

**IBM System Networking RackSwitch™ G7028/G7052**



# **Application Guide**

**For Networking OS 7.6**

**Note:** Before using this information and the product it supports, read the general information in the *Safety information and Environmental Notices and User Guide* documents on the IBM *Documentation* CD and the *Warranty Information* document that comes with the product.

**Second Edition (October 2014)**

**© Copyright IBM Corporation 2014**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Preface</b> . . . . .	<b>11</b>
Who Should Use This Guide . . . . .	11
What You'll Find in This Guide . . . . .	11
Additional References. . . . .	13
Typographic Conventions . . . . .	13
How to Get Help . . . . .	14
 <b>Part 1:. Getting Started.</b> . . . . .	 <b>15</b>
 <b>Chapter 1. Switch Administration</b> . . . . .	 <b>17</b>
Administration Interfaces . . . . .	18
Browser-Based Interface . . . . .	18
Establishing a Connection . . . . .	19
Using the Switch Management Ports . . . . .	19
Using the Switch Data Ports. . . . .	19
Using Telnet. . . . .	21
Using Secure Shell . . . . .	21
Using a Web Browser . . . . .	22
Using Simple Network Management Protocol . . . . .	24
Switch Login Levels . . . . .	25
Setup vs. the Command Line . . . . .	27
Idle Disconnect . . . . .	27
 <b>Chapter 2. Initial Setup.</b> . . . . .	 <b>29</b>
Information Needed for Setup. . . . .	30
Stopping and Restarting Setup Manually . . . . .	30
Setup Part 1: Basic System Configuration . . . . .	31
Setup Part 2: Port Configuration. . . . .	33
Setup Part 3: VLANs . . . . .	35
Setup Part 4: IP Configuration . . . . .	36
IP Interfaces . . . . .	36
Setup Part 5: Final Steps . . . . .	38
Optional Setup for Telnet Support . . . . .	39
 <b>Chapter 3. Switch Software Management</b> . . . . .	 <b>41</b>
Loading New Software to Your Switch . . . . .	42
Loading Software via the ISCLI . . . . .	42
Loading Software via BBI. . . . .	43
USB Options. . . . .	44
USB Boot . . . . .	44
USB Copy . . . . .	45
The Boot Management Menu . . . . .	46
Recovering from a Failed Upgrade . . . . .	46
Recovering a Failed Boot Image . . . . .	49

<b>Part 2:. Securing the Switch . . . . .</b>	<b>51</b>
<b>Chapter 4. Securing Administration . . . . .</b>	<b>53</b>
Secure Shell and Secure Copy . . . . .	54
Configuring SSH/SCP Features on the Switch . . . . .	54
Configuring the SCP Administrator Password. . . . .	55
Using SSH and SCP Client Commands . . . . .	55
SSH and SCP Encryption of Management Messages . . . . .	57
Generating RSA Host Key for SSH Access . . . . .	57
SSH/SCP Integration with Radius Authentication . . . . .	57
SSH/SCP Integration with TACACS+ Authentication . . . . .	57
SecurID Support . . . . .	58
End User Access Control . . . . .	59
Considerations for Configuring End User Accounts . . . . .	59
Strong Passwords . . . . .	59
User Access Control . . . . .	60
Listing Current Users . . . . .	61
Logging into an End User Account . . . . .	61
<b>Chapter 5. Authentication &amp; Authorization Protocols . . . . .</b>	<b>63</b>
RADIUS Authentication and Authorization. . . . .	64
How RADIUS Authentication Works . . . . .	64
Configuring RADIUS on the Switch. . . . .	64
RADIUS Authentication Features in IBM N/OS . . . . .	65
Switch User Accounts . . . . .	66
RADIUS Attributes for IBM N/OS User Privileges . . . . .	66
TACACS+ Authentication . . . . .	68
How TACACS+ Authentication Works. . . . .	68
TACACS+ Authentication Features in IBM N/OS . . . . .	69
Command Authorization and Logging. . . . .	70
Configuring TACACS+ Authentication on the Switch . . . . .	71
LDAP Authentication and Authorization. . . . .	72
<b>Chapter 6. 802.1X Port-Based Network Access Control . . . . .</b>	<b>75</b>
Extensible Authentication Protocol over LAN . . . . .	76
EAPoL Authentication Process . . . . .	77
EAPoL Message Exchange . . . . .	78
EAPoL Port States. . . . .	79
Guest VLAN . . . . .	79
Supported RADIUS Attributes . . . . .	80
EAPoL Configuration Guidelines . . . . .	82
<b>Chapter 7. Access Control Lists. . . . .</b>	<b>83</b>
Summary of Packet Classifiers . . . . .	84
Summary of ACL Actions . . . . .	86
Assigning Individual ACLs to a Port . . . . .	86
ACL Order of Precedence . . . . .	87
ACL Groups . . . . .	87
Assigning ACL Groups to a Port . . . . .	89
ACL Metering and Re-Marking . . . . .	89
ACL Port Mirroring . . . . .	90
Viewing ACL Statistics . . . . .	91

ACL Logging . . . . .	91
Enabling ACL Logging . . . . .	91
Logged Information . . . . .	91
Rate Limiting Behavior . . . . .	92
Log Interval . . . . .	92
ACL Logging Limitations . . . . .	92
ACL Configuration Examples . . . . .	93
Using Storm Control Filters. . . . .	95
<b>Part 3:. Switch Basics . . . . .</b>	<b>97</b>
<b>Chapter 8. VLANs . . . . .</b>	<b>99</b>
VLANs Overview. . . . .	100
VLANs and Port VLAN ID Numbers . . . . .	100
VLAN Numbers . . . . .	100
PVID/Native VLAN Numbers . . . . .	101
VLAN Tagging/Trunk Mode . . . . .	102
VLAN Topologies and Design Considerations . . . . .	106
Multiple VLANs with Tagging/Trunk Mode Adapters . . . . .	106
VLAN Configuration Example . . . . .	108
Protocol-Based VLANs . . . . .	109
Port-Based vs. Protocol-Based VLANs . . . . .	110
PVLAN Priority Levels . . . . .	110
PVLAN Tagging/Trunk Mode . . . . .	110
PVLAN Configuration Guidelines. . . . .	110
Configuring PVLAN . . . . .	111
Private VLANs . . . . .	112
Private VLAN Ports . . . . .	112
Configuration Guidelines . . . . .	113
Configuration Example. . . . .	113
<b>Chapter 9. Ports and Trunking . . . . .</b>	<b>115</b>
Trunking Overview . . . . .	116
Static Trunks . . . . .	117
Static Trunk Requirements . . . . .	117
Static Trunk Group Configuration Rules . . . . .	117
Configuring a Static Port Trunk . . . . .	118
Link Aggregation Control Protocol . . . . .	120
LACP Overview . . . . .	120
LACP Minimum Links Option . . . . .	122
Configuring LACP. . . . .	122
Configurable Trunk Hash Algorithm . . . . .	123
Packet-Based Trunk Hashing . . . . .	123
Port-Based Trunk Hashing . . . . .	124
<b>Chapter 10. Spanning Tree Protocols. . . . .</b>	<b>127</b>
Spanning Tree Protocol Modes . . . . .	128
Global STP Control. . . . .	129

PVRST Mode . . . . .	129
Port States . . . . .	130
Bridge Protocol Data Units . . . . .	130
Bridge Protocol Data Units Overview . . . . .	130
Determining the Path for Forwarding BPDUs . . . . .	130
Simple STP Configuration . . . . .	132
Per-VLAN Spanning Tree Groups . . . . .	134
Using Multiple STGs to Eliminate False Loops. . . . .	134
VLANs and STG Assignment . . . . .	135
Manually Assigning STGs. . . . .	136
Guidelines for Creating VLANs. . . . .	136
Rules for VLAN Tagged Ports . . . . .	136
Adding and Removing Ports from STGs . . . . .	137
The Switch-Centric Model. . . . .	138
Configuring Multiple STGs . . . . .	139
Rapid Spanning Tree Protocol . . . . .	141
Port States . . . . .	141
RSTP Configuration Guidelines . . . . .	141
RSTP Configuration Example . . . . .	141
Multiple Spanning Tree Protocol . . . . .	142
MSTP Region . . . . .	142
Common Internal Spanning Tree . . . . .	142
MSTP Configuration Guidelines . . . . .	143
MSTP Configuration Examples . . . . .	143
Port Type and Link Type . . . . .	146
Edge Port/Portfast . . . . .	146
Link Type. . . . .	146
<b>Chapter 11. Virtual Link Aggregation Groups . . . . .</b>	<b>147</b>
VLAG Overview . . . . .	147
VLAG Capacities . . . . .	150
VLAGs versus Port Trunks . . . . .	150
Configuring VLAGs . . . . .	151
Basic VLAG Configuration . . . . .	152
Configuring VLAGs in Multiple Layers. . . . .	155
<b>Chapter 12. Quality of Service . . . . .</b>	<b>159</b>
QoS Overview . . . . .	160
Using ACL Filters . . . . .	161
Summary of ACL Actions . . . . .	161
ACL Metering and Re-Marking . . . . .	162
Using DSCP Values to Provide QoS . . . . .	163
Differentiated Services Concepts . . . . .	163
Per Hop Behavior . . . . .	164
QoS Levels . . . . .	165
DSCP Re-Marking and Mapping. . . . .	166
DSCP Re-Marking Configuration Examples . . . . .	167
Using 802.1p Priority to Provide QoS . . . . .	169
Queuing and Scheduling . . . . .	170
WRED with ECN . . . . .	171
How WRED/ECN work together . . . . .	171
Configuring WRED/ECN . . . . .	172
WRED/ECN Configuration Example . . . . .	173

<b>Part 4:. IP Features</b>	<b>.177</b>
<b>Chapter 13. Basic IP Features</b>	<b>.179</b>
Dynamic Host Configuration Protocol.	.180
<b>Chapter 14. Internet Protocol Version 6</b>	<b>.181</b>
IPv6 Limitations	.182
IPv6 Address Format	.183
IPv6 Address Types	.184
IPv6 Address Autoconfiguration	.185
IPv6 Interfaces	.186
Supported Applications	.187
Configuration Guidelines	.188
IPv6 Configuration Example	.189
<b>Chapter 15. Internet Group Management Protocol</b>	<b>.191</b>
IGMP Terms	.192
How IGMP Works	.193
IGMP Capacity and Default Values	.194
IGMP Snooping	.195
IGMP Querier	.195
IGMP Groups	.196
IGMPv3 Snooping	.196
IGMP Snooping Configuration Guidelines	.197
IGMP Snooping Configuration Example	.198
Advanced Configuration Example: IGMP Snooping.	.199
Prerequisites	.200
Configuration	.200
Troubleshooting	.204
Additional IGMP Features	.207
FastLeave.	.207
IGMP Filtering	.207
Static Multicast Router	.209
<b>Chapter 16. Basic Redundancy</b>	<b>.211</b>
Trunking for Link Redundancy	.212
Virtual Link Aggregation	.212
Hot Links.	.213
Forward Delay	.213
Preemption	.213
FDB Update	.213
Configuration Guidelines	.213
Configuring Hot Links	.214
<b>Chapter 17. Layer 2 Failover</b>	<b>.215</b>
Monitoring Trunk Links	.215
Setting the Failover Limit	.215
Manually Monitoring Port Links	.216
L2 Failover with Other Features	.217
LACP	.217
Spanning Tree Protocol	.217
Configuration Guidelines	.218
Configuring Layer 2 Failover	.218

<b>Part 5:. Network Management.</b>	<b>219</b>
<b>Chapter 18. Link Layer Discovery Protocol</b>	<b>221</b>
LLDP Overview	222
Enabling or Disabling LLDP	223
Global LLDP Setting	223
Transmit and Receive Control.	223
LLDP Transmit Features	224
Scheduled Interval	224
Minimum Interval.	224
Time-to-Live for Transmitted Information	225
Trap Notifications	225
Changing the LLDP Transmit State.	226
Types of Information Transmitted	226
LLDP Receive Features.	228
Types of Information Received	228
Viewing Remote Device Information	228
Time-to-Live for Received Information	230
LLDP Example Configuration.	232
<b>Chapter 19. Simple Network Management Protocol.</b>	<b>233</b>
SNMP Version 1 & Version 2.	234
SNMP Version 3	235
Configuring SNMP Trap Hosts	237
SNMP MIBs	240
Switch Images and Configuration Files	242
Loading a New Switch Image	243
Loading a Saved Switch Configuration	243
Saving the Switch Configuration	244
Saving a Switch Dump	244
<b>Part 6:. Monitoring</b>	<b>245</b>
<b>Chapter 20. Remote Monitoring</b>	<b>247</b>
RMON Overview	247
RMON Group 1—Statistics	247
RMON Group 2—History	248
History MIB Object ID	248
Configuring RMON History	249
RMON Group 3—Alarms	250
Alarm MIB objects	250
Configuring RMON Alarms	250
RMON Group 9—Events	251
<b>Chapter 21. Port Mirroring</b>	<b>253</b>
<b>Part 7:. Appendices</b>	<b>255</b>
<b>Appendix A. Glossary</b>	<b>257</b>
<b>Appendix B. Getting help and technical assistance.</b>	<b>259</b>
Before you call	259

Using the documentation . . . . .	259
Getting help and information on the World Wide Web . . . . .	260
Software service and support . . . . .	260
Hardware service and support . . . . .	260
IBM Taiwan product service . . . . .	260
<b>Appendix C. Notices. . . . .</b>	<b>261</b>
Trademarks. . . . .	261
Important Notes . . . . .	262
Particulate contamination . . . . .	263
Documentation format . . . . .	264
Electronic emission notices . . . . .	264
Federal Communications Commission (FCC) statement. . . . .	264
Industry Canada Class A emission compliance statement . . . . .	264
Avis de conformité à la réglementation d'Industrie Canada . . . . .	264
Australia and New Zealand Class A statement . . . . .	264
European Union EMC Directive conformance statement. . . . .	265
Germany Class A statement . . . . .	265
Japan VCCI Class A statement . . . . .	266
Korea Communications Commission (KCC) statement . . . . .	266
Russia Electromagnetic Interference (EMI) Class A statement . . . . .	266
People's Republic of China Class A electronic emission statement . . . . .	267
Taiwan Class A compliance statement . . . . .	267
<b>Index . . . . .</b>	<b>269</b>



---

## Preface

The *IBM N/OS 7.6 Application Guide* describes how to configure and use the IBM Networking OS 7.6 software on the RackSwitch G7028/G7052 (referred to as G7028/G7052 throughout this document). For documentation on installing the switch physically, see the *Installation Guide* for your G7028/G7052.

**Note:** This *Application Guide* is applicable to G7028/G7052 product family. In some places “G7028” is used in examples. Unless explicitly mentioned, all commands and descriptions are applicable to both the G7028 and G7052.

---

## Who Should Use This Guide

This guide is intended for network installers and system administrators engaged in configuring and maintaining a network. The administrator should be familiar with Ethernet concepts, IP addressing, Spanning Tree Protocol, and SNMP configuration parameters.

---

## What You’ll Find in This Guide

This guide will help you plan, implement, and administer IBM N/OS software. Where possible, each section provides feature overviews, usage examples, and configuration instructions. The following material is included:

### Part 1: Getting Started

This material is intended to help those new to IBM N/OS products with the basics of switch management. This part includes the following chapters:

- [Chapter 1, “Switch Administration,”](#) describes how to access the G7028/G7052 to configure the switch and view switch information and statistics. This chapter discusses a variety of manual administration interfaces, including local management via the switch console, and remote administration via Telnet, a web browser, or via SNMP.
- [Chapter 2, “Initial Setup,”](#) describes how to use the built-in Setup utility to perform first-time configuration of the switch.
- [Chapter 3, “Switch Software Management,”](#) describes how to update the N/OS software operating on the switch.

### Part 2: Securing the Switch

- [Chapter 4, “Securing Administration,”](#) describes methods for using Secure Shell for administration connections, and configuring end-user access control.
- [Chapter 5, “Authentication & Authorization Protocols,”](#) describes different secure administration for remote administrators. This includes using Remote Authentication Dial-in User Service (RADIUS), as well as TACACS+ and LDAP.
- [Chapter 6, “802.1X Port-Based Network Access Control,”](#) describes how to authenticate devices attached to a LAN port that has point-to-point connection characteristics. This feature prevents access to ports that fail authentication and authorization and provides security.
- [Chapter 7, “Access Control Lists,”](#) describes how to use filters to permit or deny specific types of traffic, based on a variety of source, destination, and packet attributes.

### Part 3: Switch Basics

- [Chapter 8, “VLANs,”](#) describes how to configure Virtual Local Area Networks (VLANs) for creating separate network segments, including how to use VLAN tagging for devices that use multiple VLANs. This chapter also describes Protocol-based VLANs, and Private VLANs.
- [Chapter 9, “Ports and Trunking,”](#) describes how to group multiple physical ports together to aggregate the bandwidth between large-scale network devices.
- [Chapter 10, “Spanning Tree Protocols,”](#) discusses how Spanning Tree Protocol (STP) configures the network so that the switch selects the most efficient path when multiple paths exist. Covers Rapid Spanning Tree Protocol (RSTP), Per-VLAN Rapid Spanning Tree (PVRST), and Multiple Spanning Tree Protocol (MSTP).
- [Chapter 11, “Virtual Link Aggregation Groups,”](#) describes using Virtual Link Aggregation Groups (VLAG) to form trunks spanning multiple VLAG-capable aggregator switches.
- [Chapter 12, “Quality of Service,”](#) discusses Quality of Service (QoS) features, including IP filtering using Access Control Lists (ACLs), Differentiated Services, and IEEE 802.1p priority values.

### Part 4: IP Features

- [Chapter 13, “Basic IP Features,”](#) describes how to configure basic IP features like DHCP for G7028/G7052.
- [Chapter 14, “Internet Protocol Version 6,”](#) describes how to configure the G7028/G7052 for IPv6 host management.
- [Chapter 15, “Internet Group Management Protocol,”](#) describes how the N/OS software implements IGMP Snooping or IGMP Relay to conserve bandwidth in a multicast-switching environment.

### Part 6: High Availability Fundamentals

- [Chapter 16, “Basic Redundancy,”](#) describes how the G7028/G7052 supports redundancy through trunking and hotlinks.
- [Chapter 17, “Layer 2 Failover,”](#) describes how the G7028/G7052 supports high-availability network topologies using Layer 2 Failover.

### Part 5: Network Management

- [Chapter 18, “Link Layer Discovery Protocol,”](#) describes how Link Layer Discovery Protocol helps neighboring network devices learn about each others’ ports and capabilities.
- [Chapter 19, “Simple Network Management Protocol,”](#) describes how to configure the switch for management through an SNMP client.

### Part 6: Monitoring

- [Chapter 20, “Remote Monitoring,”](#) describes how to configure the RMON agent on the switch, so that the switch can exchange network monitoring data.
- [Chapter 21, “Port Mirroring,”](#) discusses tools how copy selected port traffic to a monitor port for network analysis.

### Part 7: Appendices

- [Appendix A, “Glossary,”](#) describes common terms and concepts used throughout this guide.

---

## Additional References

Additional information about installing and configuring the G7028/G7052 is available in the following guides:

- *IBM System Networking RackSwitch G7028/G7052 Installation Guide*
- *IBM Networking OS 7.6 ISCLI Command Reference*

---

## Typographic Conventions

The following table describes the typographic styles used in this book.

Table 1. *Typographic Conventions*

Typeface or Symbol	Meaning	Example
ABC123	This type is used for names of commands, files, and directories used within the text.  It also depicts on-screen computer output and prompts.	View the <code>readme.txt</code> file.  Main#
<b>ABC123</b>	This bold type appears in command examples. It shows text that must be typed in exactly as shown.	Main# <b>sys</b>
<ABC123>	This italicized type appears in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command. Do not type the brackets.  This also shows book titles, special terms, or words to be emphasized.	To establish a Telnet session, enter: host# <b>telnet</b> <IP address>  Read your <i>User's Guide</i> thoroughly.
[ ]	Command items shown inside brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	host# <b>ls</b> [-a]
	The vertical bar (   ) is used in command examples to separate choices where multiple options exist. Select only one of the listed options. Do not type the vertical bar.	host# <b>set left right</b>
<b>AaBbCc123</b>	This block type depicts menus, buttons, and other controls that appear in Web browsers and other graphical interfaces.	Click the <b>Save</b> button.

---

## How to Get Help

If you need help, service, or technical assistance, visit our web site at the following address:

<http://www.ibm.com/support>

The warranty card received with your product provides details for contacting a customer support representative. If you are unable to locate this information, please contact your reseller. Before you call, prepare the following information:

- Serial number of the switch unit
- Software release version number
- Brief description of the problem and the steps you have already taken
- Technical support dump information (# `show tech-support`)

# **Part 1: Getting Started**



---

## Chapter 1. Switch Administration

Your RackSwitch G7028/G7052 (G7028/G7052) is ready to perform basic switching functions right out of the box. Some of the more advanced features, however, require some administrative configuration before they can be used effectively.

The extensive IBM Networking OS switching software included in the G7028/G7052 provides a variety of options for accessing the switch to perform configuration, and to view switch information and statistics.

This chapter discusses the various methods that can be used to administer the switch.

---

## Administration Interfaces

IBM N/OS provides a variety of user-interfaces for administration. These interfaces vary in character and in the methods used to access them: some are text-based, and some are graphical; some are available by default, and some require configuration; some can be accessed by local connection to the switch, and others are accessed remotely using various client applications. For example, administration can be performed using any of the following:

- A built-in, text-based command-line interface for access via serial-port connection or an optional Telnet or SSH session
- The built-in Browser-Based Interface (BBI) available using a standard web-browser
- SNMP support for access through network management software such as IBM Director or HP OpenView

The specific interface chosen for an administrative session depends on user preferences, as well as the switch configuration and the available client tools.

In all cases, administration requires that the switch hardware is properly installed and turned on. (see the *IBM System Networking RackSwitch G7028/G7052 Installation Guide*).

## Browser-Based Interface

The Browser-based Interface (BBI) provides access to the common configuration, management and operation features of the G7028/G7052 through your Web browser.

For more information, refer to the *BBI Quick Guide*.

---

## Establishing a Connection

The factory default settings permit initial switch administration through *only* the built-in serial port. All other forms of access require additional switch configuration before they can be used.

Remote access using the network requires the accessing terminal to have a valid, routable connection to the switch interface. The client IP address may be configured manually, or an IPv4 address can be provided automatically through the switch using a service such as DHCP, or an IPv6 address can be obtained using IPv6 stateless address configuration.

**Note:** Throughout this manual, *IP address* is used in places where either an IPv4 or IPv6 address is allowed. IPv4 addresses are entered in dotted-decimal notation (for example, 10.10.10.1), while IPv6 addresses are entered in hexadecimal notation (for example, 2001:db8:85a3::8a2e:370:7334). In places where only one type of address is allowed, *IPv4 address* or *IPv6 address* is specified.

## Using the Switch Management Ports

To manage the switch through the management ports, you must configure an IP interface for each management interface. Configure the following IPv4 parameters:

- IP address/mask
  - Default gateway address
1. Log on to the switch.
  2. Enter Global Configuration mode.

```
RS G7028> enable
RS G7028# configure terminal
```

3. Configure a management IP address and mask:

```
RS G7028(config)# interface ip 4
RS G7028(config-ip-if)# ip address <management interface IPv4 address>
RS G7028(config-ip-if)# ip netmask <IPv4 subnet mask>
RS G7028(config-ip-if)# enable
RS G7028(config-ip-if)# exit
```

4. Configure the appropriate default gateway.  
IP gateway 4 is required for IF 4.

```
RS G7028(config)# ip gateway 4 address <default gateway IPv4 address>
RS G7028(config)# ip gateway 4 enable
```

Once you configure a management IP address for your switch, you can connect to a management port and use the Telnet program from an external management station to access and control the switch. The management port provides *out-of-band* management.

## Using the Switch Data Ports

To manage the switch using Telnet, SNMP, or a Web browser, you must configure an IP interface.

When a DHCP server is present in the local network for the switch, the DHCP server will be used to configure the IP interface. However, if the switch fails to renew the address obtained through DHCP, the following factory configured settings will be used for IP interface 1:

```
Data IPv4 address: 192.168.49.50
Mask:255.255.255.0
Gateway:192.168.29.255
DHCP: enabled
```

If you manually configure a static IP address, DHCP is disabled. If you manually enable DHCP, the interface will be configured by the DHCP server.

To access the switch, the following IP parameters must be configured:

1. Log on to the switch.
2. Enter IP interface mode.

```
RS G7028> enable
RS G7028# configure terminal
RS G7028(config)# interface ip <IP interface number>
```

3. Configure the management IP interface/mask.

– Using IPv4:

```
RS G7028(config-ip-if)# ip address <management interface IPv4 address>
RS G7028(config-ip-if)# ip netmask <IPv4 subnet mask>
```

– Using IPv6:

```
RS G7028(config-ip-if)# ipv6 address <management interface IPv6 address>
RS G7028(config-ip-if)# ipv6 prefixlen <IPv6 prefix length>
```

4. Configure the VLAN, and enable the interface.

```
RS G7028(config-ip-if)# vlan 1
RS G7028(config-ip-if)# enable
RS G7028(config-ip-if)# exit
```

5. Configure the default gateway.

– If using IPv4:

```
RS G7028(config)# ip gateway <gateway number> address <IPv4 address>
RS G7028(config)# ip gateway <gateway number> enable
```

– If using IPv6:

```
RS G7028(config)# ip gateway6 <gateway number> address <IPv6 address>
RS G7028(config)# ip gateway6 <gateway number> enable
```

Once you configure the IP address and have a network connection, you can use the Telnet program from an external management station to access and control the switch. Once the default gateway is enabled, the management station and your switch do not need to be on the same IP subnet.

## Using Telnet

A Telnet connection offers the convenience of accessing the switch from a workstation connected to the network. Telnet access provides the same options for user and administrator access as those available through the console port.

By default, Telnet access is enabled. Use the following commands to disable or re-enable Telnet access:

```
RS G7028(config)# [no] access telnet enable
```

Once the switch is configured with an IP address and gateway, you can use Telnet to access switch administration from any workstation connected to the management network.

To establish a Telnet connection with the switch, run the Telnet program on your workstation and issue the following Telnet command:

```
telnet <switch IPv4 or IPv6 address>
```

You will then be prompted to enter a password as explained [“Switch Login Levels” on page 25](#).

## Using Secure Shell

Although a remote network administrator can manage the configuration of a G7028/G7052 via Telnet, this method does not provide a secure connection. The Secure Shell (SSH) protocol enables you to securely log into another device over a network to execute commands remotely. As a secure alternative to using Telnet to manage switch configuration, SSH ensures that all data sent over the network is encrypted and secure.

The switch can do only one session of key/cipher generation at a time. Thus, a SSH/SCP client will not be able to login if the switch is doing key generation at that time. Similarly, the system will fail to do the key generation if a SSH/SCP client is logging in at that time.

The supported SSH encryption and authentication methods are:

- Server Host Authentication: Client RSA-authenticates the switch when starting each connection
- Key Exchange: RSA
- Encryption: 3DES-CBC, DES
- User Authentication: Local password authentication, RADIUS, TACACS+

IBM Networking OS implements the SSH version 2.0 standard and is confirmed to work with SSH version 2.0-compliant clients such as the following:

- OpenSSH\_5.4p1 for Linux
- Secure CRT Version 5.0.2 (build 1021)
- Putty SSH release 0.60

### Using SSH to Access the Switch

By default, the SSH feature is disabled. Once the IP parameters are configured and the SSH service is enabled, you can access the command line interface using an SSH connection.

To establish an SSH connection with the switch, run the SSH program on your workstation by issuing the SSH command, followed by the switch IPv4 or IPv6 address:

```
# ssh <switch IP address>
```

If SecurID authentication is required, use the following command:

```
# ssh -l ace <switch IP address>
```

You will then be prompted to enter a password as explained [“Switch Login Levels” on page 25](#).

## Using a Web Browser

The switch provides a Browser-Based Interface (BBI) for accessing the common configuration, management and operation features of the G7028/G7052 through your Web browser.

By default, BBI access via HTTP is enabled on the switch.

You can also access the BBI directly from an open Web browser window. Enter the URL using the IP address of the switch interface (for example, `http://<IPv4 or IPv6 address>`).

### Configuring HTTP Access to the BBI

By default, BBI access via HTTP is enabled on the switch.

To disable or re-enable HTTP access to the switch BBI, use the following commands:

```
RS G7028(config)# access http enable      (Enable HTTP access)
-Or-
RS G7028(config)# no access http enable   (Disable HTTP access)
```

The default HTTP web server port to access the BBI is port 80. However, you can change the default Web server port with the following command:

```
RS G7028(config)# access http port <TCP port number>
```

To access the BBI from a workstation, open a Web browser window and type in the URL using the IP address of the switch interface (for example, `http://<IPv4 or IPv6 address>`).

### Configuring HTTPS Access to the BBI

The BBI can also be accessed via a secure HTTPS connection over management and data ports.

1. Enable HTTPS.

By default, BBI access via HTTPS is disabled on the switch. To enable BBI Access via HTTPS, use the following command:

```
RS G7028(config)# access https enable
```

2. Set the HTTPS server port number (optional).

To change the HTTPS Web server port number from the default port 443, use the following command:

```
RS G7028(config)# access https port <x>
```

3. Generate the HTTPS certificate.

Accessing the BBI via HTTPS requires that you generate a certificate to be used during the key exchange. A default certificate is created the first time HTTPS is enabled, but you can create a new certificate defining the information you want to be used in the various fields.

```
RS G7028(config)# access https generate-certificate
Country Name (2 letter code) []: <country code>
State or Province Name (full name) []: <state>
Locality Name (eg, city) []: <city>
Organization Name (eg, company) []: <company>
Organizational Unit Name (eg, section) []: <org. unit>
Common Name (eg, YOUR name) []: <name>
Email (eg, email address) []: <email address>
Confirm generating certificate? [y/n]: y
Generating certificate. Please wait (approx 30 seconds)
restarting SSL agent
```

4. Save the HTTPS certificate.

The certificate is valid only until the switch is rebooted. To save the certificate so it is retained beyond reboot or power cycles, use the following command:

```
RS G7028(config)# access https save-certificate
```

The certificate is valid only until the switch is rebooted. To save the certificate so it is retained beyond reboot or power cycles, a CLI can be used. When a client (such as a web browser) connects to the switch, the client is asked to accept the certificate and verify that the fields match what is expected. Once BBI access is granted to the client, the BBI can be used as described in the *IBM Networking OS 7.6 BBI Quick Guide*.

## Using Simple Network Management Protocol

N/OS provides Simple Network Management Protocol (SNMP) version 1, version 2, and version 3 support for access through any network management software, such as IBM Director or HP-OpenView.

**Note:** SNMP read and write functions are enabled by default. For best security practices, if SNMP is not needed for your network, it is recommended that you disable these functions prior to connecting the switch to the network.

To access the SNMP agent on the G7028/G7052, the read and write community strings on the SNMP manager must be configured to match those on the switch. The default read community string on the switch is `public` and the default write community string is `private`.

The read and write community strings on the switch can be changed using the following commands:

```
RS G7028(config)# snmp-server read-community <1-32 characters>  
  
-and-  
  
RS G7028(config)# snmp-server write-community <1-32 characters>
```

The SNMP manager must be able to reach any one of the IP interfaces on the switch.

For the SNMP manager to receive the SNMPv1 traps sent out by the SNMP agent on the switch, configure the trap host on the switch with the following commands:

```
RS G7028(config)# snmp-server trap-src-if <trap source IP interface>  
RS G7028(config)# snmp-server host <IPv4 address> <trap host community string>
```

For more information on SNMP usage and configuration, see [“Simple Network Management Protocol” on page 233](#).

---

## Switch Login Levels

To enable better switch management and user accountability, three levels or *classes* of user access have been implemented on the G7028/G7052. Levels of access to CLI, Web management functions, and screens increase as needed to perform various switch management tasks. Conceptually, access classes are defined as follows:

- User interaction with the switch is completely passive—nothing can be changed on the G7028/G7052. Users may display information that has no security or privacy implications, such as switch statistics and current operational state information.
- Operators can only effect temporary changes on the G7028/G7052. These changes will be lost when the switch is rebooted/reset. Operators have access to the switch management features used for daily switch operations. Because any changes an operator makes are undone by a reset of the switch, operators cannot severely impact switch operation.
- Administrators are the only ones that may make permanent changes to the switch configuration—changes that are persistent across a reboot/reset of the switch. Administrators can access switch functions to configure and troubleshoot problems on the G7028/G7052. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.

Access to switch functions is controlled through the use of unique surnames and passwords. Once you are connected to the switch via local Telnet, remote Telnet, or SSH, you are prompted to enter a password. The default user names/password for each access level are listed in the following table.

**Note:** It is recommended that you change default switch passwords after initial configuration and as regularly as required under your network security policies.

*Table 2. User Access Levels - Default Settings*

User Account	Password	Description and Tasks Performed	Status
user	user	The User has no direct responsibility for switch management. He or she can view all switch status information and statistics, but cannot make any configuration changes to the switch.	Disabled
oper	oper	The Operator manages all functions of the switch. The Operator can reset ports, except the management ports.	Disabled
admin	admin	The superuser Administrator has complete access to all menus, information, and configuration commands on the G7028/G7052, including the ability to change both the user and administrator passwords.	Enabled

**Note:** Access to each user level (except admin account) can be disabled by setting the password to an empty value. To disable admin account, use the command: `RS G7028(config)# no access user administrator-enable`. Admin account can be disabled only if there is at least one user account enabled and configured with administrator privilege.

---

## Setup vs. the Command Line

Once the administrator password is verified, you are given complete access to the switch. If the switch is still set to its factory default configuration, the system will ask whether you wish to run Setup (see [“Initial Setup” on page 29](#)), a utility designed to help you through the first-time configuration process. If the switch has already been configured, the command line is displayed instead.

---

## Idle Disconnect

By default, the switch will disconnect your Telnet session after 10 minutes of inactivity. This function is controlled by the idle timeout parameter, which can be set from 0 to 60 minutes, where 0 means the session will never timeout.

Use the following command to set the idle timeout value:

```
RS G7028(config)# system idle <0-60>
```



---

## Chapter 2. Initial Setup

To help with the initial process of configuring your switch, the IBM Networking OS software includes a Setup utility. The Setup utility prompts you step-by-step to enter all the necessary information for basic configuration of the switch.

Whenever you log in as the system administrator under the factory default configuration, you are asked whether you wish to run the Setup utility. Setup can also be activated manually from the command line interface any time after login.

---

## Information Needed for Setup

Setup requests the following information:

- Basic system information
  - Date & time
  - Whether to use Spanning Tree Group or not
- Optional configuration for each port
  - Speed, duplex, flow control, and negotiation mode (as appropriate)
  - Whether to use VLAN tagging/trunk mode or not (as appropriate)
- Optional configuration for each VLAN
  - Name of VLAN
  - Which ports are included in the VLAN
- Optional configuration of IP parameters
  - IP address/mask and VLAN for each IP interface
  - IP addresses for default gateway

If the factory default configuration is detected, the system prompts:

```
IBM System Networking RackSwitch G7028/G7052
18:44:05 Wed Jan 3, 2009

The switch is booted with factory default configuration.
To ease the configuration of the switch, a "Set Up" facility which
will prompt you with those configuration items that are essential to the
operation of the switch is provided.
Would you like to run "Set Up" to configure the switch? [y/n]:
```

---

## Stopping and Restarting Setup Manually

### Stopping Setup

To abort the Setup utility, press <Ctrl-C> during any Setup question. When you abort Setup, the system will prompt:

```
Would you like to run from top again? [y/n]
```

Enter n to abort Setup, or y to restart the Setup program at the beginning.

### Restarting Setup

You can restart the Setup utility manually at any time by entering the following command at the administrator prompt:

```
RS G7028# setup
```

---

## Setup Part 1: Basic System Configuration

When Setup is started, the system prompts:

"Set Up" will walk you through the configuration of System Date and Time, Spanning Tree, Port Speed/Mode, VLANs, and IP interfaces. [type Ctrl-C to abort "Set Up"]

1. Enter y if you will be configuring VLANs. Otherwise enter n.

If you decide not to configure VLANs during this session, you can configure them later using the configuration menus, or by restarting the Setup facility. For more information on configuring VLANs, see the IBM Networking OS *Application Guide*.

Next, the Setup utility prompts you to input basic system information.

2. Enter the year of the current date at the prompt:

System Date:  
Enter year [2009]:

Enter the four-digits that represent the year. To keep the current year, press <Enter>.

3. Enter the month of the current system date at the prompt:

System Date:  
Enter month [1]:

Enter the month as a number from 1 to 12. To keep the current month, press <Enter>.

4. Enter the day of the current date at the prompt:

Enter day [3]:

Enter the date as a number from 1 to 31. To keep the current day, press <Enter>. The system displays the date and time settings:

System clock set to 18:55:36 Wed Jan 28, 2009.

5. Enter the hour of the current system time at the prompt:

System Time:  
Enter hour in 24-hour format [18]:

Enter the hour as a number from 00 to 23. To keep the current hour, press <Enter>.

6. Enter the minute of the current time at the prompt:

Enter minutes [55]:

Enter the minute as a number from 00 to 59. To keep the current minute, press <Enter>.

7. Enter the seconds of the current time at the prompt:

```
Enter seconds [37]:
```

Enter the seconds as a number from 00 to 59. To keep the current second, press <Enter>. The system then displays the date and time settings:

```
System clock set to 8:55:36 Wed Jan 28, 2009.
```

8. Turn Spanning Tree Protocol on or off at the prompt:

```
Spanning Tree:  
Current Spanning Tree Group 1 setting: ON  
Turn Spanning Tree Group 1 OFF? [y/n]
```

Enter y to turn off Spanning Tree, or enter n to leave Spanning Tree on.

---

## Setup Part 2: Port Configuration

**Note:** When configuring port options for your switch, some prompts and options may be different.

1. Select whether you will configure VLANs and VLAN tagging/trunk mode for ports:

```
Port Config:
Will you configure VLANs and VLAN Tagging/Trunk-Mode for ports? [y/n]
```

If you wish to change settings for VLANs, enter *y*, or enter *n* to skip VLAN configuration.

**Note:** The sample screens that appear in this document might differ slightly from the screens displayed by your system. Screen content varies based on the firmware versions and options that are installed.

2. Select the port to configure, or skip port configuration at the prompt:

If you wish to change settings for individual ports, enter the number of the port you wish to configure. To skip port configuration, press <Enter> without specifying any port and go to [“Setup Part 3: VLANs” on page 35](#).

3. Configure Gigabit Ethernet port flow parameters.

The system prompts:

```
Gig Link Configuration:
Port Flow Control:
Current Port EXT1 flow control setting:    both
Enter new value ["rx"/"tx"/"both"/"none"]:
```

Enter *rx* to enable receive flow control, *tx* for transmit flow control, *both* to enable both, or *none* to turn flow control off for the port. To keep the current setting, press <Enter>.

4. Configure Gigabit Ethernet port autonegotiation mode.

If you selected a port that has a Gigabit Ethernet connector, the system prompts:

```
Port Auto Negotiation:
Current Port EXT1 autonegotiation:        on
Enter new value ["on"/"off"]:
```

Enter *on* to enable port autonegotiation, *off* to disable it, or press <Enter> to keep the current setting.

5. If configuring VLANs, enable or disable VLAN tagging/trunk mode for the port.

If you have selected to configure VLANs back in Part 1, the system prompts:

```
Port VLAN tagging/trunk mode config (tagged/trunk mode port can be a member of
multiple VLANs)
Current VLAN tag/trunk mode support:      disabled
Enter new VLAN tag/trunk mode support [d/e]:
```

Enter *d* to disable VLAN tagging/trunk mode for the port or enter *e* to enable VLAN tagging/trunk mode for the port. To keep the current setting, press <Enter>.

6. The system prompts you to configure the next port:

Enter port (1-28):

When you are through configuring ports, press <Enter> without specifying any port. Otherwise, repeat the steps in this section.

---

## Setup Part 3: VLANs

If you chose to skip VLANs configuration back in Part 2, skip to [“Setup Part 4: IP Configuration” on page 36](#).

1. Select the VLAN to configure, or skip VLAN configuration at the prompt:

```
VLAN Config:
Enter VLAN number from 2 to 4094, NULL at end:
```

If you wish to change settings for individual VLANs, enter the number of the VLAN you wish to configure. To skip VLAN configuration, press <Enter> without typing a VLAN number and go to [“Setup Part 4: IP Configuration” on page 36](#).

2. Enter the new VLAN name at the prompt:

```
Current VLAN name: VLAN 2
Enter new VLAN name:
```

Entering a new VLAN name is optional. To use the pending new VLAN name, press <Enter>.

3. Enter the VLAN port numbers:

```
Define Ports in VLAN:
Current VLAN 2: empty
Enter ports one per line, NULL at end:
```

Enter each port, by port number or port alias, and confirm placement of the port into this VLAN. When you are finished adding ports to this VLAN, press <Enter> without specifying any port.

4. Configure Spanning Tree Group membership for the VLAN:

```
Spanning Tree Group membership:
Enter new Spanning Tree Group index [1-127]:
```

5. The system prompts you to configure the next VLAN:

```
VLAN Config:
Enter VLAN number from 2 to 4094, NULL at end:
```

Repeat the steps in this section until all VLANs have been configured. When all VLANs have been configured, press <Enter> without specifying any VLAN.

---

## Setup Part 4: IP Configuration

The system prompts for IPv4 parameters.

Although the switch supports both IPv4 and IPv6 networks, the Setup utility permits only IPv4 configuration. For IPv6 configuration, see [“Internet Protocol Version 6” on page 181](#).

### IP Interfaces

IP interfaces are used for defining the networks to which the switch belongs.

Up to 4 IP interfaces can be configured on the RackSwitch G7028/G7052 (G7028/G7052). The IP address assigned to each IP interface provides the switch with an IP presence on your network. No two IP interfaces can be on the same IP network. The interfaces can be used for connecting to the switch for remote configuration, and for routing between subnets and VLANs (if used).

1. Select the IP interface to configure, or skip interface configuration at the prompt:

```
IP Config:

IP interfaces:
Enter interface number: (1-4)
```

If you wish to configure individual IP interfaces, enter the number of the IP interface you wish to configure. To skip IP interface configuration, press <Enter> without typing an interface number and go to [“Setup Part 5: Final Steps” on page 38](#).

2. For the specified IP interface, enter the IP address in IPv4 dotted decimal notation:

```
Current IP address:    0.0.0.0
Enter new IP address:
```

To keep the current setting, press <Enter>.

3. At the prompt, enter the IPv4 subnet mask in dotted decimal notation:

```
Current subnet mask:    0.0.0.0
Enter new subnet mask:
```

To keep the current setting, press <Enter>.

4. If configuring VLANs, specify a VLAN for the interface.

This prompt appears if you selected to configure VLANs back in Part 1:

```
Current VLAN:    1
Enter new VLAN [1-4094]:
```

Enter the number for the VLAN to which the interface belongs, or press <Enter> without specifying a VLAN number to accept the current setting.

5. The system prompts you to configure another interface:

Enter interface number: (1-4)

Repeat the steps in this section until all IP interfaces have been configured. When all interfaces have been configured, press <Enter> without specifying any interface number.

---

## Setup Part 5: Final Steps

1. When prompted, decide whether to restart Setup or continue:

Would you like to run from top again? [y/n]

Enter y to restart the Setup utility from the beginning, or n to continue.

2. When prompted, decide whether you wish to review the configuration changes:

Review the changes made? [y/n]

Enter y to review the changes made during this session of the Setup utility. Enter n to continue without reviewing the changes. We recommend that you review the changes.

3. Next, decide whether to apply the changes at the prompt:

Apply the changes? [y/n]

Enter y to apply the changes, or n to continue without applying. Changes are normally applied.

