the target `syslog`, the format specified is used for the messages sent to the specified syslog host.

Timestamps

Timestamps refer to the time an event occurred, and can be output in either seconds as described in RFC 3164 (for example, “13:42:56”), hundredths of a second (for example, “13:42:56.98”), or suppressed altogether. To display timestamps as hh:mm:ss, use the `seconds` keyword, to display as hh:mm:ss.HH, use the `hundredths` keyword, or to suppress timestamps altogether, use the `none` keyword. Timestamps are displayed in hundredths by default.

Date

The date an event occurred can be output as described in RFC 3164. Dates are output in different formats, depending on the keyword chosen. The following lists the `date` keyword options, and how the date “March 26, 2005” would be output:

- `Mmm-dd`—Mar 26
- `mm-dd-yyyy`—03/26/2005
- `dd-mm-yyyy`—26-03-2005
- `yyyy-mm-dd`—2005-03-26
- `dd-Mmm-yyyy`—26-Mar-2005

Dates are suppressed altogether by specifying `none`. Dates are displayed as `mm-dd-yyyy` by default.

Severity

A four-letter abbreviation of the severity of the event can be output by specifying `severity on` or suppressed by specifying `severity off`. The default setting is `severity on`. The abbreviations are: Crit, Erro, Warn, Noti, Info, Summ, Verb, and Data. These correspond to: Critical, Error, Warning, Notice, Informational, Debug-Summary, Debug-Verbose, and Debug-Data.

Event Names

Event names can be output as the component name only by specifying `event-name component` and as component and subcomponent name with condition mnemonic by specifying `event-name condition`, or suppressed by specifying `event-name none`. The default setting is `event-name condition` to specify the complete name of the events.

Host Name

The configured SNMP name of the switch can be output as `HOSTNAME` described in RFC 3164 by specifying `host-name`. The default setting is `off`.

Process Name

For providing detailed information to technical support, the (internal) NETGEAR 8800 task names of the applications detecting the events can be displayed by specifying `process-name`. The default setting is off.

Process Slot

For providing detailed information to technical support, the slot from which the logged message was generated can be displayed by specifying `process-slot`. The default setting is off.

Process ID

For providing detailed information to technical support, the (internal) NETGEAR 8800 task identifiers of the applications detecting the events can be displayed by specifying `process-id`. The default setting is off.

Source Line

For providing detailed information to technical support, the application source file names and line numbers detecting the events can be displayed by specifying `source-line`. The default setting is off. You must enable debug mode using the `enable log debug-mode` command to view the source line information. For messages generated prior to enabling debug mode, the source line information is not displayed.

Example

In the following example, the switch generates the identical event from the component `SNTP`, using three different formats.

Using the default format for the session target, an example log message might appear as:

```
05/29/2005 12:15:25.00 <Warn:SNTP.RslvSrvrFail> The SNTP server parameter value
(TheWrongServer.example.com) can not be resolved.
```

If you set the current session format using the following command:

```
configure log target session format timestamp seconds date mm-dd-yyyy event-name component
```

The same example would appear as:

```
05/29/2005 12:16:36 <Warn:SNTP> The SNTP server parameter value (TheWrongServer.example.com)
can not be resolved.
```

To provide some detailed information to technical support, you set the current session format using the following command:

```
configure log target session format timestamp hundredths date mmm-dd event-name condition
source-line process-name
```

The same example would appear as:

```
May 29 12:17:20.11 SNTP: <Warn:SNTP.RslvSrvrFail> tSntp: (sntpLib.c:606) The SNTP server
parameter value (TheWrongServer.example.com) can not be resolved.
```

configure log target match

```
configure log target [console | memory-buffer | nvram | primary-msm | primary-node | backup-msm
| backp-node | session | syslog [all | <ipaddress> | <ipPort> {vr <vr_name>} [local0 ...
local7]]] match [any |<match-expression>]
```

Description

Associates a match expression to a target.

In a stack, this command is applicable only on a Master and Backup nodes. This command is not applicable for standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.
backup-node	Specifies the backup-node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .
local0 ... local7	Specifies the local syslog facility.
any	Specifies that any messages will match. This effectively removes a previously configured match expression.
match-expression	Specifies a regular expression. Only messages that match the regular expression will be sent.

Default

By default, targets do not have a match expression. If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command configures the specified target with a match expression. The filter associated with the target is not affected. A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified.

See the command [show log](#) on page 387 for a detailed description of simple regular expressions. By default, targets do not have a match expression.

Specifying `any` instead of `match-expression` effectively removes a match expression that had been previously configured, causing any message to be sent that has satisfied all of the other requirements.

To see the configuration of a target, use the following command:

```
show log configuration target {console | memory-buffer | nvram | primary-msm |
primary-node | backup-msm | backup-node | session | syslog [<ipaddress> | <ipPort> | vr
<vr_name>] {local0 ... local7}}
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

Example

The following command sends log messages to the current session, that pass the current filter and severity level, and contain the string `user5`:

```
configure log target session match user5
```

configure log target severity

```
configure log target [console | memory-buffer | nvram | primary-msm | primary-node |
backup-msm | backup-node | session | syslog [all | <ipaddress> | <ipPort> {vr <vr_name>}
[local0 ... local7]]] {severity <severity> {only}}
```

Description

Sets the severity level of messages sent to the target.

In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
primary-node	Specifies the primary node in a stack.

backup-msm	Specifies the backup MSM.
backup-node	Specifies the backup node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
 Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .	
local0 ... local7	Specifies the local syslog facility.
severity	Specifies the least severe level to send (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be sent.

Default

By default, targets are sent messages of the following severity level and above:

- console display—info
- memory buffer—debug-data
- NVRAM—warning
- session—info
- syslog—debug-data
- primary MSM/MM—warning
- backup MSM/MM—warning
- primary node—warning (stack only)
- backup node—warning (stack only)

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command configures the specified target with a severity level. The filter associated with the target is not affected. A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified.

See the command [show log on page 387](#) for a detailed description of severity levels.

To see the current configuration of a target, use the following command:


```
show log configuration target {console | memory-buffer | nvram | primary-msm |
primary-node | backup-msm | backup-node | session | syslog {<ipaddress> | <ipPort> | vr
<vr_name>} {[local0 ... local7]}}
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

Example

The following command sends log messages to the current session, that pass the current filter at a severity level of info or greater, and contain the string *user5*:

```
configure log target session severity info
```

configure log target syslog

```
configure log target syslog [all | <ipaddress> | <ipPort>] {vr <vr_name>} {local0 ... local7}
from <source-ip-address>
```

Description

Configures the syslog server's IP address for one or all syslog targets.

Syntax Description

syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog server's IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .
local0 ... local7	Specifies the local syslog facility.
source-ip-address	Specifies the local source IP address to use.

Default

If a virtual router is not specified, the VR-Mgmt virtual router is used.

Usage Guidelines

Use this command to identify and configure the syslog server's IP address. By configuring a source IP address, the syslog server can identify from which switch it received the log message.

Options for configuring the remote syslog server include:

- **all**—Specifies all of the remote syslog server hosts.

- `ipaddress`—The IP address of the remote syslog server host.
- `ipPort`—The UDP port.
- `vr_name`—The virtual router that can reach the syslog host.
- `local0-local7`—The syslog facility level for local use.
- `from`—The local source IP address.

If you do not configure a source IP address for the syslog target, the switch uses the IP address in the configured VR that has the closed route to the destination.

Example

The following command configures the IP address for the specified syslog target named *orange*:

```
configure log target syslog orange from 10.234.56.78
```

configure sflow agent ipaddress

```
configure sflow agent {ipaddress} <ip-address>
```

Description

Configures the sFlow agent's IP address.

Syntax Description

<code>ip-address</code>	Specifies the IP address from which sFlow data is sent on the switch.
-------------------------	---

Default

The default configured IP address is 0.0.0.0, but the effective IP address is the management port IP address.

Usage Guidelines

This command allows you to configure the IP address of the sFlow agent. Typically, you would set this to the IP address used to identify the switch in the network management tools that you use. The agent address is stored in the payload of the sFlow data, and is used by the sFlow collector to identify each agent uniquely. The default configured value is 0.0.0.0, but the switch will use the management port IP address if it exists.

The `unconfigure sflow agent` command will reset the agent parameter to the default.

Example

The following command sets the sFlow agent's IP address to 10.2.0.1:

```
configure sflow agent ipaddress 10.2.0.1
```

configure sflow collector ipaddress

```
configure sflow collector {ipaddress} <ip-address> {port <udp-port-number>} {vr <vrname>}
```

Description

Configures the sFlow collector IP address.

Syntax Description

ip-address	Specifies the IP address to send the sFlow data.
udp-port-number	Specifies the UDP port to send the sFlow data.
vrname	Specifies from which virtual router to send the sFlow data.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*.

Default

The following values are the defaults for this command:

- UDP port number—6343
- Virtual router—VR-Mgmt (previously called VR-0).

Usage Guidelines

This command allows you to configure where to send the sFlow data. You must specify an IP address for the sFlow data collector, and you may specify a particular UDP port, if your collector uses a non-standard port. You may also need to specify from which virtual router to send the data.

You can configure up to four sFlow collectors. Each unique IP address/UDP port/virtual router combination identifies a collector.

The `unconfigure sflow collector` command will reset the collector parameters to the default.

Example

The following command specifies that sFlow data should be sent to port 6343 at IP address 192.168.57.1 using the virtual router *VR-Mgmt*:

```
configure sflow collector ipaddress 192.168.57.1
```

configure sflow max-cpu-sample-limit

```
configure sflow max-cpu-sample-limit <rate>
```

Description

Configures the maximum number of sFlow samples handled by the CPU per second.

Syntax Description

rate	Specifies the maximum sFlow samples per second.
------	---

Default

The default value is 2000 samples per second.

Usage Guidelines

This command configures the maximum number of samples sent to the CPU per second. If this rate is exceeded, the internal sFlow CPU throttling mechanism kicks in to limit the load on the CPU.

Every time the limit is reached, the sample rate is halved (the value of `number` in the `configure sflow sample-rate <number>` or `configure sflow ports <portlist> sample-rate <number>` command is doubled) on the slot (modular switch) or ports (stand-alone switch) on which maximum number of packets were received during the last snapshot.

This effectively halves the sampling frequency of all the ports on that slot or stand-alone switch with a sub-sampling factor of 1. The sampling frequency of ports on that slot or stand-alone switch with a sub-sampling factor greater than 1 will not change; the sub-sampling factor is also halved so that the same rate of samples are sent from that port.

The maximum CPU sample rate is based on the total number of samples received from all the sources. The valid range is 100 to 5000 samples per second.

Example

The following command specifies that the sFlow maximum CPU sample rate should be set to 4000 samples per second:

```
configure sflow max-cpu-sample-limit 4000
```

configure sflow poll-interval

```
configure sflow poll-interval <seconds>
```

Description

Configures the sFlow counter polling interval.

Syntax Description

seconds	Specifies the number of seconds between polling each counter. The value can range from 0 to 3600 seconds.
---------	---

Default

The default polling interval is 20 seconds.

Usage Guidelines

Each sFlow statistics counter is polled at regular intervals, and this data is then sent to the sFlow collector. This command is used to set the polling interval. To manage CPU load, polling for sFlow enabled ports are distributed over the polling interval, so that all ports are not polled at the same instant. For example, if the polling interval is 20 seconds and there are twenty counters, data is collected successively every second.

Specifying a poll interval of 0 (zero) seconds disables polling.

Example

The following command sets the polling interval to 60 seconds:

```
configure sflow poll-interval 60
```

configure sflow ports sample-rate

```
configure sflow ports <portlist> sample-rate <number>
```

Description

Configures the sFlow per-port sampling rate.

Syntax Description

portlist	Specifies a list of ports.
number	Specifies the fraction (1/number) of packets to be sampled.

Default

The default number is 8192, unless modified by the `configure sflow sample-rate` command.

Usage Guidelines

This command configures the sampling rate on a particular set of ports and overrides the system-wide value set in the `configure sflow sample-rate` command. The rate is rounded off to the next power of two, so if 400 is specified, the sample rate is configured as 512. The valid range is 1 to 536870912.

All ports on the switch or same I/O module are sampled individually.

Example

The following command sets the sample rate for the ports 4:6 to 4:10 to one packet out of every 16384:

```
configure sflow ports 4:6-4:10 sample-rate 16384
```

configure sflow sample-rate

```
configure sflow sample-rate <number>
```

Description

Configures the sFlow default sampling rate.

Syntax Description

number	Specifies the fraction (1/number) of packets to be sampled.
--------	---

Default

The default number is 8192.

Usage Guidelines

This command configures the default sampling rate. This is the rate that newly enabled sFlow ports will have their sample rate set to. Changing this rate will not affect currently enabled sFlow ports. The rate is rounded off to the next power of two, so if 400 is specified, the sample rate is configured as 512. The valid range is 1 to 536870912.

Configuring a lower number for the sample rate means that more samples will be taken, increasing the load on the switch. Do not configure the sample rate to a number lower than the default unless you are sure that the traffic rate on the source is low.

The minimum rate that these platforms sample is 1 out of every 256 packets. If you configure a rate to be less than 256, the switch automatically rounds up the sample rate to 256.

Example

The following command sets the sample rate to one packet out of every 16384:

```
configure sflow sample-rate 16384
```

configure sys-health-check all level

```
configure sys-health-check all level [normal | strict]
```

Description

Configures how the NETGEAR 8800 software handles faults.

Syntax Description

normal	Upon a fault detection, the switch only sends a message to the syslog. This is the default setting.
--------	---

strict	Upon a fault detection, the switch takes the action configured by the <code>configure sys-recovery-level slot</code> command.
--------	---

Default

The default setting is normal.

Usage Guidelines

On a NETGEAR 8800 series switch, use this command in conjunction with the `configure sys-recovery-level slot [all | <slot_number>] [none | reset | shutdown]` command to implement your network's fault handling strategy.

If you configure the `strict` parameter, the switch takes the action configured by the `configure sys-recovery-level slot` command, which can include logging only or restarting, rebooting, or shutting down the suspect device.

System Behavior for the NETGEAR 8800 Series Switches

Depending on your switch configuration, [Table 9](#) shows how the 8800 series switches behave when the 8800 OS software detects a fault:

Table 9. System behavior for the NETGEAR 8800 series switches

Fault Handling Configuration	Module Recovery Configuration	Behavior
<code>configure sys-health-check all level normal</code>	<code>configure sys-recovery-level slot none</code>	The switch sends messages to the syslog.
Same as above.	<code>configure sys-recovery-level slot reset</code>	Same as above.
Same as above.	<code>configure sys-recovery-level slot shutdown</code>	Same as above.
<code>configure sys-health-check all level strict</code>	<code>configure sys-recovery-level slot none</code>	Same as above.
Same as above.	<code>configure sys-recovery-level slot reset</code>	8800 OS reboots the affected switch or module.
Same as above.	<code>configure sys-recovery-level slot shutdown</code>	8800 OS shuts down the affected switch or module.

Displaying the System Health Check Setting

To display the system health check setting, including polling and how the 8800 OS handles faults on the switch, use the following command:

```
show switch
```

The system health check setting, displayed as `SysHealth check`, shows the polling setting and how NETGEAR 8800 handles faults. The polling setting appears as `Enabled`, and the fault handling setting appears in parenthesis next to the polling setting. In the following truncated output from a NETGEAR 8800 switch, the system health check setting appears as `SysHealth check: Enabled (Normal)`:

```
SysName:          TechPubs Lab
SysName:          BD-8810Rack3
SysLocation:
SysContact:       support@netgear.com
System MAC:       00:04:96:1F:A2:60

SysHealth check:  Enabled (Normal)
Recovery Mode:    None
System Watchdog:  Enabled
```

If you use the `strict` parameter, which configures the switch to take the action configured by the `configure sys-recovery-level slot` command, `(Strict)` would appear next to `Enabled`.

Example

On a NETGEAR 8800 series switch, the following command configures the switch to forward faults to be handled by the level set by the `configure sys-recovery-level slot` command:

```
configure sys-health-check all level strict
```

configure sys-health-check interval

```
configure sys-health-check interval <interval>
```

Description

Configures the frequency of sending backplane diagnostic packets and the polling interval.

Syntax Description

interval	<p>NETGEAR 8800 series switches—Specifies the frequency of sending backplane diagnostic packets.</p> <ul style="list-style-type: none"> • If backplane diagnostic packets are enabled on a particular slot, the default value for sending diagnostic packets is 5 seconds on that slot. • If only polling occurs (this is the system default), the default value is 5 seconds. (The polling interval is not a user-configured parameter, and polling always occurs.)
----------	---

Default

Depending upon your platform, the following defaults apply:

- If backplane diagnostics are enabled on a particular slot, the default for sending packets is 5 seconds on that slot.
- The polling interval is always 5 seconds (this is not a user-configured parameter).

Usage Guidelines

Use this command with the guidance of NETGEAR Technical Support personnel.

The system health checker tests I/O modules and the backplane by forwarding backplane diagnostic packets. Use this command to configure the amount of time it takes for the packets to be forwarded and returned to the MSM.

To enable backplane diagnostic packets, use the `enable sys-health-check slot <slot>` command. With backplane diagnostic packets enabled on a specific slot, the `interval` option of the `configure sys-health-check interval` command specifies the frequency of sending backplane diagnostic packets. For example, if you specify an interval of 9, backplane diagnostic packets are sent every 9 seconds on only the enabled slot.

Note: NETGEAR does not recommend configuring an interval of less than the default interval. Doing this can cause excessive CPU utilization.

By default, the system health checker always polls the control plane health between MSMs and I/O modules, monitors memory levels on the I/O module, monitors the health of the I/O module, and checks the health of applications and processes running on the I/O module. If the system health checker detects an error, the health checker notifies the MSM.

You must enable the backplane diagnostic packets feature to send backplane diagnostic packets. If you enable this feature, the system health checker tests the data link for a specific I/O module every 5 seconds by default. The MSM sends and receives diagnostic packets from the I/O module to determine the state and connectivity. If you disable backplane diagnostics, the system health checker stops sending backplane diagnostic packets.

Example

The following examples assume that you enabled backplane diagnostic packets on a specific I/O slot.

On the NETGEAR 8800 series switches, the following command configures the backplane diagnostic packet interval to 8 seconds:

```
configure sys-health-check interval 8
```

configure sys-recovery-level

```
configure sys-recovery-level [all | none]
```

Description

Configures a recovery option for instances where a software exception occurs in NETGEAR 8800.

Syntax Description

all	Configures the NETGEAR 8800 to log an error into the syslog and reboot the system after any software task exception occurs.
none	Configures the recovery level to none. No action is taken when a software task exception occurs; there is no system reboot, which can cause unexpected switch behavior.

Note: Use this parameter only under the guidance of NETGEAR Technical Support personnel.

Default

The default setting is `all`.

Usage Guidelines

If the software fails, the switch automatically reboots or leaves the system in its current state. You must specify one of the following parameters for the system to respond to software failures:

- `all`—The system will send error messages to the syslog and reboot if any software task exception occurs.

This command sets the recovery level only for the MSMs/MMs. The MSM/MM should reboot only if there is a software exception that occurs on the MSM/MM. The MSM/MM should not reboot if a software exception occurs on an I/O module.

To set the recovery level for all slots (MSM/MM and I/O) use the `configure sys-recovery-level slot` command.

- `none`—No action is taken when a software task exception occurs. The system does not reboot, which can cause unexpected switch behavior.

Note: Use the `none` parameter only under the guidance of NETGEAR Technical Support personnel.

The default setting and behavior is `all`. NETGEAR strongly recommends using the default setting.

Displaying the System Recovery Setting

To display the software recovery setting on the switch, use the following command:

```
show switch
```

This command displays general switch information, including the software recovery level. The following truncated output displays the software recovery setting (displayed as *Recovery Mode*):

```
SysName:          TechPubs Lab
SysLocation:
SysContact:       support@netgear.com
System MAC:       00:04:96:20:B4:13

SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled
```

Note: All platforms display the software recovery setting as *Recovery Mode*.

Example

The following command configures a switch to not take an action when any software task exception occurs:

```
configure sys-recovery-level none
```

configure sys-recovery-level slot

```
configure sys-recovery-level slot [all | <slot_number>] [none | reset | shutdown]
```

Description

Configures a recovery option for instances where an exception occurs on the specified MSM/MM or I/O module.

Syntax Description

all	Specifies all slots of the MSM/MM and I/O module.
slot_number	Specifies the slot of the MSM/MM or I/O module. <ul style="list-style-type: none"> • A and B—Indicate an MSM/MM • 1 through 10—Indicate an I/O module
none	Configures the MSM/MM or I/O module to maintain its current state regardless of the detected hardware fault. The offending MSM/MM or I/O module is not reset. For more information about the states of an MSM/MM or I/O module see the show slot command.

reset	Configures the offending MSM/MM or I/O module to reset upon a hardware fault detection. For more detailed information, see the Usage Guidelines described below.
shutdown	Configures the switch to shut down all slots/modules configured for shutdown upon fault detection. On the modules configured for shutdown, all ports in the slot are taken offline in response to the reported errors; however, the MSMs/MMs remain operational for debugging purposes only. NETGEAR 8800 logs fault, error, system reset, system reboot, and system shutdown messages to the syslog.

Default

The default setting is `reset`.

Usage Guidelines

Use this command for system auto-recovery upon detection of hardware problems. You can configure the MSMs/MMs or I/O modules to take no action, automatically reset, shutdown, or if dual MSMs/MMs are installed, failover to the other MSM/MM if the switch detects a faulty MSM/MM or I/O module. This enhanced level of recovery detects faults in the ASICs as well as packet buses.

You must specify one of the following parameters for the system to respond to MSM/MM or I/O module failures:

- `none`—Configures the MSM/MM or I/O module to maintain its current state regardless of the detected fault. The offending MSM/MM or I/O module is not reset. NETGEAR 8800 logs fault and error messages to the syslog and notifies you that the errors are ignored. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
- `reset`—Configures the offending MSM/MM or I/O module to reset upon fault detection. NETGEAR 8800 logs fault, error, system reset, and system reboot messages to the syslog.
- `shutdown`—Configures the switch to shut down all slots/modules configured for shutdown upon fault detection. On the modules configured for shutdown, all ports in the slot are taken offline in response to the reported errors; however, the MSMs/MMs remain operational for debugging purposes only. You must save the configuration, using the [save configuration](#) command, for it to take effect. NETGEAR 8800 logs fault, error, system reset, system reboot, and system shutdown messages to the syslog.

Depending on your configuration, the switch resets the offending MSM/MM or I/O module if fault detection occurs. An offending MSM/MM is reset any number of times, and the MSM/MM is not permanently taken offline. An offending I/O module is reset a maximum of five times. After the maximum number of resets, the I/O module is permanently taken offline.

Messages Displayed

If you configure the hardware recovery setting to either none (ignore) or shutdown, the switch prompts you to confirm this action. The following is a sample shutdown message:

Are you sure you want to shutdown on errors? (y/n)

Enter `y` to confirm this action and configure the hardware recovery level. Enter `n` or press [Enter] to cancel this action.

Taking Ports Offline

You can configure the switch to shut down one or more modules upon fault detection by specifying the `shutdown` option. If you configure one or more slots to shut down and the switch detects a hardware fault, all ports in all of the configured shut down slots are taken offline in response to the reported errors. (MSMs are available for debugging purposes only.)

The affected module remains in the shutdown state across additional reboots or power cycles until you explicitly clear the shutdown state. If a module enters the shutdown state, the module actually reboots and the `show slot` command displays the state of the slot as `Initialized`; however, the ports are shut down and taken offline. For more information about clearing the shutdown state, see the `clear sys-recovery-level` command.

Module Recovery Actions

Table 10 describes the actions module recovery takes based on your module recovery setting. For example, if you configure a module recovery setting of `reset` for an I/O module, the module is reset a maximum of five times before it is taken permanently offline.

From left to right, the columns display the following information:

- **Module Recovery Setting**—This is the parameter used by the `configure sys-recovery-level slot` command to distinguish the module recovery behavior.
- **Hardware**—This indicates the hardware that you may have in your switch.
- **Action Taken**—This describes the action the hardware takes based on the module recovery setting.

Table 10. Module Recovery Actions for the NETGEAR 8800 Series Switches

Module Recovery Setting	Hardware	Action Taken
none		
	Single MSM	The MSM remains powered on in its current state. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
	Dual MSM	The MSM remains powered on in its current state. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
	I/O Module	The I/O module remains powered on in its current state. The switch sends error messages to the log and notifies you that the errors are ignored. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
reset		

Table 10. Module Recovery Actions for the NETGEAR 8800 Series Switches (Continued)

Module Recovery Setting	Hardware	Action Taken
	Single MSM	Resets the MSM.
	Dual MSM	Resets the primary MSM and fails over to the backup MSM.
	I/O Module	Resets the I/O module a maximum of five times. After the fifth time, the I/O module is permanently taken offline.
shutdown		
	Single MSM	The MSM is available for debugging purposes only (the I/O ports also go down); however, you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the MSM to become operational. After you clear the shutdown state, you must reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
	Dual MSM	The MSM is available for debugging purposes only (the I/O ports also go down); however, you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the MSM to become operational. After you clear the shutdown state, you must reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
	I/O Module	Reboots the I/O module. When the module comes up, the ports remain inactive because you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the I/O module to become operational. After you clear the shutdown state, you must reset each affected I/O module or reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.

Displaying the Module Recovery Setting

To display the module recovery setting, use the following command:

```
show slot
```

The `show slot` output has been modified to include the shutdown configuration. If you configure the module recovery setting to shutdown, the output displays an “E” flag that indicates any errors detected on the slot disables all ports on the slot. The “E” flag appears only if you configure the module recovery setting to shutdown.

Note: If you configure one or more slots for shut down and the switch detects a hardware fault on one of those slots, all of the configured slots enter the shutdown state and remain in that state until explicitly cleared.

If you configure the module recovery setting to none, the output displays an “e” flag that indicates no corrective actions will occur for the specified MSM/MM or I/O module. The “e” flag appears only if you configure the module recovery setting to none.

The following sample output displays the module recovery action. In this example, notice the flags identified for slot 10:

Slots	Type	Configured	State	Ports	Flags

Slot-1	XCM88P	XCM88P	Operational	48	MB S
Slot-2	XCM8824F	XCM8824F	Operational	24	MB S
Slot-3	XCM8848T	XCM8848T	Operational	48	MB S
Slot-4			Empty	0	
Slot-5	XCM8808X	XCM8808X	Operational	8	MB S
Slot-6	XCM8808X	XCM8808X	Operational	8	MB S
Slot-7			Empty	0	
Slot-8	XCM8848T	XCM8848T	Operational	48	MB S
Slot-9	XCM8848T		Operational	48	MB S
Slot-10	XCM8848T	XCM8848T	Operational	48	MB S E
MSM-A	XCM88S1		Operational	0	S
MSM-B	XCM88S1		Operational	0	S

Flags : M - Backplane link to Master MSM is Active
 B - Backplane link to Backup MSM is also Active
 D - Slot Disabled, S - Slot Secured
 I - Insufficient Power (refer to "show power budget")
 e - Errors on slot will be ignored (no corrective action initiated)
 E - Errors on slot will disable all ports on slot

Displaying Detailed Module Recovery Information

To display the module recovery setting for a specific port on a module, including the current recovery mode, use the following command:

```
show slot <slot>
```

In addition to the information displayed with show slot, this command displays the module recovery setting configured on the slot. The following truncated output displays the module recovery setting (displayed as `Recovery Mode`) for the specified slot:

```
Slot-6 information:
  State:                Operational
  Download %:          100
```

```

Flags:                M
Restart count:        0 (limit 5)
Serial number:         800421-00 000000000000
Hw Module Type:        XCM8848T(P)
SW Version:            12.4.4.0
SW Build:              v1244b0-br-SR3-1
Configured Type:        XCM8848T(P)
Ports available:        48
Recovery Mode:         Reset

```

```

Flags : M - Backplane link to Master is Active
        B - Backplane link to Backup is also Active
        D - Slot Disabled, S - Slot Secured
        I - Insufficient Power (refer to "show power budget")

```

Troubleshooting Module Failures

If you experience an I/O module failure, use the following troubleshooting methods when you can bring the switch offline to solve or learn more about the problem:

- **Restarting the I/O module**—Use the `disable slot <slot>` command followed by the `enable slot <slot>` command to restart the offending I/O module. By issuing these commands, the I/O module and its associated fail counter is reset. If the module does not restart, or you continue to experience I/O module failure, please contact NETGEAR Technical Support.
- **Running diagnostics**—Use the `run diagnostics normal <slot>` command to run operational diagnostics on the offending I/O module to ensure that you are not experiencing a hardware issue. If the module continues to enter the failed state, please contact NETGEAR Technical Support.

If you experience an MSM/MM failure, please contact NETGEAR Technical Support.

Example

The following command configures a switch to not take an action if a hardware fault occurs:

```
configure sys-recovery-level slot none
```

configure syslog add

```
configure syslog add [<ipaddress> | <ipPort>] {vr <vr_name>} [local0 ... local7] {<severity>}
```

Description

Configures the remote syslog server host address, and filters messages to be sent to the remote syslog target.

Syntax Description

ipaddress	Specifies the remote syslog server IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
<p>Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i>.</p>	
local0 ... local7	Specifies the local syslog facility.
severity	Specifies a message severity. Severities include critical, error, warning, notice, info, debug-summary, debug-verbose, and debug-data.

Default

If a severity level is not specified, all messages are sent to the remote syslog server target. If a virtual router is not specified, VR-Mgmt is used. If UDP port is not specified, 514 is used.

Usage Guidelines

Options for configuring the remote syslog server include:

- ipaddress—The IP address of the remote syslog server host.
- ipPort—The UDP port.
- local0-local7—The syslog facility level for local use.
- vr_name—The virtual router that can reach the syslog host.
- severity—Filters the messages sent to the remote syslog server target to have the selected severity or higher (more critical). Severities include critical, error, warning, notice, info, debug-summary, debug-verbose, and debug-data.

The switch log overwrites existing log messages in a wrap-around memory buffer, which may cause you to lose valuable information once the buffer becomes full. The remote syslog server does not overwrite log information, and can store messages in non-volatile files (disks, for example).

The `enable syslog` command must be issued in order for messages to be sent to the remote syslog server(s). Syslog is disabled by default. A total of four syslog servers can be configured at one time.

When a syslog server is added, it is associated with the filter *DefaultFilter*. Use the `configure log target filter` command to associate a different filter.

The syslog facility level is defined as local0 – local7. The facility level is used to group syslog data.

Example

The following command configures the remote syslog server target with a critical severity:

```
configure syslog 123.45.67.78 local1 critical
```

configure syslog delete

```
configure syslog delete [all | <ipaddress> | <ipPort>] {vr <vr_name>} {local0 ... local7}
configure syslog delete <host name/ip> {: <udp-port>} [local0 ... local7]
```

Description

Deletes a remote syslog server address.

Syntax Description

all	Specifies all remote syslog servers.
ipaddress	Specifies the remote syslog server IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .
local0 ... local7	Specifies the local syslog facility.

Default

If a virtual router is not specified, VR-Mgmt is used.

If a UDP port number is not specified, 514 is used.

Usage Guidelines

This command is used to delete a remote syslog server target.

Example

The following command deletes the remote syslog server with an IP address of 10.0.0.1:

```
configure syslog delete 10.0.0.1 local1
```

create log filter

```
create log filter <name> {copy <filter name>}
```

Description

Creates a log filter with the specified name.

Syntax Description

name	Specifies the name of the filter to create.
copy	Specifies that the new filter is to be copied from an existing one.
filter name	Specifies the existing filter to copy.

Default

N/A.

Usage Guidelines

This command creates a filter with the name specified. A filter is a customizable list of events to include or exclude, and optional parameter values. The list of events can be configured by component or subcomponent with optional severity, or individual condition, each with optional parameter values. See the commands [configure log filter events](#) and [configure log filter events match](#) for details on how to add items to the filter.

The filter can be associated with one or more targets using the [configure log target filter](#) command to control the messages sent to those targets. The system has one built-in filter named *DefaultFilter*, which itself may be customized. Therefore, the [create log filter](#) command can be used if a filter other than *DefaultFilter* is desired. As its name implies, *DefaultFilter* initially contains the default level of logging in which every NETGEAR 8800 component and subcomponent has a pre-assigned severity level.

If another filter needs to be created that will be similar to an existing filter, use the `copy` option to populate the new filter with the configuration of the existing filter. If the `copy` option is not specified, the new filter will have no events configured and therefore no incidents will pass through it.

The total number of supported filters, including *DefaultFilter*, is 20.

Example

The following command creates the filter named *fdb2*, copying its configuration from the filter *DefaultFilter*.

```
create log filter fdb2 copy DefaultFilter
```

delete log filter

```
delete log filter [<filter name> | all]
```

Description

Deletes a log filter with the specified name.

Syntax Description

filter name	Specifies the filter to delete.
all	Specifies that all filters, except DefaultFilter, are to be deleted

Default

N/A.

Usage Guidelines

This command deletes the specified filter, or all filters except for the filter *DefaultFilter*. The specified filter must not be associated with a target. To remove that association, associate the target with *DefaultFilter* instead of the filter to be deleted, using the following command:

```
configure log target <target> filter DefaultFilter
```

Example

The following command deletes the filter named *fdb2*:

```
delete log filter fdb2
```

disable cli-config-logging

```
disable cli-config-logging
```

Description

Disables the logging of CLI configuration commands to the switch Syslog.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Every command is displayed in the log window which allows you to view every command executed on the switch.

The `disable cli-config-logging` command discontinues the recording of all switch configuration changes and their sources that are made using the CLI via Telnet or the local console. After you disable configuration logging, no further changes are logged to the system log.

To view the status of configuration logging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for configuration logging.

Example

The following command disables the logging of CLI configuration command to the Syslog:

```
disable cli-config-logging
```

disable log display

```
disable log display
```

Description

Disables the sending of messages to the console display.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If the log display is disabled, log information is no longer written to the serial console.

This command setting is saved to FLASH and determines the initial setting of the console display at boot up.

You can also use the following command to control logging to different targets:

```
disable log display
```

The `disable log display` command is equivalent to `disable log target console-display` command.

Example

The following command disables the log display:

```
disable log display
```

disable log target

```
disable log target [console | memory-buffer | nvram | primary-msm | primary-node | backup-msm  
| backup-node | session | syslog [all | <ipaddress> | <ipPort>] {vr <vr_name>} [local0 ...  
local7]]]
```

Description

Stops sending log messages to the specified target.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvrn	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.
backup-node	Specifies the backup node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog host name or IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
<p>Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i>.</p>	
local0 ... local7	Specifies the local syslog facility.

Default

Enabled, for memory buffer, NVRAM, primary MSM, and backup MSM/MM; all other targets are disabled by default.

Usage Guidelines

This command stops sending messages to the specified target. By default, the memory buffer, NVRAM, primary MSM/MM, and backup MSM/MM targets are enabled. Other targets must be enabled before messages are sent to those targets.

Configuration changes to the `session` target are in effect only for the duration of the console display or telnet session, and are not saved in FLASH. Changes to the other targets are saved to FLASH.

You can also use the following command to disable displaying the log on the console:

```
disable log display
```

The `disable log display` command is equivalent to `disable log target console-display` command.

Note that the `backup-msm` target is only active on the primary MSM/MM, and the `primary-msm` target is only active on the backup MSM/MM.

Example

The following command disables log messages to the current session:

```
disable log target session
```

disable rmon

```
disable rmon
```

Description

Disables the collection of RMON statistics on the switch.

Syntax Description

This command has no arguments or variables.

Default

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events.

Usage Guidelines

The switch supports four out of nine groups of Ethernet RMON statistics. In a disabled state, the switch continues to respond queries of statistics. Collecting of history, alarms, and events is stopped; however, the switch still queries old data.

To view the status of RMON polling on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for RMON polling.

To view the RMON memory usage statistics for a specific memory type (for example, statistics, events, logs, history, or alarms) or for all memory types, use the following command:

```
show rmon memory {detail | <memoryType>}
```

Example

The following command disables the collection of RMON statistics on the switch:

```
disable rmon
```

disable sflow

```
disable sflow
```

Description

Globally disables sFlow statistical packet sampling.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables sFlow globally on the switch. When you disable sFlow globally, the individual ports are also put into the disabled state. If you later enable the global sFlow state, individual ports return to their previous state.

Example

The following command disables sFlow sampling globally:

```
disable sflow
```

disable sflow ports

```
disable sflow ports <portlist>
```

Description

Disables sFlow statistical packet sampling and statistics gathering on a particular list of ports.

Syntax Description

portlist	Specifies a list of ports.
----------	----------------------------

Default

Disabled.

Usage Guidelines

This command disables sFlow on a particular list of ports. Once sFlow is disabled on a port, sampling and polling will stop. If sFlow is disabled globally, all sampling and polling stops.

Use the following command to disable sFlow globally:

```
disable sflow
```


Example

The following command disables sFlow sampling on port 3:1:

```
disable sflow ports 3:1
```

disable sys-health-check

```
disable sys-health-check slot <slot>
```

Description

Discontinues sending backplane diagnostic packets.

Syntax Description

slot	Specifies the slot to disable sending backplane diagnostic packets.
------	---

Default

Polling is enabled, backplane diagnostic packets are disabled.

Depending upon your platform, when disabling backplane diagnostic packets, note that by default the system health checker discontinues sending backplane diagnostic packets to the specified slot. Only polling is enabled.

Usage Guidelines

When you use this command, backplane diagnostic packets are disabled and no longer sent by the system health checker.

If you modify the `interval` in the `configure sys-health-check interval <interval>` command and later disable backplane diagnostics, the configured interval for sending backplane diagnostic packets remains. The next time you enable backplane diagnostic packets, the health checker sends backplane diagnostics packets at the configured interval. For example, if you configure an interval of 8 seconds, the system health checker sends backplane diagnostic packets every 8 seconds.

To return to the "default" interval of 5 seconds, configure the frequency of sending backplane diagnostic packets to 5 seconds using the following command:

```
configure sys-health-check interval 5
```

Example

On the NETGEAR 8800 series switches, the following example assumes that you did not modify the `interval` option in the `configure sys-health-check interval <interval>` command.

The following command disables backplane diagnostics on slot 3, polling is always enabled and occurs every 5 seconds.

```
disable sys-health-check slot 3
```

disable syslog

```
disable syslog
```

Description

Disables logging to all remote syslog server targets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Disables logging to all remote syslog server targets, not to the switch targets. This setting is saved in FLASH, and will be in effect upon boot up.

Example

The following command disables logging to all remote syslog server targets:

```
disable syslog
```

enable cli-config-logging

```
enable cli-config-logging
```

Description

Enables the logging of CLI configuration commands to the Syslog for auditing purposes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

NETGEAR 8800 allows you to record all configuration changes and their sources that are made using the CLI by way of Telnet or the local console. The changes are logged to the system log. Each log entry includes the user account name that performed the changes and the source IP address of the client (if Telnet was used). Configuration logging applies only to commands that result in a configuration change.

To view the status of configuration logging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for configuration logging.

Example

The following command enables the logging of CLI configuration commands to the Syslog:

```
enable cli-config-logging
```

enable log display

```
enable log display
```

Description

Enables a running real-time display of log messages on the console display.

In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If you enable the log display on a terminal connected to the console port, your settings will remain in effect even after your console session is ended (unless you explicitly disable the log display).

You configure the messages displayed in the log using the `configure log display`, or `configure log target console-display` commands.

You can also use the following command to control logging to different targets:

```
enable log display
```

The `enable log display` command is equivalent to `enable log target console-display` command.

To change the log filter association, severity threshold, or match expression for messages sent to the console display, use the `configure log target console-display` command

Example

The following command enables a real-time display of log messages:

```
enable log display
```

enable log target

```
enable log target [console | memory-buffer | nvram | primary-msm | primary-node | backup-msm |
backup-node | session | syslog [all | <ipaddress> | <ipPort>] {vr <vr_name>} [local0 ...
local7]]]
```

Description

Starts sending log messages to the specified target.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
primary-node	Specifies the primary node of a stack.
backup-msm	Specifies the backup MSM.
backup-node	Specifies the backup node of a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .	
local0 ... local7	Specifies the local syslog facility.

Default

Enabled for memory buffer and NVRAM; all other targets are disabled by default.

Usage Guidelines

This command starts sending messages to the specified target. By default, the memory-buffer, NVRAM, primary MSM/MM, and backup MSM/MM targets are enabled. Other targets must be enabled before messages are sent to those targets.

Configuration changes to the `session` target are in effect only for the duration of the console display or Telnet session, and are not saved in FLASH. Others are saved in FLASH.

You can also use the following command to enable displaying the log on the console:

```
enable log display
```

The `enable log display` command is equivalent to the `enable log target console-display` command.

Note that the `backup-msm` target is only active on the primary MSM/MM, and the `primary-msm` target is only active on the backup MSM/MM.

Example

The following command enables log messages on the current session:

```
enable log target session
```

enable rmon

```
enable rmon
```

Description

Enables the collection of RMON statistics on the switch.

Syntax Description

This command has no arguments or variables.

Default

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events. By enabling RMON, the switch begins the processes necessary for collecting switch statistics.

Usage Guidelines

The switch supports four out of nine groups of Ethernet RMON statistics. In an enabled state, the switch responds to the following four groups:

- **Statistics**—The RMON Ethernet Statistics group provides traffic and error statistics showing packets, bytes, broadcasts, multicasts, and errors on a LAN segment or VLAN.
- **History**—The History group provides historical views of network performance by taking periodic samples of the counters supplied by the Statistics group. The group features user-defined sample intervals and bucket counters for complete customization of trend analysis.
- **Alarms**—The Alarms group provides a versatile, general mechanism for setting threshold and sampling intervals to generate events on any RMON variable. Both rising and falling thresholds are supported, and thresholds can be on the absolute value of a variable or its delta value. In addition, alarm thresholds may be auto calibrated or set manually.
- **Events**—The Events group creates entries in an event log and/or sends SNMP traps to the management workstation. An event is triggered by an RMON alarm. The action taken can be configured to ignore it, to log the event, to send an SNMP trap to the receivers

listed in the trap receiver table, or to both log and send a trap. The RMON traps are defined in RFC 1757 for rising and falling thresholds.

The switch also supports the following parameters for configuring the RMON agent, as defined in RFC 2021:

- **probeCapabilities**—If you configure the probeCapabilities object, you can view the RMON MIB groups supported on at least one interface by the probe.
- **probeSoftwareRev**—If you configure the probeSoftwareRev object, you can view the current software version of the monitored device.
- **probeHardwareRev**—If you configure the probeHardwareRev object, you can view the current hardware version of the monitored device.
- **probeDateTime**—If you configure the probeDateTime object, you can view the current date and time of the probe.
- **probeResetControl**—If you configure the probeResetControl object, you can restart a managed device that is not running normally. Depending on your configuration, you can do one of the following:
 - **Warm boot**—A warm boot restarts the device using the current configuration saved in non-volatile memory.
 - **Cold boot**—A cold boot causes the device to reset the configuration parameters stored in non-volatile memory to the factory defaults and then restarts the device using the restored factory default configuration.

Note: You can only use the RMON features of the system if you have an RMON management application and have enabled RMON on the switch.

RMON requires one probe per LAN segment, and stand-alone RMON probes have traditionally been expensive. Therefore, the approach taken by NETGEAR has been to build an inexpensive RMON probe into the agent of each system. This allows RMON to be widely deployed around the network without costing more than traditional network management. The switch accurately maintains RMON statistics at the maximum line rate of all of its ports.

For example, statistics can be related to individual ports. Also, because a probe must be able to see all traffic, a stand-alone probe must be attached to a nonsecure port. Implementing RMON in the switch means that all ports can have security features enabled.

To view the status of RMON polling on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for RMON polling.

To view the RMON memory usage statistics for a specific memory type (for example, statistics, events, logs, history, or alarms) or for all memory types, use the following command:

```
show rmon memory {detail | <memoryType>}
```

Example

The following command enables the collection of RMON statistics on the switch:

```
enable rmon
```

enable sflow

```
enable sflow
```

Description

Globally enables sFlow statistical packet sampling.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables sFlow globally on the switch.

Example

The following command enables sFlow sampling globally:

```
enable sflow
```

enable sflow ports

```
enable sflow ports <port_list>
```

Description

Enables sFlow statistical packet sampling on a particular list of ports.

Syntax Description

port_list	Specifies a list of ports.
-----------	----------------------------

Default

Disabled.

Usage Guidelines

This command enables sFlow on a particular list of ports. You also need to enable sFlow globally in order to gather statistics and send the data to the collector. Once sFlow is enabled globally, and on the ports of interest, sampling and polling begins.

Use the following command to enable sFlow globally:

```
enable sflow
```

Example

The following command enables sFlow sampling on the port 3:1:

```
enable sflow ports 3:1
```

enable sys-health-check

```
enable sys-health-check slot <slot>
```

Description

Enables backplane diagnostic packets on the specified slot.

Syntax Description

slot	Specifies the slot to participate in sending backplane diagnostic packets.
------	--

Default

Polling is enabled, backplane diagnostic packets are disabled.

Depending upon your platform, when you enable diagnostic packets, the system health checker tests the data link every 5 seconds for the specified slot.

Usage Guidelines

Configure the system health checker with guidance from NETGEAR Technical Support personnel.

The system health checker tests I/O modules and the backplane by sending diagnostic packets. By isolating faults to a specific module or backplane connection, the system health checker notifies you of a possible hardware failure.

System health check errors are reported to the syslog. Syslog output includes the slot number where the problem occurred, the loopback packet ID number, and a notification that the MSM/MM did not receive the last packet. If you see an error, please contact NETGEAR Technical Support.

Note: Enabling backplane diagnostic packets increases CPU utilization and competes with network traffic for resources.

The system health checker continues to periodically forward test packets to failed components.

To configure the frequency of the backplane diagnostic packets on the NETGEAR 8800 series switches, use the `configure sys-health-check interval` command.

Displaying the System Health Check Setting

To display the system health check polling setting on the switch, use the following command:

```
show switch
```

As previously described, polling is always enabled on the switch, which is why you see the system health check setting as Enabled. The following truncated output from a NETGEAR 8810 switch displays the system health check setting (displayed as `SysHealth check`):

```
SysName:          XCM8810
SysLocation:
SysContact:       support@netgear.com
System MAC:       00:04:96:1F:A2:60

SysHealth check:  Enabled
Recovery Mode:    None
System Watchdog:  Enabled
```

Example

The following command enables backplane diagnostic packets on slot 6:

```
enable sys-health-check slot 6
```

enable syslog

```
enable syslog
```

Description

Enables logging to all remote syslog host targets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

To enable remote logging, you must do the following:

- Configure the syslog host to accept and log messages.
- Enable remote logging by using the `enable syslog` command.
- Configure remote logging by using the `configure syslog` command.

When you use the `enable syslog` command, the exporting process of the syslog begins. This command also determines the initial state of an added remote syslog target.

Example

The following command enables logging to all remote syslog hosts:

```
enable syslog
```

show fans

```
show fans {detail}
```

Description

Displays the status of the fans in the system.

Syntax Description

detail	The detail option is reserved for future use.
--------	---

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the fans.

This status information may be useful for your technical support representative if you have a network problem.

The switch collects and displays the following fan information:

- State—The current state of the fan. Options are:
 - Empty: There is no fan installed.
 - Failed: The fan failed.
 - Operational: The fan is installed and working normally.
- NumFan—The number of fans in the fan tray.

- Fan Name, displayed as Fan-1, Fan-2, and so on (and a description of the location, for example, Upper or Upper-Right)—Specifies the individual state for each fan in a fan tray and its current speed in revolutions per minute (rpm).

The output also includes the following information:

- PartInfo—Information about the fan tray, including the:
 - Serial number—A collection of numbers and letters, that make up the serial number of the fan. This is the first series of numbers and letters in the display.
 - Part number—A collection of numbers and letters, that make up the part number of the fan. This is the second series of numbers and letters in the display.
- Revision—The revision number of the fan.
- Odometer—Specifies the power-on date and how long the fan tray has been operating since it was first powered-on.

Example

The following command displays the status of the installed fans. If a fan is not installed, the state of the fan is `Empty`.

```
show fans
```

The following is sample output from a NETGEAR 8800 series switch:

```
FanTray information:
State:                Operational
NumFan:               9
PartInfo:             0404X-00015 450102-00-01
Revision:             1.0
Odometer:             111 days 16 hours 30 minutes  since Oct-13-2004
Upper-Left  Fan-1:    Operational at 2880 RPM
Middle-Left Fan-2:    Operational at 2820 RPM
Lower-Left  Fan-3:    Operational at 2820 RPM
Upper-Center Fan-4:   Operational at 2820 RPM
Center      Fan-5:    Operational at 2820 RPM
Lower-Center Fan-6:   Operational at 2880 RPM
Upper-Right Fan-7:    Operational at 2880 RPM
Middle-Right Fan-8:   Operational at 2820 RPM
Lower-Right Fan-9:    Operational at 2880 RPM
```

show log

```
show log {messages [memory-buffer | nvram]} {events {<event-condition> | <event-component>}}
{severity <severity> {only}} {starting [date <date> time <time> | date <date> | time <time>]}
{ending [date <date> time <time> | date <date> | time <time>]} {match <regex>} {chronological}
```

Description

Displays the current log messages.

Syntax Description

messages	Specifies the target location from which to display the log messages.
memory-buffer	Show messages stored in volatile memory (default).
nvrn	Show messages stored in NVRAM.
events	Show event messages.
event-condition	Specifies the event condition to display.
event-component	Specifies the event component to display.
severity	Specifies the minimum severity level to display (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be displayed
starting	Show messages with timestamps equal to or greater than that specified
date	Specifies the date, where date is <month (1-12)> / <day (1-31)> {/ <year (yyyy)>}. <month> {<month>}
time	Specifies the time, where time is <hour (0-23)> {: <minute (0-59)> {: <seconds (0-59)> {, <hundredths>}}}
ending	Show messages with timestamps equal to or less than that specified.
regex	Specifies a regular expression. Only messages that match the regular expression will be displayed.
chronological	Specifies displaying log messages in ascending chronological order (oldest to newest).

Default

The following defaults apply:

- messages—memory buffer
- event—no restriction (displays user-specified event)
- severity—none (displays everything stored in the target)
- starting, ending—if not specified, no timestamp restriction
- match—no restriction
- chronological—if not specified, show messages in order from newest to oldest

Usage Guidelines

Switch configuration and fault information is filtered and saved to target logs, in a memory buffer, and in NVRAM. Each entry in the log contains the following information:

- Timestamp—records the month and day of the event, along with the time (hours, minutes, seconds, and hundredths).

- **Severity Level**—indicates the urgency of a condition reported in the log. [Table 11](#) describes the severity levels assigned to events.
- **Component, Subcomponent, and Condition Name**—describes the subsystem in the software that generates the event. This provides a good indication of where a fault might lie.
- **Message**—a description of the event occurrence. If the event was caused by a user, the user name is also provided.

This command displays the messages stored in either the internal memory buffer or in NVRAM. The messages shown can be limited by specifying a severity level, a time range, or a match expression. Messages stored in the target have already been filtered as events occurred, and specifying a severity or match expression on the `show log` command can only further limit the messages shown.

If the `messages` keyword is not present, the messages stored in the memory-buffer target are displayed. Otherwise, the messages stored in the specified target are displayed.

If the `only` keyword is present following the severity value, then only the events at that exact severity are included. Without the `only` keyword, events at that severity or more urgent are displayed. For example, severity `warning` implies critical, error, or warning, whereas severity `warning only` implies only warning.

Messages whose timestamps are equal or later than the starting time and are equal or earlier than the specified ending time will be shown if they also pass the severity requirements and match expression, if specified.

If a `match` phrase is specified, the formatted message must match the simple regular expression specified by `match-expression` for it to be shown.

A simple regular expression is a string of single characters including the dot character (`.`), which are optionally combined with quantifiers and constraints. A dot matches any single character while other characters match only themselves (case is significant). Quantifiers include the star character (`*`) that matches zero or more occurrences of the immediately preceding character or dot. Constraints include the caret character (`^`) that matches at the beginning of a message, and the currency character (`$`) that matches at the end of a message. Bracket expressions are not supported. There are a number of sources available on the Internet and in various language references describing the operation of regular expressions.

If the `chronological` keyword is specified, messages are shown from oldest to newest; otherwise, messages are displayed newest to oldest.

Severity Level

The severity levels are `critical`, `error`, `warning`, `notice`, and `info`, plus three severity levels for extended debugging, `debug-summary`, `debug-verbose`, and `debug-data`. In log messages, the severity levels are shown by four letter abbreviations. The abbreviated forms are:

- Critical—Crit
- Error—Erro
- Warning—Warn

- Notice—Noti
- Info—Info
- Debug-Summary—Summ
- Debug-Verbose—Verb
- Debug-Data—Data

The three severity levels for extended debugging, `debug-summary`, `debug-verbose`, and `debug-data`, require that debug mode be enabled (which may cause a performance degradation). See the command `enable log debug-mode` on page 1350. [Table 11](#) describes the security levels.

Table 11. Severity Levels Assigned by the Switch

Level	Description
Critical	A serious problem has been detected that is compromising the operation of the system and that the system cannot function as expected unless the situation is remedied. The switch may need to be reset.
Error	A problem has been detected that is interfering with the normal operation of the system and that the system is not functioning as expected.
Warning	An abnormal condition, not interfering with the normal operation of the system, has been detected that may indicate that the system or the network in general may not be functioning as expected.
Notice	A normal but significant condition has been detected, which signals that the system is functioning as expected.
Info (Informational)	A normal but potentially interesting condition has been detected, which signals that the system is functioning as expected and simply provides information or confirmation about the condition.
Debug-Summary	A condition has been detected that may interest a developer determining the reason underlying some system behavior.
Debug-Verbose	A condition has been detected that may interest a developer analyzing some system behavior at a more verbose level than provided by the debug summary information.
Debug-Data	A condition has been detected that may interest a developer inspecting the data underlying some system behavior.

Messages stored in NVRAM are in encoded format. To restore the ASCII text of a message, the version of the NETGEAR 8800 loaded must be able to interpret the data written prior to reboot. When the encoded format for a particular message cannot be interpreted by the version of the NETGEAR 8800 currently loaded, the messages are displayed in the following format:

```
03/21/2005 17:15:37.36 : NO MESSAGE DECODE; Missing component "epm" v24.2
DUMP-10: 00 14 C3 C1 00 11 00 1C 01 FF 00 08 65 70 6D 00 '.....epm.'
DUMP-20: 08 FF 00 0C 00 18 00 02 65 70 6D 00 '.....epm.'
```

Log entries remain in the NVRAM log after a switch reboot. Issuing a `clear log` command does not remove these static entries. To remove log entries from NVRAM, use the following command:

```
clear log messages nvram
```

Example

The following command displays messages with a critical severity:

```
show log severity critical
```

The following command displays messages with warning, error, or critical severity:

```
show log severity warning
```

The following is sample output:

```
11/12/2004 00:38:10.30 <Warn:dm.Warn> MSM-A: Insufficient Power to power-on Slot-7
```

```
11/12/2004 00:38:08.77 <Warn:dm.Warn> MSM-A: Slot-7 being Powered OFF due to insufficient power
```

```
11/12/2004 00:36:23.77 <Warn:dm.Warn> MSM-A: Slot-7 being Powered OFF due to insufficient power
```

```
...
```

A total of 83 log messages were displayed.

The following command displays messages containing the string “slot 2”:

```
show log match "slot 2"
```

show log components

```
show log components {<event component>} {version}
```

Description

Displays the name, description and default severity for all components.

Syntax Description

event component	Specifies the component to display.
version	Specifies the version number of the component.

Default

N/A.

Usage Guidelines

This command displays the name, description, and default severity defined for the specified components or subcomponents.

Depending on the software version running on your switch or your switch model, additional or different component information might be displayed.

Example

The following command displays the log components:

```
show log components
```

The following is sample output from this command:

Severity			
Component	Title		Threshold

AAA	Authentication, Authorization, Accounting		Info
	RADIUS	Remote Authentication Dial In User Service	Error
	TACACS	Terminal Access Controller Access Control Syst	Info
ACL	ACL		Info
	CLEARFlow	CLEARFlow	Info
	Policy	Policy actions	Info
bgp	Border Gateway Protocol		Info
	damp	BGP Route Flap Dampening related debug message	Error
	event	BGP FSM related events	Error
	inUpdt	Incoming Update related debug msgs	Warning
	keepalive	BGP keepalive message	Warning
	misc	Miscellaneous debug (Import, Aggregate, NextHop	Warning
	msgs	Debug for BGP messages (OPEN, Update, Notifica	Warning
	outUpdt	Transmit Update related debug	Warning
bootp	BOOTP, DHCP Component		Error
	relay	BOOTP Relay trace component	Error
	server	DHCP Server subcomponent	Info
cli	Command Line Interface		Info
	shell	CLI configuration shell.	Error
	subagent	CLI application subagent	Error
cm	Configuration Manager		Warning
	file	CM file operation events	Warning
	sys	CM system events	Warning
DM	Device Manager		Info
	Card	Device Manager Card State Machine	Info
dosprot	dosprot		Info
ds	Directory Services		Error
fdb	fdb module event		Error
HAL	Hardware Abstraction Layer		Error
	Card	Card State Driver	Info

	FDB	Forwarding Database Driver	Info
	IPv4ACL	IPv4 Access Control List Driver	Info
	IPv4Adj	IPv4 Adjacency Driver	Info
	IPv4FIB	IPv4 FIB Driver	Info
	IPv4Mc	IPv4 Multicast Driver	Info
	Mirror	Mirroring Driver	Error
	Msg	Message Handler	Info
	Port	I/O Port Driver	Info
	SM	Switch Manager	Info
	Sys	System Driver	Info
	VLAN	VLAN Driver	Info
IPMC		IP Multicast Main Module	Info
	Snoop	IP Multicast Snooping Module	Error
	VLAN	IP Multicast VLAN Module	Error
Kern		Kernel messages	Error
LACP		Link Aggregation Control Protocol	Info
lldp		Link Layer Discovery Protocol (IEEE 802.1AB)	Warning
log		Log server messages	Warning
netTool		netTools framework	Error
	dnsclient	Dns Client	Error
	dnsproxy	Dns Proxy	Error
	routeradv	IPv6 Router Advertisements	Warning
	sntp	Sntp client	Warning
nl		Network Login	Info
	dot1x	802.1x-based Network Login	Warning
	mac	MAC-based Network Login	Warning
	web	Web-based Network Login	Warning
NM		Node Manager	Info
ospf		open shortest path first	Error
	event	ospf events	Info
	hello	ospf hello	Error
	lsa	ospf link-state advertisement	Error
	neighbor	ospf neighbor	Error
	spf	ospf shortest path first	Error
ospfv3		OSPFv3 related EMS messages	Warning
	events	OSPF6 events related messages	Error
	lsa	LSA related messages	Warning
	nbr	OSPF6 neighbor related EMS messages	Warning
	pkt	OSPF6 Packet receive/transmit/processing relat	Warning
	route	OSPF6 route add/delete related messages	Warning
	spf	SPF computation related messages	Error
pim		Pim Protocol Events	Warning
	cache	PIM cache maintenance.	Warning
	debug	PIM debug messages	Notice
	hello	Hello messages	Warning
	mcdbg	multicast forwarding engine	Warning
	msg	Trace for pim control packtes	Notice

	nbr	Neighbor creation/deletion etc	Warning
	rpm	RP message exchange.	Warning
pm		Policy Manager	Error
	config	Policy file events	Info
POE		Inline Power	Notice
rip		RIP routing	Error
	cfg	rip configuration	Warning
	event	rip events	Warning
	inUpdt	rip - inbound route updates	Warning
	msgs	rip - socket messages in and out	Warning
	outUpdt	rip - outbound route updates	Warning
	sys	rip - exos kernel interface	Warning
ripng		RIPng Protocol Events	Warning
	debug	RIPng debug messages	Notice
	external	RIPng external interface related messages	Warning
	message	RIPng control messages	Warning
	route	Hello messages	Warning
rmon		RMON general info	Error
	alarm	RMON alarm info	Error
	estat	RMON statistics info	Error
	event	RMON event info	Error
	history	RMON history	Error
RtMgr		Route Manager	Info
	VLAN	rtmgr vlan interface	Info
sflow		Sflow Protocol Events	Warning
	debug	SFLOW debug messages	Notice
	extended	SFLOW extended data collection	Notice
	msg	SFLOW process initializaion related message	Warning
	sample	SFLOW sample collection related messages	Warning
	statistics	SFLOW port statistics related message	Warning
STP		Spanning-Tree Protocol	Error
	InBPDU	STP In Bridge Protocol Data Unit	Warning
	OutBPDU	STP Out Bridge Protocol Data Unit	Warning
	System	STP System	Error
System		XOS system related log messages	Info
telnetd		telnet server	Info
tftpd		tftp server	Info
thttpd		thttp server	Info
trace		Debug trace messages	Warning
vlan		Vlan mgr	Info
	ack	vlan ack	Error
	dbg	Debug information	Info
	err	errors	Error
	mac	Virtual MAC Debugging	Info
	msgs	Messages	Info
VRRP		Config/State messages	Warning

Advert	Subsystem description	Warning
System	System/Library messages	Warning

A total of 143 component(s) were displayed.

The following command displays the version number of the VRRP component:

```
show log components vrrp version
```

The following is sample output from this command:

Component	Title	Version
-----	-----	-----
VRRP	Config/State messages	2.4
Advert	Subsystem description	3.1
System	System/Library messages	3.2

A total of 3 component(s) were displayed.

show log configuration

```
show log configuration
```

Description

Displays the log configuration for switch log settings, and for certain targets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the log configuration for all targets. The state of the target, enabled or disabled is displayed. For the enabled targets, the associated filter, severity, match expression, and format is displayed. The debug mode state of the switch is also displayed.

Example

The following command displays the configuration of all the log targets and all existing filters:

```
show log configuration
```

The following is sample output from this command:

```
Debug-Mode: Enabled
```

```
Log Target      : memory-buffer
Enabled ?      : yes
Filter Name     : DefaultFilter
```

```

Match regex : Any
Severity    : Debug-Data (through Critical)
Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
ion>
Buffer size : 1000 messages

Log Target   : nvram
Enabled ?    : yes
Filter Name  : DefaultFilter
Match regex  : Any
Severity     : Warning (through Critical)
Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
ion>

Log Target   : console
Enabled ?    : no
Filter Name  : DefaultFilter
Match regex  : Any
Severity     : Info (through Critical)
Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Filter Name: DefaultFilter
I/                                     Severity
E Comp.   Sub-comp.   Condition       CEWNISVD
- - - - - - - - - - - - - - - - - - -
I All                                           - - - - -

Log Filter Name: myFilter
I/                                     Severity
E Comp.   Sub-comp.   Condition       CEWNISVD
- - - - - - - - - - - - - - - - - - -
I STP                                           - - - - -

Include/Exclude: I - Include, E - Exclude
Component Unreg: * - Component/Subcomponent is not currently registered
Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I - Info
Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
                  + - Debug Severities, but log debug-mode not enabled

If Match parameters present:
Parameter Flags: S - Source, D - Destination, (as applicable)
                  I - Ingress, E - Egress, B - BGP
Parameter Types: Port - Physical Port list, Slot - Physical Slot #
                  MAC  - MAC address, IP - IP Address/netmask, Mask - Netmask
                  VID  - Virtual LAN ID (tag), VLAN - Virtual LAN name
                  L4   - Layer-4 Port #, Num - Number, Str - String

```

```

Nbr - Neighbor, Rtr - Routerid
Proc - Process Name
Strict Match : Y - every match parameter entered must be present in the event
               N - match parameters need not be present in the event

```

show log configuration filter

```
show log configuration filter {<filter name>}
```

Description

Displays the log configuration for the specified filter.

Syntax Description

filter name	Specifies the filter to display.
-------------	----------------------------------

Default

If no options are specified, the command displays the configuration for all filters.

Usage Guidelines

This command displays the configuration for filters.

Example

The following command displays the configuration for the filter, *myFilter*:

```
show log configuration filter myFilter
```

The following is sample output from this command:

Log Filter Name: myFilter

I/				Severity
E	Comp.	Sub-comp.	Condition	CEWNISVD
-	-----	-----	-----	-----
I	STP			-----
I	aaa			-----

Include/Exclude: I - Include, E - Exclude

Component Unreg: * - Component/Subcomponent is not currently registered

Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I - Info
* - Pre-assigned severities in effect for specified component

Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
+ - Debug Severities, but log debug-mode not enabled

If Match parameters present:

Parameter Flags: S - Source, D - Destination, (as applicable)
I - Ingress, E - Egress, B - BGP

Parameter Types: Port - Physical Port list, Slot - Physical Slot #

MAC - MAC address, IP - IP Address/netmask, Mask - Netmask
 VID - Virtual LAN ID (tag), VLAN - Virtual LAN name
 L4 - Layer-4 Port #, Num - Number, Str - String
 Nbr - Neighbor, Rtr - Routerid
 Proc - Process Name
 Strict Match : Y - every match parameter entered must be present in the event
 N - match parameters need not be present in the event

show log configuration target

```
show log configuration target {console | memory-buffer | nvram | primary-msm | primary-node |
backup-msm | backup-node | session | syslog {<ipaddress> | <ipPort> | vr <vr_name>} {[local0
... local7]}}
```

Description

Displays the log configuration for the specified target.

Syntax Description

console	Show the log configuration for the console display.
memory-buffer	Show the log configuration for volatile memory.
nvram	Show the log configuration for NVRAM.
primary-msm	Specifies the primary MSM.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.
backup-node	Specifies the backup-node in a stack.
session	Show the log configuration for the current session (including console display).
syslog	Show the configuration for the specified syslog target.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .	
local0 ... local7	Specifies the local syslog facility.

Default

If no options are specified, the command displays the configuration for the current session and console display.

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command displays the log configuration for the specified target. The associated filter, severity, match expression, and format is displayed.

Example

The following command displays the log configuration:

```
show log configuration target
```

The following is sample output from this command:

```
Log Target      : memory-buffer
  Enabled ?    : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Debug-Data (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
  Buffer size   : 1000 messages
```

```
Log Target      : nvram
  Enabled ?    : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Warning (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
```

```
Log Target      : console
  Enabled ?    : no
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Info (through Critical)
  Format       : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
```

```
Log Target      : primary-msm
  Enabled      : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Warning (through Critical)
```

```
Log Target      : backup-msm
  Enabled      : yes
  Filter Name  : DefaultFilter
  Match regex  : Any
  Severity     : Warning (through Critical)
```

show log counters

```
show log counters {<event condition> | [all | <event component>]} {include | notified |
occurred} {severity <severity> {only}}
```

Description

Displays the incident counters for events.

Syntax Description

event condition	Specifies the event condition to display.
all	Specifies that all events are to be displayed.
event component	Specifies that all the events associated with a particular component or subcomponent should be displayed.
include	Specifies if one or more targets should be included in this event.
notified	Specifies the number of times this event has occurred.
occurred	Specifies the number of times this event has occurred since the last clear or reboot.
severity	Specifies the minimum severity level of events to display (if the keyword only is omitted).
only	Specifies that only events of the specified severity level are to be displayed

Default

If `severity` is not specified, then events of all severity are displayed.

Usage Guidelines

This command displays the incident counters for each event specified. Two incident counters are displayed. One counter displays the number of times an event has occurred, and the other displays the number of times that notification for the event was made to the system (an incident record was injected into the system for further processing). Both incident counters reflect totals accumulated since reboot or since the counters were cleared using the `clear log counters` or `clear counters` command, regardless of whether it was filtered or not.

The keywords `include`, `notified`, and `occurred` only display events with non-zero counter values for the corresponding counter.

This command also displays a reference count (the column titled `rf` in the output). The reference count is the number of enabled targets receiving notifications of this event.

See the command [show log](#) on page 387 for more information about severity levels.

To get a listing of the event conditions in the system, use the following command:

```
show log events
```


To get a listing of the components present in the system, use the following command:

```
show log components
```

Example

The following command displays the event counters for event conditions of severity debug-summary or greater in the component *STP.InBPDU*:

```
show log counters stp.inbpdu severity debug-summary
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Occurred	In	Notified
STP	InBPDU	Drop	Error	0	Y	0
STP	InBPDU	Ign	Debug-Summary	0	N	0
STP	InBPDU	Mismatch	Warning	0	Y	0

Occurred : # of times this event has occurred since last clear or reboot

Flags : (*) Not all applications responded in time with there count values

In(cluded): Set to Y(es) if one or more targets filter includes this event

Notified : # of times this event has occurred when 'Included' was Y(es)

The following command displays the event counters for the event condition *PDUDrop* in the component *STP.InBPDU*:

```
show log counters "STP.InBPDU.Drop"
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Occurred	In	Notified
STP	InBPDU	Drop	Error	0	Y	0

Occurred : # of times this event has occurred since last clear or reboot

Flags : (*) Not all applications responded in time with there count values

In(cluded): Set to Y(es) if one or more targets filter includes this event

Notified : # of times this event has occurred when 'Included' was Y(es)

show log events

```
show log events [<event condition> | [all | <event component>] {severity <severity> {only}}]
{details}
```

Description

Displays information about the individual events (conditions) that can be logged.

Syntax Description

event condition	Specifies the event condition to display.
-----------------	---

all	Specifies that all events are to be displayed.
event component	Specifies that all the events associated with a particular component should be displayed.
severity	Specifies the minimum severity level of events to display (if the keyword only is omitted).
only	Specifies that only events of the specified severity level are to be displayed.
details	Specifies that detailed information, including the message format and parameter types, be displayed.

Default

If severity is not specified, then events of all severity are displayed. If detail is not specified, then summary only information is displayed.

Usage Guidelines

This command displays the mnemonic, message format, severity, and parameter types defined for each condition in the event set specified.

See the command [show log on page 387](#) for more information about severity levels.

When the `detail` option is specified, the message format is displayed for the event conditions specified. The message format parameters are replaced by the value of the parameters when the message is generated.

To get a listing of the components present in the system, use the following command:

```
show log components
```

Example

The following command displays the event conditions of severity debug-summary or greater in the component *STP.InBPDU*:

```
show log events stp.inbpdu severity debug-summary
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Parameters
STP	InBPDU	Drop	Error	2 total
STP	InBPDU	Ign	Debug-Summary	2 total
STP	InBPDU	Mismatch	Warning	2 total

The following command displays the details of the event condition *PDUTrace* in the component *STP.InBPDU*:

```
show log events stp.inbpdu.pdutrace details
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Parameters
------	---------	-----------	----------	------------

```

-----
STP      InBPDU      Trace      Debug-Verbose  2 total
                                0 - string
                                1 - string (printf)

Port=%0%: %1%

```

show ports rxerrors

```
show ports {<port_list>} rxerrors {no-refresh}
```

Description

Displays real-time receive error statistics. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the receive errors at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, receive error statistics are displayed for all ports.

If you do not specify the `no-refresh` parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the `no-refresh` parameter, the output provides a snapshot of the real-time receive error statistics at the time you issue the command and displays the output in page-by-page mode. This setting is not saved; therefore you must specify the `no-refresh` parameter each time you want a snapshot of the port receive errors.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Receive Error Information

The switch collects the following port receive error information:

- Port Number
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.

- Ready (R)—The port is ready to accept a link.
- Not Present (NP)—The port is configured, but the module is not installed in the slot.
- Loopback (L)—The port is in Loopback mode.
- Receive Bad CRC Frames (RX CRC)—The total number of frames received by the port that were of the correct length, but contained a bad FCS value.
- Receive Oversize Frames (RX Over)—The total number of good frames received by the port greater than the supported maximum length of 1,522 bytes.
- Receive Undersize Frames (RX Under)—The total number of frames received by the port that were less than 64 bytes long.
- Receive Fragmented Frames (RX Frag)—The total number of frames received by the port were of incorrect length and contained a bad FCS value.
- Receive Jabber Frames (RX Jabber)—The total number of frames received by the port that was of greater than the support maximum length and had a Cyclic Redundancy Check (CRC) error.
- Receive Alignment Errors (RX Align)—The total number of frames received by the port that occurs if a frame has a CRC error and does not contain an integral number of octets.
- Receive Frames Lost (RX Lost)—The total number of frames received by the port that were lost because of buffer overflow in the switch.

Port Monitoring Display Keys

For information about the available port monitoring display keys, see the `show ports statistics` command.

Example

The following command displays receive error statistics for slot 5, ports 4 through 7, on the switch with auto-refresh disabled:

```
show ports 5:4-5:7 rxerrors no-refresh
```

The following is sample output from this command:

```
Port Rx Error monitor
  Port      Link   Rx    Rx    Rx    Rx    Rx    Rx    Rx
          State  Crc   Over Under Frag Jabber Align Lost
=====
   5:4      R      0     0     0     0     0     0     0
   5:5      R      0     0     0     0     0     0     0
   5:6      R      0     0     0     0     0     0     0
   5:7      R      0     0     0     0     0     0     0
=====
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
```

show ports statistics

```
show ports {<port_list>} statistics {no-refresh}
```

Description

Displays real-time port statistic information. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

stacking-port-list	Specifies one or more stacking slots and ports.
port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the port statistics at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, statistics are displayed for all ports.

If you do not specify the `no-refresh` parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the `no-refresh` parameter, the output provides a snapshot of the real-time port statistics at the time you issue the command and displays the output in page-by-page mode. This setting is not saved; therefore you must specify the `no-refresh` parameter each time you want a snapshot of the port statistics.

Jumbo frame statistics are displayed for switches only that are configured for jumbo frame support.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Statistics

The switch collects the following port statistic information:

- Port Number
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Not Present (NP)—The port is configured, but the module is not installed in the slot.
 - Loopback (L)—The port is in Loopback mode.
- Transmitted Packet Count (Tx Pkt Count)—The number of packets that have been successfully transmitted by the port.

- Transmitted Byte Count (Tx Byte Count)—The total number of data bytes successfully transmitted by the port.
- Received Packet Count (RX Pkt Count)—The total number of good packets that have been received by the port.
- Received Byte Count (RX Byte Count)—The total number of bytes that were received by the port, including bad or lost frames. This number includes bytes contained in the Frame Check Sequence (FCS), but excludes bytes in the preamble.
- Received Broadcast (RX Bcast)—The total number of frames received by the port that are addressed to a broadcast address.
- Received Multicast (RX Mcast)—The total number of frames received by the port that are addressed to a multicast address.

Port Monitoring Display Keys

Table 12 describes the keys used to control the display that appears if auto-refresh is enabled (the default behavior).

Table 12. Port Monitoring Display Keys with Auto-Refresh Enabled

Key(s)	Description
U	Displays the previous page of ports.
D	Displays the next page of ports.
[Esc]	Exits from the screen.
0	Clears all counters.

Table 13 describes the keys used to control the display that appears if you auto-refresh is disabled.

Table 13. Port Monitoring Displays Keys with Auto-Refresh Disabled

Key	Description
Q	Exits from the screen.
[Space]	Displays the next page of ports.

Example

The following command displays port statistics for slot 1, ports 1 through 2, on the switch with auto-refresh disabled:

```
show ports 1:1-1:2 statistics no-refresh
```

The following is sample output from this command:

```
Port Statistics
Port      Link      Tx Pkt    Tx Byte    Rx Pkt    Rx Byte    Rx      Rx
```

```

          State      Count      Count      Count      Count      Bcast      Mcast
=====
1:1        A        7241    2722608    14482    3968068        0        0
1:2        R          0         0         0         0         0         0
=====
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback

```

show ports txerrors

```
show ports {<port_list> | stack-ports <stacking-port-list>} txerrors {no-refresh}
```

Description

Displays real-time transmit error statistics. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the transmit errors at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, error statistics are displayed for all ports.

If you do not specify the `no-refresh` parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the `no-refresh` parameter, the output provides a snapshot of the real-time transmit error statistics at the time you issue the command and displays the output in page-by-page mode. This setting is not saved; therefore, you must specify the `no-refresh` parameter each time you want a snapshot of the port transmit errors.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Transmit Error Information

The switch collects the following port transmit error information:

- Port Number
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.

- Ready (R)—The port is ready to accept a link.
- Not Present (NP)—The port is configured, but the module is not installed in the slot.
- Loopback (L)—The port is in Loopback mode.
- Transmit Collisions (TX Coll)—The total number of collisions seen by the port, regardless of whether a device connected to the port participated in any of the collisions.
- Transmit Late Collisions (TX Late Coll)—The total number of collisions that have occurred after the port's transmit window has expired.
- Transmit Deferred Frames (TX Deferred)—The total number of frames that were transmitted by the port after the first transmission attempt was deferred by other network traffic.
- Transmit Errored Frames (TX Errors)—The total number of frames that were not completely transmitted by the port because of network errors (such as late collisions or excessive collisions).
- Transmit Lost Frames (TX Lost)—The total number of transmit frames that do not get completely transmitted because of buffer problems (FIFO underflow).
- Transmit Parity Frames (TX Parity)—The bit summation has a parity mismatch.

Port Monitoring Display Keys

For information about the available port monitoring display keys, see the `show ports statistics` command.

Example

The following command displays transmit error statistics for slot 5, ports 4 through 7, on the switch with auto-refresh disabled:

```
show ports 5:4-5:7 txerrors no-refresh
```

The following is sample output from this command:

```
Port Transmission errors
Port          Link   Tx    Tx      Tx      Tx      Tx      Tx
              State Coll  Late coll  Deferred  Errors  Lost  Parity
=====
5:4           R      0      0        0        0        0      0
5:5           R      0      0        0        0        0      0
5:6           R      0      0        0        0        0      0
5:7           R      0      0        0        0        0      0
=====
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
```

show rmon memory

```
show rmon memory {detail | <memoryType>}
```


Description

Displays RMON specific memory usage and statistics.

Syntax Description

detail	Displays detailed information.
memoryType	Specifies the type of memory usage and statistics to display.

Default

N/A.

Usage Guidelines

If you do not specify the `detailed` keyword or a enter a specific RMON memory type, the output contains usage information for all memory types.

Example

The following command displays RMON memory statistics:

```
show rmon memory
```

The following is sample output from this command:

```

RMON Memory Information
-----

Bytes Allocated: 14298032 AllocFailed: 0
Current Memory Utilization Level: GREEN

Memory Utilization Statistics
-----

```

	Size	16	32	48	64	80	96	112	128	144	176	208
256	384	5										
12	768	1024	2048	4096	8192	16384	18432	40960	64000			

```

-----
-----
-----
Used Blocks 1558 3 2490 1 0 0 0 1 1 0 63444
1 1869
0 311 0 0 0 0 0 0 0 0
rmonEstat 0 0 0 0 0 0 0 0 0 0
0 311
0 0 0 0 0 0 0 0 0
rmonOwner 1555 0 0 0 0 0 0 0 0 0
0 0
0 0 0 0 0 0 0 0 0

```

```

    rmonHisc      0      0      0      0      0      0      0      0      0      0      0      0
0  1244
    0      0      0      0      0      0      0      0      0      0      0
    rmonHist      0      0      0      0      0      0      0      0      0      0      0      0  63444
0      0
    0      0      0      0      0      0      0      0      0      0      0
    rmonAlarm      0      0      0      0      0      0      0      0      0      0      0      0
0      3
    0      0      0      0      0      0      0      0      0      0      0
rmonLogDescription      0      0      0      0      0      0      0      0      0      1      0      0
0      0
    0      0      0      0      0      0      0      0      0      0      0
    rmonLog      0      1      0      0      0      0      0      0      0      0      0      0
0      0
    0      0      0      0      0      0      0      0      0      0      0
    rmonEvent      0      0      0      0      0      0      0      0      0      1      0      0      0
0      0
    0      0      0      0      0      0      0      0      0      0      0
rmonEventDescription      0      1      0      0      0      0      0      0      0      0      0      0
0      0
    0      0      0      0      0      0      0      0      0      0      0
rmonEventCommunity      0      1      0      0      0      0      0      0      0      0      0      0
0      0
    0      0      0      0      0      0      0      0      0      0      0
    rmonCommunity      1      0      0      0      0      0      0      0      0      0      0      0
0      0
    0      0      0      0      0      0      0      0      0      0      0
    rmonDs      0      0      0      0      0      0      0      0      0      0      0      0
0      0
    0  311      0      0      0      0      0      0      0      0      0
    rmonDbx      0      0  2490      0      0      0      0      0      0      0      0      0
0      0
    0      0      0      0      0      0      0      0      0      0
    rmonOid      0      0      0      0      0      0      0      0      0      0      0      0
0  311
    0      0      0      0      0      0      0      0      0      0
rmonMdbIndexOid      2      0      0      1      0      0      0      0      0      0      0      0
0      0
    0      0      0      0      0      0      0      0      0      0
    rmonMdbString      0      0      0      0      0      0      0      0      0      0      0      0
1      0
    0      0      0      0      0      0      0      0      0      0

```

The following command displays RMON event statistics:

```
show rmon memory rmonEvent
```

The following is sample output from this command:

```

RMON Memory Information
-----

```

Bytes Allocated: 14298032 AllocFailed: 0
Current Memory Utilization Level: GREEN

Memory Utilization Statistics

Memory Statistics for rmonEvent

Size		16	32	48	64	80	96	112	128	144	176	208
256	384	512	768	1024	2048	4096	8192	16384	18432	40960	64000	
Allocated		0	0	0	0	0	0	0	1	0	0	0
0	0											
AllocatedPeak		0	0	0	0	0	0	0	0	1	0	0
0	0											
AllocSuccess		0	0	0	0	0	0	0	1	0	0	0
0	0											
FreeSuccess		0	0	0	0	0	0	0	0	0	0	0
0	0											
AllocFail		0	0	0	0	0	0	0	0	0	0	0
0	0											
FreeFail		0	0	0	0	0	0	0	0	0	0	0
0	0											
		0	0	0	0	0	0	0	0			

show sflow configuration

show sflow {configuration}

Description

Displays the current sFlow configuration.

Syntax Description

This command has no arguments or variables

Default

N/A.

Usage Guidelines

This command displays the sFlow configuration of your system.

The following fields are displayed:

- Global Status—sFlow is globally enabled or disabled
- Polling interval—How often the hardware is polled for statistics, in seconds
- Sampling rate—Packets are sampled, on average, once for every rate-number of packets
- Maximum cpu sample limit—Maximum number of packets per second sampled before sample throttling takes effect
- Agent IP—IP address inserted into the sFlow data packets to identify the sFlow switch
- Collectors—To which IP address and port, and from which virtual router, the sFlow packets are sent
- Port Status—Enabled or disabled for statistics gathering
- Port Sample-rate—Shows the sampling rate configured for the port and the actual rate if CPU throttling has taken effect
- Port Subsampling factor—See the command `configure sflow ports sample-rate` for details

Example

To display the sFlow configuration on your system, use the following command:

```
show sflow
```

The output from this command is similar to the following:

```
SFLOW Global Configuration
Global Status: enabled
Polling interval: 20
Sampling rate: 8192
Maximum cpu sample limit: 2000
SFLOW Configured Agent IP: 10.203.2.38 Operational Agent IP: 10.203.2.38
Collectors
Collector IP 10.201.6.250, Port 6343, VR "VR-Mgmt"
SFLOW Port Configuration
Port      Status      Sample-rate      Subsampling
          Config / Actual  factor
1:41      enabled      8192 / 8192      1
2:40      enabled      1024 / 1024      1
2:58      enabled      8192 / 8192      8
2:59      enabled      8192 / 8192      8
```

show sflow statistics

```
show sflow statistics
```

Description

Displays sFlow statistics.

Syntax Description

This command has no arguments or variables

Default

N/A.

Usage Guidelines

This command displays sFlow statistics for your system.

The following fields are displayed:

- Received frames—Number of frames received on sFlow enabled ports
- Sampled Frames—Number of packets that have been sampled by sFlow
- Transmitted Frames—Number of UDP packets sent to remote collector(s)
- Broadcast Frames—Number of broadcast frames received on sFlow enabled ports
- Multicast Frames—Number of multicast frames received on sFlow enabled ports
- Packet Drops—Number of samples dropped

Example

To display sFlow statistics for your system, use the following command:

```
show sflow statistics
```

The output from this command is similar to the following:

```
SFLOW Statistics
```

```
Received frames      : 1159044921
Sampled Frames       : 104944
Transmitted Frames   : 10518
Broadcast Frames     : 0
Multicast Frames     : 1055652
Packet Drops         : 0
```

show temperature

```
show temperature
```

Description

Depending on the platform, this command displays the current temperature of the I/O modules, management modules, power supply controllers, XGM-2xn card, and the switch.

On a stack, the command displays the current temperature of the modules in each slot.

Syntax Description

This command has no arguments or variables

Default

N/A.

Usage Guidelines

Depending on the software version running on your switch or your switch model, additional or different temperature information might be displayed.

Use this command to display the temperature in Celsius and the current status of the following installed components in the switch:

- Management modules (MSM/MM)
- I/O modules
- Power controllers

The switch monitors the temperature of each component and generates a warning if the temperature exceeds the normal operating range. If the temperature exceeds the minimum/maximum limits, the switch shuts down the overheated module.

Displaying the Temperature of Other Installed Components

You can also view the temperature of the power supplies and the fan trays in the switch.

To view the temperature of the power supplies installed in the switch, use the following command:

```
show power {<ps_num>} {detail}
```

Example

Depending on the platform, the following command displays the temperature of various switch components:

```
show temperature
```

The following is sample output from a NETGEAR 8806 switch:

```
XCM8806.8 # show temperature
```

Field Replaceable Units	Temp (C)	Status	Min	Normal	Max
Slot-1 : XCM8824F	30.00	Normal	-10	0-50	60
Slot-2 :					
Slot-3 : XCM888F	32.50	Normal	-10	0-50	60
Slot-4 :					
Slot-5 : XCM8808X	37.00	Normal	-10	0-50	60
Slot-6 : XCM8848T(P)	34.50	Normal	-10	0-50	60
MSM-A : XCM88S1	37.50	Normal	-10	0-50	60
MSM-B :					
PSUCTRL-1 :	38.38	Normal	-10	0-50	60
PSUCTRL-2 :	42.40	Normal	-10	0-50	60

```
(Demo)*XCM8806.9 #
```

show version

```
show version {detail | process <name> | images {partition <partition>} {slot <slotid>}}
```

Description

Displays the hardware serial and version numbers, the software version currently running on the switch, and (if applicable) the software version running on the modules and power controllers.

Syntax Description

detail	Specifies display of slot board name and chassis or platform name.
process	Specifies display of all of the processes on the switch.
name	Specifies display of a specific process on the switch.
images	Specifies the display of installed images.
partition	Specifies display of a specific partition (primary or secondary).
slotid	Specifies display of an MSM/MM in a specific slot (A or B).

Default

N/A.

Usage Guidelines

The following describes the information displayed when you execute the `show version` or `show version detail` commands:

- **Part Number**—A collection of numbers and letters that make up the part number of the switch and when applicable the hardware components installed in the switch.
- **Serial Number**—A collection of numbers and letters that make up the serial number of the switch and when applicable the hardware components installed in the switch.

Note: For information about the physical location of the serial number on your switch, refer to the section that describes your specific switch model in the hardware documentation.

- **Image**—The NETGEAR 8800 software version currently running on the switch. If you have two software images downloaded on the switch, only the currently running NETGEAR 8800 version information is displayed. The information displayed includes the major version number, minor version number, a specific patch release, and the build number. The software build date is also displayed.
- **BootROM**—The BootROM version currently running on the switch.
- **Diagnostics**—A number that corresponds to the version of the I/O module diagnostics included in the particular version of NETGEAR 8800 OS.

Depending on the model of your switch and the software running on your switch, different version information may be displayed.

Note: The information displayed does not include the I/O version number on the NETGEAR 8800 series switch. The I/O version number includes the major, minor, and I/O version number, not the patch and build numbers.

If you use the `process` option, you will see the following information about the processes running on the switch:

- Card—The location (MSM/MM) where the process is running on the switch.
- Process Name—The name of the process.
- Version—The version number of the process.
- BuiltBy—The name of the software build manager.
- Link Date—The date the executable was linked.

Example

The following command displays the hardware and software versions currently running on the switch:

```
show version
```

The following is sample output from a NETGEAR 8806 switch (the output from the NETGEAR 8810 is similar):

```
(Demo)*XCM8806.9#show version
```

```
Chassis ESN Number : 1102G-00001
```

```
Chassis      : 800418-00   1102G-00001  Rev 0.0
Slot-1       : 800423-00   000000000000 Rev 0.0 BootROM: 1.0.4.0   IMG: 12.4.4.0
Slot-2       :
Slot-3       : 800426-00   000000000000 Rev 0.0 BootROM: 1.0.4.0   IMG: 12.4.4.0
Slot-4       :
Slot-5       : 800229-00-05 1027G-00178 Rev 5.0 BootROM: 1.0.4.0   IMG: 12.4.4.0
Slot-6       : 800421-00   000000000000 Rev 0.0 BootROM: 1.0.4.0   IMG: 12.4.4.0
MSM-A        : 800420-00   000000000000 Rev 0.0 BootROM: 1.0.4.4   IMG: 12.4.4.0
MSM-B        :
PSUCTRL-1    : 450352-00   1107G-0002 Rev 0.0 BootROM: 2.18
PSUCTRL-2    : 450352-00   1107G-0002 Rev 0.0 BootROM: 2.18
PSU-1        : PS 2336 4300-00145 1049J-00188 Rev 11.0
PSU-2        : PS 2336 4300-00145 1049J-00177 Rev 11.0
PSU-3        : PS 2336 4300-00145 1049J-00176 Rev 11.0
PSU-4        :
```



```
PSU-5      :
PSU-6      :
```

```
Image      : NETGEAR version 12.4.4.0 v1244b0-br-SR3-1 by release-manager
            on Tue Feb 8 07:22:38 PST 2011
BootROM    : 1.0.4.4
Diagnostics : 1.13
```

Using the `process` option of the `show version` command produces output similar to the following:

Card	Process Name	Version	BuiltBy	Link Date
MSM-A	aaa	3.0.0.2	release-manager	Thu Mar 31 09:23:54 PST 2005
MSM-A	acl	3.0.0.2	release-manager	Thu Mar 31 09:26:46 PST 2005
MSM-A	bgp	3.0.0.2	release-manager	Thu Mar 31 09:27:54 PST 2005
MSM-A	cfgmgr	3.0.0.21	release-manager	Thu Mar 31 09:23:42 PST 2005
MSM-A	cli	3.0.0.22	release-manager	Thu Mar 31 09:23:34 PST 2005
MSM-A	devmgr	3.0.0.2	release-manager	Thu Mar 31 09:23:22 PST 2005
MSM-A	dirser	3.0.0.2	release-manager	Thu Mar 31 09:24:02 PST 2005
MSM-A	ems	3.0.0.2	release-manager	Thu Mar 31 09:35:08 PST 2005
MSM-A	epm	3.0.0.3	release-manager	Thu Mar 31 09:23:11 PST 2005
....				

If you specify the `name` option, only the process you select is displayed.

Using the `images` option in the `show version` command produces output similar to the following:

Card	Partition	Installation Date	Version	Name
MSM-A	primary	Wed Jun 30 22:30:22 UTC 2004	11.0.0.24	NG8800-12.4.3.5-1-4.xos
MSM-A	primary	Thu Jul 1 03:29:41 UTC 2004	11.0.0.24	NG8800-12.4.3.5-1-4-ssh.xmod
MSM-A	secondary	Tue Jun 29 06:09:26 UTC 2004	11.0.0.23	NG8800-12.4.3.5-1-4.xos
MSM-A	secondary	Tue Jun 29 06:29:14 UTC 2004	11.0.0.23	NG8800-12.4.3.5-1-4-ssh.xmod

If you specify the `partition` option, only images on the specified partition is shown.

unconfigure log filter

```
unconfigure log filter <filter name>
```

Description

Resets the log filter to its default values; removes all filter items.

Syntax Description

filter name	Specifies the log filter to unconfigure.
-------------	--

Default

N/A.

Usage Guidelines

If the filter name specified is *DefaultFilter*, this command restores the configuration of *DefaultFilter* back to its original settings.

If the filter name specified is not *DefaultFilter*, this command sets the filter to have no events configured and therefore, no incidents will pass. This is the configuration of a newly created filter that was not copied from an existing one.

See the `delete log filter` command for information about deleting a filter.

Example

The following command sets the log filter myFilter to stop passing any events:

```
unconfigure log filter myFilter
```

unconfigure log target format

```
unconfigure log target [console | memory-buffer | nvram | session | syslog [all | <ipaddress>
| <ipPort> {vr <vr_name>} [local0 ... local7]]] format
```

Description

Resets the log target format to its default values.

Syntax Description

console	Specifies the console display format.
memory-buffer	Specifies the switch memory buffer format.
nvram	Specifies the switch NVRAM format.
session	Specifies the current session (including console display) format.
syslog	Specifies a syslog target format.
all	Specifies all remote syslog servers.
ipaddress	Specifies the syslog IP address.
ipPort	Specifies the UDP port number for the syslog target.
vr_name	Specifies the virtual router that can reach the server IP address.
<p>Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i>.</p>	
local0 ... local7	Specifies the local syslog facility.
format	Specifies that the format for the target will be reset to the default value.

Default

When a target format is unconfigured, it is reset to the default values.

The following defaults apply to console display, memory buffer, NVRAM, and session targets:

- timestamp—hundredths
- date—mm-dd-yyyy
- severity—on
- event-name—condition
- host-name—off
- sequence-number—off
- process-name—off
- process-slot—on
- process-id—off
- source-line—off

The following defaults apply to syslog targets (per RFC 3164):

- timestamp—seconds
- date—mmm-dd
- severity—on
- event-name—none
- host-name—off
- sequence-number—off
- process-name—off
- process-slot—on
- process-id—off
- source-line—off

Usage Guidelines

Use this command to reset the target format to the default format.

Example

The following command sets the log format for the target `session` (the current session) to the default:

```
unconfigure log target session format
```

unconfigure sflow

```
unconfigure sflow
```

Description

Resets all the sFlow values to the default values.

Syntax Description

This command has no arguments or variables

Default

The default values for sFlow are as follows:

- sFlow agent IP address—0.0.0.0
- sampling frequency—sample one every 8196 packets
- polling interval—20 seconds
- maximum CPU sample limit—2000 samples per second

sFlow is unconfigured and disabled on all ports.

Usage Guidelines

This command resets sFlow values to the default values, and removes any port configurations, and any sFlow collectors configured on the switch.

Example

The following command unconfigures sFlow:

```
unconfigure sflow
```

unconfigure sflow agent

```
unconfigure sflow agent
```

Description

Resets the sFlow agent's IP address to the default value.

Syntax Description

This command has no arguments or variables.

Default

The default IP address is 0.0.0.0.

Usage Guidelines

This command resets the sFlow agent IP address to its default value.

Example

The following command resets the agent IP back to the management IP address:

```
unconfigure sflow agent
```

unconfigure sflow collector

```
unconfigure sflow collector {ipaddress} <ip-address> {port <udp-port-number>} {vr <vrname>}
```

Description

Unconfigures the sFlow collector.

Syntax Description

ip-address	Specifies the IP address of the collector to reset.
------------	---

udp-port-number	Specifies the UDP port.
-----------------	-------------------------

vrname	Specifies which virtual router.
--------	---------------------------------

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*.

Default

The following values are the defaults for this command:

- UDP port number—6343
- Virtual router—VR-Mgmt (previously called VR-0).

Usage Guidelines

This command allows you to reset the specified sFlow collector parameters to the default values.

The `unconfigure sflow collector` command will reset the collector parameters to the default.

Example

The following command removes the collector at IP address 192.168.57.1:

```
unconfigure sflow collector ipaddress 192.168.57.1
```

unconfigure sflow ports

```
unconfigure sflow ports <port_list>
```

Description

Removes the specified ports from the sFlow configuration, and stops sampling them.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
-----------	---

Default

N/A.

Usage Guidelines

This command removes the specified ports from the sFlow configuration, and stops sampling them.

Example

The following command unconfigures sFlow on the ports 2:5-2:7:

```
unconfigure sflow ports 2:5-2:7
```

upload log

```
upload log <ipaddress> {vr <vr_name>} <filename> {messages [memory-buffer | nvram] {events
{<event-condition> | <event_component>}}} {severity <severity> {only}} {match <regex>}
{chronological}}
```

Description

Uploads the current log messages to a TFTP server.

Syntax Description

ipaddress	Specifies the ipaddress of the TFTP server.
vr_name	Specifies the virtual router that can reach the TFTP server. Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .
filename	Specifies the file name for the log stored on the TFTP server.
messages	Specifies the location from which to display the log messages.
memory-buffer	Show messages stored in volatile memory.
nvram	Show messages stored in NVRAM
events	Show event messages.
event-condition	Specifies the event condition to display.
event-compone nt	Specifies the event component to display.
severity	Specifies the minimum severity level to display (if the keyword <code>only</code> is omitted).

only	Specifies that only the specified severity level is to be displayed.
regex	Specifies a regular expression. Only messages that match the regular expression will be displayed.
chronological	Specifies uploading log messages in ascending chronological order (oldest to newest).

Default

The following defaults apply:

- messages—memory buffer
- severity—none (displays everything stored in the target)
- match—no restriction
- chronological—if not specified, show messages in order from newest to oldest

Usage Guidelines

This command is similar to the [show log](#) command, but instead of displaying the log contents on the command line, this command saves the log to a file on the TFTP server you specify. For more details on most of the options of this command, see the command [show log ON](#) [page 387](#).

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)

- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a local or remote file, remember the requirements listed above.

Example

The following command uploads messages with a critical severity to the filename *switch4critical.log* on TFTP server at 10.31.8.25:

```
upload log 10.31.8.25 switch4critical.log critical
```

The following command uploads messages with warning, error, or critical severity to the filename *switch4warn.log* on TFTP server at 10.31.8.25:

```
upload log 10.31.8.25 switch4warn.log warning
```


VLAN Commands

9

This chapter describes commands for configuring and managing:

- VLANs
- Private VLANs (PVLANS)
- VLAN translation

For an introduction to VLAN features, see the *NETGEAR 8800 User Manual*.

configure private-vlan add network

```
configure private-vlan <name> add network <vlan_name>
```

Description

Adds the specified VLAN as the network VLAN on the specified PVLAN.

Syntax Description

name	Specifies the name of the PVLAN to which the VLAN is added.
vlan_name	Specifies a VLAN to add to the PVLAN.

Default

N/A.

Usage Guidelines

The VLAN must be created and configured with a tag before it is added to the PVLAN.

Example

The following command adds VLAN *sharednet* as the network VLAN for the PVLAN named *companyx*:

```
configure private-vlan companyx add network sharednet
```

configure private-vlan add subscriber

```
configure private-vlan <name> add subscriber <vlan_name> {non-isolated} {loopback-port <port>}
```

Description

Adds the specified VLAN as a subscriber VLAN on the specified PVLAN.

Syntax Description

name	Specifies the name of the PVLAN to which the VLAN is added.
vlan_name	Specifies a VLAN to add to the PVLAN.
non-isolated	Configures the subscriber VLAN as a non-isolated subscriber VLAN.
port	Specifies the port that serves as the loopback port.

Default

If the `non-isolated` option is omitted, this command adds the specified VLAN as an isolated subscriber VLAN.

Usage Guidelines

The VLAN must be created and configured with a tag before it is added to the PVLAN. If the `non-isolated` option is omitted, the VLAN is added as an *isolated* subscriber VLAN. If the `non-isolated` option is included, the VLAN is added as an *non-isolated* subscriber VLAN.

If two or more subscriber VLANs have overlapping ports (where the same ports are assigned to both VLANs), each of the subscriber VLANs with overlapping ports must have a dedicated loopback port.

Example

The following command adds VLAN *restricted* as a subscriber VLAN for the PVLAN named *companyx*:

```
configure private-vlan companyx add subscriber restricted isolated
```

configure private-vlan delete

```
configure private-vlan <name> delete [network | subscriber] <vlan_name>
```

Description

Deletes the specified VLAN from the specified PVLAN.

Syntax Description

name	Specifies the name of the PVLAN from which the VLAN is deleted.
network	Specifies that the VLAN to be deleted is a network VLAN.
subscriber	Specifies that the VLAN to be deleted is a subscriber VLAN.
vlan_name	Specifies the VLAN to delete from the PVLAN.

Default

N/A.

Usage Guidelines

This command deletes a VLAN from a PVLAN, but it does not delete the VLAN from the system—it just breaks the link between the VLAN and the PVLAN. You can use this command to delete both network and subscriber VLANs.

Example

The following command deletes network VLAN *sharednet* from the PVLAN named *companyx*:

```
configure private-vlan companyx delete network sharednet
```

configure protocol add

```
configure protocol <name> add [etype | llc | snap] <hex> {[etype | llc | snap] <hex>}
```

Description

Configures a user-defined protocol filter.

Syntax Description

name	Specifies a protocol filter name.
hex	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: <ul style="list-style-type: none"> • The Ethernet protocol type taken from a list maintained by the IEEE. • The DSAP/SSAP combination created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP). • The SNAP-encoded Ethernet protocol type.

Default

N/A.

Usage Guidelines

Supported protocol types include:

- etype – IEEE Ethertype.
- llc – LLC Service Advertising Protocol.
- snap – Ethertype inside an IEEE SNAP packet encapsulation.

A maximum of 15 protocol filters, each containing a maximum of six protocols, can be defined.

The protocol filter must already exist before you can use this command. Use the [create protocol](#) command to create the protocol filter.

No more than seven protocols can be active and configured for use.

Example

The following command configures a protocol named Fred by adding protocol type LLC SAP with a value of FFEF:

```
configure protocol fred add llc 0xfeff
```

configure protocol delete

```
configure protocol <name> delete [etype | llc | snap] <hex> {[etype | llc | snap] <hex>} ...
```

Description

Deletes the specified protocol type from a protocol filter.

Syntax Description

name	Specifies a protocol filter name.
hex	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: <ul style="list-style-type: none"> • The Ethernet protocol type taken from a list maintained by the IEEE. • The DSAP/SSAP combination created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP). • The SNAP-encoded Ethernet protocol type.

Default

N/A.

Usage Guidelines

Supported protocol types include:

- etype – IEEE Ethertype.

- llc – LLC Service Advertising Protocol.
- snap – Ethernet inside an IEEE SNAP packet encapsulation.

Example

The following command deletes protocol type LLC SAP with a value of FEFF from protocol *fred*:

```
configure protocol fred delete llc feff
```

configure vlan add ports

```
configure {vlan} <vlan_name> add ports [<port_list> | all] {tagged | untagged} {{stpd}
<stpd_name>} {dot1d | emistp | pvst-plus}}
```

Description

Adds one or more ports in a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
port_list	Specifies a list of ports or slots and ports.
all	Specifies all ports.
tagged	Specifies the ports should be configured as tagged.
untagged	Specifies the ports should be configured as untagged.
stpd_name	Specifies an STP domain name.
dot1d emistp pvst-plus	Specifies the BPDU encapsulation mode for these STP ports.

Default

Untagged.

Usage Guidelines

The VLAN must already exist before you can add (or delete) ports: use the [create vlan](#) command to create the VLAN.

If the VLAN uses 802.1Q tagging, you can specify tagged or untagged port(s). If the VLAN is untagged, the ports cannot be tagged.

Untagged ports can only be a member of a single VLAN. By default, they are members of the default VLAN (named *Default*). In order to add untagged ports to a different VLAN, you must first remove them from the default VLAN. You do not need to do this to add them to another VLAN as tagged ports. If you attempt to add an untagged port to a VLAN prior to removing it from the default VLAN, you see the following error message:

Error: Protocol conflict when adding untagged port 1:2. Either add this port as tagged or assign another protocol to this VLAN.

The ports that you add to a VLAN and the VLAN itself cannot be explicitly assigned to different virtual routers. When multiple virtual routers are defined, consider the following guidelines while adding ports to a VLAN:

- A VLAN can belong (either through explicit or implicit assignment) to only one VR.
- If a VLAN is not explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to a single VR.
- If a VLAN is explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to the same VR or to no VR.
- If a port is added to VLANs that are explicitly assigned to different VRs, the port must be explicitly assigned to no VR.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*. On switches that do not support user-created VRs, all VLANs are created in *VR-Default* and cannot be moved.

For more information on configuring Spanning Tree Domains, see [Chapter 17, STP Commands](#).

Note: If you use the same name across categories (for example, STPD names), NETGEAR recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Example

The following command assigns tagged ports 1:1, 1:2, 1:3, and 1:6 to a VLAN named *accounting*:

```
configure vlan accounting add ports 1:1, 1:2, 1:3, 1:6 tagged
```

configure vlan add ports private-vlan translated

```
configure {vlan} <vlan_name> add ports <port_list> private-vlan translated
```

Description

Adds the specified ports to the specified network VLAN and enables tag translation for all subscriber VLAN tags to the network VLAN tag. Translation from network VLAN tag to each subscriber VLAN tag is done by default in a private VLAN.

Syntax Description

vlan_name	Specifies the network VLAN to which the ports are added.
port_list	Specifies the ports to be added to the network VLAN.

Default

N/A.

Usage Guidelines

This command is allowed only when the specified VLAN is configured as a network VLAN on a PVLAN.

Example

The following command adds port 2:1 to VLAN *sharednet* and enables VLAN translation on that port:

```
configure sharednet add ports 2:1 private-vlan translated
```

configure vlan add ports tagged private-vlan end-point

```
configure {vlan} <vlan_name> add ports <port_list> tagged private-vlan end-point
```

Description

Adds the specified ports as tagged end points on the specified network VLAN.

Syntax Description

vlan_name	Specifies the network VLAN to which the ports are added.
port_list	Specifies the ports to be added to the network VLAN.

Default

N/A.

Usage Guidelines

This command is allowed only when the specified VLAN is configured as a network VLAN on a PVLAN.

An end point port defines the PVLAN boundary. The end point port can connect to other devices, but cannot be used to extend the PVLAN to other switches.

Example

The following command adds port 2:1 as a tagged end point on VLAN *sharednet*.

```
configure sharednet add ports 2:1 tagged private-vlan end-point
```

configure vlan delete ports

```
configure {vlan} <vlan_name> delete ports [all | <port_list>]
```

Description

Deletes one or more ports in a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all ports.
port_list	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes ports 1:1, 1:2, 4:3, and 5:6 on the switch from a VLAN named *accounting*.

```
configure accounting delete port 1:1, 1:2, 4:3, 5:6
```

configure vlan ipaddress

```
configure {vlan} <vlan_name> ipaddress [<ipaddress> {<ipNetmask>} |  
ipv6-link-local | {eui64} <ipv6_address_mask>]
```

Description

Assigns an IPv4 address and an optional subnet mask or an IPv6 address to the VLAN. You can assign either an IPv4 address, and IPv6 address, or both to the VLAN. You can use this command to assign an IP address to a specified vMAN and enable multicasting on that vMAN.

Syntax Description

vlan_name	Specifies a VLAN name.
ipaddress	Specifies an IPv4 address.
ipNetmask	Specifies an IPv4 subnet mask in dotted-quad notation (for example, 255.255.255.0).
ipv6-link-local	Specifies IPv6 and configures a link-local address generated by combining the standard link-local prefix with the automatically generated interface in the EUI-64 format. Using this option automatically generates an entire IPv6 address; this address is only a link-local, or VLAN-based, IPv6 address, that is, ports on the same segment can communicate using this IP address and do not have to pass through a gateway.
eui64	Specifies IPv6 and automatically generates the interface ID in the EUI-64 format using the interface's MAC address. Once you enter this parameter, you must add the following variables: <ipv6_address_mask>. Use this option when you want to enter the 64-bit prefix and use a EUI-64 address for the rest of the IPv6 address.
ipv6_address_mask	Specify the IPv6 address in the following format: x:x:x:x:x:x/prefix length, where each x is the hexadecimal value of one of the 8 16-bit pieces of the 128-bit wide address.

Default

N/A.

Usage Guidelines

Note: You can also use this command to assign an IP address to a vMAN on any NETGEAR 8800 that supports the vMAN feature. For information on which software licenses and platforms support the vMAN feature, see Appendix A in the *NETGEAR 8800 User Manual*.

The VLAN must already exist before you can assign an IP address: use the `create vlan` command to create the VLAN (also the vMAN must already exist).

Note: See Chapter 19, “IP Unicast Commands,” for information on adding secondary IP addresses to VLANs.

You can specify IPv6 addresses. See Chapter 20, “IPv6 Unicast Commands,” for information on IPv6 addresses.

Example

The following commands are equivalent; both assign an IPv4 address of 10.12.123.1 to a VLAN named accounting:

```
configure vlan accounting ipaddress 10.12.123.1/24
configure vlan accounting ipaddress 10.12.123.1 255.255.255.0
```

The following command assigns a link local IPv6 address to a VLAN named management:

```
configure vlan accounting ipaddress ipv6-link-local
```

configure vlan name

```
configure {vlan} <vlan_name> name <name>
```

Description

Renames a previously configured VLAN.

Syntax Description

vlan_name	Specifies the current (old) VLAN name.
name	Specifies a new name for the VLAN.

Default

N/A.

Usage Guidelines

You cannot change the name of the default VLAN “Default.”

For information on VLAN name requirements and a list of reserved keywords, see the section on “Object Names” of the *NETGEAR 8800 User Manual*.

Note: If you use the same name across categories (for example, STPD names), NETGEAR recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Example

The following command renames VLAN *vlan1* to *engineering*:

```
configure vlan vlan1 name engineering
```

configure vlan protocol

```
configure {vlan} <vlan_name> protocol <protocol_name>
```

Description

Configures a VLAN to use a specific protocol filter.

Syntax Description

vlan_name	Specifies a VLAN name.
protocol_name	<p>Specifies a protocol filter name. This can be the name of a predefined protocol filter, or one you have defined.</p> <p>The following protocol filters are predefined:</p> <ul style="list-style-type: none"> • IP • IPv6 • IPX • NetBIOS • DECNet • IPX_8022 • IPX_SNAP • AppleTalk <p>any indicates that this VLAN should act as the default VLAN for its member ports.</p>

Default

Protocol any.

Usage Guidelines

If the keyword `any` is specified, all packets that cannot be classified into another protocol-based VLAN are assigned to this VLAN as the default for its member ports.

Use the `configure protocol` command to define your own protocol filter.

The NETGEAR 8800 does not forward packets with a protocol-based VLAN set to AppleTalk. To ensure that AppleTalk packets are forwarded on the device, create a protocol-based VLAN set to “any” and define other protocol-based VLANs for other traffic, such as IP traffic. The AppleTalk packets pass on the “any” VLAN, and the other protocols pass traffic on their specific protocol-based VLANs.

Example

The following command configures a VLAN named accounting as an IP protocol-based VLAN:

```
configure accounting protocol ip
```

configure vlan tag

```
configure {vlan} <vlan_name> tag <tag> {remote-mirroring}
```

Description

Assigns a unique 802.1Q tag to the VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
tag	Specifies a value to use as an 802.1Q tag. The valid range is from 2 to 4095.
remote-mirroring	Specifies that the tagged VLAN is for remote mirroring.

Default

The default VLAN uses an 802.1Q tag (and an internal VLANid) of 1.

Usage Guidelines

If any of the ports in the VLAN use an 802.1Q tag, a tag must be assigned to the VLAN. The valid range is from 2 to 4094 (tag 1 is assigned to the default VLAN, and tag 4095 is assigned to the management VLAN).

The 802.1Q tag is also used as the internal VLANid by the switch.

You can specify a value that is currently used as an internal VLANid on another VLAN; it becomes the VLANid for the VLAN you specify, and a new VLANid is automatically assigned to the other untagged VLAN.

Example

The following command assigns a tag (and internal VLANid) of 120 to a VLAN named *accounting*:

```
configure accounting tag 120
```

create private-vlan

```
create private-vlan <name> {vr <vr_name>}
```

Description

Creates a PVLAN framework with the specified name.

Syntax Description

name	Specifies a name for the new PVLAN.
vr_name	Specifies the virtual router in which the PVLAN is created.

Default

N/A.

Usage Guidelines

The PVLAN is a framework that links network and subscriber VLANs; it is not an actual VLAN.

A private VLAN name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For private VLAN naming guidelines and a list of reserved names, see the section on “Object Names” in the *NETGEAR 8800 User Manual*.

If no virtual router is specified, the PVLAN is created in the default VR context.

Example

The following command creates a PVLAN named *companyx*:

```
create private-vlan companyx
```

create protocol

```
create protocol <name>
```

Description

Creates a user-defined protocol filter.

Syntax Description

name	Specifies a protocol filter name. The protocol filter name can have a maximum of 31 characters.
------	---

Default

N/A.

Usage Guidelines

Protocol-based VLANs enable you to define packet filters that the switch can use as the matching criteria to determine if a particular packet belongs to a particular VLAN.

After you create the protocol, you must configure it using the `configure protocol` command. To assign it to a VLAN, use the `configure {vlan} <vlan_name> protocol <protocol_name>` command.

Example

The following command creates a protocol named *fred*:

```
create protocol fred
```

create vlan

```
create vlan <vlan_name> {vr <vr-name>}
```

Description

Creates a named VLAN.

Syntax Description

vlan_name	Specifies a VLAN name (up to 32 characters).
vr	Specifies a virtual router.
vr-name	Specifies in which virtual router to create the VLAN.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*. On switches that do not support user-created VRs, all VLANs are created in *VR-Default* and cannot be moved.

Default

A VLAN named *Default* exists on all new or initialized NETGEAR 8800 switches:

- It initially contains all ports on a new or initialized switch, except for the management port(s), if there are any.
- It has an 802.1Q tag of 1.
- The default VLAN is untagged on all ports.
- It uses protocol filter *any*.

A VLAN named *Mgmt* exists on switches that have management modules or management ports:

- It initially contains the management port(s) the switch.
- It is assigned the next available internal VLANid as an 802.1Q tag.

If you do not specify the virtual router, the VLAN is created in the current virtual router.

Usage Guidelines

A newly-created VLAN has no member ports, is untagged, and uses protocol filter *any* until you configure it otherwise. Use the various `configure vlan` commands to configure the VLAN to your needs.

Internal VLANids are assigned automatically using the next available VLANid starting from the high end (4094) of the range.

The VLAN name can include up to 32 characters. VLAN names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are

allowed in the remainder of the name. VLAN names cannot match reserved keywords. For more information on VLAN name requirements and a list of reserved keywords, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

Note: If you use the same name across categories (for example, STPD names), NETGEAR recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

VLAN names are locally significant. That is, VLAN names used on one switch are only meaningful to that switch. If another switch is connected to it, the VLAN names have no significance to the other switch.

You must use mutually exclusive names for:

- VLANs
- vMANs
- Ipv6 tunnels
- BVLANs
- SVLANs
- CVLANs

If you do not specify a virtual router when you create a VLAN, the system creates that VLAN in the default virtual router (*VR-Default*). The management VLAN is always in the management virtual router (*VR-Mgmt*).

Once you create virtual routers, NETGEAR 8800 software allows you to designate one of these as the domain in which all your subsequent configuration commands, including VLAN commands, are applied. If you create virtual routers, ensure that you are creating the VLANs in the desired virtual-router domain.

Note: User-created VRs are supported only on the platforms listed for this feature in the *NETGEAR 8800 User Manual*, Appendix A, “NETGEAR 8800 Software Licenses.” On switches that do not support user-created VRs, all VLANs are created in VR-Default and cannot be moved.

Example

The following command creates a VLAN named *accounting* on the current virtual router:

```
create vlan accounting
```

delete private-vlan

```
delete private-vlan <name>
```

Description

Deletes the PVLAN framework with the specified name.

Syntax Description

name	Specifies the name of the PVLAN to be deleted.
------	--

Default

N/A.

Usage Guidelines

The PVLAN is a framework that links network and subscriber VLANs; it is not an actual VLAN.

This command deletes the PVLAN framework, but it does not delete the associated VLANs. If the ports in the network VLAN were set to *translate*, they are changed to *tagged*.

Example

The following command deletes the PVLAN named *companyx*:

```
delete private-vlan companyx
```

delete protocol

```
delete protocol <name>
```

Description

Deletes a user-defined protocol.

Syntax Description

name	Specifies a protocol name.
------	----------------------------

Default

N/A.

Usage Guidelines

If you delete a protocol that is in use by a VLAN, the protocol associated with that VLAN becomes `none`.

Example

The following command deletes a protocol named *fred*:

```
delete protocol fred
```

delete vlan

```
delete vlan <vlan_name>
```

Description

Deletes a VLAN.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
------------------------	------------------------

Default

N/A.

Usage Guidelines

If you delete a VLAN that has untagged port members and you want those ports to be returned to the default VLAN, you must add them back explicitly using the `configure svlan delete ports` command.

Note: The default VLAN cannot be deleted.

Example

The following command deletes the VLAN *accounting*:

```
delete accounting
```

disable loopback-mode vlan

```
disable loopback-mode vlan <vlan_name>
```

Description

Disallows a VLAN to be placed in the UP state without an external active port. This allows (disallows) the VLAN's routing interface to become active.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

Use this command to specify a stable interface as a source interface for routing protocols. This decreases the possibility of route flapping, which can disrupt connectivity.

Example

The following command disallows the VLAN *accounting* to be placed in the UP state without an external active port:

```
disable loopback-mode vlan accounting
```

disable vlan

```
disable vlan <vlan_name>
```

Description

Use this command to disable the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN you want to disable.
-----------	---

Default

Enabled.

Usage Guidelines

This command allows you to administratively disable specified VLANs. The following guidelines apply to working with disabling VLANs:

- Disabling a VLAN stops *all* traffic on all ports associated with the specified VLAN.
- You cannot disable a VLAN that is running Layer 2 protocol control traffic for protocols such as STP.

When you attempt to disable a VLAN running Layer 2 protocol control traffic, the system returns a message similar to the following:

```
VLAN accounting cannot be disabled because it is actively use by an L2 Protocol
```

- You can disable the default VLAN; ensure that this is necessary prior to disabling the default VLAN.
- You *cannot* disable the management VLAN.
- Although you can remove ports from a disabled VLAN, you *cannot* add ports to a disabled VLAN or bind Layer 2 protocols to that VLAN.

When you attempt to disable a VLAN running Layer 2 protocol traffic, the system returns a message similar to the following:

```
VLAN accounting is disabled. Enable VLAN before adding ports.
```

Example

The following command disables the VLAN named accounting:

```
disable vlan accounting
```

enable loopback-mode vlan

```
enable loopback-mode vlan <vlan_name>
```

Description

Allows a VLAN to be placed in the UP state without an external active port. This allows (disallows) the VLANs routing interface to become active.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

Use this command to specify a stable interface as a source interface for routing protocols. This decreases the possibility of route flapping, which can disrupt connectivity.

Example

The following command allows the VLAN *accounting* to be placed in the UP state without an external active port:

```
enable loopback-mode vlan accounting
```

enable vlan

```
enable vlan <vlan_name>
```

Description

Use this command to re-enable a VLAN that you previously disabled.

Syntax Description

vlan_name	Specifies the VLAN you want to disable.
-----------	---

Default

Enabled.

Usage Guidelines

This command allows you to administratively enable specified VLANs that you previously disabled.

Example

The following command enables the VLAN named accounting:

```
enable vlan accounting
```

show private-vlan

```
show private-vlan
```

Description

Displays information about all the PVLANS on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If the PVLAN is incomplete because it does not have a network or any subscriber VLAN configured, [INCOMPLETE] appears next to the PVLAN name.

Example

The following command displays all the PVLANS on the switch:

```
XCM8810.1 # show private-vlan
```

Name	VID	Protocol	Addr	Flags	Proto	Ports Active / Total	Virtual router

Engineering							
Network VLAN:							
-Engr1	10	-----			ANY	4 / 5	VR-Default
Non-Isolated Subscriber VLAN:							
-ni1	400	-----			ANY	1 / 1	VR-Default
-ni2	401	-----			ANY	1 / 1	VR-Default
Isolated Subscriber VLAN:							
-i1	500	-----			ANY	1 / 1	VR-Default
Ops							
Network VLAN:							
-Ops	20	-----			ANY	2 / 2	VR-Default
Non-Isolated Subscriber VLAN:							
-OpsNi1	901	-----			ANY	1 / 1	VR-Default
-OpsNi2	902	-----			ANY	1 / 1	VR-Default
-OpsNi3	903	-----			ANY	1 / 1	VR-Default
-OpsNi4	904	-----			ANY	1 / 1	VR-Default
Isolated Subscriber VLAN:							
-OpsI0	600	-----			ANY	1 / 1	VR-Default
-OpsI1	601	-----			ANY	1 / 1	VR-Default
-OpsI2	602	-----			ANY	1 / 1	VR-Default
-OpsI3	603	-----			ANY	1 / 1	VR-Default
-OpsI4	604	-----			ANY	1 / 1	VR-Default
Sales [INCOMPLETE]							
Network VLAN:							
-NONE							
Non-Isolated Subscriber VLAN:							
-SalesNi1	701	-----			ANY	1 / 1	VR-Default
-SalesNi2	702	-----			ANY	1 / 1	VR-Default
Isolated Subscriber VLAN:							
-SalesI0	800	-----			ANY	1 / 1	VR-Default

Flags : (d) NetLogin Dynamically created VLAN,							
(D) VLAN Admin Disabled, (f) IP Forwarding Enabled,							
(i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback Enabled,							
(l) MPLS Enabled, (m) IPmc Forwarding Enabled, (n) IP Multinetting Enabled,							
(N) Network LogIn vlan, (o) OSPF Enabled, (p) PIM Enabled,							
(r) RIP Enabled, (T) Member of STP Domain, (V) VPLS Enabled, (v) VRRP Enabled							
Total number of PVLAN(s) : 3							

show private-vlan <name>`show {private-vlan} <name>`**Description**

Displays information about the specified PVLAN.

Syntax Description

name	Specifies the name of the PVLAN to display.
------	---

Default

N/A.

Usage Guidelines

If the PVLAN is incomplete because it does not have a network or any subscriber VLAN configured, [INCOMPLETE] appears next to the PVLAN name.

Example

The following command displays information for the *companyx* PVLAN:

```
XCM8810.1 # show private-vlan "Engineering"
```

```
-----
Name                VID  Protocol Addr          Flags                Proto  Ports  Virtual
                                Active router
                                /Total
-----
Engineering
Network VLAN:
-Engr1              10  ----- ANY      4 /5  VR-Default
Non-Isolated Subscriber VLAN:
-ni1                 400  ----- ANY      1 /1  VR-Default
-ni2                 401  ----- ANY      1 /1  VR-Default
Isolated Subscriber VLAN:
-il                  500  ----- ANY      1 /1  VR-Default
-----

Flags : (d) NetLogin Dynamically created VLAN,
        (D) VLAN Admin Disabled, (f) IP Forwarding Enabled,
        (i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback Enabled,
        (l) MPLS Enabled, (m) IPmc Forwarding Enabled, (n) IP Multinetting Enabled,
        (N) Network LogIn vlan, (o) OSPF Enabled, (p) PIM Enabled,
        (r) RIP Enabled, (T) Member of STP Domain, (V) VPLS Enabled, (v) VRRP Enabled
```

show protocol

`show protocol {<name>}`

Description

Displays protocol filter definitions.

Syntax Description

name	Specifies a protocol filter name.
------	-----------------------------------

Default

Displays all protocol filters.

Usage Guidelines

Displays the defined protocol filter(s) with the types and values of its component protocols.

Example

The following is an example of the `show protocol` command:

Protocol Name	Type	Value

IP	etype	0x0800
	etype	0x0806
ANY	ANY	0xffff
ipx	etype	0x8137
decnet	etype	0x6003
	etype	0x6004
netbios	llc	0xf0f0
	llc	0xf0f1
ipx_8022	llc	0xe0e0
ipx_snap	snap	0x8137
appletalk	snap	0x809b
	snap	0x80f3

show vlan

`show vlan {detail {ipv4 | ipv6} | <vlan_name> {ipv4 | ipv6} | virtual-router <vr-router> | <vlan_name> stpd | security}`

Description

Displays information about VLANs.

Syntax Description

detail	Specifies that detailed information should be displayed for each VLAN.
vlan_name	Specifies a VLAN name.
ipv4	Specifies IPv4.
ipv6	Specifies IPv6.
vr-name	Specifies a virtual router name.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> . On switches that do not support user-created VRs, all VLANs are created in <i>VR-Default</i> and cannot be moved.
stpd	Specifies that STP domains displays for each VLAN.
security	Enables security checking

Default

Summary information for all VLANs on the device.

Usage Guidelines

Note: To display IPv6 information, you must issue either the `show vlan detail` command or `show vlan` command with the name of the specified VLAN.

Unlike many other VLAN-related commands, the keyword *vlan* is required in all forms of this command except when requesting information for a specific vlan.

Use the command `show vlan` to display summary information for all VLANs. It shows various configuration options as a series of *flags* (see the example below). VLAN and protocol names may be abbreviated in this display.

Use the command `show vlan detail` to display detailed information for all VLANs. This displays the same information as for an individual VLAN, but shows every VLAN, one-by-one. After each VLAN display you can elect to continue or quit.

Protocol *none* indicates that this VLAN was configured with a user-defined protocol that has subsequently been deleted.

Note: The NETGEAR 8800 series switches display the Mgmt VLAN in VR-Mgmt.

When an IPv6 address is configured for the VLAN, the system may display one of the following two address types in parentheses after the IPv6 address:

- Tentative
- Duplicate

Note: See the *NETGEAR 8800 User Manual* for information on IPv6 address types.

You can display additional useful information on VLANs configured with IPv6 addresses by issuing the `show ipconfig ipv6 vlan <vlan_name>`. The following is sample output from this command:

```
# show ipconfig ipv6 my_ipv6_100
Router Interface on my_ipv6_100 is enabled and up. MTU: 1500
Locally registered unicast addresses:
2001:db8::8:802:200c:417a/64
fe80::230:48ff:fe41:ed97%my_ipv6_100/64
Flags:
IPv6 Forwarding: YES Accept recvd RA: NO
Send redirects: NO Accept redirects: NO
```

When a displayed VLAN is part of a PVLAN, the display includes the PVLAN name and type (which is network, non-isolated subscriber, or isolated subscriber).

When the displayed VLAN is configured for VLAN translation, the display provides translation VLAN information. If the displayed VLAN is a translation VLAN, a list of translation VLAN members appears. If the displayed VLAN is a member VLAN, the display indicates the translation VLAN to which the member VLAN belongs.

Example

The following is an example of the `show vlan` command on the NETGEAR 8806 switch:

```
XCM8806.4 # show vlan
```

Name	VID	Protocol	Addr	Flags	Proto	Ports Active /Total	Virtual router
alan1	4094	192.18.1.1	/24	-f-----mop-----	ANY	0 /1	VR-Default
alan2	4093	192.18.2.1	/24	-f-----mop-----	ANY	0 /1	VR-Default
alan3	4092	192.18.3.1	/24	-f-----mop-----	ANY	0 /1	VR-Default
alan4	4091	192.18.4.1	/24	-f-----mop-----	ANY	0 /1	VR-Default
CISCO-OSPF	4090	111.1.1.2	/24	-f-----o-----	ANY	0 /1	VR-Default
Default	1			-----T-----	ANY	3 /90	VR-Default
Mgmt	4095	172.26.2.145	/24	-----	ANY	1 /1	VR-Mgmt
VLANRIP	4088	123.1.1.1	/24	-f-----r-----	ANY	0 /1	VR-Default

Flags : (c) 802.1ad customer VLAN (d) NetLogin Dynamically created VLAN,
 (D) VLAN Admin Disabled,
 (f) IP Forwarding Enabled, (F) Learning Disabled,
 (L) Loopback Enabled, (m) IPmc Forwarding Enabled,
 (M) Subscriber VLAN, (n) IP Multinetting Enabled,
 (N) Network Login VLAN, (o) OSPF Enabled,
 (O) Flooding Disabled, (p) PIM Enabled,
 (r) RIP Enabled, (R) Sub-VLAN IP Range Configured,
 (s) Sub-VLAN, (S) Super-VLAN, (t) Network VLAN,
 (T) Member of STP Domain, (v) VRRP Enabled,

Total number of VLAN(s) : 9

The following is an example of the `show vlan Default` command:

```
* XCM8806.5 # show vlan "Default"
VLAN Interface with name Default created by user
Admin State:      Enabled Tagging:      802.1Q Tag 1
Virtual router:   VR-Default
IPv6:             None
STPD:             s0(Disabled,Auto-bind)
Protocol:         Match all unfiltered protocols
Loopback:         Disabled
NetLogin:         Disabled
QosProfile:       None configured
Egress Rate Limit Designated Port: None configured
Flood Rate Limit QosProfile:      None configured
Ports:    90.      (Number of active ports=3)
           Untag:   1:1,    1:2,    1:7,    1:8,    1:9,    1:10,   1:11,
                   1:12,   1:13,   1:14,   1:15,   1:16,   1:17,   1:18,
                   1:19,   1:20,   1:21,   1:22,   1:23,   1:24,   3:1,
                   3:2,    3:3,    3:4,    3:5,    3:6,    3:7,    3:8,
                   4:1,    4:2,    4:3,    4:4,    4:5,    4:6,    4:7,
                   4:8,    5:1,    5:2,    5:3,    *5:4,   5:5,    5:6,
                   5:7,    5:8,    6:1,    6:2,    6:3,    6:5,    6:6,
                   6:7,    6:8,    6:9,    6:10,   6:12,   6:13,   6:14,
                   6:15,   6:16,   6:17,   6:18,   6:19,   6:20,   6:21,
                   6:22,   6:23,   *6:24,   6:25,   6:26,   6:27,   6:28,
                   6:29,   6:30,   6:31,   6:32,   6:33,   6:34,   6:35,
                   6:36,   6:37,   6:38,   6:39,   6:40,   6:41,   6:42,
                   6:43,   6:44,   6:45,   6:46,   6:47,   *6:48

Flags:  (*) Active, (!) Disabled, (g) Load Sharing port
        (b) Port blocked on the vlan, (m) Mac-Based port
        (a) Egress traffic allowed for NetLogin
        (u) Egress traffic unallowed for NetLogin
        (t) Translate VLAN tag for Private-VLAN
        (s) Private-VLAN System Port, (L) Loopback port
        (e) Private-VLAN End Point Port
```

(x) VMAN Tag Translated port

Note: The `m` flag for MAC-based ports represents network login information.

Note: The `number of active ports` line displays the number of ports presently in forwarding state on this VLAN.

The output for the `show vlan detail` command displays the same information for all VLANs configured on the switch.

Note: See [Chapter 19, “IP Unicast Commands,”](#) for information on adding secondary IP addresses to VLANs.

unconfigure vlan ipaddress

```
unconfigure {vlan} <vlan_name> ipaddress {<ipv6_address_mask>}
```

Description

Removes the IP address of the VLAN or a vMAN. With no parameters, the command removes the primary IPv4 address on the specified VLAN. Using the IPv6 parameters, you can remove specified IPv6 addresses from the specified VLAN.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
<code>ipv6_address_mask</code>	Specifies an IPv6 address using the format of IPv6-address/prefix-length, where IPv6 is the 128-bit address and the prefix length specifies the number of leftmost bits that comprise the prefix.

Default

Removes the primary IPv4 address from the specified VLAN.

Usage Guidelines

Note: You need an Advanced license to use vMANs.

If you do not specify any parameters, this command removes the primary IPv4 address from the VLAN.

Note: With IPv6, you cannot remove the last link local IPv6 address until all global IPv6 addresses are removed.

Example

The following command removes the primary IPv4 address from the VLAN *accounting*:

```
unconfigure vlan accounting ipaddress
```

The following command removes an IPv6 addresses from the VLAN *finance*:

```
unconfigure vlan finance ipaddress 3ffe::1
```

This chapter describes commands for:

- Configuring FDB entries
- Displaying FDB entries

For an introduction to FDB features, see the *NETGEAR 8800 User Manual*.

clear counters fdb mac-tracking

```
clear counters fdb mac-tracking [<mac_addr> | all]
```

Description

Clears the event counters for the FDB MAC-tracking feature.

Syntax Description

mac_addr	Specifies a MAC address, using colon-separated bytes.
all	Clears the counters for all tracked MAC addresses.

Default

N/A.

Usage Guidelines

The `clear counters` command also clears the counters for all tracked MAC addresses.

Example

The following command example clears the counters for all entries in the MAC address tracking table:

```
XCM8806.1 # clear counters fdb mac-tracking all
```

clear fdb

```
clear fdb {<mac_addr> | ports <port_list> | vlan <vlan_name> | blackhole}
```

Description

Clears dynamic FDB entries that match the filter.

Syntax Description

mac_addr	Specifies a MAC address, using colon-separated bytes.
port_list	Specifies one or more ports or slots and ports.
vlan_name	Specifies a VLAN name.
blackhole	Specifies the blackhole entries.

Default

Clears all dynamic FDB entries.

Usage Guidelines

This command clears FDB entries based on the specified criteria. When no options are specified, the command clears all dynamic FDB entries.

Examples

The following command clears any FDB entries associated with ports 4:3-4:5 on the switch:

```
clear fdb ports 4:3-4:5
```

The following command clears any FDB entries associated with VLAN *corporate*:

```
clear fdb vlan corporate
```

configure fdb agingtime

```
configure fdb agingtime <seconds>
```

Description

Configures the FDB aging time for dynamic entries.

Syntax Description

seconds	Specifies the FDB aging time in seconds. A value of 0 indicates that the entry should never be aged out. The NETGEAR 8800 can support the value 0 (no aging) and a range of 15 to 1,000,000 seconds.
---------	---

Default

300 seconds.

Usage Guidelines

If the aging time is set to zero, all dynamic entries in the database become static, nonaging entries. This means that they do not age out, but non-permanent static entries can be deleted if the switch is reset.

On NETGEAR 8800 switches, the software flushes the FDB table once the aging timeout parameter is reached, even if the switch is running traffic and populating addresses in the FDB table.

Example

The following command sets the FDB aging time to 3,000 seconds:

```
configure fdb agingtime 3000
```

configure fdb mac-tracking ports

```
configure fdb mac-tracking {[add|delete]} ports [<port_list>|all]
```

Description

Enables or disables MAC address tracking for all MAC addresses on the specified ports.

Syntax Description

add	Enables MAC address tracking for the specified ports.
delete	Disables MAC address tracking for the specified ports.
port_list	Specifies a list of ports on which MAC address tracking is to be enabled or disabled.
all	Specifies that MAC address tracking is to be enabled or disabled on all ports.

Default

No ports are enabled for MAC address tracking.

Usage Guidelines

MAC address tracking events on enabled ports generate EMS messages and can optionally generate SNMP traps.

Note: When a MAC address is configured in the tracking table, but detected on a MAC tracking enabled port, the per MAC address statistical counters are not updated.

Example

The following command enables MAC address tracking for all MAC addresses on port 2:1:

```
configure fdb mac-tracking add ports 2:1
```

create fdb mac-tracking entry

```
create fdb mac-tracking entry <mac_addr>
```

Description

Adds a MAC address to the MAC address tracking table.

Syntax Description

mac_addr	Specifies a device MAC address, using colon-separated bytes.
----------	--

Default

The MAC address tracking table is empty.

Usage Guidelines

None.

Example

The following command adds a MAC address to the MAC address tracking table:

```
create fdb mac-tracking entry 00:E0:2B:12:34:56
```

create fdbentry vlan ports

```
create fdbentry <mac_addr> vlan <vlan_name> [ports <port_list> | blackhole]
```

Description

Creates a permanent static FDB entry.

Syntax Description

mac_addr	Specifies a device MAC address, using colon-separated bytes.
vlan_name	Specifies a VLAN name associated with a MAC address.
port_list	Specifies one or more ports or slots and ports associated with the MAC address.
interface-list	Specifies one or more interfaces to associate with the MAC address.

blackhole	Enables the <i>blackhole</i> option. Any packets with either a source MAC address or a destination MAC address matching the FDB entry are dropped.
-----------	--

Default

N/A.

Usage Guidelines

Permanent entries are retained in the database if the switch is reset or a power off/on cycle occurs. A permanent static entry can either be a unicast or multicast MAC address. After they have been created, permanent static entries stay the same as when they were created. If the same MAC address and VLAN is encountered on another virtual port that is not included in the permanent MAC entry, it is handled as a blackhole entry. The static entry is not updated when any of the following take place:

- A VLAN identifier (VLANid) is changed.
- A port is disabled.
- A port enters blocking state.
- A port goes down (link down).

A permanent static FDB entry is deleted when any of the following take place:

- A VLAN is deleted.
- A port mode is changed (tagged/untagged).
- A port is deleted from a VLAN.

Permanent static entries are designated by *spm* in the flags field of the `show fdb` output. You can use the `show fdb` command to display permanent FDB entries.

If the static entry is for a PVLAN VLAN that requires more than one underlying entry, the system automatically adds the required entries. For example, if the static entry is for a PVLAN network VLAN, the system automatically adds all required extra entries for the subscriber VLANs.

You can create FDB entries to multicast MAC addresses and list one or more ports. If more than one port number is associated with a permanent MAC entry, packets are multicast to the multiple destinations.

IGMP snooping rules take precedence over static multicast MAC addresses in the IP multicast range (01:00:5e:xx:xx:xx) unless IGMP snooping is disabled.

Note: When a multiport list is assigned to a unicast MAC address, load sharing is not supported on the ports in the multiport list.

Examples

The following command adds a permanent, static entry to the FDB for MAC address 00 E0 2B 12 34 56, in VLAN marketing on slot 2, port 4 on the switch:

```
create fdbentry 00:E0:2B:12:34:56 vlan marketing port 2:4
```

The following example creates a multicast FDB entry, in VLAN black, on slot 1, ports 1, 2, and 4, on the NETGEAR 8800 switches:

```
create fdbentry 01:00:00:00:00:01 vlan black port 1:1, 1:2, 1:4
```

delete fdb mac-tracking entry

```
delete fdb mac-tracking entry [<mac_addr> | all]
```

Description

Deletes a MAC address from the MAC address tracking table.

Syntax Description

mac_addr	Specifies a device MAC address, using colon-separated bytes.
all	Specifies that all MAC addresses are to be deleted from the MAC address tracking table.

Default

The MAC address tracking table is empty.

Usage Guidelines

None.

Example

The following command deletes a MAC address from the MAC address tracking table:

```
delete fdb mac-tracking entry 00:E0:2B:12:34:56
```

delete fdbentry

```
delete fdbentry [all | <mac_address> [vlan <vlan name>]]
```

Description

Deletes one or all permanent FDB entries.

Syntax Description

all	Specifies all FDB entries.
mac_address	Specifies a device MAC address, using colon-separated bytes.
vlan_name	Specifies the specific VLAN name.

Default

N/A.

Usage Guidelines

None.

Examples

The following example deletes a permanent entry from the FDB:

```
delete fdbentry 00:E0:2B:12:34:56 vlan marketing
```

The following example deletes all permanent entries from the FDB:

```
delete fdbentry all
```

disable flooding ports

```
disable flooding [all_cast | broadcast | multicast | unicast] ports [<port_list> | all]
```

Description

Disables Layer 2 egress flooding on one or more ports. With the NETGEAR 8800 family of switches, you can further identify the type of packets for which to block flooding.

Syntax Description

all_cast	Specifies disabling egress flooding for <i>all</i> packets on specified ports.
broadcast	Specifies disabling egress flooding only for broadcast packets.
multicast	Specifies disabling egress flooding only for multicast packets.
unicast	Specifies disabling egress flooding only for unknown unicast packets.
port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled for all packet types.

Usage Guidelines

Note: If an application requests specific packets on a specific port, those packets are not affected by the `disable flooding ports` command.

You might want to disable egress flooding to do the following:

- enhance security
- enhance privacy
- improve network performance

This is particularly useful when you are working on an edge device in the network. The practice of limiting flooded egress packets to selected interfaces is also known as upstream forwarding.

Note: If you disable egress flooding with static MAC addresses, this can affect many protocols, such as IP and ARP.

The following guidelines apply to enabling and disabling egress flooding:

- Disabling multicasting egress flooding does not affect those packets within an IGMP membership group at all; those packets are still forwarded out. If IGMP snooping is disabled, multicast packets are not flooded.
- Egress flooding can be disabled on ports that are in a load-sharing group. In a load-sharing group, the ports in the group take on the egress flooding state of the master port; each member port of the load-sharing group has the same state as the master port.
- FDB learning takes place on ingress ports and is independent of egress flooding; either can be enabled or disabled independently.
- Disabling unicast or all egress flooding to a port also stops packets with unknown MAC addresses to be flooded *to* that port.
- Disabling broadcast or all egress flooding to a port also stops broadcast packets to be flooded *to* that port.

You can disable egress flooding for unicast, multicast, or broadcast MAC addresses, as well as for all packets on the ports of the NETGEAR 8800 family of switches. The default behavior for the NETGEAR 8800 family of switches is enabled egress flooding for all packet types.

Example

The following command disables egress flooding on slot 4, ports 5 and 6 on a NETGEAR 8800 switch:

```
disable flooding all_cast port 4:5-4:6
```

disable learning iparp sender-mac

```
disable learning iparp {vr <vr_name>} sender-mac
```

Description

Disables MAC address learning from the payload of IP ARP packets.

Syntax Description

vr_name	Specifies a virtual router.
---------	-----------------------------

Default

Disabled.

Usage Guidelines

To view the configuration for this feature, use the following command:

```
show iparp
```

Example

The following command disables MAC address learning from the payload of IP ARP packets:

```
disable learning iparp sender-mac
```

disable learning port

```
disable learning {drop-packets | forward-packets} port [<port_list> | all]
```

Description

Disables MAC address learning on one or more ports for security purposes.

Syntax Description

port	Specifies the port.
port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports and slots.
drop-packets	Specifies that packets with unknown source MAC addresses be dropped. If you do not specify the <code>forward-packets</code> option, this option is used.
forward-packets	Specifies that packets with unknown source MAC addresses be forwarded.

Default

Enabled.

Usage Guidelines

Use this command in a secure environment where access is granted via permanent forwarding database (FDB) entries per port.

Example

The following command disables MAC address learning on port 4:3:

```
disable learning ports 4:3
```

disable snmp traps fdb mac-tracking

```
disable snmp traps fdb mac-tracking
```

Description

Disables SNMP trap generation when MAC-tracking events occur for a tracked MAC address.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables SNMP traps for MAC-tracking events:

```
disable snmp traps fdb mac-tracking
```

enable flooding ports

```
enable flooding [all_cast | broadcast | multicast | unicast] ports [<port_list> | all]
```

Description

Enables egress flooding on one or more ports. With the NETGEAR 8800 series switches, you can further identify the type of packets to flood on the specified ports.