4. At the prompt, decide whether to make the changes permanent:

Save changes to flash? [y/n]

Enter y to save the changes to flash. Enter n to continue without saving the changes. Changes are normally saved at this point.

5. If you do not apply or save the changes, the system prompts whether to abort them:

Abort all changes? [y/n]

Enter y to discard the changes. Enter n to return to the “Apply the changes?” prompt.

**Note:** After initial configuration is complete, it is recommended that you change the default passwords.

---

## Optional Setup for Telnet Support

**Note:** This step is optional. Perform this procedure only if you are planning on connecting to the G7028/G7052 through a remote Telnet connection. Telnet is enabled by default. To change the setting, use the following command:

```
RS G7028# [no] access telnet enable
```



---

## Chapter 3. Switch Software Management

The switch software image is the executable code running on the G7028/G7052. A version of the image comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your switch. To get the latest version of software supported for your G7028/G7052, go to the following website:

<http://www.ibm.com/systems/support/>

To determine the software version currently used on the switch, use the following switch command:

```
RS G7028# show boot
```

The typical upgrade process for the software image consists of the following steps:

- Load a new software image and boot image onto an FTP or TFTP server on your network.
- Transfer the new images to your switch.
- Specify the new software image as the one which will be loaded into switch memory the next time a switch reset occurs.
- Reset the switch.

For instructions on the typical upgrade process using the CLI, ISCLI, or BBI, see ["Loading New Software to Your Switch" on page 42..](#)



**CAUTION:**

**Although the typical upgrade process is all that is necessary in most cases, upgrading from (or reverting to) some versions of IBM Networking OS requires special steps prior to or after the software installation process. Please be sure to follow all applicable instructions in the release notes document for the specific software release to ensure that your switch continues to operate as expected after installing new software.**

---

## Loading New Software to Your Switch

The G7028/G7052 can store up to two different switch software images (called `image1` and `image2`) as well as special boot software (called `boot`). When you load new software, you must specify where it is placed: either into `image1`, `image2`, or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.



### CAUTION:

**When you upgrade the switch software image, always load the new boot image and the new software image before you reset the switch. If you do not load a new boot image, your switch might not boot properly (To recover, see [“Recovering from a Failed Upgrade” on page 46](#)).**

To load a new software image to your switch, you will need the following:

- The image and boot software loaded on an FTP or TFTP server on your network.

**Note:** Be sure to download both the new boot file and the new image file.

- The hostname or IP address of the FTP or TFTP server

**Note:** The DNS parameters must be configured if specifying hostnames.

- The name of the new software image or boot file

After you have successfully loaded the new software image, the message “Valid image detected” is displayed, ensuring that the software image has been received from the trusted source.

**Note:** If the message “Failure: image contains invalid signature” is displayed while the new image is loading, the software image does not have a valid security signature for the switch. In this case, make sure you have downloaded the correct image from an authenticated source.

When the software requirements are met, use one of the following procedures to download the new software to your switch. You can use the IBM N/OS CLI, the ISCLI, USB, or the BBI to download and activate new software.

## Loading Software via the ISCLI

1. In Privileged EXEC mode, enter the following command:

```
RS G7028# copy {tftp|ftp} {image1|image2|boot-image}
```

2. Enter the hostname or IP address of the FTP or TFTP server.

```
Address or name of remote host: <name or IP address>
```

3. Enter the name of the new software file on the server.

```
Source file name: <filename>
```

The exact form of the name will vary by server. However, the file location is normally relative to the FTP or TFTP directory (for example, `tftpboot`).

4. If required by the FTP or TFTP server, enter the appropriate username and password.
5. The switch will prompt you to confirm your request.  
Once confirmed, the software will begin loading into the switch.
6. When loading is complete, use the following commands to enter Global Configuration mode to select which software image (`image1` or `image2`) you want to run in switch memory for the next reboot:

```
RS G7028# configure terminal
RS G7028(config)# boot image {image1|image2}
```

The system will then verify which image is set to be loaded at the next reset:

```
Next boot will use switch software image1 instead of image2.
```

7. Reboot the switch to run the new software:

```
Router(config)# reload
```

The system prompts you to confirm your request. Once confirmed, the switch will reboot to use the new software.

## Loading Software via BBI

You can use the Browser-Based Interface to load software onto the G7028/G7052. The software image to load can reside in one of the following locations:

- FTP server
- TFTP server
- Local computer

After you log onto the BBI, perform the following steps to load a software image:

1. Click the Configure context tab in the toolbar.
2. In the Navigation Window, select System > Config/Image Control.

The Switch Image and Configuration Management page appears.

3. If you are loading software from your computer (HTTP client), skip this step and go to the next. Otherwise, if you are loading software from a FTP/TFTP server, enter the server's information in the FTP/TFTP Settings section.
4. In the Image Settings section, select the image version you want to replace (Image for Transfer).
  - If you are loading software from a FTP/TFTP server, enter the file name and click **Get Image**.
  - If you are loading software from your computer, click **Browse**.

In the File Upload Dialog, select the file and click **OK**. Then click **Download via Browser**.

Once the image has loaded, the page refreshes to show the new software.

## USB Options

You can insert a USB drive into the USB port on the G7028/G7052 and use it to work with switch image and configuration files. You can boot the switch using files located on the USB drive, or copy files to and from the USB drive.

To safely remove the USB drive, first use the following command to un-mount the USB file system:

```
system usb-eject
```

**Command mode:** Global configuration

## USB Boot

USB Boot allows you to boot the switch with a software image file, boot file, or configuration file that resides on a USB drive inserted into the USB port. Use the following command to enable or disable USB Boot:

```
[no] boot usbboot enable
```

**Command mode:** Global configuration

When enabled, when the switch is reset/reloaded, it checks the USB port. If a USB drive is inserted into the port, the switch checks the root directory on the USB drive for software and image files. If a valid file is present, the switch loads the file and boots using the file.

**Note:** The following file types are supported: FAT32, NTFS (read-only), EXT2, and EXT3.

The following list describes the valid file names, and describes the switch behavior when it recognizes them. The file names must be exactly as shown, or the switch will not recognize them.

- RSG7028\_Boot.img (for G7028)  
RSG7052\_Boot.img (for G7052)  
The switch replaces the current boot image with the new image, and boots with the new image.
- RSG7028\_OS.img (for G7028)  
RSG7052\_OS.img (for G7052)  
The switch boots with the new software image. The existing images are not affected.
- RSG7028\_replace1\_OS.img (for G7028)  
RSG7052\_replace1\_OS.img (for G7052)  
The switch replaces the current software image1 with the new image, and boots with the new image. RSG7028\_replace1\_OS.img takes precedence over RSG7028\_OS.img; RSG7052\_replace1\_OS.img takes precedence over RSG7052\_OS.img
- RSG7028\_replace2\_OS.img (for G7028)  
RSG7052\_replace2\_OS.img (for G7052)  
The switch replaces the current software image2 with the new image, and boots with the new image. RSG7028\_replace2\_OS.img takes precedence over RSG7028.img; RSG7052\_replace2\_OS.img takes precedence over RSG7052\_OS.img
- RSG7028.cfg (for G7028)  
RSG7052.cfg (for G7052)  
The switch boots with the new configuration file. The existing configuration files (active and backup) are not affected.

- RSG7028\_replace.cfg (for G7028)  
RSG7052\_replace.cfg (for G7052)  
The switch replaces the active configuration file with the new file, and boots with the new file. This file takes precedence over any other configuration files that may be present on the USB drive.

If more than one valid file is present, the switch loads all valid files and boots with them. For example, you may simultaneously load a new boot file, image file, and configuration file from the USB drive.

The switch ignores any files that do not match the valid file names or that have the wrong format.

## USB Copy

If a USB drive is inserted into the USB port, you can copy files from the switch to the USB drive, or from the USB drive to the switch. USB Copy is available only for software image 1 and the active configuration.

### Copy to USB

Use the following command to copy a file from the switch to the USB drive (Privileged EXEC mode):

```
usbcopy tousb <filename> {boot|image1|active|syslog|crashdump}
```

In this example, the active configuration file is copied to a directory on the USB drive:

```
RS G7028(config)# usbcopy tousb a_folder/myconfig.cfg active
```

### Copy from USB

Use the following command to copy a file from the USB drive to the switch:

```
usbcopy fromusb <filename> {boot|image1|active}
```

In this example, the active configuration file is copied from a directory on the USB drive:

```
RS G7028(config)# usbcopy fromusb a_folder/myconfig.cfg active
```

The new file replaces the current file.

**Note:** Do not use two consecutive dot characters ( .. ). Do not use a slash character ( / ) to begin a filename.

---

## The Boot Management Menu

The Boot Management menu allows you to switch the software image, reset the switch to factory defaults, or to recover from a failed software download.

You can interrupt the boot process and enter the Boot Management menu from the serial console port. When the system displays Memory Test, press **<Shift B>**. The Boot Management menu appears.

```
Resetting the System ...
Memory Test .....

Boot Management Menu
1 - Change booting image
2 - Change configuration block
3 - Boot in recovery mode (tftp and xmodem download of images to recover
  switch)
4 - Xmodem download (for boot image only - use recovery mode for
  application images)
5 - Reboot
6 - Exit
Please choose your menu option: 1
Current boot image is 1. Enter image to boot: 1 or 2: 2
Booting from image 2
```

The Boot Management menu allows you to perform the following actions:

- To change the booting image, press 1 and follow the screen prompts.
- To change the configuration block, press 2, and follow the screen prompts.
- To perform a TFTP/XModem image download, press 3 and follow the screen prompts.
- To perform an Xmodem download, press 4 and follow the screen prompts.
- To reboot the switch, press 5. The booting process restarts.
- To exit the Boot Management menu, press 6. The booting process continues.

## Recovering from a Failed Upgrade

Use the following procedure to recover from a failed software upgrade.

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports XModem Download (for example, HyperTerminal, CRT, PuTTY) and select the following serial port characteristics:
  - Speed: 9600 bps
  - Data Bits: 8
  - Stop Bits: 1
  - Parity: None
  - Flow Control: None
3. Boot the switch and access the Boot Management menu by pressing **<Shift B>** while the Memory Test is in progress and the dots are being displayed.

4. Select 3 for Boot in recovery mode. You will see the following display:

```
Entering Rescue Mode.  
Please select one of the following options:  
    T) Configure networking and tftp download an image  
    X) Use xmodem 1K to serial download an image  
    R) Reboot  
    E) Exit
```

- If you choose option x (Xmodem serial download), go to step 5.
- If you choose option t (TFTP download), go to step 6.

5. **Xmodem download:** When you see the following message, change the Serial Port characteristics to 115200 bps:

```
Change the baud rate to 115200 bps and hit the <ENTER> key before  
initiating the download.
```

- a. Press **<Enter>** to set the system into download accept mode. When the readiness meter displays (a series of "C" characters), start XModem on your terminal emulator.
- b. When you see the following message, change the Serial Port characteristics to 9600 bps:

```
Change the baud rate back to 9600 bps, hit the <ESC> key.
```

- c. When you see the following prompt, enter the image number where you want to install the new software and press **<Enter>**:

```
Install image as image 1 or 2 (hit return to just boot image): 1
```

- d. The following message is displayed when the image download is complete. Continue to step 7.

```
Installing image as imagel...  
Imagel updated successfully  
Please select one of the following options:  
    T) Configure networking and tftp download an image  
    X) Use xmodem 1K to serial download an image  
    R) Reboot  
    E) Exit
```

6. **TFTP download:** The switch prompts you to enter the following information:

```
Performing TFTP rescue. Please answer the following questions (enter
'q' to
quit):
IP addr :
Server addr:
Netmask :
Gateway :
Image Filename:
```

- a. Enter the required information and press **<Enter>**.
- b. You will see a display similar to the following:

```
Host IP : 10.10.98.110
Server IP : 10.10.98.100
Netmask : 255.255.255.0
Broadcast : 10.10.98.255
Gateway : 10.10.98.254
Installing image 7.1.0_OS.img from TFTP server 10.10.98.100
```

- c. When you see the following prompt, enter the image number and press **<Enter>**:

```
Install image as image 1 or 2 (hit return to just boot image): 1
```

- d. The following message is displayed when the image download is complete. Continue to step 7.

```
Installing image as image1...
Image1 updated successfully
Please select one of the following options:
    T) Configure networking and tftp download an image
    X) Use xmodem 1K to serial download an image
    R) Reboot
    E) Exit
```

7. Image recovery is complete. Perform one of the following steps:
  - Press **r** to reboot the switch.
  - Press **e** to exit the Boot Management menu
  - Press the Escape key (**<Esc>**) to re-display the Boot Management menu.

## Recovering a Failed Boot Image

Use the following procedure to recover from a failed boot image upgrade.

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports Xmodem download (for example, HyperTerminal, CRT, PuTTY) and select the following serial port characteristics:
  - Speed: 9600 bps
  - Data Bits: 8
  - Stop Bits: 1
  - Parity: None
  - Flow Control: None
3. Boot the switch and access the Boot Management menu by pressing **<Shift B>** while the Memory Test is in progress and the dots are being displayed.
4. Select 4 for **Xmodem download**. You will see the following display:

```
Perform xmodem download
To download an image use 1K Xmodem at 115200 bps.
```

5. When you see the following message, change the Serial Port characteristics to 115200 bps:

```
Change the baud rate to 115200 bps and hit the <ENTER> key before
initiating the download.
```

- a. Press **<Enter>** to set the system into download accept mode. When the readiness meter displays (a series of “C” characters), start Xmodem on your terminal emulator. You will see a display similar to the following:

```
Extracting images ... Do *NOT* power cycle the switch.
**** RAMDISK ****
Un-Protected 38 sectors
Erasing Flash...
..... done
Erased 38 sectors
Writing to
Flash...9....8....7....6....5....4....3....2....1....done
Protected 38 sectors
**** KERNEL ****
Un-Protected 24 sectors
Erasing Flash...
..... done
Erased 24 sectors
Writing to Flash...9....8....7....6....5....4....3....2....1....
```

- b. When you see the following message, change the Serial Port characteristics to 9600 bps:

```
Change the baud rate back to 9600 bps, hit the <ESC> key.
```

Boot image recovery is complete.



# **Part 2: Securing the Switch**



---

## Chapter 4. Securing Administration

Secure switch management is needed for environments that perform significant management functions across the Internet. Common functions for secured management are described in the following sections:

- [“Secure Shell and Secure Copy” on page 54](#)
- [“End User Access Control” on page 59](#)

**Note:** SNMP read and write functions are enabled by default. For best security practices, if SNMP is not needed for your network, it is recommended that you disable these functions prior to connecting the switch to the network (see [“Using Simple Network Management Protocol” on page 24](#)).

---

## Secure Shell and Secure Copy

Because using Telnet does not provide a secure connection for managing a G7028/G7052, Secure Shell (SSH) and Secure Copy (SCP) features have been included for G7028/G7052 management. SSH and SCP use secure tunnels to encrypt and secure messages between a remote administrator and the switch.

**SSH** is a protocol that enables remote administrators to log securely into the G7028/G7052 over a network to execute management commands.

**SCP** is typically used to copy files securely from one machine to another. SCP uses SSH for encryption of data on the network. On a G7028/G7052, SCP is used to download and upload the switch configuration via secure channels.

Although SSH and SCP are disabled by default, enabling and using these features provides the following benefits:

- Identifying the administrator using Name/Password
- Authentication of remote administrators
- Authorization of remote administrators
- Determining the permitted actions and customizing service for individual administrators
- Encryption of management messages
- Encrypting messages between the remote administrator and switch
- Secure copy support

IBM Networking OS implements the SSH version 2.0 standard and is confirmed to work with SSH version 2.0-compliant clients such as the following:

- OpenSSH\_5.4p1 for Linux
- Secure CRT Version 5.0.2 (build 1021)
- Putty SSH release 0.60

## Configuring SSH/SCP Features on the Switch

SSH and SCP features are disabled by default. To change the SSH/SCP settings, using the following procedures.

**Note:** To use SCP, you must first enable SSH.

### To Enable or Disable the SSH Feature

Begin a Telnet session from the console port and enter the following commands:

```
RS G7028(config)# [no] ssh enable
```

### To Enable or Disable SCP Apply and Save

Enter the following commands from the switch CLI to enable the SCP `putcfg_apply` and `putcfg_apply_save` commands:

```
RS G7028(config)# [no] ssh scp-enable
```

## Configuring the SCP Administrator Password

To configure the SCP-only administrator password, enter the following command (the default password is admin):

```
RS G7028(config)# [no] ssh scp-password
Changing SCP-only Administrator password; validation required...
Enter current administrator password: <password>
Enter new SCP-only administrator password: <new password>
Re-enter new SCP-only administrator password: <new password>
New SCP-only administrator password accepted.
```

## Using SSH and SCP Client Commands

This section shows the format for using some client commands. The following examples use 205.178.15.157 as the IP address of a sample switch.

### To Log In to the Switch

Syntax:

```
>> ssh [-4|-6] <switch IP address>

-or-

>> ssh [-4|-6] <login name>@<switch IP address>
```

**Note:** The -4 option (the default) specifies that an IPv4 switch address will be used. The -6 option specifies IPv6.

Example:

```
>> ssh scpadmin@205.178.15.157
```

### To Copy the Switch Configuration File to the SCP Host

Syntax:

```
>> scp [-4|-6] <username>@<switch IP address>:getcfg <local filename>
```

Example:

```
>> scp scpadmin@205.178.15.157:getcfg ad4.cfg
```

### To Load a Switch Configuration File from the SCP Host

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putcfg
```

Example:

```
>> scp ad4.cfg scpadmin@205.178.15.157:putcfg
```

## To Apply and Save the Configuration

When loading a configuration file to the switch, the `apply` and `save` commands are still required for the configuration commands to take effect. The `apply` and `save` commands may be entered manually on the switch, or by using SCP commands.

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putcfg_apply
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putcfg_apply_save
```

Example:

```
>> scp ad4.cfg scpadmin@205.178.15.157:putcfg_apply
>> scp ad4.cfg scpadmin@205.178.15.157:putcfg_apply_save
```

- The CLI `diff` command is automatically executed at the end of `putcfg` to notify the remote client of the difference between the new and the current configurations.
- `putcfg_apply` runs the `apply` command after the `putcfg` is done.
- `putcfg_apply_save` saves the new configuration to the flash after `putcfg_apply` is done.
- The `putcfg_apply` and `putcfg_apply_save` commands are provided because extra `apply` and `save` commands are usually required after a `putcfg`; however, an SCP session is not in an interactive mode.

## To Copy the Switch Image and Boot Files to the SCP Host

Syntax:

```
>> scp [-4|-6] <username>@<switch IP address>:getimg1 <local filename>
>> scp [-4|-6] <username>@<switch IP address>:getimg2 <local filename>
>> scp [-4|-6] <username>@<switch IP address>:getboot <local filename>
```

Example:

```
>> scp scpadmin@205.178.15.157:getimg1 6.1.0_os.img
```

## To Load Switch Configuration Files from the SCP Host

Syntax:

```
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putimg1
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putimg2
>> scp [-4|-6] <local filename> <username>@<switch IP address>:putboot
```

Example:

```
>> scp 6.1.0_os.img scpadmin@205.178.15.157:putimg1
```

## SSH and SCP Encryption of Management Messages

The following encryption and authentication methods are supported for SSH and SCP:

- Server Host Authentication: Client RSA authenticates the switch at the beginning of every connection
- Key Exchange: RSA
- Encryption: 3DES-CBC, DES
- User Authentication: Local password authentication, RADIUS, SecurID (via RADIUS or TACACS+ for SSH only—does not apply to SCP)

## Generating RSA Host Key for SSH Access

To support the SSH host feature, an RSA host key is required. The host key is 1024 bits and is used to identify the G7028/G7052.

To configure RSA host key, first connect to the G7028/G7052 through the console port (commands are not available via external Telnet connection), and enter the following command to generate it manually.

```
RS G7028(config)# ssh generate-host-key
```

When the switch reboots, it will retrieve the host key from the FLASH memory.

**Note:** The switch will perform only one session of key/cipher generation at a time. Thus, an SSH/SCP client will not be able to log in if the switch is performing key generation at that time. Also, key generation will fail if an SSH/SCP client is logging in at that time.

## SSH/SCP Integration with Radius Authentication

SSH/SCP is integrated with RADIUS authentication. After the RADIUS server is enabled on the switch, all subsequent SSH authentication requests will be redirected to the specified RADIUS servers for authentication. The redirection is transparent to the SSH clients.

## SSH/SCP Integration with TACACS+ Authentication

SSH/SCP is integrated with TACACS+ authentication. After the TACACS+ server is enabled on the switch, all subsequent SSH authentication requests will be redirected to the specified TACACS+ servers for authentication. The redirection is transparent to the SSH clients.

## SecurID Support

SSH/SCP can also work with SecurID, a token card-based authentication method. The use of SecurID requires the interactive mode during login, which is not provided by the SSH connection.

**Note:** There is no SNMP or Browser-Based Interface (BBI) support for SecurID because the SecurID server, ACE, is a one-time password authentication and requires an interactive session.

### Using SecurID with SSH

Using SecurID with SSH involves the following tasks.

- To log in using SSH, use a special username, “ace,” to bypass the SSH authentication.
- After an SSH connection is established, you are prompted to enter the username and password (the SecurID authentication is being performed now).
- Provide your username and the token in your SecurID card as a regular Telnet user.

### Using SecurID with SCP

Using SecurID with SCP can be accomplished in two ways:

- Using a RADIUS server to store an administrator password.  
You can configure a regular administrator with a fixed password in the RADIUS server if it can be supported. A regular administrator with a fixed password in the RADIUS server can perform both SSH and SCP with no additional authentication required.
- Using an SCP-only administrator password.  
Set the SCP-only administrator password (`ssh scp-password`) to bypass checking SecurID.  
An SCP-only administrator's password is typically used when SecurID is not used. For example, it can be used in an automation program (in which the tokens of SecurID are not available) to back up (download) the switch configurations each day.

**Note:** The SCP-only administrator's password must be different from the regular administrator's password. If the two passwords are the same, the administrator using that password will not be allowed to log in as an SSH user because the switch will recognize him as the SCP-only administrator. The switch will only allow the administrator access to SCP commands.

---

## End User Access Control

IBM N/OS allows an administrator to define end user accounts that permit end users to perform operation tasks via the switch CLI commands. Once end user accounts are configured and enabled, the switch requires username/password authentication.

For example, an administrator can assign a user, who can then log into the switch and perform operational commands (effective only until the next switch reboot).

## Considerations for Configuring End User Accounts

Note the following considerations when you configure end user accounts:

- A maximum of 10 user IDs are supported on the switch.
- N/OS supports end user support for console, Telnet, BBI, and SSHv2 access to the switch.
- If RADIUS authentication is used, the user password on the Radius server will override the user password on the G7028/G7052. Also note that the password change command only modifies only the user password on the switch and has no effect on the user password on the Radius server. Radius authentication and user password cannot be used concurrently to access the switch.
- Passwords for end users can be up to 128 characters in length for TACACS, RADIUS, Telnet, SSH, Console, and Web access.

## Strong Passwords

The administrator can require use of Strong Passwords for users to access the G7028/G7052. Strong Passwords enhance security because they make password guessing more difficult.

The following rules apply when Strong Passwords are enabled:

- Minimum length: 8 characters; maximum length: 64 characters
- Must contain at least one uppercase alphabet
- Must contain at least one lowercase alphabet
- Must contain at least one number
- Must contain at least one special character:  
Supported special characters: ! " # % & ' ( ) ; < = > ? [ \ ] \* + , - . / : ^ \_ { | } ~
- Cannot be same as the username
- No consecutive four characters can be the same as in the old password

When strong password is enabled, users can still access the switch using the old password but will be advised to change to a strong password while attempting to log in.

Strong password requirement can be enabled using the following command:

```
RS G7028(config)# access user strong-password enable
```

The administrator can choose the number of days allowed before each password expires. When a strong password expires, the user is allowed to log in one last time (last time) to change the password. A warning provides advance notice for users to change the password.

## User Access Control

The end-user access control commands allow you to configure end-user accounts.

### Setting up User IDs

Up to 10 user IDs can be configured. Use the following commands to define any user name and set the user password at the resulting prompts:

```
RS G7028(config)# access user 1 name <1-8 characters>
RS G7028(config)# access user 1 password

Changing user1 password; validation required:
Enter current admin password: <current administrator password>
Enter new user1 password: <new user password>
Re-enter new user1 password: <new user password>
New user1 password accepted.
```

### Defining a User's Access Level

The end user is by default assigned to the user access level (also known as class of service, or COS). COS for all user accounts have global access to all resources except for User COS, which has access to view only resources that the user owns. For more information, see [Table 3 on page 66](#).

To change the user's level, select one of the following options:

```
RS G7028(config)# access user 1 level {user|operator|administrator}
```

### Validating a User's Configuration

```
RS G7028# show access user uid 1
```

### Enabling or Disabling a User

An end user account must be enabled before the switch recognizes and permits login under the account. Once enabled, the switch requires any user to enter both username and password.

```
RS G7028(config)# [no] access user 1 enable
```

### Locking Accounts

To protect the switch from unauthorized access, the account lockout feature can be enabled. By default, account lockout is disabled. To enable this feature, ensure the strong password feature is enabled (See ["Strong Passwords" on page 59](#)). Then use the following command:

```
RS G7028(config)# access user strong-password lockout enable
```

After multiple failed login attempts, the switch locks the user account if lockout has been enabled on the switch.

## Re-enabling Locked Accounts

The administrator can re-enable a locked account by reloading the switch or by using the following command:

```
RS G7028(config)# access user strong-password clear local user lockout <user id>
```

However, the above command cannot be used to re-enable an account disabled by the administrator.

To re-enable all locked accounts, use the following command:

```
RS G7028(config)# access user strong-password clear local user all
```

## Listing Current Users

The following command displays defined user accounts and whether or not each user is currently logged into the switch.

```
RS G7028# show access user

Usernames:
  user    - Enabled - offline
  oper    - Disabled - offline
  admin   - Always Enabled - online 1 session

Current User ID table:
1: name jane    , ena, cos user    , password valid, online 1 session
2: name john    , ena, cos user    , password valid, online 2 sessions
```

## Logging into an End User Account

Once an end user account is configured and enabled, the user can login to the switch using the username/password combination. The level of switch access is determined by the COS established for the end user account.



---

## Chapter 5. Authentication & Authorization Protocols

Secure switch management is needed for environments that perform significant management functions across the Internet. The following are some of the functions for secured IPv4 management and device access:

- [“RADIUS Authentication and Authorization” on page 64](#)
- [“TACACS+ Authentication” on page 68](#)
- [“LDAP Authentication and Authorization” on page 72](#)

**Note:** IBM Networking OS 7.6 does not support IPv6 for RADIUS, TACACS+ or LDAP.

---

## RADIUS Authentication and Authorization

IBM N/OS supports the RADIUS (Remote Authentication Dial-in User Service) method to authenticate and authorize remote administrators for managing the switch. This method is based on a client/server model. The Remote Access Server (RAS)—the switch—is a client to the back-end database server. A remote user (the remote administrator) interacts only with the RAS, not the back-end server and database.

RADIUS authentication consists of the following components:

- A protocol with a frame format that utilizes UDP over IP (based on RFC 2138 and 2866)
- A centralized server that stores all the user authorization information
- A client: in this case, the switch

The G7028/G7052—acting as the RADIUS client—communicates to the RADIUS server to authenticate and authorize a remote administrator using the protocol definitions specified in RFC 2138 and 2866. Transactions between the client and the RADIUS server are authenticated using a shared key that is not sent over the network. In addition, the remote administrator passwords are sent encrypted between the RADIUS client (the switch) and the back-end RADIUS server.

### How RADIUS Authentication Works

The RADIUS authentication process follows these steps:

1. A remote administrator connects to the switch and provides a user name and password.
2. Using Authentication/Authorization protocol, the switch sends request to authentication server.
3. The authentication server checks the request against the user ID database.
4. Using RADIUS protocol, the authentication server instructs the switch to grant or deny administrative access.

### Configuring RADIUS on the Switch

Use the following procedure to configure Radius authentication on your switch.

1. Configure the IPv4 addresses of the Primary and Secondary RADIUS servers, and enable RADIUS authentication.

```
RS G7028(config)# radius-server primary-host 10.10.1.1
RS G7028(config)# radius-server secondary-host 10.10.1.2
RS G7028(config)# radius-server enable
```

**Note:** You can use a configured loopback address as the source address so the RADIUS server accepts requests only from the expected loopback address block. Use the following command to specify the loopback interface:

```
RS G7028(config)# ip radius source-interface loopback <1-5>
```

2. Configure the RADIUS secret.

```
RS G7028(config)# radius-server primary-host 10.10.1.1 key <1-32 character secret>
RS G7028(config)# radius-server secondary-host 10.10.1.2 key <1-32 character secret>
```

3. If desired, you may change the default UDP port number used to listen to RADIUS.

The well-known port for RADIUS is 1812.

```
RS G7028(config)# radius-server port <UDP port number>
```

4. Configure the number retry attempts for contacting the RADIUS server, and the timeout period.

```
RS G7028(config)# radius-server retransmit 3
RS G7028(config)# radius-server timeout 5
```

## RADIUS Authentication Features in IBM N/OS

N/OS supports the following RADIUS authentication features:

- Supports RADIUS client on the switch, based on the protocol definitions in RFC 2138 and RFC 2866.
- Allows RADIUS secret password up to 32 bytes and less than 16 octets.
- Supports *secondary authentication server* so that when the primary authentication server is unreachable, the switch can send client authentication requests to the secondary authentication server. Use the following command to show the currently active RADIUS authentication server:

```
RS G7028# show radius-server
```

- Supports user-configurable RADIUS server retry and time-out values:
  - Time-out value = 1-10 seconds
  - Retries = 1-3

The switch will time out if it does not receive a response from the RADIUS server in 1-3 retries. The switch will also automatically retry connecting to the RADIUS server before it declares the server down.

- Supports user-configurable RADIUS application port. The default is 1812/UDP-based on RFC 2138. Port 1645 is also supported.
- Supports user-configurable RADIUS application port. The default is UDP port 1645. UDP port 1812, based on RFC 2138, is also supported.
- Allows network administrator to define privileges for one or more specific users to access the switch at the RADIUS user database.

## Switch User Accounts

The user accounts listed in [Table 3](#) can be defined in the RADIUS server dictionary file.

Table 3. User Access Levels

User Account	Description and Tasks Performed	Password
User	The User has no direct responsibility for switch management. They can view all switch status information and statistics but cannot make any configuration changes to the switch.	user
Operator	The Operator manages all functions of the switch. The Operator can reset ports, except the management port.	oper
Administrator	The super-user Administrator has complete access to all commands, information, and configuration commands on the switch, including the ability to change both the user and administrator passwords.	admin

## RADIUS Attributes for IBM N/OS User Privileges

When the user logs in, the switch authenticates his/her level of access by sending the RADIUS access request, that is, the client authentication request, to the RADIUS authentication server.

If the remote user is successfully authenticated by the authentication server, the switch will verify the *privileges* of the remote user and authorize the appropriate access. The administrator has two options: to allow *backdoor* access via Telnet, SSH, HTTP, or HTTPS; to allow *secure backdoor* access via Telnet, SSH, or BBI. Backdoor and secure backdoor provides access to the switch when the RADIUS servers cannot be reached.

The default G7028/G7052 setting for backdoor and secure backdoor access is *disabled*. Backdoor and secure backdoor access is always enabled on the console port.

Irrespective of backdoor/secure backdoor being enabled or not, you can always access the switch via the console port by using *noradius* as radius username. You can then enter the username and password configured on the switch. If you are trying to connect via SSH/Telnet/HTTP/HTTPS (not console port), there are two possibilities:

- Backdoor is enabled: The switch acts like it is connecting via console.
- Secure backdoor is enabled: You must enter the username: *noradius*. The switch checks if RADIUS server is reachable. If it is reachable, then you must authenticate via remote authentication server. Only if RADIUS server is not reachable, you will be prompted for local user/password to be authenticated against these local credentials.

All user privileges, other than those assigned to the Administrator, have to be defined in the RADIUS dictionary. RADIUS attribute 6 which is built into all RADIUS servers defines the administrator. The file name of the dictionary is RADIUS vendor-dependent. The following RADIUS attributes are defined for G7028/G7052 user privileges levels:

*Table 4. IBM N/OS-proprietary Attributes for RADIUS*

<b>User Name/Access</b>	<b>User-Service-Type</b>	<b>Value</b>
User	<i>Vendor-supplied</i>	255
Operator	<i>Vendor-supplied</i>	252
Admin	<i>Vendor-supplied</i>	6

---

## TACACS+ Authentication

N/OS supports authentication and authorization with networks using the Cisco Systems TACACS+ protocol. The G7028/G7052 functions as the Network Access Server (NAS) by interacting with the remote client and initiating authentication and authorization sessions with the TACACS+ access server. The remote user is defined as someone requiring management access to the G7028/G7052 either through a data port or management port.

TACACS+ offers the following advantages over RADIUS:

- TACACS+ uses TCP-based connection-oriented transport; whereas RADIUS is UDP-based. TCP offers a connection-oriented transport, while UDP offers best-effort delivery. RADIUS requires additional programmable variables such as re-transmit attempts and time-outs to compensate for best-effort transport, but it lacks the level of built-in support that a TCP transport offers.
- TACACS+ offers full packet encryption whereas RADIUS offers password-only encryption in authentication requests.
- TACACS+ separates authentication, authorization and accounting.

## How TACACS+ Authentication Works

TACACS+ works much in the same way as RADIUS authentication as described on [page 64](#).

1. Remote administrator connects to the switch and provides user name and password.
2. Using Authentication/Authorization protocol, the switch sends request to authentication server.
3. Authentication server checks the request against the user ID database.
4. Using TACACS+ protocol, the authentication server instructs the switch to grant or deny administrative access.

During a session, if additional authorization checking is needed, the switch checks with a TACACS+ server to determine if the user is granted permission to use a particular command.

## TACACS+ Authentication Features in IBM N/OS

Authentication is the action of determining the identity of a user, and is generally done when the user first attempts to log in to a device or gain access to its services. N/OS supports ASCII inbound login to the device. PAP, CHAP and ARAP login methods, TACACS+ change password requests, and one-time password authentication are not supported.

### Authorization

Authorization is the action of determining a user's privileges on the device, and usually takes place after authentication.

The default mapping between TACACS+ authorization levels and N/OS management access levels is shown in [Table 5](#). The authorization levels must be defined on the TACACS+ server.

*Table 5. Default TACACS+ Authorization Levels*

N/OS User Access Level	TACACS+ level
user	0
oper	3
admin	6

Alternate mapping between TACACS+ authorization levels and N/OS management access levels is shown in [Table 6](#). Use the following command to set the alternate TACACS+ authorization levels.

```
RS G7028(config)# tacacs-server privilege-mapping
```

*Table 6. Alternate TACACS+ Authorization Levels*

N/OS User Access Level	TACACS+ level
user	0 - 1
oper	6 - 8
admin	14 - 15

If the remote user is successfully authenticated by the authentication server, the switch verifies the *privileges* of the remote user and authorizes the appropriate access. The administrator has an option to allow *secure backdoor* access via Telnet/SSH. Secure backdoor provides switch access when the TACACS+ servers cannot be reached. You always can access the switch via the console port, by using `notacacs` and the administrator password, whether secure backdoor is enabled or not.

**Note:** To obtain the TACACS+ backdoor password for your G7028/G7052, contact Technical Support.

## Accounting

Accounting is the action of recording a user's activities on the device for the purposes of billing and/or security. It follows the authentication and authorization actions. If the authentication and authorization is not performed via TACACS+, there are no TACACS+ accounting messages sent out.

You can use TACACS+ to record and track software login access, configuration changes, and interactive commands.

The G7028/G7052 supports the following TACACS+ accounting attributes:

- protocol (console/Telnet/SSH/HTTP/HTTPS)
- start\_time
- stop\_time
- elapsed\_time
- disc\_cause

**Note:** When using the Browser-Based Interface, the TACACS+ Accounting Stop records are sent only if the **Logout** button on the browser is clicked.

## Command Authorization and Logging

When TACACS+ Command Authorization is enabled, N/OS configuration commands are sent to the TACACS+ server for authorization. Use the following command to enable TACACS+ Command Authorization:

```
RS G7028(config)# tacacs-server command-authorization
```

When TACACS+ Command Logging is enabled, N/OS configuration commands are logged on the TACACS+ server. Use the following command to enable TACACS+ Command Logging:

```
RS G7028(config)# tacacs-server command-logging
```

The following examples illustrate the format of N/OS commands sent to the TACACS+ server:

```
authorization request, cmd=shell, cmd-arg=interface ip
accounting request, cmd=shell, cmd-arg=interface ip
authorization request, cmd=shell, cmd-arg=enable
accounting request, cmd=shell, cmd-arg=enable
```

## Configuring TACACS+ Authentication on the Switch

1. Configure the IPv4 addresses of the Primary and Secondary TACACS+ servers, and enable TACACS authentication. Specify the interface port (optional).

```
RS G7028(config)# tacacs-server primary-host 10.10.1.1
RS G7028(config)# tacacs-server primary-host mgtb-port
RS G7028(config)# tacacs-server secondary-host 10.10.1.2
RS G7028(config)# tacacs-server secondary-host data-port
RS G7028(config)# tacacs-server enable
```

- Note:** You can use a configured loopback address as the source address so the TACACS+ server accepts requests only from the expected loopback address block. Use the following command to specify the loopback interface:
- ```
RS G7028(config)# ip tacacs source-interface loopback <I-5>
```

2. Configure the TACACS+ secret and second secret.

```
RS G7028(config)# tacacs-server primary-host 10.10.1.1 key <1-32 character secret>
RS G7028(config)# tacacs-server secondary-host 10.10.1.2 key <1-32 character secret>
```

3. If desired, you may change the default TCP port number used to listen to TACACS+.

The well-known port for TACACS+ is 49.

```
RS G7028(config)# tacacs-server port <TCP port number>
```

4. Configure the number of retry attempts, and the timeout period.

```
RS G7028(config)# tacacs-server retransmit 3
RS G7028(config)# tacacs-server timeout 5
```

---

## LDAP Authentication and Authorization

N/OS supports the LDAP (Lightweight Directory Access Protocol) method to authenticate and authorize remote administrators to manage the switch. LDAP is based on a client/server model. The switch acts as a client to the LDAP server. A remote user (the remote administrator) interacts only with the switch, not the back-end server and database.

LDAP authentication consists of the following components:

- A protocol with a frame format that utilizes TCP over IP
- A centralized server that stores all the user authorization information
- A client: in this case, the switch

Each entry in the LDAP server is referenced by its Distinguished Name (DN). The DN consists of the user-account name concatenated with the LDAP domain name. If the user-account name is John, the following is an example DN:

```
uid=John,ou=people,dc=domain,dc=com
```

### Configuring the LDAP Server

G7028/G7052 user groups and user accounts must reside within the same domain. On the LDAP server, configure the domain to include G7028/G7052 user groups and user accounts, as follows:

- User Accounts:  
Use the *uid* attribute to define each individual user account.
- User Groups:  
Use the *members* attribute in the *groupOfNames* object class to create the user groups. The first word of the common name for each user group must be equal to the user group names defined in the G7028/G7052, as follows:
  - admin
  - oper
  - user

### Configuring LDAP Authentication on the Switch

1. Turn LDAP authentication on, then configure the IPv4 addresses of the Primary and Secondary LDAP servers. Specify the interface port (optional).

```
>> # ldap-server enable
>> # ldap-server primary-host 10.10.1.1 mgta-port
>> # ldap-server secondary-host 10.10.1.2 data-port
```

2. Configure the domain name.

```
>> # ldap-server domain <ou=people,dc=my-domain,dc=com>
```

3. You may change the default TCP port number used to listen to LDAP (optional). The well-known port for LDAP is 389.

```
>> # ldap-server port <1-65000>
```

4. Configure the number of retry attempts for contacting the LDAP server, and the timeout period.

```
>> # ldap-server retransmit 3  
>> # ldap-server timeout 10
```



---

## Chapter 6. 802.1X Port-Based Network Access Control

Port-Based Network Access control provides a means of authenticating and authorizing devices attached to a LAN port that has point-to-point connection characteristics. It prevents access to ports that fail authentication and authorization. This feature provides security to ports of the RackSwitch G7028/G7052 (G7028/G7052) that connect to blade servers.

The following topics are discussed in this section:

- [“Extensible Authentication Protocol over LAN” on page 76](#)
- [“EAPoL Authentication Process” on page 77](#)
- [“EAPoL Port States” on page 79](#)
- [“Guest VLAN” on page 79](#)
- [“Supported RADIUS Attributes” on page 80](#)
- [“EAPoL Configuration Guidelines” on page 82](#)

---

## Extensible Authentication Protocol over LAN

IBM Networking OS can provide user-level security for its ports using the IEEE 802.1X protocol, which is a more secure alternative to other methods of port-based network access control. Any device attached to an 802.1X-enabled port that fails authentication is prevented access to the network and denied services offered through that port.

The 802.1X standard describes port-based network access control using Extensible Authentication Protocol over LAN (EAPoL). EAPoL provides a means of authenticating and authorizing devices attached to a LAN port that has point-to-point connection characteristics and of preventing access to that port in cases of authentication and authorization failures.

EAPoL is a client-server protocol that has the following components:

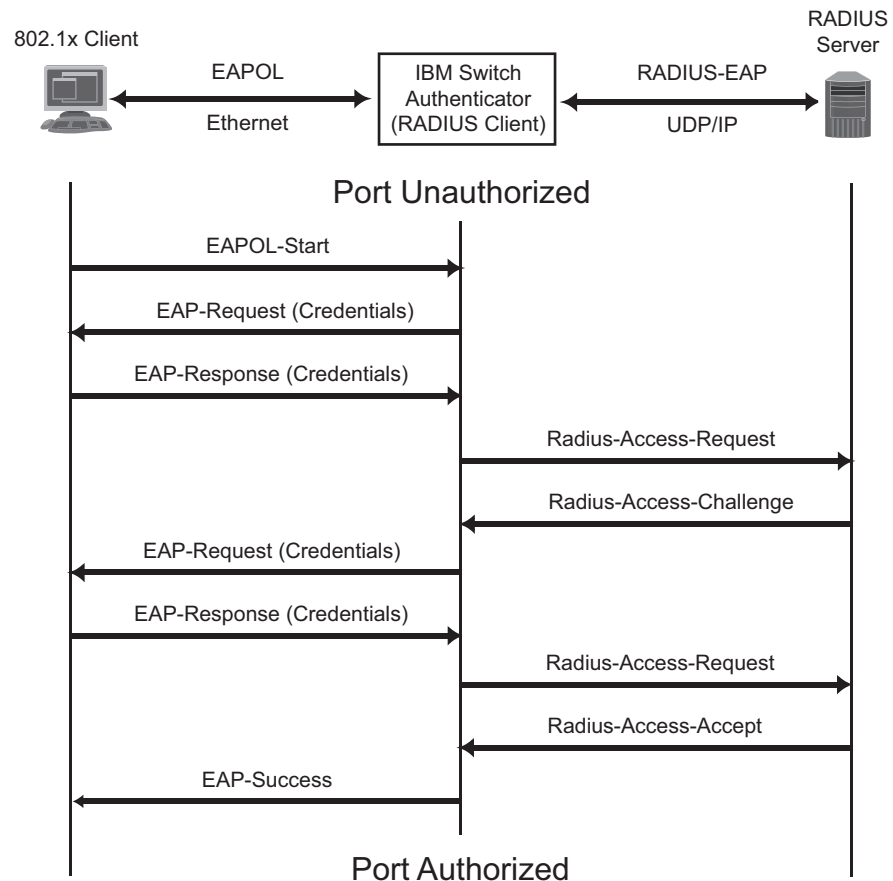
- **Supplicant or Client**  
The Supplicant is a device that requests network access and provides the required credentials (user name and password) to the Authenticator and the Authenticator Server.
- **Authenticator**  
The Authenticator enforces authentication and controls access to the network. The Authenticator grants network access based on the information provided by the Supplicant and the response from the Authentication Server. The Authenticator acts as an intermediary between the Supplicant and the Authentication Server: requesting identity information from the client, forwarding that information to the Authentication Server for validation, relaying the server's responses to the client, and authorizing network access based on the results of the authentication exchange. The G7028/G7052 acts as an Authenticator.
- **Authentication Server**  
The Authentication Server validates the credentials provided by the Supplicant to determine if the Authenticator ought to grant access to the network. The Authentication Server may be co-located with the Authenticator. The G7028/G7052 relies on external RADIUS servers for authentication.