Syntax Description

all_cast	Specifies enabling egress flooding for <i>all</i> packets on specified ports.
broadcast	Specifies enabling egress flooding only for broadcast packets.

multicast	Specifies enabling egress flooding only for multicast packets.
unicast	Specifies enabling egress flooding only for unknown unicast packets.
port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled for all packet types.

Usage Guidelines

Use this command to re-enable egress flooding that you previously disabled using the `disable flooding ports` command.

The following guidelines apply to enabling and disabling egress flooding:

- Disabling multicasting egress flooding does not affect those packets within an IGMP membership group at all; those packets are still forwarded out. If IGMP snooping is disabled, multicast packets are not flooded.
- Egress flooding can be disabled on ports that are in a load-sharing group. If that is the situation, the ports in the group take on the egress flooding state of the master port; each member port of the load-sharing group has the same state as the master port.
- FDB learning is independent of egress flooding. FDB learning and egress flooding can be enabled or disabled independently.
- Disabling unicast or all egress flooding to a port also stops packets with unknown MAC addresses to be flooded *to* that port.
- Disabling broadcast or all egress flooding to a port also stops broadcast packets to be flooded *to* that port.

You can disable egress flooding for unicast, multicast, or broadcast MAC addresses, as well as for all packets on the ports of the NETGEAR 8800 series switches. The default behavior for the NETGEAR 8800 series switches is enabled egress flooding for all packet types.

Example

The following command enables egress flooding on slot 1, ports 1 and 2 on a NETGEAR 8800 switch:

```
enable flooding all_cast port 1:1-1:2
```

enable learning iparp sender-mac

```
enable learning iparp {request | reply | both-request-and-reply} {vr <vr_name>} sender-mac
```

Description

Enables MAC address learning from the payload of IP ARP packets.

Syntax Description

request	Enables learning only for IP ARP request packets.
reply	Enables learning only for IP ARP reply packets.
both-request-and-reply	Enables learning for both request and reply packets.
vr_name	Specifies a virtual router.

Default

Disabled.

Usage Guidelines

To view the configuration for this feature, use the following command:

```
show iparp
```

Example

The following command enables MAC address learning from the payload of reply IP ARP packets:

```
enable learning iparp reply sender-mac
```

enable learning port

```
enable learning ports [all | <port_list>]
```

Description

Enables MAC address learning on one or more ports.

Syntax Description

all	Specifies all ports.
port_list	Specifies one or more ports or slots and ports.

Default

Enabled.

Example

The following command enables MAC address learning on slot 1, ports 7 and 8 on the switch:

```
enable learning ports 1:7-8
```


enable snmp traps fdb mac-tracking

```
enable snmp traps fdb mac-tracking
```

Description

Enables SNMP trap generation when MAC-tracking events occur for a tracked MAC address.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables SNMP traps for MAC-tracking events:

```
enable snmp traps fdb mac-tracking
```

show fdb

```
show fdb {blackhole {netlogin [all | mac-based-vlans]} | netlogin [all | mac-based-vlans] |
permanent {netlogin [all | mac-based-vlans]} | <mac_addr> {netlogin [all | mac-based-vlans]}
| ports <port_list> {netlogin [all | mac-based-vlans]} | vlan <vlan_name> {netlogin [all |
mac-based-vlans]}}
```

Description

Displays FDB entries.

Syntax Description

blackhole	Displays the blackhole entries. (All packets addressed to these entries are dropped.)
slot	Specifies a slot in the switch.
num_entries	Specifies the maximum number of hardware entries to display. The range is 1 to 25.
netlogin all	Displays all FDBs created as a result of the netlogin process.
netlogin mac-based-vlans	Displays all netlogin MAC-based VLAN FDB entries. See Chapter 16, “Network Login Commands,” for more information on netlogin.
permanent	Displays all permanent entries, including the ingress and egress QoS profiles.

mac_addr	Specifies a MAC address, using colon-separated bytes, for which FDB entries should be displayed.
port_list	Displays the entries for one or more ports or ports and slots.
vlan_name	Displays the entries for a specific VLAN.

Default

All.

Usage Guidelines

The `show fdb` command output displays the following information:

Mac	The MAC address that defines the entry.
Vlan	The PVLAN or VLAN for the entry.
Age	The age of the entry, in seconds (does not appear if the keyword permanent is specified). The age parameter does not display for the backup MSM/MM on the switch.
Flags	Flags that define the type of entry: <ul style="list-style-type: none"> • b - Ingress Blackhole • B - Egress Blackhole • D - Drop entry for an isolated subscriber VLAN • d - Dynamic • h - Aged in hardware • i - an entry also exists in the IP FDB • l - lockdown MAC • L - lockdown-timeout MAC • m - MAC • M - Mirror • n - NetLogin • o - IEEE 802.1ah backbone MAC • P - PVLAN created entry • p - Permanent • s - Static • v - NetLogin MAC-Based VLAN • x - an entry also exists in the IPX FDBs
Port List	The ports on which the MAC address has been learned.

Examples

The following command example shows how the FDB entries appear for all options except the `hardware` option:

```
# show fdb
```

Mac	Vlan	Age	Flags	Port / Virtual Port List
00:0c:29:4b:34:cf	v101(0101)	0041	d m	D 1:2
00:0c:29:4b:34:cf	v100(0100)	0041	d m	P 1:2
00:0c:29:d2:2d:48	v102(0102)	0045	d m	1:3
00:0c:29:d2:2d:48	v100(0100)	0045	d m	P 1:3
00:0c:29:f1:f2:f5	v100(0100)	0045	d m	1:1
00:0c:29:f1:f2:f5	v102(0102)	0045	d m	P 1:1
00:0c:29:f1:f2:f5	v101(0101)	0045	d m	P 1:1

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
 x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress
 Blackhole,
 b - Ingress Blackhole, v - MAC-Based VLAN, P - Private VLAN, T - VLAN
 translation,
 D - drop packet, h - Hardware Aging, o - IEEE 802.1ah Backbone MAC.

Total: 3 Static: 0 Perm: 0 Dyn: 3 Dropped: 0 Locked: 0 Locked with Timeout: 0
 FDB Aging time: 300
 FDB VPLS Aging time: 300

The following example shows the display format when a PVLAN is configured. Note that VLAN translation is configured on some ports (as indicated by the t flag).

XCM8806.9 # show fdb

Mac	Vlan	Age	Flags	Port / Virtual Port List
00:04:0d:f3:9b:84	Default(0001)	0048	d m	6:48
00:1a:b9:33:f8:68	Default(0001)	0000	d m	6:48
00:23:ac:da:4c:0b	Default(0001)	0044	d m	6:48
00:d0:b0:10:c7:00	Default(0001)	0028	d m	6:24
00:d0:b0:10:cb:00	Default(0001)	0005	d m	6:48
e0:91:f5:06:2c:2a	Default(0001)	0050	d m	6:48

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
 x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress Blackhole,
 b - Ingress Blackhole, v - MAC-Based VLAN, P - Private VLAN, T - VLAN translation,
 D - drop packet, h - Hardware Aging, o - IEEE 802.1ah Backbone MAC.

Total: 6 Static: 0 Perm: 0 Dyn: 6 Dropped: 0 Locked: 0 Locked with Timeout: 0
 FDB Aging time: 300
 FDB VPLS Aging time: 300

show fdb mac-tracking configuration

show fdb mac-tracking configuration

Description

Displays configuration information for the MAC address tracking feature.

Syntax Description

This command has no arguments or variables.

Default

The MAC address tracking table is empty.

Usage Guidelines

None.

Example

The following command example displays the contents of the MAC address tracking table:

```
Switch.8 # show fdb mac-tracking configuration
MAC-Tracking enabled ports: 1-3,10,20
SNMP trap notification      : Enabled
MAC address tracking table (4 entries):
  00:30:48:72:ee:88
  00:21:9b:0e:ca:32
  00:12:48:82:9c:56
  00:30:48:84:d4:16
```

show fdb mac-tracking statistics

```
show fdb mac-tracking statistics {<mac_addr>} {no-refresh}
```

Description

Displays statistics for the MAC addresses that are being tracked.

Syntax Description

mac_addr	Specifies a MAC address, using colon-separated bytes, for which FDB entries should be displayed.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.

Default

N/A.

Usage Guidelines

Use the keys listed below the display to clear the statistics counters or page up or down through the table entries.

Example

The following command example displays statistics for the entries in the MAC address tracking table:

```
XCM8810.3 # show fdb mac-tracking statistics
MAC Tracking Statistics      Fri Mar 20 15:25:01 2009
      Add      Move      Delete
MAC Address  events  events  events
=====
00:00:00:00:00:01      0      0      0
00:00:00:00:00:02      0      0      0
00:00:00:00:00:03      0      0      0
00:00:00:00:00:04      0      0      0
00:00:00:00:00:05      0      0      0
00:00:00:00:00:06      0      0      0
00:00:00:00:00:07      0      0      0
00:00:00:00:00:08      0      0      0
00:00:00:00:00:09      0      0      0
00:00:00:00:00:10      0      0      0
00:00:00:00:00:11      0      0      0
00:00:00:00:00:12      0      0      0
00:00:00:00:00:13      0      0      0
00:00:00:00:00:14      0      0      0
00:00:00:00:00:15      0      0      0
00:00:00:00:00:16      0      0      0
00:00:00:00:00:17      0      0      0
00:00:00:00:00:18      0      0      0
=====
0->Clear Counters  U->page up  D->page down  ESC->exit
```

show fdb stats

```
show fdb stats [{ports {all | <port_list>} | vlan {all} | {vlan} <vlan_name> } {no-refresh}]
```

Description

Displays FDB entry statistics for the specified ports or VLANs in either a dynamic or a static report.

Syntax Description

all	Requests statistics for all ports or all VLANs.
port_list	Specifies which ports are to be included in the statistics display.
vlan_name	Specifies a single VLAN to be included in the statistics display.
no-refresh	Specifies a static display, which is not automatically updated.

Default

Summary FDB statistics for the switch.

Usage Guidelines

The dynamic display remains visible and continues to update until you press <Esc>.

The `show fdb stats` command output displays the following information:

Port	When you chose to display statistics for ports, this column displays port numbers.
Link State	When you chose to display statistics for ports, this column displays the link states, which are described at the bottom of the display.
VLAN	When you chose to display statistics for VLANs, this column displays VLAN names.
MAC Addresses	This column displays the total number of MAC addresses for each port or VLAN.
Dynamic	This column displays the total number of MAC addresses that were learned dynamically for each port or VLAN.
Static	This column displays the total number of MAC addresses that are configured on this switch for each port or VLAN.
Dropped	This column displays the total number of dynamic MAC addresses that were discovered, but not stored in the FDB. Discovered MAC addresses might be dropped because a configured learning limit is reached, the FDB is in lockdown, or a port forwarding state is in transition. Some conditions that lead to dropped MAC addresses can produce log messages or SNMP traps.

Examples

The following command example displays summary FDB statistics for the switch:

```
torinol.1 # show fdb stats

Total: 4 Static: 3 Perm: 3 Dyn: 1 Dropped: 0
FDB Aging time: 300
FDB VPLS Aging time: 300
(pacman debug) torinol.2 #
```

The following command example displays FDB statistics for ports 1 to 16 on slot 1:

```
# show fdb stats ports 1:1-1:16

FDB Stats                               Mon Mar 15 15:30:49 2010
Port      Link   MAC
          State  Addresses      Dynamic      Static      Dropped
=====
1:1        A      2394           2389           5           2
```

1:2	A	37	37	0	0
1:3	A	122	121	1	452
1:4	R	0	0	0	0
1:5	R	0	0	0	0
1:6	A	43	43	0	0
1:7	A	118	118	0	0
1:8	R	0	0	0	0
1:9	R	0	0	0	0
1:10	A	8	8	0	0
1:11	A	2998	2990	8	1
1:12	A	486	486	0	0
1:13	R	0	0	0	0
1:14	A	42	42	0	0
1:15	A	795	795	0	0
1:16	A	23	23	0	2

=====

Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback

U->page up D->page down ESC->exit

The following command example displays FDB statistics for all VLANs:

```
# show fdb stats vlan all
```

FDB Stats		Mon Mar 15 15:30:49 2010		
VLAN	MAC Addresses	Dynamic	Static	Dropped
=====				
SV_PPPOE	2394	2389	5	2
NV_PPPOE	122	121	1	452
=====				

U->page up D->page down ESC->exit

Commands for Virtual Routers

11

This chapter describes commands for:

- Creating and deleting virtual routers
- Configuring and managing virtual routers
- Displaying information about virtual routers

For an introduction to virtual routers, see the *NETGEAR 8800 User Manual*.

configure vr add ports

```
configure vr <vr-name> add ports <portlist>
```

Description

Assigns a list of ports to the virtual router specified.

Syntax Description

vr-name	Specifies the name of the virtual router.
portlist	Specifies the ports to add to the virtual router.

Default

By default, all ports are assigned to the virtual router, *VR-Default*.

Usage Guidelines

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to assign ports to a virtual router. Since all ports are initially assigned to *VR-Default*, you might need to delete the desired ports first from the virtual router where they reside, before you add them to the desired virtual router.

If you plan to assign VR ports to a VLAN, be aware that the ports that you add to a VLAN and the VLAN itself cannot be explicitly assigned to different virtual routers. When multiple virtual routers are defined, consider the following guidelines while adding ports to a VR:

- A VLAN can belong (either through explicit or implicit assignment) to only one VR.
- If a VLAN is not explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to a single VR.
- If a VLAN is explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to the same VR or to no VR.
- If a port is added to VLANs that are explicitly assigned to different VRs, the port must be explicitly assigned to no VR.

Example

The following command adds all the ports on slot 2 to the virtual router *vr-acme*:

```
configure vr vr-acme add ports 2:*
```

configure vr add protocol

```
configure vr <vr-name> add protocol <protocol-name>
```

Description

Starts a Layer 3 protocol on a virtual router.

Syntax Description

vr-name	Specifies the name of the virtual router.
protocol-name	Specifies the Layer 3 protocol.

Default

N/A.

Usage Guidelines

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to start the Layer 3 protocol specified on the virtual router. The choices for *protocol-name* are:

- RIP
- OSPF
- BGP
- PIM

MPLS is the only protocol that you can add to or delete from the *VR-Default* virtual router. You cannot add or delete any other protocols from *VR-Default*, and you cannot add or delete any protocols from the other system virtual routers, *VR-Mgmt* and *VR-Control*.

Example

The following command starts RIP on the virtual router *vr-acme*:

```
configure vr vr-acme add protocol rip
```

configure vr delete ports

```
configure vr <vr-name> delete ports <portlist>
```

Description

Removes a list of ports from the virtual router specified.

Syntax Description

vr-name	Specifies the name of the virtual router.
portlist	Specifies the ports to remove from the virtual router.

Default

By default, all ports are assigned to the virtual router, *VR-Default*.

Usage Guidelines

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to remove ports from a virtual router. Since all ports are initially assigned to *VR-Default*, you might need to delete the desired ports first from the virtual router where they reside, before you add them to the desired virtual router.

Example

The following command removes all the ports on slot 2 from the virtual router *vr-acme*:

```
configure vr vr-acme delete ports 2:*
```

configure vr delete protocol

```
configure vr <vr-name> delete protocol <protocol-name>
```

Description

Stops and removes a Layer 3 protocol on a virtual router.

Syntax Description

vr-name	Specifies the name of the virtual router.
protocol-name	Specifies the Layer 3 protocol.

Default

N/A.

Usage Guidelines

The choices for `protocol-name` are:

- RIP
- OSPF
- BGP
- PIM

You cannot add or delete any other protocols from *VR-Default*, and you cannot add or delete any protocols from the other system virtual routers, *VR-Mgmt* and *VR-Control*.

Example

The following command shutsdowns and removes RIP from the virtual router *vr-acme*:

```
configure vr vr-acme delete protocol rip
```

create virtual-router

```
create virtual-router <vr-name>
```

Description

Creates a user virtual router.

Syntax Description

vr-name	Specifies the name of the user virtual router.
---------	--

Default

N/A.

Usage Guidelines

This command creates a new user virtual router. The three default system virtual routers, *VR-Mgmt*, *VR-Control*, and *VR-Default* always exist and cannot be deleted or renamed. For

backward compatibility, you cannot name a virtual router *VR-0*, *VR-1*, or *VR-2*, as they were the original names of the system virtual routers.

A virtual router name must begin with an alphabetical character and may contain alphanumeric characters and underscores (`_`), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. The name must be unique among the VLAN and virtual router names on the switch. Virtual router names are case insensitive. For information on virtual router name guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added.-

Example

The following command creates the virtual router *vr-acme*:

```
create virtual-router vr-acme
```

delete virtual-router

```
delete virtual-router <vr-name>
```

Description

Deletes a virtual router.

Syntax Description

vr-name	Specifies the name of the virtual router.
---------	---

Default

N/A.

Usage Guidelines

Only user virtual routers can be deleted. When a virtual router gets deleted, all of the VLANs in the virtual router are deleted. All of the ports assigned to this virtual router are deleted and made available to assign to other virtual routers. Any routing protocol that is running on the virtual router is shut down and deleted gracefully.

Example

The following command creates the virtual router *vr-acme*:

```
delete virtual-router vr-acme
```

show virtual-router

```
show virtual-router {<vr-name>}
```

Description

Displays information about the virtual routers.

Syntax Description

vr-name	Specifies the name of the virtual router.
---------	---

Default

N/A.

Usage Guidelines

During system boot up, the NETGEAR 8800 creates three system virtual routers: *VR-Mgmt*, *VR-Control*, and *VR-Default* (previous to release 11.0 these virtual routers were named *VR-0*, *VR-1*, and *VR-2*, respectively). The following defines each system virtual router:

- The management port on both the primary and backup MSMs/MMs and the VLAN *mgmt* belong to *VR-Mgmt*.
- Internal system operations use *VR-Control*.
- The default VLAN belongs to *VR-Default*.

Beginning with release 11.0, you can create additional virtual routers, called user virtual routers. User virtual routers are created without any routing protocols, so the protocols must be added. The protocols on the system virtual routers are predefined and cannot be changed.

The output displays, in tabular format, the:

- Name of the virtual router
- Number of VLANs that belong to that virtual router
- Number of ports that belong to that virtual router
- Which routing protocols have been added to that virtual router

When you specify a particular virtual router, the output displays:

- The number of ports
- A list of ports
- The protocols configured
- The name of the process supporting the protocol on that virtual router

Examples

The following command displays the virtual router configurations on the switch:

```
Switch.19 # show virtual-router
-----
Virtual Router    Number of    Number of    Flags
```

```

                Vlans      Ports
-----
region1         7         0         -----
VR-Control      0         0         -----
VR-Default      1        20        boprimOR
VR-Mgmt         1         0         -----
-----

Flags : Routing protocols configured on the virtual router
      (b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
      (O) OSPFv3, (R) RIPng
System Totals      :
Total Virtual Routers :      4      Max Virtual Routers      :      67
Total Protocols      :      8      Max Protocols              :      48

```

The following command displays the virtual router *VR-Default*:

```

Switch.20 # show virtual-router "VR-Default"
Virtual router : VR-Default
No of vlans    : 1
No of ports    : 20
Port List      : 1:1-20
Protocols Configured:
    Protocol: BGP,      Process Name: bgp
    Protocol: OSPF,     Process Name: ospf
    Protocol: RIP,      Process Name: rip
    Protocol: PIM,      Process Name: pim
    Protocol: ISIS,     Process Name: isis
    Protocol: MPLS,     Process Name: mpls
    Protocol: OSPFv3,   Process Name: ospfv3
    Protocol: RIPng,    Process Name: ripng
VLANs          : Default

Virtual Router Totals :
Total Protocols      :      8      Max Protocols              :      8

```

The following command displays information for user virtual router *region1*:

```

Switch.21 # show virtual-router region1
Virtual router : region1
No of vlans    : 7
Protocols Configured:
    Protocol: BGP,      Process Name: bgp-3
    Protocol: OSPF,     Process Name: ospf-3
VLANs          : zone1, zone2, zone3,
                zone4, zone5, zone6,
                zone7

Virtual Router Totals :
Total Protocols      :      2      Max Protocols              :      6

```

virtual-router

```
virtual-router {<vr-name>}
```

Description

Changes the virtual router context.

Syntax Description

vr-name	Specifies the name of the virtual router.
---------	---

Default

N/A.

Usage Guidelines

Use this command to change the virtual router context for subsequent commands. When you issue the command, the prompt changes to reflect the virtual router domain. Configuration commands for Layer 3 routing protocols, creating VLANs, and deleting VLANs apply only to the current virtual router context.

Under a virtual router configuration domain, any virtual router commands are applied only to that virtual router. The virtual router commands consist of all the BGP, OSPF, PIM and RIP commands, and the commands listed in [Table 14](#).

Table 14. Virtual Router Commands

[enable disable] ipforwarding
clear iparp *
clear counters iparp *
configure iparp *
configure iparp [add delete] *
[enable disable] iparp *
show iparp *
configure iproute [add delete] *
show iproute *
show ipstats *
rtlookup
create [vlan vman] <vlan-name>
[enable disable] igmp

Table 14. Virtual Router Commands (Continued)

[enable disable] igmp snooping *
[enable disable] ipmcforwarding
show igmp
show igmp snooping
show igmp group
show igmp snooping cache

* Indicates that other commands are available with these listed.

The virtual router context simplifies configuration because you do not have to specify the virtual router for each individual protocol configuration command. The current VR context is indicated in the command line interface (CLI) prompt.

For example, if you wish to configure OSPF for the user virtual router *vr-manufacturing*, you would change the virtual router context to that of *vr-manufacturing*. All the subsequent OSPF commands would apply to that virtual router, unless the context is changed again.

A virtual router is identified by a name (up to 32 characters long). The name must be unique among the VLAN and virtual router names on the switch. For backward compatibility, you cannot name a virtual router *VR-0*, *VR-1*, or *VR-2*. Virtual router names are case insensitive.

When a new virtual router is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added.

Example

The following command changes the virtual router context to *vr-acme*:

```
virtual-router vr-acme
```


Policy Manager Commands

12

This chapter describes commands for:

- Creating and configuring policy files for IP access lists (ACLs)
- Creating and configuring policy files for routing policies

Policies are a generalized category of features that impact forwarding and route forwarding decisions. Access policies are used primarily for security and quality of service (QoS) purposes.

IP access lists (also referred to as Access Lists or ACLs) consist of IP access rules and are used to perform packet filtering and forwarding decisions on traffic traversing the switch. Each packet on an interface is compared to the access list in sequential order and is either forwarded to a specified QoS profile or dropped. Additionally, packets can be metered using ACLs. Using access lists has no impact on switch performance.

Access lists are typically applied to traffic that crosses Layer 3 router boundaries, but it is possible to use access lists within a Layer 2 VLAN. NETGEAR products are capable of performing this function with no additional configuration.

Routing policies are used to control the advertisement or recognition of routes from routing protocols, such as RIP, OSPF, or BGP. Routing policies can be used to ‘hide’ entire networks or to trust only specific sources for routes or ranges of routes. The capabilities of routing policies are specific to the type of routing protocol involved, but are sometimes more efficient and easier to implement than access lists.

Note: Although the NETGEAR 8800 does not prohibit mixing ACL and routing type entries in a policy file, it is strongly recommended that you do not mix the entries, and you use separate policy files for ACL and routing policies.

check policy

```
check policy <policy-name> {access-list}
```

Description

Checks the syntax of the specified policy.

Syntax Description

policy-name	Specifies the policy to check.
access-list	Specifies that an access list specific check is performed.

Default

N/A.

Usage Guidelines

Use this command to check the policy syntax before applying it. If any errors are found, the line number and a description of the syntax error are displayed. A policy that contains syntax errors will not be applied.

This command can only determine if the syntax of the policy file is correct and can be loaded into the policy manager database. Since a policy can be used by multiple applications, a particular application may have additional constraints on allowable policies.

Example

The following example checks the syntax of the policy *zone5*:

```
check policy zone5
```

If no syntax errors are discovered, the following message is displayed:

```
Policy file check successful.
```

check policy attribute

```
check policy attribute {<attr>}
```

Description

Displays the syntax of the specified policy attribute.

Syntax Description

attr	Specifies the attribute check.
------	--------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the syntax of policy attributes. The command displays any additional keywords to use with this attribute, and the types of values expected.

Policy attributes are used in the rule entries that make up a policy file.

For each attribute, this command displays which applications use the attribute, and whether the attribute is a match condition or a set (action, action modifier) condition.

The current applications are:

- ACL—access-lists
- RT—routing profiles, route maps
- CLF—CLEAR-Flow

The syntax display does not show the text synonyms for numeric entries. For example, the `icmp-type` match condition allows you to specify either an integer or a text synonym for the condition. Specifying `icmp-type 8` or `icmp-type echo-request` are equivalent, but the syntax display shows only the numeric option.

Note: The syntax displayed is used by the policy manager to verify the syntax of policy files. The individual applications are responsible for implementing the individual attributes. Inclusion of a particular policy attribute in this command output does not imply that the attribute has been implemented by the application. See the documentation of the particular application for detailed lists of supported attributes.

Example

The following example displays the syntax of the policy attribute *icmp-type*:

```
check policy attribute icmp-type
```

The following is sample output for this command:

```
( match ) ( ACL )
icmp-type <uint32 val>
```

edit policy

```
edit policy <filename>
```

Description

Edits a policy text file.

Syntax Description

filename	Specifies the filename of the policy text file.
----------	---

Default

N/A.

Usage Guidelines

This command edits policy text files that are on the switch. All policy files use “.pol” as the filename extension, so to edit the text file for the policy *boundary* use `boundary.pol` as the filename. If you specify the name of a file that does not exist, you will be informed and the file will be created.

This command spawns a VI-like editor to edit the named file. For information on using VI, if you are not familiar with it, do a web search for “VI editor basic information”, and you should find many resources. The following is only a short introduction to the editor.

Edit operates in one of two modes; command and input. When a file first opens, you are in the command mode. To write in the file, use the keyboard arrow keys to position your cursor within the file, then press one of the following keys to enter input mode:

- `i` - To insert text ahead of the initial cursor position
- `a` - To append text after the initial cursor position

To escape the input mode and return to the command mode, press the Escape key.

There are several commands that can be used from the command mode. The following are the most commonly used:

- `dd` - To delete the current line
- `yy` - To copy the current line
- `p` - To paste the line copied
- `:w` - To write (save) the file
- `:q` - To quit the file if no changes were made
- `:q!` - To forcefully quit the file without saving changes
- `:wq` - To write and quit the file

Refresh Policy

After you have edited the text file for a policy that is currently active, you will need to refresh the policy if you want the changes to be reflected in the policy database. When you refresh the policy, the text file is read, the syntax is checked, the policy information is added to the policy manager database, and the policy then takes effect. Use the following command to refresh a policy:

```
refresh policy <policy-name>
```

If you just want to check to be sure the policy contains no syntax errors, use the following command:

```
check policy <policy-name> {access-list}
```

Example

The following command allows you to begin editing the text file for the policy *boundary*:

```
edit policy boundary.pol
```

refresh policy

```
refresh policy <policy-name>
```

Description

Refreshes the specified policy.

Syntax Description

policy-name	Specifies the policy to refresh.
-------------	----------------------------------

Default

N/A.

Usage Guidelines

Use this command when a new policy file for a currently active policy has been downloaded to the switch, or when the policy file for an active policy has been edited. This command reprocesses the text file and updates the policy database.

The policy manager uses Smart Refresh to update the ACLs. When a change is detected, only the ACL changes needed to modify the ACLs are sent to the hardware, and the unchanged entries remain. This behavior avoids having to blackhole packets because the ACLs have been momentarily cleared. Smart Refresh works well for minor changes, however, if the changes are too great, the refresh reverts to the earlier behavior. To take advantage of Smart Refresh, disable access-list refresh blackholing by using the command:

```
disable access-list refresh blackhole
```

If you attempt to refresh a policy that cannot take advantage of Smart Refresh, you will receive a message similar to the following if blackholing is enabled:

```
Incremental refresh is not possible given the configuration of policy <name>. Note, the
current setting for Access-list Refresh Blackhole is Enabled.
```

```
Would you like to perform a full refresh? (Yes/No) [No]:
```

and if blackholing is not enabled:

```
Incremental refresh is not possible given the configuration of policy <name>. Note, the
current setting for Access-list Refresh Blackhole is Disabled.
```

```
WARNING: If a full refresh is performed, it is possible packets that should be denied may be
forwarded through the switch during the time the access list is being installed.
```

```
Would you like to perform a full refresh? (Yes/No) [No]:
```

If you attempt to refresh a policy that is not currently active, you will receive an error message.

For an ACL policy, the command is rejected if there is a configuration error or hardware resources are not available.

Example

The following example refreshes the policy *zone5*:

```
refresh policy zone5
```

show policy

```
show policy {<policy-name> | detail}
```

Description

Displays the specified policy.

Syntax Description

policy-name	Specifies the policy to display.
detail	Show the policy in detail.

Default

If no policy name is specified, all policies are shown

Usage Guidelines

Use this command to display which clients are using the specified policy. The detail option displays the rules that make up the policy.

Example

The following example displays the policy *zone5*:

```
show policy zone5
```

This chapter describes commands for creating and configuring IP access lists (ACLs).

IP access lists (also referred to as Access Lists or ACLs) consist of IP access rules and are used to perform packet filtering and forwarding decisions on traffic traversing the switch. Each packet on an interface is compared to the access list in sequential order and is either forwarded to a specified QoS profile or dropped. Additionally, for the NETGEAR 8800 series switches, packets can be metered using ACLs. Using access lists has no impact on switch performance.

Access lists are typically applied to traffic that crosses Layer 3 router boundaries, but it is possible to use access lists within a Layer 2 VLAN. NETGEAR products are capable of performing this function with no additional configuration.

Note: Although the NETGEAR 8800 does not prohibit mixing ACL and routing type entries in a policy file, it is strongly recommended that you do not mix the entries, and you use separate policy files for ACL and routing policies.

clear access-list counter

```
clear access-list {dynamic} counter {<countername>} {any | ports <portlist> | vlan  
<vlanname>} {ingress | egress}
```

Description

Clears the specified access list counters.

Syntax Description

dynamic	Specifies that the counter is from a dynamic ACL.
countername	Specifies the ACL counter to clear.
any	Specifies the wildcard ACL.
portlist	Specifies to clear the counters on these ports.
vlanname	Specifies to clear the counters on the VLAN.

ingress	Clear the ACL counter for packets entering the switch on this interface.
egress	Clear the ACL counter for packets leaving the switch from this interface.

Default

The default direction is ingress; the default ACL type is non-dynamic.

Usage Guidelines

Use this command to clear the ACL counters. If you do not specify an interface, or the `any` option, you will clear all the counters.

Example

The following example clears all the counters of the ACL on port 2:1:

```
clear access-list counter port 2:1
```

The following example clears the counter *counter2* of the ACL on port 2:1

```
clear access-list counter counter2 port 2:1
```

clear access-list meter

```
clear access-list meter {<metername>} [any | ports <portlist> | vlan <vlannname>]
```

Description

Clears the specified access list meters.

Syntax Description

metername	Specifies the ACL meter to clear.
portlist	Specifies to clear the counters on these ports.
vlannname	Specifies to clear the counters on the VLAN.

Default

N/A.

Usage Guidelines

Use this command to clear the out-of-profile counters associated with the meter configuration.

Example

The following example clears all the out-of-profile counters for the meters of the ACL on port 2:1:

```
clear access-list meter port 2:1
```

The following example clears the out-of-profile counters for the meter *meter2* of the ACL on port 2:1

```
clear access-list meter meter2 port 2:1
```

configure access-list

```
configure access-list <aclname> [any | ports <portlist> | vlan <vlanname>] {ingress | egress}
```

Description

Configures an access list to the specified interface.

Syntax Description

policy-name	Specifies the ACL policy name. The name can be from 1-32 characters long.
aclname	Specifies the ACL name.
any	Specifies that this ACL is applied to all interfaces as the lowest precedence ACL.
portlist	Specifies the ingress port list on which the ACL is applied.
port_list	Specifies the egress port list.
vlanname	Specifies the VLAN on which the ACL is applied.
ingress	Apply the ACL to packets entering the switch on this interface.
egress	Apply the ACL to packets leaving the switch from this interface.

Default

The default direction is ingress.

Usage Guidelines

The access list applied in this command is contained in a text file created either externally to the switch or using the `edit policy` command. The file is transferred to the switch using TFTP before it is applied to the ports. The ACL name is the file name without its “.pol” extension. For example, the ACL *blocknetfour* would be in the file *blocknetfour.pol*. For more information on policy files, see the *NETGEAR 8800 User Manual*.

Specifying the keyword `any` applies the ACL to all the ports, and is referred to as the wildcard ACL. This ACL is evaluated for ports without a specific ACL applied to it, and is also applied to packets that do not match the ACL applied to the interface.

Example

The following command configures the ACL policy *test* to port 1:2 at ingress:

```
configure access-list test ports 1:2
```

The following command configures the ACL *mydefault* as the wildcard ACL:

```
configure access-list mydefault any
```

The following command configures the ACL policy *border* as the wildcard egress ACL:

```
configure access-list border any egress
```

configure access-list add

```
configure access-list add <dynamic_rule> [ [[first | last] {priority <p_number>} {zone <zone>}  
] | [[before | after] <rule>] | [ priority <p_number> {zone <zone>} ]] [ any | vlan  
<vlanname> | ports <portlist> ] {ingress | egress}
```

Description

Configures a dynamic ACL rule to the specified interface and sets the priority and zone for the ACL.

Syntax Description

dynamic_rule	Specifies a dynamic ACL rule.
first	Specifies that the new dynamic rule is to be added as the first rule.
last	Specifies that the new dynamic rule is to be added as the last rule.
zone	Specifies the ACL zone for the rule.
p_number	Specifies the priority number of the rule within a zone. The range is from 0 (highest priority) to 7 (lowest priority).
before <rule>	Specifies that the new dynamic rule is to be added before an existing dynamic rule.
after <rule>	Specifies that the new dynamic rule is to be added after an existing dynamic rule.
any	Specifies that this ACL is applied to all interfaces.
vlanname	Specifies the VLAN on which this ACL is applied.
portlist	Specifies the ports on which this ACL is applied.
ingress	Apply the ACL to packets entering the switch on this interface.
egress	Apply the ACL to packets leaving the switch from this interface.

Default

The default direction is ingress.

Usage Guidelines

The dynamic rule must first be created before it can be applied to an interface. Use the following command to create a dynamic rule:

```
create access-list <dynamic-rule> <conditions> <actions> {non-permanent}
```

When a dynamic ACL rule is applied to an interface, you will specify its precedence among any previously applied dynamic ACLs. All dynamic ACLs have a higher precedence than any ACLs applied through ACL policy files.

Specifying the keyword *any* applies the ACL to all the ports, and is referred to as the wildcard ACL. This ACL is evaluated for ports without a specific ACL applied to them, and is also applied to packets that do not match the ACL applied to the interface.

The *priority* keyword can be used to specify a sub-zone within an application's space. For example, to place ACLs into three sub-zones within the CLI application, you can use three priority numbers, such as 2, 4, and 7.

Configuring priority number *1* is the same as configuring *first* priority. Configuring priority number *8* is the same as configuring *last* priority.

Example

The following command applies the dynamic ACL *icmp-echo* as the first (highest precedence) dynamic ACL to port 1:2 at ingress:

```
configure access-list add icmp-echo first ports 1:2
```

The following command applies the dynamic ACL *udpacl* to port 1:2, with a higher precedence than rule *icmp-echo*:

```
configure access-list add udpacl before icmp-echo ports 1:2
```

configure access-list delete

```
configure access-list delete <ruleName> [ any | vlan <vlanname> | ports <portlist> | all ]
{ingress | egress}
```

Description

Removes a dynamic ACL rule from the specified interface.

Syntax Description

ruleName	Specifies a dynamic ACL rule name.
any	Deletes this ACL as the wildcard ACL.
vlanname	Specifies the VLAN on which this ACL is deleted.
portlist	Specifies the ports on which this ACL is deleted.
all	Deletes this ACL from all interfaces.
ingress	Deletes the ACL for packets entering the switch on this interface.

egress	Deletes the ACL for packets leaving the switch from this interface.
--------	---

Default

The default direction is ingress.

Usage Guidelines

Specifying the keyword `all` removes the ACL from all interfaces it is used on.

Example

The following command removes the dynamic ACL *icmp-echo* from the port 1:2:

```
configure access-list delete icmp-echo ports 1:2
```

configure access-list rule-compression port-counters

```
configure access-list rule-compression port-counters [shared | dedicated]
```

Description

Switches between ACL configuration modes.

Syntax Description

shared	Sharing is “on” for counter rules.
dedicated	Sharing is “off” for counter rules.

Default

Dedicated

Usage Guidelines

Use this command to switch between two ACL configuration modes. In the first mode, “port-counters shared”, similar port-based ACL rules with counters are allowed to share the same hardware entry. This uses less space but provides an inaccurate counter value. In the second mode, “port-counters dedicated”, similar port-based ACL rules with counters are not allowed to share the same hardware entry, thereby consuming more entries but providing a precise count.

Only ACLs that are entered after this command is entered are affected. The command does not affect any ACLs that are already configured.

To configure all ACLs in *shared* mode, `configure access-list rule-compression port-counters shared` must be entered before any ACLs are configured or have been saved in the configuration when a switch is booted.

This is a global setting for the switch; that is, the option does not support setting some ACL rules with shared counters and some with dedicated counters.

To view the results of the configuration use the `show access-list configuration` command.

Example

The following command configures ACL rules with counters to share the same hardware entry:

```
configure access-list rule-compression port-counters shared
```

configure access-list vlan-acl-precedence

```
configure access-list vlan-acl-precedence [dedicated | shared]
```

Description

Configures precedence mode for policy-file based ACLs that are applied on a VLAN.

Syntax Description

dedicated	Allocates exclusive precedence for VLAN-based ACLs.
shared	VLAN-based ACLs share the precedence with other ACLs.

Default

Shared

Usage Guidelines

The following feature applies to only policy-file based ACLs that are applied on a VLAN. Use this command to switch between two VLAN-based ACL configuration modes. In the shared `vlan-acl-precedence` mode, VLAN-based ACL rules share the same precedence with other types of ACL rules. This is the default mode and provides the same behavior as in the previous software releases. In the dedicated `vlan-acl-precedence` mode, VLAN-based ACL rules have different precedence compared to other types of ACL rules. The dedicated mode yields improved installation performance for VLAN-based access-lists but may affect hardware rule utilization in some configurations.

After configuring, you are prompted to reboot the system for the changes to take effect.

Example

The following command allocates exclusive precedence for VLAN-based static ACL rules:

```
configure access-list vlan-acl-precedence dedicated
```

configure access-list zone

```
configure access-list zone <name> zone-priority <number>
```

```
configure access-list zone <name> move-application <appl-name> to-zone <name>
application-priority <number>
```

```
configure access-list zone <name> {add} application <appl-name> application_priority <number>
```

```
configure access-list zone <name> delete application <appl-name>
```

Description

Configures the priority of a zone; moves an application from one zone to another at a specified priority; adds an application to a zone with a specified priority, or changes the priority of an application within a zone; deletes an application from a zone.

Syntax Description

name	Specifies a a zone name.
zone-priority <number>	Sets the priority of the zone.
move-application <appl-name>	Specifies the name of an application to be moved.
to-zone <name>	Specifies the zone to which the application is moved.
application-priority <number>	Sets the priority of the application within the zone. The range is from 0 (highest priority) to 7 (lowest priority).
add	Adds an application to a zone at a specified priority.
application <appl_name>	Specifies the application to be added to the zone.
application_priority <number>	Sets the priority of a new or existing application within a zone. The range is from 0 (highest priority) to 7 (lowest priority).

Default

N/A.

Usage Guidelines

To configure the priority of a specific zone, use the syntax:

```
configure access-list zone <name> zone-priority <number>
```

To move an application from one zone to another, and set its priority in the new zone, use the syntax:

```
configure access-list zone <name> move-application <appl-name> to-zone <name>
application-priority <number>
```

To add an application to a zone and specify its priority or to change the priority of an application within a zone, use the syntax:

```
configure access-list zone <name> {add} application <appl-name> application_priority <number>
```

To delete an application from a zone, use the syntax:

```
configure access-list zone <name> delete application <appl-name>
```

Example

The following command adds the CLI application to the zone *myzone* at a priority of 6:

```
configure access-list zone myzone add cli application-priority 6
```

configure flow-redirect add nexthop

```
configure flow-redirect <flow-redirect-name> add nexthop <ipaddress> priority <number>
```

Description

Adds a nexthop for the named flow redirection policy.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
ipaddress	Specifies the IP address of a new nexthop
number	Specifies the priority value for the nexthop.

Default

N/A.

Usage Guidelines

Use this command to add a new nexthop for the named flow redirection policy with a priority value. The priority value can range from a low of “1” to a high of “254.” The nexthop with the highest priority among multiple ones is preferred as the working nexthop. When each added nexthop has the same priority, the first one configured is preferred.

Example

The following command adds a nexthop *10.1.1.1* for the flow redirection policy *flow10* with a priority of *100*:

```
configure flow-redirect flow10 add nexthop 10.1.1.1 priority 100.
```

configure flow-redirect delete nexthop

```
configure flow-redirect <flow-redirect-name> delete nexthop <ipaddress>
```

Description

Deletes a nexthop for the named flow redirection policy.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
ip address	Specifies the IP address of the nexthop

Default

N/A.

Usage Guidelines

Use this command to delete a nexthop for the named flow redirection policy. If the deleted nexthop is the working nexthop for the policy-based routing entry, another is selected from the remaining active next hops, based on priority.

Example

The following command deletes the nexthop *10.1.1.1* from the flow redirection policy *flow10*:

```
configure flow-redirect flow10 delete nexthop 10.1.1.1
```

configure flow-redirect health-check

```
configure flow-redirect <flow-redirect-name> health-check [ping | arp]
```

Description

Configures health checking for a specific flow redirection policy.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
ping	Specifies ping health checking. This includes ARP.
arp	Specifies ARP health checking.

Default

Ping is the default

Usage Guidelines

Use this command to configure health checking for a specific named flow redirection policy. Ping includes ARP.

Example

The following command specifies *arp* health checking for the flow redirection policy *flow10*


```
configure flow-redirect flow10 health-check arp
```

configure flow-redirect nexthop

```
configure flow-redirect <flow-redirect-name> nexthop <ipaddress> ping interval <interval>
miss <miss>
```

Description

Configures the ping interval and miss count for a nexthop in the flow redirection policy.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
ip address	Specifies the IP address of the nexthop
interval	Specifies the number of seconds between pings. The default is "2".
miss	Specifies the number of misses allowed. The default is "2".

Default

N/A.

Usage Guidelines

Use this command to set a ping interval and miss count. When the ping response is not received with the interval * (miss + 1), the nexthop is considered to be dead and a new candidate is selected from the remaining active nexthops.

Example

The following command configures a ping interval of 3 and miss count of 3 for the nexthop 10.1.1.1 in the flow redirection policy *flow 3*:

```
configure flow-redirect flow3 nexthop 10.1.1.1 ping interval 3 miss 3
```

configure flow-redirect no-active

```
configure flow-redirect <flow-redirect-name> no-active [drop|forward]
```

Description

Configures packets to either follow the normal routing table or be dropped.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
drop	Specifies that the packets are to be dropped.

forward	Specifies that the packets are to follow the normal routing table.
---------	--

Default

The default is `forward`.

Usage Guidelines

Use this command to set a drop or forward configuration for packets to be applied when all configured next hops become unreachable.

Example

The following command configures packets of the flow redirection policy *flow3* to be dropped when all configured next hops become unreachable:

```
configure flow-redirect flow3 no-active drop
```

configure flow-redirect vr

```
configure flow-redirect <flow-redirect-name> vr <vr-name>
```

Description

Configures a virtual router for a flow redirection policy.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
vr-name	Specifies the name of the virtual router

Default

The default virtual router is *VR-Default*.

Usage Guidelines

Because ACLs do not recognize the virtual router concept, one policy-based routing can be used for multiple virtual routing entries when a VLAN-based virtual router is used for one port. This configuration of a VR into a flow-redirect makes a policy-based routing work for a specific VR.

Example

The following command configures virtual router *mgmt* for flow redirection policy *flow3*:

```
configure flow-redirect flow3 vr mgmt
```

create access-list

```
create access-list <dynamic-rule> <conditions> <actions> {non-permanent}
```

Description

Creates a dynamic ACL

Syntax Description

dynamic-rule	Specifies the dynamic ACL name. The name can be from 1-32 characters long.
conditions	Specifies the match conditions for the dynamic ACL.
actions	Specifies the actions for the dynamic ACLs.
non-permanent	Specifies that the ACL is not to be saved.

Default

By default, ACLs are permanent.

Usage Guidelines

This command creates a dynamic ACL rule. Use the `configure access-list add` command to apply the ACL to an interface.

The `conditions` parameter is a quoted string of match conditions, and the `actions` parameter is a quoted string of actions. Multiple match conditions or actions are separated by semi-colons. A complete listing of the match conditions and actions is in Chapter 17 in the *NETGEAR 8800 User Manual*.

Dynamic ACL rule names must be unique, but can be the same as used in a policy-file based ACL. Any dynamic rule counter names must be unique. For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

By default, ACL rules are saved when the `save` command is executed, and persist across system reboots. Configuring the optional keyword `non-permanent` means the ACL will not be saved.

Example

The following command creates a dynamic ACL that drops all ICMP echo-request packets on the interface:

```
create access-list icmp-echo "protocol icmp;icmp-type echo-request" "deny"
```

The created dynamic ACL will take effect after it has been configured on the interface. The previous example creates a dynamic ACL named `icmp-echo` that is equivalent to the following ACL policy file entry:

```
entry icmp-echo {
  if {
```

```

        protocol icmp;
        icmp-type echo-request;
    } then {
        deny;
    }
}

```

The following command creates a dynamic ACL that accepts all the UDP packets from the 10.203.134.0/24 subnet that are destined for the host 140.158.18.16, with source port 190 and a destination port in the range of 1200 to 1250:

```

create access-list udpacl "source-address 10.203.134.0/24;destination-address
140.158.18.16/32;protocol udp;source-port 190;destination-port 1200 - 1250;" "permit"

```

The previous example creates a dynamic ACL entry named *udpac1* that is equivalent to the following ACL policy file entry:

```

entry udpac1 {
    if {
        source-address 10.203.134.0/24;
        destination-address 140.158.18.16/32;
        protocol udp;
        source-port 190;
        destination-port 1200 - 1250;
    } then {
        permit;
    }
}

```

create access-list zone

```

create access-list zone <name> zone-priority <number>

```

Description

Creates a dynamic ACL zone, and sets the priority of the zone.

Syntax Description

name	Specifies the dynamic ACL zone name. The name can be from 1-32 characters long.
zone-priority <number>	Specifies priority of the zone. The range is from 1 (highest priority) to 4294967295 (lowest priority).

Default

The denial of service, system, and security zones are configured by default, and cannot be deleted.

Usage Guidelines

This command creates a dynamic ACL zone. You can configure the priority of the zone in relation to the default zones or to other configured zones.

Example

The following command creates a new zone, called myzone, with a priority of 2:

```
create access-list myzone zone-priority 2
```

create flow-redirect

```
create flow-redirect <flow-redirect-name>
```

Description

Creates a named flow redirection policy.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
--------------------	--

Default

N/A.

Usage Guidelines

Use this command to create a named flow redirection policy to which nexthop information can be added.

For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

Example

The following command creates a flow redirection policy names *flow3*:

```
create flow-redirect flow3
```

delete access-list

```
delete access-list <dynamic-rule>
```

Description

Deletes a dynamic ACL

Syntax Description

dynamic-rule	Specifies the dynamic ACL name.
--------------	---------------------------------

Default

N/A.

Usage Guidelines

This command deletes a dynamic ACL rule. Before you delete a dynamic ACL, it must be removed from any interfaces it is applied to. Use the `configure access-list delete` command to remove the ACL from an interface.

Example

The following command deletes the dynamic ACL *icmp-echo*:

```
delete access-list icmp-echo
```

delete access-list zone

```
delete access-list zone <name>
```

Description

Deletes an ACL zone.

Syntax Description

name	Specifies the zone name.
------	--------------------------

Default

N/A.

Usage Guidelines

This command deletes an ACL zone. You must remove all applications from a zone before you can delete the zone. To delete an application from a zone, use the command `configure access-list zone <name> delete application <appl-name>`

You cannot delete the default zones.

Example

The following command deletes the zone *my_zone*:

```
delete access-list zone my_zone
```

delete flow-redirect

```
delete flow-redirect <flow-redirect-name>
```

Description

Deletes the named flow redirection policy.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
--------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete a named flow-redirection policy. Before it can be deleted, all nexthop information must be deleted, otherwise an error message is displayed.

disable access-list permit to-cpu

```
disable access-list permit to-cpu
```

Description

Allows special packets to be blocked by low priority ACLs.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows ACLs to deny certain special packets from reaching the CPU, even if the packets match ACLs that would otherwise deny them. The special packets include STP BPDUs and ARP replies for the switch.

When this feature is disabled, these same packets will be denied if an ACL is applied that contains a matching entry that denies the packets. Contrary to expectations, the packets will still be denied if there is a higher precedence entry that permits the packets.

To enable this feature, use the following command:

```
enable access-list permit to-cpu
```

Example

The following command enables ACLs to deny STP BPDU packets from reaching the switch CPU:

```
disable access-list permit to-cpu
```

disable access-list refresh blackhole

```
disable access-list refresh blackhole
```

Description

Disables blackholing of packets during ACL refresh.

Syntax Description

This command has no arguments or variables.

Default

The feature is enabled.

Usage Guidelines

When access control lists (ACLs) are refreshed, this feature provides that any packets arriving during the refresh will be blackholed.

If you disable this feature, the ACLs will be refreshed as described in the [refresh policy](#) command.

To enable this feature, use the following command:

```
enable access-list refresh blackhole
```

Example

The following command disables dropping of packets during an ACL refresh:

```
disable access-list refresh blackhole
```

enable access-list permit to-cpu

```
enable access-list permit to-cpu
```

Description

Enables control packets to reach CPU, even if an ACL would deny them.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows control packets to reach the CPU, even if the packets match ACLs that would otherwise deny them. The control packets include STP BPDUs and ARP replies for the switch.

If this feature is disabled, these same packets will be denied if an ACL is applied that contains a matching entry that denies the packets. Contrary to expectations, when this feature is disabled, the packets will still be denied if there is a higher precedence entry that permits the packets.

To disable this feature, use the following command:

```
disable access-list permit to-cpu
```

Example

The following command enables STP BPDU packets to reach the switch CPU, despite any ACL:

```
enable access-list permit to-cpu
```

enable access-list refresh blackhole

```
enable access-list refresh blackhole
```

Description

Enables blackholing of packets during ACL refresh.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When access control lists (ACLs) are refreshed, this command provides that any packets arriving during the refresh will be blackholed. As the ACL is being refreshed, packets may arrive while the ACL is in an indeterminate state, and packets may be permitted that otherwise are dropped. This feature protects the switch during an ACL refresh.

To disable this feature, use the following command:

```
disable access-list refresh blackhole
```

Example

The following command enables dropping of packets during an ACL refresh:

```
enable access-list refresh blackhole
```

show access-list

```
show access-list {any | ports <portlist> | vlan <vlanname>} {ingress | egress}
```

Description

Displays the ACLs configured on an interface.

Syntax Description

aclname	Specifies the ACL name. The name can be from 1-32 characters long.
any	Specifies the wildcard ACL.
portlist	Specifies which ports' ACLs to display.
vlanname	Specifies which VLAN's ACL to display.
ingress	Display ingress ACLs.
egress	Display egress ACLs.

Default

The default is to display all interfaces, ingress.

Usage Guidelines

The ACL with the port and VLAN displayed as an asterisk (*) is the wildcard ACL.

If you do not specify an interface, the policy names for all the interfaces are displayed, except that dynamic ACL rule names are not displayed. To display dynamic ACLs use the following commands:

```
show access-list dynamic
show access-list dynamic rule <rule> {detail}
```

If you specify an interface, all the policy entries, and dynamic policy entries are displayed.

Example

The following command displays all the interfaces configured with an ACL:

```
show access-list
```

The output from this command is similar to the following:

Vlan Name	Port	Policy Name	Dir	Rules	Dyn Rules
=====					
*	3:6	TCP_flag	ingress	3	2
*	3:8	qos_hongkong	ingress	3	0
*	2:1	tc_2.4	ingress	4	0
*	2:7	tcp	ingress	1	0
v1	*	tcp	ingress	1	0

```
*          *          firewall1          ingress  2          1
```

The following command displays the ingress access list entries configured on the VLAN v1006:

```
show access-list v1006 ingress
```

The output from this command is similar to the following:

```
# RuleNo 1
entry dacl13 {          #Dynamic Entry
if match all {
    ethernet-destination-address 00:01:05:00:00:00 ;
} then {
    count c13 ;
    redirect 1.1.5.100 ;
} }

# RuleNo 2
entry dacl14 {          #Dynamic Entry
if match all {
    ethernet-source-address 00:01:05:00:00:00 ;
} then {
    count c14 ;
    qosprofile qp7 ;
} }

# RuleNo 3
entry dacl13 {
if match all {
    ethernet-destination-address 00:01:05:00:00:00 ;
} then {
    count c13 ;
    redirect 1.1.5.100 ;
} }
```

show access-list configuration

```
show access-list configuration
```

Description

Displays the ACL configuration.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

This command displays the state of the ACL configuration, set by the following commands:

```
enable access-list refresh blackhole
enable access-list permit to-cpu
configure access-list rule-compression port-counters
configure access-list vlan-acl-precedence
```

Example

The following command displays the state of the ACL configuration:

```
show access-list configuration
```

The output from this command is similar to the following:

```
Access-list Refresh Blackhole: Enabled
Access-list Permit To-CPU: Enabled

Access-list configured vlan-acl precedence mode: Dedicated or Shared
Access-list operational vlan-acl-precedence mode: Dedicated or Shared
Access-list Rule-compression Port-counters: Dedicated or Shared
```

show access-list counter

```
show access-list counter {<countername>} {any | ports <portlist> | vlan <vlanname>} {ingress
| egress}
```

Description

Displays the specified access list counters.

Syntax Description

countername	Specifies the ACL counter to display.
portlist	Specifies to display the counters on these ports.
vlanname	Specifies to display the counters on the VLAN.
ingress	Specifies to display ingress counters.
egress	Specifies to display egress counters.

Default

The default direction is ingress.

Usage Guidelines

Use this command to display the ACL counters.

Example

The following example displays all the counters for all ACLs:

```
show access-list counter
```

On a NETGEAR 8800 switch, the output of this command is similar to the following:

Policy Name	Vlan Name	Port	Direction	Counter Name	Packet Count	Byte Count
firewall11	*	*	ingress	DENY_SYN	0	0
				PERMIT_SYN_ACK	1228300404	1920048848
tc_2.4	*	2:1	ingress	arp192	3	204
				denyAll	0	0
				destIp	0	0
				destIp2	0	0
tcp	*	2:7	ingress	PERMIT_SYN_ACK	0	0
TCP_flag	*	3:6	ingress	denyAll	0	0
				ipArp	0	0
				tcpflags-syn	0	0
qos_hongkong	*	3:8	ingress	qp2cnt	0	0
				qp4cnt	0	0
				qp5cnt	0	0
tcp	v1	*	ingress	PERMIT_SYN_ACK	3759119344	2217044928

The following example displays all the counters for the ACL on port 2:1:

```
show access-list counter port 2:1
```

Policy Name	Vlan Name	Port	Direction	Counter Name	Packet Count	Byte Count
don1	*	2:1	ingress	source1111	0	0
				source2222	0	0

show access-list dynamic

```
show access-list dynamic
```

Description

Displays the names of existing dynamic ACLs and a count of how many times each is used.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

This command displays the names of existing dynamic ACLs, and how many times the ACL is used (bound to an interface).

To see the conditions and actions for a dynamic ACL, use the following command:

```
show access-list dynamic rule <rule> {detail}
```

Example

The following command displays names of all the dynamic ACLs:

```
show access-list dynamic
```

The following is sample output for this command:

```
Dynamic Rules:
Udpacl                      Bound to 1 interfaces
icmp-echo                   Bound to 1 interfaces
```

show access-list dynamic counter

```
show access-list dynamic counter {{<countername>}} any
| {{<countername>}} ports <portlist>
| {{<countername>}} vlan <vlanname>}
{ingress | egress}
```

Description

Displays the dynamic ACL counters.

Syntax Description

countername	Display the counter.
any	Specifies the wildcard ACL.
portlist	Specifies which ports' ACLs to display.
vlanname	Specifies which VLAN's ACL to display.
ingress	Display ingress ACLs.
egress	Display egress ACLs.

Default

The default is to display all interfaces, ingress.

Usage Guidelines

None.

Example

The following command displays all the dynamic ACL counters:

```
show access-list dynamic counter
```

show access-list dynamic rule

```
show access-list dynamic rule <rule> {detail}
```

Description

Displays the syntax of a dynamic ACL.

Syntax Description

rule	Specifies the rule to display.
detail	Specifies to display where the ACL has been applied.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the syntax of the dynamic ACL *udpac1*:

```
show access-list dynamic rule udpac1
```

The output of the command is similar to the following:

```
entry udpac1 {
if match all {
    source-address 10.203.134.0/24 ;
    destination-address 140.158.18.16/32 ;
    protocol udp ;
    source-port 190 ;
    destination-port 1200 - 1250 ;
} then {
```

```
    permit ;
} }
```

The following command displays where the dynamic ACL *udpac1* has been applied:

```
show access-list dynamic rule updacl
```

The output of the command is similar to the following:

```
Rule updacl has been applied to the following interfaces.
Vlan Name      Port    Direction
=====
*              1      ingress
```

show access-list interface

```
show access-list {rule <rule> {<start>} } [ any | port <port> | vlan <vlanname> ] {zone
<zone_name> { appl-name <appl_name> {priority <number> }}} {ingress | egress} {detail}
```

Description

Displays the specified ACL zones, including their priority, applications, and the application priorities.

Syntax Description

any	Displays all zones on the specified interface.
vlan <vlanname>	Displays all ACLs associated with the specified VLAN.
port <port>	Displays all ACLs associated with the specified ports.
zone <zone_name>	Specifies a zone to be displayed.
appl-name <appl_name>	Displays information by application within a zone.
priority <number>	Displays ACLs of the specified priority only, within an application area.
ingress	Displays ACLs applied to traffic in the ingress direction.
egress	Displays ACLs applied to traffic in the egress direction.
detail	Displays all ACLs applied to the specified interface.

Default

N/A.

Usage Guidelines

Use this command to display the ACL zones, applications, and priorities.

Specifying a zone will show all the ACLs installed in the particular zone. Specifying a priority within a zone will show all the ACLs installed at a particular priority within a zone.

Use the `detail` keyword to display all ACLs installed on a given interface.

Example

The following example displays the detailed view of the ACLs on port 1:1:

```
show access-list port 1:1 detail
```

The output of this command is similar to the following:

```
* BD-PC.1 # show access-list port 1:1 detail
RuleNo  Application    Zone          Sub Zone
=====
      1    CLI             myZone1
entry mac1 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:c1 ;
    destination-address 192.168.11.144/32 ;
} then {
    count mac1 ;
} }

      2    CLI             myZone5
entry mac51 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:51 ;
} then {
    count mack51;
} }

      3    CLI             myZone5
entry mac52 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:52 ;
} then {
    count mac52 ;
} }
```

The following example displays the detailed view of the priority 5 ACLs in the zone *myzone* on port 1:1:

```
* BD-PC.2 # show access-list port 1:1 zone myZone priority 5 detail
RuleNo  Application    Zone          Sub Zone
=====
      2    CLI             myZone5
entry mac51 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:51 ;
} then {
    count mack51;
} }
```

```

3      CLI          myZone5
entry mac52 {
if match all {
    ethernet-source-address 00:0c:29:e5:94:52 ;
} then {
    count mac52 ;
} }

```

The following example displays the priority 5 ACLs in the zone *myzone* on port 1:1:

```
BD-PC.2 # show access-list port 1:1 zone myZone priority 5
```

```
#Dynamic Entries ((*)- Rule is non-permanent )
```

RuleNo	Name			Application	Zone	Sub-Zone
1	mac51	CLI	myZone	5		
2	mac52	CLI	myZone	5		

show access-list usage acl-mask port

```
show access-list usage acl-mask port <port>
```

Description

Displays the number of ACL masks consumed by the ACLs on a particular port.

Syntax Description

port	Specifies to display the usage on this port.
------	--

Default

N/A.

Usage Guidelines

The NETGEAR 8800 switches have a total of 16 ACL masks per port on the switch. To avoid exhausting the masks available on the switch, you must carefully plan your use of ACL masks.

Use this command to display how many masks are currently consumed on a port.

Example

The following example displays the ACL mask usage on port 1:1:

```
show access-list usage acl-mask port 1:1
```

The output of this command is similar to the following:

Used: 3 Available: 12

show access-list usage acl-range port

`show access-list usage acl-range port <port>`

Description

Displays the number of Layer 4 port ranges consumed by the ACLs on the slices that support a particular port.

Syntax Description

port	Specifies to display the usage for the slices that support this port.
------	---

Default

N/A.

Usage Guidelines

The NETGEAR 8800 switches can support a total of 16 Layer 4 port ranges among the slices that support each group of 24 ports.

Use this command to display how many of these Layer 4 ranges are currently consumed by the ACLs on the slices that support a particular port. The output of this command also displays which ports share the same slices as the specified port.

Example

The following example displays the Layer 4 range usage on port 9:1:

```
show access-list usage acl-range port 9:1
```

The output of this command is similar to the following:

```
Ports 9:1-9:12, 9:25-9:36
L4 Port Ranges:  Used: 4  Available: 12
```

show access-list usage acl-rule port

`show access-list usage acl-rule port <port>`

Description

Displays the number of ACL rules consumed by the ACLs on a particular port or on the slices that support a particular port.

Syntax Description

port	Specifies to display the usage on this port.
------	--

Default

N/A.

Usage Guidelines

Use this command to display the rules used per slice, and also display the rule usage of the specified port.

The slice support for the NETGEAR 8800 series modules that use this mechanism is as follows:

- 8800 modules—
 - XCM888F—
 - Its 8 ports have 4 slices with each slice having enough memory for 128 egress rules.
 - Its 8 ports have 16 slices with each slice having enough memory for 256 ingress rules.
 - XCM8808X—
 - Each group of 2 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 2 ports has 16 slices with each slice having enough memory for 256 ingress rules.
 - XCM8848T/XCM8824F—
 - Each group of 24 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 24 ports has 16 slices with each slice having enough memory for 256 ingress rules.

Example

The following example displays the ACL rule usage on port 5:

```
show access-list usage acl-rule port 5
```

The following example displays the ACL ingress and egress rule usage on port 5:1.

```
show access-list usage acl-rule port 5:1
```

The output of this command on a NETGEAR 8806 series switch is similar to the following:

```
* (debug) BD-8806.5 # show access-list usage acl-rule port 5:1
Ports 5:1-5:48
Total Ingress/Egress Rules:
Used: 11 Available: 8181
Used: 1 Available: 1023
```

show access-list usage acl-slice port

```
show access-list usage acl-slice port <port>
```

Description

Displays the number of ACL slices and rules consumed by the ACLs on the slices that support a particular port.

Syntax Description

port	Specifies to display the usage for the slices that support this port.
------	---

Default

N/A.

Usage Guidelines

Use this command to display how many slices and how many rules per each slice are currently consumed by the ACLs on the slices that support a particular port. This command also displays which ports share the same slices as the specified port.

The slice support for the NETGEAR 8800 series modules that use this mechanism is as follows:

- 8800 modules—
 - XCM888F—
 - Its 8 ports have 4 slices with each slice having enough memory for 128 egress rules.
 - Its 8 ports have 16 slices with each slice having enough memory for 256 ingress rules.
 - XCM8848T/XCM8824F—
 - Each group of 24 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 24 ports has 16 slices with each slice having enough memory for 256 ingress rules.

Example

The following example displays the ACL slice usage on port 8:1:

```
show access-list usage acl-slice port 8:1
```

The output of this command is similar to the following:

```
Ports 8:1-8:12, 8:25-8:36
Slices:           Used: 3  Available: 5
Slice 5 Rules:    Used: 9  Available: 119
Slice 6 Rules:    Used: 1  Available: 127
Slice 7 Rules:    Used: 24 Available: 104
```

The following example displays the ACL ingress and egress slice usage on port 5:1:

```
show access-list usage acl-slice port 5:1
```

The output of this command on a NETGEAR 8806 series switch is similar to the following:

```
* (debug) BD-8806.6 # show access-list usage acl-slice port 5:1
Ports 5:1-5:48
Stage: INGRESS
Slices:          Used: 2  Available: 14
Slice 14 Rules:  Used: 8  Available: 504
Slice 15 Rules:  Used: 3  Available: 509
Stage: EGRESS
Slices:          Used: 1  Available: 3
Slice 3 Rules:   Used: 1  Available: 255
```

show flow-redirect

```
show flow-redirect <flow-redirect-name>
```

Description

Displays nexthop ipaddresses, up/down status, health-checking (ping/ARP) and ACL bindings.

Syntax Description

flow-redirect-name	Specifies the name of the flow redirection policy.
--------------------	--

Default

N/A.

Usage Guidelines

None

unconfigure access-list

```
unconfigure access-list <policy-name> {any | ports <portlist> | vlan <vlannname>} {ingress | egress}
```

Description

Removes a policy file ACL from the specified interface.

Syntax Description

policy-name	Specifies the ACL policy name. The name can be from 1-32 characters long.
aclname	Specifies the ACL name.

portlist	Specifies the ingress port list on which the ACL is applied.
port_list	Specifies the ports egress port list.
vlanname	Specifies the VLAN on which the ACL is applied.
ingress	Remove the ACL for packets entering the switch on this interface.
egress	Remove the ACL for packets leaving the switch from this interface.

Default

The default direction is ingress.

Usage Guidelines

This command removes ACLs that are contained in ACL policy files. To remove dynamic ACLs, use the following command:

```
configure access-list delete <ruleName> [ any | vlan <vlanname> | ports <portlist> | all ]
{ingress | egress}
```

To remove all non-dynamic ACLs from all interfaces, do not specify any ports or VLANs.

Example

The following command removes the ACL from port 1:2:

```
unconfigure access-list ports 1:2
```

The following command removes the ACLs from ports 1:2-6:3 and 7:1:

```
unconfigure access-list ports 1:2-6:3,7:1
```

The following command removes the wildcard ACL:

```
unconfigure access-list any
```

The following command removes all ACLs from all the interfaces, including the wildcard ACL:

```
unconfigure access-list
```

This chapter describes commands for:

- Configuring Quality of Service (QoS) profiles
- Creating traffic groupings and assigning the groups to QoS profiles
- Configuring, enabling, and disabling explicit class-of-service traffic groupings (802.1p and DiffServ)
- Configuring traffic grouping priorities
- Metering using ACLs
- Verifying configuration and performance
- Egress traffic rate limiting

For an introduction to QoS features, see the *NETGEAR 8800 User Manual*.

configure diffserv examination code-point qosprofile

The syntax is:

```
configure diffserv examination code-point <code_point> {qosprofile} <qosprofile>
```

Description

Configures the default ingress DiffServ code point (DSCP) to QoS profile mapping.

Syntax Description

code-point	Specifies a DiffServ code point (a 6-bit value in the IP-TOS byte in the IP header). Supported values are 0 to 63.
qosprofile	Specifies the QoS profile to which the DiffServ code point is mapped.

Default

See [Table 15](#).

Usage Guidelines

You can specify up to 64 different code points for each port. Code point values are grouped and assigned to the default QoS profiles as shown in [Table 15](#).

Table 15. Default DiffServ Code Point-to-QoS Profile Mapping

Code Point	NETGEAR 8800 Switches QoS Profile
0-7	QP1
8-15	QP1
16-23	QP1
24-31	QP1
32-39	QP1
40-47	QP1
48-55	QP1
56-63	QP8

Example

The following command specifies that code point 25 be assigned to QP2:

```
configure diffserv examination code-point 25 qosprofile qp2
```

configure diffserv replacement code-point

The syntax is:

```
configure diffserv replacement [{qosprofile} <qosprofile> | priority <priority>] code-point <code_point>
```

Description

Configures the egress Diffserv replacement mapping for either a QoS profile or an 802.1p priority value.

Syntax Description

qosprofile	Specifies a QoS profile.
value	Specifies an 802.1p priority value to map to a code point.
code_point	Specifies a 6-bit value to be used as the replacement DSCP in the DiffServ (IP-TOS byte) of the IP header.

Default

N/A.

Usage Guidelines

Note: NETGEAR recommends that you use the `qosprofile <qosprofile>` value to configure this parameter.

Egress packets contain the DSCP assigned to the QoS profile, which is can be selected by the 802.1p code point or by an ACL. The default 802.1p priority value to QoS profile to DSCP mapping is shown in [Table 16](#).

Table 16. Default QoS Profile-to-802.1p Priority Value-to-Code Point

802.1p Priority Value	NETGEAR 8800 Switches QoS Profile	DSCP
0	QP1	0
1	QP1	8
2	QP1	16
3	QP1	24
4	QP1	32
5	QP1	40
6	QP1	48
7	QP8	56

Example

The following command specifies that a code point value of 5 should be used to replace the DiffServ (TOS) bits in packets in QP2:

```
configure diffserv replacement qosprofile qp2 code-point 5
```

configure dot1p replacement

```
configure dot1p replacement {qosprofile} <qosprofile> priority <vpri> {ports <port_list>}
```

Description

Configures an 802.1p priority replacement configuration to override the 802.1p priority value configured for the specified QoS profile on the specified ports.

Syntax Description

qosprofile	Specifies a specific QoS profile. The value range is QP1 to QP8.
vpri	Specifies the 802.1p priority override value. The value is an integer between 0 and 7.
port_list	Specifies a list of slots and ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures *QP1* on slot 1, port 5 to replace the 802.1p value in egress frames with the value 4:

```
configure dot1p replacement QP1 priority 4 ports 1:5
```

configure dot1p type

The syntax is:

```
configure dot1p type <dot1p_priority> {qosprofile} <qosprofile>
```

Description

Configures an 802.1p priority to QoS profile mapping for the specified ports.

Syntax Description

dot1p_priority	Specifies the 802.1p priority value. The value is an integer between 0 and 7.
qosprofile	Specifies a specific QoS profile. The value range is QP1 to QP8.

Default

The default mapping of each 802.1p priority value to QoS profile is shown in [Table 17](#).

Table 17. Default 802.1p Priority Value-to-QoS Profile Mapping

802.1p Priority Value	NETGEAR 8800 Switches Default QoS Profile
0	QP1
1	QP1
2	QP1
3	QP1
4	QP1
5	QP1
6	QP1
7	QP8

Usage Guidelines

An 802.1p priority value seen on ingress can be mapped to a particular QoS profile and with specific bandwidth management and priority behavior.

You must create the QoS profile first, using the `create qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]` command, to map the 802.1p information to QoS profile 2 through 7.

Example

The following commands reassign (from the default) the QoS profiles associated with 802.1p priority values 1 and 2:

```
configure dot1p type 2 qosprofile qp2
configure dot1p type 1 qosprofile qp3
```

configure meter

```
configure meter <metername> {max-burst-size <burst-size> [Kb | Mb]} {committed-rate <cir> [Gbps | Mbps | Kbps]} {out-actions [drop | set-drop-precedence {dscp [none | <dscp-value>]}]}
```

Description

Configures an ACL meter to provide ingress traffic rate shaping on NETGEAR 8800 series switches. You can use this command to configure meters for ingress and egress rate limiting.

Syntax Description

metername	Specifies the ACL meter name.
max-burst-size	Specifies the maximum burst size or peak burst size in kilobytes (Kb) or megabytes (Mb).

committed-rate	Specifies the committed information rate in gigabits per second (Gbps), megabits per second (Mbps), or kilobits per second (Kbps).
out-actions	Specifies actions to take if traffic exceeds the profile.
drop	Specifies to drop out of profile traffic.
set-drop-precedence	Specifies to mark packet for high drop precedence.
dscp	Specifies to set DSCP.
none	Specifies to leave the DSCP value unchanged.

Default

By default, a newly committed meter has no maximum burst size, no committed rate, and a default action of drop.

Usage Guidelines

The meter configured with this command is associated with an ACL rule by specifying the meter name using the `meter` action modifier within the rule.

The `committed-rate` keyword specifies the traffic rate allowed for this meter, and the configured rate operates as described in [Table 18](#). The rate you specify is rounded up to the next granularity increment value (see [Table 18](#)). For example, if you configure a 1 Mbps committed rate for a platform with a 64Kbps granularity increment, this value falls between the increment values of 960 Kbps and 1024 Kbps, so the effective committed rate is set to 1024 Kbps. Also, note that some platforms listed in [Table 18](#) require an adjustment to the expected rate to calculate the configured rate.

Table 18. Rate Configuration Notes

Platform	Granularity	Notes
NETGEAR 8800 switches	64Kbps	Specify the traffic rate in Kbps, Mbps, or Gbps. The range is 64Kbps to 1 Gbps for GE ports and 1 Mbps to 10 Gbps for 10GE ports. Add 20 bytes per frame to the expected rate to determine the configured rate.

The `max-burst-size` keyword specifies the maximum number of consecutive bits that are allowed to be in-profile at wire-speed. The `max-burst-size` parameter can be specified in Kb, Mb, or Gb. The specified `max-burst-size` is rounded down to the nearest supported size. The `max-burst-size` range is 32Kb to 128Mb.

The keyword `out-actions` specifies the action that is taken when a packet is out-of-profile. The supported actions include dropping the packet, marking the drop precedence for the packet, or setting the DSCP value in the packet. The keyword `drop` indicates that any out-of-profile packet is immediately dropped. The keyword `set-drop-precedence` marks out-of-profile packets with high drop precedence. If the optional keyword `set-dscp` is specified, the DSCP value, as specified by the parameter `<dscp-value>`, is written into the

out-of-profile packet. Setting the DSCP value to `none` leaves the DSCP value in the packet unchanged.

Example

The following command configures the ACL meter *maximum_bandwidth*, assigns it a rate of 10 Mbps, and sets the out of profile action to `drop`:

```
configure meter maximum_bandwidth committed-rate 10 Mbps out-action drop
```

configure ports qosprofile

```
configure ports <port_list> {qosprofile} <qosprofile>
```

Description

Creates a port-based traffic group, which configures one or more ingress ports to use a particular egress QoS profile.

Syntax Description

<code>port_list</code>	Specifies a list of ports or slots and ports.
<code>qosprofile</code>	Specifies a QoS profile.

Default

All ingress ports have the default qosprofile of QP1.

Usage Guidelines

This command assigns traffic ingressing the specified port to a specified egress QoS profile. NETGEAR 8800 switches support eight egress QoS profiles (QP1 to QP8) for each port.

Example

The following command configures port 5 on slot 5 of the switch to use QoS profile QP3:

```
configure ports 5:5 qosprofile QP3
```

configure ports rate-limit egress

```
configure ports <port_list> rate-limit egress [no-limit | <cir-rate> [Kbps | Mbps | Gbps]
{max-burst-size <burst-size> [Kb | Mb]}]
```

Description

Configures an egress traffic rate limit for a port or groups of ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
no-limit	Specifies traffic be transmitted without limit; use to reconfigure or unconfigure previous rate-limiting parameters.
cir-rate	Specifies the desired rate limit in Kbps, Mbps, or Gbps.
max-burst-size	Specifies the maximum burst size or peak burst size in kilobits (Kb) or megabits (Mb).

Default

No-limit.

Usage Guidelines

Port speed limits the egress traffic, as follows:

- 1 Gbps port—64 Kbps increments
- 10 Gbps port—1 Mbps increments

If the specified egress limit (`cir-rate`) is not a multiple of 64 Kbps for a 1 Gbps port or 1 Mbps for a 10 Gbps port, the specified value is rounded down to the nearest appropriate multiple based on the port type.

Use the `no-limit` parameter to:

- Unconfigure egress rate limiting on the port(s)
- Reconfigure existing egress rate limiting on the port(s)

The `max-burst-size` parameter is the amount of traffic above the value in the `cir-rate` parameter that is allowed to burst from the port(s) for a short duration.

Example

The following command configures egress rate-limiting on slot 3 port 1 of the switch for 3 Mbps and a maximum burst size of 5 M bits:

```
configure port 3:1 rate-limit egress 3 Mbps max-burst-size 5 Mb
```

configure qosprofile

```
configure qosprofile <qosprofile> {committed_rate <committed_bps> [k | m]} {maxbuffer
<percent>} {maxbw <maxbw_number>} {minbw <minbw_number>} {peak_rate <peak_rate> [k | m]}
{use-strict-priority} {weight <value>}
```

Description

Modifies the rate-shaping parameters for QoS profiles on NETGEAR 8800 switches.

Syntax Description

qosprofile	Specifies a QoS profile name. Range is QP1 to QP8; the defaults are QP1 and QP8.
committed_rate	Specifies a committed information rate in Kbps (k) bits or Mbps (m).
maxbw	The maximum bandwidth (maxbw) option specifies the peak rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 100. When autonegotiation is off, the peak rate is the specified percentage of the configured port speed. When autonegotiation is on, the peak rate is the specified percentage of the maximum port speed (the switch does not detect the negotiated port speed).
minbw	The minimum bandwidth (minbw) option specifies the committed information rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 0. When autonegotiation is off, the CIR is the specified percentage of the configured port speed. When autonegotiation is on, the CIR is the specified percentage of the maximum port speed.
peak_rate	Specifies a peak rate in Kbps (k) bits or Mbps (m).
percent	Specifies the percentage of the total buffer you are reserving for this QoS profile. The range is 1 to 100; the default setting is 100.
use-strict-priority	When the global qosscheduler configuration (<code>configure qosscheduler</code> command) is set to <code>weighted-round-robin</code> , this option overrides the global configuration for the specified QoS profile, so that it operates in strict-priority-mode. This enables hybrid strict-priority and weighted-round-robin scheduling operation.
value	Specifies the weight value used for queue service weighting in the weighted-round-robin scheduler for this QoS profile. The range is 1 to 16; default is 1.

Default

- QoS profiles—QP1 through QP8 on NETGEAR 8800 series switches.
- Minimum bandwidth—0%
- Maximum bandwidth—100%
- Maximum buffer—100%
- Weight—1

Usage Guidelines

On NETGEAR 8800 series switches, QoS profiles QP1 through QP8 are preconfigured and cannot be deleted.

The `maxbuffer` parameter configures the maximum amount of packet buffer, by percentage, that the packets associated with the specified QoS profile can consume. Regardless of the setting for this parameter, the system does not drop any packets as long as packet buffer memory remains available and the current buffer use of the specified QoS profile is below the specified `maxbuffer` setting.

The `weight` parameter does not apply when the switch is configured for strict priority scheduling, which is the default configuration. To configure the type of scheduling you want to use for the entire switch, use the `configure qoscheduler [strict-priority | weighted-round-robin]` command.

The `weight` parameter configures the relative weighting for each QoS profile. Because each QoS profile has a default weight of 1, all QoS profiles have equal weighting. If you configure a QoS profile with a weight of 4, that specified QoS profile is serviced 4 times as frequently as the remaining QoS profiles, which still have a weight of 1. If you configure all QoS profiles with a weight of 16, each QoS profile is serviced equally but for a longer period.

When the switch is configured for weighted-round-robin mode, the `use-strict-priority` option overrides the switch configuration for the specified QoS profile on all ports. Among QoS profiles configured with the `use-strict-priority` option, QoS profile QP8 has the highest priority and QP1 has the lowest priority. All strict-priority QoS profiles are serviced first according to their priority level, and then all other QoS profiles are serviced based on their configured weight.

Note: If you specify `use-strict-priority`, lower-priority queues and weighted-round-robin queues are not serviced at all as long as higher-priority queues have any remaining packets.

Example

The following command configures the QoS profile parameters of QoS profile *qp1*:

```
configure qosprofile qp1 maxbuffer 75 weight 4
```

configure qosprofile egress

```
configure qosprofile {egress} <qosprofile> [{minbw <minbw_number>} {maxbw <maxbw_number>}] |
[{committed_rate <committed_bps> [K | M]} {peak_rate <peak_bps> [K | M]}] ] [ports
[<port_list> | all]]
```

Description

Modifies the default egress QoS profile parameters.

Syntax Description

<code>qosprofile</code>	Specifies a QoS profile name. Range is QP1 to QP8.
<code>minbw</code>	The minimum bandwidth (minbw) option specifies the committed information rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 0. When autonegotiation is off, the CIR is the specified percentage of the configured port speed. When autonegotiation is on, the CIR is the specified percentage of the maximum port speed.

maxbw	The maximum bandwidth (maxbw) option specifies the peak rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 100. When autonegotiation is off, the peak rate is the specified percentage of the configured port speed. When autonegotiation is on, the peak rate is the specified percentage of the maximum port speed (the switch does not detect the negotiated port speed).
committed_rate	Specifies a committed information rate in Kbps (k) bits or Mbps (m).
peak_rate	Specifies a peak rate in Kbps (k) bits or Mbps (m).
priority_number	Specifies a number that selects the service priority setting for the QoS profile. The accepted values are: <ul style="list-style-type: none"> • 1 (priority Low) • 2 (priority LowHi) • 3 (priority Normal) • 4 (priority NormalHi) • 5 (priority Medium) • 6 (priority MediumHi) • 7 (priority High) • 8 (priority HighHi)
port_list	Specifies a list of slots and ports to which the parameters apply. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
all	Specifies this applies to all ports on the device.

Default

- Minimum bandwidth—0%
- Maximum bandwidth—100%
- Priority—By default, each qosprofile is assigned a different priority level:
 - QP1 - 1, Low (the lowest priority)
 - QP2 - 2, LowHi
 - QP3 - 3, Normal
 - QP4 - 4, NormalHi
 - QP5 - 5, Medium
 - QP6 - 6, MediumHi
 - QP7 - 7, High
 - QP8 - 8, HighHi (highest priority)

Usage Guidelines

The maximum bandwidth value can be configured as either:

- an absolute percentage of the total maximum link speed, regardless of the currently configured or negotiated speed

OR

- an absolute peak rate in Mbps or Kbps

Example

The following command configures the egress QoS profile parameters of QoS profile *QP5* for specific ports on a NETGEAR 8800 series switch:

```
configure qosprofile egress qp5 minbw 10 maxbw 80 ports 5:5-5:7
```

configure qosprofile ingress

```
configure qosprofile ingress <iqp> [{committed_rate <committed_bps>
[k | m]} {maxbw <maxbw_number>} {minbw <minbw_number>} {peak_rate <peak_bps> [k | m]
{priority [<priority> | <priority_number>}}] ports [<port_list> | all]
```

Description

Sets the ingress rate shaping parameters, which is an ingress QoS profile.

Syntax Description

iqp	Specifies an ingress QoS profile: <ul style="list-style-type: none"> • for 1G I/O modules—iqp1 and iqp2 • for 10G I/O modules—iqp1 to iqp8
committed_rate	Specifies a committed information rate in Kbps (k) bits or Mbps (m).
maxbw	The maximum bandwidth (maxbw) option specifies the peak rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 100. When autonegotiation is off, the peak rate is the specified percentage of the configured port speed. When autonegotiation is on, the peak rate is the specified percentage of the maximum port speed (the switch does not detect the negotiated port speed).
minbw	The minimum bandwidth (minbw) option specifies the committed information rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 0. When autonegotiation is off, the CIR is the specified percentage of the configured port speed. When autonegotiation is on, the CIR is the specified percentage of the maximum port speed.
peak_rate	Specifies a peak rate in Kbps (k) bits or Mbps (m).
priority	Specifies a text service priority setting for the specified ingress QoS profile. The supported values are as follows: <ul style="list-style-type: none"> • 1G I/O module—2 queues and 2 priorities available; values are Low and LowHi. • 10G module—8 queues and 8 priorities available; values are Low, LowHi, Normal, NormalHi, Medium, MediumHi, High, and HighHi.
priority_number	Specifies a numerical service priority setting for the specified ingress QoS profile. The supported values are as follows: <ul style="list-style-type: none"> • 1G I/O module—2 queues and 2 priorities available; values are 1 (takes 1-4) or 2 (takes 5-8). • 10G module—8 queues and 8 priorities available; values are 1 to 8, with 8 being the highest priority.

port_list	Specifies a list of slots and ports to which the parameters apply. Specify ports using the following formats: 3-5, 2:5, 2:6-2:8.
all	Specifies this applies to all ports on the device.

Default

Disabled by default.

- Minimum bandwidth—0%
- Maximum bandwidth—100%
- Priority—By default, each qosprofile is assigned a different priority level, which varies by I/O module:
 - 1G I/O module:
 - IQP1 - 1, Low
 - IQP2 - 2, LowHi
 - 10G I/O module:
 - IQP1 - 1, Low
 - IQP2 - 2, LowHi
 - IQP3 - 3, Normal
 - IQP4 - 4, NormalHi
 - IQP5 - 5, Medium
 - IQP6 - 6, MediumHi
 - IQP7 - 7, High
 - IQP8 - 8, HighHi (highest priority)

Usage Guidelines

The number of ingress queues per port varies between the 1G I/O module and the 10G module.

On the 1G module, you have two ingress queues per port. The priority values of 1 to 4 map to the first queue, and the priority values of 5 to 8 map to the second queue.

On the 10G module, you have eight ingress queues per port. The priority values of 1 to 8 map one to each of the eight queues.

Example

The following command configures the ingress rate shaping parameters of QoS profile IQP3 for specified ports, using bandwidth percentages:

```
configure qosprofile ingress iqp3 minbw 27 maxbw 57 priority 4 ports 3:2
```

The following command configures the ingress rate shaping parameters for QoS profile IQP3 for all ports, using absolute values for committed rate and peak rate:

```
configure qosprofile ingress iqp3 committed-rate 64 k peak-rate 1000 k priority 4 ports all
```

configure qosscheduler

```
configure qosscheduler [strict-priority | weighted-round-robin]
```

Description

Specifies the method the switch uses to service QoS profiles.

Syntax Description

strict-priority	Specifies the switch services the higher-priority QoS profiles first.
weighted-round-robin	Specifies the switch services all QoS profiles based on the configured weighting for each QoS profile.

Default

Strict-priority.

Usage Guidelines

The configured QoS scheduling algorithm applies to all switch ports, but you can override this configuration for a QoS profile using the following command:

```
configure qosprofile <qosprofile> use-strict-priority
```

In strict-priority mode, QoS profile QP8 has the highest priority and QP1 has the lowest priority.

Note: If you specify strict-priority, lower-priority queues are not serviced at all as long as higher-priority queues have any remaining packets. If you specify weighted-round-robin, the switch services higher-weighted queues more frequently but continues to service lower-weighted queues (even when packets remain in the higher-weighted queues).

Example

The following command configures the switch for weighted-round-robin servicing:

```
configure qosscheduler weighted-round-robin
```

configure vlan qosprofile

```
configure vlan <vlan_name> {qosprofile} <qosprofile>
```

Description

Configures a VLAN traffic group, which links all the ingress ports in the specified VLAN to the specified egress QoS profile.

Syntax Description

vlan_name	Specifies a VLAN name.
qosprofile	Specifies an egress QoS profile. The supported values are: <code>qp1</code> to <code>qp8</code> and <code>none</code> .

Default

The default is none.

Usage Guidelines

The NETGEAR 8800 switches support eight egress QoS profiles (QP1 to QP8) for each port.

Example

The following command configures VLAN *accounting* to use QoS profile QP3:

```
configure vlan accounting qosprofile qp3
```

create meter

```
create meter <meter-name>
```

Description

This command creates a meter for ingress traffic rate limiting.

Syntax Description

meter-name	Specifies the meter name.
------------	---------------------------

Default

N/A.

Usage Guidelines

Meter names must begin with an alphabetical character and may contain alphanumeric characters and underscores (`_`), but they cannot contain spaces. The maximum allowed length for a name is 32 characters. For meter name guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

Example

The following command creates the meter *maximum_bandwidth*:

```
create meter maximum_bandwidth
```

create qosprofile

```
create qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]
```

Description

Creates a QoS profile.

Syntax Description

QP1....QP7	Specifies the QoS profile you want to create.
------------	---

Default

N/A.

Usage Guidelines

The NETGEAR 8800 series switches allow dynamic creation and deletion of QoS profiles QP2 to QP7. Creating a QoS profile dynamically does not cause loss of traffic.

QoS profiles QP1 and QP8 are part of the default configuration and cannot be deleted. You must create a QoS profile in the range of QP2 to QP7 before you can configure it or assign it to traffic groups.

Example

The following command creates QoS profile QP3:

```
create qosprofile qp3
```

delete meter

```
delete meter <metername>
```

Description

Deletes a meter.

Syntax Description

metername	Specifies the meter name.
-----------	---------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the meter *maximum_bandwidth*:

```
delete meter maximum_bandwidth
```

delete qosprofile

```
delete qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]
```

Description

Deletes a user-created QoS profile.

Syntax Description

QP1....QP7	Specifies the user-created QoS profile you want to delete.
------------	--

Default

N/A.

Usage Guidelines

You cannot delete the default QoS profiles of QP1 and QP8. If you attempt to delete QoS profile QP7, the system returns an error.

All configuration information associated with the specified QoS profile is removed.

Example

The following command deletes the user-created QoS profile QP3:

```
delete qosprofile qp3
```

delete traffic queue

```
delete traffic queue <queue_name>
```


Description

Deletes the specified traffic queue and removes all rate limiting resources associated with this queue from the hardware. This command does not delete any associated meters, which remain available for use with other traffic queues.

Syntax Description

queue_name	Specifies the traffic queue you are deleting.
------------	---

Default

N/A.

Usage Guidelines

Prior to deleting a traffic queue, you must remove all ACL policy file associations; you cannot delete a traffic queue that is currently associated with one or more ACL policy files.

When you delete any type of traffic queue, the associated meters are *not* deleted. Rather, those meters remain and can be associated with other traffic queues. To display the configured meters, issue the `show meters` command.

Example

The following command deletes the traffic queue named *test*:

```
delete traffic queue test
```

disable diffserv examination ports

```
disable diffserv examination ports [<port_list> | all]
```

Description

Disables the examination of the DiffServ field in an IP packet.

Syntax Description

port_list	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ examination should be disabled for all ports.

Default

Disabled.

Usage Guidelines

The diffserv examination feature is disabled by default.

Example

The following command disables DiffServ examination on the specified ports:

```
disable diffserv examination ports 5:3,5:5,6:6
```

disable diffserv replacement ports

```
disable diffserv replacement ports [<port_list> | all]
```

Description

Disables the replacement of DiffServ code points in packets transmitted by the switch.

Syntax Description

port_list	Specifies a list of ports or slots and ports on which Diffserv replacement will be disabled.
all	Specifies that DiffServ replacement should be disabled for all ports.

Default

N/A.

Usage Guidelines

The DiffServ replacement feature is disabled by default.

Note: The specified ports are the *ingress* ports.

Example

The following command disables DiffServ replacement on selected ports:

```
disable diffserv replacement ports 1:2,5:5,6:6
```

disable dot1p examination ports

```
disable dot1p examination ports [<port_list> | all]
```

Description

Prevents examination of the 802.1p priority field as part of the QoS configuration.

Syntax Description

port_list	Specifies a list of ports or slots and ports.
-----------	---

all	Specifies that dot1p replacement should be disabled for all ports.
-----	--

Default

Enabled.

Usage Guidelines

The 802.1p examination feature is enabled by default. To free ACL resources, disable this feature whenever another QoS traffic grouping is configured. (See [Chapter 13, “ACL Commands,”](#) for information on available ACL resources.)

Note: If you disable this feature when no other QoS traffic grouping is in effect, 802.1p priority enforcement of 802.1q tagged packets continues.

Example

The following command disables 802.1p value examination on ports 1 to 5:

```
disable dot1p examination ports 1-5
```

disable dot1p replacement ports

```
disable dot1p replacement ports [<port_list> | all]
```

Description

Disables the ability to overwrite 802.1p priority values for a given set of ports.

Syntax Description

port_list	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that 802.1p replacement should be disabled for all ports.

Default

N/A.

Usage Guidelines

The dot1p replacement feature is disabled by default.

On the 1 Gigabit Ethernet ports, 802.1p replacement always happens when you configure the DiffServ traffic grouping.

Note: The specified ports are ingress ports.

Example

The following command disables 802.1p value replacement on all ports:

```
disable dot1p replacement ports all
```

enable diffserv examination ports

```
enable diffserv examination ports [<port_list> | all]
```

Description

Enables the DiffServ field of an IP packet to be examined in order to select a QoS profile.

Syntax Description

port_list	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ examination is enabled for all ports.

Default

Disabled.

Usage Guidelines

The Diffserv examination feature is disabled by default.

If you are using DiffServ for QoS parameters, NETGEAR recommends that you also configure 802.1p or port-based QoS parameters to ensure that high-priority traffic is not dropped prior to reaching the MSM/MM on the switch.

Example

The following command enables DiffServ examination on selected ports:

```
enable diffserv examination ports 1:1,5:5,6:2
```

enable diffserv replacement ports

```
enable diffserv replacement ports [<port_list> | all]
```

Description

Enables the DiffServ code point to be overwritten in IP packets transmitted by the switch.

Syntax Description

port_list	Specifies a list of ingress ports or slots and ports on which to enable Diffserv replacement.
all	Specifies that DiffServ replacement should be enabled for all ports.

Default

N/A.

Usage Guidelines

The Diffserv replacement feature is disabled by default.

Note: The port in this command is the ingress port.
This command affects only that traffic in traffic groupings based on explicit packet class of service information and physical/logical configuration.

Example

The following command enables DiffServ replacement on specified ports:

```
enable diffserv replacement ports 5:3,5:5,6:2
```

enable dot1p examination ports

```
enable dot1p examination ports [<port_list> | all]
```

Description

Enables egress QoS profile selection based on the 802.1p bits in the incoming frame.

Syntax Description

port_list	Specifies a list of ports on which to enable the dot1p examination feature.
all	Specifies that dot1p examination should be enabled for all ports.

Default

Enabled.

Usage Guidelines

To increase available ACLs, you can disable the 802.1p examination feature if you are not running QoS or are running QoS using DiffServ. See *NETGEAR 8800 User Manual* for information on ACL limitations on these platforms.

Use this command to re-enable the 802.1p examination feature.

Example

The following command enables dot1p examination on ports 1 to 5:

```
enable dot1p examination ports 1-5
```

enable dot1p replacement ports

```
enable dot1p replacement ports [<port_list> | all]
```

Description

Allows the 802.1p priority field to be overwritten on egress according to the QoS profile to 802.1p priority mapping for a given set of ports.

Syntax Description

port_list	Specifies a list of ports or slots and ports.
all	Specifies that dot1p replacement should be enabled for all ports.

Default

N/A.

Usage Guidelines

The dot1p replacement feature is disabled by default.

By default, 802.1p priority information is not replaced or manipulated, and the information observed on ingress is preserved when transmitting the packet.

Note: The port in this command is the ingress port.

If 802.1p replacement is enabled, the 802.1p priority information that is transmitted is determined by the hardware queue that is used when transmitting the packet.

Note: This command affects only that traffic in traffic groupings based on explicit packet class of service information and physical/logical configuration.

On the 1 Gigabit Ethernet ports, 802.1p replacement always happens when you configure the DiffServ traffic grouping.

Example

The following command enables dot1p replacement on all ports:

```
enable dot1p replacement ports all
```

show access-list meter

```
show access-list meter {<metername>} [any | ports <portlist> | vlan <vlanname>]
```

Description

Displays the specified access list meter statistics and configurations.

Syntax Description

metername	Specifies the ACL meter to display.
portlist	Specifies to display the meters on these ports.
vlanname	Specifies to display the meters on the VLAN.

Default

N/A.

Usage Guidelines

Use this command to display the ACL meters.

Example

The following example displays access list meter information for port 7:1

```
Switch.8 # show access-list meter mtr1 port 7:1
Policy Name      Vlan Name      Port
                Committed  Committed Burst Peak Rate  Peak Burst
Out-of-Profile
  Meter          Rate (Kbps) Size (Kb)      (Kbps)      Size(kb)      Packet
Count
=====
```

```

irl1          *          7:1
a.    mtr1      10      20      10      20      0

```

show diffserv examination

The syntax is:

```
show diffserv examination
```

Description

Displays the DiffServ-to-QoS profile mapping.

Default

N/A.

Usage Guidelines

Once you alter the default mappings, the “->” in the display (shown below) becomes “* >”.

Examples

Because the NETGEAR 8800 series switches have 8 default QoS profiles, you see different displays depending on the platform.

The following is sample output from a NETGEAR 8800 switch:

```

show diffserv examination
CodePoint->QOSProfile mapping:
    00->QP1 01->QP1 02->QP1 03->QP1 04->QP1 05->QP1 06->QP1 07->QP1
    08->QP1 09->QP1 10->QP1 11->QP1 12->QP1 13->QP1 14->QP1 15->QP1
    16->QP1 17->QP1 18->QP1 19->QP1 20->QP1 21->QP1 22->QP1 23->QP1
    24->QP1 25->QP1 26->QP1 27->QP1 28->QP1 29->QP1 30->QP1 31->QP1
    32->QP1 33->QP1 34->QP1 35->QP1 36->QP1 37->QP1 38->QP1 39->QP1
    40->QP1 41->QP1 42->QP1 43->QP1 44->QP1 45->QP1 46->QP1 47->QP1
    48->QP1 49->QP1 50->QP1 51->QP1 52->QP1 53->QP1 54->QP1 55->QP1
    56->QP8 57->QP8 58->QP8 59->QP8 60->QP8 61->QP8 62->QP8 63->QP8

```

show diffserv replacement

The syntax is:

```
show diffserv replacement
```

Description

Displays the DiffServ replacement code-point values assigned to each QoS profile. These values are placed in egress packets when DiffServ replacement is enabled.

Default

N/A.

Usage Guidelines

Once you alter the default mappings, the “->” in the display (shown below) becomes “* >”.

Examples

The following is sample output from a NETGEAR 8810 switch:

```
show diffserv replacement
QOSProfile->CodePoint mapping:
    QP1->00
    QP8->56
```

show dot1p

The syntax is:

```
show dot1p
```

Description

Displays the 802.1p-to-QoS profile mappings.

Default

N/A.

Example

Following is sample output from the `show dot1p` command on the NETGEAR 8810 switch:

```
show dot1p
802.1p Priority Value      QOS Profile
      0                  QP1
      1                  QP1
      2                  QP1
      3                  QP1
      4                  QP1
      5                  QP1
      6                  QP1
      7                  QP8
```

show meter

```
show meter
```

Description

Displays the configured meters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays meters on the switch:

```
show meter
```

The following is sample output from this command:

```

-----
Name      Committed Rate(Kbps) Peak Rate(Kbps)
-----
peggy          1000000          --

```

Note: When using a NETGEAR 8800 series switch, you configure a peak rate for QoS meters using the `configure meter <metername> {max-burst-size <burst-size> [Kb | Mb]} {committed-rate <cir> [Gbps | Mbps | Kbps]} {out-actions [drop | set-drop-precedence {dscp [none | <dscp-value>}]}` command.

show ports congestion

```
show ports <port_list> congestion {no-refresh}
```

Description

Displays the port egress congestion statistics (dropped packets) for the specified ports on the front panel.

Syntax Description

port_list	Specifies one or more slots and ports.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.

Default

Displays the port congestion statistics for all ports in real-time.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, you can clear the counters or page up or down through the list of ports.

Note: If you are displaying congestion statistics in real time and another CLI session resets the counters for a port you are monitoring, the counters displayed in your session for that port are also reset.

If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

Note: Packets can be dropped at multiple locations along the path through the hardware. The per-port congestion counters count all dropped packets for all ports except the 10 GB ports. On the ports with hardware limitations, the dropped-packet counts are approximate and can be lower than the actual dropped packet counts.

If you do not specify a port number or range in the command, dropped packet counts are displayed for all ports.

Note: To display the congestion statistics for the QoS profiles on a port, use the `show ports <port_list> qosmonitor {congestion} {no-refresh}` command.

Examples

The following example shows the packets dropped due to congestion for all ports in real time:

```
BD-8810.1 # show ports congestion
Port Congestion Monitor                                     Tue May 27 13:02:37 2008
Port      Link      Packet
          State     Drop
=====
1:1       R         0
1:2       R         0
1:3       A        96
1:4       R         0
```

```

2:1      R      0
2:2      A     28513
2:3      R      0
2:4      R      0
2:5      R      0
2:6      R      0
2:7      R      0
2:8      R      0
3:1      R      0
3:2      R      0
3:3      R      0
3:4      R      0

```

```

=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->clear counters  U->page up  D->page down  ESC->exit

```

The following example shows a snapshot display of the packets dropped due to congestion for all ports:

```
BD-8810.1 # show ports congestion no-refresh
```

```

Port      Link      Packet
          State     Drop

```

```

=====
1:1      R      0
1:2      R      0
1:3      A     96
1:4      R      0
2:1      R      0
2:2      A     28513
2:3      R      0
2:4      R      0
2:5      R      0
2:6      R      0
2:7      R      0
2:8      R      0
3:1      R      0
3:2      R      0
3:3      R      0
3:4      R      0
5:1      R      0

```

```

=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback

```

show ports qosmonitor {congestion}

```
show ports <port_list> qosmonitor {congestion} {no-refresh}
```

Description

Displays egress packet counts or dropped-traffic counts for each QoS profile on the specified ports.

Syntax Description

<code>port_list</code>	Specifies one or more slots and ports.
<code>congestion</code>	Specifies the display of packets dropped at ingress due to port congestion.
<code>no-refresh</code>	Specifies a static snapshot of data instead of the default dynamic display.

Default

Displays egress packet counts in real-time.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, the spacebar toggles the display between egress packet counts and ingress dropped-packet counts.

Note: This command does not work properly if another CLI session is displaying congestion statistics in real time.

If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

Note: Packets can be dropped at multiple locations along the path through the hardware. Due to hardware limitations, the dropped-packet counters for QoS profiles cannot count dropped packets from all possible locations. Because of these limitations, the sum of all dropped packets for all QoS profiles can be less than the per port count displayed with the command: `show ports <port_list> congestion {no-refresh}`.

You can display packet counts for one port per slot or module at a time. You can simultaneously display packet counts for multiple ports, but they must be from different slots or modules. The dropped packet display is limited to the 8 most-significant digits.

When you display the packet counts for a port, this action configures the hardware to monitor that port. If the slot or module hardware was previously configured to monitor a different port,

the counters are reset for the new port. If the selected port is the last port displayed on the module, the counters are not reset.

Examples

The following example shows the egress packet counts for the specified ports:

```
# show ports 2:1, 3:6 qosmonitor
Qos Monitor Req Summary                               Thu Mar  2 10:58:23 2006
Port      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
          Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt
          Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts
=====
2:1        0        0        0        0        0        0        0        0
3:6        0        0        0        0        0        0        0        0
=====
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down ESC->exit
```

The next example shows the dropped packet counts for the specified ports:

```
# show ports 2:1, 3:6 qosmonitor congestion
QoS Monitor Req Summary                               Thu Jun 12 01:17:14 2008
Port      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
          Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt
          Cong      Cong      Cong      Cong      Cong      Cong      Cong      Cong
=====
2:1        0        0        0        0        0        0        0        0
3:6      8745        0       129        0        0        0        0        0
=====
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down ESC->exit
```

show qosprofile

```
show qosprofile {ingress} ports [ all | <port_list>]
```

Description

Displays QoS information on the switch.

Syntax Description

ingress	Specifies ingress queues.
ports	Specifies to display information for specified ports.
port_list	Specifies a list of slots and ports.
all	Specifies all ports.

Default

Displays egress QoS information for all ports.

Usage Guidelines

The displayed QoS profile information differs depending on the platform you are running on. The following section shows examples for different platforms.

Example

The display varies depending on your platform.

The following shows the information that appears when you omit the optional `port` parameter:

```
BD-8810Rack3.3 # show qosprofile
QP1    Weight = 1      Max Buffer Percent = 100
QP2    Weight = 1      Max Buffer Percent = 100
QP8    Weight = 1      Max Buffer Percent = 100
```

The following example shows how the display appears when the switch is configured for weighted-round-robin mode and some QoS profiles are configured for strict priority mode:

```
BD-8810.7 # show qosprofile
QP1    Weight = 1      Max Buffer Percent = 100
QP2    Weight = 1      Max Buffer Percent = 100
QP3    Weight = 1      Max Buffer Percent = 100
QP5    Strict-Priority Max Buffer Percent = 100
QP8    Strict-Priority Max Buffer Percent = 100
```

When you add the optional `port` parameter, the switch displays the following sample output:

```
BD-8810Rack3.6 # show qosprofile port 8:1
```

```
Port: 8:1
QP1    MinBw =      0% MaxBw =    100%
QP2    MinBw =      0% MaxBw =    100%
QP8    MinBw =      0% MaxBw =    100%
```

unconfigure diffserv examination

The syntax is:

```
unconfigure diffserv examination
```

Description

Disables DiffServ traffic groups.

Default

Disabled.

Example

The following command disables DiffServ code point examination:

```
unconfigure diffserv examination
```

unconfigure diffserv replacement

The syntax is:

```
unconfigure diffserv replacement
```

Description

Resets all DiffServ replacement mappings to the default values.

Default

The default code point to QoS profile mappings are shown in [Table 19](#).

Table 19. NETGEAR 8800 Series Switch Default DiffServ Code Point-to-QoS Profile Mapping

Code Point	NETGEAR 8800 Series Switch QoS profile
0-55	QP1
56-63	QP8

Example

The following command resets the DiffServ replacement mappings to their default values:

```
unconfigure diffserv examination
```

unconfigure qosprofile

```
unconfigure qosprofile {ingress | egress} {ports [<port_list>|all]}
```

Description

Returns the rate-shaping parameters for all QoS profiles on the specified ports to the default values.

Syntax Description

ingress	Specifies all ingress QoS profiles for the specified ports. If you do not specify <i>ingress</i> , the command returns all egress QoS profile values to the default values.
egress	Specifies all egress QoS profiles for the specified ports.

port_list	Specifies the ports on which to unconfigure QoS profiles.
all	Specifies that this command applies to all ports on the device.

Default

The default values for egress bandwidth on all supported platforms are:

- Minimum bandwidth—0%
- Maximum bandwidth—100%

The default values for egress priority and ingress QoS profiles differ by platform as described in the following sections.

The platform-specific default values for the two default egress QoS profiles (QP1 and QP8) on the NETGEAR 8800 series switches are:

- Maximum buffer—100%
- Weight—1

Usage Guidelines

None.

Example

The following command resets the QoS profiles for all ports to default settings:

```
unconfigure qosprofile
```

This chapter describes commands for:

- Managing the switch using SSH2
- Configuring switch user authentication through a RADIUS client
- Configuring switch user authentication through TACACS+
- Protecting the switch from Denial of Service attacks

SSH

Secure Shell 2 (SSH2) is a feature of the NETGEAR 8800 that allows you to encrypt session data between a network administrator using SSH2 client software and the switch. Configuration and policy files may also be transferred to the switch using the Secure Copy Program 2 (SCP2).

SSL

Secure Socket Layer (SSL) allows users to connect using a more secure HTTPS connection.

Note: If you cannot find SSH or SSL commands, your image probably did not come with SSH or SSL preinstalled. To download and install the SSH/SSL module, go to <http://kbserver.netgear.com/products/xcm8806.asp> or <http://kbserver.netgear.com/products/xcm8810.asp>.

User Authentication

Remote Authentication Dial In User Service (RADIUS, RFC 2138) is a mechanism for authenticating and centrally administrating access to network nodes. The NETGEAR 8800 RADIUS client implementation allows authentication for SSH2, Telnet or console access to the switch.

NETGEAR 8800 switches are also capable of sending RADIUS accounting information. You can configure RADIUS accounting servers to be the same as the authentication servers, but this is not required.

Terminal Access Controller Access Control System Plus (TACACS+) is a mechanism for providing authentication, authorization, and accounting on a centralized server, similar in function to the RADIUS client. The NETGEAR 8800 version of TACACS+ is used to authenticate prospective users who are attempting to administer the switch. TACACS+ is used to communicate between the switch and an authentication database.

Note: You cannot use RADIUS and TACACS+ at the same time.

Denial of Service

You can configure the NETGEAR 8800 to protect your NETGEAR switches in the event of a denial of service attack. During a typical denial of service attack, the CPU on the switch gets flooded with packets from multiple attackers, potentially causing the switch to fail. To protect against this type of attack, you can configure the software so that when the number of packets received is more than the configured threshold limit of packets per second, a hardware ACL is enabled.

clear ip-security anomaly-protection notify cache

```
clear ip-security anomaly-protection notify cache {slot [<slot> | all ]}
```

Description

Clear the local protocol anomaly event cache.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

N/A.

Usage Guidelines

This command clears the local protocol anomaly event cache.

clear ip-security arp validation violations

```
clear ip-security arp validation violations
```

Description

Clear the violation counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command clears the ARP validation violation counters.

clear ip-security dhcp-snooping entries

```
clear ip-security dhcp-snooping entries { vlan } <vlan_name>
```

Description

Clears the DHCP binding entries present on a VLAN.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server.
-----------	--

Default

N/A.

Usage Guidelines

Use this command to clear the DHCP binding entries present on a VLAN. When an entry is deleted, all its associated entries (such as source IP lockdown, secured ARP, and so on) and their associated ACLs, if any, are also deleted.

Example

The following command clears the DHCP binding entry temporary from the VLAN:

```
clear ip-security dhcp-snooping entries temporary
```

clear ip-security source-ip-lockdown entries ports

```
clear ip-security source-ip-lockdown entries ports [ <ports> | all ]
```

Description

Clears locked-down source IP addresses on a per-port basis.

Syntax Description

ports	Specifies the port or ports to be cleared.
all	Specifies that all ports are to be cleared.

Default

N/A.

Usage Guidelines

Use this command to clear locked-down source IP addresses on a per port basis. This command deletes the entries on the indicated ports and clears the associated ACLs.

clear vlan dhcp-address-allocation

```
clear vlan <vlan_name> dhcp-address-allocation [[all {offered | assigned | declined |
expired}] | <ipaddress>]
```

Description

Removes addresses from the DHCP allocation table.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server.
all	Specifies all IP addresses, or all IP addresses in a particular state.
offered	Specifies IP addresses offered to clients.
assigned	Specifies IP addresses offered to and accepted by clients.
declined	Specifies IP addresses declined by clients
expired	Specifies IP addresses whose lease has expired and not renewed by the DHCP server.
ipaddress	Specifies a particular IP address.

Default

N/A.

Usage Guidelines

You can delete either a single entry, using the IP address, or all entries. If you use the `all` option, you can additionally delete entries in a specific state.

Example

The following command removes all the declined IP addresses by hosts on the VLAN *temporary*:

```
clear vlan temporary dhcp-address-allocation all declined
```

configure dos-protect acl-expire

```
configure dos-protect acl-expire <seconds>
```

Description

Configures the denial of service protection ACL expiration time.

Syntax Description

seconds	Specifies how long the ACL is in place.
---------	---

Default

The default is 5 seconds.

Usage Guidelines

This command configures how long the DoS protection ACL remains in place.

Example

This example sets the ACL expiration time to 15 seconds:

```
configure dos-protect acl-expire 15
```

configure dos-protect interval

```
configure dos-protect interval <seconds>
```

Description

Configures the denial of service protection interval.

Syntax Description

seconds	Specifies how often the DoS protection counter is monitored.
---------	--

Default

The default is one second.

Usage Guidelines

This command configures how often the DoS protection counter is monitored.

Example

This example sets the interval to 5 seconds:

```
configure dos-protect interval 5
```

configure dos-protect trusted ports

```
configure dos-protect trusted-ports
[ports [<ports> | all]
| add-ports [<ports-to-add> | all]
| delete-ports [<ports-to-delete> | all]
]
```

Description

Configures the list of trusted ports.

Syntax Description

ports	Specifies the trusted ports list.
ports-to-add	Specifies the ports to add to the trusted ports list.
all	Specifies all the ports.
ports-to-delete	Specifies the ports to delete from the trusted ports list.

Default

N/A.

Usage Guidelines

Traffic from trusted ports will be ignored when DoS protect counts the packets to the CPU. If we know that a machine connected to a certain port on the switch is a safe "trusted" machine, and we know that we will not get a DoS attack from that machine, the port where this machine is connected to can be configured as a trusted port, even though a large amount of traffic is going through this port.

Example

This example sets the trusted port list to 3:1-3:7:

```
configure dos-protect trusted-ports ports 3:1-3:7
```

This example adds the trusted port 3:8 to the current list (use this command with a network administrator machine not connected to the internet that is attached to port 3:8):

```
configure dos-protect trusted-ports add-ports 3:8
```

configure dos-protect type l3-protect alert-threshold

```
configure dos-protect type l3-protect alert-threshold <packets>
```

Description

Configures the denial of service protection alert threshold.

Syntax Description

packets	Specifies how many packets in an interval will cause an alert.
---------	--

Default

The default is 4000 packets.

Usage Guidelines

This command configures how many packets received in an interval will cause a DoS protection alert. When an alert occurs, the packets are analyzed, and a temporary ACL is applied to the switch.

Example

This example sets the alert threshold to 8000 packets:

```
configure dos-protect type l3-protect alert-threshold 8000
```

configure dos-protect type l3-protect notify-threshold

```
configure dos-protect type l3-protect notify-threshold <packets>
```

Description

Configures the denial of service protection notification threshold.

Syntax Description

packets	Specifies how many packets in an interval will cause a notification.
---------	--

Default

The default is 3500 packets.

Usage Guidelines

This command configures how many packets received in an interval will cause a DoS protection notification.

Example

This example sets the notification threshold to 7500 packets:

```
configure dos-protect type l3-protect notify-threshold 7500
```

configure ip-security anomaly-protection icmp ipv4-max-size

```
configure ip-security anomaly-protection icmp ipv4-max-size <size> {slot [ <slot> | all ]}
```

Description

Configures the maximum IPv4 ICMP allowed size.

Syntax Description

size	Specifies the size of the IPv4 ICMP in bytes.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default size is 512 bytes.

Usage Guidelines

This command configures the IPv4 ICMP allowed size. The absolute maximum is 1023 bytes.

configure ip-security anomaly-protection icmp ipv6-max-size

```
configure ip-security anomaly-protection icmp ipv6-max-size <size> {slot [ <slot> | all ]}
```

Description

Configures the maximum ipv6 ICMP allowed size.

Syntax Description

size	Specifies the size of the IPv6 ICMP in bytes.
slot	Specifies the slot to be used.

all	Specifies all IP addresses, or all IP addresses in a particular state.
-----	--

Default

The default size is 512 bytes.

Usage Guidelines

This command configures the IPv6 ICMP allowed size. The absolute maximum is 16K bytes.

You can use this command to configure the maximum IPv6 ICMP packet size for detecting IPv6 ICMP anomalies. If the next header in the IPv6 ICMP packet is not 0x3A:ICMP, this anomaly is not detected. For example, an IPv6 ICMP packet with packet header 0x2c: Fragment Header is not detected.

configure ip-security anomaly-protection notify cache

```
configure ip-security anomaly-protection notify cache <size> {slot [<slot> | all ]}
```

Description

Configures the size of local notification cache.

Syntax Description

size	Specifies the size of the local notification cache.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1000 events.

Usage Guidelines

This command configures the size of local notification cache. Cached events are stored in local memory. The range is between 1 and 1000 events per second. If the cache is full, newer events replace older events.

configure ip-security anomaly-protection notify rate limit

```
configure ip-security anomaly-protection notify rate limit <value> {slot [<slot> | all ]}
```

Description

Configures the rate limiting for protocol anomaly notification.

Syntax Description

value	Specifies the period of the rate limit.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 10 events per second.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify rate window` that configures the rate limiting for protocol anomaly notification. When the anomaly notification is enabled, in order to avoid overloading CPU, the system generates only the number of limited notifications in a period of window seconds. The range is from 1 to 100 events.

configure ip-security anomaly-protection notify rate window

```
configure ip-security anomaly-protection notify rate window <value> {slot [<slot> | all ]}
```

Description

Configures the rate limiting for protocol anomaly notification.

Syntax Description

value	Specifies the period of the rate limit.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1 second.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify rate limit` that configures the rate limiting for protocol anomaly notification. When the anomaly notification is enabled, in order to avoid overloading CPU, the system generates only the number of limited notifications in a period of window seconds. The range is between 1 and 300 seconds.

configure ip-security anomaly-protection notify trigger off

```
configure ip-security anomaly-protection notify trigger off <value> { slot [<slot> | all ] }
```

Description

Configures an anomaly rate-based notification feature.

Syntax Description

value	Specifies the number of events for the trigger.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify trigger on` that configures an anomaly rate-based notification feature. The anomaly notification is automatically triggered if the rate of anomaly events is greater than the configured ON value, and the notification is disabled if the rate falls below the value set in the `configure ip-security anomaly-protection notify trigger off` command.

The command takes effects after the anomaly notification is enabled.

Note: The value set in ON must be greater than or equal to the value set in OFF.

configure ip-security anomaly-protection notify trigger on

```
configure ip-security anomaly-protection notify trigger on <value> { slot [<slot> | all ] }
```

Description

Configures an anomaly rate-based notification feature.

Syntax Description

value	Specifies the number of events for the trigger.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify trigger off` that configures an anomaly rate-based notification feature. The anomaly notification is automatically triggered if the rate of anomaly events is greater than the configured ON value, and the notification is disabled if the rate falls below the value set in the `configure ip-security anomaly-protection notify trigger off` command.

The command takes effects after the anomaly notification is enabled.

Note: The value set in ON must be greater than or equal to the value set in OFF.

configure ip-security anomaly-protection tcp

```
configure ip-security anomaly-protection tcp min-header-size <size> {slot [ <slot> | all ]}
```

Description

Configures the minimum TCP header allowed.

Syntax Description

size	Specifies the size of the header in bytes.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default value is 20 bytes.

Usage Guidelines

This command configures the minimum TCP header allowed. It takes effect for both IPv4 and IPv6 TCP packets.

The range of the minimum TCP header may be between 8 and 255 bytes.

configure ip-security dhcp-snooping information check

```
configure ip-security dhcp-snooping information check
```

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking in the server-originated packets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command enables the checking of the server-originated packets for the presence of option 82. In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to prevent DHCP reply packets with invalid or missing relay agent options from being forwarded to the client. With checking enabled, the following checks and actions are performed:

- When the option 82 is present in the packet, the MAC address specified in the remote-ID sub-option is the switch system MAC address. If the check fails, the packet is dropped.
- When option 82 is not present in the packet, the DHCP packet is forwarded with no modification.

To disable this check, use the following command:

```
unconfigure ip-security dhcp-snooping information check
```

Example

The following command enables DHCP relay agent option checking:

```
configure ip-security dhcp-snooping information check
```

configure ip-security dhcp-snooping information circuit-id port-information port

```
configure ip-security dhcp-snooping information circuit-id port-information <port_info> port <port>
```

Description

Configures the port information portion of the circuit ID.

Syntax Description

port_info	Specifies the circuit ID port information in the format of <VLAN Info> - <Port Info>; maximum length is 32 bytes.
port	Specifies the port for which DHCP Snooping should be enabled.

Default

The default value is the ASCII representation of the ingress port's SNMP ifIndex.

Usage Guidelines

This command allows you to configure the port information portion of the circuit ID whose format is <vlan info> - <port info> for each port. The parameter <port info> is a string of up to 32 bytes in length. When a specific value is not configured for port information, the `port_info` defaults to the ASCII representation of the ingress ports's SNMP ifIndex.

configure ip-security dhcp-snooping information circuit-id vlan-information

```
configure ip-security dhcp-snooping information circuit-id vlan-information <vlan_info>
{vlan} [<vlan_name> | all]
```

Description

Configures the VLAN info portion of the circuit ID of a VLAN.

Syntax Description

<code>vlan_info</code>	Specifies the circuit ID VLAN information for each VLAN in the format of <VLAN Info>-<Port Info>; maximum length is 32 bytes.
<code>vlan_name</code>	Specifies the VLAN for which DHCP should be enabled.
<code>all</code>	Specifies all VLANs.

Default

The default value is the ASCII representation of the ingress VLAN's ID.

Usage Guidelines

This command allows you to configure the VLAN information portion of the circuit ID of a VLAN. The VLAN info is a string of characters of up to 32 bytes in length, and is entered in the format of <VLAN Info><Port Info>. When a specific value is not configured for a VLAN, `vlan_info` defaults to the ASCII representation of the ingress VLAN's ID.

configure ip-security dhcp-snooping information option

```
configure ip-security dhcp-snooping information option
```

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

The default is `unconfigured`.

Usage Guidelines

This command enables the DHCP relay agent option (option 82), which is inserted into client-originated DHCP packets before they are forwarded to the server.

To disable the DHCP relay agent option (option 82), use the following command:

```
unconfigure ip-security dhcp-snooping information option
```

Example

The following command enable the DHCP relay agent option:

```
configure ip-security dhcp-snooping information information option
```

configure ip-security dhcp-snooping information policy

```
configure ip-security dhcp-snooping information policy [drop | keep | replace]
```

Description

Configures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

drop	Specifies to drop the packet.
keep	Specifies to keep the existing option 82 information in place.
replace	Specifies to replace the existing data with the switch's own data.

Default

The default value is `replace`.

Usage Guidelines

Use this command to set a policy for the relay agent. Packets can be dropped, the option 82 information can be replaced (the default), or the packet can be forwarded with the information unchanged.

Example

The following command configures the DHCP relay agent option 82 policy to keep:

```
configure ip-security dhcp-snooping information information policy keep
```


configure ip-security dhcp-bindings add

```
configure ip-security dhcp-binding add ip <ip_address> mac <mac_address> {vlan} <vlan_name>
server-port <server_port> client-port <client_port> lease-time <seconds>
```

Description

Creates a DHCP binding

Syntax Description

ip_address	Specifies the IP address for the DHCP binding.
mac_address	Specifies the MAC address for the DHCP binding.
vlan_name	Specifies the name of the VLAN for the DHCP binding.
server_port	Specifies the server port for the DHCP binding.
client_port	Specifies the client port for the DHCP binding.
seconds	Specifies the number of seconds for the lease.

Default

N/A.

Usage Guidelines

This commands allows you to add a DHCP binding in order to re-create the bindings after reboot and to allow IP Security features to work with clients having static IP addresses.

Note: Setting the lease-time to 0 causes the DHCP binding to be static; in other words, it is not aged-out if no DHCP renew occurs. This is for use with clients using static IP addresses.

configure ip-security dhcp-bindings delete

```
configure ip-security dhcp-binding delete ip <ip_address> {vlan} <vlan_name>
```

Description

Deletes a DHCP binding.

Syntax Description

ip_address	Specifies the IP address for the DHCP binding.
vlan_name	Specifies the name of the VLAN for the DHCP binding.

Default

N/A.

Usage Guidelines

This commands allows you to delete a DHCP binding created with the command `configure ip-security dhcp-binding add ip <ip_address> mac <mac_address> {vlan} <vlan_name> server-port <server_port> client-port <client_port> lease-time <seconds>`.

configure ip-security dhcp-binding storage filename

```
configure ip-security dhcp-bindings storage filename <name>
```

Description

Creates a storage file for DHCP binding information.

Syntax Description

name	Specifies the name of the DHCP binding storage file.
------	--

Default

N/A.

Usage Guidelines

This commands allows you to configure the filename with which the DHCP bindings storage file is created on the external server when it is uploaded to the external server. The text file resides on an external server. You can configure the server with the command `configure ip-security dhcp-bindings storage location server [primary | secondary] <ip_address> | <hostname>] tftp`.

The bindings file must have a `.xsf` extension. If the input filename doesn't already have a `.xsf` extension, one is added automatically.

configure ip-security dhcp-binding storage location

```
configure ip-security dhcp-bindings storage location server [primary | secondary]
<ip_address> | <hostname>] tftp
```

Description

Specifies the server location for the DHCP bindings storage file.

Syntax Description

ip_address	Specifies the IP address location for the bindings storage file.
------------	--

Default

N/A.

Usage Guidelines

This commands allows you to specify where you want to store the DHCP storage file that you created with the command `configure ip-security dhcp-bindings storage filename <name>`.

configure ip-security dhcp-bindings storage

```
configure ip-security dhcp-bindings storage [write-interval <minutes> | write-threshold
<num_changed_entries>]
```

Description

Configures DHCP bindings file storage upload variables.

Syntax Description

minutes	Specifies the number of minutes for the write interval.
num_changed_entries	Specifies the limit for the write threshold.

Default

The default write threshold is 50 entries; the default write interval is 30 minutes;

Usage Guidelines

This commands allows you to configure the upload variables for the DHCP bindings file that you created with the command `configure ip-security dhcp-bindings storage filename <name>` and specified the location of with the command `configure ip-security dhcp-bindings storage location server [primary | secondary] <ip_address> | <hostname>] tftp`.

For redundancy, the DHCP bindings file is uploaded to both the primary and the secondary server. The failure of one upload (for example, due to a TFTP server timeout) does not affect the upload of any other.

When the maximum file size limit is reached, no additional DHCP bindings can be uploaded until one of the older bindings is removed.

The point at which DHCP bindings can be uploaded can be configured to work in one of the following ways:

- **Periodic upload:** Upload every *N* minutes, provided that DHCP bindings have changed since the last upload.
- **Upload based on number of yet-to-be uploaded entries:** Allows you to configure the maximum number of changed entries that are allowed to accumulate before being uploaded.

The write interval is configurable from 5 minutes to 1 day, with a default value of 30 minutes. The default value of the write threshold is 50 entries, with a minimum of 25 and maximum of 200.

Additions and deletions are considered changes, but updates are not, which means that DHCP renewals of existing leases are not counted.

By default, the write interval is in effect, but not the write-threshold. You may change whichever of these you wish by explicitly configuring the value.

configure mac-lockdown-timeout ports aging-time

```
configure mac-lockdown-timeout ports [all | <port_list>] aging-time <seconds>
```

Description

Configures the MAC address lock down timeout value in seconds for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.
seconds	Configures the length of the time out value in seconds. The default is 15 seconds; the range is 15 to 2,000,000 seconds.

Default

The default is 15 seconds.

Usage Guidelines

This timer overrides the FDB aging time.

This command only sets the duration of the MAC address lock down timer. To enable the lock down timeout feature, use the following command:

```
enable mac-lockdown-timeout ports [all | <port_list>]
```

Example

The following command configures the MAC address lock down timer duration for 300 seconds for ports 2:3, 2:4, and 2:6:

```
configure mac-lockdown-timeout ports 2:3, 2:4, 2:6 aging-time 300
```

configure ports rate-limit flood

```
configure ports <port_list> rate-limit flood [broadcast | multicast | unknown-destmac]
[no-limit | <pps>]
```

Description

Limits the amount of ingress flooded traffic; minimizes network impact of broadcast loops.

Syntax Description

port_list	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
broadcast	Specifies all broadcast packets.
multicast	Specifies all flooded multicast packets (known IP multicast caches are still forwarded at line rate).
unknown-destmac	Specifies all packets with unknown MAC DAs.
no-limit	Specifies unlimited rate.
pps	Packets per second allowed; range is from 0 to 262,144.

Default

No limit.

Usage Guidelines

Use this command to limit the amount of ingress flooding traffic and to minimize the network impact of broadcast loops.

To display results, use the `show ports rate-limit flood` command.

Example

The following command rate limits broadcast packets on port 3 on a stand-alone switch to 500 pps:

```
configure ports 3 rate-limit flood broadcast 500
```

configure ports vlan

```
configure ports <portlist> vlan <vlan_name> [limit-learning <number> {action [blackhole | stop-learning]} | lock-learning | unlimited-learning | unlock-learning]
```

Description

Configures virtual ports for limited or locked MAC address learning.

Syntax Description

portlist	Specifies one or more ports or slots and ports.
vlan_name	Specifies the name of the VLAN.

limit-learning <number>	Specifies a limit on the number of MAC addresses that can be dynamically learned on the specified ports.
blackhole	Specifies that blackhole entries are allowed.
stop-learning	Specifies that the learning be halted to protect the switch from exhausting FDB resources by not creating blackhole entries.
lock-learning	Specifies that the current FDB entries for the specified ports should be made permanent static, and no additional learning should be allowed.
unlimited-learning	Specifies that there should not be a limit on MAC addresses that can be learned.
unlock-learning	Specifies that the port should be unlocked (allow unlimited, dynamic learning).

Default

Unlimited, unlocked learning.

Usage Guidelines

N/A

Limited learning

The limited learning feature allows you to limit the number of dynamically-learned MAC addresses per VLAN. When the learned limit is reached, all new source MAC addresses are blackholed at both the ingress and egress points. This prevent these MAC addresses from learning and responding to Internet control message protocol (ICMP) and address resolution protocol (ARP) packets.

If the limit you configure is greater than the current number of learned entries, all the current learned entries are purged.

Dynamically learned entries still get aged, and can be cleared. If entries are cleared or aged out after the learning limit has been reached, new entries will then be able to be learned until the limit is reached again.

Permanent static and permanent dynamic entries can still be added and deleted using the `create fdbentry` and `delete fdbentry` commands. These override any dynamically learned entries.

For ports that have a learning limit in place, the following traffic still flows to the port:

- Packets destined for permanent MACs and other non-blackholed MACs
- Broadcast traffic

Traffic from the permanent MAC and any other non-blackholed MACs will still flow from the virtual port.

Stop learning

When `stop-learning` is enabled with `learning-limit` configured, the switch is protected from exhausting FDB resources by not creating blackhole entries. Any additional learning and forwarding is prevented, but packet forwarding from FDB entries is not impacted.

Port lockdown

The port lockdown feature allows you to prevent any additional learning on the virtual port, keeping existing learned entries intact. This is equivalent to making the dynamically-learned entries permanent static, and setting the learning limit to zero. All new source MAC addresses are blackholed.

Locked entries do not get aged, but can be deleted like any other permanent FDB entries. The maximum number of permanent lockdown entries is 1024. Any FDB entries above will be flushed and blackholed during lockdown.

For ports that have lockdown in effect, the following traffic still flows to the port:

- Packets destined for the permanent MAC and other non-blackholed MACs
- Broadcast traffic

Traffic from the permanent MAC will still flow from the virtual port.

Once the port is locked down, all the entries become permanent and will be saved across reboot.

When you remove the lockdown using the `unlock-learning` option, the `learning-limit` is reset to unlimited, and all associated entries in the FDB are flushed.

To display the locked entries on the switch, use the following command:

```
show fdb
```

Locked MAC address entries have the “I” flag.

To verify the MAC security configuration for the specified VLAN or ports, use the following commands:

```
show vlan <vlan name> security
show ports <portlist> info detail
```

Example

The following command limits the number of MAC addresses that can be learned on ports 1, 2, 3, and 6 in a VLAN named *accounting*, to 128 addresses:

```
configure ports 1, 2, 3, 6 vlan accounting learning-limit 128
```

The following command locks ports 4 and 5 of VLAN *accounting*, converting any FDB entries to static entries, and prevents any additional address learning on these ports:

```
configure ports 4,5 vlan accounting lock-learning
```

The following command removes the learning limit from the specified ports:

```
configure ports 1, 2, vlan accounting unlimited-learning
```

The following command unlocks the FDB entries for the specified ports:

```
configure ports 4,5 vlan accounting unlock-learning
```

configure radius server client-ip

```
configure radius {mgmt-access | netlogin} [primary | secondary] server [<ipaddress> |  
<hostname>] {<udp_port>} client-ip [<ipaddress>] {vr <vr_name>}
```

Description

Configures the primary and secondary RADIUS authentication server.

Syntax Description

mgmt-access	Specifies the RADIUS authentication server for switch management.
netlogin	Specifies the RADIUS authentication server for network login.
primary	Configures the primary RADIUS authentication server.
secondary	Configures the secondary RADIUS authentication server.
ipaddress	The IP address of the server being configured.
hostname	The host name of the server being configured.
udp_port	The UDP port to use to contact the RADIUS authentication server.
ipaddress	The IP address used by the switch to identify itself when communicating with the RADIUS authentication server.
vr_name	Specifies the virtual router on which the client IP is located.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*.

Default

The following lists the default behavior of this command:

- The UDP port setting is 1812
- The virtual router used is VR-Mgmt, the management virtual router
- Switch management and network login use the same primary and secondary RADIUS servers for authentication.

Usage Guidelines

Use this command to specify RADIUS server information.

Use of the <hostname> parameter requires that DNS be enabled.

The RADIUS server defined by this command is used for user name authentication and CLI command authentication.

You can specify one pair of RADIUS authentication servers for switch management and another pair for network login. To specify RADIUS authentication servers for switch management (Telnet, SSH, and console sessions), use the `mgmt-access` keyword. To specify RADIUS authentication servers for network login, use the `netlogin` keyword. If you do not specify a keyword, switch management and network login use the same pair of RADIUS authentication servers.

Example

The following command configures the primary RADIUS server on host `radius1` using the default UDP port (1812) for use by the RADIUS client on switch `10.10.20.30` using a virtual router interface of `VR-Default`:

```
configure radius primary server radius1 client-ip 10.10.20.30 vr vr-Default
```

The following command configures the primary RADIUS server for network login authentication on host `netlog1` using the default UDP port for use by the RADIUS client on switch `10.10.20.31` using, by default, the management virtual router interface:

```
configure radius netlogin primary server netlog1 client-ip 10.10.20.31
```

configure radius shared-secret

```
configure radius {mgmt-access | netlogin} [primary | secondary] shared-secret {encrypted}
<string>
```

Description

Configures the authentication string used to communicate with the RADIUS authentication server.

Syntax Description

<code>mgmt-access</code>	Specifies the switch management RADIUS authentication server.
<code>netlogin</code>	Specifies the network login RADIUS authentication server.
<code>primary</code>	Configures the authentication string for the primary RADIUS server.
<code>secondary</code>	Configures the authentication string for the secondary RADIUS server.
<code>encrypted</code>	Indicates that the string is already encrypted.
<code>string</code>	The string to be used for authentication.

Default

Unconfigured.

Usage Guidelines

The secret must be the same between the client switch and the RADIUS server.

The RADIUS server must first be configured for use with the switch as a RADIUS client.

The `mgmt-access` keyword specifies the RADIUS server used for switch management authentication.

The `netlogin` keyword specifies the RADIUS server used for network login authentication.

If you do not specify the `mgmt-access` or `netlogin` keywords, the secret applies to both the primary or secondary switch management and `netlogin` RADIUS servers.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purplegreen” on the primary RADIUS server for both switch management and network login:

```
configure radius primary shared-secret purplegreen
```

The following command configures the shared secret as “redblue” on the primary switch management RADIUS server:

```
configure radius mgmt-access primary shared-secret redblue
```

configure radius timeout

```
configure radius {mgmt-access | netlogin} timeout <seconds>
```

Description

Configures the timeout interval for RADIUS authentication requests.

Syntax Description

<code>mgmt-access</code>	Specifies the switch management RADIUS authentication server.
<code>netlogin</code>	Specifies the network login RADIUS authentication server.
<code>seconds</code>	Specifies the number of seconds for authentication requests. Range is 3 to 120 seconds

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for RADIUS authentication requests. When the timeout has expired, another authentication attempt will be made. After three failed attempts to authenticate, the alternate server will be used. After six failed attempts, local user authentication will be used.

The `mgmt-access` keyword specifies the RADIUS server used for switch management authentication.

The `netlogin` keyword specifies the RADIUS server used for network login authentication.

If you do not specify the `mgmt-access` or `netlogin` keywords, the timeout interval applies to both switch management and netlogin RADIUS servers.

Example

The following command configures the timeout interval for RADIUS authentication to 10 seconds. After 30 seconds (three attempts), the alternate RADIUS server will be used. After 60 seconds (six attempts) local user authentication is used.

```
configure radius timeout 10
```

configure radius-accounting server client-ip

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] server  
[<ipaddress> | <hostname>] {<tcp_port>} client-ip [<ipaddress>] {vr <vr_name>}
```

Description

Configures the RADIUS accounting server.

Syntax Description

mgmt-access	Specifies the RADIUS accounting server for switch management.
netlogin	Specifies the RADIUS accounting server for network login.
primary	Configure the primary RADIUS accounting server.
secondary	Configure the secondary RADIUS accounting server.
ipaddress	The IP address of the accounting server being configured.
hostname	The host name of the accounting server being configured.
tcp_port	The UDP port to use to contact the RADIUS accounting server.
ipaddress	The IP address used by the switch to identify itself when communicating with the RADIUS accounting server.
vr_name	Specifies the virtual router on which the client IP is located.
Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .	

Default

The following lists the default behavior of this command:

- The UDP port setting is 1813
- The virtual router used is VR-Mgmt, the management virtual router
- Switch management and network login use the same RADIUS accounting server.

Usage Guidelines

Use this command to specify the radius accounting server.

The accounting server and the RADIUS authentication server can be the same.

Use of the `<hostname>` parameter requires that DNS be enabled.

You can specify one pair of RADIUS accounting servers for switch management and another pair for network login. To specify RADIUS accounting servers for switch management (Telnet, SSH, and console sessions), use the `mgmt-access` keyword. To specify RADIUS accounting servers for network login, use the `netlogin` keyword. If you do not specify a keyword, switch management and network login use the same pair of RADIUS accounting servers.

Example

The following command configures RADIUS accounting on host radius1 using the default UDP port (1813) for use by the RADIUS client on switch 10.10.20.30 using a virtual router interface of VR-Default for both management and network login:

```
configure radius-accounting primary server radius1 client-ip 10.10.20.30 vr vr-Default
```

The following command configures RADIUS accounting for network login on host netlog1 using the default UDP port for use by the RADIUS client on switch 10.10.20.31 using the default virtual router interface:

```
configure radius-accounting netlogin primary server netlog1 client-ip 10.10.20.31
```

configure radius-accounting shared-secret

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] shared-secret {encrypted} <string>
```

Description

Configures the authentication string used to communicate with the RADIUS accounting server.

Syntax Description

<code>mgmt-access</code>	Specifies the switch management RADIUS accounting server.
<code>netlogin</code>	Specifies the network login RADIUS accounting server.
<code>primary</code>	Configures the authentication string for the primary RADIUS accounting server.
<code>secondary</code>	Configures the authentication string for the secondary RADIUS accounting server.
<code>encrypted</code>	Indicates that the string is already encrypted.
<code>string</code>	The string to be used for authentication.

Default

Unconfigured.

Usage Guidelines

The secret must be the same between the client switch and the RADIUS accounting server.

The `mgmt-access` keyword specifies the RADIUS accounting server used for switch management.

The `netlogin` keyword specifies the RADIUS accounting server used for network login.

If you do not specify the `mgmt-access` or `netlogin` keywords, the secret applies to both the primary or secondary switch management and netlogin RADIUS accounting servers.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purpleaccount” on the primary RADIUS accounting server for both management and network login:

```
configure radius primary shared-secret purpleaccount
```

The following command configures the shared secret as “greenaccount” on the primary management RADIUS accounting server:

```
configure radius mgmt-access primary shared-secret greenaccount
```

configure radius-accounting timeout

```
configure radius-accounting {mgmt-access | netlogin} timeout <seconds>
```

Description

Configures the timeout interval for RADIUS-Accounting authentication requests.

Syntax Description

<code>mgmt-access</code>	Specifies the switch management RADIUS accounting server.
<code>netlogin</code>	Specifies the network login RADIUS accounting server.
<code>seconds</code>	Specifies the number of seconds for accounting requests. Range is 3 to 120 seconds.

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for RADIUS-Accounting authentication requests. When the timeout has expired, another authentication attempt will be made. After three failed attempts to authenticate, the alternate server will be used.

The `mgmt-access` keyword specifies the RADIUS accounting server used for switch management.

The `netlogin` keyword specifies the RADIUS accounting server used for network login.

If you do not specify the `mgmt-access` or `netlogin` keywords, the timeout interval applies to both switch management and netlogin RADIUS accounting servers.

Example

This example configures the timeout interval for RADIUS-Accounting authentication to 10 seconds. After 30 seconds (three attempts), the alternate RADIUS server will be used:

```
configure radius-accounting timeout 10
```

configure ssh2 key

```
configure ssh2 key {pregenerated}
```

Description

Generates the Secure Shell 2 (SSH2) host key.

Syntax Description

pregenerated	Indicates that the SSH2 authentication key has already been generated. The user will be prompted to enter the existing key.
--------------	---

Default

The switch generates a key for each SSH2 session.

Usage Guidelines

Secure Shell 2 (SSH2) is a feature of the NETGEAR 8800 that allows you to encrypt session data between a network administrator using SSH2 client software and the switch or to send encrypted data from the switch to an SSH2 client on a remote system. Configuration, policy, image, and public key files may also be transferred to the switch using the Secure Copy Program 2 (SCP2).

Before you use SSH2, you must generate a host key and enable SSH2. To generate an SSH2 host key, use the `configure ssh2 key` command. To enable SSH2, use the `enable ssh2` command.

An authentication key must be generated before the switch can accept incoming SSH2 sessions. This can be done automatically by the switch, or you can enter a previously generated key.

If you elect to have the key generated, the key generation process can take up to ten minutes, and cannot be canceled after it has started. Once the key has been generated, you should save your configuration to preserve the key.

To use a key that has been previously created, use the `pregenerated` keyword. Use the `show ssh2 private-key` command to list and copy the previously generated key. Then use the `configure ssh2 key {pregenerated}` command where “pregenerated” represents the key that you paste.

The key generation process generates the SSH2 private host key. The SSH2 public host key is derived from the private host key, and is automatically transmitted to the SSH2 client at the beginning of an SSH2 session.

To view the status of SSH2 on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2 sessions, whether a valid key is present, and the TCP port and virtual router that is being used.

Example

The following command generates an authentication key for the SSH2 session:

```
configure ssh2 key
```

The command responds with the following messages:

```
WARNING: Generating new server host key
This will take approximately 10 minutes and cannot be canceled.
Continue? (y/n)
```

If you respond yes, the command begins the process.

To configure an SSH2 session using a previously generated key, use the following command:

```
configure ssh2 key pregenerated <pre-generated key>
```

Enter the previously-generated key (you can copy and paste it from the saved configuration file; a part of the key pattern is similar to `2d:2d:2d:2d:20:42:45:47:`).

configure sshd2 user-key add user

```
configure sshd2 user-key <key_name> add user <user_name>
```

Description

Associates a user to a key.