Upon a successful authentication of the client by the server, the 802.1X-controlled port transitions from unauthorized to authorized state, and the client is allowed full access to services through the port. When the client sends an EAP-Logoff message to the authenticator, the port will transition from authorized to unauthorized state.

## EAPoL Authentication Process

The clients and authenticators communicate using Extensible Authentication Protocol (EAP), which was originally designed to run over PPP, and for which the IEEE 802.1X Standard has defined an encapsulation method over Ethernet frames, called EAP over LAN (EAPOL). [Figure 1](#) shows a typical message exchange initiated by the client.

Figure 1. Authenticating a Port Using EAPoL



---

## EAPoL Message Exchange

During authentication, EAPoL messages are exchanged between the client and the G7028/G7052 authenticator, while RADIUS-EAP messages are exchanged between the G7028/G7052 authenticator and the RADIUS server.

Authentication is initiated by one of the following methods:

- The G7028/G7052 authenticator sends an EAP-Request/Identity packet to the client
- The client sends an EAPoL-Start frame to the G7028/G7052 authenticator, which responds with an EAP-Request/Identity frame.

The client confirms its identity by sending an EAP-Response/Identity frame to the G7028/G7052 authenticator, which forwards the frame encapsulated in a RADIUS packet to the server.

The RADIUS authentication server chooses an EAP-supported authentication algorithm to verify the client's identity, and sends an EAP-Request packet to the client via the G7028/G7052 authenticator. The client then replies to the RADIUS server with an EAP-Response containing its credentials.

Upon a successful authentication of the client by the server, the 802.1X-controlled port transitions from unauthorized to authorized state, and the client is allowed full access to services through the controlled port. When the client later sends an EAPoL-Logoff message to the G7028/G7052 authenticator, the port transitions from authorized to unauthorized state.

If a client that does not support 802.1X connects to an 802.1X-controlled port, the G7028/G7052 authenticator requests the client's identity when it detects a change in the operational state of the port. The client does not respond to the request, and the port remains in the unauthorized state.

**Note:** When an 802.1X-enabled client connects to a port that is not 802.1X-controlled, the client initiates the authentication process by sending an EAPoL-Start frame. When no response is received, the client retransmits the request for a fixed number of times. If no response is received, the client assumes the port is in authorized state, and begins sending frames, even if the port is unauthorized.

---

## EAPoL Port States

The state of the port determines whether the client is granted access to the network, as follows:

- **Unauthorized**  
While in this state the port discards all ingress and egress traffic except EAP packets.
- **Authorized**  
When the client is successfully authenticated, the port transitions to the authorized state allowing all traffic to and from the client to flow normally.
- **Force Unauthorized**  
You can configure this state that denies all access to the port.
- **Force Authorized**  
You can configure this state that allows full access to the port.

Use the 802.1X global configuration commands (`dot1x`) to configure 802.1X authentication for all ports in the switch. Use the 802.1X port commands to configure a single port.

---

## Guest VLAN

The guest VLAN provides limited access to unauthenticated ports. The guest VLAN can be configured using the following commands:

```
RS G7028(config)# dot1x guest-vlan ?
```

Client ports that have not received an EAPOL response are placed into the Guest VLAN, if one is configured on the switch. Once the port is authenticated, it is moved from the Guest VLAN to its configured VLAN.

When Guest VLAN enabled, the following considerations apply while a port is in the unauthenticated state:

- The port is placed in the guest VLAN.
- The Port VLAN ID (PVID) is changed to the Guest VLAN ID.
- Port tagging is disabled on the port.

## Supported RADIUS Attributes

The 802.1X Authenticator relies on external RADIUS servers for authentication with EAP. [Table 7](#) lists the RADIUS attributes that are supported as part of RADIUS-EAP authentication based on the guidelines specified in Annex D of the 802.1X standard and RFC 3580.

*Table 7. Support for RADIUS Attributes*

| #  | Attribute          | Attribute Value                                                                                                                                                                                                                              | A-R | A-A | A-C | A-R |
|----|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|-----|-----|
| 1  | User-Name          | The value of the Type-Data field from the supplicant's EAP-Response/ Identity message. If the Identity is unknown (for example, Type-Data field is zero bytes in length), this attribute will have the same value as the Calling-Station-Id. | 1   | 0-1 | 0   | 0   |
| 4  | NAS-IP-Address     | IPv4 address of the authenticator used for Radius communication.                                                                                                                                                                             | 1   | 0   | 0   | 0   |
| 5  | NAS-Port           | Port number of the authenticator port to which the supplicant is attached.                                                                                                                                                                   | 1   | 0   | 0   | 0   |
| 24 | State              | Server-specific value. This is sent unmodified back to the server in an Access-Request that is in response to an Access-Challenge.                                                                                                           | 0-1 | 0-1 | 0-1 | 0   |
| 30 | Called-Station-ID  | The MAC address of the authenticator encoded as an ASCII string in canonical format, such as 000D5622E3 9F.                                                                                                                                  | 1   | 0   | 0   | 0   |
| 31 | Calling-Station-ID | The MAC address of the supplicant encoded as an ASCII string in canonical format, such as 00034B436206.                                                                                                                                      | 1   | 0   | 0   | 0   |
| 64 | Tunnel-Type        | Only VLAN (type 13) is currently supported (for 802.1X RADIUS VLAN assignment). The attribute must be untagged (the Tag field must be 0).                                                                                                    | 0   | 0-1 | 0   | 0   |
| 65 | Tunnel-Medium-Type | Only 802 (type 6) is currently supported (for 802.1X RADIUS VLAN assignment). The attribute must be untagged (the Tag field must be 0).                                                                                                      | 0   | 0-1 | 0   | 0   |

Table 7. Support for RADIUS Attributes (continued)

| #                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Attribute               | Attribute Value                                                                                                                                                                                                                                                                                                                                                                           | A-R | A-A | A-C | A-R |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|-----|-----|
| 81                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Tunnel-Private-Group-ID | VLAN ID (1-4094). When 802.1X RADIUS VLAN assignment is enabled on a port, if the RADIUS server includes the tunnel attributes defined in RFC 2868 in the Access-Accept packet, the switch will automatically place the authenticated port in the specified VLAN. Reserved VLANs (such as for management) may not be specified. The attribute must be untagged (the Tag field must be 0). | 0   | 0-1 | 0   | 0   |
| 79                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | EAP-Message             | Encapsulated EAP packets from the supplicant to the authentication server (Radius) and vice-versa. The authenticator relays the decoded packet to both devices.                                                                                                                                                                                                                           | 1+  | 1+  | 1+  | 1+  |
| 80                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Message-Authenticator   | Always present whenever an EAP-Message attribute is also included. Used to integrity-protect a packet.                                                                                                                                                                                                                                                                                    | 1   | 1   | 1   | 1   |
| 87                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | NAS-Port-ID             | Name assigned to the authenticator port, e.g. Server1_Port3                                                                                                                                                                                                                                                                                                                               | 1   | 0   | 0   | 0   |
| <p><b>Legend:</b> RADIUS Packet Types: A-R (Access-Request), A-A (Access-Accept), A-C (Access-Challenge), A-R (Access-Reject)</p> <p>RADIUS Attribute Support:</p> <ul style="list-style-type: none"> <li>• 0 This attribute MUST NOT be present in a packet.</li> <li>• 0+ Zero or more instances of this attribute MAY be present in a packet.</li> <li>• 0-1 Zero or one instance of this attribute MAY be present in a packet.</li> <li>• 1 Exactly one instance of this attribute MUST be present in a packet.</li> <li>• 1+ One or more of these attributes MUST be present.</li> </ul> |                         |                                                                                                                                                                                                                                                                                                                                                                                           |     |     |     |     |

---

## EAPoL Configuration Guidelines

When configuring EAPoL, consider the following guidelines:

- The 802.1X port-based authentication is currently supported only in point-to-point configurations, that is, with a single supplicant connected to an 802.1X-enabled switch port.
- When 802.1X is enabled, a port has to be in the authorized state before any other Layer 2 feature can be operationally enabled. For example, the STG state of a port is operationally disabled while the port is in the unauthorized state.
- The 802.1X supplicant capability is not supported. Therefore, none of its ports can successfully connect to an 802.1X-enabled port of another device, such as another switch, that acts as an authenticator, unless access control on the remote port is disabled or is configured in forced-authorized mode. For example, if a G7028/G7052 is connected to another G7028/G7052, and if 802.1X is enabled on both switches, the two connected ports must be configured in force-authorized mode.
- Unsupported 802.1X attributes include Service-Type, Session-Timeout, and Termination-Action.
- RADIUS accounting service for 802.1X-authenticated devices or users is not currently supported.
- Configuration changes performed using SNMP and the standard 802.1X MIB will take effect immediately.

---

## Chapter 7. Access Control Lists

Access Control Lists (ACLs) are filters that permit or deny traffic for security purposes. They can also be used with QoS to classify and segment traffic to provide different levels of service to different traffic types. Each filter defines the conditions that must match for inclusion in the filter, and also the actions that are performed when a match is made.

IBM Networking OS 7.6 supports the following ACLs:

- IPv4 ACLs

Up to 512 ACLs are supported for networks that use IPv4 addressing. IPv4 ACLs are configured using the following ISCLI command path:

```
RS G7028(config)# access-control list <IPv4 ACL number> ?
```

- IPv6 ACLs

Up to 128 ACLs are supported for networks that use IPv6 addressing. IPv6 ACLs are configured using the following ISCLI command path:

```
RS G7028(config)# access-control list6 <IPv6 ACL number> ?
```

---

## Summary of Packet Classifiers

ACLs allow you to classify packets according to a variety of content in the packet header (such as the source address, destination address, source port number, destination port number, and others). Once classified, packet flows can be identified for more processing.

IPv4 ACLs and IPv6 ACLs allow you to classify packets based on the following packet attributes:

- Ethernet header options (for IPv4 ACLs only)
  - Source MAC address
  - Destination MAC address
  - VLAN number and mask
  - Ethernet type (ARP, IP, IPv6, MPLS, RARP, etc.)
  - Ethernet Priority (the IEEE 802.1p Priority)
- IPv4 header options (for IPv4 ACLs only)
  - Source IPv4 address and subnet mask
  - Destination IPv4 address and subnet mask
  - Type of Service value
  - IP protocol number or name as shown in [Table 8](#):

*Table 8. Well-Known Protocol Types*

| Number | Protocol Name |
|--------|---------------|
| 1      | icmp          |
| 2      | igmp          |
| 6      | tcp           |
| 17     | udp           |
| 89     | ospf          |
| 112    | vrrp          |

- IPv6 header options (for IPv6 ACLs only)
  - Source IPv6 address and prefix length
  - Destination IPv6 address and prefix length
  - Next Header value
  - Flow Label value
  - Traffic Class value

- TCP/UDP header options (for all ACLs)
  - TCP/UDP application source port and mask as shown in [Table 9](#)
  - TCP/UDP application destination port as shown in [Table 9](#)

*Table 9. Well-Known Application Ports*

| Port | TCP/UDP Application | Port | TCP/UDP Application | Port      | TCP/UDP Application |
|------|---------------------|------|---------------------|-----------|---------------------|
| 20   | ftp-data            | 79   | finger              | 179       | bgp                 |
| 21   | ftp                 | 80   | http                | 194       | irc                 |
| 22   | ssh                 | 109  | pop2                | 220       | imap3               |
| 23   | telnet              | 110  | pop3                | 389       | ldap                |
| 25   | smtp                | 111  | sunrpc              | 443       | https               |
| 37   | time                | 119  | nntp                | 520       | rip                 |
| 42   | name                | 123  | ntp                 | 554       | rtsp                |
| 43   | whois               | 143  | imap                | 1645/1812 | Radius              |
| 53   | domain              | 144  | news                | 1813      | Radius              |
| 69   | tftp                | 161  | snmp                | 1985      | Accounting          |
| 70   | gopher              | 162  | snmptrap            |           | hsrp                |

- TCP/UDP flag value as shown in [Table 10](#)

*Table 10. Well-Known TCP flag values*

| Flag | Value  |
|------|--------|
| URG  | 0x0020 |
| ACK  | 0x0010 |
| PSH  | 0x0008 |
| RST  | 0x0004 |
| SYN  | 0x0002 |
| FIN  | 0x0001 |

- Packet format (for IPv4 ACLs only)
  - Ethernet format (eth2, SNAP, LLC)
  - Ethernet tagging format
  - IP format (IPv4, IPv6)
- Egress port packets (for all ACLs)

---

## Summary of ACL Actions

Once classified using ACLs, the identified packet flows can be processed differently. For each ACL, an *action* can be assigned. The action determines how the switch treats packets that match the classifiers assigned to the ACL. G7028/G7052 ACL actions include the following:

- Pass or Drop the packet
- Re-mark the packet with a new DiffServ Code Point (DSCP)
- Re-mark the 802.1p field
- Set the COS queue

---

## Assigning Individual ACLs to a Port

Once you configure an ACL, you must assign the ACL to the appropriate ports. Each port can accept multiple ACLs, and each ACL can be applied for multiple ports. ACLs can be assigned individually, or in groups.

To assign an individual ACLs to a port, use the following IP Interface Mode commands:

```
RS G7028(config)# interface port <port>
RS G7028(config-ip)# access-control list <IPv4 ACL number>
RS G7028(config-ip)# access-control list6 <IPv6 ACL number>
```

When multiple ACLs are assigned to a port, higher-priority ACLs are considered first, and their action takes precedence over lower-priority ACLs. ACL order of precedence is discussed in the next section.

To create and assign ACLs in groups, see [“ACL Groups” on page 87](#).

---

## ACL Order of Precedence

When multiple ACLs are assigned to a port, the order in which the ACLs are applied to port traffic (or whether they are applied at all) depends on the following factors:

- The precedence group in which the ACL resides;
- The ACL number;
- Whether a prior ACL in the precedence group is also matched;
- And whether the ACL action is compatible with preceding ACLs.

ACLs are automatically divided into precedence groups as follows:

Precedence Group 1 includes ACL 1—128.  
Precedence Group 2 includes ACL 129—256.  
Precedence Group 3 includes ACL 257—384.  
Precedence Group 4 includes ACL 385—512.

The switch processes each precedence group in numeric sequence; Precedence group 1 is evaluated first, followed by precedence group 2, and so on.

Within each precedence group, ACLs assigned to the port are processed in numeric sequence, based on ACL number. Lower-numbered ACLs take precedence over higher-numbered ACLs. For example, ACL 1 (if assigned to the port) is evaluated first and has top priority within precedence group 1.

For each precedence group, only the first assigned ACL that matches the port traffic is considered. If multiple ACLs in the precedence group match the traffic, only the one with the lowest ACL number is considered. The others in the precedence group are ignored.

One ACL match from each precedence group is permitted, meaning that up to [four] ACL matches may be considered for action: one from precedence group 1, one from precedence group 2, and so on.

Of the matching ACLs permitted, each configured ACL action is applied in sequence, based on ACL number, with the lowest-numbered ACL's action applied first. If an ACL action contradicts a preceding ACL (one with a lower ACL number), the action of the higher-numbered ACL is ignored.

If no assigned ACL matches the port traffic, no ACL action is applied.

---

## ACL Groups

To assist in organizing multiple ACLs and assigning them to ports, you can place ACLs into ACL Groups, thereby defining complex traffic profiles. ACLs and ACL Groups can then be assigned on a per-port basis. Any specific ACL can be assigned

to multiple ACL Groups, and any ACL or ACL Group can be assigned to multiple ports. If, as part of multiple ACL Groups, a specific ACL is assigned to a port multiple times, only one instance is used. The redundant entries are ignored.

- **Individual ACLs**

The G7028/G7052 supports up to 512 ACLs. Each ACL defines one filter rule for matching traffic criteria. Each filter rule can also include an action (permit or deny the packet). For example:

|                                                                                  |
|----------------------------------------------------------------------------------|
| <b>ACL 1:</b><br>VLAN = 1<br>SIP = 10.10.10.1 (255.255.255.0)<br>Action = permit |
|----------------------------------------------------------------------------------|

- **Access Control List Groups**

An Access Control List Group (ACL Group) is a collection of ACLs. For example:

| ACL Group 1                                                                          |
|--------------------------------------------------------------------------------------|
| <b>ACL 1:</b><br>VLAN = 1<br>SIP = 10.10.10.1 (255.255.255.0)<br>Action = permit     |
| <b>ACL 2:</b><br>VLAN = 2<br>SIP = 10.10.10.2 (255.255.255.0)<br>Action = deny       |
| <b>ACL 3:</b><br>Priority = 7<br>DIP = 10.10.10.3 (255.255.255.0)<br>Action = permit |

ACL Groups organize ACLs into traffic profiles that can be more easily assigned to ports. The G7028/G7052 supports up to 512 ACL Groups.

**Note:** ACL Groups are used for convenience in assigning multiple ACLs to ports. ACL Groups have no effect on the order in which ACLs are applied (see [“ACL Order of Precedence” on page 87](#)). All ACLs assigned to the port (whether individually assigned or part of an ACL Group) are considered as individual ACLs for the purposes of determining their order of precedence.

---

## Assigning ACL Groups to a Port

To assign an ACL Group to a port, use the following command:

```
RS G7028(config-if)# access-control group <ACL group number>
RS G7028(config-if)# exit
```

---

## ACL Metering and Re-Marking

You can define a profile for the aggregate traffic flowing through the G7028/G7052 by configuring a QoS meter (if desired) and assigning ACLs to ports.

**Note:** When you add ACLs to a port, make sure they are ordered correctly in terms of precedence (see [“ACL Order of Precedence” on page 87](#)).

Actions taken by an ACL are called *In-Profile* actions. You can configure additional In-Profile and Out-of-Profile actions on a port. Data traffic can be metered, and re-marked to ensure that the traffic flow provides certain levels of service in terms of bandwidth for different types of network traffic.

### Metering

QoS metering provides different levels of service to data streams through user-configurable parameters. A meter is used to measure the traffic stream against a traffic profile which you create. Thus, creating meters yields In-Profile and Out-of-Profile traffic for each ACL, as follows:

- **In-Profile** If there is no meter configured or if the packet conforms to the meter, the packet is classified as In-Profile.
- **Out-of-Profile** If a meter is configured and the packet does not conform to the meter (exceeds the committed rate or maximum burst rate of the meter), the packet is classified as Out-of-Profile.

Using meters, you set a Committed Rate in Kbps (1000 bits per second in each Kbps). All traffic within this Committed Rate is In-Profile. Additionally, you can set a Maximum Burst Size that specifies an allowed data burst larger than the Committed Rate for a brief period. These parameters define the In-Profile traffic.

Meters keep the sorted packets within certain parameters. You can configure a meter on an ACL, and perform actions on metered traffic, such as packet re-marking.

### Re-Marking

Re-marking allows for the treatment of packets to be reset based on new network specifications or desired levels of service. You can configure the ACL to re-mark a packet as follows:

- Change the DSCP value of a packet, used to specify the service level that traffic receives.
- Change the 802.1p priority of a packet.

---

## ACL Port Mirroring

For IPv4 ACLs, packets that match the filter can be mirrored to another switch port for network diagnosis and monitoring.

The source port for the mirrored packets cannot be a portchannel, but may be a member of a portchannel.

The destination port to which packets are mirrored must be a physical port.

The action (permit, drop, etc.) of the ACL must be configured before assigning it to a port.

Use the following commands to add mirroring to an ACL:

```
RS G7028(config)# access-control list <ACL number> mirror port <destination port>
```

The ACL must be also assigned to its target ports as usual (see [“Assigning Individual ACLs to a Port” on page 86](#), or [“Assigning ACL Groups to a Port” on page 89](#)).

---

## Viewing ACL Statistics

ACL statistics display how many packets have “hit” (matched) each ACL. Use ACL statistics to check filter performance or to debug the ACL filter configuration.

You must enable statistics for each ACL that you wish to monitor:

```
RS G7028(config)# access-control list <ACL number> statistics
```

---

## ACL Logging

ACLs are generally used to enhance port security. Traffic that matches the characteristics (source addresses, destination addresses, packet type, etc.) specified by the ACLs on specific ports is subject to the actions (chiefly permit or deny) defined by those ACLs. Although switch statistics show the number of times particular ACLs are matched, the ACL logging feature can provide additional insight into actual traffic patterns on the switch, providing packet details in the system log for network debugging or security purposes.

## Enabling ACL Logging

By default, ACL logging is disabled. Enable or disable ACL logging on a per-ACL basis as follows:

```
RS G7028(config)# [no] access-control list <IPv4 ACL number> log
RS G7028(config)# [no] access-control list6 <IPv6 ACL number> log
```

## Logged Information

When ACL logging is enabled on any particular ACL, the switch will collect information about packets that match the ACL. The information collected depends on the ACL type:

- For IP-based ACLs, information is collected regarding
  - Source IP address
  - Destination IP address
  - TCP/UDP port number
  - ACL action
  - Number of packets logged

For example:

```
Sep 27  4:20:28 DUT3 NOTICE  ACL-LOG: %IP ACCESS LOG: list
ACL-IP-12-IN denied tcp 1.1.1.1 (0) -> 200.0.1.2 (0), 150
packets.
```

- For MAC-based ACLs, information is collected regarding
  - Source MAC address
  - Source IP address
  - Destination IP address
  - TCP/UDP port number
  - ACL action
  - Number of packets logged

For example:

```
Sep 27  4:25:38 DUT3 NOTICE  ACL-LOG: %MAC ACCESS LOG: list
ACL-MAC-12-IN permitted tcp 1.1.1.2 (0) (12,
00:ff:d7:66:74:62) -> 200.0.1.2 (0) (00:18:73:ee:a7:c6), 32
packets.
```

## Rate Limiting Behavior

Because ACL logging can be CPU-intensive, logging is rate-limited. By default, the switch will log only 10 matching packets per second. This pool is shared by all log-enabled ACLs. The global rate limit can be changed as follows:

```
RS G7028(config)# access-control log rate-limit <1-1000>
```

Where the limit is specified in packets per second.

## Log Interval

For each log-enabled ACL, the first packet that matches the ACL initiates an immediate message in the system log. Beyond that, additional matches are subject to the log interval. By default, the switch will buffer ACL log messages for a period of 300 seconds. At the end of that interval, all messages in the buffer are written to the system log. The global interval value can be changed as follows:

```
RS G7028(config)# access-control log interval <5-600>
```

Where the interval rate is specified in seconds.

In any given interval, packets that have identical log information are condensed into a single message. However, the packet count shown in the ACL log message represents only the logged messages, which due to rate-limiting, may be significantly less than the number of packets actually matched by the ACL.

Also, the switch is limited to 64 different ACL log messages in any interval. Once the threshold is reached, the oldest message will be discarded in favor of the new message, and an overflow message will be added to the system log.

## ACL Logging Limitations

ACL logging reserves packet queue 1 for internal use. Features that allow remapping packet queues (such as CoPP) may not behave as expected if other packet flows are reconfigured to use queue 1.

---

## ACL Configuration Examples

### ACL Example 1

Use this configuration to block traffic to a specific host. All traffic that ingresses on port 1 is denied if it is destined for the host at IP address 100.10.1.1

1. Configure an Access Control List.

```
RS G7028(config)# access-control list 1 ipv4 destination-ip-address 100.10.1.1
RS G7028(config)# access-control list 1 action deny
```

2. Add ACL 1 to port 1.

```
RS G7028(config)# interface port 1
RS G7028(config-if)# access-control list 1
RS G7028(config-if)# exit
```

### ACL Example 2

Use this configuration to block traffic from a network destined for a specific host address. All traffic that ingresses in port 2 with source IP from class 100.10.1.0/24 and destination IP 200.20.2.2 is denied.

1. Configure an Access Control List.

```
RS G7028(config)# access-control list 2 ipv4 source-ip-address 100.10.1.0
255.255.255.0
RS G7028(config)# access-control list 2 ipv4 destination-ip-address 200.20.2.2
255.255.255.255
RS G7028(config)# access-control list 2 action deny
```

2. Add ACL 2 to port 2.

```
RS G7028(config)# interface port 2
RS G7028(config-if)# access-control list 2
RS G7028(config-if)# exit
```

### ACL Example 3

Use this configuration to block traffic from a specific IPv6 source address. All traffic that ingresses in port 2 with source IP from class 2001:0:0:5:0:0:0:2/128 is denied.

1. Configure an Access Control List.

```
RS G7028(config)# access-control list6 3 ipv6 source-address 2001:0:0:5:0:0:0:2
128
RS G7028(config)# access-control list6 3 action deny
```

2. Add ACL 2 to port 2.

```
RS G7028(config)# interface port 2
RS G7028(config-if)# access-control list6 3
RS G7028(config-if)# exit
```

## ACL Example 4

Use this configuration to deny all ARP packets that ingress a port.

1. Configure an Access Control List.

```
RS G7028(config)# access-control list 2 ethernet ethernet-type arp
RS G7028(config)# access-control list 2 action deny
```

2. Add ACL 2 to port EXT2.

```
RS G7028(config)# interface port 2
RS G7028(config-if)# access-control list 2
RS G7028(config-if)# exit
```

## ACL Example 5

Use the following configuration to permit access to hosts with destination MAC address that matches 11:05:00:10:00:00 FF:F5:FF:FF:FF:FF and deny access to all other hosts.

1. Configure Access Control Lists.

```
RS G7028(config)# access-control list 30 ethernet destination-mac-address
11:05:00:10:00:00 FF:F5:FF:FF:FF:FF
RS G7028(config)# access-control list 30 action permit
RS G7028(config)# access-control list 100 ethernet destination-mac-address
00:00:00:00:00:00 00:00:00:00:00:00
RS G7028(config)# access-control list 100 action deny
```

2. Add ACLs to a port.

```
RS G7028(config)# interface port 2
RS G7028(config-if)# access-control list 30
RS G7028(config-if)# access-control list 100
RS G7028(config-if)# exit
```

## ACL Example 6

This configuration blocks traffic from a network that is destined for a specific egress port. All traffic that ingresses port 1 from the network 100.10.1.0/24 and is destined for port 3 is denied.

1. Configure an Access Control List.

```
RS G7028(config)# access-control list 4 ipv4 source-ip-address 100.10.1.0
255.255.255.0
RS G7028(config)# access-control list 4 egress-port 3
RS G7028(config)# access-control list 4 action deny
```

2. Add ACL 4 to port 1.

```
RS G7028(config)# interface port 1
RS G7028(config-if)# access-control list 4
RS G7028(config-if)# exit
```

---

## Using Storm Control Filters

The G7028/G7052 provides filters that can limit the number of the following packet types transmitted by switch ports:

- Broadcast packets
- Multicast packets
- Unknown unicast packets (destination lookup failure)

### Broadcast Storms

Excessive transmission of broadcast or multicast traffic can result in a broadcast storm. A broadcast storm can overwhelm your network with constant broadcast or multicast traffic, and degrade network performance. Common symptoms of a broadcast storm are slow network response times and network operations timing out.

Unicast packets whose destination MAC address is not in the Forwarding Database are *unknown unicasts*. When an unknown unicast is encountered, the switch handles it like a broadcast packet and floods it to all other ports in the VLAN (broadcast domain). A high rate of unknown unicast traffic can have the same negative effects as a broadcast storm.

### Configuring Storm Control

Configure broadcast filters on each port that requires broadcast storm control. Set a threshold that defines the total number of broadcast packets transmitted (0-2097151), in packets per second. When the threshold is reached, no more packets of the specified type are transmitted.

To filter broadcast packets on a port, use the following commands:

```
RS G7028(config)# interface port 1
RS G7028(config-if)# storm-control broadcast level pps <packets per second>
```

To filter multicast packets on a port, use the following commands:

```
RS G7028(config-if)# storm-control multicast level pps <packets per second>
```

To filter unknown unicast packets on a port, use the following commands:

```
RS G7028(config-if)# storm-control unicast level pps <packets per second>
RS G7028(config-if)# exit
```



# Part 3: Switch Basics

This section discusses basic switching functions:

- VLANs
- Port Trunking
- Spanning Tree Protocols (Spanning Tree Groups, Rapid Spanning Tree Protocol, and Multiple Spanning Tree Protocol)
- Virtual Link Aggregation Groups
- Quality of Service



---

## Chapter 8. VLANs

This chapter describes network design and topology considerations for using Virtual Local Area Networks (VLANs). VLANs commonly are used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments. The following topics are discussed in this chapter:

- [“VLANs and Port VLAN ID Numbers” on page 100](#)
- [“VLAN Tagging/Trunk Mode” on page 102](#)
- [“VLAN Topologies and Design Considerations” on page 106](#)

This section discusses how you can connect users and segments to a host that supports many logical segments or subnets by using the flexibility of the multiple VLAN system.

- [“Protocol-Based VLANs” on page 109](#)
- [“Private VLANs” on page 112](#)

**Note:** VLANs can be configured from the Command Line Interface (see “VLAN Configuration” as well as “Port Configuration” in the *Command Reference*).

---

## VLANs Overview

Setting up virtual LANs (VLANs) is a way to segment networks to increase network flexibility without changing the physical network topology. With network segmentation, each switch port connects to a segment that is a single broadcast domain. When a switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belong to one broadcast domain.

Ports are grouped into broadcast domains by assigning them to the same VLAN. Frames received in one VLAN can only be forwarded within that VLAN, and multicast, broadcast, and unknown unicast frames are flooded only to ports in the same VLAN.

The RackSwitch G7028/G7052 (G7028/G7052) supports jumbo frames with a Maximum Transmission Unit (MTU) of 12,288 bytes. Within each frame, 18 bytes are reserved for the Ethernet header and CRC trailer. The remaining space in the frame (up to 12,200 bytes) comprise the packet, which includes the payload of up to 12,002 bytes and any additional overhead, such as 802.1q or VLAN tags. Jumbo frame support is automatic: it is enabled by default, requires no manual configuration, and cannot be manually disabled.

---

## VLANs and Port VLAN ID Numbers

### VLAN Numbers

The G7028/G7052 supports up to 4095 VLANs per switch. Each can be identified with any number between 1 and 4094. VLAN 1 is the default VLAN for the data ports. VLAN 4095 is used by the management network, which includes the management port.

Use the following command to view VLAN information:

```
RS G7028# show vlan
```

| VLAN | Name         | Status | Ports  |
|------|--------------|--------|--------|
| 1    | Default VLAN | ena    | 1-XGE4 |
| 2    | VLAN 2       | dis    | empty  |
| 4095 | Mgmt VLAN    | ena    | MGMT   |

## PVID/Native VLAN Numbers

Each port in the switch has a configurable default VLAN number, known as its *PVID/Native VLAN*. By default, the PVID/Native VLAN for all non-management ports is set to 1, which correlates to the default VLAN ID. The PVID/Native VLAN for each port can be configured to any VLAN number between 1 and 4094.

Use the following command to view PVIDs/Native VLANs:

```
RS G7028# show interface information
```

| Alias | Port | Tag | RMON | Ln | Fld | PVID  | DESCRIPTION | VLAN(s) |
|-------|------|-----|------|----|-----|-------|-------------|---------|
|       |      | Trk |      |    |     | NVLAN |             |         |
| 1     | 1    | n   | d    | e  | e   | 1     |             | 1       |
| 2     | 2    | n   | d    | e  | e   | 1     |             | 1       |
| 3     | 3    | n   | d    | e  | e   | 1     |             | 1       |
| 4     | 4    | n   | d    | e  | e   | 1     |             | 1       |
| ...   | ...  |     |      |    |     |       | ...         |         |
| XGE3  | 27   | n   | d    | e  | e   | 1     |             | 1       |
| XGE4  | 28   | n   | d    | e  | e   | 1     |             | 1       |

\* = PVID/Native-VLAN is tagged.  
Trk = Trunk mode  
NVLAN = Native-VLAN

Use the following command to set the port PVID/Native VLAN:

```
Access Mode Port

RS G7028(config)# interface port <port number>
RS G7028(config-if)# switchport access vlan <VLAN ID>

For Trunk Mode Port

RS G7028(config)# interface port <port number>
RS G7028(config-if)# switchport trunk native vlan <VLAN ID>
```

Each port on the switch can belong to one or more VLANs, and each VLAN can have any number of switch ports in its membership. Any port that belongs to multiple VLANs, however, must have VLAN *tagging/trunk mode* enabled (see [“VLAN Tagging/Trunk Mode” on page 102](#)).

---

## VLAN Tagging/Trunk Mode

IBM Networking OS software supports 802.1Q VLAN *tagging/trunk mode*, providing standards-based VLAN support for Ethernet systems.

Tagging/trunk mode places the VLAN identifier in the frame header of a packet, allowing each port to belong to multiple VLANs. When you add a port to multiple VLANs, you also must enable tagging/trunk mode on that port.

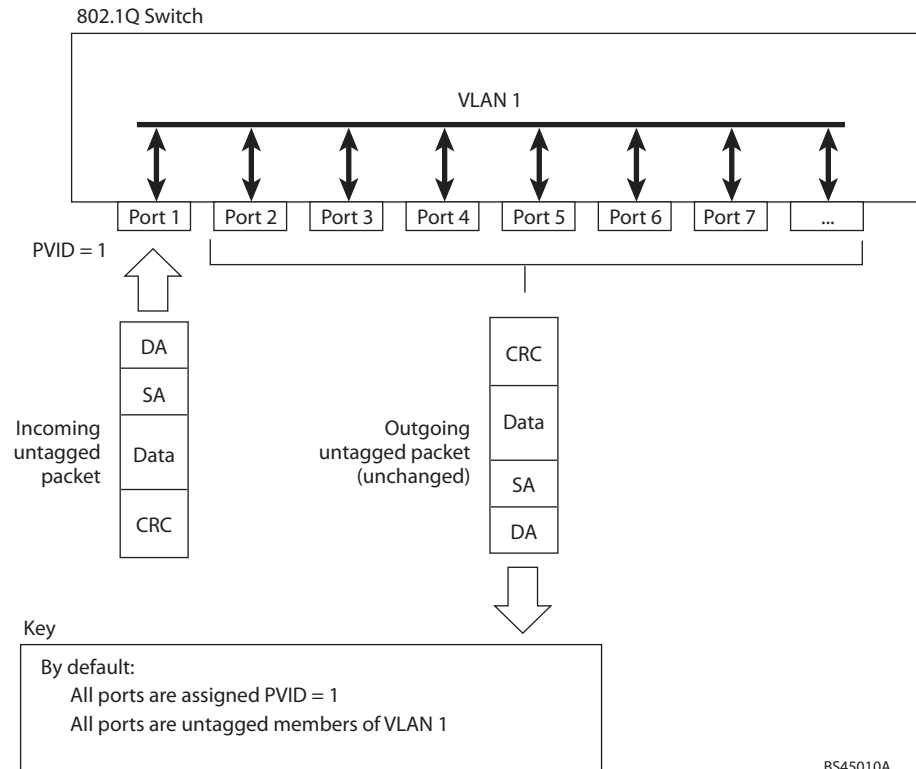
Since tagging/trunk mode fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged/trunk mode frames from being transmitted to devices that do not support 802.1Q VLAN tags, or devices where tagging/trunk mode is not enabled.

Important terms used with the 802.1Q tagging/trunk mode feature are:

- VLAN identifier (VID)—the 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN.
- Port VLAN identifier (PVID)/Native VLAN—a classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID/Native VLAN of 3 (PVID =3) assigns all untagged/access mode frames received on this port to VLAN 3. Any untagged/access mode frames received by the switch are classified with the PVID/native VLAN of the receiving port.
- Tagged/Trunk Mode frame—a frame that carries VLAN tagging/trunk mode information in the header. This VLAN tagging/trunk mode information is a 32-bit field (VLAN tag) in the frame header that identifies the frame as belonging to a specific VLAN. Untagged/access mode frames are marked (tagged/trunk mode) with this classification as they leave the switch through a port that is configured as a tagged/trunk mode port.
- Untagged/Access Mode frame—a frame that does not carry any VLAN tagging/trunk mode information in the frame header.
- Untagged/Access Mode member—a port that has been configured as an untagged/access mode member of a specific VLAN. When an untagged/access mode frame exits the switch through an untagged/access mode member port, the frame header remains unchanged. When a tagged/trunk mode frame exits the switch through an untagged/access mode member port, the tag is stripped and the tagged/trunk mode frame is changed to an untagged/access mode frame.
- Tagged/Trunk Mode member—a port that has been configured as a tagged/trunk mode member of a specific VLAN. When an untagged/access mode frame exits the switch through a tagged/trunk mode member port, the frame header is modified to include the 32-bit tag associated with the PVID/native VLAN. When a tagged/trunk mode frame exits the switch through a tagged/trunk mode member port, the frame header remains unchanged (original VID remains).

**Note:** If a 802.1Q tagged/trunk mode frame is received by a port that has VLAN-tagging/trunk mode disabled and the port VLAN ID (PVID/native VLAN) is different than the VLAN ID of the packet, then the frame is dropped at the ingress port.

Figure 2. Default VLAN settings



**Note:** The port numbers specified in these illustrations may not directly correspond to the physical port configuration of your switch model.

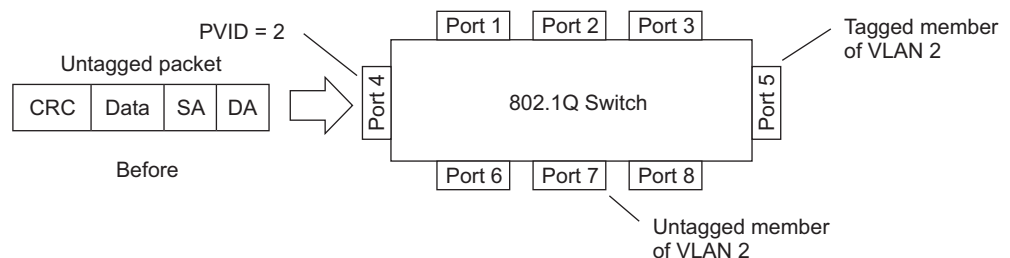
When a VLAN is configured, ports are added as members of the VLAN, and the ports are defined as either *tagged* or *untagged* (see [Figure 3](#) through [Figure 6](#)).

The default configuration settings for the G7028/G7052 has all ports set as untagged members of VLAN 1 with all ports configured as PVID = 1. In the default configuration example shown in [Figure 2](#), all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID = 1).

[Figure 3](#) through [Figure 6](#) illustrate generic examples of VLAN tagging/trunk mode. In [Figure 3](#), untagged incoming packets are assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

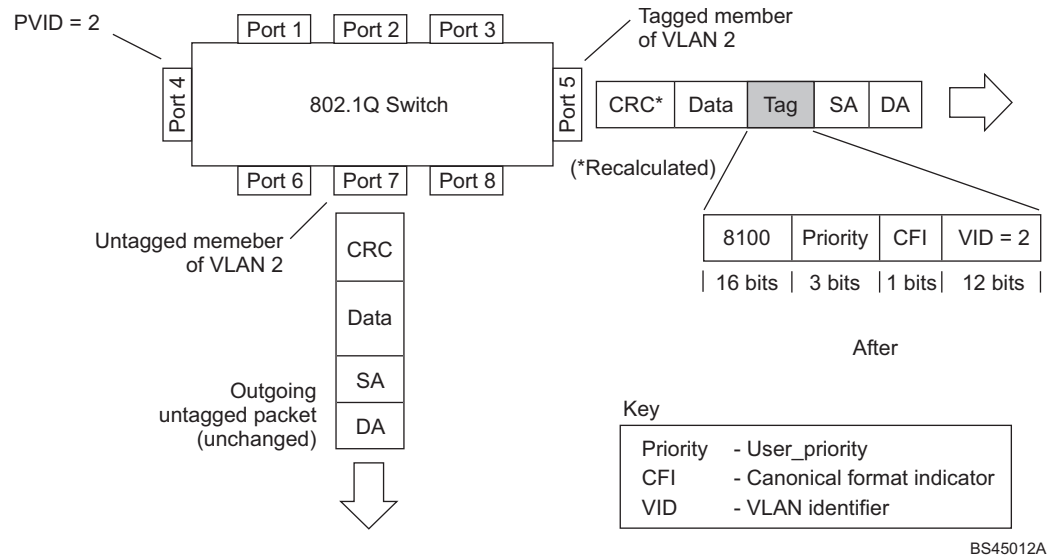
**Note:** The port assignments in the following figures are not meant to match the G7028/G7052.

Figure 3. Port-based VLAN assignment



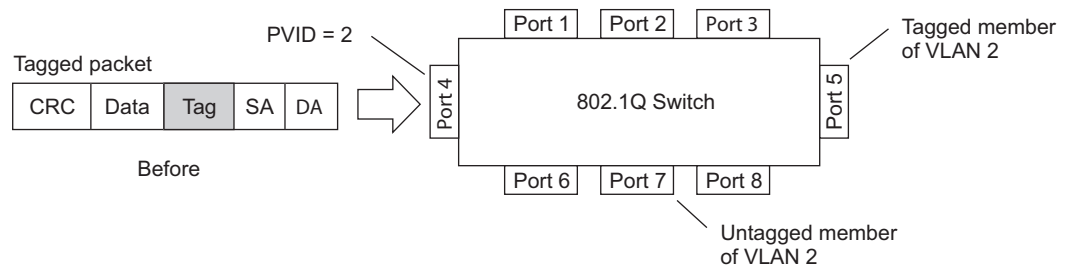
As shown in Figure 4, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 4. 802.1Q tagging/trunk mode (after port-based VLAN assignment)



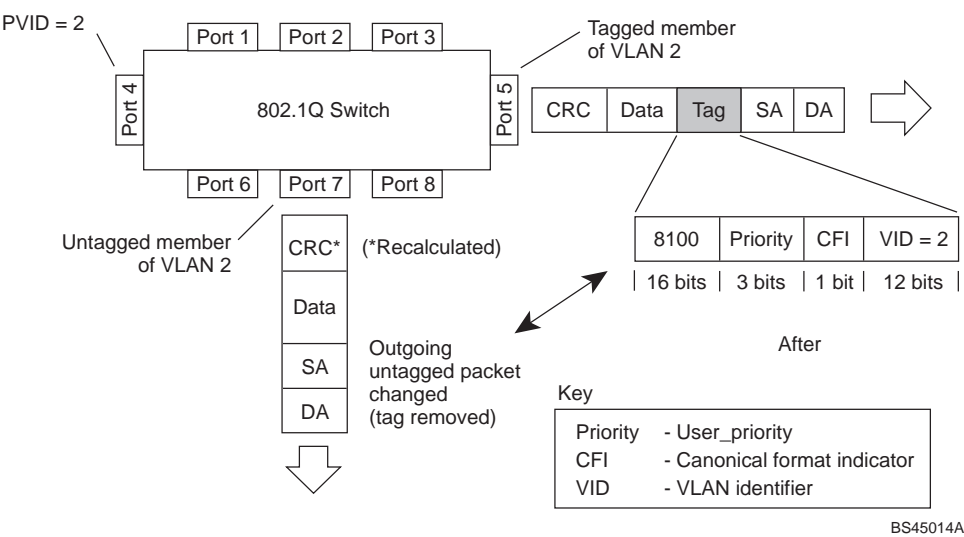
In Figure 5, tagged incoming packets are assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

Figure 5. 802.1Q tag assignment



As shown in Figure 6, the tagged packet remains unchanged as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 6. 802.1Q tagging/trunk mode (after 802.1Q tag assignment)



---

## VLAN Topologies and Design Considerations

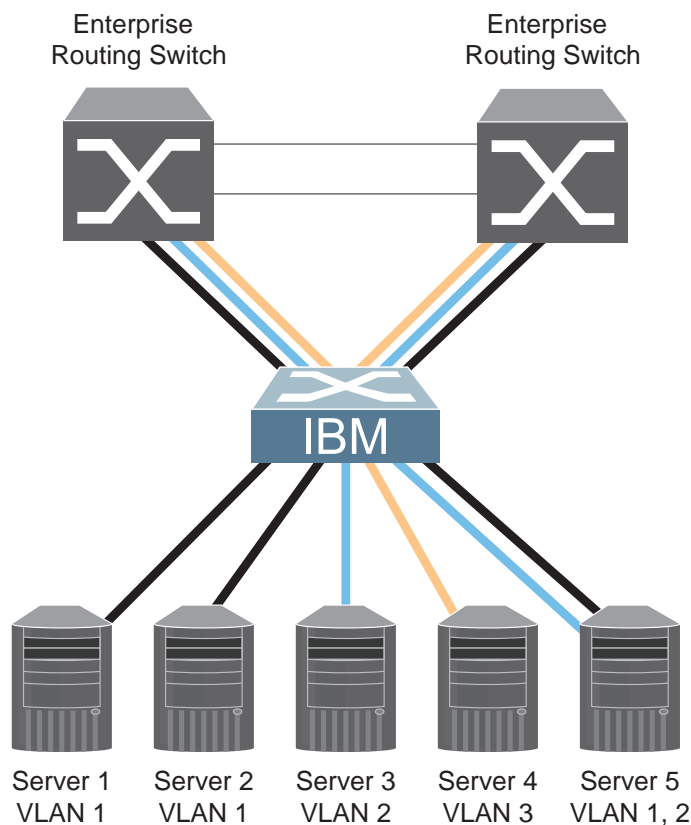
Note the following when working with VLAN topologies:

- By default, the G7028/G7052 software is configured so that tagging/trunk mode is disabled on all ports.
- By default, the G7028/G7052 software is configured so that all data ports are members of VLAN 1.
- By default, the IBM N/OS software is configured so that the management port is a member of VLAN 4095 (the management VLAN).
- STG 128 is reserved for switch management.
- When using Spanning Tree, STG 2-128 may contain only one VLAN unless Multiple Spanning-Tree Protocol (MSTP) mode is used. With MSTP mode, STG 1 to 32 can include multiple VLANs.
- All ports involved in both trunking and port mirroring must have the same VLAN configuration. If a port is on a trunk with a mirroring port, the VLAN configuration cannot be changed. For more information trunk groups, see [“Ports and Trunking” on page 115](#) and [“Port Mirroring” on page 253](#).