Syntax Description

<code>key_name</code>	Specifies the name of the public key.
-----------------------	---------------------------------------

user_name	Specifies the name of the user.
-----------	---------------------------------

Default

N/A.

Usage Guidelines

This command associates (or *binds*) a user to a key.

Example

The following example binds the key `id_dsa_2048` to user `admin`.

```
configure sshd2 user-key id_dsa_2048 add user admin
```

configure sshd2 user-key delete user

```
configure sshd2 user-key <key_name> delete user <user_name>
```

Description

Disassociates a user to a key.

Syntax Description

key_name	Specifies the name of the public key.
user_name	Specifies the name of the user.

Default

N/A.

Usage Guidelines

This command disassociates (or *unbinds*) a user to a key.

Example

The following example unbinds the key `id_dsa_2048` from user `admin`.

```
configure sshd2 user-key id_dsa_2048 delete user admin
```

configure ssl certificate pregenerated

```
configure ssl certificate pregenerated
```


Description

Obtains the pre-generated certificate from the user.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You must upload or generate a certificate for SSL server use. With this command, you copy and paste the certificate into the command line followed by a blank line to end the command. The following security algorithms are supported:

- RSA for public key cryptography (generation of certificate and public-private key pair, certificate signing). RSA key size between 1024 and 4096 bits.
- Symmetric ciphers (for data encryption): RC4, DES, and 3DES.
- Message Authentication Code (MAC) algorithms: MD5 and SHA.

This command is also used when downloading or uploading the configuration. Do not modify the certificate stored in the uploaded configuration file because the certificate is signed using the issuer's private key.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Example

The following command obtains the pre-generated certificate from the user:

```
configure ssl certificate pregenerated
```

Next, you open the certificate and then copy and paste the certificate into the console/Telnet session, followed by a blank line to end the command.

configure ssl certificate privkeylen

```
configure ssl certificate privkeylen <length> country <code> organization <org_name>  
common-name <name>
```

Description

Creates a self signed certificate and private key that can be saved in the EEPROM.

Syntax Description

length	Specifies the private key length in bytes. Valid values are between 1024 and 4096.
--------	--

code	Specifies the country code in 2-character form.
org_name	Specifies the organization name. The organization name can be up to 64 characters long.
name	Specifies the common name. The common name can be up to 64 characters long.

Default

N/A.

Usage Guidelines

This command creates a self signed certificate and private key that can be saved in the EEPROM. The certificate generated is in the PEM format.

Any existing certificate and private key is overwritten.

The size of the certificate depends on the RSA key length (`privkeylen`) and the length of the other parameters (`country`, `organization name`, and so forth) supplied by the user. If the RSA key length is 1024, then the certificate is approximately 1 kb. For an RSA key length of 4096, the certificate length is approximately 2 kb, and the private key length is approximately 3 kb.

Example

The following command creates an SSL certificate in the USA for a website called bigcats:

```
configure ssl certificate privkeylen 2048 country US organization IEEE common-name bigcats
```

configure ssl privkey pregenerated

```
configure ssl privkey pregenerated
```

Description

Obtains the pre-generated private key from the user.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command is also used when downloading or uploading the configuration. The private key is stored in the EEPROM, and the certificate is stored in the configuration file.

With this command, you copy and paste the private key into the command line followed by a blank line to end the command. The following security algorithms are supported:

- RSA for public key cryptography (generation of certificate and public-private key pair, certificate signing). RSA key size between 1024 and 4096 bits.
- Symmetric ciphers (for data encryption): RC4, DES, and 3DES.
- Message Authentication Code (MAC) algorithms: MD5 and SHA.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Example

The following command obtains the pre-generated private key from the user:

```
configure ssl privkey pregenerated
```

Next, you open the certificate and then copy and paste the certificate into the console/Telnet session, followed by a RETURN to end the command.

configure tacacs server client-ip

```
configure tacacs [primary | secondary] server [<ipaddress> | <hostname>] {<tcp_port>}
client-ip <ipaddress> {vr <vr_name>}
```

Description

Configures the server information for a TACACS+ authentication server.

Syntax Description

primary	Configures the primary TACACS+ server.
secondary	Configures the secondary TACACS+ server.
ipaddress	The IP address of the TACACS+ server being configured.
hostname	The host name of the TACACS+ server being configured.
tcp_port	The TCP port to use to contact the TACACS+ server.
ipaddress	The IP address used by the switch to identify itself when communicating with the TACACS+ server.
vr_name	Specifies the virtual router on which the client IP is located.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*.

Default

TACACS+ uses TCP port 49. The default virtual router is VR-Mgmt, the management virtual router.

Usage Guidelines

Use this command to configure the server information for a TACACS+ server.

To remove a server, use the following command:

```
unconfigure tacacs server [primary | secondary]
```

Use of the <hostname> parameter requires that DNS be enabled.

Example

The following command configures server tacacs1 as the primary TACACS+ server for client switch 10.10.20.35 using a virtual router interface of VR-Default:

```
configure tacacs primary server tacacs1 client-ip 10.10.20.35 vr vr-Default
```

configure tacacs shared-secret

```
configure tacacs [primary | secondary] shared-secret {encrypted} <string>
```

Description

Configures the shared secret string used to communicate with the TACACS+ authentication server.

Syntax Description

primary	Configures the authentication string for the primary TACACS+ server.
secondary	Configures the authentication string for the secondary TACACS+ server.
encrypted	Indicates that the string is already encrypted.
string	The string to be used for authentication.

Default

N/A.

Usage Guidelines

The secret must be the same between the client switch and the TACACS+ server.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purplegreen” on the primary TACACS+ server:

```
configure tacacs-accounting primary shared-secret purplegreen
```

configure tacacs timeout

```
configure tacacs timeout <seconds>
```

Description

Configures the timeout interval for TACAS+ authentication requests.

Syntax Description

seconds	Specifies the number of seconds for authentication requests. Range is 3 to 120 seconds.
---------	---

Default

The default is 3 seconds.

Usage Guidelines

Use this command to configure the timeout interval for TACACS+ authentication requests.

To detect and recover from a TACACS+ server failure when the timeout has expired, the switch makes one authentication attempt before trying the next designated TACACS+ server or reverting to the local database for authentication. In the event that the switch still has IP connectivity to the TACACS+ server, but a TCP session cannot be established, (such as a failed TACACS+ daemon on the server), failover happens immediately regardless of the configured timeout value.

For example, if the timeout value is set for 3 seconds (the default value), it will take 3 seconds to fail over from the primary TACACS+ server to the secondary TACACS+ server. If both the primary and the secondary servers fail or are unavailable, it takes approximately 6 seconds to revert to the local database for authentication.

Example

The following command configures the timeout interval for TACACS+ authentication to 10 seconds:

```
configure tacacs timeout 10
```

configure tacacs-accounting server

```
configure tacacs-accounting [primary | secondary] server [<ipaddress> | <hostname>]
{<udp_port>} client-ip <ipaddress> {vr <vr_name>}
```

Description

Configures the TACACS+ accounting server.

Syntax Description

primary	Configures the primary TACACS+ accounting server.
secondary	Configures the secondary TACACS+ accounting server.
ipaddress	The IP address of the TACACS+ accounting server being configured.
hostname	The host name of the TACACS+ accounting server being configured.
tcp_port	The TCP port to use to contact the TACACS+ server.
ipaddress	The IP address used by the switch to identify itself when communicating with the TACACS+ accounting server.
vr_name	Specifies the virtual router on which the client IP is located.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*.

Default

Unconfigured. The default virtual router is VR-Mgmt, the management virtual router.

Usage Guidelines

You can use the same TACACS+ server for accounting and authentication.

To remove a server, use the following command:

```
unconfigure tacacs server [primary | secondary]
```

Example

The following command configures server tacacs1 as the primary TACACS+ accounting server for client switch 10.10.20.35 using a virtual router interface of VR-Default:

```
configure tacacs-accounting primary server tacacs1 client-ip 10.10.20.35 vr vr-Default
```

configure tacacs-accounting shared-secret

```
configure tacacs-accounting [primary | secondary] shared-secret {encrypted} <string>
```

Description

Configures the shared secret string used to communicate with the TACACS+ accounting server.

Syntax Description

primary	Configures the authentication string for the primary TACACS+ accounting server.
---------	---

secondary	Configures the authentication string for the secondary TACACS+ accounting server.
string	The string to be used for authentication.

Default

N/A.

Usage Guidelines

Secret needs to be the same as on the TACACS+ server.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “tacacsaccount” on the primary TACACS+ accounting server:

```
configure tacacs-accounting primary shared-secret tacacsaccount
```

configure tacacs-accounting timeout

```
configure tacacs-accounting timeout <seconds>
```

Description

Configures the timeout interval for TACACS+ accounting authentication requests.

Syntax Description

seconds	Specifies the number of seconds for accounting requests. Range is 3 to 120 seconds
---------	--

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for TACACS+ accounting authentication requests.

To detect and recover from a TACACS+ accounting server failure when the timeout has expired, the switch makes one authentication attempt before trying the next designated TACACS+ accounting server or reverting to the local database for authentication. In the event that the switch still has IP connectivity to the TACACS+ accounting server, but a TCP

session cannot be established, (such as a failed TACACS+ daemon on the accounting server), failover happens immediately regardless of the configured timeout value.

For example, if the timeout value is set for 3 seconds (the default value), it takes 3 seconds to fail over from the primary TACACS+ accounting server to the secondary TACACS+ accounting server. If both the primary and the secondary servers fail or are unavailable, it takes approximately 6 seconds to revert to the local database for authentication.

Example

The following command configures the timeout interval for TACACS+ accounting authentication to 10 seconds:

```
configure tacacs-accounting timeout 10
```

configure trusted-ports trust-for dhcp-server

```
configure trusted-ports [<ports>|all] trust-for dhcp-server
```

Description

Configures one or more trusted DHCP ports.

Syntax Description

ports	Specifies one or more ports to be configured as trusted ports.
all	Specifies all ports to be configured as trusted ports.

Default

N/A.

Usage Guidelines

To configure trusted DHCP ports, you must first enable DHCP snooping on the switch. To enable DHCP snooping, use the following command:

```
enable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>]  
violation-action [drop-packet {[block-mac | block-port] [duration <duration_in_seconds>  
| permanently] | none}}] {snmp-trap}
```

Trusted ports do not block traffic; rather, the switch forwards any DHCP server packets that appear on trusted ports. Depending on your DHCP snooping configuration, the switch drops packets and can disable the port temporarily, disable the port permanently, blackhole the MAC address temporarily, blackhole the MAC address permanently, and so on.

If you configure one or more trusted ports, the switch assumes that all DHCP server packets on the trusted port are valid.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display any violations that occur, including those on DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command configures ports 2:2 and 2:3 as trusted ports:

```
configure trusted-ports 2:2-2:3 trust-for dhcp-server
```

configure trusted-servers add server

```
configure trusted-servers {vlan} <vlan_name> add server <ip_address> trust-for dhcp-server
```

Description

Configures and enables a trusted DHCP server on the switch.

Syntax Description

vlan_name	Specifies the VLAN name.
ip_address	Specifies the IP address of the trusted DHCP server.

Default

N/A.

Usage Guidelines

If you configured trusted DHCP server, the switch forwards only DHCP packets from the trusted servers. The switch drops DHCP packets from other DHCP snooping-enabled ports.

You can configure a maximum of eight trusted DHCP servers on the switch.

If you configure a port as a trusted port, the switch assumes that all DHCP server packets on that port are valid.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display any violations that occur, including those on the DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command configures a trusted DHCP server on the switch:

```
configure trusted-servers vlan purple add server 10.10.10.10 trust-for dhcp-server
```

configure trusted-servers delete server

```
configure trusted-servers vlan <vlan_name> delete server <ip_address> trust-for dhcp-server
```

Description

Deletes a trusted DHCP server from the switch.

Syntax Description

vlan_name	Specifies the VLAN name.
ip_address	Specifies the IP address of the trusted DHCP server.

Default

N/A.

Usage Guidelines

Use this command to delete a trusted DHCP server from the switch.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display any violations that occur, including those on the DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command deletes a trusted DHCP server from the switch:

```
configure trusted-servers vlan purple delete server 10.10.10.10 trust-for dhcp-server
```

configure vlan dhcp-address-range

```
configure vlan <vlan_name> dhcp-address-range <ipaddress1> - <ipaddress2>
```

Description

Configures a set of DHCP addresses for a VLAN.

Syntax Description

vlan_name	Specifies the VLAN on whose ports DHCP will be enabled.
ipaddress1	Specifies the first IP address in the DHCP address range to be assigned to this VLAN.
ipaddress2	Specifies the last IP address in the DHCP address range to be assigned to this VLAN.

Default

N/A.

Usage Guidelines

The following error conditions are checked: ipaddress2 >= ipaddress1, the range must be in the VLAN's network, the range does not contain the VLAN's IP address, and the VLAN has an IP address assigned.

Example

The following command allocates the IP addresses between 192.168.0.20 and 192.168.0.100 for use by the VLAN *temporary*:

```
configure temporary dhcp-address-range 192.168.0.20 - 192.168.0.100
```

configure vlan dhcp-lease-timer

```
configure vlan <vlan_name> dhcp-lease-timer <lease-timer>
```

Description

Configures the timer value in seconds returned as part of the DHCP response.

Syntax Description

name	Specifies the VLAN on whose ports netlogin should be disabled.
lease-timer	Specifies the timer value, in seconds.

Default

N/A.

Usage Guidelines

The timer value is specified in seconds. The timer value range is 0 - 4294967295, where 0 indicates the default (not configured) value of 7200 second.

Example

The following command configures the DHCP lease timer value for VLAN *corp*:

```
configure vlan corp dhcp-lease-timer <lease-timer>
```

configure vlan dhcp-options

```
configure {vlan} <vlan_name> dhcp-options [default-gateway | dns-server {primary | secondary}  
| wins-server] <ipaddress>
```

Description

Configures the DHCP options returned as part of the DHCP response by a switch configured as a DHCP server.

Syntax Description

vlan_name	Specifies the VLAN on which to configure DHCP
default-gateway	Specifies the router option.
dns-server	Specifies the Domain Name Server (DNS) option.
primary	Specifies the primary DNS option.
secondary	Specifies the secondary DNS option.
wins-server	Specifies the NetBIOS name server (NBNS) option.
ipaddress	The IP address associated with the specified option.

Default

N/A.

Usage Guidelines

This command configures the DHCP options that can be returned to the DHCP client. For the `default-gateway` option you are only allowed to configure an IP address that is in the VLAN's network range. For the other options, any IP address is allowed.

The options below represent the following BOOTP options specified by RFC2132:

- `default-gateway`—Router option, number 3

- `dns-server`—Domain Name Server option, number 6
- `wins-server`—NetBIOS over TCP/IP Name Server option, number 44

Example

The following command configures the DHCP server to return the IP address 10.10.20.8 as the router option:

```
configure vlan <name> dhcp-options default-gateway 10.10.20.8
```

create sshd2 key-file

```
create sshd2 key-file {host-key | user-key} <key_name>
```

Description

Creates a file for the user-key or host-key.

Syntax Description

<code>host-key</code>	Specifies the name of the host-key
<code>user-key</code>	Specifies the name of the user-key.
<code>key_name</code>	Specifies the name of the public key.

Default

N/A.

Usage Guidelines

This command is used to write the user or the host public key in a file. The key files will be created with a .ssh file extension; this enables the administrator to copy the public key files to another server.

create sshd2 user-key

```
create sshd2 user-key <key_name> <key> {subject <subject>} {comment <comment>}
```

Description

Creates a user key.

Syntax Description

<code>key_name</code>	Specifies the name of the public key.
<code>key</code>	Specifies the key.

Note: The key cannot have any spaces in it.

subject	Specifies the subject.
comment	Specifies the comment (an optional field)

Default

N/A.

Usage Guidelines

This command is used to enter, or cut and paste, your public key. You can also enter the public key into the switch by using the SCP or SFTP client that is connected to the switch.

delete sshd2 user-key

```
delete sshd2 user-key <key_name>
```

Description

Deletes a user key.

Syntax Description

key_name	Specifies the name of the public key to be deleted.
----------	---

Default

N/A.

Usage Guidelines

This command is used to delete a user key. The key is deleted regardless of whether or not it is bound to a user.

Note: If a user is bound to the key, they are first unbound or unassociated, and then the key is deleted

Example

The following example shows the SSH user key `id_dsa_2048` being deleted:

```
delete sshd2 user-key id_dsa_2048
```

disable dhcp ports vlan

```
disable dhcp ports <portlist> vlan <vlan_name>
```

Description

Disables DHCP on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which DHCP should be disabled.
vlan_name	Specifies the VLAN on whose ports DHCP should be disabled.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables DHCP for port 6:9 in VLAN *corp*:

```
disable dhcp ports 6:9 vlan corp
```

disable dos-protect

```
disable dos-protect
```

Description

Disables denial of service protection.

Syntax Description

There are no arguments or variables for this command.

Default

Default is disabled.

Usage Guidelines

None.

Example

The following command disables denial of service protection.

```
disable dos-protect
```

disable iparp gratuitous protect vlan

```
disable iparp gratuitous protect vlan <vlan-name>
```

Description

Disables gratuitous ARP protection on the specified VLAN.

Syntax Description

vlan-name	Specifies the VLAN.
-----------	---------------------

Default

Disabled.

Usage Guidelines

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

This command disables gratuitous ARP protection.

Example

The following command disables gratuitous ARP protection for VLAN *corp*:

```
disable iparp gratuitous protect vlan corp
```

disable ip-security anomaly-protection

```
disable ip-security anomaly-protection {slot [ <slot> | all ]}
```

Description

Disables all anomaly checking options.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This commands disables all anomaly checking options, including IP address, UDP/TCP port, TCP flag and fragment, and ICMP anomaly checking.

disable ip-security anomaly-protection ip

```
disable ip-security anomaly-protection ip { slot [ <slot> | all ] }
```

Description

Disables source and destination IP address checking.

Syntax Description

slot	Specifies the slot.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command disables source and destination IP addresses checking. This checking takes effect for both IPv4 and IPv6 packets. When enabled, the switch drops IPv4/IPv6 packets if its source IP address are the same as the destination IP address. In most cases, the condition of source IP address being the same as the destination IP address indicates a Layer 3 protocol error. (These kind of errors are found in LAND attacks.)

disable ip-security anomaly-protection l4port

```
disable ip-security anomaly-protection l4port {slot [ <slot> | all ]}
```

Description

Disables TCP and UDP ports checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command disables TCP and UDP ports checking. This checking takes effect for both IPv4 and IPv6 TCP and UDP packets. When enabled, the switch drops TCP and UDP packets if its source port is the same as its destination port. In most cases, when the condition of source port is the same as that of the destination port, it indicates a Layer 4 protocol error. (This type of error can be found in a BALT attack.)

disable ip-security anomaly-protection tcp flags

```
disable ip-security anomaly-protection tcp flags {slot [ <slot> | all ]}
```

Description

Disables TCP flag checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command disables TCP flag checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops TCP packets if one of following condition is true:

- TCP SYN flag==1 and the source port<1024
- TCP control flag==0 and the sequence number==0
- TCP FIN, URG, and PSH bits are set, and the sequence number==0
- TCP SYN and FIN both are set.

disable ip-security anomaly-protection tcp fragment

```
disable ip-security anomaly-protection tcp fragment {slot [ <slot> | all ]}
```

Description

Disables TCP fragment checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command disables TCP fragment checking. This checking takes effect for IPv4/IPv6. When it is enabled, the switch drops TCP packets if one of following condition is true:

- For the first IPv4 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv4 TCP header allowed size
- If its IP offset field==1 (for IPv4 only)

disable ip-security anomaly-protection icmp

```
disable ip-security anomaly-protection icmp {slot [ <slot> | all ]}
```

Description

Disables ICMP size and fragment checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command disables ICMP size and fragment checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops ICMP packets if one of following condition is true:

- Fragmented ICMP packets for IPv4 packets.
- IPv4 ICMP pings packets with payload size greater than the maximum IPv4 ICMP-allowed size. (The maximum allowed size is configurable.)
- IPv6 ICMP ping packets with payload size > the maximum IPv6 ICMP-allowed size. (The maximum allowed size is configurable.)

disable ip-security anomaly-protection notify

```
disable ip-security anomaly-protection notify [log | snmp | cache] {slot [ <slot> | all ]}
```

Description

Disables protocol anomaly notification.

Syntax Description

log	Specifies the switch to send the notification to a log file.
snmp	Specifies the switch to send an SNMP trap when an event occurs.
cache	Specifies the switch to send the notification to cache.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command disables anomaly notification. When enabled, any packet failed to pass enabled protocol checking is sent to XOS Host CPU and notifies the user. There are three different types of notifications:

- `log`: log anomaly events in the switch log system; you can view and manage this log with the `show log` and `configure log` commands
- `snmp`: the anomaly events generate SNMP traps
- `cache`: logs the most recent and unique anomaly events in memory; rebooting the switch will cause all the logged events to be lost (the number of cached events is configured by command)

When disabled, the switch drops all violating packets silently.

disable ip-security arp gratuitous-protection

```
disable ip-security arp gratuitous-protection {vlan} [all | <vlan_name>]
```

Description

Disables gratuitous ARP protection on one or all VLANs on the switch.

Syntax Description

all	Specifies all VLANs configured on the switch.
-----	---

vlan-name	Specifies the VLAN.
-----------	---------------------

Default

By default, gratuitous ARP protection is disabled.

Usage Guidelines

This command replaces the `disable iparp gratuitous protect vlan` command.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

This command disables gratuitous ARP protection.

Example

The following command disables gratuitous ARP protection for VLAN *corp*:

```
disable ip-security arp gratuitous-protection vlan corp
```

disable ip-security arp learning learn-from-arp

```
disable ip-security arp learning learn-from-arp {vlan} <vlan_name> ports [all | <ports>]
```

Description

Disables ARP learning on the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
ports	Specifies one or more ingress ports.

Default

By default, ARP learning is enabled.

Usage Guidelines

You can disable ARP learning so that the only entries in the ARP table are either manually added or those created by DHCP secured ARP; the switch does not add entries by tracking

ARP requests and replies. By disabling ARP learning and adding a permanent entry or configuring DHCP secured ARP, you can centrally manage and allocate client IP addresses and prevent duplicate IP addresses from interrupting network operation.

To manually add a permanent entry to the ARP table, use the following command:

```
configure iparp add <ip_addr> {vr <vr_name>} <mac>
```

To configure DHCP secure ARP as a method to add entries to the ARP table, use the following command:

```
enable ip-security arp learning learn-from-dhcp vlan <vlan_name> ports [all | <ports>]
{poll-interval <interval_in_seconds>} {retries <number_of_retries>}
```

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} <vlan_name>
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {<ip_addr> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

Note: DHCP secured ARP entries are stored as static entries in the ARP table.

Example

The following command disables ARP learning on port 1:1 of the VLAN learn:

```
disable ip-security arp learning learn-from-arp vlan learn ports 1:1
```

disable ip-security arp learning learn-from-dhcp

```
disable ip-security arp learning learn-from-dhcp {vlan} <vlan_name> ports [all | <ports>]
```

Description

Disables DHCP secured ARP learning for the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
ports	Specifies one or more ingress ports.

Default

By default, DHCP secured ARP learning is disabled.

Usage Guidelines

Use this command to disable DHCP secured ARP learning.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} <vlan_name>
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {<ip_addre> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

Example

The following command disables DHCP secured ARP learning on port 1:1 of the VLAN learn:

```
disable ip-security arp learning learn-from-dhcp vlan learn ports 1:1
```

disable ip-security arp validation

```
disable ip-security arp validation {vlan} <vlan_name> [all | <ports>]
```

Description

Disables ARP validation for the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ports.
ports	Specifies one or more ports.

Default

By default, ARP validation is disabled.

Usage Guidelines

Use this command to disable ARP validation.

Displaying ARP Validation Information

To display information about ARP validation, use the following command:

```
show ip-security arp validation {vlan} <vlan_name>
```

Example

The following command disables ARP validation on port 1:1 of the VLAN valid:

```
disable ip-security arp validation vlan valid ports 1:1
```

disable ip-security dhcp-bindings restoration

```
disable ip-security dhcp-bindings restoration
```

Description

Disables the download and upload of DHCP bindings.

Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The command allows you to disable the download and upload of the DHCP bindings, essentially disabling the DHCP binding functionality. The default is disabled.

disable ip-security dhcp-snooping

```
disable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>]
```

Description

Disables DHCP snooping on the switch.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN.
all	Specifies all ports to stop receiving DHCP packets.
ports	Specifies one or more ports to stop receiving DHCP packets.

Default

By default, DHCP snooping is disabled

Usage Guidelines

Use this command to disable DHCP snooping on the switch.

Example

The following command disables DHCP snooping on the switch:

```
disable ip-security dhcp-snooping vlan snoop ports 1:1
```

disable ip-security source-ip-lockdown ports

```
disable ip-security source-ip-lockdown ports [all | <ports>]
```

Description

Disables the source IP lockdown feature on one or more ports.

Syntax Description

all	Specifies all ports for which source IP lockdown should be disabled.
ports	Specifies one or more ports for which source IP lockdown should be disabled.

Default

By default, source IP lockdown is disabled on the switch.

Usage Guidelines

To display the source IP lockdown configuration on the switch, use the following command:

```
show ip-security source-ip-lockdown
```

Example

The following command disables source IP lockdown on ports 1:1 and 1:4:

```
disable ip-security source-ip-lockdown ports 1:1, 1:4
```

disable mac-lockdown-timeout ports

```
disable mac-lockdown-timeout ports [all | <port_list>]
```

Description

Disables the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.

Default

By default, the MAC address lock down feature is disabled.

Usage Guidelines

If you disable the MAC lock down timer on a port, existing MAC address entries for the port will time out based on the FDB aging period.

Example

The following command disables the MAC address lock down timer set for ports 2:3 and 2:4:

```
disable mac-lockdown-timeout ports 2:3, 2:4
```

disable radius

```
disable radius {mgmt-access | netlogin}
```

Description

Disables the RADIUS client.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

Default

RADIUS authentication is disabled for both switch management and network login by default.

Usage Guidelines

Use the `mgmt-access` keyword to disable RADIUS authentication for switch management functions.

Use the `netlogin` keyword to disable RADIUS authentication for network login.

If you do not specify a keyword, RADIUS authentication is disabled on the switch for both management and network login.

Example

The following command disables RADIUS authentication on the switch for both management and network login:

```
disable radius
```

The following command disables RADIUS authentication on the switch for network login:

```
disable radius netlogin
```

disable radius-accounting

```
disable radius-accounting {mgmt-access | netlogin}
```

Description

Disables RADIUS accounting.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.

Default

RADIUS accounting is disabled for both switch management and network login by default.

Usage Guidelines

Use the `mgmt-access` keyword to disable RADIUS accounting for switch management functions.

Use the `netlogin` keyword to disable RADIUS accounting for network login.

If you do not specify a keyword, RADIUS accounting is disabled on the switch for both management and network login.

Example

The following command disables RADIUS accounting on the switch for both management and network login:

```
disable radius-accounting
```

The following command disables RADIUS accounting on the switch for network login:

```
disable radius-accounting netlogin
```

disable ssh2

```
disable ssh2
```

Description

Disables the SSH2 server for incoming SSH2 sessions to switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

SSH2 options (non-default port setting) are not saved when SSH2 is disabled.

To view the status of SSH2 on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2.

Example

The following command disables the SSH2 server:

```
disable ssh2
```

disable tacacs

```
disable tacacs
```

Description

Disables TACACS+ authentication.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables TACACS+ authentication for the switch:

```
disable tacacs
```

disable tacacs-accounting

```
disable tacacs-accounting
```

Description

Disables TACACS+ accounting.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables TACACS+ accounting:

```
disable tacacs-accounting
```

disable tacacs-authorization

```
disable tacacs-authorization
```

Description

Disables TACACS+ authorization.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This disables CLI command authorization but leaves user authentication enabled.

Example

The following command disables TACACS+ CLI command authorization:

```
disable tacacs-authorization
```

disable web http

```
disable web http
```

Description

Disables the hypertext transfer protocol (HTTP) access to the switch on the default port (80).

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to disallow users from connecting with HTTP. Disabling HTTP access forces user to use a secured HTTPS connection if web HTTPS is enabled.

Use the following command to enable web HTTPS:

```
enable web https
```

Example

The following command disables HTTP on the default port:

```
disable web http
```

disable web https

```
disable web https
```

Description

Disables the secure socket layer (SSL) access to the switch on the default port (443).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable SSL before changing the certificate or private key.

Example

The following command disables SSL on the default port:

```
disable web https
```

download ssl certificate

```
download ssl <ip_address> certificate <cert file>
```

Description

Permits downloading of a certificate key from files stored in a TFTP server.

Syntax Description

ip_address	Specifies the IP address of the TFTP server.
cert file	Specifies the name of the certificate key.

Default

N/A.

Usage Guidelines

If the download operation is successful, any existing certificate is overwritten. After a successful download, the software attempts to match the public key in the certificate against the private key stored. If the private and public keys do not match, the switch displays a warning message similar to the following: `Warning: The Private Key does not match with the Public Key in the certificate.` This warning acts as a reminder to also download the private key.

Note: You can only download a certificate key in the VR-Mgmt virtual router.

Downloaded certificates and keys are not saved across switch reboots unless you save your current switch configuration. Once you issue the save command, the downloaded certificate is stored in the configuration file and the private key is stored in the EEPROM.

Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for remote IP addresses.

When specifying a remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)

- Numerals (0-9)
- Period (.)
- Colon (:)

When configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command downloads a certificate from a TFTP server with the IP address of 123.45.6.78:

```
download ssl 123.45.6.78 certificate g0ethner1
```

download ssl privkey

```
download ssl <ip_address> privkey <key file>
```

Description

Permits downloading of a private key from files stored in a TFTP server.

Syntax Description

ip_address	Specifies the IP address of the TFTP server.
key file	Specifies the name of the private key file.

Default

N/A.

Usage Guidelines

If the operation is successful, the existing private key is overwritten.

After a successful download, a check is performed to find out whether the private key downloaded matches the public key stored in the certificate. If the private and public keys do not match, the switch displays a warning similar to the following: `Warning: The Private Key does not match with the Public Key in the certificate.` This warning acts as a reminder to also download the corresponding certificate.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Downloaded certificates and keys are not saved across switch reboots unless you save your current switch configuration. Once you issue the `save` command, the downloaded certificate is stored in the configuration file and the private key is stored in the EEPROM.

Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for remote IP addresses.

When specifying a remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Colon (:)

When configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command downloads a private key from a TFTP server with the IP address of 123.45.6.78:

```
download ssl 123.45.6.78 privkey t00Tsle
```

enable dhcp ports vlan

```
enable dhcp ports <portlist> vlan <vlan_name>
```

Description

Enables DHCP on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which DHCP should be enabled.
vlan_name	Specifies the VLAN on whose ports DHCP should be enabled.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables DHCP for port 5:9 in VLAN *corp*:

```
enable dhcp ports 5:9 vlan corp
```

enable dos-protect

```
enable dos-protect
```

Description

Enables denial of service protection.

Syntax Description

This command has no arguments or variables.

Default

The default is disabled.

Usage Guidelines

None.

Example

The following command enables denial of service protection.

```
enable dos-protect
```

enable dos-protect simulated

```
enable dos-protect simulated
```

Description

Enables simulated denial of service protection.

Syntax Description

This command has no arguments or variables.

Default

The default is disabled.

Usage Guidelines

If simulated denial of service is enabled, no ACLs are created. This mode is useful to gather information about normal traffic levels on the switch. This will assist in configuring denial of service protection so that legitimate traffic is not blocked.

Example

The following command enables simulated denial of service protection.

```
enable dos-protect simulated
```

enable iparp gratuitous protect

```
enable iparp gratuitous protect vlan <vlan-name>
```

Description

Enables gratuitous ARP protection on the specified VLAN.

Syntax Description

vlan-name	Specifies the VLAN.
-----------	---------------------

Default

By default, gratuitous ARP is disabled.

Usage Guidelines

The `enable ip-security arp gratuitous-protection` command replaces this command for configuring gratuitous ARP.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

Example

The following command enables gratuitous ARP protection for VLAN *corp*:

```
enable iparp gratuitous protect vlan corp
```

enable ip-option loose-source-route

```
enable ip-option loose-source-route
```

Description

Enables processing of the loose source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This enables the switch to forward IP packets that have the loose source route IP option (0x83) enabled.

Source routing is used when a sending host specifies the router interfaces that the packet must traverse on its way to its destination.

With loose source routing enabled, the packet is forwarded if the routing table has a reverse path to the source IP address of the packet.

Example

The following command enables processing of the loose source route IP option:

```
enable ip-option loose-source-route
```

enable ip-security anomaly-protection

```
enable ip-security anomaly-protection {slot [ <slot> | all ]}
```

Description

Enables all anomaly checking options.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This commands enables all anomaly checking options, including IP address, UDP/TCP port, TCP flag and fragment, and ICMP anomaly checking.

enable ip-security anomaly-protection icmp

```
enable ip-security anomaly-protection icmp {slot [ <slot> | all ]}
```

Description

Enables ICMP size and fragment checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables ICMP size and fragment checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops ICMP packets if one of following condition is true:

- Fragmented ICMP packets.
- IPv4 ICMP pings packets with payload size greater than the maximum IPv4 ICMP-allowed size. (The maximum allowed size is configurable.)
- IPv6 ICMP ping packets with payload size > the maximum IPv6 ICMP-allowed size. (The maximum allowed size is configurable.)

enable ip-security anomaly-protection ip

```
enable ip-security anomaly-protection ip { slot [ <slot> | all ] }
```

Description

Enables source and destination IP address checking.

Syntax Description

slot	Specifies the slot.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command enables source and destination IP addresses checking. This checking takes effect for both IPv4 and IPv6 packets. When enabled, the switch drops IPv4/IPv6 packets if its source IP address are the same as the destination IP address. In most cases, the condition of source IP address being the same as the destination IP address indicates a Layer 3 protocol error. (These kind of errors are found in LAND attacks.)

enable ip-security anomaly-protection l4port

```
enable ip-security anomaly-protection l4port {slot [ <slot> | all ]}
```

Description

Enables TCP and UDP ports checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command enabled TCP and UDP ports checking. This checking takes effect for both IPv4 and IPv6 TCP and UDP packets. When enabled, the switch drops TCP and UDP packets if its source port is the same as its destination port. In most cases, when the condition of source port is the same as that of the destination port, it indicates a Layer 4 protocol error. (This type of error can be found in a BALT attack.)

enable ip-security anomaly-protection notify

```
enable ip-security anomaly-protection notify [log | snmp | cache] {slot [ <slot> | all ]}
```

Description

Enables protocol anomaly notification.

Syntax Description

log	Specifies the switch to send the notification to a log file.
snmp	Specifies the switch to send an SNMP trap when an event occurs.
cache	Specifies the switch to send the notification to cache.
slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command enables anomaly notification. When enabled, any packet failed to pass enabled protocol checking is sent to XOS Host CPU and notifies the user. There are three different types of notifications:

- `log`: The anomaly events are logged into EMS log.
- `snmp`: The anomaly events generate SNMP traps.

- `cache`: The most recent and unique anomaly events are stored in memory for review and investigation.

When disabled, the switch drops all violating packets silently.

enable ip-security anomaly-protection tcp flags

```
enable ip-security anomaly-protection tcp flags {slot [ <slot> | all ]}
```

Description

Enables TCP flag checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command Enables TCP flag checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops TCP packets if one of following condition is true:

- TCP SYN flag==1 and the source port<1024
- TCP control flag==0 and the sequence number==0
- TCP FIN, URG, and PSH bits are set, and the sequence number==0
- TCP SYN and FIN both are set.

enable ip-security anomaly-protection tcp fragment

```
enable ip-security anomaly-protection tcp fragment {slot [ <slot> | all ]}
```

Description

Enables TCP fragment checking.

Syntax Description

slot	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is `disabled`.

Usage Guidelines

This command enables TCP fragment checking. This checking takes effect for IPv4/IPv6. When it is enabled, the switch drops TCP packets if one of following condition is true:

- For the first IPv4 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv4 TCP header allowed size
- For the first IPv6 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv6 TCP header allowed size
- If its IP offset field==1 (for IPv4 only)

enable ip-security arp gratuitous-protection

```
enable ip-security arp gratuitous-protection {vlan} [all | <vlan_name>]
```

Description

Enables gratuitous ARP protection on one or all VLANs on the switch.

Syntax Description

<code>all</code>	Specifies all VLANs configured on the switch.
<code>vlan-name</code>	Specifies the VLAN.

Default

By default, gratuitous ARP protection is disabled.

Usage Guidelines

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

If you enable both DHCP secured ARP and gratuitous ARP protection, the switch protects its own IP address and those of the hosts that appear as secure entries in the ARP table.

To protect the IP addresses of the hosts that appear as secure entries in the ARP table, use the following commands to enable DHCP snooping, DHCP secured ARP, and gratuitous ARP on the switch:

- `enable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>]
violation-action [drop-packet {[block-mac | block-port] [duration
<duration_in_seconds> | permanently] | none}] {snmp-trap}`
- `enable ip-security arp learning learn-from-dhcp {vlan} <vlan_name> ports [all | <ports>]`
- `enable ip-security arp gratuitous-protection {vlan} [all | <vlan_name>]`

Displaying Gratuitous ARP Information

To display information about gratuitous ARP, use the following command:

```
show ip-security arp gratuitous-protection
```

Example

The following command enables gratuitous ARP protection for VLAN *corp*:

```
enable ip-security arp gratuitous-protectection vlan corp
```

enable ip-security arp learning learn-from-arp

```
enable ip-security arp learning learn-from-arp {vlan} <vlan_name> ports [all | <ports>]
```

Description

Enables ARP learning for the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
ports	Specifies one or more ingress ports.

Default

By default, ARP learning is enabled.

Usage Guidelines

ARP is part of the TCP/IP suite used to associate a device's physical address (MAC address) with its logical address (IP address). The switch broadcasts an ARP request that contains the IP address, and the device with that IP address sends back its MAC address so that traffic can be transmitted across the network. The switch maintains an ARP table (also known as an ARP cache) that displays each MAC address and its corresponding IP address.

By default, the switch builds its ARP table by tracking ARP requests and replies, which is known as ARP learning.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} <vlan_name>
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {<ip_addre> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

Example

The following command enables ARP learning on port 1:1 of the VLAN learn:

```
enable ip-security arp learning learn-from-arp vlan learn ports 1:1
```

enable ip-security arp learning learn-from-dhcp

```
enable ip-security arp learning learn-from-dhcp {vlan} <vlan_name> ports [all | <ports>]
```

Description

Enables DHCP secured ARP learning for the specified VLAN and member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
ports	Specifies one or more ingress ports.

Default

By default, DHCP secured ARP learning is disabled.

Usage Guidelines

Use this command to configure the switch to add the MAC address and its corresponding IP address to the ARP table as a secure ARP entry. The switch does not update secure ARP entries, regardless of the ARP requests and replies seen by the switch. DHCP secured ARP is linked to the “DHCP snooping” feature. The same DHCP bindings database created when you enabled DHCP snooping is also used by DHCP secured ARP to create secure ARP entries. The switch only removes secure ARP entries when the corresponding DHCP entry is removed from the trusted DHCP bindings database.

Note: If you enable DHCP secured ARP on the switch, ARP learning continues, which allows insecure entries to be added to the ARP table.

The default ARP timeout (`configure iparp timeout`) and ARP refresh (`enable iparp refresh`) settings do not apply to DHCP secured ARP entries. The switch removes DHCP secured ARP entries upon any DHCP release packet received from the DHCP client.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} <vlan_name>
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {<ip_addre> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

Example

The following command enables DHCP secured ARP learning on port 1:1 of the VLAN learn and uses the default polling and retry intervals:

```
enable ip-security arp learning learn-from-dhcp vlan learn ports 1:1
```

enable ip-security arp validation violation-action

```
enable ip-security arp validation {destination-mac} {source-mac} {ip} {vlan} <vlan_name> [all  
| <ports>] violation-action [drop-packet {[block-port] [duration <duration_in_seconds> |  
permanently]]] {snmp-trap}
```

Description

Enables ARP validation for the specified VLAN and member ports.

Syntax Description

destination-mac	Specifies that the switch checks the ARP payload for the MAC destination address in the Ethernet header and the receiver's host address in the ARP response.
source-mac	Specifies that the switch checks ARP requests and responses for the MAC source address in the Ethernet header and the sender's host address in the ARP payload.
ip	Specifies the switch checks the IP address in the ARP payload and compares it to the DHCP bindings database. If the IP address does exist in the DHCP bindings table, the switch verifies that the MAC address is the same as the sender hardware address in the ARP request. If not, the packet is dropped.

vlan_name	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ports to participate in ARP validation.
ports	Specifies one or more ports to participate in ARP validation.
drop-packet	Specifies that the switch drops the invalid ARP packet.
block-port	Indicates that the switch blocks invalid ARP requests on the specified port.
permanently	Specifies the switch to permanently disable the port upon receiving an invalid ARP request.
duration_in_seconds	Specifies the switch to temporarily disable the specified port upon receiving an invalid ARP request. The range is seconds.
snmp-trap	Specifies the switch to send an SNMP trap when an event occurs.

Default

By default, ARP validation is disabled.

Usage Guidelines

The violation action setting determines what action(s) the switch takes when an invalid ARP is received.

Depending on your configuration, the switch uses the following methods to check the validity of incoming ARP packets:

- **Drop packet**—The switch confirms that the MAC address and its corresponding IP address are in the DHCP binding database built by DHCP snooping. This is the default behavior when you enable ARP validation. If the MAC address and its corresponding IP address are in the DHCP bindings database, the entry is valid. If the MAC address and its corresponding IP address are not in the DHCP bindings database, the entry is invalid, and the switch drops the ARP packet.
- **IP address**—The switch checks the IP address in the ARP payload. If the switch receives an IP address in the ARP payload that is in the DHCP binding database, the entry is valid. If the switch receives an IP address that is not in the DHCP binding database, for example 255.255.255.255 or an IP multicast address, the entry is invalid or unexpected.
- **Source MAC address**—The switch checks ARP requests and responses for the source MAC address in the Ethernet header and the sender's host address in the ARP payload. If the source MAC address and sender's host address are the same, the entry is valid. If the source MAC source and the sender's host address are different, the entry is invalid.
- **Destination MAC address**—The switch checks the ARP payload for the destination MAC address in the Ethernet header and the receiver's host address. If the destination MAC address and the target's host address are the same, the entry is valid. If the destination MAC address and the target's host address are different, the entry is invalid.

Any violation that occurs causes the switch to generate an Event Management System (EMS) log message. You can configure to suppress the log messages by configuring EMS

log filters. For more information about EMS, see the EMS commands in [Chapter 8, “Commands for Status Monitoring and Statistics.”](#)

Displaying ARP Validation Information

To display information about ARP validation, use the following command:

```
show ip-security arp validation {vlan} <vlan_name>
```

Example

The following command enables ARP validation on port 1:1 of the VLAN valid:

```
enable ip-security arp validation vlan valid ports 1:1 drop-packet
```

enable ip-security dhcp-bindings restoration

```
enable ip-security dhcp-bindings restoration
```

Description

Enables download and upload of DHCP bindings.

Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The command allows you to enable the download and upload of the DHCP bindings, essentially enabling the DHCP binding functionality. The default is disabled.

enable ip-security dhcp-snooping

```
enable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>] violation-action
[drop-packet {[block-mac | block-port] [duration <duration_in_seconds> | permanently] |
none}}] {snmp-trap}
```

Description

Enables DHCP snooping for the specified VLAN and ports.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN. Create and configure the VLAN before enabling DHCP snooping.
all	Specifies all ports to receive DHCP packets.

ports	Specifies one or more ports to receive DHCP packets.
drop-packet	Indicates that the switch drop the rogue DHCP packet received on the specified port.
block-mac	Indicates that the switch blocks rogue DHCP packets from the specified MAC address on the specified port. The MAC address is added to the DHCP bindings database.
block-port	Indicates that the switch blocks rogue DHCP packets on the specified port. The port is added to the DHCP bindings database.
duration_in_seconds	Specifies that the switch temporarily disable the specified port upon receiving a rogue DHCP packet. The range is seconds.
permanently	Specifies that the switch to permanently disable the specified port upon receiving a rogue DHCP packet.
none	Specifies that the switch takes no action when receiving a rogue DHCP packet; the switch does not drop the packet.
snmp-trap	Specifies the switch to send an SNMP trap when an event occurs.

Default

By default, DHCP snooping is disabled.

Usage Guidelines

Use this command to enable DHCP snooping on the switch.

Note: Snooping IP fragmented DHCP packets is not supported.

The violation action setting determines what action(s) the switch takes when a rouge DHCP server packet is seen on an untrusted port or the IP address of the originating server is not among those of the configured trusted DHCP servers. The DHCP server packets are DHCP OFFER, ACK and NAK. The following list describes the violation actions:

- **block-mac**—The switch automatically generates an ACL to block the MAC address on that port. The switch does not blackhole that MAC address in the FDB. The switch can either temporarily or permanently block the MAC address.
- **block-port**—The switch blocks all incoming rogue DHCP packets on that port. The switch disables the port either temporarily or permanently to block the traffic on that port.
- **none**—The switch takes no action to drop the rogue DHCP packet or block the port, and so on. In this case, DHCP snooping continues to build and manage the DHCP bindings database and DHCP forwarding will continue in hardware as before.

Any violation that occurs causes the switch to generate an Event Management System (EMS) log message. You can configure to suppress the log messages by configuring EMS

log filters. For more information about EMS, see the EMS commands in [Chapter 8, “Commands for Status Monitoring and Statistics.”](#)

Displaying DHCP Snooping Information

To display the DHCP snooping configuration settings, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display the DHCP bindings database, use the following command:

```
show ip-security dhcp-snooping entries {vlan} <vlan_name>
```

To display any violations that occur, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command enables DHCP snooping on the switch and has the switch block DHCP packets from port 1:1:

```
enable ip-security dhcp-snooping vlan snoop ports 1:1 violation-action drop-packet block-port
```

enable ip-security source-ip-lockdown ports

```
enable ip-security source-ip-lockdown ports [all | <ports>]
```

Description

Enables the source IP lockdown feature on one or more ports.

Syntax Description

all	Specifies all ports for which source IP lockdown should be enabled.
ports	Specifies one or more ports for which source IP lockdown should be enabled.

Default

By default, source IP lockdown is disabled on the switch.

Usage Guidelines

Source IP lockdown prevents IP address spoofing by automatically placing source IP address filters on specified ports. If configured, source IP lockdown allows only traffic from a valid DHCP-assigned address obtained by a DHCP snooping-enabled port or an authenticated static IP address to enter the network.

To configure source IP lockdown, you must enable DHCP snooping on the ports connected to the DHCP server and DHCP client before you enable source IP lockdown. You must enable source IP lockdown on the ports connected to the DHCP client, not on the ports connected to the DHCP server. The same DHCP bindings database created when you enable DHCP

snooping is also used by the source IP lockdown feature to create ACLs that permit traffic from DHCP clients. All other traffic is dropped. In addition, the DHCP snooping violation action setting determines what action(s) the switch takes when a rouge DHCP server packet is seen on an untrusted port.

To enable DHCP snooping, use the following command:

```
enable ip-security dhcp-snooping {vlan} <vlan_name> ports [all | <ports>]
violation-action [drop-packet {[block-mac | block-port] [duration <duration_in_seconds>
| permanently] | none}]] {snmp-trap}
```

Displaying Source IP Lockdown Information

To display the source IP lockdown configuration on the switch, use the following command:

```
show ip-security source-ip-lockdown
```

Example

The following command enables source IP lockdown on ports 1:1 and 1:4:

```
enable ip-security source-ip-lockdown ports 1:1, 1:4
```

enable mac-lockdown-timeout ports

```
enable mac-lockdown-timeout ports [all | <port_list>]
```

Description

Enables the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.

Default

By default, the MAC address lock down timeout feature is disabled.

Usage Guidelines

You cannot enable the MAC lock down timer on a port that also has the lock learning feature enabled.

Example

The following command enables the MAC address lock down timeout feature for ports 2:3, 2:4, and 2:6:

```
enable mac-lockdown-timeout ports 2:3, 2:4, 2:6
```

enable radius

```
enable radius {mgmt-access | netlogin}
```

Description

Enables the RADIUS client on the switch.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

Default

RADIUS authentication is disabled for both switch management and network login by default.

Usage Guidelines

Before you enable RADIUS on the switch, you must configure the servers used for authentication and configure the authentication string (shared secret) used to communicate with the RADIUS authentication server.

To configure the RADIUS authentication servers, use the following command:

```
configure radius {mgmt-access | netlogin} [primary | secondary] server [<ipaddress> | <hostname>] [<udp_port>] client-ip [<ipaddress>] {vr <vr_name>}
```

To configure the shared secret, use the following command:

```
configure radius {mgmt-access | netlogin} [primary | secondary] shared-secret {encrypted} <string>
```

If you do not specify a keyword, RADIUS authentication is enabled on the switch for both management and network login. When enabled, all web, Telnet, and SSH logins are sent to the RADIUS servers for authentication. When used with a RADIUS server that supports NETGEAR 8800 CLI authorization, each CLI command is sent to the RADIUS server for authorization before it is executed.

Use the `mgmt-access` keyword to enable RADIUS authentication for switch management functions.

Use the `netlogin` keyword to enable RADIUS authentication for network login.

Example

The following command enables RADIUS authentication on the switch for both management and network login:

```
enable radius
```

The following command enables RADIUS authentication on the switch for network login:

```
enable radius netlogin
```

enable radius-accounting

```
enable radius-accounting {mgmt-access | netlogin}
```

Description

Enables RADIUS accounting.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.

Default

RADIUS accounting is disabled for both switch management and network login by default.

Usage Guidelines

The RADIUS client must also be enabled.

Before you enable RADIUS accounting on the switch, you must configure the servers used for accounting and configure the authentication string (shared secret) used to communicate with the RADIUS accounting server.

To configure the RADIUS accounting servers, use the following command:

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] server
[<ipaddress> | <hostname>] {<tcp_port>} client-ip [<ipaddress>] {vr <vr_name>}
```

To configure the shared secret, use the following command:

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] shared-secret
{encrypted} <string>
```

If you do not specify a keyword, RADIUS accounting is enabled on the switch for both management and network login.

Use the `mgmt-access` keyword to enable RADIUS accounting for switch management functions.

Use the `netlogin` keyword to enable RADIUS accounting for network login.

Example

The following command enables RADIUS accounting on the switch for both management and network login:

```
enable radius-accounting
```

The following command enables RADIUS accounting for network login:

```
enable radius-accounting netlogin
```

enable ssh2

```
enable ssh2 {access-profile [<access_profile> | none]} {port <tcp_port_number>} {vr
[<vr_name> | all | default]}
```

Description

Enables SSH2 server to accept incoming sessions from SSH2 clients.

Syntax Description

access_profile	Specifies an ACL policy.
none	Cancels a previously configured ACL policy.
port	Specifies a TCP port number. The default is port 22.
vr_name	Specifies a virtual router name.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .
all	Specifies that SSH is enabled on all virtual routers.
default	Specifies that SSH is enabled on the default virtual router.

Default

The SSH2 feature is disabled by default.

Usage Guidelines

SSH2 enables the encryption of session data. You must be logged in as an administrator to enable SSH2.

Before you use SSH2, you must generate a host key and enable SSH2. To generate an SSH2 host key, use the `configure ssh2 key` command. To enable SSH2, use the `enable ssh2` command.

Use the `port` option to specify a TCP port number other than the default port of 22. You can only specify ports 22 and 1024 through 65535.

Using ACLs to Control SSH Access

You can specify a list of predefined clients that are allowed SSH2 access to the switch. To do this, you configure an ACL policy to permit or deny a specific list of IP addresses and subnet masks for the SSH port. You must create an ACL policy file before you can use the `access-profile` option. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

Use the `none` option to cancel a previously configured ACL.

In the ACL policy file for SSH2, the `source-address` field is the only supported match condition. Any other match conditions are ignored.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see the chapters on Policy Manager and ACLs in the *NETGEAR 8800 User Manual*.

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile_2.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `ls` command. If the policy does not exist, create the ACL policy file.

Viewing SSH Information

To view the status of SSH2 sessions on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2 sessions and whether a valid key is present.

Example

The following command enables the SSH2 feature:

```
enable ssh2
```

The next example assumes you have already created an ACL to apply to SSH.

The following command applies the ACL `MyAccessProfile_2` to SSH:

```
enable ssh2 access-profile MyAccessProfile_2
```

enable tacacs

```
enable tacacs
```

Description

Enables TACACS+ authentication.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

After they have been enabled, all web and Telnet logins are sent to one of the two TACACS+ servers for login name authentication.

Example

The following command enables TACACS+ user authentication:

```
enable tacacs
```

enable tacacs-accounting

```
enable tacacs-accounting
```

Description

Enables TACACS+ accounting.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If accounting is used, the TACACS+ client must also be enabled.

Example

The following command enables TACACS+ accounting for the switch:

```
enable tacacs-accounting
```

enable tacacs-authorization

```
enable tacacs-authorization
```

Description

Enables CLI command authorization.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When enabled, each command is transmitted to the remote TACACS+ server for authorization before the command is executed. TACACS+ authentication must also be enabled to use TACACS+ authorization. Use the following command to enable authentication:

```
enable tacacs
```

Example

The following command enables TACACS+ command authorization for the switch:

```
enable tacacs-authorization
```

enable web http

```
enable web http
```

Description

Enables hypertext transfer protocol (HTTP) access to the switch on the default HTTP port (80).

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

If HTTP access has been disabled, use this command to enable HTTP access to the switch.

Example

The following command enables HTTP on the default port:

```
enable web http
```

enable web https

```
enable web https
```

Description

Enables secure socket layer (SSL) access to the switch on the default port (443).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to allow users to connect using a more secure HTTPS connection.

To use secure HTTP access (HTTPS) for web-based login connections, you must specify HTTPS as the protocol when configuring the redirect URL. For more information about configuring the redirect URL, see the [configure netlogin redirect-page](#) command.

Example

The following command enables SSL on the default port:

```
enable web https
```

scp2

```
scp2 {vr <vr_name>} {cipher [3des | blowfish]} {port <portnum>} <user>@ [<hostname> | <ipaddress>]:<remote_file> <local_file>
```

or

```
scp2 {vr <vr_name>} {cipher [3des | blowfish]} {port <portnum>} <local_file> <user>@ [<hostname> | <ipaddress>]:<remote_file>
```

Description

The first command initiates an SCP2 client session to a remote SCP2 server and copies a configuration or policy file from the remote system to the switch.

The second command initiates an SCP2 client session to a remote SCP2 server and copies a configuration or policy file from the switch to a remote system.

Syntax Description

vr_name	Specifies the virtual router. The default virtual router is VR-Mgmt. Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .
3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
portnum	Specifies the TCP port number to be used for communicating with the SSH2 client. The default is port 22.
user	Specifies a login name for the remote host.
hostname	Specifies the name of the remote host.
ipaddress	Specifies the IP address of the remote host.

remote_file	Specifies the name of the remote file (configuration file, policy file, image file, public key file) to be transferred.
local_file	Specifies the name of the local file (configuration file, policy file, image file, public key file) to be transferred.

Default

The default settings for SSH2 parameters are as follows:

- cipher—3des encryption
- port—22
- compression—off
- vr_name—VR-Mgmt

Usage Guidelines

You must be running the SSH2 module (ssh.xmod), which is under Export Control, in order to use the SCP2 command.

SSH2 does not need to be enabled on the switch in order to use this command.

This command logs into the remote host as `<user>` and accesses the file `<remote_file>`. You will be prompted for a password from the remote host, if required.

Host Name, User Name, and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name, user name, or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted for host and user names
- Underscore (_) Permitted for host and user names
- Colon (:)
- At symbol (@) Permitted only for user names
- Slash (/) Permitted only for user names

When naming the host, creating a user name, or configuring the IP address, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command copies the configuration file test.cfg on host system1 to the switch:

```
scp2 admin@system1:test.cfg localtest.cfg
```

The following command copies the configuration file engineering.cfg from the switch to host system1:

```
scp2 engineering.cfg admin@system1:engineering.cfg
```

The following command copies the file Anna5.xsf from the default virtual router to 150.132.82.140:

```
scp2 vr vr-default Anna5.xsf root@150.132.82.140:Anna5.xsf
Upload /config/Anna5.xsf to
Connecting to 150.132.82.140...
```

show dhcp-server

```
show dhcp-server {vlan <vlan_name>}
```

Description

Displays the DHCP server's configuration and address allocation on a specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server of interest.
-----------	--

Default

N/A.

Usage Guidelines

If no VLAN is specified, the configuration and address allocation for the servers on all the VLANs is displayed.

Example

The following command displays the configuration and address allocation for the DHCP server for the VLAN *test*:

```
show dhcp-server vlan test
```

The following is sample output from this command:

```
DHCP Address Range    : 10.10.10.100->10.10.10.200
Netlogin Lease Timer  : Not configured (Default = 10 seconds)
DHCP Lease Timer      : Not configured (Default = 7200 seconds)
Primary DNS Server    : 1.1.1.1
Secondary DNS Server  : 2.2.2.2
Ports DHCP Enabled    : 23
```

show dos-protect

```
show dos-protect {detail}
```

Description

Displays DoS protection configuration and state.

Syntax Description

detail	Specifies to display statistics in addition to configuration and state.
--------	---

Default

N/A.

Usage Guidelines

Use this command to display the DoS protection settings. Using the `detail` option will also display the following cumulative statistics:

- trusted
- notify
- alerts

Example

The following command displays the DoS protection settings for the switch:

```
show dos-protect
```

The following is sample output from this command:

```
dos-protect is disabled
```

```
dos-protect settings:
```

```
interval:          1 (measurement interval secs)
acl expire time:   5 (secs)
```

```
trusted ports:
    no trusted ports configured
```

```
type L3-Protect:
    notify threshold: 3500 (level to log a message)
    alert threshold:  4000 (level to generate an ACL)
```

The following command displays detailed DoS protection settings for the switch:

```
show dos-protect detail
```

The following is sample output from this command:

```
dos-protect is enabled
```

```
dos-protect settings:
interval:          1 (measurement interval secs)
acl expire time:   5 (secs)
```

```
trusted ports:
    1:2
```

```
type L3-Protect:
    notify threshold: 3500 (level to log a message)
    alert threshold:  4000 (level to generate an ACL)
```

```
dos-protect statistics:
    trusted: 1301
    notify:  0
    alerts:  0
```

show ip-security anomaly-protection notify cache ports

```
show ip-security anomaly-protection notify cache ports <port list>
```

Description

Displays most anomaly notification caches.

Syntax Description

port list	Specifies one or more ports or slots and ports.
-----------	---

Default

N/A.

Usage Guidelines

This command displays most anomaly notification caches.

show ip-security arp gratuitous-protection

```
show ip-security arp gratuitous-protection
```

Description

If configured for gratuitous ARP, displays the gratuitous ARP protection configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The switch displays the name of each VLAN configured for gratuitous ARP.

If you do not have gratuitous ARP configured, the switch does not display any VLAN information.

Example

The following command displays the gratuitous ARP configuration on the switch:

```
show ip-security arp gratuitous-protection
```

The following is sample output from this command:

```
Gratuitous ARP Protection enabled on following VLANs:
Default, test
```

show ip-security arp learning

```
show ip-security arp learning {vlan} <vlan_name>
```

Description

Displays how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports.

Syntax Description

vlan_name	Specifies the name of the VLAN.
-----------	---------------------------------

Default

N/A.

Usage Guidelines

The switch displays the following ARP learning information:

- Port—The member port of the VLAN.
- Learn-from—The method the port uses to build the ARP table. The methods are:
 - ARP—ARP learning is enabled. The switch uses a series of requests and replies to build the ARP table.
 - DHCP—DHCP secured ARP is enabled. The switch uses DHCP snooping to build the ARP table.
 - None—Both DHCP secured ARP and ARP learning are disabled.

Example

The following command displays how the switch builds its ARP table for the VLAN learn:

```
show ip-security arp learning vlan learn
```

The following is sample output from this command:

Port	Learn-from
2:1	ARP
2:2	DHCP, poll 300 sec, retries 3
2:3	ARP
2:4	None
2:5	ARP
2:6	ARP
2:7	ARP
2:8	ARP

show ip-security arp validation

```
show ip-security arp validation {vlan} <vlan_name>
```

Description

Displays ARP validation information for the specified VLAN.

Syntax Description

vlan_name	Specifies the name of the VLAN.
-----------	---------------------------------

Default

N/A.

Usage Guidelines

The switch displays the following ARP validation information:

- Port—Indicates the port that received the ARP entry.
- Validation—Indicates how the entry is validated.
- Violation-action—Determines what action(s) the switch takes when an invalid ARP is received.

Example

The following command displays ARP validation on for the VLAN valid:

```
show ip-security arp validation vlan valid
```

The following is sample output from this command:

Port	Validation	Violation-action
7	DHCP	drop-packet, block-port for 120 seconds, snmp-trap
23	DHCP	drop-packet, block-port for 120 seconds, snmp-trap

show ip-security arp validation violations

```
show ip-security arp validation violations {vlan} <vlan_name> ports [<ports> | all]
```

Description

Displays the violation count on an ARP validation.

Syntax Description

vlan_name	Specifies the name of the VLAN.
ports	Specifies the name of the port.
all	Specifies all ports.

Default

N/A.

Usage Guidelines

The switch displays the following ARP validation information:

- Port—Indicates the port that received the ARP entry.

- Validation—Indicates how the entry is validated.
- Violation count—Indicates the number of violations for each port.

Example

The following command displays ARP validation violation counts on all ports:

```
show ip-security arp validation violations ragu ports all
```

The following is sample output from this command:

Port	Validation	Violation Count
1:1	ip,DHCP	1233
1:3	ip,DHCP	3425
1:4	ip,DHCP	5654
1:5	ip,DHCP	0
1:6	ip,DHCP	3645

show ip-security dhcp-snooping entries

```
show ip-security dhcp-snooping entries {vlan} <vlan_name>
```

Description

Displays the DHCP bindings database on the switch.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN.
-----------	---

Default

N/A.

Usage Guidelines

The switch displays the following DHCP bindings database information:

- VLAN—The name of the DHCP-snooping VLAN
- IP Addr—The IP address of the untrusted interface or client
- MAC Addr—The MAC address of the untrusted interface or client
- Port—The port number where the untrusted interface or client attempted to access the network

Example

The following command displays the DHCP bindings database on the switch:


```
show ip-security dhcp-snooping entries vlan dhcpVlan
```

The following is sample output from this command:

```
-----
Vlan: dhcpVlan
-----

```

		Server	Client
IP Addr	MAC Addr	Port	Port
-----	-----	-----	-----
172.16.100.9	00:90:27:c6:b7:65	1:1	1:2

show ip-security dhcp-snooping information-option

```
show ip-security dhcp-snooping information-option
```

Description

Displays the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays DHCP relay agent option (option 82) settings. For example, the following command:

```
show ip-security dhcp-snooping information-option
```

Generates the following output:

```
Information option insertion: Enabled
Information option checking : Disabled
Information option policy   : Drop
```

The following command:

```
show ip-security dhcp-snooping information-option
```

Generates the following output:

```
Information option insertion: Enabled
Information option checking : Enabled
Information option policy   : Keep
```

show ip-security dhcp-snooping information circuit-id port-information

```
show ip-security dhcp-snooping information circuit-id port-information ports [<portlist> |
all ]
```

Description

Displays the port information portion of the circuit ID for the indicated port(s).

Syntax Description

portlist	Specifies one or more ports.
all	Specifies all ports

Default

N/A.

Usage Guidelines

This command displays the port information portion of the circuit ID for the indicated ports.

Example

The following command:

```
x250e-48t.7 # show ip-security dhcp-snooping information circuit-id port-information ports
1-7
```

Displays the following output:

```
Port          Circuit-ID Port information string
----          -
1             portinfostring1
2             portinfostring2
3             portinfostring3
4             portinfostring4
5             portinfostring5

Port          Circuit-ID Port information string
----          -
6             1006
7             1007
```

Note: The full Circuit ID string has the form '<Vlan Info>--<Port Info>'

show ip-security dhcp-snooping information-option circuit-id vlan-information

```
show ip-security dhcp-snooping information-option circuit-id vlan-information {{vlan}
<vlan_name>}
```

Description

Displays the VLAN information portion of the circuit ID for the indicated VLAN.

Syntax Description

vlan_name	Specifies a vlan_name
-----------	-----------------------

Default

N/A.

Usage Guidelines

This command displays the VLAN information portion of the circuit ID for the indicated VLAN. When a VLAN is not specified, the circuit ID information for all the VLANs is displayed

Example

The following command:

```
show ip-security dhcp-snooping information-option circuit-id vlan-information vlan Mktg
```

Displays the following output:

```
Vlan          Circuit-ID vlan information string
----          -
Mktg          DSLAM1
```

Note: The full Circuit ID string has the form <Vlan Info>-<Port ifIndex>.

show ip-security dhcp-snooping

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

Description

Displays the DHCP snooping configurations on the switch.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN.
-----------	---

Default

N/A.

Usage Guidelines

The switch displays the following DHCP snooping information:

- DHCP snooping enabled on ports—The ports that have DHCP snooping enabled
- Trusted ports—The ports configured as trusted ports
- Trusted DHCP servers—The servers configured as trusted DHCP servers
- Port—The specific port that has DHCP snooping enabled
- Violation-action—The action the switch takes upon detecting a rogue DHCP packet on the port

Example

The following command displays the DHCP snooping settings for the switch:

```
show ip-security dhcp-snooping vlan dhcpVlan
```

The following is sample output from this command:

```
DHCP Snooping enabled on ports: 1:2, 1:3, 1:4, 1:7, 1:9
```

```
Trusted Ports: 1:7
```

```
Trusted DHCP Servers: None
```

```

-----
Port      Violation-action
-----
1:2       none
1:3       drop-packet
1:4       drop-packet, block-mac permanently
1:7       none
1:9       drop-packet, snmp-trap

```

show ip-security dhcp-snooping violations

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Description

Displays the MAC addressed from which the rouge DHCP packet was received by the switch.

Syntax Description

vlan_name	Specifies the name of the DHCP-snooping VLAN.
-----------	---

Default

N/A.

Usage Guidelines

The switch displays the following DHCP snooping information:

- Port—The specific port that received the rouge DHCP packet
- Violating MAC—The MAC address from which the rouge DHCP was received by the switch

Example

The following command displays the DHCP snooping violations for the VLAN green:

```
show ip-security dhcp-snooping violations green
```

The following is sample output from this command:

Violations seen on following ports

```

-----
Port           Violating MAC
-----
2:3           00-0c-11-a0-3e-12

```

show ip-security source-ip-lockdown

```
show ip-security source-ip-lockdown
```

Description

Displays the source IP lockdown configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The switch displays the following source IP lockdown information:

- Port—Indicates the port that has DHCP snooping enabled and is configured for source IP lockdown
- Locked IP Address—Indicates a valid DHCP-assigned address obtained by a DHCP snooping-enabled port or an authenticated static IP address

Example

The following command displays the source IP configuration on the switch:

```
show ip-security source-ip-lockdown
```

The following is sample output from this command:

```
Ports          Locked IP Address
23             10.0.0.101
```

show mac-lockdown-timeout fdb ports

```
show mac-lockdown-timeout fdb ports [all | <port_list>]
```

Description

Displays the MAC entries that are learned on the specified port or group of ports or for all ports on the switch along with the aging time of each port.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

If a port is down, the command displays all of the MAC entries that are maintained locally in the software.

The MAC entries learned on the specified port are displayed only if the MAC lock down timeout feature is enabled on the port. If you specify a port on which this feature is disabled, the MAC entries learned on that port are not displayed.

The switch displays the following information:

- Mac—The MAC address that defines the entry
- Vlan—The VLAN name and ID for the entry
- Age—The age of the entry, in seconds
- Flags—Flags that define the type of entry:
 - B—Egress Blackhole
 - b—Ingress Blackhole
 - F—Entry in the hardware FDB
 - L—Entry in the software

- Port—The port on which the MAC address has been learned

Example

The following command displays information about the MAC address lock down timeout settings for ports 2:3 and 2:4:

```
show mac-lockdown-timeout fdb ports 2:3, 2:4
```

The following is sample output from this command:

Mac	Vlan	Age	Flags	Port
00:00:01:02:03:04	v1(4094)	0010	F	2:3
00:00:01:00:00:02	v1(4094)	0030	FB b	2:3
00:00:0A:02:03:04	v2(4093)	0050	L	2:4
00:00:0B:02:03:04	v2(4093)	0090	F	2:4

Flags : (F) Entry as in h/w FDB, (L) Entry in s/w and not in h/w
(B) Egress Blackhole, (b) Ingress Blackhole

Total: 4 Entries in FDB: 3 Entries in s/w: 1

show mac-lockdown-timeout ports

```
show mac-lockdown-timeout ports [all | <port_list>]
```

Description

Displays information about the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
port_list	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The switch displays the following MAC address timeout information:

- Port—Indicates the port number that you specified in the command
- MAC Lockdown Timeout—Specifies the enabled/disabled state of the MAC address lock down timeout feature.

- Timeout (in seconds)—Specifies the timeout value for the specified ports. By default, the timeout value is 15 seconds. Even if MAC address lock down is disabled, the default timeout value is displayed.

Example

The following command displays information about the MAC address lock down timeout settings for ports 2:3, 2:4, and 2:6:

```
show mac-lockdown-timeout ports 2:3, 2:4, 2:6
```

The following is sample output from this command:

```
Ports    MAC Lockdown Timeout    Timeout (in seconds)
=====
2:3      Enabled                   300
2:4      Enabled                   300
2:6      Disabled                  15
```

show ports rate-limit flood

```
show ports {<port_list>} rate-limit flood {no-refresh}
```

Description

Displays rate-limit discard statistics.

Syntax

list	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.

Default

N/A.

Usage Guidelines

This command displays the per port ingress rate-limit flood traffic counter as well as information about received packets that have not been discarded due to rate-limiting.

It is used to show the results of the `configure ports <port_list> rate-limit flood [broadcast | multicast | unknown-destmac] [no-limit | <pps>]` command.

Note: As part of the system health check, the system polls the Rate-limit Flood Counters every 5 minutes and looks for non-zero counters on a port. A HAL.RateLimit.Info log message is logged when this is first detected on a port to alert the user that something in the network has triggered the rate limiting to occur. The message is not be logged again unless the counters are cleared.

Example

The following command displays information for port 1:1 without a screen refresh on a NETGEAR 8800 switch.

```
show port 1:1 rate-limit flood no-refresh
```

Following is sample output from this command.

```
BD-8810.1 # show port 1:1 rate-limit flood no-refresh
Port Rate-Limit Discard Monitor                               Tue May 27 13:02:37 2008
Port      Link      Rx Pkt      Rx Byte Rx Pkt Rx Pkt      Flood Rate
          State      Count      Count   Bcast Mcast      Exceeded
=====
1:1       R          5225       65230   2112    0          2112
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
```

show radius

```
show radius {mgmt-access | netlogin}
```

Description

Displays the current RADIUS client configuration and statistics.

Syntax Description

mgmt-access	Specifies configuration and statistics for the switch management RADIUS authentication server.
netlogin	Specifies configuration and statistics for the network login RADIUS authentication server.

Default

N/A.

Usage Guidelines

If you do not specify a keyword, configuration details related to both management and network login are displayed. The output from this command displays the status of RADIUS and RADIUS accounting (enabled or disabled) and the primary and secondary servers for RADIUS and RADIUS accounting.

Use the `mgmt-access` keyword to display only RADIUS configuration details related to management access.

Use the `netlogin` keyword to only RADIUS configuration details related to network login.

Example

The following command displays the current RADIUS client configuration and statistics for both management and network login:

```
show radius
```

The following is sample output from this command:

```
Switch Management Radius: enabled
Switch Management Radius server connect time out: 3 seconds
Switch Management Radius Accounting: disabled
Switch Management Radius Accounting server connect time out: 3 seconds
Netlogin Radius: enabled
Netlogin Radius server connect time out: 3 seconds
Netlogin Radius Accounting: disabled
Netlogin Radius Accounting server connect time out: 3 seconds
```

```
Primary Switch Management Radius server:
```

```
  Server name      :
  IP address       : 10.100.1.100
  Server IP Port    : 1812
  Client address    : 10.116.3.101 (VR-Mgmt)
  Shared secret     : g~`#uovpkpvi~`
Access Requests    : 0                Access Accepts      : 0
Access Rejects     : 0                Access Challenges   : 0
Access Retransmits : 0                Client timeouts     : 0
Bad authenticators : 0                Unknown types       : 0
Round Trip Time    : 0
```

```
Secondary Switch Management Radius server:
```

```
  Server name      :
  IP address       : 10.100.1.101
  Server IP Port    : 1812
  Client address    : 10.116.3.101 (VR-Mgmt)
  Shared secret     : g~`#uovpkpvi~`
Access Requests    : 0                Access Accepts      : 0
Access Rejects     : 0                Access Challenges   : 0
```

```

Access Retransmits: 0           Client timeouts   : 0
Bad authenticators: 0           Unknown types    : 0
Round Trip Time   : 0

```

Primary Netlogin Radius server:

```

Server name      :
IP address       : 10.100.1.200
Server IP Port   : 1812
Client address   : 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkpvi~`
Access Requests  : 0           Access Accepts   : 0
Access Rejects   : 0           Access Challenges : 0
Access Retransmits: 0           Client timeouts   : 0
Bad authenticators: 0           Unknown types    : 0
Round Trip Time  : 0

```

Secondary Netlogin Radius server:

```

Server name      :
IP address       : 10.100.1.201
Server IP Port   : 1812
Client address   : 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkpvi~`
Access Requests  : 0           Access Accepts   : 0
Access Rejects   : 0           Access Challenges : 0
Access Retransmits: 0           Client timeouts   : 0
Bad authenticators: 0           Unknown types    : 0
Round Trip Time  : 0

```

show radius-accounting

```
show radius-accounting {mgmt-access | netlogin}
```

Description

Displays the current RADIUS accounting client configuration and statistics.

Syntax Description

mgmt-access	Specifies configuration and statistics for the switch management RADIUS accounting server.
netlogin	Specifies configuration and statistics for the network login RADIUS accounting server.

Default

N/A.

Usage Guidelines

If you do not specify a keyword, configuration details related to both management and network login are displayed. The output from this command displays information about the status and configuration of RADIUS accounting.

Use the `mgmt-access` keyword to display only RADIUS accounting configuration details related to management access.

Use the `netlogin` keyword to display only RADIUS accounting configuration details related to network login.

Example

The following command displays RADIUS accounting client configuration and statistics for both management and network login:

```
show radius-accounting
```

The following is sample output from this command:

```
Switch Management Radius Accounting: disabled
Switch Management Radius Accounting server connect time out: 3 seconds
Netlogin Radius Accounting: disabled
Netlogin Radius Accounting server connect time out: 3 seconds
```

Primary Switch Management Accounting server:

```
Server name      :
IP address       : 10.100.1.100
Server IP Port: 1813
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Acct Requests    : 0                      Acct Responses    : 0
Acct Retransmits : 0                      Timeouts           : 0
```

Secondary Switch Management Accounting server:

```
Server name      :
IP address       : 10.100.1.101
Server IP Port: 1813
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Acct Requests    : 0                      Acct Responses    : 0
Acct Retransmits : 0                      Timeouts           : 0
```

Primary Netlogin Accounting server:

```
Server name      :
IP address       : 10.100.1.200
Server IP Port: 1813
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
```

```

Acct Requests      : 0                Acct Responses    : 0
Acct Retransmits   : 0                Timeouts         : 0

```

Secondary Netlogin Accounting server:

```

Server name       :
IP address        : 10.100.1.201
Server IP Port    : 1813
Client address    : 10.116.3.101 (VR-Mgmt)
Shared secret     : g~`#uovpkpvi~`
Acct Requests     : 0                Acct Responses    : 0
Acct Retransmits  : 0                Timeouts         : 0

```

show ssh2 private-key

```
show ssh2 private-key
```

Description

Displays the ssh2 server's private key.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the ssh server's private key which can be used to configure the key later or on another switch by using the `configure ssh2 key {pregenerated}` command. The key is saved in the switch's EEPROM.

To erase the key from the EEPROM, use the `unconfigure switch` command.

show sshd2 user-key

```
show sshd2 user-key {<key_name> {users}}
```

Description

Displays the user names bound to a key.

Syntax Description

key_name	Specifies the name of the public key.
users	Specifies the name of the users.

Default

N/A.

Usage Guidelines

This command displays the names of the users that are bound to a public key.

show ssl

```
show ssl {detail}
```

Description

Displays the secure socket layer (SSL) configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following information:

- HTTPS port configured. This is the port on which the clients will connect.
- Length of the RSA key (the number of bits used to generate the private key).
- Basic information about the stored certificate.

If you attempt to use this command before installing the SSH module, the switch displays a message similar to the following:

```
SSL Module: Not Installed.
```

Note: The switch utilizes the SSH module for SSL functionality. You do not install an SSL module, only the SSH module.

Example

The following command displays the SSL configuration:

```
show ssl
```

The following is sample output from this command:

```
HTTPS Port Number: 443
```

Private Key matches with the Public Key in certificate. (or Private key does not match with the Public Key in the certificate)

RSA Key Length: 1024

Certificate:

Data:

Version: 1 (0x0)

Serial Number: 6 (0x6)

Signature Algorithm: md5WithRSAEncryption

Issuer: C=AU, O=CryptSoft Pty Ltd, CN=Test CA (1024 bit)

Validity

Not Before: Oct 16 22:31:03 2000 GMT

Not After : Jan 14 22:31:03 2003 GMT

Subject: C=AU, O=CryptSoft Pty Ltd, CN=Server test cert (512 bit)

show tacacs

show tacacs

Description

Displays the current TACACS+ configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:

- TACACS+—The current state of TACACS+, enabled or disabled.
- TACACS+ Authorization—The current state of TACACS+ authorization, enabled or disabled.
- TACACS+ Accounting—The current state of TACACS+ accounting, enabled or disabled.
- TACACS+ Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ server failure.
- Primary TACACS+ Server—Describes information about the primary TACACS+ server, including:
 - The name of the primary TACACS+ server
 - The IP address of the primary TACACS+ server
 - The TCP port to use to contact the primary TACACS+ server
 - The IP address and VR used by the switch
 - The shared secret configured for the primary TACACS+ server

- Secondary TACACS+ Server—Contains the same type of output as the primary TACACS+ server for the secondary TACACS+ server, if configured.
- TACACS+ Acct Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ accounting server failure.
- TACACS+ Accounting Server parameters, if configured. Contains the same type of output as the TACACS+ server for the TACACS+ accounting server(s), if configured.

Example

The following command displays TACACS+ client configuration and statistics:

```
show tacacs
```

The following is sample output from this command:

```
TACACS+: enabled
TACACS+ Authorization: enabled
TACACS+ Accounting : enabled
TACACS+ Server Connect Timeout sec: 3
Primary TACACS+ Server:
  Server name      :
  IP address       : 10.201.31.238
  Server IP Port: 49
  Client address: 10.201.31.65 (VR-Default)
  Shared secret : qijxou
Secondary TACACS+ Server:
  Server name      :
  IP address       : 10.201.31.235
  Server IP Port: 49
  Client address: 10.201.31.65 (VR-Default)
  Shared secret : qijxou
TACACS+ Acct Server Connect Timeout sec: 3
Primary TACACS+ Accounting Server:
  Server name      :
  IP address       : 10.201.31.238
  Server IP Port: 49
  Client address: 10.201.31.65 (VR-Default)
  Shared secret : qijxou
Secondary TACACS+ Accounting Server:
  Server name      :
  IP address       : 10.201.31.235
  Server IP Port: 49
  Client address: 10.201.31.65 (VR-Default)
  Shared secret : qijxou
```

show tacacs-accounting

```
show tacacs-accounting
```


Description

Displays the current TACACS+ accounting client configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:

- TACACS+ Accounting—The current state of TACACS+ accounting, enabled or disabled.
- TACACS+ Accounting Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ server failure.
- Primary TACACS+ Accounting Server—Describes information about the primary TACACS+ accounting server, including:
 - The name of the primary TACACS+ accounting server
 - The IP address of the primary TACACS+ accounting server
 - The TCP port to use to contact the primary TACACS+ accounting server
 - The IP address and VR used by the switch
 - The shared secret configured for the primary TACACS+ accounting server
- Secondary TACACS+ Accounting Server—Contains the same type of output as the primary TACACS+ accounting server for the secondary TACACS+ accounting server, if configured.

Example

The following command displays TACACS+ accounting client configuration and statistics:

```
show tacacs-accounting
The following is sample output of this command:
TACACS+ Accounting : enabled
TACACS+ Acct Server Connect Timeout sec: 3
Primary TACACS+ Accounting Server:
  Server name      :
  IP address       : 10.201.31.238
  Server IP Port: 49
  Client address: 10.201.31.85 (VR-Default)
  Shared secret : qijxou
Secondary TACACS+ Accounting Server:Not configured
```

show vlan dhcp-address-allocation

```
show vlan <vlan_name> dhcp-address-allocation
```

Description

Displays the DHCP server's address allocation on a specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server of interest.
-----------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration of the DHCP for the VLAN *corp*:

```
show vlan corp dhcp-address-allocation
```

The following is sample output from this command:

```
=====
IP                MAC                State      Lease Renewal Time
=====
    10.0.0.2      00:02:03:04:05:00   Offered    0000:00:10
    10.0.0.3      00:08:03:04:05:00   Assigned   0000:59:09
    10.0.0.4      ee:1c:00:04:05:00   Assigned   0000:59:09
=====
```

show vlan dhcp-config

```
show {vlan} <vlan_name> dhcp-config
```

Description

Displays the DHCP server's configuration for the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN of the DHCP server of interest.
-----------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration of the DHCP server for the VLAN *test*:

```
show vlan test dhcp-config
```

The following is sample output from this command:

```
DHCP Address Range      : 10.10.10.100->10.10.10.200
Netlogin Lease Timer    : Not configured (Default = 10 seconds)
DHCP Lease Timer        : Not configured (Default = 7200 seconds)
Primary DNS Server      : 1.1.1.1
Secondary DNS Server    : 2.2.2.2
Ports DHCP Enabled      : 23
```

show vlan security

```
show vlan <vlan_name> security
```

Description

Displays the MAC limit-learning and lock-learning information for the specified VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

The switch displays the following information:

- Port—Indicates the port on which the MAC address has been learned
- Limit—Indicates that there is either a limited or unlimited amount of learned entries
- State—Indicates that the current FDB entries for the port are permanent, no additional entries are learned, or that the port allows unlimited, dynamic learning
- Learned—Specifies the number of learned entries
- Blackholed—Specifies the number of blackholed entries
- Locked—Specifies the number of locked entries

Example

The following command displays the security setting of the DHCP server for the VLAN *corp*:

```
show vlan blue security
```

The following is sample output from this command:

Port	Limit	State	Learned	Blackholed	Locked
24	Unlimited	Unlocked	0	0	0

ssh2

```
ssh2 {cipher [3des | blowfish]} {port <portnum>} {compression [on | off]} {user <username>}
{<username>} [ <host> | <ipaddress>] {<remote command>} {vr <vr_name>}
```

Description

Initiates an SSH2 client session to a remote SSH2 server.

Syntax Description

3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
portnum	Specifies the TCP port number to be used for communicating with the SSH2 client. The default is port 22.
on	Specifies that the data is to be compressed.
off	Specifies that compression is not to be used. This is the default.
username	Specifies a login name for the remote host, as an alternate to the <code>username@host</code> parameter. Can be omitted if it is the same as the username on the switch.
host	Specifies the name of the remote host.
ipaddress	Specifies the IP address of the remote host.
remote command	Specifies a command to be passed to the remote system for execution. The switch does not support remote commands. The option is only valid if the remote system is a system, such as a UNIX workstation, that accepts remote commands.
vr_name	Specifies the virtual router. The default virtual router is VR-Mgmt.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*.

Default

The default settings for SSH2 parameters are as follows:

- cipher—3des encryption
- port—22
- compression—off
- vr_name—VR-Mgmt

Usage Guidelines

SSH2 does not need to be enabled on the switch in order to use this command.

Typically, this command is used to establish a secure session to a remote switch. You are prompted for your password. Once you have logged in successfully, all NETGEAR 8800 commands you enter are executed on the remote switch. When you terminate the remote session, commands will then resume being executed on the original switch.

Host Name, User Name, and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name, user name, or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted for host and user names
- Underscore (_) Permitted for host and user names
- Colon (:) Permitted for host names and remote IP addresses
- At symbol (@) Permitted only for user names

When naming the host, creating a user name, or configuring the IP address, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command establishes an SSH2 session on switch engineering1:

```
ssh2 admin@engineering1
```

The following command establishes an SSH2 session with the switch named NETGEAR8810 over TCP port 2050 with compression enabled:

```
ssh2 compression on port 2050 admin@NETGEAR8810
```

unconfigure ip-security dhcp-snooping information check

```
unconfigure ip-security dhcp-snooping information check
```

Description

Disables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking in the server-originated packets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command disables the checking of the server-originated packets for the presence of option 82 so the packets will be forwarded normally.

unconfigure ip-security dhcp-snooping information circuit-id port-information ports

```
unconfigure ip-security dhcp-snooping information circuit-id port-information ports  
[<portlist> | all]
```

Description

Unconfigures the port information portion of the circuit ID.

Syntax Description

portlist	Specifies the port(s) for which port information of the circuit-ID is unconfigured.
all	Specifies all ports.

Default

The default is `all`.

Usage Guidelines

This command unconfigures the port information portion of the circuit ID string for the indicated ports thereby restoring it to the default (ifIndex value).

unconfigure ip-security dhcp-snooping information circuit-id vlan-information

```
unconfigure ip-security dhcp-snooping information circuit-id vlan-information {vlan}
[<vlan_name>|all]
```

Description

Unconfigures the VLAN info portion of the circuit ID of a VLAN.

Syntax Description

vlan_name	Specifies the VLAN for which VLAN information of the circuit-ID is unconfigured.
all	Specifies all VLANs.

Default

The default is `all`.

Usage Guidelines

This command unconfigures the VLAN information portion of the circuit ID of a VLAN, restoring it to the default.

unconfigure ip-security dhcp-snooping information option

```
unconfigure ip-security dhcp-snooping information option
```

Description

Disables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command disables the DHCP relay agent option (option 82), which is inserted into client-originated DHCP packets before they are forwarded to the server.

unconfigure ip-security dhcp-snooping information policy

```
unconfigure ip-security dhcp-snooping information policy
```

Description

Unconfigures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command unconfigures the DHCP relay agent option information policy to the default value of `replace`.

unconfigure radius

```
unconfigure radius {mgmt-access | netlogin} {server [primary | secondary]}
```

Description

Unconfigures the RADIUS client configuration.

Syntax Description

<code>mgmt-access</code>	Specifies the switch management RADIUS authentication server.
<code>netlogin</code>	Specifies the network login RADIUS authentication server.
<code>primary</code>	Unconfigures the primary RADIUS server.
<code>secondary</code>	Unconfigures the secondary RADIUS server.

Default

Unconfigures both primary and secondary servers for management and network login.

Usage Guidelines

If you do not specify any keywords, this command unconfigures both the primary and secondary servers for management and network login.

The following list describes the available keywords:

- `mgmt-access`—Use this keyword to unconfigure only the server(s) for management functions.
- `netlogin`—Use this keyword to unconfigure only the server(s) for network login.
- `primary`—Use this keyword to specify only the primary RADIUS sever.
- `secondary`—Use this keyword to specify only the secondary RADIUS server.

Example

The following command unconfigures the secondary RADIUS server settings for both management and network login:

```
unconfigure radius server secondary
```

The following command unconfigures the secondary RADIUS server settings for only network login:

```
unconfigure radius netlogin server secondary
```

The following command unconfigures all RADIUS server settings for only management functions:

```
unconfigure radius mgmt-access
```

unconfigure radius-accounting

```
unconfigure radius-accounting {mgmt-access | netlogin} {server [primary | secondary]}
```

Description

Unconfigures the RADIUS accounting server configuration.

Syntax Description

<code>mgmt-access</code>	Specifies the switch management RADIUS accounting server.
<code>netlogin</code>	Specifies the network login RADIUS accounting server.
<code>primary</code>	Unconfigures the primary RADIUS accounting server.
<code>secondary</code>	Unconfigures the secondary RADIUS accounting server.

Default

Unconfigures both the primary and secondary accounting servers for management and network login.

Usage Guidelines

If you do not specify any keywords, this command unconfigures both the primary and secondary accounting servers for management and network login.

The following list describes the available keywords:

- `mgmt-access`—Use this keyword to unconfigure only the accounting server(s) for management functions.
- `netlogin`—Use this keyword to unconfigure only the accounting server(s) for network login.
- `primary`—Use this keyword to specify only the primary RADIUS accounting sever.
- `secondary`—Use this keyword to specify only the secondary RADIUS accounting server.

Example

The following command unconfigures the secondary RADIUS accounting server settings for both management and network login:

```
unconfigure radius-accounting server secondary
```

The following command unconfigures the secondary RADIUS accounting server settings for only network login:

```
unconfigure radius-accounting netlogin server secondary
```

The following command unconfigures all RADIUS accounting server settings for only management functions:

```
unconfigure radius-accounting mgmt-access
```

unconfigure tacacs

```
unconfigure tacacs {server [primary | secondary]}
```

Description

Unconfigures the TACACS+ server configuration.

Syntax Description

<code>primary</code>	Unconfigures the primary TACACS+ server.
<code>secondary</code>	Unconfigures the secondary TACACS+ server.

Default

Unconfigures both the primary and secondary TACACS+ servers.

Usage Guidelines

None.

Example

The following command unconfigures all TACACS+ servers settings:

```
unconfigure tacacs
```

unconfigure tacacs-accounting

```
unconfigure tacacs-accounting {server [primary | secondary]}
```

Description

Unconfigures the TACACS+ accounting server configuration.

Syntax Description

primary	Unconfigures the primary TACACS+ accounting server.
secondary	Unconfigures the secondary TACACS+ accounting server.

Default

Unconfigures both the primary and secondary TACACS+ accounting servers.

Usage Guidelines

None.

Example

The following command unconfigures all TACACS+ accounting servers settings:

```
unconfigure tacacs-accounting
```

unconfigure trusted-ports trust-for dhcp-server

```
unconfigure trusted-ports [<ports>|all] trust-for dhcp-server
```

Description

Unconfigures, disables one or more DHCP trusted ports.

Syntax Description

ports	Specifies one or more trusted ports.
all	Specifies all trusted ports.

Default

N/A.

Usage Guidelines

Use this command to disable one or more DHCP trusted ports.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} <vlan_name>
```

To display any violations that occur, including those on DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} <vlan_name>
```

Example

The following command unconfigures ports 2:2 and 2:3 as trusted ports:

```
unconfigure trusted-ports 2:2-2:3 trust-for dhcp-server
```

unconfigure vlan dhcp

```
unconfigure vlan <vlan_name> dhcp
```

Description

Unconfigure all the DHCP configuration information for the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN on which to unconfigure DHCP.
-----------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP server for the VLAN *temporary*:

```
unconfigure temporary dhcp
```

unconfigure vlan dhcp-address-range

```
unconfigure vlan <vlan_name> dhcp-address-range
```

Description

Unconfigure the DHCP address range information for the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN on which to unconfigure DHCP.
-----------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP address range for the VLAN *temporary*:

```
unconfigure temporary dhcp-address-range
```

unconfigure vlan dhcp-options

```
unconfigure {vlan} <vlan_name> dhcp-options {[ default-gateway | dns-server {primary | secondary} | wins-server]}
```

Description

Unconfigure the DHCP option information for the specified VLAN.

Syntax Description

vlan_name	Specifies the VLAN on which to unconfigure DHCP.
default-gateway	Specifies the router option.
dns-server	Specifies the Domain Name Server (DNS) option.
primary	Specifies the primary DNS option.
secondary	Specifies the secondary DNS option.
wins-server	Specifies the NetBIOS name server (NBNS) option.

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP options for the VLAN *temporary*:

```
unconfigure temporary dhcp-options
```

upload dhcp-bindings

```
upload dhcp-bindings
```

Description

Upload the DHCP bindings immediately on demand.

Syntax

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This commands enables the functionality to allow you to upload DCHP bindings on demand.

This chapter describes commands for configuring network login.

Network login is a feature designed to control the admission of user packets into a network by giving network access only to users that have been properly authenticated. Network login is controlled by an administrator on a per port, per VLAN basis and uses an integration of DHCP, user authentication over the web interface, user authentication by MAC address, or 802.1x client software, and a RADIUS server to provide a user database or specific configuration details.

Network login has two modes of operation:

- Campus mode, used when a port in a VLAN will move to another VLAN when authentication has been completed successfully. This mode is for the roaming user who will not always be using the same port for authentication. Campus mode requires a DHCP server and a RADIUS server configured for Network Login.
- ISP mode, used when the port and VLAN used will remain constant. All network settings are configured for that VLAN.

A DHCP server is included to support network login functionality.

clear netlogin state

```
clear netlogin state {port <portlist>}
```

Description

Clears and initializes the network login sessions on a VLAN port.

Syntax Description

portlist	Specifies the ports to clear.
----------	-------------------------------

Default

None.

Usage Guidelines

Clear the states of every MAC learned on this VLAN port and put the port back to unauthenticated state. The port will be moved to its original VLAN if configured in campus mode.

Example

The following command clears the Network Login state of port 2:9:

```
clear netlogin state port 2:9
```

clear netlogin state mac-address

```
clear netlogin state mac-address <mac>
```

Description

Initialize/reset the network login sessions for a specified supplicant.

Syntax Description

mac	Specifies the MAC address of the supplicant.
-----	--

Default

N/A.

Usage Guidelines

This command is essentially equivalent to a particular supplicant logging out. The MAC address will be cleared from the FDB, the port is put back to its original VLAN (for campus mode), and the port state is set to unauthenticated, if this was the last authenticated MAC on this port.

Example

The following command resets the Network Login session for the supplicant with the MAC address of 00:e0:18:01:32:1f:

```
clear netlogin state mac-address 00:e0:18:01:32:1f
```

configure netlogin add mac-list

```
configure netlogin add mac-list [<mac> {<mask>} | default] {encrypted} {<password>} {ports <port_list>}
```

Description

Adds an entry to the MAC address list for MAC-based network login.

Syntax Description

mac	Specifies the MAC address to add.
mask	Specifies the number of bits to use for the mask.
default	Specifies the default entry.
encrypted	Used to display encrypted form of password in configuration files. Do not use.
password	Specifies the password to send for authentication.
ports	Specifies the port or port list to use for authentication.

Default

If no password is specified, the MAC address will be used.

Usage Guidelines

Use this command to add an entry to the MAC address list used for MAC-based network login.

If no match is found in the table of MAC entries, and a default entry exists, the default will be used to authenticate the client. All entries in the list are automatically sorted in longest prefix order.

Associating a MAC Address to a Port

You can configure the switch to accept and authenticate a client with a specific MAC address. Only MAC addresses that have a match for the specific ports are sent for authentication. For example, if you associate a MAC address with one or more ports, only authentication requests for that MAC addresses received on the port(s) are sent to the RADIUS server. The port(s) block all other authentication requests that do not have a matching entry. This is also known as secure MAC.

To associate a MAC address with one or more ports, specify the `ports` option when using the `configure netlogin add mac-list [<mac> [<mask>] | default] {encrypted} [<password>] {ports <port_list>}` command.

You must enable MAC-based network login on the switch and the specified ports before using this command. If MAC-based network login is not enabled on the specified port(s), the switch displays a warning message similar to the following:

```
WARNING: Not all specified ports have MAC-Based NetLogin enabled.
```

If this occurs, make sure to enable MAC-based network login.

Example

The following command adds the MAC address 10:20:30:40:50:60 with the password *foo* to the list:

```
configure netlogin add mac-list 10:20:30:40:50:60 password foo
```

The following command associates MAC address 10:20:30:40:50:70 with ports 2:2 and 2:3. This means authentication requests from MAC address 10:20:30:40:50:70 are only accepted on ports 2:2 and 2:3:

```
configure netlogin add mac-list mac 10:20:30:40:50:70 ports 2:2-2:3
```

configure netlogin add proxy-port

```
configure netlogin add proxy-port <tcp_port> {http | https}
```

Description

Configure the ports that will be hijacked and redirected for HTTP or HTTPS traffic.

Syntax Description

tcp_port	Specifies the port to be hijacked.
----------	------------------------------------

Default

HTTP traffic.

Usage Guidelines

This command allows you to configure the ports that will be hijacked and redirected for HTTP or HTTPS traffic. For each hijacked proxy port, you must specify whether the port is to be used for HTTP or HTTPS traffic.

No more than 5 such ports are supported in addition to ports 80 and ports 443. Attempts to add more than 5 ports generate an error.

configure netlogin agingtime

```
configure netlogin agingtime <minutes>
```

Description

Lets you configure network login aging.

Syntax Description

minutes	Specifies the aging time in minutes.
---------	--------------------------------------

Default

The default value is 5.

Usage Guidelines

Use this command to configure the aging time for network login. The aging time is the time after which learned clients that failed authentication or did not attempt to authenticate are removed from the system. This prevents the switch from keeping all clients ever seen on a network-login-enabled port.

The range can be from 0 to 3000, where 0 indicates no age out.

Example

The following command specifies an aging time of 15 minutes:

```
configure netlogin agingtime 15
```

configure netlogin allowed-refresh-failures

```
configure netlogin allowed-refresh-failures <num_failures>
```

Description

Sets the number refresh failures.

Syntax Description

num_failures	Specifies the number of refresh failures. The range is from 0 to 5.
--------------	---

Default

The default is 0.

Usage Guidelines

This command allows you to set the number of refresh failures allowed. You can set the number of failures to be from between 0 to 5. The default value is 0.

configure netlogin authentication database-order

```
configure netlogin [mac | web-based] authentication database-order
[[radius] | [local] | [radius local] | [local radius]]
```

Description

Configures the order of database authentication protocols to use.

Syntax Description

mac	Specifies MAC-based authentication.
web-based	Specifies Web-based authentication.
radius	Specifies an authentication order from only the RADIUS database.
local	Specifies an authentication order from only the local database.
radius local	Specifies an authentication order of RADIUS database first, followed by local.
local radius	Specifies an authentication order of local database first, followed by RADIUS.

Default

By default, the authentication order is RADIUS, local-user database.

Usage Guidelines

Use this command in situations where, when both a network login RADIUS server and a local-user database are configured, you want to have control over which database to use first. If one authentication fails, the other database is tried; if that authentication is successful, the switch authenticates the network login user.

Example

The following command sets the database authentication order to local-user database, RADIUS:

```
configure netlogin mac authentication database-order local radius
```

configure netlogin authentication failure vlan

```
configure netlogin authentication failure vlan <vlan_name> {ports <port_list>}
```

Description

Configures authentication failure VLAN on network login enabled ports.

Syntax Description

vlan_name	Specifies the name of the authentication failure VLAN.
port_list	Specifies one or more ports or slots and ports. If the <code>ports</code> keyword is not used, the command applies to all ports.

Default

By default, authentication failure VLAN is configured on all network login enabled ports if no port is specifically configured.

Usage Guidelines

Use this command to configure authentication failure VLAN on network login enabled ports. When a supplicant fails authentication, it is moved to the authentication failure VLAN and is given limited access until it passes the authentication either through RADIUS or local. Depending on the authentication database order for that particular network login method (MAC, web or dot1x), the other database is used to authenticate the client. If the final result is an authentication failure and if the authentication failure VLAN is configured and enabled on that port, the client is moved to that location.

There four different authentication orders which can be configured per authentication method currently. They are:

- RADIUS
- local
- RADIUS, local
- local, RADIUS

In each case, you must consider the end result in deciding whether to authenticate the client in authentication failure VLAN or authentication service unavailable VLAN (if configured).

For example, when `netlogin mac authentication database order` is `local, radius`, if the authentication of a MAC client fails through a local database, RADIUS is used for authentication. If RADIUS also fails authentication, the client is moved to authentication failure VLAN. The same is true for all authentication database orders (`radius, local`; `local, radius`; `radius`; `local`).

If authentication through local fails but passes through RADIUS, the client is moved to the appropriate destination VLAN.

If the local authentication fails and the RADIUS server is not available, the client is not moved to authentication failure VLAN.

configure netlogin authentication service-unavailable vlan

```
configure netlogin authentication service-unavailable vlan <vlan_name>
{ports <port_list>}
```

Description

Configures authentication service unavailable VLAN on network login enabled ports.

Syntax Description

vlan_name	Specifies the name of the service-unavailable VLAN.
port_list	Specifies one or more ports or slots and ports. If the <code>ports</code> keyword is not used, the command applies to all ports.

Default

Defaults to all network login enabled ports.

Usage Guidelines

This command configures authentication service unavailable VLAN on the specified network login enabled ports. Authentication service unavailable VLAN is configured on all the network login enabled ports, if no port is specifically mentioned. When an authentication service is not available to authenticate the network login clients, they are moved to the authentication service-unavailable VLAN and are given limited access until the authentication service is available either through RADIUS or local. Depending on the authentication database order for that particular network login method (MAC, web or dot1x), the other database is used to authenticate the client. If the final result is an authentication failure and if the authentication failure VLAN is configured and enabled on that port, the client is moved to that location.

Note: The local database can be configured for MAC and Web authentication method only, not for dot1x.

There are four different authentication orders which can be configured per authentication method currently. They are:

- RADIUS
- Local
- RADIUS, local
- Local, RADIUS

In each case, you must consider the end result in deciding whether to authenticate the client in authentication failure VLAN or authentication service unavailable VLAN (if configured).

For example, when `netlogin mac authentication database order` is `local, radius`, if the authentication of a MAC client fails through a local database, RADIUS is used for authentication. If RADIUS also fails authentication, the client is moved to authentication failure VLAN. The same is true for all authentication database orders (`radius, local`; `local, radius`; `radius, local`).

If authentication through local fails but passes through RADIUS, the client is moved to appropriate destination VLAN.

If the local authentication fails and the RADIUS server is not available, the client is not moved to authentication failure VLAN.

Authentication service is considered to be unavailable in the following cases:

- For local authentication if the user entry is not present in the local database.
- For RADIUS in the following cases:
 - the RADIUS server is not running.
 - the RADIUS server is not configured on the switch
 - the RADIUS server is configured but not enabled on the switch.

Note: If web is enabled on a port where dot1x or MAC are also enabled, the authentication failure/service-unavailable VLAN configuration is not applicable to those clients where dot1x or MAC clients which fail authentication or where authentication service is not available.

configure netlogin banner

```
configure netlogin banner <banner>
```

Description

Configures the network login page banner.

Syntax Description

banner	Specifies the HTML code for the banner.
--------	---

Default

The default banner is the NETGEAR logo.

Usage Guidelines

The banner is a quoted, HTML string, that will be displayed on the network login page. The string is limited to 1024 characters.

This command applies only to the web-based authentication mode of network login.

Example

The following command configures the network login page banner:

```
configure netlogin banner "<html><head>Please Login</head></html>"
```

configure netlogin base-url

```
configure netlogin base-url <url>
```

Description

Configures the base URL for network login.

Syntax Description

url	Specifies the base URL for network login.
-----	---

Default

The base URL default value is “network-access.com.”

Usage Guidelines

When you login using a web browser, you are redirected to the specified base URL, which is the DNS name for the switch.

You must configure a DNS name of the type “www.xx...xx.xxx” or “xx...xx.xxx”.

This command applies only to the web-based authentication mode of network login.

Example

The following command configures the network login base URL as `access.net`:

```
configure netlogin base-url access.net
```

configure netlogin delete mac-list

```
configure netlogin delete mac-list [<mac> {<mask>} | default]
```

Description

Deletes an entry from the MAC address list for MAC-based network login.

Syntax Description

mac	Specifies the MAC address to delete.
mask	Specifies the number of bits to use for the mask.
default	Specifies the default entry.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the MAC address list used for MAC-based network login.

Example

The following command deletes the MAC address 10:20:30:40:50:60 from the list:

```
configure netlogin delete mac-list 10:20:30:40:50:60
```

configure netlogin delete proxy-port

```
configure netlogin delete proxy-port <tcp_port>
```

Description

Configure the ports that are to be hijacked and redirected for HTTP or HTTPS traffic.

Syntax Description

tcp_port	Specifies the port to be hijacked.
----------	------------------------------------

Default

N/A.

Usage Guidelines

This command allows you to unconfigure the ports that will be hijacked and redirected for HTTP or HTTPS traffic.

configure netlogin dot1x eapol-transmit-version

```
configure netlogin dot1x eapol-transmit-version <eapol-version>
```

Description

Configures the default EAPOL version sent in transmitted packets for network login.

Syntax Description

eapol-version	Specifies the EAPOL version. Choices are "v1" or "v2".
---------------	--

Default

The default is "v1".

Usage Guidelines

Although the NETGEAR 8800 software supports EAPOL version 2, some clients do not yet accept the version 2 EAPOL packets. The packet format for the two versions is the same.

Example

The following command changes the EAPOL version to 2:

```
configure netlogin dot1x eapol-transmit-version v2
```

configure netlogin dot1x guest-vlan

```
configure netlogin dot1x guest-vlan <vlan_name> {ports <port_list>}
```

Description

Configures a guest VLAN for 802.1x authentication network login.

Syntax Description

vlan_name	Specifies the name of the guest VLAN.
port_list	Specifies one or more ports or slots and ports. If the <code>ports</code> keyword is not used, the command applies to all ports.

Default

N/A.

Usage Guidelines

This command configures the guest VLAN for 802.1x on the current virtual router (VR).

Note: You can configure guest VLANs on a per port basis, which allows you to configure more than one guest VLAN per VR.

If you do not specify any ports, the guest VLAN is configured for all ports.

Each port can have a different guest VLAN.

A guest VLAN provides limited or restricted network access if a supplicant connected to a port does not respond to the 802.1x authentication requests from the switch. A port always moves untagged into the guest VLAN.

Keep in mind the following when configuring guest VLANs:

- You must create a VLAN and configure it as a guest VLAN before enabling the guest VLAN feature.

- Configure guest VLANs only on network login ports with 802.1x enabled.
- Movement to guest VLANs is not supported on network login ports with MAC-based or web-based authentication.
- 802.1x must be the only authentication method enabled on the port for movement to guest VLAN.
- No supplicant on the port has 802.1x capability.
- You configure only one guest VLAN per virtual router interface.

Note: The supplicant does not move to a guest VLAN if it fails authentication after an 802.1x exchange; the supplicant moves to the guest VLAN only if it does not respond to an 802.1x authentication request.

Modifying the Supplicant Timer

By default, the switch attempts to authenticate the supplicant every 30 seconds for a maximum of three tries. If the supplicant does not respond to the authentication requests, the client moves to the guest VLAN. The number of authentication attempts is not a user-configured parameter.

To modify the supplicant response timer, use the following command and specify the `supp-resp-timeout` parameter:

```
configure netlogin dot1x timers [{server-timeout <server_timeout>} {quiet-period <quiet_period>} {reauth-period <reauth_period>} {reauth-max <max_num_reauths>}] {supp-resp-timeout <supp_resp_timeout>}]
```

If a supplicant on a port in the guest VLAN becomes 802.1x-capable, the switch starts processing the 802.1x responses from the supplicant. If the supplicant is successfully authenticated, the port moves from the guest VLAN to the destination VLAN specified by the RADIUS server.

Enabling Guest VLANs

To enable the guest VLAN, use the following command:

```
enable netlogin dot1x guest-vlan ports [all | <ports>]
```

Example

The following command creates a guest VLAN for 802.1x named guest for all ports:

```
configure netlogin dot1x guest-vlan guest
```

The following command creates a guest VLAN named guest for ports 2 and 3:

```
configure netlogin dot1x guest-vlan guest ports 2,3
```

configure netlogin dot1x timers

```
configure netlogin dot1x timers [{server-timeout <server_timeout>} {quiet-period
<quiet_period>} {reauth-period <reauth_period>} {reauth-max <max_num_reauths>}}
{supp-resp-timeout <supp_resp_timeout>}]
```

Description

Configures the 802.1x timers for network login.

Syntax Description

server-timeout	Specifies the timeout period for a response from the RADIUS server. The range is 1 to 120 seconds.
quiet-period	Specifies the time for which the switch will not attempt to communicate with the supplicant after authentication has failed. The range is 0 to 65535 seconds.
reauth-period	Specifies time after which the switch will attempt to re-authenticate an authenticated supplicant. The range is 0, 30 to 7200 seconds.
reauth-max	Specifies the maximum reauthentication counter value. The range is 1 to 10.
supp-resp-timeout	Specifies the time for which the switch will wait for a response from the supplicant. The range is 1 to 120 seconds.

Default

The defaults are as follows:

- server-timeout—30 seconds
- quiet-period—60 seconds
- reauth-period—3600 seconds
- reauth-max—3
- supp-resp-timeout—30 seconds

Usage Guidelines

To disable re-authentication, specify 0 for the `reauth-period` parameter. (If `reauth-period` is set to 0, `reauth-max` value doesn't apply.)

If you attempt to configure a timer value that is out of range (not supported), the switch displays an error message. The following is a list of sample error messages:

- server-timeout—ERROR: RADIUS server response timeout out of range (1..120 sec)
- quiet-period—%% Invalid number detected at '^' marker.
%% Input number must be in the range [0, 65535].
- reauth-period—ERROR: Re-authentication period out of range (0, 30..7200 sec)
- reauth-counter—ERROR: Re-authentication counter value out of range (1..10)

- `supp-resp-timeout`—ERROR: Supplicant response timeout out of range (1..120 sec)

To display the 802.1x timer settings, use the `show netlogin` and `show netlogin dot1x` commands.

Example

The following command changes the 802.1x server-timeout to 10 seconds:

```
configure netlogin dot1x timers server-timeout 10
```

configure netlogin dynamic-vlan

```
configure netlogin dynamic-vlan [disable | enable]
```

Description

Configures the switch to automatically and dynamically create a VLAN after receiving authentication requests from one or more supplicants (clients).

Syntax Description

disable	Specifies that the switch does not automatically create dynamic VLANs. This is the default behavior.
enable	Specifies that the switch automatically create dynamic VLANs.

Default

The default is disabled.

Usage Guidelines

Use this command to configure the switch to dynamically create a VLAN. If configured for dynamic VLAN creation, the switch automatically creates a supplicant VLAN that contains both the supplicant's physical port and one or more uplink ports.

A dynamically created VLAN is only a Layer 2 bridging mechanism; this VLAN does not work with routing protocols to forward traffic. After the switch unauthenticates all of the supplicants from the dynamically created VLAN, the switch deletes that VLAN.

Note: Dynamically created VLANs do not support the session refresh feature of web-based network login because dynamically created VLANs do not have an IP address. Also, dynamic VLANs are not supported on ports when STP and network login are both configured on the ports.

By dynamically creating and deleting VLANs, you minimize the number of active VLANs configured on your edge switches. In addition, the RADIUS server forwards VSA information to dynamically create the VLAN thereby simplifying switch management. A key difference between dynamically created VLANs and other VLANs is that the switch does not save dynamically created VLANs. Even if you use the `save` command, the switch does not save a dynamically created VLAN.

Supported Vendor Specific Attributes

To prevent conflicts with existing VLANs on the switch, the RADIUS server uses Vendor Specific Attributes (VSAs) to forward VLAN information, including VLAN ID, to the switch. The following list specifies the supported VSAs for configuring dynamic network login VLANs:

- Netlogin-VLAN-ID (VSA 209)
- Netlogin-Extended-VLAN (VSA 211)

Note: If the ASCII string only contains numbers, it is interpreted as the VLAN ID. Dynamic VLANs only support numerical VLAN IDs; VLAN names are not supported.

- IETF: Tunnel-Private-Group-ID (VSA 81)

The switch automatically generates the VLAN name in the following format: `SYS_NLD_<TAG>` where `<TAG>` specifies the VLAN ID. For example, a dynamic network login VLAN with an ID of 10 has the name `SYS_NLD_0010`. >

Specifying the Uplink Ports

To specify one or more ports as tagged uplink ports that are added to the dynamically created VLAN, use the following command:

```
configure netlogin dynamic-vlan uplink-ports [<port_list> | none]
```

The uplink ports send traffic to and from the supplicants from the core of the network.

By default the setting is none. For more information about this command, see the usage guidelines for `configure netlogin dynamic-vlan uplink-ports`.

Viewing Status Information

To display summary information about all of the VLANs on the switch, including any dynamic VLANs currently operating on the switch, use the following command:

```
show vlan
```

If the switch dynamically creates a VLAN, the VLAN name begins with `SYS_NLD_` and the output contains a `d` flag for the dynamically created VLAN.

To display the status of dynamic VLAN configuration on the switch, use the following command:

```
show netlogin
```

The switch displays the current state of dynamic VLAN creation (enabled or disabled) and the uplink port(s) associated with the dynamic VLAN.

Example

The following command automatically adds ports 1:1-1:2 to the dynamically created VLAN as uplink ports:

```
configure netlogin dynamic-vlan uplink-ports 1:1-1:2
```

configure netlogin dynamic-vlan uplink-ports

```
configure netlogin dynamic-vlan uplink-ports [<port_list> | none]
```

Description

Specifies which port(s) are added as tagged, uplink ports to the dynamically created VLANs for network login.

Syntax Description

port_list	Specifies one or more ports to add to the dynamically created VLAN for network login.
none	Specifies that no ports are added. This is the default setting.

Default

The default setting is `none`.

Usage Guidelines

Use this command to specify which port(s) are used as uplink ports and added to the dynamically created VLAN for network login. The uplink ports send traffic to and from the supplicants from the core of the network.

Uplink ports should not be configured for network login (network login is disabled on uplink ports). If you specify an uplink port with network login enabled, the configuration fails and the switch displays an error message similar to the following:

```
ERROR: The following ports have NetLogin enabled: 1, 2
```

If this occurs, select a port with network login disabled.

Enabling Dynamic Network Login VLANs

To configure the switch to dynamically create a VLAN upon receiving an authentication response, use the following command:

```
configure netlogin dynamic-vlan [disable | enable]
```

By default, the setting is disabled. For more detailed information about this command, see the usage guidelines `configure netlogin dynamic-vlan uplink-ports`.

Viewing Status Information

To display summary information about all of the VLANs on the switch, including any dynamic VLANs currently operating on the switch, use the following command:

```
show vlan
```

If the switch dynamically creates a VLAN, the VLAN name begins with `SYS_NLD_` and the output contains a `d` flag for the dynamically created VLAN.

To display the status of dynamic VLAN configuration on the switch, use the following command:

```
show netlogin
```

The switch displays the current state of dynamic VLAN creation (enabled or disabled) and the uplink port(s) associated with the dynamic VLAN.

Example

The following command configures the switch to add ports 1:1-1:2 to the dynamically created network login VLAN:

```
configure netlogin dynamic-vlan uplink-ports 1:1-1:2
```

configure netlogin local-user

```
configure netlogin local-user <user-name> {vlan-vsa [{tagged | untagged} [<vlan_name> | <vlan_tag>]] | none}}
```

Description

Configures an existing local network login account.

Syntax Description

<code>user-name</code>	Specifies the name of an existing local network login account.
<code>tagged</code>	Specifies that the client be added as tagged.
<code>untagged</code>	Specifies that the client be added as untagged.
<code>vlan_name</code>	Specifies the name of the destination VLAN.
<code>vlan_tag</code>	Specifies the VLAN ID, tag, of the destination VLAN.
<code>none</code>	Specifies that the VSA 211 wildcard (*) is applied, only if you do not specify tagged or untagged

Default

N/A.

Usage Guidelines

Use this command to modify the attributes of an existing local network login account. You can update the following attributes associated with a local network login account:

- Password of the local network login account
- Destination VLAN attributes including: adding clients tagged or untagged, the name of the VLAN, and the VLAN ID

Note: Passwords are case-sensitive and must have a minimum of 1 character and a maximum of 32 characters.

You must create a local network login account before using this command. To create a local network login user name and password, use the following command:

```
create netlogin local-user <user-name> {encrypted} {<password>} {vlan-vsa [{tagged |
untagged} [<vlan_name>] | <vlan_tag>]]} {security-profile <security_profile>}
```

If the switch displays a message similar to the following:

```
* XCM8810.14 # configure netlogin local-user purplenet
                                     ^
%% Invalid input detected at '^' marker.
```

You might be attempting to modify a local network login account that is not present on the switch, or you might have incorrectly entered the account name. To confirm the names of the local network login accounts on your switch, use the following command:

```
show netlogin local-users
```

Additional Requirements

This command applies only to the web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.

You must have administrator privileges to use this command. If you do not have administrator privileges, the switch displays a message similar to the following:

```
This user does not have permissions for this command.
```

Passwords are case-sensitive. Passwords must have a minimum of 0 characters and a maximum of 32 characters. If you attempt to create a password with more than 32 characters, the switch displays the following message after you re-enter the password:

```
Password cannot exceed 32 characters
```

Example

This section contains the following examples:

- Updating the password
- Modifying destination VLAN attributes

Updating the Password

The following command updates the password of an existing local network login account:

```
configure netlogin local-user megtest
```

After you enter the local network login user name, press [Enter]. The switch prompts you to enter a password; however, the switch does not display the password. At the prompt enter the new password:

```
password:
```

After you enter the new password, press [Enter]. The switch then prompts you to re-enter the password:

```
Reenter password:
```

Updating VLAN Attributes

You can add a destination VLAN, change the destination VLAN, or remove the destination from an existing local network login account. This example changes the destination VLAN for the specified local network login account:

```
configure netlogin local-user megtest vlan-vsa green
```

configure netlogin local-user security-profile

```
configure netlogin local-user <user-name> security-profile <security_profile>
```

Description

Changes a previously associated security profile.

Syntax Description

user-name	Specifies the name of an existing local network login account.
security_profile	Specifies a security profile string during account creation.

Default

N/A.

Usage Guidelines

Use this command to change any previously associated security profiles on the switch.

configure netlogin mac timers reauth-period

```
configure netlogin mac timers reauth-period <reauth_period>
```

Description

Configures the reauthentication period for network login MAC-based authentication.

Syntax Description

reauth-period	Specifies time after which the switch will attempt to re-authenticate an authenticated supplicant. The range is 0, 30 to 7200 seconds.
---------------	--

Default

The default is 0 (disabled).

Usage Guidelines

This command allows you to configure the reauth-period for network login MAC-based authentication. The session-timeout configuration on the RADIUS server overrides the reauth-period if it has been configured.

configure netlogin move-fail-action

```
configure netlogin move-fail-action [authenticate | deny]
```

Description

Configures the action network login takes if a VLAN move fails. This can occur if two clients attempt to move to an untagged VLAN on the same port.

Syntax Description

authenticate	Specifies that the client is authenticated.
deny	Specifies that the client is not authenticated. This is the default setting.

Default

The default setting is `deny`.

Usage Guidelines

Use this command to specify how network login behaves if a VLAN move fails. Network login can either authenticate the client on the current VLAN or deny the client.

The following describes the parameters of this command if two clients want to move to a different untagged VLAN on the same port:

- **authenticate**—Network login authenticates the first client that requests a move and moves that client to the requested VLAN. Network login authenticates the second client but does not move that client to the requested VLAN. The second client moves to the first client's authenticated VLAN.
- **deny**—Network login authenticates the first client that requests a move and moves that client. Network login does not authenticate the second client.

To view the current move-fail-action setting on the switch, use the `show netlogin` command.

Example

The following command configures network login to authenticate the client on the current VLAN:

```
configure netlogin move-fail-action authenticate
```

configure netlogin port allow egress-traffic

```
configure netlogin ports [<port_list> | all] allow egress-traffic [none | unicast | broadcast | all_cast]
```

Description

Configures the egress traffic in an unauthenticated state.

Syntax Description

all	Specifies all network login ports.
port_list	Specifies one or more network login ports.
none	Specifies that no traffic is sent out if if no authenticated clients exist on the VLAN.
unicast	Specifies that the unicast flooding traffic for the VLANs on the network login enabled port be sent.
broadcast	Specifies that the broadcast traffic for the VLANs on the network login enabled port be sent.
all_cast	Specifies that the broadcast and unicast flooding traffic for the VLANs on the network login enabled port be sent.

Default

The default is `none`.

Usage Guidelines

This command allows you to configure the egress traffic in an unauthenticated state on a per-port basis.

configure netlogin ports mode

```
configure netlogin ports [all | <port_list>] mode [mac-based-vlans | port-based-vlans]
```

Description

Configures the network login port's mode of operation.

Syntax Description

<code>all</code>	Specifies all netlogin ports.
<code>port_list</code>	Specifies one or more network login ports.
<code>mac-based-vlans</code>	Allows more than one untagged VLAN.
<code>port-based-vlans</code>	Allows only one untagged VLAN. This is the default behavior.

Default

The default setting is `port-based-vlans`.

Usage Guidelines

Use this command to configure network login MAC-based VLANs on a network login port.

If you modify the mode of operation to `mac-based-vlans` and later disable all network login protocols on that port, the mode of operation automatically returns to `port-based-vlans`.

When you change the network login port's mode of operation, the switch deletes all currently known supplicants from the port and restores all VLANs associated with that port to their original state. In addition, by selecting `mac-based-vlans`, you are unable to manually add or delete untagged VLANs from this port. Network login now controls these VLANs.

With network login MAC-based operation, every authenticated client has an additional FDB flag that indicates a translation MAC address. If the supplicant's requested VLAN does not exist on the port, the switch adds the requested VLAN.

Important Rules and Restrictions

This section summarizes the rules and restrictions for configuring network login MAC-based VLANs:

- If you attempt to configure the port's mode of operation before enabling network login, the switch displays an error message similar to the following:

```
ERROR: The following ports do not have NetLogin enabled: 1
```

To enable network login on the switch, use the following command to enable network login and to specify an authentication method (for example, 802.1x—identified as `dot1x` in the CLI):

```
enable netlogin dot1x
```

To enable network login on the ports, use the following command to enable network login and to specify an authentication method (for example, 802.1x—identified as dot1x in the CLI):

```
enable netlogin ports 1:1 dot1x
```

Displaying FDB Information

To view network login-related FDB entries, use the following command:

```
show fdb netlogin [all | mac-based-vlans]
```

The following is sample output from the `show fdb netlogin mac-based-vlans` command:

Mac	Vlan	Age	Use	Flags	Port List
00:04:96:10:51:80	VLONE(0021)	0086	0000	n m	v 1:11
00:04:96:10:51:81	VLTWO(0051)	0100	0000	n m	v 1:11
00:04:96:10:51:91	VLTWO(0051)	0100	0000	n m	v 1:11

Flags: d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
 x - IPX, l - lockdown MAC, M - Mirror, B - Egress Blackhole,
 b - Ingress Blackhole, v - NetLogin MAC-Based VLAN.

The flags associated with network login include:

- **v**—Indicates the FDB entry was added because the port is part of a MAC-based virtual port/VLAN combination.
- **n**—Indicates the FDB entry was added by network login.

Displaying Port and VLAN Information

To view information about the VLANs that are temporarily added in MAC-based mode for network login, use the following command

```
show ports <port_list> information detail
```

The following is sample output from this command:

```
Port: 1
Virtual-router: VR-Default
Type:          UTP
Random Early drop:      Disabled
Admin state:    Enabled with auto-speed sensing auto-duplex
Link State:     Active, 100Mbps, full-duplex
Link Counter: Up      1 time(s)
VLAN cfg:
  Name: Default, Internal Tag = 1(MAC-Based), MAC-limit = No-limit
...<truncated output>
Egress 802.1p Replacement:      Disabled
NetLogin:                      Enabled
NetLogin authentication mode:   Mac based
```

```

NetLogin port mode:          MAC based VLANs
Smart redundancy:           Enabled
Software redundant port:    Disabled
auto-polarity:              Enabled

```

The added output displays information about the mode of operation for the network login port.

- **VLAN cfg**—The term MAC-based appears next to the tag number.
- **Netlogin port mode**—This output was added to display the port mode of operation. *Mac based* appears as the network login port mode of operation.

To view information about the ports that are temporarily added in MAC-based mode for network login, due to discovered MAC addresses, use the following command:

```
show vlan detail
```

The following is sample output from this command:

```

VLAN Interface with name Default created by user
Tagging:          802.1Q Tag 1
Priority:          802.1P Priority 0
Virtual router:   VR-Default
STPD:             s0(Disabled,Auto-bind)
Protocol:         Match all unfiltered protocols
Loopback:         Disable
NetLogin:         Disabled
Rate Shape:       Disabled
QosProfile:       None configured
Ports:  26.      (Number of active ports=2)
  Untag:  *1um,    *2,      3,      4,      5,      6,      7,
           8,      9,      10,     11,     12,     13,     14,
           15,     16,     17,     18,     19,     20,     21,
           22,     23,     24,     25,     26
Flags:  (*) Active, (!) Disabled, (g) Load Sharing port
        (b) Port blocked on the vlan, (a) Authenticated NetLogin Port
        (u) Unauthenticated NetLogin port, (m) Mac-Based port

```

The flags associated with network login include:

- **a**—Indicates an authenticated network login port.
- **u**—Indicates an unauthenticated network login port.
- **m**—Indicates that the network login port operates in MAC-based mode.

Example

The following command configures the network login ports mode of operation:

```
configure netlogin ports 1:1-1:10 mode mac-based-vlans
```

configure netlogin ports no-restart

```
configure netlogin ports [all | <port_list>] no-restart
```

Description

Disables the network login port restart feature.

Syntax Description

all	Specifies all network login ports.
port_list	Specifies one or more network login ports.

Default

The default setting is no-restart; the network login port restart feature is disabled.

Usage Guidelines

Use this command to disable the network login port restart feature on a network login port.

Configure network login port restart on ports with directly attached supplicants. If you use a hub to connect multiple supplicants, only the last unauthenticated supplicant causes the port to restart.

Displaying the Port Restart Configuration

To display the network login settings on the port, including the configuration for port restart, use the following command:

```
show netlogin port <port_list>
```

Output from this command includes the enable/disable state for network login port restart.

Example

The following command disables network login port restart on port 1:1:

```
configure netlogin ports 1:1 no-restart
```

configure netlogin ports restart

```
configure netlogin ports [all | <port_list>] restart
```

Description

Enables the network login port restart feature.

Syntax Description

all	Specifies all network login ports.
port_list	Specifies one or more network login ports.

Default

The default setting is no-restart; the network login port restart feature is disabled.

Usage Guidelines

Use this command to enable the network login port restart feature on a network login port. This allows network login to restart specific network login-enabled ports when the last authenticated supplicant releases, regardless of the configured protocols on the port.

Configure network login port restart on ports with directly attached supplicants. If you use a hub to connect multiple supplicants, only the last unauthenticated supplicant causes the port to restart.

Displaying the Port Restart Configuration

To display the network login settings on the port, including the configuration for port restart, use the following command:

```
show netlogin port <port_list>
```

Output from this command includes the enable/disable state for network login port restart.

Example

The following command enables network login port restart on port 1:1:

```
configure netlogin ports 1:1 restart
```

configure netlogin redirect-page

```
configure netlogin redirect-page <url>
```

Description

Configures the redirect URL for Network Login.

Syntax Description

url	Specifies the redirect URL for Network Login.
-----	---

Default

The redirect URL default value is `http://www.netgear.com`; the default port value is 80.

Usage Guidelines

In ISP mode, you can configure network login to be redirected to a base page after successful login using this command. If a RADIUS server is used for authentication, then base page redirection configured on the RADIUS server takes priority over this configuration.

You must configure a complete URL starting with `http://` or `https://`

You can also configure a specific port location at a specific target URL location. For example, you can configure a target port 8080 at `netgear.com` with the following command:

```
configure netlogin redirect-page "www.netgear.com:8080"
```

This command applies only to the web-based authentication mode of Network Login.

Example

The following command configures the redirect URL as `http://www.netgear.com/support`:

```
configure netlogin redirect-page http://www.netgear.com/support
```

configure netlogin session-refresh

```
configure netlogin session-refresh {<refresh_seconds>}
```

Description

Configures network login session refresh.

Syntax Description

<code>refresh_seconds</code>	Specifies the session refresh time for network login in seconds.
------------------------------	--

Default

Enabled, with a value of 180 seconds for session refresh.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the Logout link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to 360 seconds by default. The value can range from 1 to 3600 seconds. When you configure the network login session refresh for the logout window, ensure that the FDB aging timer is greater than the network login session refresh timer.

This command applies only to the web-based authentication mode of network login.

Example

The following command enables network login session refresh and sets the refresh time to 100 seconds:

```
configure netlogin session-refresh 100
```

configure netlogin vlan

```
configure netlogin vlan <vlan_name>
```

Description

Configures the VLAN for Network Login.

Syntax Description

vlan	Specifies the VLAN for Network Login.
------	---------------------------------------

Default

N/A.

Usage Guidelines

This command will configure the VLAN used for unauthenticated clients. One VLAN needs to be configured per VR. To change the VLAN, network login needs to be disabled. Network login can only be enabled when a VLAN is assigned (and no ports are configured for it).

By default no VLAN is assigned for network login.

Example

The following command configures the VLAN *login* as the network login VLAN:

```
configure netlogin vlan login
```

configure vlan netlogin-lease-timer

```
configure vlan <vlan name> netlogin-lease-timer <seconds>
```

Description

Configures the timer value returned as part of the DHCP response for clients attached to network login-enabled ports.

Syntax Description

vlan name	Specifies the VLAN to which this timer value applies.
seconds	Specifies the timer value, in seconds.

Default

10 seconds.

Usage Guidelines

The timer value is specified in seconds.
This command applies only to the web-based authentication mode of network login.

Example

The following command sets the timer value to 15 seconds for VLAN *corp*:

```
configure vlan corp netlogin-lease-timer 15
```

create netlogin local-user

```
create netlogin local-user <user-name> {encrypted} {<password>} {vlan-vsa [{tagged |  
untagged} [<vlan_name>] | <vlan_tag>]]} {security-profile <security_profile>}
```

Description

Creates a local network login user name and password.

Syntax Description

user-name	Specifies a new local network login user name. User names must have a minimum of 1 character and a maximum of 32 characters.
encrypted	The encrypted option is used by the switch to encrypt the password. Do not use this option through the command line interface (CLI).
password	Specifies a local network login user password. Passwords must have a minimum of 0 characters and a maximum of 32 characters.
tagged	Specifies that the client be added as tagged.
untagged	Specifies that the client be added as untagged.
vlan_name	Specifies the name of the destination VLAN.
vlan_tag	Specifies the VLAN ID, tag, of the destination VLAN.
security_profile	Specifies a security profile string during account creation.

Default

N/A.

Usage Guidelines

Use this command to create a local network login account and to configure the switch to use its local database for network login authentication. This method of authentication is useful in the following situations:

- If both the primary and secondary (if configured) RADIUS servers timeout or are unable to respond to authentication requests.
- If no RADIUS servers are configured.
- If the RADIUS server used for network login authentication is disabled.

If any of the above conditions are met, the switch checks for a local user account and attempts to authenticate against that local account.

NETGEAR recommends creating a maximum of 64 local accounts. If you need more than 64 local accounts, NETGEAR recommends using RADIUS for authentication. For more information about RADIUS authentication, see the *NETGEAR 8800 User Manual*.

You can also specify the destination VLAN to enter upon a successful authentication.

Note: If you do not specify a password or the keyword `encrypted`, you are prompted for one.

Additional Requirements

This command applies only to the web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.

You must have administrator privileges to use this command. If you do not have administrator privileges, the switch displays a message similar to the following:

```
This user does not have permissions for this command.
```

Both user names and passwords are case-sensitive. User names must have a minimum of 1 character and a maximum of 32 characters. Passwords must have a minimum of 0 characters and a maximum of 32 characters. If you use RADIUS for authentication, NETGEAR recommends that you use the same user name and password for both local authentication and RADIUS authentication.

If you attempt to create a user name with more than 32 characters, the switch displays the following messages:

```
%% Invalid name detected at '^' marker.
%% Name cannot exceed 32 characters.
```

If you attempt to create a password with more than 32 characters, the switch displays the following message after you re-enter the password:

```
Password cannot exceed 32 characters
```

Modifying an Existing Account

To modify an existing local network login account, use the following command:

```
configure netlogin local-user <user-name> {vlan-vsa [{tagged | untagged} [<vlan_name> |
<vlan_tag>]] | none}}
```

Displaying Local Network Login Accounts

To display a list of local network login accounts on the switch, including VLAN information, use the following command:

```
show netlogin local-users
```

Example

The following command creates a local network login user name and password:

```
create netlogin local-user megtest
```

After you enter the local network login user name, press [Enter]. The switch prompts you to enter a password (the switch does not display the password):

```
password:
```

After you enter the password, press [Enter]. The switch then prompts you to re-enter the password:

```
Reenter password:
```

The following command creates a local network login user name, password, and associates a destination VLAN with this account:

```
create netlogin local-user accounting vlan-vsa blue
```

As previously described, the switch prompts you to enter and confirm the password.

delete netlogin local-user

```
delete netlogin local-user <user-name>
```

Description

Deletes a specified local network login user name and its associated password.

Syntax Description

user-name	Specifies a local network login user name.
-----------	--

Default

N/A.

Usage Guidelines

Use the `show netlogin local-users` command to determine which local network login user name you want to delete from the system. The `show netlogin local-users` output displays the user name and password in a tabular format.

This command applies only to web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.

You must have administrator privileges to use this command.

Example

The following command deletes the local network login megtest along with its associated password:

```
delete netlogin local-user megtest
```

disable netlogin

```
disable netlogin [{dot1x} {mac} {web-based}]
```

Description

Disables network login modes.

Syntax Description

dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All types of authentication are disabled.

Usage Guidelines

Any combination of authentication types can be disabled on the same switch. To enable an authentication mode, use the following command:

```
enable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command disables MAC-based network login:

```
disable netlogin mac
```

disable netlogin authentication failure vlan ports

```
disable netlogin authentication failure vlan ports [<ports> | all]
```

Description

Disables the configured authentication failure VLAN on the specified ports.

Syntax Description

all	Specifies all ports included in the authentication failure VLAN.
ports	Specifies one or more ports or slots and ports on which the authentication failure VLAN is enabled.

Default

All ports.

Usage Guidelines

Use this command to disable the configured authentication failure VLAN on either the specified ports, or all ports.

disable netlogin authentication service-unavailable vlan ports

```
disable netlogin authentication service-unavailable vlan ports [ports | all]
```

Description

Disable the configured authentication service-unavailable VLAN on the specified ports.

Syntax Description

all	Specifies all ports included in the authentication service-unavailable VLAN.
ports	Specifies one or more ports or slots and ports on which the authentication service-unavailable VLAN is enabled.

Default

All ports.

Usage Guidelines

Use this command to disable the configured authentication service-unavailable VLAN on the specified ports, or on all ports.

disable netlogin dot1x guest-vlan ports

```
disable netlogin dot1x guest-vlan ports [all | <ports>]
```

Description

Disables the guest VLAN on the specified 802.1x network login ports.

Syntax Description

all	Specifies all ports included in the guest VLAN.
ports	Specifies one or more ports included in the guest VLAN.

Default

Disabled.

Usage Guidelines

Use this command to disable the guest VLAN feature.

Enabling Guest VLANs

To enable the guest VLAN, use the following command:

```
enable netlogin dot1x guest-vlan ports [all | <ports>]
```

Example

The following command disables the guest VLAN on all ports:

```
disable netlogin dot1x guest-vlan ports all
```

disable netlogin logout-privilege

```
disable network login logout-privilege
```

Description

Disables network login logout window pop-up.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command controls the logout window pop-up on the web-based network client. This command applies only to the web-based authentication mode of network login. When disabled, the logout window pop-up will no longer appear. However, if session refresh is enabled, the login session will be terminated after the session refresh timeout.

Example

The following command disables network login logout-privilege:

```
disable netlogin logout-privilege
```

disable netlogin ports

```
disable netlogin ports <ports> [{dot1x} {mac} {web-based}]
```

Description

Disables network login on a specified port for a particular method.

Syntax Description

ports	Specifies the ports for which network login should be disabled.
dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

Network login is disabled by default.

Usage Guidelines

Network login must be disabled on a port before you can delete a VLAN that contains that port.

This command applies to the MAC-based, web-based, and 802.1x mode of network login. To control which authentication mode is used by network login, use the following commands:

```
enable netlogin [{dot1x} {mac} {web-based}]
disable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command disables dot1x and web-based network login on port 2:9:

```
disable netlogin ports 2:9 dot1x web-based
```

disable netlogin reauthenticate-on-refresh

```
disable netlogin reauthenticate-on-refresh
```

Description

Disables network login reauthentication on refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The web-based Netlogin client's session is periodically refreshed by sending an HTTP request which acts as a keep-alive without actually re-authenticating the user's credentials with the back-end RADIUS server or local database. If `reauthenticate-on-refresh` is enabled, re-authentication occurs with the session refresh.

disable netlogin redirect-page

```
disable netlogin redirect-page
```

Description

Disables the network login redirect page function.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command disables the network login redirect page so that the client is sent to the originally requested page.

disable netlogin session-refresh

```
disable netlogin session-refresh
```

Description

Disables network login session refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the LogOut link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to three minutes by default.

This command applies only to the web-based authentication mode of network login.

Example

The following command disables network login session refresh:

```
disable netlogin session-refresh
```

enable netlogin

```
enable netlogin [{dot1x} {mac} {web-based}]
```

Description

Enables network login authentication modes.

Syntax Description

dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All types of authentication are disabled.

Usage Guidelines

Any combination of types of authentication can be enabled on the same switch. At least one of the authentication types must be specified on the command line.

To disable an authentication mode, use the following command:

```
disable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command enables web-based network login:

```
enable netlogin web-based
```

enable netlogin authentication failure vlan ports

```
enable netlogin authentication failure vlan ports [<ports> | all]
```

Description

Enables the configured authentication failure VLAN on the specified ports.

Syntax Description

all	Specifies all ports included in the authentication failure VLAN.
ports	Specifies one or more ports or slots and ports on which the authentication failure VLAN is enabled.

Default

All ports.

Usage Guidelines

Use this command to enable the configured authentication failure VLAN on either the specified ports, or all ports.

enable netlogin authentication service-unavailable vlan ports

```
enable netlogin authentication service-unavailable vlan ports [ports | all]
```

Description

Enables the configured authentication service-unavailable VLAN on the specified ports.

Syntax Description

all	Specifies all ports included in the service-unavailable VLAN.
ports	Specifies one or more ports or slots and ports on which the service-unavailable VLAN is enabled.

Default

All ports.

Usage Guidelines

Use this command to enable the configured authentication service-unavailable VLAN on the specified ports, or on all ports.

enable netlogin dot1x guest-vlan ports

```
enable netlogin dot1x guest-vlan ports [all | <ports>]
```

Description

Enables the guest VLAN on the specified 802.1x network login ports.

Syntax Description

all	Specifies all ports included in the guest VLAN.
ports	Specifies one or more ports or slots and ports on which the guest VLAN is enabled.

Default

Disabled.

Usage Guidelines

A guest VLAN provides limited or restricted network access if a supplicant connected to a port does not respond to the 802.1x authentication requests from the switch. A port always moves untagged into the guest VLAN.

Modifying the Supplicant Timer

By default, the switch attempts to authenticate the supplicant every 30 seconds for a maximum of three tries. If the supplicant does not respond to the authentication requests, the client moves to the guest VLAN. The number of authentication attempts is a user-configured parameter with allowed values in the range of 1 to 10.

To modify the supplicant response timer, use the following command and specify the `supp-resp-timeout` parameter:

```
configure netlogin dot1x timers [{server-timeout <server_timeout>} {quiet-period <quiet_period>} {reauth-period <reauth_period>} {reauth-max <max_num_reauths>}] {supp-resp-timeout <supp_resp_timeout>}]
```

Creating the Guest VLAN

Before you can enable the guest VLAN on the specified ports, you must create the guest VLAN. To create the guest VLAN, use the following command:

```
configure netlogin dot1x guest-vlan <vlan_name> {ports <port_list>}
```

Example

The following command enables the guest VLAN on all ports:

```
enable netlogin dot1x guest-vlan ports all
```

The following command enables the guest VLAN on ports 2 and 3:

```
enable netlogin dot1x guest-vlan ports 2,3
```

enable netlogin logout-privilege

```
enable netlogin logout-privilege
```

Description

Enables network login logout pop-up window.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command controls the logout window pop-up on the web-based network client. This command applies only to the web-based authentication mode of network login.

Example

The following command enables network login logout-privilege:

```
enable netlogin logout-privilege
```

enable netlogin ports

```
enable netlogin ports <ports> [{dot1x} {mac} {web-based}]
```

Description

Enables network login on a specified port for a particular authentication method.

Syntax Description

ports	Specifies the ports for which network login should be enabled.
dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All methods are disabled on all ports.

Usage Guidelines

For campus mode network login with web-based clients, the following conditions must be met:

- A DHCP server must be available, and a DHCP range must be configured for the port or ports in the VLAN on which you want to enable Network Login.
- The switch must be configured as a RADIUS client, and the RADIUS server must be configured to enable the network login capability.

For ISP mode login, no special conditions are required. A RADIUS server must be used for authentication.

Network login is used on a per port basis. A port that is tagged can belong to more than one VLAN. In this case, network login can be enabled on one port for each VLAN.

Windows authentication is not supported via network login.

Example

The following command configures network login on port 2:9 using web-based authentication:

```
enable netlogin ports 2:9 web-based
```

enable netlogin reauthentication-on-refresh

```
enable netlogin reauthentication-on-refresh
```

Description

Enables network login reauthentication on refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The web-based Netlogin client's session is periodically refreshed by sending a HTTP request which acts as a keep-alive without actually re-authenticating the user's credentials with the back-end RADIUS server or local database. If `reauthenticate-on-refresh` is enabled, re-authentication occurs with the session refresh.

enable netlogin redirect-page

```
enable netlogin redirect-page
```

Description

Enables the network login redirect page function.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command enables the network login redirect page so that the client is sent to the redirect page rather than the original page.

enable netlogin session-refresh

```
enable netlogin session-refresh {<refresh_minutes>}
```

Description

Enables network login session refresh.

Syntax Description

refresh_minutes	Specifies the session refresh time for network login in minutes.
-----------------	--

Default

Enabled, with a value of three minutes for session refresh.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the Logout link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to three minutes by default. The value can range from 1 to 255 minutes. When you configure the network login session refresh for the logout window, ensure that the FDB aging timer is greater than the network login session refresh timer.

This command applies only to the web-based authentication mode of network login.

To reset the session refresh value to the default behavior, use this command without the `minutes` parameter.

Example

The following command enables network login session refresh and sets the refresh time to ten minutes:

```
enable netlogin session-refresh 10
```

show banner netlogin

```
show banner netlogin
```

Description

Displays the user-configured banner string for network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed on the network login page.

Example

The following command displays the network login banner:

```
show banner netlogin
```

If a custom banner web page exists, `show banner netlogin` generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****
NOTE: Banner is not in use. Overridden since custom login page "netlogin_login_page.html" is
present.
```

If a custom banner web page does not exist, `show banner netlogin` generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****
```

show netlogin

```
show netlogin {port <portlist> vlan <vlan_name>} {dot1x {detail}} {mac} {web-based}
```

Description

Shows status information for network login.

Syntax Description

portlist	Specifies one or more ports or slots and ports.
vlan_name	Specifies the name of a VLAN.
dot1x	Specifies 802.1x information.
mac	Specifies MAC-based information.
web-based	Specifies web-based information.
detail	Shows detailed information.

Default

N/A.

Usage Guidelines

Depending on your configuration, software version, and the parameters you choose to display, the information reported by this command may include some or all of the following:

- Whether network login is enabled or disabled.
- The base-URL.
- The default redirect page.
- The logout privileges setting.
- The network login session-refresh setting and time.
- The MAC and IP address of supplicants.
- The type of authentication, 802.1x, MAC-based, or HTTP (web-based).
- The guest VLAN configurations, if applicable.
- The dynamic VLAN state and uplink ports, if configured.
- Whether network login port restart is enabled or disabled.
- Which order of authentication protocols is currently being used.

If you do not specify the authentication method, the switch displays information for all network login authentication methods.

Example

The following command shows the summary network login information:

```
show netlogin
```

The following is sample output from this command:

```
NetLogin Authentication Mode : web-based ENABLED; 802.1x ENABLED; mac-based ENABLED
NetLogin VLAN                : "nvlan"
NetLogin move-fail-action    : Authenticate
NetLogin Client Aging Time   : 5 minutes
```

```
Dynamic VLAN Creation      : Enabled
Dynamic VLAN Uplink Ports  : 12
```

```
-----
Web-based Mode Global Configuration
-----
```

```
Base-URL                  : network-access.com
Default-Redirect-Page     : http://www.yahoo.com
Logout-privilege          : YES
Netlogin Session-Refresh  : ENABLED; 3 minutes
Authentication Database    : Radius, Local-User database
-----
```

```
-----
802.1x Mode Global Configuration
-----
```

```
Quiet Period              : 60
Supplicant Response Timeout : 30
Re-authentication period   : 200
RADIUS server timeout      : 30
EAPOL MPDU version to transmit : v1
Authentication Database     : Radius
-----
```

```
-----
MAC Mode Global Configuration
-----
```

MAC Address/Mask	Password (encrypted)	Port(s)
00:00:86:3F:1C:35/48	yaqu	any
00:01:20:00:00:00/24	yaqu	any
00:04:0D:28:45:CA/48	=4253C5;500@	any
00:10:14:00:00:00/24	yaqu	any
00:10:A4:A9:11:3B/48	yaqu	any
00:10:A4:00:00:00/24	yaqu	any
Default	yaqu	any

```
Authentication Database      : Radius, Local-User database
-----
```

```
Port: 5, Vlan: nvlan, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled
```

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
-----	------------	---------------	------	--------------	------

Port: 9, Vlan: nvlan, State: Enabled, Authentication: web-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
-----	------------	---------------	------	--------------	------

Port: 10, Vlan: nvlan, State: Enabled, Authentication: 802.1x, mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
-----	------------	---------------	------	--------------	------

Port: 17, Vlan: engr, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
-----	------------	---------------	------	--------------	------

Port: 17, Vlan: mktg, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
-----	------------	---------------	------	--------------	------

Port: 19, Vlan: corp, State: Enabled, Authentication: 802.1x, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	No		0	00040D50E13A
00:10:dc:98:54:00	10.201.31.113	Yes, Radius	802.1x	24	md5isp7

Port: 19, Vlan: nvlan, State: Enabled, Authentication: 802.1x, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	No	802.1x	0	

Port: 19, Vlan: voice-ip, State: Enabled, Authentication: 802.1x, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	Yes, Radius	802.1x	75	00040D50E13A

The following command shows more detailed information, including the configured authentication methods:

```
show netlogin port 3:2 vlan "Default"
```

The following is sample output from this command:

```
Port: 2:1          Vlan: Default
Authentication: Web-Based, 802.1x
Port State:       Unauthenticated
Guest VLAN:       Not Enabled
DHCP:             Not Enabled
```

MAC	IP address	Auth	Type	ReAuth-Timer	User
00:0C:F1:E8:4E:13	0.0.0.0	No	802.1x	0	Unknown
00:01:30:F3:EA:A0	10.0.0.1	Yes	802.1x	0	testUser

The following command shows information about a specific port configured for network login:

```
show netlogin port 1:1
```

The following is sample output from this command:

```
Port          : 1:1
Port Restart  : Enabled
Vlan          : Default
Authentication: mac-based
Port State    : Enabled
Guest Vlan    : Disabled
```

MAC	IP address	Auth	Type	ReAuth-Timer	User

The following command shows the details of the 802.1x mode:

```
show netlogin dot1x detail
```

The following is sample output from this command:

```
NetLogin Authentication Mode : web-based DISABLED; 802.1x ENABLED; mac-based DISABLED
NetLogin VLAN                : "nl"
NetLogin move-fail-action    : Deny
```

```
-----
```

802.1x Mode Global Configuration

```
-----
```

```
Quiet Period           : 30
Supplicant Response Timeout : 30
Re-authentication period : 3600
RADIUS server timeout   : 30
EAPOL MPDU version to transmit : v1
Guest VLAN              : destVlan
```

```
-----
Port: 1:1, Vlan: Default, State: Enabled, Authentication: 802.1x, Guest Vlan: destVlan
```

```
MAC
00:00:86:53:c3:14 : IP=0.0.0.0      Auth=Yes User= testUser
                  : AuthPAE state=AUTHENTICATED BackAuth state=IDLE
                  : ReAuth time left=3595      ReAuth count=1
                  : Quiet time left=37
00:02:03:04:04:05 : IP=0.0.0.0      Auth=No  User=
                  : AuthPAE state=CONNECTING  BackAuth state=IDLE
                  : ReAuth time left=0        ReAuth count=2
                  : Quiet time left=37
-----
```

For 802.1x, if re-authentication is disabled, the re-authentication period appears as follows:

```
Re-authentication period      : 0 (Re-authentication disabled)
```

The following command:

```
show netlogin port 5:4 dot1x
```

Generates the following sample output:

```
Port          : 5:4
Port Restart  : Disabled
Vlan          : corp
Authentication : 802.1x
Port State    : Enabled
Guest Vlan    : Enabled
```

```
MAC           IP address   Authenticated  Type    ReAuth-Timer  User
00:10:dc:92:53:2d  10.201.31.119 Yes,Radius     802.1x  14            md5isp4
-----
```

The following command:

```
sh netlogin port 5:4 dot1x detail
```

Generates the following sample output:

```
Port          : 5:4
Port Restart  : Disabled
Vlan          : corp
Authentication : 802.1x
Port State    : Enabled
Guest Vlan    : Enabled
```

```
MAC
00:10:dc:92:53:2d : IP=10.201.31.119 Auth=Yes User=md5isp4
```

```

: AuthPAE state=AUTHENTICATED BackAuth state=IDLE
: ReAuth time left=8          ReAuth count=0
: Quiet time left=0
-----

```

show netlogin authentication failure vlan

```
show netlogin authentication failure vlan {<vlan_name>}
```

Description

Displays the authentication failure VLAN related configuration details.

Syntax Description

vlan_name	Specifies the name of a failure VLAN.
-----------	---------------------------------------

Default

N/A.

Usage Guidelines

Use this command to display configuration details for the authentication failure VLAN.

Example

The following command displays :

```
show netlogin authentication failure vlan
```

The following is sample output from this command:

```

-----
Authentication Service unavailable Vlanport          Status
-----
corp 1:2                      Disabled

```

show netlogin authentication service-unavailable vlan

```
show netlogin authentication service-unavailable vlan {<vlan_name>}
```

Description

Displays the authentication service-unavailable VLAN related configuration details.

Syntax Description

vlan_name	Specifies the name of the authentication service-unavailable VLAN.
-----------	--

Default

N/A.

Usage Guidelines

Use this command to display configuration details for the service-unavailable VLAN.

Example

The following command displays:

```
show netlogin authentication service-unavailable vlan
```

The following is sample output from this command:

server-unavailable Vlan	port	Status

xyz	2:1	Disabled
abc	3:1	Enabled

show netlogin guest-vlan

```
show netlogin guest-vlan {vlan_name}
```

Description

Displays the configuration for the guest VLAN feature.

Syntax Description

vlan_name	Specifies the name of a guest VLAN.

Default

N/A.

Usage Guidelines

Use this command to display the guest VLANs configured on the switch.

If you specify the `vlan_name`, the switch displays information for only that guest VLAN.

The output displays the following information in a tabular format:

- `Port`—Specifies the 802.1x enabled port configured for the guest VLAN.
- `Guest-vlan`—Displays the enabled/disabled state of the guest VLAN feature.
- `Vlan`—Specifies the name of the guest VLAN.

Example

The following command displays the local network login list:

```
show netlogin guest-vlan
```

The following is sample output from this command:

Port	Guest-vlan	Vlan

5:1	Disabled	gv11
5:2	Enabled	gv12
5:3	Disabled	gv13
5:4	Enabled	gv14

show netlogin local-users

```
show netlogin local-users
```

Description

Displays the local network login users configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the list of local network login users and associated VLANs.

If you associated a VLAN with a local network login user, the output displays the name of the VLAN. If you have not associated a VLAN with a local network login user, the output displays not configured.

The Extended-VLAN VSA column displays the name of the VLAN and the following information:

- `<not configured>`—Specifies that you have not associated a VLAN with a local network login user.
- `*`—Specifies the movement based on the incoming port's traffic. For example, the VLAN behaves like VSA 203 if identified with a VLAN name or VSA 209 if identified with a VLAN ID.
- `T`—Specifies a tagged client.
- `U`—Specifies an untagged client.

In addition, this output is useful to determine which local network login user you want to modify or delete from the system.

Example

The following command displays the local network login list:

```
show netlogin local-users
```

The following is sample output from this command:

Netlogin Local User Name	Password (encrypted)	Extended-VLAN VSA
000000000012	Iqyydz\$MP7AG.VAmwOoqiKX2u13H1	U hallo
00008653C314	BoO28L\$oRVvKv8.wmxcorhhXxQY40	* default
megtest	w7iMbp\$1BL34/dLx4G4M8aAdiCvI	<not configured>
testUser	/Jhouw\$iHE15steebwhOibgj6pZq.	T testVlan

show netlogin mac-list

```
show netlogin mac-list
```

Description

Displays the MAC address list for MAC-based network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the MAC address list used for MAC-based network login.

MAC-based authentication is VR aware, so there is one MAC list per VR.

Example

The following command displays the MAC address list:

```
show netlogin mac-list
```

The following is sample output from this command:

MAC Address/Mask	Password (encrypted)	Port(s)
00:00:00:00:00:10/48	<not configured>	1:1-1:5
00:00:00:00:00:11/48	<not configured>	1:6-1:10
00:00:00:00:00:12/48	<not configured>	any
00:01:30:70:0C:00/48	yaqu	any
00:01:30:32:7D:00/48	ravdqsr	any
00:04:96:00:00:00/24	<not configured>	any

unconfigure netlogin allowed-refresh-failures

```
unconfigure netlogin allowed-refresh-failures
```

Description

Restores the number refresh failures to the default value.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command allows you to restore the number of refresh failures allowed to the default value of 0.

unconfigure netlogin authentication database-order

```
unconfigure netlogin [mac | web-based] authentication database-order
```

Description

Restores the default order of database authentication protocols to use.

Syntax Description

mac	Specifies the MAC address to add.
mask	Specifies the number of bits to use for the mask.
default	Specifies the default entry.
encrypted	Used to display encrypted form of password in configuration files. Do not use.
password	Specifies the password to send for authentication.
ports	Specifies the port or port list to use for authentication.

Default

By default, the authentication order is RADIUS, local-user database.

Usage Guidelines

Use this command to restore the default configuration order for the database authentication protocols. For details see [configure netlogin authentication database-order](#).

Example

The following command sets the database authentication order to RADIUS, local user database for MAC-based authentication:

```
unconfigure netlogin mac authentication database-order
```

unconfigure netlogin authentication failure vlan

```
unconfigure netlogin authentication failure vlan <vlan_name> {ports <port_list>}
```

Description

Disables authentication failure VLAN on network login enabled ports.

Syntax Description

vlan_name	Specifies the name of the authentication failure VLAN.
port_list	Specifies one or more ports or slots and ports. If the <code>ports</code> keyword is not used, the command applies to all ports.

Default

N/A.

Usage Guidelines

Use this command to disable authentication failure VLAN on network login enabled ports. When a supplicant fails authentication, it is moved to the authentication failure vlan and is given limited access until it passes the authentication.

unconfigure netlogin authentication service-unavailable vlan

```
unconfigure netlogin authentication service-unavailable vlan <vlan_name>
{ports <port_list>}
```

Description

Unconfigures authentication service unavailable VLAN on network login enabled ports.

Syntax Description

vlan_name	Specifies the name of the authentication service-unavailable VLAN.
port_list	Specifies one or more ports or slots and ports. If the <code>ports</code> keyword is not used, the command applies to all ports.

Default

Defaults to all network login enabled ports.

Usage Guidelines

This command unconfigures authentication service unavailable VLAN on the specified network login enabled ports. Authentication service unavailable VLAN is unconfigured on all the network login enabled ports, if no port is specifically mentioned.

unconfigure netlogin dot1x guest-vlan

```
unconfigure netlogin dot1x guest-vlan {ports <port_list> | <vlan_name>}
```

Description

Unconfigures the guest VLAN feature for 802.1x authentication.

Syntax Description

ports_list	Specifies one or more ports included in the guest VLAN.
vlan_name	Specifies all ports included in the guest VLAN.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the guest VLAN for 802.1x authentication.

If you do not specify one or more ports or the VLAN name, this command unconfigures all of the 802.1x ports configured for the guest VLAN feature.

If you specify one or more ports, this command unconfigures the specified 802.1x ports for the guest VLAN feature.

If you specify the VLAN name, this command unconfigures all of the 802.1x ports configured for the specified guest VLAN.

Example

The following command unconfigures the guest VLAN feature for 802.1x authentication:

```
unconfigure netlogin dot1x guest-vlan
```

unconfigure netlogin local-user security-profile

```
unconfigure netlogin local-user <user-name> security-profile
```

Description

Clears a previously associated security profile.

Syntax Description

user-name	Specifies the name of an existing local network login account.
-----------	--

Default

N/A.

Usage Guidelines

Use this command to clear any previously associated security profiles on the switch.

unconfigure netlogin session-refresh

```
unconfigure netlogin session-refresh
```

Description

Restores the session refresh value to the default.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command allows you to restore the session refresh to the default value of 180 seconds.

unconfigure netlogin vlan

```
unconfigure netlogin vlan
```

Description

Unconfigures the VLAN for network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command unconfigures the VLAN used for unauthenticated clients. One VLAN needs to be configured per VR. To change the VLAN, network login needs to be disabled.

Example

The following command unconfigures the network login VLAN:

```
unconfigure netlogin vlan
```


This chapter describes commands for:

- Creating, configuring, enabling, and disabling Spanning Tree Protocol (STP) on the switch
- Enabling and disabling Rapid Spanning Tree Protocol (RSTP) on the switch
- Enabling and disabling Multiple Spanning Tree Protocol (MSTP) on the switch
- Displaying and resetting STP settings on the switch

STP

STP is a bridge-based mechanism for providing fault tolerance on networks. STP is a part of the 802.1D bridge specification defined by the IEEE Computer Society. To explain STP in terms used by the 802.1D specification, the switch is referred to as a bridge.

STP allows you to implement parallel paths for network traffic and ensure that redundant paths are:

- Disabled when the main paths are operational.
- Enabled if the main path fails.

RSTP

The Rapid Spanning Tree Protocol (RSTP) IEEE 802.1w provides an enhanced spanning tree algorithm that improves the convergence speed of bridged networks. RSTP takes advantage of point-to-point links in the network and actively confirms that a port can safely transition to the forwarding state without relying on any timer configurations. If a network topology change or failure occurs, RSTP rapidly recovers network connectivity by confirming the change locally before propagating that change to other devices across the network. For broadcast links, there is no difference in convergence time between STP and RSTP.

RSTP supersedes legacy STP protocols, supports the existing STP parameters and configurations, and allows for seamless interoperability with legacy STP.

MSTP

MSTP logically divides a Layer 2 network into regions. Each region has a unique identifier and contains multiple spanning tree instances (MSTIs). An MSTI is a spanning tree domain that operates within and is bounded by a region. MSTIs control the topology inside the regions. The Common and Internal Spanning Tree (CIST) is a single spanning tree domain that interconnects MSTP regions. The CIST is responsible for creating a loop-free topology by exchanging and propagating BPDUs across regions to form a Common Spanning Tree (CST).

MSTP uses RSTP as its converging algorithm and is interoperable with the legacy STP protocols: STP (802.1D) and RSTP (802.1w).

Spanning Tree Domains

The switch can be partitioned into multiple virtual bridges. Each virtual bridge can run an independent spanning tree instance. Each spanning tree instance is called a Spanning Tree Domain (STPD). Each STPD has its own root bridge and active path. After an STPD is created, one or more VLANs can be assigned to it.

A port can belong to multiple STPDs. In addition, a VLAN can span multiple STPDs.

The key points to remember when configuring VLANs and STP are:

- Each VLAN forms an independent broadcast domain.
- STP blocks paths to create a loop-free environment.
- Within any given STPD, all VLANs belonging to it use the same spanning tree.

Member VLANs

When you add a VLAN to an STPD, that VLAN becomes a member of the STPD. The two types of member VLANs in an STPD are:

- Carrier
- Protected

Carrier VLAN

A carrier VLAN defines the scope of the STPD, which includes the physical and logical ports that belong to the STPD and if configured, the 802.1Q tag used to transport Multiple Instance Spanning Tree Protocol (EMISTP) or Per VLAN Spanning Tree (PVST+) encapsulated Bridge Protocol Data Units (BPDUs). Only one carrier VLAN can exist in a given STPD, although some of its ports can be outside the control of any STPD at the same time.

Note: If you use EMISTP or PVST+, the STPD ID must be identical to the VLAN ID of the carrier VLAN in that STPD.

If you have an 802.1D configuration, NETGEAR recommends that you configure the StpdID to be identical to the VLAN ID of the carrier VLAN in that STPD.

If you configure MSTP, you do not need carrier VLANs for MSTP operation. With MSTP, you configure a CIST that controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate the status of MSTP regions. All VLANs participating in the MSTP region have the same privileges.

Protected VLAN

Protected VLANs are all other VLANs that are members of the STPD. These VLANs “piggyback” on the carrier VLAN. Protected VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes and inherit the state of the carrier VLAN. Protected VLANs can participate in multiple STPD, but any particular port in the VLAN can belong to only *one* STPD. Also known as non-carrier VLANs.

If you configure MSTP, all member VLANs in an MSTP region are protected VLANs. These VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes communicated by the CIST to the MSTP regions. MSTIs cannot share the same protected VLAN; however, any port in a protected VLAN can belong to multiple MSTIs.

STPD Modes

An STPD has three modes of operation:

- 802.1D mode

Use this mode for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1D. When configured in this mode, all rapid configuration mechanisms are disabled.

- 802.1w mode

Use this mode for compatibility with Rapid Spanning Tree (RSTP). When configured in this mode, all rapid configuration mechanisms are enabled. The benefit of this mode is available on point-to-point and edge ports only.

You enable or disable RSTP on a per STPD basis only. You do not enable RSTP on a per port basis.

- MSTP mode

Use this mode for compatibility with Multiple Spanning Tree (MSTP, 802.1s). MSTP is an extension of RSTP and offers the benefit of better scaling with fast convergence. When configured in this mode, all rapid configuration mechanisms are enabled. The benefit of MSTP is available only on point-to-point links and when you configure the peer in MSTP

or 802.1w mode. If you do not select point-to-point links and the peer is not configured in 802.1w mode, the STPD fails back to 802.1D mode.

You can create only one MSTP region on the switch, and all switches that participate in the region must have the same regional configurations. You enable or disable an MSTP on a per STPD basis only. You do not enable MSTP on a per port basis.

By default, the:

- STPD operates in 802.1D mode.
- Default device configuration contains a single STPD called *s0*.
- Default VLAN is a member of STPD *s0* with autobind enabled.

All STP parameters default to the IEEE 802.1D values, as appropriate.

Encapsulation Modes

You can configure ports within an STPD to accept and transmit specific BPDU encapsulations. This STP port encapsulation is separate from the STP mode of operation. For example, you can configure a port to accept the PVST+ BPDU encapsulation while running in 802.1D mode.

An STP port has three possible encapsulation modes:

- 802.1D mode

This mode is used for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1D. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- Multiple Instance Spanning Tree Protocol (EMISTP) mode

EMISTP mode is proprietary to NETGEAR and is an extension of STP that allows a physical port to belong to multiple STPDs by assigning the port to multiple VLANs. EMISTP adds significant flexibility to STP network design. BPDUs are sent with an 802.1Q tag having an STPD instance Identifier (STPD ID) in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- Per VLAN Spanning Tree (PVST+) mode

This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains to which it belongs.

MSTP STPDs use only 802.1D BPDU encapsulation mode. The switch prevents you from configuring EMISTP or PVST+ encapsulation mode for MSTP STPDs.

STP Rules and Restrictions

This section summarizes the rules and restrictions for configuring STP as follows:

- The carrier VLAN must span all ports of the STPD. (This is not applicable to MSTP.)
- The STPD ID must be the VLAN ID of the carrier VLAN; the carrier VLAN cannot be partitioned. (This is not applicable to MSTP.)
- A default VLAN cannot be partitioned. If a VLAN traverses multiple STPDs, the VLAN must be tagged.
- An STPD can carry, at most, one VLAN running in PVST+ mode, and its STPD ID must be identical with that VLAN ID. In addition, the PVST+ VLAN cannot be partitioned.
- The default VLAN of a PVST+ port must be identical with the native VLAN on the PVST+ device connected to that port.
- If an STPD contains both PVST+ and non-PVST+ ports, that STPD must be enabled. If that STPD is disabled, the BPDUs are flooded in the format of the incoming STP port, which may be incompatible with those of the connected devices.
- The 802.1D ports must be untagged; and the EMISTP/PVST+ ports must be tagged in the carrier VLAN.
- An STPD with multiple VLANs must contain only VLANs that belong to the same virtual router instance.
- STP and network login operate on the same port as follows:
 - STP (802.1D), RSTP (802.1W), and MSTP (802.1S) support both network login and STP on the same port.
 - At least one VLAN on the intended port should be configured both for STP and network login.
 - When STP blocks a port, network login does not process authentication requests and BPDUs are the only traffic in and out of the port. All user data forwarding stops.
 - When STP places a port in forwarding state, network login operates and BPDUs and user data flow in and out of the port. The forwarding state is the only STP state that allows network login and user data forwarding.
 - When RSTP is used with network login campus mode, autobind must be enabled on all VLANs that support RSTP and network login campus mode.
 - When RSTP is used with network login campus mode on a port, dynamic VLANs cannot be supported.
- STP cannot be configured on the following ports:
 - A mirroring target port.

- A software-controlled redundant port.
- MSTP and 802.1D STPDs cannot share a physical port.
- Only one MSTP region can be configured on a switch.
- In an MSTP environment, A VLAN can belong to either a CIST or one of the MSTIs.
- A VLAN can belong to only one MSTP domain.
- MSTP is not interoperable with PVST+.
- The CIST can operate without any member VLANs.

clear counters stp

```
clear counters stp {[all | diagnostics | domains | ports]}
```

Description

Clears, resets all STP statistics and counters.

Syntax Description

all	Specifies all STP domain, port, and diagnostics counters.
diagnostics	Specifies STP diagnostics counters.
domains	Specifies STP domain counters.
ports	Specifies STP port counters.

Default

N/A.

Usage Guidelines

If you do not enter a parameter, the result is the same as specifying the `all` parameter: the counters for all domains, ports, and diagnostics are reset.

Enter one of the following parameters to reset the STP counters on the switch:

- `all`—Specifies the counters for all STPDs and ports, and clears all STP counters
- `diagnostics`—Clears the internal diagnostic counters
- `domains`—Clears the domain level counters
- `ports`—Clears the counters for all ports and leaves the domain level counters

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counters, you can see fresh statistics for the time period that you are monitoring.

Example

The following command clears all of the STP domain, port, and diagnostic counters:

```
clear counters stp
```

configure mstp format

```
configure mstp format <format_identifier>
```

Description

Configures the number used to identify the MSTP BPDUs sent in the MSTP region.

Syntax Description

format_identifier	Specifies a number that MSTP uses to identify all BPDUs sent in the MSTP region. The default is 0. The range is 0 to 255.
-------------------	---

Default

The default value used to identify the MSTP BPDU is 0.

Usage Guidelines

For a switch to be part of an MSTP region, you must configure each switch in the region with the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- **Region Name**—The name of the MSTP region.
- **Format Selector**—The number used to identify the format of MSTP BPDUs. The default is 0.
- **Revision Level**—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

You can configure only one MSTP region on the switch at any given time.

The switches contained in a region transmit and receive BPDUs that contain information relevant to only that MSTP region. By having devices look at the region identifiers, MSTP discovers the logical boundary of a region.

If you have an active MSTP region, NETGEAR recommends that you disable all active STPDs in the region before modifying the value used to identify MSTP BPDUs on all participating switches.

Example

The following command configures the number 2 to identify the MSTP BPDUs sent within an MSTP region:

```
configure mstp format 2
```

configure mstp region

```
configure mstp region <regionName>
```

Description

Configures the name of an MSTP region on the switch.

Syntax Description

regionName	Specifies a user-defined name for the MSTP region. May be up to 32 characters.
------------	--

Default

By default, the switch uses the MAC address of the switch to generate an MSTP region.

Before you configure the MSTP region, it also has the following additional defaults:

- MSTP format Identifier—0
- MSTP Revision Level—3

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (_) but cannot be any reserved keywords, for example, mstp. Names must start with an alphabetical character, for example, a, Z.

By default, the switch uses the unique MAC address of the switch to generate an MSTP region. Since each MAC address is unique, every switch is in its own region by default.

For multiple switches to be part of an MSTP region, you must configure each switch in the region with the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

You can configure only one MSTP region on the switch at any given time.

The switches inside a region exchange BPDUs that contain information for MSTIs. The switches connected outside of the region exchange CIST information. By having devices look at the region identifiers, MSTP discovers the logical boundary of a region.

If you have an active MSTP region, NETGEAR recommends that you disable all active STPDs in the region before renaming the region on all of the participating switches.

Viewing MSTP Information

To view the MSTP configuration on the switch, use the `show stpd` command. Output from this command contains global MSTP settings, including the name of the MSTP region, the number or tag that identifies all of the BPDUs sent in the MSTP region, and the reserved MSTP revision level. If configured, the output also displays the name of the Common and Internal Spanning Tree (CIST), and the number of Multiple Spanning Tree Instances (MSTIs).

Example

The following command creates an MSTP region named purple:

```
configure mstp region purple
```

configure mstp revision

```
configure mstp revision <revision>
```

Description

Configures the revision number of the MSTP region.

Syntax Description

revision	This parameter is reserved for future use.
----------	--

Default

The default value of the revision level is 3.

Usage Guidelines

Although this command is displayed in the CLI, it is reserved for future use. Please do not use this command.

If you accidentally configure this command, remember that each switch in the region must have the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

Example

The following command returns the MSTP revision number to 3, the default revision number:

```
configure mstp revision 3
```

configure stpd add vlan

```
configure stpd <stpd_name> add vlan <vlan_name> ports [all | <port_list>] {[dot1d | emistp | pvst-plus]}
```

Description

Adds all ports or a list of ports within a VLAN to a specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
vlan_name	Specifies a VLAN name.
all	Specifies all of the ports in the VLAN to be included in the STPD.
port_list	Specifies the port or ports to be included in the STPD.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1D.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are in `dot1d` mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

To create an STP domain, use the `create stpd` command. To create a VLAN, use the `create vlan` command.

In an EMISTP or PVST+ environment, this command adds a list of ports within a VLAN to a specified STPD provided the carrier VLAN already exists on the same set of ports. You can also specify the encapsulation mode for those ports.

In an MSTP environment, you do not need a carrier VLAN. A CIST controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate region status. You must use the `dot1d` encapsulation mode in an MSTP environment.

You cannot configure STP on the following ports:

- Mirroring target ports
- Software-controlled redundant ports

If you see an error similar to the following:

```
Error: Cannot add VLAN default port 3:5 to STP domain
```

You might be attempting to add:

- A carrier VLAN port to a different STP domain than the carrier VLAN belongs

- A VLAN/port for which the carrier VLAN does not yet belong

Note: This restriction is only enforced in an active STP domain and when you enable STP to make sure you have a legal STP configuration.

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.

By default, when the switch boots for the first time, it automatically creates a VLAN named *default* with a tag value of 1 and STPD *s0*. The switch associates VLAN *default* to STPD *s0*. All ports that belong to this VLAN and STPD are in 802.1D encapsulation mode with autobind enabled. If you disable autobind on the VLAN *default*, that configuration is saved across a reboot.

Naming Conventions

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keywords `stpd` and `vlan` are optional.

STP Encapsulations Modes

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

STPD Identifier

An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STPD and that VLAN cannot belong to another STPD.

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the Common and Internal Spanning Tree (CIST). The switch assigns this ID automatically when you configure the CIST STPD. A multiple spanning tree instance identifier identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that participates in the MSTP region. In an MSTP region, MSTI IDs only have local significance. You can reuse MSTI IDs across MSTP regions.

Automatically Inheriting Ports—MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Example

Create a VLAN named *marketing* and an STPD named *STPD1* as follows:

```
create vlan marketing
create stpd stpd1
```

The following command adds the VLAN named *marketing* to the STPD *STPD1*, and includes all the ports of the VLAN in *STPD1*:

```
configure stpd stpd1 add vlan marketing ports all
```

configure stpd default-encapsulation

```
configure stpd <stpd_name> default-encapsulation [dot1d | emistp | pvst-plus]
```

Description

Configures the default encapsulation mode for all ports added to the specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.

`pvst-plus`

Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (`s0`) are `dot1d` mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.

By default, when the switch boots for the first time, it automatically creates a VLAN named *default* with a tag value of 1 and STPD *s0*. The switch associates VLAN *default* to STPD *s0*. All ports that belong to this VLAN and STPD are in 802.1d encapsulation mode with autobind enabled. If you disable autobind on the VLAN *default*, that configuration is saved across a reboot.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

Naming Conventions

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional. For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

STP Encapsulation Modes

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

Note: These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

STPD Identifier

An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STP domain, and that VLAN cannot belong to another STPD.

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the Common and Internal Spanning Tree (CIST). The switch assigns this ID automatically when you configure the CIST STPD. A multiple spanning tree instance identifier identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that participates in the MSTP region. In an MSTP region, MSTI IDs only have local significance. You can reuse MSTI IDs across MSTP regions.

Example

The following command specifies that all ports subsequently added to the STPD *STPD1* be in PVST+ encapsulation mode unless otherwise specified or manually changed:

```
configure stpd stpd1 default-encapsulation pvst-plus
```

configure stpd delete vlan

```
configure stpd <stpd_name> delete vlan <vlan_name> ports [all | <port_list>]
```

Description

Deletes one or more ports in the specified VLAN from an STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
vlan_name	Specifies a VLAN name.
all	Specifies that all of the ports in the VLAN are to be removed from the STPD.
port_list	Specifies the port or ports to be removed from the STPD.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keywords `stpd` and `vlan` are optional.

In EMISTP and PVST+ environments, if the specified VLAN is the carrier VLAN, all protected VLANs on the same set of ports are also removed from the STPD.

You also use this command to remove autobind ports from a VLAN. The NETGEAR 8800 records the deleted ports so that the ports are not automatically added to the STPD after a system restart.

When a port is deleted on the MSTI, it is automatically deleted on the CIST as well.

Example

The following command removes all ports of a VLAN named *Marketing* from the STPD *STPD1*:

```
configure stpd stpd1 delete vlan marketing ports all
```

configure stpd forwarddelay

```
configure stpd <stpd_name> forwarddelay <seconds>
```

Description

Specifies the time (in seconds) that the ports in this STPD spend in the listening and learning states when the switch is the root bridge.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>seconds</code>	Specifies the forward delay time in seconds. The default is 15 seconds, and the range is 4 to 30 seconds.

Default

The default forward delay time is 15 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the `<seconds>` parameter is 4 through 30 seconds.

Example

The following command sets the forward delay from *STPD1* to 20 seconds:

```
configure stpd stpd1 forwarddelay 20
```

configure stpd hellotime

```
configure stpd <stpd_name> hellotime <seconds>
```

Description

Specifies the time delay (in seconds) between the transmission of BPDUs from this STPD when it is the root bridge.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>seconds</code>	Specifies the hello time in seconds. The default is 2 seconds, and the range is 1 to 10 seconds.

Default

The default hello time is 2 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

In an MSTP environment, configure the hello timer only on the CIST, not on the MSTIs.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the `<seconds>` parameter is 1 through 10 seconds.

Example

The following command sets the time delay from *STPD1* to 10 seconds:

```
configure stpd stpd1 hellotime 10
```

configure stpd maxage

```
configure stpd <stpd_name> maxage <seconds>
```


Description

Specifies the maximum age of a BPDU in the specified STPD.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>seconds</code>	Specifies the maxage time in seconds. The default is 20 seconds, and the range is 6 to 40 seconds.

Default

The default maximum age of a BPDU is 20 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

In an MSTP environment, configure the maximum age of a BPDU only on the CIST, not on the MSTIs.

The range for the `<seconds>` parameter is 6 through 40 seconds.

Note that the time must be greater than, or equal to $2 * (\text{Hello Time} + 1)$ and less than, or equal to $2 * (\text{Forward Delay} - 1)$.

Example

The following command sets the maximum age of *STPD1* to 30 seconds:

```
configure stpd stpd1 maxage 30
```

configure stpd max-hop-count

```
configure stpd <stpd_name> max-hop-count <hopcount>
```

Description

Specifies the maximum hop count of a BPDU until the BPDU is discarded in the specified MSTP STP domain.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
------------------------	---------------------------------------

hopcount	Specifies the number of hops required to age out information and notify changes in the topology. The default is 20 hops, and the range is 6 to 40 hops.
----------	---

Default

The default hop count of a BPDU is 20 hops.

Usage Guidelines

This command is applicable only in an MSTP environment.

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the `<hopcount>` parameter is 6 through 40 hops.

In an MSTP environment, the hop count has the same purpose as the maxage timer for 802.1D and 802.1w environments.

The main responsibility of the CIST is to exchange or propagate BPDUs across regions. The switch assigns the CIST an instance ID of 0, which allows the CIST to send BPDUs for itself in addition to all of the MSTIs within an MSTP region. Inside a region, the BPDUs contain CIST records and piggybacked M-records. The CIST records contain information about the CIST, and the M-records contain information about the MSTIs. Boundary ports only exchange CIST record BPDUs.

On boundary ports, only CIST record BPDUs are exchanged. In addition, if the other end is an 802.1D or 802.1w bridge, the maxage timer is used for interoperability between the protocols.

Example

The following command sets the hop of the MSTP STPD, STPD2, to 30 hops:

```
configure stpd stpd2 max-hop-count 30
```

configure stpd mode

```
configure stpd <stpd_name> mode [dot1d | dot1w | mstp [cist | msti <instance>]]
```

Description

Configures the operational mode for the specified STP domain.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
dot1d	Specifies the STPD mode of operation to be 802.1D.
dot1w	Specifies the STPD mode of operation to be 802.1w, and rapid configuration is enabled.
mstp	Specifies the STPD mode of operation to be 802.1s, and rapid configuration is enabled.
cist	Configures the specified STPD as the common instance spanning tree for the MSTP region.
msti	Configures the specified STPD as a multiple spanning tree instance for the MSTP region.
instance	Specifies the Id of the multiple spanning tree instance. The range is 1 to 4,094.

Default

The STPD operates in 802.1D mode.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

If you configure the STP domain in 802.1D mode, the rapid reconfiguration mechanism is disabled.

If you configure the STP domain in 802.1w mode, the rapid reconfiguration mechanism is enabled. You enable or disable RSTP on a per STPD basis only. You do not enable RSTP on a per port basis.

If you configure the STP domain in MSTP mode, the rapid reconfiguration mechanism is enabled. You enable or disable MSTP on a per STPD basis only. You do not enable MSTP on a per port basis. MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

You must first configure a Common and Internal Spanning Tree (CIST) before configuring any multiple spanning tree instances (MSTIs) in the region. You cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command configures STPD `s1` to enable the rapid reconfiguration mechanism and operate in 802.1w mode:

```
configure stpd s1 mode dot1w
```

The following command configures STPD s2 to operate as an MSTI in an MSTP domain:

```
configure stpd s2 mode mstp msti 3
```

configure stpd ports bpdu-restrict

```
configure {stpd} <stpd_name> ports bpdu-restrict [enable | disable] <port_list>
{recovery-timeout {<seconds>}}
```

Description

Configures BPDU Restrict.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
port_list	Specifies one or more ports or slots and ports.
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
seconds	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

The default is disabled.

Usage Guidelines

Before using this command, the port(s) should be configured for edge-safeguard.

Example

The following command enables bpdu-restrict on port 2 of STPD s1:

```
configure stpd s1 ports bpdu-restrict enable 2
```

configure stpd ports cost

```
configure stpd <stpd_name> ports cost [auto | <cost>] <port_list>
```

Description

Specifies the path cost of the port in the specified STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
-----------	---------------------------------------

auto	Specifies the switch to remove any user-defined port cost value(s) and use the appropriate default port cost value(s).
cost	Specifies a numerical port cost value. The range is 1 through 200,000,000.
port_list	Specifies one or more ports or slots and ports.

Default

The switch automatically assigns a default path cost based on the speed of the port, as follows:

- 10 Mbps port—the default cost is 2,000,000
- 100 Mbps port—the default cost is 200,000
- 1000 Mbps port—the default cost is 20,000
- 10000 Mbps ports—the default cost is 2,000

The default port cost for trunked ports is dynamically calculated based on the available bandwidth.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The 802.1D-2004 standard modified the default port path cost value to allow for higher link speeds. If you have a network with both 802.1D-2004 and 802.1D-1998 compliant bridges, a higher link speed can create a situation whereby an 802.1D-1998 compliant bridge could become the most favorable transit path and possibly cause the traffic to span more bridges. To prevent this situation, configure the port path cost to make links with the same speed use the same path host value. For example, if you have 100 Mbps links on all bridges, configure the port path cost for the 802.1D-2004 compliant bridges to 19 instead of using the default 200,000.

Note: You cannot configure the port path cost on 802.1D-1998 compliant bridges to 200,000 because the path cost range setting is 1 to 65,535.

The range for the `cost` parameter is 1 through 200,000,000. If you configure the port cost, a setting of 1 indicates the highest priority.

If you configured a port cost value and specify the `auto` option, the switch removes the user-defined port cost value and returns to the default, automatically assigned, port cost value.

The `auto` port cost of a trunk port is calculated based on number member ports in the trunk port. Link up and down of the member port does not affect the trunk port cost, thus it does not trigger topology change. Only adding or removing a member port to/from the trunk port causes `auto` trunk port cost to change. Also, by so configuring a static trunk port cost, the value is frozen regardless of the number of member ports in the trunk port.

Example

The following command configures a cost of 100 to slot 2, ports 1 through 5 in STPD `s0`:

```
configure stpd s0 ports cost 100 2:1-2:5
```

configure stpd ports edge-safeguard disable

```
configure {stpd} <stpd_name> ports edge-safeguard disable <port_list> {bpdu-restrict}
{recovery-timeout {<seconds>}}
```

Description

Disables the edge safeguard loop prevention on the specified RSTP or MSTP edge port.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>port_list</code>	Specifies one or more edge ports.
<code>bpdu-restrict</code>	Disables port as soon as a BPDU is received.
<code>recovery-timeout</code>	Time after which the port will be re-enabled.
<code>seconds</code>	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

By default, this feature is disabled.

Usage Guidelines

This command applies only to ports that have already been configured as edge ports.

Loop prevention and detection on an edge port configured for RSTP or MSTP is called *edge safeguard*. An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs.

If you disable this feature, the edge port enters the forwarding state but no longer transmits BPDUs unless a BPDU is received by that edge port. This is the default behavior.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDU restrict can be disabled using the `configure stpd <stpd_name> ports bpdu-restrict disable <port_list>` command.

If edge safeguard is disabled, BPDU restrict is also disabled.

To view the status of the edge safeguard feature use the `show {stpd} <stpd_name> ports {[detail | <port_list> {detail}]}` command. You can also use the `show stpd {<stpd_name> | detail}` command to display the STPD configuration on the switch, including the enable/disable state for edge safeguard.

Note: In MSTP, configuring edge safeguard at CIST will be inherited in all MSTI.

To enable or re-enable edge safeguard, use one of the following commands:

- `configure {stpd} <stpd_name> ports edge-safeguard enable <port_list> {bpdu-restrict} {recovery-timeout {<seconds>}}`
- `configure stpd <stpd_name> ports link-type [[auto | broadcast | point-to-point] <port_list> | edge <port_list> {edge-safeguard [enable | disable] {bpdu-restrict} {recovery-timeout <seconds>}}]`

Example

The following command disables edge safeguard on RSTP edge port 4 in STPD *s1* on a stand-alone switch:

```
configure stpd s1 ports edge-safeguard disable 4
```

The following command disables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD *s1* on the switch:

```
configure stpd s1 ports edge-safeguard disable 2:3
```

configure stpd ports edge-safeguard enable

```
configure {stpd} <stpd_name> ports edge-safeguard enable <port_list> {bpdu-restrict} {recovery-timeout {<seconds>}}
```

Description

Enables the edge safeguard loop prevention on the specified RSTP or MSTP edge port.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
port_list	Specifies one or more edge ports.
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
seconds	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

By default, this feature is disabled.

Usage Guidelines

This command applies only to ports that have already been configured as edge ports.

Loop prevention and detection on an edge port configured for RSTP or MSTP is called *edge safeguard*. You configure edge safeguard on RSTP or MSTP edge ports to prevent accidental or deliberate misconfigurations (loops) resulting from connecting two edge ports together or by connecting a hub or other non-STP switch to an edge port. Edge safeguard also limits the impact of broadcast storms that might occur on edge ports.

An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs. This advanced loop prevention mechanism improves network resiliency but does not interfere with the rapid convergence of edge ports.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDUs can be disabled using the `configure {stpd} <stpd_name> ports bpdu-restrict [enable | disable] <port_list> {recovery-timeout {<seconds>}}` command and selecting `disable`.

If edge safeguard is disabled, BPDU restrict is also disabled.

To view the status of the edge safeguard feature use the `show {stpd} <stpd_name> ports {[detail | <port_list> {detail}]}` command. You can also use the `show stpd {<stpd_name> | detail}` command to display the STPD configuration on the switch, including the enable/disable state for edge safeguard.

Note: In MSTP, configuring edge safeguard at CIST will be inherited in all MSTI.

To disable edge safeguard, use one of the following commands:

- `configure {stpd} <stpd_name> ports edge-safeguard disable <port_list> {bpdu-restrict} {recovery-timeout {<seconds>}}`
- `configure stpd <stpd_name> ports link-type [[auto | broadcast | point-to-point] <port_list> | edge <port_list> {edge-safeguard [enable | disable] {bpdu-restrict} {recovery-timeout <seconds>}}]`

Example

The following command enables edge safeguard on RSTP edge port 4 in STPD `s1` on a stand-alone switch:

```
configure stpd s1 ports edge-safeguard enable 4
```


The following command enables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD *s1* on the switch:

```
configure stpd s1 ports edge-safeguard enable 2:3
```

configure stpd ports link-type

```
configure stpd <stpd_name> ports link-type [[auto | broadcast | point-to-point] <port_list> |  
edge <port_list> {edge-safeguard [enable | disable] {bpdu-restrict} {recovery-timeout  
<seconds>}}]
```

Description

Configures the ports in the specified STPD as auto, broadcast, edge, or point-to-point link types.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
auto	Specifies the switch to automatically determine the port link type. An auto link behaves like a point-to-point link if the link is in full-duplex mode or if link aggregation is enabled on the port. Used for 802.1w configurations.
broadcast	Specifies a port attached to a LAN segment with more than two bridges. Used for 802.1D configurations. A port with broadcast link type cannot participate in rapid reconfiguration using RSTP or MSTP. By default, all STP.1D ports are broadcast links.
point-to-point	Specifies a port attached to a LAN segment with only two bridges. A port with point-to-point link type can participate in rapid reconfiguration. Used for 802.1w and MSTP configurations. By default, all 802.1w and MSTP ports are point-to-point link types.
port_list	Specifies one or more ports or slots and ports.
edge	Specifies a port that does not have a bridge attached. An edge port is placed and held in the STP forwarding state unless a BPDU is received by the port. Used for 802.1w and MSTP configurations.
edge-safeguard	Specifies that the edge port be configured with edge safeguard, a loop prevention and detection mechanism. Used for 802.1w and MSTP configurations
enable	Specifies that edge safeguard be enabled on the edge port(s).
disable	Specifies that edge safeguard be disabled on the edge port(s).
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
seconds	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

STP.1D ports are broadcast link types
802.1w and MSTP ports are point-to-point link types

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

The default, broadcast links, supports legacy STP (802.1D) configurations. If the switch operates in 802.1D mode, any configured port link type will behave the same as the broadcast link type.

RSTP rapidly moves the designated ports of a point-to-point link type into the forwarding state. This behavior is supported by RSTP and MSTP only.

In an MSTP environment, configure the same link types for the CIST and all MSTIs.

Auto Link Type

An auto link behaves like a point-to-point link if the link is in full duplex mode or if link aggregation is enabled on the port; otherwise, an auto link behaves like a broadcast link. If a non-STP switch exists between several switches operating in 802.1w mode with auto links, the non-STP switch may negotiate full-duplex even though the broadcast domain extends over several STP devices.

Edge Link Type

RSTP does not send any BPDUs from an edge port nor does it generate topology change events when an edge port changes its state.

If you configure a port to be an edge port, the port immediately enters the forwarding state. Edge ports remain in the forwarding state unless the port receives a BPDU. In that case, edge ports enter the blocking state. The edge port remains in the blocking state until it stops receiving BPDUs and the message age timer expires.

Edge Safeguard

Loop prevention and detection on an edge port configured for RSTP or MSTP is called *edge safeguard*. You configure edge safeguard on RSTP or MSTP edge ports to prevent accidental or deliberate misconfigurations (loops) resulting from connecting two edge ports together or by connecting a hub or other non-STP switch to an edge port. Edge safeguard also limits the impact of broadcast storms that might occur on edge ports.

An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs. This advanced loop prevention mechanism improves network resiliency but does not interfere with the rapid convergence of edge ports.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDU restrict can be disabled using the `configure stpd <stpd_name> ports bpdu-restrict disable <port_list>` command.

If edge safeguard is disabled, BPDU restrict is also disabled.

To configure a port as an edge port and enable edge safeguard on that port, use the `configure stpd <stpd_name> ports link-type edge <port_list> edge-safeguard` command and specify `enable`.

To disable edge safeguard on the edge port, use the `configure stpd <stpd_name> ports link-type edge <port_list> edge-safeguard` command and specify `disable`.

Two other commands are also available to enable and disable edge safeguard:

```
configure stpd ports edge-safeguard enable
configure stpd ports edge-safeguard disable
```

In MSTP, configuring edge safeguard at CIST will be inherited in all MSTI.

Example

The following command configures slot 2, ports 1 through 4 to be point-to-point links in STPD s1:

```
configure stpd s1 ports link-type point-to-point 2:1-2:4
```

The following command enables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD s1 configured for RSTP:

```
configure stpd s1 ports link-type edge 2:3 edge-safeguard enable
```

configure stpd ports mode

```
configure stpd <stpd_name> ports mode [dot1d | emistp | pvst-plus] <port_list>
```

Description

Configures the encapsulation mode for the specified port list.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>dot1d</code>	Specifies the STP encapsulation mode of operation to be 802.1d.
<code>emistp</code>	Specifies the STP encapsulation mode of operation to be EMISTP.
<code>pvst-plus</code>	Specifies the STP encapsulation mode of operation to be PVST+.
<code>port_list</code>	Specifies one or more ports or slots and ports.

Default

Ports in the default STPD (s0) are `dot1d` mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

Example

The following command configures STPD `s1` with PVST+ packet formatting for slot 2, port 1:

```
configure stpd s1 ports mode pvst-plus 2:1
```

configure stpd ports port-priority

```
configure stpd <stpd_name> ports port-priority <priority> <port_list>
```

Description

Specifies the port priority of the port in the specified STPD.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>priority</code>	Specifies a numerical port priority value. The range is 0 through 240 and is subject to the multiple of 16 restriction.
<code>port_list</code>	Specifies one or more ports or slots and ports.

Default

The default is 128.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the priority of the port, you can make it more or less likely to become the root port or a designated port.

A setting of 0 indicates the highest priority.

The range for the `priority` parameter is 0 through 240 and is subject to the multiple of 16 restriction.

Example

The following command assigns a priority of 32 to slot 2, ports 1 through 5 in STPD `s0`:

```
configure stpd s0 ports port-priority 32 2:1-2:5
```

configure stpd ports priority

```
configure stpd <stpd_name> ports priority <priority> <port_list>
```

Description

Specifies the port priority of the port in the specified STPD.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>priority</code>	Specifies a numerical port priority value. The range is 0 through 31 for STP and 0 through 15 for MSTP and RSTP.
<code>port_list</code>	Specifies one or more ports or slots and ports.

Default

The default is 128.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the priority of the port, you can make it more or less likely to become the root port or a designated port.

A setting of 0 indicates the highest priority.

The range for the `priority` parameter is 0 through 31 for STP and 0 through 15 for MSTP and RSTP.

NETGEAR 8800 introduces support for a new ports priority command: `configure stpd ports port-priority`. The priority range of this command is 0 through 240 and is subject to the multiple of 16 restriction. For more information see [configure stpd ports port-priority](#) on page 764.

Example

The following command assigns a priority of 1 to slot 2, ports 1 through 5 in STPD s0:

```
configure stpd s0 ports priority 1 2:1-2:5
```

configure stpd ports restricted-role disable

```
configure stpd <stpd_name> ports restricted-role disable <port_list>
```

Description

Disables *restricted role* on the specified port inside the core network.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>port_list</code>	Specifies one or more ports or slots and ports.

Default

NA

Usage Guidelines

The restricted role is disabled by default. If set, it can cause a lack of spanning tree connectivity. A network administrator enables the restricted role to prevent bridges external to a core region of the network from influencing the spanning tree active topology, possibly because those bridges are not under the full control of the administrator.

Note: Disabling Restricted Role at CIST is inherited by all MSTI.

Example

The following command disables restricted role for `s1` on port 6:3.

```
configure stpd s1 ports restricted-role disable 6:3
```

configure stpd ports restricted-role enable

```
configure stpd <stpd_name> ports restricted-role enable <port_list>
```

Description

Enables *restricted role* on the specified port inside the core network.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>port_list</code>	Specifies one or more ports or slots and ports.

Default

NA

Usage Guidelines

Enabling restricted role causes the port not to be selected as a root port even if it has the best spanning tree priority vector. Such a port is selected as an alternate port after the root port has been selected.

The restricted role is disabled by default. If set, it can cause a lack of spanning tree connectivity. A network administrator enables the restricted role to prevent bridges external to a core region of the network from influencing the spanning tree active topology, possibly because those bridges are not under the full control of the administrator.

Note: Restricted role should not be enabled with edge mode.

Note: Enabling Restricted Role at CIST is inherited by all MSTI.

Example

The following command enables restricted role on port 6:3.

```
configure stpd s1 ports restricted-role enable 6:3
```

configure stpd priority

```
configure stpd <stpd_name> priority <priority>
```

Description

Specifies the bridge priority of the STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
priority	Specifies the bridge priority of the STPD. The range is 0 through 61,440 and is subject to the multiple of 4,096 restriction.

Default

The default priority is 32,768.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the bridge priority of the STPD, you can make it more or less likely to become the root bridge.

The range for the `<priority>` parameter is 0 through 61,440 and is subject to the multiple of 4,096 restriction. A setting of 0 indicates the highest priority.

For example, to lower the numerical value of the priority (which gives the priority a higher precedence), you subtract 4,096 from the default priority: $32,768 - 4,096 = 28,672$. If you modify the priority by a value other than 4,096, the switch rejects the entry.

Example

The following command sets the bridge priority of *STPD1* to 16,384:

```
configure stpd stpd1 priority 16384
```

configure stpd tag

```
configure stpd <stpd_name> tag <stpd_tag>
```

Description

Assigns an StpdID to an STPD.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>stpd_tag</code>	Specifies the VLAN ID of the carrier VLAN that is owned by the STPD.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

An STPD ID is used to identify each STP domain. You assign the StpdID when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STP domain, and that VLAN cannot belong to another STPD. Unless all ports are running in 802.1D mode, an STPD with ports running in either EMISTP mode or PVST+ mode must be configured with an STPD ID.

You must create and configure the VLAN, along with the tag, before you can configure the STPD tag. To create a VLAN, use the `create vlan` command. To configure the VLAN, use the `configure vlan` commands.

MSTP Only

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the CIST. The switch assigns this ID automatically when you configure the CIST STPD. To configure the CIST STPD, use the `configure stpd <stpd_name> mode [dot1d | dot1w | mstp [cist | msti <instance>]]` command.

An MSTI identifier (MSTI ID) identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that participates in the MSTP region. Each STPD that participates in a particular MSTP region must have the same MSTI ID. To configure the MSTI ID, use the `configure stpd <stpd_name> mode [dot1d | dot1w | mstp [cist | msti <instance>]]` command.

Example

The following command assigns an StpdID to the `purple_st` STPD:

```
configure stpd purple_st tag 200
```

configure vlan add ports stpd

```
configure vlan <vlan_name> add ports [all | <port_list>] {tagged | untagged} stpd <stpd_name>
{[dot1d | emistp | pvst-plus]}
```

Description

Adds one or more ports in a VLAN to a specified STPD.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
<code>all</code>	Specifies all of the ports to be included in the STPD.
<code>port_list</code>	Specifies the port or ports to be included in the STPD.
<code>tagged</code>	Specifies the ports should be configured as tagged.
<code>untagged</code>	Specifies the ports should be configured as untagged.
<code>stpd_name</code>	Specifies an STPD name on the switch.
<code>dot1d</code>	Specifies the STP encapsulation mode of operation to be 802.1d.
<code>emistp</code>	Specifies the STP encapsulation mode of operation to be EMISTP.
<code>pvst-plus</code>	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are in `dot1d` mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

To create a VLAN, use the `create vlan` command. To create an STP domain, use the `create stpd` command.

In an EMISTP or PVST+ environment, this command adds a list of ports to a VLAN and a specified STPD at the same time provided the carrier VLAN already exists on the same set of ports. You can also specify the encapsulation mode for those ports.

In an MSTP environment, you do not need a carrier VLAN. A CIST controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate region status. You must use the `dot1d` encapsulation mode in an MSTP environment.

You cannot configure STP on the following ports:

- Mirroring target ports
- Software-controlled redundant ports

If you see an error similar to the following:

```
Error: Cannot add VLAN default port 3:5 to STP domain
```

You might be attempting to add:

- A carrier VLAN port to a different STP domain than the carrier VLAN belongs

- A VLAN/port for which the carrier VLAN does not yet belong

Note: This restriction is only enforced in an active STP domain and when you enable STP to ensure you have a legal STP configuration.

Naming Conventions

If your VLAN has the same name as another component, for example an STPD, NETGEAR recommends that you specify the identifying keyword as well as the name. If your VLAN has a name unique only to that VLAN, the keywords `vlan` and `stpd` are optional.

STP Encapsulation Modes

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only *one* STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

MSTP STPDs use only 802.1D BPDU encapsulation mode. The switch prevents you from configuring EMISTP or PVST+ encapsulation mode for MSTP STPDs.

Automatically Inheriting Ports—MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Example

The following command adds slot 1, port 2 and slot 2, port 3, members of a VLAN named *Marketing*, to the STPD named *STPD1*, and specifies that they be in *EMISTP* mode:

```
configure vlan marketing add ports 1:2, 2:3 tagged stpd stpd1 emistp
```

create stpd

```
create stpd <stpd_name>
```

Description

Creates a user-defined STPD.

Syntax Description

stpd_name	Specifies a user-defined STPD name to be created. May be up to 32 characters in length.
-----------	---

Default

The default device configuration contains a single STPD called *s0*.

When an STPD is created, the STPD has the following default parameters:

- State—disabled
- StpdID—none
- Assigned VLANs—none
- Bridge priority—32,768
- Maximum BPDU age—20 seconds
- Hello time—2 seconds
- Forward delay—15 seconds
- Operational mode—802.1D
- Rapid Root Failover—disabled
- Default Binding Mode (encapsulation mode)—Ports in the default STPD (*s0*) are in *802.1d* mode. Ports in user-created STPDs are in *emistp* mode.
- Maximum hop count (when configured for MSTP)—20 hops

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (`_`) but cannot be any reserved keywords, for example, *stp* or *stpd*. Names must start with an alphabetical character, for example, *a*, *Z*. For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

Each STPD name must be unique and cannot duplicate any other named STPDs on the switch. If you are uncertain about the STPD names on the switch, use the `show stpd` command to view the STPD names.

You can, however, re-use names across multiple categories of switch configuration. For example, you can use the name *Test* for an STPD and a VLAN. If you use the same name, NETGEAR recommends that you specify the appropriate keyword when configuring the STPD. If you do not specify the appropriate keyword, the switch displays a message similar to the following:

```
%% Ambiguous command: "configure Test"
```

To view the names of the STPDs on the switch, enter `configure` and press TAB. Scroll to the end of the output to view the names.

Each STPD has its own Root Bridge and active path. After the STPD is created, one or more VLANs can be assigned to it.

Example

The following example creates an STPD named *purple_st*:

```
create stpd purple_st
```

delete stpd

```
delete stpd <stpd_name>
```

Description

Removes a user-defined STPD from the switch.

Syntax Description

<code>stpd_name</code>	Specifies a user-defined STPD name on the switch.
------------------------	---

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If you do not specify the `stpd` keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "delete Test"
```

In this example, to delete the STPD *Test*, enter `delete stpd Test`.

If you created an STPD with a name unique only to that STPD, the keyword `stpd` is optional.

The default STPD, *s0*, cannot be deleted.

In an MSTP environment, you cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command deletes an STPD named *purple_st*:

```
delete stpd purple_st
```

disable stpd

```
disable stpd {<stpd_name>}
```

Description

Disables the STP protocol on a particular STPD or for all STPDs.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
-----------	---------------------------------------

Default

Disabled.

Usage Guidelines

After you have created the STPD with a unique name, the keyword `stpd` is optional.

If you want to disable the STP protocol for all STPDs, do not specify an STPD name.

In an MSTP environment, you cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command disables an STPD named *purple_st*:

```
disable stpd purple_st
```

The following command disables the STP protocol for all STPDs on the switch:

```
disable stpd
```

disable stpd auto-bind

```
disable stpd <stpd_name> auto-bind vlan <vlan_name>
```

Description

Disables the ability to automatically add ports to an STPD when they are added to a member VLAN.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
vlan_name	Specifies the name of a member VLAN with autobind enabled.

Default

The autobind feature is disabled on user-created STPDs. The autobind feature is enabled on the default VLAN that participates in the default STPD S0.

Usage Guidelines

Note: Ports already in the STPD remain in that domain (as if they were added manually).

If you create an STPD and a VLAN with unique names, the keywords `stpd` and `vlan` are optional.

Ports added to the STPD automatically when autobind is enabled are not removed when autobind is disabled. The ports are present after a switch reboot.

To view STP configuration status of the ports in a VLAN, use the following command:

```
show vlan <vlan_name> stpd
```

Example

The following example disables autobind on an STPD named s8:

```
disable stpd s8 auto-bind v5
```

disable stpd ports

```
disable stpd <stpd_name> ports [all | <port_list>]
```

Description

Disables STP on one or more ports for a given STPD.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
all	Specifies all ports for a given STPD.
port_list	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

If you create the STPD with a unique name, the keyword `stpd` is optional.

Disabling STP on one or more ports puts those ports in the *forwarding* state; all BPDUs received on those ports are disregarded and dropped.

Use the `all` keyword to specify that all ports of a given STPD are disabled.

Use the `port_list` parameter to specify a list of ports of a given STPD are disabled.

If you do not use the default STPD, you must create one or more STPDs and configure and enable the STPD before you can use the `disable stpd ports` command.

Example

The following command disables slot 2, port 4 on an STPD named *Backbone_st*:

```
disable stpd backbone_st ports 2:4
```

disable stpd rapid-root-failover

```
disable stpd <stpd_name> rapid-root-failover
```

Description

Disables rapid root failover for STP recovery times.

Syntax Description

<code>stpd_name</code>	Specifies an STPD name on the switch.
------------------------	---------------------------------------

Default

Disabled.

Usage Guidelines

This command is applicable for STPDs operating in 802.1D.

After you have created the STPD with a unique name, the keyword `stpd` is optional.

To view the status of rapid root failover on the switch, use the `show stpd` command. The `show stpd` command displays information about the STPD configuration on the switch including the enable/disable state for rapid root failover.

Example

The following command disables rapid root fail over on STPD *Backbone_st*:

```
disable stpd backbone_st rapid-root-failover
```

enable stpd

```
enable stpd {<stpd_name>}
```

Description

Enables the STP protocol for one or all STPDs.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
-----------	---------------------------------------

Default

Disabled.

Usage Guidelines

If you want to enable the STP protocol for all STPDs, do not specify an STPD name.

Example

The following command enables an STPD named *Backbone_st*:

```
enable stpd backbone_st
```

enable stpd auto-bind

```
enable stpd <stpd_name> auto-bind vlan <vlan_name>
```

Description

Automatically adds ports to an STPD when ports are added to a member VLAN.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
vlan_name	Specifies the name of the VLAN to have autobind enabled.

Default

The autobind feature is disabled on user-created STPDs. The autobind feature is enabled on the default VLAN that participates in the default STPD S0.

If you enable autobind and add ports to a member VLAN, those ports are automatically added to the STPD.

Usage Guidelines

If you create an STPD and a VLAN with unique names, the keywords `stpd` and `vlan` are optional.

You cannot configure the autobind feature on a network login VLAN.

In an EMISTP or PVST+ environment, when you issue this command, any port or list of ports that you add to the carrier VLAN are automatically added to the STPD with autobind enabled. In addition, any port or list of ports that you remove from a carrier VLAN are automatically removed from the STPD. This allows the STPD to increase or decrease its span as you add ports to or remove ports from a carrier VLAN.

For MSTP, when you issue this command, any port or list of ports that gets automatically added to an MSTI are automatically inherited by the CIST. In addition, any port or list of ports that you remove from an MSTI protected VLAN are automatically removed from the CIST. For more information see the section. For more information, see [Automatically Inheriting Ports—MSTP Only](#) on page 779.

Carrier VLAN

A carrier VLAN defines the scope of the STPD, which includes the physical and logical ports that belong to the STPD and the 802.1Q tag used to transport STP BPDUs in the encapsulation mode is EMISTP or PVST+. Only one carrier VLAN can exist in a given STPD, although some of its ports can be outside the control of any STPD at the same time.

Note: The STPD ID must be identical to the VLAN ID of the carrier VLAN in that STPD.

If you configure MSTP, you do not need a carrier VLAN. With MSTP, you configure a CIST that controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate the status of MSTP regions. All VLANs participating in the MSTP region have the same privileges.

Protected VLAN

Protected VLANs are all other VLANs that are members of the STPD. These VLANs “piggyback” on the carrier VLAN. Protected VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes and inherit the state of the carrier VLAN. Protected VLANs can participate in multiple STPDs, but any particular port in the VLAN can belong to only one STPD.

Enabling autobind on a protected VLAN does not expand the boundary of the STPD. However, the VLAN and port combinations are added to or removed from the STPD subject to the boundaries of the carrier VLAN.

If you configure MSTP, all member VLANs in an MSTP region are protected VLANs. These VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes communicated by the CIST to the MSTP regions. MSTIs cannot share the same protected VLAN; however, any port in a protected VLAN can belong to multiple MSTIs.

Automatically Inheriting Ports—MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Displaying STP Information

To view STP configuration status of the ports on a VLAN, use the following command:

```
show vlan <vlan_name> stpd
```

Example

The examples in this section assume that you have already removed the ports from the *Default* VLAN.

To automatically add ports to an STPD running 802.1D, EMISTP, or PVST+ and to expand the boundary of the STPD, you must complete the following tasks:

- Create the carrier VLAN.
- Assign a VLAN ID to the carrier VLAN.
- Add ports to the carrier VLAN.
- Create an STPD (or use the default, *S0*).
- Enable autobind on the STPDs carrier VLAN.
- Configure the STPD tag (the STPD ID must be identical to the VLAN ID of the carrier VLAN in the STP domain).
- Enable STP.

The following example enables autobind on an STPD named *s8* after creating a carrier VLAN named *v5*:

```
create vlan v5
configure vlan v5 tag 100
configure vlan v5 add ports 1:1-1:20 tagged
create stpd s8
enable stpd s8 auto-bind v5
configure stpd s8 tag 100
enable stpd s8
```

To automatically add ports to the CIST STPD and to expand the boundary of the STPD, you must complete the following tasks:

- Create a VLAN or use the *Default* VLAN. (In this example, the *Default* VLAN is used.)

- Create the MSTP region.
- Create the STPD to be used as the CIST, and configure the mode of operation for the STPD.
- Specify the priority for the CIST.
- Enable the CIST.

The following example enables autobind on the VLAN *Default* for the CIST STPD named s1:

```
configure mstp region 1
create stpd s1
configure stpd s1 mode mstp cist
configure stpd s1 priority 32768
enable stpd s1
```

The following example enables autobind on the VLAN *math* for the MSTI STPD named s2:

```
create vlan math
configure vlan math tag 2
configure vlan math add ports 2-3
configure mstp region 1
create stpd s2
configure stpd s2 mode mstp msti 1
configure stpd s2 priority 32768
enable stpd s2 auto-bind vlan math
configure stpd s2 ports link-type point-to-point 5-6
enable stpd s2
```

enable stpd ports

```
enable stpd <stpd_name> ports [all | <port_list>]
```

Description

Enables the STP protocol on one or more ports.

Syntax Description

stpd_name	Specifies an STPD on the switch.
all	Specifies all ports for a given STPD.
port_list	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

If you create an STPD with a unique name, the keyword `stpd` is optional.

If STP is enabled for a port, BPDUs are generated and processed on that port if STP is enabled for the associated STPD.

You must configure one or more STPDs before you can use the `enable stpd ports` command. To create an STPD, use the `create stpd <stpd_name>` command. If you have considerable knowledge and experience with STP, you can configure the STPD using the `configure stpd` commands. However, the default STP parameters are adequate for most networks.

Example

The following command enables slot 2, port 4 on an STPD named *Backbone_st*:

```
enable stpd backbone_st ports 2:4
```

enable stpd rapid-root-failover

```
enable stpd <stpd_name> rapid-root-failover
```

Description

Enables rapid root failover for faster STP recovery times.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
-----------	---------------------------------------

Default

Disabled.

Usage Guidelines

This command is applicable for STPDs operating in 802.1D.

If you create an STPD with a unique name, the keyword `stpd` is optional.

To view the status of rapid root failover on the switch, use the `show stpd` command. The `show stpd` command displays information about the STPD configuration on the switch including the enable/disable state for rapid root failover.

Example

The following command enables rapid root fail over on STPD *Backbone_st*:

```
enable stpd backbone_st rapid-root-failover
```

show stpd

```
show stpd {<stpd_name> | detail}
```

Description

Displays STPD settings on the switch.

Syntax Description

<code>stpd_name</code>	Specifies an STPD on the switch.
<code>detail</code>	Specifies that STPD settings should be shown for each STPD.

Default

N/A.

Usage Guidelines

If you specify the command without any options, the following STPD information appears:

- Name—The name of the STPD.
- Tag—The StpdID of the domain, if configured.
- Flags—The following flags communicate information about the current state of the STPD:
 - (C) Topology Change—A network topology change has occurred in the network.
 - (D) Disable—The STPD is disabled.
 - (E) Enable—The STPD is enabled.
 - (R) Rapid Root Failover—The STPD has been configured for rapid root failover
 - (T) Topology Change Detected—The STPD has detected a change in the network topology.
 - (M) MSTP CIST—The STPD has been configured for MSTP, and the STPD is the common and internal spanning tree.
 - (I) MSTP MSTI—The STPD has been configured for MSTP, and the STPD is a multiple instance spanning tree.
- Ports—The number of ports that are part of the STPD.
- Bridge ID—The MAC addresses of the switch.
- Designated Root—The MAC address of the switch that is the designated root bridge.
- Rt Port—The root port.
- Rt Cost—The path cost to the root port.
- Total Number of STPDs—The total number of STPDs configured on the switch.

If you have an MSTP region and associated spanning trees configured on the switch, the command also displays the following global MSTP information:

- MSTP Region—The name of the MSTP region configured on the switch.
- Format Identifier—The number used by BPDUs to communicate within an MSTP region.
- Revision Level—This number is reserved for future use.

- Common and Internal Spanning Tree (CIST)—The name of the CIST that controls the connectivity of interconnecting MSTP regions.
- Total number of MST Instances (MSTI)—The number of MSTIs running in the MSTP region.

If you use the `show stpd` command and specify the name of an STPD, in addition to the data previously described, the command displays more detailed information about the STPD. If you specify the `detail` option, the switch displays the same type of information for all of the STPDs configured on the switch.

The additional output includes the following:

- STPD mode of operation
- Autobind mode
- Active VLANs
- Timer information
- Topology change information

If you have MSTP configured, the command also displays the following information:

- Bridge role
- CIST root
- CIST regional root
- MSTI instances
- Master port (Displayed only on MSTI STPDs)

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If you do not specify the `stpd` keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "show Test"
```

In this example, to view the settings of the STPD *Test*, enter `show stpd Test`.

If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

Example

The following command displays the STPD settings on a switch that has MSTP configured:

```
show stpd
```

The following is sample output from this command:

```
MSTP Global Configuration:
MSTP Region Name           : 00304841ed97
MSTP format Identifier     : 0
MSTP Revision Level        : 3
Common and Internal Spanning Tree (CIST) : ----
Total Number of MST Instances (MSTI)   : 0
```

```

Name Tag  Flags  Ports Bridge ID          Designated Root  Rt Port Rt Cost
s0      0000D----- 08000001030f99dc0 0000000000000000 ----- 0
Total number of STPDs: 1
Flags: (C) Topology Change, (D) Disable, (E) Enable, (R) Rapid Root Failover (T)
Topology Change Detected, (M) MSTP CIST, (I) MSTP MSTI

```

The following command displays STPD settings on an STPD named *Backbone_st*:

```
show stpd backbone_st
```

The following is sample output from this command:

```

Stpd: backbone_st Stp: ENABLED          Number of Ports: 51
Rapid Root Failover: Disabled
Operational Mode: 802.1W                Default Binding Mode: 802.1D
802.1Q Tag: (none)
Ports: 1:1,1:2,2:1,2:2,3:1,3:2,4:1,4:2,5:1,5:2,
       5:3,5:4,5:5,5:6,5:7,5:8,5:9,5:10,5:11,5:12,
       5:13,5:14,5:15,5:16,5:17,5:18,5:19,5:20,5:21,5:22,
       5:23,5:24,5:25,5:26,5:27,5:28,5:29,5:30,5:31,5:32,
       5:33,5:34,5:35,5:36,5:37,5:38,5:39,5:40,5:41,5:42,
       5:43
Participating Vlans: Default
Auto-bind Vlans: Default
Bridge Priority: 5000
BridgeID:          13:88:00:01:30:f4:06:80
Designated root:   0a:be:00:01:30:28:b7:00
RootPathCost: 19   Root Port: 28
MaxAge: 20s        HelloTime: 2s          ForwardDelay: 15s
CfgBrMaxAge: 20s   CfgBrHelloTime: 2s        CfgBrForwardDelay: 15s
Topology Change Time: 35s          Hold time: 1s
Topology Change Detected: FALSE    Topology Change: FALSE
Number of Topology Changes: 7
Time Since Last Topology Change: 4967s

```

The following is sample output for an STPD configured as the CIST (the output is similar for an STPD configured as an MSTI):

```

Stpd: s0              Stp: DISABLED          Number of Ports: 0          Rapid Root
Failover: Disabled    Operational Mode:
MSTP                  Default Binding Mode: 802.1d
MSTP Instance :      CIST                    CIST : s0
802.1Q Tag:          (none)
Ports:               (none)
Participating Vlan Count: 1
Auto-bind Vlans Count: 1

Bridge Priority:      32768                    BridgeID:
                    80:00:00:10:30:f9:9d:c0Bridge
Role :               CIST Regional Root
CIST Root            80:00:00:10:30:f9:9d:c0CIST
Regional Root:       80:00:00:10:30:f9:9d:c0

```



```

Designated root:          00:00:00:00:00:00:00:00          RootPathCost:
0          External RootPathCost: 0          Root Port:      ----
MaxAge:          0s          HelloTime:          0s          ForwardDelay:          0s
CfgBrMaxAge:    20s          CfgBrHelloTime:      2s          CfgBrForwardDelay: 15s MaxHopCount:
20          CfgBrMaxHopCount :          20

```

```

Topology Change Time:      35s          Hold time:          1s
Topology Change Detected:  FALSE          Topology Change:      FALSE
Number of Topology Changes: 0          Time Since
Last Topology Change: 0s

```

```

Participating Vlans      :          (none)
Auto-bind Vlans          :          Default

```

show stpd ports

```
show {stpd} <stpd_name> ports {[detail | <port_list> {detail}]}
```

Description

Displays the STP state of a port.

Syntax Description

stpd_name	Specifies an STPD name.
port_list	Specifies one or more ports or slots and ports.
detail	Specifies more detailed information about one or more ports of the STPD.

Default

N/A.

Usage Guidelines

This command displays the following:

- STPD port configuration
- STPD port encapsulation mode
- STPD path cost
- STPD priority
- STPD state (root bridge, and so on)
- Port role (root designated, alternate and so on)
- STPD port state (forwarding, blocking, and so on)
- Configured port link type
- Operational port link type
- Edge port settings (inconsistent behavior, edge safeguard setting)

- Restricted role (enabled, disabled)
- MSTP port role (internal or boundary)

To display more detailed information for one or more ports in the specified STPD, including participating VLANs, specify the `detail` option.

If you have MSTP configured and specify the `detail` option, this command displays additional information:

- MSTP internal path cost
- MSTP timers

If your STPD has the same name as another component, for example a VLAN, NETGEAR recommends that you specify the identifying keyword as well as the name. If you do not specify the `stpd` keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "show Test ports"
```

In this example, to view all of the port settings of STPD *Test*, enter `show stpd Test ports`.

If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

Example

The following command displays the state of slot 3, ports 1 through 3 on an STPD named *s0*:

```
show stpd S0 ports 3:1-3:3
```

The following is sample output from this command:

Port	Mode	State	Cost	Flags	Priority	Port ID	Designated Bridge
3:1	802.1D	FORWARDING	100	e-----	16	16641	00:00:00:00:00:00:00:00
3:1	802.1D	BLOCKING	20000	eApp-w---r	128	8001	10:00:00:04:96:26:5f:47
3:2	802.1D	FORWARDING	100	e-----	16	16642	00:00:00:00:00:00:00:00
3:2	802.1D	FORWARDING	20000	eRppaw----	128	8003	10:00:00:04:96:26:5f:47
3:3	802.1D	FORWARDING	100	e-----	16	16643	00:00:00:00:00:00:00:00
3:3	802.1D	FORWARDING	20000	eDpp-w----	128	8004	20:00:00:04:96:26:5c:48

```
Total Ports: 3
```

```
----- Flags: -----
1:                e=Enable, d=Disable
2: (Port role)    R=Root, D=Designated, A=Alternate, B=Backup, M=Master
3: (Config type) b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type)  b=broadcast, p=point-to-point, e=edge
5:                p=proposing, a=agree
6: (partner mode) d = 802.1d, w = 802.1w, m = mstp
7:                i = edgeport inconsistency
8:                S = edgeport safe guard active
                  s = edgeport safe guard configured but inactive
9:                B = Boundary, I = Internal
10:               r = Restricted Role
```

The following command displays the detailed information for the ports in STPD S0:

```
show stpd S0 ports detail
```

The following is sample output from this command:

```
Stpd: s0          Port: 1:1          PortId: 8002    Stp: ENABLED    Path Cost: 19
External Path Cost:
Port Mode: 802.1D
Port State: FORWARDING          Topology Change Ack: FALSE
Port Priority: 16
Designated Root:  00:00:00:00:00:00:00:00    Designated Cost: 0
Designated Bridge: 00:00:00:00:00:00:00:00    Designated Port Id: 0
Partner STP version: Dot1d
Restricted Role: Disabled
Edge Port Safe Guard: Disabled
Participating Vlans: Default
```

```
Stpd: s0          Port: 1:2          PortId: 8002    Stp: ENABLED    Path Cost: 19
External Path Cost:
Port Mode: 802.1D
Port State: FORWARDING          Topology Change Ack: FALSE
Port Priority: 16
Designated Root:  00:00:00:00:00:00:00:00    Designated Cost: 0
Designated Bridge: 00:00:00:00:00:00:00:00    Designated Port Id: 0
Partner STP version: Dot1d
Restricted Role: Enabled
Edge Port Safe Guard: Disabled
Participating Vlans: Default
```

The following command displays the detailed information for the ports in STPD s1 configured for MSTP:

```
show stpd s1 ports detail
```

The following is sample output from this command:

```
Stpd: s1          Port: 1 PortId: 8001    Stp: ENABLED    Path Cost: 4
Port Mode: 802.1D
Port State: FORWARDING          Topology Change Ack: FALSE
Port Priority: 16
Designated Root:  80:00:00:04:96:1f:a8:44    Designated Cost: 0, IntCost: 0
Designated Bridge: 80:00:00:04:96:1f:a8:44    Designated Port Id: 8001
Partner STP version: MSTP
Restricted Role: Disabled
Edge Port Safe Guard: Disabled
maxAge: 20      msgAge: 0      fwdDelay: 15    helloTime: 2    maxHops: 20
Participating Vlans: v1

Stpd: s1          Port: 2 PortId: 8002    Stp: ENABLED    Path Cost: 4
Port Mode: 802.1D
```

```

Port State: BLOCKING           Topology Change Ack: FALSE
Port Priority: 16
Designated Root: 80:00:00:04:96:1f:a8:44      Designated Cost: 0, IntCost: 0
Designated Bridge: 80:00:00:04:96:1f:a8:44      Designated Port Id: 8002
Partner STP version: Dot1d
Restricted Role: Enabled
Edge Port Safe Guard: Disabled
maxAge: 20      msgAge: 0      fwdDelay: 15      helloTime: 2      maxHops: 20
Participating Vlan: v1

```

The following command displays information for port 9 in STPD *s1* configured with a *bpdu-restrict recovery-timeout* of 400:

```
X250e-48p.1 # show s1 ports
```

The following is sample output from this command:

Port	Mode	State	Cost	Flags	Priority	Port ID	Designated Bridge
9	EMISTP	FORWARDING	20000	eDeepw-G--	128	8009	80:00:00:04:96:1f:a8:48

```
Total Ports: 1
```

```

----- Flags: -----
1:          e=Enable, d=Disable
2: (Port role)  R=Root, D=Designated, A=Alternate, B=Backup, M=Master
3: (Config type) b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type)  b=broadcast, p=point-to-point, e=edge
5:          p=proposing, a=agree
6: (partner mode) d = 802.1d, w = 802.1w, m = mstp
7:          i = edgeport inconsistency

8:          S = edgeport safe guard active
          s = edgeport safe guard configured but inactive
          G = edgeport safe guard bpdu restrict active
          g = edgeport safe guard bpdu restrict configured but inactive only dot1w,
mstp
9:          B = Boundary, I = Internal
10:         r = Restricted Role

```

show vlan stpd

```
show vlan <vlan_name> stpd
```

Description

Displays the STP configuration of the ports assigned to a specific VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

If you have a VLAN that spans multiple STPDs, use this command to display the STP configuration of the ports assigned to that specific VLAN.

This command displays the following:

- STPD port configuration
- STPD port mode of operation
- STPD path cost
- STPD priority
- STPD state (root bridge, and so on)
- Port role (root designated, alternate and so on)
- STPD port state (forwarding, blocking, and so on)
- Configured port link type
- Operational port link type

If your VLAN has the same name as another component, for example an STPD, NETGEAR recommends that you specify the identifying keyword as well as the name. If you do not specify the `vlan` keyword, the switch displays an error message similar to the following:

```
%% Ambiguous command: "show Test stpd"
```

In this example, to view the STPD state of VLAN *Test*, enter `show vlan Test stpd`.

If you enter a VLAN name that is not associated with an STPD or does not exist, the switch displays an error message similar to the following:

```
Failed to find vlan 'vlan1' or it has no STP domains configured on it
```

If this happens, check to make sure you typed the correct name of the VLAN and that the VLAN is associated with an STPD.

If your VLAN has a name unique only to that VLAN, the keyword `vlan` is optional.

Example

The following command displays the spanning tree configurations for the VLAN *Default*:

```
show vlan default stpd
```

The following is sample output from this command:

```
s0(enabled) Tag: (none) Ports: 8 Root/P/C: 80:00:00:01:30:94:79:00/-----/0
```

Port	Mode	State	Cost	Flags	Priority	Port ID	Designated Bridge
1:1	802.1D	LEARNING	19	eDbb-d-	16	8001	80:00:00:01:30:94:79:00
1:2	802.1D	DISABLED	4	e-----	16	8002	00:00:00:00:00:00:00:00
1:3	802.1D	DISABLED	4	e-----	16	8003	00:00:00:00:00:00:00:00
1:4	802.1D	LEARNING	4	eDbb-d-	16	8004	80:00:00:01:30:94:79:00
1:5	802.1D	LEARNING	4	eDbb-d-	16	8005	80:00:00:01:30:94:79:00
1:6	802.1D	DISABLED	4	e-----	16	8006	00:00:00:00:00:00:00:00
1:7	802.1D	DISABLED	4	e-----	16	8007	00:00:00:00:00:00:00:00
1:8	802.1D	DISABLED	4	e-----	16	8008	00:00:00:00:00:00:00:00

```

----- Flags: -----
1:                e=Enable, d=Disable
2: (Port role)    R=Root, D=Designated, A=Alternate, B=Backup, M=Master, Y=Boundary
3: (Config type)  b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type)   b=broadcast, p=point-to-point, e=edge
5:                p=proposing, a=agree
6: (partner mode) d=802.1d, w=802.1w, m=mstp
7:                i=edgeport inconsistency
8:                B = Boundary, I = Internal

```

unconfigure mstp region

```
unconfigure mstp region
```

Description

Unconfigures the MSTP region on the switch and returns all MSTP settings to their default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Before you unconfigure an MSTP region, NETGEAR recommends that you disable all active STPDs in the region. This includes the CIST and any active MSTIs.

After you issue this command, all of the MSTP settings return to their default values, as described below:

- **Region Name**—This indicates the name of the MSTP region. In the NETGEAR implementation, the maximum length of the name is 32 characters and can be a combination of alphanumeric characters and underscores (_).

- **Format Selector**—This indicates a number to identify the format of MSTP BPDUs. The default is 0.
- **Revision Level**—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

Example

The following command unconfigures the MSTP region on the switch:

```
unconfigure mstp region
```

unconfigure stpd

```
unconfigure stpd {<stpd_name>}
```

Description

Restores default STP values to a particular STPD or all STPDs.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
-----------	---------------------------------------

Default

N/A.

Usage Guidelines

If you create an STPD with a unique name, the keyword `stpd` is optional.

Use this command to restore default STP values to a particular STPD. If you want to restore default STP values on all STPDs, do not specify a spanning tree name.

Example

The following command restores default values to an STPD named *Backbone_st*.

```
unconfigure stpd backbone_st
```

unconfigure stpd ports link-type

```
unconfigure stpd <stpd_name> ports link-type <port_list>
```

Description

Returns the specified port to the factory default setting of broadcast link.

Syntax Description

stpd_name	Specifies an STPD name on the switch.
port_list	Specifies one or more ports or slots and ports.

Default

All ports are broadcast link types.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, you must enter the `stpd` keyword to specify the STPD. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

If the switch operates in 802.1D mode, any configured port link type will behave the same as the broadcast link type.

In an MSTP environment, configure the same link types for the CIST and all MSTIs.

Example

The following command configures slot 2, ports 1 through 4 to return to the factory default of broadcast links in STPD `s1`:

```
unconfigure stpd s1 ports link-type 2:1-2:4
```


This chapter describes commands for:

- Enabling and disabling Virtual Router Redundancy Protocol (VRRP)
- Performing basic VRRP configuration
- Displaying VRRP information

For an introduction to VRRP, see the *NETGEAR 8800 User Manual*.

clear counters vrrp

```
clear counters vrrp {{vlan <vlan_name>} {vrid <vridval>}}
```

Description

Clears, resets all VRRP statistics and counters.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

Use this command to reset the VRRP statistics on the switch. Statistics are not reset when you disable and re-enable VRRP.

If you do not enter a parameter, statistics for all VRRP VLANs are cleared.

If you specify only VLAN name, statistics for all VRRP VRIDs on that VLAN are cleared.

If you specify VLAN name and VRRP VRID, only statistics for that particular VRID are cleared.

Example

The following command clears the VRRP statistics on VRRP VLAN *v1*:

```
clear counters vrrp vlan v1
```

The following command clears the VRRP statistics for VRID *1* on VRRP VLAN *v1*:

```
clear counters vrrp vlan v1 vrid 1
```

configure vrrp vlan vrid add ipaddress

```
configure vrrp vlan <vlan_name> vrid <vridval> add <ipaddress>
```

Description

Associates a virtual IP address with a specific VRRP virtual router.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
ipaddress	Specifies the IP address of the virtual router in which this device participates.

Usage Guidelines

The restrictions on this command are as follows:

- If the priority of the VRRP virtual router is 255, the IP address to be added must be owned by the VLAN on which the VRRP virtual router exists. If the priority is not 255, the IP address must not be owned by that VLAN.
- When a VRRP virtual router is enabled, it must have at least one virtual IP address.

Example

Create a VLAN named *vlan-1* with an ipaddress of 10.1.2.2 and a VRRP VRID of 1:

```
create vlan vlan-1
configure vlan vlan-1 ipaddress 10.1.2.2
create vrrp vlan-1 vrid 1
```

The following example associates IP address 10.1.2.3 to VLAN *vlan-1*:

```
configure vrrp vlan vlan-1 vrid 1 add 10.1.2.3
```

configure vrrp vlan vrid add track-iproute

```
configure vrrp vlan <vlan_name> vrid <vridval> add track-iproute
<ipaddress>/<masklength>
```

Description

Creates a tracking entry for the specified route. When this route becomes unreachable, this entry is considered to be failing.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the virtual router ID of the target VRRP virtual router. Value can be in the range of 1-255.
ipaddress	Specifies the prefix of the route to track.
masklength	Specifies the length of the route's prefix.

Default

N/A.

Usage Guidelines

The route specified in this command might not exist in the IP routing table. When you create the entry for a route, an immediate VRRP failover might occur.

Example

The following command enables IP route failure tracking for routes to the specified subnet:

```
configure vrrp vlan vlan-1 vrid 1 add track-iproute 3.1.0.0/24
```

configure vrrp vlan vrid add track-ping

```
configure vrrp vlan <vlan_name> vrid <vridval> add track-ping <ipaddress> frequency <seconds>
miss <misses>
```

Description

Creates a tracking entry for the specified IP address. The entry is tracked via pings to the IP address, sent at the specified frequency.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
ipaddress	Specifies the IP address to be tracked.
seconds	Specifies the number of seconds between pings to the target IP address. The range is 1 to 600 seconds.

misses	Specifies the number of misses allowed before this entry is considered to be failing. The range is 1 to 255 pings.
--------	--

Default

N/A.

Usage Guidelines

Adding an entry with the same IP address as an existing entry causes the new values to overwrite the existing entry's frequency and miss number.

Example

The following command enables ping tracking for the external gateway at 3.1.0.1, pinging every 3 seconds, and considering the gateway to be unreachable if no response is received to 5 consecutive pings:

```
configure vrrp vlan vlan-1 vrid 1 add track-ping 3.1.0.1 frequency 3 miss 5
```

configure vrrp vlan vrid add track-vlan

```
configure vrrp vlan <vlan_name> vrid <vridval> add track-vlan <target_vlan_name>
```

Description

Configures a VRRP VLAN to track port connectivity to a specified VLAN. When this VLAN is in the “down” state, this entry is considered to be failing.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
target_vlan_name	Specifies the name of the VLAN to track.

Default

N/A.

Usage Guidelines

Up to eight VLANs can be tracked.

Deleting a tracked VLAN does not constitute a failover event for the VR tracking it, and the tracking entry is deleted.

Example

The following command enables VRRP VLAN *vlan-1* to track port connectivity to VLAN *vlan-2*:

```
configure vrrp vlan vlan-1 vrid 1 add track-vlan vlan-2
```

configure vrrp vlan vrid advertisement-interval

```
configure vrrp vlan <vlan_name> vrid <vridval> advertisement-interval <interval> [{seconds} | milliseconds]
```

Description

Configures the time between VRRP advertisements (pings) in seconds or milliseconds.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
interval	Specifies the time interval between advertisements in seconds unless otherwise specified as milliseconds. The default is 1 second. The range is 1 through 255 seconds or 100 through 999 milliseconds. Note: If you must configure the range in milliseconds, specify the <code>milliseconds</code> keyword. If you enter a number from 100 through 255 and do not specify the <code>milliseconds</code> keyword, the interval defaults to seconds.

Default

The advertisement interval is 1 second.

Usage Guidelines

The advertisement interval specifies the interval between advertisements sent by the master router to inform the backup routers that its alive. You must use whole integers when configuring the advertisement interval.

An extremely busy CPU can create a short dual master situation. To avoid this, increase the advertisement interval.

The seconds range is 1 through 255. The switch uses the seconds interval by default. You do not need to specify the `seconds` keyword to configure the advertisement interval in seconds.

The milliseconds range is 100 through 999. Since the switch uses the seconds interval by default, you must specify the `milliseconds` keyword to configure the advertisement interval in milliseconds.

To view your VRRP configuration, including the configured advertisement interval, use one of the following commands:

- `show vrrp {detail}`
- `show vrrp vlan <vlan_name> {stats}`

Out of Range Error Messages

If you enter a number that is out of the seconds or milliseconds range, the switch displays an error message.

In the following example, the configured advertisement interval is set to 999; however, the `milliseconds` keyword is missing. Since the switch uses the seconds interval by default, and 999 is out of the seconds range, the switch displays an error message similar to the following:

```
configure vrrp blue vrid 250 advertisement-interval 999
Error: Advertisement interval must be between 1 and 255 seconds. 999 out of range
```

If you configure an out of range milliseconds interval, the switch displays an error message similar to the following:

```
configure vrrp blue vrid 2 advertisement-interval 1000 milliseconds
Error: Advertisement interval must be between 100 and 999 milliseconds. 1000 out of range.
```

Example

The following command configures the advertisement interval of 15 seconds:

```
configure vrrp vlan vrrp-1 vrid 1 advertisement-interval 15
```

The following command configures the advertisement interval of 200 milliseconds:

```
configure vrrp vlan vrrp-1 vrid 1 advertisement -interval 200 milliseconds
```

configure vrrp vlan vrid authentication

```
configure vrrp vlan <vlan_name> vrid <vridval> authentication [none | simplepassword
<password>]
```

Description

Configures the authentication type for a specific VRRP virtual router.

Syntax Description

<code>vlan_name</code>	Specifies the name of a VRRP VLAN.
<code>vridval</code>	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
<code>password</code>	Specifies the user-defined password for authentication.

Default

Authentication is set to *none*.

Usage Guidelines

This command can add a modest amount of security to VRRP advertisements. All VRRP routers using the same VRID must use the same password when using this feature.

A simple password must be between 1 and 8 characters long.

Example

The following command configures authentication for VRRP VLAN *vrrp-1* with the password *newvrrp*:

```
configure vrrp vlan vrrp-1 vrid 1 authentication simple-password newvrrp
```

configure vrrp vlan vrid delete

```
configure vrrp vlan <vlan_name> vrid <vridval> delete <ipaddress>
```

Description

Deletes a virtual IP address from a specific VRRP virtual router.

Syntax Description

<code>vlan_name</code>	Specifies the name of a VRRP VLAN.
<code>vridval</code>	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
<code>ipaddress</code>	Specifies the IP address of the virtual router in which this device participates.

Usage Guidelines

The restrictions on this command are as follows:

- If the priority of the VR is 255, the IP address to be added must be owned by the VLAN on which the VR exists. If the priority is not 255, the IP address must not be owned by that VLAN.
- When a VR is enabled, it must have at least one virtual IP address. When the VR is not enabled, there are no restrictions on deleting the IP address.

Example

The following command removes IP address 10.1.2.3 from VLAN *vlan-1*:

```
configure vrrp vlan vlan-1 vrid 1 delete 10.1.2.3
```

configure vrrp vlan vrid delete track-iproute

```
configure vrrp vlan <vlan_name> vrid <vridval> delete track-iproute
<ipaddress>/<masklength>
```

Description

Deletes a tracking entry for the specified route.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
ipaddress	Specifies the prefix of the route to track.
masklength	Specifies the length of the route's prefix.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

Example

The following command disables tracking of routes to the specified subnet for VLAN *vlan-1*:

```
configure vrrp vlan vlan-1 vrid 1 delete track-iproute 3.1.0.0/24
```

configure vrrp vlan vrid delete track-ping

```
configure vrrp vlan <vlan_name> vrid <vridval> delete track-ping <ipaddress>
```

Description

Deletes a tracking entry for the specified IP address.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
ipaddress	Specifies the IP address to be tracked.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

A VRRP node with a priority of 255 might not recover from a ping-tracking failure if there is a Layer 2 switch between it and another VRRP node. In cases where a Layer 2 switch is used to connect VRRP nodes, NETGEAR recommends that those nodes have priorities of less than 255.

Example

The following command disables ping tracking for the external gateway at 3.1.0.1:

```
configure vrrp vlan vlan-1 vrid 1 delete track-ping 3.1.0.1
```

configure vrrp vlan vrid delete track-vlan

```
configure vrrp vlan <vlan_name> vrid <vridval> delete track-vlan <target_vlan_name>
```

Description

Deletes the tracking of port connectivity to a specified VLAN.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies the VRRP virtual router ID of the target virtual router. Value can be in the range of 1-255.
target_vlan_name	Specifies the name of the tracked VLAN.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

Example

The following command disables the tracking of port connectivity to VLAN *vlan-2*:

```
configure vrrp vlan vlan-1 vrid 1 delete track-vlan vlan-2
```

configure vrrp vlan vrid dont-preempt

```
configure vrrp vlan <vlan_name> vrid <vridval> dont-preempt
```

Description

Specifies that a higher priority backup router does not preempt a lower priority master.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

The default setting is preempt.

Usage Guidelines

The preempt mode controls whether a higher priority backup router preempts a lower priority master. `dont-preempt` prohibits preemption. The router that owns the virtual IP address always preempts, independent of the setting of this parameter.

Example

The following command disallows preemption:

```
configure vrrp vlan vlan-1 vrid 1 dont-preempt
```

configure vrrp vlan vrid preempt

```
configure vrrp vlan <vlan_name> vrid <vridval> preempt {delay <seconds>}
```

Description

Specifies that a higher priority backup router preempts a lower priority master.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
seconds	Specifies a preempt delay period in seconds. The value range is 1 to 3600 seconds, or 0, which selects the original preempt delay period.

Default

Preempt enabled.

Delay configuration: 0.

Usage Guidelines

The preempt option enables a higher-priority backup router to preempt a master with a lower priority. When a VRRP enabled router receives a lower priority VRRP advertisement and preemption is enabled, the higher-priority VRRP enabled router takes over as master. The new master starts sending VRRP advertisements and the old, lower-priority master relinquishes mastership.

Note: The router that owns the virtual IP address always preempts, independent of the setting of this parameter.

When a VRRP enabled router preempts the master, it does so in one of the following ways:

- If the preempt delay timer is configured for between 1 and 3600 seconds and the lower-priority master is still operating, the router preempts the master when the timer expires.
- If the preempt delay timer is configured for 0, the router preempts the master after 3 times the hello interval.
- If the higher priority router stops receiving advertisements from the current master for 3 times the hello interval, it takes over mastership immediately.

Note: The preempt feature can be disabled with the `configure vrrp vlan vrid dont-preempt` command.

Example

The following command allows preemption:

```
configure vrrp vlan vlan-1 vrid 1 preempt
```

configure vrrp vlan vrid priority

```
configure vrrp vlan <vlan_name> vrid <vridval> priority <priorityval>
```

Description

Configures the priority value of a VRRP virtual router.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.
priorityval	Specifies the priority value of the router. The default is 100. The priority range is 1-255.

Default

The default priority is 100.

Usage Guidelines

This command changes the priority of a VRRP VR. If the VR is assigned an IP address that is physically owned by the switch, the VR's priority is 255 and cannot be changed. If the IP address is not owned by switch, the priority cannot be changed to 255.

To change the priority in either of the described cases, disable VRRP and remove the virtual IP address(es) first.

To disable VRRP, use the `disable vrrp {vlan <vlan_name> vrid <vridval>}` command. To remove the virtual IP address(es), use the `configure vrrp vlan <vlan_name> vrid <vridval> delete <ipaddress>` command.

Example

The following command configures a priority of 150 for VLAN *vrrp-1*:

```
configure vrrp vlan vrrp-1 vrid 1 priority 150
```

configure vrrp vlan vrid track-mode

```
configure vrrp vlan <vlan_name> vrid <vridval> track-mode [all | any]
```

Description

Defines the conditions under which the router automatically relinquishes master status when the tracked entities fail.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

all	Specifies that the mastership is relinquished when one of the following events occur: <ul style="list-style-type: none"> • All of the tracked VLANs fail • All of the tracked routes fail • All of the tracked PINGs fail
any	Specifies that the mastership is relinquished when any of the tracked VLANs, routes, or PINGs fail.

Default

The default setting is all.

Usage Guidelines

None.

Example

The following command configures the track mode to `any`:

```
configure vrrp vlan vrrp-1 vrid 1 track-mode any
```

create vrrp vlan vrid

```
create vrrp vlan <vlan_name> vrid <vridval>
```

Description

Creates a VRRP virtual router on the switch.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

VRRP Virtual Router IDs can be used across multiple VLANs. You can create multiple virtual routers on different VLANs. Virtual Router IDs need not be unique to a specific VLAN.

Before configuring any virtual router parameters, you must first create the VRRP instance on the switch. If you define VRRP parameters before creating the VRRP, you might see an error similar to the following:

Error: VRRP VR for vlan vrrp1, vrid 1 does not exist.
 Please create the VRRP VR before assigning parameters.
 Configuration failed on backup MSM, command execution aborted!

If this happens, create the VRRP instance and then configure its parameters.

Example

The following command creates a VRRP router on VLAN *vrrp-1*, with a VRRP virtual router ID of 1:

```
create vrrp vlan vrrp-1 vrid 1
```

delete vrrp vlan vrid

```
delete vrrp vlan <vlan_name> vrid <vridval>
```

Description

Deletes a specified VRRP virtual router.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the VRRP virtual router on the VLAN *vrrp-1* identified by VRID 2:

```
delete vrrp vlan vrrp-1 vrid 2
```

disable vrrp vrid

```
disable vrrp {vlan <vlan_name> vrid <vridval>}
```

Description

Disables a specific VRRP virtual router or all VRRP virtual routers.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vidval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

This disables a specific virtual router on the switch. If no VRRP VLAN is specified, all virtual routers on the switch are disabled.

Example

The following command disables all VRRP virtual routers on the switch:

```
disable vrrp
```

enable vrrp vrid

```
enable vrrp {vlan <vlan_name> vrid <vridval>}
```

Description

Enables a specific VRRP virtual router or all VRRP virtual routers on the switch.

Syntax Description

vlan_name	Specifies the name of a VRRP VLAN.
vridval	Specifies a VRRP Virtual Router ID (VRID). Value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

This enables a specific virtual router on the device. If you do not specify a VRRP virtual router, all VRRP virtual routers on this device are enabled.

Example

The following command enables all VRRP virtual routers on the switch:

```
enable vrrp
```

show vrrp

`show vrrp {detail}`

Description

Displays VRRP configuration information for all VRRP VLANs.

Syntax Description

detail	Specifies more detailed VRRP information.
--------	---

Default

N/A.

Usage Guidelines

Depending on the software version running on your switch or your switch model, additional or different VRRP information might be displayed.

If you specify the command without the `detail` keyword, the following VRRP information appears:

- VLAN Name—The name of the VRRP VLAN and whether VRRP is enabled or disabled on the VLAN. The enable/disable state appears as follows:
 - En—VRRP is enabled on this VLAN.
 - Ds—VRRP is disabled on this VLAN.
- VRID—The VRRP Virtual Router Identification number for the VRRP VLAN.
- Pri—The priority value of the VRRP VLAN.
- Virtual IP Addr—If configured, the virtual IP address associated with the VRRP VLAN.
- State—The current state of the VRRP router. The state includes the following:
 - Init—The VRRP router is in the initial state.
 - Backup—The VRRP router is a backup router.
 - Master—The VRRP router is the master router.
- Master Mac Address—The MAC address of the master VRRP router.
- TP—Indicates the number of tracked pings.
- TR—Indicates the number of tracked routes.
- TV—Indicates the number of tracked VLANs.
- P—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- T—Indicates the configured advertisement interval.

If you specify the command with the `detail` keyword, the switch displays similar information in a bulleted format, including:

- VLAN—The name of the VRRP VLAN.
- VRID—The VRRP Virtual Router Identification number for the VRRP VLAN.
- VRRP—The enabled/disabled state of VRRP on the VLAN.
- State—The current state of the VRRP router. The state includes the following:
 - Init—The VRRP router is in the initial state.
 - Backup—The VRRP router is a backup router.
 - Master—The VRRP router is the master router.
- Priority—The priority value of the VRRP VLAN.
- Advertisement Interval—Indicates the configured advertisement interval.
- Preempt—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- Authentication—If configured, identifies the VRRP simple password.
- Virtual IP Addresses—If configured, the virtual IP address associated with the VRRP VLAN.
- Tracked Pings—If configured, displays the:
 - Target IP address you are pinging.
 - Number of seconds between pings to the target IP address.
 - Number of misses allowed before this entry is considered to be failing.
- Tracked IP Routes—If configured, displays the IP address and subnet mask of the tracked route(s).
- Tracked VLANs—If configured, displays the name of the tracked VLAN(s).

Example

The following command displays a summary of status information for VRRP:

```
show vrrp
```

The following is sample output from this command:

```
VLAN Name VRID Pri Virtual IP Addr State Master Mac Address TP/TR/TV/P/T
v1(En)    0001 255 1.1.1.1          MSTR 00:00:5e:00:01:01  0 0 0 Y 1
```

En-Enabled, Ds-Disabled, Pri-Priority, T-Advert Timer, P-Preempt
TP-Tracked Pings, TR-Tracked Routes, TV-Tracked VLANs

The following command displays detailed status information for VRRP on all platforms:

```
show vrrp detail
```

The following is sample output from this command:

```
VLAN: v1      VRID: 1      VRRP: Enabled State: MASTER
Virtual Router: VR-Default
Priority: 255(master) Advertisement Interval: 1 sec
Preempt: Yes Authentication: simple-password key: foo
Virtual IP Addresses: 1.1.1.1
```

```

Tracked Pings: 10.10.10.3 / 10 second intervals / 5 misses allowed
Tracked IP Routes: -
Tracked VLANs: green, purple
* indicates a tracking condition has failed

```

show vrrp virtual-router

```
show vrrp virtual-router
```

Description

Displays the VRRP configuration information for all VRRP instances that belong to all created virtual routers.

Default

N/A

Usage Guidelines

This displays the following VRRP information:

- **VLAN Name**—The name of the VRRP VLAN and whether VRRP is enabled or disabled on the VLAN. The enable/disable state appears as follows
 - **En**—VRRP is enabled on this VLAN.
 - **DS**—VRRP is disabled on this VLAN.
- **VRID**—The VRRP Virtual Router Identification number for the VRRP VLAN.
- **Pri**—The priority value of the VRRP VLAN.
- **Virtual IP Addr**—The virtual IP address associated with the VRRP VLAN.
- **State**—The current state of the VRRP router. The state appears as follows:
 - **INIT**—The VRRP router is in the initial state.
 - **BKUP**—The VRRP router is a backup router.
 - **MSTR**—The VRRP router is the master router.
- **Master Mac Address**—The MAC address of the master VRRP router.
- **Virtual-Router**—The virtual router to which a particular VLAN belongs.

Example

This command produces output similar to the following:

VLAN Name	VRID	Pri	Virtual IP Addr	State	Master Mac Address	Virtual-Router
vlan_3(En)	0003	200	30.1.2.1	MSTR	00:00:5e:00:01:03	vir_3
vlan_4(En)	0003	200	40.1.2.1	MSTR	00:00:5e:00:01:03	vir_4
vlan_5(En)	0005	100	50.1.2.1	BKUP	00:00:5e:00:01:05	vir_3
vlan_6(En)	0005	100	60.1.2.1	BKUP	00:00:5e:00:01:05	vir_4

show vrrp virtual-router`show vrrp virtual-router <vr-name>`**Description**

Displays the VRRP that matches the specified virtual router.

Syntax Description

vr-name	Specifies the name of the virtual router
---------	--

Default

N/A

Usage Guidelines

The command displays the VRRP information only for a specific virtual router on the device. If you do not specify a virtual router (refer to `show vrrp virtual-router`), all VRRP virtual routers on this device are displayed.

This displays the following VRRP information:

- **VLAN Name**—The name of the VRRP VLAN and whether VRRP is enabled or disabled on the VLAN. The enable/disable state appears as follows:
 - **En**—VRRP is enabled on this VLAN.
 - **Ds**—VRRP is disabled on this VLAN.
- **VRID**—The VRRP Virtual Router Identification number for the VRRP VLAN.
- **Pri**—The priority value of the VRRP VLAN.
- **Virtual IP Addr**—If configured, the virtual IP address associated with the VRRP VLAN.
- **State**—The current state of the VRRP router. The state includes the following:
 - **Init**—The VRRP router is in the initial state.
 - **Backup**—The VRRP router is a backup router.
 - **Master**—The VRRP router is the master router.
- **Master Mac Address**—The MAC address of the master VRRP router.
- **TP**—Indicates the number of tracked pings.
- **TR**—Indicates the number of tracked routes.
- **TV**—Indicates the number of tracked VLANs.
- **P**—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- **T**—Indicates the configured advertisement interval.

Example

The following command displays VRRP information for virtual router #3:

```
show vrrp vir_3
```

This command produces output similar to the following:

```
VLAN Name VRID Pri Virtual IP Addr State Master Mac Address TP/TR/TV/P/T
vlan_3(En) 0003 200 30.1.2.1 MSTR 00:00:5e:00:01:03 0 0 0 Y 1
vlan_5(En) 0005 100 50.1.2.1 BKUP 00:00:5e:00:01:05 0 0 0 Y 1
```

show vrrp vlan

```
show vrrp vlan <vlan_name> {stats}
```

Description

Displays VRRP information for a particular VLAN.

Syntax Description

<code>vlan_name</code>	Specifies the name of a VRRP VLAN.
<code>stats</code>	Specifies statistics for a particular VLAN.

Default

N/A.

Usage Guidelines

Depending on the software version running on your switch or your switch model, additional or different VRRP information might be displayed.

If you specify the command without the `stats` keyword, the following VRRP information appears:

- **VLAN**—The name of the VRRP VLAN.
- **VRID**—The VRRP Virtual Router Identification number for the VRRP VLAN.
- **VRRP**—The enabled/disabled state of VRRP on the VLAN.
- **State**—The current state of the VRRP router. The state includes the following:
 - **Init**—The VRRP router is in the initial state.
 - **Backup**—The VRRP router is a backup router.
 - **Master**—The VRRP router is the master router.
- **Priority**—The priority value of the VRRP VLAN.
- **Advertisement Interval**—Indicates the configured advertisement interval.

- **Preempt**—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- **Authentication**—If configured, identifies the VRRP simple password.
- **Virtual IP Addresses**—If configured, the virtual IP address associated with the VRRP VLAN.
- **Tracked Pings**—If configured, displays the:
 - Target IP address you are pinging.
 - Number of seconds between pings to the target IP address.
 - Number of misses allowed before this entry is considered to be failing.
- **Tracked IP Routes**—If configured, displays the IP address and subnet mask of the tracked route(s).
- **Tracked VLANs**—If configured, displays the name of the tracked VLAN(s).

If you specify the `stats` keyword, you see counter and statistics information for the specified VRRP VLAN.

Example

The following command displays configuration information for the specified VRRP VLAN:

```
show vrrp vlan blue
```

The following is sample output from this command:

```
VLAN:  blue      VRID:  2          VRRP:  Disabled State:  INIT
Priority:  1(backup)  Advertisement Interval:  1 sec
Preempt:  Yes    Authentication:  None
Virtual IP Addresses:
Tracked Pings:  -
Tracked IP Routes:  -
Tracked VLANs:  -
* indicates a tracking condition has failed
```

The following command displays statistics for VLAN *vrrp-1*:

```
show vrrp vlan vrrp-1 stats
```

The following is sample output from this command:

```
VLAN vrrp-1, VR ID 25
  Chksum Err:0, Ver Err:0, VRID Err:0, Auth Mismatch:0, Pkt-len Err:0
  Become Master:0, Adv rcv:0, Adv Err:0, Auth Fail:0, TTL Err:0
  Pri-0-rcv:0, Pri-0-snt:0, Addr-List Err:0, Invalid Auth Err:0
```

This chapter describes commands for configuring and managing the following IP protocols and functions:

- DHCP and BOOTP relay
- IP ARP
- IP routing
- IP multinetting
- IP broadcast handling
- Broadcast UDP packet forwarding
- Static routes
- ICMP
- IRDP
- VLAN aggregation

For an introduction to these IP protocols and functions, see the *NETGEAR 8800 User Manual*.

clear iparp

```
clear iparp {<ip_addr> {vr <vr_name>} | vlan <vlan_name> | vr <vr_name>}
```

Description

Removes dynamic entries in the IP ARP table.

Syntax Description

ip_addr	Specifies an IP address.
vlan_name	Specifies a VLAN name.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Permanent IP ARP entries are not affected.

This command is specific to a single virtual router, and it applies to the current virtual router if you do not specify a virtual router.

Example

The following command removes a dynamically created entry from the IPARP table:

```
clear iparp 10.1.1.5
```

configure bootprelay add

```
configure bootprelay add <ip_address> {vr <vrid>}
```

Description

Configures the addresses to which BOOTP requests should be directed.

Syntax Description

ip_address	Specifies an IP address.
vrid	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets.

To configure the relay function, follow these steps:

1. Configure VLANs and IP unicast routing.
2. Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

```
configure bootprelay add <ip_address>
```

3. Enable the DHCP or BOOTP relay function, using the following command:

```
enable bootprelay
```

Example

The following command configures BOOTP requests to be directed to 123.45.67.8:

```
configure bootprelay add 123.45.67.8
```

configure bootprelay delete

```
configure bootprelay delete [<ip_address> | all] {vr <vrid>}
```

Description

Removes one or all IP destination addresses for forwarding BOOTP packets.

Syntax Description

ip_address	Specifies an IP address.
vrid	Specifies a VR name.
all	Specifies all IP address entries.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

None.

Example

The following command removes the destination address:

```
configure bootprelay delete 123.45.67.8
```

configure bootprelay dhcp-agent information check

```
configure bootprelay dhcp-agent information check
```

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to prevent DHCP reply packets with invalid or missing relay agent options from being forwarded to the client.

To disable this check, use the following command:

```
unconfigure bootprelay dhcp-agent information check
```

Example

The following command configures the DHCP relay agent option check:

```
configure bootprelay dhcp-agent information check
```

configure bootprelay dhcp-agent information circuit-id port-information

```
configure bootprelay dhcp-agent information circuit-id port-information <port_info> port
<port>
```

Description

Configures the circuit ID sub-option that identifies the port for an incoming DHCP request.

Syntax Description

port_info	Specifies a text string that becomes the circuit ID sub-option for the specified port. Specify a text string composed of 1 to 32 characters.
port	Specifies the port to which the circuit ID sub-option is assigned.

Default

The default port_info is encoded as ((slot_number * 1000) + port_number/portlindex). For example, if the DHCP request is received on port 3:12, the default circuit ID port_info value is 3012.

Usage Guidelines

The full circuit ID string uses the format <vlan_info>-<port_info>. To configure the vlan_info portion of the circuit ID string, use the following command:

```
configure bootprelay dhcp-agent information circuit-id vlan-information <vlan_info>
{vlan} [<vlan_name>|all]
```

To display the port_info information, use the following command:

```
show bootprelay dhcp-agent information circuit-id port-information ports all
```

Example

The following command configures the circuit ID port_info value *slot1port3* for port 1:3:

```
configure bootprelay dhcp-agent information circuit-id port-information slot1port3 port 1:3
```

configure bootprelay dhcp-agent information circuit-id vlan-information

```
configure bootprelay dhcp-agent information circuit-id vlan-information <vlan_info> {vlan}
[<vlan_name>|all]
```

Description

Configures the circuit ID sub-option that identifies the VLAN for an incoming DHCP request.

Syntax Description

vlan_info	Specifies a text string that becomes the circuit ID sub-option for the specified VLAN. Specify a text string composed of 1 to 32 characters.
vlan_name	Specifies the VLAN to which the circuit ID sub-option is assigned.
all	Specifies that the vlan_info entered is to be used in the circuit ID sub-option for all VLANs.

Default

The default vlan_info for each VLAN is the VLAN ID or tag.

Usage Guidelines

The full circuit ID string uses the format <vlan_info>-<port_info>. To configure the port_info portion of the circuit ID string, use the following command:

```
configure bootprelay dhcp-agent information circuit-id port-information <port_info> port
<port>
```

To display the vlan_info information, use the following command:

```
show bootprelay dhcp-agent information circuit-id vlan-information
```

Example

The following command configures the circuit ID vlan_info value *VLANblue* for VLAN *blue*:

```
configure bootprelay dhcp-agent information circuit-id vlan-information VLANblue blue
```

configure bootprelay dhcp-agent information option

```
configure bootprelay dhcp-agent information option
```

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward DHCP or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets.

To configure the relay function, follow these steps:

1. Configure VLANs and IP unicast routing.
2. Enable the DHCP or BOOTP relay function, using the following command:


```
enable bootprelay [{vlan} [<vlan_name>] | [{vr} <vr_name>] | all [{vr} <vr_name>]}
```
3. Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

```
configure bootprelay add <ip_address> {vr <vrid>}
```

4. Configure the DHCP relay agent option (option 82), using the following command:

```
configure bootprelay dhcp-agent information option
```

To disable the DHCP relay agent option (option 82), use the following command:

```
unconfigure bootprelay dhcp-agent information option
```

Example

The following command configures the DHCP relay agent option:

```
configure bootprelay dhcp-agent information option
```

configure bootprelay dhcp-agent information policy

```
configure bootprelay dhcp-agent information policy [drop | keep | replace]
```

Description

Configures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

drop	Specifies to drop the packet.
keep	Specifies to keep the existing option 82 information in place.

replace	Specifies to replace the existing data with the switch's own data.
---------	--

Default

Replace.

Usage Guidelines

Use this command to set a policy for the relay agent. Packets can be dropped, the option 82 information can be replaced (the default), or the packet can be forwarded with the information unchanged.

Example

The following command configures the DHCP relay agent option 82 policy to keep:

```
configure bootprelay dhcp-agent information policy keep
```

configure bootprelay dhcp-agent information remote-id

```
configure bootprelay dhcp-agent information remote-id [<remote_id> | system-name] {vr
<vrid>}
```

Description

Configures the remote ID sub-option that identifies the relaying switch for DHCP requests and replies.

Syntax Description

remote_id	Specifies a text string that becomes the remote ID sub-option for the switch. Specify a text string composed of 1 to 32 characters.
system-name	Specifies that the switch name is used as the remote ID sub-option for the switch.
vrid	Specifies the VR on which to configure the remote ID sub-option.

Default

The switch MAC address.

Usage Guidelines

To display the remote-ID, use the following command:

```
show bootprelay
```

Example

The following command configures the remote ID sub-option to specify the switch name in DHCP requests and replies:

```
configure bootprelay dhcp-agent information remote-id system-name
```

configure forwarding sharing

```
configure forwarding sharing [L3 | L3_L4]
```

Description

Identifies the fields that are used to select ECMP routes and load-sharing group ports.

Syntax Description

L3	Uses only Layer 3 IP addresses to select ECMP routes and load-sharing ports.
L3_L4	Uses Layer 3 IP addresses and Layer 4 TCP/UDP port numbers, if present, to select ECMP routes and load-sharing ports.

Default

L3_L4

Usage Guidelines

This command configures the criteria used to select ECMP routes and load-sharing group ports.

For ECMP routes, the configured criteria selects the next hop gateway. The `L3` option uses only the source and destination IP addresses to select the next hop gateway. The `L3_L4` option uses the Layer 4 TCP or UDP port and the source and destination IP addresses to select the next hop gateway.

For load-sharing groups (link aggregation groups), the configured criteria selects the load-sharing group port. The load-sharing groups can be configured to use the following address-based algorithms:

- `L2`—Specifies port selection based on Layer 2 information.
- `L3`—Specifies port selection based on Layer 3 information.
- `L3_L4`—Specifies port selection based on Layer 3 and Layer 4 information.

This command affects all the load-sharing groups that use either the `L3` or `L3_L4` link aggregation algorithm. If the `L3` option is specified, all the load-sharing groups that are configured with either the `L3` or the `L3_L4` address-based link aggregation algorithm use just the Layer 3 IP addresses for the egress port selection. Similarly if the `L3_L4` option is specified, all the load-sharing groups that are configured with either `L3` or `L3_L4`

address-based link aggregation algorithm use the Layer 3 IP addresses and Layer 4 port number for the egress port selection.

Selecting the `L3` option over `L3_L4` can be useful in a network where IP fragments are present, since only the first fragment contains the Layer 4 TCP or UDP port number. If the `L3` option is selected, all IP fragments in a given TCP or UDP session use the same ECMP gateway or load-sharing group port, potentially avoiding inefficient packet reordering by the destination. If IP fragments are not prevalent, better traffic distribution can be achieved by selecting `L3_L4`.

To display the forwarding sharing feature configuration, enter the command:

```
show forwarding configuration
```

Example

The following command modifies the sharing selection criteria to use just the Layer 3 IP addresses:

```
configure forwarding sharing L3
```

The following command modified the sharing selection criteria to use the Layer 3 and Layer 4 information:

```
configure forwarding sharing L3_L4
```

configure iparp add

```
configure iparp add <ip_addr> {vr <vr_name>}} <mac>
```

Description

Adds a permanent entry to the ARP table. Specify the IP address and MAC address of the entry.

Syntax Description

<code>ip_addr</code>	Specifies an IP address.
<code>mac</code>	Specifies a MAC address.
<code>vr_name</code>	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

None.

Example

The following command adds a permanent IP ARP entry to the switch for IP address 10.1.2.5:

```
configure iparp add 10.1.2.5 00:11:22:33:44:55
```

configure iparp add proxy

```
configure iparp add proxy [<ipNetmask> | <ip_addr> {<mask>}] {vr <vr_name>} {<mac> | vrrp} {always}
```

Description

Configures the switch to respond to ARP Requests on behalf of devices that are incapable of doing so.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_addr	Specifies an IP address.
mask	Specifies a subnet mask.
mac	Specifies a MAC address to use in the ARP reply.
vrrp	Specifies a MAC address to use in the ARP reply. For VLANs running VRRP, the switch replies with the VRRP virtual MAC. For non-VRRP VLANs, the switch replies with the switch MAC.
always	Specifies that the switch responds regardless of the VLAN that the request arrives from.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

When `mask` is not specified, an address with the mask 255.255.255.255 is assumed. When neither `mac` nor `vrrp` is specified, the MAC address of the switch is used in the ARP Response. When `always` is specified, the switch answers ARP Requests without filtering requests that belong to the same subnet of the receiving router interface.

After IP ARP is configured, the system responds to ARP requests on behalf of the device as long as the following conditions are satisfied:

- The valid IP ARP request is received on a router interface.
- The target IP address matches the IP address configured in the proxy ARP table.

- The source IP address is not on the same subnet as the target address (unless the always flag is set).

After all the proxy ARP conditions have been met, the switch formulates an ARP response using the configured MAC address in the packet.

The default maximum number of proxy entries is 256, but can be increased to 4096 by using the following command:

```
configure iparp max_proxy_entries {vr <vr_name>} <max_proxy_entries>
```

Example

The following command configures the switch to answer ARP requests for all devices with the address range of 100.101.45.1 to 100.101.45.255:

```
configure iparp add proxy 100.101.45.0/24
```

configure iparp delete

```
configure iparp delete <ip_addr> {vr <vr_name>}
```

Description

Deletes an entry from the ARP table. Specify the IP address of the entry to delete.

Syntax Description

ip_addr	Specifies an IP address.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Removes any IP ARP entry (dynamic or permanent) from the table.

Example

The following command deletes an IP address entry from the ARP table:

```
configure iparp delete 10.1.2.5
```

configure iparp delete proxy

```
configure iparp delete proxy [[<ipNetmask> | <ip_addr> {<mask>}]] {vr <vr_name>} | all]
```


Description

Deletes one or all proxy ARP entries.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_addr	Specifies an IP address.
mask	Specifies a subnet mask.
all	Specifies all ARP entries.
vr_name	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

When the `mask` is not specified, the software assumes a host address (that is, a 32-bit mask).

Example

The following command deletes the IP ARP proxy entry *100.101.45.0/24*:

```
configure iparp delete proxy 100.101.45.0/24
```

configure iparp max_entries

```
configure iparp max_entries {vr <vr_name>} <max_entries>
```

Description

Configures the maximum allowed IP ARP entries.

Syntax Description

vr_name	Specifies a VR name.
max_entries	Specifies a number of maximum IP ARP entries. The range is 1 to 20480.

Default

8192 for the Default virtual router and all user-created virtual routers.

4096 for the VR-Mgmt virtual router.

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

The maximum IP ARP entries include dynamic, static, and incomplete IP ARP entries.

The NETGEAR 8800 software allows you to configure up to 32768 maximum entries on each VR, but when the hardware limit is reached, the switch displays a message and can no longer store additional ARP entries.

Example

The following command sets the maximum IP ARP entries to 2000 entries:

```
configure iparp max-entries 2000
```

configure iparp max_pending_entries

```
configure iparp max_pending_entries {vr <vr_name>} <max_pending_entries>
```

Description

Configures the maximum allowed incomplete IP ARP entries.

Syntax Description

vr_name	Specifies a VR name.
max_pending_entries	Specifies a number of maximum IP ARP entries.

Default

256.

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Range: 1 - 4096.

Example

The following command sets the maximum pending IP ARP entries to 500 entries:

```
configure iparp max_pending_entries 500
```

configure iparp max_proxy_entries

```
configure iparp max_proxy_entries {vr <vr_name>} <max_proxy_entries>
```

Description

Configures the maximum allowed IP ARP proxy entries.

Syntax Description

vr_name	Specifies a VR name.
max_proxy_entries	Specifies maximum number of IP ARP proxy entries.

Default

256.

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Range: 0 - 4096.

Example

The following command sets the maximum IP ARP proxy entries to 500 entries:

```
configure iparp max_proxy_entries 500
```

configure iparp timeout

```
configure iparp timeout {vr <vr_name>} <minutes>
```

Description

Configures the IP ARP timeout period.

Syntax Description

vr_name	Specifies which virtual router IP ARP setting to change.
minutes	Specifies a time in minutes.

Default

20 minutes.

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

The range is 0-32,767. A setting of 0 disables timeout.

When the switch learns an ARP entry, it begins the timeout for that entry. When the timer reaches 0, the entry is aged out, unless IP ARP refresh is enabled. If ARP refresh is enabled, the switch sends an ARP request for the address before the timer expires. If the switch receives a response, it resets the timer for that address.

Example

The following command sets the IP ARP timeout period to 10 minutes:

```
configure iparp timeout 10
```

configure ipforwarding originated-packets

```
configure ipforwarding originated-packets [ require-ipforwarding | dont-require-ipforwarding ]
```

Description

Configures the ipforwarding behavior.

Syntax Description

This command has no arguments or variables.

Default

```
dont-require-ipforwarding.
```

Usage Guidelines

You can configure IP forwarding to be either required or not required.

Example

The following command configures required ipforwarding:

```
configure ipforwarding originated-packets require-ipforwarding
```

configure iproute add (IPv4)

```
configure iproute add [<ipNetmask> | <ip_addr> <mask>] <gateway> {metric} {multicast | multicast-only | unicast | unicast-only} {vr <vrname>}
```

Description

Adds a static route to the specified routing table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_addr	Specifies an IP address.
mask	Specifies a subnet mask.
gateway	Specifies a gateway IP address.
metric	Specifies a cost metric.

vrname	Specifies the virtual router to which the route is added.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a mask value of 255.255.255.255 to indicate a host entry.

The gateway address must be present on a directly attached subnet, or the following message appears:

```
ERROR: Gateway is not on directly attached subnet
```

The gateway address must be different from the VLAN address, or the following message appears:

```
ERROR: Gateway cannot be own address (x.x.x.x) #where x.x.x.x is the IP address specified
```

Note: Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the `multicast` option.

Example

The following command adds a static address to the routing table:

```
configure iproute add 10.1.1.0/24 123.45.67.1 5
```

configure iproute add blackhole

```
configure iproute add blackhole [<ipNetmask> | <ipaddress> <mask>] {multicast |
multicast-only | unicast | unicast-only} {vr <vrname>}
```

Description

Adds a blackhole address to the routing table. All traffic destined for a configured blackhole IP address is silently dropped, and no Internet Control Message Protocol (ICMP) message is generated.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ipaddress	Specifies an IP address.
mask	Specifies a subnet mask.
vrname	Specifies the virtual router to which the route is added.
multicast	Adds the blackhole route to the multicast routing table.
multicast-only	Adds the blackhole route to the multicast routing table.
unicast	Adds the blackhole route to the unicast routing table.
unicast-only	Adds the blackhole route to the unicast routing table.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

A blackhole entry configures packets with the specified destination IP subnet to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination IP subnet must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

Example

The following command adds a blackhole address to the routing table for packets with a destination address of 100.101.145.4:

```
configure iproute add blackhole 100.101.145.4/32
```

configure iproute add blackhole ipv4 default

```
configure iproute add blackhole ipv4 default {multicast | multicast-only | unicast |
unicast-only} {vr <vrname>}
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IP destination is silently dropped, and no Internet Control Message Protocol (ICMP) message is generated.

Syntax Description

vrname	Specifies the virtual router to which the route is added.
multicast	Adds the default blackhole route to the multicast routing table.

multicast-only	Adds the default blackhole route to the multicast routing table.
unicast	Adds the default blackhole route to the unicast routing table.
unicast-only	Adds the default blackhole route to the unicast routing table.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

While a default route is for *forwarding* traffic destined to an unknown IP destination, and a blackhole route is for *discarding* traffic destined to a specified IP destination, a *default blackhole* route is for *discarding* traffic to the unknown IP destination.

Using this command, all traffic with an unknown destination is discarded.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IP address for this route is 0.0.0.0.

Example

The following command adds a blackhole default route into the routing table:

```
configure iproute add blackhole default
```

configure iproute add default

```
configure iproute add default <gateway> {<metric>} {multicast | multicast-only | unicast |
unicast-only} {vr <vrname>}
```

Description

Adds a default gateway to the routing table.

Syntax Description

gateway	Specifies a VLAN gateway.
metric	Specifies a cost metric. If no metric is specified, the default of 1 is used.
multicast	Adds the default route to the multicast routing table.
multicast-only	Adds the default route to the multicast routing table.
unicast	Adds the default route to the unicast routing table.
unicast-only	Adds the default route to the unicast routing table.
vrname	Specifies the virtual router to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface. Use the `unicast-only` or `multicast-only` options to specify a particular traffic type. If not specified, both unicast and multicast traffic uses the default route.

Example

The following command configures a default route for the switch:

```
configure iproute add default 123.45.67.1
```

configure iproute delete

```
configure iproute delete [<ipNetmask> | <ipaddress> <mask>] <gateway> {multicast |
multicast-only | unicast | unicast-only} {vr <vrname>}
```

Description

Deletes a static address from the routing table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ipaddress	Specifies an IP address.
mask	Specifies a subnet mask.
gateway	Specifies a VLAN gateway.
multicast	Specifies a multicast route to delete.
multicast-only	Specifies a multicast route to delete.
unicast	Specifies a unicast route to delete.
unicast-only	Specifies a unicast route to delete.
vrname	Specifies the virtual router to which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a value of 255.255.255.255 or /32 for mask to indicate a host entry.

Example

The following command deletes an address from the gateway:

```
configure iproute delete 10.101.0.0/24 10.101.0.1
```

configure iproute delete blackhole

```
configure iproute delete blackhole [<ipNetmask> | <ipaddress> <mask>] {multicast |  
multicast-only | unicast | unicast-only} {vr <vrname>}
```

Description

Deletes a blackhole address from the routing table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ipaddress	Specifies an IP address.
mask	Specifies a netmask.
multicast	Specifies a blackhole multicast route to delete.
multicast-only	Specifies a blackhole multicast-only route to delete.
unicast	Specifies a blackhole unicast route to delete.
unicast-only	Specifies a blackhole unicast-only route to delete.
vrname	Specifies the virtual router from which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

None.

Example

The following command removes a blackhole address from the routing table:

```
configure iproute delete blackhole 100.101.145.4
```

configure iproute delete blackhole ipv4 default

```
configure iproute delete blackhole ipv4 default {multicast | multicast-only | unicast |  
unicast-only} {vr <vrname>}
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

multicast	Specifies a default blackhole multicast route to delete.
multicast-only	Specifies a default blackhole multicast route to delete. This option is provided for backward compatibility with releases prior to NETGEAR 8800 Release 12.1.
unicast	Specifies a default blackhole unicast route to delete.
unicast-only	Specifies a default blackhole unicast-only route to delete. This option is provided for backward compatibility with releases prior to NETGEAR 8800 Release 12.1.
vrname	Specifies a VR name.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

None.

Example

The following command deletes a blackhole default route from the routing table:

```
configure iproute delete blackhole default
```

configure iproute delete default

```
configure iproute delete default <gateway> {multicast | multicast-only | unicast |  
unicast-only} {vr <vrname>}
```

Description

Deletes a default gateway from the routing table.

Syntax Description

gateway	Specifies a VLAN gateway.
multicast	Specifies a default multicast route to delete.
multicast-only	Specifies a default multicast route to delete.
unicast	Specifies a default unicast route to delete.
unicast-only	Specifies a default unicast route to delete.
vrname	Specifies the virtual router to which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface.