### Multiple VLANs with Tagging/Trunk Mode Adapters

[Figure 7](#) illustrates a network topology described in [Note: on page 107](#) and the configuration example on [page 108](#).

Figure 7. Multiple VLANs with VLAN-Tagged Gigabit Adapters



The features of this VLAN are described in the following table.

Table 11. Multiple VLANs Example

| Component                   | Description                                                                                                                                                                                                                                                                                          |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| G7028/G7052 switch          | This switch is configured with three VLANs that represent three different IP subnets. Five ports are connected downstream to servers. Two ports are connected upstream to routing switches. Uplink ports are members of all three VLANs, with VLAN tagging/trunk mode enabled.                       |
| Server 1                    | This server is a member of VLAN 1 and has presence in only one IP subnet. The associated switch port is only a member of VLAN 1, so tagging/trunk mode is disabled.                                                                                                                                  |
| Server 2                    | This server is a member of VLAN 1 and has presence in only one IP subnet. The associated switch port is only a member of VLAN 1, so tagging/trunk mode is disabled.                                                                                                                                  |
| Server 3                    | This server belongs to VLAN 2, and it is logically in the same IP subnet as Server 5. The associated switch port has tagging/trunk mode disabled.                                                                                                                                                    |
| Server 4                    | A member of VLAN 3, this server can communicate only with other servers via a router. The associated switch port has tagging/trunk mode disabled.                                                                                                                                                    |
| Server 5                    | A member of VLAN 1 and VLAN 2, this server can communicate only with Server 1, Server 2, and Server 3. The associated switch port has tagging/trunk mode enabled.                                                                                                                                    |
| Enterprise Routing switches | These switches must have all three VLANs (VLAN 1, 2, 3) configured. They can communicate with Server 1, Server 2, and Server 5 via VLAN 1. They can communicate with Server 3 and Server 5 via VLAN 2. They can communicate with Server 4 via VLAN 3. Tagging/trunk mode on switch ports is enabled. |

**Note:** VLAN tagging/trunk mode is required only on ports that are connected to other switches or on ports that connect to tag-capable end-stations, such as servers with VLAN-tagging/trunk mode adapters.

## VLAN Configuration Example

Use the following procedure to configure the example network shown in [Figure 7 on page 106](#).

1. Enable VLAN tagging/trunk mode on server ports that support multiple VLANs.

```
RS G7028(config)# interface port 5
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 2
RS G7028(config-if)# exit
```

2. Enable tagging/trunk mode on uplink ports that support multiple VLANs.

```
RS G7028(config)# interface port 19
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 2,3
RS G7028(config-if)# exit
RS G7028(config)# interface port 20
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 2,3
RS G7028(config-if)# exit
```

3. Configure server ports that belong to a single VLAN.

```
RS G7028(config)# interface port 4
RS G7028(config-if)# switchport mode access
RS G7028(config-if)# switchport access vlan 2
RS G7028(config-if)# exit
```

By default, all ports are members of VLAN 1, so configure only those ports that belong to other VLANs.

---

## Protocol-Based VLANs

Protocol-based VLANs (PVLANS) allow you to segment network traffic according to the network protocols in use. Traffic for supported network protocols can be confined to a particular port-based VLAN. You can give different priority levels to traffic generated by different network protocols.

With PVLAN, the switch classifies incoming packets by Ethernet protocol of the packets, not by the configuration of the ingress port. When an untagged or priority-tagged frame arrives at an ingress port, the protocol information carried in the frame is used to determine a VLAN to which the frame belongs. If a frame's protocol is not recognized as a pre-defined PVLAN type, the ingress port's PVID is assigned to the frame. When a tagged frame arrives, the VLAN ID in the frame's tag is used.

Each VLAN can contain up to eight different PVLANS. You can configure separate PVLANS on different VLANs, with each PVLAN segmenting traffic for the same protocol type. For example, you can configure PVLAN 1 on VLAN 2 to segment IPv4 traffic, and PVLAN 8 on VLAN 100 to segment IPv4 traffic.

To define a PVLAN on a VLAN, configure a PVLAN number (1-8) and specify the frame type and the Ethernet type of the PVLAN protocol. You must assign at least one port to the PVLAN before it can function. Define the PVLAN frame type and Ethernet type as follows:

- Frame type—consists of one of the following values:
  - Ether2 (Ethernet II)
  - SNAP (Subnetwork Access Protocol)
  - LLC (Logical Link Control)
- Ethernet type—consists of a 4-digit (16 bit) hex value that defines the Ethernet type. You can use common Ethernet protocol values, or define your own values. Following are examples of common Ethernet protocol values:
  - IPv4 = 0800
  - IPv6 = 86dd
  - ARP = 0806

## Port-Based vs. Protocol-Based VLANs

Each VLAN supports both port-based and protocol-based association, as follows:

- The default VLAN configuration is port-based. All data ports are members of VLAN 1, with no PVLAN association.
- When you add ports to a PVLAN, the ports become members of both the port-based VLAN and the PVLAN. For example, if you add port 1 to PVLAN 1 on VLAN 2, the port also becomes a member of VLAN 2.
- When you delete a PVLAN, its member ports remain members of the port-based VLAN. For example, if you delete PVLAN 1 from VLAN 2, port 1 remains a member of VLAN 2.
- When you delete a port from a VLAN, the port is deleted from all corresponding PVLANS.

## PVLAN Priority Levels

You can assign each PVLAN a priority value of 0-7, used for Quality of Service (QoS). PVLAN priority takes precedence over a port's configured priority level. If no priority level is configured for the PVLAN (priority = 0), each port's priority is used (if configured).

All member ports of a PVLAN have the same PVLAN priority level.

## PVLAN Tagging/Trunk Mode

When PVLAN tagging/trunk mode is enabled, the switch tags frames that match the PVLAN protocol. For more information about tagging/trunk mode, see [“VLAN Tagging/Trunk Mode” on page 102](#).

Untagged ports must have PVLAN tagging/trunk mode disabled. Tagged ports can have PVLAN tagging/trunk mode either enabled or disabled.

PVLAN tagging/trunk mode has higher precedence than port-based tagging/trunk mode. If a port is tag/trunk mode enabled, and the port is a member of a PVLAN, the PVLAN tags egress frames that match the PVLAN protocol.

Use the tag list command (`protocol-vlan <x> tag-pvlan`) to define the complete list of tag-enabled ports in the PVLAN. Note that all ports not included in the PVLAN tag list will have PVLAN tagging/trunk mode disabled.

## PVLAN Configuration Guidelines

Consider the following guidelines when you configure protocol-based VLANs:

- Each port can support up to 16 VLAN protocols.
- The G7028/G7052 can support up to 16 protocols simultaneously.
- Each PVLAN must have at least one port assigned before it can be activated.
- The same port within a port-based VLAN can belong to multiple PVLANS.
- An untagged port can be a member of multiple PVLANS.
- A port cannot be a member of different VLANs with the same protocol association.

## Configuring PVLAN

Follow this procedure to configure a Protocol-based VLAN (PVLAN).

1. Configure VLAN tagging/trunk mode for ports.

```
RS G7028(config)# interface port 1, 2
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# exit
```

2. Create a VLAN and define the protocol type(s) supported by the VLAN.

```
RS G7028(config)# vlan 2
RS G7028(config-vlan)# protocol-vlan 1 frame-type ether2 0800
```

3. Configure the priority value for the protocol.

```
RS G7028(config-vlan)# protocol-vlan 1 priority 2
```

4. Add member ports for this PVLAN.

```
RS G7028(config-vlan)# protocol-vlan 1 member 1, 2
```

**Note:** If VLAN tagging/trunk mode is turned on and the port being added to the VLAN has a different default VLAN (PVID), you will be asked to confirm changing the PVID to the current VLAN.

5. Enable the PVLAN.

```
RS G7028(config-vlan)# protocol-vlan 1 enable
RS G7028(config-vlan)# exit
```

6. Verify PVLAN operation.

```
RS G7028(config)# show vlan
```

| VLAN | Name         | Status | Ports           |
|------|--------------|--------|-----------------|
| 1    | Default VLAN | ena    | 1-48, XGE1-XGE4 |
| 2    | VLAN 2       | ena    | 1 2             |

| PVLAN | Protocol | FrameType | EtherType | Priority | Status  | Ports |
|-------|----------|-----------|-----------|----------|---------|-------|
| 2     | 1        | Ether2    | 0800      | 2        | enabled | 1 2   |

| PVLAN | PVLAN-Tagged Ports |
|-------|--------------------|
| none  | none               |

---

## Private VLANs

Private VLANs provide Layer 2 isolation between the ports within the same broadcast domain. Private VLANs can control traffic within a VLAN domain, and provide port-based security for host servers.

Use Private VLANs to partition a VLAN domain into sub-domains. Each sub-domain is comprised of one primary VLAN and one or more secondary VLANs, as follows:

- **Primary VLAN**—carries unidirectional traffic downstream from promiscuous ports. Each Private VLAN configuration has only one primary VLAN. All ports in the Private VLAN are members of the primary VLAN.
- **Secondary VLAN**—Secondary VLANs are internal to a private VLAN domain, and are defined as follows:
  - **Isolated VLAN**—carries unidirectional traffic upstream from the host servers toward ports in the primary VLAN and the gateway. Each Private VLAN configuration can contain only one isolated VLAN.
  - **Community VLAN**—carries upstream traffic from ports in the community VLAN to other ports in the same community, and to ports in the primary VLAN and the gateway. Each Private VLAN configuration can contain multiple community VLANs.

After you define the primary VLAN and one or more secondary VLANs, you map the secondary VLAN(s) to the primary VLAN.

## Private VLAN Ports

Private VLAN ports are defined as follows:

- **Promiscuous**—A promiscuous port is a port that belongs to the primary VLAN. The promiscuous port can communicate with all the interfaces, including ports in the secondary VLANs (Isolated VLAN and Community VLANs). Each promiscuous port can belong to only one Private VLAN.
- **Isolated**—An isolated port is a host port that belongs to an isolated VLAN. Each isolated port has complete layer 2 separation from other ports within the same private VLAN (including other isolated ports), except for the promiscuous ports.
  - Traffic sent to an isolated port is blocked by the Private VLAN, except the traffic from promiscuous ports.
  - Traffic received from an isolated port is forwarded only to promiscuous ports.
- **Community**—A community port is a host port that belongs to a community VLAN. Community ports can communicate with other ports in the same community VLAN, and with promiscuous ports. These interfaces are isolated at layer 2 from all other interfaces in other communities and from isolated ports within the Private VLAN.

## Configuration Guidelines

The following guidelines apply when configuring Private VLANs:

- The default VLAN 1 cannot be a Private VLAN.
- The management VLAN 4095 cannot be a Private VLAN. Management ports cannot be members of Private VLANs.
- IGMP Snooping must be disabled on isolated VLANs.
- Each secondary port's (isolated port and community ports) PVID/Native VLAN must match its corresponding secondary VLAN ID.
- Ports within a secondary VLAN cannot be members of other VLANs.
- All VLANs that comprise the Private VLAN must belong to the same Spanning Tree Group.
- LACP cannot be enabled on ports that are members of a Private VLAN.

## Configuration Example

Follow this procedure to configure a Private VLAN.

1. Select a VLAN and define the Private VLAN type as primary.

```
RS G7028(config)# vlan 700
RS G7028(config-vlan)# private-vlan primary
RS G7028(config-vlan)# exit
```

2. Configure a promiscuous port for VLAN 700.

```
RS G7028(config)# interface port 1
RS G7028(config-if)# switchport mode private-vlan promiscuous
RS G7028(config-if)# switchport private-vlan mapping 700
RS G7028(config-if)# exit
```

3. Configure two secondary VLANs: isolated VLAN and community VLAN.

```
RS G7028(config)# vlan 701
RS G7028(config-vlan)# private-vlan isolated
RS G7028(config-vlan)# exit
RS G7028(config)# vlan 702
RS G7028(config-vlan)# private-vlan community
RS G7028(config-vlan)# exit
```

4. Map secondary VLANs to primary VLAN.

```
RS G7028(config)# vlan 700
RS G7028(config-vlan)# private-vlan association 701,702
RS G7028(config-vlan)# exit
```

5. Configure host ports for secondary VLANs.

```
RS G7028(config)# interface port 2
RS G7028(config-if)# switchport mode private-vlan host
RS G7028(config-if)# switchport private-vlan association 700 701
RS G7028(config-if)# exit

RS G7028(config)# interface port 3
RS G7028(config-if)# switchport mode private-vlan host
RS G7028(config-if)# switchport private-vlan association 700 702
RS G7028(config-if)# exit
```

6. Verify the configuration.

```
RS G7028(config)# show vlan private-vlan
```

| Private-VLAN | Type      | Mapped-To | Status | Ports |
|--------------|-----------|-----------|--------|-------|
| 700          | primary   | 701,702   | ena    | 1     |
| 701          | isolated  | 700       | ena    | 2     |
| 702          | community | 700       | ena    | 3     |

---

## Chapter 9. Ports and Trunking

Trunk groups can provide super-bandwidth, multi-link connections between the RackSwitch G7028/G7052 (G7028/G7052) and other trunk-capable devices. A trunk group is a group of ports that act together, combining their bandwidth to create a single, larger virtual link. This chapter provides configuration background and examples for trunking multiple ports together:

- [“Trunking Overview” on page 116](#)
- [“Configuring a Static Port Trunk” on page 118](#)
- [“Configurable Trunk Hash Algorithm” on page 123](#)
- [“Link Aggregation Control Protocol” on page 120](#)

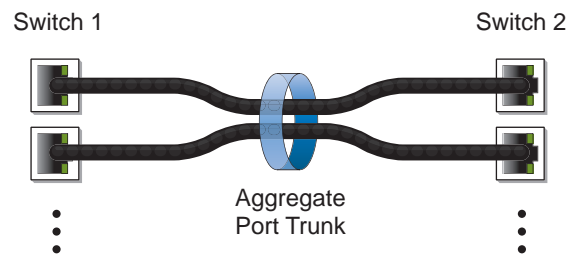
---

## Trunking Overview

When using port trunk groups between two switches, as shown in [Figure 8](#), you can create a virtual link between the switches, operating with combined throughput levels that depends on how many physical ports are included.

Two trunk types are available: static trunk groups (portchannel), and dynamic LACP trunk groups. Each G7028/G7052 supports up to 64 trunk groups—64 static and 64 dynamic. Each type can contain up to 16 member ports, depending on the port type and availability.

Figure 8. Port Trunk Group



Trunk groups are also useful for connecting a G7028/G7052 to third-party devices that support link aggregation, such as Cisco routers and switches with EtherChannel technology (*not* ISL trunking technology) and Sun's Quad Fast Ethernet Adapter. Trunk Group technology is compatible with these devices when they are configured manually.

Trunk traffic is statistically distributed among the ports in a trunk group, based on a variety of configurable options.

Also, since each trunk group is comprised of multiple physical links, the trunk group is inherently fault tolerant. As long as one connection between the switches is available, the trunk remains active and statistical load balancing is maintained whenever a port in a trunk group is lost or returned to service.

---

## Static Trunks

### Static Trunk Requirements

When you create and enable a static trunk, the trunk members (switch ports) take on certain settings necessary for correct operation of the trunking feature.

Before you configure your trunk, you must consider these settings, along with specific configuration rules, as follows:

1. Read the configuration rules provided in the section, [“Static Trunk Group Configuration Rules” on page 117](#).
2. Determine which switch ports (up to 16) are to become *trunk members* (the specific ports making up the trunk).
3. Ensure that the chosen switch ports are set to `enabled`. Trunk member ports must have the same VLAN and Spanning Tree configuration.
4. Consider how the existing Spanning Tree will react to the new trunk configuration. See [Chapter 10, “Spanning Tree Protocols,”](#) for Spanning Tree Group configuration guidelines.
5. Consider how existing VLANs will be affected by the addition of a trunk.

### Static Trunk Group Configuration Rules

The trunking feature operates according to specific configuration rules. When creating trunks, consider the following rules that determine how a trunk group reacts in any network topology:

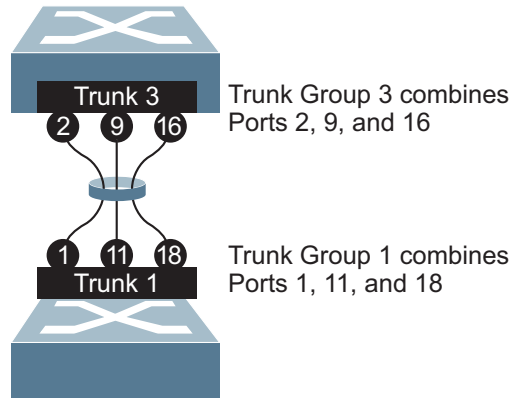
- All trunks must originate from one logical device, and lead to one logical destination device. Usually, a trunk connects two physical devices together with multiple links. However, in some networks, a single logical device may include multiple physical devices when using VLAGs (see [“Virtual Link Aggregation Groups” on page 147](#)). In such cases, links in a trunk are allowed to connect to multiple physical devices because they act as one logical device.
- Any physical switch port can belong to only one trunk group.
- Trunking from third-party devices must comply with Cisco® EtherChannel® technology.
- All trunk member ports must be assigned to the same VLAN configuration before the trunk can be enabled.
- When an active port is configured in a trunk, the port becomes a *trunk member* when you enable the trunk. The Spanning Tree parameters for the port then change to reflect the new trunk settings.
- All trunk members must be in the same Spanning Tree Group (STG) and can belong to only one Spanning Tree Group (STG). However if all ports are *tagged*, then all trunk ports can belong to multiple STGs.
- If you change the Spanning Tree participation of any trunk member to `enabled` or `disabled`, the Spanning Tree participation of all trunk members changes similarly.
- When a trunk is enabled, the trunk’s Spanning Tree participation setting takes precedence over that of any trunk member.
- You cannot configure a trunk member as a monitor port in a port-mirroring configuration.

- Trunks cannot be monitored by a monitor port; however, trunk members can be monitored.
- All ports in static trunks must have the same link configuration (speed, duplex, flow control).

## Configuring a Static Port Trunk

In the following example, three ports are trunked between two switches.

Figure 9. Port Trunk Group Configuration Example



Prior to configuring each switch in this example, you must connect to the appropriate switches as the administrator.

**Note:** For details about accessing and using any of the commands described in this example, see the *IBM System Networking RackSwitch G7028/G7052 ISCLI Reference*.

1. Follow these steps on the G7028/G7052:
  - a. Define a trunk group.

```
RS G7028(config)# portchannel 3 port 2,9,16
RS G7028(config)# portchannel 3 enable
```

- b. Verify the configuration.

```
# show portchannel information
```

Examine the resulting information. If any settings are incorrect, make appropriate changes.

2. Repeat the process on the other switch.

```
RS G7028(config)# portchannel 1 port 1,11,18
RS G7028(config)# portchannel 1 enable
```

3. Connect the switch ports that will be members in the trunk group.

Trunk group 3 (on the G7028/G7052) is now connected to trunk group 1 (on the other switch).

**Note:** In this example, two G7028/G7052 switches are used. If a third-party device supporting link aggregation is used (such as Cisco routers and switches with EtherChannel technology or Sun's Quad Fast Ethernet Adapter), trunk groups on the third-party device must be configured manually. Connection problems could arise when using automatic trunk group negotiation on the third-party device.

4. Examine the trunking information on each switch.

```
# show portchannel information
PortChannel 3: Enabled
Protocol-Static
port state:
  2: STG 1 forwarding
  9: STG 1 forwarding
 16: STG 1 forwarding
```

Information about each port in each configured trunk group is displayed. Make sure that trunk groups consist of the expected ports and that each port is in the expected state.

The following restrictions apply:

- Any physical switch port can belong to only one trunk group.
- Up to 16 ports can belong to the same trunk group.
- All ports in static trunks must have the same link configuration (speed, duplex, flow control).
- Trunking from third-party devices must comply with Cisco® EtherChannel® technology.

---

# Link Aggregation Control Protocol

## LACP Overview

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard for grouping several physical ports into one logical port (known as a dynamic trunk group or Link Aggregation group) with any device that supports the standard. Please refer to IEEE 802.3ad-2002 for a full description of the standard.

The 802.3ad standard allows standard Ethernet links to form a single Layer 2 link using the Link Aggregation Control Protocol (LACP). Link aggregation is a method of grouping physical link segments of the same media type and speed in full duplex, and treating them as if they were part of a single, logical link segment. If a link in a LACP trunk group fails, traffic is reassigned dynamically to the remaining link(s) of the dynamic trunk group.

**Note:** LACP implementation in the IBM Networking OS does not support the Churn machine, an option used to detect if the port is operable within a bounded time period between the actor and the partner. Only the Marker Responder is implemented, and there is no marker protocol generator.

A port's Link Aggregation Identifier (LAG ID) determines how the port can be aggregated. The Link Aggregation ID (LAG ID) is constructed mainly from the *system ID* and the port's *admin key*, as follows:

- **System ID:** an integer value based on the switch's MAC address and the system priority assigned in the CLI.
- **Admin key:** a port's Admin key is an integer value (1-65535) that you can configure in the CLI. Each switch port that participates in the same LACP trunk group must have the same *admin key* value. The Admin key is *local significant*, which means the partner switch does not need to use the same Admin key value.

For example, consider two switches, an Actor (the G7028/G7052) and a Partner (another switch), as shown in [Table 12](#).

Table 12. Actor vs. Partner LACP configuration

| Actor Switch             | Partner Switch 1        |
|--------------------------|-------------------------|
| Port 7 (admin key = 100) | Port 1 (admin key = 50) |
| Port 8 (admin key = 100) | Port 2 (admin key = 50) |

In the configuration shown in [Table 12](#), Actor switch port 7 and port 8 aggregate to form an LACP trunk group with Partner switch port 1 and port 2.

LACP automatically determines which member links can be aggregated and then aggregates them. It provides for the controlled addition and removal of physical links for the link aggregation.

Each port on the switch can have one of the following LACP modes.

- off (default)  
The user can configure this port in to a regular static trunk group.
- active  
The port is capable of forming an LACP trunk. This port sends LACPDU packets to partner system ports.
- passive  
The port is capable of forming an LACP trunk. This port only responds to the LACPDU packets sent from an LACP *active* port.

Each active LACP port transmits LACP data units (LACPDUs), while each passive LACP port listens for LACPDUs. During LACP negotiation, the admin key is exchanged. The LACP trunk group is enabled as long as the information matches at both ends of the link. If the admin key value changes for a port at either end of the link, that port's association with the LACP trunk group is lost.

When the system is initialized, all ports by default are in LACP *off* mode and are assigned unique *admin keys*. To make a group of ports aggregatable, you assign them all the same *admin key*. You must set the port's LACP mode to *active* to activate LACP negotiation. You can set other port's LACP mode to *passive*, to reduce the amount of LACPDU traffic at the initial trunk-forming stage.

Use the following command to check whether the ports are trunked:

```
RS G7028 # show lacp information
```

**Note:** If you configure LACP on ports with 802.1X network access control, make sure the ports on both sides of the connection are properly configured for both LACP and 802.1X.

## LACP Minimum Links Option

For dynamic trunks that require a guaranteed amount of bandwidth to be considered useful, you can specify the minimum number of links for the trunk. If the specified minimum number of ports are not available, the trunk link will not be established. If an active LACP trunk loses one or more component links, the trunk will be placed in the down state if the number of links falls to less than the specified minimum. By default, the minimum number of links is 1, meaning that LACP trunks will remain operational as long as at least one link is available.

The LACP minimum links setting can be configured as follows:

- Via interface configuration mode:

```
RS G7028# interface port <port number or range>
RS G7028(config-if)# port-channel min-links <minimum links>
RS G7028(config-if)# exit
```

- Or via portchannel configuration mode:

```
RS G7028# interface portchannel lacp <LACP key>
RS G7028(config-PortChannel)# port-channel min-links <minimum links>
RS G7028(config-if)# exit
```

## Configuring LACP

Use the following procedure to configure LACP for port 7 and port 8 to participate in link aggregation.

1. Configure port parameters. All ports that participate in the LACP trunk group must have the same settings, including VLAN membership.
2. Select the port range and define the admin key. Only ports with the same admin key can form an LACP trunk group.

```
RS G7028(config)# interface port 7-8
RS G7028(config-if)# lacp key 100
```

3. Set the LACP mode.

```
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit
```

---

## Configurable Trunk Hash Algorithm

### Packet-Based Trunk Hashing

Traffic in a trunk group is statistically distributed among member ports using a *hash* process where various address and attribute bits from each transmitted frame are recombined to specify the particular trunk port the frame will use.

The switch can be configured to use a variety of hashing options. To achieve the most even traffic distribution, select options that exhibit a wide range of values for your particular network. Avoid hashing on information that is not usually present in the expected traffic, or which does not vary.

The G7028/G7052 supports the following hashing options, which can be used in any combination:

- Frame MAC and IP information. One of the following combinations is required:
  - Source MAC address (*s mac*)

```
RS G7028(config)# portchannel hash source-mac-address
```

- Destination MAC address (*d mac*)

```
RS G7028(config)# portchannel hash destination-mac-address
```

- Both source and destination MAC address

```
RS G7028(config)# portchannel hash source-destination-mac
```

- IPv4/IPv6 source IP address (*s ip*)

```
RS G7028(config)# portchannel hash source-ip-address
```

- IPv4/IPv6 destination IP address (*d ip*)

```
RS G7028(config)# portchannel hash destination-ip-address
```

- Both source and destination IPv4/IPv6 address (enabled by default)

```
RS G7028(config)# portchannel hash source-destination-ip
```

- Ingress port number (disabled by default)

```
RS G7028(config)# portchannel hash ingress
```

- Layer 4 port information (disabled by default)

```
RS G7028(config)# portchannel hash L4port
```

When enabled, Layer 4 port information (TCP, UDP, etc.) is added to the hash if available. The `L4port` option is ignored when Layer 4 information is not included in the packet (such as for Layer 2 packets).

## Port-Based Trunk Hashing

The standard hashing process that determines where traffic in a trunk group is distributed among member ports uses address and attribute bits from each transmitted frame to specify which trunk port the frame will use for egress.

With port-based trunk hashing, packet data is disregarded. Instead, you configure each ingress port to use a specific egress port. When you have many ingress ports with traffic to distribute over a smaller number of egress ports with equal bandwidth, such as 40 10Gb ingress ports and four 100Gb egress ports, this may enable you to make better use of those egress ports.

**Note:** Port-based trunk hashing is disabled by default and the standard hashing algorithm is used.

To enable port-based trunk hashing:

1. Enable or disable the feature on a per-trunk basis.

– For static trunks:

```
RS G7028(config)# [no] portchannel hash static-distribution in-port
portchannel <trunk ID>
```

– Or for dynamic trunks:

```
RS G7028(config)# [no] portchannel hash static-distribution in-port lacp-key
<admin key>
```

2. For each port, specify its member hash index.

```
RS G7028(config)# interface port <port, list, or range>
RS G7028(config-if)# portchannel distribution trunk-member-index <1-8>
RS G7028(config-if)# exit
```

The trunk-member index will only be used for traffic destined for a trunk where port-based trunk hashing has been enabled. Otherwise, the value is ignored.

When applicable, the index specifies which member link within the trunk will be used. Traffic from ports that have the same member index will occupy the same links when destined for the same trunk.

**Note:** The administrator must ensure that traffic is properly balanced among the individual links within the trunk.

At any time, if a port's member index is greater than the number of links available within a trunk where port-based hashing is enabled, a modulus approach is used. For instance, if there are four members in the trunk, traffic from a port with a trunk-member index of 1 will use the first link in the trunk, index 2 will use the second link, 3 will use the third, 4 will start again at the first link, 5 at the second, 6 at the third, and so on.

Each time a port is added to or removed from a trunk, the mapping between its member index and physical links is updated.

The trunk-member index can be configured on a per-port basis. The index assigned a given port is used with all trunks where port-based trunk hashing is enabled. The same port cannot use different trunk-member indexes for different trunks.



---

## Chapter 10. Spanning Tree Protocols

When multiple paths exist between two points on a network, Spanning Tree Protocol (STP), or one of its enhanced variants, can prevent broadcast loops and ensure that the RackSwitch G7028/G7052 (G7028/G7052) uses only the most efficient network path.

This chapter covers the following topics:

- [“Spanning Tree Protocol Modes” on page 128](#)
- [“Global STP Control” on page 129](#)
- [“PVRST Mode” on page 129](#)
- [“Rapid Spanning Tree Protocol” on page 141](#)
- [“Multiple Spanning Tree Protocol” on page 142](#)
- [“Port Type and Link Type” on page 146](#)

---

## Spanning Tree Protocol Modes

IBM Networking OS 7.6 supports the following STP modes:

- Rapid Spanning Tree Protocol (RSTP)

IEEE 802.1D (2004) RSTP allows devices to detect and eliminate logical loops in a bridged or switched network. When multiple paths exist, STP configures the network so that only the most efficient path is used. If that path fails, STP automatically configures the best alternative active path on the network to sustain network operations. RSTP is an enhanced version of IEEE 802.1D (1998) STP, providing more rapid convergence of the Spanning Tree network path states on STG 1.

See [“Rapid Spanning Tree Protocol” on page 141](#) for details.

- Per-VLAN Rapid Spanning Tree (PVRST)

PVRST mode is based on RSTP to provide rapid Spanning Tree convergence, but supports instances of Spanning Tree, allowing one STG per VLAN. PVRST mode is compatible with Cisco R-PVST/R-PVST+ mode.

PVRST is the default Spanning Tree mode on the G7028/G7052. See [“PVRST Mode” on page 129](#) for details.

- Multiple Spanning Tree Protocol (MSTP)

IEEE 802.1Q (2003) MSTP provides both rapid convergence and load balancing in a VLAN environment. MSTP allows multiple STGs, with multiple VLANs in each.

See [“Multiple Spanning Tree Protocol” on page 142](#) for details.

---

## Global STP Control

By default, the Spanning Tree feature is globally enabled on the switch, and is set for PVRST mode. Spanning Tree (and thus any currently configured STP mode) can be globally disabled using the following command:

```
RS G7028(config)# spanning-tree mode disable
```

Spanning Tree can be re-enabled by specifying the STP mode:

```
RS G7028(config)# spanning-tree mode {pvrst|rstp|mst}
```

where the command options represent the following modes:

- rstp: RSTP mode
- pvrst: PVRST mode
- mst: MSTP mode

---

## PVRST Mode

**Note:** Per-VLAN Rapid Spanning Tree (PVRST) is enabled by default on the G7028/G7052.

Using STP, network devices detect and eliminate logical loops in a bridged or switched network. When multiple paths exist, Spanning Tree configures the network so that a switch uses only the most efficient path. If that path fails, Spanning Tree automatically sets up another active path on the network to sustain network operations.

N/OS PVRST mode is based on IEEE 802.1w RSTP. Like RSTP, PVRST mode provides rapid Spanning Tree convergence. However, PVRST mode is enhanced for multiple instances of Spanning Tree. In PVRST mode, each VLAN may be automatically or manually assigned to one of 127 available STGs. If you need to assign a VLAN to STG 128, it must be done manually. Each STG acts as an independent, simultaneous instance of STP. PVRST uses IEEE 802.1Q tagging to differentiate STP BPDUs and is compatible with Cisco R-PVST/R-PVST+ modes.

The relationship between ports, trunk groups, VLANs, and Spanning Trees is shown in [Table 13](#).

*Table 13. Ports, Trunk Groups, and VLANs*

| Switch Element     | Belongs To                                                                                                                                              |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port               | Trunk group or one or more VLANs                                                                                                                        |
| Trunk group        | One or more VLANs                                                                                                                                       |
| VLAN (non-default) | <ul style="list-style-type: none"><li>• PVRST: One VLAN per STG</li><li>• RSTP: All VLANs are in STG 1</li><li>• MSTP: Multiple VLANs per STG</li></ul> |

## Port States

The port state controls the forwarding and learning processes of Spanning Tree. In PVST, the port state has been consolidated to the following: discarding, learning, and forwarding.

Due to the sequence involved in these STP states, considerable delays may occur while paths are being resolved. To mitigate delays, ports defined as *edge* ports ([“Port Type and Link Type” on page 146](#)) may bypass the discarding and learning states, and enter directly into the forwarding state.

## Bridge Protocol Data Units

### Bridge Protocol Data Units Overview

To create a Spanning Tree, the switch generates a configuration Bridge Protocol Data Unit (BPDU), which it then forwards out of its ports. All switches in the Layer 2 network participating in the Spanning Tree gather information about other switches in the network through an exchange of BPDUs.

A bridge sends BPDU packets at a configurable regular interval (2 seconds by default). The BPDU is used to establish a path, much like a hello packet in IP routing. BPDUs contain information about the transmitting bridge and its ports, including bridge MAC addresses, bridge priority, port priority, and path cost. If the ports are tagged, each port sends out a special BPDU containing the tagged information.

The generic action of a switch on receiving a BPDU is to compare the received BPDU to its own BPDU that it will transmit. If the received BPDU is better than its own BPDU, it will replace its BPDU with the received BPDU. Then, the switch adds its own bridge ID number and increments the path cost of the BPDU. The switch uses this information to block any necessary ports.

**Note:** If STP is globally disabled, BPDUs from external devices will transit the switch transparently. If STP is globally enabled, for ports where STP is turned off, inbound BPDUs will instead be discarded.

### Determining the Path for Forwarding BPDUs

When determining which port to use for forwarding and which port to block, the G7028/G7052 uses information in the BPDU, including each bridge ID. A technique based on the “lowest root cost” is then computed to determine the most efficient path for forwarding.

#### Bridge Priority

The bridge priority parameter controls which bridge on the network is the STG root bridge. To make one switch become the root bridge, configure the bridge priority lower than all other switches and bridges on your network. The lower the value, the higher the bridge priority. Use the following command to configure the bridge priority:

```
RS G7028(config)# spanning-tree stp <x> bridge priority <0-65535>
```

## Port Priority

The port priority helps determine which bridge port becomes the root port or the designated port. The case for the root port is when two switches are connected using a minimum of two links with the same path-cost. The case for the designated port is in a network topology that has multiple bridge ports with the same path-cost connected to a single segment, the port with the lowest port priority becomes the designated port for the segment. Use the following command to configure the port priority:

```
RS G7028(config-if)# spanning-tree stp <STG> priority <port priority>
```

where *priority value* is a number from 0 to 240, in increments of 16 (such as 0, 16, 32, and so on). If the specified priority value is not evenly divisible by 16, the value will be automatically rounded down to the nearest valid increment whenever manually changed in the configuration, or whenever a configuration file from a release prior to N/OS 6.5 is loaded.

## Port Path Cost

The port path cost assigns lower values to high-bandwidth ports, such as 10 Gigabit Ethernet, to encourage their use. The objective is to use the fastest links so that the route with the lowest cost is chosen. A value of 0 (the default) indicates that the default cost will be computed for an auto-negotiated link or trunk speed.

Use the following command to modify the port path cost:

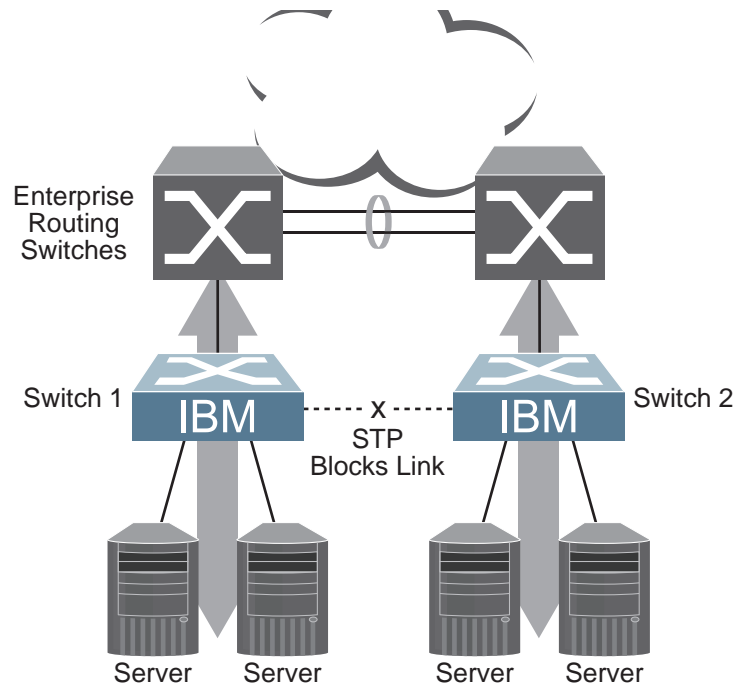
```
RS G7028(config)# interface port <port number>
RS G7028(config-if)# spanning-tree stp <STG> path-cost <path cost value>
RS G7028(config-if)# exit
```

The port path cost can be a value from 1 to 200000000. Specify 0 for automatic path cost.

## Simple STP Configuration

[Figure 10](#) depicts a simple topology using a switch-to-switch link between two G7028/G7052 1 and 2.

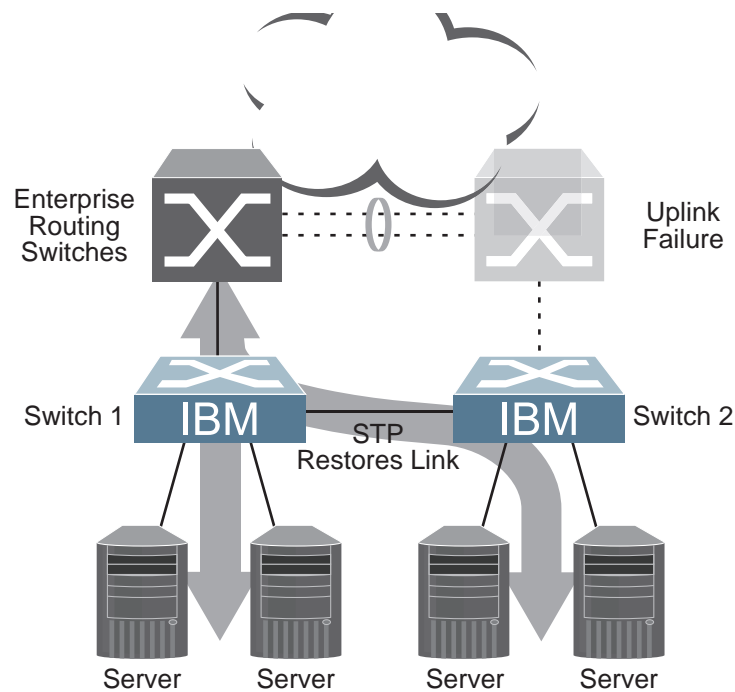
Figure 10. Spanning Tree Blocking a Switch-to-Switch Link



To prevent a network loop among the switches, STP must block one of the links between them. In this case, it is desired that STP block the link between the IBM switches, and not one of the G7028/G7052 uplinks or the Enterprise switch trunk.

During operation, if one G7028/G7052 experiences an uplink failure, STP will activate the IBM switch-to-switch link so that server traffic on the affected G7028/G7052 may pass through to the active uplink on the other G7028/G7052, as shown in [Figure 11](#).

Figure 11. Spanning Tree Restoring the Switch-to-Switch Link



In this example, port 10 on each G7028/G7052 is used for the switch-to-switch link. To ensure that the G7028/G7052 switch-to-switch link is blocked during normal operation, the port path cost is set to a higher value than other paths in the network. To configure the port path cost on the switch-to-switch links in this example, use the following commands on each G7028/G7052.

```
RS G7028(config)# interface port 10
RS G7028(config-if)# spanning-tree stp 1 path-cost 60000
RS G7028(config-if)# exit
```

## Per-VLAN Spanning Tree Groups

PVRST mode supports a maximum of 128 STGs, with each STG acting as an independent, simultaneous instance of STP.

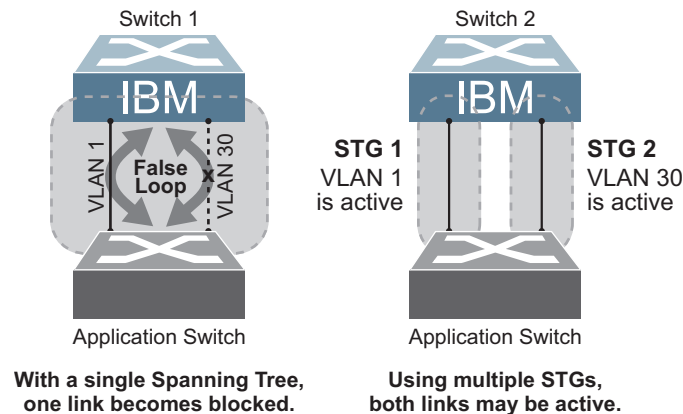
Multiple STGs provide multiple data paths which can be used for load-balancing and redundancy. To enable load balancing between two G7028/G7052s using multiple STGs, configure each path with a different VLAN and then assign each VLAN to a separate STG. Since each STG is independent, they each send their own IEEE 802.1Q tagged Bridge Protocol Data Units (BPDUs).

Each STG behaves as a bridge group and forms a loop-free topology. The default STG 1 may contain multiple VLANs (typically until they can be assigned to another STG). STGs 2-128 may contain only one VLAN each.