Example

The following command deletes a default gateway:

```
configure iproute delete default 123.45.67.1
```

configure iproute priority

```
configure iproute {ipv4} priority [blackhole | bootp | ebgp | ibgp | icmp | ospf-as-external  
| ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static] <priority> {vr  
<vrname>}
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

blackhole	Specifies the blackhole route.
bootp	Specifies BOOTP.
ebgp	Specifies E-BGP routes
ibgp	Specifies I-BGP routes
icmp	Specifies ICMP.
ospf-as-external	Specifies OSPF as External routing.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPFInter routing.
ospf-intra	Specifies OSPFIntra routing.
rip	Specifies RIP.
static	Specifies static routes.
priority	Specifies a priority number in the range of 11 to 65534.
vrname	Specifies a virtual router name.

Default

Table 20 lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 20. Relative Route Priorities

Route Origin	Priority
Direct	10
MPLS	20
Blackhole	50
Static	1100
ICMP	1200
EBGP	1700
IBGP	1900
OSPFIntra	2200
OSPFInter	2300
RIP	2400
OSPFAsExt	3100
OSPF External 1	3200
OSPF External 2	3300
BOOTP	5000

Usage Guidelines

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences. If you change the route priority, you must save the configuration and reboot the system.

Note: The priority for a blackhole route cannot overlap with the priority of any other route origin.

Example

The following command sets IP route priority for static routing to 1200:

```
configure iproute priority static 1200
```

configure iproute reserved-entries

```
configure iproute reserved-entries [ <num_routes_needed> | maximum | default ] slot [all | <slot_num>]
```

Description

Reserves storage space for IPv4 and IPv6 routes in the Longest Prefix Match (LPM) hardware tables, allowing individual local and remote IPv4 unicast hosts to occupy the unused portions of the tables.

Syntax Description

num_routes_needed	Specifies a specific number of routes to reserve.
maximum	Reserves the maximum amount of space for IP route entries. No IPv4 hosts are stored in the LPM and External tables.
default	Reserves the default amount of space for IP route entries.
all	This option applies the reservation to all applicable slots.
slot_num	This option applies the reservation to the specified slot.

Default

The default value is 12240.

Usage Guidelines

Demand on the Layer 3 Hash table can be reduced by allowing IPv4 hosts to be stored in the LPM tables instead. This command allows you to reserve a portion of the LPM tables for routes, and this creates an unreserved portion that can be used to store IPv4 hosts. For more information, see the “Extended IPv4 Host Cache” section in the *NETGEAR 8800 User Manual*.

The default setting can support most networks, but if more than a few hundred local IP hosts and IP multicast entries are present, you can improve switch performance by calculating and configuring the reserved space for route entries to allow unreserved space for IPv4 hosts. Changing the number of reserved route entries does not require a reboot of the affected slots or switches.

You can view the current LPM hardware table usage by entering the `show iproute reserved-entries statistics` command. The LPM table statistics are in the columns under the *In HW Route Table* heading.

If the switch contains fewer routes than the capacity of the LPM tables, the number of route entries to reserve for a slot or switch should be the number of routes currently used in the hardware tables, plus an additional cushion for anticipated growth. Because each IPv6 route takes up the space of two IPv4 routes, the number of route entries to reserve is two times the value in the IPv6 routes column, plus the value in the IPv4 routes column, plus room for

anticipated growth. For example, if you want to reserve space for 100 IPv4 routes and 20 IPv6 routes, the required number of route entries is 140 (100 + 2*20).

Note: On a NETGEAR 8800 switch, the capacity of the LPM table is 4,096 higher than the capacity for local IPv4 or IPv6 hosts. Therefore, on such hardware, there is no need to configure fewer than 4096 reserved route entries.

The maximum value for `num_routes_needed` is 12256.

When `maximum` is specified, IPv4 hosts do not occupy LPM table space. Note that when `maximum` is specified, software forwarding can result, depending on the utilization and addresses in the Layer 3 Hash table, and is therefore not recommended.

If the switch contains more routes than the capacity of the LPM tables, say 700 routes on an 8800 module, a trade-off can be made. You can choose to reserve 400 iproute entries, for example. The 400 IPv4 routes with the longest length network masks will be installed in the LPM table, and the remainder of the LPM table can be used for cache space for local and remote hosts. The remote host entries are only required for IPv4 addresses matching one of the 300 routes not installed in the LPM table. Since not all 700 routes can be stored on an 8800 module anyway, leaving appropriate room for individual remote hosts can result in more fast-path forwarding.

Depending on the actual routes present, IP route compression can be enabled to reduce the number of routes required in the LPM tables. For more information, see the description for the following command:

```
enable iproute compression {vr <vrname>}
```

Example

The following command reserves up to 140 IPv4 routes or 70 IPv6 routes, or any combination in between:

```
configure iproute reserved-entries 140 slot all
```

For details on the configuration changes, see the command descriptions for the following commands:

```
show iproute reserved-entries
show iproute reserved-entries statistics
```

configure iproute sharing max-gateways

```
configure iproute sharing max-gateways <max_gateways>
```

Description

Depending on the platform, this command configures one of the following:

- The maximum number of gateways in each gateway set in the equal-cost multipath (ECMP) hardware table.
- The maximum number of gateways per subnet in the ECMP hardware table.

Syntax Description

max-gateways	Specifies the maximum number of ECMP gateways in a gateway set or for a subnet. The only values allowed are 2, 4, and 8.
--------------	--

Default

4 gateways

Usage Guidelines

When IP route sharing is enabled, the *maximum number of gateways* value represents the maximum number of next-hop gateways that can be used for communications with a destination subnet. Each gateway represents an alternative path to a subnet. The gateways can be defined with static routes, or they can be learned through the OSPF, BGP, or IS-IS protocols.

The *NETGEAR 8800 Release Notes* lists the total number of route destinations and the total combinations of gateway sets that each platform can support with the different `max-gateways` option selections. For more information on selecting the maximum number of gateways and how this affects different platforms, see the section “ECMP Hardware Table” in the *NETGEAR 8800 User Manual*.

You must save the configuration and reboot the switch for the new value to take effect. To see the current and configured value, use the following command:

```
show ipconfig
```

Example

The following command changes the maximum number of ECMP gateways per subnet or gateway set to eight:

```
configure iproute sharing max-gateways 8
```

configure irdp

```
configure irdp [multicast | broadcast | <mininterval> <maxinterval> <lifetime> <preference>]
```

Description

Configures the destination address of the router advertisement messages.

Syntax Description

multicast	Specifies multicast setting.
-----------	------------------------------

broadcast	Specifies broadcast setting.
mininterval	Specifies the minimum time between advertisements.
maxinterval	Specifies the maximum time between advertisements. Default is 600.
lifetime	Specifies the lifetime of the advertisement. Default is 1800.
preference	Specifies the router preference level. Default is 0.

Default

Broadcast (255.255.255.255). The default mininterval is 450.

Usage Guidelines

ICMP Router Discovery Protocol allows client machines to determine what default gateway address to use. The switch sends out IP packets at the specified intervals identifying itself as a default router. IRDP enabled client machines use this information to determine which gateway address to use for routing data packets to other networks.

Example

The following command sets the address of the router advertiser messages to multicast:

```
configure irdp multicast
```

configure vlan add secondary-ipaddress

```
configure vlan <vlan_name> add secondary-ipaddress [<ipaddress> {<netmask>} | <ipNetmask>]
```

Description

Configures secondary IP addresses on a VLAN to support multinetting.

Syntax Description

vlan_name	Specifies a VLAN name.
ipaddress	Specifies an IP address.
netmask	Specifies a network mask.
ipNetmask	Specifies an IP address with network mask.

Default

N/A.

Usage Guidelines

Adding a secondary IP address to a VLAN enables multinetting. Secondary addresses are added to support legacy stub IP networks.

Once you have added a secondary IP address to a VLAN, you cannot unconfigure the primary IP address of that VLAN until you delete all the secondary addresses. Delete secondary address with the following command:

```
configure vlan <vlan_name> delete secondary-ipaddress [<ipaddress> | all]
```

Example

The following command configures the VLAN *multi* to support the 10.1.1.0/24 subnet in addition to its primary subnet:

```
configure vlan multi add secondary-ipaddress 10.1.1.1/24
```

configure vlan delete secondary-ipaddress

```
configure vlan <vlan_name> delete secondary-ipaddress [<ipaddress> | all]
```

Description

Removes secondary IP addresses on a VLAN that were added to support multinetting.

Syntax Description

vlan_name	Specifies a VLAN name.
ipaddress	Specifies an IP address.
all	Specifies all secondary IP addresses.

Default

N/A.

Usage Guidelines

Once you have added a secondary IP address to a VLAN, you cannot unconfigure the primary IP address of that VLAN until you delete all the secondary addresses. Use the `all` keyword to delete all the secondary IP addresses from a VLAN.

Example

The following command removes the 10.1.1.0 secondary IP address from the VLAN *multi*:

```
configure vlan multi delete secondary-ipaddress 10.1.1.1
```

configure vlan subvlan

```
configure vlan <vlan_name> [add | delete] subvlan <sub_vlan_name>
```

Description

Adds or deletes a subVLAN to a superVLAN.

Syntax Description

vlan_name	Specifies a superVLAN name.
add	Specifies to add the subVLAN to the superVLAN.
delete	Specifies to delete the subVLAN from the superVLAN.
sub_vlan_name	Specifies a subVLAN name.

Default

N/A.

Usage Guidelines

The following properties apply to VLAN aggregation operation:

- All broadcast and unknown traffic remains local to the subVLAN and does not cross the subVLAN boundary. All traffic within the subVLAN is switched by the subVLAN, allowing traffic separation between subVLANs (while using the same default router address among the subVLANs).
- Hosts can be located on the superVLAN or on subVLANs. Each host can assume any IP address within the address range of the superVLAN router interface. Hosts on the subVLAN are expected to have the same network mask as the superVLAN and have their default router set to the IP address of the superVLAN.
- All IP unicast traffic between subVLANs is routed through the superVLAN. For example, no ICMP redirects are generated for traffic between subVLANs, because the superVLAN is responsible for subVLAN routing. Unicast IP traffic across the subVLANs is facilitated by the automatic addition of an ARP entry (similar to a proxy ARP entry) when a subVLAN is added to a superVLAN. This feature can be disabled for security purposes.

Example

The following command adds the subVLAN vsub1 to the superVLAN vsuper:

```
configure vlan vsuper add subvlan vsub1
```

configure vlan subvlan-address-range

```
configure vlan <vlan_name> subvlan-address-range <ip address1> - <ip address2>
```

Description

Configures subVLAN address ranges on each subVLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

Syntax Description

vlan_name	Specifies a subVLAN name.
ip address1	Specifies an IP address.
ip address2	Specifies another IP address.

Default

N/A.

Usage Guidelines

There is no error checking to prevent the configuration of overlapping subVLAN address ranges between multiple subVLANs. Doing so can result in unexpected behavior of ARP within the superVLAN and associated subVLANs.

Example

The following command configures the subVLAN *vsuper* to prohibit the entry of IP addresses from hosts outside of the configured range of IP addresses:

```
configure vlan vsuper subvlan-address-range 10.1.1.1 - 10.1.1.255
```

configure vlan udp-profile

```
configure vlan <vlan_name> udp-profile [<profilename> | none]
```

Description

Associates a UDP forwarding profile to a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
profilename	Specifies a policy file to use for the UDP forwarding profile.
none	Removes any UDP forwarding profile from the VLAN.

Default

No UDP profiles are associated with the VLAN.

Usage Guidelines

You can apply a UDP forwarding policy only to an L3 VLAN (a VLAN having at least one IP address configured on it). If there is no IP address configured on the VLAN, then the command is rejected.

A UDP forwarding policy must contain only the following attributes. Unrecognized attributes are ignored.

- Match attributes
 - Destination UDP port number (destination-port)
 - Source IP address (source-ipaddress)
- Action modified (set) attributes
 - Destination IP address (destination-ipaddress)
 - VLAN name (vlan)

Policy files used for UDP forwarding are processed differently from standard policy files. Instead of terminating when an entry's match clause becomes true, each entry in the policy file is processed and the corresponding action is taken for each true match clause.

For example, if the following policy file is used as a UDP forwarding profile, any packets destined for UDP port 67 are sent to IP address 20.0.0.5 AND flooded to VLAN *to7*:

```
entry one {
  if match all {
    destination-port 67 ;
  } then {
    destination-ipaddress 20.0.0.5 ;
  }
}

entry two {
  if match all {
    destination-port 67 ;
  } then {
    vlan "to7" ;
  }
}
```

If you include more than one VLAN set attribute or more than one destination-ipaddress set attribute in one policy entry, the last one is accepted and the rest are ignored.

Note: Although the XOS Policy manager allows you to set a range for the destination-port, you should not specify the range for the destination-port attribute in the match clause of the policy statement for the UDP profile. If a destination-port range is configured, the last port in the range is accepted and the rest are ignored.

You can have two valid set statements in each entry of a UDP forwarding policy; one a destination-ipaddress and one a VLAN. The NETGEAR 8800 currently allows a maximum of

eight entries in a UDP forwarding policy, so you can define a maximum of sixteen destinations for one inbound broadcast UDP packet: eight IP addresses and eight VLANs.

Note: It is strongly advised to have no more than eight entries in a UDP forwarding profile. The UDP forwarding module processes those entries even if the entries do not contain any attributes for UDP forwarding. Having more than eight entries drastically reduces system performance. If the inbound UDP traffic rate is very high, having more than eight entries could cause the system to freeze or become locked.

Note: If you rename a VLAN referred to in your UDP forwarding profile, you must manually edit the policy to reflect the new name, and refresh the policy.

You can also validate whether the UDP profile has been successfully associated with the VLAN by using the command `show policy {<policy-name> | detail}`. UDP forwarding is implemented as part of the netTools process, so the command does display netTools as a user of the policy.

Example

The following command associates the UDP forwarding profile *port123_to_corporate* to the VLAN *to-sales*:

```
configure vlan to-sales udp-profile port123_to_corporate
```

disable bootp vlan

```
disable bootp vlan [<vlan> | all]
```

Description

Disables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server.

Syntax Description

vlan	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the generation and processing of BOOTP packets on a VLAN named *accounting*:

```
disable bootp vlan accounting
```

disable bootprelay

```
disable bootprelay [{vlan} [<vlan_name>] | [{vr} <vr_name>] | all [{vr} <vr_name>]]
```

Description

Disables the BOOTP relay function on one or all VLANs for the specified virtual router.

Syntax Description

vlan_name	Specifies a single VLAN on which to disable the BOOTP relay feature.
vr_name	Specifies a single VR on which to disable the BOOTP relay feature.
all	Specifies that BOOTP relay is to be disabled for all VLANs on the specified virtual router.

Default

The BOOTP relay function is disabled on all VLANs and virtual routers.

Usage Guidelines

Because VLAN names are unique on the switch, you can specify only a VLAN name (and omit the VR name) to disable BOOTP relay. When you disable BOOTP relay on a virtual router, BOOTP relay is disabled on all VLANs for that virtual router. If you enter the command without specifying a VLAN or a virtual router, the functionality is disabled for all VLANs in the current VR context.

Examples

The following command disables the forwarding of BOOTP requests on all VLANs in the current VR context:

```
disable bootprelay
```

You can use either of the following commands to disable the forwarding of BOOTP requests on VLAN *unit2*:

```
disable bootprelay unit2
disable bootprelay vlan unit2
```

You can use any one of the following commands to disable the forwarding of BOOTP requests on all VLANs in virtual router *zone3*:

```
disable bootprelay zone3
disable bootprelay vr zone3
disable bootprelay all zone3
disable bootprelay all vr zone3
```

disable icmp address-mask

```
disable icmp address-mask {vlan <name>}
```

Description

Disables the generation of an ICMP address-mask reply on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Disabled.

Usage Guidelines

Disables the generation of an ICMP address-mask reply (type 18, code 0) when an ICMP address mask request is received. The default setting is disabled. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP address-mask reply on VLAN *accounting*:

```
disable icmp address-mask vlan accounting
```

disable icmp parameter-problem

```
disable icmp parameter-problem {vlan <name>}
```

Description

Disables the generation of an ICMP parameter-problem message on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP parameter-problem message on VLAN *accounting*:

```
disable icmp parameter-problem vlan accounting
```

disable icmp port-unreachables

```
disable icmp port-unreachables {vlan <name>}
```

Description

Disables the generation of ICMP port unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of ICMP port unreachable messages (type 3, code 3) when a TCP or UDP request is made to the switch and no application is waiting for the request, or an access

policy denies the request. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables ICMP port unreachable messages on VLAN *accounting*:

```
disable icmp port-unreachables vlan accounting
```

disable icmp redirects

```
disable icmp redirects {vlan <name>}
```

Description

Disables generation of ICMP redirect messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of ICMP redirects (Type 5) to hosts who direct routed traffic to the switch where the switch detects that there is another router in the same subnet with a better route to the destination.

Example

The following command disables ICMP redirects from VLAN *accounting*:

```
disable icmp redirects vlan accounting
```

disable icmp time-exceeded

```
disable icmp time-exceeded {vlan <name>}
```

Description

Disables the generation of ICMP time exceeded messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP time exceeded message (type 11) when the TTL field expires during forwarding. IP multicast packets do not trigger ICMP time exceeded messages. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of ICMP time exceeded messages on VLAN *accounting*:

```
disable icmp time-exceeded vlan accounting
```

disable icmp timestamp

```
disable icmp timestamp {vlan <name>}
```

Description

Disables the generation of an ICMP timestamp response on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP timestamp response on VLAN *accounting*:

```
disable icmp timestamp vlan accounting
```

disable icmp unreachable

```
disable icmp unreachable {vlan <name>}
```

Description

Disables the generation of ICMP unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP timestamp response (type 3, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of ICMP unreachable messages on all VLANs:

```
disable icmp unreachable
```

disable icmp useredirects

```
disable icmp useredirects
```

Description

Disables the modification of route table information when an ICMP redirect message is received.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This option only applies to the switch when the switch is not in routing mode.

If the switch has a route to a destination network, the switch uses that router as the gateway to forward the packets to. If that router knows about a better route to the destination, and the next hop is in the same subnet as the originating router, the second router sends an ICMP redirect message to the first router. If ICMP userredirects is disabled, the switch disregards these messages and continues to send the packets to the second router.

Example

The following command disables the changing of routing table information:

```
disable icmp userredirects
```

disable iparp checking

```
disable iparp {vr <vr_name>} checking
```

Description

Disable checking if the ARP request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

vr_name	Specifies a virtual router.
---------	-----------------------------

Default

Enabled.

Usage Guidelines

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command disables IP ARP checking:

```
disable iparp checking
```

disable iparp refresh

```
disable iparp {vr <vr_name>} refresh
```

Description

Disables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

vr_name	Specifies a virtual router.
---------	-----------------------------

Default

Enabled.

Usage Guidelines

The purpose of disabling ARP refresh is to reduce ARP traffic in a high node count Layer 2 switching only environment.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command disables IP ARP refresh:

```
disable iparp refresh
```

disable ipforwarding

```
disable ipforwarding {broadcast} {vlan <vlan_name>}
```

Description

Disables routing (or routing of broadcasts) for one or all VLANs. If no argument is provided, disables routing for all VLANs.

Syntax Description

broadcast	Specifies broadcast IP forwarding.
vlan_name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

Disabling IP forwarding also disables broadcast forwarding. Broadcast forwarding can be disabled without disabling IP forwarding. When new IP interfaces are added, IP forwarding (and IP broadcast forwarding) is disabled by default.

Other IP related configuration is not affected.

Example

The following command disables forwarding of IP broadcast traffic for a VLAN named *accounting*:

```
disable ipforwarding broadcast vlan accounting
```

disable ip-option loose-source-route

```
disable ip-option loose-source-route
```

Description

Disables processing of the loose source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from forwarding IP packets with the IP option for loose source routing turned on. Packets with the `loose-source-route` option enabled are dropped by the switch.

Example

The following command disables processing of the loose source route IP option:

```
disable ip-option loose-source-route
```

disable ip-option record-route

```
disable ip-option record-route
```

Description

Disables processing of the record route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from adding itself into the IP options header when the record route IP option is enabled in a packet that is transiting the switch.

Example

The following command disables processing of the record route IP option:

```
disable ip-option record-route
```

disable ip-option record-timestamp

```
disable ip-option record-timestamp
```

Description

Disables processing of the record timestamp IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from adding a timestamp into the IP options header when it receives a packet with the record timestamp IP option.

Example

The following command disables processing of the record timestamp IP option:

```
disable ip-option record-timestamp
```

disable ip-option router-alert

```
disable ip-option router-alert
```

Description

Disables processing of the router alert IP option in IPv4 packet headers.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables processing of the router alert IP option:

```
disable ip-option router-alert
```

disable ip-option strict-source-route

```
disable ip-option strict-source-route
```

Description

Disables processing the strict source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from forwarding IP packets that have the strict source routing IP option turned on. The switch drops packets that have the strict source routing IP option enabled.

Example

The following command disables processing of the strict source route IP option:

```
disable ip-option strict-source-route
```

disable iproute compression

```
disable iproute compression {vr <vrname>}
```

Description

Disables IPv4 route compression.

Syntax Description

vrname	Virtual router name for which the IP route compression is being disabled. If the virtual router name is not specified, route compression is disabled for the VR context from which CLI command is being issued.
--------	---

Default

Disabled.

Usage Guidelines

Disables IPv4 route compression for a specified virtual router.

Example

The following example disables IP route compression:

```
disable iproute compression
```

disable iproute sharing

```
disable iproute {ipv4} sharing
```

Description

Disables IPv4 route sharing.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables load sharing for multiple routes:

```
disable iproute sharing
```

disable irdp

```
disable irdp {vlan <name>}
```

Description

Disables the generation of ICMP router advertisement messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Disabled.

Usage Guidelines

If no optional argument is specified, all the IP interfaces are affected.

Example

The following command disables IRDP on VLAN *accounting*:

```
disable irdp vlan accounting
```

disable subvlan-proxy-arp vlan

```
disable subvlan-proxy-arp vlan [<vlan-name> | all]
```

Description

Disables the automatic entry of subVLAN information in the proxy ARP table.

Syntax Description

vlan-name	Specifies a superVLAN name.
all	Specifies all VLANs.

Default

Enabled.

Usage Guidelines

To facilitate communication between subVLANs, by default, an entry is made in the IP ARP table of the superVLAN that performs a proxy ARP function. This allows clients on one subVLAN to communicate with clients on another subVLAN. In certain circumstances, intra-subVLAN communication may not be desired for isolation reasons.

Note: The isolation option works for normal, dynamic, ARP-based client communication.

Example

The following command disables the automatic entry of subVLAN information in the proxy ARP table of the superVLAN *vsuper*.

```
disable subvlan-proxy-arp vlan vsuper
```

disable udp-echo-server

```
disable udp-echo-server {vr <vrid>}
```

Description

Disables UDP echo server support.

Syntax Description

vrid	Specifies a virtual router.
------	-----------------------------

Default

Disabled.

Usage Guidelines

UDP Echo packets are used to measure the transit time for data between the transmitting and receiving end.

Example

The following command disables UDP echo server support:

```
disable udp-echo-server
```

enable bootp vlan

```
enable bootp vlan [<vlan> | all]
```

Description

Enables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server.

Syntax Description

vlan	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the generation and processing of BOOTP packets on a VLAN named *accounting*:

```
enable bootp vlan accounting
```

enable bootprelay

```
enable bootprelay [{vlan} [<vlan_name>] | [{vr} <vr_name>} | all [{vr} <vr_name>}]
```

Description

Enables the BOOTP relay function on one or all VLANs for the specified virtual router.

Syntax Description

vlan_name	Specifies a single VLAN on which to enable the BOOTP relay feature.
vr_name	Specifies a single VR on which to enable the BOOTP relay feature.
all	Specifies that BOOTP relay is to be enabled for all VLANs on the specified virtual router.

Default

The BOOTP relay function is disabled on all VLANs and virtual routers.

Usage Guidelines

Because VLAN names are unique on the switch, you can specify only a VLAN name (and omit the VR name) to enable BOOTP relay on a particular VLAN. When you enable BOOTP relay on a virtual router, BOOTP relay is enabled on all VLANs for that virtual router. If you enter the command without specifying a VLAN or a virtual router, the functionality is enabled for all VLANs in the current VR context.

Examples

The following command enables the forwarding of BOOTP requests for all VLANs in the current VR context:

```
enable bootprelay
```

You can use either of the following commands to enable the forwarding of BOOTP requests for VLAN *client1*:

```
enable bootprelay "client1"
enable bootprelay vlan "client1"
```

You can use any one of the following commands to enable the forwarding of BOOTP requests for all VLANs on VR *zone3*:

```
enable bootprelay zone3
```

```
enable bootprelay vr zone3
enable bootprelay all zone3
enable bootprelay all vr zone3
```

enable icmp address-mask

```
enable icmp address-mask {vlan <name>}
```

Description

Enables the generation of an ICMP address-mask reply on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Disabled.

Usage Guidelines

Enables the generation of an ICMP address-mask reply (type 18, code 0) when an ICMP address mask request is received. The default setting is disabled. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP address-mask reply on VLAN *accounting*:

```
enable icmp address-mask vlan accounting
```

enable icmp parameter-problem

```
enable icmp parameter-problem {vlan <name>}
```

Description

Enables the generation of an ICMP parameter-problem message on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP parameter-problem message on VLAN *accounting*:

```
enable icmp parameter-problem vlan accounting
```

enable icmp port-unreachables

```
enable icmp port-unreachables {vlan <name>}
```

Description

Enables the generation of ICMP port unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of ICMP port unreachable messages (type 3, code 3) when a TCP or UDP request is made to the switch and no application is waiting for the request, or when an access policy denies the request. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables ICMP port unreachable messages on VLAN *accounting*:

```
enable icmp port-unreachables vlan accounting
```

enable icmp redirects

```
enable icmp redirects {vlan <name>}
```

Description

Enables generation of ICMP redirect messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

This option only applies to the switch when the switch is in routing mode.

ICMP redirects are used in the situation where there are multiple routers in the same subnet. If a host sends a packet to its default gateway, the gateway router looks at its route table to find the best route to the destination. If it sees that the best route is through a router in the same subnet as the originating host, the switch sends an ICMP redirect (type 5) message to the host that originated the packet, telling it to use the other router with the better route. The switch also forwards the packet to the destination.

Example

The following command enables the generation of ICMP redirect messages on all VLANs:

```
enable icmp redirects
```

enable icmp time-exceeded

```
enable icmp time-exceeded {vlan <name>}
```

Description

Enables the generation of ICMP time exceeded messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP time exceeded message (type 11) when the TTL field expires during forwarding. IP multicast packets do not trigger ICMP time exceeded messages. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of ICMP time exceeded messages on VLAN *accounting*:

```
enable icmp time-exceeded vlan accounting
```

enable icmp timestamp

```
enable icmp timestamp {vlan <name>}
```

Description

Enables the generation of an ICMP timestamp response on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP timestamp response on VLAN *accounting*:

```
enable icmp timestamp vlan accounting
```

enable icmp unreachable

```
enable icmp unreachable {vlan <name>}
```


Description

Enables the generation of ICMP unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP timestamp response (type 3, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of ICMP unreachable messages on all VLANs:

```
enable icmp unreachable
```

enable icmp useredirects

```
enable icmp useredirects
```

Description

Enables the modification of route table information when an ICMP redirect message is received.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If the switch has a route to a destination network, the switch uses that router as the gateway to forward the packets to. If that router knows about a better route to the destination, and the next hop is in the same subnet as the originating router, the second router sends an ICMP

redirect message to the originating router. If ICMP `useredirects` is enabled, the switch adds a route to the destination network using the third router as the next hop and starts sending the packets to the third router.

Example

The following command enables the modification of route table information:

```
enable icmp useredirects
```

enable iparp checking

```
enable iparp {vr <vr_name>} checking
```

Description

Enables checking if the ARP request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

vr_name	Specifies a virtual router.
---------	-----------------------------

Default

Enabled.

Usage Guidelines

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command enables IP ARP checking:

```
enable iparp checking
```

enable iparp refresh

```
enable iparp {vr <vr_name>} refresh
```

Description

Enables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

vr_name	Specifies a virtual router.
---------	-----------------------------

Default

Enabled.

Usage Guidelines

If ARP refresh is enabled, the switch resends ARP requests for the host at 3/4 of the configured ARP timer value.

For example: If the ARP timeout is set to 20 minutes, the switch attempts to resend an ARP request for the host when the host entry is at 15 minutes. If the host replies, the ARP entry is reset back to 0, and the timer starts again.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command enables IP ARP refresh:

```
enable iparp refresh
```

enable ipforwarding

```
enable ipforwarding {ipv4 | broadcast} {vlan <vlan_name>}
```

Description

Enables IPv4 routing or IPv4 broadcast forwarding for one or all VLANs. If no argument is provided, enables IPv4 routing for all VLANs that have been configured with an IP address on the current virtual router.

Syntax Description

ipv4	Specifies IPv4 forwarding
broadcast	Specifies broadcast IP forwarding.
vlan_name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

IP forwarding must first be enabled before IP broadcast forwarding can be enabled. When new IP interfaces are added, IP forwarding (and IP broadcast forwarding) is disabled by default. Currently, NETGEAR switches only have a single hardware control per VLAN for IP forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv4 forwarding on a VLAN also enables IPv6 hardware forwarding on that VLAN. Future switches may have independent controls per-VLAN for forwarding of IPv4 and IPv6 unicast packets.

The `broadcast` option prompts with a warning message when executed while the IP forwarding on the corresponding VLAN is disabled. The hardware and software are NOT programmed until IP forwarding is enabled on the VLAN.

Example

The following command enables forwarding of IP traffic for all VLANs in the current virtual router context with IP addresses:

```
enable ipforwarding
```

The following command enables forwarding of IP broadcast traffic for a VLAN named *accounting*:

```
enable ipforwarding broadcast vlan accounting
```

enable ip-option record-route

```
enable ip-option record-route
```

Description

Enables processing of the record route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

IP option record-route (IP option 7) means that each router along the path should add its IP address into the options data.

Enabling means that the switch adds itself into the IP options header when the record route IP option is enabled in a packet that is transiting the switch.

Example

The following command enables processing of the record route IP option:

```
enable ip-option record-route
```

enable ip-option record-timestamp

```
enable ip-option record-timestamp
```

Description

Enables processing of the record timestamp IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Enables the switch to use the timestamp IP option (0x44). When the switch receives an IP packet with the timestamp option turned on, it inserts the timestamp into the IP options header before forwarding the packet to the destination.

Example

The following command enables processing of the record timestamp IP option:

```
enable ip-option record-timestamp
```

enable ip-option strict-source-route

```
enable ip-option strict-source-route
```

Description

Enables processing of the strict source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This enables the switch to forward IP packets that have the strict source route IP option (0x89) enabled.

Source routing is used when a sending host specifies the router interfaces that the packet must traverse on its way to its destination.

When strict source routing is used, it means that the packet must use the exact path of routers that lie in the designated router path.

With strict source routing enabled, the switch forwards IP packets with the strict source route option enabled, only if the switch's IP is in the designated list and as long as the next hop in the list is directly attached to one of the router's interfaces.

Example

The following command enables processing of the strict source route IP option:

```
enable ip-option strict-source-route
```

enable ip-option router-alert

```
enable ip-option router-alert
```

Description

Enables processing of the router alert IP option in IPv4 packet headers.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables processing of the router alert IP option:

```
enable ip-option router-alert
```

enable iproute compression

```
enable iproute compression {vr <vrname>}
```

Description

Enables IPv4 route compression.

Syntax Description

vrname	Virtual router name for which the IP route compression is being enabled.
--------	--

Default

Disabled.

Usage Guidelines

Enables IPv4 route compression for the specified virtual router. If the virtual router name is not specified, route compression is enabled for the VR context from which the CLI command is issued.

The command applies a compression algorithm on each of the IP prefixes in the routing table. Essentially, routes with longer network masks might not be necessary if they are a subset of other routes with shorter network masks using the same gateway(s). When IP route compression is enabled, these unnecessary routes are not provided to the Forwarding Information Base (FIB).

Example

The following example enables IP route compression:

```
enable iproute compression
```

enable iproute sharing

```
enable iproute {ipv4} sharing
```

Description

Enables load sharing if multiple routes to the same destination are available. When multiple routes to the same destination are available, load sharing can be enabled to distribute the traffic to multiple destination gateways. Only paths with the same lowest cost are shared.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

IP route sharing allows multiple equal-cost routes to be used concurrently. IP route sharing can be used with static routes or with OSPF or BGP routes. In OSPF or BGP, this capability is referred to as *equal cost multipath* (ECMP) routing.

Configure static routes and OSPF or BGP as you would normally. The NETGEAR 8800 software supports route sharing across up to 8 static routes or ECMP routes for OSPF or BGP. However, on the NETGEAR 8800 family switches, by default, up to 4 routes are supported. To support 2, 4, or 8 routes on these switches, use the following command:

```
configure iproute sharing max-gateways <max_gateways>
```

Using route sharing makes router troubleshooting more difficult because of the complexity in predicting the path over which the traffic travels.

Example

The following command enables load sharing for multiple routes:

```
enable iproute sharing
```

enable irdp

```
enable irdp {vlan <name>}
```

Description

Enables the generation of ICMP router advertisement messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Disabled.

Usage Guidelines

ICMP Router Discovery Protocol allows client machines to determine what default gateway address to use. The switch sends out IP packets at the specified intervals identifying itself as a default router. IRDP enabled client machines use this information to determine which gateway address to use for routing data packets to other networks.

If no optional argument is specified, all the IP interfaces are affected.

Example

The following command enables IRDP on VLAN *accounting*:

```
enable irdp vlan accounting
```

enable subvlan-proxy-arp vlan

```
enable subvlan-proxy-arp vlan [<vlan-name> | all]
```

Description

Enables the automatic entry of subVLAN information in the proxy ARP table.

Syntax Description

vlan-name	Specifies a superVLAN name.
all	Specifies all VLANs.

Default

Enabled.

Usage Guidelines

To facilitate communication between subVLANs, by default, an entry is made in the IP ARP table of the superVLAN that performs a proxy ARP function. This allows clients on one subVLAN to communicate with clients on another subVLAN. In certain circumstances, intra-subVLAN communication may not be desired for isolation reasons.

Note: The isolation option works for normal, dynamic, ARP-based client communication.

Example

The following command enables the automatic entry of subVLAN information in the proxy ARP table of the superVLAN *vsuper*:

```
enable subvlan-proxy-arp vlan vsuper
```

enable udp-echo-server

```
enable udp-echo-server {vr <vrid>}{udp-port <port>}
```

Description

Enables UDP echo server support.

Syntax Description

port	Specifies the UDP port.
vrid	Specifies the virtual router.

Default

Disabled.

Usage Guidelines

UDP Echo packets are used to measure the transit time for data between the transmitting and receiving ends.

Example

The following command enables UDP echo server support:

```
enable udp-echo-server
```

rtlookup

```
rtlookup [<ipaddress> | <ipv6address>] { unicast | multicast | vr <vr_name> }
```

Description

Looks up and displays routes to the specified IP address.

Syntax Description

ipaddress	Specifies an IPv4 address.
ipv6address	Specifies an IPv6 address.
unicast	Displays the routes from the unicast routing table in the current router context.
multicast	Selects the multicast routing table for the search.
vr_name	Displays the available routes in the specified router context.

Default

N/A.

Usage Guidelines

When IP Route sharing is enabled, the `rtlookup` command displays all ECMP routes for the specified IP address.

When IP route sharing is disabled and there are multiple ECMP routes for the specified IP address, the `rtlookup` command displays only one route, which is the route with lowest value gateway IP address.

Example

The following command looks up IP address 10.0.0.0 in the *VR-Mgmt* router and displays the available routes:

```
BD-12804.4 # rtlookup 66.6.6.6
Ori Destination      Gateway      Mtr  Flags      VLAN      Duration
#s  66.6.6.6/32      80.1.10.58   1    UG---S-um--f v8      0d:0h:18m:58s
```

```
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMP, (el) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (il) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
```

(#) Preferred unicast and multicast route

Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route

(L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast

(P) LPM-routing, (R) Modified, (S) Static, (s) Static LSP

(T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up

(f) Provided to FIB (c) Compressed Route

show bootprelay

show bootprelay

Description

Displays the DHCP/BOOTP relay statistics and the configuration for the virtual routers.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

The fields displayed in the *DHCP Information Option 82* section depend on the configuration defined by the `configure bootprelay dhcp-agent information policy [drop | keep | replace]` command. If the policy configured is `keep`, the *Requests unmodified* counter appears. If the policy configured is `replace`, the *Requests replaced* counter appears. And if the `drop` policy is configured, the *Requests dropped* counter appears.

The *Opt82 added to Requests* counter indicates the number of DHCP requests to which the bootprelay agent (the switch) has added its own option 82 information.

Example

The following example displays the DHCP/BOOTP relay statistics for existing virtual routers:

```
Switch.1 # show bootprelay
Bootprelay : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Option : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Check : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Policy : Replace
DHCP Relay Agent Information Remote-ID : "default"

Bootprelay servers for virtual router "VR-Default":
  Destination: 10.127.8.1

DHCP/BOOTP relay statistics for virtual router "VR-Default"
```

```

Received from client =          2  Received from server =          2
Requests relayed     =          2  Responses relayed     =          2
DHCP Discover        =          1  DHCP Offer         =          1
DHCP Request         =          1  DHCP Ack           =          1
DHCP Decline         =          0  DHCP NACK          =          0
DHCP Release         =          0
DHCP Inform          =          0

```

DHCP Information Option 82 packets statistics for virtual router "VR-Default"

```

Received from client   =          0  Received from server =          2
Requests replaced      =          0  Responses dropped    =          0
Opt82 added to Requests =          2

```

Note: Default Remote-ID : System MAC Address

show bootprelay configuration

```
show bootprelay configuration [{vlan} <vlan_name> | {vr} <vr_name>]
```

Description

Displays the enabled/disabled configuration of BOOTP relay on one or all VLANs for the specified VR.

Syntax Description

vlan_name	Specifies a single VLAN for which to display BOOTP relay configuration information.
vr_name	Specifies a single VR for which to display BOOTP relay configuration information.

Default

None.

Usage Guidelines

If a VR is not specified, this command displays the specified VLANs for the current VR context.

Examples

The following example displays the BOOTP relay configuration for all VLANs on the *VR-Default* virtual router:

```

Switch.88 # show bootprelay configuration vr "VR-Default"
BOOTP Relay : Enabled on virtual router "VR-Default"

```

VLAN	BOOTP Relay
-----	-----
Default	Disabled
client1	Enabled
serv	Enabled

The following example displays the BOOTP relay configuration for all VLANs in the current VR context:

```
Switch.95 # show bootprelay configuration
BOOTP Relay : Enabled on virtual router "VR-Default"
```

VLAN	BOOTP Relay
-----	-----
Default	Disabled
client1	Disabled
serv	Disabled

The following example displays the BOOTP relay configuration for VLAN *client1*:

```
Switch.87 # show bootprelay configuration vlan "client1"
BOOTP Relay : Enabled on virtual router "VR-Default"
```

VLAN	BOOTP Relay
-----	-----
client1	Disabled

show bootprelay dhcp-agent information circuit-id port-information

```
show bootprelay dhcp-agent information circuit-id port-information ports all
```

Description

Displays the circuit ID sub-option that identifies the port for an incoming DHCP request.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the circuit ID port_info value for all ports:

```
Switch.12 # show bootprelay dhcp-agent information circuit-id port-information ports all
```

Port	Circuit-ID Port information string
----	-----
1	1001
2	1002
3	netgear1
4	1004
5	1005
6	1006
7	1007
8	1008
9	1009
10	1010
:	
:	
11	1011
12	1012
:	
:	
48	1048
49	1049
50	1050

Note: The full Circuit ID string has the form '<Vlan Info>--<Port Info>'

show bootprelay dhcp-agent information circuit-id vlan-information

```
show bootprelay dhcp-agent information circuit-id vlan-information
```

Description

Displays the circuit ID sub-options that identify the VLANs on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the circuit ID vlan_info for all VLANs:

```
X250e-48t.8 # show bootprelay dhcp-agent information circuit-id vlan-information
Vlan          Circuit-ID vlan information string
----          -
Default       1
Mgmt          4095
v1            4094
v2            netgear123
Note: The full Circuit ID string has the form '<Vlan Info>--<Port Info>'
```

show iparp

```
show iparp {<ip_addr> | <mac> | vlan <vlan_name> | permanent} {vr <vr_name>}
```

Description

Displays the IP Address Resolution Protocol (ARP) table. You can filter the display by IP address, MAC address, VLAN, or permanent entries.

Syntax Description

ip_addr	Specifies an IP address.
mac	Specifies a MAC address.
vlan_name	Specifies a VLAN name.
permanent	Specifies permanent entries.
vr_name	Specifies a virtual router.

Default

Show all entries, except for proxy entries.

Usage Guidelines

Displays the IP ARP table, including:

- IP address
- MAC address
- Aging timer value
- VLAN name, VLAN ID and port number
- Flags

If you do not specify a virtual router, the command applies to the current virtual router.

The show output displays the following information:

Dynamic entries	The number of dynamic (learned ARP) entries in the table.
Static entries	The number of configured (static ARP) entries in the table.

Pending entries	The number of sent ARP requests that have not yet received a response.
In Request	The number of ARP request packets received (by this virtual router).
In Response	The number of ARP reply packets received (by this virtual router).
Out Request	The number of ARP request packets sent (by this virtual router).
Out Response	The number of ARP reply packets sent (by this virtual router).
Failed requests	The number of failed ARP requests sent (by this virtual router).
Proxy Answered	The number of ARP requests answered by the ARP proxy.
RX Error	The number of incorrect ARP request and reply packets received. The malformed packets include the following errors: incorrect ARP op code, hardware address type is not Ethernet, the protocol address is not IP, and similar errors.
Dup IP Addr	IP addresses that have been used by other hosts on the network.
Rejected Count	The number of rejected ARP request packets.
Rejected IP	The source address for the last rejected ARP request. An example reason for an ARP request packet to be rejected is if the source address of the packet is not in the subnet.
Rejected Port	The port on which the last rejected ARP request packet arrived.
Rejected I/F	The VLAN on which the last rejected ARP request packet arrived.
Max ARP entries	Maximum ARP table size for the virtual router (each virtual router has its own ARP table).
Max ARP pending entries	Maximum number of incomplete (pending) ARP entries allowed in the table.
ARP address check	Whether IP ARP checking is enabled or disabled. IP ARP checking verifies if the ARP request's source address is in the receiving interface's subnet.
ARP refresh	Whether ARP refresh is enabled or disabled. ARP refresh is performed when an ARP entry's age is three-fourths of the timeout value.
Timeout	Timeout value for a dynamic (learned) ARP entry.

Example

The following command displays the IP ARP table for the current virtual router:

```
show iparp
```

The following is sample output for the command:

VR	Destination	Mac	Age	Static	VLAN	VID	Port
VR-Default	10.10.10.6	00:04:96:1f:a5:71	8	NO	bluered	4092	1
VR-Default	10.128.32.1	00:01:30:ba:6a:a0	0	NO	Default	4095	
VR-Default	10.128.32.2	00:01:03:1c:ae:b0	5	NO	Default	4095	
VR-Default	10.128.32.4	00:d0:59:17:74:83	3	NO	Default	4095	
VR-Default	10.128.32.5	00:02:a5:c2:5c:dd	0	NO	Default	4095	
VR-Default	10.128.32.6	00:12:3f:1c:f8:fb	5	NO	Default	4095	

VR-Default	10.128.32.7	00:11:11:80:9c:b9	7	NO	Default	4095
VR-Default	10.128.32.8	00:11:43:53:8e:f1	0	NO	Default	4095
VR-Default	10.128.32.9	00:02:a5:bf:ac:70	7	NO	Default	4095
VR-Default	10.128.32.10	00:11:43:44:18:68	10	NO	Default	4095
VR-Default	10.128.32.11	00:12:3f:1c:e9:f2	0	NO	Default	4095
VR-Default	10.128.32.12	00:02:a5:bf:af:79	8	NO	Default	4095
VR-Default	10.128.32.13	00:11:43:40:89:91	0	NO	Default	4095
VR-Default	10.128.32.16	00:0f:1f:c9:2d:80	2	NO	Default	4095
VR-Default	10.128.32.17	00:06:5b:b1:6a:91	1	NO	Default	4095
VR-Default	10.128.32.19	00:11:43:3a:96:1d	10	NO	Default	4095
VR-Default	10.128.32.20	00:08:02:d5:c5:b7	6	NO	Default	4095
VR-Default	10.128.32.24	00:12:3f:0a:44:92	14	NO	Default	4095
VR-Default	10.128.32.26	00:50:04:ad:36:5e	6	NO	Default	4095
VR-Default	10.128.32.30	00:b0:d0:23:f2:9a	11	NO	Default	4095
VR-Default	10.128.32.54	00:b0:d0:59:e4:e2	6	NO	Default	4095
VR-Default	10.128.32.55	00:a0:c9:0c:41:de	3	NO	Default	4095
VR-Default	10.128.32.59	00:b0:d0:7c:d6:07	14	NO	Default	4095
VR-Default	10.128.32.99	00:04:96:05:00:03	13	NO	Default	4095
VR-Default	10.128.32.101	00:04:96:1f:a8:48	0	NO	Default	4095
VR-Default	10.128.32.104	00:30:48:41:ed:45	0	NO	Default	4095
VR-Default	10.128.32.105	00:30:48:41:ed:97	0	NO	Default	4095
VR-Default	10.128.32.106	00:01:30:23:c1:00	0	NO	Default	4095
VR-Default	10.128.32.108	00:04:96:1f:a5:71	0	NO	Default	4095
VR-Default	10.128.32.116	00:04:96:1f:a4:0e	0	NO	Default	4095

Dynamic Entries :	1	Static Entries :	0
Pending Entries :	0		
In Request :	111	In Response :	3
Out Request :	110	Out Response :	111
Failed Requests :	0		
Proxy Answered :	0		
Rx Error :	0	Dup IP Addr :	0.0.0.0
Rejected Count :		Rejected IP :	
Rejected Port :		Rejected I/F :	
Max ARP entries :	4096	Max ARP pending entries :	256
ARP address check:	Enabled	ARP refresh :	Enabled
Timeout :	20 minutes	ARP Sender-Mac Learning:	Request and Reply

show iparp proxy

```
show iparp proxy [{<ipNetmask> | <ip_addr> <mask>}] {vr <vr_name>}
```

Description

Displays the proxy ARP table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_address	Specifies an IP address.
mask	Specifies a subnet mask.
vr_name	Specifies a virtual router.

Default

N/A.

Usage Guidelines

If no argument is specified, then all proxy ARP entries are displayed.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command displays the proxy ARP table:

```
show iparp proxy 10.1.1.5/24
```

show iparp stats

```
show iparp stats [[ <vr_name> | vr {all | <vr_name>} ] {no-refresh} | {vr} summary]
show iparp stats [vlan {all {vr <vr_name>}} | {vlan} <vlan_name>] {no-refresh}
show iparp stats ports {all | <port_list>} {no-refresh}
```

Description

Displays the IP ARP statistics for one or more virtual routers, VLANs, or ports.

Syntax Description

vr_name	Specifies a virtual router for which to display statistics.
vlan_name	Specifies a VLAN for which to display statistics.
port_list	Specifies a list of ports for which to display statistics.
no-refresh	Requests a static display for statistics.

Default

Shows all VLAN ARP statistics in a dynamic display.

Usage Guidelines

VLAN statistics and totals are displayed for a single VR. When you display IP ARP statistics for one or all VLANs, the display includes the specified VLANs for the specified VR. If you do not specify a VR for a VLAN report, the display includes the specified VLANs for the current VR context.

Examples

The following command displays ARP table statistics for all virtual routers:

```
Switch.1 # show iparp stats vr all
```

```
IP ARP VR Statistics                                     Wed Apr 07 15:30:49 2010
  ARP Total      Dynamic      Static      Pending      Unneeded      Failed (Rejected)
=====
VR-Default
      96          89          5          0          0          2          0
VR-Mgmt
      4           2          2          0          0          0          0
VR-SV_PPPOE
     287         286          1          0          0          0      2785
VR-NV_PPPOE
      19          19          0          0          0          0          0
chicago2
      50          44          5          0          1          0          0

Total for all VRs
     456         440         13          0          1          2      2785
=====
U->page up  D->page down  ESC->exit
```

The following command displays ARP table statistics for all VLANs in the current VR context:

```
Switch.2 # show iparp stats vlan all
```

```
IP ARP VLAN Statistics                                   Wed Apr 07 15:30:49 2010
VLAN                                     ARP Total      Dynamic      Static
=====
VLAN_06-AAR                             94             89           5
VLAN_07-AAR                            122            121           1
VLAN_02-BOT                             43             42           1
```

```

=====
Totals for VR U3c-South.                      Total Entries :          455
  Dynamic :           440      Static   :           13   Pending   :           0
  Failed   :             2      Unneeded :            0   (Rejected):        5639
Last Rejected ARP :
  IP: 10.66.118.243      Port: 1:23      Vlan: VLAN_02-BOT
=====
U->page up  D->page down  ESC->exit

```

The following command displays ARP table statistics for ports 1:1 to 1:17:

```
Switch.3 # show iparp stats ports 1:1-1:17
```

```

IP ARP Port Statistics                               Wed Apr 07 15:30:49 2010
Port           Link State           ARP Total           Dynamic           Static
=====
1:1             A                   94                  89                5
1:2             A                   37                  37                0
1:3             A                  122                 121               1
1:4             R                   0                   0                0
1:5             R                   0                   0                0
1:6             A                   43                  43                0
1:7             A                  118                 118               0
1:8             R                   0                   0                0
1:9             R                   0                   0                0
1:10            A                   8                   8                0
1:11            A                   8                   6                2
1:12            A                  41                  41                0
1:13            A                   17                  17                0
1:14            R                   0                   0                0
1:15            R                   0                   0                0
1:16            A                   8                   8                0
1:17            A                   8                   6                2
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
U->page up  D->page down  ESC->exit

```

show iparp security

```
show iparp security {{vlan} <vlan_name>}
```

Description

Displays the IP ARP security violation information for one or all VLANs.

Syntax Description

vlan_name	Specifies a single VLAN for which to display security violation information.
-----------	--

Default

Shows security violation information for all VLANs except *Mgmt*.

Usage Guidelines

None.

Examples

The following command displays IP ARP security violation information for all VLANs:

```
Switch.4 # show iparp security
```

```

                                     Most Recent Violation
                                     =====
Vlan          Security  Violations  Type  IP address      MAC      Port
=====
Default       ----
test          ----
```

Security Setting: (G) Gratuitous ARP Protection

Violation Type : (g) Gratuitous ARP Violation

The following command displays IP ARP security violation information for VLAN *Default*:

```
Switch.5 # show iparp security "Default"
```

```

                                     Most Recent Violation
                                     =====
Vlan          Security  Violations  Type  IP address      MAC      Port
=====
Default       ----
```

Security Setting: (G) Gratuitous ARP Protection

Violation Type : (g) Gratuitous ARP Violation

show ipconfig

```
show ipconfig {ipv4} {vlan <vlan_name>}
```

Description

Displays configuration information for one or more VLANs.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

If no VLAN information is specified, then global IP configuration is displayed. Otherwise, specific VLAN information is displayed.

Example

The following example displays global IP configuration information and summary information on VLANs for a NETGEAR 8800 series switch:

```
XCM8806.1 # show ipconfig vlan accounting
```

```
Router Interface on VLAN alan1 is enabled and up.
inet 192.18.2.1/24 broadcast 192.18.2.255 Mtu 1500
```

Flags:

```
AddrMaskRly NO   BOOTP Host NO   DirBcstHwFwd NO   Fwd Bcast NO
IgnoreBcast NO   IP Fwding YES   IPmc Fwd YES     Multinetted VLAN NO
IRDP Advert NO   SendParam YES   SendPortUn YES   Send Redir YES
SendTimxceed YES SendUnreach YES   TimeStampRly NO   VRRP NO
Unicast RPF NO
Minimum Interval: 450      Lifetime: 1800 Preference: 0
```

show iproute

```
show iproute {ipv4} {priority | vlan <vlan_name> | permanent | <ip_address> <netmask> |
summary} {multicast | unicast} {vr <vrname>}}
```

Description

Displays the contents of the IP routing table or the route origin priority.

Syntax Description

priority	Displays the priority values for each route origin type.
vlan_name	Specifies a VLAN name.
permanent	Specifies permanent routing.
ip_address	Specifies an IP address.

netmask	Specifies a subnet mask.
vrname	Specifies a virtual router.

Default

N/A.

Usage Guidelines

A `c` flag in the Flags column indicates a compressed route resulting from enabling compression using the `enable iproute compression` command. The total number of compressed routes is also shown.

All routes that are provided to the FIB display the `f` flag.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command example displays detailed information about all IP routing:

```
Switch.3 # show iproute
```

Ori	Destination	Gateway	Mtr	Flags	VLAN	Duration
d	2.2.0.0/16	2.2.2.3	1	-----um---	v2	2d:10h:17m:41s
#d	3.2.2.0/24	3.2.2.23	1	U-----um--f	jim_igmp	1d:19h:49m:49s
d	3.3.3.0/24	3.3.3.1	1	-----um---	v3	2d:10h:17m:49s
d	4.4.4.0/24	4.4.4.2	1	-----um---	v4	2d:10h:17m:5s

```
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
              (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
              (e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
              (is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
              (mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
              (oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
              (r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
              (*) Preferred unicast route (@) Preferred multicast route
              (#) Preferred unicast and multicast route
```

```
Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
        (L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
        (P) LPM-routing, (R) Modified, (S) Static, (s) Static LSP
        (T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
        (f) Provided to FIB (c) Compressed Route
```

```
Mask distribution:
```

```
    1 routes at length 16          3 routes at length 24
```

Route Origin distribution:

4 routes from Direct

Total number of routes = 4

Total number of compressed routes = 0

The following command displays the IP route origin priority:

```
*XCM8806.11 # show iproute priority
```

```
Direct          10
MPLS            20
Blackhole       50
Static          1100
ICMP            1200
EBGP            1700
IBGP            1900
OSPFIntra       2200
OSPFInter       2300
Isis            2350
IsisL1          2360
IsisL2          2370
RIP             2400
OSPFAsExt       3100
OSPFExt1        3200
OSPFExt2        3300
IsisL1Ext       3400
IsisL2Ext       3500
Bootp           5000
```

show iproute origin

```
show iproute origin [bgp | blackhole | bootp | direct | ebgp | embgp | ibgp | icmp | imbgp |
mbgp | ospf | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static ] {unicast}
{vr <vrname>}
```

Description

Displays the contents of the IP routing table for routes with the specified origin.

Syntax Description

bgp	Specifies BGP routes.
blackhole	Specifies blackhole routes.
bootp	Specifies BOOTP routes.
direct	Specifies direct routes.
ebgp	Specifies E-BGP routes.

embgp	Specifies EMBGP routes.
ibgp	Specifies I-BGP routes.
icmp	Specifies ICMP routes.
imbgp	Specifies IMBGP routes.
mbgp	Specifies MBGP routes.
ospf	Specifies OSPF routes.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPFInter routing.
ospf-intra	Specifies OSPFIntra routing.
rip	Specifies RIP routes.
static	Specifies static routes.
unicast	Displays unicast routes with the specified origin.
vrname	Specifies a virtual router.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays all the BGP routes:

```
show iproute origin bgp
```

show iproute reserved-entries

```
show iproute reserved-entries {slot <slot_num>}
```

Description

Displays the configured number of IPv4 and IPv6 routes reserved in the Longest Prefix Match (LPM) hardware table.

Syntax Description

slot_num	This option displays the reservations for the specified slot.
----------	---

Default

N/A.

Usage Guidelines

The IPv4 Routes column in the command output shows whether IPv4 routes are stored in internal or external LPM tables.

Use the following command to modify the configuration that the `show iproute reserved-entries` command displays:

```
configure iproute reserved-entries [ <num_routes_needed> | maximum | default ] slot [all  
| <slot_num>]
```

Example

The following command displays the reserved space for IP routes:

```
(Demo) XCM8806.6 # show iproute reserved-entries
```

Slot	Type	IPv4 Routes	# Reserved Routes IPv4 (or IPv6)	Minimum # IPv4 Hosts
1	XCM8824F	Internal	12240 (6120) [default]	16
2	XCM8824F	Internal	12240 (6120) [default]	16
3	XCM888F	Internal	12240 (6120) [default]	16
4				
5	XCM8808X	Internal	12240 (6120) [default]	16
6	XCM8848T(P)	Internal	12240 (6120) [default]	16

Maximum supported # IPv4 (or IPv6) Reserved Routes:

```
"XCM"-series" Internal 12256 ( 6128)
```

Note: IPv4 Hosts can occupy unused HW Route table space,
unless # Reserved Routes is "maximum".

show iproute reserved-entries statistics

```
show iproute reserved-entries statistics { slot <slot_num> }
```

Description

Displays the current usage statistics of the Longest Prefix Match (LPM) hardware table and the Layer 3 hardware hash table by resource type.

Syntax Description

slot_num	For NETGEAR 8800 series switches, this option displays the statistics for the specified slot.
----------	---

Default

N/A.

Usage Guidelines

This command shows the current number of IP routes and local and remote IPv4 hosts in the LPM hardware table. It also shows the number of IPv4 unicast, multicast, and IPv6 unicast entries in the Layer 3 hardware hash table. The theoretical maximums for each individual resource type are shown at the bottom of the output. These maximum values cannot all be achieved simultaneously, and individual values might not be reached depending on the addresses or routes in use.

The NETGEAR 8800 software supports the coexistence of higher- and lower-capacity hardware in the same NETGEAR 8800 chassis. To allow for coexistence and increased hardware forwarding, when the number of IPv4 routes exceeds 25,000, the lower-capacity hardware automatically transitions from using LPM routing to forwarding of individual remote hosts, also known as IP Forwarding Database (IP FDB) mode. Higher-capacity hardware continues using LPM routing. Lower capacity hardware operating in IP FDB mode will be indicated with a *d* flag in the output of `show iproute reserved-entries statistics` command, indicating that only *direct* routes are installed. For more information, see the section “Coexistence of Higher- and Lower-Capacity Hardware” in the *NETGEAR 8800 User Manual*.

Example

The following command displays usage statistics for the LPM and Layer 3 hardware tables:

```
(Demo) XCM8806.6 # show iproute reserved-entries statistics
      |-----In HW Route Table-----| |--In HW L3 Hash Table--|
      # Used Routes # IPv4 Hosts IPv4 IPv4 IPv6 IPv4
Slot Type          IPv4   IPv6  Local Remote Local Rem.  Loc. MCast
-----
1   XCM8824F          18    0    0    0    0    0    0    0
2   XCM8824F          18    0    0    0    0    0    0    0
3   XCM888F           18    0    0    0    0    0    0    0
4                                     -    -    -    -    -    -
5   XCM8808X          18    0    0    0    0    0    0    0
6   XCM8848T(P)       18    0    0    0    0    0    0    0

Theoretical maximum for each resource type:
      0 0 0 0 0 0 0 0 * 0
      0 0 0 0 0 0 0 0 * 0
"XCM"-series 12256 6128 8189 12288 8189 8192 4096 *6000
      0 0 0 0 0 0 0 0 * 0
      0 0 0 0 0 0 0 0 * 0
```

Flags: (!) Indicates all reserved route entries in use.

(d) Indicates only direct IPv4 routes are installed.

(>) Some IPv6 routes with mask > 64 bits are installed and do not use

entries in HW Route Table.
 (*) Assumes IP Multicast compression is on.

show ipstats

```
show ipstats {ipv4} {vlan <name> | vr <vrname>}
```

Description

Displays IP statistics for the switch CPU or for a particular VLAN.

Syntax Description

name	Specifies a VLAN name.
vrname	Specifies a virtual router.

Default

N/A.

Usage Guidelines

This command only shows statistics of the CPU-handled packets. Not all packets are handled by the CPU.

If you do not specify a virtual router, the command applies to the current virtual router.

The fields displayed in the `show ipstats` command are defined in [Table 21](#) through [Table 24](#).

Table 21. Global IP Statistics Field Definitions

Field	Definition
InReceives	Total number of incoming IP packets processed by the CPU.
InUnicast	Total number of unicast IP packets processed by the CPU.
InBcast	Total number of broadcast IP packets processed by the CPU.
InMcast	Total number of multicast IP packets processed by the CPU.
InHdrEr	Total number of packets with an IP Header Error forwarded to the CPU.
Bad vers	Total number of packets with a version other than IPv4 in the IP version field.
Bad chksum	Total number of packets with a bad IP checksum forwarded to the CPU.
Short pkt	IP packets that are too short.
Short hdr	IP packets with a header that is too short.
Bad hdrlen	IP packets with a header length that is less than the length specified.
Bad length	IP packets with a length less than that of the header.

Table 21. Global IP Statistics Field Definitions (Continued)

Field	Definition
InDelivers	IP packets passed to upper layer protocols.
Bad Proto	IP packets with unknown (not standard) upper layer protocol.
OutRequest	IP packets sent from upper layers to the IP stack.
OutDiscard	IP packets that are discarded due to lack of buffer space or the router interface being down, or broadcast packets with broadcast forwarding disabled.
OutNoRoute	IP packets with no route to the destination.
Forwards	ForwardOK and Fwd Err aggregate count.
ForwardOK	Total number of IP packets forwarded correctly.
Fwd Err	Total number of IP packets that cannot be forwarded.
NoFwding	Aggregate number of IP packets not forwarded due to errors.
Redirects	IP packets forwarded on the same network.
No route	Not used.
Bad TTL	IP packets with a bad time-to-live.
Bad MC TTL	IP packets with a bad multicast time-to-live.
Bad IPdest	IP packets with an address that does not comply with the IPv4 standard.
Blackhole	IP packets with a destination that is a blackhole entry.
Output err	Not used. This is the same as Fwd Err.
MartianSrc	IP packets with an invalid source address.

Table 22. Global ICMP Statistics Field Definitions

Field	Definition
OutResp	Echo replies sent from the CPU.
OutError	Redirect from broadcast or multicast source addresses.
InBadcode	Incoming ICMP packets with an invalid CODE value.
InTooshort	Incoming ICMP packets that are too short.
Bad chksum	Incoming ICMP packets with checksum errors.
In Badlen	Incoming ICMP packets with length errors.
echo reply (In/Out):	ICMP “echo reply” packets that are received and transmitted.
destination unreachable (In/Out):	ICMP packets with destination unreachable that are received and transmitted.

Table 22. Global ICMP Statistics Field Definitions (Continued)

Field	Definition
port unreachable (In/Out):	ICMP packets with port unreachable that are received and transmitted.
echo (In/Out):	ICMP echo packets that are received and transmitted.

Table 23. Global IGMP Statistics Field Definitions

Field	Definition
Out Query	Number of IGMP query messages sent by the router.
Out Report	Number of reports sent on an active multicast route interface for reserved multicast addresses and for regular IGMP reports forwarded by the query router.
Out Leave	Number of IGMP out leave messages forwarded for IP multicast router interfaces.
In Query	Number of IGMP query messages received.
In Report	Number of IGMP report messages received (mostly from hosts).
In Leave	Number of IGMP leave messages received (mostly from hosts).
In Error	Number of IGMP packets with bad header fields or checksum failures.

Table 24. Router Interface Statistics Field Definitions

Field	Definition
Packets IN/OUT	Total number of IP packets received or transmitted on a VLAN router interface.
Octets IN/OUT	Total number of octets received or transmitted on a VLAN router interface.
Mcast packets IN/OUT	Total number of multicast packets received or transmitted on a VLAN router interface.
Bcast packets IN/OUT	Total number of broadcast packets received or transmitted on a VLAN router interface.
Errors IN/OUT	Total number of IP packets with errors received or transmitted on a VLAN router interface.
Discards IN/OUT	Total number of IP packets that cannot travel up to the CPU due to lack of buffer space.
Unknown Protocols IN/OUT	Total number of IP packets with unknown upper layer protocols received by the router interface.

Example

The following command displays IP statistics for the VLAN *accounting*:

```
show ipstats vlan accounting
```

show udp-profile

```
show udp-profile {vlan <vlan-name> | {policy} <policy-name>}
```

Description

Displays UDP forwarding profiles.

Syntax Description

vlan-name	Specifies a VLAN.
policy-name	Specifies a UDP forwarding profile.

Default

If no VLAN or policy is specified, all configured profiles are displayed.

Usage Guidelines

UDP profiles can also be displayed by using the policy manager command `show policy {<policy-name> | detail}`. However, the format of the policy display is different than that for this command.

Example

The following command displays all the configured UDP forwarding profiles on the switch:

```
show udp-profile
```

The following is sample output:

```
UDP Profile Name: move_to7
  Number of datagram forwarded: 181
    Dest UDP Port: 67 Fwd to IP Addr: 20.0.0.5
    Dest UDP Port: 67 Fwd to VLAN: to7
  Applied to incoming traffic on VLANs:
    to-mariner
```

unconfigure bootprelay dhcp-agent information check

```
unconfigure bootprelay dhcp-agent information check
```

Description

Disables Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to disable the switch from preventing DHCP reply packets with invalid or missing relay agent options from being forwarded to the client.

To enable this check, use the following command:

```
configure bootprelay dhcp-agent information check
```

Example

The following command disables the DHCP relay agent option check:

```
unconfigure bootprelay dhcp-agent information check
```

unconfigure bootprelay dhcp-agent information circuit-id port-information

```
unconfigure bootprelay dhcp-agent information circuit-id port-information ports [<portlist> | all]
```

Description

Configures the circuit ID sub-option that identifies the specified ports to use the default value.

Syntax Description

portlist	Specifies a list of one or more ports that are to be configured to use the default value.
all	Specifies that all ports are to be configured to use the default value.

Default

The port_info is encoded as ((slot_number * 1000) + port_number). For example, if the DHCP request is received on port 3:12, the default circuit ID port_info value is 3012. On non-slot-based switches, the default circuit ID port_info value is simply the port number.

Usage Guidelines

None.

Example

The following command configures port 1:3 to use the default circuit ID port information value:

```
unconfigure bootprelay dhcp-agent information circuit-id port-information ports 1:3
```

unconfigure bootprelay dhcp-agent information circuit-id vlan-information

```
unconfigure bootprelay dhcp-agent information circuit-id vlan-information {vlan}
[<vlan_name>|all]
```

Description

Configures the circuit ID sub-option that identifies the specified VLANs to use the default value.

Syntax Description

vlan_name	Names a VLAN to be configured to use the default value.
all	Specifies that all VLANs are to be configured to use the default value.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures VLAN *blue* to use the default VLAN information for the circuit ID sub-option:

```
unconfigure bootprelay dhcp-agent information circuit-id vlan-information blue
```

unconfigure bootprelay dhcp-agent information option

```
unconfigure bootprelay dhcp-agent information option
```

Description

Disables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

To enable the DHCP relay agent option (option 82), use the following command:

```
configure bootprelay dhcp-agent information option
```

Example

The following command disables the DHCP relay agent option:

```
unconfigure bootprelay dhcp-agent information option
```

unconfigure bootprelay dhcp-agent information policy

```
unconfigure bootprelay dhcp-agent information policy
```

Description

Unconfigures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

This command has no arguments or variables.

Default

Replace.

Usage Guidelines

Use this command to unconfigure the policy for the relay agent.

Example

The following command unconfigures the DHCP relay agent option 82 policy:

```
unconfigure bootprelay dhcp-agent information policy
```

unconfigure bootprelay dhcp-agent information remote-id

```
unconfigure bootprelay dhcp-agent information remote-id {vr <vrid>}
```

Description

Configures the remote ID sub-option to the default value.

Syntax Description

vrid	Specifies the VR on which to configure the remote ID sub-option to the default value.
------	---

Default

The switch MAC address.

Usage Guidelines

None.

Example

The following command configures the remote ID sub-option to use the default value on the current VR:

```
configure bootprelay dhcp-agent information remote-id
```

unconfigure icmp

```
unconfigure icmp
```

Description

Resets all ICMP settings to the default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all ICMP settings to the default values.

```
unconfigure icmp
```

unconfigure iparp

```
unconfigure iparp
```

Description

Resets the following to their default values:

- IP ARP timeout
- Maximum ARP entries
- Maximum ARP pending entries
- ARP checking
- ARP refresh

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets IP ARP timeout to its default value:

```
unconfigure iparp
```

unconfigure iproute priority

```
unconfigure iproute {ipv4} priority [all | blackhole | bootp | ebgp | ibgp | icmp |
ospf-as-external | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static] {vr
<vrname>}
```

Description

Unconfigures the priority for all IP routes from one or all route origin types.

Syntax Description

all	Specifies all route origins.
blackhole	Specifies the blackhole route.
bootp	Specifies BOOTP.
ebgp	Specifies E-BGP routes
ibgp	Specifies I-BGP routes
icmp	Specifies ICMP.

ospf-as-external	Specifies OSPF as External routing.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPFInter routing.
ospf-intra	Specifies OSPFIntra routing.
rip	Specifies RIP.
static	Specifies static routes.
vrname	Specifies a virtual router name.

Default

N/A

Usage Guidelines

Table 25 lists the default priorities that apply after you enter this command.

Table 25. Default Route Priorities

Route Origin	Priority
Direct	10
MPLS	20
Blackhole	50
Static	1100
ICMP	1200
EBGP	1700
IBGP	1900
OSPFIntra	2200
OSPFInter	2300
IS-IS	2350
IS-IS L1	2360
IS-IS L2	2370
RIP	2400
OSPFAExt	3100
OSPF External 1	3200
OSPF External 2	3300

Table 25. Default Route Priorities (Continued)

Route Origin	Priority
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500
BOOTP	5000

Example

The following command returns the IP route priority for all route origins to the default values:

```
unconfigure iproute priority all
```

unconfigure irdp

```
unconfigure irdp
```

Description

Resets all router advertisement settings to the default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all router advertisement settings to the default values.

```
unconfigure irdp
```

unconfigure vlan subvlan-address-range

```
unconfigure vlan <vlan_name> subvlan-address-range
```

Description

Unconfigures subVLAN address ranges on each subVLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

Syntax Description

vlan_name	Specifies a subVLAN name.
-----------	---------------------------

Default

N/A.

Usage Guidelines

This command removes a subVLAN address range. There is no error checking to prevent the configuration of overlapping subVLAN address ranges between multiple subVLANs. Doing so can result in unexpected behavior of ARP within the superVLAN and associated subVLANs.

unconfigure vlan udp-profile

```
unconfigure vlan <vlan_name> udp-profile
```

Description

Removes any UDP forwarding profile from a VLAN.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

No UDP profiles are associated with the VLAN.

Usage Guidelines

None.

Example

The following command removes any UDP forwarding profile from the VLAN *to-sales*:

```
unconfigure vlan to-sales udp-profile
```

This chapter describes commands for configuring and managing the following IPv6 features:

- IPv6 unicast routing
- Route sharing
- Route compression
- IP multinetting

For an introduction to these IPv6 features, see the *NETGEAR 8800 User Manual*.

clear neighbor-discovery cache

```
clear neighbor-discovery cache ipv6 {<ipv6address> {vr <vr_name>} | vlan <vlan_name> | vr <vr_name>}
```

Description

Deletes a dynamic entry from the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
ipv6address	Specifies an IPv6 address.
vlan_name	Specifies an IPv6 configured VLAN.

Default

N/A.

Usage Guidelines

This command clears dynamic entries from the neighbor cache. The `vr` option is used to specify the virtual router on which the operation is performed. When this option is omitted it applies to *VR-Default*.

When the `ipv6address` or `vlan` options are specified, only the entries with matching IPv6 addresses or that correspond to that VLAN are cleared.

Example

The following command clears all entries from the neighbor cache:

```
clear neighbor-discovery cache
```

configure iproute add (IPv6)

```
configure iproute add <ipv6Netmask> [<ipv6Gateway> | <ipv6ScopedGateway>] {<metric>} {vr
<vrname>} {multicast | multicast-only | unicast | unicast-only}
```

Description

Adds an IPv6 static route to the routing table.

Syntax Description

ipv6Netmask	Specifies an IPv6 address/prefix length.
ipv6Gateway	Specifies a gateway.
ipv6ScopedGateway	Specifies a scoped gateway.
metric	Specifies a cost metric.
vrname	Specifies the virtual router to which the route is added.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table.

Default

If you do not specify a virtual router, the current virtual router context is used. If you do not specify a metric, then the default metric of 1 is used.

Usage Guidelines

Use a prefix length of 128 to indicate a host entry.

Note: Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the `multicast` option.

Example

The following command adds a static route to the routing table:

```
configure iproute add 2001:db8:0:1111::/64 fe80::1111%default
```

configure iproute add blackhole

```
configure iproute add blackhole {ipv6} [<ipv6Netmask>] {vr <vrname>} {multicast-only | unicast-only}
```

Description

Adds a blackhole address to the routing table. All traffic destined for a configured blackhole IP address is silently dropped.

Syntax Description

ipv6Netmask	Specifies an IPv6 address/prefix length.
vrname	Specifies the virtual router to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

A blackhole entry directs packets with a matching specified address prefix to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

The packets are silently discarded. In other words, no ICMP message is sent to indicate that the packets are discarded.

Example

The following command causes packets with a destination address of 2001:db8::3452 to be silently discarded:

```
configure iproute add blackhole 2001:db8::3452/128
```

configure iproute add blackhole ipv6 default

```
configure iproute add blackhole ipv6 default {vr <vrname>} {multicast-only | unicast-only}
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IP destination is silently dropped.

Syntax Description

vr_name	Specifies the virtual router to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

While a default route is for *forwarding* traffic destined to an unknown IP destination, and a blackhole route is for *discarding* traffic destined to a specified IP destination, a *default blackhole* route is for *discarding* traffic to the unknown IP destination.

Using this command, all traffic with an unknown destination is discarded.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IP address for this route is ::.

The packets are silently discarded. In other words, no ICMP message is sent to indicate that the packets are discarded.

Example

The following command adds a blackhole default route into the routing table:

```
configure iproute add blackhole ipv6 default
```

configure iproute add default

```
configure iproute add default [<ipv6Gateway> | <ipv6ScopedGateway>] {metric} {vr <vrname>}
{multicast-only | unicast-only}
```

Description

Adds a default gateway to the routing table.

Syntax Description

ipv6Gateway	Specifies a VLAN gateway IPv6 address.
metric	Specifies a cost metric. If no metric is specified, the default of 1 is used.

ipv6ScopedGateway	Specifies a scoped gateway.
vrname	Specifies the virtual router to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface. Use the `unicast-only` or `multicast-only` options to specify a particular traffic type. If not specified, both unicast and multicast traffic uses the default route.

Example

The following command configures a default route for the switch:

```
configure iproute add default 2001:db8::1234:5678
```

configure iproute delete

```
configure iproute delete <ipv6Netmask> [<ipv6Gateway> | <ipv6ScopedGateway>] {vr <vrname>}
```

Description

Deletes an IPv6 static route from the routing table.

Syntax Description

ipv6Netmask	Specifies an IPv6 address/prefix length.
ipv6Gateway	Specifies a gateway.
ipv6ScopedGateway	Specifies a scoped gateway.
vrname	Specifies the virtual router to which the route is added.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a prefix length of 128 to indicate a host entry.

Example

The following command deletes a static address from the routing table:

```
configure iproute delete 2001:db8:0:1111::/64 fe80::1111
```

configure iproute delete blackhole

```
configure iproute delete blackhole [<ipv6Netmask>] {vr <vrname>}
```

Description

Deletes a blackhole route from the routing table.

Syntax Description

ipv6Netmask	Specifies an IPv6 address/prefix length.
vrname	Specifies the virtual router to which the route is added.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

A blackhole entry directs packets with a specified destination address to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

Example

The following command deletes a blackhole route from the routing table for packets with a destination address of 2001:db8::3452, so the packets are no longer discarded:

```
configure iproute delete blackhole 2001:db8::3452/128
```

configure iproute delete blackhole ipv6 default

```
configure iproute delete blackhole ipv6 default {vr <vrname>}
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

vrname	Specifies the virtual router to which the route is added.
--------	---

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

While a default route is for *forwarding* traffic destined to an unknown IP destination, and a blackhole route is for *discarding* traffic destined to a specified IP destination, a *default blackhole* route is for *discarding* traffic to the unknown IP destination.

Using this command, all traffic with an unknown destination is discarded.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IP address for this route is ::.

Example

The following command deletes a blackhole default route from the routing table:

```
configure iproute delete blackhole default
```

configure iproute delete default

```
configure iproute delete default [<ipv6Gateway> | <ipv6ScopedGateway>] {vr <vrname>}
```

Description

Deletes a default gateway from the routing table.

Syntax Description

ipv6Gateway	Specifies a VLAN gateway IPv6 address.
ipv6ScopedGateway	Specifies a scoped gateway.
vrname	Specifies the virtual router to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface.

Example

The following command deletes a default route from the switch:

```
configure iproute delete default 2001:db8::1234:5678
```

configure iproute ipv6 priority

```
configure iproute ipv6 priority [ripng | blackhole | icmp | static | ospfv3-intra |
ospfv3-inter | ospfv3-as-external | ospfv3-extern1 | ospfv3-extern2] <priority> {vr <vrname>}
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

ripng	Specifies RIPng.
icmp	Specifies ICMP.
blackhole	Specifies the blackhole route.
static	Specifies static routes.
ospfv3-intra	Specifies OSPFv3 Intra routing.
ospfv3-inter	Specifies OSPFv3 Inter routing.
ospfv3-as-external	Specifies OSPFv3 AS External routing.
ospfv3-extern1	Specifies OSPFv3 External 1 routing.
ospfv3-extern2	Specifies OSPFv3 External 2 routing.
priority	Specifies a priority number in the range of 11 to 65534.
vrname	Specifies a virtual router name.

Default

Table 26 lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 26. Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPF3Intra	2200
OSPF3Inter	2300
IS-IS L1	2360
IS-IS L2	2370

Table 26. Route Priorities (Continued)

Route Origin	Priority
RIPg	2400
OSPFv3 ASExt	3100
OSPFv3 Extern1	3200
OSPFv3 Extern2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500

Usage Guidelines

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences. If you change the route priority, you must save the configuration and reboot the system.

Note: The priority for a blackhole route can not overlap with the priority of any other route origin.

Example

The following command sets the IPv6 route priority for static routing to 1200:

```
configure iproute ipv6 priority static 1200
```

configure neighbor-discovery cache add

```
configure neighbor-discovery cache {vr <vr_name>} add [<ipv6address> | <scoped_link_local>]  
<mac>
```

Description

Adds a static entry to the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
ipv6address	Specifies an IPv6 address.
scoped_link_local	Specifies a scoped, link-local address.
mac	Specifies a MAC address.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

This command adds static entries to the neighbor cache.

Example

The following command adds a static entry to the neighbor cache:

```
configure neighbor-discovery cache add fe80::2315%default 00:11:22:33:44:55
```

configure neighbor-discovery cache delete

```
configure neighbor-discovery cache {vr <vr_name>} delete [<ipv6address> |  
<scoped_link_local>]
```

Description

Deletes a static entry from the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
ipv6address	Specifies an IPv6 address.
scoped_link_local	Specifies a scoped, link-local address.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

This command deletes static entries from the neighbor cache.

Example

The following command deletes a static entry from the neighbor cache:

```
configure neighbor-discovery cache delete fe80::2315%default
```

configure neighbor-discovery cache max_entries

```
configure neighbor-discovery cache {vr <vr_name>} max_entries <max_entries>
```

Description

Configures the maximum allowed IPv6 neighbor entries.

Syntax Description

vr_name	Specifies a virtual router.
max_entries	Specifies the maximum allowed IPv6 neighbor entries. The range is 1 to 20480.

Default

4096.

Usage Guidelines

None.

Example

The following command sets the maximum allowed IPv6 neighbor entries to 512:

```
configure neighbor-discovery cache max_entries 512
```

configure neighbor-discovery cache max_pending_entries

```
configure neighbor-discovery cache {vr <vr_name>} max_pending_entries <max_pending_entries>
```

Description

Configures the maximum number of pending IPv6 neighbor entries.

Syntax Description

vr_name	Specifies a virtual router.
max_entries	Specifies the maximum number of pending IPv6 neighbor entries. The range is 1 to 4096.

Default

1024.

Usage Guidelines

None.

Example

The following command sets the maximum number of pending IPv6 neighbor entries to 2056:

```
configure neighbor-discovery cache max_pending_entries 2056
```

configure neighbor-discovery cache timeout

```
configure neighbor-discovery cache {vr <vr_name>} timeout <timeout>
```

Description

Configures a timeout value for entries in the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
timeout	Specifies a timeout value for neighbor cache entries.

Default

20 minutes

Usage Guidelines

None.

Example

The following command configures the neighbor cache timeout for 30 minutes:

```
configure neighbor-discovery cache timeout 30
```

configure vlan router-discovery add prefix

```
configure vlan <vlan_name> router-discovery {ipv6} add prefix <prefix>
```

Description

Adds a prefix to the router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
prefix	Specifies the prefix to add.

Default

N/A.

Usage Guidelines

This command adds a prefix to the router advertisement messages for the VLAN. Prefixes defined with this command are only included in the router advertisement messages and have no operational impact on VLANs.

To configure the parameters for this prefix, use the following command:

```
configure vlan <vlan_name> router-discovery {ipv6} set prefix <prefix> [autonomous-flag
<auto_on_off> | onlink-flag <onlink_on_off> | preferred-lifetime <preflife>
|valid-lifetime <validlife>]
```

Example

The following command adds the prefix 2001:db8:3456::/64 for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery add prefix 2001:db8:3456::/64
```

configure vlan router-discovery delete prefix

```
configure vlan <vlan_name> router-discovery {ipv6} delete prefix [<prefix> | all]
```

Description

Deletes prefixes from the router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
prefix	Specifies the prefix to delete.
all	Specifies to delete all prefixes.

Default

N/A.

Usage Guidelines

This command deletes previously defined router advertisement prefixes.

Example

The following command deletes the prefix 2001:db8:3161::/64 for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery delete 2001:db8:3161::/64
```

configure vlan router-discovery default-lifetime

```
configure vlan <vlan_name> router-discovery {ipv6} default-lifetime <defaultlifetime>
```

Description

Configures the router lifetime value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
defaultlifetime	Specifies the router lifetime. Range is 0, max-interval to 9000 seconds.

Default

1800 seconds

Usage Guidelines

This command configures the router lifetime value to be included in the router advertisement messages.

The value is specified in seconds and is either 0, or between max-interval and 9000 seconds. A value of 0 indicates that the router is not to be used as a default router.

After a host sends a router solicitation, and receives a valid router advertisement with a non-zero router lifetime, the host must desist from sending additional solicitations on that interface, until an event such as re-initialization takes place.

Example

The following command configures the default-lifetime to be 3600 seconds for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery default-lifetime 3600
```

configure vlan router-discovery hop-limit

```
configure vlan <vlan_name> router-discovery {ipv6} hop-limit <currenthoplimit>
```

Description

Configures the current hop limit value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
currenthoplimit	Specifies the current hop limit. Range is 0 to 255.

Default

64.

Usage Guidelines

Configures the current hop limit value to be included in the router advertisement messages. Hosts then use the current hop limit when they send packets.

A value of 0 means unspecified by this router. The default value is 64. The maximum value is 255.

Example

The following command configures the current hop limit to be 32 for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery hop-limit 32
```

configure vlan router-discovery link-mtu

```
configure vlan <vlan_name> router-discovery {ipv6} link-mtu <linkmtu>
```

Description

Configures the link MTU value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
linkmtu	Specifies the link MTU. Range is 0 to 9216.

Default

0, meaning that no link MTU information is sent.

Usage Guidelines

This command configures the link MTU placed into the router advertisement messages. Advertisement of the MTU helps ensure use of a consistent MTU by hosts on the VLAN.

The minimum value is 0. The maximum value is 9216. The default value is 0, which means that no link MTU information is included in the router discovery messages.

Example

The following command configures the link MTU to be 5126 for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery link-mtu 5126
```

configure vlan router-discovery managed-config-flag

```
configure vlan <vlan_name> router-discovery {ipv6} managed-config-flag <on_off>
```

Description

Configures the managed address configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
on_off	Specifies setting the flag to on or off.

Default

Off.

Usage Guidelines

This command configures the contents of the managed address configuration flag in the router advertisement messages.

A value of on tells hosts to use the administered (stateful) protocol (DHCP) for address autoconfiguration in addition to any addresses autoconfigured using stateless address autoconfiguration. A value of off tells hosts to use stateless address autoconfiguration. If this command is not entered, the default value is off.

Example

The following command configures the managed address configuration flag to be on for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery managed-config-flag on
```

configure vlan router-discovery max-interval

```
configure vlan <vlan_name> router-discovery {ipv6} max-interval <maxinterval>
```

Description

Configures the maximum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
maxinterval	Specifies the maximum time between advertisements, in seconds. Range is 4 to 1800

Default

600 seconds

Usage Guidelines

This command configures the maximum amount of time before an unsolicited router advertisement message is advertised over the links corresponding to the VLAN.

Example

The following command configures the max-interval to be 300 seconds for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery max-interval 300
```

configure vlan router-discovery min-interval

```
configure vlan <vlan_name> router-discovery {ipv6} min-interval <mininterval>
```

Description

Configures the minimum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
mininterval	Specifies the minimum time between advertisements, in seconds. Range is 3 to 1350 (see guidelines).

Default

200 seconds, or max-interval * .33 (see guidelines)

Usage Guidelines

This command configures the minimum amount of time before an unsolicited router advertisement message is advertised over the links corresponding to the VLAN.

The minimum value is 3 seconds. The maximum time is (.75 * max-interval) seconds. If you do not explicitly set this value, the min-interval value is reset whenever the max-interval is configured. Min-interval will then be dynamically adjusted to .33 times the max-interval.

Example

The following command configures the min-interval to be 300 seconds for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery min-interval 300
```


configure vlan router-discovery other-config-flag

```
configure vlan <vlan_name> router-discovery {ipv6} other-config-flag <on_off>
```

Description

Configures the other stateful configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
on_off	Specifies setting the flag to on or off.

Default

Off.

Usage Guidelines

This command configures the contents of the other stateful configuration flag in the router advertisement messages.

When set to on, hosts use the administered (stateful) protocol (DHCP) for autoconfiguration of other (non-address) information. If this command is not entered, the default value is off.

Example

The following command configures the other stateful configuration flag to be on for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery other-config-flag on
```

configure vlan router-discovery reachable-time

```
configure vlan <vlan_name> router-discovery {ipv6} reachable-time <reachabletime>
```

Description

Configures the reachable time value in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
reachabletime	Specifies the reachable time value in advertisements, in milliseconds. Range is 0 to 3,600,000 (one hour).

Default

30,000 milliseconds.

Usage Guidelines

The reachable time is the time, in milliseconds, that a node assumes a neighbor is reachable after having received a reachability confirmation. A value of 0 means the time is unspecified by this router. The maximum value is 3,600,000 (1 hour).

Example

The following command configures the reachable time to be 3,600,000 milliseconds for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery reachable-time 3600000
```

configure vlan router-discovery retransmit-time

```
configure vlan <vlan_name> router-discovery {ipv6} retransmit-time <retransmittime>
```

Description

Configures the retransmit time value in router discovery advertisements on the VLAN.

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
<code>retransmittime</code>	Specifies the reachable time value in advertisements, in milliseconds. Range is 0 to 4,294,967,295 (approximately 50 days).

Default

1,000 milliseconds.

Usage Guidelines

This command configures the retransmit time value in the router advertisement messages.

The retransmit time, in milliseconds, is the time between retransmitted neighbor solicitation messages. A value of 0 means the value is unspecified by this router. The maximum value is 4,294,967,295.

Example

The following command configures the retransmit time to be 604,800,000 milliseconds (one week) for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery retransmit-time 604800000
```

configure vlan router-discovery set prefix

```
configure vlan <vlan_name> router-discovery {ipv6} set prefix <prefix> [autonomous-flag
<auto_on_off> | onlink-flag <onlink_on_off> | preferred-lifetime <preflife> | valid-lifetime
<validlife>]
```

Description

Sets the parameters for a prefix in the router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
prefix	Specifies which prefix's parameters to set.
auto_on_off	Specifies the autonomous flag.
onlink_on_off	Specifies the on link flag.
preflife	Specifies the preferred lifetime in seconds. Maximum value is 4,294,967,295.
validlife	Specifies the valid lifetime in seconds. Maximum value is 4,294,967,295.

Default

The prefix parameter defaults are:

- Valid lifetime—2,592,000 seconds (30 days)
- On-link flag—on
- Preferred lifetime—604,800 seconds (7 days)
- Autonomous flag—on

Usage Guidelines

This command configures the attributes associated with the specified prefix.

The autonomous flag option modifies the autonomous flag of the prefix. The autonomous flag value specifies whether the prefix can be used for autonomous address configuration (on) or not (off).

The onlink flag option modifies the on link flag of the prefix. The on link flag specifies whether the prefix can be used for on link determination (on) or not (off). The default value of the on link flag is on.

The preferred lifetime option modifies the preferred lifetime of a prefix. The preferred lifetime value is the time (from when the packet is sent) that addresses generated from the prefix via stateless address autoconfiguration remain preferred. The maximum value is 4,294,967,295. The default value is 604,800 seconds (7 days).

The valid lifetime option modifies the valid lifetime of a prefix. The valid lifetime value is the time (from when the packet was sent) that the prefix is valid for the purpose of on-link

determination. The maximum value is a 4,294,967,295. The default value is 2,592,000 seconds (30 days).

Example

The following command sets the on link parameter of the prefix 3aaa:3161::/64 to off, for the VLAN *top_floor*:

```
configure vlan top_floor router-discovery set prefix 3aaa:3161::/64 onlink-flag off
```

configure tunnel ipaddress

```
configure tunnel <tunnel_name> ipaddress [ipv6-link-local | {eui64} <ipv6_address_mask> ]
```

Description

Configures an IPv6 address/prefix on a tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
eui64	Specifies an EUI64 interface identifier for the lower 64 bits of the address.
ipv6_address_mask	Specifies an IPv6 address / IPv6 prefix length.
ipv6-link-local	Specifies the link-local address for a tunnel.

Default

N/A.

Usage Guidelines

This command will configure an IPv6 address/prefix route on the specified tunnel.

6to4 tunnels must follow the standard address requirement. The address must be of the form 2002:<IPv4_source_endpoint>::/16, where <IPv4_source_endpoint> is replaced by the IPv4 source address of the endpoint, in hexadecimal, colon separated form. For example, for a tunnel endpoint located at IPv4 address 10.20.30.40, the tunnel address would be 2002:a14:1e28::/16. In hex, 10 is a, 20 is 14, 30 is 1e and 40 is 28.

6in4 tunnels have no restrictions on their address format or prefix allocations.

Example

The following command configures the 6in4 tunnel link39 with the IPv6 link-local address:

```
configure tunnel link39 ipaddress ipv6-link-local
```

create tunnel 6to4

```
create tunnel <tunnel_name> 6to4 source <source-address>
```

Description

Creates an IPv6-to-IPv4 (6to4) tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
source-address	Specifies an IPv4 address for the tunnel.

Default

N/A.

Usage Guidelines

This command will create a new IPv6-to-IPv4 (also known as a 6to4 tunnel), and add it to the system. A maximum of 1 6to4 tunnel can be configured on any particular virtual router.

The tunnel name must be unique and cannot overlap the same name space as VLANs, other tunnels, or virtual routers. The name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

The source address of the tunnel must be one of the IPv4 addresses already configured on the switch. You cannot remove an IPv4 address from the switch if a tunnel that uses it still exists.

Example

The following command creates the 6to4 tunnel link35 with source address 192.168.10.1:

```
create tunnel link35 6to4 source 192.168.10.1
```

create tunnel ipv6-in-ipv4

```
create tunnel <tunnel_name> ipv6-in-ipv4 destination <destination-address> source  
<source-address>
```

Description

Creates an IPv6-in-IPv4 (6in4) tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
source-address	Specifies an IPv4 address for the tunnel.

Default

N/A.

Usage Guidelines

This command will create a new IPv6-in-IPv4 (otherwise known as a configured tunnel or a 6in4 tunnel) and add it to the system. A maximum of 255 tunnels (including one 6to4 tunnel) can be configured on the system.

The tunnel name must be unique and cannot overlap the same name space as VLANs, other tunnels, or virtual routers. The name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

The source address of the tunnel must be one of the IPv4 addresses already configured on the switch. You cannot remove an IPv4 address from the switch if a tunnel is still exists that uses it.

Example

The following command creates the 6in4 tunnel link39 with destination address 10.10.10.10 and source address 192.168.10.15:

```
create tunnel link39 ipv6-in-ipv4 destination 10.10.10.10 source 192.168.10.15
```

delete tunnel

```
delete tunnel <tunnel_name>
```

Description

Deletes an IPv6 tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
-------------	---------------------------

Default

N/A.

Usage Guidelines

This command will destroy a previously created tunnel. The command acts on either a 6to4 or a 6in4 tunnel. When the tunnel interface is removed, all dynamic routes through that interface are purged from the system. The configured static routes are removed from the hardware tables and become inactive.

Example

The following command deletes the tunnel link39:

```
delete tunnel link39
```

disable ipforwarding ipv6

```
disable ipforwarding ipv6 {vlan <vlan_name> | tunnel <tunnel_name> | vr <vr_name>}
```

Description

Disables routing for one or all interfaces. If no argument is provided, disables routing for all interfaces on the current virtual router.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
tunnel_name	Specifies an IPv6 tunnel.
vr_name	Specifies a virtual router.

Default

Disabled.

Usage Guidelines

When new IPv6 interfaces are added, IP forwarding is disabled by default.

NETGEAR switches have a single hardware control per VLAN for IP forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv6 forwarding on a VLAN also enables IPv4 hardware forwarding on that VLAN.

Example

The following command disables forwarding of IPv6 traffic for a VLAN named *accounting*:

```
disable ipforwarding ipv6 vlan accounting
```

disable iproute ipv6 sharing

```
disable iproute ipv6 sharing
```

Description

Disables IPv6 route sharing.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables load sharing for multiple routes:

```
disable iproute ipv6 sharing
```

disable iproute ipv6 compression

```
disable iproute ipv6 compression {vr <vrname>}
```

Description

This command disables IPv6 route compression.

Syntax Description

vrname	Specifies the virtual router to route from.
--------	---

Default

Vrname—current CLI context VR

By default, IP route compression is disabled for all address families and VRs.

Usage Guidelines

This command disables IP route compression for the IPv6 address family and VR. This command decompresses previously compressed prefixes in the IPv6 prefix database.

Example

The following example disables IP route compression for the IPv6 address family and the VR of the current CLI context.

```
disable iproute ipv6 compression
```

disable neighbor-discovery refresh

```
disable neighbor-discovery {vr <vr_name>} refresh
```

Description

Prevents the IPv6 neighbor cache from refreshing an entry before the timeout period expires.

Syntax Description

vr_name	Specifies a virtual router.
---------	-----------------------------

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables the refresh of neighbor discovery cache entries:

```
disable neighbor-discovery refresh
```

disable router-discovery

```
disable router-discovery {ipv6} vlan <vlan_name>
```

Description

Disables router discovery advertisements on the VLAN and the processing of router discovery messages.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

None

Example

The following command disables router discovery for the VLAN *top_floor*:

```
disable router-discovery vlan top_floor
```

enable ipforwarding ipv6

```
enable ipforwarding ipv6 {vlan <vlan_name> | tunnel <tunnel-name> | vr <vr_name>}
```

Description

Enables IP routing VLANs. If no argument is provided, enables IP routing for all VLANs and tunnels that have been configured with an IPv6 address on the current virtual router.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
tunnel_name	Specifies an IPv6 tunnel.
vr_name	Specifies a virtual router.

Default

Disabled.

Usage Guidelines

When new IPv6 interfaces are added, IP forwarding is disabled by default.

NETGEAR switches have a single hardware control per VLAN for IP forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv6 forwarding on a VLAN also enables IPv4 hardware forwarding on that VLAN.

Example

The following command enables forwarding of IPv6 traffic for all VLANs in the current virtual router context with IPv6 addresses:

```
enable ipforwarding ipv6
```

enable iproute ipv6 compression

```
enable iproute ipv6 compression {vr <vrname>}
```

Description

This command enables IPv6 route compression.

Syntax Description

vrname	Specifies the virtual router to route from.
--------	---

Default

Vrname—current CLI context VR

Usage Guidelines

This command enables IP route compression for the IPv6 address family and VR. This command applies a compression algorithm to each IPv6 prefix in the IPv6 prefix database.

Example

The following example enables IP route compression for the IPv6 address family and the VR of the current CLI context.

```
enable iproute ipv6 compression
```

enable iproute ipv6 sharing

```
enable iproute ipv6 sharing
```

Description

Enables load sharing if multiple routes to the same destination are available. When multiple routes to the same destination are available, load sharing can be enabled to distribute the traffic to multiple destination gateways. Only paths with the same lowest cost are shared.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Configure static routes and OSPF as you would normally. NETGEAR 8800 supports route sharing across up to 8 static routes or ECMP routes for OSPF. However, on the NETGEAR 8800 family, by default up to 4 routes are supported. To support up to 8 routes on these switches, use the following command:

```
configure iproute sharing max-gateways <max_gateways>
```

Example

The following command enables load sharing for multiple routes:

```
enable iproute ipv6 sharing
```

enable neighbor-discovery refresh

```
enable neighbor-discovery {vr <vr_name>} refresh
```

Description

Enables the IPv6 neighbor cache to refresh each entry before the timeout period expires.

Syntax Description

vr_name	Specifies a virtual router.
---------	-----------------------------

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the refresh of neighbor discovery cache entries:

```
enable neighbor-discovery refresh
```

enable router-discovery

```
enable router-discovery {ipv6} vlan <vlan_name>
```

Description

Enables router discovery advertisements on the VLAN and the processing of router discovery messages.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command is only valid when the specified VLAN has an IPv6 address associated with it. After IPv6 Router Discovery is enabled on a VLAN, router advertisement messages are regularly sent on all ports associated with the VLAN.

Example

The following command enables router discovery for the VLAN *top_floor*:

```
enable router-discovery vlan top_floor
```

rtlookup

```
rtlookup [<ipaddress> | <ipv6address>] { unicast | multicast | vr <vr_name> }
```

Description

Displays the available routes to the specified IP address.

Syntax Description

<code>ipaddress</code>	Specifies an IPv4 address.
<code>ipv6address</code>	Specifies an IPv6 address.
<code>unicast</code>	Displays the routes from the unicast routing table in the current router context.
<code>multicast</code>	Displays the routes from the multicast routing table in the current router context.
<code>vr_name</code>	Specifies the virtual router for which to display the route.

Default

N/A.

Usage Guidelines

None.

Example

The following command performs a look up in the route table to determine the best way to reach the specified IP address:

```
rtlookup 3aaa:5643::ef80:2525:1023:5213 unicast
```

rtlookup rpf

```
rtlookup [<ipaddress> | <ipv6address>] rpf {vr <vr_name>}
```

Description

Displays the RPF for a specified multicast source.

Syntax Description

<code>ipaddress</code>	Specifies an IPv4 address.
<code>ipv6address</code>	Specifies an IPv6 address.
<code>rpf</code>	Selects the RPF for the specified multicast source.

vr_name	Specifies the virtual router for which to display the route.
---------	--

Default

vr_name is the VR of the current CLI context.

Usage Guidelines

None.

Example

The following example displays the RPF lookup for multicast source 12.1.20.12 in *VR-Default*:

```
rtlookup 3aaa:5643::ef80:2525:1023:5213 rpf vr vr-default
```

show ipconfig ipv6

```
show ipconfig ipv6 {vlan <vlan_name> | tunnel <tunnelname>}
```

Description

Displays configuration information for one or more interfaces.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
tunnelname	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

If no interface is specified, then global IP configuration is displayed. Otherwise, specific interface(s) will be displayed. Global IP configuration information includes:

- IP address/netmask/etc.
- IP forwarding information / IP multicast forwarding information

Example

The following command displays configuration information on a VLAN named *accounting*:

```
show ipconfig ipv6 vlan accounting
```

show iproute ipv6

```
show iproute ipv6 {priority | vlan <vlan_name> | tunnel <tunnel-name> | <ipv6Netmask> |
summary {multicast | unicast}} {vr <vrname>}}
```

Description

Displays the contents of the IPv6 routing table.

Syntax Description

priority	Displays the priority values for each route origin type.
vlan_name	Specifies a VLAN name.
tunnel	Specifies a tunnel name.
ipv6Netmask	Specifies an IPv6 address/prefix length.
summary	Specifies summary information
vrname	Specifies a virtual router.

Default

N/A.

Usage Guidelines

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command displays detailed information about all IPv6 routing:

```
Switch.18 # show iproute ipv6
Ori Destination                                     Mtr  Flags          Duration
Gateway                                             Interface
#d  3001::/64                                       1    U-----um--f  0d:0h:5m:31s
    3001::52                                       ixia
#or 3020:1:2:78::/64                               50   UG-D---um--f  0d:0h:0m:1s
    fe80::200:40ff:feba:a38e                       ixia
#or 3020:1:2:79::/64                               50   UG-D---um--f  0d:0h:0m:1s
    fe80::200:40ff:feba:a38e                       ixia
#or 3020:1:2:7a::/64                               50   UG-D---um--f  0d:0h:0m:1s
    fe80::200:40ff:feba:a38e                       ixia
#or 3020:1:2:7b::/64                               50   UG-D---um--f  0d:0h:0m:1s
    fe80::200:40ff:feba:a38e                       ixia
#or 3020:1:2:7c::/64                               50   UG-D---um--f  0d:0h:0m:1s
    fe80::200:40ff:feba:a38e                       ixia
#or 3020:1:2:7d::/64                               50   UG-D---um--f  0d:0h:0m:1s
    fe80::200:40ff:feba:a38e                       ixia
```

```

#or 3020:1:2:7e::/64          50    UG-D---um--f 0d:0h:0m:1s
    fe80::200:40ff:feba:a38e    ixia
#or 3020:1:2:7f::/64          50    UG-D---um--f 0d:0h:0m:1s
    fe80::200:40ff:feba:a38e    ixia
#or 3020:1:2:80::/64          50    UG-D---um--f 0d:0h:0m:1s
    fe80::200:40ff:feba:a38e    ixia
#or 3020:1:2:81::/64          50    UG-D---um--f 0d:0h:0m:1s
    fe80::200:40ff:feba:a38e    ixia
#d fe80::%ixia/64             1     U-----um--f 0d:0h:5m:31s
    fe80::204:96ff:fe27:8697    ixia

```

Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
 (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
 (e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
 (is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (ma) MPLSIntra
 (mr) MPLSInter, (mo) MOSPF (o) OSPFv3, (o1) OSPFv3Ext1, (o2) OSPFv3Ext2
 (oa) OSPFv3Intra, (oe) OSPFv3AsExt, (or) OSPFv3Inter, (pd) PIM-DM, (ps) PIM-SM
 (r) RIPng, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) Unknown
 (*) Preferred unicast route (@) Preferred multicast route
 (#) Preferred unicast and multicast route

Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
 (L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
 (P) LPM-
 routing, (R) Modified, (S) Static, (s) Static LSP
 (T) Matching RSVP-TE LSP (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
 (f) Provided to FIB (c) Compressed Route

Mask distribution:
 12 routes at length 64

Route Origin distribution:
 2 routes from Direct 10 routes from OSPFv3Inter

Total number of routes = 12
 Total number of compressed routes = 0

The following command displays the IPv6 route origin priority:

```

XCM8806.12 # show iproute ipv6 priority
Direct          10
Blackhole        50
Static          1100
ICMP            1200
OSPFv3Intra     2200
OSPFv3Inter     2300

```


Isis	2350
IsisL1	2360
IsisL2	2370
RIPng	2400
OSPFv3AsExt	3100
OSPFv3Ext1	3200
OSPFv3Ext2	3300
IsisL1Ext	3400
IsisL2Ext	3500

show iproute ipv6 origin

```
show iproute ipv6 origin [direct | static | blackhole | ripng | ospfv3 | ospfv3-intra |
ospfv3-inter | ospfv3-extern1 | ospfv3-extern2] {vr <vrname>}
```

Description

Displays the contents of the IP routing table for routes with the specified origin.

Syntax Description

direct	Specifies direct routes.
static	Specifies static routes.
blackhole	Specifies blackhole routes.
ripng	Specifies RIPng routes.
ospfv3	Specifies OSPFv3 routes.
ospfv3-intra	Specifies OSPFv3 Intra routing.
ospfv3-inter	Specifies OSPFv3 Inter routing.
ospfv3-extern1	Specifies OSPFv3 External 1 routing.
ospfv3-extern2	Specifies OSPFv3 External 2 routing.
vrname	Specifies a virtual router.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the RIPng routes:

```
show iproute ipv6 origin ripng
```

show ipstats ipv6

```
show ipstats ipv6 {vlan <name> | tunnel <tunnelname> | vr <vrname>}
```

Description

Displays IPv6 statistics for the CPU for the switch or for a particular VLAN.

Syntax Description

name	Specifies a VLAN name.
tunnelname	Specifies a tunnel name.
vrname	Specifies a virtual router.

Default

N/A.

Usage Guidelines

This command only shows statistics of the CPU-handled packets. Not all packets are handled by the CPU. For example, packets forwarded in hardware do not increment the statistics counters.

If you do not specify a virtual router, the command applies to the current virtual router.

Example

The following command displays IPv6 statistics for the VLAN *accounting*:

```
show ipstats ipv6 vlan accounting
```

show neighbor-discovery cache ipv6

```
show neighbor-discovery {cache {ipv6}} {[<ipv6_addr> | <mac> | permanent] {vr <vr_name>}} |  
vlan <vlan_name> | vr <vr_name>}
```

Description

This command displays all the entries from the neighbor cache.

Syntax Description

vr_name	Specifies a virtual router.
mac	Specifies a MAC address.
permanent	Specifies static entries.
ipv6_addr	Specifies an IPv6 address.

vlan_name	Specifies an IPv6 configured VLAN.
vr_name	Specifies a virtual router.

Default

N/A.

Usage Guidelines

This command displays the entries present in the neighbor cache.

The entries displayed can be filtered by IPv6 address, MAC address, or by VLAN. The `permanent` keyword filters the output to display static entries.

The `vr_name` indicates the virtual router on which the operation is performed. In its absence, the operation applies to *VR-Default*.

Example

The following command shows all entries from the neighbor cache:

```
show neighbor-discovery cache ipv6
```

The following is sample output:

VR	Destination	Mac	Age	Static	VLAN	VID	Port
VR-Default	3ffe:100::7	00:01:30:00:6b:00	0	NO	gtag100	100	1:2
VR-Default	3ffe:100::99	00:01:02:33:33:33	0	YES	gtag100	100	
VR-Default	3ffe:99::99	00:01:02:01:01:01	0	YES	gtag99	99	
Total Entries : 0							
Dynamic Entries : 0 Static Entries : 0							
Pending Entries : 0							
Max Entries : 1024 Max Pending entries : 1024							
Timeout : 20 minutes Refresh : Enable							

show router-discovery

```
show router-discovery {ipv6} {vlan <vlan_name>}
```

Description

Displays the router discovery settings.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

If no VLAN is specified, the settings are displayed for all IPv6 configured VLANs.

Example

The following command displays router discovery settings for the VLAN *top_floor*:

```
show router-discovery vlan top_floor
```

The following is sample output:

```
Router Advertisements disabled on vl
  Minimum/Maximum Interval: 200 / 600
  Managed / Other Info Flags: Off / Off
  Link MTU: 0
  Reachable Time: 0
  Retrans Timer: 0
  Current Hop Limit: 64
  Default Lifetime: 1800
  Number of Prefixes: 1, Prefix List:
                                Valid      Preferred
      Prefix  Lifetime      Auto  Lifetime OnLink
11::1/64
                                2592000    On    604800    On
```

show tunnel

```
show [{tunnel}] {<tunnel_name>}}
```

Description

Displays system tunnel information for a specified tunnel or for all tunnels.

Syntax Description

tunnel_name	Specifies an IPv6-in-IPv4 or IPv6-to-IPv4 tunnel name.
-------------	--

Default

N/A.

Usage Guidelines

The `tunnel` keyword is optional only when you specify a valid IPv6-in-IPv4 or IPv6-to-IPv4 tunnel name. The *Total tunnels* count in the display represents all tunnels on the switch.

Examples

The following command displays system tunnel information for all tunnels:

```
Switch.1 # show tunnel
```

Name	Type	Flags
tunfour	6in4 10.20.30.40 => 10.10.10.10	U
tunfive2	6to4 10.20.30.40 => *.*.*.*	D

Total tunnels: 2

Flags: (U) Up / (D) Down

The following command displays system tunnel information for tunnel *tunfour*:

```
Switch.3 # show "tunfour"
```

Name	Type	Flags
tunfour	6in4 10.20.30.40 => 10.10.10.10	U

Total tunnels: 2

Flags: (U) Up / (D) Down

unconfigure iproute ipv6 priority

```
unconfigure iproute ipv6 priority [all | blackhole | icmp | ospfv3-as-external |
ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter | ospfv3-intra | ripng | static] {vr <vrname>}
```

Description

Resets the priority for all IPv6 routes from one or all route origin types to the default values.

Syntax Description

all	Specifies all route origins.
blackhole	Specifies the blackhole route.
icmp	Specifies ICMP.
ospf-as-external	Specifies OSPF as External routing.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPFInter routing.

ospf-intra	Specifies OSPFIntra routing.
ripng	Specifies RIP.
static	Specifies static routes.
vrname	Specifies a virtual router name.

Default

N/A

Usage Guidelines

Table 27 lists the default values that apply after you enter this command.

Table 27. Default Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPF3Intra	2200
OSPF3Inter	2300
IS-IS L1	2360
IS-IS L2	2370
RIPg	2400
OSPFv3 AExt	3100
OSPFv3 Extern1	3200
OSPFv3 Extern2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500

Example

The following command returns the IPv6 route priority for all route origins to the default values:

```
unconfigure iproute ipv6 priority all
```

unconfigure neighbor-discovery cache

```
unconfigure neighbor-discovery cache {vr <vr_name>}
```

Description

Resets the neighbor-discovery cache configuration parameters to their default values.

Syntax Description

vr_name	Specifies a virtual router.
---------	-----------------------------

Default

IPv6 neighbor timeout: 20 minutes

Maximum IPv6 neighbor entries: 1024

Maximum IPv6 neighbor pending entries: 1024

IPv6 neighbor refresh: Enabled

Usage Guidelines

None.

Example

The following command resets the neighbor-discovery cache configuration:

```
unconfigure neighbor-discovery cache
```

unconfigure vlan router-discovery

```
unconfigure vlan <vlan_name> router-discovery {ipv6}
```

Description

Unconfigures all the router-discovery parameters and resets them to their respective default values.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

Each of the router-discovery parameters is set to the default value. For example, the default-lifetime parameter is set to 1800 seconds. The default value for each of the router-discovery parameters is listed in the corresponding `configure vlan router-discovery` command description.

Example

The following command unconfigures all the router-discovery parameters for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery
```

unconfigure vlan router-discovery default-lifetime

```
unconfigure vlan <vlan_name> router-discovery {ipv6} default-lifetime
```

Description

Unconfigures the router lifetime value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the default-lifetime parameter to the default value of 1800 seconds.

Example

The following command unconfigures the default-lifetime for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery default-lifetime
```

unconfigure vlan router-discovery hop-limit

```
unconfigure vlan <vlan_name> router-discovery {ipv6} hop-limit
```

Description

Unconfigures the current hop limit value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the hop-limit parameter to the default value of 64.

Example

The following command unconfigures the current hop limit for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery hop-limit
```

unconfigure vlan router-discovery link-mtu

```
unconfigure vlan <vlan_name> router-discovery {ipv6} link-mtu
```

Description

Unconfigures the link MTU value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the link-mtu parameter to the default value of 0.

Example

The following command unconfigures the link MTU for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery link-mtu
```

unconfigure vlan router-discovery managed-config-flag

```
unconfigure vlan <vlan_name> router-discovery {ipv6} managed-config-flag
```

Description

Unconfigures the managed address configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the managed-config-flag parameter to the default value *off*.

Example

The following command unconfigures the managed address configuration flag for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery managed-config-flag
```

unconfigure vlan router-discovery max-interval

```
unconfigure vlan <vlan_name> router-discovery {ipv6} max-interval
```

Description

Unconfigures the maximum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the max-interval parameter to the default value of 600 seconds.

Example

The following command unconfigures the max-interval for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery max-interval
```

unconfigure vlan router-discovery min-interval

```
unconfigure vlan <vlan_name> router-discovery {ipv6} min-interval
```

Description

Unconfigures the minimum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the min-interval parameter to the default value of (max-interval * .33 seconds).

Example

The following command unconfigures the min-interval for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery min-interval
```

unconfigure vlan router-discovery other-config-flag

```
unconfigure vlan <vlan_name> router-discovery {ipv6} other-config-flag
```

Description

Unconfigures the other stateful configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the other-config-flag parameter to the default value *off*.

Example

The following command unconfigures the other stateful configuration flag for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery other-config-flag
```

unconfigure vlan router-discovery reachable-time

```
unconfigure vlan <vlan_name> router-discovery {ipv6} reachable-time
```

Description

Unconfigures the reachable time value in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the reachable-time parameter to the default value of 30,000 milliseconds.

Example

The following command unconfigures the reachable time for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery reachable-time
```

unconfigure vlan router-discovery retransmit-time

```
unconfigure vlan <vlan_name> router-discovery {ipv6} retransmit-time
```

Description

Unconfigures the retransmit time value in router discovery advertisements on the VLAN.

Syntax Description

vlan_name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the retransmit-time parameter to the default value of 1000 milliseconds.

Example

The following command unconfigures the retransmit time for the VLAN *top_floor*:

```
unconfigure vlan top_floor router-discovery retransmit-time
```

unconfigure tunnel

```
unconfigure tunnel <tunnel_name> ipaddress <ipv6_address_mask>
```

Description

Unconfigures an IPv6 address/prefix route from a tunnel.

Syntax Description

tunnel_name	Specifies an IPv6 tunnel.
ipv6_address_mask	Specifies an IPv6 address / IPv6 prefix length.

Default

N/A.

Usage Guidelines

This command will unconfigure an IPv6 address/prefix route from the specified tunnel.

Example

The following command unconfigures the 6in4 tunnel link39 with the address 3aaa::1111/64

```
unconfigure tunnel link39 3aaa::1111/64
```

This chapter describes commands used for the interior gateway protocol RIP.

Routing Information Protocol (RIP) is an Interior Gateway Protocol (IGP) first used in computer routing in the Advanced Research Projects Agency Network (ARPAnet) as early as 1969. It is primarily intended for use in homogeneous networks of moderate size.

To determine the best path to a distant network, a router using RIP always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

The routing table in a router using RIP contains an entry for every known destination network. Each routing table entry contains the following information:

- IP address of the destination network
- Metric (hop count) to the destination network
- IP address of the next router
- Timer that tracks the amount of time since the entry was last updated

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

A new version of RIP, called RIP version 2 (RIPv2), expands the functionality of RIP version 1 to include:

- Variable-Length Subnet Masks (VLSMs)
- Next-hop addresses
- Support for next-hop addresses allows for optimization of routes in certain environments
- Multicasting

If you are using RIP with supernetting/Classless Inter-Domain Routing (CIDR), you must use RIPv2 only, and RIP route aggregation must be turned off.

clear rip counters

```
clear rip counters
```

Description

Clears the RIP counters (statistics).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the RIP statistics counters:

```
clear rip counters
```

configure rip add vlan

```
configure rip add vlan [<vlan-name> | all]
```

Description

Configures RIP on an IP interface.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled, the parameters are not reset to default automatically.

Example

The following command configures RIP on the VLAN *finance*:

```
configure rip add finance
```

configure rip delete vlan

```
configure rip delete vlan [<vlan-name> | all]
```

Description

Disables RIP on an IP interface.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled by this command, the parameters are not reset to default automatically.

Example

The following command deletes RIP on a VLAN named *finance*:

```
configure rip delete finance
```

configure rip garbagetime

```
configure rip garbagetime {<seconds>}
```

Description

Configures the RIP garbage time.

Syntax Description

seconds	Specifies a time in seconds.
---------	------------------------------

Default

120 seconds.

Usage Guidelines

None.

Example

The following command configures the RIP garbage time to have a 60-second delay:

```
configure rip garbagetime 60
```

configure rip import-policy

```
configure rip import-policy [<policy-name> | none]
```

Description

Configures the import policy for RIP.

Syntax Description

policy-name	Specifies the policy.
-------------	-----------------------

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding RIP routes to the IP route table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the `none` option to remove an import policy.

Example

The following example applies the policy *campuseast* to RIP routes:

```
configure rip import-policy campuseast
```

configure rip routetimeout

```
configure rip routetimeout <seconds>
```

Description

Configures the route timeout period.

Syntax Description

seconds	Specifies a time in seconds.
---------	------------------------------

Default

180 seconds.

Usage Guidelines

If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

Example

The following example sets the route timeout period to 120 seconds:

```
configure rip routetimeout 120
```

configure rip updatetime

```
configure rip updatetime <seconds>
```

Description

Specifies the time interval in seconds within which RIP sends update packets.

Syntax Description

seconds	Specifies a time in seconds. The range is 10 to 180.
---------	--

Default

30 seconds.

Usage Guidelines

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). The timer granularity is 10 seconds. Timer minimum is 10 seconds and maximum is 180 seconds.

Example

The following command sets the update timer to 60 seconds:

```
configure rip updatetime 60
```

configure rip vlan cost

```
configure rip vlan [<vlan-name> | all] cost <cost>
```

Description

Configures the cost (metric) of the interface.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
cost	Specifies a cost metric.

Default

The default setting is 1.

Usage Guidelines

The specified interface cost is added to the cost of the route received through this interface.

Example

The following command configures the cost for the VLAN *finance* to a metric of 3:

```
configure rip vlan finance cost 3
```

configure rip vlan route-policy

```
configure rip vlan [<vlan-name> | all] route-policy [in | out] [<policy-name> | none]
```

Description

Configures RIP to ignore certain routes received from its neighbor, or to suppress certain routes when performing route advertisements.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
policy-name	Specifies a policy.
none	Removes any policy from the VLAN.

Default

N/A.

Usage Guidelines

Use the `in` option to configure an input route policy, which determines which RIP routes are accepted as valid routes. This policy can be combined with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors.

Use the `out` option to configure an output route policy, which determines which RIP routes are advertised on the VLAN.

Example

The following command configures the VLAN *backbone* to accept selected routes from the policy *nosales*:

```
configure rip vlan backbone route-policy in nosales
```

The following command uses the policy *nosales* to determine which RIP routes are advertised into the VLAN *backbone*:

```
configure rip vlan backbone route-policy out nosales
```

configure rip vlan rxmode

```
configure rip [vlan <vlan-name> | all] rxmode [none | v1only | v2only | any]
```

Description

Changes the RIP receive mode for one or all VLANs.

Syntax Description

none	Specifies to drop all received RIP packets.
v1only	Specifies to accept only RIP version 1 format packets.
v2only	Specifies to accept only RIP version 2 format packets.
any	Specifies to accept RIP version 1 and RIP version 2 packets.
vlan-name	Specifies to apply settings to specific VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the receive mode for the VLAN *finance* to accept only RIP version 1 format packets:

```
configure rip finance rxmode vlonly
```

configure rip vlan trusted-gateway

```
configure rip vlan [<vlan-name> | all] trusted-gateway [<policy-name> | none]
```

Description

Configures a trusted neighbor policy to determine trusted RIP router neighbors for the VLAN on the switch running RIP.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
policy-name	Specifies a policy.
none	Removes any trusted-gateway policy from the VLAN.

Default

N/A.

Usage Guidelines

Use this command to set a policy to determine trusted neighbors. A neighbor is defined by its IP address. Only the RIP control packets from trusted neighbors will be processed.

Example

The following command configures RIP to use the policy *nointernet* to determine from which RIP neighbor to receive (or reject) the routes to the VLAN *backbone*:

```
configure rip vlan backbone trusted-gateway nointernet
```

configure rip vlan txmode

```
configure rip [vlan <vlan-name> | all] txmode [none | vlonly | vlcomp | v2only]
```

Description

Changes the RIP transmission mode for one or all VLANs.

Syntax Description

none	Specifies to not transmit any packets on this interface.
v1only	Specifies to transmit RIP version 1 format packets to the broadcast address.
v1comp	Specifies to transmit RIP version 2 format packets to the broadcast address.
v2only	Specifies to transmit RIP version 2 format packets to the RIP multicast address.
vlan-name	Specifies to apply settings to a specific VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the transmit mode for the VLAN *finance* to transmit version 2 format packets to the broadcast address:

```
configure rip finance txmode vlcomp
```

disable rip

```
disable rip
```

Description

Disables RIP for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

RIP has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table

- Slow convergence
- Routing decisions based on hop count; no concept of link costs or delay
- Flat networks; no concept of areas or boundaries

Example

The following command disables RIP for the whole router:

```
disable rip
```

disable rip aggregation

```
disable rip aggregation
```

Description

Disables the RIP aggregation of subnet information on a RIP version 2 (RIPv2) router.

Syntax Description

This command has no arguments or variables.

Default

RIP aggregation is disabled by default.

Usage Guidelines

The disable RIP aggregation command disables the RIP aggregation of subnet information on a switch configured to send RIPv2-compatible traffic. The switch summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation:

- Within a class boundary, no routes are aggregated.
- If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary.

Example

The following command disables RIP aggregation on the interface:

```
disable rip aggregation
```

disable rip export

```
disable rip export [bgp | direct | e-bgp | i-bgp | ospf | ospf-extern1 | ospf-extern2 |  
ospf-inter | ospf-intra | static | isis ]
```

Description

Disables RIP from redistributing routes from other routing protocols.

Syntax Description

static	Specifies static routes.
bgp	Specifies BGP routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
e-bgp	Specifies external BGP routes.
i-bgp	Specifies internal BGP routes.
ospf	Specifies all OSPF routes.
ospf-intra	Specifies OSPF-intra area routes.
ospf-inter	Specifies OSPF-inter area routes.
ospf-extern1	Specifies OSPF external route type 1.
ospf-extern2	Specifies OSPF external route type 2.

Default

Disabled.

Usage Guidelines

This command disables the exporting of BGP, static, direct, and OSPF-learned routes into the RIP domain.

Example

The following command disables RIP from redistributing any routes learned from OSPF:

```
disable rip export ospf
```

disable rip originate-default

```
disable rip originate-default
```

Description

Disables the advertisement of a default route.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command unconfigures a default route to be advertised by RIP if no other default route is advertised:

```
disable rip originate-default
```

disable rip poisonreverse

```
disable rip poisonreverse
```

Description

Disables poison reverse algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Example

The following command disables the split horizon with poison reverse algorithm for RIP:

```
disable rip poisonreverse
```

disable rip splithorizon

```
disable rip splithorizon
```

Description

Disables the split horizon algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command disables the split horizon algorithm for RIP:

```
disable rip splithorizon
```

disable rip triggerupdates

```
disable rip triggerupdates
```

Description

Disables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

Example

The following command disables the trigger update mechanism:

```
disable rip triggerupdate
```

disable rip use-ip-router-alert

```
disable rip use-ip-router-alert
```

Description

Disables router alert IP option in outgoing RIP control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the RIP router alert IP option:

```
disable rip use-ip-router-alert
```

enable rip

```
enable rip
```

Description

Enables RIP for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

RIP has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table
- Slow convergence
- Routing decisions based on hop count; no concept of link costs or delay
- Flat networks; no concept of areas or boundaries

Example

The following command enables RIP for the whole router:

```
enable rip
```

enable rip aggregation

```
enable rip aggregation
```

Description

Enables the RIP aggregation of subnet information on a RIP version 2 (RIPv2) interface.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The enable (disable) rip aggregation command enables (disables) the RIP aggregation of subnet information on an interface configured to send RIPv1 or RIPv2-compatible traffic. The switch summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation:

- Subnet routes are aggregated to the nearest class network route when crossing a class boundary.
- Within a class boundary, no routes are aggregated.
- If aggregation is enabled, the behavior is the same as in RIPv1.
- If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary.

Example

The following command enables RIP aggregation on the interface:

```
enable rip aggregation
```

enable rip export

```
enable rip export [bgp | direct | e-bgp | i-bgp | ospf | ospf-extern1 | ospf-extern2 |  
ospf-inter | ospf-intra | static ] [cost <number> {tag <number>} | policy <policy-name>]
```

Description

Enables RIP to redistribute routes from other routing functions.

Syntax Description

bgp	Specifies BGP routes.
-----	-----------------------

<code>direct</code>	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
<code>e-bgp</code>	Specifies E-BGP routes.
<code>I-bgp</code>	Specifies I-BGP routes.
<code>ospf</code>	Specifies all OSPF routes.
<code>ospf-intra</code>	Specifies OSPF-intra area routes.
<code>ospf-inter</code>	Specifies OSPF-inter area routes.
<code>ospf-extern1</code>	Specifies OSPF external route type 1.
<code>ospf-extern2</code>	Specifies OSPF external route type 2.
<code>static</code>	Specifies static routes.
<code>cost <number></code>	Specifies the <code>cost</code> metric, from 0-15. If set to 0, RIP uses the route metric obtained from the route origin.
<code>tag <number></code>	Specifies a tag number.
<code><policy-name></code>	Specifies a policy.

Default

Disabled.

Usage Guidelines

This command enables the exporting of BGP, static, direct, and OSPF-learned routes into the RIP domain. You can choose which types of OSPF routes are injected, or you can simply choose `ospf`, which will inject all learned OSPF routes regardless of type.

The cost metric is inserted for all RIP-learned, static, and direct routes injected into RIP. If the cost metric is set to 0, the cost is inserted from the route. For example, with BGP, the cost could be the MED or the length of the BGP path. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag.

Each protocol can have a policy associated with it to control or modify the exported routes.

Example

The following command enables RIP to redistribute routes from all OSPF routes:

```
enable rip export ospf cost 0
```

enable rip originate-default cost

```
enable rip originate-default {always} cost <number> {tag<number>}
```

Description

Configures a default route to be advertised by RIP.

Syntax Description

<code>always</code>	Specifies to always advertise the default route.
<code>cost <number></code>	Specifies a cost metric. The range is 1 - 15.
<code>tag <number></code>	Specifies a tag number.

Default

Disabled.

Usage Guidelines

If `always` is specified, RIP always advertises the default route to its neighbors. If `always` is not specified, RIP advertises a default route only if a reachable default route is in the system route table.

The default route advertisement is filtered using the out policy.

The cost metric is inserted for all RIP-learned, static, and direct routes injected into RIP. The tag value is used only by special routing applications.

Example

The following command configures a default route to be advertised by RIP if there is a default route in the system routing table:

```
enable rip originate-default cost 7
```

enable rip poisonreverse

```
enable rip poisonreverse
```

Description

Enables poison reverse algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Example

The following command enables the split horizon with poison reverse algorithm for RIP:

```
enable rip poisonreverse
```

enable rip splithorizon

```
enable rip splithorizon
```

Description

Enables the split horizon algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command enables the split horizon algorithm for RIP:

```
enable rip splithorizon
```

enable rip triggerupdates

```
enable rip triggerupdates
```

Description

Enables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

Example

The following command enables the trigger update mechanism:

```
enable rip triggerupdate
```

enable rip use-ip-router-alert

```
enable rip use-ip-router-alert
```

Description

Enables the router alert IP option in the outgoing RIP control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the RIP router alert IP option:

```
enable rip use-ip-router-alert
```

show rip

```
show rip
```

Description

Displays RIP specific configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific configuration:

```
show rip
```

show rip interface

```
show rip interface {detail}
```

Description

Displays RIP-specific configuration and statistics for all VLANs.

Syntax Description

detail	Specifies detailed display.
--------	-----------------------------

Default

Show summary output for all interfaces.

Usage Guidelines

Summary includes the following information per interface:

- VLAN name
- IP address and mask
- interface status
- packets transmitted
- packets received
- number of triggered updates
- cost

Detail includes the following per interface:

- VLAN name
- IP address and mask
- tx mode
- rx mode

- cost
- peer information (for each peer)
 - age
 - version
 - received packets
 - received updates
 - received bad packets
 - received bad routes
- in policy
- out policy
- trusted gateway policy
- packets transmitted
- sent triggered updates
- packets received
- bad packets received
- bad routes received

Example

The following command displays the RIP configuration for all VLANs:

```
show rip interface
```

The following command displays RIP-specific statistics for all VLANs:

```
show rip interface detail
```

show rip interface vlan

```
show rip interface vlan <vlan-name>
```

Description

Displays RIP specific statistics and configuration for a VLAN in detail.

Syntax Description

vlan-name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific statistics for the VLAN *accounting*:

```
show rip interface accounting
```

show rip memory

```
show rip memory {detail | <memoryType>}
```

Description

Displays RIP specific memory usage.

Syntax Description

detail	Displays detail information.
memoryType	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific memory for all types:

```
show rip memory detail
```

show rip routes

```
show rip routes {detail} {network <ripNetworkPrefix>}
```

Description

Displays routes advertised by RIP.

Syntax Description

detail	Displays all available information from the RIP routing table.
ripNetworkPrefix	Specifies the route prefix for the routes to show.

Default

N/A.

Usage Guidelines

The routes displayed include all routes advertised by RIP, including routes exported from the system routing table and originated by other protocols, for example BGP.

Example

The following command displays a summary of RIP specific routes for the networks 10.0.0.0/8:

```
show rip routes network 10.0.0.0/8
```

unconfigure rip

```
unconfigure rip {vlan <vlan-name> | all}
```

Description

Resets all RIP parameters to the default for all VLANs or for the specified VLAN.

Syntax Description

vlan-name	Specifies a VLAN name.
-----------	------------------------

Default

All.

Usage Guidelines

Does not change the enable/disable state of the RIP settings.

Example

The following command resets the RIP configuration to the default for the VLAN *finance*:

```
unconfigure rip finance
```

This chapter describes commands used for the IPv6 interior gateway protocol RIPng.

To determine the best path to a distant network, a router using RIPng always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

The routing table in a router using RIPng contains an entry for every known destination network. Each routing table entry contains the following information:

- IP address and prefix length of the destination network
- Metric (hop count) to the destination network
- IP address of the next hop router, if the destination is not directly connected
- Interface for the next hop
- Timer that tracks the amount of time since the entry was last updated
- A flag that indicates if the entry is a new one since the last update
- The source of the route, for example, static, RIPng, OSPFv3, etc.

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

clear ripng counters

```
clear ripng counters {vlan <vlan-name> | tunnel <tunnel-name>}
```

Description

Clears the RIPng global or interface-specific counters (statistics).

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the RIPng statistics counters:

```
clear ripng counters
```

configure ripng add

```
configure ripng add [vlan <vlan-name> | tunnel <tunnel-name>
| [vlan | tunnel] all]
```

Description

Configures RIPng on an IP interface.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or tunnels.

Default

N/A.

Usage Guidelines

For RIPng to be active on the interface, it must also be globally enabled using the command `enable ripng`. If the keyword `all` is specified, all IPv6 configured VLANs or tunnels will be configured for RIPng.

Example

The following command configures RIPng on the VLAN *finance*:

```
configure ripng add finance
```

configure ripng cost

```
configure ripng [vlan <vlan-name> | tunnel <tunnel-name>] cost <metric>
```

Description

Configures the cost (metric) of the interface.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
metric	Specifies a cost metric. Range is 1 to 15.

Default

The default setting is 1.

Usage Guidelines

The specified interface cost is added to the cost of the route received through this interface.

Example

The following command configures the cost for the VLAN *finance* to a metric of 3:

```
configure ripng vlan finance cost 3
```

configure ripng delete

```
configure ripng delete [vlan <vlan-name> | tunnel <tunnel-name>
| [vlan | tunnel] all]
```

Description

Removes an interface from RIPng routing.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or tunnels.

Default

N/A.

Usage Guidelines

This command removes an interface from RIPng routing. However, the RIPng-specific interface configuration will be preserved, even if RIPng is unconfigured on the interface. The

interface configuration information is removed only when the IPv6 interface itself gets deleted by, for example, by unconfiguring all the IPv6 addresses on the interface.

Example

The following command removes the VLAN *finance* from RIPng routing:

```
configure ripng delete finance
```

configure ripng garbagemtime

```
configure ripng garbagemtime {<seconds>}
```

Description

Configures the RIPng garbage time.

Syntax Description

seconds	Specifies a time in seconds. Range is 10 to 2400 seconds.
---------	---

Default

120 seconds.

Usage Guidelines

This command configures the time interval after which a route in the RIPng routing database that has expired will be removed. The value is rounded off to nearest multiple of 10.

Example

The following command configures the RIPng garbage time to have a 60-second delay:

```
configure ripng garbagemtime 60
```

configure ripng import-policy

```
configure ripng import-policy [<policy-name> | none]
```

Description

Configures the import policy for RIPng.

Syntax Description

policy-name	Specifies the policy.
-------------	-----------------------

Default

No policy.

Usage Guidelines

Use this command to configure the policy to be applied to RIPng routes installed into the system routing table from the RIPng routing process. This policy can be used to modify parameters associated with routes installed into the routing table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the `none` option to remove the import policy.

The following is a sample policy file that can be used with RIPng. It changes the metric to 12 for any routes from the subnets 2001:db8:2ccc::/64 and 2001:db8:2ccd::/64:

```
entry filter_routes {
  if match any{
    nlri 2001:db8:2ccc:: /64;
    nlri 2001:db8:2ccd:: /64;
  }
  then {
    cost 12;
  }
}
```

Example

The following example applies the policy *campuseast* to RIPng routes:

```
configure ripng import-policy campuseast
```

configure ripng route-policy

```
configure ripng [vlan <vlan-name> | tunnel <tunnel-name>] route-policy [in | out]
[<policy-name> | none]
```

Description

Configures RIPng to ignore or modify certain routes received from its neighbors, or to suppress certain routes when performing route advertisements.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
policy-name	Specifies a policy.
none	Removes any policy from the VLAN.

Default

N/A.

Usage Guidelines

Use the `in` option to configure an input route policy, which determines which RIPng routes are accepted as valid routes from RIPng neighbors. This policy can be combined with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors.

Use the `out` option to configure an output route policy, which determines which RIPng routes are advertised to other RIPng neighbors.

The following is a sample policy file that could be used with RIPng. It will drop any routes from the subnets 2001:db8:2ccc::/64 and 2001:db8:2ccd::/64:

```

entry filter_routes {
    If match any{
        nlri 2001:db8:2ccc:: /64;
        nlri 2001:db8:2ccd:: /64;
    }
    then {
        deny;
    }
}

```

Example

The following command configures the VLAN *backbone* to accept routes from its neighbor as specified by the policy *nosales*:

```
configure ripng vlan backbone route-policy in nosales
```

The following command uses the policy *nosales* to determine which RIP routes are advertised into the VLAN *backbone*:

```
configure rip vlan backbone route-policy out nosales
```

configure ripng routetimeout

```
configure ripng routetimeout <seconds>
```

Description

Configures the route timeout period for RIPng.

Syntax Description

seconds	Specifies a time in seconds. Range is 10 to 3600.
---------	---

Default

180 seconds.

Usage Guidelines

If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

The configured value is rounded off to the nearest multiple of 10.

Example

The following example sets the route timeout period to 120 seconds:

```
configure ripng routetimeout 120
```

configure ripng trusted-gateway

```
configure ripng [vlan <vlan-name> | tunnel <tunnel-name>] trusted-gateway [<policy-name> | none]
```

Description

Configures a trusted neighbor policy to determine trusted RIPng router neighbors for the interfaces on the switch running RIPng.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
policy-name	Specifies a policy.
none	Removes any trusted-gateway policy from the VLAN.

Default

None. Control packets from all of the neighbors are processed.

Usage Guidelines

Use this command to set a policy to determine trusted neighbors. A neighbor is defined by its IP address. Only the RIPng control packets from trusted neighbors will be processed.

The following policy designates neighbors from the fe80:202:b3ff:fe4a:6ada:: /64 subnet and the neighbor at fe80:203::b3ff:fe4a:6ada as trusted gateways:

```
entry filter_gateways {
    If match any{
        nlri fe80:202:b3ff:fe4a:6ada:: /64;
```

```

        nlri fe80:203::b3ff:fe4a:6ada:: /64;
    }
    then {
        permit;
    }
}

```

Example

The following command configures RIPng to use the policy *nointernet* to determine from which RIPng neighbor to receive (or reject) the routes to the VLAN *backbone*:

```
configure ripng vlan backbone trusted-gateway nointernet
```

configure ripng updatetime

```
configure ripng updatetime <seconds>
```

Description

Specifies the time interval in seconds within which RIPng sends update packets.

Syntax Description

seconds	Specifies a time in seconds. The range is 10 to 3600.
---------	---

Default

30 seconds.

Usage Guidelines

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). The timer granularity is 10 seconds. Timer minimum is 10 second and maximum is 3600 seconds.

Example

The following command sets the update timer to 60 seconds:

```
configure ripng updatetime 60
```

disable ripng

```
disable ripng
```

Description

Disables RIPng for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables RIPng for the whole router:

```
disable ripng
```

disable ripng export

```
disable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter |  
ospfv3-intra | static ]
```

Description

Disables RIPng from redistributing routes from other routing protocols.

Syntax Description

static	Specifies user configured static routes.
direct	Specifies directly reachable subnets from the router (only interfaces that have IP forwarding enabled are exported).
ospfv3	Specifies all OSPFv3 routes.
ospfv3-intra	Specifies OSPFv3-intra area routes.
ospfv3-inter	Specifies OSPFv3-inter area routes.
ospfv3-extern1	Specifies OSPFv3 external route type 1.
ospfv3-extern2	Specifies OSPFv3 external route type 2.

Default

Disabled.

Usage Guidelines

This command disables the exporting of static, direct, IS-IS, and OSPF-learned routes from the switch routing table into the RIPng domain.

Example

The following command disables RIPng from redistributing any routes learned from OSPFv3:

```
disable ripng export ospfv3
```

disable ripng originate-default

```
disable ripng originate-default
```

Description

Disables the advertisement of a default route to the neighbors.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command unconfigures a default route to be advertised by RIPng if no other default route is advertised:

```
disable ripng originate-default
```

disable ripng poisonreverse

```
disable ripng poisonreverse {vlan <vlan-name> | tunnel <tunnel_name> | [vlan | tunnel] all}
```

Description

Disables poison reverse algorithm for RIPng.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Example

The following command disables the split horizon with poison reverse algorithm for RIPng:

```
disable ripng poisonreverse
```

disable ripng splithorizon

```
disable ripng splithorizon {vlan <vlan-name> | tunnel <tunnel_name> | [vlan | tunnel] all}
```

Description

Disables the split horizon algorithm for RIPng.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command disables the split horizon algorithm for RIPng:

```
disable rip splithorizon
```

disable ripng triggerupdate

```
disable ripng triggerupdate {vlan <vlan-name> | tunnel <tunnel_name> | [vlan | tunnel] all}
```

Description

Disables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIPng-related traffic.

When this feature is disabled, any metric change on the interface, or an interface going down will not be communicated until the next periodic update. To configure how often periodic updates are sent, use the following command:

```
configure ripng trusted-gateway
```

Example

The following command disables the trigger update mechanism:

```
disable ripng triggerupdate
```

enable ripng

```
enable ripng
```

Description

Enables RIPng for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Although RIPng is useful in small networks, it has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table
- Slow convergence
- Routing decisions based on hop count; no concept of link costs or delay
- Flat networks; no concept of areas or boundaries

For larger networks, consider OSPFv3 as an alternative IGP.

Example

The following command enables RIPng for the whole router:

```
enable ripng
```

enable ripng export

```
enable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter |  
ospfv3-intra | static] [cost <number> {tag <number>} | policy <policy-name>]
```

Description

Enables RIPng to redistribute routes from other routing functions.

Syntax Description

direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
ospfv3	Specifies all OSPFv3 routes.
ospfv3-intra	Specifies OSPFv3-intra area routes.
ospfv3-inter	Specifies OSPFv3-inter area routes.
ospfv3-extern1	Specifies OSPFv3 external route type 1.
ospfv3-extern2	Specifies OSPFv3 external route type 2.
static	Specifies static routes.
cost <number>	Specifies the <code>cost</code> metric, from 0-15. If set to 0, RIPng uses the route metric obtained from the route origin.
tag <number>	Specifies a tag number.
<policy-name>	Specifies a policy.

Default

Disabled. However, direct routes will always be advertised for all the interfaces where RIPng is enabled. For those interfaces where RIPng is not enabled, the corresponding direct route could be redistributed if direct route export is enabled through this command.

Default tag is 0.

Usage Guidelines

This command enables the exporting of static, direct, and OSPFv3-learned routes from the routing table into the RIPng domain. You can choose which types of OSPFv3 routes are injected, or you can simply choose `ospfv3`, which will inject all learned routes (of all types) for the selected protocol.

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into RIPng. If the cost metric is set to 0, the cost is inserted from the route table. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag.

Each protocol can have a policy associated with it to control or modify the exported routes. The following is sample policy file which modifies the cost of redistributed routes from OSPFv3 and statically configured routes:

```
entry filter_rt {
    If match any {
        Route-origin ospfv3;
        Route-origin static;
    }
    then {
        cost 10;
    }
}
```

Example

The following command enables RIPng to redistribute routes from all OSPFv3 routes:

```
enable ripng export ospfv3 cost 0
```

enable ripng originate-default

```
enable ripng originate-default {always} cost <metric> {tag <number>}
```

Description

Configures a default route to be advertised by RIPng.

Syntax Description

<code>always</code>	Specifies to advertise the default route in addition to learned default route.
<code>cost <metric></code>	Specifies a cost metric. The range is 1 - 15.
<code>tag <number></code>	Specifies a tag number.

Default

Disabled.

Usage Guidelines

If `always` is specified, RIPng always advertises the default route to its neighbors. If `always` is not specified, RIPng advertises a default route only if a reachable default route is in the system route table (the route is learned from other neighbors).

The default route advertisement is filtered using the out policy. Use the command, `configure ripng route-policy`, to specify the out policy.

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into RIPng. The tag value is used only by special routing applications.

Example

The following command configures a default route to be advertised by RIPng if there is a default route in the system routing table:

```
enable ripng originate-default cost 7
```

enable ripng poisonreverse

```
enable ripng poisonreverse {vlan <vlan-name> | tunnel <tunnel_name> | [vlan | tunnel] all}
```

Description

Enables the split horizon with poison reverse algorithm for RIPng on specified interfaces.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Used with split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

If both split horizon and poison reverse are enabled, poison reverse takes precedence.

Example

The following command enables split horizon with poison reverse for RIPng on all IPv6 interfaces in the virtual router:

```
enable ripng poisonreverse
```

The following command enables split horizon with poison reverse for all the IPv6 configured VLANs in the virtual router:

```
enable ripng poisonreverse vlan all
```

enable ripng splithorizon

```
enable ripng splithorizon {vlan <vlan-name> | tunnel <tunnel_name> | [vlan | tunnel] all}
```

Description

Enables the split horizon algorithm for RIPng.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command enables the split horizon algorithm for RIPng on all IPv6 configured interfaces:

```
enable ripng splithorizon
```

enable ripng triggerupdates

```
enable ripng triggerupdates {vlan <vlan-name> | tunnel <tunnel_name> | [vlan | tunnel] all}
```

Description

Enables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
-----------	------------------------------------

tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIPv6-related traffic.

Example

The following command enables the trigger update mechanism on all IPv6 configured interfaces:

```
enable ripng triggerupdate
```

show ripng

```
show ripng
```

Description

Displays RIPv6 global configuration and runtime information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIPv6 global configuration and runtime information:

```
show ripng
```

The following is sample output from this command:

```
Ripng Routing :Enabled
Triggered Updates: Enabled           Aggregation           : Disabled
```

```

Update Interval   : 30                Route Timeout      : 180
Garbage Timeout   : 120               Hold Down Time     : 0
Originate Default: Disabled
Sys Import-Policy: None
Redistribute:

```

Protocol	Status	Cost	Tag	Policy

Direct	Enabled	0	0	none
Static	Disabled	0	0	none
OSPFIntra	Disabled	0	0	none
OSPFInter	Disabled	0	0	none
OSPFExt1	Disabled	0	0	none
OSPFExt2	Disabled	0	0	none

show ripng interface

```
show ripng interface {detail | vlan <vlan-name> | tunnel <tunnel-name>}
```

Description

Displays RIPng-specific configuration and statistics for the specified interface.

Syntax Description

detail	Specifies detailed display.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

Show summary output for all interfaces.

Usage Guidelines

Displays the RIPng interface configuration and runtime information. If no interface is specified, only the summary data for all the configured interfaces is displayed. If an interface is specified, only the data for that interface is displayed in detail. If the keyword `detail` is specified, detailed data for all interfaces is displayed.

Example

The following command displays the RIPng configuration summary for all interfaces:

```
show ripng interface
```

The following is sample output from this command:

VLAN	IP Address	Flags	Sent	Rcvd	Triggered Cost
------	------------	-------	------	------	----------------

				Packets	Packets	Updates	
v1	22cc::3	/64	rif-pst	106349	106349	3	15
v2	22bb::1	/64	rif-pst	106349	106095	3	1
v3	2abc::1	/120	rif-pst	106351	0	4	1
v4	3ffe::1	/64	rif-pst	106349	139124	3	1

Flags: (f) Interface Forwarding Enabled, (i) Interface RIPng Enabled
 (n) Multinetted Interface, (r) Router RIPng Enabled
 (p) Poison Reverse Enabled, (s) Split Horizon Enabled
 (t) Triggerred Update Enabled.

The following command displays RIPng-specific statistics for the VLAN v1:

```
show ripng interface v1
```

The following is sample output from this command:

```
VLAN           : v1           Interface       : 22cc::3/64
Router RIPng    : Enabled      Cost            : 15
Input Policy    : None        Output Policy   : None
Trusted GW Policy : gw6       Poison Reverse  : Enabled
Split Horizon   : Enabled     Triggerred Updates : Enabled
Rcvd Packets    : 106358      Sent Packets    : 106358
Sent Trig. Updates : 3        Rcvd Bad Packets : 0
Rcvd Bad Routes : 0
```

```
Neighbor Addresses : fe80::201:30ff:fe94:f400
```

```
Interface Addresses : 22cc::3/64, fe80::280:c8ff:feb9:2855/64
```

show ripng routes

```
show ripng routes {detail} {network <ripngNetworkPrefix>}
```

Description

Displays all matching routes in the RIPng routing database.

Syntax Description

detail	Displays all available information from the RIPng routing table.
ipv6-prefix	Specifies the route prefix for the routes to show.
prefix-length	Specifies the address mask of the IPv6 prefix.

Default

N/A.

Usage Guidelines

The routes displayed include all routes advertised by RIPng, including routes exported from the system routing table and originated by other protocols, for example OSPFv3 (also called redistributed routes).

Example

The following command displays a summary of RIPng specific routes:

```
show ripng routes
```

The following is sample output from this command:

Network	Next Hop	Mtr	VLAN
*> 2aaa::/64	fe80::201:30ff:fe4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	2	v2
*> 2bbb::/64	fe80::201:30ff:fe4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	3	v2
*> 2ccc::/64	(local)	1	(direct)
*	fe80::201:30ff:fe4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	3	v2
*> 2ddd::/64	(local)	1	(direct)
*	fe80::201:30ff:fe94:f400%v2	2	v2

The following command displays the detailed RIPng route information:

```
show ripng routes detail
```

The following is sample output from this command:

```
IPv6 RIPng routing table entry for 2aaa::/64
Paths: (2 available, best #1)
  fe80::201:30ff:fe4:5ca0%v1 from fe80::201:30ff:fe4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid, best
  fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 2, tag 0, timeout in 02:44, valid

IPv6 RIPng routing table entry for 2bbb::/64
Paths: (2 available, best #1)
  fe80::201:30ff:fe4:5ca0%v1 from fe80::201:30ff:fe4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid, best
  fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 3, tag 0, timeout in 02:44, valid

IPv6 RIPng routing table entry for 2ccc::/64
Paths: (3 available, best #1)
  Local from direct
    Metric 1, tag 0, no timeout, valid, best
  fe80::201:30ff:fe4:5ca0%v1 from fe80::201:30ff:fe4:5ca0%v1 (v1)
```



```
Metric 2, tag 0, timeout in 02:38, valid
fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
Metric 3, tag 0, timeout in 02:44, valid
```

unconfigure ripng

```
unconfigure ripng {vlan <vlan-name> | tunnel <tunnel-name> | vlan all | tunnel all}
```

Description

Resets RIPng parameters to the default value.

Syntax Description

vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies either all IPv6 configured VLANs or all IPv6 tunnels.

Default

N/A.

Usage Guidelines

Issuing the command `unconfigure ripng` resets all the interfaces and the global configuration to the defaults, and disables RIPng, as that is the default.

Example

The following command resets the RIPng configuration to the default for the VLAN *finance*:

```
unconfigure rip finance
```

OSPF Commands

23

This chapter describes commands used for the interior gateway protocol OSPF.

Open Shortest Path First (OSPF) is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an *autonomous system* (AS). In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router has an identical database maintained from the perspective of that router.

From the link-state database (LSDB), each router constructs a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can be distributed among them. The cost of a route is described by a single metric.

OSPF allows parts of a network to be grouped together into *areas*. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in link-state advertisement (LSA) traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPF are as follows:

- **Internal Router (IR)**—An internal router has all of its interfaces within the same area.
- **Area Border Router (ABR)**—An ABR has interfaces belonging to two or more areas. It is responsible for exchanging summary advertisements with other ABRs.
- **Autonomous System Border Router (ASBR)**—An ASBR acts as a gateway between OSPF and other routing protocols, or other autonomous systems.

Each switch that is configured to run OSPF must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPF, instead of having the switch automatically choose its router ID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older LSDB remaining in use.

Note: Do not set the router ID to 0.0.0.0.

For information about licensing requirements, see Appendix A in the *NETGEAR 8800 User Manual*.

OSPF Edge Mode

OSPF Edge Mode is a subset of OSPF available on platforms with an Advanced Core license. There are two restrictions on OSPF Edge Mode:

- At most, four Active OSPF VLAN interfaces are permitted. There is no restriction on the number of Passive interfaces.
- The OSPF Priority on VLANs is zero, and is not configurable. This prevents the system from acting as a DR or BDR.

clear ospf counters

```
clear ospf counters
{ interfaces [all | vlan <vlan-name> | area <area-identifier>]
| area [all | <area-identifier>]
| virtual-link [all | <router-identifier> <area-identifier>]
| neighbor [all | routerid [<ip-address> {<ip-mask> | <ipNetmask>]
| vlan <vlan-name>]| system}
```

Description

Clears the OSPF counters (statistics).

Syntax Description

vlan-name	Specifies a VLAN name.
router-identifier	Specifies a router interface number.
area-identifier	Specifies an OSPF area.
ip-address	Specifies an IP address
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask
system	Specifies the OSPF system counters.

Default

N/A.

Usage Guidelines

The global command `clear counters` also clears all OSPF counters. This global command is the equivalent of `clear ospf counters` for OSPF.

Example

The following command clears the OSPF counters for area 1.1.1.1:

```
clear ospf counters area 1.1.1.1
```

configure ospf add virtual-link

```
configure ospf add virtual-link <router-identifier> <area-identifier>
```

Description

Adds a virtual link connected to another ABR.

Syntax Description

router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- router-identifier—Far-end router interface number.
- area-identifier—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command configures a virtual link between the two interfaces:

```
configure ospf add virtual-link 10.1.2.1 10.1.0.0
```

configure ospf add vlan area

```
configure ospf add vlan [<vlan-name> | all] area <area-identifier> {passive}
```

Description

Enables OSPF on one or all VLANs (router interfaces).

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
area-identifier	Specifies the area to which the VLAN is assigned.

passive	Specifies to stop sending and receiving hello packets on this interface.
---------	--

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables OSPF on a VLAN named *accounting*:

```
configure ospf add vlan accounting area 0.0.0.1
```

configure ospf add vlan area link-type

```
configure ospf add vlan <vlan-name> area <area-identifier> link-type [auto | broadcast |  
point-to-point] {passive}
```

Description

Configures the OSPF link type.

Syntax Description

vlan-name	Specifies a VLAN name.
area-identifier	Specifies the area to which the VLAN is assigned.
auto	Specifies to automatically determine the OSPF link type based on the interface type.
broadcast	Specifies a broadcast link, such as Ethernet. Routers must elect a DR and a BDR during synchronization.
point-to-point	Specifies a point-to-point link type, such as PPP.
passive	Specifies to stop sending and receiving packets on this interface.

Default

Auto.

Usage Guidelines

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Example

The following command configures the OSPF link type as automatic on a VLAN named *accounting*:

```
configure ospf add vlan accounting area 0.0.0.1 link-type auto
```

configure ospf area external-filter

```
configure ospf area <area-identifier> external-filter [<policy-map> | none]
```

Description

Configures an external filter policy.