### Using Multiple STGs to Eliminate False Loops

Figure 12 shows a simple example of why multiple STGs are needed. In the figure, two ports on a G7028/G7052 are connected to two ports on an application switch. Each of the links is configured for a different VLAN, preventing a network loop. However, in the first network, since a single instance of Spanning Tree is running on all the ports of the G7028/G7052, a physical loop is assumed to exist, and one of the VLANs is blocked, impacting connectivity even though no actual loop exists.

Figure 12. Using Multiple Instances of Spanning Tree Group



In the second network, the problem of improper link blocking is resolved when the VLANs are placed into different Spanning Tree Groups (STGs). Since each STG has its own independent instance of Spanning Tree, each STG is responsible only for the loops within its own VLAN. This eliminates the false loop, and allows both VLANs to forward packets between the switches at the same time.

## VLANs and STG Assignment

In PVRST mode, up to 128 STGs are supported. Ports cannot be added directly to an STG. Instead, ports must be added as members of a VLAN, and the VLAN must then be assigned to the STG.

STG 1 is the default STG. Although VLANs can be added to or deleted from default STG 1, the STG itself cannot be deleted from the system. By default, STG 1 is enabled and includes VLAN 1, which by default includes all switch ports (except for management VLANs and management ports).

STG 128 is reserved for switch management. By default, STG 128 is disabled, but includes management VLAN 4095 and the management ports.

By default, all other STGs (STG 2 through 128) are enabled, though they initially include no member VLANs. VLANs must be assigned to STGs. By default, this is done automatically using VLAN Automatic STG Assignment (VASA), though it can also be done manually (see [“Manually Assigning STGs” on page 136](#)).

**Note:** VLAN assignment to STG 128 must be done manually.

When VASA is enabled (as by default), each time a new VLAN is configured, the switch will automatically assign that new VLAN to its own STG. Conversely, when a VLAN is deleted, if its STG is not associated with any other VLAN, the STG is returned to the available pool.

The specific STG number to which the VLAN is assigned is based on the VLAN number itself. For low VLAN numbers (1 through 127), the switch will attempt to assign the VLAN to its matching STG number. For higher numbered VLANs, the STG assignment is based on a simple modulus calculation; the attempted STG number will “wrap around,” starting back at the top of STG list each time the end of the list is reached. However, if the attempted STG is already in use, the switch will select the next available STG. If an empty STG is not available when creating a new VLAN, the VLAN is automatically assigned to default STG 1.

If ports are tagged, each tagged port sends out a special BPDU containing the tagged information. Also, when a tagged port belongs to more than one STG, the egress BPDUs are tagged to distinguish the BPDUs of one STG from those of another STG.

VASA is enabled by default, but can be disabled or re-enabled using the following commands:

```
RS G7028(config)# [no] spanning-tree stg-auto
```

If VASA is disabled, when you create a new VLAN, that VLAN automatically belongs to default STG 1. To place the VLAN in a different STG, assign it manually.

VASA applies only to PVRST mode and is ignored in RSTP and MSTP modes.

## Manually Assigning STGs

The administrator may manually assign VLANs to specific STGs, whether or not VASA is enabled.

1. If no VLANs exist (other than default VLAN 1), see [“Guidelines for Creating VLANs” on page 136](#) for information about creating VLANs and assigning ports to them.
2. Assign the VLAN to an STG using one of the following methods:
  - From the global configuration mode:

```
RS G7028(config)# spanning-tree stp <STG number> vlan <VLAN>
```

- Or from within the VLAN configuration mode:

```
RS G7028(config)# vlan <VLAN number>
RS G7028(config-vlan)# stg <STG number>
RS G7028(config-vlan)# exit
```

When a VLAN is assigned to a new STG, the VLAN is automatically removed from its prior STG.

**Note:** For proper operation with switches that use Cisco PVST+, it is recommended that you create a separate STG for each VLAN.

## Guidelines for Creating VLANs

Follow these guidelines when creating VLANs:

- When you create a new VLAN, if VASA is enabled (the default), that VLAN is automatically assigned its own STG. If VASA is disabled, the VLAN automatically belongs to STG 1, the default STG. To place the VLAN in a different STG, see [“Manually Assigning STGs” on page 136](#). The VLAN is automatically removed from its old STG before being placed into the new STG.
- Each VLANs must be contained *within* a single STG; a VLAN cannot span multiple STGs. By confining VLANs within a single STG, you avoid problems with Spanning Tree blocking ports and causing a loss of connectivity within the VLAN. When a VLAN spans multiple switches, it is recommended that the VLAN remain within the same STG (be assigned the same STG ID) across all the switches.
- If ports are tagged, all trunked ports can belong to multiple STGs.
- A port cannot be directly added to an STG. The port must first be added to a VLAN, and that VLAN added to the desired STG.

## Rules for VLAN Tagged Ports

The following rules apply to VLAN tagged ports:

- Tagged ports can belong to more than one STG, but untagged ports can belong to only one STG.
- When a tagged port belongs to more than one STG, the egress BPDUs are tagged to distinguish the BPDUs of one STG from those of another STG.

## Adding and Removing Ports from STGs

The following rules apply when you add ports to or remove ports from STGs:

- When you add a port to a VLAN that belongs to an STG, the port is also added to that STG. However, if the port you are adding is an untagged port and is already a member of another STG, that port will be removed from its current STG and added to the new STG. An untagged port cannot belong to more than one STG.

For example: Assume that VLAN 1 belongs to STG 1, and that port 1 is untagged and does not belong to any STG. When you add port 1 to VLAN 1, port 1 will automatically become part of STG 1.

However, if port 5 is untagged and is a member of VLAN 3 in STG 2, then adding port 5 to VLAN 1 in STG 1 will change the port PVID from 3 to 1:

"Port 5 is an UNTAGGED port and its PVID changed from 3 to 1.

- When you remove a port from a VLAN that belongs to an STG, that port will also be removed from the STG. However, if that port belongs to another VLAN in the same STG, the port remains in the STG.

As an example, assume that port 2 belongs to only VLAN 2, and that VLAN 2 belongs to STG 2. When you remove port 2 from VLAN 2, the port is moved to default VLAN 1 and is removed from STG 2.

However, if port 2 belongs to both VLAN 1 and VLAN 2, and both VLANs belong to STG 2, removing port 2 from VLAN 2 does not remove port 2 from STG 1, because the port is still a member of VLAN 1, which is still a member of STG 1.

- An STG cannot be deleted, only disabled. If you disable the STG while it still contains VLAN members, Spanning Tree will be off on all ports belonging to that VLAN.

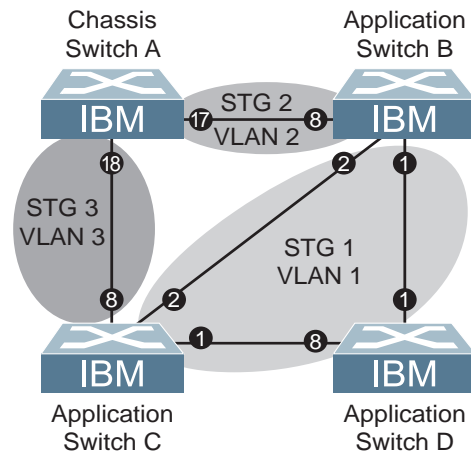
The relationship between port, trunk groups, VLANs, and Spanning Trees is shown in [Table 13 on page 129](#).

## The Switch-Centric Model

PVRST is switch-centric: STGs are enforced only on the switch where they are configured. PVRST allows only one VLAN per STG, except for the default STG 1 to which multiple VLANs can be assigned. The STG ID is not transmitted in the Spanning Tree BPDU. Each Spanning Tree decision is based entirely on the configuration of the particular switch.

For example, in [Figure 13](#), each switch is responsible for the proper configuration of its own ports, VLANs, and STGs. Switch A identifies its own port 17 as part of VLAN 2 on STG 2, and the Switch B identifies its own port 8 as part of VLAN 2 on STG 2.

Figure 13. Implementing Multiple Spanning Tree Groups



The VLAN participation for each Spanning Tree Group in [Figure 13 on page 138](#) is as follows:

- **VLAN 1 Participation**  
Assuming Switch B to be the root bridge, Switch B transmits the BPDU for STG 1 on ports 1 and 2. Switch C receives the BPDU on port 2, and Switch D receives the BPDU on port 1. Because there is a network loop between the switches in VLAN 1, either Switch D will block port 8 or Switch C will block port 1, depending on the information provided in the BPDU.
- **VLAN 2 Participation**  
Switch B, the root bridge, generates a BPDU for STG 2 from port 8. Switch A receives this BPDU on port 17, which is assigned to VLAN 2, STG 2. Because switch B has no additional ports participating in STG 2, this BPDU is not forwarded to any additional ports and Switch B remains the designated root.
- **VLAN 3 Participation**  
For VLAN 3, Switch A or Switch C may be the root bridge. If Switch A is the root bridge for VLAN 3, STG 3, then Switch A transmits the BPDU from port 18. Switch C receives this BPDU on port 8 and is identified as participating in VLAN 3, STG 3. Since Switch C has no additional ports participating in STG 3, this BPDU is not forwarded to any additional ports and Switch A remains the designated root.

## Configuring Multiple STGs

This configuration shows how to configure the three instances of STGs on the switches A, B, C, and D illustrated in [Figure 13 on page 138](#).

Because VASA is enabled by default, each new VLAN is automatically assigned its own STG.

1. Set the Spanning Tree mode on each switch to PVRST.

```
RS G7028(config)# spanning-tree mode pvrst
```

**Note:** PVRST is the default mode on the G7028/G7052. This step is not required unless the STP mode has been previously changed, and is shown here merely as an example of manual configuration.

2. Configure the following on Switch A:  
Enable VLAN 2 and VLAN 3.

```
RS G7028(config)# vlan 2
RS G7028(config-vlan)# exit
RS G7028(config)# vlan 3
RS G7028(config-vlan)# exit
```

If VASA is disabled, enter the following commands:

```
RS G7028(config)# spanning-tree stp 2 vlan 2
RS G7028(config)# spanning-tree stp 3 vlan 3
```

Add port 17 to VLAN 2, port 18 to VLAN 3.

```
RS G7028(config)# interface port 17
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 2
RS G7028(config-if)# exit

RS G7028(config)# interface port 18
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 3
RS G7028(config-if)# exit
```

VLAN 2 and VLAN 3 are removed from STG 1.

**Note:** In PVRST mode, each instance of STG is enabled by default.

3. Configure the following on Switch B:

Add port 8 to VLAN 2. Ports 1 and 2 are by default in VLAN 1 assigned to STG 1.

```
RS G7028(config)# vlan 2
RS G7028(config-vlan)# stg 2
RS G7028(config-vlan)# exit
RS G7028(config)# interface port 8
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 2
RS G7028(config-if)# exit
```

If VASA is disabled, enter the following command:  
RS G7028(config)# spanning-tree stp 2 vlan 2

VLAN 2 is automatically removed from STG 1. By default VLAN 1 remains in STG 1.

4. Configure the following on application switch C:

Add port 8 to VLAN 3. Ports 1 and 2 are by default in VLAN 1 assigned to STG 1.

```
RS G7028(config)# vlan 3
RS G7028(config-vlan)# stg 3
RS G7028(config-vlan)# exit
RS G7028(config)# interface port 8
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 3
RS G7028(config-if)# exit
```

If VASA is disabled, enter the following command:  
RS G7028(config)# spanning-tree stp 3 vlan 3

VLAN 3 is automatically removed from STG 1. By default VLAN 1 remains in STG 1.

5. Switch D does not require any special configuration for multiple Spanning Trees. Switch D uses default STG 1 only.

---

## Rapid Spanning Tree Protocol

RSTP provides rapid convergence of the Spanning Tree and provides the fast re-configuration critical for networks carrying delay-sensitive traffic such as voice and video. RSTP significantly reduces the time to reconfigure the active topology of the network when changes occur to the physical topology or its configuration parameters. RSTP reduces the bridged-LAN topology to a single Spanning Tree.

RSTP was originally defined in IEEE 802.1w (2001) and was later incorporated into IEEE 802.1D (2004), superseding the original STP standard.

RSTP parameters apply only to Spanning Tree Group (STG) 1. The PVRST mode STGs 2-128 are not used when the switch is placed in RSTP mode.

RSTP is compatible with devices that run IEEE 802.1D (1998) Spanning Tree Protocol. If the switch detects IEEE 802.1D (1998) BPDUs, it responds with IEEE 802.1D (1998)-compatible data units. RSTP is not compatible with Per-VLAN Rapid Spanning Tree (PVRST) protocol.

### Port States

RSTP port state controls are the same as for PVRST: discarding, learning, and forwarding.

Due to the sequence involved in these STP states, considerable delays may occur while paths are being resolved. To mitigate delays, ports defined as *edge* ports ([“Port Type and Link Type” on page 146](#)) may bypass the discarding and learning states, and enter directly into the forwarding state.

### RSTP Configuration Guidelines

This section provides important information about configuring RSTP. When RSTP is turned on, the following occurs:

- STP parameters apply only to STG 1.
- Only STG 1 is available. All other STGs are turned off.
- All VLANs are moved to STG 1.

### RSTP Configuration Example

This section provides steps to configure RSTP.

1. Configure port and VLAN membership on the switch.
2. Set the Spanning Tree mode to Rapid Spanning Tree.

```
RS G7028(config)# spanning-tree mode rstp
```

---

## Multiple Spanning Tree Protocol

Multiple Spanning Tree Protocol (MSTP) extends Rapid Spanning Tree Protocol (RSTP), allowing multiple Spanning Tree Groups (STGs) which may each include multiple VLANs. MSTP was originally defined in IEEE 802.1s (2002) and was later included in IEEE 802.1Q (2003).

In MSTP mode, the G7028/G7052 supports up to 32 instances of Spanning Tree, corresponding to STGs 1-32, with each STG acting as an independent, simultaneous instance of RSTP.

MSTP allows frames assigned to different VLANs to follow separate paths, with each path based on an independent Spanning Tree instance. This approach provides multiple forwarding paths for data traffic, thereby enabling load-balancing, and reducing the number of Spanning Tree instances required to support a large number of VLANs.

Due to Spanning Tree's sequence of discarding, learning, and forwarding, lengthy delays may occur while paths are being resolved. Ports defined as *edge* ports (["Port Type and Link Type" on page 146](#)) bypass the Discarding and Learning states, and enter directly into the Forwarding state.

**Note:** In MSTP mode, Spanning Tree for the management ports is turned off by default.

## MSTP Region

A group of interconnected bridges that share the same attributes is called an MST region. Each bridge within the region must share the following attributes:

- Alphanumeric name
- Revision number
- VLAN-to STG mapping scheme

MSTP provides rapid re-configuration, scalability and control due to the support of regions, and multiple Spanning-Tree instances support within each region.

## Common Internal Spanning Tree

The Common Internal Spanning Tree (CIST) or MST0 provides a common form of Spanning Tree Protocol, with one Spanning-Tree instance that can be used throughout the MSTP region. CIST allows the switch to interoperate with legacy equipment, including devices that run IEEE 802.1D (1998) STP.

CIST allows the MSTP region to act as a virtual bridge to other bridges outside of the region, and provides a single Spanning-Tree instance to interact with them.

CIST port configuration includes Hello time, path-cost, and interface priority. These parameters do not affect Spanning Tree Groups 1-32. They apply only when the CIST is used.

## MSTP Configuration Guidelines

This section provides important information about configuring Multiple Spanning Tree Groups:

- When MSTP is turned on, the switch automatically moves all VLANs to the CIST. When MSTP is turned off, the switch moves all VLANs from the CIST to STG 1.
- When you enable MSTP, you must configure the Region Name. A default revision number of 1 is configured automatically.
- Each bridge in the region must have the same name, version number, and VLAN mapping.

## MSTP Configuration Examples

### Example 1

This section provides steps to configure MSTP on the G7028/G7052.

1. Configure port and VLAN membership on the switch.
2. Configure MSTP region parameters, and set the mode to Multiple Spanning Tree.

```
RS G7028(config)# spanning-tree mst name <name>
RS G7028(config)# spanning-tree mode mst
```

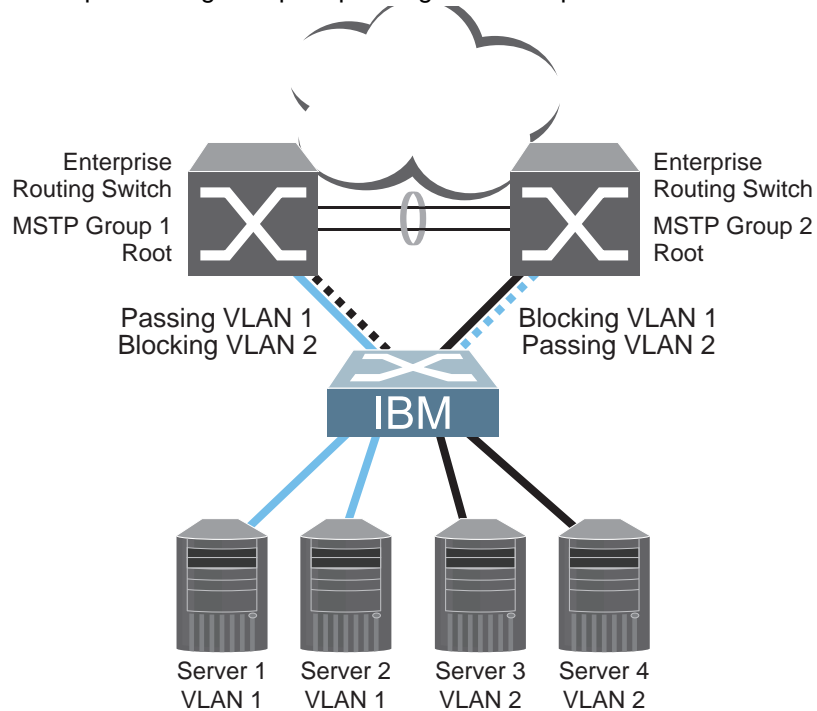
3. Assign VLANs to Spanning Tree Groups.

```
RS G7028(config)# spanning-tree mst <instance ID> vlan <VLAN ID>
```

## Example 2

This configuration shows how to configure MSTP Groups on the switch, as shown in Figure 13.

Figure 14. Implementing Multiple Spanning Tree Groups



This example shows how multiple Spanning Trees can provide redundancy without wasting any uplink ports. In this example, the server ports are split between two separate VLANs. Both VLANs belong to two different MSTP groups. The Spanning Tree *priority* values are configured so that each routing switch is the root for a different MSTP instance. All of the uplinks are active, with each uplink port backing up the other.

1. Configure port membership and define the STGs for VLAN 1. Enable tagging on uplink ports that share VLANs. Port 19 and port 20 connect to the Enterprise Routing switches.

```
RS G7028(config)# interface port 19
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# exit
RS G7028(config)# interface port 20
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# exit
```

2. Add server ports 1 and 2 to VLAN 1. Add uplink ports 19 and port 20 to VLAN 1.

```
RS G7028(config)# spanning-tree mst 1 vlan 1
RS G7028(config)# spanning-tree mst 1 enable
RS G7028(config)# interface port 1,2,19,20
RS G7028(config-if)# switchport trunk allowed vlan add 1
RS G7028(config-if)# exit
```

3. Configure MSTP: Spanning Tree mode, region name, and revision.

```
RS G7028(config)# spanning-tree mst name MyRegion
RS G7028(config)# spanning-tree mode mst
RS G7028(config)# spanning-tree mst revision 100
```

4. Configure port membership and define the STGs for VLAN 2. Add server ports 3, 4, and 5 to VLAN 2. Add uplink ports 19 and 20 to VLAN 2. Assign VLAN 2 to STG 2.

```
RS G7028(config)# spanning-tree mst 2 vlan 2
RS G7028(config)# interface port 3,4,5,19,20
RS G7028(config-if)# switchport trunk allowed vlan add 2
RS G7028(config-if)# exit
```

**Note:** Each STG is enabled by default.

---

## Port Type and Link Type

### Edge Port/Portfast

A port that does not connect to a bridge is called an *edge/portfast port*. Since edge/portfast ports are assumed to be connected to non-STP devices (such as directly to hosts or servers), they are placed in the forwarding state as soon as the link is up.

Edge ports send BPDUs to upstream STP devices like normal STP ports, but do not receive BPDUs. If a port with `edge` enabled does receive a BPDU, it immediately begins working as a normal (non-edge) port, and participates fully in Spanning Tree.

Use the following commands to define or clear a port as an edge/portfast port:

```
RS G7028(config)# interface port <port>
RS G7028(config-if)# [no] spanning-tree portfast
RS G7028(config-if)# exit
```

### Link Type

The link type determines how the port behaves in regard to Rapid Spanning Tree. Use the following commands to define the link type for the port:

```
RS G7028(config)# interface port <port>
RS G7028(config-if)# [no] spanning-tree link-type <type>
RS G7028(config-if)# exit
```

where *type* corresponds to the duplex mode of the port, as follows:

- `p2p`            A full-duplex link to another device (point-to-point)
- `shared`        A half-duplex link is a shared segment and can contain more than one device.
- `auto`           The switch dynamically configures the link type.

**Note:** Any STP port in full-duplex mode can be manually configured as a shared port when connected to a non-STP-aware shared device (such as a typical Layer 2 switch) used to interconnect multiple STP-aware devices.

---

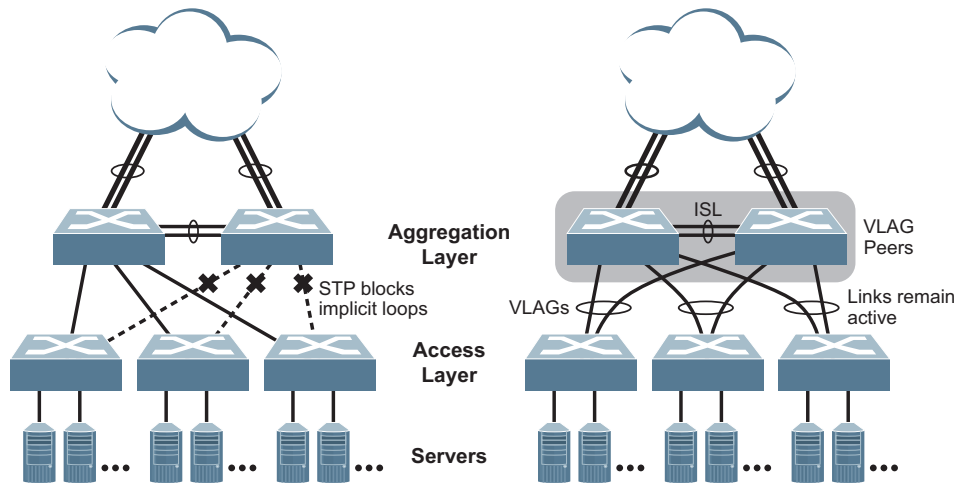
## Chapter 11. Virtual Link Aggregation Groups

---

### VLAG Overview

In many data center environments, downstream servers or switches connect to upstream devices which consolidate traffic. For example, see [Figure 15](#).

Figure 15. Typical Data Center Switching Layers with STP vs. VLAG



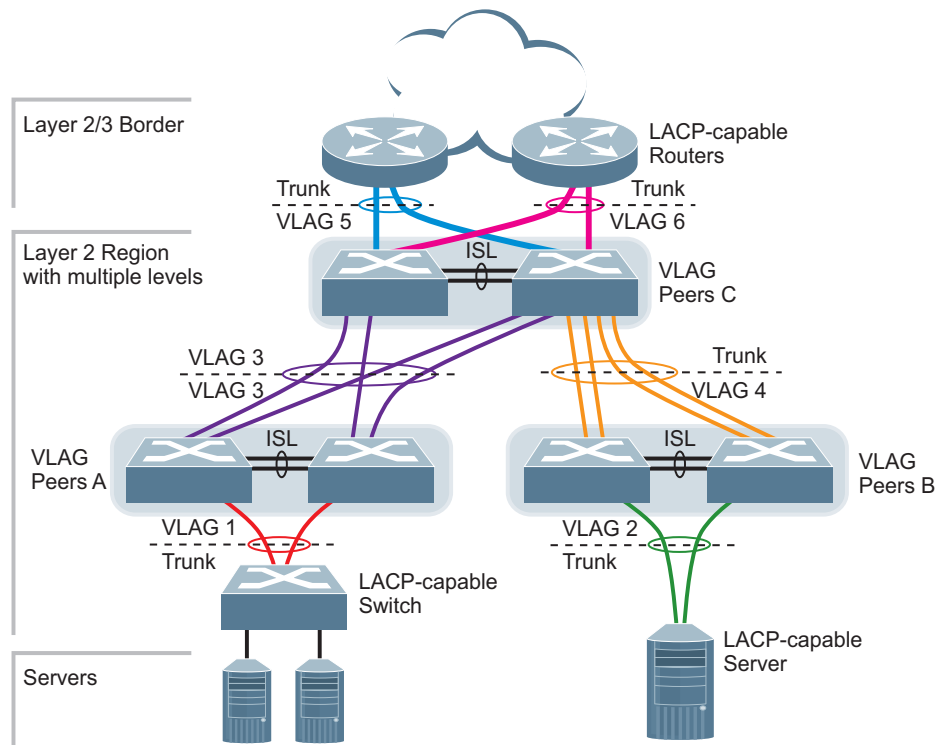
As shown in the example, a switch in the access layer may be connected to more than one switch in the aggregation layer to provide for network redundancy. Typically, Spanning Tree Protocol (RSTP, PVRST, or MSTP—see [“Spanning Tree Protocols” on page 127](#)) is used to prevent broadcast loops, blocking redundant uplink paths. This has the unwanted consequence of reducing the available bandwidth between the layers by as much as 50%. In addition, STP may be slow to resolve topology changes that occur during a link failure, and can result in considerable MAC address flooding.

Using Virtual Link Aggregation Groups (VLGs), the redundant uplinks remain active, utilizing all available bandwidth.

Two switches are paired into VLAG peers, and act as a single virtual entity for the purpose of establishing a multi-port trunk. Ports from both peers can be grouped into a VLAG and connected to the same LAG-capable target device. From the perspective of the target device, the ports connected to the VLAG peers appear to be a single trunk connecting to a single logical device. The target device uses the configured Tier ID to identify the VLAG peers as this single logical device. It is important that you use a unique Tier ID for each VLAG pair you configure. The VLAG-capable switches synchronize their logical view of the access layer port structure and internally prevent implicit loops. The VLAG topology also responds more quickly to link failure and does not result in unnecessary MAC flooding.

VLGs are also useful in multi-layer environments for both uplink and downlink redundancy to any regular LAG-capable device. For example:

Figure 16. VLAG Application with Multiple Layers

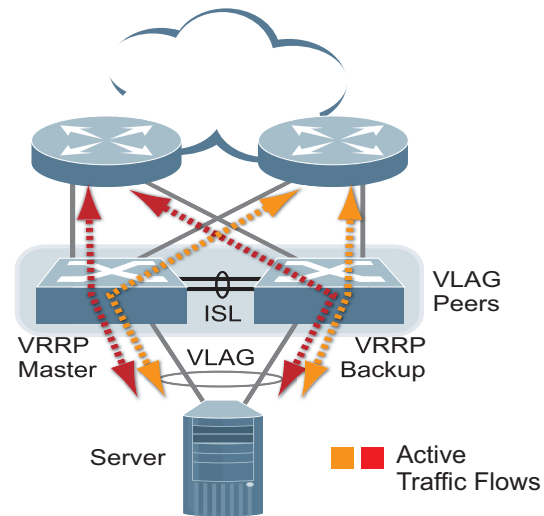


Wherever ports from *both* peered switches are trunked to another device, the trunked ports must be configured as a VLAG. For example, VLAGs 1 and 3 must be configured for both VLAG Peer A switches. VLAGs 2 and 4 must be configured for both VLAG Peer B switches. VLAGs 3, 5, and 6 must be configured on both VLAG Peer C switches. Other devices connecting to the VLAG peers are configured using regular static or dynamic trunks.

**Note:** Do not configure a VLAG for connecting only one switch in the peer set to another device or peer set. For instance, in VLAG Peer C, a regular trunk is employed for the downlink connection to VLAG Peer B because only one of the VLAG Peer C switches is involved.

In addition, when used with VRRP, VLAGs can provide seamless active-active failover for network links. For example

Figure 17. VLAG Application with VRRP:



---

## VLAG Capacities

Servers or switches that connect to the VLAG peers using a multi-port VLAG are considered VLAG clients. VLAG clients are not required to be VLAG-capable. The ports participating in the VLAG are configured as regular port trunks on the VLAG client end.

On the VLAG peers, the VLAGs are configured similarly to regular port trunks, using many of the same features and rules. See [“Ports and Trunking” on page 115](#) for general information concerning all port trunks.

Each VLAG begins as a regular port trunk on each VLAG-peer switch. The VLAG may be either a static trunk group (portchannel) or dynamic LACP trunk group, and consumes one slot from the overall port trunk capacity pool. The trunk type must match that used on VLAG client devices. Additional configuration is then required to implement the VLAG on both VLAG peer switches.

You may configure up to 64 trunk groups on the switch, with all types (regular or VLAG, static or LACP) sharing the same pool.

Each trunk type can contain up to 16 member ports, depending on the port type and availability.

---

## VLAGs versus Port Trunks

Though similar to regular port trunks in many regards, VLAGs differ from regular port trunks in a number of important ways:

- A VLAG can consist of multiple ports on two VLAG peers, which are connected to one logical client device such as a server, switch, or another VLAG device.
- The participating ports on the client device are configured as a regular port trunk.
- The VLAG peers must be the same model, and run the same software version.
- VLAG peers require a dedicated inter-switch link (ISL) for synchronization. The ports used to create the ISL must have the following properties:
  - ISL ports must have VLAN tagging turned on.
  - ISL ports must be configured for all VLAG VLANs.
  - ISL ports must be placed into a regular port trunk group (dynamic or static).
  - A minimum of two ports on each switch are recommended for ISL use.
- Dynamic routing protocols, such as OSPF, cannot terminate on VLAGs.
- Routing over VLAGs is not supported. However, IP forwarding between subnets served by VLAGs can be accomplished using VRRP.
- VLAGs are configured using additional commands.
- It is recommended that end-devices connected to VLAG switches use NICs with dual-homing. This increases traffic efficiency, reduces ISL load, and provides faster link failover.

---

## Configuring VLAGs

When configuring VLAG or making changes to your VLAG configuration, consider the following VLAG behavior:

- When adding a static Mrouter on VLAG links, ensure that you also add it on the ISL link to avoid VLAG link failure. If the VLAG link fails, traffic cannot be recovered through the ISL.
- When you enable VLAG on the switch, if a MSTP region mismatch is detected with the VLAG peer, the ISL will shut down. In such a scenario, correct the region on the VLAG peer and manually enable the ISL.
- If you have enabled VLAG on the switch, and you need to change the STP mode, ensure that you first disable VLAG and then change the STP mode.
- When VLAG is enabled, you may see two root ports on the secondary VLAG switch. One of these will be the actual root port for the secondary VLAG switch and the other will be a root port synced with the primary VLAG switch.
- The LACP key used must be unique for each VLAG in the entire topology.

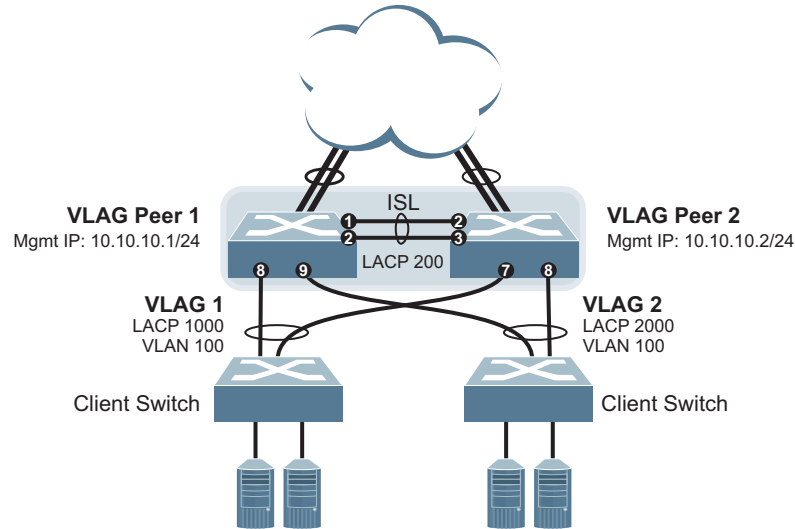
The following parameters must be identically configured on the VLAG ports of both the VLAG peers:

- VLANs
- Native VLAN tagging
- STP mode
- BPDU Guard setting
- STP port setting
- MAC aging timers
- Static MAC entries
- ACL configuration parameters
- QoS configuration parameters

## Basic VLAG Configuration

Figure 18 shows an example configuration where two VLAG peers are used for aggregating traffic from downstream devices.

Figure 18. Basic VLAGs



In this example, each client switch is connected to both VLAG peers. On each client switch, the ports connecting to the VLAG peers are configured as a dynamic LACP port trunk. The VLAG peer switches share a dedicated ISL for synchronizing VLAG information. On the individual VLAG peers, each port leading to a specific client switch (and part of the client switch's port trunk) is configured as a VLAG.

In the following example configuration, only the configuration for VLAG 1 on VLAG Peer 1 is shown. VLAG Peer 2 and all other VLAGs are configured in a similar fashion.

### Configure the ISL

The ISL connecting the VLAG peers is shared by all their VLAGs. The ISL needs to be configured only once on each VLAG peer.

1. If STP is desired on the switch, use PVRST or MSTP mode only:

```
RS G7028(config)# spanning-tree mode pvrst
```

2. Configure the ISL ports and place them into a port trunk group:

```
RS G7028(config)# interface port 1-2
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# lacp key 200
RS G7028(config-if)# exit
```

**Note:** In this case, a dynamic trunk group is shown. A static trunk (portchannel) could be configured instead.

3. Configure VLAG Tier ID. This is used to identify the VLAG switch in a multi-tier environment.

```
RS G7028(config)# vlag tier-id 10
```

4. Configure the ISL for the VLAG peer.

Make sure you configure the VLAG peer (VLAG Peer 2) using the same ISL trunk type (dynamic or static), the same VLAN, and the same STP mode and tier ID used on VLAG Peer 1.

## Configure the VLAG

1. Configure the VLAN for VLAG 1. Make sure members include the ISL and VLAG 1 ports.

```
RS G7028(config)# interface port 1-2,8
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 100
RS G7028(config-if)# exit
```

2. Place the VLAG 1 port(s) in a port trunk group:

```
RS G7028(config)# interface port 8
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# lacp key 1000
RS G7028(config-if)# exit
```

3. Assign the trunk to the VLAG:

```
RS G7028(config)# vlag adminkey 1000 enable
```

4. Enable VLAG globally.

```
RS G7028(config)# vlag enable
```

5. Continue by configuring all required VLAGs on VLAG Peer 1, and then repeat the configuration for VLAG Peer 2.

For each corresponding VLAG on the peer, the port trunk type (dynamic or static), VLAN, and STP mode and ID must be the same as on VLAG Peer 1.

6. Verify the completed configuration:

```
# show vlag
```

## Configuring Health Check

We strongly recommend that you configure the G7028/G7052 to check the health status of its VLAG peer. Although the operational status of the VLAG peer is generally determined via the ISL connection, configuring a network health check provides an alternate means to check peer status in case the ISL links fail. Use an independent link between the VLAG switches to configure health check.

**Note:** Configuring health check on an ISL VLAN interface or on a VLAG data port may impact the accuracy of the health check status.

1. Configure a management interface for the switch.

**Note:** If the switch does not have a dedicated management interface, configure a VLAN for the health check interface:

```
RS G7028(config)# interface ip 1
RS G7028(config-ip-if)# ip address 10.10.10.1 255.255.255.0
RS G7028(config-ip-if)# enable
RS G7028(config-ip-if)# exit
```

**Note:** Configure a similar interface on VLAG Peer 2. For example, use IP address 10.10.10.2.

2. Specify the IP address of the VLAG Peer:

```
RS G7028(config)# vlag hlthchk peer-ip 10.10.10.2
```

**Note:** For VLAG Peer 2, the management interface would be configured as 10.10.10.2, and the health check would be configured for 10.10.10.1, pointing back to VLAG Peer 1.

## Configuring VLAGs in Multiple Layers

Figure 19. VLAG in Multiple Layers

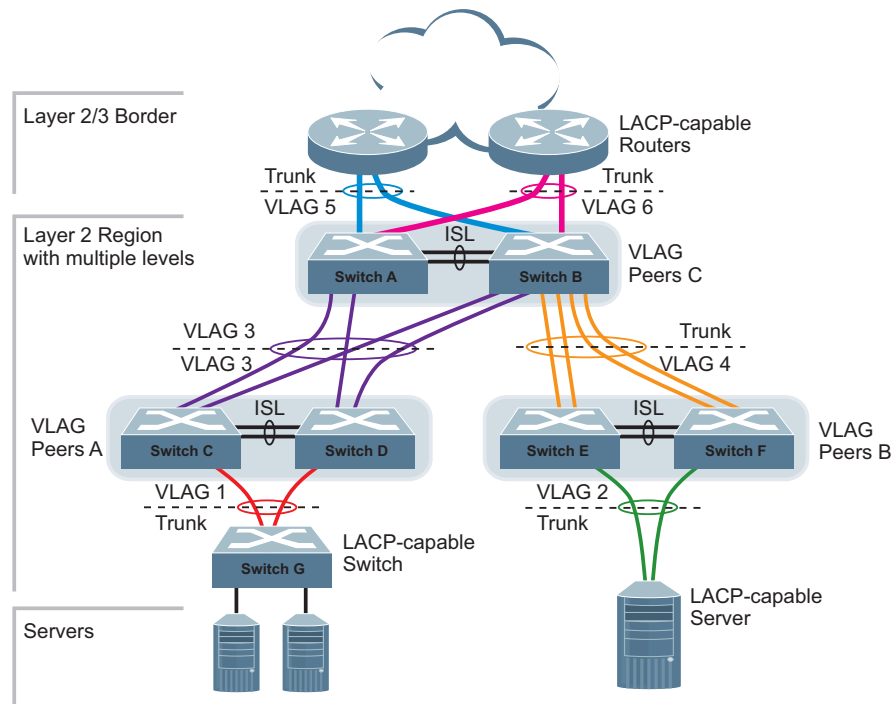


Figure 19 shows an example of VLAG being used in a multi-layer environment. Following are the configuration steps for the topology.

### Task 1: Configure Layer 2/3 border switches.

Configure ports on border switch as follows:

```
RS G7028(config)# interface port 1,2
RS G7028(config-if)# lacp key 100
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit
```

Repeat the above steps for the second border switch.

### Task 2: Configure switches in the Layer 2 region.

Consider the following:

- ISL ports on switches A and B - ports 1, 2
- Ports connecting to Layer 2/3 - ports 5, 6
- Ports on switches A and B connecting to switches C and D: ports 10, 11
- Ports on switch B connecting to switch E: ports 15, 16
- Ports on switch B connecting to switch F: ports 17, 18

1. Configure VLAG tier ID and enable VLAG globally.

```
RS G7028(config)# vlag tier-id 10
RS G7028(config)# vlag enable
```

2. Configure ISL ports on Switch A.

```
RS G7028(config)# interface port 1,2
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# lacp key 200
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit

RS G7028(config)# vlag isl adminkey 200
RS G7028(config-vlan)# exit
```

3. Configure port on Switch A connecting to Layer 2/3 router 1.

```
RS G7028(config)# interface port 1,2,5
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 10
RS G7028(config-if)# exit

RS G7028(config)# interface port 5
RS G7028(config-if)# lacp key 400
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit

RS G7028(config)# vlag adminkey 400 enable
```

Repeat the above steps on Switch B for ports connecting to Layer 2/3 router 1.

4. Configure port on Switch A connecting to Layer 2/3 router 2.

```
RS G7028(config)# interface port 1,2,6
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 20
RS G7028(config-if)# exit

RS G7028(config)# interface port 6
RS G7028(config-if)# lacp key 500
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit

RS G7028(config)# vlag adminkey 500 enable
```

Repeat the above steps on Switch B for ports connecting to Layer 2/3 router 2.

5. Configure ports on Switch A connecting to downstream VLAG switches C and D.

```
RS G7028(config)# interface port 10,11
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 20
RS G7028(config-if)# lacp key 600
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit

RS G7028(config)# vlag adminkey 600 enable
```

Repeat the above steps on Switch B for ports connecting to downstream VLAG switch C and D.

6. Configure ports on Switch B connecting to downstream switches E and F.

```
RS G7028(config)# interface port 15-18
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan add 30
RS G7028(config-if)# lacp key 700
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit
```

7. Configure ISL between switches C and D, and between E and F as shown in Step 1.
8. Configure the Switch G as shown in Step 2.



---

## Chapter 12. Quality of Service

Quality of Service features allow you to allocate network resources to mission-critical applications at the expense of applications that are less sensitive to such factors as time delays or network congestion. You can configure your network to prioritize specific types of traffic, ensuring that each type receives the appropriate Quality of Service (QoS) level.

The following topics are discussed in this section:

- [“QoS Overview” on page 160](#)
- [“Using ACL Filters” on page 161](#)
- [“Using DSCP Values to Provide QoS” on page 163](#)
- [“Using 802.1p Priority to Provide QoS” on page 169](#)
- [“Queuing and Scheduling” on page 170](#)

---

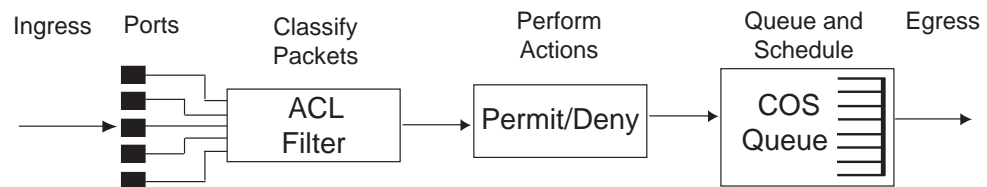
## QoS Overview

QoS helps you allocate guaranteed bandwidth to the critical applications, and limit bandwidth for less critical applications. Applications such as video and voice must have a certain amount of bandwidth to work correctly; using QoS, you can provide that bandwidth when necessary. Also, you can put a high priority on applications that are sensitive to timing out or that cannot tolerate delay, by assigning their traffic to a high-priority queue.

By assigning QoS levels to traffic flows on your network, you can ensure that network resources are allocated where they are needed most. QoS features allow you to prioritize network traffic, thereby providing better service for selected applications.

Figure 20 shows the basic QoS model used by the switch.

Figure 20. QoS Model



The basic QoS model works as follows:

- Classify traffic:
  - Read DSCP value.
  - Read 802.1p priority value.
  - Match ACL filter parameters.
- Perform actions:
  - Define bandwidth and burst parameters
  - Select actions to perform on in-profile and out-of-profile traffic
  - Deny packets
  - Permit packets
  - Mark DSCP or 802.1p Priority
  - Set COS queue (with or without re-marking)
- Queue and schedule traffic:
  - Place packets in one of the COS queues.
  - Schedule transmission based on the COS queue.

---

## Using ACL Filters

Access Control Lists (ACLs) are filters that allow you to classify and segment traffic, so you can provide different levels of service to different traffic types. Each filter defines the conditions that must match for inclusion in the filter, and also the actions that are performed when a match is made.

IBM Networking OS 7.6 supports up to 512 ACLs.

The G7028/G7052 allows you to classify packets based on various parameters. For example:

- Ethernet: source MAC, destination MAC, VLAN number/mask, Ethernet type, priority.
- IPv4: Source IP address/mask, destination address/mask, type of service, IP protocol number.
- TCP/UDP: Source port, destination port, TCP flag.
- Packet format

For ACL details, see [“Access Control Lists” on page 83](#).

## Summary of ACL Actions

Actions determine how the traffic is treated. The G7028/G7052 QoS actions include the following:

- Pass or Drop
- Re-mark a new DiffServ Code Point (DSCP)
- Re-mark the 802.1p field
- Set the COS queue

## ACL Metering and Re-Marking

You can define a profile for the aggregate traffic flowing through the G7028/G7052 by configuring a QoS meter (if desired) and assigning ACLs to ports. When you add ACLs to a port, make sure they are ordered correctly in terms of precedence.

Actions taken by an ACL are called *In-Profile* actions. You can configure additional In-Profile and Out-of-Profile actions on a port. Data traffic can be metered, and re-marked to ensure that the traffic flow provides certain levels of service in terms of bandwidth for different types of network traffic.

### Metering

QoS metering provides different levels of service to data streams through user-configurable parameters. A meter is used to measure the traffic stream against a traffic profile, which you create. Thus, creating meters yields In-Profile and Out-of-Profile traffic for each ACL, as follows:

- **In-Profile**—If there is no meter configured or if the packet conforms to the meter, the packet is classified as In-Profile.
- **Out-of-Profile**—If a meter is configured and the packet does not conform to the meter (exceeds the committed rate or maximum burst rate of the meter), the packet is classified as Out-of-Profile.

Using meters, you set a Committed Rate in Kbps (multiples of 64 Mbps). All traffic within this Committed Rate is In-Profile. Additionally, you set a Maximum Burst Size that specifies an allowed data burst larger than the Committed Rate for a brief period. These parameters define the In-Profile traffic.

Meters keep the sorted packets within certain parameters. You can configure a meter on an ACL, and perform actions on metered traffic, such as packet re-marking.

### Re-Marking

Re-marking allows for the treatment of packets to be reset based on new network specifications or desired levels of service. You can configure the ACL to re-mark a packet as follows:

- Change the DSCP value of a packet, used to specify the service level traffic receives.
- Change the 802.1p priority of a packet.

---

## Using DSCP Values to Provide QoS

The switch uses the Differentiated Services (DiffServ) architecture to provide QoS functions. DiffServ is described in IETF RFCs 2474 and 2475.

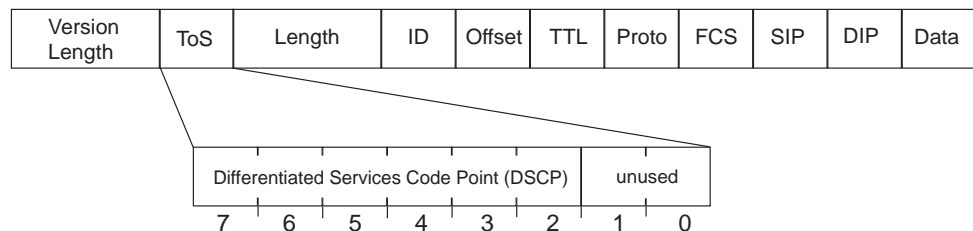
The six most significant bits in the TOS byte of the IP header are defined as DiffServ Code Points (DSCP). Packets are marked with a certain value depending on the type of treatment the packet must receive in the network device. DSCP is a measure of the Quality of Service (QoS) level of the packet.

The switch can classify traffic by reading the DiffServ Code Point (DSCP) or IEEE 802.1p priority value, or by using filters to match specific criteria. When network traffic attributes match those specified in a traffic pattern, the policy instructs the switch to perform specified actions on each packet that passes through it. The packets are assigned to different Class of Service (COS) queues and scheduled for transmission.

## Differentiated Services Concepts

To differentiate between traffic flows, packets can be classified by their DSCP value. The Differentiated Services (DS) field in the IP header is an octet, and the first six bits, called the DS Code Point (DSCP), can provide QoS functions. Each packet carries its own QoS state in the DSCP. There are 64 possible DSCP values (0-63).

Figure 21. Layer 3 IPv4 packet



The switch can perform the following actions to the DSCP:

- Read the DSCP value of ingress packets.
- Re-mark the DSCP value to a new value
- Map the DSCP value to a Class of Service queue (COSq).

The switch can use the DSCP value to direct traffic prioritization.

With DiffServ, you can establish policies to direct traffic. A policy is a traffic-controlling mechanism that monitors the characteristics of the traffic, (for example, its source, destination, and protocol) and performs a controlling action on the traffic when certain characteristics are matched.

## Trusted/Untrusted Ports

By default, all ports on the G7028/G7052 are trusted. To configure untrusted ports, re-mark the DSCP value of the incoming packet to a lower DSCP value using the following commands:

```
RS G7028(config)# interface port 1
RS G7028(config-if)# dscp-marking
RS G7028(config-if)# exit
RS G7028(config)# qos dscp dscp-mapping <DSCP value (0-63)> <new value>
RS G7028(config)# qos dscp re-marking
```

## Per Hop Behavior

The DSCP value determines the Per Hop Behavior (PHB) of each packet. The PHB is the forwarding treatment given to packets at each hop. QoS policies are built by applying a set of rules to packets, based on the DSCP value, as they hop through the network.

The default settings are based on the following standard PHBs, as defined in the IEEE standards:

- Expedited Forwarding (EF)—This PHB has the highest egress priority and lowest drop precedence level. EF traffic is forwarded ahead of all other traffic. EF PHB is described in RFC 2598.
- Assured Forwarding (AF)—This PHB contains four service levels, each with a different drop precedence, as shown in the following table. Routers use drop precedence to determine which packets to discard last when the network becomes congested. AF PHB is described in RFC 2597.

| Drop Precedence | Class 1        | Class 2        | Class 3        | Class 4        |
|-----------------|----------------|----------------|----------------|----------------|
| Low             | AF11 (DSCP 10) | AF21 (DSCP 18) | AF31 (DSCP 26) | AF41 (DSCP 34) |
| Medium          | AF12 (DSCP 12) | AF22 (DSCP 20) | AF32 (DSCP 28) | AF42 (DSCP 36) |
| High            | AF13 (DSCP 14) | AF23 (DSCP 22) | AF33 (DSCP 30) | AF43 (DSCP 38) |

- Class Selector (CS)—This PHB has eight priority classes, with CS7 representing the highest priority, and CS0 representing the lowest priority, as shown in the following table. CS PHB is described in RFC 2474.

| Priority | Class Selector | DSCP |
|----------|----------------|------|
| Highest  | CS7            | 56   |
|          | CS6            | 48   |
|          | CS5            | 40   |
|          | CS4            | 32   |
|          | CS3            | 24   |
|          | CS2            | 16   |
|          | CS1            | 8    |
| Lowest   | CS0            | 0    |

## QoS Levels

[Table 14](#) shows the default service levels provided by the switch, listed from highest to lowest importance:

*Table 14. Default QoS Service Levels*

| Service Level   | Default PHB           | 802.1p Priority |
|-----------------|-----------------------|-----------------|
| Critical        | CS7                   | 7               |
| Network Control | CS6                   | 6               |
| Premium         | EF, CS5               | 5               |
| Platinum        | AF41, AF42, AF43, CS4 | 4               |
| Gold            | AF31, AF32, AF33, CS3 | 3               |
| Silver          | AF21, AF22, AF23, CS2 | 2               |
| Bronze          | AF11, AF12, AF13, CS1 | 1               |
| Standard        | DF, CS0               | 0               |

## DSCP Re-Marking and Mapping

The switch can use the DSCP value of ingress packets to re-mark the DSCP to a new value, and to set an 802.1p priority value. Use the following command to view the default settings.

```
RS G7028# show qos dscp
Current DSCP Remarking Configuration: OFF
```

| DSCP | New DSCP | New 802.1p Prio |
|------|----------|-----------------|
| 0    | 0        | 0               |
| 1    | 1        | 0               |
| 2    | 2        | 0               |
| 3    | 3        | 0               |
| 4    | 4        | 0               |
| 5    | 5        | 0               |
| 6    | 6        | 0               |
| 7    | 7        | 0               |
| 8    | 8        | 1               |
| 9    | 9        | 0               |
| 10   | 10       | 1               |
| ...  |          |                 |
| 54   | 54       | 0               |
| 55   | 55       | 0               |
| 56   | 56       | 7               |
| 57   | 57       | 0               |
| 58   | 58       | 0               |
| 59   | 59       | 0               |
| 60   | 60       | 0               |
| 61   | 61       | 0               |
| 62   | 62       | 0               |
| 63   | 63       | 0               |

Use the following command to turn on DSCP re-marking globally:

```
RS G7028(config)# qos dscp re-marking
```

Then you must enable DSCP re-marking on any port that you wish to perform this function (Interface Port mode).

**Note:** If an ACL meter is configured for DSCP re-marking, the meter function takes precedence over QoS re-marking.

## DSCP Re-Marking Configuration Examples

### Example 1

The following example includes the basic steps for re-marking DSCP value and mapping DSCP value to 802.1p.

1. Turn DSCP re-marking on globally, and define the DSCP-DSCP-802.1p mapping. You can use the default mapping.

```
RS G7028(config)# qos dscp re-marking
RS G7028(config)# qos dscp dscp-mapping <DSCP value (0-63)> <new value>
RS G7028(config)# qos dscp dot1p-mapping <DSCP value (0-63)> <802.1p value>
```

2. Enable DSCP re-marking on a port.

```
RS G7028(config)# interface port 1
RS G7028(config-if)# qos dscp re-marking
RS G7028(config-if)# exit
```

### Example 2

The following example assigns strict priority to VoIP traffic and a lower priority to all other traffic.

1. Create an ACL to re-mark DSCP value and COS queue for all VoIP packets.

```
RS G7028(config)# access-control list 2 tcp-udp source-port 5060 0xffff
RS G7028(config)# access-control list 2 meter committed-rate 10000000
RS G7028(config)# access-control list 2 meter enable
RS G7028(config)# access-control list 2 re-mark in-profile dscp 56
RS G7028(config)# access-control list 2 re-mark dot1p 7
RS G7028(config)# access-control list 2 action permit
```

2. Create an ACL to set a low priority to all other traffic.

```
RS G7028(config)# access-control list 3 action set-priority 1
RS G7028(config)# access-control list 3 action permit
```

3. Apply the ACLs to a port and enable DSCP marking.

```
RS G7028(config)# interface port 5
RS G7028(config-if)# access-control list 2
RS G7028(config-if)# access-control list 3 ethernet source-mac-address
00:00:00:00:00:00 00:00:00:00:00:00
RS G7028(config-if)# dscp-marking
RS G7028(config-if)# exit
```

4. Enable DSCP re-marking globally.

```
RS G7028(config)# qos dscp re-marking
```

5. Assign the DSCP re-mark value.

```
RS G7028(config)# qos dscp dscp-mapping 40 9
RS G7028(config)# qos dscp dscp-mapping 46 9
```

6. Assign strict priority to VoIP COS queue.

```
RS G7028(config)# qos transmit-queue weight-cos 7 0
```

7. Map priority value to COS queue for non-VoIP traffic.

```
RS G7028(config)# qos transmit-queue mapping 1 1
```

8. Assign weight to the non-VoIP COS queue.

```
RS G7028(config)# qos transmit-queue weight-cos 1 2
```

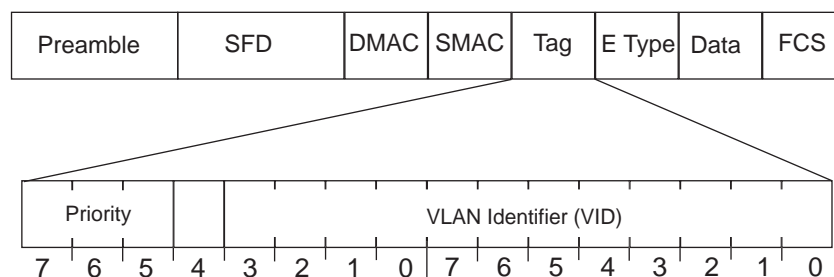
---

## Using 802.1p Priority to Provide QoS

The G7028/G7052 provides Quality of Service functions based on the priority bits in a packet's VLAN header. (The priority bits are defined by the 802.1p standard within the IEEE 802.1Q VLAN header.) The 802.1p bits, if present in the packet, specify the priority to be given to packets during forwarding. Packets with a numerically higher (non-zero) priority are given forwarding preference over packets with lower priority value.

The IEEE 802.1p standard uses eight levels of priority (0-7). Priority 7 is assigned to highest priority network traffic, such as OSPF or RIP routing table updates, priorities 5-6 are assigned to delay-sensitive applications such as voice and video, and lower priorities are assigned to standard applications. A value of 0 (zero) indicates a “best effort” traffic prioritization, and this is the default when traffic priority has not been configured on your network. The switch can filter packets based on the 802.1p values.

Figure 22. Layer 2 802.1q/802.1p VLAN tagged packet



Ingress packets receive a priority value, as follows:

- **Tagged packets**—switch reads the 802.1p priority in the VLAN tag.
- **Untagged packets**—switch tags the packet and assigns an 802.1p priority value, based on the port's default 802.1p priority.

Egress packets are placed in a COS queue based on the priority value, and scheduled for transmission based on the scheduling weight of the COS queue.

To configure a port's default 802.1p priority value, use the following commands.

```
RS G7028(config)# interface port 1
RS G7028(config-if)# dot1p <802.1p value (0-7)>
RS G7028(config-if)# exit
```

---

## Queuing and Scheduling

The G7028/G7052 can be configured to have either 2 or 8 output Class of Service (COS) queues per port, into which each packet is placed. Each packet's 802.1p priority determines its COS queue, except when an ACL action sets the COS queue of the packet.

You can configure the following attributes for COS queues:

- Map 802.1p priority value to a COS queue
- Define the scheduling weight of each COS queue

You can map 802.1p priority value to a COS queue, as follows:

```
RS G7028(config)# qos transmit-queue mapping <802.1p priority value (0-7)> <COS queue (0-7)>
```

To set the COS queue scheduling weight, use the following command.

```
RS G7028(config)# qos transmit-queue weight-cos <COSq number> <COSq weight (0-15)>
```

The scheduling weight can be set from 0 to 15. Weight values from 1 to 15 set the queue to use weighted round-robin (WRR) scheduling, which distributes larger numbers of packets to queues with the highest weight values. For distribution purposes, each packet is counted the same, regardless of the packet's size.

A scheduling weight of 0 (zero) indicates strict priority. Traffic in strict priority queue has precedence over other all queues. If more than one queue is assigned a weight of 0, the strict queue with highest queue number will be served first. Once all traffic in strict queues is delivered, any remaining bandwidth will be allocated to the WRR queues, divided according to their weight values.

**Note:** Use caution when assigning strict scheduling to queues. Heavy traffic in queues assigned with a weight of 0 can starve lower priority queues.

---

## WRED with ECN

Weighted Random Early Detection (WRED) is a congestion avoidance algorithm that helps prevent a TCP collapse, where a congested port indiscriminately drops packets from all sessions. The transmitting hosts wait to retransmit resulting in a dramatic drop in throughput. Often times, this TCP collapse repeats in a cycle, which results in a saw-tooth pattern of throughput. WRED selectively drops packets before the queue gets full, allowing majority of the traffic to flow smoothly.

WRED discards packets based on the CoS queues. Packets marked with lower priorities are discarded first.

Explicit Congestion Notification (ECN) is an extension to WRED. For packets that are ECN-aware, the ECN bit is marked to signal impending congestion instead of dropping packets. The transmitting hosts then slow down sending packets.

### How WRED/ECN work together

For implementing WRED, you must define a profile with minimum threshold, maximum threshold, and a maximum drop probability. The profiles can be defined on a port or a CoS.

For implementing ECN, you require ECN-specific field that has two bits—the ECN-capable Transport (ECT) bit and the CE (Congestion Experienced) bit—in the IP header. ECN is identified and defined by the values in these bits in the Differentiated Services field of IP Header. [Table 15](#) shows the combination values of the ECN bits.

*Table 15. ECN Bit Setting*

| ECT Bit | CE Bit | Description                                         |
|---------|--------|-----------------------------------------------------|
| 0       | 0      | Not ECN-capable                                     |
| 0       | 1      | Endpoints of the transport protocol are ECN-capable |
| 1       | 0      | Endpoints of the transport protocol are ECN-capable |
| 1       | 1      | Congestion experienced                              |

WRED and ECN work together as follows:

- If the number of packets in the queue is less than the minimum threshold, packets are transmitted. This happens irrespective of the ECN bit setting, and on networks where only WRED (without ECN) is enabled.
- If the number of packets in the queue is between the minimum threshold and the maximum threshold, one of the following occurs:
  - If the ECN field on the packet indicates that the endpoints are ECN-capable and the WRED algorithm determines that the packet has likely been dropped based on the drop probability, the ECT and CE bits for the packet are changed to 1, and the packet is transmitted.
  - If the ECN field on the packet indicates that neither endpoint is ECN-capable, the packet may be dropped based on the WRED drop probability. This is true even in cases where only WRED (without ECN) is enabled.
  - If the ECN field on the packet indicates that the network is experiencing congestion, the packet is transmitted. No further marking is required.
- If the number of packets in the queue is greater than the maximum threshold, packets are dropped based on the drop probability. This is the identical treatment a packet receives when only WRED (without ECN) is enabled.

## Configuring WRED/ECN

For configuring WRED, you must define a TCP profile and a non-TCP profile. WRED prioritizes TCP traffic over non-TCP traffic.

For configuring ECN, you must define a TCP profile. You don't need a non-TCP profile as ECN can be enabled only for TCP traffic.

If you do not configure the profiles, the profile thresholds are set to the maximum value to avoid drops.

**Note:** WRED/ECN can be configured only on physical ports and not on trunks. WRED and ECN are applicable only to unicast traffic.

Consider the following guidelines for configuring WRED/ECN:

- Profiles can be configured globally or per port. Global profiles are applicable to all ports.
- Always enable the global profile before applying the port-level profile.

**Note:** You can enable the global profile and disable the port-level profile. However, you must not enable the port-level profile if the global profile is disabled.

- WRED settings are dependent on Memory Management Unit (MMU) Settings. If you change the MMU setting, it could impact WRED functionality.
- You cannot enable WRED if you have QoS buffer settings such as Converged Enhanced Ethernet (CEE), Priority-based Flow Control (PFC), or Enhanced Transmission Selection (ETS).
- The number of WRED profiles per-port must match the total number of COS Queues configured in the system.
- If you have configured a TCP profile and enabled ECN, ECN marking happens when traffic experiencing congestion is TCP, or a mix of TCP and non-TCP traffic.
- Configure a TCP profile only after enabling ECN on the interface.

- You can apply TCP and non-TCP profile configurations irrespective of ECN status (enabled/disabled).

## WRED/ECN Configuration Example

Follow these steps to enable WRED/ECN and configure a global and/or port-level profile. If you configure global and port-level profile, WRED/ECN uses the port-level profile to make transmit/drop decisions when experiencing traffic congestion.

### Configure Global Profile for WRED

1. Enable WRED globally.

```
RS G7028(config)# qos random-detect enable
```

2. Enable a transmit queue.

```
RS G7028(config)# qos random-detect transmit-queue 0 enable
```

3. Configure WRED thresholds (minimum, maximum, and drop rate) for TCP traffic.

```
RS G7028(config)# qos random-detect transmit-queue 0 tcp min-threshold 1
max-threshold 2 drop-rate 3
```

**Note:** Percentages are of Average Queue available in hardware and not percentages of traffic.

4. Configure WRED thresholds (minimum, maximum, and drop rate) for non-TCP traffic.

```
RS G7028(config)# qos random-detect transmit-queue 0 non-tcp min-threshold 4
max-threshold 5 drop-rate 6
```

5. Select the port.

```
RS G7028(config)# interface port 1
```

6. Enable WRED for the port.

```
RS G7028(config-if)# random-detect enable
RS G7028(config-if)# exit
```

### Configure Port-level Profile for WRED

1. Enable WRED globally.

```
RS G7028(config)# qos random-detect enable
```

2. Select the port.

```
RS G7028(config)# interface port 1
```

3. Enable WRED for the port .

```
RS G7028(config-if)# random-detect enable
```

4. Enable a transmit queue.

```
RS G7028(config-if)# random-detect transmit-queue 0 enable
```

5. Configure WRED thresholds (minimum, maximum, and drop rate) for TCP traffic.

```
RS G7028(config-if)# random-detect transmit-queue 0 tcp min-threshold 11  
max-threshold 22 drop-rate 33
```

**Note:** Percentages are of Average Queue available in hardware and not percentages of traffic.

6. Configure WRED thresholds (minimum, maximum, and drop rate) for non-TCP traffic.

```
RS G7028(config-if)# random-detect transmit-queue 0 non-tcp min-threshold 44  
max-threshold 55 drop-rate 66  
RS G7028(config-if)# exit
```

## Configure Global Profile for ECN

1. Enable ECN globally.

```
RS G7028(config)# qos random-detect ecn enable
```

2. Enable a transmit queue.

```
RS G7028(config)# qos random-detect transmit-queue 0 enable
```

3. Configure ECN thresholds (minimum, maximum, and drop rate) for TCP traffic.

```
RS G7028(config)# qos random-detect transmit-queue 0 tcp min-threshold 1  
max-threshold 2 drop-rate 3
```

**Note:** Percentages are of Average Queue available in hardware and not percentages of traffic.

4. Select the port.

```
RS G7028(config)# interface port 1
```

5. Enable ECN for the port.

```
RS G7028(config-if)# random-detect ecn enable  
RS G7028(config-if)# exit
```

## Configure Port-level Profile for ECN

1. Enable ECN globally.

```
RS G7028(config)# qos random-detect ecn enable
```

2. Select the port.

```
RS G7028(config)# interface port 1
```

3. Enable ECN for the port.

```
RS G7028(config-if)# random-detect ecn enable
```

4. Enable a transmit queue.

```
RS G7028(config-if)# random-detect transmit-queue 0 enable
```

5. Configure ECN thresholds (minimum, maximum, and drop rate) for TCP traffic.

```
RS G7028(config-if)# random-detect transmit-queue 0 tcp min-threshold 11
max-threshold 22 drop-rate 33
RS G7028(config-if)# exit
```

**Note:** Percentages are of Average Queue available in hardware and not percentages of traffic.

## Verifying WRED/ECN

Use the following command to view global WRED/ECN information.

```
RS G7028(config)# show qos random-detect
Current wred and ecn configuration:
Global ECN: Enable
Global WRED: Enable
TQ0: -WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
Ena 10 20 30 10 20 30
TQ1: -WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
Dis 0 0 0 0 0 0
TQ2: -WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
Dis 0 0 0 0 0 0
TQ3: -WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
Dis 0 0 0 0 0 0
TQ4: -WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
Dis 0 0 0 0 0 0
TQ5: -WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
Dis 0 0 0 0 0 0
TQ6: -WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
Dis 0 0 0 0 0 0
TQ7: -WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
Dis 0 0 0 0 0 0
```

Use the following command to view port-level WRED/ECN information.

```
RS G7028(config)# show interface port 1 random-detect
Port: 1
    ECN: Enable
    WRED: Enable
TQ0:-WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
    Dis    0        0        0        0        0        0
TQ1:-WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
    Ena    4        5        6        1        2        3
TQ2:-WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
    Dis    0        0        0        0        0        0
TQ3:-WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
    Dis    0        0        0        0        0        0
TQ4:-WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
    Dis    0        0        0        0        0        0
TQ5:-WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
    Dis    0        0        0        0        0        0
TQ6:-WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
    Dis    0        0        0        0        0        0
TQ7:-WRED-TcpMinThr-TcpMaxThr-TcpDrate-NonTcpMinThr-NonTcpMaxThr-NonTcpDrate-
    Dis    0        0        0        0        0        0
```

# Part 4: IP Features

This section discusses basic and advanced IP features:

- Basic IP Features
- IPv6 Host Management
- Internet Group Management Protocol (IGMP)



---

## Chapter 13. Basic IP Features

This chapter provides configuration background and examples for the following topics:

- [“Dynamic Host Configuration Protocol” on page 180](#)

---

## Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) is a transport protocol that provides a framework for automatically assigning IP addresses and configuration information to other IP hosts or clients in a large TCP/IP network. Without DHCP, the IP address must be entered manually for each network device. DHCP allows a network administrator to distribute IP addresses from a central point and automatically send a new IP address when a device is connected to a different place in the network.

The switch accepts gateway configuration parameters if they have not been configured manually. The switch ignores DHCP gateway parameters if the gateway is configured.

DHCP is an extension of another network IP management protocol, Bootstrap Protocol (BOOTP), with an additional capability of being able to allocate reusable network addresses and configuration parameters for client operation.

Built on the client/server model, DHCP allows hosts or clients on an IP network to obtain their configurations from a DHCP server, thereby reducing network administration. The most significant configuration the client receives from the server is its required IP address; (other optional parameters include the “generic” file name to be booted, the address of the default gateway, and so forth).

To enable DHCP on a switch interface, use the following command:

```
RS G7028(config)# system dhcp
```

---

## Chapter 14. Internet Protocol Version 6

Internet Protocol version 6 (IPv6) is a network layer protocol intended to expand the network address space. IPv6 is a robust and expandable protocol that meets the need for increased physical address space. The switch supports the following RFCs for IPv6-related features:

- |            |            |            |            |
|------------|------------|------------|------------|
| • RFC 1981 | • RFC 3306 | • RFC 4007 | • RFC 4443 |
| • RFC 2404 | • RFC 3307 | • RFC 4213 | • RFC 4552 |
| • RFC 2410 | • RFC 3411 | • RFC 4291 | • RFC 4718 |
| • RFC 2451 | • RFC 3412 | • RFC 4292 | • RFC 4835 |
| • RFC 2460 | • RFC 3413 | • RFC 4293 | • RFC 4861 |
| • RFC 2474 | • RFC 3414 | • RFC 4301 | • RFC 4862 |
| • RFC 2526 | • RFC 3484 | • RFC 4302 | • RFC 5095 |
| • RFC 2711 | • RFC 3602 | • RFC 4303 | • RFC 5114 |
| • RFC 2740 | • RFC 3810 | • RFC 4306 | • RFC 5340 |
| • RFC 3289 | • RFC 3879 | • RFC 4307 |            |

This chapter describes the basic configuration of IPv6 addresses and how to manage the switch via IPv6 host management.

---

## IPv6 Limitations

The following IPv6 features are not supported in this release.

- Dynamic Host Control Protocol for IPv6 (DHCPv6)
- Routing Information Protocol for IPv6 (RIPng)

Most other IBM Networking OS 7.6 features permit IP addresses to be configured using either IPv4 or IPv6 address formats. However, the following switch features support IPv4 only:

- Bootstrap Protocol (BOOTP) and DHCP
- RADIUS, TACACS+ and LDAP
- QoS metering and re-marking ACLs for out-profile traffic
- Internet Group Management Protocol (IGMP)

---

## IPv6 Address Format

The IPv6 address is 128 bits (16 bytes) long and is represented as a sequence of eight 16-bit hex values, separated by colons.

Each IPv6 address has two parts:

- Subnet prefix representing the network to which the interface is connected
- Local identifier, either derived from the MAC address or user-configured

The preferred hexadecimal format is as follows:

XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX

Example IPv6 address:

FEDC:BA98:7654:BA98:FEDC:1234:ABCD:5412

Some addresses can contain long sequences of zeros. A single contiguous sequence of zeros can be compressed to :: (two colons). For example, consider the following IPv6 address:

FE80:0:0:0:2AA:FF:FA:4CA2

The address can be compressed as follows:

FE80::2AA:FF:FA:4CA2

Unlike IPv4, a subnet mask is not used for IPv6 addresses. IPv6 uses the subnet prefix as the network identifier. The prefix is the part of the address that indicates the bits that have fixed values or are the bits of the subnet prefix. An IPv6 prefix is written in address/prefix-length notation. For example, in the following address, 64 is the network prefix:

21DA:D300:0000:2F3C::/64

IPv6 addresses can be either user-configured or automatically configured. Automatically configured addresses always have a 64-bit subnet prefix and a 64-bit interface identifier. In most implementations, the interface identifier is derived from the switch's MAC address, using a method called EUI-64.

Most IBM N/OS 7.6 features permit IP addresses to be configured using either IPv4 or IPv6 address formats. Throughout this manual, *IP address* is used in places where either an IPv4 or IPv6 address is allowed. In places where only one type of address is allowed, the type (*IPv4* or *IPv6*) is specified.

---

## IPv6 Address Types

IPv6 supports three types of addresses: unicast (one-to-one), multicast (one-to-many), and anycast (one-to-nearest). Multicast addresses replace the use of broadcast addresses.

### Unicast Address

Unicast is a communication between a single host and a single receiver. Packets sent to a unicast address are delivered to the interface identified by that address. IPv6 defines the following types of unicast address:

- **Global Unicast address:** An address that can be reached and identified globally. Global Unicast addresses use the high-order bit range up to FF00, therefore all non-multicast and non-link-local addresses are considered to be global unicast. A manually configured IPv6 address must be fully specified. Autoconfigured IPv6 addresses are comprised of a prefix combined with the 64-bit EUI. RFC 4291 defines the IPv6 addressing architecture.

The interface ID must be unique within the same subnet.

- **Link-local unicast address:** An address used to communicate with a neighbor on the same link. Link-local addresses use the format FE80 : : EUI

Link-local addresses are designed to be used for addressing on a single link for purposes such as automatic address configuration, neighbor discovery, or when no routers are present.

Routers must not forward any packets with link-local source or destination addresses to other links.

### Multicast

Multicast is communication between a single host and multiple receivers. Packets are sent to all interfaces identified by that address. An interface may belong to any number of multicast groups.

A multicast address (FF00 - FFFF) is an identifier for a group interface. The multicast address most often encountered is a solicited-node multicast address using prefix FF02 : : 1 : FF00 : 0000 / 104 with the low-order 24 bits of the unicast or anycast address.

The following well-known multicast addresses are pre-defined. The group IDs defined in this section are defined for explicit scope values, as follows:

FF00 : : : : : 0 through FF0F : : : : : 0

### Anycast

Packets sent to an anycast address or list of addresses are delivered to the nearest interface identified by that address. Anycast is a communication between a single sender and a list of addresses.

Anycast addresses are allocated from the unicast address space, using any of the defined unicast address formats. Thus, anycast addresses are syntactically indistinguishable from unicast addresses. When a unicast address is assigned to more than one interface, thus turning it into an anycast address, the nodes to which the address is assigned must be explicitly configured to know that it is an anycast address.

---

## IPv6 Address Autoconfiguration

IPv6 supports the following types of address autoconfiguration:

- **Stateful address configuration**

Address configuration is based on the use of a stateful address configuration protocol, such as DHCPv6, to obtain addresses and other configuration options.

- **Stateless address configuration**

Address configuration is based on the receipt of Router Advertisement messages that contain one or more Prefix Information options.

N/OS 7.6 supports stateless address configuration.

Stateless address configuration allows hosts on a link to configure themselves with link-local addresses and with addresses derived from prefixes advertised by local routers. Even if no router is present, hosts on the same link can configure themselves with link-local addresses and communicate without manual configuration.

---

## IPv6 Interfaces

Each IPv6 interface supports multiple IPv6 addresses. You can manually configure up to two IPv6 addresses for each interface, or you can allow the switch to use stateless autoconfiguration.

You can manually configure two IPv6 addresses for each interface, as follows:

- Initial IPv6 address is a global unicast or anycast address.

```
RS G7028(config)# interface ip <interface number>
RS G7028(config-ip-if)# ipv6 address <IPv6 address>
```

Note that you cannot configure both addresses as anycast. If you configure an anycast address on the interface you must also configure a global unicast address on that interface.

- Second IPv6 address can be a unicast or anycast address.

```
RS G7028(config-ip-if)# ipv6 secaddr6 <IPv6 address>
RS G7028(config-ip-if)# exit
```

You cannot configure an IPv4 address on an IPv6 management interface. Each interface can be configured with only one address type: either IPv4 or IPv6, but not both. When changing between IPv4 and IPv6 address formats, the prior address settings for the interface are discarded.

Each IPv6 interface can belong to only one VLAN. Each VLAN can support only one IPv6 interface. Each VLAN can support multiple IPv4 interfaces.

Use the following commands to configure the IPv6 gateway:

```
RS G7028(config)# ip gateway6 1 address <IPv6 address>
RS G7028(config)# ip gateway6 1 enable
```

IPv6 gateway 1 is reserved for IPv6 data interfaces. IPv6 gateway 4 is the default IPv6 management gateway.

---

## Supported Applications

The following applications have been enhanced to provide IPv6 support.

- **Ping**

The ping command supports IPv6 addresses. Use the following format to ping an IPv6 address:

```
ping <host name>|<IPv6 address> [-n <tries (0-4294967295)>]
[-w <msec delay (0-4294967295)>] [-l <length (0/32-65500/2080)>]
[-s <IP source>] [-v <TOS (0-255)>] [-f] [-t]
```

To ping a link-local address (begins with FE80), provide an interface index, as follows:

```
ping <IPv6 address>%<Interface index> [-n <tries (0-4294967295)>]
[-w <msec delay (0-4294967295)>] [-l <length (0/32-65500/2080)>]
[-s <IP source>] [-v <TOS (0-255)>] [-f] [-t]
```

- **Traceroute**

The traceroute command supports IPv6 addresses (but not link-local addresses).

Use the following format to perform a traceroute to an IPv6 address:

```
traceroute <host name>|<IPv6 address> [<max-hops (1-32)>]
[<msec delay (1-4294967295)>]]
```

- **Telnet server**

The telnet command supports IPv6 addresses (but not link-local addresses).

Use the following format to Telnet into an IPv6 interface on the switch:

```
telnet <host name>|<IPv6 address> [<port>]
```

- **Telnet client**

The telnet command supports IPv6 addresses (but not link-local addresses).

Use the following format to Telnet to an IPv6 address:

```
telnet <host name>|<IPv6 address> [<port>]
```

- **HTTP/HTTPS**

The HTTP/HTTPS servers support both IPv4 and IPv6 connections.

- **SSH**

Secure Shell (SSH) connections over IPv6 are supported (but not link-local addresses). The following syntax is required from the client:

```
ssh -u <IPv6 address>
```

Example:

```
ssh -u 2001:2:3:4:0:0:0:142
```

- **TFTP**

The TFTP commands support both IPv4 and IPv6 addresses. Link-local addresses are not supported.

- **FTP**

The FTP commands support both IPv4 and IPv6 addresses. Link-local addresses are not supported.

- **DNS client**

DNS commands support both IPv4 and IPv6 addresses. Link-local addresses are not supported. Use the following command to specify the type of DNS query to be sent first:

```
RS G7028(config)# ip dns ipv6 request-version {ipv4|ipv6}
```

If you set the request version to `ipv4`, the DNS application sends an `A` query first, to resolve the hostname with an IPv4 address. If no `A` record is found for that hostname (no IPv4 address for that hostname) an `AAAA` query is sent to resolve the hostname with a IPv6 address.

If you set the request version to `ipv6`, the DNS application sends an `AAAA` query first, to resolve the hostname with an IPv6 address. If no `AAAA` record is found for that hostname (no IPv6 address for that hostname) an `A` query is sent to resolve the hostname with an IPv4 address.

---

## Configuration Guidelines

When you configure an interface for IPv6, consider the following guidelines:

- Support for subnet router anycast addresses is not available.
- A single interface can accept either IPv4 or IPv6 addresses, but not both IPv4 and IPv6 addresses.
- A single interface can accept multiple IPv6 addresses.
- A single interface can accept only one IPv4 address.
- If you change the IPv6 address of a configured interface to an IPv4 address, all IPv6 settings are deleted.
- A single VLAN can support only one IPv6 interface.
- Health checks are not supported for IPv6 gateways.
- IPv6 interfaces support Path MTU Discovery. The CPU's MTU is fixed at 1500 bytes.
- Support for jumbo frames (1,500 to 12,288 byte MTUs) is limited. Any jumbo frames intended for the CPU must be fragmented by the remote node. The switch can re-assemble fragmented packets up to 12k. It can also fragment and transmit jumbo packets received from higher layers.

---

## IPv6 Configuration Example

This section provides steps to configure IPv6 on the switch. Use the following example to manually configure IPv6 on an interface.

1. Assign an IPv6 address and prefix length to the interface.

```
RS G7028(config)# interface ip 3
RS G7028(config-ip-if)# ipv6 address 2001:BA98:7654:BA98:FEDC:1234:ABCD:5214
RS G7028(config-ip-if)# ipv6 prefixlen 64
RS G7028(config-ip-if)# ipv6 seccaddr6 2003::1 32
RS G7028(config-ip-if)# vlan 2
RS G7028(config-ip-if)# enable
RS G7028(config-ip-if)# exit
```

The secondary IPv6 address is compressed, and the prefix length is 32.

2. Configure the IPv6 default gateway.

```
RS G7028(config)# ip gateway6 1 address 2001:BA98:7654:BA98:FEDC:1234:ABCD:5412
RS G7028(config)# ip gateway6 1 enable
```

3. Verify the configuration.

```
RS G7028(config-ip-if)# show layer3
```



---

## Chapter 15. Internet Group Management Protocol

Internet Group Management Protocol (IGMP) is used by IPv4 Multicast routers (Mrouters) to learn about the existence of host group members on their directly attached subnet. The IPv4 Mrouters get this information by broadcasting IGMP Membership Queries and listening for IPv4 hosts reporting their host group memberships. This process is used to set up a client/server relationship between an IPv4 multicast source that provides the data streams and the clients that want to receive the data. The switch supports three versions of IGMP:

- IGMPv1: Defines the method for hosts to join a multicast group. However, this version does not define the method for hosts to leave a multicast group. See RFC 1112 for details.
- IGMPv2: Adds the ability for a host to signal its desire to leave a multicast group. See RFC 2236 for details.
- IGMPv3: Adds support for source filtering by which a host can report interest in receiving packets only from specific source addresses, or from all but specific source addresses, sent to a particular multicast address. See RFC 3376 for details.

The G7028/G7052 can perform IGMP Snooping, and connect to static Mrouters. The G7028/G7052 can act as a Querier, and participate in the IGMP Querier election process.

The following topics are discussed in this chapter:

- [“IGMP Terms” on page 192](#)
- [“How IGMP Works” on page 193](#)
- [“IGMP Capacity and Default Values” on page 194](#)
- [“IGMP Snooping” on page 195](#)
- [“Additional IGMP Features” on page 207](#)

---

## IGMP Terms

The following are commonly used IGMP terms:

- Multicast traffic: Flow of data from one source to multiple destinations.
- Group: A multicast stream to which a host can join. Multicast groups have IP addresses in the range: 224.0.1.0 to 239.255.255.255.
- IGMP Querier: A router or switch in the subnet that generates *Membership Queries*.
- IGMP Snooper: A networking device that forwards multicast traffic only to hosts that are interested in receiving multicast data. This device can be a router or a Layer 2/3 switch.
- Multicast Router: A router configured to make routing decisions for multicast traffic. The router identifies the type of packet received (unicast or multicast) and forwards the packet to the intended destination.
- IGMP Proxy: A device that filters Join messages and Leave messages sent upstream to the Mrouter to reduce the load on the Mrouter.
- Membership Report: A report sent by the host that indicates an interest in receiving multicast traffic from a multicast group.
- Leave: A message sent by the host when it wants to leave a multicast group.
- FastLeave: A process by which the switch stops forwarding multicast traffic to a port as soon as it receives a Leave message.
- Membership Query: Message sent by the Querier to verify if hosts are listening to a group.
- General Query: A *Membership Query* sent to all hosts. The Group address field for general queries is 0.0.0.0 and the destination address is 224.0.0.1.
- Group-specific Query: A *Membership Query* sent to all hosts in a multicast group.

---

## How IGMP Works

When IGMP is not configured, switches forward multicast traffic through all ports, increasing network load. When IGMPv2 is configured on a switch, multicast traffic flows as follows:

- A server sends multicast traffic to a multicast group.
- The Mrouter sends *Membership Queries* to the switch, which forwards them to all ports in a given VLAN.
- Hosts respond with *Membership Reports* if they want to join a group. The switch forwards these reports to the Mrouter.
- The switch forwards multicast traffic only to hosts that have joined a group.
- The Mrouter periodically sends *Membership Queries* to ensure that a host wants to continue receiving multicast traffic. If a host does not respond, the IGMP Snooper stops sending traffic to the host.
- The host can initiate the Leave process by sending an IGMP Leave packet to the IGMP Snooper.
- When a host sends an IGMPv2 Leave packet, the IGMP Snooper queries to find out if any other host connected to the port is interested in receiving the multicast traffic. If it does not receive a Join message in response, the IGMP Snooper removes the group entry and passes on the information to the Mrouter.

The G7028/G7052 supports the following:

- IGMP version 1, 2, and 3
- 128 Mrouters

**Note:** Unknown multicast traffic is sent to all ports if the flood option is enabled and no Membership Report was learned for that specific IGMP group. If the flood option is disabled, unknown multicast traffic is discarded if no hosts or Mrouters are learned on a switch.

To enable or disable IGMP flood, use the following command:

```
RS G7028(config)# vlan <vlan ID>
RS G7028(config-vlan)# [no] flood
```

---

## IGMP Capacity and Default Values

The following table lists the maximum and minimum values of the G7028/G7052 variables.

Table 16. G7028/G7052 Capacity Table

| Variable               | Maximum |
|------------------------|---------|
| IGMP Entries - Snoop   | 256     |
| VLANs - Snoop          | 512     |
| Static Mrouters        | 128     |
| Dynamic Mrouters       | 128     |
| Number of IGMP Filters | 16      |

The following table lists the default settings for IGMP features and variables.

Table 17. IGMP Default Configuration Settings

| Field                        | Default Value                                                                                                       |
|------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Global IGMP State            | Disabled                                                                                                            |
| IGMP Querier                 | Disabled                                                                                                            |
| IGMP Snooping                | Disabled                                                                                                            |
| IGMP Filtering               | Disabled                                                                                                            |
| IP Multicast (IPMC) Flood    | Enabled                                                                                                             |
| IGMP FastLeave               | Disabled for all VLANs                                                                                              |
| IGMP Mrouter Timeout         | 255 Seconds                                                                                                         |
| IGMP Report Timeout Variable | 10 Seconds                                                                                                          |
| IGMP Query-Interval Variable | 125 Seconds                                                                                                         |
| IGMP Robustness Variable     | 2                                                                                                                   |
| IGMPv3                       | Disabled                                                                                                            |
| IGMPv3 number of sources     | 8 (The switch processes only the first eight sources listed in the IGMPv3 group record.)<br><br>Valid range: 1 - 64 |
| IGMPv3 - allow v1v2 Snooping | Enabled                                                                                                             |

---

## IGMP Snooping

IGMP Snooping allows a switch to listen to the IGMP conversation between hosts and Mrouters. By default, a switch floods multicast traffic to all ports in a broadcast domain. With IGMP Snooping enabled, the switch learns the ports interested in receiving multicast data and forwards it only to those ports. IGMP Snooping conserves network resources.

The switch can sense IGMP *Membership Reports* from attached hosts and acts as a proxy to set up a dedicated path between the requesting host and a local IPv4 Mrouter. After the path is established, the switch blocks the IPv4 multicast stream from flowing through any port that does not connect to a host member, thus conserving bandwidth.

## IGMP Querier

For IGMP Snooping to function, you must have an Mrouter on the network that generates IGMP Query packets. Enabling the IGMP Querier feature on the switch allows it to participate in the Querier election process. If the switch is elected as the Querier, it will send IGMP Query packets for the LAN segment.

### Querier Election

If multiple Mrouters exist on the network, only one can be configured as a Querier. The Mrouters elect the one with the lowest source IPv4 address or MAC address as the Querier. The Querier performs all periodic membership queries. All other Mrouters (non-Queriers) do not send IGMP Query packets.

**Note:** When IGMP Querier is enabled on a VLAN, the switch performs the role of an IGMP Querier only if it meets the IGMP Querier election criteria.

Each time the Querier switch sends an IGMP Query packet, it initializes a *general query timer*. If a Querier receives a General Query packet from an Mrouter with a lower IP address or MAC address, it transitions to a non-Querier state and initializes an *other querier present timer*. When this timer expires, the Mrouter transitions back to the Querier state and sends a General Query packet.

Follow this procedure to configure IGMP Querier.