Syntax Description

area-identifier	Specifies the OSPF target area.
policy-map	Specifies a policy.
none	Specifies not to apply an external filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPF areas (an ABR function), a policy can be applied to an OSPF area that filters a set of OSPF external routes from being advertised into that area.

Using the none mode specifies that no external filter is applied.

Example

The following command configures an external filter policy, *nosales*:

```
configure ospf area 1.2.3.4 external-filter nosales
```

configure ospf area interarea-filter

```
configure ospf area <area-identifier> interarea-filter [<policy-map> | none]
```

Description

Configures a global inter-area filter policy.

Syntax Description

area-identifier	Specifies the OSPF target area.
policy-map	Specifies a policy.
none	Specifies not to apply an interarea filter.

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPF areas (an ABR function), a policy can be applied to an OSPF area that filters a set of OSPF inter-area routes from being sourced from any other areas.

Example

The following command configures an inter-area filter policy, *nosales*:

```
configure ospf area 0.0.0.6 interarea-filter nosales
```

configure ospf area add range

```
configure ospf area <area-identifier> add range [<ip-address> <ip-mask> | <ipNetmask>]  
[advertise | noadvert] [type-3 | type-7]
```

Description

Configures a range of IP addresses in an OSPF area to be aggregated.

Syntax Description

area-identifier	Specifies an OSPF area.
ip-address	Specifies an IP address
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask.
advertise	Specifies to advertise the aggregated range of IP addresses.
noadvertise	Specifies not to advertise the aggregated range of IP addresses.
type-3	Specifies type 3 LSA, summary LSA.
type-7	Specifies type 7 LSA, NSSA external LSA.

Default

N/A.

Usage Guidelines

If advertised, the aggregated IP range is exported as a single LSA by the ABR.

Example

The following command is used to summarize a certain range of IP addresses within an area and export them out as a single address:

```
configure ospf area 1.2.3.4 add range 10.1.2.0/24 advertise type-3
```

configure ospf area delete range

```
configure ospf area <area-identifier> delete range [<ip-address> <ip-mask> | <ipNetmask>]
```

Description

Deletes a range of aggregated IP addresses in an OSPF area.

Syntax Description

area-identifier	Specifies an OSPF area.
ip-address	Specifies an IP address.
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes an aggregated IP address range:

```
configure ospf area 1.2.3.4 delete range 10.1.2.0/24
```

configure ospf area normal

```
configure ospf area <area-identifier> normal
```


Description

Configures an OSPF area as a normal area.

Syntax Description

area-identifier	Specifies an OSPF area.
-----------------	-------------------------

Default

Normal.

Usage Guidelines

A normal area is an area that is not any of the following:

- Stub area
- NSSA

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Example

The following command configures an OSPF area as a normal area:

```
configure ospf area 10.1.0.0 normal
```

configure ospf area nssa stub-default-cost

```
configure ospf area <area-identifier> nssa [summary | nosummary] stub-default-cost <cost>
{translate}
```

Description

Configures an OSPF area as an NSSA.

Syntax Description

area-identifier	Specifies an OSPF area.
summary	Specifies that type-3 can be propagated into the area.
nosummary	Specifies that type-3 cannot be propagated into the area.
cost	Specifies a cost metric.
translate	Specifies whether type-7 LSAs are translated into type-5 LSAs.

Default

N/A.

Usage Guidelines

NSSAs are similar to the existing OSPF stub area configuration option, but have the following two additional capabilities:

- External routes originating from an ASBR connected to the NSSA can be advertised within the NSSA.
- External routes originating from the NSSA can be propagated to other areas, including the backbone area, if translated to type 5 LSAs.

When configuring an OSPF area as an NSSA, the translate option should only be used on NSSA border routers, where translation is to be enforced. If translate is not used on any NSSA border router in a NSSA, one of the ABRs for that NSSA is elected to perform translation (as indicated in the NSSA specification). The option should not be used on NSSA internal routers. Doing so inhibits correct operation of the election algorithm.

Example

The following command configures an OSPF area as an NSSA:

```
configure ospf area 10.1.1.0 nssa summary stub-default-cost 10 translate
```

configure ospf area stub stub-default-cost

```
configure ospf area <area-identifier> stub [summary | nosummary] stub-default-cost <cost>
```

Description

Configures an OSPF area as a stub area.

Syntax Description

area-identifier	Specifies an OSPF area.
summary	Specifies that type-3 can be propagated into the area.
nosummary	Specifies that type-3 cannot be propagated into the area.
cost	Specifies a cost metric.

Default

N/A.

Usage Guidelines

A stub area is connected to only one other area. The area that connects to a stub area can be the backbone area. External route information is not distributed into stub areas. Stub areas are used to reduce memory and computation requirements on OSPF routers.

Example

The following command configures an OSPF area as a stub area:

```
configure ospf area 0.0.0.6 stub nosummary stub-default-cost 10
```

configure ospf area timer

```
configure ospf area <area-identifier> timer <retransmit-interval> <transit-delay>
<hello-interval> <dead-interval> {<wait-timer-interval>}
```

Description

Configures the timers for all interfaces in the same OSPF area.

Syntax Description

area-identifier	Specifies an OSPF area.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1- 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.
wait-timer-interval	Specifies the interval between the interface coming up and the election of the DR and BDR. Usually equal to the dead timer interval.

Default

- retransmit interval—Default: 5
- transit delay—Default: 1
- hello interval—Default: 10
- dead interval—Default: 40
- wait timer interval—Default: dead interval

Usage Guidelines

Configuring OSPF timers on a per-area basis is a shorthand for applying the timers and authentication to each VLAN in the area at the time of configuration. If you add more VLANs to the area, you must configure the timers and authentication for the new VLANs explicitly.

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.
- wait timer interval—This interval is required by the OSPF standard to be equal to the router dead interval. Under some circumstances, setting the wait interval to smaller values can help OSPF routers on a broadcast network to synchronize more quickly at the expense of possibly electing an incorrect DR or BDR. This value should not be set to less than the hello interval. The default value is equal to the router dead interval.

Example

The following command sets the timers in area 0.0.0.2:

```
configure ospf area 0.0.0.2 timer 10 1 20 200
```

configure ospf ase-limit

```
configure ospf ase-limit <number> {timeout <seconds>}
```

Description

Configures the AS-external LSA limit and overflow duration associated with OSPF database overflow handling.

Syntax Description

number	Specifies the number of external routes that can be held in a link-state database.
seconds	Specifies a duration for which the system has to remain in the overflow state.

Default

The default for timeout is 0, which indicates that once the router goes into overflow state, it stays there until OSPF is disabled and then re-enabled.

Usage Guidelines

None.

Example

The following command configures the AS-external LSA limit and overflow duration:

```
configure ospf ase-limit 50000 timeout 1800
```

configure ospf ase-summary add

```
configure ospf ase-summary add [<ip-address> <ip-mask> | <ipNetmask>] cost <cost>
{tag <number>}
```

Description

Aggregates AS-external routes in a specified address range.

Syntax Description

ip-address	Specifies an IP address.
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask.
cost	Specifies a metric that will be given to the summarized route.
tag	Specifies an OSPF external route tag.

Default

N/A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command summarizes AS-external routes:

```
configure ospf ase-summary add 175.1.0.0/16 cost 10
```

configure ospf ase-summary delete

```
configure ospf ase-summary delete [<ip-address> <ip-mask> | <ipNetmask>]
```

Description

Deletes an aggregated OSPF external route.

Syntax Description

ip-address	Specifies an IP address.
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask.

Default

N/A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command deletes the aggregated AS-external route:

```
configure ospf ase-summary delete 175.1.0.0/16
```

configure ospf authentication

```
configure ospf [vlan [<vlan-name> | all] | area <area-identifier> | virtual-link
<router-identifier> <area-identifier>] authentication [{encrypted} simple-password
<simple-password> | {encrypted} md5 <md5_key_id> <md5_key>| none]
```

Description

Specifies the authentication password (up to eight characters) or Message Digest 5 (MD5) key for one or all interfaces in a specific area or a virtual link.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs
area-identifier	Specifies an OSPF area.
router-identifier	Specifies the router ID of the remote router.
encrypted	Indicates that the password (or key) is already encrypted (do not use this option).
simple-password	Specifies an authentication password (up to 8 ASCII characters).
md5-key_id	Specifies a Message Digest 5 key, from 0-255.
md5_key	Specifies a numeric value from 0-65,536. Can also be alphanumeric, up to 26 characters.
none	Disables authentication.

Default

N/A.

Usage Guidelines

The `md5_key` is a numeric value with the range 0 to 65,536 or alphanumeric. When the OSPF area is specified, authentication information is applied to all OSPF interfaces within the area.

The `encrypted` option is used by the switch when generating a configuration file and when parsing a switch-generated configuration file. Do not select the `encrypted` option in the CLI.

Example

The following command configures MD5 authentication on the VLAN *subnet_26*:

```
configure ospf vlan subnet_26 authentication md5 32 test
```

configure ospf cost

```
configure ospf [area <area-identifier> | vlan [<vlan-name> | all]] cost [automatic | <cost>]
```

Description

Configures the cost metric of one or all interface(s) or an area.

Syntax Description

area-identifier	Specifies an OSPF area.
vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
automatic	Determine the advertised cost from the OSPF metric table.
cost	Specifies the cost metric.

Default

The default cost is automatic.

Usage Guidelines

The range is 1 through 65535.

Example

The following command configures the cost metric of the VLAN *accounting*:

```
configure ospf vlan accounting cost 10
```

configure ospf delete virtual-link

```
configure ospf delete virtual-link <router-identifier> <area-identifier>
```

Description

Removes a virtual link.

Syntax Description

router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a virtual link:

```
configure ospf delete virtual-link 10.1.2.1 10.1.0.0
```

configure ospf delete vlan

```
configure ospf delete vlan [<vlan-name> | all]
```

Description

Disables OSPF on one or all VLANs (router interfaces).

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPF on VLAN *accounting*:

```
configure ospf delete vlan accounting
```

configure ospf import-policy

```
configure ospf import-policy [<policy-map> | none]
```

Description

Configures the import policy for OSPF.

Syntax Description

policy-map	Specifies the policy.
------------	-----------------------

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding OSPF routes to the IP route table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the `none` option to remove an import policy.

Example

The following example applies the policy *campuseast* to OSPF routes:

```
configure ospf import-policy campuseast
```

configure ospf lsa-batch-interval

```
configure ospf lsa-batch-interval <seconds>
```

Description

Configures the OSPF LSA batching interval.

Syntax Description

seconds	Specifies a time in seconds.
---------	------------------------------

Default

The default setting is 30 seconds.

Usage Guidelines

The range is between 0 (disabled) and 600 seconds, using multiples of 5 seconds. The LSAs added to the LSDB during the interval are batched together for refresh or timeout.

Example

The following command configures the OSPF LSA batch interval to a value of 100 seconds:

```
configure ospf lsa-batch-interval 100
```

configure ospf metric-table

```
configure ospf metric-table 10M <cost_10m> 100M <cost_100m> 1G <cost_1g> {10G <cost_10g>}
```

Description

Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces, and optionally, the 10 Gbps interface.

Syntax Description

cost	Specifies the interface cost for the indicated interfaces.
------	--

Default

- 10 Mbps—The default cost is 10.
- 100 Mbps—The default cost is 5.
- 1 Gbps—The default cost is 4.
- 10 Gbps—The default cost is 2.

Usage Guidelines

None.

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces:

```
configure ospf metric-table 10m 20 100m 10 1g 2
```

configure ospf priority

```
configure ospf [area <area-identifier> | vlan [<vlan-name> | all]] priority <priority>
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for one or all OSPF interface(s) or for all the interfaces within the area.

Syntax Description

area-identifier	Specifies an OSPF area.
vlan-name	Specifies a VLAN name.
all	Specifies all VLANs.
priority	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

The range is 0 through 255, and the default setting is 1. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets all the interfaces in area 1.2.3.4 to not be selected as the designated router:

```
configure ospf area 1.2.3.4 priority 0
```

configure ospf restart

```
configure ospf restart [none | planned | unplanned | both]
```

Description

Configures the router as a graceful OSPF restart router.

Syntax Description

none	Do not act as a graceful OSPF restart router.
planned	Only act as a graceful OSPF restart router for planned restarts.
unplanned	Only act as a graceful OSPF restart router for unplanned restarts.
both	Act as a graceful OSPF restart router for both planned and unplanned restarts.

Default

The default is none.

Usage Guidelines

This command configures the router as a graceful OSPF router. When configured for planned restarts, it will advertise Grace-LSAs before restarting (for example, during an upgrade of the OSPF module). When configured for unplanned restarts, it will advertise Grace-LSAs after restarting but before sending any Hellos. When configured for both, the router will advertise restarting regardless of whether the restart was planned or unplanned.

Example

The following command configures a router to perform graceful OSPF restarts only for planned restarts:

```
configure ospf restart planned
```

configure ospf restart grace-period

```
configure ospf restart grace-period <seconds>
```

Description

Configures the grace period sent out in Grace-LSAs and used by a restarting router.

Syntax Description

seconds	Grace period, in seconds. The default value is 120 seconds. Range is 1 to 1800 seconds.
---------	---

Default

The default is 120 seconds.

Usage Guidelines

This command configures the grace period sent out to helper neighbor routers and used by the restarting router. The value of the grace period must be greater than the dead interval, and less than the LSA refresh time.

Example

The following command configures a router to send LSAs with a 240 second grace period during graceful OSPF restarts:

```
configure ospf restart grace-period 240
```

configure ospf restart-helper

```
configure ospf [vlan [all | <vlan-name>] | area <area-identifier> | virtual-link
<router-identifier> <area-identifier>] restart-helper [none | planned | unplanned | both]
```

Description

Configures the router as a graceful OSPF restart helper router.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs
area-identifier	Specifies an OSPF area.
router-identifier	Specifies the router ID of the remote router of the virtual link.
none	Do not act as a graceful OSPF restart helper router.
planned	Only act as a graceful OSPF restart helper router for planned restarts.
unplanned	Only act as a graceful OSPF restart helper router for unplanned restarts.
both	Act as a graceful OSPF restart helper router for both planned and unplanned restarts.

Default

The router default is none.

Usage Guidelines

This command configures the router as a graceful OSPF restart helper router for a single or multiple routers. When the router is acting as a helper, it will continue to advertise the restarting router as if it was fully adjacent.

One OSPF interface may not help more than one restarting router. An OSPF interface may not enter helper mode when the router is performing a graceful restart. All the interfaces to a neighbor router must be configured as graceful restart helpers, or the router will not support graceful restart for its neighbor.

Example

The following command configures a router to be a graceful OSPF helper router for planned restarts for all routers in area 10.20.30.40:

```
configure ospf area 10.20.30.40 restart-helper planned
```

configure ospf routerid

```
configure ospf routerid [automatic | <router-identifier>]
```

Description

Configures the OSPF router ID. If automatic is specified, the switch uses the highest IP interface address as the OSPF router ID.

Syntax Description

automatic	Specifies to use automatic addressing.
router-identifier	Specifies a router address.

Default

Automatic.

Usage Guidelines

Each switch that is configured to run OSPF must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPF, instead of having the switch automatically choose its router ID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older link-state database remaining in use.

Note: Do not set the router ID to 0.0.0.0.

Example

The following command sets the router ID:

```
configure ospf routerid 10.1.6.1
```

configure ospf spf-hold-time

```
configure ospf spf-hold-time <seconds>
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

seconds	Specifies a time in seconds. The range is 0 to 300 seconds.
---------	---

Default

3 seconds.

Usage Guidelines

None.

Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
configure ospf spf-hold-time 6
```

configure ospf virtual-link timer

```
configure ospf virtual-link <router-identifier> <area-identifier> timer <retransmit-interval>
<transit-delay> <hello-interval> <dead-interval>
```

Description

Configures the timers for a virtual link.

Syntax Description

router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies an OSPF area.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.

Default

- retransmit interval—Default: 5
- transit delay—Default: 1
- hello interval—Default: 10
- dead interval—Default: 40
- wait timer interval—Default: dead interval

Usage Guidelines

Configuring OSPF timers on a per-area basis is a shorthand for applying the timers and authentication to each VLAN in the area at the time of configuration. If you add more VLANs to the area, you must configure the timers and authentication for the new VLANs explicitly.

Example

The following command sets the timers on the virtual link in area 0.0.0.2 and remote router ID 6.6.6.6:

```
configure ospf virtual-link 6.6.6.6 0.0.0.2 timer 10 1 20 200
```

configure ospf vlan area

```
configure ospf vlan <vlan-name> area <area-identifier>
```

Description

Associates a VLAN (router interface) with an OSPF area. By default, all router interfaces are associated with area 0.0.0.0.

Syntax Description

vlan-name	Specifies a VLAN name.
area-identifier	Specifies an OSPF area.

Default

Area 0.0.0.0

Usage Guidelines

Any OSPF network that contains more than one area is required to have an area configured as area 0, also called the *backbone*. All areas in an autonomous system must be connected to the backbone. When designing networks, you should start with area 0, and then expand into other areas.

The backbone allows summary information to be exchanged between ABRs. Every ABR hears the area summaries from all other ABRs. The ABR then forms a picture of the distance to all networks outside of its area by examining the collected advertisements, and adding in the backbone distance to each advertising router.

When a VLAN is configured to run OSPF, by default you must assign it to an area.

Example

The following command associates the VLAN *accounting* with an OSPF area:

```
configure ospf vlan accounting area 0.0.0.6
```


configure ospf vlan neighbor add

```
configure ospf vlan <vlan-name> neighbor add <ip-address>
```

Description

Configures the IP address of a point-to-point neighbor.

Syntax Description

vlan-name	Specifies a VLAN name.
ip-address	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the IP address of a point-to-point neighbor:

```
configure ospf vlan accounting neighbor add 10.0.0.1
```

configure ospf vlan neighbor delete

```
configure ospf vlan <vlan-name> neighbor delete <ip-address>
```

Description

Deletes the IP address of a point-to-point neighbor.

Syntax Description

vlan-name	Specifies a VLAN name.
ip-address	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the IP address of a point-to-point neighbor:

```
configure ospf vlan accounting neighbor delete 10.0.0.1
```

configure ospf vlan timer

```
configure ospf vlan [<vlan-name> | all] timer <retransmit-interval> <transit-delay>
<hello-interval> <dead-interval> {<wait-timer-interval>}
```

Description

Configures the OSPF wait interval for a VLAN or all VLANs.

Syntax Description

vlan-name	Specifies a VLAN name.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647.
wait-timer-interval	Specifies the interval between the interface coming up and the election of the DR and BDR. Usually equal to the dead timer interval.

Default

- retransmit interval—5 seconds.
- transit delay—1 second.
- hello interval—10 seconds.
- dead interval—40 seconds.
- wait timer interval—dead interval.

Usage Guidelines

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.

- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.
- wait timer interval—This interval is required by the OSPF standard to be equal to the router dead interval. Under some circumstances, setting the wait interval to smaller values can help OSPF routers on a broadcast network to synchronize more quickly at the expense of possibly electing an incorrect DR or BDR. This value should not be set to less than the hello interval. The default value is equal to the router dead interval.

Example

The following command configures the OSPF wait interval on the VLAN *accounting*:

```
configure ospf vlan accounting timer 10 15 20 60 60
```

create ospf area

```
create ospf area <area-identifier>
```

Description

Creates an OSPF area.

Syntax Description

area-identifier	Specifies an OSPF area.
-----------------	-------------------------

Default

Area 0.0.0.0

Usage Guidelines

Area 0.0.0.0 does not need to be created. It exists by default.

Example

The following command creates an OSPF area:

```
create ospf area 1.2.3.4
```

delete ospf area

```
delete ospf area [<area-identifier> | all]
```

Description

Deletes an OSPF area or all OSPF areas.

Syntax Description

area-identifier	Specifies an OSPF area.
all	Specifies all areas.

Default

N/A.

Usage Guidelines

An OSPF area cannot be deleted if it has an associated interface. Also, area 0.0.0.0 cannot be deleted.

Example

The following command deletes an OSPF area:

```
delete ospf area 1.2.3.4
```

disable ospf

```
disable ospf
```

Description

Disables the OSPF process for the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables the OSPF process for the router:

```
disable ospf
```

disable ospf capability opaque-lsa

```
disable ospf capability opaque-lsa
```

Description

Disables opaque LSAs across the entire system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Opaque LSAs are a generic OSPF mechanism used to carry auxiliary information in the OSPF database. Opaque LSAs are most commonly used to support OSPF traffic engineering.

Normally, support for opaque LSAs is auto-negotiated between OSPF neighbors. In the event that you experience interoperability problems, you can disable opaque LSAs.

If your network uses opaque LSAs, all routers on your OSPF network should support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.

For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

Example

The following command disables opaque LSAs across the entire system:

```
disable ospf capability opaque-lsa
```

disable ospf export

```
disable ospf export [bgp | direct | e-bgp | i-bgp | rip | static]
```

Description

Disables redistribution of routes to OSPF.

Syntax Description

bgp	Specifies BGP routes.
-----	-----------------------

direct	Specifies direct routes.
i-bgp	Specifies I-BGP routes.
e-bgp	Specifies E-BGP routes.
rip	Specifies RIP routes.
static	Specifies static routes.

Default

The default setting is disabled.

Usage Guidelines

Use this command to stop OSPF from exporting routes derived from other protocols.

Example

The following command disables OSPF to export BGP-related routes to other OSPF routers:

```
disable ospf export bgp
```

disable ospf originate-default

```
disable ospf originate-default
```

Description

Disables the generation of a default external LSA.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables generating a default external LSA:

```
disable ospf originate-default
```

disable ospf restart-helper-lsa-check

```
disable ospf [vlan [all | <vlan-name>] | area <area-identifier> | virtual-link
<router-identifier> <area-identifier>] restart-helper-lsa-check
```

Description

Disables the restart helper router from terminating graceful OSPF restart when received LSAs would affect the restarting router.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs
area-identifier	Specifies an OSPF area.
router-identifier	Specifies the router ID of the remote router of the virtual link.

Default

The default is enabled.

Usage Guidelines

This command disables the restart helper router from terminating graceful OSPF restart when received LSAs would affect the restarting router.

Example

The following command disables a router from terminating graceful OSPF restart for all routers in area 10.20.30.40 if it receives an LSA that would affect routing:

```
disable ospf area 10.20.30.40 restart-helper-lsa-check
```

disable ospf use-ip-router-alert

```
disable ospf use-ip-router-alert
```

Description

Disables the router alert IP option in outgoing OSPF control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the OSPF router alert IP option:

```
disable ospf use-ip-router-alert
```

enable ospf

```
enable ospf
```

Description

Enables the OSPF process for the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command enables the OSPF process for the router:

```
enable ospf
```

enable ospf capability opaque-lsa

```
enable ospf capability opaque-lsa
```

Description

Enables opaque LSAs across the entire system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Opaque LSAs are a generic OSPF mechanism used to carry auxiliary information in the OSPF database. Opaque LSAs are most commonly used to support OSPF traffic engineering.

Normally, support for opaque LSAs is auto-negotiated between OSPF neighbors. In the event that you experience interoperability problems, you can disable opaque LSAs.

If your network uses opaque LSAs, all routers on your OSPF network should support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.

For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

Example

The following command enables opaque LSAs across the entire system:

```
enable ospf capability opaque-lsa
```

enable ospf export

```
enable ospf export [bgp | direct | e-bgp | i-bgp | rip | static] [cost <cost> type [ase-type-1 | ase-type-2] {tag <number>} | <policy-map>]
```

Description

Enables redistribution of routes to OSPF.

Syntax Description

bgp	Specifies BGP routes.
i-bgp	Specifies I-BGP routes.
direct	Specifies direct routes.
e-bgp	Specifies E-BGP routes.
rip	Specifies RIP routes.
static	Specifies static routes.
cost	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.

number	Specifies a tag value.
policy-map	Specifies a policy.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After OSPF export is enabled, the OSPF router is considered to be an ASBR. Interface routes that correspond to the interface that has OSPF enabled are ignored.

The cost metric is inserted for all BGP, RIP-learned, static, and direct routes injected into OSPF. If the cost metric is set to 0, the cost is inserted from the route. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.

The same cost, type, and tag values can be inserted for all the export routes, or a policy can be used for selective insertion. When a policy is associated with the export command, the policy is applied on every exported route. The exported routes can also be filtered using a policy.

Example

The following command enables OSPF to export BGP-related routes using LSAs to other OSPF routers:

```
enable ospf export bgp cost 1 ase-type-1 tag 0
```

enable ospf originate-default

```
enable ospf originate-default {always} cost <cost> type [ase-type-1 | ase-type-2] {tag <number>}
```

Description

Enables a default external LSA to be generated by OSPF, if no other default route is originated by OSPF by way of RIP and static route re-distribution.

Syntax Description

always	Specifies for OSPF to always advertise the default route.
cost	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.

Default

N/A.

Usage Guidelines

If always is specified, OSPF always advertises the default route. If always is not specified, OSPF adds the default LSA if a reachable default route is in the route table.

Example

The following command generates a default external type-1 LSA:

```
enable ospf originate-default cost 1 ase-type-1 tag 0
```

enable ospf restart-helper-lsa-check

```
enable ospf [vlan [all | <vlan-name>] | area <area-identifier> | virtual-link  
<router-identifier> <area-identifier>] restart-helper-lsa-check
```

Description

Enables the restart helper router to terminate graceful OSPF restart when received LSAs would affect the restarting router.

Syntax Description

vlan-name	Specifies a VLAN name.
all	Specifies all VLANs
area-identifier	Specifies an OSPF area.
router-identifier	Specifies the router ID of the remote router of the virtual link.

Default

The default is enabled.

Usage Guidelines

This command configures the restart helper router to terminate graceful OSPF restart when received LSAs would affect the restarting router. This will occur when the restart-helper receives an LSA that will be flooded to the restarting router or when there is a changed LSA on the restarting router's retransmission list when graceful restart is initiated.

Example

The following command configures a router to terminate graceful OSPF restart for all routers in area 10.20.30.40 if it receives an LSA that would affect routing:

```
enable ospf area 10.20.30.40 restart-helper-lsa-check
```

enable ospf use-ip-router-alert

```
enable ospf use-ip-router-alert
```

Description

Enables the generation of the OSPF router alert IP option.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the OSPF router alert IP option:

```
enable ospf use-ip-router-alert
```

show ospf

```
show ospf
```

Description

Displays global OSPF information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays global OSPF information:

```
show ospf
```

show ospf area

```
show ospf area {<area-identifier>}
```

Description

Displays information about OSPF areas.

Syntax Description

area-identifier	Specifies an OSPF area.
-----------------	-------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about OSPF area 1.2.3.4:

```
show ospf area 1.2.3.4
```

show ospf area detail

```
show ospf area detail
```

Description

Displays information about all OSPF areas.

Syntax Description

detail	Specifies to display the information in detailed format.
--------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about all OSPF areas:

```
show ospf area detail
```

show ospf ase-summary

```
show ospf ase-summary
```

Description

Displays the OSPF external route aggregation configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the OSPF external route aggregation configuration:

```
show ospf ase-summary
```

show ospf interfaces

```
show ospf interfaces {vlan <vlan-name> | area <area-identifier> | enabled}
```

Description

Displays information about one or all OSPF interfaces.

Syntax Description

vlan-name	Specifies a VLAN name.
area-identifier	Specifies an OSPF area.
enabled	Displays only OSPF enabled interfaces.

Default

If no argument is specified, all OSPF interfaces are displayed.

Usage Guidelines

None.

Example

The following command displays information about one or all OSPF interfaces on the VLAN *accounting*:

```
show ospf interfaces vlan accounting
```

show ospf interfaces detail

```
show ospf interfaces detail
```

Description

Displays detailed information about all OSPF interfaces.

Syntax Description

detail	Specifies to display the information in detailed format.
--------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about all OSPF interfaces:

```
show ospf interfaces detail
```

show ospf lsdb

```
show ospf lsdb {detail | stats} {area [<area-identifier> | all]}
{{lstype} [<lstype> | all]} {lsid <lsid-address>{<lsid-mask>}}
{routerid <routerid-address> {<routerid-mask>}} {interface[ [<ip-address>{<ip-mask>} |
<ipNetmask>] | vlan <vlan-name>]}
```

Description

Displays a table of the current Link-State Database (LSDB).

Syntax Description

detail	Specifies to display all fields of matching LSAs in a multi-line format.
stats	Specifies to display the number of matching LSAs, but not any of their contents.
area-identifier	Specifies an OSPF area.

all	Specifies all OSPF areas.
lstype	Specifies an LS type
lsid	Specifies an LS ID.
lsid-mask	Specifies an LS ID mask
interface	Specifies to display interface types.
routerid-address	Specifies a LSA router ID address.
vlan-name	Specifies a VLAN name.

Default

Display in summary format.

Usage Guidelines

The NETGEAR 8800 provides several filtering criteria for the `show ospf lsdb` command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospf lsdb
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is `all` with no detail. If detail is specified, each entry includes complete LSA information.

Example

The following command displays all areas and all types in a summary format:

```
show ospf lsdb
```

show ospf memory

```
show ospf memory {detail | <memoryType>}
```

Description

Displays OSPF specific memory usage.

Syntax Description

detail	Displays detail information.
memoryType	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays OSPF specific memory for all types:

```
show ospf memory detail
```

show ospf neighbor

```
show ospf neighbor {routerid [<ip-address> {<ip-mask>} | <ipNetmask>]} {vlan <vlan-name>}
{detail}
```

Description

Displays information about an OSPF neighbor.

Syntax Description

ip-address	Specifies an IP address
ip-mask	Specifies a subnet mask.
ipNetmask	Specifies IP address / Netmask
vlan-name	Specifies a VLAN name.
detail	Specifies detail information.

Default

If no argument is specified, all OSPF neighbors are displayed.

Usage Guidelines

None.

Example

The following command displays information about the OSPF neighbors on the VLAN *accounting*:

```
show ospf neighbor vlan accounting
```

show ospf virtual-link

```
show ospf virtual-link {<router-identifier> <area-identifier>}
```

Description

Displays virtual link information about a particular router or all routers.

Syntax Description

router-identifier	Specifies a router interface number.
area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

The area-identifier refer to the transit area used for connecting the two end-points. The transit area cannot have an area identifier of 0.0.0.0 and cannot be a stub or NSSA area.

Example

The following command displays virtual link information about a particular router:

```
show ospf virtual-link 1.2.3.4 10.1.6.1
```

unconfigure ospf

```
unconfigure ospf {vlan <vlan-name> | area <area-identifier>}
```

Description

Resets one or all OSPF interfaces to the default settings.

Syntax Description

vlan-name	Specifies a VLAN name.
area-identifier	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

The NETGEAR 8800 OSPF allows you to change certain configurable OSPF parameters on the fly. This command selectively resets the configurable parameters to their default values. Following is the list of parameters whose values will be reset to their defaults:

Interface

- Hello interval
- Dead interval
- Transmit delay
- Retransmit interval
- Priority
- Cost
- OSPF graceful restart helper mode

Area

- All the parameters of interfaces associated with this area
- Inter-Area-Prefix_LSA Filter
- AS-External-LSA Filter

OSPF Global

- All parameters of all areas in this OSPF domain
- SPF delay interval
- Interface cost metric table
- Route redistribution
- OSPF graceful restart

Example

The following command resets the OSPF interface to the default settings on the VLAN *accounting*:

```
unconfigure ospf accounting
```

This chapter describes commands used for the IPv6 interior gateway protocol OSPFv3.

Open Shortest Path First (OSPFv3) is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an *autonomous system* (AS). In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router in an area has an identical database maintained from the perspective of that router.

OSPFv3 supports IPv6, and uses commands only slightly modified from that used to support IPv4. OSPFv3 has retained the use of the four-byte, dotted decimal numbers for router IDs, LSA IDs, and area IDs.

From the link-state database (LSDB), each router constructs a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can be distributed among them. The cost of a route is described by a single metric.

OSPFv3 allows parts of a network to be grouped together into *areas*. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in link-state advertisement (LSA) traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPFv3 are as follows:

- **Internal Router (IR)**—An internal router has all of its interfaces within the same area.
- **Area Border Router (ABR)**—An ABR has interfaces belonging to two or more areas. It is responsible for exchanging summary advertisements with other OSPFv3 routers.
- **Autonomous System Border Router (ASBR)**—An ASBR acts as a gateway between OSPFv3 and other routing protocols, or other autonomous systems.

Each switch that is configured to run OSPFv3 must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPFv3, instead of having the switch automatically choose its router ID based on the highest interface IPv4 address, since your router may not have an IPv4 address. Not performing this configuration in larger, dynamic environments could result in an older LSDB remaining in use.

Note: Do not set the router ID to 0.0.0.0.

For information about licensing requirements, see Appendix A in the *NETGEAR 8800 User Manual*.

OSPF Edge Mode

OSPF Edge Mode is a subset of OSPF available on platforms with an Advanced Edge license. There are two restrictions on OSPF Edge Mode:

- At most, four Active OSPF VLAN interfaces are permitted. There is no restriction on the number of Passive interfaces.
- The OSPF Priority on VLANs is zero, and is not configurable. This prevents the system from acting as a DR or BDR.

clear ospfv3 counters

```
clear ospfv3 {domain <domainName>} counters
{ interfaces [[vlan | tunnel] all | vlan <vlan-name> |
tunnel <tunnel-name> | area <area-identifier>]
| area [all | <area-identifier>]
| virtual-link [all | {routerid} <router-identifier>]
{area} <area-identifier>]
| neighbor [all | routerid <router-identifier> | vlan <vlan-name> |
tunnel <tunnel-name>]
| system
}
```

Description

Clears the OSPFv3 counters (statistics).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
all	Specifies all VLANs, tunnels, areas, neighbors, or virtual-links.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
router-identifier	Specifies a router identifier, a four-byte, dotted decimal number.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
system	Specifies the OSPFv3 system/global counters.

Default

N/A.

Usage Guidelines

The global command `clear counters` also clears all OSPFv3 counters. This global command is the equivalent of `clear ospfv3 counters` for OSPFv3.

This command can be used to clear various OSPFv3 counters (Interface, Area, Virtual-Link, System etc.). The following is the list of various counters that would be reset to zero by this command:

- Neighbor specific counters
 - Number of state changes
 - Number of events
- Interface/VLAN/Virtual-link/Tunnel specific counters
 - Number of Hellos rxed
 - Number of Hellos txed
 - Number of DB Description rxed
 - Number of DB description txed
 - Number of LS request rxed
 - Number of LS request txed
 - Number of LS update rxed
 - Number of LS update txed
 - Number of LS ack rxed
 - Number of LS ack txed
 - Number of rxed OSPFv3 packet discarded
 - Number of state changes
 - Number of events
- Area Specific counters
 - All counters of interfaces associated with an area
 - Number of SPF runs
- Domain (global)/system specific counters
 - Number of self originated LSAs
 - Number of received LSAs

Example

The following command clears the OSPFv3 counters for area 1.1.1.1:

```
clear ospfv3 counters area 1.1.1.1
```

The following command clears all the OSPFv3 counters for the neighbor 192.168.0.1 in the domain *ospf-core*:

```
clear ospfv3 domain ospf-core counters neighbor routerid 192.168.0.1
```

configure ospfv3 add interface

```
configure ospfv3 {domain <domainName>} add [vlan <vlan-name> | tunnel <tunnel-name>]  
{instance-id <instanceId>} area <area-identifier> link-type [auto | broadcast |  
point-to-point] {passive}
```

Description

Enables OSPFv3 on an interface.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
instanceId	Specifies the instance ID for this interfaces. Range is 0 to 255.
area-identifier	Specifies the area to which the VLAN is assigned.
auto	Specifies to automatically determine the OSPFv3 link type based on the interface type.
broadcast	Specifies a broadcast link, such as Ethernet. Routers must elect a DR and a BDR during synchronization.
point-to-point	Specifies a point-to-point link type, such as PPP.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

The default link-type is Auto.

The default instance ID is 0.

Usage Guidelines

This command is used to enable the OSPFv3 protocol on an IPv6 configured VLAN or an IPv6 tunnel. The instance ID is used to control the selection of other routers as neighbors. The router will become a neighbor only with routers that have the same instance ID.

An interface can have only one instance ID associated with it in one OSPFv3 domain. However, the same interface can be associated with another OSPFv3 domain with a different instance ID. An interface associated with two OSPFv3 domains cannot have same instance ID.

To change the instance ID associated with an interface, you must first remove the interface from the OSPFv3 area and then add it back with a different instance ID.

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Enable IPv6 forwarding before enabling OSPFv3, otherwise, you will receive a warning message.

Note: Configuration of the link-type parameter is not supported. OSPFv3 will always consider the link-type to be broadcast.

Example

The following command adds the VLAN *accounting* (enabling OSPFv3 on the interface), to the area 0.0.0.1 with an instance ID of 2:

```
configure ospfv3 add vlan accounting instance-id 2 area 0.0.0.1 link-type auto
```

configure ospfv3 add interface all

```
configure ospfv3 {domain <domainName>} add [vlan | tunnel] all {instance-id <instanceId>} area <area-identifier> {passive}
```

Description

Enables OSPFv3 on all VLANs or all tunnels (router interfaces).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
instanceId	Specifies the instance ID for these interfaces. Range is 0 to 255.
area-identifier	Specifies the area to which the interfaces are assigned.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

OSPFv3 is disabled on the interfaces.

The default instance ID is 0.

Usage Guidelines

This command is used to enable the OSPFv3 protocol on all IPv6 configured VLANs or all IPv6 tunnels. The instance ID is used to control the selection of other routers as neighbors. The router will become a neighbor only with routers that have the same instance ID.

An interface can have only one instance ID associated with it in one OSPFv3 domain. However, the same interface can be associated with another OSPFv3 domain with a different instance ID. An interface associated with two OSPFv3 domains cannot have same instance ID.

To change the instance ID associated with an interface, you must first remove the interface from the OSPFv3 area and then add it back with a different instance ID.

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Example

The following command enables OSPFv3 on all IPv6 tunnels:

```
configure ospfv3 add tunnel all area 0.0.0.1
```

configure ospfv3 add virtual-link

```
configure ospfv3 {domain <domainName>} add virtual-link {routerid} <router-identifier> {area}
<area-identifier>
```

Description

Adds a virtual link connected to another ABR.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies the transit area identifier, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- router-identifier—Far-end router identifier, a four-byte, dotted decimal number.

- **area-identifier**—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command configures a virtual link with router ID 10.1.2.1 through the transit area 10.1.0.0:

```
configure ospfv3 add virtual-link 10.1.2.1 10.1.0.0
```

configure ospfv3 area add range

```
configure ospfv3 {domain <domainName>} area <area-identifier> add range <ipv6netmask>
[advertise | noadvert] inter-prefix
```

Description

Configures a range of IP addresses in an OSPFv3 area to be aggregated.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
ipv6netmask	Specifies an IPv6 address / prefix length.
advertise	Specifies to advertise the aggregated range of IP addresses.
noadvert	Specifies not to advertise the aggregated range of IP addresses.
inter-prefix	Specifies aggregate, inter-area-prefix LSAs.

Default

No OSPFv3 inter-area-prefix LSAs are configured.

Usage Guidelines

If advertised, the aggregated IP range is exported as a single LSA by the ABR.

Example

The following command is used to summarize a certain range of IP addresses within an area and export them out as a single address to area 0.0.0.1:

```
configure ospfv3 area 0.0.0.1 add range 2aaa:456:3ffe::/64 advertise inter-prefix
```

configure ospfv3 area cost

```
configure ospfv3 {domain <domainName>} area <area-identifier> cost [automatic | <cost>]
```

Description

Configures the cost of sending a packet to all interfaces belonging to an area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
automatic	Determine the advertised cost from the OSPFv3 metric table.
cost	Specifies the cost metric. Range is 1 to 65535.

Default

The default cost is automatic. The default domain is OSPF-Default.

Usage Guidelines

Use this command to set the cost of the links belonging to area manually, if the default cost needs to be overwritten. The interface cost is advertised as the link cost in router-LSA.

Example

The following command configures the cost of area 0.0.0.1 to 10. All the links of this area will inherit the area's cost value of 10.

```
configure ospfv3 domain ospf-enterprise area 0.0.0.1 cost 10
```

configure ospfv3 area delete range

```
configure ospfv3 {domain <domainName>} area <area-identifier> delete range <ipv6netmask>
```

Description

Removes a range of IP addresses in an OSPFv3 area to be aggregated.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
ipv6netmask	Specifies an IPv6 address / prefix length.

Default

No OSPFv3 inter-area-prefix LSAs are configured.

Usage Guidelines

If you attempt to delete a range that was not configured, you will receive an error message.

Example

The following command is used to delete a summary network from area 0.0.0.1:

```
configure ospfv3 area 0.0.0.1 delete range 2aaa:456:3ffe::/64
```

configure ospfv3 area external-filter

```
configure ospfv3 {domain <domainName>} area <area-identifier> external-filter [<policy-map>
|none]
```

Description

Configures an external filter policy.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies the OSPFv3 target area.
policy-map	Specifies a policy.
none	Specifies not to apply an external filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPFv3 areas (an ABR function), a policy can be applied to an OSPFv3 area that filters a set of OSPFv3 external routes from being advertised into that area, in other words, filtering some of the inbound AS-external-LSAs.

OSPFv3 routers that do not have enough memory to hold the entire AS-external-LSAa should configure an external area filter to drop part of the external-LSAs. Configuring this policy will enable routers with limited resources to be put into an OSPFv3 network.

Using the none mode specifies that no external filter is applied.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - nlri <IPv6-address>/<mask-len>
- Action (set) attributes
 - permit

- deny

Any other policy attribute will not be recognized and will be ignored.

The following is an example of an external area policy file:

```
entry one {
  if match any{
    nlri 2001:db8:3e5c::/48;
    nlri 2001:db8:2146:2341::/64;
  } then {
    deny;
  }
}
```

Example

The following command configures an external filter policy, *nosales* for area 1.2.3.4:

```
configure ospfv3 area 1.2.3.4 external-filter nosales
```

configure ospfv3 area interarea-filter

```
configure ospfv3 {domain <domainName>} area <area-identifier> interarea-filter [<policy-map>
|none]
```

Description

Configures an inter-area filter policy.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies the OSPFv3 target area.
policy-map	Specifies a policy.
none	Specifies not to apply an inter-area filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

The NETGEAR 8800 OSPFv3 can apply an inter-area policy to filter some inter-area-prefix-LSAs and inter-area-router-LSAs from other areas. This can reduce the size of link state database of routers belonging to the area.

Using the none mode specifies that no external filter is applied.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - nlri <IPv6-address>/<mask-len>
- Action (set) attributes
 - permit
 - deny

Any other policy attribute will not be recognized and will be ignored.

The following is an example of an external area policy file:

```
entry one {
    if match any{
        nlri 2001:db8:3e5c::/48;
        nlri 2001:db8:2146:2341::/64;
    } then {
        deny;
    }
}
entry two {
    if match any{
        nlri 2001:db8:444::/48;
        nlri 2001:db8:541f:65bd::/64;
    } then {
        permit;
    }
}
```

Example

The following command configures an inter-area filter policy, *nosales* for area 1.2.3.4:

```
configure ospfv3 area 1.2.3.4 interarea-filter nosales
```

configure ospfv3 area normal

```
configure ospfv3 {domain <domainName>} area <area-identifier> normal
```

Description

Configures an OSPFv3 area as a normal area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

Normal.

Usage Guidelines

A normal area is an area that is not any of the following:

- Stub area
- NSSA

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Example

The following command configures an OSPFv3 area as a normal area:

```
configure ospfv3 area 10.1.0.0 normal
```

configure ospfv3 area priority

```
configure ospfv3 {domain <domainName>} area <area-identifier> priority <priority>
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for all the interfaces within the area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
priority	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

When two routers are attached to a network, both attempt to become the designated router. The one with the higher priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets all the interfaces in area 1.2.3.4 to not be selected as the designated router:

```
configure ospfv3 area 1.2.3.4 priority 0
```

configure ospfv3 area stub

```
configure ospfv3 {domain <domainName>} area <area-identifier> stub [summary | nosummary]
stub-default-cost <cost>
```

Description

Configures an OSPFv3 area as a stub area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
summary	Specifies that inter-area LSAs can be propagated into the area.
nosummary	Specifies that inter-area LSAs cannot be propagated into the area.
cost	Specifies a cost metric.

Default

N/A.

Usage Guidelines

A stub area is connected to only one other area. The area that connects to a stub area can be the backbone area. External route information is not distributed into stub areas. Stub areas are used to reduce memory consumption requirements on OSPFv3 routers.

Example

The following command configures an OSPFv3 area as a stub area:

```
configure ospfv3 area 0.0.0.6 stub nosummary stub-default-cost 10
```

configure ospfv3 area timer

```
configure ospfv3 {domain <domainName>} area <area-identifier> timer {retransmit-interval}
<retransmit-interval> {transit-delay} <transit-delay> {hello-interval} <hello-interval>
{dead-interval} <dead-interval>
```


Description

Configures the timers for all interfaces in the same OSPFv3 area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1- 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.

Default

- retransmit interval—Default: 5 seconds
- transit delay—Default: 1 second
- hello interval—Default: 10 seconds
- dead interval—Default: 40 seconds

Usage Guidelines

Configuring OSPFv3 timers on a per-area basis is a shorthand for applying the timers to each VLAN and tunnel in the area at the time of configuration. If you add more VLANs or tunnels to the area, you must configure the timers for them explicitly.

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.

The value of the dead interval and the hello interval must be same for all OSPFv3 routers connected to a common link. The value of the dead interval and the hello interval are advertised by OSPFv3 in Hello packets. The shorter the hello interval, the earlier topological changes will be detected, but more routing traffic will ensue.

The retransmit interval must be greater than the expected round trip delay between any two routers on the attached network. The setting of this parameter must be conservative, or needless retransmission will result.

Note: The wait interval for the interface is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers in area 0.0.0.2:

```
configure ospfv3 area 0.0.0.2 timer 10 1 20 200
```

configure ospfv3 delete interface

```
configure ospfv3 {domain <domainName>} delete [vlan <vlan-name> | tunnel <tunnel-name> | [vlan  
| tunnel] all]
```

Description

Disables OSPFv3 on one or all VLANs or tunnels (router interfaces).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all VLANs, or tunnels.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPFv3 on VLAN *accounting*:

```
configure ospfv3 delete vlan accounting
```

configure ospfv3 delete virtual-link

```
configure ospfv3 {domain <domainName>} delete virtual-link {routerid} <router-identifier>
{area} <area-identifier>
```

Description

Deletes a virtual link connected to another ABR.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies the transit area identifier, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- router-identifier—Far-end router identifier, a four-byte, dotted decimal number.
- area-identifier—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command deletes a virtual link with router ID 10.1.2.1 through the transit area 10.1.0.0:

```
configure ospfv3 delete virtual-link 10.1.2.1 10.1.0.0
```

configure ospfv3 import-policy

```
configure ospfv3 {domain <domainName>} import-policy [<policy-map> | none]
```

Description

Configures the import policy for OSPFv3.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
------------	---

policy-map	Specifies the policy.
------------	-----------------------

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding OSPFv3 routes to the IP route table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the `none` option to remove the policy association.

Policy files for this command will recognize only the following policy attributes:

- Match attributes
 - `nlri <IPv6-address>/<mask-len>`
 - `route-origin [ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter | ospfv3-intra]`
- Action (set) attributes
 - `cost <cost>`
 - `tag <number>`

Any other policy attribute will not be recognized and will be ignored.

Example

The following example applies the policy *campuseast* to OSPFv3 routes:

```
configure ospfv3 import-policy campuseast
```

configure ospfv3 interface area

```
configure ospfv3 {domain <domainName>} [vlan <vlan-name> | tunnel <tunnel-name>] area  
<area-identifier>
```

Description

Moves an interface from one OSPFv3 area to another.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

Area 0.0.0.0

Usage Guidelines

Use this command to move an already configured interface from one area to another. The instance ID associated with the interface will be unchanged.

Example

The following command moves the VLAN *accounting* to the OSPFv3 area 0.0.0.6:

```
configure ospfv3 vlan accounting area 0.0.0.6
```

configure ospfv3 interface cost

```
configure ospfv3 {domain <domainName>} [vlan <vlan-name> | tunnel <tunnel-name> | [vlan | tunnel] all]] cost [automatic | <cost>]
```

Description

Configures the cost of one or all interface(s).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
automatic	Determine the advertised cost from the OSPFv3 metric table.
cost	Specifies the cost metric. Range is 1 to 65535.

Default

The default cost is automatic.

Usage Guidelines

Use this command to set the cost of an interface (a VLAN or tunnel) manually, if the default cost needs to be overwritten. The interface cost is advertised as the link cost in router-LSA.

Example

The following command configures the cost metric of the VLAN *accounting*:

```
configure ospfv3 vlan accounting cost 10
```

configure ospfv3 interface priority

```
configure ospfv3 {domain <domainName>} [vlan <vlan-name> | tunnel <tunnel-name> | [vlan |
tunnel] all] priority <priority>
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for one or all OSPFv3 interface(s).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
priority	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

When two routers are attached to a network, both attempt to become the designated router. The one with the higher priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets the priority of the interface VLAN corporate to 10:

```
configure ospfv3 domain ospf-internal vlan corporate priority 10
```

configure ospfv3 interface timer

```
configure ospfv3 {domain <domainName>} [vlan <vlan-name> | tunnel <tunnel-name> | [vlan |
tunnel] all] timer {retransmit-interval} <retransmit-interval> {transit-delay}
<transit-delay> {hello-interval} <hello-interval> {dead-interval} <dead-interval>
```

Description

Configures the timers for all interfaces in the same OSPFv3 area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.

Default

- retransmit interval—Default: 5 seconds
- transit delay—Default: 1 second
- hello interval—Default: 10 seconds
- dead interval—Default: 40 seconds

Usage Guidelines

Use this command to configure the OSPFv3 timers on a per-interface basis.

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.

The value of the dead interval and the hello interval must be same for all OSPFv3 routers connected to a common link. The value of the dead interval and the hello interval are advertised by OSPFv3 in Hello packets. The shorter the hello interval, the earlier topological changes will be detected, but more routing traffic will ensue.

The retransmit interval must be greater than the expected round trip delay between any two routers on the attached network. The setting of this parameter must be conservative, or needless retransmission will result.

Note: The wait interval for the interface is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers for the VLAN *corporate*:

```
configure ospfv3 domain ospf-default vlan corporate timer retransmit-interval 10
transit-delay 2 hello-interval 20 dead-interval 80
```

configure ospfv3 metric-table

```
configure ospfv3 {domain <domainName>} metric-table 10M <cost_10m> 100M <cost_100m> 1G
<cost_1g> {10G <cost_10g>}
```

Description

Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces, and optionally, the 10 Gbps interface.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
cost_x	Specifies the interface cost for the indicated interfaces. Range is 1 to 65535.

Default

- 10 Mbps—The default cost is 100.
- 100 Mbps—The default cost is 50.
- 1 Gbps—The default cost is 40.
- 10 Gbps—The default cost is 20.

Usage Guidelines

The value of the costs cannot be greater for higher speed interfaces. In other words, the following condition must be true:

```
cost_10m >= cost_100m >= cost_1g >= cost_10g
```

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces:

```
configure ospfv3 metric-table 10m 200 100m 100 1g 20
```


configure ospfv3 routerid

```
configure ospfv3 {domain <domainName>} routerid [automatic | <router-identifier>]
```

Description

Configures the OSPFv3 router ID. If automatic is specified, the switch uses the highest IPv4 interface address as the OSPFv3 router ID.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
automatic	Specifies to use automatic addressing.
router-identifier	Specifies a router identifier, a four-byte, dotted decimal number.

Default

Automatic.

Usage Guidelines

Each switch that is configured to run OSPFv3 must have a unique router ID. The router ID is a four-byte, dotted decimal number, like an IPv4 address. Even though the IP address format has changed from IPv4 to IPv6, the router ID format has not. It is recommended that you manually set the router ID of the switches participating in OSPFv3, instead of having the switch automatically choose its router ID based on the highest interface IPv4 address (if it exists). Not performing this configuration in larger, dynamic environments could result in an older link-state database remaining in use.

This command is accepted only when OSPFv3 is globally disabled.

Note: Do not set the router ID to 0.0.0.0.

Example

The following command sets the router ID to 10.1.6.1:

```
configure ospfv3 routerid 10.1.6.1
```

configure ospfv3 spf-hold-time

```
configure ospfv3 {domain <domainName>} spf-hold-time <seconds>
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
seconds	Specifies a time in seconds. The range is 0 to 300 seconds.

Default

3 seconds.

Usage Guidelines

Setting the interval too high will force OSPFv3 to run SPF calculations less frequently. This will reduce the CPU load, but will cause delay in routes getting updated in the IP routing table. Setting the interval too low will decrease the interval between SPF calculations, but will increase the processing load on CPU.

Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
configure ospfv3 spf-hold-time 6
```

configure ospfv3 virtual-link timer

```
configure ospfv3 {domain <domainName>} virtual-link {routerid} <router-identifier> {area}
<area-identifier> timer {retransmit-interval} <retransmit-interval> {transit-delay}
<transit-delay> {hello-interval} <hello-interval> {dead-interval} <dead-interval>
```

Description

Configures the timers for a virtual link.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
router-identifier	Specifies the router ID of the other end of the link.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
retransmit-interval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600 seconds.
transit-delay	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
hello-interval	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.

dead-interval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.
---------------	---

Default

- retransmit interval—Default: 5 seconds
- transit delay—Default: 1 second
- hello interval—Default: 10 seconds
- dead interval—Default: 40 seconds

Usage Guidelines

In OSPFv3, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link.

The smaller the hello interval, the faster topological changes will be detected, but more routing traffic will ensue.

The setting of the retransmit interval should be conservative, or needless retransmissions will result. The value should be larger for serial lines and virtual links.

The transmit delay value should take into account the transmission and propagation delays for the interface.

Note: The wait interval is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers on the virtual link to router 6.6.6.6 transiting area 0.0.0.2:

```
configure ospfv3 virtual-link 6.6.6.6 area 0.0.0.2 timer 10 transit-delay 1 hello-interval 20
dead-interval 200
```

create ospfv3 area

```
create ospfv3 {domain <domainName>} area <area-identifier>
```

Description

Creates an OSPFv3 area.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

Area 0.0.0.0

Usage Guidelines

Area 0.0.0.0 does not need to be created. It exists by default.

Example

The following command creates a non-backbone OSPFv3 area:

```
create ospfv3 area 1.2.3.4
```

delete ospfv3 area

```
delete ospfv3 {domain <domainName>} area [<area-identifier> | all]
```

Description

Deletes an OSPFv3 area or all OSPFv3 areas.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all areas.

Default

N/A.

Usage Guidelines

An OSPFv3 area cannot be deleted if it has an associated interface. Also, area 0.0.0.0 cannot be deleted.

Example

The following command deletes an OSPFv3 area:

```
delete ospfv3 area 1.2.3.4
```

disable ospfv3

```
disable ospfv3 {domain <domainName>}
```

Description

Disables OSPFv3 for the router.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPFv3 for the router:

```
disable ospfv3
```

disable ospfv3 export

```
disable ospfv3 {domain <domainName>} export [direct | ripng | static]
```

Description

Disables redistribution of routes to OSPFv3.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
direct	Specifies direct routes.
ripng	Specifies RIP routes.
static	Specifies static routes.

Default

The default setting is disabled.

Usage Guidelines

Use this command to stop OSPFv3 from exporting routes derived from other protocols.

Example

The following command disables OSPFv3 to export RIPng routes to other OSPFv3 routers:

```
disable ospfv3 export ripng
```

enable ospfv3

```
enable ospfv3 {domain <domainName>}
```

Description

Enables OSPFv3 for the router.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
------------	---

Default

N/A.

Usage Guidelines

When OSPFv3 is enabled, it will start exchanging Hellos on all of its active interfaces. It will also start exporting routes into OSPFv3 routing domain from other protocols, if enabled.

When OSPFv3 is disabled, it will release all the run-time allocated resources like adjacencies, link state advertisements, run-time memory, etc.

OSPFv3 can be enabled successfully if and only if:

- At least one of the VLANs in the current virtual router has one IPv4 address configured
- OR—
- You explicitly configure the OSPFv3 router ID, a four-byte, dotted decimal number

Example

The following command enables OSPFv3 for the router:

```
enable ospfv3
```

enable ospfv3 export

```
enable ospfv3 {domain <domainName>} export [direct | ripng | static] [cost <cost> type
[ase-type-1 | ase-type-2] | <policy-map>]
```

Description

Enables redistribution of routes to OSPFv3.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
direct	Specifies direct routes.
ripng	Specifies RIPng routes.
static	Specifies static routes.
cost	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
policy-map	Specifies a policy.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into OSPFv3. If the cost metric is set to 0, the cost is inserted from the route. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.

Note: Setting the tag value is not supported in this release.

The same cost, type, and tag values can be inserted for all the export routes, or a policy can be used for selective insertion. When a policy is associated with the export command, the policy is applied on every exported route. The exported routes can also be filtered using a policy.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - nlri <IPv6-address>/<mask-len>
- Action (set) attributes
 - cost <cost>
 - tag <number>

- cost-type [ase-type-1 | ase-type-2]
- permit
- deny

Any other policy attribute will not be recognized and will be ignored.

The following is an example OSPFv3 export policy file:

```
entry first {
    if match any{
        nlri 2001:db8:200:300::/64;
        nlri 2001:db8:2146:23d1::/64;
        nlri 2001:db8:af31:3d0::/64;
        nlri 2001:db8:f6:2341::/64;
    } then {
        deny;
    }
}
entry second {
    if match any{
        nlri 2001:db8:304::/48;
        nlri 2001:db8:ca11::/48;
        nlri 2001:db8:da36::/48;
        nlri 2001:db8:f6a6::/48;
    } then {
        cost 220;
        cost-type ase-type-2;
        permit;
    }
}
```

Example

The following command enables OSPFv3 to export RIPv3-related routes and associates a policy *redist*:

```
enable ospfv3 export ripng redist
```

show ospfv3

```
show ospfv3 {domain <domainName>}
```

Description

Displays global OSPFv3 information.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command displays global OSPFv3 information:

```
show ospfv3
```

The following is sample output:

```

OSPF Domain Name      : OSPF-Default
OSPF                  : Enabled           RouterId           : 20.0.0.1
RouterId Selection    : Automatic         ASBR                 : Yes
ABR                   : Yes                ExtLSAs              : 3
ExtLSAChecksum        : 0x19420           OriginateNewLSAs     : 37
ReceivedNewLSAs       : 12                SpfHoldTime          : 10s
Num of Areas          : 5                  10M Cost              : 100
100M Cost              : 50                1000M Cost (1G)      : 40
10000M Cost (10G)     : 20                Router Alert         : Disabled
ASExternal LSALimit   : Disabled           Timeout (Count)      : Disabled (0)
Originate Default     : Disabled
Import Policy File    : none

Redistribute:
  Protocol  Status   Cost   Type  Tag   Policy
  direct    Disabled  20     2     0     none
  ripng     Disabled  20     2     0     none
  static    Enabled   20     2     0     IPv6

```

show ospfv3 area

```
show ospfv3 {domain <domainName>} area {<area-identifier> | detail}
```

Description

Displays information about OSPFv3 areas.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays summary information about the OSPFv3 areas:

```
show ospfv3 area
```

The following is sample output:

AREA ID	Type	Summ	Def	Num	Num	Num	SPF	Num	LSA
			Metric	ABR	ASBR	Intf	Runs	LSAs	Checksum
0.0.0.0	NORM	----	-----	0	0	1	7	7	0x3155b
1.0.0.0	NORM	----	-----	1	1	1	6	9	0x4793d
2.0.0.0	NORM	----	-----	0	0	1	5	10	0x47174
3.0.0.0	NORM	----	-----	1	0	1	3	12	0x420cf
5.0.0.0	NORM	----	-----	1	0	1	4	10	0x3b5b1

The following command displays information about OSPFv3 area 1.0.0.0:

```
show ospfv3 area 1.0.0.0
```

The following is sample output:

Area Identifier	: 1.0.0.0	Type	: NORM
Router ID	: 20.0.0.1	Num of Interfaces	: 1
Spf Runs	: 6	Num ABRs	: 1
Num ASBRs	: 1	Num DC-Bit LSAs	: 1
Num Indication LSAs	: 1	Num of DoNotAge LSAs	: 1
Num LSAs	: 9	LSA Chksum	: 0x4793d
Num of Nbrs	: 1	Num of Virtual Nbrs	: 0

Interfaces:

Interface Name	Ospf	State	DR ID	BDR ID
to65	E	BDR	0.0.0.65	20.0.0.1
accounts	E	DR	80.0.0.5	0.0.0.0
finance	E	BDR	90.0.0.7	66.0.0.4

```

engineering          E    ODR    192.168.0.1    165.0.0.3
Corporate            E    ODR    201.0.16.6    204.0.0.1
Inter-Area route Filter: ospfSummPolicy
External route Filter: ospfExtPolicy
Configured Address Ranges:
  Addr: fffe:408:1449::/48 Type: 3 Advt: Yes
  Addr: ffe0:930:2781::/40 Type: 7 Advt: No

```

show ospfv3 interfaces

```
show ospfv3 {domain <domainName>} interfaces {vlan <vlan-name> | tunnel <tunnel-name> | area
<area-identifier> | detail}
```

Description

Displays information about one or all OSPFv3 interfaces.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
detail	Specifies to display the information in detailed format.

Default

If no argument is specified, all OSPFv3 interfaces are displayed.

Usage Guidelines

None.

Example

The following command shows a summary of the OSPFv3 interfaces:

```
show ospfv3 interfaces
```

The following is sample output from the command:

VLAN/Tunnel	IPv6 Address	AREA ID	Flags	Cost	State	Neighbors
ixia	4:5:6:7::1/64	2.0.0.0	-rif	5/A	DR	1
to-jmpr	111:222:333::7/48	5.0.0.0	-rif	5/A	BDR	1
to-Zebra	3ffe:506::4/48	0.0.0.0	-rif	5/A	BDR	1
to5	234:567::7/48	3.0.0.0	-rif	5/A	BDR	1
to65	10:203:134:7::7/48	1.0.0.0	-rif	5/A	BDR	1

Flags : (f) Interface Forwarding Enabled, (i) Interface OSPF Enabled,
(p) Passive Interface, (r) Router OSPF Enable.

The following command displays information about the OSPFv3 interfaces on the VLAN *to5*:

```
show ospfv3 interfaces vlan to5
```

The following is sample output:

```
Interface           : to5                Enabled           : ENABLED
Router              : ENABLED             AreaID            : 3.0.0.0
RouterID            : 20.0.0.1           Link Type         : broadcast
Passive             : No                  Cost              : 40A
Priority             : 1                   Transit Delay     : 1s
Hello Interval      : 10s                 Rtr Dead Time    : 40s
Retransmit Interval : 5s                  Wait Timer       : 40s
Interface ID        : 63                  Instance ID      : 0
State               : BDR                  Number of state chg : 2
Hello due in        : 3s                  Number of events  : 3
Total Num of Nbrs   : 1                   Nbrs in FULL State : 1
Hellos Rxed         : 94                  Hellos Txed      : 94
DB Description Rxed : 4                   DB Description Txed : 3
LSA Request Rxed    : 1                   LSA Request Txed  : 1
LSA Update Rxed     : 8                   LSA Update Txed   : 7
LSA Ack Rxed        : 6                   LSA Ack Txed      : 5
In Discards         : 0
DR RtId             : 10.0.0.5            BDR RtId          : 20.0.0.1
DR Interface addr    : fe80::280:c8ff:feb9:1cf1
BDR Interface addr   : fe80::280:c8ff:feb9:2089
```

Neighbors:

```
RtrId: 10.0.0.5  IpAddr: fe80::280:c8ff:feb9:1cf1  Pri: 1  Type: Auto
State: FULL  DR: 10.0.0.5  BDR: 20.0.0.1  Dead Time: 00:00:36
Options: 0x13 (-|R|-|-|E|V6)  Opaque LSA: No
```

show ospfv3 lsdb

```
show ospfv3 {domain <domainName>} lsdb {detail} {area [<area-identifier> | all] {lstype
[router | network | inter-prefix | inter-router | intra-prefix]} | [vlan [<vlan-name> | all]
| tunnel [<tunnel-name> | all]} {lstype link} | lstype [as-external | router | network |
inter-prefix | inter-router | intra-prefix | link]} {lsid <lsid-address>} {adv-router
<router-identifier>}}
```

Description

Displays a table of the current Link-State Database (LSDB).

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
detail	Specifies to display all fields of matching LSAs in a multi-line format.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all OSPFv3 areas, IPv6 configured VLANs, or IPv6 tunnels.
link	Link LSA
router	Router LSA
network	Network LSA
inter-prefix	Inter Area Prefix LSA
inter-router	Inter Area Router LSA
intra-prefix	Intra Area Prefix LSA
as-external	AS External LSA
lsid-address	Specifies the link state ID of the LSA.
routerid-identifier	Specifies the router identifier of the advertising router.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

Display in summary format.

Usage Guidelines

The NETGEAR 8800 provides several filtering criteria for the `show ospfv3 lsd` command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospfv3 lsd
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is `all` with no detail. If detail is specified, each entry includes complete LSA information.

Example

The following command displays all areas and all types in a summary format:

```
show ospfv3 lsdb
```

show ospfv3 lsdb stats

```
show ospfv3 {domain <domainName>} lsdb stats {area [<area-identifier> | all] {lstype [router  
| network | inter-prefix | inter-router | intra-prefix]} | [vlan [<vlan-name> | all] | tunnel  
[<tunnel-name> | all]} {lstype link} | lstype [as-external | router | network | inter-prefix  
| inter-router | intra-prefix | link]} {lsid <lsid-address>} {adv-router <router-identifier>}
```

Description

Displays a table of the current Link-State Database (LSDB) statistics.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all OSPFv3 areas, IPv6 configured VLANs, or IPv6 tunnels.
link	Link LSA
router	Router LSA
network	Network LSA
inter-prefix	Inter Area Prefix LSA
inter-router	Inter Area Router LSA
intra-prefix	Intra Area Prefix LSA
as-external	AS External LSA
lsid-address	Specifies the link state ID of the LSA.
routerid-identifier	Specifies the router identifier of the advertising router.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

Display in summary format.

Usage Guidelines

The NETGEAR 8800 provides several filtering criteria for the `show ospfv3 lsdb stats` command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospfv3 lsdb stats
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is `all`.

Example

The following command displays all areas and all types in a summary format:

```
show ospfv3 lsdb stats
```

show ospfv3 memory

```
show ospfv3 memory {detail | <memoryType>}
```

Description

Displays OSPFv3 specific memory usage.

Syntax Description

<code>detail</code>	Displays detail information.
<code>memoryType</code>	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays OSPFv3 specific memory for all types:

```
show ospfv3 memory detail
```

show ospfv3 neighbor

```
show ospfv3 {domain <domainName>} neighbor {routerid <ip-address>} {vlan <vlan-name> | tunnel <tunnel-name>} {detail}
```

Description

Displays information about an OSPFv3 neighbor.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
ip-address	Specifies a neighbor router ID.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.
detail	Specifies detail information.

Default

If no argument is specified, all OSPFv3 neighbors are displayed.

Usage Guidelines

None.

Example

The following command displays information about the OSPFv3 neighbors on the VLAN *accounting*:

```
show ospfv3 neighbor vlan accounting
```

show ospfv3 virtual-link

```
show ospfv3 {domain <domainName>} virtual-link {{routerid} <router-identifier> {area}
<area-identifier>}
```

Description

Displays virtual link(s) information.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
router-identifier	Specifies a router identifier, a four-byte, dotted decimal number.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

router-identifier—Router ID for the other end of the link.

area-identifier—Transit area used for connecting the two end-points. The transit area cannot have an area identifier of 0.0.0.0 and cannot be a stub or NSSA area.

Example

The following command displays information about the virtual link to a particular router:

```
show ospfv3 virtual-link 1.2.3.4 10.1.6.1
```

unconfigure ospfv3

```
unconfigure ospfv3 {domain <domainName>} {area <area-identifier> | vlan <vlan-name> | tunnel <tunnel-name>}
```

Description

Resets one or all OSPFv3 interfaces to the default settings.

Syntax Description

domainName	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
area-identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
vlan-name	Specifies an IPv6 configured VLAN.
tunnel-name	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

The NETGEAR 8800 OSPFv3 allows you to change certain configurable OSPFv3 parameters on the fly. This command selectively resets the configurable parameters to their default values. The following is the list of parameters whose values will be reset to their defaults:

- Interface
 - Hello Interval
 - Dead Interval
 - Transmit Delay
 - Retransmit Interval
 - Priority
 - Cost

- Area
 - All the parameters of Interfaces associated with this area
 - Inter-Area-Prefix-LSA Filter
 - AS-External-LSA Filter
- OSPF Global
 - All parameters of all areas in this OSPF domain
 - SPF Delay interval
 - Interface Cost metric Table
 - Route Redistribution

Example

The following command resets the OSPFv3 interface to the default settings on the VLAN *accounting*:

```
unconfigure ospfv3 accounting
```

The following command unconfigures the parameters of the area 0.0.0.1 (and all its associated interfaces):

```
unconfigure ospfv3 domain ospf-default area 0.0.0.1
```

This chapter describes commands for doing the following:

- Configuring BGP
- Displaying BGP information
- Managing BGP

For an introduction to the BGP feature, see the *NETGEAR 8800 User Manual*.

clear bgp flap-statistics

```
clear bgp {neighbor} <remoteaddr> {address-family [ipv4-unicast | ipv4-multicast]}
flap-statistics [all | as-path <path expression>
| community [no-advertise | no-export | no-export-subconfed
| number <community_num> | <AS_Num>:<Num>]
| network [any / <netMaskLen> | <networkPrefixFilter>] {exact}]
```

Description

Clears flap statistics for routes to specified neighbors.

Syntax Description

all	Specifies flap statistics for all routes.
remoteaddr	Specifies an IP address that identifies a BGP neighbor.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_num	Specifies a community number.
AS_Num	Specifies an autonomous system ID (0-65535).
Num	Specifies a community number.
any	Specifies all routes with a given or larger mask length.

netMaskLen	Specifies a subnet mask length (number of bits).
networkPrefixFilter	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Use this command to clear flap statistics for a specified BGP neighbor.

The option `network any / <netMaskLen>` clears the statistics for all BGP routes whose mask length is equal to or greater than `<maskLength>`, irrespective of their network address.

The option `network any / <netMaskLen> exact` clears the statistics for all BGP routes whose mask length is exactly equal to `<maskLength>`, irrespective of their network address.

Example

The following command clears the flap statistics for a specified neighbor:

```
clear bgp neighbor 10.10.10.10 flap-statistics all
```

clear bgp neighbor counters

```
clear bgp neighbor [<remoteaddr> | all] counters
```

Description

Resets the BGP counters for one or all BGP neighbor sessions to zero.

Syntax Description

remoteaddr	Specifies the IP address of a specific BGP neighbor.
all	Specifies that counters for all BGP neighbors should be reset.

Default

N/A.

Usage Guidelines

This command resets the following counters:

- In-total-msgs
- Out-total-msgs
- In-updates

- Out-updates
- Last-error
- FsmTransitions

The command `clear counters` also resets all counter for all BGP neighbors. For BGP, the `clear counters` command is equivalent to the following BGP command:

```
clear bgp neighbor all counters
```

Example

The following command resets the counters for the BGP neighbor at 10.20.30.55:

```
clear bgp neighbor 10.20.30.55 counters
```

configure bgp add aggregate-address

```
configure bgp add aggregate-address {address-family [ipv4-unicast | ipv4-multicast]}
<ipaddress> {as-match | as-set} {summary-only} {advertise-policy <policy>}
{attribute-policy <policy>}
```

Description

Configures a BGP aggregate route.

Syntax Description

address-family	The address family. BGP supports two address families: IPv4 unicast and IPv4 multicast.
ipaddress	Specifies an IP network address and mask length.
as-match	Generates autonomous system sequence path information (order of AS numbers in AS_PATH is preserved).
as-set	Generates autonomous system set path information (order of AS numbers in AS_PATH is not preserved).
summary-only	Specifies to send only aggregated routes to the neighbors.
advertise-policy	Specifies the policy used to select routes for this aggregated route.
attribute-policy	Specifies the policy used to set the attributes of the aggregated route.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Before you can create an aggregate route, you must enable BGP aggregation using the following command:

```
enable bgp aggregation
```

BGP supports overlapping routes. For example, you can configure both of the following aggregate addresses:

- 192.0.0.0/8
- 192.168.0.0/16

After you create an aggregate route, the aggregate route remains inactive until BGP receives a route with an IP address and mask that conforms to an aggregate route. When a conforming route is received, the aggregate route becomes active and is advertised to BGP neighbors. If the `summary-only` option is specified, only the aggregate route becomes active and is advertised. If the `summary-only` option is omitted, any conforming aggregate routes and the received route are advertised to BGP neighbors.

Example

The following command configures a BGP aggregate route:

```
configure bgp add aggregate-address 192.1.1.4/30
```

configure bgp add confederation-peer sub-AS-number

```
configure bgp add confederation-peer sub-AS-number <number>
```

Description

Adds a sub-AS to a confederation.

Syntax Description

number	Specifies a sub-AS number of the confederation. The range is 1 to 4294967295.
--------	---

Default

N/A.

Usage Guidelines

Invoke this command multiple times to add multiple sub-ASs.

IBGP requires networks to use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, all BGP speakers in each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in *RFC 5396, Textual Representation of Autonomous System (AS) Numbers*.

Example

The following command adds one sub-AS to a confederation using the ASPLAIN 4-byte AS number format:

```
configure bgp add confederation-peer sub-AS-number 6553700
```

The following command adds one sub-AS to a confederation using the ASDOT 4-byte AS number format:

```
configure bgp add confederation-peer sub-AS-number 100.100
```

configure bgp add network

```
configure bgp add network {address-family [ipv4-unicast | ipv4-multicast]}
<ipaddr>/<mask_len> {network-policy <policy>}
```

Description

Adds a network to be originated from this router.

Syntax Description

address-family	The address family to which the network routes are exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
ipaddr	Specifies an IP network address.
mask_len	Specifies a netmask length.
policy-name	Name of policy to be associated with network export. Policy can filter and/or change the route parameters.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The network must be present in the routing table.

Using the `export` command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command adds a network to be originated from this router:

```
configure bgp add network 192.1.1.16/32
```

configure bgp as-display-format

```
configure bgp as-display-format [asdot | asplain]
```

Description

Configures the AS number format displayed in `show` commands.

Syntax Description

asdot	Specifies the ASDOT format.
asplain	Specifies the ASPLAIN format.

Default

N/A.

Usage Guidelines

The ASPLAIN and ASDOT formats are described in *RFC 5396, Textual Representation of Autonomous System (AS) Numbers*.

Examples

The following command selects the ASDOT 4-byte AS number format:

```
configure bgp as-display-format asdot
```

configure bgp as-number

```
configure bgp AS-number <number>
```

Description

Changes the local AS number used by BGP.

Syntax Description

number	Specifies a local AS number. The range is 1 to 4294967295.
--------	--

Default

N/A.

Usage Guidelines

BGP must be disabled before the AS number can be changed.

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in *RFC 5396, Textual Representation of Autonomous System (AS) Numbers*.

Examples

The following command specifies a local AS number using the ASPLAIN 4-byte AS number format:

```
configure bgp AS-number 100000
```

The following command specifies a local AS number using the ASDOT 4-byte AS number format:

```
configure bgp AS-number 45776.24064
```

configure bgp cluster-id

```
configure bgp cluster-id <cluster-id>
```

Description

Configures the local cluster ID.

Syntax Description

cluster-id	Specifies a 4 byte field used by a route reflector to recognize updates from other route reflectors in the same cluster. The range is 0 - 4294967295.
------------	---

Default

N/A.

Usage Guidelines

Used when multiple route reflectors are used within the same cluster of clients.

NETGEAR recommends disabling BGP before configuring the cluster ID.

Example

The following command appends a BGP route reflector cluster ID to the cluster list of a route:

```
configure bgp cluster-id 40000
```

configure bgp confederation-id

```
configure bgp confederation-id <number>
```

Description

Specifies a BGP routing confederation ID.

Syntax Description

confederation-id	Specifies a routing confederation identifier, which is a 4-byte AS number in the range of 1 to 4294967295.
------------------	--

Default

N/A.

Usage Guidelines

IBGP requires that networks use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

Use a confederation ID of 0 to indicate no confederation.

The confederation ID is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in *RFC 5396, Textual Representation of Autonomous System (AS) Numbers*.

Example

The following command specifies a BGP routing confederation ID using the ASPLAIN 4-byte AS number format:

```
configure bgp confederation-id 1000000
```

The following command specifies a BGP routing confederation ID using the ASDOT 4-byte AS number format:

```
configure bgp confederation-id 15.16960
```

configure bgp delete aggregate-address

```
configure bgp delete aggregate-address {address-family [ipv4-unicast | ipv4-multicast]} [<ip address/masklength> | all]
```

Description

Deletes one or all BGP aggregated route.

Syntax Description

address-family	The address family. BGP supports two address families: IPv4 unicast and IPv4 multicast.
----------------	---

ip address/mask length	Specifies an IP network address and netmask length.
all	Specifies all aggregated routes.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Example

The following command deletes a BGP aggregate route:

```
configure bgp delete aggregate-address 192.1.1.4/30
```

configure bgp delete confederation-peer sub-AS-number

```
configure bgp delete confederation-peer sub-AS-number <number>
```

Description

Specifies a sub-AS that should be deleted from a confederation.

Syntax Description

sub-AS-number	Specifies a sub-AS.
---------------	---------------------

Default

N/A.

Usage Guidelines

BGP requires that networks use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

Example

The following command deletes a sub-AS from a confederation using the ASPLAIN 4-byte AS number format:

```
configure bgp delete confederation-peer sub-AS-number 6553700
```

The following command deletes a sub-AS from a confederation using the ASDOT 4-byte AS number format:

```
configure bgp delete confederation-peer sub-AS-number 100.100
```

configure bgp delete network

```
configure bgp delete network {address-family [ipv4-unicast | ipv4-multicast]} [all |  
<ipaddress/mask length>]
```

Description

Deletes a network to be originated from this router.

Syntax Description

address-family	The address family to which the network routes are exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
all	Specifies all networks.
ipaddress	Specifies an IP network address and a netmask length.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

None.

Example

The following command deletes a network to be originated from this router:

```
configure bgp delete network 192.1.1.12/30
```

configure bgp export shutdown-priority

```
configure bgp export [direct | isis | ospf | ospf-extern1 | ospf-extern2 | ospf-inter |  
ospf-intra | rip | static {address-family [{ipv4-unicast | ipv4-multicast}] shutdown-priority  
<number>
```

Description

Configures the shutdown priority for IGP export.

Syntax Description

direct	Specifies direct routing.
--------	---------------------------

ospf	Specifies OSPF routing.
ospf-extern1	Specifies OSPF-extern1 routing.
ospf-extern2	Specifies OSPF-extern2 routing.
ospf-inter	Specifies OSPF-inter routing.
ospf-intra	Specifies OSPF-intra routing.
rip	Specifies RIP routing.
static	Specifies static routing.
address-family	The address family to which the IGP routes are exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
number	Specifies the shutdown priority. The range is 0 - 65,535.

Default

The default value is 2048.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Note: This command is not currently supported, and is not recommended for use.

Higher priority values lower the chance of an IGP export to be automatically disabled in case BGP or the system goes to a low memory condition.

Example

The following command configures the shutdown priority of BGP exported OSPF routes to 1000:

```
configure bgp export ospf shutdown-priority 1000
```

configure bgp import-policy

```
configure bgp import-policy [<policy-name> | none]
```

Description

Configures the import policy for BGP.

Syntax Description

policy-name	Specifies the policy.
-------------	-----------------------

none	Specifies no policy.
------	----------------------

Default

N/A.

Usage Guidelines

Use the none keyword to remove a BGP import policy.

An import policy is used to modify route attributes while adding BGP routes to the IP route table.

Example

The following command configures a policy *imprt_plcy* for BGP:

```
configure bgp import-policy imprt_plcy
```

The following command unconfigures the import policy for BGP:

```
configure bgp import-policy none
```

configure bgp local-preference

```
configure bgp local-preference <number>
```

Description

Changes the default local preference attribute.

Syntax Description

number	Specifies a value used to advertise this router's degree of preference to other routers within the AS. Range is 0 to 2147483647.
--------	--

Default

100.

Usage Guidelines

The range is 0 to 2,147,483,647.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code

- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command changes the default local preference attribute to *500*:

```
configure bgp local-preference 500
```

configure bgp maximum-paths

```
configure bgp maximum-paths <max-paths>
```

Description

Enables or disables the BGP ECMP feature and specifies the maximum number of paths supported on the current VR.

Syntax Description

max-paths	Specifies the maximum number of paths. The range is 1 to 8. The value 1 disables BGP ECMP. A value greater than 1 enables BGP ECMP and specifies the maximum number of paths.
-----------	---

Default

One. BGP ECMP is disabled.

Usage Guidelines

This command triggers the BGP decision process, causing BGP to re-install the entire BGP routing table into the IP forwarding table. This activity requires a significant amount of switch processor resources, so we recommend that you enable or disable the BGP ECMP feature before enabling the BGP protocol globally on a VR. To ensure that BGP ECMP routes are programmed in the hardware, enter the `enable iproute sharing` command.

Example

The following command enables BGP ECMP and sets the maximum number of paths to *4*:

```
configure bgp maximum-paths 4
```

configure bgp med

```
configure bgp med [none | <bgp_med>]
```

Description

Configures the metric to be included in the Multi-Exit-Discriminator (MED) path attribute. The MED path attribute is included in route updates sent to external peers if a value is configured.

Syntax Description

none	Specifies not to use a multi-exist-discriminator number.
bgp_med	Specifies a multi-exist-discriminator number. The range is 0-2147483647.

Default

N/A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command configures the metric to be included in the MED path attribute:

```
configure bgp med 3
```

configure bgp neighbor allowas-in

```
configure bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast | ipv4-multicast]}
allowas-in {max-as-occurrence <as-count>}
```

Description

Configures BGP to receive and accept a looped BGP route from the specified neighbor, provided the number of occurrences of local AS number in AS-Path is less than or equal to the value of `as-count`.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
as-count	The maximum number of occurrences of local AS number in the received route AS-Path. If the number of occurrences of local AS number in AS-Path is more than as-count, the route is not accepted. The valid range is from 1-16.

Default

This feature is disabled by default. as-count defaults to 3. If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound BGP route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled. This feature can also be enabled for both IBGP and EBGP neighbors, wherever necessary.

All BGP routes with looped AS-Path are silently discarded by default.

Example

The following example enables BGP to accept looped BGP routes that contains a maximum of 6 occurrences of receiver's AS-number in AS-Path attribute:

```
configure bgp neighbor 192.162.17.54 allowas-in max-as-occurrence 6
```

configure bgp neighbor dampening

```
configure bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast |
ipv4-multicast]} dampening {{half-life <half-life-minutes> {reuse-limit <reuse-limit-number>
suppress-limit <suppress-limit-number> max-suppress <max-suppress-minutes>}} | policy-filter
[<policy-name> | none]}
```

Description

Configures

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

half-life	Specifies the dampening half life. Range is 1 to 45 minutes.
reuse	Specifies the reuse limit. Range is 1 to 20000.
suppress	Specifies the suppress limit. Range is 1 to 20000.
max-suppress	Specifies the maximum hold down time. Range is 1 to 255 minutes.
policy-filter	Specifies a policy.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The half life is the period of time, in minutes, during which the accumulated penalty of a route is reduced by half. The range is 1 to 45 minutes, and the default is 15 minutes.

The reuse limit is the penalty value below which a route is used again. The range is 1-20,000, and the default is 750.

The suppress limit is the penalty value above which a route is suppressed. The range is 1-20,000, and the default is 2,000.

The maximum hold down time is the maximum time a route can be suppressed, no matter how unstable it has been, as long as it no longer flaps. The range is 1-255 minutes, and the default is 4 * the half life.

Instead of explicitly configuring the dampening parameters using the command line, you can associate a policy using the `policy-filter` option. Multiple sets of parameters can be supplied using a policy.

Use the following command to disable route flap dampening for BGP neighbors:

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast |
ipv4-multicast]} no-dampening
```

Example

The following command configures route flap dampening to the BGP neighbor at 192.168.1.22 to the default values:

```
configure bgp neighbor 192.168.1.22 dampening
```

configure bgp neighbor description

```
configure bgp neighbor [all | <remoteaddr>] description {<description>}
```

Description

Configures a description for a BGP neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
description	Specifies a string used to describe the neighbor.

Default

The description is a NULL string by default.

Usage Guidelines

Use this command to attach a description to a BGP neighbor. This description is displayed in the output of the `show bgp neighbor` command when you specify the `detail` option, or when you specify a particular neighbor. Enclose the string in double quotes if there are any blank spaces in the string. The maximum length of the string is 56 characters.

If you do not specify the `<description>` parameter, the description is reset to the default.

Example

The following command configures the description for the BGP neighbor 192.168.1.22 to *Toledo_5*:

```
configure bgp neighbor 192.168.1.22 description Toledo_5
```

configure bgp neighbor dont-allowas-in

```
configure bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast | ipv4-multicast]}  
dont-allowas-in
```

Description

Disables BGP from receiving and accepting a looped BGP route from the specified neighbor, provided the number of occurrences of local AS number in AS-Path is less than or equal to the value of `as-count`.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

This feature is disabled by default. as-count defaults to 3. If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound BGP route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled. This feature can also be enabled for both IBGP and EBGp neighbors, wherever necessary.

All BGP routes with looped AS-Path are silently discarded by default.

configure bgp neighbor maximum-prefix

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast | ipv4-multicast]}
maximum-prefix <number> {{threshold <percent>} {teardown {holddown-interval <seconds>}}}
{send-traps}
```

Description

Configures the maximum number of IP prefixes accepted from a BGP neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
number	Specifies the maximum number of prefixes that can be accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
percent	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log (and console), and/or a trap is sent to the SNMP manager.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
seconds	Specifies the length of time before the session is re-established, if the session is torn down due to maximum prefix exceeded. If the hold-down interval is zero or not specified, it is kept down until the peer is enabled. The range is 30 to 86400 seconds.
send-traps	Specifies sending "number of prefix reached threshold" and "number of prefix exceed the max-prefix limit" SNMP traps.

Default

This feature is disabled by default.

The default threshold is 75%.

By default, `teardown` is not specified.

By default, `send-traps` is not specified.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Configure the peer group before configuring the neighbors. To configure the peer group, use the following command:

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast |
ipv4-multicast]} maximum-prefix <number> {{threshold <percent>} {teardown
{holddown-interval <seconds>}}} {send-traps}
```

Example

The following command configures the maximum number of IP prefixes accepted from all neighbors to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
configure bgp neighbor all maximum-prefix 5000 threshold 60 send-traps
```

configure bgp neighbor next-hop-self

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast | ipv4-multicast]}
[next-hop-self | no-next-hop-self]
```

Description

Configures the next hop address used in the outgoing updates to be the address of the BGP connection originating the update.

Syntax Description

<code>remoteaddr</code>	Specifies an IP address.
<code>all</code>	Specifies all neighbors.
<code>address-family</code>	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
<code>next-hop-self</code>	Specifies that the next hop address used in the updates be the address of the BGP connection originating it.
<code>no-next-hop-self</code>	Specifies that the next hop address used in the updates not be the address of the BGP connection originating it (lets BGP decide what would be the next hop).

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
configure bgp neighbor 172.16.5.25 next-hop-self
```

configure bgp neighbor no-dampening

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast | ipv4-multicast]} no-dampening
```

Description

Configures no route flap dampening over BGP peer sessions (disables route flap dampening).

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Use the following command to enable route flap dampening for BGP neighbors:

```
configure bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast | ipv4-multicast]} dampening {{half-life <half-life-minutes> {reuse-limit <reuse-limit-number> suppress-limit <suppress-limit-number> max-suppress <max-suppress-minutes>}} | policy-filter [<policy-name> | none]}
```

Example

The following command disables route flap dampening to the BGP neighbor at 192.168.1.22:

```
configure bgp neighbor 192.168.1.22 no-dampening
```

configure bgp neighbor password

```
configure bgp neighbor [all | <remoteaddr>] password [none | {encrypted} <tcpPassword>]
```

Description

Configures an MD5 secret password for a neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.
none	Specifies not to use a password
encrypted	Specifies an encrypted string; do not use.
tcpPassword	Specifies a password string.

Default

N/A.

Usage Guidelines

Disable the BGP neighbor or the BGP protocol before changing the password.

When a password is configured, TCP MD5 authentication is enabled on the TCP connection that is established with the neighbor.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

To change any one of the following parameters you must disable and re-enable the peer session:

- timer
- source-interface
- soft-in-reset
- password

Changing a route reflector client automatically disables and enables the peer session.

The `encrypted` option is used by the switch when generating a configuration file, and when parsing a switch-generated configuration file. Do not select the encrypted option in the CLI.

Example

The following command configures the password for a neighbor as *netgear*:

```
configure bgp neighbor 192.168.1.5 password netgear
```

configure bgp neighbor peer-group

```
configure bgp neighbor [all | <remoteaddr>] peer-group [<peer-group-name> | none]
{acquire-all}
```

Description

Configures an existing neighbor as the member of a peer group.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.
peer-group-name	Specifies a peer group name.
none	Removes the neighbor from the peer group.
acquire-all	Specifies that all parameters should be inherited by the neighbor from the peer group.

Default

By default, remote AS (if configured for the peer group), source-interface, outbound route policy, send-community and next-hop-self settings are inherited.

Usage Guidelines

If `acquire-all` is not specified, only the default parameters are inherited by the neighbor.

When you remove a neighbor from a peer group, it retains the parameter settings of the group. The parameter values are not reset to those the neighbor had before it inherited the peer group values.

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
```

The new neighbor is created as part of the peer group and inherits all of the existing parameters of the peer group. The peer group must have remote AS configured.

Example

The following command configures an existing neighbor as the member of the peer group *outer*:

```
configure bgp neighbor 192.1.1.22 peer-group outer
```

configure bgp neighbor route-policy

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast | ipv4-multicast]}
route-policy [in | out] [none | <policy>]
```


Description

Configures a route map filter for a neighbor.

Syntax Description

remoteaddr	Specifies an IP address.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
policy	Specifies a policy.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The policy can be installed on the input or output side of the router. The policy is used to modify or filter the NLRI information and the path attributes associated with it when exchanging updates with the neighbor.

Example

The following command configures the route-policy filter for a neighbor based on the policy *nosales*:

```
configure bgp neighbor 192.168.1.22 route-policy in nosales
```

configure bgp neighbor route-reflector-client

```
configure bgp neighbor [<remoteaddr> | all] [route-reflector-client |  
no-route-reflector-client]
```

Description

Configures a BGP neighbor to be a route reflector client.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.

<code>route-reflector-client</code>	Specifies for the BGP neighbor to be a route reflector client.
<code>no-route-reflector-client</code>	Specifies for the BGP neighbor not to be a route reflector client.

Default

N/A.

Usage Guidelines

Another way to overcome the difficulties of creating a fully-meshed AS is to use *route reflectors*. Route reflectors allow a single router to serve as a central routing point for the AS or sub-AS.

Use this command to implicitly define the router to be a route reflector. The neighbor must be in the same AS as the router.

When changing the route reflector status of a peer, the peer is automatically disabled and re-enabled and a warning message appears on the console and in the log.

A *cluster* is formed by the route reflector and its client routers. Peer routers that are not part of the cluster must be fully meshed according to the rules of BGP.

Example

The following command configures a BGP neighbor to be a route reflector client:

```
configure bgp neighbor 192.168.1.5 route-reflector-client
```

configure bgp neighbor send-community

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast | ipv4-multicast]}
[send-community | dont-send-community] {both | extended | standard}
```

Description

Configures whether the community path attribute associated with a BGP NLRI should be included in the route updates sent to the BGP neighbor.

Syntax Description

<code>remoteaddr</code>	Specifies an IP address of a BGP neighbor.
<code>all</code>	Specifies all neighbors.
<code>address-family</code>	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
<code>send-community</code>	Specifies to include the community path attribute.
<code>dont-send-community</code>	Specifies not to include the community path attribute.

both	Send both standard and extended community attributes to this BGP neighbor, or neighbors in peer group
extended	Send only extended communities to this BGP neighbor or neighbors in peer group
standard	Send only standard communities to this BGP neighbor or neighbors in peer group

Default

If no address family is specified, IPv4 unicast is the default. If no optional keyword (`both`, `standard` or `extended`) is specified, `standard` is assumed.

Usage Guidelines

A BGP community is a group of BGP destinations that require common handling. The NETGEAR 8800 supports the following well-known BGP community attributes:

- `no-export`
- `no-advertise`
- `no-export-subconfed`

The command is additive; that is, if the command is executed twice with the `standard` or `extended` option, both the `extended` and `standard` communities are sent to the BGP neighbor.

Example

The following command includes the community path attribute associated with a BGP NLRI in the route updates sent to all BGP neighbors:

```
configure bgp neighbor all send-community
```

configure bgp neighbor shutdown-priority

```
configure bgp neighbor [all | <remoteaddr>] shutdown-priority <number>
```

Description

Configures the shutdown priority for a BGP neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
number	Specifies the shutdown priority. The range is 0 - 65,535.

Default

The default value is 1024.

Usage Guidelines

Note: This command is not currently supported, and is not recommended for use.

Higher priority values lower the chance of a BGP neighbor to be automatically disabled in case BGP or the system goes to a low memory condition.

Example

The following command configures the shutdown priority of the BGP neighbor 10.0.20.1 to 500:

```
configure bgp neighbor 10.0.20.1 shutdown-priority 1000
```

configure bgp neighbor soft-reset

```
configure bgp neighbor [<remoteaddr> | all] {address-family [ipv4-unicast | ipv4-multicast]}
soft-reset {in | out}
```

Description

Applies the current input or output routing policy to the routing information already exchanged with the neighbor.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
soft-reset	Do a soft reconfiguration for the BGP neighbor.
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The input/output policy is determined by the route policy configured for the neighbor on the input and/or output side of the router. This command does not affect the switch configuration.

If both the local BGP neighbor and the neighbor router support the route refresh capability (NETGEAR 8800 does not support this feature), a dynamic soft input reset can be performed. The `configure bgp neighbor soft-reset` command triggers the generation of a *Route-Refresh* message to the neighbor. As a response to the *Route-Refresh* message, the neighbor sends the entire BGP routing table in updates and the switch applies the appropriate routing policy to the updates.

If the route-refresh capability is not supported by the neighbor (like NETGEAR 8800), the `configure bgp neighbor soft-reset` command reprocesses the BGP route database using the policy configured for that neighbor.

Example

The following command applies the current input routing policy to the routing information already exchanged with the neighbor:

```
configure bgp neighbor 192.168.1.5 soft-reset in
```

configure bgp neighbor source-interface

```
configure bgp neighbor [<remoteaddr> | all] source-interface [any | ipaddress <ipAddr>]
```

Description

Changes the BGP source interface for TCP connections.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
all	Specifies all neighbors.
any	Specifies any source interface.
ipAddr	Specifies the IP address of a source interface.

Default

Any.

Usage Guidelines

The source interface IP address must be a valid IP address of any VLAN configured on the switch.

Example

The following command changes the BGP source interface to 10.43.55.10:

```
configure bgp neighbor 192.168.1.5 source-interface ipaddress 10.43.55.10
```

configure bgp neighbor timer

```
configure bgp neighbor [<remoteaddr> | all] timer keep-alive <keepalive> hold-time <holdtime>
```

Description

Configures the BGP neighbor timers.

Syntax Description

remoteaddr	Specifies an IP address.
all	Specifies all neighbors.
keepalive	Specifies a BGP neighbor timer keepalive time in seconds. The range is 0 to 21,845 seconds.
holdtime	Specifies a BGP neighbor timer hold time in seconds. The range is 0 and 3 to 65,535 seconds.

Default

The default keepalive setting is 60 seconds. The default hold time is 180 seconds.

Usage Guidelines

The BGP neighbor or BGP protocol must be disabled before changing the timer values.

Example

The following command configures the BGP neighbor timers:

```
configure bgp neighbor 192.168.1.5 timer keep-alive 120 hold-time 360
```

configure bgp neighbor weight

```
configure bgp neighbor [<remoteaddr> | all] weight <weight>
```

Description

Assigns a locally-used weight to a neighbor connection for the route selection algorithm.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
all	Specifies all neighbors.
weight	Specifies a BGP neighbor weight.

Default

By default, the weight is 0.

Usage Guidelines

All routes learned from this peer are assigned the same weight. The route with the highest weight is more preferable when multiple routes are available to the same network. The range is 0 to 65,535.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command assigns a locally used weight of 10 to a neighbor connection:

```
configure bgp neighbor 192.168.1.5 weight 10
```

configure bgp peer-group allowin-as

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
allowas-in {max-as-occurrence <as-count>}
```

Description

Configures BGP to receive and accept a looped BGP route from the neighbors of the specified peer group, provided the number of occurrences of local AS number in AS-Path is less than or equal to that specified in `as-count`.

Syntax Description

peer-group-name	Specifies a peer group
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
as-count	The maximum number of occurrences of local AS number in the received route AS-Path. If the number of occurrences of local AS number in AS-Path is more than as-count, the route is not accepted. The valid range is from 1-16.

Default

This feature is disabled by default. as-count defaults to 3. If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound BGP route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled.

This feature can also be enabled for both IBGP and EBGP neighbors, wherever necessary.

Note: BGP neighbors do not inherit the `allowas-in` configuration from their peer group unless you explicitly specify the `acquire-all` option when adding a neighbor to a peer-group.

Example

The following example enables BGP to accept looped BGP routes that contains a maximum of 8 occurrences of receiver's AS-number in AS-Path attribute:

```
configure bgp peer-group internal allowas-in max-as-occurrence 8
```

configure bgp peer-group dampening

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
dampening {{half-life <half-life-minutes> {reuse-limit <reuse-limit-number> suppress-limit
<suppress-limit-number> max-suppress <max-suppress-minutes>}} | policy-filter [<policy-name>
| none}}
```

Description

Configures route flap dampening for a BGP peer group.

Syntax Description

peer-group-name	Specifies a peer group
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
half-life-minutes	Specifies the dampening half life.
reuse-limit-number	Specifies the reuse limit.
suppress-limit-number	Specifies the suppress limit.
max-suppress-minutes	Specifies the maximum hold down time.
policy-name	Specifies a policy

none	Removes any policy association.
------	---------------------------------

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The half life is the period of time, in minutes, during which the accumulated penalty of a route is reduced by half. The range is 1 to 45 minutes, and the default is 15 minutes.

The reuse limit is the penalty value below which a route is used again. The range is 1-20,000, and the default is 750.

The suppress limit is the penalty value above which a route is suppressed. The range is 1-20,000, and the default is 2,000.

The maximum hold down time is the maximum time a route can be suppressed, no matter how unstable it has been, as long as it no longer flaps. The range is 1-255 minutes, and the default is 4 * the half life.

Instead of explicitly configuring the dampening parameters using the command line, you can associate a policy using the `policy-filter` option. Multiple sets of parameters can be supplied using a policy.

Use the following command to disable route flap dampening for a BGP peer-group:

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast |
ipv4-multicast]} no-dampening
```

Example

The following command configures route flap dampening for the BGP peer group *outer*:

```
configure bgp peer-group outer dampening
```

configure bgp peer-group dont-allowin-as

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
dont-allowas-in
```

Description

Disables BGP from receiving and accepting a looped BGP route from the neighbors of the specified peer group, provided the number of occurrences of local AS number in AS-Path is less than or equal to that specified in `as-count`.

Syntax Description

peer-group-name	Specifies a peer group
-----------------	------------------------

address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
----------------	--

Default

This feature is disabled by default. as-count defaults to 3. If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound BGP route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled.

This feature can also be enabled for both IBGP and EBGP neighbors, wherever necessary.

Note: BGP neighbors do not inherit the `allowas-in` configuration from their peer group unless you explicitly specify the `acquire-all` option when adding a neighbor to a peer-group.

configure bgp peer-group maximum-prefix

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
maximum-prefix <number> [{threshold <percent>} {teardown {holddown-interval <seconds>}}]
{send-traps}
```

Description

Configures the maximum number of IP prefixes accepted for all neighbors in the peer group.

Syntax Description

name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
number	Specifies the maximum number of prefixes that can be accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
percent	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log (and on the console). An SNMP trap can also be sent.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
seconds	Specifies the length of time before the session is re-established, if the session has been torn down due to exceeding the max limit. If the hold down interval is 0 or not specified, it is kept down until the peer is enabled. The range is 30 to 86400 seconds.

send-traps	Specifies sending “number of prefix reached threshold” and “number of prefix exceed the max-prefix limit” SNMP traps.
------------	---

Default

This feature is disabled by default.

The default threshold is 75%.

By default, `teardown` is not specified.

By default, `send-traps` is not specified.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Configure the peer group before configuring the neighbors. To configure the neighbors, use the following command:

```
configure bgp neighbor 192.168.1.1 maximum-prefix
```

Example

The following command configures the maximum number of IP prefixes accepted from the peer group *outer* to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
configure bgp peer-group outer maximum-prefix 5000 threshold 60 send-traps
```

configure bgp peer-group next-hop-self

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
[next-hop-self | no-next-hop-self]
```

Description

Configures the next hop address used in the updates to be the address of the BGP connection originating the update.

Syntax Description

peer-group-name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
next-hop-self	Specifies that the next hop address used in the updates be the address of the BGP connection originating it.
no-next-hop-self	Specifies that the next hop address used in the updates not be the address of the BGP connection originating it (Let the BGP protocol decide the next hop).

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
configure bgp peer-group outer next-hop-self
```

configure bgp peer-group no-dampening

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}  
no-dampening
```

Description

Configures no route flap dampening for a BGP peer group (disables route flap dampening).

Syntax Description

peer-group-name	Specifies a BGP peer group.
ipv4-unicast	Specifies the address family IPV4 unicast
ipv4-multicast	Specifies the address family IPV4 multicast.

Default

This feature is disabled by default.

Usage Guidelines

Use the following command to enable route flap dampening for a BGP peer-group:

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast |  
ipv4-multicast]} dampening {{half-life <half-life-minutes> {reuse-limit  
<reuse-limit-number> supress-limit <suppress-limit-number> max-suppress  
<max-suppress-minutes>}} | policy-filter [<policy-name> | none]}
```

Example

The following command disables route flap dampening to the BGP peer group *outer*:

```
configure bgp peer-group outer no-dampening
```

configure bgp peer-group password

```
configure bgp peer-group <peer-group-name> password [none | <tcpPassword>]
```

Description

Configures the TCP MD5 secret password for a peer group and all neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
none	Specifies no password.
tcpPassword	Specifies a password.

Default

N/A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the password as *netgear* for the peer group *outer* and its neighbors:

```
configure bgp peer-group outer password netgear
```

configure bgp peer-group remote-AS-number

```
configure bgp peer-group <peer-group-name> remote-AS-number <number>
```

Description

Configures the remote AS number for a peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
number	Specifies a remote AS number. The range is 1 to 4294967295.

Default

N/A.

Usage Guidelines

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in *RFC 5396, Textual Representation of Autonomous System (AS) Numbers*.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the remote AS number for the peer group *outer* and its neighbors using the ASPLAIN 4-byte AS number format:

```
configure bgp peer-group outer remote-AS-number 300000000
```

The following command configures the remote AS number for the peer group *abc* and its neighbors using the ASDOT 4-byte AS number format:

```
configure bgp peer-group abc remote-AS-number 9.10176
```

configure bgp peer-group route-policy

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
route-policy [in | out] [none | <policy>]
```

Description

Configures the policy for a peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
ipv4-unicast	Specifies the address family IPV4 unicast
ipv4-multicast	Specifies the address family IPV4 multicast.
in	Specifies to install the policy on the input side.
out	Specifies to install the policy on the output side.
none	Specifies to remove the filter.
policy	Specifies a policy.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the route policy for the peer group *outer* and its neighbors using the policy *nosales*:

```
configure bgp peer-group outer route-policy in nosales
```

configure bgp peer-group route-reflector-client

```
configure bgp peer-group <peer-group-name> [route-reflector-client |  
no-route-reflector-client]
```

Description

Configures all the peers in a peer group to be a route reflector client.

Syntax Description

peer-group-name	Specifies a peer group.
route-reflector-client	Specifies that all the neighbors in the peer group be a route reflector client.
no-route-reflector-client	Specifies that all the neighbors in the peer group not be a route reflector client.

Default

N/A.

Usage Guidelines

This command implicitly defines this router to be a route reflector.

The peer group must be in the same AS of this router.

Example

The following command configures the peer group *outer* as a route reflector client:

```
configure bgp peer-group outer route-reflector-client
```

configure bgp peer-group send-community

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
[send-community | dont-send-community] {both | extended | standard}
```

Description

Configures whether communities should be sent to neighbors as part of route updates.

Syntax Description

peer-group-name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
send-community	Specifies that communities are sent to neighbors as part of route updates.
dont-send-community	Specifies that communities are not sent to neighbors as part of route updates.
both	Send both standard and extended community attributes to this BGP neighbor, or neighbors in peer group
extended	Send only extended communities to this BGP neighbor or neighbors in peer group
standard	Send only standard communities to this BGP neighbor or neighbors in peer group

Default

If no address family is specified, IPv4 unicast is the default. If no optional keyword (*both*, *standard* or *extended*) is specified, *standard* is assumed.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

The command is additive; that is, if the command is executed twice with the *standard* or *extended* option, both the *extended* and *standard* communities are sent to the BGP neighbor.

Example

The following command configures communities to be sent to neighbors as part of route updates:

```
configure bgp peer-group outer send-community
```

configure bgp peer-group soft-reset

```
configure bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
soft-reset {in | out}
```

Description

Applies the current input/output routing policy to the neighbors in the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
ipv4-unicast	Specifies the address family IPV4 unicast
ipv4-multicast	Specifies the address family IPV4 multicast.
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

N/A.

Usage Guidelines

The input/output routing policy is determined by the route policy configured for the neighbors in the peer group on the input/output side of the router. This command does not affect configuration of the switch.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command applies the current input routing policy to the neighbors in the peer group *outer*:

```
configure bgp peer-group outer soft-reset in
```

configure bgp peer-group source-interface

```
configure bgp peer-group <peer-group-name> source-interface [any | ipaddress <ipAddr>]
```

Description

Configures the source interface for a peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
any	Specifies any source interface.
ipAddr	Specifies an interface.

Default

N/A.

Usage Guidelines

The source interface IP address must be a valid IP address of a VLAN configured on the switch.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the source interface for the peer group *outer* and its neighbors on 10.34.25.10:

```
configure bgp peer-group outer source-interface ipaddress 10.34.25.10
```

configure bgp peer-group timer

```
configure bgp peer-group <peer-group-name> timer keep-alive <seconds> hold-time <seconds>
```

Description

Configures the keepalive timer and hold timer values for a peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
keep-alive seconds	Specifies a keepalive time in seconds. Range is 0 to 21845.
hold-time seconds	Specifies a hold-time in seconds. Range is 0 and 3 to 65535.

Default

N/A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the keepalive timer and hold timer values for the peer group *outer* and its neighbors:

```
configure bgp peer-group outer timer keep-alive 30 hold-time 90
```

configure bgp peer-group weight

```
configure bgp peer-group <peer-group-name> weight <number>
```

Description

Configures the weight for the peer group and all the neighbors of the peer group.

Syntax Description

peer-group-name	Specifies a peer group.
number	Specifies a BGP peer group weight. Range is 0 to 65535.

Default

N/A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command configures the weight for the peer group *outer* and its neighbors:

```
configure bgp peer-group outer weight 5
```

configure bgp restart

```
configure bgp restart [none | planned | unplanned | both | aware-only]
```

Description

Configures the router as a graceful BGP restart router.

Syntax Description

none	Do not act as a graceful BGP restart router.
planned	Only act as a graceful BGP restart router for planned restarts.
unplanned	Only act as a graceful BGP restart router for unplanned restarts.
both	Act as a graceful BGP restart router for both planned and unplanned restarts.
aware-only	Only act as a graceful BGP receiver (helper) router.

Default

The default is none; graceful restart is disabled.

Usage Guidelines

This command configures the router as a graceful BGP router. You can decide to configure a router to enter graceful restart for only planned restarts, for only unplanned restarts, or for both. Also, you can decide to configure a router to be a receiver only (which helps a restarting BGP router to perform the graceful restart process), and not to do graceful restarts itself.

This command cannot be used while BGP is enabled globally on the switch.

Example

The following command configures a router to perform graceful BGP restarts only for planned restarts:

```
configure bgp restart planned
```

configure bgp restart address-family

```
configure bgp restart [add | delete] address-family [ipv4-unicast | ipv4-multicast]
```

Description

Configures the address family used with graceful BGP restart.

Syntax Description

add	Add the address family.
delete	Remove the address family.
ipv4-unicast	IPv4 unicast addresses.
ipv4-multicast	IPv4 multicast addresses.

Default

The default is IPv4 unicast.

Usage Guidelines

This command configures the address family participating in graceful BGP restart. An address family can be added or deleted. By adding an address family, BGP instructs the local hardware and software to preserve BGP routes of that address family during a graceful restart. The local OPEN message contains all the added address families.

This command cannot be used while BGP is enabled globally on the switch.

Note: For BGP graceful restart to inter-operate with Cisco routers, any restarting routers connected to Cisco routers must be configured with the command, `enable bgp neighbor capability`, in the following form, `enable bgp neighbor <remoteaddr> capability ipv4-unicast`. The command must be executed before BGP is enabled globally on the switch.

Example

The following command configures a router to add IPv4 unicast addresses to graceful BGP restarts:

```
configure bgp restart add address-family ipv4-unicast
```

configure bgp restart restart-time

```
configure bgp restart restart-time <seconds>
```

Description

Configures the restart time used with graceful BGP restart. This is the maximum time a receiver router waits for a restarting router to come back up.

Syntax Description

seconds	Specifies the restart time. The range is 1 to 3600 seconds.
---------	---

Default

The default is 120 seconds

Usage Guidelines

This command configures the restart timer. This timer is started on the receiver router when it detects the neighbor router is restarting (usually when the peer TCP session is reset). At that time, routes from the restarting router are marked as stale, but are preserved in the routing table. The timer is stopped when the restarting BGP neighbor goes to the ESTABLISHED state (it has finished restarting). If the timer expires, the stale routes are deleted.

Example

The following command configures the graceful BGP restart timer:

```
configure bgp restart restart-time 200
```

configure bgp restart stale-route-time

```
configure bgp restart stale-route-time <seconds>
```

Description

Configures the stale route timer used with graceful BGP restart. This is the maximum time to hold stale paths on receiver routers while its neighbor gracefully restarts.

Syntax Description

seconds	Specifies the stale route time. The range is 1 to 3600 seconds.
---------	---

Default

The default is 360 seconds

Usage Guidelines

This command configures the stale route timer. This timer is started when the restarting BGP peer goes to the ESTABLISHED state after it restarts. The timer is stopped when the restarting BGP peer sends EOR messages for all address families. When the timer is stopped, or it expires, the stale routes are deleted.

Example

The following command configures the graceful BGP stale route timer:

```
configure bgp restart stale-route-time 400
```

configure bgp restart update-delay

```
configure bgp restart update-delay <seconds>
```

Description

Configures the update delay timer used with graceful BGP restart. This is the maximum time to delay updating BGP routes to the local IP route table.

Syntax Description

seconds	Specifies the stale route time. The range is 1 to 3600 seconds.
---------	---

Default

The default is 600 seconds

Usage Guidelines

This command configures the update delay timer. Usually, a restarting router waits to receive EOR messages from all the receiving BGP neighbors before it starts the route update. Otherwise, it does the route selection when the timer expires.

Example

The following command configures the graceful BGP update delay timer:

```
configure bgp restart update-delay 800
```

configure bgp routerid

```
configure bgp routerid <router identifier>
```

Description

Changes the router identifier.

Syntax Description

router identifier	Specifies a router identifier in the IPv4 address format.
-------------------	---

Default

N/A.

Usage Guidelines

BGP must be disabled before changing the router ID.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest router ID

Example

The following command changes the router ID:

```
configure bgp routerid 192.1.1.13
```


configure bgp soft-reconfiguration`configure bgp soft-reconfiguration`**Description**

Immediately applies the route policy associated with the network command, aggregation, import, and redistribution.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command does not affect the switch configuration.

Example

The following command applies the route policy associated with the network command, aggregation, import, and redistribution:

`configure bgp soft-reconfiguration`***create bgp neighbor peer-group***`create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}`**Description**

Creates a new neighbor and makes it part of the peer group.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
peer-group-name	Specifies a peer group.
multi-hop	Specifies to allow connections to EBGp peers that are not directly connected.

Default

N/A.

Usage Guidelines

All the parameters of the neighbor are inherited from the peer group. The peer group should have the remote AS configured.

To add an existing neighbor to a peer group, use the following command:

```
configure bgp neighbor [all | <remoteaddr>] peer-group [<peer-group-name> | none]
{acquire-all}
```

If you do not specify `acquire-all`, only the mandatory parameters are inherited from the peer group. If you specify `acquire-all`, all of the parameters of the peer group are inherited. This command disables the neighbor before adding it to the peer group.

Example

The following command creates a new neighbor and makes it part of the peer group *outer*:

```
create bgp neighbor 192.1.1.22 peer-group outer
```

create bgp neighbor remote-AS-number

```
create bgp neighbor <remoteaddr> remote-AS-number <number> {multi-hop}
```

Description

Creates a new BGP peer.

Syntax Description

<code>remoteaddr</code>	Specifies an IP address of the BGP neighbor.
<code>number</code>	Specifies a remote AS number. The range is 1 to 4294967295.
<code>multi-hop</code>	Specifies to allow connections to EBGP peers that are not directly connected.

Default

N/A.

Usage Guidelines

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in *RFC 5396, Textual Representation of Autonomous System (AS) Numbers*.

If the AS number is the same as the AS number provided in the `configure bgp as` command, then the peer is consider an IBGP peer, otherwise the neighbor is an EBGP peer. The BGP session to a newly created peer is not started until the `enable bgp neighbor` command is issued.

Examples

The following command specifies a BGP peer AS number using the ASPLAIN 4-byte AS number format:

```
create bgp neighbor 14.0.0.1 remote-AS-number 4000000000
```

The following command specifies a BGP peer AS number using the ASDOT 4-byte AS number format:

```
create bgp neighbor 16.0.0.1 remote-as-number 1.20
```

create bgp peer-group

```
create bgp peer-group <peer-group-name>
```

Description

Creates a new peer group.

Syntax Description

peer-group-name	Specifies a peer group.
-----------------	-------------------------

Default

N/A.

Usage Guidelines

You can use BGP peer groups to group together up to 512 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-route-policy
- send-community
- next-hop-self

The BGP peer group name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

Example

The following command creates a new peer group named *outer*:

```
create bgp peer-group outer
```

delete bgp neighbor

```
delete bgp neighbor [<remoteaddr> | all]
```

Description

Deletes one or all BGP neighbors.

Syntax Description

<code>remoteaddr</code>	Specifies the IP address of the BGP neighbor to be deleted.
<code>all</code>	Specifies all neighbors.

Default

N/A.

Usage Guidelines

Use this command to delete one or all BGP neighbors.

Example

The following command deletes the specified BGP neighbor:

```
delete bgp neighbor 192.168.1.17
```

delete bgp peer-group

```
delete bgp peer-group <peer-group-name>
```

Description

Deletes a peer group.

Syntax Description

<code>peer-group-name</code>	Specifies a peer group.
------------------------------	-------------------------

Default

N/A.

Usage Guidelines

Use this command to delete a specific BGP peer group.

Example

The following command deletes the peer group named *outer*:

```
delete bgp peer-group outer
```

disable bgp`disable bgp`**Description**

Disables BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable BGP on the router.

Example

The following command disables BGP:

```
disable bgp
```

disable bgp advertise-inactive-route

```
disable bgp {address-family [ipv4-unicast | ipv4-multicast]} advertise-inactive-route
```

Description

Disables advertisement of BGP inactive routes, which are defined as those routes that rated *best* by BGP and *not best* in the IP routing table.

Syntax Description

ipv4-unicast	Disables inactive route advertisement for IPv4 unicast routes. If you do not specify an address family, the command applies to IPv4 unicast routes.
ipv4-multicast	Disables inactive route advertisement for IPv4 multicast routes.

Default

Disabled.

Usage Guidelines

This command can be successfully executed only when BGP is globally disabled. If you want to disable inactive route advertisement and BGP is enabled, you must disable BGP (`disable bgp`), disable this feature, and then enable BGP (`enable bgp`).

Example

The following command disables inactive route advertisement for IPv4 unicast traffic:

```
disable bgp address-family ipv4-unicast advertise-inactive-route
```

disable bgp aggregation

```
disable bgp aggregation
```

Description

Disables BGP route aggregation.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Use this command to disable BGP route aggregation.

Example

The following command disables BGP route aggregation:

```
disable bgp aggregation
```

disable bgp always-compare-med

```
disable bgp always-compare-med
```

Description

Disables BGP from comparing Multi Exit Discriminators (MEDs) for paths from neighbors in different Autonomous Systems (AS).

Syntax Description

This command has no arguments or variables.

Default

The NETGEAR 8800 does not compare MEDs for paths from neighbors in different AS.

Usage Guidelines

The MED is one of the parameters that is considered when selecting the best path among many alternative paths. The path with a lower MED is preferred over a path with a higher MED. By default, during the best path selection process, MED comparison is done only among paths from the same AS.

Example

The following command disables MED from being used in comparison among paths from different AS:

```
disable bgp always-compare-med
```

disable bgp community format

```
disable bgp community format AS-number : number
```

Description

Disables the AS-number:number format of display for communities in the output of show and upload commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Using this command, communities are displayed as a single decimal value.

Example

The following command disables the AS-number:number format of display for communities:

```
disable bgp community format AS-number : number
```

disable bgp export

```
disable bgp export [blackhole | direct | ospf | ospf-extern1 | ospf-extern2 | ospf-inter |  
ospf-intra | rip | static {address-family [{ipv4-unicast | ipv4-multicast}]}
```

Description

Disables BGP from exporting routes from other protocols to BGP peers.

Syntax Description

blackhole	Specifies blackhole routes.
direct	Specifies direct routes.
ospf	Specifies OSPF routes.
ospf-extern1	Specifies OSPF-extern1 routes.
ospf-extern2	Specifies OSPF-extern2 routes.
ospf-inter	Specifies OSPF-inter routes.
ospf-intra	Specifies OSPF-intra routes.
rip	Specifies RIP routes.
static	Specifies static routes.
address-family	The address family to which the IGP routes is exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The exporting of routes between any two routing protocols is a discrete configuration function. For example, you must configure the switch to export routes from OSPF to BGP and, if desired, you must configure the switch to export routes from BGP to OSPF. You must first configure both protocols and then verify the independent operation of each. Then you can configure the routes to export from OSPF to BGP, and the routes to export from BGP to OSPF. Similarly for BGP and RIP.

You can use policies to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. Policies can also be used to filter out exported routes.

Using the `export` command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command disables BGP from exporting routes from the OSPF protocol to BGP peers:


```
disable bgp export ospf
```

disable bgp fast-external-fallover

```
disable bgp fast-external-fallover
```

Description

Disables BGP fast external fallover functionality.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the BGP fast external fallover on the router. This command applies to all directly-connected external BGP neighbors.

When BGP fast external fallover is enabled, the directly-connected EBGP neighbor session is immediately reset when the connecting link goes down.

If BGP fast external fallover is disabled, BGP waits until the default hold timer expires (3 keepalives) to reset the neighboring session. In addition, BGP might teardown the session somewhat earlier than hold timer expiry if BGP detects that the TCP session and it's directly connected link is broken (BGP detects this while sending or receiving data from TCP socket).

Example

The following command disables BGP fast external fallover:

```
disable bgp fast-external-fallover
```

disable bgp neighbor

```
disable bgp neighbor [<remoteaddr> | all]
```

Description

Disables the BGP session.

Syntax Description

remoteaddr	Specifies an IP address of the BGP neighbor.
all	Specifies all neighbors.

Default

Disabled.

Usage Guidelines

After the session has been disabled, all the information in the route information base (RIB) for the neighbor is flushed.

Example

The following command disables the BGP session:

```
disable bgp neighbor 192.1.1.17
```

disable bgp neighbor capability

```
disable bgp neighbor [all | <remoteaddr>] capability [ipv4-unicast | ipv4-multicast | route-refresh]
```

Description

This command disables BGP Multiprotocol (MP) and route-refresh capabilities for neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of the BGP neighbor.
ipv4-unicast	Specifies BGP MP unicast capabilities.
ipv4-multicast	Specifies BGP MP multicast capabilities.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are enabled by default.

Usage Guidelines

Note: This command is not generally supported, and is not recommended for general use. However, to inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

This command disables BGP Multiprotocol and route-refresh capabilities for one or all neighbors. After the capabilities have been enabled, the BGP neighbor announces its capabilities to neighbors in an OPEN message.

Example

The following command disables the route-refresh feature for all neighbors:

```
disable bgp neighbor all capability route-refresh
```

disable bgp neighbor originate-default

```
disable bgp [{neighbor} <remoteaddr> | neighbor all] {address-family [ipv4-unicast |  
ipv4-multicast]} originate-default
```

Description

Removes a default route to a single BGP neighbor or to all BGP neighbors.

Syntax Description

neighbor remoteaddr	Specifies the IP address of a BGP neighbor for which the default route is removed.
neighbor all	Specifies that default routes are to be removed for all BGP neighbors.
ipv4-unicast	Specifies that the removed default routes apply to IPv4 unicast routes. If you do not specify an address family, the command applies to IPv4 unicast routes.
ipv4-multicast	Specifies that the removed default routes apply to IPv4 multicast routes.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peer is enabled or disabled.

Example

The following command removes default routes for IPv4 unicast traffic for all BGP peer nodes:

```
disable bgp neighbor all originate-default
```

disable bgp neighbor remove-private-AS-numbers

```
disable bgp neighbor [<remoteaddr> | all] remove-private-AS-numbers
```

Description

Disables the removal of private AS numbers from the AS path in route updates sent to EBGp peers.

Syntax Description

remoteaddr	Specifies an IP address.
all	Specifies all neighbors.

Default

Disabled.

Usage Guidelines

Private AS numbers are AS numbers in the range 64512 through 65534. You can remove private AS numbers from the AS path attribute in updates that are sent to external BGP (EBGP) neighbors. Possible reasons for using private AS numbers include:

- The remote AS does not have officially allocated AS numbers.
- You want to conserve AS numbers if you are multi-homed to the local AS.

Private AS numbers should not be advertised on the Internet. Private AS numbers can only be used locally within an administrative domain. Therefore, when routes are advertised out to the Internet, the private AS number can be stripped out from the AS paths of the advertised routes using this feature.

Example

The following command disables the removal of private AS numbers from the AS path in route updates sent to the EBGp peers:

```
disable bgp neighbor 192.168.1.17 remove-private-AS-numbers
```

disable bgp neighbor soft-in-reset

```
disable bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast | ipv4-multicast]}
soft-in-reset
```

Description

Disables the soft input reset feature.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address.

address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
----------------	---

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

This command can be issued only when both BGP and the BGP neighbor is disabled.

Example

The following command disables the soft input reset for the neighbor at 192.168.1.17:

```
disable bgp neighbor 192.168.1.17 soft-in-reset
```

disable bgp neighbor use-ip-router-alert

```
disable bgp neighbor [all | <remoteaddr>] use-ip-router-alert
```

Description

Disables the router alert IP option in outgoing BGP messages to the specified neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of the BGP neighbor.

Default

Disabled.

Usage Guidelines

The IP router alert option in a BGP message forces intermediate routers to examine the packet very closely and therefore, indirectly, gives greater reliability that a packet is delivered to its destination.

Example

The following command disables the feature:

```
disable bgp neighbor 192.168.1.17 use-ip-router-alert
```

disable bgp peer-group

```
disable bgp peer-group <peer-group-name>
```

Description

Disables a BGP peer group and all its BGP neighbors.

Syntax Description

peer-group-name	Specifies a peer group.
-----------------	-------------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the BGP peer group *outer*.

```
disable bgp peer-group outer
```

disable bgp peer-group capability

```
disable bgp peer-group <peer-group-name> capability [ipv4-unicast | ipv4-multicast |  
route-refresh]
```

Description

Disables BGP Multiprotocol (MP) and route-refresh capabilities for a peer-group.

Syntax Description

peer-group-name	Specifies a peer group.
ipv4-unicast	Specifies BGP MP unicast capabilities.
ipv4-multicast	Specifies BGP MP multicast capabilities.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are disabled by default.

Usage Guidelines

Note: This command is not generally supported, and is not recommended for general use. However, to inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

This command disables BGP Multiprotocol and route-refresh capabilities for a peer group. Once the capabilities are enabled, the BGP peer announces its capabilities to neighbors in an OPEN message

Example

The following command disables the route-refresh feature for the peer group *outer*:

```
disable bgp peer-group outer route-refresh
```

disable bgp peer-group originate-default

```
disable bgp {peer-group} <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}  
originate-default
```

Description

Removes default routes to all BGP neighbors in the specified peer group.

Syntax Description

peer-group-name	Specifies the BGP peer group for which the default routes are removed.
ipv4-unicast	Specifies that the default routes apply to IPv4 unicast routes. If you do not specify an address family, the command applies to IPv4 unicast routes.
ipv4-multicast	Specifies that the default routes apply to IPv4 multicast routes.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peers are enabled or disabled.

Example

The following command removes default routes for IPv4 unicast traffic for all nodes in the *test* BGP peer group:

```
disable bgp peer-group test originate-default
```

disable bgp peer-group remove-private-AS-numbers

```
disable bgp peer-group <peer-group-name> remove-private-AS-numbers
```

Description

Disables the removal of private autonomous system (AS) numbers from the AS_Path attribute of outbound updates.

Syntax Description

peer-group-name	Specifies a peer group.
-----------------	-------------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the BGP peer group *outer* from removing private AS numbers:

```
disable bgp peer-group outer remove-private-AS-numbers
```

disable bgp peer-group soft-in-reset

```
disable bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]} soft-in-reset
```

Description

Disables the soft input reset feature.

Syntax Description

peer-group-name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

Example

The following command disables the soft input reset feature:

```
disable bgp peer-group outer soft-in-reset
```

disable bgp peer-group use-ip-router-alert

```
disable bgp peer-group <peer-group-name> use-ip-router-alert
```

Description

Disables the router alert IP option in outgoing BGP messages to the specified peer group.

Syntax Description

peer-group-name	Specifies a peer group.
-----------------	-------------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the feature for the peer group *outer*:

```
disable bgp peer-group outer use-ip-router-alert
```

enable bgp

```
enable bgp
```

Description

Enables BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the Border Gateway Protocol (BGP) on the router. Before invoking this command, the local AS number and BGP router ID must be configured.

Example

The following command enables BGP:

```
enable bgp
```

enable bgp advertise-inactive-route

```
enable bgp {address-family [ipv4-unicast | ipv4-multicast]} advertise-inactive-route
```

Description

Enables advertisement of BGP inactive routes, which are defined as those routes that rated *best* by BGP and *not best* in the IP routing table.

Syntax Description

ipv4-unicast	Enables inactive route advertisement for IPv4 unicast routes. If you do not specify an address family, the command applies to IPv4 unicast routes.
ipv4-multicast	Enables inactive route advertisement for IPv4 multicast routes.

Default

Disabled.

Usage Guidelines

This command can be successfully executed only when BGP is globally disabled. It is best to enable this feature before you enable BGP (`enable bgp`). If BGP is enabled, you must disable BGP (`disable bgp`), enable this feature, and then enable BGP.

Example

The following command enables inactive route advertisement for IPv4 unicast traffic:

```
enable bgp address-family ipv4-unicast advertise-inactive-route
```

enable bgp aggregation

```
enable bgp aggregation
```

Description

Enables BGP route aggregation.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

To use BGP route aggregation, follow these steps:

1. Enable aggregation using the following command:

```
enable bgp aggregation
```

2. Create an aggregate route using the following command:

```
configure bgp add aggregate-address {address-family [ipv4-unicast |  
ipv4-multicast]} <ipaddress> {as-match | as-set} {summary-only}  
{advertise-policy <policy>} {attribute-policy <policy>}
```

Example

The following command enables BGP route aggregation:

```
enable bgp aggregation
```

enable bgp always-compare-med

```
enable bgp always-compare-med
```

Description

Enables BGP to use the Multi Exit Discriminator (MED) from neighbors in different autonomous systems (ASs) in the route selection algorithm.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

MED is only used when comparing paths from the same AS, unless `always-compare-med` is enabled. When this command is issued, MEDs from different AS are used in comparing paths. A MED value of zero is treated as the lowest MED and therefore the most preferred route.

Example

The following command enables BGP to use the Multi Exit Discriminator (MED) from neighbors in different autonomous systems in the route selection algorithm:

```
enable bgp always-compare-med
```

enable bgp community format

```
enable bgp community format AS-number : number
```

Description

Enables the `as-number:number` format of display for the communities in the output of `show` and `upload` commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If not enabled, the communities are displayed as a single decimal value.

Example

The following command enables the `AS-number:number` format of display for communities:

```
enable bgp community format AS-number : number
```

enable bgp export

```
enable bgp export [blackhole | direct | ospf | ospf-extern1 | ospf-extern2 | ospf-inter |  
ospf-intra | rip | static {address-family [{ipv4-unicast | ipv4-multicast}] } {export-policy  
<policy-name>}]
```

Description

Enables BGP to export routes from other protocols to BGP peers.

Syntax Description

blackhole	Specifies blackhole routes.
direct	Specifies direct routes.
ospf	Specifies OSPF routes.
ospf-extern1	Specifies OSPF-extern1 routes.
ospf-extern2	Specifies OSPF-extern2 routes.
ospf-inter	Specifies OSPF-inter routes.
ospf-intra	Specifies OSPF-intra routes.
rip	Specifies RIP routes.
static	Specifies static routes.
address-family	The address family to which the IGP routes are exported. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
policy-name	Name of policy to be associated with network export. Policy can filter and/or change the route parameters.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The exporting of routes between any two routing protocols is a discrete configuration function. For example, you must configure the switch to export routes from OSPF to BGP and, if desired, you must configure the switch to export routes from BGP to OSPF. You must first configure both protocols and then verify the independent operation of each. Then, you can configure the routes to export from OSPF to BGP, and the routes to export from BGP to OSPF. Similarly for BGP and RIP.

You can use a policy to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. A policy can also be used to filter out exported routes.

Using the `export` command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command enables BGP to export routes from the OSPF protocol to BGP peers:

```
enable bgp export ospf
```

enable bgp fast-external-fallover

```
enable bgp fast-external-fallover
```

Description

Enables BGP fast external fallover functionality.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the BGP fast external fallover on the router. This command applies to all directly-connected external BGP neighbors.

When BGP fast external fallover is enabled, the directly-connected EBGP neighbor session is immediately reset when the connecting link goes down.

If BGP fast external fallover is disabled, BGP waits until the default hold timer expires (3 keepalives) to reset the neighboring session. In addition, BGP might teardown the session somewhat earlier than hold timer expiry if BGP detects that the TCP session and its directly connected link is broken (BGP detects this while sending or receiving data from TCP socket).

Example

The following command enables BGP fast external fallover:

```
enable bgp fast-external-fallover
```

enable bgp neighbor

```
enable bgp neighbor [<remoteaddr> | all]
```

Description

Enables the BGP session. The neighbor must be created before the BGP neighbor session can be enabled.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.

Default

Disabled.

Usage Guidelines

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
```

Example

The following command enables the BGP neighbor session:

```
enable bgp neighbor 192.168.1.17
```

enable bgp neighbor capability

```
enable bgp neighbor [all | <remoteaddr>] capability [ipv4-unicast | ipv4-multicast |  
route-refresh]
```

Description

This command enables multi protocol BGP (MBGP) and route-refresh capabilities for one or all BGP neighbors.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.
ipv4-unicast	Specifies BGP MP unicast capabilities.
ipv4-multicast	Specifies BGP MP multicast capabilities.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are enabled by default.

Usage Guidelines

Note: This command is not generally supported, and is not recommended for general use. However, to inter-operate with Cisco routers for BGP graceful restart, you must enable the IPv4 unicast address capability.

After the capabilities have been enabled, the BGP neighbor announces its capabilities to neighbors in an OPEN message.

If IPv4 unicast, IPv4 multicast, or both capabilities are enabled, MBGP extension is enabled and the routes from the specified address families are updated, accepted, and installed. (Thus to enable MBGP, you must configure at least one or both address family capabilities).

When both IPv4 unicast and IPv4 multicast are disabled, MBGP is disabled and BGP peering defaults to its backward-compatible behavior of carrying IPv4 unicast updates. In this configuration, the BGP peers accept and install IPv4 unicast routes.

Example

The following command enables the route-refresh feature for all neighbors:

```
enable bgp neighbor all capability route-refresh
```

enable bgp neighbor originate-default

```
enable bgp [{neighbor} <remoteaddr> | neighbor all] {address-family [ipv4-unicast |  
ipv4-multicast]} originate-default {policy <policy-name>}
```

Description

Enables the origination and advertisement of a default route to a single BGP neighbor or to all BGP neighbors.

Syntax Description

neighbor remoteaddr	Specifies the IP address of a BGP neighbor for which the default route is originated and advertised.
neighbor all	Specifies that default routes are to be originated and advertised for all BGP neighbors.
ipv4-unicast	Specifies that the default routes apply to IPv4 unicast routes. If you do not specify an address family, the command applies to IPv4 unicast routes.
ipv4-multicast	Specifies that the default routes apply to IPv4 multicast routes.
policy-name	Specifies a policy to be applied to the default route origination.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peer is enabled or disabled. The default route or routes are created regardless of whether or not there are matching entries in the IP route table.

When a BGP neighbor is added to a peer group, it does not inherit the default route origination configuration from the peer group. Also, default route origination for a neighbor and the associated peer group can be different.

If a policy is configured and specified in the command, a default route can be originated only if there is a route in the local BGP RIB that matches the policy's match rules. The default route's attribute can be modified using the same policy file by including statements in the *set* block of the policy.

Example

The following command enables the origination and advertisement of default routes for IPv4 unicast traffic for all BGP peer nodes:

```
enable bgp neighbor all originate-default
```

enable bgp neighbor remove-private-AS-numbers

```
enable bgp neighbor [<remoteaddr> | all] remove-private-AS-numbers
```

Description

Enables the removal of private AS numbers from the AS path in route updates sent to EBGp peers.

Syntax Description

remoteaddr	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.

Default

Disabled.

Usage Guidelines

Private AS numbers are AS numbers in the range 64512 through 65534. You can remove private AS numbers from the AS path attribute in updates that are sent to external BGP (EBGP) neighbors. Possible reasons for using private AS numbers include:

- The remote AS does not have officially allocated AS numbers.
- You want to conserve AS numbers if you are multi-homed to the local AS.

Private AS numbers should not be advertised on the Internet. Private AS numbers can only be used locally within an administrative domain. Therefore, when routes are advertised out to the Internet, the routes can be stripped out from the AS paths of the advertised routes using this feature.

Example

The following command enables the removal of private AS numbers from the AS path in route updates sent to the EBGp peers:

```
enable bgp neighbor 192.168.1.17 remove-private-AS-numbers
```

enable bgp neighbor soft-in-reset

```
enable bgp neighbor [all | <remoteaddr>] {address-family [ipv4-unicast | ipv4-multicast]}
soft-in-reset
```

Description

Enables the soft input reset feature.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

This command can be issued only when both BGP and the BGP neighbor is disabled.

Example

The following command enables the soft recognition feature:

```
enable bgp neighbor 192.168.1.17 soft-in-reset
```

enable bgp neighbor use-ip-router-alert

```
enable bgp neighbor [all | <remoteaddr>] use-ip-router-alert
```

Description

Enables the router alert IP option in outgoing BGP messages to the specified neighbor.

Syntax Description

all	Specifies all neighbors.
remoteaddr	Specifies an IP address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

This command forces the IP layer of the NETGEAR 8800 to insert the IP Router Alert Option field in all the outbound BGP messages. IP packets with IP Router Alert option in them examined closely by all the intermediate routers in the transit path, thereby causing transmit delays.

Example

The following command enables the feature:

```
enable bgp neighbor 192.168.1.17 use-ip-router-alert
```

enable bgp peer-group

```
enable bgp peer-group <peer-group-name>
```

Description

Enables a peer group and all the neighbors of a peer group.

Syntax Description

peer-group-name	Specifies a peer group.
-----------------	-------------------------

Default

Disabled.

Usage Guidelines

You can use BGP peer groups to group together up to 200 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command enables the BGP peer group *outer* and all its neighbors:

```
enable bgp peer-group outer
```

enable bgp peer-group capability

```
enable bgp peer-group <peer-group-name> capability [ipv4-unicast | ipv4-multicast |  
route-refresh]
```

Description

This command enables BGP Multiprotocol (MP) and route-refresh capabilities for a peer-group.

Syntax Description

peer-group-name	Specifies a peer group.
ipv4-unicast	Specifies BGP MP unicast capabilities.
ipv4-multicast	Specifies BGP MP multicast capabilities.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are disabled by default.

Usage Guidelines

Note: This command is not generally supported, and is not recommended for general use. However, to inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

This command enables BGP Multiprotocol and route-refresh capabilities for a peer group. After the capabilities have been enabled, the BGP peer announces its capabilities to neighbors in an OPEN message.

Example

The following command enables the route-refresh feature for the peer group *outer*:

```
enable bgp peer-group outer capability route-refresh
```

enable bgp peer-group originate-default

```
enable bgp {peer-group} <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}  
originate-default {policy <policy-name>}
```

Description

Enables the origination and advertisement of default routes to all BGP neighbors in the specified peer group.

Syntax Description

peer-group peer-group-name	Specifies the BGP peer group for which the default routes are originated and advertised.
ipv4-unicast	Specifies that the default routes apply to IPv4 unicast routes. If you do not specify an address family, the command applies to IPv4 unicast routes.
ipv4-multicast	Specifies that the default routes apply to IPv4 multicast routes.
policy-name	Specifies a policy to be applied to the default routes during origination.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peers are enabled or disabled. The default routes are created regardless of whether or not there are matching entries in the IGP route table.

When a BGP neighbor is added to a peer group, it does not inherit the default route origination configuration from the peer group. Also, default route origination for a neighbor and the associated peer group can be different.

If a policy is configured and specified in the command, a default route can be originated only if there is a route in the local BGP RIB that matches the policy's match rules. The default route's attribute can be modified using the same policy file by including statements in the *set* block of the policy.

Example

The following command enables the origination and advertisement of default routes for IPv4 unicast traffic for all nodes in the *test* BGP peer group:

```
enable bgp peer-group test originate-default
```

enable bgp peer-group remove-private-AS-numbers

```
enable bgp peer-group <peer-group-name> remove-private-AS-numbers
```

Description

Enables the removal of private autonomous system (AS) numbers from the AS_Path attribute of outbound updates.

Syntax Description

peer-group-name	Specifies a peer group.
-----------------	-------------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the BGP peer group *outer* from removing private AS numbers:

```
enable bgp peer-group outer remove-private-AS-numbers
```

enable bgp peer-group soft-in-reset

```
enable bgp peer-group <peer-group-name> {address-family [ipv4-unicast | ipv4-multicast]}
soft-in-reset
```

Description

Enables the soft input reset feature.

Syntax Description

peer-group-name	Specifies a peer group.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

Example

The following command enables the soft input reset feature:

```
enable bgp peer-group outer soft-in-reset
```

enable bgp peer-group use-ip-router-alert

```
enable bgp peer-group <peer-group-name> use-ip-router-alert
```

Description

Enables the router alert IP option in outgoing BGP messages to the specified peer group.

Syntax Description

peer-group-name	Specifies a peer group.
-----------------	-------------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the feature for the peer group *outer*:

```
enable bgp peer-group outer use-ip-router-alert
```

show bgp

```
show bgp
```

Description

Displays BGP configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays BGP configuration information:

```
* Switch.4 # show bgp
```

Enabled	: Yes	OperStatus	: Up
RouterId	: 10.203.134.55	AS	: 65535.65535
LocalPref	: 100	MED	: None
Always-Compare-MED	: Disabled	Aggregation	: Disabled
Route Reflector	: No	RR ClusterId	: 0
IGP Synchronization	: Disabled	New Community Format	: Disabled
Routes from EBGp	: 0	Routes from IBGP	: 0
Routes redistributed	: 0	Out Updates queued	: 0
Fast Ext Fallover	: Disabled	MPLS LSP as Next-Hop	: No
AS Disp Format	: Asdot	Maximum ECMP Paths	: 1
ConfedId	: 0		
Confed Peers	:		
Networks	: 0		
Aggregate Networks	: 0		
Redistribute:			
ipv4	Admin	Operational	Shutdown Policy
unicast	Status	Status	Priority

Direct	Disabled	Down	2048	None
Static	Disabled	Down	2048	None
RIP	Disabled	Down	2048	None
BlackHole	Disabled	Down	2048	None
OSPFIntra	Disabled	Down	2048	None
OSPFInter	Disabled	Down	2048	None
OSPFExt1	Disabled	Down	2048	None
OSPFExt2	Disabled	Down	2048	None
ISISL1	Disabled	Down	2048	None
ISISL2	Disabled	Down	2048	None
ISISL1Ext	Disabled	Down	2048	None
ISISL2Ext	Disabled	Down	2048	None

ipv4	Admin	Operational	Shutdown	Policy
multicast	Status	Status	Priority	

Direct	Disabled	Down	2048	None
Static	Disabled	Down	2048	None
RIP	Disabled	Down	2048	None
BlackHole	Disabled	Down	2048	None
OSPFIntra	Disabled	Down	2048	None
OSPFInter	Disabled	Down	2048	None
OSPFExt1	Disabled	Down	2048	None
OSPFExt2	Disabled	Down	2048	None
ISISL1	Disabled	Down	2048	None
ISISL2	Disabled	Down	2048	None
ISISL1Ext	Disabled	Down	2048	None
ISISL2Ext	Disabled	Down	2048	None

Advertise Inactive Routes:

```

ipv4-unicast    : Disabled
ipv4-multicast  : Disabled

```

show bgp memory

```
show bgp memory {detail | <memoryType>}
```

Description

Displays BGP specific memory usage.