1. Enable IGMP and configure the source IPv4 address for IGMP Querier on a VLAN.

```
RS G7028(config)# ip igmp enable
RS G7028(config)# ip igmp querier vlan 2 source-ip 10.10.10.1
```

2. Enable IGMP Querier on the VLAN.

```
RS G7028(config)# ip igmp querier vlan 2 enable
```

3. Configure the querier election type and define the address.

```
RS G7028(config)# ip igmp querier vlan 2 election-type ipv4
```

#### 4. Verify the configuration.

```
RS G7028# show ip igmp querier vlan 2

Current IGMP snooping Querier information:
IGMP Querier information for vlan 2:
Other IGMP querier - none
Switch-querier enabled, current state: Querier
Switch-querier type: Ipv4, address 10.10.10.1,
Switch-querier general query interval: 125 secs,
Switch-querier max-response interval: 100 'tenths of secs',
Switch-querier startup interval: 31 secs, count: 2
Switch-querier robustness: 2
IGMP configured version is v3
IGMP Operating version is v3
```

## IGMP Groups

One IGMP entry is allocated for each unique join request, based on the VLAN and IGMP group address. If multiple ports join the same IGMP group using the same VLAN, only a single IGMP entry is used.

## IGMPv3 Snooping

IGMPv3 includes new Membership Report messages that extend IGMP functionality. The switch provides snooping capability for all types of IGMPv3 *Membership Reports*.

IGMPv3 supports Source-Specific Multicast (SSM). SSM identifies session traffic by both source and group addresses.

The IGMPv3 implementation keeps records on the multicast hosts present in the network. If a host is already registered, when it receives a new `IS_INC`, `TO_INC`, `IS_EXC`, or `TO_EXC` report from same host, the switch makes the correct transition to new (port-host-group) registration based on the IGMPv3 RFC. The registrations of other hosts for the same group on the same port are not changed.

The G7028/G7052 supports the following IGMPv3 filter modes:

- **INCLUDE mode:** The host requests membership to a multicast group and provides a list of IPv4 addresses from which it wants to receive traffic.
- **EXCLUDE mode:** The host requests membership to a multicast group and provides a list of IPv4 addresses from which it does not want to receive traffic. This indicates that the host wants to receive traffic only from sources that are not part of the Exclude list. To disable snooping on EXCLUDE mode reports, use the following command:

```
RS G7028(config)# no ip igmp snoop igmpv3 exclude
```

By default, the G7028/G7052 snoops the first eight sources listed in the IGMPv3 Group Record. Use the following command to change the number of snooping sources:

```
RS G7028(config)# ip igmp snoop igmpv3 sources <1-64>
```

IGMPv3 Snooping is compatible with IGMPv1 and IGMPv2 Snooping. To disable snooping on version 1 and version 2 reports, use the following command:

```
RS G7028(config)# no ip igmp snoop igmpv3 v1v2
```

## IGMP Snooping Configuration Guidelines

Consider the following guidelines when you configure IGMP Snooping:

- IGMP operation is independent of the routing method. You can use RIP, OSPF, or static routes for Layer 3 routing.
- When multicast traffic flood is disabled, the multicast traffic sent by the multicast server is discarded if no hosts or Mrouter are learned on the switch.
- The Mrouter periodically sends IGMP Queries.
- The switch learns the Mrouter on the port connected to the router when it sees Query messages. The switch then floods the IGMP queries on all other ports including a Trunk Group, if any.
- Multicast hosts send IGMP Reports as a reply to the IGMP Queries sent by the Mrouter.
- The switch can also learn an Mrouter when it receives a PIM hello packet from another device. However, an Mrouter learned from a PIM packet has a lower priority than an Mrouter learned from an IGMP Query. A switch overwrites an Mrouter learned from a PIM packet when it receives an IGMP Query on the same port.
- A host sends an IGMP Leave message to its multicast group. The expiration timer for this group is updated to IGMP timeout variable (the default is 10 seconds). The Layer 2/3 switch sends IGMP Group-Specific Query to the host that had sent the Leave message. If the host does not respond with an IGMP Report during the timeout interval, all the groups expire and the switch deletes the host from the IGMP groups table. The switch then proxies the IGMP Leave messages to the Mrouter.

## IGMP Snooping Configuration Example

This section provides steps to configure IGMP Snooping on the G7028/G7052.

1. Configure port and VLAN membership on the switch.
2. Add VLANs to IGMP Snooping.

```
RS G7028(config)# ip igmp snoop vlan 1
```

3. Enable IGMP Snooping.

```
RS G7028(config)# ip igmp snoop enable
```

4. Enable IGMPv3 Snooping (optional).

```
RS G7028(config)# ip igmp snoop igmpv3 enable
```

5. Enable the IGMP feature.

```
RS G7028(config)# ip igmp enable
```

6. View dynamic IGMP information.

```
RS G7028# show ip igmp groups
```

Total entries: 2 Total IGMP groups: 1

Note: The <Total IGMP groups> number is computed as the number of unique (Group, Vlan) entries!

Note: Local groups (224.0.0.x) are not snooped/relayed and will not appear.

| Source      | Group     | VLAN | Port | Version | Mode | Expires | Fwd |
|-------------|-----------|------|------|---------|------|---------|-----|
| 10.1.1.1    | 232.1.1.1 | 2    | 4    | V3      | INC  | 4:16    | Yes |
| 10.1.1.5    | 232.1.1.1 | 2    | 4    | V3      | INC  | 4:16    | Yes |
| *           | 232.1.1.1 | 2    | 4    | V3      | INC  | -       | No  |
| 10.10.10.43 | 235.0.0.1 | 9    | 1    | V3      | INC  | 2:26    | Yes |
| *           | 236.0.0.1 | 9    | 1    | V3      | EXC  | -       | Yes |

```
RS G7028# show ip igmp mrouter
```

Total entries: 1 Total number of dynamic mrouter: 1

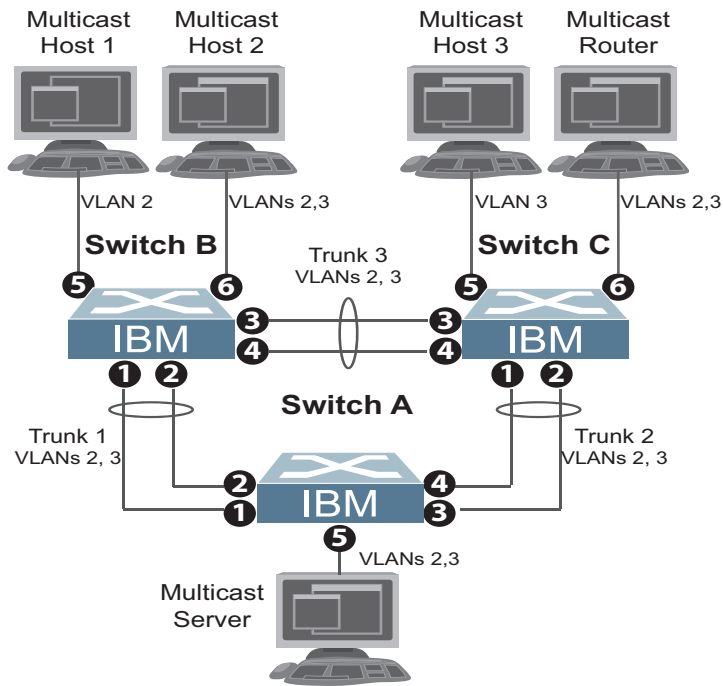
| SrcIP       | VLAN | Port | Version | Expires | MRT     | QRV | QQIC |
|-------------|------|------|---------|---------|---------|-----|------|
| 10.1.1.1    | 2    | 21   | V3      | 4:09    | 128     | 2   | 125  |
| 10.1.1.5    | 2    | 23   | V2      | 4:09    | 125     | -   | -    |
| 10.10.10.43 | 9    | 24   | V2      | static  | unknown | -   | -    |

These commands display information about IGMP Groups and Mrouters learned by the switch.

## Advanced Configuration Example: IGMP Snooping

Figure 23 shows an example topology. Switches B and C are configured with IGMP Snooping.

Figure 23. Topology



Devices in this topology are configured as follows:

- STG2 includes VLAN2; STG3 includes VLAN3.
- The multicast server sends IP multicast traffic for the following groups:
  - VLAN 2, 225.10.0.11 – 225.10.0.12, Source: 22.10.0.11
  - VLAN 2, 225.10.0.13 – 225.10.0.15, Source: 22.10.0.13
  - VLAN 3, 230.0.2.1 – 230.0.2.2, Source: 22.10.0.1
  - VLAN 3, 230.0.2.3 – 230.0.2.5, Source: 22.10.0.3
- The Mrouter sends IGMP Query packets in VLAN 2 and VLAN 3. The Mrouter's IP address is 10.10.10.10.
- The multicast hosts send the following IGMP Reports:
  - IGMPv2 Report, VLAN 2, Group: 225.10.0.11, Source: \*
  - IGMPv2 Report, VLAN 3, Group: 230.0.2.1, Source: \*
  - IGMPv3 IS\_INCLUDE Report, VLAN 2, Group: 225.10.0.13, Source: 22.10.0.13
  - IGMPv3 IS\_INCLUDE Report, VLAN 3, Group: 230.0.2.3, Source: 22.10.0.3

- The hosts receive multicast traffic as follows:
  - Host 1 receives multicast traffic for groups (\*, 225.10.0.11), (22.10.0.13, 225.10.0.13)
  - Host 2 receives multicast traffic for groups (\*, 225.10.0.11), (\*, 230.0.2.1), (22.10.0.13, 225.10.0.13), (22.10.0.3, 230.0.2.3)
  - Host 3 receives multicast traffic for groups (\*, 230.0.2.1), (22.10.0.3, 230.0.2.3)
- The Mrouter receives all the multicast traffic.

## Prerequisites

Before you configure IGMP Snooping, ensure you have performed the following actions:

- Configured VLANs.
- Enabled IGMP.
- Configured a switch or Mrouter as the Querier.
- Identified the IGMP version(s) you want to enable.
- Disabled IGMP flooding.

## Configuration

This section provides the configuration details of the switches shown in [Figure 23](#).

### Switch A Configuration

1. Configure VLANs and tagging.

```
RS G7028(config)# interface port 1-5
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# switchport trunk allowed vlan remove 1
RS G7028(config-if)# exit

RS G7028(config)# interface port 1-5
RS G7028(config-if)# switchport trunk allowed vlan add 2,3
RS G7028(config-if)# exit
```

2. Configure an IP interface with IPv4 address, and assign a VLAN.

```
RS G7028(config)# interface ip 1
RS G7028(config-ip-if)# ip address 10.10.10.1 enable
RS G7028(config-ip-if)# vlan 2
RS G7028(config-ip-if)# exit
```

3. Assign a bridge priority lower than the default bridge priority to enable the switch to become the STP root in STG 2 and 3.

```
RS G7028(config)# spanning-tree stp 2 bridge priority 4096
RS G7028(config)# spanning-tree stp 3 bridge priority 4096
```

#### 4. Configure LACP dynamic trunk groups (portchannels).

```
RS G7028(config)# interface port 1
RS G7028(config-if)# lacp key 100
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit

RS G7028(config)# interface port 2
RS G7028(config-if)# lacp key 100
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit

RS G7028(config)# interface port 3
RS G7028(config-if)# lacp key 200
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit

RS G7028(config)# interface port 4
RS G7028(config-if)# lacp key 200
RS G7028(config-if)# lacp mode active
```

### Switch B Configuration

#### 1. Configure VLANs and tagging.

```
RS G7028(config)# interface port 1-6
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# exit

RS G7028(config)# interface port 1-6
RS G7028(config-if)# switchport trunk allowed vlan add 2
RS G7028(config-if)# exit

RS G7028(config)# interface port 1-4,6
RS G7028(config-if)# switchport trunk allowed vlan add 3
RS G7028(config-if)# exit

RS G7028(config)# interface port 1-5
RS G7028(config-if)# switchport trunk allowed vlan remove 1
RS G7028(config-if)# exit
```

#### 2. Configure an IP interface with IPv4 address, and assign a VLAN.

```
RS G7028(config)# interface ip 1
RS G7028(config-ip-if)# ip address 10.10.10.2 enable
RS G7028(config-ip-if)# vlan 2
RS G7028(config-ip-if)# exit
```

#### 3. Configure STP. Reset the ports to make the edge configuration operational.

```
RS G7028(config)# interface port 5,6
RS G7028(config-if)# spanning-tree edge
RS G7028(config-if)# shutdown
RS G7028(config-if)# no shutdown
RS G7028(config-if)# exit
```

#### 4. Configure an LACP dynamic trunk group (portchannel).

```
RS G7028(config)# interface port 1,2
RS G7028(config-if)# lacp key 300
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit
```

5. Configure a static trunk group (portchannel).

```
RS G7028(config)# portchannel 1 port 3,4 enable
```

6. Configure IGMP Snooping.

```
RS G7028(config)# ip igmp enable
RS G7028(config)# ip igmp snoop vlan 2,3
RS G7028(config)# ip igmp snoop source-ip 10.10.10.2
RS G7028(config)# ip igmp snoop igmpv3 enable
RS G7028(config)# ip igmp snoop igmpv3 sources 64
RS G7028(config)# ip igmp snoop enable

RS G7028(config)# vlan 2
RS G7028(config-vlan)# no flood
RS G7028(config-vlan)# exit

RS G7028(config)# vlan 3
RS G7028(config-vlan)# no flood
RS G7028(config-vlan)# exit
```

## Switch C Configuration

1. Configure VLANs and tagging.

```
RS G7028(config)# interface port 1-4,6
RS G7028(config-if)# switchport mode trunk
RS G7028(config-if)# exit

RS G7028(config)# interface port 1-4,6
RS G7028(config-if)# switchport trunk allowed vlan add 2
RS G7028(config-if)# exit

RS G7028(config)# interface port 1-6
RS G7028(config-if)# switchport trunk allowed vlan add 3
RS G7028(config-if)# exit

RS G7028(config)# interface 1-6
RS G7028(config-if)# switchport trunk allowed vlan remove 1
RS G7028(config-if)# exit
```

2. Configure an IP interface with IPv4 address, and assign a VLAN.

```
RS G7028(config)# interface ip 1
RS G7028(config-ip-if)# ip address 10.10.10.3 enable
RS G7028(config-ip-if)# vlan 2
RS G7028(config-ip-if)# exit
```

3. Configure STP. Reset the ports to make the edge configuration operational.

```
RS G7028(config)# interface port 5,6
RS G7028(config-if)# spanning-tree edge
RS G7028(config-if)# shutdown
RS G7028(config-if)# no shutdown
RS G7028(config-if)# exit
```

4. Configure an LACP dynamic trunk group (portchannel).

```
RS G7028(config)# interface port 1,2
RS G7028(config-if)# lacp key 400
RS G7028(config-if)# lacp mode active
RS G7028(config-if)# exit
```

5. Configure a static trunk group (portchannel).

```
RS G7028(config)# portchannel 1 port 3,4 enable
```

6. Configure IGMP Snooping.

```
RS G7028(config)# ip igmp enable
RS G7028(config)# ip igmp snoop vlan 2,3
RS G7028(config)# ip igmp snoop source-ip 10.10.10.3
RS G7028(config)# ip igmp snoop igmpv3 enable
RS G7028(config)# ip igmp snoop igmpv3 sources 64
RS G7028(config)# ip igmp snoop enable

RS G7028(config)# vlan 2
RS G7028(config-vlan)# no flood
RS G7028(config-vlan)# exit

RS G7028(config)# vlan 3
RS G7028(config-vlan)# no flood
RS G7028(config-vlan)# exit
```

## Troubleshooting

This section provides the steps to resolve common IGMP Snooping configuration issues. The topology described in [Figure 23](#) is used as an example.

### Multicast traffic from non-member groups reaches the host or Mrouter

- Check if traffic is unregistered. For unregistered traffic, an IGMP entry is not displayed in the IGMP groups table.

```
RS G7028# show ip igmp groups
```

- Ensure IPMC flooding is disabled and CPU is enabled.

```
RS G7028(config)# vlan <vlan id>
RS G7028(config-vlan)# no flood
RS G7028(config-vlan)# cpu
```

- Check the egress port's VLAN membership. The ports to which the hosts and Mrouter are connected must be used only for VLAN 2 and VLAN 3.

```
RS G7028# show vlan
```

**Note:** To avoid such a scenario, disable IPMC flooding for all VLANs enabled on the switches (if this is an acceptable configuration).

- Check IGMP Reports on switches B and C for information about the IGMP groups.

```
RS G7028# show ip igmp groups
```

If the non-member IGMP groups are displayed in the table, close the application that may be sending the IGMP Reports for these groups.

Identify the traffic source by using a sniffer on the hosts and reading the source IP/MAC address. If the source IP/MAC address is unknown, check the port statistics to find the ingress port.

```
RS G7028# show interface port <port id> interface-counters
```

- Ensure no static multicast MACs, static multicast groups, or static Mrouters are configured.
- Ensure IGMP Relay and PIM are not configured.

### Not all multicast traffic reaches the appropriate receivers.

- Ensure hosts are sending IGMP Reports for all the groups. Check the VLAN on which the groups are learned.

```
RS G7028# show ip igmp groups
```

If some of the groups are not displayed, ensure the multicast application is running on the host device and the generated IGMP Reports are correct.

- Ensure multicast traffic reaches the switch to which the host is connected.

Close the application sending the IGMP Reports. Clear the IGMP groups by flapping (disabling, then re-enabling) the port.

**Note:** To clear all IGMP groups, use the following command:

```
RS G7028(config)# clear ip igmp groups
```

However, this will clear all the IGMP groups and will influence other hosts.

Check if the multicast traffic reaches the switch.

```
RS G7028# show ip igmp ipmcgrp
```

If the multicast traffic group is not displayed in the table, check the link state, VLAN membership, and STP convergence.

- Ensure multicast server is sending all the multicast traffic.
- Ensure no static multicast MACs, static multicast groups, or static multicast routes are configured.

### IGMP queries sent by the Mrouter do not reach the host.

- Ensure the Mrouter is learned on switches B and C.

```
RS G7028# show ip igmp mrouter
```

If it is not learned on switch B but is learned on switch C, check the link state of the trunk group, VLAN membership, and STP convergence.

If it is not learned on any switch, ensure the multicast application is running and is sending correct IGMP Query packets.

If it is learned on both switches, check the link state, VLAN membership, and STP port states for the ports connected to the hosts.

### IGMP Reports/Leaves sent by the hosts do not reach the Mrouter

- Ensure IGMP Queries sent by the Mrouter reach the hosts.
- Ensure the Mrouter is learned on both switches. Note that the Mrouter may not be learned on switch B immediately after a trunk group failover/failback.

```
RS G7028# show ip igmp mrouter
```

- Ensure the host's multicast application is started and is sending correct IGMP Reports/Leaves.

```
RS G7028# show ip igmp groups
RS G7028# show ip igmp counters
```

**A host receives multicast traffic from the incorrect VLAN**

- Check port VLAN membership.
- Check IGMP Reports sent by the host.
- Check multicast data sent by the server.

**The Mrouter is learned on the incorrect trunk group**

- Check link state. Trunk group 1 might be down or in STP discarding state.
- Check STP convergence.
- Check port VLAN membership.

**Hosts receive multicast traffic at a lower rate than normal**

**Note:** This behavior is expected if IPMC flood is disabled and CPU is enabled. As soon as the IGMP/IPMC entries are installed on ASIC, the IPMC traffic recovers and is forwarded at line rate. This applies to unregistered IPMC traffic.

- Ensure a multicast threshold is not configured on the trunks.

```
RS G7028(config)# interface port <port id>  
RS G7028(config-ip)# no storm-control multicast
```

- Check link speeds and network congestion.

---

## Additional IGMP Features

The following topics are discussed in this section:

- [“FastLeave” on page 207](#)
- [“IGMP Filtering” on page 207](#)
- [“Static Multicast Router” on page 209](#)

### FastLeave

In normal IGMP operation, when the switch receives an IGMPv2 Leave message, it sends a Group-Specific Query to determine if any other devices in the same group (and on the same port) are still interested in the specified multicast group traffic. The switch removes the affiliated port from that particular group, if the switch does not receive an IGMP Membership Report within the query-response-interval.

With FastLeave enabled on the VLAN, a port can be removed immediately from the port list of the group entry when the IGMP Leave message is received.

**Note:** Only IGMPv2 supports FastLeave. Enable FastLeave on ports that have only one host connected. If more than one host is connected to a port, you may lose some hosts unexpectedly.

Use the following command to enable FastLeave.

### IGMP Filtering

With IGMP filtering, you can allow or deny certain IGMP groups to be learned on a port.

If access to a multicast group is denied, IGMP *Membership Reports* from the port are dropped, and the port is not allowed to receive IPv4 multicast traffic from that group. If access to the multicast group is allowed, Membership Reports from the port are forwarded for normal processing.

To configure IGMP filtering, you must globally enable IGMP filtering, define an IGMP filter, assign the filter to a port, and enable IGMP filtering on the port. To define an IGMP filter, you must configure a range of IPv4 multicast groups, choose whether the filter will allow or deny multicast traffic for groups within the range, and enable the filter.

## Configuring the Range

Each IGMP filter allows you to set a start and end point that defines the range of IPv4 addresses upon which the filter takes action. Each IPv4 address in the range must be between 224.0.0.0 and 239.255.255.255.

## Configuring the Action

Each IGMP filter can allow or deny IPv4 multicasts to the range of IPv4 addresses configured. If you configure the filter to deny IPv4 multicasts, then IGMP *Membership Reports* from multicast groups within the range are dropped. You can configure a secondary filter to allow IPv4 multicasts to a small range of addresses within a larger range that a primary filter is configured to deny. The two filters work together to allow IPv4 multicasts to a small subset of addresses within the larger range of addresses.

**Note:** Lower-numbered filters take precedence over higher-number filters. For example, the action defined for IGMP filter 1 supersedes the action defined for IGMP filter 2.

## Configure IGMP Filtering

1. Enable IGMP filtering on the switch.

```
RS G7028(config)# ip igmp filtering
```

2. Define an IGMP filter with IPv4 information.

```
RS G7028(config)# ip igmp profile 1 range 224.0.0.0 226.0.0.0
RS G7028(config)# ip igmp profile 1 action deny
RS G7028(config)# ip igmp profile 1 enable
```

3. Assign the IGMP filter to a port.

```
RS G7028(config)# interface port 3
RS G7028(config-if)# ip igmp profile 1
RS G7028(config-if)# ip igmp filtering
```

## Static Multicast Router

A static Mrouter can be configured for a particular port on a particular VLAN. A static Mrouter does not have to be learned through IGMP Snooping. Any data port can accept a static Mrouter.

When you configure a static Mrouter on a VLAN, it replaces any dynamic Mrouters learned through IGMP Snooping.

### Configure a Static Multicast Router

1. For each Mrouter, configure a port, VLAN, and IGMP version.

```
RS G7028(config)# ip igmp mrouter 5 1 2
```

The IGMP version is set for each VLAN, and cannot be configured separately for each Mrouter.

2. Verify the configuration.

```
RS G7028(config)# show ip igmp mrouter
```



---

## Chapter 16. Basic Redundancy

IBM Networking OS 7.6 includes various features for providing basic link or device redundancy:

- [“Trunking for Link Redundancy” on page 212](#)
- [“Virtual Link Aggregation” on page 212](#)
- [“Hot Links” on page 213](#)

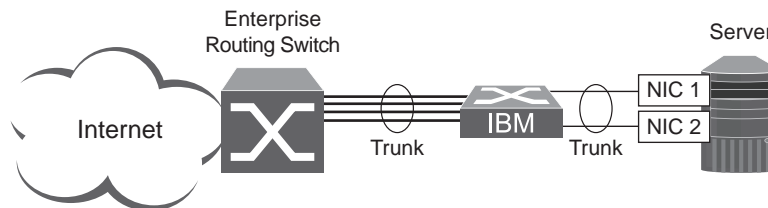
---

## Trunking for Link Redundancy

Multiple switch ports can be combined together to form robust, high-bandwidth trunks to other devices. Since trunks are comprised of multiple physical links, the trunk group is inherently fault tolerant. As long as one connection between the switches is available, the trunk remains active.

In [Figure 24](#), four ports are trunked together between the switch and the enterprise routing device. Connectivity is maintained as long as one of the links remain active. The links to the server are also trunked, allowing the secondary NIC to take over in the event that the primary NIC link fails.

Figure 24. Trunking Ports for Link Redundancy



For more information on trunking, see [“Ports and Trunking” on page 115](#).

---

## Virtual Link Aggregation

Using the VLAG feature, switches can be paired as VLAG peers. The peer switches appear to the connecting device as a single virtual entity for the purpose of establishing a multi-port trunk. The VLAG-capable switches synchronize their logical view of the access layer port structure and internally prevent implicit loops. The VLAG topology also responds more quickly to link failure and does not result in unnecessary MAC flooding.

VLAGs are useful in multi-layer environments for both uplink and downlink redundancy to any regular LAG-capable device. They can also be used in for active-active VRRP connections.

For more information on VLAGs, see [“Virtual Link Aggregation Groups” on page 147](#).

---

## Hot Links

For network topologies that require Spanning Tree to be turned off, Hot Links provides basic link redundancy with fast recovery.

Hot Links consists of up to 25 triggers. A trigger consists of a pair of layer 2 interfaces, each containing an individual port, trunk, or LACP adminkey. One interface is the Master, and the other is a Backup. While the Master interface is set to the active state and forwards traffic, the Backup interface is set to the standby state and blocks traffic until the Master interface fails. If the Master interface fails, the Backup interface is set to active and forwards traffic. Once the Master interface is restored, it regains its active state and forwards traffic. The Backup transitions to standby state.

You may select a physical port, static trunk, or an LACP adminkey as a Hot Link interface.

## Forward Delay

The Forward Delay timer allows Hot Links to monitor the Master and Backup interfaces for link stability before selecting one interface to transition to the active state. Before the transition occurs, the interface must maintain a stable link for the duration of the Forward Delay interval.

For example, if you set the Forward delay timer to 10 seconds, the switch will select an interface to become active only if a link remained stable for the duration of the Forward Delay period. If the link is unstable, the Forward Delay period starts again.

## Preemption

You can configure the Master interface to resume the active state whenever it becomes available. With Hot Links preemption enabled, the Master interface transitions to the active state immediately upon recovery. The Backup interface immediately transitions to the standby state. If Forward Delay is enabled, the transition occurs when an interface has maintained link stability for the duration of the Forward Delay period.

## FDB Update

Use the FDB update option to notify other devices on the network about updates to the Forwarding Database (FDB). When you enable FDB update, the switch sends multicasts of addresses in the forwarding database (FDB) over the active interface, so that other devices on the network can learn the new path. The Hot Links FDB update option uses the station update rate to determine the rate at which to send FDB packets.

## Configuration Guidelines

The following configuration guidelines apply to Hot links:

- When Hot Links is turned on, MSTP, RSTP, and PVRST must be turned off.
- A port that is a member of the Master interface cannot be a member of the Backup interface. A port that is a member of one Hot Links trigger cannot be a member of another Hot Links trigger.
- An individual port that is configured as a Hot Link interface cannot be a member of a trunk.

## Configuring Hot Links

Use the following commands to configure Hot Links.

```
RS G7028(config)# hotlinks trigger 1 enable  (Enable Hot Links Trigger 1)  
RS G7028(config)# hotlinks trigger 1 master port 1 (Add port to Master interface)  
RS G7028(config)# hotlinks trigger 1 backup port 2 (Add port to Backup interface)  
RS G7028(config)# hotlinks enable           (Turn on Hot Links)
```

---

## Chapter 17. Layer 2 Failover

The primary application for Layer 2 Failover is to support Network Adapter Teaming. With Network Adapter Teaming, all the NICs on each server share the same IP address, and are configured into a team. One NIC is the primary link, and the other is a standby link. For more details, refer to the documentation for your Ethernet adapter.

**Note:** Only two links per server can be used for Layer 2 Trunk Failover (one primary and one backup). Network Adapter Teaming allows only one backup NIC for each server blade.

---

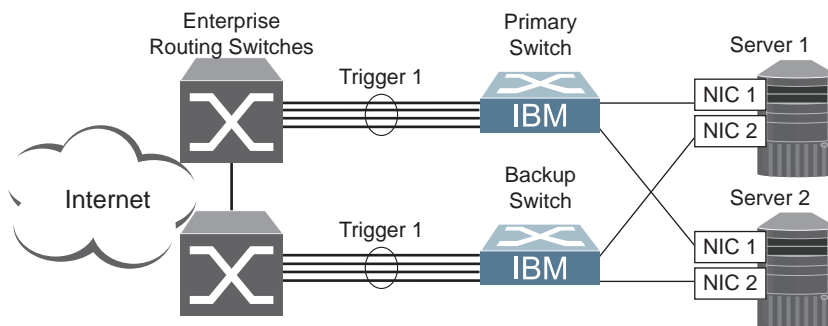
### Monitoring Trunk Links

Layer 2 Failover can be enabled on any trunk group in the G7028/G7052, including LACP trunks. Trunks can be added to failover trigger groups. Then, if some specified number of monitor links fail, the switch disables all the control ports in the switch. When the control ports are disabled, it causes the NIC team on the affected servers to failover from the primary to the backup NIC. This process is called a failover event.

When the appropriate number of links in a monitor group return to service, the switch enables the control ports. This causes the NIC team on the affected servers to fail back to the primary switch (unless Auto-Fallback is disabled on the NIC team). The backup switch processes traffic until the primary switch's control links come up, which can take up to five seconds.

Figure 25 is a simple example of Layer 2 Failover. One G7028/G7052 is the primary, and the other is used as a backup. In this example, all ports on the primary switch belong to a single trunk group, with Layer 2 Failover enabled, and Failover Limit set to 2. If two or fewer links in trigger 1 remain active, the switch temporarily disables all control ports. This action causes a failover event on Server 1 and Server 2.

Figure 25. Basic Layer 2 Failover



---

### Setting the Failover Limit

The failover limit lets you specify the minimum number of operational links required within each trigger before the trigger initiates a failover event. For example, if the limit is two, a failover event occurs when the number of operational links in the trigger is two or fewer. When you set the limit to zero, the switch triggers a failover event only when no links in the trigger are operational.

---

## Manually Monitoring Port Links

The Manual Monitor allows you to configure a set of ports and trunks to monitor for link failures (a monitor list), and another set of ports and trunks to disable when the trigger limit is reached (a control list). When the switch detects a link failure on the monitor list, it automatically disables the items in control list. When server ports are disabled, the corresponding server's network adapter can detect the disabled link, and trigger a network-adapter failover to another port or trunk on the switch, or another switch.

The switch automatically enables the control list items when the monitor list items return to service.

### Monitor Port State

A monitor port is considered operational as long as the following conditions are true:

- The port must be in the `Link Up` state.
- If STP is enabled, the port must be in the `Forwarding` state.
- If the port is part of an LACP trunk, the port must be in the `Aggregated` state.

If any of these conditions is false, the monitor port is considered to have failed.

### Control Port State

A control port is considered Operational if the monitor trigger is up. As long as the trigger is up, the port is considered operational from a teaming perspective, even if the port itself is actually in the `Down` state, `Blocking` state (if STP is enabled on the port), or `Not Aggregated` state (if part of an LACP trunk).

A control port is considered to have failed only if the monitor trigger is in the `Down` state.

To view the state of any port, use one of the following commands:

|                                                    |                         |
|----------------------------------------------------|-------------------------|
| >> # show interface link                           | (View port link status) |
| >> # show interface port <x> spanning-tree stp <x> | (View port STP status)  |
| >> # show lacp information                         | (View port LACP status) |

---

## L2 Failover with Other Features

L2 Failover works together with Link Aggregation Control Protocol (LACP) and with Spanning Tree Protocol (STP), as described in the next sections.

### LACP

Link Aggregation Control Protocol allows the switch to form dynamic trunks using automatic monitoring. When you add an *admin key* to a trigger, any LACP trunk with that *admin key* becomes a member of the trigger.

### Spanning Tree Protocol

If Spanning Tree Protocol (STP) is enabled on the ports in a failover trigger, the switch monitors the port STP state rather than the link state. A port failure results when STP is not in a Forwarding state (such as Learning, Discarding, or No Link) in all the Spanning Tree Groups (STGs) to which the port belongs. The switch automatically disables the appropriate control ports.

When the switch determines that ports in the trigger are in STP Forwarding state in any one of the STGs it belongs to, then it automatically enables the appropriate control ports. The switch *fails back* to normal operation.

For example, if a monitor port is a member of STG1, STG2, and STG3, a failover will be triggered only if the port is not in a forwarding state in all the three STGs. When the port state in any of the three STGs changes to forwarding, then the control port is enabled and normal switch operation is resumed.

---

## Configuration Guidelines

This section provides important information about configuring Layer 2 Failover.

- Any specific failover trigger can monitor ports only, static trunks only, or LACP trunks only. The different types cannot be combined in the same trigger.
- A maximum of 52 LACP keys can be added per trigger.
- Port membership for different triggers must not overlap. Any specific port must be a member of only one trigger.

---

## Configuring Layer 2 Failover

Use the following procedure to configure a Layer 2 Failover Manual Monitor.

1. Specify the links to monitor.

```
>> # failover trigger 1 mmon monitor member 1-5
```

2. Specify the links to disable when the failover limit is reached.

```
>> # failover trigger 1 mmon control member 6-10
```

3. Configure general Failover parameters.

```
>> # failover enable  
>> # failover trigger 1 enable  
>> # failover trigger 1 limit 2
```

# **Part 5: Network Management**



---

## Chapter 18. Link Layer Discovery Protocol

The IBM Networking OS software support Link Layer Discovery Protocol (LLDP). This chapter discusses the use and configuration of LLDP on the switch:

- [“LLDP Overview” on page 222](#)
- [“Enabling or Disabling LLDP” on page 223](#)
- [“LLDP Transmit Features” on page 224](#)
- [“LLDP Receive Features” on page 228](#)
- [“LLDP Example Configuration” on page 232](#)

---

## LLDP Overview

Link Layer Discovery Protocol (LLDP) is an IEEE 802.1AB-2005 standard for discovering and managing network devices. LLDP uses Layer 2 (the data link layer), and allows network management applications to extend their awareness of the network by discovering devices that are direct neighbors of already known devices.

With LLDP, the G7028/G7052 can advertise the presence of its ports, their major capabilities, and their current status to other LLDP stations in the same LAN. LLDP transmissions occur on ports at regular intervals or whenever there is a relevant change to their status. The switch can also receive LLDP information advertised from adjacent LLDP-capable network devices.

In addition to discovery of network resources, and notification of network changes, LLDP can help administrators quickly recognize a variety of common network configuration problems, such as unintended VLAN exclusions or mis-matched port aggregation membership.

The LLDP transmit function and receive function can be independently configured on a per-port basis. The administrator can allow any given port to transmit only, receive only, or both transmit and receive LLDP information.

The LLDP information to be distributed by the G7028/G7052 ports, and that which has been collected from other LLDP stations, is stored in the switch's Management Information Base (MIB). Network Management Systems (NMS) can use Simple Network Management Protocol (SNMP) to access this MIB information. LLDP-related MIB information is read-only.

Changes, either to the local switch LLDP information or to the remotely received LLDP information, are flagged within the MIB for convenient tracking by SNMP-based management systems.

For LLDP to provide expected benefits, all network devices that support LLDP must be consistent in their LLDP configuration.

---

## Enabling or Disabling LLDP

### Global LLDP Setting

By default, LLDP is enabled on the G7028/G7052. To turn LLDP on or off, use the following command:

|                                    |                                |
|------------------------------------|--------------------------------|
| RS G7028(config)# [no] lldp enable | (Turn LLDP on or off globally) |
|------------------------------------|--------------------------------|

### Transmit and Receive Control

The G7028/G7052 can also be configured to transmit or receive LLDP information on a port-by-port basis. By default, when LLDP is globally enabled on the switch, G7028/G7052 ports transmit and receive LLDP information (see the `tx_rx` option in the following example). To change the LLDP transmit and receive state, the following commands are available:

|                                                |                              |
|------------------------------------------------|------------------------------|
| RS G7028(config)# interface port 1             | (Select a switch port)       |
| RS G7028(config-if)# lldp admin-status tx_rx   | (Transmit and receive LLDP)  |
| RS G7028(config-if)# lldp admin-status tx_only | (Only transmit LLDP)         |
| RS G7028(config-if)# lldp admin-status rx_only | (Only receive LLDP)          |
| RS G7028(config-if)# no lldp admin-status      | (Do not participate in LLDP) |
| RS G7028(config-if)# exit                      | (Exit port mode)             |

To view the LLDP transmit and receive status, use the following commands:

|                                                |                           |
|------------------------------------------------|---------------------------|
| RS G7028(config)# show lldp port               | (status of all ports)     |
| RS G7028(config)# show interface port <n> lldp | (status of selected port) |

---

## LLDP Transmit Features

Numerous LLDP transmit options are available, including scheduled and minimum transmit interval, expiration on remote systems, SNMP trap notification, and the types of information permitted to be shared.

### Scheduled Interval

The G7028/G7052 can be configured to transmit LLDP information to neighboring devices once each 5 to 32768 seconds. The scheduled interval is global; the same interval value applies to all LLDP transmit-enabled ports. However, to help balance LLDP transmissions and keep them from being sent simultaneously on all ports, each port maintains its own interval clock, based on its own initialization or reset time. This allows switch-wide LLDP transmissions to be spread out over time, though individual ports comply with the configured interval.

The global transmit interval can be configured using the following command:

```
RS G7028(config)# lldp refresh-interval <interval>
```

where *interval* is the number of seconds between LLDP transmissions. The range is 5 to 32768. The default is 30 seconds.

### Minimum Interval

In addition to sending LLDP information at scheduled intervals, LLDP information is also sent when the G7028/G7052 detects relevant changes to its configuration or status (such as when ports are enabled or disabled). To prevent the G7028/G7052 from sending multiple LLDP packets in rapid succession when port status is in flux, a transmit delay timer can be configured.

The transmit delay timer represents the minimum time permitted between successive LLDP transmissions on a port. Any interval-driven or change-driven updates will be consolidated until the configured transmit delay expires.

The minimum transmit interval can be configured using the following command:

```
RS G7028(config)# lldp transmission-delay <interval>
```

where *interval* is the minimum number of seconds permitted between successive LLDP transmissions on any port. The range is 1 to one-quarter of the scheduled transmit interval (`lldp refresh-interval <value>`), up to 8192. The default is 2 seconds.

## Time-to-Live for Transmitted Information

The transmitted LLDP information is held by remote systems for a limited time. A time-to-live parameter allows the switch to determine how long the transmitted data is held before it expires. The hold time is configured as a multiple of the configured transmission interval.

```
RS G7028(config)# lldp holdtime-multiplier <multiplier>
```

where *multiplier* is a value between 2 and 10. The default value is 4, meaning that remote systems will hold the port's LLDP information for 4 x the 30-second `msgtxint` value, or 120 seconds, before removing it from their MIB.

## Trap Notifications

If SNMP is enabled on the G7028/G7052 (see [“Using Simple Network Management Protocol” on page 24](#)), each port can be configured to send SNMP trap notifications whenever LLDP transmissions are sent. By default, trap notification is disabled for each port. The trap notification state can be changed using the following commands (Interface Port mode):

```
RS G7028(config)# interface port 1
RS G7028(config-if)# [no] lldp trap-notification
RS G7028(config-if)# exit
```

In addition to sending LLDP information at scheduled intervals, LLDP information is also sent when the G7028/G7052 detects relevant changes to its configuration or status (such as when ports are enabled or disabled). To prevent the G7028/G7052 from sending multiple trap notifications in rapid succession when port status is in flux, a global trap delay timer can be configured.

The trap delay timer represents the minimum time permitted between successive trap notifications on any port. Any interval-driven or change-driven trap notices from the port will be consolidated until the configured trap delay expires.

The minimum trap notification interval can be configured using the following command:

```
RS G7028(config)# lldp trap-notification-interval <interval>
```

where *interval* is the minimum number of seconds permitted between successive LLDP transmissions on any port. The range is 1 to 3600. The default is 5 seconds.

If SNMP trap notification is enabled, the notification messages can also appear in the system log. This is enabled by default. To change whether the SNMP trap notifications for LLDP events appear in the system log, use the following command:

```
RS G7028(config)# [no] logging log lldp
```

## Changing the LLDP Transmit State

When the port is disabled, or when LLDP transmit is turned off for the port using the LLDP admin-status command options (see [“Transmit and Receive Control” on page 223](#)), a final LLDP packet is transmitted with a time-to-live value of 0. Neighbors that receive this packet will remove the LLDP information associated with the G7028/G7052 port from their MIB.

In addition, if LLDP is fully disabled on a port and then later re-enabled, the G7028/G7052 will temporarily delay resuming LLDP transmissions on the port to allow the port LLDP information to stabilize. The reinitialization delay interval can be globally configured for all ports using the following command:

```
RS G7028(config)# lldp reinit-delay <interval>
```

where *interval* is the number of seconds to wait before resuming LLDP transmissions. The range is between 1 and 10. The default is 2 seconds.

## Types of Information Transmitted

When LLDP transmission is permitted on the port (see [“Enabling or Disabling LLDP” on page 223](#)), the port advertises the following required information in type/length/value (TLV) format:

- Chassis ID
- Port ID
- LLDP Time-to-Live

LLDP transmissions can also be configured to enable or disable inclusion of optional information, using the following command (Interface Port mode):

```
RS G7028(config)# interface port 1
RS G7028(config-if)# [no] lldp tlv <type>
RS G7028(config-if)# exit
```

where *type* is an LLDP information option from [Table 18](#):

Table 18. LLDP Optional Information Types

| Type     | Description                          | Default  |
|----------|--------------------------------------|----------|
| portdesc | Port Description                     | Enabled  |
| sysname  | System Name                          | Enabled  |
| sysdescr | System Description                   | Enabled  |
| syscap   | System Capabilities                  | Enabled  |
| mgmtaddr | Management Address                   | Enabled  |
| portvid  | IEEE 802.1 Port VLAN ID              | Disabled |
| portprot | IEEE 802.1 Port and Protocol VLAN ID | Disabled |
| vlanname | IEEE 802.1 VLAN Name                 | Disabled |
| protid   | IEEE 802.1 Protocol Identity         | Disabled |

Table 18. LLDP Optional Information Types (continued)

| Type     | Description                                                                                                                               | Default  |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------|----------|
| macphy   | IEEE 802.3 MAC/PHY Configuration/Status, including the auto-negotiation, duplex, and speed status of the port.                            | Disabled |
| powermdi | IEEE 802.3 Power via MDI, indicating the capabilities and status of devices that require or provide power over twisted-pair copper links. | Disabled |
| linkaggr | IEEE 802.3 Link Aggregation status for the port.                                                                                          | Disabled |
| framesz  | IEEE 802.3 Maximum Frame Size for the port.                                                                                               | Disabled |
| all      | Select all optional LLDP information for inclusion or exclusion.                                                                          | Disabled |

---

## LLDP Receive Features

### Types of Information Received

When the LLDP receive option is enabled on a port (see [“Enabling or Disabling LLDP” on page 223](#)), the port may receive the following information from LLDP-capable remote systems:

- Chassis Information
- Port Information
- LLDP Time-to-Live
- Port Description
- System Name
- System Description
- System Capabilities Supported/Enabled
- Remote Management Address

The G7028/G7052 stores the collected LLDP information in the MIB. Each remote LLDP-capable device is responsible for transmitting regular LLDP updates. If the received updates contain LLDP information changes (to port state, configuration, LLDP MIB structures, deletion), the switch will set a change flag within the MIB for convenient notification to SNMP-based management systems.