Syntax Description

detail	Displays detail information.
memoryType	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays BGP specific memory for all types:

```
show bgp memory detail
```

show bgp neighbor

```
show bgp [neighbor {detail} | {neighbor} <remoteaddr>]
```

Description

Displays information about a specified neighbor.

Syntax Description

remoteaddr	Specifies an IP address that identifies a BGP neighbor.
detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

Use this command to display information about a specific BGP neighbor. If you do not specify a neighbor, information about all neighbors is displayed.

Example

The following command displays information about all neighbors:

```
* Switch.18 # show bgp neighbor
```

Peer	AS	Weight	State	InMsgs	OutMsgs(InQ)	Up/Down

Ee-- 14.0.0.1	4000000000	1	ESTABLISHED	3	5 (0)	0:0:00:04

Flags: (d) disabled, (e) enabled, (E) external peer, (I) internal peer
(m) EBGP multihop, (r) route reflector client

BGP Peer Statistics

Total Peers	: 1		
EBGP Peers	: 1	IBGP Peers	: 0
RR Client	: 0	EBGP Multihop	: 0
Enabled	: 1	Disabled	: 0

The following command shows the detail report for all neighbors:

```
* Switch.21 # show bgp neighbor detail
```

```
EBGP Peer      : 14.0.0.1      AS      : 4000000000
Enabled        : Yes           OperStatus : Up
Weight        : 1             Shutdown-Priority : 1024
ConnectRetry   : 120          MinAsOrig  : 15
HoldTimeCfgr   : 180          KeepaliveCfgr : 60
Source Interface : Not configured RRClient   : No
EBGP-Multihop  : No           Remove Private AS : No
Capabilities Config : ipv4-unicast, ipv4-multicast, 4-Byte-As, route-refresh
Policy for NLRI Type ipv4-unicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled      Send Communities : No
  Soft Input Recfg : Disabled     Allow Looped AS-Path: No
Policy for NLRI Type ipv4-multicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled      Send Communities : No
  Soft Input Recfg : Disabled     Allow Looped AS-Path: No
State           : ESTABLISHED
FSM Up since    : Mon Mar 23 15:15:20 2009 (Duration: 0:0:06:36)
Remote Addr     : 14.0.0.1      Local Addr    : 14.0.0.2
Remote Port     : 179           Local Port    : 42885
Remote RouterId : 10.203.134.8  Local RouterId : 8.203.134.55
HoldTimeNegotiated : 180        KeepAliveNegotiated : 60
FsmTransitions  : 1
InUpdateElapsedTime : 00:00:06:35 InMsgElapsedTime : 0:0:00:35
InUpdates       : 2             OutUpdates (in TxQ) : 3 (0)
InTotalMsgs     : 9             OutTotalMsgs      : 11
InRouteRefreshes : 0            OutRouteRefreshes : 0
Route Statistics for NLRI Type ipv4-unicast
  Received       : 0             Accepted        : 0
  Rejected       : 0             Active           : 0
  Suppressed     : 0
Route Statistics for NLRI Type ipv4-multicast
  Received       : 0             Accepted        : 0
  Rejected       : 0             Active           : 0
  Suppressed     : 0
```

```

Capabilities Tx      : ipv4-unicast, ipv4-multicast, 4-Byte-AS, route-refresh (old &
                        new)
Capabilities Rx      : ipv4-unicast, ipv4-multicast, 4-Byte-AS, route-refresh (old &
                        new)
NLRI for the session: ipv4-unicast, ipv4-multicast
Last State          : ESTABLISHED          Last Event          : RX_KEEP
LastError           : 'None'

```

BGP Peer Statistics

```

Total Peers        : 1
EBGP Peers         : 1                    IBGP Peers          : 0
RR Client          : 0                    EBGP Multihop       : 0
Enabled            : 1                    Disabled            : 0

```

show bgp neighbor flap-statistics

```

show bgp neighbor <remoteaddr> {address-family [ipv4-unicast | ipv4-multicast]}
flap-statistics {detail}
[all
| as-path <path-expression>
| community [no-advertise | no-export | no-export-subconfed
| number <community_num> | <AS_Num>:<Num>
]
| network [any / <netMaskLen> | <networkPrefixFilter>] {exact}
]

```

Description

Displays information about neighbor route flap dampening statistics.

Syntax Description

remoteaddr	Specifies an IP address that identifies a BGP neighbor.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
flap-statistics	Specifies that only flap-statistics should be displayed (for route flap dampening enabled routes).
detail	Specifies to display the information in detailed format.
all	Specifies all routes.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_num	Specifies a community number.
AS_Num	Specifies an autonomous system ID (0-65535).

Num	Specifies the BGP community number.
any	Specifies all routes with a given or larger mask length.
netMaskLen	Specifies a subnet mask length (number of bits).
networkPrefixFilter	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Use this command to display information about BGP neighbor route flap dampening.

The option `network any / <netMaskLen>` displays all BGP routes whose mask length is equal to or greater than `<maskLength>`, irrespective of their network address.

The option `network any / <netMaskLen> exact` displays all BGP routes whose mask length is exactly equal to `<maskLength>`, irrespective of their network address.

Example

The following command displays information about a specified neighbor's dampened routes:

```
show bgp neighbor 10.10.10.10 flap-statistics all
```

show bgp neighbor routes

```
show bgp neighbor <remoteaddr> {address-family [ipv4-unicast | ipv4-multicast]}
[accepted-routes | received-routes | rejected-routes | suppressed-routes |
transmitted-routes] {detail} [all | as-path <path-expression> | community [no-advertise |
no-export | no-export-subconfed | number <community_num> | <AS_Num>:<Num>] | network [any /
<netMaskLen> | <networkPrefixFilter>] {exact}]
```

Description

Displays information about specified neighbor routes.

Syntax Description

remoteaddr	Specifies an IP address that identifies a BGP neighbor.
address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast
accepted-routes	Specifies that only accepted routes should be displayed.
received-routes	Specifies that only received routes should be displayed.
rejected-routes	Specifies that only rejected routes should be displayed.

suppressed-routes	Specifies that only suppressed routes should be displayed (for route flap dampening enabled routes).
transmitted-routes	Specifies that only transmitted routes should be displayed.
detail	Specifies to display the information in detailed format.
all	Specifies all routes.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_num	Specifies a community number.
AS_Num	Specifies an autonomous system ID (0-65535).
Num	Specifies the BGP community number.
any	Specifies all routes with a given or larger mask length.
netMaskLen	Specifies a subnet mask length (number of bits).
networkPrefixFilter	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Use this command to display information about a specific BGP neighbor routes.

The option `network any / <netMaskLen>` displays all BGP routes whose mask length is equal to or greater than `<maskLength>`, irrespective of their network address.

The option `network any / <netMaskLen> exact` displays all BGP routes whose mask length is exactly equal to `<maskLength>`, irrespective of their network address.

Example

The following command displays information about a specified neighbor's received routes:

```
show bgp neighbor 10.10.10.10 received-routes all
```

show bgp peer-group

```
show bgp peer-group {detail | <peer-group-name> {detail}}
```

Description

Displays the peer groups configured in the system.

Syntax Description

detail	Specifies to display the information in detailed format.
peer-group-name	Specifies a peer group.

Default

N/A.

Usage Guidelines

If the `detail` keyword is specified then the parameters of the neighbors in the peer group, which are different from the ones that are configured in the peer group, are displayed.

If no peer group name is specified, all the peer group information is displayed.

Example

The following command displays information for the *outer* peer group:

```
XCM8810 # show bgp peer-group "outer"
```

```
Peer Group           : outer
Enabled              : No           AS              : 300000000
Router Enabled       : Yes          Weight         : 1
ConnectRetry        : 120          MinAsOrig      : 15
HoldTimeCfg         : 180          KeepaliveCfg   : 60
Source Interface     : Not configured RRClient       : No
Remove Private AS   : No           Router-Alert    : Disabled
Capabilities Config  : ipv4-unicast ipv4-multicast route-refresh 4-Byte-AS
Policy for NLRI Type ipv4-unicast
  In Policy          : None
  Out Policy         : None
  NextHopSelf        : Disabled      Send Communities : No
  Soft Input Recfg   : Disabled      Allow Looped AS-Path: No
Policy for NLRI Type ipv4-multicast
  In Policy          : None
  Out Policy         : None
  NextHopSelf        : Disabled      Send Communities : No
  Soft Input Recfg   : Disabled      Allow Looped AS-Path: No
Peers                :
```

```
BGP Peer Group Statistics
```

```
Total Peer Groups : 1
Enabled           : 0
Disabled          : 1
```

show bgp routes

```
show bgp routes {address-family [ipv4-unicast | ipv4-multicast]} {detail} [all | as-path
<path-expression> | community [no-advertise | no-export | no-export-subconfed | number
<community_num> | <AS_Num>:<Num>] | network [any / <netMaskLen> | <networkPrefixFilter>]
{exact}]
```

Description

Displays the BGP route information base (RIB).

Syntax Description

address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
all	Specifies all routes.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_num	Specifies a community number.
AS_Num	Specifies an autonomous system ID (0-65535).
Num	Specifies the BGP community number.
any	Specifies all routes with a given or larger mask length.
netMaskLen	Specifies a subnet mask length (number of bits).
networkPrefixFilter	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The option `network any / <netMaskLen>` displays all BGP routes whose mask length is equal to or greater than `<maskLength>`, irrespective of their network address.

The option `network any / <netMaskLen> exact` displays all BGP routes whose mask length is exactly equal to `<maskLength>`, irrespective of their network address.

Example

The following command displays detailed information about all BGP routes:

```
* Switch.5 # show bgp routes all
```


Feasible Routes

Destination	Peer	Next-Hop	LPref	Weight	MED	AS-Path
*>i 10.0.0.0/8	16.0.0.1	16.0.0.1	100	1		1.20 100
*>i 11.0.0.0/8	16.0.0.1	16.0.0.1	100	1		1.20 100
*>i 12.0.0.0/8	16.0.0.1	16.0.0.1	100	1		1.20 100

Unfeasible Routes

Destination	Peer	Next-Hop	LPref	Weight	MED	AS-Path
-------------	------	----------	-------	--------	-----	---------

Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
 (s) Stale, (m) Multipath

Origin: (?) Incomplete, (e) EGP, (i) IGP

BGP Route Statistics

Total Rxed Routes : 3
 Feasible Routes : 3
 Unfeasible Routes : 0
 Active Routes : 3

Route Statistics on Session Type

Routes from Int Peer: 0
 Routes from Ext Peer: 3

The following command displays detailed information about all BGP routes:

* Switch.21 # show bgp routes detail all

Feasible Routes

Route: 20.20.20.0/24, Peer 16.0.0.1, BEST
 Origin IGP, Next-Hop 16.0.0.1, LPref 100
 Weight 1,
 AS-Path: 3 33 333 3333 33333 333333 3333333 33333333 333333333 3333333333 3333333333 1000000000

Unfeasible Routes

BGP Route Statistics

Total Rxed Routes : 1
 Feasible Routes : 1
 Unfeasible Routes : 0

```

Active Routes      : 0
Route Statistics on Session Type
Routes from Int Peer: 0
Routes from Ext Peer: 1

```

show bgp routes summary

```
show bgp routes {address-family [ipv4-unicast | ipv4-multicast]} summary
```

Description

Displays a summary the BGP route information base (RIB).

Syntax Description

address-family	The address family. BGP supports two address families: IPv4 Unicast and IPv4 Multicast.
----------------	---

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

None.

Example

The following command displays a summary of the BGP route information base (RIB) for IPv4 multicast:

```
show bgp routes address-family ipv4-multicast summary
```

This chapter describes commands for configuring and managing the following IPv4 multicast features:

- Multicast routing
- Protocol Independent Multicast (PIM)
- Internet Group Management Protocol (IGMP)
- Multicast VLAN Registration (MVR)

For an introduction to these features, see the *NETGEAR 8800 User Manual*.

clear igmp group

```
clear igmp group {<grpipaddress>} [{vlan} <name>}
```

Description

Removes one or all IGMP groups.

Syntax Description

name	Specifies a VLAN name.
grpipaddress	Specifies the group IP address.

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove learned IGMP group entries instantly. Traffic is impacted until the IGMP groups are relearned. Use this command for diagnostic purposes only.

Example

The following command clears all IGMP groups from VLAN *accounting*:

```
clear igmp group accounting
```

clear igmp snooping

```
clear igmp snooping [{vlan} <name>]
```

Description

Removes one or all IGMP snooping entries.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove IGMP snooping entries instantly. However, removing an IGMP snooping entry can disrupt the normal forwarding of multicast traffic, until the snooping entries are learned again.

The dynamic IGMP snooping entries are removed, then recreated upon the next general query. The static router entry and static group entries are removed and recreated immediately.

This command clears both the IGMPv2 and IGMPv3 snooping entries.

Example

The following command clears IGMP snooping from VLAN *accounting*:

```
clear igmp snooping accounting
```

clear pim cache

```
clear pim cache {<group_addr> {<source_addr>}}
```

Description

Resets the IP multicast cache table.

Syntax Description

group_addr	Specifies a group address.
source_addr	Specifies a source IP address.

Default

If no options are specified, all IP multicast cache entries are flushed.

Usage Guidelines

This command can be used by network operators to manually remove IPMC software and hardware forwarding cache entries instantly. If the stream is available, caches are re-created, otherwise caches are removed permanently. This command can disrupt the normal forwarding of multicast traffic.

Example

The following command resets the IP multicast table for group 224.1.2.3:

```
clear pim cache 224.1.2.3
```

clear pim snooping

```
clear pim snooping {vlan} <name>}
```

Description

Clears all PIM snooping neighbors, joins received on the VLAN, and the VLAN forwarding entries.

Syntax Description

name	Specifies the VLAN to which this command applies.
------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the PIM snooping database for the Default VLAN:

```
clear pim snooping "Default"
```

configure forwarding ipmc compression

```
configure forwarding ipmc compression {group-table | off}
```

Description

Enables or disables compression of entries in the IP multicast group table to facilitate improved IP multicast scaling.

Syntax Description

group-table	Enables compression.
off	Disables compression.

Default

group-table

Usage Guidelines

Compression of IP multicast group table entries allows the switch to process more multicast traffic using the faster switch hardware instead of the relatively slower switch software. Compression requires additional processing. Disable this feature if you suspect a problem exposed by IP multicast compression.

When you enable or disable this feature, all IP multicast entries are flushed, and this can result in a temporary loss of multicast traffic while the IP multicast entries are relearned.

Note: On NETGEAR 8800 series switches, all IP multicast forwarding entries utilizing the same IP multicast group table entry share a single backplane link, limiting the total throughput to 12Gbps.

To display the compression feature configuration, enter the command:

```
show forwarding configuration
```

Example

The following command disables compression:

```
configure forwarding ipmc compression off
```

configure igmp

```
configure igmp <query_interval> <query_response_interval> <last_member_query_interval>
{<robustness>}
```

Description

Configures the Internet Group Management Protocol (IGMP) timers.

Syntax Description

query_interval	Specifies the interval (in seconds) between general queries.
query_response_interval	Specifies the maximum query response time (in seconds).
last_member_query_interval	Specifies the maximum group-specific query response time (in seconds).
robustness	Specifies the degree of robustness for the network.

Default

- query interval—125 seconds
- query response interval—10 seconds
- last member query interval—1 second
- robustness—2

Usage Guidelines

Timers are based on RFC2236. Specify the following:

- query interval—The amount of time, in seconds, the system waits between sending out general queries. The range is 1 to 429,496,729 seconds.
- query response interval—The maximum response time inserted into the periodic general queries. The range is 1 to 25 seconds.
- last member query interval—The maximum response time inserted into a group-specific query sent in response to a leave group message. The range is 1 to 25 seconds.
- robustness—The degree of robustness of the network. The range is 2 to 7. This parameter allows tuning for the expected packet loss on a link. If a link is expected to have packet loss, this parameter can be increased.
- The group timeout is defined by the formula: $\text{group_timeout} = (\text{query_interval} \times \text{robustness}) + \text{query_response_interval}$, according to RFC 2236. You can explicitly define the host timeout using the `configure igmp snooping timer <router_timeout> <host_timeout> {vr <vrname>}` command. The effective host_timeout is the lesser value of the group_timeout and the configured host_timeout.

Example

The following command configures the IGMP timers:

```
configure igmp 100 5 1 3
```

configure igmp snooping filters

```
configure igmp snooping filters [per-port | per-vlan]
```

Description

Selects the type of IGMP snooping filters that are installed.

Syntax Description

per-port	Installs the per-port IGMP snooping filters.
per-vlan	Installs the per-VLAN IGMP snooping filters.

Default

per-port

Usage Guidelines

Use the `per-vlan` option when the number of VLANs configured on the switch is lower than the maximum number (2000). This option conserves usage of the hardware Layer 3 multicast forwarding table.

When the number of configured VLANs is larger than 2000, select the `per-port` option. Each VLAN requires additional interface hardware ACL resources. The per-port option conserves usage of the interface hardware ACL resources.

The maximum number of VLANs supported by per VLAN IGMP snooping filters is 2000.

To display the IGMP snooping filters configuration, use the `show igmp` command.

Example

The following command configures the switch to install the per-VLAN IGMP snooping filters:

```
configure igmp snooping filters per-vlan
```

configure igmp snooping flood-list

```
configure igmp snooping flood-list [<policy> | none] {vr <vrname>}
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

policy	Specifies a policy file with a list of multicast addresses to be handled.
none	Specifies no policy file is to be used.
vrname	Specifies a virtual router.

Default

None.

Usage Guidelines

With this command, a user can configure certain multicast addresses to be slow path flooded within the VLAN, which otherwise are fast path forwarded according to IGMP and/or Layer 3 multicast protocol.

A policy file is a text file with the extension, .pol. It can be created or edited with any text editor. The specified policy file `<policy file>` should contain a list of addresses which determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with a destination address which is in the `<policy file>` in 'permit' mode, that stream is software flooded and no hardware entry is installed.

When adding an IP address into the policy file, a 32-bit host address is recommended.

This feature is meant to solve the multicast connectivity problem for unknown destination addresses within system reserved ranges. Specifically this feature was introduced to solve the problem of recognizing certain streams as control packets.

To create a policy file for the snooping flood-list, use the following template:

```
# This is a template for IGMP Snooping Flood-list Policy File
# Add your group addresses between "Start" and "End"
# Do not touch the rest of the file!!!!
entry igmpFlood {
    if match any {
#----- Start of group addresses -----
        nlri 234.1.1.1/32;
        nlri 239.1.1.1/32;
#----- end of group addresses -----
    } then {
        permit;
    }
}

entry catch_all {
    if {

    } then {
        deny;
    }
}
```

Note: The switch does not validate any IP address in the policy file used in this command. Therefore, slow-path flooding should be used only for streams which are very infrequent, such as control packets. It should not be used for multicast data packets. This option overrides any default mechanism of hardware forwarding (with respect to IGMP, PIM, or DVMRP) so it should be used with caution.

Slow path flooding is done within the L2 VLAN only.

Use the `none` option to effectively disable slow path flooding.

You can use the `show igmp` command to see the configuration of slow path flooding.

Example

The following command configures the multicast data stream specified in `access1` for slow path flooding:

```
configure igmp snooping flood-list access1
```

The following command specifies that no policy file is to be used, this effectively disabling slow path flooding:

```
configure igmp snooping flood-list none
```

configure igmp snooping forwarding-mode

```
configure igmp snooping forwarding-mode [group-vlan | source-group-vlan]
```

Description

Configures the format for stored multicast entries in hardware as either the default mode (S, G) or the optional mode (*, G).

Syntax Description

group-vlan	Stores multicast entries in the format (* <AnySourceIP>, GroupIP, VlanId), which is also referred to as (*, G).
source-group-vlan	Stores multicast entries in the format (SourceIP, GroupIP, VlanId), which is also referred to as (S, G).

Default

source-group-vlan (S, G).

Usage Guidelines

In networks where there are many sources for each multicast address, the default (S, G) format can consume storage space. To use less storage space for multicast entries, specify the (*, G) format with the `group-vlan` format.

Note: Once the entries are programmed as (*, G), any multicast traffic for the group is forwarded based on the (*, G) entries and does not come to the CPU. Use the `group-vlan` option only with IGMPv2 networks.

Example

The following command configures the group-vlan format:

```
configure igmp snooping forwarding-mode group-vlan
```

configure igmp snooping leave-timeout

```
configure igmp snooping leave-timeout <leave_timeout_ms> {vr <vrname>}
```

Description

Configures the IGMP snooping leave timeout.

Syntax Description

leave_timeout_ms	Specifies an IGMP leave timeout value in milliseconds.
vrname	Specifies a virtual router.

Default

1000 ms.

Usage Guidelines

The leave-timeout is the IGMP leave override interval. If no other hosts override the IGMP leave by the end of this interval, the receiver port is removed.

The range is 0 - 10000 ms (10 seconds). For timeout values of one second or less, you must set the leave-timeout to a multiple of 100 ms. For values of more than one second, you must set the leave-timeout to a multiple of 1000 ms (one second).

Example

The following command configures the IGMP snooping leave timeout:

```
configure igmp snooping leave-timeout 10000
```

configure igmp snooping timer

```
configure igmp snooping timer <router_timeout> <host_timeout> {vr <vrname>}
```

Description

Configures the IGMP snooping timers.

Syntax Description

router_timeout	Specifies the time in seconds before removing a router snooping entry.
host_timeout	Specifies the time in seconds before removing a host's group snooping entry.

vrname	Specifies a virtual router.
--------	-----------------------------

Default

The router timeout default setting is 260 seconds. The host timeout setting is 260 seconds.

Usage Guidelines

Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following:

- router timeout—The maximum time, in seconds, that a router snooping entry can remain in the IGMP snooping table without receiving a router report. If a report is not received, the entry is deleted. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.
- host timeout—The maximum time, in seconds, that a group snooping entry can remain in the IGMP snooping table without receiving a group report. If a report is not received, the entry is deleted. The range is 10 to 214,748,364 seconds. The default setting is 260 seconds.

Note: The `host_timeout` value should be less than or equal to the query timeout value, which is defined by the following `configure igmp` command timers as follows: `query_interval X robustness) + query_response_interval`.

IGMP snooping expects at least one device on every VLAN to periodically generate IGMP query messages. Without an IGMP querier, the switch eventually stops forwarding IP multicast packets to any port, because the IGMP snooping entries time out, based on the value specified in `host_timeout` or `router_timeout`.

Example

The following command configures the IGMP snooping timers:

```
configure igmp snooping timer 600 600
```

configure igmp snooping vlan ports add static group

```
configure igmp snooping {vlan} <vlannname> ports <portlist> add static group <ip address>
```

Description

Configures VLAN ports to receive the traffic from a multicast group, even if no IGMP joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports.
ip address	Specifies the multicast group IP address.

Default

None.

Usage Guidelines

Use this command to forward a particular multicast group to VLAN ports. In effect, this command emulates a host on the port that has joined the multicast group. As long as the port is configured with the static entry, multicast traffic for that multicast group is forwarded to that port.

This command is for IGMPv2 only.

The switch sends proxy IGMP messages in place of those generated by a real host. The proxy messages use the VLAN IP address for source address of the messages. If the VLAN has no IP address assigned, the proxy IGMP message uses 0.0.0.0 as the source IP address.

The multicast group should be in the class-D multicast address space, but should not be in the multicast control subnet range (224.0.0.x/24).

If the ports also have an IGMP filter configured, the filter entries take precedence. IGMP filters are configured using the command:

```
configure igmp snooping vlan <vlan name> ports <portlist> filter <policy file>
```

Example

The following command configures a static IGMP entry so the multicast group 224.34.15.37 is forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 add static group 224.34.15.37
```

configure igmp snooping vlan ports add static router

```
configure igmp snooping {vlan} <vlanname> ports <portlist> add static router
```

Description

Configures VLAN ports to forward the traffic from all multicast groups, even if no IGMP joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports.

Default

None.

Usage Guidelines

Use this command to forward all multicast groups to the specified VLAN ports. In effect, this command emulates a multicast router attached to those ports. As long as the ports are configured with the static entry, all available multicast traffic is forwarded to those ports.

Example

The following command configures a static IGMP entry so all multicast groups are forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 add static router
```

configure igmp snooping vlan ports delete static group

```
configure igmp snooping {vlan} <vlanname> ports <portlist> delete static group [<ip_address>
| all]
```

Description

Removes the port configuration that causes multicast group traffic to be forwarded, even if no IGMP leaves have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports.
ip address	Specifies the multicast group IP address.
all	Delete all the static groups.

Default

None.

Usage Guidelines

Use this command to remove a static group entry created by the following command:

```
configure igmp snooping vlan <vlan name> ports <portlist> add static group <ip address>
```

Example

The following command removes a static IGMP entry that forwards the multicast group 224.34.15.37 to the VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 delete static group 224.34.15.37
```

configure igmp snooping vlan ports delete static router

```
configure igmp snooping vlan <vlanname> ports <portlist> delete static router
```

Description

Removes the configuration that causes VLAN ports to forward the traffic from all multicast groups, even if no IGMP joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to remove an entry created by the following command:

```
configure igmp snooping vlan <vlanname> ports <portlist> add static router
```

Example

The following command removes the static IGMP entry that caused all multicast groups to be forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 delete static router
```

configure igmp snooping vlan ports filter

```
configure igmp snooping vlan <vlanname> ports <portlist> filter [<policy> | none]
```

Description

Configures an IGMP snooping policy file filter on VLAN ports.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
policy	Specifies the policy file for the filter.

Default

None.

Usage Guidelines

Use this command to filter multicast groups to the specified VLAN ports.

The policy file used by this command is a text file that contains the class-D addresses of the multicast groups that you wish to block.

To remove IGMP snooping filtering from a port, use the `none` keyword version of the command.

Use the following template to create a snooping filter policy file:

```
#
# Add your group addresses between "Start" and "end"
# Do not touch the rest of the file!!!!

entry igmpFilter {
    if match any {
#----- Start of group addresses -----
        nlri 239.11.0.0/16;
        nlri 239.10.10.4/32;
#----- end of group addresses -----
    } then {
        deny;
    }
}

entry catch_all {
    if {

    } then {
        permit;
    }
}
```


Example

The following command configures the policy file *ap_multicast* to filter multicast packets forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 filter ap_multicast
```

configure igmp snooping vlan ports set join-limit

```
configure igmp snooping {vlan} <vlaname> ports <portlist> set join-limit {<num>}
```

Description

Configures VLAN ports to support a maximum number of IGMP joins.

Syntax Description

vlaname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports.
num	Specifies the maximum number of joins permitted on the ports. The range is 1 to 500.

Default

No limit.

Usage Guidelines

None.

Example

The following command configures port 2:1 in the Default VLAN to support a maximum of 100 IGMP joins:

```
configure igmp snooping "Default" ports 2:1 set join-limit 100
```

configure igmp ssm-map add

```
configure igmp ssm-map add <group_ip> [/<prefix> | <mask>] <source_ip> {vr <vr-name>}
```

Description

Configures an IGMP SSM mapping.

Syntax Description

group_ip	Specifies the multicast IP address for the group mapping.
----------	---

prefix	Specifies a prefix length for the multicast group IP address. The range is 4 to 32.
mask	Specifies the network mask for the group multicast IP address.
source_ip	The IP address for a multicast group source.
vr-name	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

IGMP SSM mapping operates only with IPv4.

Example

The following command configures an IGMP-SSM mapping for the range of multicast IP addresses at 232.1.1.0/24 to IP host 172.16.8.1:

```
configure igmp ssm-map add 232.1.1.0/24 172.16.8.1
```

configure igmp ssm-map delete

```
configure igmp ssm-map delete <group_ip> [/<prefix>} | <mask>] [<source_ip> | all] <vr  
<vr-name>}
```

Description

Unconfigures an SSM mapping.

Syntax Description

group_ip	Specifies the multicast IP address for the group mapping.
prefix	Specifies a prefix length for the multicast group IP address. The range is 4 to 32.
mask	Specifies the network mask for the group multicast IP address.
source_ip	The IP address for a multicast group source.
all	Specifies that all sources for the specified group or mask are deleted.
vr-name	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes an IGMP-SSM mapping for the range of multicast IP addresses at 232.1.1.0/24 to IP host 172.16.8.1:

```
configure igmp ssm-map delete 232.1.1.0/24 172.16.8.1
```

configure ipmcforwarding

```
configure ipmcforwarding to-cpu [auto | off] ports <port_list>
```

Description

Configure whether IP multicast CPU filters are installed automatically

Syntax Description

auto	The software will automatically program IP multicast processing based on configuration.
off	IP multicast packets received on this port are always flooded with no CPU processing.
port_list	Specifies on or more ports.

Default

N/A.

Usage Guidelines

IP forwarding and IPMC forwarding must be enabled for the configuration to operate.

Example

The following example configures automatic operation for port 2.1:

```
configure ipmcforwarding to-cpu auto ports 2.1
```

configure ipmroute add

```
configure ipmroute add [<source-net>/<mask-len> | <source-net> <mask>] [{<protocol>}<br><protocol>] <rpf-address> [<metric>] [vr <vr-name>]
```

Description

Adds a static multicast route to the multicast routing table.

Syntax Description

source-net	Specifies an IP address/mask length.
mask-len	Mask length for the IP multicast source's subnet. Range is [1-32].
mask	Specifies a subnet mask.
protocol	Unicast routing protocol that is to be used for route learning.
rpf-address	Next hop through which the multicast source can be reached.
metric	Specifies a cost metric.
vr-name	Specifies the virtual router to which the route is added.

Default

The following defaults apply:

- metric—1
- vr-name—VR of the current CLI context
- protocol—none

Usage Guidelines

This command allows you to statically configure where multicast sources are located (even though the unicast routing table has different entries). It allows you to configure a multicast static route in such a way as to have non-congruent topology for Unicast and Multicast topology and traffic.

Example

The following command configures a multicast static route for all multicast sources within network subnet 192.168.0.0/16. Those sources are reachable through the gateway 192.75.0.91.

```
configure ipmroute add 192.168.0.0/16 192.75.0.91
```

The following example configures multicast static route for all sources via a single gateway with a metric of 100:

```
configure ipmroute add 0.0.0.0/0 192.75.0.91 100
```

configure ipmroute delete

```
configure ipmroute delete [<source-net>/<mask-len> | <source-net> <mask>] {{protocol}
<protocol>} <rpf-address> {vr <vr-name>}
```

Description

Deletes a static multicast address from the multicast routing table.

Syntax Description

source-net	Specifies an IP address/mask length.
mask-len	Mask length for the IP multicast source's subnet. Range is [1-32].
mask	Specifies a subnet mask.
protocol	Unicast routing protocol that is to be used for route learning.
rpf-address	Next hop through which the multicast source can be reached.
vr-name	Specifies the virtual router to which the route is added.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

This command allows you to delete an existing multicast static route. It allows you to configure congruent topology for unicast and multicast packets and traffic.

Example

The following command deletes a multicast static route:

```
configure ipmroute delete 192.168.0.0/16 192.75.0.91
```

configure iproute add (Multicast)

```
configure iproute add [<ipNetmask> | <ip_addr> <mask>] <gateway> {metric} {multicast |
multicast-only | unicast | unicast-only} {vr <vrname>}
```

Description

Adds a static route to the routing table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ip_addr	Specifies an IP address.
mask	Specifies a subnet mask.
gateway	Specifies a VLAN gateway.
metric	Specifies a cost metric.
vrname	Specifies the virtual router to which the route is added.
multicast	Adds the specified route to the multicast routing table.

multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a mask value of 255.255.255.255 to indicate a host entry.

Note: Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the `multicast` option.

Example

The following command adds a static address to the multicast routing table:

```
configure iproute add 10.1.1.0/24 123.45.67.1 5 multicast
```

configure iproute delete

```
configure iproute delete [<ipNetmask> | <ipaddress> <mask>] <gateway> {multicast |
multicast-only | unicast | unicast-only} {vr <vrname>}
```

Description

Deletes a static address from the routing table.

Syntax Description

ipNetmask	Specifies an IP address/mask length.
ipaddress	Specifies an IP address.
mask	Specifies a subnet mask.
gateway	Specifies a VLAN gateway.
multicast	Specifies a multicast route to delete.
multicast-only	Specifies a multicast route to delete.

unicast	Specifies a unicast route to delete.
unicast-only	Specifies a unicast route to delete.
vrname	Specifies the virtual router to which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a value of 255.255.255.255 or /32 for mask to indicate a host entry.

Example

The following command deletes an address from the multicast routing table:

```
configure iproute delete 10.101.0.0/24 10.101.0.1 multicast
```

configure mvr add receiver

```
configure mvr vlan <vlan-name> add receiver port <port-list>
```

Description

Configures a port to receive MVR multicast streams.

Syntax Description

vlan-name	Specifies a VLAN name.
port-list	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

This command is used to add a group of virtual ports for multicast forwarding through MVR. By default, some ports on non-MVR VLANs (router ports) are excluded from the MVR cache egress list. This command is used to override these rules, so that if valid IGMP memberships are received, or a router is detected, streams are forwarded out on the ports.

Example

The following command adds the ports 1:1 and 1:2 of VLAN v1 to MVR for forwarding:

```
configure mvr vlan v1 add receiver port 1:1-1:2
```

configure mvr add vlan

```
configure mvr add vlan <vlan-name>
```

Description

Configures a VLAN as an MVR VLAN.

Syntax Description

vlan-name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

Configures MVR on the specified VLAN. When a multicast stream in the specified MVR address range is received on the VLAN, it is leaked to all other VLAN ports where the corresponding IGMP join message is received. By default, the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24 is used for MVR. To change the MVR address range, use the following command:

```
configure mvr vlan <vlan-name> mvr-address {<policy-name> | none}
```

Example

The following command configures VLAN v1 as an MVR VLAN:

```
configure mvr add vlan v1
```

configure mvr delete receiver

```
configure mvr vlan <vlan-name> delete receiver port <port-list>
```

Description

Configures a port not to receive MVR multicast streams.

Syntax Description

vlan-name	Specifies a VLAN name.
port-list	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

This command is used to delete a group of virtual ports for multicast forwarding through MVR. After using this command, the ports revert to the default forwarding rules.

Example

The following command deletes the ports 1:1 and 1:2 of VLAN *v1* to MVR for forwarding:

```
configure mvr vlan v1 delete receiver port 1:1-1:2
```

configure mvr delete vlan

```
configure mvr delete vlan <vlan-name>
```

Description

Deletes a VLAN from MVR.

Syntax Description

vlan-name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

Removes MVR from the specified VLAN.

Example

The following command configures VLAN *v1* as a non-MVR VLAN:

```
configure mvr delete vlan v1
```

configure mvr mvr-address

```
configure mvr vlan <vlan-name> mvr-address {<policy-name> | none}
```

Description

Configures the MVR address range on a VLAN.

Syntax Description

vlan-name	Specifies a VLAN name.
policy-name	Specifies a policy file.

Default

The default address range is 224.0.0.0/4 (all multicast addresses), but excluding 224.0.0.0/24 (the multicast control range).

Usage Guidelines

If no policy file is specified (the `none` option), the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24 is used for MVR.

MVR must first be configured on the VLAN before using this command.

If the policy is later refreshed, groups denied and newly allowed groups in the policy are flushed from fast path forwarding. This allows synching existing channels with the new policy, without disturbing existing channels.

The following is a sample policy file `mvrpol.pol`. This policy configures 236.1.1.0/24 as the MVR address range. Any address outside this range has the standard switching behavior on an MVR VLAN.

```
Entry netgear1 {
    if match any {
        nlri 236.1.1.0/24 ;
    }
    then {
        permit ;
    }
}
```

Example

The following command configures the MVR address range specified in the policy file `mvrpol.pol` for the VLAN `v1`:

```
configure mvr vlan v1 mvr-address mvrpol
```

configure mvr static group

```
configure mvr vlan <vlan-name> static group {<policy-name> | none}
```

Description

Configures the MVR static group address range on a VLAN.

Syntax Description

<code>vlan-name</code>	Specifies a VLAN name.
<code>policy-name</code>	Specifies a policy file.

Default

By default, all the MVR group addresses work in static mode.

Usage Guidelines

If no policy file is specified (the `none` option), the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24, is used for static groups in MVR.

MVR must first be configured on the VLAN before using this command.

The following is a sample policy file `mvrpol.pol`. This policy configures 236.1.1.0/24 as the MVR static group address range. Any MVR addresses outside this range are dynamically registered through IGMP. An MVR VLAN will proxy join only for addresses that are not in the static group. If you want all the multicast groups to be dynamic, use a policy file with this command that denies all multicast addresses.

```
Entry netgear1 {
    if match any {
        nlri 236.1.1.0/24 ;
    }
    then {
        permit ;
    }
}
```

Example

The following command configures the MVR static group address range specified in the policy file `mvrpol.pol` for the VLAN `v1`:

```
configure mvr vlan v1 static group mvrpol
```

configure pim add vlan

```
configure pim add vlan [<vlan-name> | all] {dense | sparse} {passive}
```

Description

Configures an IP interface for PIM.

Syntax Description

<code>vlan-name</code>	Specifies a VLAN name.
<code>all</code>	Specifies all VLANs.
<code>dense</code>	Specifies PIM dense mode (PIM-DM).
<code>sparse</code>	Specifies PIM sparse mode (PIM-SM).
<code>passive</code>	Specifies a passive interface.

Default

Dense.

Usage Guidelines

When an IP interface is created, per-interface PIM configuration is disabled by default.

The switch supports both dense mode and sparse mode operation. You can configure dense mode or sparse mode on a per-interface basis. After they are enabled, some interfaces can run dense mode, while others run sparse mode.

Passive interfaces are host only interfaces that allow a multicast stream from other VLANs to be forwarded to edge hosts. Since they do not peer with other PIM routers, do not connect a multicast router to a passive interface.

In order for the interface to participate in PIM, PIM must be enabled on the switch using the following command:

```
enable pim
```

Example

The following command enables PIM-DM multicast routing on VLAN *accounting*:

```
configure pim add vlan accounting dense
```

configure pim cbsr

```
configure pim cbsr [{vlan} <vlan_name> {<priority [0-254]} | none]
```

Description

Configures a candidate bootstrap router for PIM sparse-mode operation.

Syntax Description

vlan_name	Specifies a VLAN name.
priority	Specifies a priority setting. The range is 0 - 254.
none	Specifies to delete a CBSR.

Default

The default setting for priority is 0, and indicates the lowest priority.

Usage Guidelines

The VLAN specified for CBSR must have PIM enabled for it to take effect. After PIM is enabled, CBSRs advertise themselves in the PIM domain. A bootstrap router (BSR) is

elected among all the candidates based on CBSR priority. To break the tie among routers with the same priority setting, the router with the numerically higher IP address is chosen.

A NETGEAR 8800 switch can support up to 145 RPs per group when it is configured as a PIM BSR (bootstrap router). If more than 145 RPs are configured for a single group, the BSR ignores the group and does not advertise the RPs. Non-BSR switches can process more than 145 RPs in the BSR message.

Example

The following command configures a candidate bootstrap router on the VLAN *accounting*:

```
configure pim cbsr vlan accounting 30
```

configure pim crp static

```
configure pim crp static <ip_address> [none | <policy>] {<priority> [0-254]}
```

Description

Configures a rendezvous point and its associated groups statically, for PIM sparse mode operation.

Syntax Description

<code>ip_address</code>	Specifies a static CRP address.
<code>none</code>	Deletes the static rendezvous point.
<code>policy</code>	Specifies a policy file name.
<code>priority</code>	Specifies a priority setting. The range is 0 - 254.

Default

The default setting for priority is 0, which indicates highest priority.

Usage Guidelines

In PIM-SM, the router sends a join message to the rendezvous point (RP). The RP is a central multicast router that is responsible for receiving and distributing multicast packets. If you use a static RP, all switches in your network must be configured with the same RP address for the same group (range).

NETGEAR 8800 switches support up to 50 RPs in a switch, and up to 180 groups within a single RP. If you configure more than 180 groups for a single RP, the switch does not process the groups added after the first 180.

The policy file contains a list of multicast group addresses served by this RP.

This policy file is not used for filtering purposes. As used with this command, the policy file is just a container for a list of addresses. So a typical policy file used for RP configuration looks a little different from a policy used for other purposes.

If routers have different group-to-RP mappings, due to misconfiguration of the static RP (or any other reason), traffic is disrupted.

Example

The following command statically configures an RP and its associated groups defined in policy file *rp-list*:

```
configure pim crp static 10.0.3.1 rp-list
```

The following is a sample policy file:

```
entry netgear1 {
    if match any { }
    then { nlri 224.0.0.0/4 ;
           nlri 239.255.0.0/24 ;
           nlri 232.0.0.0/8 ;
           nlri 238.1.0.0/16 ;
           nlri 232.232.0.0/20 ;
    }
}
```

configure pim crp timer

```
configure pim crp timer <crp_adv_interval>
```

Description

Configures the candidate rendezvous point advertising interval in PIM sparse mode operation.

Syntax Description

<code>crp_adv_interval</code>	Specifies a candidate rendezvous point advertising interval in seconds. The range is 1 to 1,717,986,918.
-------------------------------	--

Default

The default is 60 seconds.

Usage Guidelines

Increasing this time results in increased convergence time for CRP information to the PIM routers.

Example

The following command configures the candidate rendezvous point advertising interval to 120 seconds:

```
configure pim crp timer 120
```

configure pim crp vlan

`configure pim crp vlan <vlan_name> [none | <policy>] {<priority>}`

Description

Configures the dynamic candidate rendezvous point (CRP) for PIM sparse-mode operation.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
<code>none</code>	Specifies no policy file.
<code>policy</code>	Specifies a policy file name.
<code>priority</code>	Specifies a priority setting. The range is 0 - 254.

Default

The default setting for priority is 0 and indicates the highest priority.

Usage Guidelines

NETGEAR 8800 switches support up to 50 RPs in a switch, and up to 180 groups within a single RP. If you configure more than 180 groups for a single RP, the switch does not process the groups added after the first 180.

The policy file contains the list of multicast group addresses serviced by this RP. This set of group addresses are advertised as candidate RPs. Each router then elects the common RP for a group address based on a common algorithm. This group to RP mapping should be consistent on all routers.

This policy file is not used for filtering purposes. As used with this command, the policy file is just a container for a list of addresses. So a typical policy file used for RP configuration looks a little different from a policy used for other purposes. The following is a sample policy file which configures the CRP for the address ranges 239.0.0.0/24 and 232.144.27.0/24:

```
entry netgear1 {
    if match any {
    }
    then {
        nlri 239.0.0.0/24 ;
        nlri 232.144.27.0/24 ;
    }
}
```

The VLAN specified for a CRP must have PIM configured.

To delete a CRP, use the keyword `none` as the access policy.

Example

The following command configures the candidate rendezvous point for PIM sparse-mode operation on the VLAN *HQ_10_0_3* with the policy *rp-list* and priority set to 30:

```
configure pim crp HQ_10_0_3 rp-list 30
```

configure pim delete vlan

```
configure pim delete vlan [<vlanname> | all]
```

Description

Disables PIM on a router interface.

Syntax Description

vlan name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables PIM on VLAN *accounting*:

```
configure pim delete vlan accounting
```

configure pim register-rate-limit-interval

```
configure pim register-rate-limit-interval <interval>
```

Description

Configures the initial PIM-SM periodic register rate.

Syntax Description

interval	Specifies an interval time in seconds. Range is 0 - 60. Default is 0.
----------	---

Default

Default is 0.

Usage Guidelines

Configuring a non-zero interval time can reduce the CPU load on the first hop switch, in case register stop messages are not received normally.

When a non-zero value is configured, the first hop switch sends a few register messages and then waits for a corresponding register stop from the RP for `<time>` seconds. The process is repeated until the register stop is received. This command should be used when the (S,G) tree between the first hop router and the RP is not converging quickly.

The default value is zero in default mode, the switch sends continuous register messages until the register stop is received.

Example

The following command configures the initial PIM register rate limit interval:

```
configure pim register-rate-limit-interval 2
```

configure pim register-suppress-interval register-probe-interval

```
configure pim register-suppress-interval <reg-interval> register-probe-interval
<probe_interval>
```

Description

Configures an interval for periodically sending null-registers.

Syntax Description

reg-interval	Specifies an interval time in seconds. Range is 30 - 200 seconds. Default is 60.
probe-interval	Specifies an interval time in seconds. Default is 5.

Default

The following defaults apply:

- register-suppress-interval—60
- register-probe-interval—5

Usage Guidelines

The register-probe-interval time should be set less than the register-suppress-interval time. By default, a null register is sent every 55 seconds (*register-suppress-interval - register-probe-interval*). A response to the null register is expected within register probe interval. By specifying a larger interval, a CPU peak load can be avoided because the null-registers are generated less frequently. The register probe time should be less than half of the register suppress time, for best results.

Example

The following command configures the register suppress interval and register probe time:

```
configure pim register-suppress-interval 90 register-probe time 10
```

configure pim register-checksum-to

```
configure pim register-checksum-to [include-data | exclude-data]
```

Description

Configures the checksum computation to either include data (for compatibility with Cisco Systems products) or to exclude data (for RFC-compliant operation), in the register message.

Syntax Description

include-data	Specifies to include data.
exclude-data	Specifies to exclude data.

Default

Include data.

Usage Guidelines

None.

Example

The following command configures the checksum mode to include data for compatibility with Cisco Systems products:

```
configure pim register-checksum-to include-data
```

configure pim spt-threshold

```
configure pim spt-threshold <leaf-threshold> {<rp_threshold>}
```

Description

Configures the threshold, in kbps, for switching to SPT. On leaf routers, this setting is based on data packets. On the RP, this setting is based on register packets.

Syntax Description

leaf-threshold	Specifies the rate of traffic per (s,g,v) group in kbps for the last hop. Range is 0 - 419403.
rp_threshold	Specifies an RP threshold. Range is 0 - 419403.

Default

The default setting is 0 for both parameters.

Usage Guidelines

For the best performance, use default value of 0.

Example

The following command sets the threshold for switching to SPT:

```
configure pim spt-threshold 4 16
```

configure pim ssm range

```
configure pim ssm range [default | policy <policy-name>]
```

Description

Configures the range of multicast addresses for PIM SSM.

Syntax Description

default	Specifies the default address range, 232.0.0.0/8.
policy-name	Specifies a policy that defines the SSM address range.

Default

By default, no SSM range is configured. Using this command with the `default` keyword sets the range to 232.0.0.0/8. To reset the switch to the initial state, use the `unconfigure pim ssm range` command.

Usage Guidelines

You must disable PIM before configuring or unconfiguring a PIM-SSM range. Use the `disable pim` command.

Initially, no range is configured for SSM. After a range is configured, you can remove the range with the `unconfigure pim ssm range` command. If you wish to change the PIM SSM range, you must first unconfigure the existing range, and then configure the new range.

SSM requires that hosts use IGMPv3 messages to register to receive multicast group packets. When a range is configured for SSM, any IGMPv2 messages for an address in the range are ignored. Also, any IGMPv3 Exclude messages are ignored.

Note: If a PIM-SSM range is configured, IGMPv2 messages and IGMPv3 exclude messages within the PIM-SSM range are ignored on all IP interfaces, whether or not PIM-SSM is configured on the interfaces.

To specify a range different from the default PIM SSM range, create a policy file. The match statement of the policy file contains the group addresses to be treated as PIM SSM addresses. For example, to specify the PIM SSM address range as 232.0.0.0/8 and 233.0.0.0/8, use the following policy file:

```
Entry netgear1 {
    if match any {
        nlri 232.0.0.0/8 ;
        nlri 233.0.0.0/8 ;
    }
    then {
        permit ;
    }
}
```

Example

The following command sets the PIM SSM range to 232.0.0.0/8 and 233.0.0.0/8, if the policy file `ssmrange.pol` contains the policy example used above:

```
configure pim ssm range policy ssmrange.pol
```

configure pim state-refresh

```
configure pim state-refresh {vlan} [<vlanname> | all] [on | off]
```

Description

Enables or disables the PIM-DM state refresh feature on one or all VLANs.

Syntax Description

<code>vlanname</code>	Specifies a VLAN on which to enable or disable the PIM-DM state refresh feature.
<code>on</code>	Enables the PIM-DM state refresh feature on the specified VLANs.
<code>off</code>	Disables the PIM-DM state refresh feature on the specified VLANs.

Default

Disabled.

Usage Guidelines

When this feature is disabled on an interface, the interface behaves as follows:

- State refresh messages are not originated.
- State refresh messages received on the interface are dropped without processing.
- State refresh messages received on other interfaces are not forwarded to the disabled interface.

Example

The following command enables the PIM-DM state refresh feature on VLAN *blue*:

```
configure pim state-refresh blue on
```

configure pim state-refresh timer origination-interval

```
configure pim state-refresh timer origination-interval <interval>
```

Description

Configures the interval at which state refresh messages are originated.

Syntax Description

interval	Specifies a refresh interval in seconds. The range is 30 to 90 seconds.
----------	---

Default

60 seconds.

Usage Guidelines

None.

Example

The following command configures the interval to 45 seconds:

```
configure pim state-refresh timer origination-interval 45
```

configure pim state-refresh timer source-active-timer

```
configure pim state-refresh timer source-active-timer <interval>
```

Description

Defines how long a multicast source (S,G) is considered active after a packet is received from the source.

Syntax Description

interval	Specifies a source-active timer interval in seconds. The range is 90 to 300 seconds.
----------	--

Default

210 seconds.

Usage Guidelines

None.

Example

The following command configures the interval to 45 seconds:

```
configure pim state-refresh timer source-active-timer 180
```

configure pim state-refresh ttl

```
configure pim state-refresh ttl <ttlvalue>
```

Description

Configures a time-to-live (TTL) value for PIM-DM state refresh messages.

Syntax Description

ttl	Specifies a TTL value. The range is 1 to 64.
-----	--

Default

16.

Usage Guidelines

None.

Example

The following command the TTL value for 24:

```
configure pim state-refresh ttl 24
```

configure pim timer vlan

```
configure pim timer <hello_interval> <jp_interval> [{vlan} <vlan_name> | vlan all]
```

Description

Configures the global PIM timers on the specified router interfaces.

Syntax Description

hello_interval	Specifies the amount of time before a hello message is sent out by the PIM router. The range is 1 to 65,535 seconds.
jp_interval	Specifies the join/prune interval. The range is 1 to 65,535 seconds.
vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

- hello_interval—30 seconds.
- jp_interval—60 seconds.

Usage Guidelines

These default timers should only be adjusted when excess PIM control packets are observed on the interface.

Example

The following command configures the PIM timers on the VLAN *accounting*:

```
configure pim timer 150 300 vlan accounting
```

configure pim vlan trusted-gateway

```
configure pim [{vlan} <vlan_name> | vlan all] trusted-gateway [<policy> | none]
```

Description

Configures a trusted neighbor policy.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.
policy	Specifies an policy file name.
none	Specifies no policy file, so all gateways are trusted.

Default

No policy file, so all gateways are trusted.

Usage Guidelines

Because PIM leverages the unicast routing capability that is already present in the switch, the access policy capabilities are, by nature, different. When the PIM protocol is used for routing IP multicast traffic, the switch can be configured to use a policy file to determine trusted PIM router neighbors for the VLAN on the switch running PIM. This is a security feature for the PIM interface.

Example

The following command configures a trusted neighbor policy on the VLAN *backbone*:

```
configure pim vlan backbone trusted-gateway nointernet
```

disable igmp

```
disable igmp {vlan <name>}
```

Description

Disables IGMP on a router interface. If no VLAN is specified, IGMP is disabled on all router interfaces.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, hosts respond to the query, and group registration is maintained.

IGMP is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of IGMP packets. IGMP should be enabled when the switch is configured to perform IP multicast routing.

This command disables IGMPv2 and IGMPv3.

Example

The following command disables IGMP on VLAN *accounting*:


```
disable igmp vlan accounting
```

disable igmp snooping

```
disable igmp snooping {forward-mcrouter-only | with-proxy | vlan <name>}
```

Description

Disables IGMP snooping.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
with-proxy	Disables the IGMP snooping proxy.
name	Specifies a VLAN.

Default

IGMP snooping and the with-proxy option are enabled by default, but forward-mcrouter-only option is disabled by default.

Usage Guidelines

If a VLAN is specified, IGMP snooping is disabled only on that VLAN, otherwise IGMP snooping is disabled on all VLANs.

This command applies to both IGMPv2 and IGMPv3.

If the switch is in the `forward-mcrouter-only` mode, then the command `disable igmp snooping forward-mcrouter-only` changes the mode so that all multicast traffic is forwarded to any IP router. If not in the `forward-mcrouter-mode`, the command `disable igmp snooping forward-mcrouter-only` has no effect.

To change the snooping mode you must disable IP multicast forwarding. Use the command:

```
disable ipmcforwarding
```

The with-proxy option can be used for troubleshooting purpose. It should be enabled for normal network operation.

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary IGMP leave messages so that they are forwarded only when the last member leaves the group.

Example

The following command disables IGMP snooping on the VLAN *accounting*:

```
disable igmp snooping accounting
```

disable igmp snooping vlan fast-leave

```
disable igmp snooping {vlan} <name> fast-leave
```

Description

Disables the IGMP snooping fast leave feature on the specified VLAN.

Syntax Description

name	Specifies a VLAN.
------	-------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the IGMP snooping fast leave feature on the default VLAN:

```
disable igmp snooping "Default" fast-leave
```

disable igmp ssm-map

```
disable igmp ssm-map {vr <vr-name>}
```

Description

Disables IGMP SSM mapping.

Syntax Description

vr-name	Specifies a virtual router name. If the VR name is omitted, the switch disables mapping on the VR specified by the current CLI VR context.
---------	--

Default

Disabled on all interfaces.

Usage Guidelines

None.

Example

The following command disables IGMP-SSM mapping on the VR in the current CLI VR context:

```
disable igmp ssm-map
```

disable ipmcforwarding

```
disable ipmcforwarding {vlan <name>}
```

Description

Disables IP multicast forwarding on a router interface.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IP multicast forwarding is disabled by default.

IP forwarding must be enabled before enabling IP multicast forwarding.

Disabling IP multicast forwarding disables any Layer 3 multicast routing for the streams coming to the interface.

Example

The following command disables IP multicast forwarding on the VLAN *accounting*:

```
disable ipmcforwarding vlan accounting
```

disable mvr

```
disable mvr
```

Description

Disables MVR on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables MVR on the system:

```
disable mvr
```

disable pim

```
disable pim
```

Description

Disables PIM on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables PIM on the system:

```
disable pim
```

disable pim snooping

```
disable pim snooping {{vlan} <name>}
```

Description

Disables PIM snooping and clears all the snooping PIM neighbors, joins received on the VLAN, and the forwarding entries belonging to one or all VLANs.

Syntax Description

name	Specifies a VLAN.
------	-------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables PIM snooping for all VLANs on the switch:

```
disable pim snooping
```

disable pim ssm vlan

```
disable pim ssm vlan [<vlan_name> | all]
```

Description

Disables PIM SSM on a router interface.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled on all interfaces.

Usage Guidelines

This command disables PIM-SSM on the specified Layer 3 VLAN.

IGMPv3 include messages for multicast addresses in the SSM range is only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following command disables PIM-SSM multicast routing on VLAN *accounting*:

```
disable pim ssm vlan accounting
```

enable igmp

```
enable igmp {vlan <vlan name>} {IGMPv1 | IGMPv2 | IGMPv3}
```

Description

Enables IGMP on a router interface. If no VLAN is specified, IGMP is enabled on all router interfaces.

Syntax Description

vlan name	Specifies a VLAN name.
IGMPv1	Specifies the compatibility mode as IGMPv1.
IGMPv2	Specifies the compatibility mode as IGMPv2.
IGMPv3	Specifies the compatibility mode as IGMPv3.

Default

Enabled, set to IGMPv2 compatibility mode.

Usage Guidelines

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, IP hosts respond to the query, and group registration is maintained.

IGMPv2 is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of IGMP packets. IGMP should be enabled when the switch is configured to perform IP multicast routing.

Example

The following command enables IGMPv2 on the VLAN *accounting*:

```
enable igmp vlan accounting
```

The following command enables IGMPv3 on the VLAN *finance*:

```
enable igmp vlan finance igmpv3
```

enable igmp snooping

```
enable igmp snooping {forward-mcrouter-only | {vlan} <name> | with-proxy vr <vrname>}
```

Description

Enables IGMP snooping on one or all VLANs.

Syntax Description

<code>forward-mcrouter-only</code>	Specifies that the switch forward all multicast traffic to the multicast router only.
<code>name</code>	Specifies a VLAN or vMAN on which to enable IGMP snooping.
<code>with-proxy vr vname</code>	Controls how join and leave messages are forwarded from the specified virtual router. If this option is specified, one join message per query is forwarded, and a leave message is forwarded only if it is from the last receiver on the VLAN.

Default

Enabled.

Usage Guidelines

This command applies to both IGMPv2 and IGMPv3.

IGMP snooping is enabled by default on the switch. If you are using multicast routing, IGMP snooping can be enabled or disabled. If IGMP snooping is disabled, all IGMP and IP multicast traffic floods within a given VLAN or vMAN.

The `forward-mcrouter-only`, `vlan`, and `with-proxy` options control three separate and independent features. You can manage one feature at a time with the `enable igmp snooping` command, and you can enter the command multiple times as needed to control each feature. For example, you can enter the command twice to enable both the `forward-mcrouter-only` and `with-proxy` options.

If a VLAN or vMAN is specified with the `enable igmp snooping` command, IGMP snooping is enabled only on that VLAN or vMAN. If no options are specified, IGMP snooping is enabled on all VLANs.

The `with-proxy` option enables the IGMP snooping proxy feature, which reduces the number of join and leave messages forwarded on the virtual router as described in the table above. This feature is enabled by default.

An optional optimization for IGMP snooping is the strict recognition of routers only if the remote devices are running a multicast protocol. Two IGMP snooping modes are supported:

- The `forward-mcrouter-only` mode forwards all multicast traffic to the multicast router (that is, the router running PIM, DVMRP or CBT).
- When not in the `forward-mcrouter-only` mode, the switch forwards all multicast traffic to any IP router (multicast or not), and any active member port to the local network that has one or more subscribers.

Note: The forward-mcrouter-only mode for IGMP snooping is enabled/disabled on a switch-wide basis, not on a per-VLAN basis. In other words, all the interfaces enabled for IGMP snooping are either in the forward-mcrouter-only mode or in the non-forward-mcrouter-only mode, and not a mixture of the two modes.

To change the snooping mode you must disable IP multicast forwarding. To disable IP multicast forwarding, use the command:

```
disable ipmcforwarding {vlan <name>}
```

To change the IGMP snooping mode from the non-forward-mcrouter-only mode to the forward-mcrouter-only mode, use the commands:

```
disable ipmcforwarding
enable igmp snooping forward-mcrouter-only
enable ipmcforwarding {vlan <name>}
```

To change the IGMP snooping mode from the forward-mcrouter-only mode to the non-forward-mcrouter-only mode, use the commands:

```
disable ipmcforwarding
disable igmp snooping forward-mcrouter-only
enable ipmcforwarding {vlan <name>}
```

Example

The following command enables IGMP snooping on the switch:

```
enable igmp snooping
```

enable igmp snooping vlan fast-leave

```
enable igmp snooping {vlan} <name> fast-leave
```

Description

Enables the IGMP snooping fast leave feature on the specified VLAN.

Syntax Description

name	Specifies a VLAN.
------	-------------------

Default

Disabled.

Usage Guidelines

The fast leave feature operates only with IGMPv2.

To view the fast leave feature configuration, use the `show configuration msmgr` command. This `show` command displays the fast leave configuration only when the feature is enabled.

Example

The following command enables the IGMP snooping fast leave feature on the default VLAN:

```
enable igmp snooping "Default" fast-leave
```

enable igmp snooping with-proxy

```
enable igmp snooping with-proxy
```

Description

Enables the IGMP snooping proxy. The default setting is enabled.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary IGMP leave messages so that they are forwarded only when the last member leaves the group.

This command can be used for troubleshooting purpose. It should be enabled for normal network operation. The command does not alter the snooping setting.

This feature can be enabled when IGMPv3 is enabled; however, it is not effective for IGMPv3.

Example

The following command enables the IGMP snooping proxy:

```
enable igmp snooping with-proxy
```

enable igmp ssm-map

```
enable igmp ssm-map {vr <vr-name>}
```

Description

Enables IGMP SSM mapping on a VR.

Syntax Description

vr-name	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.
---------	---

Default

Disabled on all interfaces.

Usage Guidelines

Configure the range of multicast addresses for PIM SSM with the `configure pim ssm range [default | policy <policy-name>]` command before you enable IGMP SSM mapping. IGMP SSM mapping operates only with IPv4.

Example

The following command enables IGMP-SSM mapping on the VR in the current CLI VR context:

```
enable igmp ssm-map
```

enable ipmcforwarding

```
enable ipmcforwarding {vlan <name>}
```

Description

Enables IP multicast forwarding on an IP interface.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IPMC forwarding is disabled by default.

IP forwarding must be enabled before enabling IPMC forwarding.

Example

The following command enables IPMC forwarding on the VLAN *accounting*:

```
enable ipmcforwarding vlan accounting
```

enable mvr

```
enable mvr
```

Description

Enables MVR on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables MVR on the system:

```
enable mvr
```

enable pim

```
enable pim
```

Description

Enables PIM on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables PIM on the system:

```
enable pim
```

enable pim snooping

```
enable pim snooping [{vlan} <name>]
```

Description

Enables PIM snooping on one or all VLANs.

Syntax Description

name	Specifies a VLAN.
------	-------------------

Default

Disabled.

Usage Guidelines

PIM snooping does not require PIM to be enabled. However, IGMP snooping must be disabled on VLANs that use PIM snooping. PIM snooping and MVR cannot be enabled simultaneously on a switch. PIM snooping should not be enabled on a VLAN that supports PIM-DM neighbors.

Example

The following command enables PIM snooping on the default VLAN:

```
enable pim snooping default
```

enable pim ssm vlan

```
enable pim ssm vlan [<vlan_name> | all]
```

Description

Enables PIM SSM on an IP interface.

Syntax Description

vlan_name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled on all interfaces.

Usage Guidelines

This command enables PIM-SSM on the specified Layer 3 VLAN.

PIM-SM must also be configured on the interface for PIM to begin operating (which includes enabling IP multicast forwarding).

IGMPv3 include messages for multicast addresses in the SSM range are only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 include messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following command enables PIM-SSM multicast routing on VLAN *accounting*:

```
enable pim ssm vlan accounting
```

mrinfo

```
mrinfo {<router_address>} {from <from_address>} {timeout <seconds>}  
{multiple-response-timeout <multi_resp_timeout>} {vr <vrname>}
```

Description

Requests information from a multicast router.

Syntax Description

router_address	Specifies the unicast IP address of the router for which you want information.
from_address	Specifies the unicast IP address of the interface where the mrinfo request is generated.
seconds	Specifies a maximum time to wait for a response. The range is 1 to 30 seconds.
multi_resp_timeout	Specifies a maximum time to wait for additional responses after the first response is received. The range is 0 to 3 seconds.
vrname	Specifies a VR name.

Default

router_address: One of the local interface addresses.

from: IP address of interface from which the mrinfo query is generated.

timeout: 3 seconds

multiple-response-timeout: 1 second

vr: DefaultVR

Usage Guidelines

The last column of the `mrinfo` command output displays information in the following format:

```
[Metric/threshold/type/flags]
```

This information is described in [Table 28](#).

Table 28. mroute Command Display Data

Data	Description
Metric	This should always be 1 because mroute queries the directly connected interfaces of a device.
Threshold	This should always be 0 because the threshold feature is not supported.
Type	The type specifies the multicast protocol type. Because the NETGEAR 8800 only supports PIM, this value is always <code>pim</code> .
querier	The <code>querier</code> flag indicates that the queried router is the IGMP querier.
leaf	The <code>leaf</code> flag indicates that the IP interface has no neighbor router.
down	The <code>down</code> flag indicates that the interface link status is <i>down</i> .

Example

The following command requests information from multicast router `1.1.1.1`:

```
Switch.1 # mroute 1.1.1.1
1.1.1.1 [Flags:PGM]
2.2.2.1 -> 2.2.2.2 [1/0/pim/querier]
1.1.1.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
8.8.8.1 -> 8.8.8.4 [1/0/pim/querier]
3.3.3.1 -> 0.0.0.0 [1/0/pim/down]
```

mtrace

```
mtrace source <src_address> {destination <dest_address>} {group <grp_address>} {from
<from_address>} {gateway <gw_address>} {timeout <seconds>} {maximum-hops <number>} {vr
<vrname>}
```

Description

Traces multicast traffic from the receiver back to the source.

Syntax Description

<code>src_address</code>	Specifies the unicast IP address of the multicast source.
<code>dest_address</code>	Specifies the unicast IP address of the multicast group receiver.
<code>grp_address</code>	Specifies the multicast IP address of the group.
<code>from_address</code>	Specifies the unicast IP address of the interface where the mtrace request originates. This is used as the IP destination address of the mtrace response packet.

gw_address	Specifies the gateway router IP address of the multicast router to which the unicast mtrace query is sent.
seconds	Specifies a maximum time to wait for the mtrace response before making the next attempt. The range is 1 to 30 seconds.
number	Specifies the maximum number of hops for the trace. The range is 1 to 255.
vrname	Specifies a VR name.

Default

destination: IP address of interface from which mtrace query is generated.

group: 0.0.0.0

from: IP address of interface from which mtrace query is generated.

gateway: 224.0.0.2 when the destination is in the same subnet as one of the IP interfaces. For a non-local destination address, it is mandatory to provide a valid multicast router address.

timeout: 3 seconds

maximum-hops: 32

vr: DefaultVR

Usage Guidelines

The multicast traceroute initiator node generates a multicast query and waits for timeout period to expire. If there is no response for the timeout period, the initiator node makes 2 more attempts. If no response is received after 3 attempts, the initiator node moves to a hop-by-hop trace by manipulating the maximum hop fields to perform a linear search.

The multicast trace response data contains the following fields:

- Incoming interface address—Interface on which traffic is expected from the specific source and group
- Outgoing interface address—Interface on which traffic is forwarded from the specified source and group towards the destination
- Previous hop router address
- Input packet count on incoming interface
- Output packet count on outgoing interface
- Total number of packets for this source-group pair
- Multicast routing protocol
- Forwarding code

NETGEAR switches set the packet count statistics field to 0xffffffff to indicate that this field is not supported.

The last column of the `mtrace` command output displays forwarding codes, which are described in [Table 29](#).

Table 29. mtrace Command Forwarding Codes

Forwarding Code	Description
Wrong interface	mtrace request arrived on an interface to which this router would not forward for this source and group.
Prune sent upstream	This router has sent a prune request upstream for the source and group in the mtrace request.
Output pruned	This router has stopped forwarding for this source and group in response to a prune request from the next hop router.
Hit scope boundary ^a	The group is subject to administrative scoping at this hop.
No route	This router has no route for the source or group and no way to determine a potential route.
Wrong Last Hop	This router is not the proper last-hop router.
Not forwarding ^a	This router is not forwarding for this source and group on the outgoing interface for an unspecified reason.
Reached RP/Core	Reached rendezvous point or core.
RPF Interface	mtrace request arrived on the expected RPF interface (upstream interface) for this source and group.
Multicast disabled	mtrace request arrived on an interface which is not enabled for multicast.
Info. Hidden ^a	One or more hops have been hidden from this trace.
No space in packet	There was not enough room to insert another response data block in the packet.
Next router no mtrace ^a	The previous hop router does not understand mtrace requests.
Admin. Prohibited ^a	mtrace is administratively prohibited.

a. NETGEAR 8800 switches along the mtrace path do not provide this forwarding code.

Example

The following command initiates an mtrace for group 225.1.1.1 at IP address 1.1.1.100:

```
Switch.6 # mtrace source 1.1.1.100 group 225.1.1.1
Mtrace from 1.1.1.100 to Self via 225.1.1.1
  0          34.2.2.4
-1          34.2.2.4  PIM thresh^ 0          1.1.1.100/32  RPF Interface
-2          34.2.2.3  PIM thresh^ 0          1.1.1.100/32
-3          23.1.1.2  PIM thresh^ 0          1.1.1.100/32
-4          2.2.2.1   PIM thresh^ 0          1.1.1.100/32
Round trip time 9 ms; total ttl of 4 required.
```


rtlookup

```
rtlookup [<ipaddress> | <ipv6address>] { unicast | multicast | rpf } { vr <vr_name> }
```

Description

Displays the available routes to the specified IP address.

Syntax Description

ipaddress	Specifies an IPv4 address.
ipv6address	Specifies an IPv6 address.
unicast	Displays the routes from the unicast routing table in the current router context.
multicast	Displays the routes from the multicast routing table in the current router context.
rpf	Displays the RPF route to the specified destination.
vr-name	Specifies the virtual router for which to display the route.

Default

vr-name is the VR of the current CLI context.

When no option (`unicast` or `multicast`) is provided, this command displays the route in the unicast routing table.

Usage Guidelines

None.

Example

The following example displays the route lookup for 12.1.20.12 in the multicast routing table for the default VR:

```
# rtlookup 12.1.20.12 multicast vr vr-default
@mbe 12.1.0.0/16          50.1.10.21      1      UG---S--m--- toronto      0d:0h:41m:1s

Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
             (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
             (e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
             (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
             (mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
             (oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
             (r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
             (*) Preferred unicast route (@) Preferred multicast route
             (#) Preferred unicast and multicast route
```

Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
 (L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
 (P) LPM-routing, (R) Modified, (S) Static, (s) Static LSP
 (T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
 (c) Compressed Route

```
Mariner # rtlookup 12.1.20.12 multicast vr vr-default
No route to 12.1.10.12
```

rtlookup rpf

```
rtlookup [<ipaddress> | <ipv6address>] rpf {vr <vr_name>}
```

Description

Displays the RPF for a specified multicast source.

Syntax Description

ipaddress	Specifies an IPv4 address.
ipv6address	Specifies an IPv6 address.
rpf	Selects the RPF for the specified multicast source.
vr-name	Specifies the virtual router for which to display the route.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

None.

Example

The following example displays the RPF lookup for multicast source 12.1.20.12 in the default VR:

```
# rtlookup 12.1.20.12 rpf vr vr-default
```

```

Ori Prefix      Route      Gateway      VLAN
@d  12.1.10.22  12.1.10.0/24  12.1.10.10   v1

```

```
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
              (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
              (e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
```

```
(mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or)OSPFIInter, (pd) PIM-DM,(ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) unicast route (@) multicast route
```

show igmp

```
show igmp {vlan} {<vlan name>}
```

Description

This command can be used to display an IGMP-related configuration and group information, per VLAN.

Syntax Description

vlan name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

The output of this command shows:

- The VLAN name.
- The router interface IP address and subnet mask.
- If the interface is active (up), by the letter U.
- If IP forwarding is enabled for the interface, by the letter f.
- If multicast forwarding is enabled, by the letter M.
- If IGMP is enabled, by the letter i.
- If IGMP snooping is enabled, by the letter z.

Example

The following command displays the IGMP configuration:

```
show igmp
```

The following is sample output from this command:

```
IGMP:
  Query Interval: 125 sec
  Max Response Time: 10 sec
  Last Member Query: 1 sec
  Robustness: 2
```

IGMP Snooping:

```

Router Timeout: 260 sec
Host Timeout: 260 sec
Igmp Snooping Fast Leave Time: 1000 ms
Igmp Snooping Flag: forward-all-router
Igmp Snooping Flood-list: none
Igmp Snooping Proxy: Enable
Igmp Snooping Filters: per-port

```

VLAN	IP Address		Flags	nLRMA	nLeMA	IGMPver
default	0.0.0.0	/ 0	----z	0	0	2
gho	0.0.0.0	/ 0	----z	0	0	2
hguo_fo	0.0.0.0	/ 0	----z	0	0	2
sqa_east	1.1.1.1	/24	-fmiz	3	0	2
vcs1	12.1.1.115	/24	Ufmiz	6	0	2
vcs2	12.1.2.115	/24	Ufmiz	6	0	2
vcs3	12.2.3.115	/24	-fmiz	3	0	2
vcs4	12.2.4.115	/24	Ufmiz	6	1	2
vcs5	12.2.5.115	/24	-fmiz	3	0	2
vcs6	12.2.6.115	/24	-fmiz	3	0	2
vcs7	12.2.7.115	/24	-fmiz	3	0	2
vcs8	12.2.8.115	/24	-fmiz	3	0	2
vhs1	0.0.0.0	/ 0	U---z	0	4	2
vhs2	117.2.2.115	/24	-fmiz	3	0	2
vhs3	117.2.3.115	/24	-fmiz	3	0	2
vhs4	117.2.4.115	/24	-fmiz	3	0	2
vms1	111.1.1.115	/24	Ufmiz	6	7	2

Flags: (E) Interface Enabled, (i) IGMP Enabled
 (f) Forwarding Enabled, (m) Multicast Forwarding Enabled
 (nLeMA) Number of Learned Multicast Addressess
 (nLRMA) Number of Locally registered Multicast Addresses
 (U) Interface Up, (z) IGMP Snooping Enabled
 (p) IGMP Proxy Query Enabled

show igmp group

```
show igmp group {{vlan} {<name>} | {<grpipaddress>}} {IGMPv3}
```

Description

Lists the IGMP group membership for the specified VLAN.

Syntax Description

grpipaddress	Specifies a group IP address.
--------------	-------------------------------

name	Specifies a VLAN name.
IGMPv3	Display the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise display in earlier format).

Default

IGMPv2.

Usage Guidelines

If no VLAN is specified all VLANs are displayed. You can also filter the display by group address and by multicast stream sender address.

The output of this command shows:

- The multicast group address received.
- The version of the IGMP group.
- The name of the VLAN where the group address is being received.
- The physical port where the group address is being received. If multiple ports within the VLAN have subscribers for the group, all the ports are listed.
- The age since the last IGMP report for this group was received.

Note: The show igmp group command output is populated on the router that is the PIM Rendezvous Point.

Example

The following command lists the IGMP group membership:

```
show igmp group
```

The following is sample output from this command:

Group Address	Ver	Vlan	Port	Age
239.2.4.70	2	banana	7	101
224.0.1.24	2	banana	7	107
239.255.255.254	2	banana	7	103

Total: 3

show igmp snooping

```
show igmp snooping {detail {IGMPv3}}
```

Description

Displays IGMP snooping registration information for all VLANs.

Syntax Description

detail	Displays the information in detailed format.
IGMPv3	Display the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise display in earlier format).

Default

IGMPv2.

Usage Guidelines

None.

Example

The following command displays IGMP snooping registration information for all VLANs:

```
show igmp snooping
```

Vlan	Vid	Port	#Senders	#Receivers	Router	Enable
default	1		0			Yes
vhs3	4090		0			Yes
vhs4	4089		0			Yes
vcs5	15		0			Yes
vcs6	16		0			Yes
vcs3	4086		0			Yes
vcs4	1014		0			Yes
		5:7		5	No	
		5:9		5	No	
		5:10		5	No	
		5:11		1	No	
		5:12		5	No	
		5:37		5	No	
		5:39		5	No	
		5:41		5	No	
		5:42		5	No	
vcs7	4084		0			Yes
vcs8	4083		0			Yes
vhs2	4082		0			Yes
hguo_fo	200		0			Yes
vcs1	12		8			Yes
		4:16		0	Yes	
vcs2	22		8			Yes
		4:16		0	Yes	
vhs1	1717		14			Yes
		4:32		0	Yes	

vms1	111	2		Yes
		4:10	5	Yes
gho	4061	0		Yes
sqa_east	4059	0		Yes

show igmp snooping cache

This command is provided for backward compatibility. The recommended command is:

```
show mcast cache {{vlan} <name>} {[group <grpaddressMask> | <grpaddressMask>] {source <sourceIP> | <sourceIP>}} {type [snooping | pim | mvr]} | {summary}}
```

The syntax for the original form of this command is:

```
show igmp snooping cache {{vlan} <name>} {{group} <grpaddressMask>}
```

Description

Displays multicast cache entries added by IGMP snooping for all VLANs and groups. The display can be limited to specific VLANs or groups.

Syntax Description

name	Specifies a VLAN name.
grpaddressMask	Specifies a multicast group address and mask.

Default

Displays information for all VLANs and groups.

Usage Guidelines

None.

Example

The following command displays IGMP snooping cache information for all VLANs and groups:

```
XCM8806.2 # show igmp snooping cache
```

This command display is the same as for the following preferred command:

```
show mcast cache {{vlan} <name>} {[group <grpaddressMask> | <grpaddressMask>] {source <sourceIP> | <sourceIP>}} {type [snooping | pim | mvr]} | {summary}}
```

show igmp snooping vlan

```
show igmp snooping {vlan} <name> {port <port>} {IGMPv3}
```

Description

Displays IGMP snooping registration information for a specific VLAN. The display can be further limited to a specific port or to only IGMPv3 entries.

Syntax Description

name	Specifies a VLAN name.
port	Specifies a single port for which information is displayed.
IGMPv3	Display the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise display in earlier format).

Default

IGMPv2.

Usage Guidelines

The two types of IGMP snooping entry are sender entry and subscribed entry.

The following information is displayed in a per-interface format:

- Group membership information
- Router entry
- Timeout information
- Sender entry

Example

The following command displays IGMP snooping registration information on VLAN *accounting*:

```
show igmp snooping vlan accounting
VLAN accounting (4093) Snooping=Enabled
  Port Host          Subscribed  Age
  5:10 192.206.38.103 235.1.1.1(s) 0
  5:12 192.206.38.103 235.1.1.1(s) 0
  5:13 192.206.38.103 235.1.1.1(s) 0
  5:14 192.206.38.103 235.1.1.1(s) 0
s = static igmp member
```

The following command displays IGMP snooping registration information for port 2:1 on VLAN *test*:

```
show igmp snooping test port 2:1
```

show igmp snooping vlan filter

```
show igmp snooping {vlan} <name> filter
```


Description

Displays IGMP snooping filters.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

None.

Usage Guidelines

Use this command to display IGMP snooping filters configured on the specified VLAN. When no VLAN is specified, all the filters are displayed.

Example

The following command displays the IGMP snooping filter configured on VLAN *vlan101*:

```
show igmp snooping vlan101 filter
Filter          Port Flags
igmpppermit0   5:10 a
```

Flags: (a) Active

show igmp snooping vlan static

```
show igmp snooping {vlan} <name> static [group | router]
```

Description

Displays static IGMP snooping entries.

Syntax Description

name	Specifies a VLAN name.
group	Displays static multicast groups.
router	Displays static router entries.

Default

None.

Usage Guidelines

Use this command to display the IGMP snooping static groups or router ports configured on the specified VLAN. When no VLAN is specified, all the static groups or router ports are displayed.

Example

The following command displays the IGMP snooping static groups configured on VLAN *vlan101*:

```
show igmp snooping vlan101 static group
```

```
VLAN vlan101 (4094)
```

Group	Port	Flags
239.1.1.2	29	s-
239.1.1.2	30	s-
239.1.1.2	31	sa
239.1.1.2	32	s-
239.1.1.2	34	s-

```
Total number of configured static IGMP groups = 5
```

```
Flags: (s) Static, (a) Active
```

show igmp ssm-map

```
show igmp ssm-map {<group_ip>} {vr <vr-name>}
```

Description

Displays the IGMP SSM feature status (enabled or disabled), the mappings for the specified multicast group IP address, and the total count of maps.

Syntax Description

group_ip	Specifies an IP multicast group, for which all mappings in the PIM SSM range are to be displayed. If no group address is specified, the switch displays all IGMP-SSM mappings.
vr-name	Specifies a virtual router name. If the VR name is omitted, the switch displays the mappings on the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

When a target group is specified, this command displays all mapping entries for the configured range in which the group IP address resides.

Examples

The following command displays the mappings for the multicast group IP address 232.1.1.2:

```
show igmp ssm-map 232.1.1.2
```

show ipmroute

```
show ipmroute {<source-net>/<mask-len> | <source-net> <mask> | summary} {vr <vr-name>}
```

Description

Displays the contents of the IP multicast routing table or the route origin priority.

Syntax Description

source-net	Specifies an IP address/mask length.
mask-len	Mask length for the IP multicast source's subnet. Range is [1-32].
mask	Specifies a subnet mask.
summary	Displays the statistics of multicast static routes.
vr-name	Specifies the virtual router to which the route is added.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

This command allows you to view the configured multicast static routes. You can specify the filtering criteria on this CLI to view only the desired route. The multicast static routes are displayed in ascending order of their prefix (same order as `show iproute` displays).

Example

The following example displays a multicast static route from a default virtual router:

```
XCM8806.19 # show ipmroute
```

Destination	Gateway	Mtr	Flags	Protocol	VLAN
Default Route	20.20.20.1	255	UG	None	pc4-1
*1.1.0.0/16	20.20.20.1	10	UG	bgp	pc4-1
*11.0.0.0/8	30.30.30.1	12	U-	None	pc5-3
11.22.0.0/16	20.20.20.1	10	UG	None	pc4-1
*11.22.33.0/24	30.30.30.1	8	U-	None	pc5-3
11.22.33.44/32	20.20.20.1	4	UG	None	pc4-1
*12.0.0.0/8	20.20.20.1	0	UG	None	pc4-1
12.24.0.0/16	30.30.30.1	0	U-	None	pc5-3
*12.24.48.96/32	30.30.30.1	2	U-	ospf-extern1	pc5-3

```
44.66.0.0/16      30.30.30.1      0      U-      None      pc5-3
```

Flags: (*) Active, (G) Gateway, (U) Up

Mask distribution:

```
1 default routes      2 routes at length 8
4 routes at length 16  1 routes at length 24
2 routes at length 32
```

Total number of multicast static routes = 10

show iproute multicast

```
show iproute {ipv4} [{vlan} <name> | [<ipaddress> <netmask> | <ipNetmask>] | origin [direct |
static | mbgp | imbgp | embgp]] multicast {vr <vr_name>}
```

Description

Displays all or a filtered set of multicast routes in the IP multicast routing table.

Syntax Description

ipv4	Selects only IPv4 multicast routes.
name	Specifies a VLAN for which to display multicast routes.
ipaddress netmask	Specifies an IP address and network mask (in dotted decimal notation) for which to display multicast routes.
ipNetmask	Specifies the IP address and network mask in classless inter domain routing (CIDR) notation.
origin	Limits the displayed multicast routes to those generated by the specified origin. Origin options select direct routes, static routes, and routes created by the MBGP, IMBGP, and EMBGP protocols.
vr_name	Specifies the virtual router for which to display multicast routes.

Default

vr_name is the VR of the current CLI context.

Usage Guidelines

This command does not display unicast routes, which can be used for multicast traffic.

Example

The following example displays all the routes in multicast routing table:

```
# show iproute multicast
Ori  Destination      Gateway      Mtr  Flags      VLAN      Duration
@d   3.3.3.3/32        3.3.3.3      1    U-----m--- lpbk      12d:1h:30m:36s
```

```

@d      28.0.0.0/24      28.0.0.15      1      U-----m--- trunk28      12d:1h:30m:36s
@mbe    77.0.0.0/24      50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.1.0/24      50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.2.0/24      50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.3.0/24      50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.4.0/24      50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.5.0/24      50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.6.0/24      50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.10.0/24     50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.11.0/24     50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.12.0/24     50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.13.0/24     50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@mbe    77.0.14.0/24     50.1.10.21     1      UG---S--m--- toronto      0d:0h:41m:1s
@d      82.0.0.0/24      82.0.0.15      1      U-----m--- trunk28-2    12d:1h:30m:36s

```

Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
 (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
 (e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
 (is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
 (mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
 (oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
 (r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
 (*) Preferred unicast route (@) Preferred multicast route
 (#) Preferred unicast and multicast route

Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
 (L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
 (P) LPM-routing, (R) Modified, (S) Static, (s) Static LSP
 (T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
 (f) Provided to FIB (c) Compressed Route

Mask distribution:

```

14 routes at length 24      1 routes at length 32

```

Route Origin distribution:

```

3 routes from Direct

```

Total number of routes = 15

Total number of compressed routes = 0

show L2stats

```
show L2stats {vlan <vlan_name>}
```

Description

Displays the counters for the number of packets bridged, switched, and snooped (Layer 2 statistics).

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the counters for the number of packets bridged, switched, and snooped (Layer 2 statistics) for the VLAN *accounting*:

```
show L2stats accounting
```

Note: You can also enter the command as show l2stats. We use the uppercase letter here to avoid confusion with the numeral 1.

show mcast cache

```
show mcast cache {{vlan} <name>}} {[group <grpaddressMask> | <grpaddressMask>]} {source <sourceIP> | <sourceIP>}} {type [snooping | pim | mvr]} | {summary}}
```

Description

Displays multicast cache information. The display can be limited to entries for specific VLANs or groups, and it can be limited to specific types of entries, such as those created by snooping protocols, PIM, or MVR.

Syntax Description

name	Specifies a VLAN name.
grpaddressMask	Specifies a multicast group address and mask.
sourceIP	Specifies the source IP address for a multicast group.
snooping	Limits the display to cache entries created by PIM or IGMP snooping.

pim	Limits the display to cache entries created by PIM.
mvr	Limits the display to cache entries created by MVR.
summary	Specifies the summary display format.

Default

Displays information for all entries in the multicast cache.

Usage Guidelines

None.

Example

The following command displays all multicast cache information:

```
XCM8806.4 # show mcast cache
```

Type	Group	Sender	Age	InVlan
snoop	225.1.1.1	222.222.222.222	17	snvlan
	Vlan	Port	Vid	
	snvlan	2	400	
		23	400	
snoop	224.0.0.5	100.1.2.2	2	pmvlan2
	Vlan	Port	Vid	
	pmvlan2	4	402	
snoop	224.0.0.5	100.1.3.3	17	pmvlan3
	Vlan	Port	Vid	
	pmvlan3	23	403	
snoop	224.0.0.13	100.1.2.2	11	pmvlan2
	Vlan	Port	Vid	
	pmvlan2	4	402	
snoop	224.0.0.13	100.1.3.3	14	pmvlan3
	Vlan	Port	Vid	
	pmvlan3	23	403	
pim	226.1.1.1	100.1.1.12	0	pmvlan1
	Vlan	Port	Vid	
	pmvlan2	4	402	
	pmvlan3	23	403	

Multicast cache distribution:

```
5 entries from Snooping          0 entries from MVR          1 entries from PIM
```

Total Cache Entries: 6

The following command displays summary cache information for VLAN *pmvlan1*:

```
XCM8806.3 # show mcast cache vlan pmvlan1 summary
```

```
=====MULTICAST CACHE SUMMARY=====
```

Multicast cache distribution:

```
5 entries from Snooping          0 entries from MVR          1 entries from PIM
```

pmvlan1: Multicast cache distribution:

```
0 entries from Snooping          0 entries from MVR          1 entries from PIM
```

Total Cache Entries: 6

Total Cache Entries for VLAN pmvlan1: 1

show mvr

```
show mvr {vlan <vlan_name>}
```

Description

Displays the MVR configuration on the switch.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

If a VLAN is specified, information for the VLAN is displayed.

Example

The following command displays the MVR configuration for the VLAN *accounting*:

```
show mvr accounting
```

show mvr cache

This command is provided for backward compatibility. The recommended command is:


```
show mcast cache {{vlan} <name>} {[group <grpaddressMask> | <grpaddressMask>] {source
<sourceIP> | <sourceIP>}} {type [snooping | pim | mvr]} | {summary}}
```

The syntax for the original form of this command is:

```
show mvr cache {vlan <vlan_name>}
```

Description

Displays the multicast cache entries added by MVR.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

If no VLAN is specified, information for all the VLANs is displayed.

Example

The following command displays the multicast cache in the MVR range for the VLAN *vlan110*:

```
Switch.78 # show mvr cache vlan110
```

This command display is the same as for the following preferred command:

```
show mcast cache {{vlan} <name>} {[group <grpaddressMask> | <grpaddressMask>] {source
<sourceIP> | <sourceIP>}} {type [snooping | pim | mvr]} | {summary}}
```

show pim

```
show pim {detail | rp-set {<group_addr>} | vlan <vlan_name>}
```

Description

Displays the PIM configuration and statistics.

Syntax Description

detail	Specifies to display the detailed format.
group_addr	Specifies an IP multicast group, for which the RP is to be displayed.
vlan_name	Specifies a VLAN name.

Default

If no VLAN is specified, the configuration is displayed for all PIM interfaces.

If no multicast group is specified for the `rp-set` option (Rendezvous Point set), all RPs are displayed.

Usage Guidelines

The `detail` version of this command displays the global statistics for PIM, as well as the details of each PIM enabled VLAN.

Examples

The following command displays the global PIM configuration and statistics:

```
Switch. 14 # show pim
```

```
PIM Enabled, Version 2
PIM CRP Disabled
BSR state           : ACCEPT_ANY ; BSR Hash Mask Length: 255.255.255.252
Current BSR Info    : 0.0.0.0 (Priority 0)
Configured BSR Info : 0.0.0.0 (Priority 0)
CRP Adv Interval    : 60 sec ; CRP Holdtime: 150
BSR Interval        : 60 sec ; BSR Timeout : 130
Cache Timer         : 210 sec ; Prune Timer : 210
Assert Timeout      : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id       : 0x4a01a1a6
PIM-DM State Refresh Source Active Timer : 210 sec
PIM-DM State Refresh TTL                   : 16
PIM-DM State Refresh Origination Interval: 60 sec
Threshold for Last Hop Routers: 0 kbps
Threshold for RP                          : 0 kbps
Register-Rate-Limit-Interval : Always active
PIM SSM address range                  : None
Register Checksum to include data
Active Sparse Ckts 0 Dense Ckts 2
```

Global Packet Statistics (In/Out)

C-RP-Advs	0	0
Registers	0	0
RegisterStops	0	0

VLAN	Cid	IP Address	Designated Router	Flags	Hello Int	J/P Int	Nbrs
v36	2	36.36.36.3	/16 36.36.36.6	rifmd-----R	30	60	1
vixia	2	64.1.1.1	/16 64.1.1.1	rifmd-----R	30	60	0

Legend: J/P Int: Join/Prune Interval

Flags : r - Router PIM Enabled, i - Interface PIM Enabled, f - Interface,
Forwarding Enabled, m - Interface Multicast Forwarding Enabled,
s - Sparse mode, d - Dense mode, c - CRP enabled,
t - Trusted Gateway configured, n - Multinetted VLAN,
p - Passive Mode, S - Source Specific Multicast, b - Border,
R - State Refresh Enabled.

The following command displays the detailed PIM configuration and statistics:

```
Switch.22 # show pim detail
```

PIM Enabled, Version 2

PIM CRP Disabled

BSR state : ACCEPT_ANY ; BSR Hash Mask Length: 255.255.255.252

Current BSR Info : 0.0.0.0 (Priority 0)

Configured BSR Info : 0.0.0.0 (Priority 0)

CRP Adv Interval : 60 sec ; CRP Holdtime: 150

BSR Interval : 60 sec ; BSR Timeout : 130

Cache Timer : 210 sec ; Prune Timer : 210

Assert Timeout : 210 sec ; Register Suppression Timeout,Probe: 60, 5

Generation Id : 0x4a01a1a6

PIM-DM State Refresh Source Active Timer : 210 sec

PIM-DM State Refresh TTL : 16

PIM-DM State Refresh Origination Interval: 60 sec

Threshold for Last Hop Routers: 0 kbps

Threshold for RP : 0 kbps

Register-Rate-Limit-Interval : Always active

PIM SSM address range : None

Register Checksum to include data

Active Sparse Ckts 0 Dense Ckts 1

Global Packet Statistics (In/Out)

C-RP-Advs	0	0
-----------	---	---

Registers	0	0
-----------	---	---

RegisterStops	0	0
---------------	---	---

PIM DENSE Interface[2] on VLAN v36 is enabled and up

IP adr: 36.36.36.3 mask: 255.255.0.0 DR of the net: 36.36.36.6

Passive : No

Hello Interval : 30 sec

Neighbor Time out : 105 sec

Join/Prune Interval : 60 sec

Join/Prune holdtime : 210 sec

Trusted Gateway : none

CRP group List : none with priority 0

Shutdown priority : 1024

```

Source Specific Multicast : Disabled
State Refresh             : On
State Refresh Capable     : Yes
Border                    : No
Neighbor IP address       Generation Id   Expires State Refresh
36.36.36.6                0x4a01a39d   90      On

```

Packet Statistics (In/Out)

```

Hellos          20          20  Bootstraps          0          0
Join/Prunes     0          0  Asserts           0          0
Grafts          0          0  GraftAcks         0          0
State Refresh   0          0

```

The following command displays the elected, active RP for the group 239.255.255.1:

```
show pim rp-set 239.255.255.1
```

```

Group          Mask          C-RP          Origin      Priority
224.0.0.0      240.0.0.0      10.10.10.2    Bootstrap 0
224.0.0.0      240.0.0.0      124.124.124.124 Bootstrap 0
224.0.0.0      240.0.0.0      124.124.124.124 static    0
239.255.255.0  255.255.255.0  124.124.124.124 Bootstrap 0

```

Elected RP is 124.124.124.124

The following command displays the PIM configuration for VLAN v36:

```
Switch.12 # show pim vlan v36
```

```

PIM DENSE Interface[2] on VLAN v36 is enabled and up
IP adr: 36.36.36.3   mask: 255.255.0.0   DR of the net: 36.36.36.6
Passive              : No
Hello Interval       : 30 sec
Neighbor Time out    : 105 sec
Join/Prune Interval  : 60 sec
Join/Prune holdtime  : 210 sec
Trusted Gateway      : none
CRP group List       : none with priority 0
Shutdown priority    : 1024
Source Specific Multicast : Disabled
State Refresh        : On
State Refresh Capable : Yes
Border               : No
Neighbor IP address   Generation Id   Expires StateRefresh
36.36.36.6            0x4a01a39d   95      On

```

Packet Statistics (In/Out)

```

Hellos          33          32  Bootstraps          0          0
Join/Prunes     0          0  Asserts           0          0
Grafts          0          0  GraftAcks         0          0

```

State Refresh 0 0

show pim cache

```
show pim cache [{detail} | {state-refresh}]{<group_addr> {<source_addr>}}
```

Description

Displays the multicast cache entries created by PIM.

Syntax Description

detail	Specifies to display the information in detailed format.
group_addr	Specifies an IP group address.
source_addr	Specifies an IP source address.
state-refresh	Specifies to display the PIM cache entries with state refresh parameters.

Default

N/A.

Usage Guidelines

Displays the following information:

- IP group address
- IP source address / source mask
- Upstream neighbor (RPF neighbor)
- Interface (VLAN) to upstream neighbor
- Cache expire time
- Egress and prune interface list

When the detail option is specified, the switch displays the egress VLAN list and the pruned VLAN list.

Examples

The following command displays the PIM cache entry for group 239.255.255.1:

```
Switch.33 # show pim cache 239.255.255.1
```

```

Index  Dest Group      Source                InVlan  Origin
[0000] 239.255.255.1 124.124.124.124 (WR) v4    Sparse
      Entry timer is not run; UpstNbr: 200.124.124.24
      EgressIfList = vbs15(0)(FW)(SM)(I)

[0001] 239.255.255.1 118.5.1.1 (S)         vbs5    Sparse
```

```
Expires after 186 secs UpstNbr: 0.0.0.0
RP: 124.124.124.124 via 200.124.124.24 in v4
EgressIfList = vbs15(0)(FW)(SM)(I) , vpim5(170)(FW)(SM)(S)
PrunedIfList = v4(0)(SM)
```

Number of multicast cache = 20

Entry flags :-

R: RP tree. S: Source tree. W: Any source.

Egress/Pruned interface flags :-

```
SM: Sparse Mode          DM: Dense Mode
Fw: Forwarding          PP: Prune pending
AL: Assert Loser        N: Neighbor present
I: IGMP member present  S: (s,g) join received
Z: (*,g) join received  Y: (*,*,rp) join received
```

The following command displays the PIM-DM cache entry with state-refresh information for group 225.0.0.1:

```
Switch.5 # show pim cache state-refresh 225.0.0.1
```

```
Index  Dest Group      Source              InVlan  Origin
[0001] 225.0.0.1      64.1.1.100 (S)     vixia   Dense   Not Pruned
Expires after 204 secs UpstNbr: 0.0.0.0
Refresh State: Originator(20), TTL: 16
EgressIfList = v36(0)(FW)(DM)(N)
[0001] 225.0.0.1      65.1.1.100 (S)     vixia   Dense   Not Pruned
Expires after 195 secs UpstNbr: 65.1.1.200
Refresh State: Not-Originator(25), TTL: 8
EgressIfList = v36(0)(FW)(DM)(N)
```

show pim snooping

```
show pim snooping {vlan} <name>
```

Description

Displays the PIM snooping configuration for a VLAN.

Syntax Description

name	Specifies a VLAN.
------	-------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command displays the PIM snooping configuration for the default VLAN:

```
BD-8810Rack3.8 # show pim snooping default
Global PIM Snooping DISABLED
Default          Snooping DISABLED
```

unconfigure igmp

```
unconfigure igmp
```

Description

Resets all IGMP settings to their default values and clears the IGMP group table.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all IGMP settings to their default values and clears the IGMP group table:

```
unconfigure igmp
```

unconfigure igmp snooping vlan ports set join-limit

```
unconfigure igmp snooping {vlan} <vlanname> ports <portlist> set join-limit
```

Description

Removes the join limit set on VLAN ports.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports.

Default

No limit.

Usage Guidelines

None.

Example

The following command removes the join limit for port 2:1 in the Default VLAN:

```
unconfigure igmp snooping "Default" ports 2:1 set join-limit
```

unconfigure igmp ssm-map

```
unconfigure igmp ssm-map {<vr <vr-name>}
```

Description

Unconfigures all SSM mappings on the virtual router.

Syntax Description

vr-name	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.
---------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes all IGMP-SSM mappings on the virtual router xyz:

```
unconfigure igmp ssm-map vr xyz
```

unconfigure pim

```
unconfigure pim {vlan <vlan_name>}
```

Description

Resets all PIM settings on one or all VLANs to their default values.

Syntax Description

<code>vlan_name</code>	Specifies the VLAN from which PIM is to be unconfigured.
------------------------	--

Default

If no VLAN is specified, the configuration is reset for all PIM interfaces.

Usage Guidelines

If you unconfigure PIM, you also unconfigure PIM-SSM, removing the PIM-SSM range.

Example

The following command resets all PIM settings on the VLAN *accounting*:

```
unconfigure pim vlan accounting
```

unconfigure pim ssm range

```
unconfigure pim ssm range
```

Description

Unconfigures the range of multicast addresses for PIM SSM.

Syntax Description

This command has no arguments or variables.

Default

By default, no SSM range is configured.

Usage Guidelines

You must disable PIM before configuring or unconfiguring a PIM-SSM range. Use the `disable pim` command.

Initially, no range is configured for SSM. After a range is configured, you can remove the range with the `unconfigure pim ssm range` command.

When no range is configured for PIM SSM, the switch does not use PIM SSM for any multicast groups.

Example

The following command removes the PIM SSM range:

```
unconfigure pim ssm range
```

This chapter describes commands for doing the following:

- Configuring IPv6 multicast routing
- Displaying IPv6 multicast information

For an introduction to the IPv6 multicast feature, see the *NETGEAR 8800 User Manual*.

clear mld group

```
clear mld group {<v6grpipaddress>} {{vlan} <name>}
```

Description

Removes one or all MLD groups.

Syntax Description

name	Specifies a VLAN name.
v6grpipaddress	Specifies the group IP address.

Default

N/A.

Usage Guidelines

This command is used to manually remove learned MLD group entries instantly.

Example

The following command clears all MLD groups from VLAN *accounting*:

```
clear mld group accounting
```

clear mld snooping

```
clear mld snooping {{vlan} <name>}
```

Description

Removes one or all MLD snooping entries.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove MLD snooping entries instantly. However, removing an MLD snooping entry can disrupt the normal forwarding of multicast traffic, until the snooping entries are learned again.

The static and dynamic MLD snooping entries are removed, then recreated upon the next general query. The static router entry is removed and recreated immediately.

Example

The following command clears MLD snooping from VLAN *accounting*:

```
clear mld snooping accounting
```

configure mld

```
configure mld <query_interval> <query_response_interval> <last_member_query_interval>
{<robustness>}
```

Description

Configures the Multicast Listener Discovery (MLD) timers.

Syntax Description

query_interval	Specifies the interval (in seconds) between general queries.
query_response_interval	Specifies the maximum query response time (in seconds).
last_member_query_interval	Specifies the maximum group-specific query response time (in seconds).
robustness	Specifies the degree of robustness for the network.

Default

- query interval—125 seconds
- query response interval—10 seconds

- last member query interval—1 second
- robustness—2

Usage Guidelines

Timers are based on RFC2710. Specify the following:

- query interval—The amount of time, in seconds, the system waits between sending out general queries. The range is 1 to 429,496,729 seconds.
- query response interval—The maximum response time inserted into the periodic general queries. The range is 1 to 25 seconds.
- last member query interval—The maximum response time inserted into a group-specific query sent in response to a leave group message. The range is 1 to 25 seconds.
- robustness—The degree of robustness of the network. The range is 2 to 7.

Example

The following command configures the MLD timers:

```
configure mld 100 5 1 3
```

configure mld snooping vlan ports add static group

```
configure mld snooping {vlan} <vlanname> ports <portlist> add static group <v6grpipaddress>
```

Description

Configures VLAN ports to receive the traffic from a multicast group, even if no MLD joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. In the form 1, 2, 3-5, 2:5, 2:6-2:8.
v6grpipaddress	Specifies the multicast group IPv6 address.

Default

None.

Usage Guidelines

Use this command to forward a particular multicast group to VLAN ports. In effect, this command emulates a host on the port that has joined the multicast group. As long as the port is configured with the static entry, multicast traffic for that multicast group is forwarded to that port.

The switch sends proxy MLD messages in place of those generated by a real host. The proxy messages use the VLAN IPv6 address for source address of the messages. If the VLAN has no IPv6 address assigned, the proxy MLD message uses 0::0 as the source IP address.

Note: In the current implementation, multicast traffic is flooded to the VLAN.

Example

The following command configures a static MLD entry so the multicast group ff02::1:1 is forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add static group ff02::1:1
```

configure mld snooping vlan ports delete static group

```
configure mld snooping {vlan} <vlanname> ports <portlist> delete static group [all | <v6grpipaddress>]
```

Description

Removes the configuration that causes VLAN ports to receive the traffic from a multicast group, even if no MLD joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. In the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all multicast groups.
v6grpipaddress	Specifies the multicast group IPv6 address.

Default

None.

Usage Guidelines

Use this command to delete a static group from a particular VLAN port.

To add a static group, use the following command:

```
configure mld snooping {vlan} <vlanname> ports <portlist> add static group <v6grpipaddress>
```

Example

The following command removes a static MLD entry so the multicast group ff02::a:b is not forwarded to VLAN *marketing* on ports 2:1-2:4, unless an MLD join message is received on the port:

```
configure mld snooping marketing ports 2:1-2:4 delete static group ff02::a:b
```

configure mld snooping vlan ports add static router

```
configure mld snooping {vlan} <vlanname> ports <portlist> add static router
```

Description

Configures VLAN ports to forward the traffic from all multicast groups, even if no MLD joins have been received on the port.

Syntax Description

vlanname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to forward all multicast groups to the specified VLAN ports. In effect, this command emulates a multicast router attached to those ports. As long as the ports are configured with the static entry, all available multicast traffic is forwarded to those ports.

Example

The following command configures a static MLD entry so all multicast groups are forwarded to VLAN *marketing* on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add static router
```

configure mld snooping vlan ports delete static router

```
configure mld snooping {vlan} <vlanname> ports <portlist> delete static router
```

Description

Configures VLAN ports to stop forwarding the traffic from all multicast groups, unless MLD joins have been received on the port.

Syntax Description

vlannname	Specifies a VLAN name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to remove the configuration that forwards all multicast groups to the specified VLAN ports.

Example

The following command removes a static MLD entry so all multicast groups are not forwarded to VLAN *marketing* on ports 2:1-2:4, unless an MLD join is received on the port:

```
configure mld snooping marketing ports 2:1-2:4 delete static router
```

configure mld snooping flood-list

```
configure mld snooping flood-list [<policy> | none]
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

policy	Specifies a policy file with a list of multicast addresses to be handled.
none	Specifies no policy file is to be used.

Default

None.

Usage Guidelines

With this command, a user can configure certain multicast addresses to be slow path flooded within the VLAN, instead of fast path forwarded according to MLD and/or Layer 3 multicast protocol.

A policy file is a text file with the extension .pol. It can be created or edited with any text editor. The specified policy file `<policy file>` should contain a list of addresses that

determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with a destination address which is in the <policy file> in 'permit' mode, that stream is software flooded and no hardware entry is installed.

When adding an IPv6 address into the policy file, a 128-bit host address is recommended.

This feature is meant to solve the multicast connectivity problem for unknown destination addresses within system reserved ranges. Specifically this feature was introduced to solve the problem of recognizing a certain stream as control packets.

To create a policy file for the snooping flood-list, use the following template:

```
# This is a template for MLD Snooping Flood-list Policy File
# Add your group addresses between "Start" and "End"
# Do not touch rest of file!!!!
entry mldFlood {
    if match any {
#----- Start of group addresses -----
        nlri ff05::100:1/128;
        nlri ff05::100:15/128;
#----- end of group addresses -----
    } then {
        permit;
    }
}

entry catch_all {
    if {

    } then {
        deny;
    }
}
```

Note: The switch does not validate any IP address in the policy file used in this command. Therefore, slow-path flooding should be used only for streams which are very infrequent, such as control packets. It should not be used for multicast data packets. This option overrides any default mechanism of hardware forwarding (with respect to MLD or PIM) so it should be used with caution.

Slow path flooding occurs within the L2 VLAN only.

Use the `none` option to effectively disable slow path flooding.

You can use the `show mld` command to see the configuration of slow path flooding.

Note: This command has no effect in the current release, since IPv6 multicast traffic floods on all platforms.

Example

The following command configures the multicast data stream specified in `access1` for slow path flooding:

```
configure mld snooping flood-list access1
```

The following command specifies that no policy file is to be used, this effectively disabling slow path flooding:

```
configure mld snooping flood-list none
```

configure mld snooping leave-timeout

```
configure mld snooping leave-timeout <leave_timeout_ms>
```

Description

Configures the MLD snooping leave timeout.

Syntax Description

<code>leave_timeout_ms</code>	Specifies an MLD leave timeout value in milliseconds, upon receiving an MLD done message.
-------------------------------	---

Default

1000 ms.

Usage Guidelines

The range is 0 - 10000 ms (10 seconds). For timeout values of one second or less, you must set the leave-timeout to a multiple of 100 ms. For values of more than one second, you must set the leave-timeout to a multiple of 1000 ms (one second).

The specified time is the maximum leave timeout value. The switch could leave sooner if an MLD done message is received before the timeout occurs.

Example

The following command configures the MLD snooping leave timeout:

```
configure mld snooping leave-timeout 10000
```

configure mld snooping timer

```
configure mld snooping timer <router_timeout> <host_timeout>
```

Description

Configures the MLD snooping timers.

Syntax Description

<code>router_timeout</code>	Specifies the time in seconds before removing a router snooping entry.
<code>host_timeout</code>	Specifies the time in seconds before removing a host's group snooping entry.

Default

The router timeout default setting is 260 seconds. The host timeout setting is 260 seconds.

Usage Guidelines

Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following:

- **router timeout**—The maximum time, in seconds, that a router snooping entry can stay without receiving a router report. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.
- **host timeout**—The maximum time, in seconds, that a group snooping entry can stay without receiving a group report. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.

MLD snooping is a Layer 2 function of the switch. It does not require multicast routing to be enabled. The feature reduces the flooding of IPv6 multicast traffic. On the VLAN, MLD snooping optimizes the usage of network bandwidth and prevents multicast traffic from being flooded to parts of the network that do not need it. The switch does not reduce any IP multicast traffic in the local multicast domain (FF02::x).

MLD snooping is enabled by default on the switch. MLD snooping expects at least one device on every VLAN to periodically generate MLD query messages. Without an MLD querier, the switch eventually stops forwarding IPv6 multicast packets to any port, because the MLD snooping entries times out, based on the value specified in `host timeout`.

Example

The following command configures the MLD snooping timers:

```
configure mld snooping timer 600 600
```

disable mld

```
disable mld {vlan <name>}
```

Description

Disables MLD on a router interface. If no VLAN is specified, MLD is disabled on all router interfaces.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

Enabled.

Usage Guidelines

MLD is a protocol used by an IPv6 host to register its IPv6 multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, hosts respond to the query, and group registration is maintained.

MLD is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of MLD packets. MLD should be enabled when the switch is configured to perform IPv6 unicast or IPv6 multicast routing.

This command disables all MLD versions.

Example

The following command disables MLD on VLAN *accounting*:

```
disable mld vlan accounting
```

disable mld snooping

```
disable mld snooping {forward-mcrouter-only | with-proxy | vlan <name>}
```

Description

Disables MLD snooping.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
with-proxy	Disables the MLD snooping proxy.
name	Specifies a VLAN.

Default

MLD snooping and the with-proxy option are enabled by default, but forward-mcrouter-only option is disabled by default.

Usage Guidelines

If a VLAN is specified, MLD snooping is disabled only on that VLAN, otherwise MLD snooping is disabled on all VLANs.

If the switch is in the `forward-mcrouter-only` mode, then the command `disable mld snooping forward-mcrouter-only` changes the mode so that all multicast traffic is forwarded to any IP router. If not in the `forward-mcrouter-mode`, the command `disable mld snooping forward-mcrouter-only` has no effect.

The with-proxy option can be used for troubleshooting purpose. It should be enabled for normal network operation.

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary MLD done messages so that they are forwarded only when the last member leaves the group.

Example

The following command disables MLD snooping on the VLAN *accounting*:

```
disable mld snooping accounting
```

enable mld

```
enable mld {vlan <vlan name>} {MLDv1 | MLDv2}
```

Description

Enables MLD on a router interface. If no VLAN is specified, MLD is enabled on all router interfaces.

Syntax Description

vlan name	Specifies a VLAN name.
MLDv1	Sets the compatibility mode to MLDv1.
MLDv2	Sets the compatibility mode to MLDv2. This option is not supported in this release.

Default

Enabled, set to MLDv1 compatibility mode.

Usage Guidelines

MLD is a protocol used by an IPv6 host to register its IPv6 multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, IPv6 hosts respond to the query, and group registration is maintained.

MLDv1 is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of MLD packets.

A VLAN must have an IPv6 address to support MLD.

Example

The following command enables MLDv1 on the VLAN *accounting*:

```
enable mld vlan accounting
```

enable mld snooping

```
enable mld snooping {forward-mcrouter-only | vlan <name>}
```

Description

Enables MLD snooping on the switch.

Note: MLD snooping is not supported in this software release.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
name	Specifies a VLAN.

Default

Disabled.

Usage Guidelines

If a VLAN is specified, MLD snooping is enabled only on that VLAN, otherwise MLD snooping is enabled on all VLANs.

A VLAN must have an IPv6 address to support MLD.

Two MLD snooping modes are supported:

- The `forward-mcrouter-only` mode forwards all multicast traffic to the multicast router (that is, the router running PIM).

- When not in the `forward-mcrouter-only` mode, the switch forwards all multicast traffic to any IP router (multicast or not), and any active member port to the local network that has one or more subscribers.

Note: The `forward-mcrouter-only` mode for MLD snooping is enabled/disabled on a switch-wide basis, not on a per-VLAN basis. In other words, all the interfaces enabled for MLD snooping are either in the `forward-mcrouter-only` mode or in the `non-forward-mcrouter-only` mode, and not a mixture of the two modes.

To change the MLD snooping mode from the `non-forward-mcrouter-only` mode to the `forward-mcrouter-only` mode, use the command:

```
enable mld snooping forward-mcrouter-only
```

To change the MLD snooping mode from the `forward-mcrouter-only` mode to the `non-forward-mcrouter-only` mode, use the command:

```
disable mld snooping forward-mcrouter-only
```

Example

The following command enables MLD snooping on the switch:

```
enable mld snooping
```

enable mld snooping with-proxy

```
enable mld snooping with-proxy
```

Description

Enables the MLD snooping proxy. The default setting is enabled.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary MLD leave messages so that they are forwarded only when the last member leaves the group.

This command can be used for troubleshooting purpose. It should be enabled for normal network operation. The command does not alter the snooping setting.

Example

The following command enables the MLD snooping proxy:

```
enable mld snooping with-proxy
```

show mld

```
show mld {vlan} {<vlan name>}
```

Description

This command can be used to display an MLD-related configuration and group information, per VLAN or for the switch as a whole.

Syntax Description

vlan name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

If you do not specify a VLAN, the command displays the switch configuration.

Example

The following command displays the MLD configuration:

```
show mld
```

The following is sample output from this command:

```
MLD:
  Query Interval: 125 sec
  Max Response Time: 10 sec
  Last Member Query: 1 sec
  Robustness: 2

MLD Snooping:
  Router Timeout: 260 sec
  Host Timeout: 260 sec
  MLD Snooping Fast Leave Time: 1000 ms
  MLD Snooping Flag: forward-all-router
  MLD Snooping Flood-list: none
  MLD Snooping Proxy: Enable
```

VLAN	IP Address	Flags	nLRMA	nLeMA	MLDver
Default	::/0	U--iz	0	0	1
loopb	::/0	U--iz	0	0	1
red	fe80::2e0:81ff:fe22:5724/64	U--iz	0	0	2
yellow	::/0	U--iz	0	0	1

Flags: (E) Interface Enabled, (i) MLD Enabled
 (f) Forwarding Enabled, (m) Multicast Forwarding Enabled
 (nLeMA) Number of Learned Multicast Addressess
 (nLRMA) Number of Locally registered Multicast Addresses
 (U) Interface Up, (z) MLD Snooping Enabled

The following command displays the MLD configuration for the VLAN *red*:

```
show mld red
```

The following is sample output from this command:

```
Interface on VLAN red is enabled and up.
inet6 fe80::2e0:81ff:fe22:5724/64
Locally registered multicast addresses:

Learned multicast addresses (Last Querier=fe80::2e0:81ff:fe22:5724):

s = static igmp member

Flags:
  IP Fwding  NO      IPmc Fwding  NO      MLD YES
  MLD Ver   v1      Snooping YES
```

show mld group

```
show mld group {{vlan} {<name>} | {<v6grpaddress>}} {MLDv2}
```

Description

Lists the MLD group membership for the specified VLAN or group.

Syntax Description

v6

grpaddress	Specifies a group IPv6 address.
name	Specifies a VLAN name.
MLDv2	Display the MLD group in MLDv2 format (if group record is MLDv2 compatible, otherwise display in earlier format). This option is not supported in this release.

Default

MLDv1.

Usage Guidelines

If no VLAN is specified all VLANs are displayed. You can also filter the display by group address and by multicast stream sender address.

Example

The following command lists the MLD group membership for the VLAN *accounting*:

```
show mld group vtest3
```

Output from this command looks similar to the following:

Group Address	Ver	Vlan	Port	Age
ff03::1:1	1	vtest3	4:5	25
ff03::1:2	1	vtest3	4:5	25
ff02::1:ff22:124	1	vtest3	4:45	26
ff05::a:abcd	1	vtest3	4:15	23
ff05::a:abce	1	vtest3	4:15	23
ff02::1:ff22:112	1	vtest3	4:45	26
ff02::1:ff1f:a418	1	vtest3	4:45	26

show mld snooping

```
show mld snooping {vlan <name> | detail} {MLDv2}
```

Description

Displays MLD snooping registration information and a summary of all MLD timers and states.

Note: MLD snooping is not supported in this software release.

Syntax Description

name	Specifies a VLAN name.
detail	Displays the information in detailed format.
MLDv2	Display the MLD group in MLDv2 format (if group record is MLDv2 compatible, otherwise display in earlier format). This option is not supported in this release.

Default

MLDv1.

Usage Guidelines

The two types of MLD snooping entries are sender entry and subscribed entry.

The following information is displayed in a per-interface format:

- Group membership information
- Router entry
- Timeout information
- Sender entry

Example

The following command displays MLD snooping registration information for the VLAN *red*:

```
show mld snooping vlan test3
```

Output from this command looks similar to the following:

```
VLAN vtest3          (133) Snooping=Enabled
```

Port	Host	Subscribed	Age
4:5	fe80::200:ff:fe00:1	ff03::1:1	10
4:5	fe80::200:ff:fe00:1	ff03::1:2	10
4:15	fe80::200:ff:fe00:201	ff05::a:abcd	8
4:15	fe80::200:ff:fe00:201	ff05::a:abce	8
4:45	fe80::204:96ff:fe1f:a418	ff02::1:ff1f:a418	11
4:45	fe80::204:96ff:fe1f:a418	ff02::1:ff22:112	11
4:45	fe80::201:30ff:fef9:9b90	ff02::1:ff22:124	11
4:45	fe80::201:30ff:fef9:9b90	All Groups	11

s = static igmp member

show mld snooping vlan static

```
show mld snooping vlan <name> static [group | router]
```

Description

Displays static MLD snooping entries.

Syntax Description

name	Specifies a VLAN name.
------	------------------------

Default

None.

Usage Guidelines

Use this command to display the MLD snooping static groups or router ports configured on the specified VLAN. When no VLAN is specified, all the static groups or router ports are displayed.

Example

The following command displays the MLD snooping static groups configured on VLAN *vlan101*:

```
show mld snooping vlan101 static group
```

The following is sample output for this command:

Group	Port	Flags
ff03::1:1:1	7	sa
ff03::1:1:1	15	sa

Flags: (s) Static, (a) Active

unconfigure mld

```
unconfigure mld
```

Description

Resets all MLD settings to their default values and clears the MLD group table.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all MLD settings to their default values and clears the MLD group table:

```
unconfigure mld
```

This chapter describes commands for doing the following:

- Configuring MSDP
- Displaying MSDP information

For an introduction to the MSDP feature, see the *NETGEAR 8800 User Manual*.

clear msdp counters

```
clear msdp counters {peer <remoteaddr> | peer all | system} {vr <vrname>}
```

Description

This command resets the MSDP counters to zero.

Syntax Description

peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
system	Clears the global MSDP counters.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The `clear msdp counters` command clears the following MSDP counters:

- Per peer counters
 - Number of SA messages received
 - Number of SA messages transmitted
 - Number of SA request messages received

- Number of SA request messages transmitted
- Number of SA response messages received
- Number of SA response messages transmitted
- Number of SA messages received without encapsulated data
- Number of SA messages transmitted without encapsulated data
- Number of SA messages received with encapsulated data
- Number of SA messages transmitted with encapsulated data
- Number of times the MSDP peer attained an “ESTABLISHED” state
- Number of times the peer-RPF check failed
- Number of times the TCP connection attempt failed
- Total number of received messages
- Total number of transmitted messages
- Global counters
 - None defined

The `clear counters` command will also clear all MSDP counters, but it clears the counters for all other applications too.

Example

The following command clears the counters for an MSDP peer with the IP address 192.168.45.43:

```
clear msdp counters peer 192.168.45.43
```

The following command clears the all peer and global counters:

```
clear msdp counters
```

The following command clears all counters for a particular peer:

```
clear msdp counters peer 192.168.32.45
```

The following command clears the counters of all MSDP peers:

```
clear msdp counters peer all
```

The following command clears the global counters:

```
clear msdp counters system
```

clear msdp sa-cache

```
clear msdp sa-cache [{peer} <remoteaddr> | peer all] {group-address <grp-addr>} {vr <vrname>}
```

Description

This command purges all SA cache entries and notifies the PIM that the SA cache is empty.

Syntax Description

grp-addr	Specifies the IP address and subnet mask of the multicast group you want to clear. All SA cache entries that match the specified group address are removed from the database.
peer all	Specifies all MSDP peers. All matching SA cache entries from all peers are removed from the database.
remoteaddr	Specifies the IP address of the MSDP peer. All matching SA cache entries learned from the specified peer are removed from the database.
vname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

MSDP receives SA messages periodically. So, after clearing SA cache entries from the local database, MSDP relearns those entries during the next advertisement from its peer.

Example

The following command clears SA cache records for an MSDP peer with the IP address 192.168.45.43:

```
clear msdp sa-cache peer 192.168.45.43
```

configure msdp as-display-format

```
configure msdp as-display-format [asdot | asplain]
```

Description

Configures the AS number format displayed in `show` commands.

Syntax Description

asdot	Specifies the ASDOT format.
asplain	Specifies the ASPLAIN format.

Default

N/A.

Usage Guidelines

The ASPLAIN and ASDOT formats are described in *RFC 5396, Textual Representation of Autonomous System (AS) Numbers*.

Examples

The following command selects the ASDOT 4-byte AS number format:

```
configure msdp as-display-format asdot
```

configure msdp max-rejected-cache

```
configure msdp max-rejected-cache <max-cache> {vr <vrname>}
```

Description

Configures the maximum limit on rejected SA cache entries that an MSDP router will store in its database.

Syntax Description

max-cache	Specifies the maximum number of rejected SA cache entries that the MSDP router will store in its database. To remove the limit, enter 0 (zero) for the <max-cache> value.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the maximum cache entries stored is zero. That is, rejected SA cache entries are not stored. Any SA cache entries that are stored and not refreshed for six minutes are removed.

Usage Guidelines

SA cache are rejected because of:

- Peer-RPF failure
- Policy denied

When a previously rejected SA cache entry is accepted because of an RP reachability change or policy rule change, the rejected SA cache entry is moved to the accepted SA cache list.

By default, rejected SA cache entries are discarded. You can configure a limit for rejected cache entries to store them, which will help debug/diagnose some issues; however, it consumes extra memory.

Example

The following command sets the maximum rejected cache limit to 100 for an MSDP router:

```
configure mdp max-rejected-cache 100
```

configure mdp originator-id

```
configure mdp originator-id <ip-address> {vr <vrname>}
```

Description

Configures the originator ID for an MSDP router. The originator ID is the RP address you want to use (instead of the default) in locally originated SA messages.

Syntax Description

ip-address	Specifies the RP address to use in locally originated SA messages. To unconfigure an originator ID (that is, to use the default RP address), enter the IP address 0.0.0.0.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the RP address is used as the originator ID in locally originated SA messages.

Usage Guidelines

Use this command to override the default RP address used in SA messages. Because only RPs and MSDP border routers originate SAs, there are times when it is necessary to change the ID used for this purpose. The originator ID address must be one of the interface addresses on the MSDP router.

You can configure the MSDP originator ID only when MSDP is disabled globally.

To remove an originator ID, enter the IP address 0.0.0.0.

Example

The following command configures the originator ID for an MSDP router:

```
configure mdp originator-id 10.203.134.1
```

The following command unconfigures the originator ID for an MSDP router:

```
configure mdp originator-id 0.0.0.0
```

configure mdp peer default-peer

```
configure mdp peer [<remoteaddr> | all] default-peer {default-peer-policy <filter-name>} {vr <vrname>}
```

Description

This command configures a default or static RPF peer from which all MSDP SA messages are accepted. To remove the default peer, enter the `configure msdp peer no-default-peer` command.

Syntax Description

filter-name	Specifies the name of the policy filter associated with the default peer. The peer will be the default peer for all SA entries that are permitted by the policy filter. If an SA message is allowed by the policy filter, it will be accepted. Otherwise, the SA message has to go through the regular RPF-check. The static peer RPF check is the last step in peer RPF algorithm. So, if an SA message is denied by the default peer policy, ultimately the SA message will be rejected by MSDP.
peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no static RPF peer is configured.

The “default-peer-policy” keyword specifies the name of the policy filter associated with the default peer. You can configure multiple default peers with different policies. If no policy is specified, then the current peer is the default RPF peer for all SA messages.

Usage Guidelines

Configuring a default peer simplifies peer-RPF checking of SA messages. If the peer-RPF check fails, the default peer rule is applied to see if the SA messages should be accepted or rejected.

If a default peer policy is specified, the peer is the default peer only for the (Source, Group), or (S, G), that satisfies the policy. If the policy is not specified, then the default peer is used for all (S, G, RP).

You can configure multiple default peers on an MSDP router; however all default peers must either have a default policy or not. A mix of default peers, with a policy and without a policy, is not allowed.

When configuring multiple default peer rules, follow these guidelines:

- When you enter multiple `default-peer` commands with the `default-peer-policy` keyword, you can use all the default peers at the same time for different RP prefixes.
- When you enter multiple `default-peer` commands without the `default-peer-policy` keyword, you can use a single active peer to accept all SA messages. If that peer goes down, then the next configured default peer accepts all SA messages. This configuration is typically used at a stub site.

You can use the following policy attributes in a default peer policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following command configures an MSDP peer with the IP address 192.168.45.43 as the default peer policy for “sales”:

```
configure msdp peer 192.168.45.43 default-peer default-peer-policy sales
```

configure msdp peer description

```
configure msdp peer <remoteaddr> description {<peer-description>} {vr <vrname>}
```

Description

Configures a name or description for an MSDP peer. This text is for display purposes only.

Syntax Description

remoteaddr	Specifies the IP address of the MSDP peer.
peer-description	Specifies the name or description of the MSDP peer. The maximum is 63 characters.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no name or description is specified.

Usage Guidelines

Use this command to configure a name or description to make an MSDP peer easier to identify. The description is visible in the output of the `show msdp peer` command.

To remove the description, use this command without a description string.

Example

The following command configures the name “internal_peer” to an MSDP peer:

```
configure msdp peer 192.168.45.43 description internal_peer
```

The following command removes the description from an MSDP peer:

```
configure msdp peer 192.168.45.43 description
```

configure msdp peer mesh-group

```
configure msdp peer [<remoteaddr> | all] mesh-group [<mesh-group-name> | none] {vr <vrname>}
```

Description

This command configures an MSDP peer to become a member of a mesh-group. To remove a peer from a mesh-group, enter the `none` CLI keyword for the mesh-group.

Syntax Description

<code>mesh-group-name</code>	Specifies the name of the MSDP mesh-group.
<code>none</code>	Removes a peer from a mesh-group.
<code>peer all</code>	Specifies all MSDP peers.
<code>remoteaddr</code>	Specifies the IP address of the MSDP peer.
<code>vrname</code>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Any SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group.

Mesh-groups achieve two goals:

- Reduce SA message flooding
- Simplify peer-RPF flooding

Example

The following command configures an MSDP peer with the IP address 192.168.45.43 to become a member of a mesh-group called “intra”:

```
configure msdp peer 192.168.45.43 mesh-group intra
```

configure msdp peer no-default-peer

```
configure msdp peer [<remoteaddr> | all] no-default-peer {vr <vrname>}
```

Description

This command removes a default peer.

Syntax Description

no-default-peer	Removes a default peer.
peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command removes all MSDP peers:

```
configure msdp peer all no-default-peer
```

configure msdp peer password

```
configure msdp peer [<remoteaddr> | all] password [none | {encrypted} <tcpPassword>] {vr  
<vrname>}
```

Description

This command configures a TCP MD5 password for an MSDP peer. This command enables TCP MD5 authentication for a MSDP peer. When a password is configured, MSDP receives only authenticated MSDP messages from its peers. All MSDP messages that fail TCP MD5 authentication are dropped.

Syntax Description

encrypted	Encrypts the password for MD5 authentication. To improve security, the password displays in encrypted format and cannot be seen as simple text. Additionally, the password is saved in encrypted format.
none	Removes the previously configured password.
peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.

tcpPassword	Specifies the password to use for MD5 authentication at the TCP level. The password must be an ASCII string with a maximum of 31 characters.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Defaults

By default, TCP MD5 authentication is disabled for the MSDP peer.

Usage Guidelines

NETGEAR recommends that you enable TCP MD5 authentication for all MSDP peers to protect MSDP sessions from attacks. You can execute this command only when the MSDP peer is disabled or when MSDP is globally disabled on that VR.

Example

The following command configures a password for the MSDP peer with the IP address 192.168.45.43, which automatically enables TCP MD5 authentication:

```
configure msdp peer 192.168.45.43 password test123
```

The following command removes the password:

```
configure msdp peer 192.168.45.43 password none
```

configure msdp peer sa-filter

```
configure msdp peer [<remoteaddr> | all] sa-filter [in | out] [<filter-name> | none] {vr
<vrname>}
```

Description

This command configures an incoming or outgoing policy filter for SA messages.

Syntax Description

filter-name	Specifies the name of the policy associated with an SA filter. To remove an SA filter, enter the "none" CLI keyword for <filter-name>.
in	Associates the SA filter with inbound SA messages.
out	Associates the SA filter with outbound SA messages.
peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no SA filter is configured for an MSDP peer. That is, incoming and outgoing SA messages are not filtered.

Usage Guidelines

This command configures an SA filter such that only a specified set of SA messages are accepted or sent to a peer. Note that an SA filter does not adversely impact the flow of SA request and response messages.

To remove an SA filter, enter the “none” CLI keyword for <filter-name>.

You can use the following policy attributes in an SA filter policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following command configures an incoming SA messages filter on an MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 sa-filter in allow_229
```

configure msdp peer sa-limit

```
configure msdp peer [<remoteaddr> | all] sa-limit <max-sa> {vr <vrname>}
```

Description

This command allows you to limit the number of SA entries from an MSDP peer that the router will allow in the SA cache. To allow an unlimited number of SA entries, use 0 (zero) as the value for <max-sa>.

Syntax Description

max-sa	Specifies the maximum number of SA entries from an MSDP peer allowed in the SA cache. To specify an unlimited number of SA entries, use 0 (zero) as the value for <max-sa>.
peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.

vname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
-------	---

Default

By default, no SA entry limit is set. The router can receive an unlimited number of SA entries from an MSDP peer.

Usage Guidelines

You can use this command to prevent a distributed denial of service (DOS) attack. NETGEAR recommends that you configure an MSDP SA limit on all MSDP peer sessions. Note that a rejected SA cache entry is not included in the number of SA cache entries received from a peer.

Example

The following command configures the SA entry limit of 500 for the MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 sa-limit 500
```

configure msdp peer source-interface

```
configure msdp peer [<remoteaddr> | all] source-interface [<ipaddress> | any] {vr <vname>}
```

Description

This command configures the source interface for the MSDP peer TCP connection.

Syntax Description

any	Specifies to use any interface as one end of the TCP connection. The source interface is selected based on the IP route entry used to reach the MSDP peer. The egress interface that reaches the MSDP peer is used as the source interface for the TCP connection. Basically, this command removes the previously configured source interface of the MSDP peer.
ipaddress	Specifies the IP address of the MSDP router interface to use on one end of a TCP connection. The <ipaddress> must be one of the MSDP router interface addresses; otherwise, the command fails and an error message displays.
peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
vname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Defaults

By default, the source interface is selected based on the IP route entry used to reach the MSDP peer. The egress interface that reaches the MSDP peer is used as the source interface for the TCP connection.

Usage Guidelines

You must first disable MSDP or the MSDP peer before using this command. NETGEAR recommends that you configure a source interface for MSDP peers that are not directly connected. We also recommend using the loopback address as the MSDP peer connection endpoint.

Example

The following command configures a source interface for an MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 source-interface 60.0.0.5
```

configure msdp peer timer

```
configure msdp peer [<remoteaddr> | all] timer keep-alive <keep-alive-sec> hold-time  
<hold-time-sec> {vr <vrname>}
```

Description

The command configures the keep-alive and hold timer intervals of the MSDP peers.

Syntax Description

hold-time-sec	Specifies the hold timer interval in seconds, in the range of 3 through 75 seconds.
keep-alive-sec	Specifies the keep-alive timer interval in seconds, in the range of 1 through 60 seconds.
peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the:

- Keep-alive timer interval is 60 seconds.
- Hold timer interval is 75 seconds.
- SA timer interval is 60 seconds.

Usage Guidelines

You can use this command only when either MSDP or the MSDP peer is disabled. The hold timer interval must be greater than the keep-alive timer interval.

Example

The following command configures the keep-alive and hold timer intervals for the MSDP peer 55.0.0.83:

```
configure msdp peer 55.0.0.83 timer keep-alive 30 hold-time 60
```

configure msdp peer ttl-threshold

```
configure msdp peer [<remoteaddr> | all] ttl-threshold <ttl> {vr <vrname>}
```

Description

Configures the limit to which multicast data packets are sent in SA messages to an MSDP peer. If the time-to-live (TTL) in the IP header of an encapsulated data packet exceeds the TTL threshold configured, encapsulated data is not forwarded to MSDP peers.

Syntax Description

all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer on which to configure a TTL threshold.
ttl	Specifies the TTL value. The range is 0 through 255. To restore the default value, enter a TTL value of 0 (zero).
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

The default value is zero, meaning all multicast data packets are forwarded to the peer regardless of the TTL value in the IP header of the encapsulated data packet.

Usage Guidelines

This command allows you to configure a TTL value to limit multicast data traffic.

Example

The following command configures a TTL threshold of 5:

```
configure msdp peer 192.168.45.43 ttl-threshold 5
```

configure msdp sa-cache-server

```
configure msdp sa-cache-server <remoteaddr> {vr <vrname>}
```

Description

Configures the MSDP router to send SA request messages to the MSDP peer when a new member becomes active in a group.

Syntax Description

remoteaddr	Specifies the IP address of the MSDP peer from which the local router requests SA messages when a new member becomes active in a group, and MSDP has no cache entry for the group in the local database.
vname	Specifies the name of the virtual router on which the MSDP cache server is configured. If a virtual router name is not specified, it is extracted from the current CLI context.

Default

By default, the router does not send SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member simply waits to receive SA messages, which eventually arrive.

Usage Guidelines

You can use this command to force a new member of a group to learn the current active multicast sources in a connected PIM-SM domain that are sending to a group. The router will send SA request messages to the specified MSDP peer when a new member joins a group and MSDP doesn't have a cache entry for that group in the local database. The peer replies with the information in an SA cache response message.

Note: An MSDP peer must exist before it can be configured as an SA cache server. The `configure msdp sa-cache-server` command accepts the value for `<remoteaddr>` only if it is an existing peer's IP address.

Example

The following command configures an MSDP cache server:

```
configure msdp sa-cache-server 172.19.34.5
```

configure pim border

```
configure pim <vlan_name> border
```

Description

Configures a PIM VLAN as a border VLAN, which is used to demarcate a PIM domain when using MSDP.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

MSDP is used to connect multiple multicast routing domains. A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN.

Example

The following command configures a PIM border on a VLAN called “vlan_border”:

```
configure pim vlan_border border
```

create msdp mesh-group

```
create msdp mesh-group <mesh-group-name> {vr <vrname>}
```

Description

Creates an MSDP mesh-group.

Syntax Description

mesh-group-name	Specifies the name for the MSDP mesh-group.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Create a mesh-group to:

- Reduce SA message flooding
- Simplify peer-RPF flooding

SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group, which reduces SA message flooding.

A mesh group name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see the section “Object Names” in the *NETGEAR 8800 User Manual*.

Example

The following command creates a mesh-group called “verizon”:

```
create msdp mesh-group verizon
```

create msdp peer

```
create msdp peer <remoteaddr> {remote-as <remote-AS>} {vr <vrname>}
```

Description

Creates an MSDP peer.

Syntax Description

remoteaddr	Specifies the IP address of the MSDP router to configure as an MSDP peer.
remote-AS	Specifies the autonomous system (AS) number of the MSDP peer. This optional parameter is deprecated, though the option is still available in the CLI for backward compatibility. The software ignores this parameter.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The BGP route database is used by MSDP to determine the AS number for the peer. You can display the AS number (which can be a 2-byte for 4-byte AS number) using the command:

```
show msdp [peer {detail} | {peer} <remoteaddr>] {vr <vrname>}.
```

Example

The following command creates an MSDP peer:

```
create msdp peer 192.168.45.43 remote-as 65001
```

delete msdp mesh-group

```
delete msdp mesh-group <mesh-group-name> {vr <vrname>}
```

Description

Removes an MSDP mesh-group.

Syntax Description

mesh-group-name	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Mesh-groups are used to achieve two goals:

- Reduce SA message flooding
- Simplify peer-RPF flooding

SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group.

Use the `delete mdp mesh-group` command only if you created a mesh-group that you want to remove. By default, there is no MSDP mesh-group.

Example

The following command removes a mesh-group called “verizon”:

```
delete mdp mesh-group verizon
```

delete mdp peer

```
delete mdp peer [all | <remoteaddr>] {vr <vrname>}
```

Description

Deletes an MSDP peer.

Syntax Description

all	Deletes all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP router to configure as an MSDP peer.

vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
--------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes an MSDP peer:

```
delete msdp peer 192.168.45.43
```

disable msdp

```
disable msdp {vr <vrname>}
```

Description

Disables MSDP on a virtual router.

Syntax Description

vrname	Specifies the name of the virtual router on which MSDP is being enabled or disabled. If a name is not specified, it is extracted from the current CLI context.
--------	--

Default

MSDP is disabled by default.

Usage Guidelines

Use this command to disable MSDP on a virtual router.

Example

The following command disables MSDP on a virtual router:

```
disable msdp
```

disable msdp data-encapsulation

```
disable msdp data-encapsulation {vr <vrname>}
```

Description

Disables the encapsulation of locally originated SA messages with multicast data (if available).

Syntax Description

vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
--------	---

Default

By default, multicast data packet encapsulation is enabled for locally originated SA messages.

Usage Guidelines

N/A.

Example

The following command disables multicast data packet encapsulation:

```
disable msdp data-encapsulation
```

disable msdp export local-sa

```
disable msdp export local-sa {vr <vrname>}
```

Description

Disables the advertisement of local sources to groups for which the router is an RP.

Syntax Description

vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
--------	---

Default

By default, the export of local sources is enabled. All sources are advertised if the router is an RP for the groups. Use this command to disable it.

Usage Guidelines

You can create a policy to filter out some of the local sources so that they are not advertised to MSDP peers and exposed to the external multicast domain. To configure an export filter, you must first disable the export of local sources (with the `disable msdp export local-sa`

command), and then re-enable it with an export filter (with the `enable msdp export local-sa export-filter` command).

Example

The following command disables the advertisement of local sources:

```
disable msdp export local-sa
```

disable msdp peer

```
disable msdp [{peer} <remoteaddr> | peer all] {vr <vrname>}
```

Description

Configures the administrative state of an MSDP peer.

Syntax Description

all	Disables all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer to disable.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, MSDP peers are disabled.

Usage Guidelines

Use this command to administratively disable MSDP peers to stop exchanging SA messages.

Example

The following command disables an MSDP peer:

```
disable msdp peer 192.168.45.43
```

disable msdp process-sa-request

```
disable msdp [{peer} <remoteaddr> | peer all] process-sa-request {vr <vrname>}
```

Description

This command configures a router to reject SA request messages from a specified peer or all peers.

Syntax Description

peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, all SA request messages are accepted from all peers.

Usage Guidelines

Use this command to configure the router to reject SA request messages from a specified peer or all peers.

You cannot change an SA request filter while SA request processing is enabled for an MSDP peer. You must first disable SA request processing for a peer and then re-enable it with an SA request filter.

You can use the following policy attributes in an SA request policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following command disables processing of SA request messages received from a peer with the IP address 192.168.45.43:

```
disable msdp peer 192.168.45.43 process-sa-request
```

enable msdp

```
enable msdp {vr <vrname>}
```

Description

Enables MSDP on a virtual router.

Syntax Description

vname	Specifies the name of the virtual router on which MSDP is being enabled or disabled. If a name is not specified, it is extracted from the current CLI context.
-------	--

Default

MSDP is disabled by default.

Usage Guidelines

Use this command to enable MSDP on a virtual router.

Example

The following command enables MSDP on a virtual router:

```
enable mdp
```

enable mdp data-encapsulation

```
enable mdp data-encapsulation {vr <vname>}
```

Description

Enables the encapsulation of locally originated SA messages with multicast data (if available).

Syntax Description

vname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
-------	---

Default

By default, multicast data packet encapsulation is enabled for locally originated SA messages. Multicast data packets with a packet size of up to 8 KB are encapsulated in SA messages.

Usage Guidelines

Enable data encapsulation to handle bursty sources.

Example

The following command enables multicast data packet encapsulation:

```
enable mdp data-encapsulation
```

enable msdp export local-sa

```
enable msdp export local-sa {export-filter <filter-name>} {vr <vrname>}
```

Description

Enables the advertisement of local sources to groups for which the router is an RP.

Syntax Description

filter-name	Specifies the policy to associate with the export of local sources. No policy is specified by default.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the export of local sources is enabled. All sources are advertised if the router is an RP for the groups.

Usage Guidelines

You can create a policy to filter out some of the local sources so that they are not advertised to MSDP peers and exposed to the external multicast domain. To configure an export filter, you must first disable the export of local sources (with the `disable msdp export local-sa` command), and then re-enable it with an export filter (with the `enable msdp export local-sa export-filter` command).

You can use the following policy attributes in an export policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Please note that the syntax for “multicast-group”, “multicast-source,” and “pim-rp” are the same as for the “nlri” policy attribute.

```
[multicast-group | multicast-source | pim-rp] [<ipaddress> | any]/<mask-length> {exact}
[multicast-group | multicast-source | pim-rp] [<ipaddress> | any] mask <mask> {exact}
```

An example of an MSDP policy file follows:

```
entry allow_internal_rp {
    if match any {
        multicast-group 234.67.89.0/24;
```

```

        multicast-source 23.123.45.0/24;
        pim-rp 10.203.134.5/32;
    } then {
        permit;
    }
}
entry deny_local_group239 {
    if match any {
        multicast-group 239.0.0.0/8;
        multicast-source 23.123.45.0/24;
    } then {
        deny;
    }
}
entry allow_external_rp_172 {
    if {
        multicast-group 234.172.0.0/16;
    } then {
        permit
    }
}
# deny remaining entries

```

Example

The following command enables the advertisement of local sources:

```
enable msdp export local-sa
```

enable msdp peer

```
enable msdp [{peer} <remoteaddr> | peer all] {vr <vrname>}
```

Description

Configures the administrative state of an MSDP peer.

Syntax Description

all	Enables all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer to configure.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, MSDP peers are disabled.

Usage Guidelines

You must use this command to administratively enable the MSDP peers before they can establish peering sessions and start exchanging SA messages.

Example

The following command enables an MSDP peer:

```
enable msdp peer 192.168.45.43
```

enable msdp process-sa-request

```
enable msdp [{peer} <remoteaddr> | peer all] process-sa-request {sa-request-filter  
<filter-name> } {vr <vrname>}
```

Description

This command configures MSDP to receive and process SA request messages from a specified peer or all peers. If an SA request filter is specified, only SA request messages from those groups permitted are accepted. All others are ignored.

Syntax Description

filter-name	Specifies the name of the policy filter associated with SA request processing.
peer all	Specifies all MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, all SA request messages are accepted from peers.

Usage Guidelines

Use this command to configure the router to accept all or just some SA request messages from peers. If no policy is specified, all SA request messages are accepted. If a policy is specified, only SA request messages from those groups permitted are accepted, and all others are ignored.

You cannot change an SA request filter while SA request processing is enabled for an MSDP peer. You must first disable SA request processing for a peer and then re-enable it with an SA request filter.

You can use the following policy attributes in an SA request policy. All other attributes are ignored.

- Match:

- multicast-group
- multicast-source
- pim-rp
- Set:
 - permit
 - deny

Example

The following command enables processing of SA request messages received from a peer with the IP address 192.168.45.43:

```
enable msdp peer 192.168.45.43 process-sa-request sa-request-filter intra_domain
```

show msdp

```
show msdp {vr <vrname>}
```

Description

This command displays global configuration and run-time parameters for MSDP.

Syntax Description

vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
--------	---

Default

N/A.

Usage Guidelines

Use this command to verify the global configuration parameters of MSDP.

Example

The following command displays global configuration and run-time parameters for MSDP:

```
Switch.2 # show msdp
```

MSDP Enabled	: No	VR-Name	: VR-Default
Originator RP Addr	: not configured	SA Cache ageout time	: 360
Store SA Cache	: Yes	SA Cache Server	: not configured
Export Local SAs	: Yes	Export SA filter	: not configured
Max Rejected Cache	: not configured	Encapsulate data	: Yes
Num of Rejected SAs	: 0	Total Num of SAs	: 0
Num of Local SAs	: 0	AS Disp Format	: Asdot

show msdp memory

```
show msdp memory {detail | <memoryType>}
```

Description

This command displays current memory utilization of the MSDP process, including all virtual router instances of the MSDP process.

Syntax Description

detail	Displays detailed statistics for all memory types.
memoryType	Displays statistics for a particular memory type.

Default

N/A.

Usage Guidelines

Use this command to view and diagnose the memory utilization of the MSDP process.

Example

The following displays current memory utilization of the MSDP process, including all virtual router instances of the MSDP process:

```
show msdp memory
```

The following is sample output from this command:

MSDP Memory Information

Bytes Allocated: 79792 AllocFailed: 0 OversizeAlloc: 0

Current Memory Utilization Level: GREEN

Memory Utilization Statistics

Size	16	32	48	64	80	96	128	256	1024	4096	8192	12288

Used Blocks0	0	256	263	3	0	2	0	1	0	0	4	
peer0	0	0	0	0	0	0	0	0	0	0	0	4
mesh-group0	0	0	3	0	0	0	0	0	0	0	0	0
sa-node0	0	0	255	0	0	0	0	0	0	0	0	0
sa-entry0	0	255	0	0	0	0	0	0	0	0	0	0
vr-node0	0	0	0	0	0	0	0	1	0	0	0	0
rt-cache0	0	0	5	0	0	0	0	0	0	0	0	0

rp-node0	0	1	0	0	0	0	0	0	0	0	0
client0	0	0	0	0	0	2	0	0	0	0	0
misc0	0	0	0	3	0	0	0	0	0	0	0

show msdp mesh-group

```
show msdp [mesh-group {detail} | {mesh-group} <mesh-group-name>] {vr <vrname>}
```

Description

This command displays configuration information about MSDP mesh-groups.

Syntax Description

detail	Displays detailed information about MSDP mesh-groups.
mesh-group-name	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to display configuration information about MSDP mesh-groups, as follows:

- For summary information, enter the `show msdp mesh-group` command.
- For detailed information, enter the `show msdp mesh-group detail` command.
- For detailed information about a specific mesh-group, enter the `show msdp mesh-group <name>` command.

Example

The following command displays the peer count for a mesh-group:

```
show msdp mesh-group
```

The following is sample output from this command:

MeshGroupName	PeerCount
external	0
internal	0
msdp_mesh	4

The following command displays detailed information about a mesh-group called “msdp_mesh”:

```
show msdp mesh-group "msdp_mesh"
```

The following is sample output from this command:

```
Mesh Group Name      : msdp_mesh      Num of Peers : 4
Peers                : 54.172.168.97  55.0.0.83    124.56.78.90
                    221.160.90.228
```

show msdp peer

```
show msdp [peer {detail} | {peer} <remoteaddr>] {vr <vrname>}
```

Description

This command displays configuration and run-time parameters about MSDP peers.

Syntax Description

detail	Displays detailed information about MSDP peers.
remoteaddr	Specifies the IP address of the MSDP peer.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to verify the configuration and run-time parameters for MSDP peers, as follows:

- For summary information, enter the `show msdp peer` command.
- For detailed information for all peers, enter the `show msdp peer detail` command.
- For detailed information for a specific peer, enter the `show msdp peer <remoteaddr>` command.

Example

The following command displays configuration and run-time parameters for MSDP peers:

```
show msdp peer
```

The following is sample output from this command:

```
Peer Address   AS      State      Up/Down   Resets  SA_Cnt  Name
-----
-d 54.172.168.97 14490 DISABLED  00:31:36  0        0      test
*e 55.0.0.83     100    ESTABLISHED 00:21:04  1        0      to-Hawaii
-d 124.56.78.90  2345  DISABLED  00:31:36  0        0
```

```
-d 221.160.90.228 23456 DISABLED 00:31:36 0 0
```

Flags: (*) default peer, (d) disabled, (e) enabled

The following command displays output from an MSDP peer with the IP address 16.0.0.2:

```
* Switch.8 # show mdp peer 16.0.0.2
```

```
MSDP Peer           : 16.0.0.2
Enabled             : No                AS Number          : 100.100
Keepalive Interval  : 60                Holdtimer Interval   : 75
Source Address      : not known          TTL Threshold        : 0
Default Peer        : No                Default Peer Filter   : not configured
Process In Request  : Yes               In Request filter     : not configured
Maximum SA Limit    : not configured     Mesh Group           : not configured
Input SA Filter     : not configured     Output SA Filter      : not configured
State               : DISABLED           Uptime/Downtime      : 00:00:02
Local Port          : 0                  Remote Port           : 0
In Total Msgs       : 0                  Out Total Msgs        : 0
In SA Msgs          : 0                  Out SA Msgs           : 0
In SA Req Msgs      : 0                  Out SA Req Msgs       : 0
In SA Resp Msgs     : 0                  Out SA Resp Msgs      : 0
Time since Last Msg : 00:00:02           Hold Tmr Exp in       : 00:00:00
Connection Attempts : 0                  Entered Established    : 0
RPF Fails           : 0                  Output Queue Size     : 0
```

show mdp sa-cache

```
show mdp [sa-cache | rejected-sa-cache] {group-address <grp-addr>} {source-address
<src-addr>} {as-number <as-num>} {originator-rp <originator-rp-addr>} {local} {peer
<remoteaddr>}} {vr <vrname>}
```

Description

This command displays the SA cache database. The following quadruplet per SA cache entry displays: {Group, Source, originating RP, and peer}. In addition, information about the following displays: the cache uptime, aging, whether sources are local or remote, etc.

Syntax Description

as-num	Displays all SA cache that originated from the specified Autonomous System (AS) number.
grp-addr	Displays the SA cache within the specified group address range.
originator-rp-addr	Displays all SA cache entries that were originated by the specified rendezvous point.
local	Displays locally originated SA cache entries only.
remoteaddr	Displays the SA cache entries received from the MSDP peer with the specified IP address.
src-addr	Displays the SA cache within the specified source address range.

vname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
-------	---

Default

N/A.

Usage Guidelines

Use this command to view and troubleshoot the SA cache database. There are various filtering criteria you can use to display just a subset of the SA cache database. The following are some of the criteria, which you can use together or separately, to display information about the SA cache:

- Filtering on the group address range
- Filtering on the source address range
- Filtering on the originator rendezvous point address
- Filtering of the advertising MSDP peer
- Locally originated SA cache
- Rejected SA cache

Example

The following command displays the SA cache database:

```
show mshp sa-cache
```

The following is sample output from this command:

Group Address	Source Address	Originator	Peer Address	Age/Ageout In

235.100.200.1	10.20.30.1	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.2	10.20.30.2	60.0.0.5	192.0.0.16	00:44:24/05:16
235.100.200.3	10.20.30.3	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.4	10.20.30.4	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.5	10.20.30.5	60.0.0.5	55.0.0.5	00:44:24/05:01
235.100.200.6	10.20.30.6	60.0.0.5	178.54.67.23	00:44:24/05:17
235.100.200.7	10.20.30.7	60.0.0.5	112.234.213.12	00:44:24/05:43
235.100.200.8	10.20.30.8	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.9	10.20.30.9	60.0.0.5	10.0.0.1	00:44:24/05:10
235.100.200.10	10.20.30.10	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.11	10.20.30.11	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.12	10.20.30.12	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.13	10.20.30.13	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.14	10.20.30.14	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.15	10.20.30.15	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.16	10.20.30.16	60.0.0.5	0.0.0.0	00:44:24/00:00
235.100.200.17	10.20.30.17	60.0.0.5	0.0.0.0	00:44:24/00:00

```

235.100.200.18  10.20.30.18      60.0.0.5      0.0.0.0      00:44:24/00:00
235.100.200.19  10.20.30.19      60.0.0.5      0.0.0.0      00:44:25/00:00

```

```

Number of accepted SAs      : 255
Number of rejected SAs     : 0

```

Flags: (a) Accepted, (f) Filtered by policy, (r) RPF check failed

unconfigure msdp sa-cache-server

```
unconfigure msdp sa-cache-server {vr <vrname>}
```

Description

Removes the MSDP SA cache server.

Syntax Description

vrname	Specifies the name of the virtual router on which the MSDP cache server is configured. If a virtual router name is not specified, it is extracted from the current CLI context.
---------------	---

Default

By default, the router does not send SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member simply waits to receive SA messages, which eventually arrive.

Usage Guidelines

Use this command to remove the MSDP SA cache server you specified with the [configure msdp sa-cache-server](#) command.

Example

The following command removes the MSDP SA cache server:

```
unconfigure msdp sa-cache-server
```

unconfigure pim border

```
unconfigure pim <vlan_name> border
```

Description

Unconfigures a PIM VLAN that has been configured as a border VLAN, which is used to demarcate a PIM domain when using MSDP.

Syntax Description

vlan_name	Specifies a VLAN name.
-----------	------------------------

Default

By default, no PIM VLANs are configured as border VLANs.

Usage Guidelines

A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN. Use the `unconfigure pim border` command to remove the border functionality of the specified PIM VLAN.

Example

The following command unconfigures a PIM border on a VLAN called “vlan_border”:

```
unconfigure pim vlan_border border
```


This chapter describes commands for managing provider bridge networks (PBNs—also known as vMANs), a Layer 2 feature.

For an introduction to these features, see the *NETGEAR 8800 User Manual*.

configure port ethertype

```
configure port <port_list> ethertype {primary | secondary}
```

Description

Assigns the primary or secondary ethertype value to the specified ports.

Syntax Description

port_list	Specifies the list of ports to be configured.
primary	Assigns the primary ethertype value to the specified ports.
secondary	Assigns the secondary ethertype value to the specified ports.

Default

N/A.

Usage Guidelines

None.

configure vman add ports

```
configure vman <vman-name> add ports [ all | <port_list> ] {untagged | tagged}
```

Description

Adds one or more ports/interfaces to a vMAN.

Syntax Description

vman-name	Specifies the vMAN to configure.
all	Specifies all switch ports.
port_list	Specifies a list of ports.
untagged	Configures the specified ports as vMAN access ports.
tagged	Configures the specified ports as vMAN network ports.

Default

N/A.

Usage Guidelines

If you do not specify a parameter, the default value is untagged.

The vMAN begins at the ingress customer access port, passes through network (switch-to-switch) ports, and terminates at the egress customer access port. Ensure that all the network ports in the vMAN are configured as *tagged* ports. Configure each vMAN access port as an *untagged* port.

Note: vMAN access ports can accept and transmit tagged VLAN frames. You must configure the vMAN *egress access port as untagged so that the vMAN header is stripped from the frame.*

The vMAN must already exist before you can add (or delete) ports. vMAN ports can belong to load-sharing groups.

When a port has untagged and tagged vMAN membership, it inspects received packets to determine whether the vMAN ethertype is matched. Packets with a matching ethertype are treated as tagged and switched across the associated vMAN. Packets with a non-matching ethertype are treated as untagged and forwarded into the associated vMAN.

When a port is only an untagged vMAN member, whether the vMAN ethertype is 0x8100 or otherwise, all received packets ingress the associated vMAN regardless of the packet's tagging.

Note: If you use the same name across categories (for example, STPD names), NETGEAR recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Note the following guidelines:

- You must enable or disable jumbo frames before configuring vMANs. You can enable or disable jumbo frames on individual ports or modules, or on the entire switch. See the chapter on configuring slots and ports on a switch in the *NETGEAR 8800 User Manual* for more information on configuring jumbo frames.
- Platforms support different combinations of tagged and untagged vMANs and VLANs as shown in [Table 30](#).

Table 30. Port Support for Combined vMANs and VLANs

Platform	Combined Untagged vMAN, Tagged VLAN	Combined Tagged vMAN, Untagged VLAN	Combined Tagged vMAN, Tagged VLAN	Combined Tagged vMAN, Untagged vMAN
NETGEAR 8800	X	X	X ¹	X ²

1. The vMAN ethertype must be set to 0x8100, which is different from the default value (0x88a8).

2. A tagged vMAN cannot be added to a port to which an untagged vMAN has previously been added when the selected ethertype is 0x8100.

All vMAN ports are automatically enabled for jumbo frames to accommodate the extra vMAN tag. If any port in the load-sharing group is enabled for a vMAN, all ports in the group are automatically enabled to handle jumbo size frames. Also, the vMAN is automatically enabled on all ports of the untagged load-sharing group.

All ports added to a specified vMAN must be in the same virtual router. For more information on displaying, configuring, and using virtual routers, see [Chapter 11, “Commands for Virtual Routers.”](#)

Example

The following command assigns ports 1:1, 1:2, 1:3, and 1:6 to a vMAN named *accounting*:

```
configure vman accounting add ports 1:1, 1:2, 1:3, 1:6 tag 100
```

configure vman delete ports

```
configure vman <vman-name> delete ports [all | <port_list>]}
```

Description

Deletes one or more ports from a vMAN.

Syntax Description

vman_name	Specifies a vMAN name.
all	Specifies all ports in the vMAN.
port_list	Specifies a list of ports.

Default

N/A.

Usage Guidelines

The vMAN must already exist before you can delete ports.

Example

The following command deletes ports *1:1*, *1:2*, *1:3*, and *1:6* on a modular switch for a vMAN named *accounting*:

```
configure vman accounting delete ports 1:1, 1:2, 1:3, 1:6
```

configure vman ethertype

```
configure vman ethertype <value> [primary | secondary]
```

Description

Changes the default ethertype for the vMAN header.

Syntax Description

value	Specifies an ethertype value in the format of 0xffff.
primary	Assigns the ethertype as the primary Ethernet value.
secondary	Assigns the ethertype as the secondary Ethernet value.

Default

Ethertype value of 0x88a8 and type primary.

Usage Guidelines

The software supports two vMAN ethertype values, a primary value and a secondary value. By default, the primary ethertype applies to all vMANs. To use the secondary ethertype, define the ethertype with this command, and then assign the secondary ethertype to ports with the following command:

```
configure port <port_list> ethertype {primary | secondary}
```

If your vMAN transits a third-party device (other than a NETGEAR device), you must configure the ethertype for the vMAN tag as the ethertype that the third-party device uses. If you configure both primary and secondary ethertypes, you can connect to devices that use either of the two values assigned.

The system supports all vMAN ethertypes, including the standard ethertype of 0x8100.

Example

The following command changes the vMAN ethertype value to *8100*:

```
configure vman ethertype 0x8100
```

configure vman tag

```
configure vman <vman-name> tag <tag>
```

Description

Assigns a tag to a vMAN.

Syntax Description

vman_name	Specifies a vMAN name.
tag	Specifies a value to use as the vMAN tag. The valid range is from 2 to 4094.

Default

N/A.

Usage Guidelines

Every vMAN requires a unique tag.

You can specify a value that is currently used as an internal VLAN ID on another VLAN; it becomes the VLAN ID for the VLAN you specify, and a new VLAN ID is automatically assigned to the other untagged VLAN.

Example

The following command assigns a tag of *120* to a vMAN named *accounting*:

```
configure vman accounting tag 120
```

create vman

```
create vman <vman-name> {vr <vr_name>}
```

Description

Creates a vMAN.

Syntax Description

vman-name	Specifies a vMAN name using up to 32 characters.
vr	Specifies a virtual router.

vr_name	Specifies a virtual router name.
---------	----------------------------------

Default

N/A

Usage Guidelines

For information on vMAN name requirements and a list of reserved keywords, see the section on object names in the *NETGEAR 8800 User Manual*. You must use mutually exclusive names for:

- VLANs
- vMANs
- IPv6 tunnels

If you do not specify the virtual router, the vMAN is created in the current virtual router.

After you create the vMAN, you must configure the vMAN tag and add the ports that you want.

Example

The following command creates a vMAN named *fred*:

```
create vman fred
```

delete vman

```
delete vman <vman-name>
```

Description

Deletes a previously created vMAN.

Syntax Description

vman-name	Specifies a vMAN name.
-----------	------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the vMAN *accounting*:

```
delete vman accounting
```

disable dot1p examination ports

```
disable dot1p examination ports [all | <port_list>]
```

Description

Used with vMANs, this command instructs the switch to examine the 802.1p value of the inner tag header of the original packet to determine the correct egress queue on the egress port.

Syntax Description

all	Specifies all ports.
port_list	Specifies a list of ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to instruct the system to refer to the 802.1p value contained in the inner tag header of the original packet when assigning the packet to an egress queue at the egress port of the vMAN.

Example

The following command uses the 802.1p value on the inner tag header of the original packet to put the packet in the egress queue on the vMAN egress port:

```
disable dot1p examination port 3:2
```

enable dot1p examination port

```
enable dot1p examination port [all | <port_list>]
```

Description

Used with vMANs, and instructs the switch to examine the 802.1p value of the inner tag, or header of the original packet, to determine the correct egress queue on the egress port.

Syntax Description

all	Specifies all ports.
port_list	Specifies a list of ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to instruct the system to refer to the 802.1p value contained in the inner, or original, tag when assigning the packet to an egress queue at the egress port of the vMAN.

Example

The following command puts the packets in the egress queue of the vMAN egress port according to the 802.1p value on the inner tag:

```
enable dot1p examination port 3:2
```

show vman

```
show vman {<vman_name> {ipv6} | etherType | detail {ipv6}}
```

Description

Displays vMAN configuration information.

Note: The information displayed for this command depends on the platform and configuration you are using.

Syntax Description

vman_name	Specifies that information is displayed for the specified vMAN.
ipv6	Specifies IPv6.
ethertype	Displays ethertype information for VLANs and vMAN.
detail	Specifies that all information is displayed for each vMAN.

Default

Summary information for all vMANs on the device.

Usage Guidelines

The information displayed with this command depends on the platform and configuration you are using.

Example

The following example displays a list of all the vMANs on the switch:

```
* XCM8806.22 # show vman
```

Name	VID	Protocol	Addr	Flags	Proto	Ports	Virtual Active router /Total
vman1	4089	-----	-----	----	ANY	0 / 0	VR-Default

```
Flags : (a) Learning Domain, (B) 802.1ah Backbone VMAN,
        (c) 802.1ad customer VLAN, (f) IP Forwarding Enabled,
        (F) Learning Disabled, (i) ISIS Enabled,
        (I) IP Forwarding lpm-routing Enabled, (L) Loopback Enabled,
        (m) IPmc Forwarding Enabled, (n) IP Multinetting Enabled, (N) Network LogIn vlan,
        (N) Network LogIn vlan, (o) OSPF Enabled, (O) Flooding Disabled,
        (p) PIM Enabled, (r) RIP Enabled,
        (S) 802.1ah Service VMAN, (T) Member of STP Domain, (v) VRRP Enabled
```

The following example displays information on a single vMAN named *vman1*:

```
XCM8806.23 # show vman vman1
```

VMAN Interface with name vman1 created by user

```
Admin State:      Enabled      Tagging:Untagged (Internal tag 4089)
Virtual router: VR-Default
IPv6:             None
STPD:             None
Protocol:         Match all unfiltered protocols
Loopback:         Disabled
NetLogin:         Disabled
QosProfile:       None configured
Egress Rate Limit Designated Port: None configured
Flood Rate Limit QosProfile:      None configured
Ports:           0.           (Number of active ports=0)
```

The `show vman detail` command shows all the information shown in the `show vman <vlan_name>` command, but displays information for all configured vMANs.

The following example shows the display from the `show vman etherType` command on switches that support only vMANs:

```
vMan Ethertype: 0x88a8
```

The following example shows the display from the `show vman etherType` command:

```
XCM8810.3 # show vman etherType
Vman Primary EtherType : 0x88a8
```

The following example shows the display from the `show vman etherType` command when a secondary ethertype is configured:

```
XCM8806.29 # show vman etherType
Vman Primary EtherType   : 0x9100
Vman Secondary EtherType : 0x8100
```


The letter *g* in the port list indicates that the port is a LAG/Trunk port, the details of which can be seen using the `show port sharing` command.

unconfigure vman ethertype

```
unconfigure vman ethertype {secondary}
```

Description

Restores the default primary vMAN ethertype value of 0x88A8 or deletes the secondary ethertype value.

Syntax Description

secondary	Deletes the secondary ethertype value.
-----------	--

Default

N/A.

Usage Guidelines

When you enter this command without the `secondary` option, the primary vMAN ethertype returns to the default value of 0x88A8. If you specify the `secondary` option, the secondary vMAN ethertype value is deleted (no value is assigned).

Note: Before unconfiguring the secondary vMAN ethertype, any secondary vMAN port must be changed to the primary vMAN ethertype; otherwise this command fails.

Examples

The following command restores the primary vMAN ethertype to the default value:

```
unconfigure vman ethertype
```

The following command restores the secondary vMAN ethertype to the default value:

```
unconfigure vman ethertype secondary
```

Configuration and Image Commands



This appendix describes commands for:

- Downloading and using a new switch software image
- Saving, uploading, and downloading switch configuration information
- Downloading and installing a new BootROM image and switch rebooting

The switch software *image* contains the executable code that runs on the switch. An image comes preinstalled from the factory. The image can be upgraded by downloading a new version from a Trivial File Transfer Protocol (TFTP) server on the network. You can also download a new version from the external compact flash memory card installed in the external compact flash slot of the Management Module (MSM/MM).

A switch can store up to two images; a primary and a secondary image. You can download a new image into either one of these, and you can select which image will load on the next switch reboot.

The *configuration* is the customized set of parameters that you have selected to run on the switch. As you make configuration changes, the new settings are stored in run-time memory. To retain the settings, and have them load when you reboot the switch, you must save the configuration to nonvolatile storage.

The switch can store multiple user-defined configuration files, each with its own file name. By default, the switch has two pre-named configurations: a primary and a secondary configuration. You can select to which configuration you want the changes saved, or you can save the changes to a new configuration file. You can also select which configuration will be used on the next switch reboot.

The BootROM initializes certain important switch variables during the switch boot process. On the NETGEAR 8800 series switches, you can upgrade the firmware, including the BootROM, when you upgrade the software image.

configure firmware

```
configure firmware [auto-install | install-on-demand]
```

Description

Configures the way a NETGEAR 8800 series switch performs a system firmware upgrade.

Syntax Description

auto-install	Specifies NETGEAR 8800 to automatically upgrade the firmware if the software detects a newer firmware image is available. The switch does not prompt you to confirm the firmware upgrade.
install-on-demand	Specifies the switch to prompt you to upgrade the firmware when the NETGEAR 8800 determines that a newer firmware image is available. This is the default behavior.

Default

The default is install-on-demand.

Usage Guidelines

Use the `configure firmware [auto-install | install-on-demand]` and `install firmware {force}` commands to upgrade the BootROM images on the MSM and I/O modules and the firmware on the PSU controllers installed in a NETGEAR 8800 series switch.

Firmware images are bundled with NETGEAR 8800 software images. NETGEAR 8800 automatically compares the existing firmware image flashed into the hardware with the firmware image bundled with the NETGEAR 8800 image when you:

- Download a new version of firmware to the alternate (inactive) partition.
- Install a new module into an active chassis.

After a firmware image upgrade, messages are sent to the log.

If you select the `auto-install` parameter, you are not prompted to confirm the firmware upgrade. Whenever NETGEAR 8800 determines a newer firmware image is available, the firmware is automatically upgraded.

If you use the default configuration `install-on-demand`, you have the opportunity to cancel the firmware upgrade. If you install a new software image, and a new firmware image is available, the switch prompts you to upgrade the firmware. Enter `y` to upgrade the firmware image. Enter `n` to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.

The following command downloads the switch software image. The secondary partition must be the alternate partition.

```
download image NG8800-12.4.3.5-1-4.xos secondary
```

If you download a new image and new firmware images are available, you see messages similar to the following:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel) Yes
```

```
Downloading to
MSM-A.....
.....
Saving configuration on secondary MSM ..... done!
```

Installing to secondary partition!

Installing to

```
MSM-A.....
.....
.....
.....
.....
```

Image installed successfully

Installing version 1.0.0.16 of the MSM bootrom(s). Do you want to continue? (y/n) Yes

Installing version 1.0.0.24 of the IO module bootrom(s). Do you want to continue? (y/n) Yes

Installing version 2.4 of the PSU control module firmware. Do you want to continue? (y/n) Yes

Installing bootrom...

MSM bootrom(s) installed successfully

Installing bootrom...

IO module bootrom(s) installed successfully

Installing firmware...

PSU controller firmware installed successfully

...

Displaying BootROM and Firmware Versions

To display the BootROM (firmware) version for all modules and PSU controllers installed in the switch, use the `show version` command.

Recovering From a Corrupted BootROM

If your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a pen into the Alternate (A) and Reset (R) holes on the NETGEAR 8800 MSM and applying pressure. For more information, please refer to the hardware documentation.

Example

The following command automatically upgrades the firmware when a newer firmware image is present without prompting you to confirm the upgrade:

```
configure firmware auto-install
```

download image

Using TFTP

```
download image [[<hostname> | <ipaddress>] <filename> {{vr} <vrname>}
| memorycard <filename>] {<partition>} {msm <slotid>}
```

Description

Downloads a new version of the NETGEAR 8800 software image.

The image file can be downloaded using TFTP which is not a secure method or SFTP and SCP2 which are secure methods. The procedure using TFTP begins above and using SFTP/SCP2 [on page 1310](#).

Note: A NETGEAR 8800 core image must be downloaded and installed on the alternate (non-active) partition. If a user tries to download to an active partition, the error message “Error: Image can only be installed to the non-active partition.” is displayed.

Syntax Description

hostname	Specifies the hostname of the TFTP server from which the image should be obtained.
ipaddress	Specifies the IP address of TFTP server from which the image should be obtained.
memorycard	Specifies that the image should be obtained from the external compact flash memory card.
filename	Specifies the filename of the new image.
vrname	Specifies the name of the virtual router.
	Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the <i>NETGEAR 8800 User Manual</i> .
partition	Specifies which partition the image should be saved to: primary or secondary. Select primary to save the image to the primary partition and secondary to save the image to the secondary partition.
slotid	Specifies the MSM/MM where the software image should be downloaded. <ul style="list-style-type: none"> • A specifies the MSM/MM installed in slot A. • B specifies the MSM/MM installed in slot B.
slot number	Specifies the slot where the software image should be downloaded. The value may be from 1 to 8.
	Note: This parameter is available only on stackable switches in a stack.

Default

Stores the downloaded image in the alternate (inactive) partition.

Using SFTP and SCP2

SFTP and SCP2 provide secure methods of downloading the NETGEAR 8800 software image files, *.xos or *.xmod. You can use one of three procedures:

- From the switch, running the command SCP2. connect to and “get” from a remote server. This is similar to the `download image` command.
- From outside the switch, connect to the switch which is acting as the server and “put” from the remote server. There is no TFTP equivalent for this method.
 - Using SFTP, or
 - Using SCP2.

Examples of these procedures are included in the Examples section that starts on [page 1317](#).

Usage Guidelines

Prior to downloading an image on the switch, you must download the image you received from NETGEAR to a TFTP server on your network. You can also download the image to the external compact flash memory card.

Note: The `download image` command in the NETGEAR 8800 causes the switch to use the newly downloaded software image during the next switch reboot. To modify or reset the software image used during a switch reboot, use the `use image` command.

Specify the `ipaddress` or `hostname` parameters to download an image from a TFTP server on the network. Use of the `hostname` parameter requires that DNS be enabled.

Core Software Images

The switch can store up to two core images: a primary and a secondary. When downloading a new image, you select which partition (primary or secondary) to install the new image. The NETGEAR 8800 core image must be downloaded and installed to the alternate partition.

Image Filenames

The software image file can be an .xos file, which contains a NETGEAR 8800 core image, or an .xmod file, which contains a NETGEAR 8800 modular software package. Modular software packages have additional functionality that supplement a core image.

Displaying the Software Image Versions

To display the software image version running on the switch, use the `show version` or `show switch` commands.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Local and Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a local or remote file, remember the requirements listed above.

Messages Displayed by the Switch

When you download a new image, you see the following message:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel)
```

Do one of the following:

- Enter `y` if you want to install the image after download.
- Enter `n` if you want to install the image at a later time.
- Press [Enter] if you want to cancel the download.

Core Dump Messages

If you configure the switch to write core dump (debug) files to the internal memory card and attempt to download a new software image, you might have insufficient space to complete

the image download. If this occurs, you must decide whether to continue the software download or move or delete the core dump files from the internal memory. For example, if your switch has an external memory card installed with space available, transfer the files to the external memory card. This frees up space on the internal memory card while keeping the core dump files.

The switch displays a message similar to the following and prompts you to take action:

```
Core dumps are present in internal-memory and must be removed before this download can
continue. (Please refer to documentation for the "configure debug core-dumps" command for
additional information)
```

```
Do you want to continue with download and remove existing core dumps? (y/n)
```

Enter `y` to remove the core dump files and download the new software image. Enter `n` to cancel this action and transfer the files before downloading the image.

For information about configuring and sending core dump information, see the [configure debug core-dumps](#) and [save debug tracefiles memorycard](#) commands.

Specify `memorycard` to download a an image from the external compact flash memory card installed in the external compact flash slot of the MSM/MM. Use a PC with appropriate hardware such as a compact flash reader/writer and follow the manufacturer's instructions to access the compact flash card and place the image onto the card. For more information about installing the external compact flash memory card into the MSM/MM, see the hardware documentation.

Downloading a New Image

The following assumes you have already downloaded the image to a network TFTP server or external memory card. The information in this section provides more detailed information for downloading a new image to your switch.

Note: Always refer to the most recent version of the release notes for the most current download instructions.

Step 1—Verifying the Virtual Router

If you loaded the image onto an external compact flash, proceed to step 2.

If you loaded the image onto a TFTP server, use one of the following ping commands to confirm which virtual router reaches your TFTP server:

```
ping vr vr-Mgmt <host>
ping vr vr-Default <host>
```

At least one of these commands must successfully reach your TFTP server for you to download the image. After verifying the virtual router that reaches your TFTP server, specify that virtual router when you download the image.

Step 2—Viewing the Partition

To view your selected and booted partition, use the following command:

```
show switch
```

Output from this command includes the selected and booted images and if they are in the primary or the secondary partition.

Step 3—Selecting the Partition

The image must be downloaded and installed to the alternate (inactive) partition. To specify the partition when downloading and installing the image, use one of the following commands:

```
download image [[<hostname> | <ipaddress>] <filename> [{vr} <vrname>} | memorycard  
<filename>] [<partition>] {msm <slotid>}
```

Step 4—Downloading and Installing the Image

To download the image, use the appropriate, previously described, `download image` command.

Downloading a NETGEAR 8800 core image

A NETGEAR 8800 core image uses the file extension `.xos`.

- Before the download begins, the switch asks if you want to install the image immediately after the download is finished.

Enter `y` to install the image after download. Enter `n` to install the image at a later time.

When you install the image to the alternate (inactive) partition; you do not need to reboot the switch until you use the image.

If you install the image at a later time, the image is still downloaded and saved to the switch, but you must use the following command to install the software and reboot the switch:

```
install image <fname> [<partition>] {msm <slotid>} {reboot}
```

Where `fname` specifies the filename of the new, downloaded image.

Downloading a NETGEAR 8800 module image

A NETGEAR 8800 module image has functionality that supplements a core image. You download and install a module onto an already installed core image. The version number of the core image and the module must match. For example, the module `NG8800-12.4.3.5-1-4-ssh.xmod` can be installed only onto the core image `NG8800-12.4.3.5-1-4.xos`.

To install a module to the inactive partition, use the `download image` command to download the module to the inactive partition, and use the `install image` command to install it, if you did not choose to install when the image was downloaded. Remember, the core image on the inactive partition must be of the same version as the module. When you make the inactive partition

active, by issuing the `use image` command and rebooting the switch, the module is also activated at boot time.

To install a module to the active partition, use the `download image` command to download the module to the active partition, and use the `?` command to install it, if you did not choose to install when the image was downloaded. Remember, the core image on the active partition must be of the same version as the module. If you reboot the switch, the module will also be activated, but you can activate the module without rebooting the switch by issuing the `run update` command. After issuing that command, all the functionality, and command line interface (CLI) commands, of the module will be available.

Performing a Hitless Upgrade

Hitless upgrade is a mechanism that allows you to upgrade the NETGEAR 8800 software running on the switch without taking the switch out of service. Some additional benefits of using hitless upgrade include:

- Minimizing network downtime
- Reducing the amount of traffic lost

Although any method of upgrading software can have an impact on network operation, including interrupting Layer 2 network operation, performing a hitless upgrade can decrease that impact.

You must have two MSMs installed in your switch to perform a hitless upgrade. With two MSMs installed in the switch, one assumes the role of primary and the other assumes the role of backup. The primary MSM provides all of the switch management functions including bringing up and programming the I/O modules, running the bridging and routing protocols, and configuring the switch. The primary MSM also synchronizes its configurations with the backup MSM which allows the backup to take over the management functions of the primary.

Note: If you download an image to the backup MSM, the image passes through the primary MSM before the image is downloaded to the backup MSM.

To perform a hitless upgrade, do the following:

1. View current switch information using the following command:

```
show switch
```

Determine your selected and booted partition, verify which MSM is the primary and which is the backup, and confirm that the MSMs are synchronized.

Output from this command indicates, for each MSM, the selected and booted images and if they are in the primary or the secondary partition. The selected image partition indicates which image will be used at the next reboot. The booted image partition indicates the image used at the last reboot. It is the active partition.

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

2. Select the partition to download the image to and download and install the new NETGEAR 8800 software on the backup MSM using the following command:

```
download image [[<hostname> | <ipaddress>] <filename> {{vr} <vrname>} |
memorycard <filename>] {<partition>} {msm <slotid>}
```

Note: If the backup MSM is installed in slot B, specify msm B. If the backup MSM is installed in slot A, specify msm A.

- Before the download begins, the switch asks if you want to install the image immediately after the download is finished.

Enter `y` to install the image after download. Enter `n` to install the image at a later time.

When you install the image after download to the alternate partition, you then need to reboot the switch.

If you install the image at a later time, use the following command to install the software:

```
install image <fname> {<partition>} {msm <slotid>} {reboot}
```

3. Verify that the backup MSM comes up correctly and that the MSMs are synchronized using the following command:

```
show switch
```

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

4. Initiate failover from the primary MSM to the backup MSM using the following command:

```
run msm-failover
```

When you failover from the primary MSM to the backup MSM, the backup becomes the new primary, runs the newly downloaded software, and provides all of the switch management functions.

If you have a NETGEAR 8800 series switch and the new NETGEAR 8800 image supports hitless upgrade but is not compatible with the current running I/O module image (the I/O version numbers do not match), you cannot perform a hitless upgrade.

The switch displays a warning message similar to the following:

```
WARNING: The other MSM operates with a different version of I/O module image.
If you continue with the MSM failover, all I/O modules will be reset.
Are you sure you want to failover? (y/n)
```

You can either continue the upgrade or cancel the action. If you continue the upgrade, the primary MSM downloads the new image to the I/O module and reboots.

5. Verify that the backup MSM comes up correctly and that the MSMs are synchronized using the following command:

```
show switch
```

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

6. Select the partition to download the image to and download and install the new NETGEAR 8800 software on the new backup MSM (this was the original primary MSM) using the following command:

```
download image [<hostname> | <ipaddress>] <filename> {vr <vrname>} msm <slotid>
```

Note: If the new backup MSM is installed in slot A, specify msm A. If the new backup MSM/MM is installed in slot B, specify msm B.

Before the download begins, the switch asks if you want to install the image immediately after the download is finished.

- When you download and install the software image on the alternate partition, you then need to reboot the switch.
- If you install the image at a later time, use the following command to install the software:

```
install image <fname> {<partition>} {msm <slotid>} {reboot}
```

7. Verify that the new backup MSM comes up correctly and that the MSMs are synchronized using the following command:

```
show switch
```

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

8. Optionally, initiate failover from the new primary MSM to the new backup MSM using the following command:

```
run msm-failover
```

When you failover from the new primary MSM to the new backup MSM, this optional step restores the switch to the original primary and backup MSM.

9. Optionally, confirm that the failover is successful by checking the current state of the MSMs using the following command:

```
show switch
```

You can also perform a hitless upgrade on NETGEAR 8800 modular software packages (.xmod files). To perform a hitless upgrade of a software package, you must install the core software image first, and the version number of the modular software package must match the version number of the core image that it will be running with. For more information about hitless upgrade, see the *NETGEAR 8800 User Manual*.

Examples

Using TFTP

The following command downloads the switch software image from the TFTP server at 10.10.15.4, from the file named NG8800-12.4.3.5-1-4.xos without specifying the desired partition:

```
download image 10.10.15.4 NG8800-12.4.3.5-1-4.xos
```

Note: The inactive partition (secondary) will be used for installation.

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel) Yes
```

```
Downloading to MSM-A.....
```

```
Installing to secondary partition!
```

```
Installing to MSM-A.....
```

```
.....  
.....  
.....
```

The following command downloads the switch software image from the TFTP server at 10.10.15.4, from the file named NG8800-12.4.3.5-1-4.xos specifying the desired partition. The secondary partition is the alternate partition in this example.

```
download image 10.10.15.4 NG8800-12.4.3.5-1-4.xos secondary
```

When you download an image into the alternate partition, you see output similar to the following:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel) Yes
```

```
Downloading to MSM-A.....
```

```
Downloading to MSM-B.....
```

```
Installing to secondary partition!
```

```
Installing to MSM-B .....
```

```
.....  
.....  
.....
```

```
Installing to MSM-A.....
```

```
.....  
.....  
.....
```

If you answer *yes* to installing the image, the switch reboots upon completion of the installation.

Using SFTP and SCP2

The following commands show the three procedures

- From the Switch using SCP2

```
XCM8806.5 # scp2 root@10.203.133.13:/tmp/NG8800-12.4.3.5-1-4-ssh.xmod
NG8800-12.4.3.5-1-4-ssh.xmod
Download to /scratch/NG8800-12.4.3.5-1-4-ssh.xmod on switch
Connecting to 10.203.133.13...
root@10.203.133.13's password:
Fetching /tmp/NG8800-12.4.3.5-1-4-ssh.xmod to
/scratch/NG8800-12.4.3.5-1-4-ssh.xmod
/tmp/NG8800-12.4.3.5-1-4-ssh.xmod 100% 997KB 996.9KB/s 00:00
Please wait validating image NG8800-12.4.3.5-1-4-ssh.xmod, this will take
approximately 30 seconds ...
XCM8806.6 #
```

- From Outside the Switch using SFTP

```
sftp admin@10.203.3.231
Connecting to 10.203.3.231...
admin@10.203.3.231's password:
sftp> put NG8800-12.4.3.5-1-4-ssh.xmod Examplesftp.xmod
Uploading NG8800-12.4.3.5-1-4-ssh.xmod to /config/Examplesftp.xmod
sftp>
```

- From Outside the Switch using SCP2

```
scp2 NG8800-12.4.3.5-1-4-ssh.xmod admin@10.203.3.231:Examplescp.xmod
Keyboard-interactive:
Keyboard-interactive authentication
Enter password for admin:
NG8800-12.4.3.5-1-4-ssh.xmod | 997kB | 997kB/s | TOC: 00:00:01 | 100%
[~/cvs/branch-fixes_v1200b4.c4/exos/cougar/release]$
```

Following download, the output from a `show log` command is similar to the following for the last two procedures—the two from outside the switch.

```
show log

04/03/2007 14:32:29.31 <Info:AAA.LogSsh> Got Image file Example.xmod
04/03/2007 14:32:29.31 <Info:AAA.LogSsh> Validating Image file, this could take
approximately 30 seconds.. Example.xmod
04/03/2007 14:32:30.89 <Info:AAA.LogSsh> Image file Example.xmod successfully
validated
can now install image
```

install firmware

```
install firmware {force}
```

Description

Installs the firmware bundled with the NETGEAR 8800 image on the NETGEAR 8800 series switches.

Syntax Description

force	Specifies that a new image is installed without a version check.
-------	--

Default

N/A.

Usage Guidelines

On NETGEAR 8800 series switches, use the `install firmware` command to upgrade the BootROM images on the MSM and I/O modules and the firmware on the PSU controllers installed in the switch.

Firmware images are bundled with NETGEAR 8800 software images.

On NETGEAR 8800 series switches, the NETGEAR 8800 software automatically compares the existing firmware image flashed into the hardware with the firmware image bundled with the NETGEAR 8800 image. You can also use the `install firmware` command to compare the firmware images.

Before using the `install firmware` command, wait until the `show slot` command indicates the MSMs and I/O modules are operational. When the modules are operational, use the `install firmware` command.

NETGEAR 8800 Series Switches

The switch scans the I/O and MSM modules and the PSU controllers for a possible firmware upgrade. If the bundled firmware image is newer than the existing firmware image, the switch prompts you to confirm the upgrade.

- Enter `y` to upgrade the firmware.
- Enter `n` to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.
- Enter `<cr>` to cancel the upgrade. After a firmware image upgrade, messages are sent to the log.

The PSU controller firmware is used immediately after it is installed without rebooting the switch. The new BootROM and firmware overwrite the older versions flashed into the hardware. Use the `reboot` command to reboot the switch and activate the new BootROM and firmware.

During the firmware upgrade, do not cycle down or disrupt the power to the switch. If a power interruption occurs, the firmware may be corrupted and need to be recovered. NETGEAR 8800 automatically recovers corrupted firmware; however, the time it takes for the switch to boot-up may increase.

The switch displays status messages after you use the `install firmware` command. The output varies depending upon your platform and the software version running on your system.

Note: If the information in the most current version of the NETGEAR 8800 Installation and Release Notes differs from the information in this section, follow the release notes.

Sample Output—NETGEAR 8800 Series Switch

The following is sample output from a NETGEAR 8800 series switch:

```
Installing version 1.0.4.2 of the bootrom for MSMs. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Do you want to save configuration changes to primary.cfg? (y or n) Yes
Saving configuration on primary MSM ..... done!
Installing version 1.0.4.0 of the bootrom for I/O modules.
Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing version 1.0.1.0 of the bootrom for newer (e.g. 8900-series)
I/O modules. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing version 2.13 of the firmware for PSU control modules.
Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes

Installing bootrom...

MSM bootrom(s) installed successfully
MSM bootrom(s) will be activated upon next MSM reboot

Installing bootrom...

IO module bootrom(s) installed successfully
IO module bootrom(s) will be activated upon next IO module reboot

Installing firmware...
PSU controller firmware installed successfully
```

Additional Behavior

During a firmware upgrade, the switch prompts you to save your configuration changes to the current, active configuration. Enter `y` to save your configuration changes to the current, active configuration. Enter `n` if you do not want to save your changes.

Use the `configure firmware [auto-install | install-on-demand]` command to configure how the switch performs a system firmware upgrade. If you select the `auto-install` parameter, you are not prompted to confirm the firmware upgrade. If you use the default configuration `install-on-demand`, you can cancel the firmware upgrade.

Power over Ethernet (PoE) firmware is always automatically upgraded or downgraded to match the operational code image. This configuration is not applicable to PoE firmware.

Recovering From a Corrupted BootROM

If your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a pen into the Alternate (A) and Reset (R) holes on the NETGEAR 8800 MSM and applying pressure. For more information, please refer to the hardware documentation.

Displaying BootROM and Firmware Versions

To display the BootROM (firmware) version for all modules and PSU controllers installed in the switch, use the `show version` command.

Example

The following command installs the newer firmware image(s):

```
install firmware
```

install image

```
install image <fname> {<partition>} {msm <slotid>} {reboot}
```

Description

Installs a new version of the NETGEAR 8800 software image.

Note: A NETGEAR 8800 core image must be installed on the alternate (non-active) partition. If a user tries to install on an active partition, the error message “Error: Image can only be installed to the non-active partition.” is displayed.

Syntax Description

fname	Specifies the software image file.
partition	Specifies which partition the image should be saved to: primary or secondary. Select primary to save the image to the primary partition and secondary to save the image to the secondary partition.
reboot	Reboots the switch after the image is installed.

Default

N/A.

Usage Guidelines

When you download a software image, you are asked if you want to install the image immediately after the download is finished. If you choose to install the image at a later time, use this command to install the software on the switch.

The software image file can be an .xos file, which contains a NETGEAR 8800 core image, or an .xmod file, which contains additional functionality to supplement a core image.

Displaying the Software Image Version

To display the software image version running on the switch, use the `show version` or `show switch` commands.

Displaying the Downloaded Software Image Version.

To display a software image version that has been downloaded but not installed, use the `install image ?` command.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local, remember the requirements listed above.

Installing a NETGEAR 8800 core image

Install the software image on the alternate partition. You can continue to run the currently booted image, but to run the newly installed image, you will need to set the boot partition with the `use image {partition} <partition> {msm <slotid>}` command and reboot the switch.

Installing a NETGEAR 8800 module image

A NETGEAR 8800 module image has functionality that supplements a core image. You will install a module onto an already installed core image. The version number of the core image and the module must match. For example, the module *NG8800-12.4.3.5-1-4-ssh.xmod* can only be installed onto the core image *NG8800-12.4.3.5-1-4.xos*.

To install a module to the alternate partition, use the `install firmware` command to install the module. Remember, the core image on the alternate partition must be of the same version as

the module. When you make the alternate partition active, by issuing the `use image` command and rebooting the switch, the module is also activated at boot time.

To install a module to the active partition, use the `install firmware` command to install the module. Remember, the core image on the active partition must be of the same version as the module. If you reboot the switch, the module will also be activated, but you can activate the module without rebooting the switch by issuing the `run update` command. After issuing that command, all the functionality, and CLI commands, of the module will be available.

Performing a Hitless Upgrade

If you specify the `msm` parameter on a NETGEAR 8800 series switch, you can initiate hitless upgrade between the primary and backup MSMs installed in the switch.

Hitless upgrade is a mechanism that allows you to upgrade the NETGEAR 8800 software running on the switch without taking the switch out of service. Some additional benefits of using hitless upgrade include:

- Minimizing network downtime
- Reducing the amount of traffic lost

Although any method of upgrading software can have an impact on network operation, including interrupting Layer 2 network operation, performing a hitless upgrade can decrease that impact.

Regardless of how you upgrade the software, you must:

- View the current switch information to determine your selected and booted image partitions, verify which MSM is the primary and which is the backup, and confirm that the MSMs are synchronized using the `show switch` command
- Select the partition to use when downloading the image using the `download image` `[[<hostname> | <ipaddress>] <filename> [{vr} <vrname>] | memorycard <filename>] {<partition>} {msm <slotid>}` command.

When performing a hitless upgrade, you must first download the software to the backup MSM. Download the image to the alternate partition. Use the `install image <fname> {<partition>} {msm <slotid>} {reboot}` command to install the software image at a later time. After the software is downloaded and installed on the switch, use the `run msm-failover` command to initiate failover from the primary MSM to the backup MSM. The original primary MSM becomes the new backup MSM.

After failover is complete, download the software to the new backup MSM. Again, download the image to the alternate partition, and use the `install image <fname> {<partition>} {msm <slotid>} {reboot}` command to install the software image at a later time.

For more detailed information about hitless upgrade, see the `download image` command.

Example

The following command installs the software image file NG8800-12.4.3.5-1-4.xos on a NETGEAR 8810 switch:

```
install image NG8800-12.4.3.5-1-4.xos
```

load script

```
load script <filename> {arg1} {arg2} ... {arg9}
```

Description

Loads (plays back) an ASCII-formatted configuration file or a user-written script file on the switch.

Syntax Description

filename	Specifies the user-defined name of the ASCII-formatted configuration file or a user-written script file. The script file is known as the XOS script file and uses the .xsf file extension.
arg	Specifies up to nine variable values that can be specified by the user. The variables are created with the names CLI.ARGV1, CLI.ARGV2, ... CLI.ARGV9.

Default

N/A.

Usage Guidelines

Use this command to load an ASCII-formatted configuration file or a user-written script file.

Configuration File: After downloading the configuration file from the TFTP server, this command loads and restores the ASCII-formatted configuration file to the switch.

An ASCII-formatted configuration file uses the .xsf file extension, not the .cfg file extension. The .xsf file extension (known as the XOS script file) saves the XML-based configuration in an ASCII format readable by a text editor.

For more detailed information about the ASCII configuration file, including the steps involved to upload, download, and save the configuration, see the [upload configuration \[<hostname> | <ipaddress>\] <filename> {vr <vr-name>}](#) command.

User-Written Script File: After writing a script, this command executes the script and passes arguments to it. As with the configuration files, these files use the .xsf file extension that is automatically added.

The command allows up to nine optional variable values to be passed to the script. These are created with the names CLI.ARGV1, CLI.ARGV2, CLI.ARGV3, ... CLI.ARGV9.

In addition, two other variables are always created. CLI.ARGV0 gives the count of the number of parameters passed, and CLI.ARGV1 contains the name of the script that is being executed.

To check the variable values use the command, `show var`.

Note: Only the .xsf extension is used. The `load script` command assumes an .xsf extension and retries opening the file if the file cannot be found with the original specified name or no extension is provided.

Example

The following command loads the ASCII-formatted **configuration** named configbackup.xsf:

```
load script configbackup.xsf
```

After issuing this command, the ASCII configuration quickly scrolls across the screen. The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```

The following command loads a **user-written script** file with the filename “createVlan.xsf” using the two arguments, “red” and “24.” The variables are then checked using `show var`.

```
XCM8810.16 # load script createVlan.xsf red 24
```

The output resembles the following:

```
Executing EXSH script file: createVlan.xsf ...
```

```
1 # set var vlanName $CLI.ARGV1
2 # set var tag $CLI.ARGV2
3 # if ($CLI.ARGV != 0) then
4 # create vlan $vlanName
5 # configure $vlanName tag $tag
6 # else
7 # create log entry "Could not create vlan"
8 # endif
* XCM8810.17 # show var
```

```
-----
Count : 8
-----
```

```
-----
variableName          variableValue
-----
CLI.ARGV              2
CLI.ARGV0             createVlan.xsf
CLI.ARGV1             red
CLI.ARGV2             24
```

```

CLI.SESSION_TYPE      telnet
CLI.USER               admin
tag                   24
vlanName               red

```

```

-----
run update
run update

```

Description

Activates a newly installed modular software package.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

After you install a modular software package to the active partition, use this command to make the update active. This command causes the NETGEAR 8800 to start the newly installed processes contained in the package, without rebooting the switch.

If you installed the package to the inactive partition, you need to reboot the switch to activate the package.

Example

The following command activates any newly installed modular software packages installed on the active partition:

```
run update
```

save configuration

```
save configuration {primary | secondary | <existing-config> | <new-config>}
```

Description

Saves the current configuration from the switch's runtime memory to non-volatile memory.

Syntax Description

primary	Specifies the primary saved configuration.
secondary	Specifies the secondary saved configuration.
existing-config	Specifies an existing user-defined configuration.

new-config	Specifies a new user-defined configuration.
------------	---

Default

Saves the current configuration to the location used on the last reboot.

Usage Guidelines

The configuration takes effect on the next reboot.

Each file name must be unique and can be up to 32 characters long but cannot include any spaces, commas, or special characters.

Configuration files have a .cfg file extension. When you enter the name of the file in the CLI, the system automatically adds the .cfg file extension. Do not use this command with ASCII-formatted configuration files. Those configuration files have an .xsf file extension. For more information about using ASCII-formatted configuration files see the [upload configuration \[<hostname> | <ipaddress>\] <filename> {vr <vr-name>} and the load script <filename> {arg1} {arg2} ... {arg9} commands](#).

This command also displays in alphabetical order a list of available configurations. The following is sample output that displays the primary, secondary, and user-created and defined configurations ("test" and "XOS1" are the names of the user-created and defined configurations):

```
exsh.9 # save configuration
<cr>                Execute the command
primary             Primary configuration file
secondary           Secondary configuration file
<existing-config>    Existing configuration file name
    "test"  "XOS1"
<new-config>        New configuration file name
```

The switch prompts you to save your configuration changes. Enter **y** to save the changes or **n** to cancel the process.

If you enter **n**, the switch displays a message similar to the following:

```
Save configuration cancelled.
```

If you enter **y**, the switch saves the configuration and displays a series of messages. The following sections provide information about the messages displayed when you save a configuration on your switch.

Note: Configuration files are forward compatible only and not backward compatible. That is, configuration files created in a newer release, such as firmware 12.4, might contain commands that do not work properly in an older release, such as firmware 12.1.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Saving a New Configuration

If you create and save a configuration with a new file name, the switch saves the new configuration and then prompts you to select the newly created configuration as the switch's default configuration.

The following sample output is similar to the message displayed:

```
Do you want to save configuration to test1.cfg? (y/n) Yes
Saving configuration on primary MSM ..... done!
Configuration saved to test1.cfg successfully.
```

The switch then prompts you to select which configuration to use to bootup the system. The following sample output is similar to the message displayed:

```
The current selected default configuration database to boot up the system
(primary.cfg) is different than the one just saved (test.cfg).
Do you want to make test.cfg the default database? (y/n)
```

Enter **y** to use the new configuration as the default configuration. Enter **n** to cancel the operation and keep using the current default, active configuration.

Saving an Existing Configuration

If you make and save changes to an existing configuration, the switch prompts you to save and override the existing configuration.

The following sample output is similar to the message displayed:

```
The configuration file test.cfg already exists.
Do you want to save configuration to test.cfg and overwrite it? (y/n) Yes
Saving configuration on primary MSM ..... done!
Configuration saved to test.cfg successfully.
```

If you override an existing configuration that is not the current default, active configuration, the switch prompts you to select which configuration to use to bootup the system. The following sample output is similar to the message displayed:


```
The current selected default configuration database to boot up the system
(primary.cfg) is different than the one just saved (test.cfg).
Do you want to make test.cfg the default database? (y/n) No
Default configuration database selection cancelled.
```

Enter **y** to use the updated configuration as the default configuration. Enter **n** to cancel the operation and keep using the current default, active configuration.

Example

The following command saves the current switch configuration to the configuration file named **XOS1**:

```
save configuration XOS1
```

The following command save the current switch configuration to the secondary configuration file:

```
save configuration secondary
```

save configuration as-script

```
save configuration as-script <script-name>
```

Description

Saves the running configuration as a script.

Syntax Description

script-name	Specifies the name of the file to save the configuration to. The script file is known as the XOS script file and uses the .xsf file extension.
-------------	--

Default

N/A

Usage Guidelines

The `save configuration as-script` command allows the user to save the current configuration as a script and export it out of the box for later use.

Example

The following command saves a running ASCII-formatted configuration named **primary.xsf**.

```
save configuration as-script primary.xsf
```

show configuration

```
show configuration {<module-name>} {detail}
```

Description

Displays the current configuration for the system or the specified module.

Syntax Description

module-name	Specifies the name of configuration module. The term <i>configuration module</i> refers to a feature in the firmware. By displaying a module, you can view the commands used to configure that feature. For example, to display all of the configurations that you made for only STP, specify the stp as the module-name.
detail	Displays configuration data including default. If the detail option is not specified, only the configuration changes you made to the factory defaults are shown.

Default

N/A.

Usage Guidelines

If the output scrolls off the top of the screen, you can use the `enable clipaging` command to pause the display when the output fills the screen. The default for clipaging is enabled.

NETGEAR recommends using the `show configuration` command to view on the CLI your currently running switch configuration. These files have the .cfg file extension. Do not use a text editor to view or modify your XML-based switch configuration files.

To save the configuration file as an ASCII-formatted file, and to view it with a text editor, see the `upload configuration [<hostname> | <ipaddress>] <filename> {vr <vr-name>}` and the `load script <filename> {arg1} {arg2} ... {arg9}` commands.

When you specify `show configuration` only, the switch displays configuration information for each of the switch modules excluding the default data.

You can display only the configuration of a module of interest by using the `module-name` keyword. For example, some of the modules are AAA, ACL, BGP, FDB, SNMP, and VLAN. Use TAB-completion to see a list. For example, to display the OSPF configuration of a user VR, the “module-name” is the name of the process, that is, `show configuration ospf-3`.

You must have administrator access to view the output of the `show configuration` command.

Depending on the software version running on your switch, the configurations on your switch, and the type of switch you have, additional or different configuration information may be displayed.

Example

This command shows the current configuration of the OSPF module in the switch:

```
show configuration ospf
```

The following is sample output from this command:

```
# Module ospf configuration.
#
configure ospf routerid automatic
configure ospf spf-hold-time 3
configure ospf metric-table 10M 10 100M 5 1G 4 10G 2
configure ospf lsa-batch-interval 30
configure ospf import-policy none
configure ospf ase-limit 0
disable ospf originate-default
disable ospf use-ip-router-alert
disable ospf
configure ospf restart none
configure ospf restart grace-period 120
disable ospf export direct
disable ospf export static
disable ospf export rip
disable ospf export e-bgp
disable ospf export i-bgp
configure ospf area 0.0.0.0 external-filter none
configure ospf area 0.0.0.0 interarea-filter none
```

show memorycard

```
show memorycard
```

Description

Displays whether an external memory card is present in the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to verify if there is an external memory card in the switch.

Example

This command shows whether a memory card is present on the switch:

```
show memorycard
```

If you do not have an external memory card installed, the output is similar to the following:

```
Memorycard is not present.
```

show script output autoexec

```
show script output autoexec
```

Description

Shows the results of executing the autoexec script.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

Use this command to show results when a autoexec.xsf file is executed. The file is not executed when a default.xsf file has been executed.

The CLI script file autoexec.xsf is executed after the configuration has been loaded. Its purpose is to run some commands after every reboot. It can also be used to revert to the original configuration following changes made by UPM executed persistent commands.

Example

This command shows the results of executing the autoexec script

```
show script output autoexec
```

When there is no autoexec.xsf file, there is no response.

show script output default

```
show script output default
```

Description

Shows the results of executing default.xsf on bootup.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

Use this command to show results when a default.xsf file is loaded.

An existing default.xsf file is executed if the switch comes up in an unconfigured state because the configuration file is missing, or the configuration file cannot be determined due to a corrupt NVRAM or other problems. This returns the switch to some basic configuration. When default.xsf is executed, the `show switch` command shows default.xsf as the booted configuration file.

Example

This command shows the results of executing the autoexec script

```
show script output default
```

When there is no default.xsf file, there is no response

synchronize

```
synchronize
```

Description

The `synchronize` command replicates all saved images and configurations from the primary MSM/MM or node to the backup MSM/MM or target node on a switch.

Default

N/A.

Usage Guidelines

This command:

1. Reboots the backup MSM/MM or target node to prepare it for synchronizing with the primary MSM/MM or node
2. Performs a binary copy of the primary MSM/MM or node to the backup MSM/MM or target node, including the primary and secondary software images, all configurations and policies, and temporary files
3. Reboots the backup MSM/MM or target node after replication is complete

During a synchronization, half of the switch fabric is lost. When the primary MSM/MM or node finishes replicating its configurations and images to the backup MSM/MM or target node, the full switch fabric is restored.

When you install a backup MSM/MM, you are not prompted to synchronize the images and the configurations from the primary. If not synchronized, the backup uses its image and the primary's configuration. This image/configuration mismatch will likely cause the switch to operate differently after failover. Use the `synchronize` command to replicate all saved images and configurations from the primary to the backup.

If you have not saved your runtime configuration, you are prompted to save it when you use the `synchronize` command. A message similar to the following appears:

```
Do you want to save configuration changes to primary.cfg? (y or n)
```

Enter `y` to save the configuration and continue with synchronizing the MSMs/MMs. Enter `n` to cancel the operation. If you enter `y`, messages similar to the following appear:

```
Saving configuration on primary MSM ..... done!
Synchronizing configuration to backup MSM .. done!
```

After the configuration has been saved and replicated to the backup MSM/MM, synchronization begins.

After the initial reboot, if the backup MSM/MM is not available or does not respond within 120 seconds, the synchronize operation fails.

Use the `show switch {detail}` command to verify that the backup MSM/MM is in sync with the primary MSM/MM.

On a NETGEAR 8800 series switch, the I/O ports on the backup MSM go down when you synchronize the MSMs. When the primary MSM finishes replicating its configurations and images to the backup MSM, the I/O ports on the backup MSM come back up.

Example

The following example assumes you have already saved your runtime configuration.

The following command replicates all saved images and configurations from the master MSM to the backup MSM:

```
synchronize
```

After you enter `synchronize`, status messages similar to the following appear:

```
Synchronize will reboot the backup MSM, then overwrite all code images
and configs with a copy from the master MSM.
Synchronize currently requires NETGEAR 8800 version 11 or greater on
the backup MSM
```

```
DO NOT interrupt synchronize, the backup MSM may become unbootable!
```

```
OK to continue? (y/n) Yes
```

```
Rebooting Backup MSM...
```

```
NOTE: The command line is locked during synchronize
```

```
synchronizing...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
```

```
synchronizing XOS...
[=====] 100% XOS
```

Synchronize complete - rebooting backup MSM...

unconfigure switch

```
unconfigure switch {all}
```

Description

Returns the switch configuration to its factory default settings and reboots the switch.

Syntax Description

all	Specifies that the entire configuration should be changed to the default values, including the management IP address, failsafe, and the switch rebooted.
-----	--

Default

N/A.

Usage Guidelines

Use `unconfigure switch` to reset the configuration to factory defaults, but without erasing the configuration. This preserves users account information, date and time settings, configuration, and so on.

Include the parameter `all` to clear the entire current configuration, including all switch and NETGEAR 8800 parameters, and reboot using the last used image and factory default configuration.

The command `unconfigure switch all` does not clear licensing information. The license cannot be disabled once it is enabled on the switch.

Example

The following command preserves the entire current configuration (but does not reload the current configuration after the switch reboots) and reboots the switch using the last specified saved image and factory default configuration:

```
unconfigure switch all
```

uninstall image

```
uninstall image <fname> <partition> {msm <slotid>} {reboot}
```

Description

Uninstalls a NETGEAR 8800 software package.

Syntax Description

fname	Specifies the software package to uninstall.
partition	Specifies which partition the package was installed to: primary or secondary. Select primary to remove it from the primary partition and secondary to remove it from the secondary partition.
slotid	Specifies the MSM where the package should be uninstalled. <ul style="list-style-type: none"> • A specifies the MSM/MM installed in slot A. • B specifies the MSM/MM installed in slot B.
reboot	Reboots the switch after the package is uninstalled.

Default

N/A.

Usage Guidelines

Use this command to uninstall a software package previously installed on the switch.

When you uninstall a software package, the switch prompts you to save your changes to your currently active configuration file:

```
Uninstallation of the EXOS module
Do you want to save configuration changes to primary.cfg? (y or n)
```

Enter *y* to save the changes to your configuration file. Enter *n* to not save the changes to your configuration file.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements previously described.

Example

The following command uninstalls the software package *NG8800-12.4.3.5-1-4.xos* from the secondary partition:

```
uninstall image NG8800-12.4.3.5-1-4.xos secondary
```


upload configuration

```
upload configuration [<hostname> | <ipaddress>] <filename> {vr <vr-name>}
```

Description

Uploads the current configuration in ASCII format to a TFTP server on your network.

Syntax Description

hostname	Specifies the hostname of the TFTP server where you want to download the configuration file. You must have DNS enabled
ipaddress	Specifies the IP address of the TFTP server where you want to download the configuration file.
filename	Specifies a user-defined name for the configuration file. You must use the .xsf file extension when naming an ASCII-formatted configuration file.
vr-name	Specifies the name of the virtual router. By default the switch uses VR-Mgmt for this command.

Note: User-created VRs are supported only on the platforms listed for this feature in Appendix A in the *NETGEAR 8800 User Manual*.

Default

Uploads the current configuration in ASCII format immediately to a TFTP server.

Usage Guidelines

Specify the `ipaddress` or `hostname` parameters to upload the current, active configuration file from the switch to a TFTP server on the network. Use of the `hostname` parameter requires that DNS be enabled.

The uploaded ASCII file retains the CLI format. This allows you to do the following:

- Modify the configuration using a text editor, and later download a copy of the file to the same switch or to one or more different switches.
- Send a copy of the configuration file to NETGEAR Technical Support for problem-solving purposes.

This command is not applicable to XML-based configurations. Those files use the .cfg file extension.

If you want to view your configuration in ASCII format, use the .xsf file extension (known as the XOS script file) when you save the configuration file on the switch. This saves the XML-based configuration in an ASCII format readable by a text editor.

If you successfully upload the active configuration to the network TFTP server, the switch displays a message similar to the following:

```
Uploading meg_upload_config1.xsf to 10.10.10.10 ... done!
```

If the switch displays a timeout error message similar to the following:

```
failed!
Error: timeout
```

Make sure you entered the correct host name or IP address of the TFTP server

If the switch displays an unreachable network error similar to the following:

```
failed!
Error: Network is unreachable
```

Make sure you entered the correct virtual router. By default the switch uses *VR-Mgmt* for this command.

Summary of Steps

The following summary only describes the CLI involved to transfer the configuration and load it on the switch; it is assumed that you know how to modify the configuration file with a text editor. As previously described, to use these commands, use the .xsf file extension. These steps are not applicable to configurations that use the .cfg file extension.

To work with an ASCII-formatted configuration file, complete the following tasks:

1. Upload the configuration to a network TFTP server using the following command:

```
upload configuration [<hostname> | <ipaddress>] <filename> {vr <vr-name>}
```

After the configuration file is on the TFTP server, use a text editor to the desired changes.

2. Download the configuration from the TFTP server to the switch using one of the following commands:

```
tftp [<host-name> | <ip-address>] -g -r <remote-file>
```

```
tftp get [<host-name> | <ip-address>] <remote-file>
```

3. Verify the configuration file is on the switch using the following command:

```
ls
```

4. Load and restore the new configuration file on the switch using the following command:

```
load script <filename> {arg1} {arg2} ... {arg9}
```

5. Save the configuration to the configuration database so the switch can reapply the configuration after switch reboot using the following command:

```
save configuration {primary | secondary | <existing-config> | <new-config>}
```

When you save the configuration file, the switch automatically adds the .cfg file extension to the filename. This saves the ASCII configuration as an XML-based configuration file.

The following describes the steps in more detail.

Uploading the ASCII Configuration File To a TFTP Server

To upload the current switch configuration as an ASCII-based file to the TFTP server, use the `upload configuration` command and save the configuration with the .xsf file extension.

For example, to transfer the current switch configuration as an ASCII-based file named `meg_upload_config1.xsf` to the TFTP server with an IP address of 10.10.10.10, do the following:

```
upload configuration 10.10.10.10 meg_upload_config1.xsf
```

If you successfully upload the configuration to the TFTP server, the switch displays a message similar to the following:

```
Uploading meg_upload_config1.xsf to 10.10.10.10 ... done!
```

Downloading the ASCII Configuration File to the Switch

To download the configuration from the TFTP server to the switch, use the `tftp` command. For example, to retrieve the configuration file named `meg-upload_config1.xsf` from a TFTP server with an IP address of 10.10.10.10, you can use one of the following commands:

```
tftp 10.10.10.10 -g -r meg_upload_config1.xsf
tftp get 10.10.10.10 meg_upload_config1.xsf
```

If you successfully download the configuration to the switch, the switch displays a message similar to the following:

```
Downloading meg_upload_config1.xsf to switch... done!
```

Verifying that the ASCII Configuration File is on the Switch

To confirm that the ASCII configuration file is on the switch, use the `ls` command. The file with an .xsf extension is the ASCII configuration.

The following sample output contains an ASCII configuration file:

```
-rw-r--r--  1 root    0          98362 Nov  2 13:53 Nov022005.cfg
-rw-r--r--  1 root    0        117136 Dec 12 12:56 epicenter.cfg
-rw-r--r--  1 root    0           68 Oct 26 11:17 mcastgroup.pol
-rw-r--r--  1 root    0         21203 Dec 13 15:40 meg_upload_config1.xsf
-rw-r--r--  1 root    0        119521 Dec  6 14:35 primary.cfg
-rw-r--r--  1 root    0         96931 Nov 11 11:01 primary_11_11_05.cfg
-rw-r--r--  1 root    0         92692 Jul 19 16:42 secondary.cfg
```

Loading the ASCII Configuration File

After downloading the configuration file, you must load the new configuration on the switch. To load and restore the ASCII configuration file, use the `load script <filename> {arg1} {arg2} ... {arg9}` command. After issuing this command, the ASCII configuration quickly scrolls across the screen.

The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```

Instead of entering each command individually, the script runs and loads the CLI on the switch.

Saving the Configuration

After you load the configuration, save it to the configuration database for use by the switch. This allows the switch to reapply the configuration after a switch reboot. To save the configuration, use the `save configuration {primary | secondary | <existing-config> | <new-config>}` command.

When you save the configuration file, the switch automatically adds the .cfg file extension to the filename. This saves the ASCII configuration as an XML-based configuration file.

You can use any name for the configuration. For example, after loading the file `meg_upload_config1.xsf`, you need to save it to the switch. To save the configuration as `configuration1.cfg`, do the following:

```
save configuration configuration1
```

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a remote file, remember the requirements previously described.

Example

The following command uploads the current switch configuration as an ASCII-based file named configbackup.xsf to the TFTP server with an IP address of 10.10.10.10:

```
upload configuration 10.10.10.10 configbackup.xsf
```

If you successfully upload the configuration to the TFTP server, the switch displays a message similar to the following:

```
Uploading configbackup.xsf to 10.10.10.10 ... done!
```

use configuration

```
use configuration [primary | secondary | <file_name>]
```

Description

Configures the switch to use a previously saved configuration on the next reboot.

Syntax Description

primary	Specifies the configuration file named <i>primary.cfg</i> .
secondary	Specifies the configuration file named <i>secondary.cfg</i> .
file_name	Specifies an existing user-defined configuration file name (displays a list of available user-defined configuration files).

Default

N/A.

Usage Guidelines

XML-based configuration files have a .cfg file extension. When you enter the name of the file in the CLI, the system automatically adds the .cfg file extension.

Do not use this command with ASCII-formatted configuration files. Those configuration files have an .xsf file extension. For more information about using and saving ASCII-formatted configuration files see the `upload configuration [<hostname> | <ipaddress>] <filename> {vr <vr-name>}` and the `load script <filename> {arg1} {arg2} ... {arg9}` commands.

There is no special significance to the primary and secondary configurations. They are just conveniences to specify the files *primary.cfg* and *secondary.cfg*.

When you configure the switch to use a previously saved configuration, the switch displays the following message:

```
The selected configuration will take effect after the next switch reboot.
```

You can create a new configuration file by saving your current switch configurations and using that file on the next reboot. For example, to create a new configuration named *test1* based on your current CLI session and switch configurations, use the following command:

```
save configuration test1
```

Tracking and Displaying Switch Configuration Files

To keep track of your configuration file names, use the `ls` command to display the files saved on your switch. Files with the `.cfg` extension are configuration files. In addition, you can see a list of available configuration files when you use the `use configuration` command.

The following is sample output from this command (“test” and “XOS1” are the names of the user-created and defined configurations):

```
exsh.1 # use configuration
primary      Primary configuration file
secondary    Secondary configuration file
<file-name>  Configuration file name
"test"      "XOS1"
```

You can also use the `ls` command to display a list of the current configuration and policy files in the system.

Displaying the Active Configuration

To view the currently active, running configuration, use the `show switch` command.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Example

The following command specifies that the next reboot should use the saved configuration file named *XOS1.cfg*:

```
use configuration XOS1
```

The following command specifies that the next reboot should use the configuration saved in the primary partition:

```
use configuration primary
```

use image

```
use image {partition} <partition> {msm <slotid>}
```

Description

Configures the switch to use a saved image on the next reboot.

Syntax Description

partition	Specifies which image to use on the next reboot, the one stored on the primary partition, or the one stored on the secondary partition.
slotid	Specifies the MSM/MM where the package should be uninstalled. <ul style="list-style-type: none"> • A specifies the MSM/MM installed in slot A. • B specifies the MSM/MM installed in slot B.

Default

The currently booted image.

Usage Guidelines

This command specifies which image to use on the next reboot. Two images can be stored, one on the primary partition, one on the secondary partition. To view your current (active) partition and the selected partition for the next reboot or installation, use the following command:

```
show switch
```

Output from this command includes the selected and booted images and if they are in the primary or the secondary partition. Primary indicates the saved image in the primary partition; secondary indicates the saved image in the secondary partition.

Example

Using TFTP

The following command configures the switch to use the image stored in the primary partition on the next reboot:

```
use image partion primary
```

A message similar to the following is displayed:

```
To take effect of partition change please reboot the switch!
```


Troubleshooting Commands

B

If you encounter problems when using your switch, NETGEAR 8800 provides troubleshooting commands. Use these commands only under the guidance of NETGEAR technical personnel.

This appendix describes commands for troubleshooting your switch, including:

- Running diagnostics and displaying diagnostic test results
- Enabling and disabling debug mode for Event Management System (EMS) components

You can contact NETGEAR Technical Support at 1-888-NETGEAR (US and Canada only; for other countries see the Support information card).

Event Management System

The Event Management System (EMS) provides enhanced features to filter and capture information generated on a switch. Details of using EMS are discussed in the *NETGEAR 8800 User Manual*, in the “Status Monitoring and Statistics” chapter, and the commands used for EMS are detailed in this document in [Chapter 8, “Commands for Status Monitoring and Statistics.”](#)

Included in this chapter are the EMS commands to enable and disable debug mode for EMS components.

configure debug core-dumps

```
configure debug core-dumps [internal-memory | memorycard | off]
```

Description

Enables or disables sending debug information to the specified memory card.

Syntax Description

internal-memory	Specifies that saving debug information to the internal memory card is enabled. This is the default behavior. Use this parameter only under the guidance of NETGEAR Technical Support personnel.
-----------------	--

memorycard	Specifies that saving debug information to the external memory card is enabled. Use this parameter only under the guidance of NETGEAR Technical Support personnel.
------------	--

Note: This parameter is available only on the NETGEAR 8800.

off	Specifies that saving debug information to the memory card is disabled.
-----	---

Default

By default, `configure debug core-dumps internal-memory` is enabled.

Usage Guidelines

Note: Use this command only under the guidance of NETGEAR Technical Support personnel to troubleshoot the switch.

The switch sends debug information to the preinstalled internal memory card. On the 8800 switch, you can also send debug information to the external memory card installed in the external compact flash slot of the MSM/MM.

The switch only generates core dump files and writes them to the specified memory card in the following situations:

- If a NETGEAR 8800 process fails.
- When forced under the guidance of NETGEAR Technical Support.

If you configure the switch to write core dump files to the internal memory and attempt to download a new software image, you might have insufficient space to complete the image download. If this occurs, move or delete the core dump files from the internal memory. For example, if you have an 8800 with an external memory card installed with space available, transfer the files to the external memory card. Transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

Before you can enable and save debug information to the external memory card, you must install an external compact flash memory card into the external compact flash slot of the MSM/MM. For more information about installing an external compact flash memory card, refer to the hardware documentation.

After you use the `eject memorycard` command and manually remove the card from the external compact flash slot of the MSM/MM, this setting is automatically changed to `off`.

Example

The following example enables an 8800 switch to save debug information to the external memory card:

```
configure debug core-dumps memorycard
```

The following example enables the switch to save debug information to the internal memory card:

```
configure debug core-dumps internal-memory
```

configure forwarding hash-algorithm

```
configure forwarding hash-algorithm [crc16 | crc32] {dual-hash [on | off]}
```

Description

Modifies hardware table utilization by configuring the hash algorithm or dual-hash settings.

Syntax Description

crc16	Specifies the CRC16 hash algorithm.
crc32	Specifies the CRC32 hash algorithm. This is the default setting.
on	Specifies that the dual-hash feature be turned on for the L3 Hash Table for hardware with dual-hash capability. With dual-hash on, each hash bucket is divided into two half-buckets with independent hash algorithms. One half-bucket uses the configured hash algorithm (CRC16 or CRC32), and the other half-bucket uses an alternate hash algorithm. This is the default setting.
off	Specifies that the dual-hash feature be turned off, even for hardware with dual-hash capability.

Default

The default hash algorithm is crc32.

The dual-hash default is “on.”

Usage Guidelines

Note: Modify the hardware table hash algorithm only with the guidance of NETGEAR technical personnel.

The switch uses a hash algorithm to decide where to store the addresses in the hardware table. The standard, default hash algorithm works well for most systems; however, for some addresses with certain patterns, the hardware may attempt to store address information in the same section of the hardware.

Example

The following command modifies the hardware table hash algorithm to crc16:

```
configure forwarding hash-algorithm crc16
```

The switch displays the following message to describe the change and to prompt you to save your configuration and reboot the switch:

```
Configured hash alorithm has been changed to 'crc16' with L3 dual-hash support 'on' for applicable HW.
```

```
Warning: This command will only take effect after a save and reboot
```

The following command disables dual-hashing on NETGEAR 8800 I/O modules.

```
configure forwarding hash-algorithm crc32 dual-hash off
```

The switch displays the following message:

```
Configured hash algorithm has been changed to 'crc32' with L3 dual-hash support 'off' for applicable HW.
```

```
Warning: This command will only take effect after a save and reboot.
```

To display the results, use the `show forwarding configuration` command.

configure forwarding hash-recursion-level

```
configure forwarding hash-recursion-level <0-3>
```

Description

Modifies hardware table utilization by configuring the dual hashing recursion level.

Syntax Description

0-3	Sets the maximum number of L3 hash buckets to modify to make room for a new entry.
-----	--

Default

The default is “1.”

Usage Guidelines

This command allows you to select the dual hashing “recursion level” for hardware with the dual-hash feature. The setting applies only if dual-hash is configured or defaulted to “on” using the `configure forwarding hash-algorithm` command.

The configured recursion level is the maximum number of existing hash entries to move in an attempt to add a new hash entry. A higher recursion level may provide better hash utilization at the expense of additional CPU processing. This command does not require a system reboot. However, the new recursion level takes effect only for addresses added after the command is issued.

Example

The following command modifies the dual-hash recursion level to modify up to two L3 hash buckets in an attempt to add a new entry:

```
configure forwarding hash-recursion-level 2
```

disable log debug-mode

```
disable log debug-mode
```

Description

Disables debug mode. The switch stops generating debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of `debug-summary`, `debug-verbose`, or `debug-data` when configuring filters
- Target format options `process-name`, `process-id`, `source-function`, and `source-line`

Example

The following command disables debug mode:

```
disable log debug-mode
```

eject memorycard

```
eject memorycard
```

Description

Ensures that the external memory card can be safely and manually removed from the external compact flash slot on the MSM/MM.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

After the switch writes to the external memory card, and before you can view the contents on the card, you must ensure it is safe to remove the card from the external compact flash slot on the MSM/MM. Use the `eject memorycard` command to prepare the card for removal. After you issue the `eject memorycard` command, you can manually remove the card from the external compact flash slot on the MSM/MM and read the data on the card.

If the `configure debug coredumps memorycard` command is in effect when you issue the `eject memorycard` command, the behavior is similar to issuing the `configure debug coredumps off` command.

For more information about removing the external memory card, please refer to the hardware documentation.

To access and read the data on the card, use a PC with appropriate hardware such as a compact flash reader/writer and follow the manufacturer's instructions to access the compact flash card and read the data.

Example

The following command prepares the external memory card to be removed from the external compact flash slot on the MSM/MM:

```
eject memorycard
```

enable log debug-mode

```
enable log debug-mode
```

Description

Enables debug mode. The switch generates debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by

technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of `debug-summary`, `debug-verbose`, or `debug-data` when configuring filters
- Target format options `process-name`, `process-id`, `source-function`, and `source-line`

Example

The following command enables debug mode:

```
enable log debug-mode
```

When you enable debug mode, the following message appears:

```
WARNING: Debug mode should only be enabled when advised by technical support,
or when advanced diagnosis is required. Performance degradation is possible.
Debug mode now enabled.
```

nslookup

```
nslookup {IPv4 | IPv6} <host-name>
```

Description

Displays the IP address of the requested host.

Syntax Description

IPv4	Lookup only IPv4 address(es)
IPv6	Lookup only IPv6 address(es)
host-name	Specifies the hostname.

Default

Lookup both IPv4 and IPv6 addresses.

Usage Guidelines

For `nslookup` to work, you must configure the DNS client, and the switch must be able to reach the DNS server.

By default, the command looks for both IPv4 and IPv6 addresses and reports an error only when neither an IPv4 address nor an IPv6 address is found for the host.

If the IPv4 or IPv6 option is specified, DNS lookup happens only for that address type, and an error is reported when no address of that type is found.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Example

The following command looks up the IP addresses of a computer with the name `myhost.mydomain` that has 2 IPv4 addresses and 1 IPv6 address:

```
nslookup myhost.mydomain
```

The following is sample output from the command on an XCM8800 switch:

```
Host "myhost.mydomain" has the IPv4 address 192.168.1.1
Host "myhost.mydomain" has the IPv4 address 192.168.1.2
Host "myhost.mydomain" has the IPv6 address 2000::1
```

run diagnostics

```
run diagnostics [extended | normal] {slot [<slot> | A | B]}
```

Description

Runs normal or extended diagnostics on the switch, slot, node, or management module.

Syntax Description

extended	Runs an extended diagnostic routine. Takes the ports offline, and performs extensive ASIC and packet loopback tests on all of the ports. If you have a Power over Ethernet (PoE) module installed, the switch also performs an extended PoE test, which tests the functionality of the inline power adapter.
normal	Runs a normal diagnostic routine. Takes the ports offline, and performs a simple ASIC and packet loopback test on all of the ports.
slot	Specifies the slot number of an I/O module.

A B	Specifies which MSM/MM to run diagnostics on.
	<ul style="list-style-type: none"> • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

Depending on your platform, use this command to run diagnostics on the switch, slot, management module, or stack port.

On an I/O module, the extended diagnostic routine can require significantly more time to complete, depending on the number of ports on the module.

On a management module, the module is taken offline while the diagnostics test is performed. Once the diagnostic test is completed, the MSM reboots, and becomes operational again.

Note: NETGEAR 8810 switch—If you run diagnostics on slots 5 and 6 with an MSM installed in those slots, the diagnostic routine tests the I/O subsystem of the MSM.

NETGEAR 8806 switch—if you run diagnostics on slots 3 and 4 with an MSM installed in those slots, the diagnostic routine tests the I/O subsystem of the MSM.

NETGEAR 8800 series switches—To run diagnostics on the management portion of the master MSM, specify slot A or B.

Viewing Diagnostics

To view results of the last diagnostics test run, use the following command:

```
show diagnostics {slot [<slot> | A | B]}
```

If the results indicate that the diagnostic failed on a module, replace the module with another module of the same type.

If the results indicate that the diagnostic failed on the switch, contact NETGEAR Technical Support.

Example

NETGEAR 8800 series switch example

The following command runs normal diagnostics on the I/O module installed in slot 2 of the NETGEAR 8800 series switch:

```
run diagnostics normal slot 2
```

The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter `y` to continue and run the diagnostics. Enter `n` to cancel the operation.

save debug tracefiles memorycard

```
save debug tracefiles memorycard
```

Description

Copies debug information to the external memory card.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Note: Use this command only under the guidance of NETGEAR Technical Support to troubleshoot the switch.

Use this command to copy debug information to the specified memory card. The debug information includes log files and trace files.

Progress messages are displayed that indicate the file being copied and when the copying is finished.

You can use the `upload debug [<hostname> | <ipaddress>] {{vr} <vrname>}` command to copy debug information to a network TFTP server.

Example

The following command copies debug information to the installed external memory card:

```
save debug tracefiles memorycard
```

show debug

```
show debug
```

Description

This command displays the status of writing core dump files to the specified memory card.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

By default, the switch does not write core dump files to the memory card. Use this command to verify if you are writing core dump files to the memory card. By default, the switch does not write core dump files to the memory card.

To configure the switch to write core dump files to the internal memory card, use the `configure debug core-dumps internal-memory` command.

To configure a switch to write core dump files to the external memory card, use the `configure debug core-dumps memorycard` command.

Example

The following example shows if the switch is sending core dump files to the specified memory card:

```
show debug
```

The following sample output shows that the switch is sending core dump files to the internal memory card:

```
Debug Settings:
  Core dumps: Enabled (internal-memory)
```

The following sample output shows that the switch is not sending core dump files to the specified memory card (this is the default behavior):

```
Debug Settings:
  Core dumps: Disabled
```

The following sample output shows that the switch is sending core dump files to the external memory card:

```
Debug Settings:
  Core dumps: Enabled (memorycard)
```

show diagnostics

```
show diagnostics {slot [<slot> | A | B]}
```

Description

Displays the status of the last diagnostic test run on the switch.

Syntax Description

slot	Specifies which I/O module to display diagnostic status information on.
A B	Specifies which MSM/MM to display diagnostic status information on: <ul style="list-style-type: none"> • A specifies the MSM installed in slot A. • B specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

Use this command to display information from the last diagnostic test run on the switch.

Output on the NETGEAR 8800 Series Switches

The switch displays the following diagnostic information:

If you have run diagnostics, a brief summary of the overall diagnostic test is displayed.

Options are:

- Date the test was run—The month, date, and year.
- Last Test Version—The firmware version associated with the results.
- Summary—A brief summary of the overall diagnostic test. Options are:
 - Diagnostics Pass—The diagnostic test has passed.
 - Diagnostics Fail—One or more diagnostic test has failed.
 - Diagnostics Interrupted—The diagnostic test was interrupted on the I/O module due to initiating an MSM failover.

If you have never run diagnostics on a specific slot, the switch displays a message similar to the following:

```
Slot-6: XCM8808X
No Diagnostics Data
```

If you attempt to view diagnostics information for a slot that does not have a module installed, the switch displays a message similar to the following:

```
Slot-7: No Module Present
```

Additional Guidelines Applicable to NETGEAR 8800

If you use the `show diagnostics {slot [<slot> | A | B]}` command on a slot where diagnostics have not been run, the switch displays messages similar to the following:

```
No Diagnostics Data
```

or

```
Diagnostics never run
```

If you try to display diagnostic test information on a slot where no module is installed, the switch displays messages similar to the following:

```
No Module Present
```

or

```
No card in slot
```

Running Diagnostics

To run diagnostics on an I/O module or an MSM installed in a NETGEAR 8800 series switch, use the following command:

```
run diagnostics [extended | normal] {slot [<slot> | A | B]}
```

Depending on the software version running on your switch or your switch model, additional or different diagnostic information might be displayed. For more information, see the command `run diagnostics` on page 1352.

Example

The following command displays the results of module diagnostics for the I/O module in slot 2:

```
show diagnostics slot 2
```

The following is sample output from a NETGEAR 8800 series switch:

```
Slot-2: XCM8824F
Last Test Date: Wed May  6 23:57:17 2009
Last Test Version: 12.3.0.1
Summary: Diagnostics Pass
```

When the version is unknown, Last Test Version reads "Unknown."

show forwarding configuration

```
show forwarding configuration
```

Description

Displays the configured selection criteria for ECMP routes and load-sharing group ports and the hardware table settings, including the configured and current hash algorithm and dual-hash settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:

- **Configured hash algorithm**—The hash algorithm configured on the switch. After the configuration is saved and the switch is rebooted, the switch uses this hash algorithm.
- **Current hash algorithm**—The hash algorithm currently used by the switch.
- **Configured dual-hash setting**—Whether the dual-hash feature is configured 'on' or 'off' on the switch. After the configuration is saved and the switch is rebooted, the switch uses this setting.
- **Current dual-hash setting**—Whether the dual-hash feature is currently 'on' or 'off' on the switch.
- **Dual-Hash recursion level**—The current dual-hash recursion level; default is '1.'
- **Sharing criteria**—Current selection criterion used for ECMP route sharing as well as for load-sharing groups. Specifies which Layer 3 and Layer 4 information is used in the sharing hash algorithm. For more information, see the description for the [configure forwarding sharing \[L3 | L3_L4\]](#) command.
- **Group Table Compression**—Whether the group table compression is currently 'on' or 'off' on the switch.
- **Switching mode**—Whether the switching mode is currently set to 'cut-through' or 'store-and-forward.'

It is possible for the values of the configured and the current hash, or the configured and current dual-hash settings to be different. For example, if you modified the hash algorithm and have not saved the configuration and rebooted the switch, the values might be different. In this situation, the switch also displays the following message:

NOTE: A save and reboot are required before the configured hash will take effect

Example

The following command displays the hardware forwarding algorithm configured on the switch:

```
show forwarding configuration
```

The following is sample output from this command on a NETGEAR 8800 series switch:

```
BD-8810.43 # show forwarding configuration
```

```
L2 and L3 Forwarding table hash algorithm:
  Configured hash algorithm:          crc32
  Current hash algorithm:             crc32
```

L3 Dual-Hash configuration: (Applies only to "c" and "xl"-series HW)

```
Configured setting:      on
Current setting:        on
Dual-Hash Recursion Level: 1
```

Hash criteria for IP unicast traffic for L2 load sharing and ECMP route sharing

```
Sharing criteria:      L3_L4
```

IP multicast:

```
Group Table Compression:  on
```

External lookup tables:

```
Configured Setting:      none
Current Setting:         none
```

Switch Settings:

```
Switching mode:          store-and-forward
```

Switch-fabric Settings:

```
Protocol:                enhanced
```

show tech

```
show tech {all | <area>} {detail} {logto [file]}
```

Description

Displays the output of various show commands to assist in monitoring and troubleshooting the switch; use only in conjunction with NETGEAR Technical support.

Syntax Description

all	Indicates all available show command output to be displayed.
area	Specifies one tech support area. For example, if you want to view STP information, enter <code>stp</code> .
detail	Specifies more detailed information.
logto [file]	Instructs the switch to log the show tech output into a file located in the switch's internal memory. The default file name is <code>show_tech.log.tgz</code> .

Default

N/A.

Usage Guidelines

Note: Use this command only under the guidance of NETGEAR Technical Support personnel to view your switch configurations and to troubleshoot the switch.

The `show tech` command displays the output of the following `show` commands, among others:

- `show bootprelay`
- `show configuration`
- `show dhcp-client state`
- `show diagnostics`
- `show memory`
- `show odometers`
- `show policy`
- `show port rxerror`
- `show port txerror`
- `show power`
- `show power budget`
- `show power controller`
- `show process`
- `show radius`
- `show session`
- `show switch`
- `show tacacs`
- `show version`
- `show vlan`

Information about the following areas is also displayed, among others:

- `aaa`
- `bootp`
- `cli`
- `stp`

If you enter the `detail` keyword, the following `show` output is displayed, among others:

- `show log`
- `show log configuration`
- `show log counters all`

- `show process detail`

This information can be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch, the configurations running on your switch, and the type of switch you have, additional or different `show` command and configuration output may be displayed.

Example

The following command displays the show command output on the switch:

```
show tech
```

top

```
top
```

Description

Displays real-time CPU utilization information by process.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show the percentage of CPU processing devoted to each process, sampled every 5 seconds.

You can change the display by typing a character while the display is active. [Table 31](#) displays the supported commands.

Table 31. TOP Interactive Command Display Options

Key	Action
P	Sort process list by CPU utilization
T	Sort process list by time usage
N	Sort process list by number (process ID)
M	Sort process list by memory usage
q [Ctrl] + c	Exit the top program

For more detailed information about the top command including display options, command fields, and command usage, please refer to your UNIX documentation.

Example

The following command displays the real-time CPU utilization information by process:

```
top
```

upload debug

```
upload debug [<hostname> | <ipaddress>] [{vr} <vrname>]
```

Description

Uploads debug information files to a tftp server. On a platform that has both primary and backup MSMs/MMs, debug information files are uploaded from both the backup and primary MSMs/MMs.

Syntax Description

hostname	Specifies the host name of the TFTP server to which the debug files will be uploaded to.
ipaddress	Specifies the IP address of the TFTP server to which the debug files will be uploaded to.
vrname	Specifies the name of the virtual router.

Default

By default, the virtual router *VR-Mgmt* will be used.

Usage Guidelines

Note: Use this command only under the guidance of NETGEAR Technical Support personnel to troubleshoot the switch.

Use this command to copy, upload debug information (for example, core, trace, show tech, configuration, and policy files) to the specified TFTP server.

Progress messages are displayed that indicate the file being copied and when the copying is finished. Depending on your platform, the switch displays a message similar to the following:

```
The following files on have been uploaded:
Tarball Name: TechPubsLab_C_09271428.tgz
./primary.cfg
```

You can also use this command in conjunction with the `show tech` command. Prior to uploading debug information files, the switch prompts you with the following message to run the `show tech` command with the `logto` file option:

```
Do you want to run show tech logto file first? (y/n)
```

Enter `y` to run the `show tech` command before uploading debug information. If you enter `y`, the `show_tech.log.tgz` file is included during the upload. Enter `n` to upload debug information without running the `show tech` command.

After you upload the debug information, you should see a compressed TAR file, which contains the debug information.

The TAR file naming convention is

```
<SysName>_<{<slot#>|A|B}I|C>_<Current Time>.tgz
- Current Time = mmddhhmm ( month(01-12), date(01-31), hour(0-24), minute(00-59) ).
```

Example

The following command uploads debug files to a network TFTP server:

```
upload debug 10.10.10.10
```

Command List

check policy attribute.....	482
check policy.....	481
clear access-list counter.....	487
clear access-list meter.....	488
clear account logout.....	14
clear bgp flap-statistics.....	1075
clear bgp neighbor counters	1076
clear counters fdb mac-tracking.....	453
clear counters ports.....	191
clear counters stp.....	742
clear counters vrrp.....	793
clear counters	332
clear cpu-monitoring.....	158
clear fdb.....	453
clear igmp group	1163
clear igmp snooping.....	1164
clear inline-power stats ports.....	300
clear iparp.....	814
clear ip-security anomaly-protection notify cache.....	555
clear ip-security arp validation violations.....	555
clear ip-security dhcp-snooping entries	556
clear ip-security source-ip-lockdown entries ports.....	556
clear lacp counters.....	192
clear license-info.....	14
clear log counters.....	333
clear log.....	333
clear mld group	1242
clear mld snooping.....	1242
clear msdp counters.....	1261
clear msdp sa-cache.....	1262
clear neighbor-discovery cache.....	904
clear netlogin state mac-address.....	680
clear netlogin state.....	679
clear ospf counters.....	995
clear ospfv3 counters.....	1037
clear pim cache	1164
clear pim snooping.....	1165
clear rip counters.....	951
clear ripng counters.....	973
clear session	15
clear slot.....	192

clear sys-recovery-level.....	334
clear vlan dhcp-address-allocation.....	557
configure access-list add.....	490
configure access-list delete	491
configure access-list rule-compression port-counters	492
configure access-list vlan-acl-precedence	493
configure access-list zone	493
configure access-list.....	489
configure account encrypted	17
configure account password-policy char-validation.....	18
configure account password-policy history	19
configure account password-policy lockout-on-login-failures.....	20
configure account password-policy max-age	21
configure account password-policy min-length.....	22
configure account	16
configure banner	23
configure bgp add aggregate-address.....	1077
configure bgp add confederation-peer sub-AS-number.....	1078
configure bgp add network.....	1079
configure bgp as-display-format	1080
configure bgp as-number.....	1080
configure bgp cluster-id.....	1081
configure bgp confederation-id	1081
configure bgp delete aggregate-address.....	1082
configure bgp delete confederation-peer sub-AS-number	1083
configure bgp delete network	1084
configure bgp export shutdown-priority.....	1084
configure bgp import-policy	1085
configure bgp local-preference	1086
configure bgp maximum-paths	1087
configure bgp med	1087
configure bgp neighbor allowas-in.....	1088
configure bgp neighbor dampening.....	1089
configure bgp neighbor description.....	1090
configure bgp neighbor dont-allowas-in.....	1091
configure bgp neighbor maximum-prefix.....	1092
configure bgp neighbor next-hop-self.....	1093
configure bgp neighbor no-dampening.....	1094
configure bgp neighbor password.....	1095
configure bgp neighbor peer-group.....	1096
configure bgp neighbor route-policy.....	1096
configure bgp neighbor route-reflector-client	1097
configure bgp neighbor send-community.....	1098
configure bgp neighbor shutdown-priority	1099
configure bgp neighbor soft-reset.....	1100
configure bgp neighbor source-interface	1101
configure bgp neighbor timer	1102
configure bgp neighbor weight	1102

configure bgp peer-group allowin-as.....	1103
configure bgp peer-group dampening.....	1104
configure bgp peer-group dont-allowin-as	1105
configure bgp peer-group maximum-prefix	1106
configure bgp peer-group next-hop-self	1107
configure bgp peer-group no-dampening.....	1108
configure bgp peer-group password.....	1109
configure bgp peer-group remote-AS-number	1109
configure bgp peer-group route-policy.....	1110
configure bgp peer-group route-reflector-client.....	1111
configure bgp peer-group send-community	1112
configure bgp peer-group soft-reset.....	1113
configure bgp peer-group source-interface	1114
configure bgp peer-group timer.....	1115
configure bgp peer-group weight.....	1115
configure bgp restart address-family.....	1117
configure bgp restart restart-time.....	1118
configure bgp restart stale-route-time	1119
configure bgp restart update-delay.....	1119
configure bgp restart.....	1116
configure bgp routerid.....	1120
configure bgp soft-reconfiguration.....	1121
configure bootprelay add.....	815
configure bootprelay delete	816
configure bootprelay dhcp-agent information check.....	816
configure bootprelay dhcp-agent information circuit-id port-information	817
configure bootprelay dhcp-agent information circuit-id vlan-information	818
configure bootprelay dhcp-agent information option	818
configure bootprelay dhcp-agent information policy	819
configure bootprelay dhcp-agent information remote-id.....	820
configure cli max-failed-logins.....	25
configure cli max-sessions	24
configure debug core-dumps	1345
configure diffserv examination code-point qosprofile	520
configure diffserv replacement code-point	521
configure dns-client add.....	25
configure dns-client default-domain.....	26
configure dns-client delete	27
configure dos-protect acl-expire.....	558
configure dos-protect interval.....	558
configure dos-protect trusted ports.....	559
configure dos-protect type l3-protect alert-threshold	560
configure dos-protect type l3-protect notify-threshold.....	560
configure dot1p replacement	522
configure dot1p type.....	523
configure failsafe-account	28
configure fdb agingtime.....	454
configure fdb mac-tracking ports.....	455

configure firmware	1306
configure flow-redirect add nexthop.....	495
configure flow-redirect delete nexthop.....	495
configure flow-redirect health-check.....	496
configure flow-redirect nexthop.....	497
configure flow-redirect no-active.....	497
configure flow-redirect vr	498
configure forwarding hash-algorithm.....	1347
configure forwarding hash-recursion-level	1348
configure forwarding ipmc compression.....	1165
configure forwarding sharing	821
configure idletimeout.....	29
configure igmp snooping filters.....	1167
configure igmp snooping flood-list.....	1168
configure igmp snooping forwarding-mode	1170
configure igmp snooping leave-timeout.....	1171
configure igmp snooping timer	1171
configure igmp snooping vlan ports add static group	1172
configure igmp snooping vlan ports add static router	1173
configure igmp snooping vlan ports delete static group	1174
configure igmp snooping vlan ports delete static router	1175
configure igmp snooping vlan ports filter	1175
configure igmp snooping vlan ports set join-limit.....	1177
configure igmp ssm-map add	1177
configure igmp ssm-map delete	1178
configure igmp	1166
configure inline-power budget	301
configure inline-power disconnect-precedence	302
configure inline-power label ports.....	303
configure inline-power operator-limit ports.....	304
configure inline-power priority ports.....	305
configure inline-power usage-threshold.....	306
configure iparp add proxy.....	823
configure iparp add.....	822
configure iparp delete proxy	824
configure iparp delete.....	824
configure iparp max_entries	825
configure iparp max_pending_entries.....	826
configure iparp max_proxy_entries.....	826
configure iparp timeout.....	827
configure ipforwarding originated-packets	828
configure ipmcforwarding.....	1179
configure ipmroute add.....	1179
configure ipmroute delete.....	1180
configure ip-mtu vlan.....	193
configure iproute add (IPv4)	828
configure iproute add (IPv6)	905
configure iproute add (Multicast).....	1181

configure iproute add blackhole ipv4 default.....	830
configure iproute add blackhole ipv6 default.....	906
configure iproute add blackhole.....	829
configure iproute add blackhole.....	906
configure iproute add default.....	831
configure iproute add default.....	907
configure iproute delete blackhole ipv4 default.....	833
configure iproute delete blackhole ipv6 default.....	909
configure iproute delete blackhole.....	833
configure iproute delete blackhole.....	909
configure iproute delete default.....	834
configure iproute delete default.....	910
configure iproute delete.....	1182
configure iproute delete.....	832
configure iproute delete.....	908
configure iproute ipv6 priority.....	911
configure iproute priority	835
configure iproute reserved-entries.....	837
configure iproute sharing max-gateways	838
configure ip-security anomaly-protection notify cache	562
configure ip-security anomaly-protection notify rate limit.....	562
configure ip-security anomaly-protection notify rate window.....	563
configure ip-security anomaly-protection notify trigger off.....	564
configure ip-security anomaly-protection notify trigger on.....	564
configure ip-security anomaly-protection tcp.....	565
configure ip-security anomaly-protection	
icmp ipv4-max-size.....	561
configure ip-security anomaly-protection	
icmp ipv6-max-size.....	561
configure ip-security dhcp-binding storage filename	570
configure ip-security dhcp-binding storage location	570
configure ip-security dhcp-bindings add	569
configure ip-security dhcp-bindings delete	569
configure ip-security dhcp-bindings storage.....	571
configure ip-security dhcp-snooping information check	565
configure ip-security dhcp-snooping information circuit-id port-information port	566
configure ip-security dhcp-snooping information circuit-id vlan-information	
567	
configure ip-security dhcp-snooping information option.....	567
configure ip-security dhcp-snooping information policy.....	568
configure irdp	839
configure jumbo-frame-size	194
configure lacp member-port priority.....	195
configure lldp med fast-start repeat-count.....	269
configure lldp ports management-address	270
configure lldp ports port-description.....	271
configure lldp ports system-capabilities	271

configure lldp ports system-description	272
configure lldp ports system-name	273
configure lldp ports vendor-specific dot1 port-protocol-vlan-ID	274
configure lldp ports vendor-specific dot1 port-vlan-ID	273
configure lldp ports vendor-specific dot1 vlan-name	275
configure lldp ports vendor-specific dot3 link-aggregation	276
configure lldp ports vendor-specific dot3 mac-phy	277
configure lldp ports vendor-specific dot3 max-frame-size	278
configure lldp ports vendor-specific dot3 power-via-mdi	278
configure lldp ports vendor-specific med capabilities	280
configure lldp ports vendor-specific med location-identification	281
configure lldp ports vendor-specific med policy application	282
configure lldp ports vendor-specific med power-via-mdi	284
configure lldp reinitialize-delay	286
configure lldp snmp-notification-interval	286
configure lldp transmit-delay	287
configure lldp transmit-hold	287
configure lldp transmit-interval	288
configure log display	336
configure log filter events match	340
configure log filter events	337
configure log target filter	343
configure log target format	346
configure log target match	350
configure log target severity	351
configure log target syslog	353
configure mac-lockdown-timeout ports aging-time	572
configure meter	524
configure mirror add ports anomaly	196
configure mirroring add	196
configure mirroring delete	199
configure mirroring mode	200
configure mld snooping flood-list	1247
configure mld snooping leave-timeout	1249
configure mld snooping timer	1250
configure mld snooping vlan ports add static group	1244
configure mld snooping vlan ports add static router	1246
configure mld snooping vlan ports delete static group	1245
configure mld snooping vlan ports delete static router	1246
configure mld	1243
configure msdp as-display-format	1263
configure msdp max-rejected-cache	1264
configure msdp originator-id	1265
configure msdp peer default-peer	1265
configure msdp peer description	1267
configure msdp peer mesh-group	1268
configure msdp peer no-default-peer	1268
configure msdp peer password	1269

configure msdp peer sa-filter	1270
configure msdp peer sa-limit	1271
configure msdp peer source-interface.....	1272
configure msdp peer timer	1273
configure msdp peer ttl-threshold.....	1274
configure msdp sa-cache-server	1274
configure mstp format.....	743
configure mstp region.....	744
configure mstp revision.....	745
configure mvr add receiver	1183
configure mvr add vlan.....	1184
configure mvr delete receiver	1184
configure mvr delete vlan	1185
configure mvr mvr-address	1185
configure mvr static group	1186
configure neighbor-discovery cache add	912
configure neighbor-discovery cache delete	913
configure neighbor-discovery cache max_entries.....	913
configure neighbor-discovery cache max_pending_entries	914
configure neighbor-discovery cache timeout.....	915
configure netlogin add mac-list.....	680
configure netlogin add proxy-port.....	682
configure netlogin agingtime	682
configure netlogin allowed-refresh-failures.....	683
configure netlogin authentication database-order.....	683
configure netlogin authentication failure vlan.....	684
configure netlogin authentication service-unavailable vlan.....	685
configure netlogin banner	687
configure netlogin base-url	688
configure netlogin delete mac-list.....	688
configure netlogin delete proxy-port.....	689
configure netlogin dot1x eapol-transmit-version.....	689
configure netlogin dot1x guest-vlan.....	690
configure netlogin dot1x timers.....	692
configure netlogin dynamic-vlan uplink-ports.....	695
configure netlogin dynamic-vlan.....	693
configure netlogin local-user security-profile.....	698
configure netlogin local-user	696
configure netlogin mac timers reauth-period.....	699
configure netlogin move-fail-action.....	699
configure netlogin port allow egress-traffic.....	700
configure netlogin ports mode	701
configure netlogin ports no-restart.....	704
configure netlogin ports restart.....	704
configure netlogin redirect-page.....	705
configure netlogin session-refresh.....	706
configure netlogin vlan.....	707
configure node priority.....	59

configure ospf add virtual-link	996
configure ospf add vlan area link-type	997
configure ospf add vlan area	996
configure ospf area add range	999
configure ospf area delete range	1000
configure ospf area external-filter	998
configure ospf area interarea-filter	998
configure ospf area normal	1000
configure ospf area nssa stub-default-cost	1001
configure ospf area stub stub-default-cost	1002
configure ospf area timer	1003
configure ospf ase-limit	1004
configure ospf ase-summary add	1005
configure ospf ase-summary delete	1005
configure ospf authentication	1006
configure ospf cost	1007
configure ospf delete virtual-link	1008
configure ospf delete vlan	1008
configure ospf import-policy	1009
configure ospf lsa-batch-interval	1009
configure ospf metric-table	1010
configure ospf priority	1010
configure ospf restart grace-period	1012
configure ospf restart	1011
configure ospf restart-helper	1013
configure ospf routerid	1013
configure ospf spf-hold-time	1014
configure ospf virtual-link timer	1015
configure ospf vlan area	1016
configure ospf vlan neighbor add	1017
configure ospf vlan neighbor delete	1017
configure ospf vlan timer	1018
configure ospfv3 add interface all	1040
configure ospfv3 add interface	1039
configure ospfv3 add virtual-link	1041
configure ospfv3 area add range	1042
configure ospfv3 area cost	1042
configure ospfv3 area delete range	1043
configure ospfv3 area external-filter	1044
configure ospfv3 area interarea-filter	1045
configure ospfv3 area normal	1046
configure ospfv3 area priority	1047
configure ospfv3 area stub	1048
configure ospfv3 area timer	1048
configure ospfv3 delete interface	1050
configure ospfv3 delete virtual-link	1051
configure ospfv3 import-policy	1051
configure ospfv3 interface area	1052

configure ospfv3 interface cost.....	1053
configure ospfv3 interface priority.....	1054
configure ospfv3 interface timer.....	1054
configure ospfv3 metric-table.....	1056
configure ospfv3 routerid.....	1057
configure ospfv3 spf-hold-time.....	1057
configure ospfv3 virtual-link timer.....	1058
configure pim add vlan.....	1187
configure pim border.....	1275
configure pim cbsr.....	1188
configure pim crp static.....	1189
configure pim crp timer.....	1190
configure pim crp vlan.....	1191
configure pim delete vlan.....	1192
configure pim register-checksum-to.....	1194
configure pim register-rate-limit-interval.....	1192
configure pim register-suppress-interval register-probe-interval ...	1193
configure pim spt-threshold.....	1194
configure pim ssm range.....	1195
configure pim state-refresh timer origination-interval.....	1197
configure pim state-refresh timer source-active-timer.....	1197
configure pim state-refresh ttl.....	1198
configure pim state-refresh.....	1196
configure pim timer vlan.....	1198
configure pim vlan trusted-gateway.....	1199
configure port ethertype.....	1296
configure ports auto off.....	200
configure ports auto on.....	202
configure ports auto-polarity.....	202
configure ports display-string.....	203
configure ports qosprofile.....	526
configure ports rate-limit egress.....	526
configure ports rate-limit flood.....	572
configure ports redundant.....	204
configure ports vlan.....	573
configure power supply.....	60
configure private-vlan add network.....	425
configure private-vlan add subscriber.....	426
configure private-vlan delete.....	426
configure protocol add.....	427
configure protocol delete.....	428
configure qosprofile egress.....	529
configure qosprofile ingress.....	531
configure qosprofile.....	527
configure qosscheduler.....	533
configure radius server client-ip.....	576
configure radius shared-secret.....	577
configure radius timeout.....	578

configure radius-accounting server client-ip	579
configure radius-accounting shared-secret	580
configure radius-accounting timeout.....	581
configure rip add vlan.....	951
configure rip delete vlan.....	952
configure rip garbagetime.....	952
configure rip import-policy	953
configure rip routetimeout	953
configure rip updatetime.....	954
configure rip vlan cost.....	954
configure rip vlan route-policy	955
configure rip vlan rxmode.....	956
configure rip vlan trusted-gateway.....	957
configure rip vlan txmode.....	957
configure ripng add.....	974
configure ripng cost.....	974
configure ripng delete.....	975
configure ripng garbagetime	976
configure ripng import-policy	976
configure ripng route-policy	977
configure ripng routetimeout	978
configure ripng trusted-gateway	979
configure ripng updatetime	980
configure safe-default-script	30
configure sflow agent ipaddress	354
configure sflow collector ipaddress.....	355
configure sflow max-cpu-sample-limit.....	355
configure sflow poll-interval	356
configure sflow ports sample-rate.....	357
configure sflow sample-rate	358
configure sharing add ports	205
configure sharing address-based custom.....	207
configure sharing delete ports	208
configure sharing health-check member-port add tcp-tracking.....	209
configure sharing health-check member-port delete tcp-tracking.....	210
configure sharing health-check member-port tcp-tracking	210
configure sharing lacp activity-mode.....	211
configure sharing lacp defaulted-state-action.....	212
configure sharing lacp system-priority.....	213
configure sharing lacp timeout	214
configure slot module.....	215
configure slot restart-limit	216
configure snmp access-profile	62
configure snmp add community	63
configure snmp add trapreceiver.....	64
configure snmp delete community.....	66
configure snmp delete trapreceiver.....	67
configure snmp sysContact.....	67

configure snmp sysLocation	68
configure snmp sysName.....	69
configure snmpv3 add access	70
configure snmpv3 add community.....	71
configure snmpv3 add filter	72
configure snmpv3 add filter-profile.....	73
configure snmpv3 add group user.....	74
configure snmpv3 add mib-view.....	75
configure snmpv3 add notify	76
configure snmpv3 add target-addr.....	77
configure snmpv3 add target-params.....	79
configure snmpv3 add user clone-from.....	82
configure snmpv3 add user	80
configure snmpv3 delete access.....	82
configure snmpv3 delete community.....	84
configure snmpv3 delete filter.....	84
configure snmpv3 delete filter-profile	85
configure snmpv3 delete group user.....	86
configure snmpv3 delete mib-view.....	87
configure snmpv3 delete notify.....	88
configure snmpv3 delete target-addr.....	89
configure snmpv3 delete target-params.....	89
configure snmpv3 delete user	90
configure snmpv3 engine-boots	91
configure snmpv3 engine-id	91
configure snmp-client update-interval.....	93
configure snmp-client.....	92
configure ssh2 key.....	582
configure sshd2 user-key add user.....	583
configure sshd2 user-key delete user.....	584
configure ssl certificate pregenerated	584
configure ssl certificate privkeylen.....	585
configure ssl privkey pregenerated.....	586
configure stpd add vlan.....	746
configure stpd default-encapsulation.....	748
configure stpd delete vlan	750
configure stpd forwarddelay	751
configure stpd hellotime.....	752
configure stpd maxage.....	752
configure stpd max-hop-count	753
configure stpd mode.....	754
configure stpd ports bpdu-restrict.....	756
configure stpd ports cost	756
configure stpd ports edge-safeguard disable.....	758
configure stpd ports edge-safeguard enable.....	759
configure stpd ports link-type.....	761
configure stpd ports mode	763
configure stpd ports port-priority.....	764

configure stpd ports priority	765
configure stpd ports restricted-role disable	766
configure stpd ports restricted-role enable	767
configure stpd priority	768
configure stpd tag	768
configure sys-health-check all level	358
configure sys-health-check interval	360
configure syslog add	368
configure syslog delete	370
configure sys-recovery-level slot	363
configure sys-recovery-level	361
configure tacacs server client-ip	587
configure tacacs shared-secret	588
configure tacacs timeout	589
configure tacacs-accounting server	589
configure tacacs-accounting shared-secret	590
configure tacacs-accounting timeout	591
configure telnet access-profile	94
configure telnet port	95
configure telnet vr	96
configure time	31
configure timezone	32
configure trusted-ports trust-for dhcp-server	592
configure trusted-servers add server	593
configure trusted-servers delete server	594
configure tunnel ipaddress	924
configure vlan add ports private-vlan translated	430
configure vlan add ports stpd	769
configure vlan add ports tagged private-vlan end-point	431
configure vlan add ports	429
configure vlan add secondary-ipaddress	840
configure vlan delete ports	432
configure vlan delete secondary-ipaddress	841
configure vlan dhcp-address-range	595
configure vlan dhcp-lease-timer	595
configure vlan dhcp-options	596
configure vlan ipaddress	432
configure vlan name	434
configure vlan netlogin-lease-timer	707
configure vlan protocol	434
configure vlan qosprofile	533
configure vlan router-discovery add prefix	915
configure vlan router-discovery default-lifetime	916
configure vlan router-discovery delete prefix	916
configure vlan router-discovery hop-limit	917
configure vlan router-discovery link-mtu	918
configure vlan router-discovery managed-config-flag	918
configure vlan router-discovery max-interval	919

configure vlan router-discovery min-interval	920
configure vlan router-discovery other-config-flag	921
configure vlan router-discovery reachable-time	921
configure vlan router-discovery retransmit-time	922
configure vlan router-discovery set prefix	923
configure vlan subvlan	841
configure vlan subvlan-address-range	842
configure vlan tag	435
configure vlan udp-profile	843
configure vman add ports	1296
configure vman delete ports	1298
configure vman ethertype	1299
configure vman tag	1300
configure vr add ports	472
configure vr add protocol	473
configure vr delete ports	474
configure vr delete protocol	474
configure vrrp vlan vrid add ipaddress	794
configure vrrp vlan vrid add track-iproute	794
configure vrrp vlan vrid add track-ping	795
configure vrrp vlan vrid add track-vlan	796
configure vrrp vlan vrid advertisement-interval	797
configure vrrp vlan vrid authentication	798
configure vrrp vlan vrid delete track-iproute	800
configure vrrp vlan vrid delete track-ping	800
configure vrrp vlan vrid delete track-vlan	801
configure vrrp vlan vrid delete	799
configure vrrp vlan vrid dont-preempt	802
configure vrrp vlan vrid preempt	802
configure vrrp vlan vrid priority	803
configure vrrp vlan vrid track-mode	804
cp	158
create access-list zone	500
create access-list	499
create account	36
create bgp neighbor peer-group	1121
create bgp neighbor remote-AS-number	1122
create bgp peer-group	1123
create fdb mac-tracking entry	456
create fdbentry vlan ports	456
create flow-redirect	501
create log filter	370
create meter	534
create msdp mesh-group	1276
create msdp peer	1277
create netlogin local-user	708
create ospf area	1019
create ospfv3 area	1059

create private-vlan	436
create protocol	437
create qosprofile	535
create snmp trap	97
create sshd2 key-file	597
create sshd2 user-key	597
create stpd	772
create tunnel 6to4	924
create tunnel ipv6-in-ipv4	925
create virtual-router	475
create vlan	438
create vman	1300
create vrrp vlan vrid	805
delete access-list zone	502
delete access-list	501
delete account	38
delete bgp neighbor	1123
delete bgp peer-group	1124
delete fdb mac-tracking entry	458
delete fdbentry	458
delete flow-redirect	502
delete log filter	371
delete meter	535
delete msdp mesh-group	1277
delete msdp peer	1278
delete netlogin local-user	710
delete ospf area	1019
delete ospfv3 area	1060
delete private-vlan	440
delete protocol	440
delete qosprofile	536
delete sshd2 user-key	598
delete stpd	773
delete traffic queue	536
delete tunnel	926
delete virtual-router	476
delete vlan	441
delete vman	1301
delete vrrp vlan vrid	806
disable access-list permit to-cpu	503
disable access-list refresh blackhole	504
disable bgp advertise-inactive-route	1125
disable bgp aggregation	1126
disable bgp always-compare-med	1126
disable bgp community format	1127
disable bgp export	1127
disable bgp fast-external-fallover	1129
disable bgp neighbor capability	1130

disable bgp neighbor originate-default	1131
disable bgp neighbor remove-private-AS-numbers.....	1131
disable bgp neighbor soft-in-reset.....	1132
disable bgp neighbor use-ip-router-alert	1133
disable bgp neighbor.....	1129
disable bgp peer-group capability.....	1134
disable bgp peer-group originate-default	1135
disable bgp peer-group remove-private-AS-numbers.....	1136
disable bgp peer-group soft-in-reset.....	1136
disable bgp peer-group use-ip-router-alert	1137
disable bgp peer-group.....	1134
disable bgp.....	1125
disable bootp vlan.....	845
disable bootprelay.....	846
disable cli space-completion	38
disable cli-config-logging	372
disable clipaging.....	39
disable cpu-monitoring.....	161
disable dhcp ports vlan.....	598
disable dhcp vlan.....	97
disable diffserv examination ports.....	537
disable diffserv replacement ports.....	538
disable dos-protect.....	599
disable dot1p examination ports.....	1302
disable dot1p examination ports.....	538
disable dot1p replacement ports.....	539
disable flooding ports.....	459
disable flow-control rx-pause ports.....	216
disable flow-control tx-pause ports.....	217
disable icmp address-mask	847
disable icmp parameter-problem.....	847
disable icmp port-unreachables	848
disable icmp redirects.....	849
disable icmp time-exceeded	849
disable icmp timestamp.....	850
disable icmp unreachable	851
disable icmp userredirects	851
disable idletimeout.....	40
disable igmp snooping vlan fast-leave.....	1202
disable igmp snooping.....	1201
disable igmp ssm-map.....	1202
disable igmp	1200
disable inline-power legacy slot.....	308
disable inline-power legacy	308
disable inline-power ports	309
disable inline-power slot	310
disable inline-power.....	307
disable iparp checking.....	852

disable iparp gratuitous protect vlan.....	600
disable iparp refresh.....	852
disable ipforwarding ipv6.....	927
disable ipforwarding.....	853
disable ipmcforwarding.....	1203
disable ip-option loose-source-route.....	854
disable ip-option record-route.....	854
disable ip-option record-timestamp.....	855
disable ip-option router-alert.....	855
disable ip-option strict-source-route.....	856
disable iproute compression.....	856
disable iproute ipv6 compression.....	928
disable iproute ipv6 sharing.....	927
disable iproute sharing.....	857
disable ip-security anomaly-protection icmp.....	603
disable ip-security anomaly-protection ip.....	601
disable ip-security anomaly-protection l4port.....	601
disable ip-security anomaly-protection notify.....	604
disable ip-security anomaly-protection tcp flags.....	602
disable ip-security anomaly-protection tcp fragment.....	602
disable ip-security anomaly-protection.....	600
disable ip-security arp gratuitous-protection.....	604
disable ip-security arp learning learn-from-arp.....	605
disable ip-security arp learning learn-from-dhcp.....	606
disable ip-security arp validation.....	607
disable ip-security dhcp-bindings restoration.....	608
disable ip-security dhcp-snooping.....	608
disable ip-security source-ip-lockdown ports.....	609
disable irdp.....	857
disable jumbo-frame ports.....	218
disable learning iparp sender-mac.....	461
disable learning port.....	218
disable learning port.....	461
disable lldp ports.....	289
disable log debug-mode.....	1349
disable log display.....	373
disable log target.....	373
disable loopback-mode vlan.....	441
disable mac-lockdown-timeout ports.....	609
disable mirroring.....	219
disable mld snooping.....	1251
disable mld.....	1250
disable msdp data-encapsulation.....	1279
disable msdp export local-sa.....	1280
disable msdp peer.....	1281
disable msdp process-sa-request.....	1281
disable msdp.....	1279
disable mvr.....	1203

disable neighbor-discovery refresh.....	928
disable netlogin authentication failure vlan ports	712
disable netlogin authentication service-unavailable vlan ports.....	712
disable netlogin dot1x guest-vlan ports	713
disable netlogin logout-privilege.....	713
disable netlogin ports.....	714
disable netlogin reauthenticate-on-refresh.....	715
disable netlogin redirect-page.....	715
disable netlogin session-refresh.....	715
disable netlogin.....	711
disable ospf capability opaque-lsa.....	1020
disable ospf export.....	1021
disable ospf originate-default.....	1022
disable ospf restart-helper-lsa-check.....	1023
disable ospf use-ip-router-alert.....	1023
disable ospf	1020
disable ospfv3 export.....	1061
disable ospfv3	1061
disable pim snooping.....	1204
disable pim ssm vlan.....	1205
disable pim.....	1204
disable port	219
disable radius	610
disable radius-accounting	611
disable rip aggregation.....	959
disable rip export.....	959
disable rip originate-default.....	960
disable rip poisonreverse	961
disable rip splithorizon.....	961
disable rip triggerupdates	962
disable rip use-ip-router-alert.....	962
disable rip.....	958
disable ripng export.....	981
disable ripng originate-default.....	982
disable ripng poisonreverse	982
disable ripng splithorizon	983
disable ripng triggerupdate	983
disable ripng	980
disable rmon	375
disable router-discovery.....	929
disable sflow ports.....	376
disable sflow	375
disable sharing	220
disable slot	221
disable smartredundancy.....	222
disable snmp access vr.....	99
disable snmp access.....	98
disable snmp community.....	99

disable snmp traps fdb mac-tracking.....	462
disable snmp traps lldp.....	289
disable snmp traps lldp-med	290
disable snmp traps port-up-down ports.....	222
disable snmp traps	100
disable snmpv3	100
disable snmp-client	101
disable ssh2.....	611
disable stpd auto-bind.....	774
disable stpd ports	775
disable stpd rapid-root-failover	776
disable stpd.....	774
disable subvlan-proxy-arp vlan	858
disable sys-health-check.....	377
disable syslog	378
disable tacacs	612
disable tacacs-accounting.....	613
disable tacacs-authorization	613
disable telnet	102
disable udp-echo-server.....	859
disable vlan.....	442
disable vrrp vrid	806
disable watchdog	102
disable web http	614
disable web https	614
disable xml-mode	162
download image	1308
download ssl certificate.....	615
download ssl privkey.....	616
edit policy.....	483
eject memorycard	1349
enable access-list permit to-cpu.....	504
enable access-list refresh blackhole.....	505
enable bgp advertise-inactive-route.....	1138
enable bgp aggregation.....	1139
enable bgp always-compare-med	1139
enable bgp community format	1140
enable bgp export	1140
enable bgp fast-external-fallover.....	1142
enable bgp neighbor capability	1143
enable bgp neighbor originate-default.....	1144
enable bgp neighbor remove-private-AS-numbers	1145
enable bgp neighbor soft-in-reset.....	1146
enable bgp neighbor use-ip-router-alert	1147
enable bgp neighbor	1142
enable bgp peer-group capability.....	1148
enable bgp peer-group originate-default	1149
enable bgp peer-group remove-private-AS-numbers.....	1150

enable bgp peer-group soft-in-reset.....	1150
enable bgp peer-group use-ip-router-alert	1151
enable bgp peer-group.....	1147
enable bgp.....	1137
enable bootp vlan.....	859
enable bootprelay.....	860
enable cli space-completion	40
enable cli-config-logging	378
enable clipaging	41
enable cpu-monitoring.....	162
enable dhcp ports vlan.....	618
enable dhcp vlan.....	103
enable diffserv examination ports.....	540
enable diffserv replacement ports.....	540
enable dos-protect simulated.....	619
enable dos-protect.....	618
enable dot1p examination port	1302
enable dot1p examination ports.....	541
enable dot1p replacement ports.....	542
enable flooding ports.....	462
enable flow-control rx-pause ports.....	223
enable flow-control tx-pause ports.....	224
enable icmp address-mask.....	861
enable icmp parameter-problem.....	861
enable icmp port-unreachables	862
enable icmp redirects.....	863
enable icmp time-exceeded	863
enable icmp timestamp.....	864
enable icmp unreachable.....	864
enable icmp userredirects.....	865
enable idletimeout.....	42
enable igmp snooping vlan fast-leave.....	1208
enable igmp snooping with-proxy.....	1209
enable igmp snooping.....	1206
enable igmp ssm-map.....	1209
enable igmp.....	1206
enable inline-power legacy slot.....	313
enable inline-power legacy	312
enable inline-power ports	313
enable inline-power slot.....	314
enable inline-power.....	311
enable iparp checking.....	866
enable iparp gratuitous protect.....	619
enable iparp refresh.....	866
enable ipforwarding ipv6.....	929
enable ipforwarding.....	867
enable ipmcforwarding.....	1210
enable ip-option loose-source-route.....	620

enable ip-option record-route	868
enable ip-option record-timestamp	868
enable ip-option router-alert	870
enable ip-option strict-source-route	869
enable iproute compression	870
enable iproute ipv6 compression	930
enable iproute ipv6 sharing	931
enable iproute sharing	871
enable ip-security anomaly-protection icmp	621
enable ip-security anomaly-protection ip	622
enable ip-security anomaly-protection l4port	622
enable ip-security anomaly-protection notify	623
enable ip-security anomaly-protection tcp flags	624
enable ip-security anomaly-protection tcp fragment	624
enable ip-security anomaly-protection	621
enable ip-security arp gratuitous-protection	625
enable ip-security arp learning learn-from-arp	626
enable ip-security arp learning learn-from-dhcp	627
enable ip-security arp validation violation-action	628
enable ip-security dhcp-bindings restoration	630
enable ip-security dhcp-snooping	630
enable ip-security source-ip-lockdown ports	632
enable irdp	872
enable jumbo-frame ports	224
enable learning iparp sender-mac	463
enable learning port	225
enable learning port	464
enable license file	44
enable license software	42
enable lldp ports	290
enable log debug-mode	1350
enable log display	379
enable log target	380
enable loopback-mode vlan	443
enable mac-lockdown-timeout ports	633
enable mirroring to port	226
enable mld snooping with-proxy	1254
enable mld snooping	1253
enable mld	1252
enable msdp data-encapsulation	1283
enable msdp export local-sa	1284
enable msdp peer	1285
enable msdp process-sa-request	1286
enable msdp	1282
enable mvr	1211
enable neighbor-discovery refresh	931
enable netlogin authentication failure vlan ports	717
enable netlogin authentication service-unavailable vlan ports	717

enable netlogin dot1x guest-vlan ports	718
enable netlogin logout-privilege.....	719
enable netlogin ports.....	719
enable netlogin reauthentication-on-refresh.....	720
enable netlogin redirect-page.....	721
enable netlogin session-refresh.....	721
enable netlogin	716
enable ospf capability opaque-lsa.....	1024
enable ospf export.....	1025
enable ospf originate-default	1026
enable ospf restart-helper-lsa-check.....	1027
enable ospf use-ip-router-alert.....	1028
enable ospf.....	1024
enable ospfv3 export.....	1062
enable ospfv3	1062
enable pim snooping.....	1212
enable pim ssm vlan.....	1212
enable pim.....	1211
enable port.....	228
enable radius	634
enable radius-accounting.....	635
enable rip aggregation.....	964
enable rip export	964
enable rip originate-default cost.....	965
enable rip poisonreverse.....	966
enable rip splithorizon.....	967
enable rip triggerupdates	967
enable rip use-ip-router-alert.....	968
enable rip.....	963
enable ripng export.....	985
enable ripng originate-default.....	986
enable ripng poisonreverse	987
enable ripng splithorizon	988
enable ripng triggerupdates	988
enable ripng	984
enable rmon.....	381
enable router-discovery.....	932
enable sflow ports.....	383
enable sflow	383
enable sharing grouping.....	228
enable slot.....	231
enable smartredundancy.....	233
enable snmp access vr.....	105
enable snmp access.....	104
enable snmp community.....	106
enable snmp traps fdb mac-tracking.....	465
enable snmp traps lldp.....	291
enable snmp traps lldp-med	292

enable snmp traps port-up-down ports.....	233
enable snmp traps	106
enable snmpv3.....	107
enable snmp-client	107
enable ssh2.....	636
enable stpd auto-bind.....	777
enable stpd ports	780
enable stpd rapid-root-failover	781
enable stpd.....	777
enable subvlan-proxy-arp vlan	872
enable sys-health-check.....	384
enable syslog.....	385
enable tacacs.....	637
enable tacacs-accounting.....	638
enable tacacs-authorization	638
enable telnet.....	108
enable udp-echo-server.....	873
enable vlan.....	443
enable vrrp vrid	807
enable watchdog	109
enable web http	639
enable web https	639
enable xml-mode	163
exit	110
history.....	45
install firmware	1318
install image.....	1321
load script.....	1324
logout	110
ls	164
mrinfo.....	1213
mtrace	1214
mv	166
nslookup.....	1351
ping	45
quit	111
reboot.....	47
refresh policy	485
Reserved Keywords	10
reset inline-power ports.....	315
restart ports.....	234
restart process	169
rm	171
rtlookup rpf.....	1218
rtlookup rpf.....	933
rtlookup.....	1217
rtlookup.....	874
rtlookup.....	933

run diagnostics	1352
run failover	234
run msm-failover	235
save configuration as-script	1329
save configuration.....	1326
save debug tracefiles memorycard.....	1354
scp2	640
show access-list configuration.....	507
show access-list counter.....	508
show access-list dynamic counter.....	510
show access-list dynamic rule.....	511
show access-list dynamic.....	509
show access-list interface	512
show access-list meter.....	543
show access-list usage acl-mask port.....	514
show access-list usage acl-range port.....	515
show access-list usage acl-rule port.....	515
show access-list usage acl-slice port.....	516
show access-list	506
show accounts password-policy.....	49
show accounts	48
show banner netlogin.....	722
show banner.....	50
show bgp memory	1153
show bgp neighbor flap-statistics.....	1156
show bgp neighbor routes.....	1157
show bgp neighbor	1154
show bgp peer-group.....	1158
show bgp routes summary.....	1162
show bgp routes	1160
show bgp.....	1152
show bootprelay configuration.....	876
show bootprelay dhcp-agent information circuit-id port-information ..	877
show bootprelay dhcp-agent information circuit-id vlan-information ..	878
show bootprelay	875
show checkpoint-data.....	112
show configuration.....	1329
show cpu-monitoring.....	173
show debug.....	1355
show dhcp-client state.....	113
show dhcp-server	642
show diagnostics	1356
show diffserv examination	544
show diffserv replacement	544
show dns-client	51
show dos-protect	643
show dot1p.....	545
show failsafe-account.....	52

show fans.....	386
show fdb mac-tracking configuration.....	467
show fdb mac-tracking statistics.....	468
show fdb stats	469
show fdb.....	465
show flow-redirect	518
show forwarding configuration	1357
show heartbeat process.....	175
show igmp group	1220
show igmp snooping cache.....	1223
show igmp snooping vlan filter	1224
show igmp snooping vlan static	1225
show igmp snooping vlan.....	1223
show igmp snooping.....	1221
show igmp ssm-map	1226
show igmp.....	1219
show inline-power configuration ports.....	317
show inline-power info ports	318
show inline-power slot.....	321
show inline-power stats ports	323
show inline-power stats slot	325
show inline-power stats.....	323
show inline-power	316
show iparp proxy	881
show iparp security.....	884
show iparp stats	882
show iparp.....	879
show ipconfig ipv6	934
show ipconfig.....	885
show ipmroute.....	1227
show iproute ipv6 origin.....	937
show iproute ipv6	935
show iproute multicast.....	1228
show iproute origin.....	888
show iproute reserved-entries statistics	890
show iproute reserved-entries	889
show iproute.....	886
show ip-security anomaly-protection notify cache ports	644
show ip-security arp gratuitous-protection	645
show ip-security arp learning	645
show ip-security arp validation violations	647
show ip-security arp validation.....	646
show ip-security dhcp-snooping entries.....	648
show ip-security dhcp-snooping information circuit-id port-information	650
show ip-security dhcp-snooping information-option circuit-id vlan-information	651
show ip-security dhcp-snooping information-option.....	649
show ip-security dhcp-snooping violations	652

show ip-security dhcp-snooping.....	651
show ip-security source-ip-lockdown.....	653
show ipstats ipv6.....	938
show ipstats	892
show L2stats	1229
show lacp counters.....	237
show lacp lag	238
show lacp member-port.....	243
show lacp.....	236
show licenses	53
show lldp neighbors.....	295
show lldp statistics.....	297
show lldp.....	293
show log components.....	391
show log configuration filter.....	397
show log configuration target.....	398
show log configuration.....	395
show log counters.....	400
show log events	401
show log.....	387
show mac-lockdown-timeout fdb ports.....	654
show mac-lockdown-timeout ports.....	655
show management	114
show mcast cache	1230
show memory process.....	180
show memory.....	176
show memorycard	1331
show meter.....	545
show mirroring	246
show mld group	1256
show mld snooping vlan static.....	1258
show mld snooping.....	1257
show mld.....	1255
show msdp memory	1288
show msdp mesh-group.....	1289
show msdp peer	1290
show msdp sa-cache.....	1291
show msdp.....	1287
show mvr cache	1232
show mvr.....	1232
show neighbor-discovery cache ipv6.....	938
show netlogin authentication failure vlan.....	728
show netlogin authentication service-unavailable vlan.....	728
show netlogin guest-vlan.....	729
show netlogin local-users	730
show netlogin mac-list.....	731
show netlogin	722
show node.....	116

show odometers	118
show ospf area detail.....	1029
show ospf area	1029
show ospf ase-summary.....	1030
show ospf interfaces detail	1031
show ospf interfaces.....	1030
show ospf lsdb	1031
show ospf memory	1032
show ospf neighbor.....	1033
show ospf virtual-link.....	1033
show ospf.....	1028
show ospfv3 area	1065
show ospfv3 interfaces.....	1067
show ospfv3 lsdb stats.....	1070
show ospfv3 lsdb	1068
show ospfv3 memory	1071
show ospfv3 neighbor.....	1071
show ospfv3 virtual-link.....	1072
show ospfv3.....	1064
show pim cache	1237
show pim snooping	1238
show pim.....	1233
show policy.....	486
show port transceiver information detail	259
show port transceiver information.....	258
show ports anomaly.....	247
show ports collisions.....	249
show ports configuration.....	250
show ports congestion.....	546
show ports information.....	251
show ports packet	255
show ports qosmonitor {congestion}.....	548
show ports rate-limit flood	656
show ports redundant.....	256
show ports rxerrors.....	403
show ports sharing.....	256
show ports statistics.....	404
show ports txerrors.....	407
show ports utilization.....	261
show ports.....	246
show power budget	121
show power controller.....	124
show power.....	119
show private-vlan <name>.....	446
show private-vlan	444
show process.....	181
show protocol.....	447
show qosprofile	550

show radius.....	657
show radius-accounting.....	659
show rip interface vlan.....	970
show rip interface.....	969
show rip memory.....	971
show rip routes.....	971
show rip.....	968
show ripng interface.....	990
show ripng routes.....	991
show ripng.....	989
show rmon memory.....	408
show router-discovery.....	939
show script output autoexec.....	1332
show script output default.....	1332
show session.....	125
show sflow configuration.....	411
show sflow statistics.....	412
show sharing health-check.....	264
show slot.....	265
show snmp vr_name.....	128
show snmp.....	127
show snmpv3 access.....	129
show snmpv3 community.....	131
show snmpv3 context.....	132
show snmpv3 counters.....	133
show snmpv3 engine-info.....	134
show snmpv3 filter.....	135
show snmpv3 filter-profile.....	136
show snmpv3 group.....	137
show snmpv3 mib-view.....	139
show snmpv3 notify.....	141
show snmpv3 target-addr.....	142
show snmpv3 target-params.....	143
show snmpv3 user.....	144
show snmp-client.....	146
show ssh2 private-key.....	661
show sshd2 user-key.....	661
show ssl.....	662
show stpd ports.....	785
show stpd.....	781
show switch.....	53
show tacacs.....	663
show tacacs-accounting.....	664
show tech.....	1359
show temperature.....	413
show tunnel.....	940
show udp-profile.....	895
show version.....	414

show virtual-router	476
show vlan dhcp-address-allocation	666
show vlan dhcp-config	666
show vlan security	667
show vlan stpd	788
show vlan	447
show vman	1303
show vrrp virtual-router	810
show vrrp virtual-router	811
show vrrp vlan	812
show vrrp	808
ssh2	668
start process	186
synchronize	1333
telnet msm	148
telnet	147
terminate process	187
tftp get	153
tftp put	155
tftp	149
top	1361
traceroute	55
unconfigure access-list	518
unconfigure bootprelay dhcp-agent information check	895
unconfigure bootprelay dhcp-agent information circuit-id port-information	896
unconfigure bootprelay dhcp-agent information circuit-id vlan-information	897
unconfigure bootprelay dhcp-agent information option	897
unconfigure bootprelay dhcp-agent information policy	898
unconfigure bootprelay dhcp-agent information remote-id	898
unconfigure diffserv examination	551
unconfigure diffserv replacement	552
unconfigure icmp	899
unconfigure igmp snooping vlan ports set join-limit	1239
unconfigure igmp ssm-map	1240
unconfigure igmp	1239
unconfigure inline-power budget slot	326
unconfigure inline-power disconnect-precedence	326
unconfigure inline-power operator-limit ports	327
unconfigure inline-power priority ports	328
unconfigure inline-power usage-threshold	328
unconfigure iparp	899
unconfigure iproute ipv6 priority	941
unconfigure iproute priority	900
unconfigure ip-security dhcp-snooping information check	670
unconfigure ip-security dhcp-snooping information circuit-id port-information ports	670
unconfigure ip-security dhcp-snooping information circuit-id vlan-information 671	

unconfigure ip-security dhcp-snooping information option.....	671
unconfigure ip-security dhcp-snooping information policy.....	672
unconfigure irdp.....	902
unconfigure lldp.....	299
unconfigure log filter.....	417
unconfigure log target format.....	418
unconfigure mld.....	1259
unconfigure msdp sa-cache-server.....	1293
unconfigure mstp region.....	790
unconfigure neighbor-discovery cache.....	943
unconfigure netlogin allowed-refresh-failures.....	732
unconfigure netlogin authentication database-order	732
unconfigure netlogin authentication failure vlan.....	733
unconfigure netlogin authentication service-unavailable vlan.....	733
unconfigure netlogin dot1x guest-vlan.....	734
unconfigure netlogin local-user security-profile.....	734
unconfigure netlogin session-refresh.....	735
unconfigure netlogin vlan	735
unconfigure ospf.....	1034
unconfigure ospfv3.....	1073
unconfigure pim border.....	1293
unconfigure pim ssm range	1241
unconfigure pim.....	1240
unconfigure ports display string.....	267
unconfigure ports redundant	268
unconfigure qosprofile.....	552
unconfigure radius.....	672
unconfigure radius-accounting	673
unconfigure rip.....	972
unconfigure ripng.....	993
unconfigure sflow agent.....	420
unconfigure sflow collector	421
unconfigure sflow ports.....	421
unconfigure sflow.....	419
unconfigure stpd ports link-type.....	791
unconfigure stpd.....	791
unconfigure switch.....	1335
unconfigure tacacs.....	674
unconfigure tacacs-accounting	675
unconfigure trusted-ports trust-for dhcp-server.....	675
unconfigure tunnel.....	949
unconfigure vlan dhcp.....	676
unconfigure vlan dhcp-address-range.....	676
unconfigure vlan dhcp-options	677
unconfigure vlan ipaddress	451
unconfigure vlan router-discovery default-lifetime	944
unconfigure vlan router-discovery hop-limit.....	944
unconfigure vlan router-discovery link-mtu.....	945

unconfigure vlan router-discovery managed-config-flag	945
unconfigure vlan router-discovery max-interval.....	946
unconfigure vlan router-discovery min-interval.....	947
unconfigure vlan router-discovery other-config-flag	947
unconfigure vlan router-discovery reachable-time.....	948
unconfigure vlan router-discovery retransmit-time.....	948
unconfigure vlan router-discovery.....	943
unconfigure vlan subvlan-address-range.....	902
unconfigure vlan udp-profile	903
unconfigure vman ethertype	1305
uninstall image	1335
upload configuration.....	1337
upload debug.....	1362
upload dhcp-bindings.....	678
upload log.....	422
use configuration.....	1341
use image.....	1343
virtual-router	479