### Viewing Remote Device Information

LLDP information collected from neighboring systems can be viewed in numerous ways:

- Using a centrally-connected LLDP analysis server
- Using an SNMP agent to examine the G7028/G7052 MIB
- Using the G7028/G7052 Browser-Based Interface (BBI)
- Using CLI or isCLI commands on the G7028/G7052

Using the CLI the following command displays remote LLDP information:

```
RS G7028(config)# show lldp remote-device [<index number>]
```

To view a summary of remote information, omit the *Index number* parameter. For example:

```
RS G7028(config)# show lldp remote-device
LLDP Remote Devices Information
```

| LocalPort | Index | Remote Chassis ID | Remote Port | Remote System Name |
|-----------|-------|-------------------|-------------|--------------------|
| 3         | 1     | 00 18 b1 33 1d 00 | 23          |                    |

To view detailed information for a remote device, specify the *Index number* as found in the summary. For example, in keeping with the sample summary, to list details for the first remote device (with an Index value of 1), use the following command:

```
RS G7028(config)# show lldp remote-device 1
Local Port Alias: 3
    Remote Device Index      : 1
    Remote Device TTL        : 99
    Remote Device RxChanges  : false
    Chassis Type              : Mac Address
    Chassis Id                : 00-18-b1-33-1d-00
    Port Type                 : Locally Assigned
    Port Id                   : 23
    Port Description          : 7

    System Name               :
    System Description: IBM Networking Operating System IBM System Networking
RackSwitch G7028/G7052, IBM Networking OS: version 7.6, Boot Image: version
6.9.1.14

    System Capabilities Supported : bridge, router
    System Capabilities Enabled   : bridge, router

    Remote Management Address:
        Subtype                  : IPv4
        Address                   : 10.100.120.181
        Interface Subtype        : ifIndex
        Interface Number         : 128
        Object Identifier        :
```

**Note:** Received LLDP information can change very quickly. When using show commands, it is possible that flags for some expected events may be too short-lived to be observed in the output.

To view detailed information of all remote devices, use the following command:

```
RS G7028# show lldp remote-device detail

Local Port Alias: MGTA
  Remote Device Index      : 1
  Remote Device TTL        : 4678
  Remote Device RxChanges  : false
  Chassis Type             : Mac Address
  Chassis Id               : 08-17-f4-a1-db-00
  Port Type               : Locally Assigned
  Port Id                 : 25
  Port Description        : MGTA

  System Name              :
  System Description       : IBM Networking Operating System IBM
System Networking RackSwitch G7028/G7052, IBM Networking OS: version 7.6, Boot
Image: version 6.9.1.14
  System Capabilities Supported : bridge, router
  System Capabilities Enabled  : bridge, router

  Remote Management Address:
    Subtype                 : IPv4
    Address                 : 10.38.22.23
    Interface Subtype       : ifIndex
    Interface Number        : 127
    Object Identifier       :

Local Port Alias: 2
  Remote Device Index      : 2
  Remote Device TTL        : 4651
  Remote Device RxChanges  : false
  Chassis Type             : Mac Address
  Chassis Id               : 08-17-f4-a1-db-00
  Port Type               : Locally Assigned
  Port Id                 : 2
  Port Description        : 2

  System Name              :
  System Description       : IBM Networking Operating System IBM
System Networking RackSwitch G7028/G7052, IBM Networking OS: version 7.6, Boot
Image: version 6.9.1.14
  System Capabilities Supported : bridge, router
  System Capabilities Enabled  : bridge, router

  Remote Management Address:
    Subtype                 : IPv4
    Address                 : 10.38.22.23
    Interface Subtype       : ifIndex
    Interface Number        : 127
    Object Identifier       :

Total entries displayed: 2
```

## Time-to-Live for Received Information

Each remote device LLDP packet includes an expiration time. If the switch port does not receive an LLDP update from the remote device before the time-to-live clock expires, the switch will consider the remote information to be invalid, and will remove all associated information from the MIB.

Remote devices can also intentionally set their LLDP time-to-live to 0, indicating to the switch that the LLDP information is invalid and must be immediately removed.

---

## LLDP Example Configuration

1. Turn LLDP on globally.

```
RS G7028(config)# lldp enable
```

2. Set the global LLDP timer features.

```
RS G7028(config)# lldp refresh-interval 30 (Transmit each 30 seconds)
RS G7028(config)# lldp transmission-delay 2 (No more often than 2 sec.)
RS G7028(config)# lldp holdtime-multiplier 4 (Remote hold 4 intervals)
RS G7028(config)# lldp reinit-delay 2 (Wait 2 sec. after reinit.)
RS G7028(config)# lldp trap-notification-interval 5 (Minimum 5 sec. between)
```

3. Set LLDP options for each port.

```
RS G7028(config)# interface port <n> (Select a switch port)
RS G7028(config-if)# lldp admin-status tx_rx (Transmit and receive LLDP)
RS G7028(config-if)# lldp trap-notification (Enable SNMP trap notifications)
RS G7028(config-if)# lldp tlv all (Transmit all optional information)
RS G7028(config-if)# exit
```

4. Enable syslog reporting.

```
RS G7028(config)# logging log lldp
```

5. Verify the configuration settings:

```
RS G7028(config)# show lldp
```

6. View remote device information as needed.

```
RS G7028(config)# show lldp remote-device
or
RS G7028(config)# show lldp remote-device <index number>
or
RS G7028(config)# show lldp remote-device detail
```

---

## Chapter 19. Simple Network Management Protocol

IBM Networking OS provides Simple Network Management Protocol (SNMP) version 1, version 2, and version 3 support for access through any network management software, such as IBM Director or HP-OpenView.

**Note:** SNMP read and write functions are enabled by default. For best security practices, if SNMP is not needed for your network, it is recommended that you disable these functions prior to connecting the switch to the network.

---

## SNMP Version 1 & Version 2

To access the SNMP agent on the G7028/G7052, the read and write community strings on the SNMP manager must be configured to match those on the switch. The default read community string on the switch is `public` and the default write community string is `private`.

The read and write community strings on the switch can be changed using the following commands on the CLI:

```
RS G7028(config)# snmp-server read-community <1-32 characters>  
  
-and-  
  
RS G7028(config)# snmp-server write-community <1-32 characters>
```

The SNMP manager must be able to reach the management interface or any one of the IP interfaces on the switch.

For the SNMP manager to receive the SNMPv1 traps sent out by the SNMP agent on the switch, configure the trap host on the switch with the following command:

```
RS G7028(config)# snmp-server trap-source <trap source IP interface>  
RS G7028(config)# snmp-server host <IPv4 address> <trap host community string>
```

**Note:** You can use a loopback interface to set the source IP address for SNMP traps. Use the following command to apply a configured loopback interface:

```
RS G7028(config)# snmp-server trap-source loopback <1-5>
```

---

## SNMP Version 3

SNMP version 3 (SNMPv3) is an enhanced version of the Simple Network Management Protocol, approved by the Internet Engineering Steering Group in March, 2002. SNMPv3 contains additional security and authentication features that provide data origin authentication, data integrity checks, timeliness indicators and encryption to protect against threats such as masquerade, modification of information, message stream modification and disclosure.

SNMPv3 allows clients to query the MIBs securely.

SNMPv3 configuration is managed using the following command path menu:

```
RS G7028(config)# snmp-server ?
```

For more information on SNMP MIBs and the commands used to configure SNMP on the switch, see the *IBM Networking OS 7.6 Command Reference*.

### Default Configuration

IBM N/OS has two SNMPv3 users by default. Both of the following users have access to all the MIBs supported by the switch:

- User 1 name is `adminmd5` (password `adminmd5`). Authentication used is MD5.
- User 2 name is `adminsha` (password `adminsha`). Authentication used is SHA.

Up to 16 SNMP users can be configured on the switch. To modify an SNMP user, enter the following commands:

```
RS G7028(config)# snmp-server user <1-16> name <1-32 characters>
```

Users can be configured to use the authentication/privacy options. The G7028/G7052 support two authentication algorithms: MD5 and SHA, as specified in the following command:

```
RS G7028(config)# snmp-server user <1-16> authentication-protocol {md5|sha}  
authentication-password  
-or-  
RS G7028(config)# snmp-server user <1-16> authentication-protocol none
```

## User Configuration Example

1. To configure a user with name “admin,” authentication type MD5, and authentication password of “admin,” privacy option DES with privacy password of “admin,” use the following CLI commands.

```
RS G7028(config)# snmp-server user 5 name admin
RS G7028(config)# snmp-server user 5 authentication-protocol md5
authentication-password
Changing authentication password; validation required:
Enter current admin password:          <admin. password>
Enter new authentication password:     <auth. password>
Re-enter new authentication password:  <auth. password>
New authentication password accepted.

RS G7028(config)# snmp-server user 5 privacy-protocol des privacy-password
Changing privacy password; validation required:
Enter current admin password:          <admin. password>
Enter new privacy password:            <privacy password>
Re-enter new privacy password:         <privacy password>
New privacy password accepted.
```

2. Configure a user access group, along with the views the group may access. Use the access table to configure the group’s access level.

```
RS G7028(config)# snmp-server access 5 name admingrp
RS G7028(config)# snmp-server access 5 level authpriv
RS G7028(config)# snmp-server access 5 read-view iso
RS G7028(config)# snmp-server access 5 write-view iso
RS G7028(config)# snmp-server access 5 notify-view iso
```

Because the read view, write view, and notify view are all set to “iso,” the user type has access to all private and public MIBs.

3. Assign the user to the user group. Use the group table to link the user to a particular access group.

```
RS G7028(config)# snmp-server group 5 user-name admin
RS G7028(config)# snmp-server group 5 group-name admingrp
```

---

## Configuring SNMP Trap Hosts

### SNMPv1 Trap Host

1. Configure a user with no authentication and password.

```
RS G7028(config)# snmp-server user 10 name v1trap
```

2. Configure an access group and group table entries for the user. Use the following menu to specify which traps can be received by the user:

```
RS G7028(config)# snmp-server access <user number>
```

In the following example the user will receive the traps sent by the switch.

```
RS G7028(config)# snmp-server access 10          (Access group to view SNMPv1 traps)
                    name v1trap
                    security snmpv1
                    notify-view iso
RS G7028(config)# snmp-server group 10          (Assign user to the access group)
                    security snmpv1
                    user-name v1trap
                    group-name v1trap
```

3. Configure an entry in the notify table.

```
RS G7028(config)# snmp-server notify 10 name v1trap
RS G7028(config)# snmp-server notify 10 tag v1trap
```

4. Specify the IPv4 address and other trap parameters in the `targetAddr` and `targetParam` tables. Use the following commands to specify the user name associated with the `targetParam` table:

```
RS G7028(config)# snmp-server target-address 10 name v1trap address 10.70.70.190
RS G7028(config)# snmp-server target-address 10 parameters-name vlparam
RS G7028(config)# snmp-server target-address 10 taglist vlparam
RS G7028(config)# snmp-server target-parameters 10 name vlparam
RS G7028(config)# snmp-server target-parameters 10 user-name vlonly
RS G7028(config)# snmp-server target-parameters 10 message snmpv1
```

**Note:** N/OS 7.6 supports only IPv4 addresses for SNMP trap hosts.

5. Use the community table to specify which community string is used in the trap.

```
RS G7028(config)# snmp-server community 10 (Define the community string)
                    index v1trap
                    name public
                    user-name v1trap
```

## SNMPv2 Trap Host Configuration

The SNMPv2 trap host configuration is similar to the SNMPv1 trap host configuration. Wherever you specify the model, use `snmpv2` instead of `snmpv1`.

```
RS G7028(config)# snmp-server user 10 name v2trap

RS G7028(config)# snmp-server group 10 security snmpv2
RS G7028(config)# snmp-server group 10 user-name v2trap
RS G7028(config)# snmp-server group 10 group-name v2trap
RS G7028(config)# snmp-server access 10 name v2trap
RS G7028(config)# snmp-server access 10 security snmpv2
RS G7028(config)# snmp-server access 10 notify-view iso

RS G7028(config)# snmp-server notify 10 name v2trap
RS G7028(config)# snmp-server notify 10 tag v2trap

RS G7028(config)# snmp-server target-address 10 name v2trap
address 100.10.2.1
RS G7028(config)# snmp-server target-address 10 taglist v2trap
RS G7028(config)# snmp-server target-address 10 parameters-name
v2param
RS G7028(config)# snmp-server target-parameters 10 name v2param
RS G7028(config)# snmp-server target-parameters 10 message snmpv2c
RS G7028(config)# snmp-server target-parameters 10 user-name v2trap
RS G7028(config)# snmp-server target-parameters 10 security snmpv2

RS G7028(config)# snmp-server community 10 index v2trap
RS G7028(config)# snmp-server community 10 user-name v2trap
```

**Note:** N/OS 7.6 supports only IPv4 addresses for SNMP trap hosts.

## SNMPv3 Trap Host Configuration

To configure a user for SNMPv3 traps, you can choose to send the traps with both privacy and authentication, with authentication only, or without privacy or authentication.

This is configured in the access table using the following commands:

```
RS G7028(config)# snmp-server access <1-32> level
RS G7028(config)# snmp-server target-parameters <1-16>
```

Configure the user in the user table accordingly.

It is not necessary to configure the community table for SNMPv3 traps because the community string is not used by SNMPv3.

The following example shows how to configure a SNMPv3 user v3trap with authentication only:

```
RS G7028(config)# snmp-server user 11 name v3trap
RS G7028(config)# snmp-server user 11 authentication-protocol md5
authentication-password
Changing authentication password; validation required:
Enter current admin password: <admin.password>
Enter new authentication password: <auth.password>
Re-enter new authentication password: <auth.password>
New authentication password accepted.
RS G7028(config)# snmp-server access 11 notify-view iso
RS G7028(config)# snmp-server access 11 level authnopriv
RS G7028(config)# snmp-server group 11 user-name v3trap
RS G7028(config)# snmp-server group 11 tag v3trap
RS G7028(config)# snmp-server notify 11 name v3trap
RS G7028(config)# snmp-server notify 11 tag v3trap
RS G7028(config)# snmp-server target-address 11 name v3trap address 47.81.25.66
RS G7028(config)# snmp-server target-address 11 taglist v3trap
RS G7028(config)# snmp-server target-address 11 parameters-name v3param
RS G7028(config)# snmp-server target-parameters 11 name v3param
RS G7028(config)# snmp-server target-parameters 11 user-name v3trap
RS G7028(config)# snmp-server target-parameters 11 level authNoPriv
```

**Note:** N/OS 7.6 supports only IPv4 addresses for SNMP trap hosts.

---

## SNMP MIBs

The N/OS SNMP agent supports SNMP version 3. Security is provided through SNMP community strings. The default community strings are “public” for SNMP GET operation and “private” for SNMP SET operation. The community string can be modified only through the Command Line Interface (CLI). Detailed SNMP MIBs and trap definitions of the N/OS SNMP agent are contained in the N/OS enterprise MIB document.

The N/OS SNMP agent supports the following standard MIBs:

- rfc1213.mib
- rfc1215.mib
- rfc1493.mib
- rfc1573.mib
- rfc1643.mib
- rfc1757.mib
- rfc1907.mib
- rfc2571.mib
- rfc2572.mib
- rfc2573.mib
- rfc2574.mib
- rfc2575.mib
- rfc2576.mib
- rfc2790.mib
- rfc4001.mib
- rfc4133.mib

The N/OS SNMP agent supports the following generic traps as defined in RFC 1215:

- ColdStart
- WarmStart
- LinkDown
- LinkUp
- AuthenticationFailure

The SNMP agent also supports two Spanning Tree traps as defined in RFC 1493:

- NewRoot
- TopologyChange

The following are the enterprise SNMP traps supported in N/OS:

*Table 19. IBM N/OS-Supported Enterprise SNMP Traps*

| Trap Name        | Description                                             |
|------------------|---------------------------------------------------------|
| SwDefGwUp        | Signifies that the default gateway is alive.            |
| SwDefGwDown      | Signifies that the default gateway is down.             |
| SwDefGwInService | Signifies that the default gateway is up and in service |

Table 19. IBM N/OS-Supported Enterprise SNMP Traps (continued)

| Trap Name             | Description                                                                                                                                                                                                              |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SwDefGwNotInService   | Signifies that the default gateway is alive but not in service                                                                                                                                                           |
| SwVrrpNewMaster       | Indicates that the sending agent has transitioned to "Master" state.                                                                                                                                                     |
| SwVrrpNewBackup       | Indicates that the sending agent has transitioned to "Backup" state.                                                                                                                                                     |
| SwVrrpAuthFailure     | Signifies that a packet has been received from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. Implementation of this trap is optional. |
| SwLoginFailure        | Signifies that someone failed to enter a valid username/password combination.                                                                                                                                            |
| SwTempExceedThreshold | Signifies that the switch temperature has exceeded maximum safety limits.                                                                                                                                                |
| SwTempReturnThreshold | Signifies that the switch temperature has returned to under maximum safety limits.                                                                                                                                       |
| SwStgNewRoot          | Signifies that the bridge has become the new root of the STG.                                                                                                                                                            |
| SwStgTopologyChanged  | Signifies that there was a STG topology change.                                                                                                                                                                          |
| SwStgBlockingState    | An SwStgBlockingState trap is sent when port state is changed in blocking state.                                                                                                                                         |
| SwCistNewRoot         | Signifies that the bridge has become the new root of the CIST.                                                                                                                                                           |
| SwCistTopologyChanged | Signifies that there was a CIST topology change.                                                                                                                                                                         |
| SwHotlinksMasterUp    | Signifies that the Master interface is active.                                                                                                                                                                           |
| SwHotlinksMasterDn    | Signifies that the Master interface is not active.                                                                                                                                                                       |
| SwHotlinksBackupUp    | Signifies that the Backup interface is active.                                                                                                                                                                           |
| SwHotlinksBackupDn    | Signifies that the Backup interface is not active.                                                                                                                                                                       |
| SwHotlinksNone        | Signifies that there are no active interfaces.                                                                                                                                                                           |

---

## Switch Images and Configuration Files

This section describes how to use MIB calls to work with switch images and configuration files. You can use a standard SNMP tool to perform the actions, using the MIBs listed in [Table 20](#).

[Table 20](#) lists the MIBs used to perform operations associated with the Switch Image and Configuration files.

*Table 20. MIBs for Switch Image and Configuration Files*

| MIB Name                   | MIB OID                               |
|----------------------------|---------------------------------------|
| agTransferServer           | 1.3.6.1.4.1.20301.2.7.17.1.1.7.1.0    |
| agTransferImage            | 1.3.6.1.4.1.20301.2.7.17.1.1.7.2.0    |
| agTransferImageFileName    | 1.3.6.1.4.1.20301.2.7.17.1.1.7.3.0    |
| agTransferCfgFileName      | 1.3.6.1.4.1.20301.2.7.17.1.1.7.4.0    |
| agTransferDumpFileName     | 1.3.6.1.4.1.20301.2.7.17.1.1.7.5.0    |
| agTransferAction           | 1.3.6.1.4.1.20301.2.7.17.1.1.7.6.0    |
| agTransferLastActionStatus | 1.3.6.1.4.1.20301.2.7.17.1.1.7.7.0    |
| agTransferUserName         | 1.3.6.1.4.1.20301.2.7.17.1.1.7.9.0    |
| agTransferPassword         | 1.3.6.1.4.1.1.20301.2.7.17.1.1.7.10.0 |
| agTransferTSDumpFileName   | 1.3.6.1.4.1.1.20301.2.7.17.1.1.7.11.0 |

The following SNMP actions can be performed using the MIBs listed in [Table 20](#).

- Load a new Switch image (boot or running) from a FTP/TFTP server
- Load a previously saved switch configuration from a FTP/TFTP server
- Save the switch configuration to a FTP/TFTP server
- Save a switch dump to a FTP/TFTP server

## Loading a New Switch Image

To load a new switch image with the name "MyNewImage-1.img" into image2, follow these steps. This example shows an FTP/TFTP server at IPv4 address 192.168.10.10, though IPv6 is also supported.

1. Set the FTP/TFTP server address where the switch image resides:  
`Set agTransferServer.0 "192.168.10.10"`
2. Set the area where the new image will be loaded:  
`Set agTransferImage.0 "image2"`
3. Set the name of the image:  
`Set agTransferImageFileName.0 "MyNewImage-1.img"`
4. If you are using an FTP server, enter a username:  
`Set agTransferUserName.0 "MyName"`
5. If you are using an FTP server, enter a password:  
`Set agTransferPassword.0 "MyPassword"`
6. Initiate the transfer. To transfer a switch image, enter 2 (gting):  
`Set agTransferAction.0 "2"`

## Loading a Saved Switch Configuration

To load a saved switch configuration with the name "MyRunningConfig.cfg" into the switch, follow these steps. This example shows a TFTP server at IPv4 address 192.168.10.10, though IPv6 is also supported.

1. Set the FTP/TFTP server address where the switch Configuration File resides:  
`Set agTransferServer.0 "192.168.10.10"`
2. Set the name of the configuration file:  
`Set agTransferCfgFileName.0 "MyRunningConfig.cfg"`
3. If you are using an FTP server, enter a username:  
`Set agTransferUserName.0 "MyName"`
4. If you are using an FTP server, enter a password:  
`Set agTransferPassword.0 "MyPassword"`
5. Initiate the transfer. To restore a running configuration, enter 3:  
`Set agTransferAction.0 "3"`

## Saving the Switch Configuration

To save the switch configuration to a FTP/TFTP server follow these steps. This example shows a FTP/TFTP server at IPv4 address 192.168.10.10, though IPv6 is also supported.

1. Set the FTP/TFTP server address where the configuration file is saved:  
`Set agTransferServer.0 "192.168.10.10"`
2. Set the name of the configuration file:  
`Set agTransferCfgFileName.0 "MyRunningConfig.cfg"`
3. If you are using an FTP server, enter a username:  
`Set agTransferUserName.0 "MyName"`
4. If you are using an FTP server, enter a password:  
`Set agTransferPassword.0 "MyPassword"`
5. Initiate the transfer. To save a running configuration file, enter 4:  
`Set agTransferAction.0 "4"`

## Saving a Switch Dump

To save a switch dump to a FTP/TFTP server, follow these steps. This example shows an FTP/TFTP server at 192.168.10.10, though IPv6 is also supported.

1. Set the FTP/TFTP server address where the configuration will be saved:  
`Set agTransferServer.0 "192.168.10.10"`
2. Set the name of dump file:  
`Set agTransferDumpFileName.0 "MyDumpFile.dmp"`
3. If you are using an FTP server, enter a username:  
`Set agTransferUserName.0 "MyName"`
4. If you are using an FTP server, enter a password:  
`Set agTransferPassword.0 "MyPassword"`
5. Initiate the transfer. To save a dump file, enter 5:  
`Set agTransferAction.0 "5"`

# Part 6: Monitoring

The ability to monitor traffic passing through the G7028/G7052 can be invaluable for troubleshooting some types of networking problems. This sections cover the following monitoring features:

- Remote Monitoring (RMON)
- Port Mirroring



---

## Chapter 20. Remote Monitoring

Remote Monitoring (RMON) allows network devices to exchange network monitoring data.

RMON allows the switch to perform the following functions:

- Track events and trigger alarms when a threshold is reached.
- Notify administrators by issuing a syslog message or SNMP trap.

---

### RMON Overview

The RMON MIB provides an interface between the RMON agent on the switch and an RMON management application. The RMON MIB is described in RFC 1757.

The RMON standard defines objects that are suitable for the management of Ethernet networks. The RMON agent continuously collects statistics and proactively monitors switch performance. RMON allows you to monitor traffic flowing through the switch.

The switch supports the following RMON Groups, as described in RFC 1757:

- Group 1: Statistics
- Group 2: History
- Group 3: Alarms
- Group 9: Events

---

### RMON Group 1—Statistics

The switch supports collection of Ethernet statistics as outlined in the RMON statistics MIB, in reference to etherStatsTable. You can configure RMON statistics on a per-port basis.

RMON statistics are sampled every second, and new data overwrites any old data on a given port.

**Note:** RMON port statistics must be enabled for the port before you can view RMON statistics.

## Example Configuration

1. Enable RMON on a port.

```
RS G7028(config)# interface port 1
RS G7028(config-if)# rmon
```

2. View RMON statistics for the port.

```
RS G7028(config-if)# show interface port 1 rmon-counters
-----
RMON statistics for port 3:
etherStatsDropEvents:                NA
etherStatsOctets:                    7305626
etherStatsPkts:                      48686
etherStatsBroadcastPkts:             4380
etherStatsMulticastPkts:             6612
etherStatsCRCAlignErrors:            22
etherStatsUndersizePkts:              0
etherStatsOversizePkts:              0
etherStatsFragments:                 2
etherStatsJabbers:                   0
etherStatsCollisions:                0
etherStatsPkts64Octets:              27445
etherStatsPkts65to127Octets:         12253
etherStatsPkts128to255Octets:        1046
etherStatsPkts256to511Octets:        619
etherStatsPkts512to1023Octets:       7283
etherStatsPkts1024to1518Octets:      38
```

---

## RMON Group 2—History

The RMON History Group allows you to sample and archive Ethernet statistics for a specific interface during a specific time interval. History sampling is done per port.

**Note:** RMON port statistics must be enabled for the port before an RMON History Group can monitor the port.

Data is stored in *buckets*, which store data gathered during discreet sampling intervals. At each configured interval, the History index takes a sample of the current Ethernet statistics, and places them into a bucket. History data buckets reside in dynamic memory. When the switch is re-booted, the buckets are emptied.

Requested buckets are the number of buckets, or data slots, requested by the user for each History Group. Granted buckets are the number of buckets granted by the system, based on the amount of system memory available. The system grants a maximum of 50 buckets.

You can use an SNMP browser to view History samples.

## History MIB Object ID

The type of data that can be sampled must be of an *ifIndex* object type, as described in RFC 1213 and RFC 1573. The most common data type for the History sample is as follows:

```
1.3.6.1.2.1.2.2.1.1.<x>
```

The last digit (x) represents the number of the port to monitor.

## Configuring RMON History

Perform the following steps to configure RMON History on a port.

1. Enable RMON on a port.

```
RS G7028(config)# interface port 1
RS G7028(config-if)# rmon
RS G7028(config-if)# exit
```

2. Configure the RMON History parameters for a port.

```
RS G7028(config)# rmon history 1 interface-oid
1.3.6.1.2.1.2.2.1.1.<x>
RS G7028(config)# rmon history 1 requested-buckets 30
RS G7028(config)# rmon history 1 polling-interval 120
RS G7028(config)# rmon history 1 owner "rmon port 1 history"
```

where <x> is the number of the port to monitor. For example, the full OID for port 1 would be:

1.3.6.1.2.1.2.2.1.1.1

3. View RMON history for the port.

```
RS G7028(config)# show rmon history
RMON History group configuration:

Index          IF OID          Interval    Rbnum    Gbnum
-----
1  1.3.6.1.2.1.2.2.1.1.1    120        30        30

Index          Owner
-----
1  rmon port 1 history
```

---

## RMON Group 3—Alarms

The RMON Alarm Group allows you to define a set of thresholds used to determine network performance. When a configured threshold is crossed, an alarm is generated. For example, you can configure the switch to issue an alarm if more than 1,000 CRC errors occur during a 10-minute time interval.

Each Alarm index consists of a variable to monitor, a sampling time interval, and parameters for rising and falling thresholds. The Alarm Group can be used to track rising or falling values for a MIB object. The object must be a counter, gauge, integer, or time interval.

Use one of the following commands to correlate an Alarm index to an Event index:

```
RS G7028(config)# rmon alarm <alarm number> rising-crossing-index
<event number>
RS G7028(config)# rmon alarm <alarm number> falling-crossing-index
<event number>
```

When the alarm threshold is reached, the corresponding event is triggered.

### Alarm MIB objects

The most common data types used for alarm monitoring are `ifStats`: errors, drops, bad CRCs, and so on. These MIB Object Identifiers (OIDs) correlate to the ones tracked by the History Group. An example statistic follows:

```
1.3.6.1.2.1.5.1.0 - mgmt.icmp.icmpInMsgs
```

This value represents the alarm's MIB OID, as a string. Note that for non-tables, you must supply a `.0` to specify end node.

### Configuring RMON Alarms

Configure the RMON Alarm parameters to track ICMP messages.

```
RS G7028(config)# rmon alarm 1 oid 1.3.6.1.2.1.5.8.0
RS G7028(config)# rmon alarm 1 alarm-type rising
RS G7028(config)# rmon alarm 1 rising-crossing-index 110
RS G7028(config)# rmon alarm 1 interval-time 60
RS G7028(config)# rmon alarm 1 rising-limit 200
RS G7028(config)# rmon alarm 1 sample delta
RS G7028(config)# rmon alarm 1 owner "Alarm for icmpInEchos"
```

This configuration creates an RMON alarm that checks `icmpInEchos` on the switch once every minute. If the statistic exceeds 200 within a 60 second interval, an alarm is generated that triggers event index 110.

---

## RMON Group 9—Events

The RMON Event Group allows you to define events that are triggered by alarms. An event can be a log message, an SNMP trap, or both.

When an alarm is generated, it triggers a corresponding event notification. Use the following commands to correlate an Event index to an alarm:

```
RS G7028(config)# rmon alarm <alarm number> rising-crossing-index  
                  <event number>  
RS G7028(config)# rmon alarm <alarm number> falling-crossing-index  
                  <event number>
```

RMON events use SNMP and syslogs to send notifications. Therefore, an SNMP trap host must be configured for trap event notification to work properly.

RMON uses a syslog host to send syslog messages. Therefore, an existing syslog host must be configured for event log notification to work properly. Each log event generates a syslog of type RMON that corresponds to the event.

For example, to configure the RMON event parameters.

```
RS G7028(config)# rmon event 110 type log  
RS G7028(config)# rmon event 110 description "SYSLOG_this_alarm"  
RS G7028(config)# rmon event 110 owner "log icmpInEchos alarm"
```

This configuration creates an RMON event that sends a syslog message each time it is triggered by an alarm.

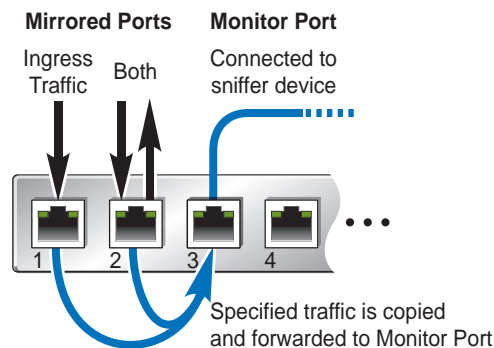


## Chapter 21. Port Mirroring

The IBM Networking OS port mirroring feature allows you to mirror (copy) the packets of a target port, and forward them to a monitoring port. Port mirroring functions for all Layer 2 and Layer 3 traffic on a port. This feature can be used as a troubleshooting tool or to enhance the security of your network. For example, an IDS server or other traffic sniffer device or analyzer can be connected to the monitoring port to detect intruders attacking the network.

The G7028/G7052 supports a “many to one” mirroring model. As shown in [Figure 26](#), selected traffic for ports 1 and 2 is being monitored by port 3. In the example, both ingress traffic and egress traffic on port 2 are copied and forwarded to the monitor. However, port 1 mirroring is configured so that only ingress traffic is copied and forwarded to the monitor. A device attached to port 3 can analyze the resulting mirrored traffic.

Figure 26. Mirroring Ports



The G7028/G7052 supports two monitor ports. Each monitor port can receive mirrored traffic from any number of target ports.

IBM N/OS does not support “one to many” or “many to many” mirroring models where traffic from a specific port is copied to multiple monitor ports. For example, port 1 traffic cannot be monitored by both port 3 and 4 at the same time, nor can port 2 ingress traffic be monitored by a different port than its egress traffic.

Ingress and egress traffic is duplicated and sent to the monitor port after processing.

### Configuring Port Mirroring

The following procedure may be used to configure port mirroring for the example shown in [Figure 26 on page 253](#):

1. Specify the monitoring port, the mirroring port(s), and the port-mirror direction.

```
RS G7028(config)# port-mirroring monitor-port 3 mirroring-port 1 in
RS G7028(config)# port-mirroring monitor-port 3 mirroring-port 2 both
```

2. Enable port mirroring.

```
RS G7028(config)# port-mirroring enable
```

### 3. View the current configuration.

```
RS G7028# show port-mirroring

Port Monitoring : Enabled

Monitoring Ports      Mirrored Ports
3                     (1, in)
                     (2, both)
...
```

# Part 7: Appendices

- Glossary
- Getting help and technical assistance
- Notices



---

## Appendix A. Glossary

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CNA</b>              | Converged Network Adapter. A device used for I/O consolidation such as that in Converged Enhanced Ethernet (CEE) environments implementing Fibre Channel over Ethernet (FCoE). The CNA performs the duties of both a Network Interface Card (NIC) for Local Area Networks (LANs) and a Host Bus Adapter (HBA) for Storage Area Networks (SANs).                                                                                                                                                             |
| <b>DIP</b>              | The destination IP address of a frame.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Dport</b>            | The destination port (application socket: for example, http-80/https-443/DNS-53)                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>HBA</b>              | Host Bus Adapter. An adapter or card that interfaces with device drivers in the host operating system and the storage target in a Storage Area Network (SAN). It is equivalent to a Network Interface Controller (NIC) from a Local Area Network (LAN).                                                                                                                                                                                                                                                     |
| <b>NAT</b>              | Network Address Translation. Any time an IP address is changed from one source IP or destination IP address to another address, network address translation can be said to have taken place. In general, half NAT is when the destination IP or source IP address is changed from one address to another. Full NAT is when both addresses are changed from one address to another. No NAT is when neither source nor destination IP addresses are translated.                                               |
| <b>Preemption</b>       | In VRRP, preemption will cause a Virtual Router that has a lower priority to go into backup if a peer Virtual Router starts advertising with a higher priority.                                                                                                                                                                                                                                                                                                                                             |
| <b>Priority</b>         | In VRRP, the value given to a Virtual Router to determine its ranking with its peer(s). Minimum value is 1 and maximum value is 254. Default is 100. A higher number will win out for master designation.                                                                                                                                                                                                                                                                                                   |
| <b>Proto (Protocol)</b> | The protocol of a frame. Can be any value represented by a 8-bit value in the IP header adherent to the IP specification (for example, TCP, UDP, OSPF, ICMP, and so on.)                                                                                                                                                                                                                                                                                                                                    |
| <b>SIP</b>              | The source IP address of a frame.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>SPort</b>            | The source port (application socket: for example, HTTP-80/HTTPS-443/DNS-53).                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Tracking</b>         | <p>In VRRP, a method to increase the priority of a virtual router and thus master designation (with preemption enabled). Tracking can be very valuable in an active/active configuration.</p> <p>You can track the following:</p> <ul style="list-style-type: none"><li>• Active IP interfaces on the Web switch (increments priority by 2 for each)</li><li>• Active ports on the same VLAN (increments priority by 2 for each)</li><li>• Number of virtual routers in master mode on the switch</li></ul> |
| <b>VIR</b>              | Virtual Interface Router. A VRRP address is an IP interface address shared between two or more virtual routers.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Virtual Router</b>   | A shared address between two devices utilizing VRRP, as defined in RFC 2338. One virtual router is associated with an IP interface. This is one of the IP interfaces that the switch is assigned. All IP interfaces on the G7028/G7052s must be in a VLAN. If there is more than one VLAN defined on the Web switch, then the VRRP broadcasts will only be sent out on the VLAN of which the associated IP interface is a member.                                                                           |

---

|             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>VRID</b> | <p>Virtual Router Identifier. In VRRP, a numeric ID is used by each virtual router to create its MAC address and identify its peer for which it is sharing this VRRP address. The VRRP MAC address as defined in the RFC is 00-00-5E-00-01-<i>&lt;VRID&gt;</i>.</p> <p>If you have a VRRP address that two switches are sharing, then the VRID number needs to be identical on both switches so each virtual router on each switch knows with whom to share.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>VRRP</b> | <p>Virtual Router Redundancy Protocol. A protocol that acts very similarly to Cisco's proprietary HSRP address sharing protocol. The reason for both of these protocols is so devices have a next hop or default gateway that is always available. Two or more devices sharing an IP interface are either advertising or listening for advertisements. These advertisements are sent via a broadcast message to an address such as 224.0.0.18.</p> <p>With VRRP, one switch is considered the master and the other the backup. The master is always advertising via the broadcasts. The backup switch is always listening for the broadcasts. If the master stops advertising, the backup will take over ownership of the VRRP IP and MAC addresses as defined by the specification. The switch announces this change in ownership to the devices around it by way of a Gratuitous ARP, and advertisements. If the backup switch didn't do the Gratuitous ARP the Layer 2 devices attached to the switch would not know that the MAC address had moved in the network. For a more detailed description, refer to RFC 2338.</p> |

---

---

## Appendix B. Getting help and technical assistance

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your system, and whom to call for service, if it is necessary.

---

### Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system and any optional devices are turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Problem Determination and Service Guide* on the IBM Documentation CD that comes with your system.
- Go to the IBM support website at <http://www.ibm.com/systems/support/> to check for technical information, hints, tips, and new device drivers or to submit a request for information.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the documentation that is provided with your IBM product. The documentation that comes with IBM systems also describes the diagnostic tests that you can perform. Most systems, operating systems, and programs come with documentation that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the documentation for the operating system or program.

---

### Using the documentation

Information about your IBM system and pre-installed software, if any, or optional device is available in the documentation that comes with the product. That documentation can include printed documents, online documents, ReadMe files, and Help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to <http://www.ibm.com/systems/support/> and follow the instructions. Also, some documents are available through the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

---

## Getting help and information on the World Wide Web

On the World Wide Web, the IBM website has up-to-date information about IBM systems, optional devices, services, and support. The address for IBM System x<sup>®</sup> and xSeries<sup>®</sup> information is <http://www.ibm.com/systems/x/>. The address for IBM BladeCenter information is <http://www.ibm.com/systems/bladecenter/>. The address for IBM IntelliStation<sup>®</sup> information is <http://www.ibm.com/intellistation/>.

You can find service information for IBM systems and optional devices at <http://www.ibm.com/systems/support/>.

---

## Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with System x and x Series servers, BladeCenter products, IntelliStation workstations, and appliances. For information about which products are supported by Support Line in your country or region, see <http://www.ibm.com/services/sl/products/>.

For more information about Support Line and other IBM services, see <http://www.ibm.com/services/>, or see <http://www.ibm.com/planetwide/> for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

---

## Hardware service and support

You can receive hardware service through your IBM reseller or IBM Services. To locate a reseller authorized by IBM to provide warranty service, go to <http://www.ibm.com/partnerworld/> and click **Find Business Partners** on the right side of the page. For IBM support telephone numbers, see <http://www.ibm.com/planetwide/>. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

---

## IBM Taiwan product service

台灣 IBM 產品服務聯絡方式：  
台灣國際商業機器股份有限公司  
台北市松仁路 7 號 3 樓  
電話：0800-016-888

IBM Taiwan product service contact information:

IBM Taiwan Corporation  
3F, No 7, Song Ren Rd.  
Taipei, Taiwan  
Telephone: 0800-016-888

---

## Appendix C. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product, and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

---

## Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

Intel, Intel Xeon, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

---

## Important Notes

Processor speed indicates the internal clock speed of the microprocessor; other factors also affect application performance.

CD or DVD drive speed is the variable read rate. Actual speeds vary and are often less than the possible maximum.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for 1024 bytes, MB stands for 1 048 576 bytes, and GB stands for 1 073 741 824 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity can vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives that are available from IBM.

Maximum memory might require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.

---

## Particulate contamination

**Attention:** Airborne particulates (including metal flakes or particles) and reactive gases acting alone or in combination with other environmental factors such as humidity or temperature might pose a risk to the device that is described in this document. Risks that are posed by the presence of excessive particulate levels or concentrations of harmful gases include damage that might cause the device to malfunction or cease functioning altogether. This specification sets forth limits for particulates and gases that are intended to avoid such damage. The limits must not be viewed or used as definitive limits, because numerous other factors, such as temperature or moisture content of the air, can influence the impact of particulates or environmental corrosives and gaseous contaminant transfer. In the absence of specific limits that are set forth in this document, you must implement practices that maintain particulate and gas levels that are consistent with the protection of human health and safety. If IBM determines that the levels of particulates or gases in your environment have caused damage to the device, IBM may condition provision of repair or replacement of devices or parts on implementation of appropriate remedial measures to mitigate such environmental contamination. Implementation of such remedial measures is a customer responsibility.

| Contaminant | Limits                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Particulate | <ul style="list-style-type: none"><li>• The room air must be continuously filtered with 40% atmospheric dust spot efficiency (MERV 9) according to ASHRAE Standard 52.2<sup>1</sup>.</li><li>• Air that enters a data center must be filtered to 99.97% efficiency or greater, using high-efficiency particulate air (HEPA) filters that meet MIL-STD-282.</li><li>• The deliquescent relative humidity of the particulate contamination must be more than 60%<sup>2</sup>.</li><li>• The room must be free of conductive contamination such as zinc whiskers.</li></ul> |
| Gaseous     | <ul style="list-style-type: none"><li>• Copper: Class G1 as per ANSI/ISA 71.04-1985<sup>3</sup></li><li>• Silver: Corrosion rate of less than 300 Å in 30 days</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                 |

<sup>1</sup> ASHRAE 52.2-2008 - *Method of Testing General Ventilation Air-Cleaning Devices for Removal Efficiency by Particle Size*. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

<sup>2</sup> The deliquescent relative humidity of particulate contamination is the relative humidity at which the dust absorbs enough water to become wet and promote ionic conduction.

<sup>3</sup> ANSI/ISA-71.04-1985. *Environmental conditions for process measurement and control systems: Airborne contaminants*. Instrument Society of America, Research Triangle Park, North Carolina, U.S.A.

---

## Documentation format

The publications for this product are in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties when you use the PDF files and want to request a web-based format or accessible PDF document for a publication, direct your mail to the following address:

Information Development  
IBM Corporation  
205/A0153039 E. Cornwallis Road  
P.O. Box 12195  
Research Triangle Park, North Carolina 27709-2195  
U.S.A.

In the request, be sure to include the publication part number and title.

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

---

## Electronic emission notices

### Federal Communications Commission (FCC) statement

**Note:** This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

### Industry Canada Class A emission compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

### Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

### Australia and New Zealand Class A statement

**Attention:** This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

## European Union EMC Directive conformance statement

This product is in conformity with the protection requirements of EU Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

**Attention:** This is an EN 55022 Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Responsible manufacturer:

International Business Machines Corp.  
New Orchard Road  
Armonk, New York 10504  
914-499-1900

European Community contact:

IBM Technical Regulations, Department M456  
IBM-Allee 1, 71137 Ehningen, Germany  
Telephone: +49 7032 15-2937  
E-mail: tjahn@de.ibm.com

## Germany Class A statement

**Deutschsprachiger EU Hinweis:**

### **Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit**

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: "Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

### **Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten**

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 2004/108/EG in der Bundesrepublik Deutschland.

**Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC EG Richtlinie 2004/108/EG) für Geräte der Klasse A**

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:

International Business Machines Corp.  
New Orchard Road  
Armonk, New York 10504  
914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

IBM Deutschland  
Technical Regulations, Department M456  
IBM-Allee 1, 71137 Ehningen, Germany  
Telephone: +49 7032 15-2937  
E-mail: tjahn@de.ibm.com

**Generelle Informationen:**

**Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.**

## Japan VCCI Class A statement

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用する  
と電波妨害を引き起こすことがあります。この場合には使用者が適切な対策  
を講ずるよう要求されることがあります。 VCCI-A

This is a Class A product based on the standard of the Voluntary Control Council for Interference (VCCI). If this equipment is used in a domestic environment, radio interference may occur, in which case the user may be required to take corrective actions.

## Korea Communications Commission (KCC) statement

이 기기는 업무용으로 전자파 적합등록을 받은 기기  
이오니, 판매자 또는 사용자는 이점을 주의하시기  
바라며, 만약 잘못 구입하셨을 때에는 구입한 곳에  
서 비업무용으로 교환하시기 바랍니다.

Please note that this equipment has obtained EMC registration for commercial use. In the event that it has been mistakenly sold or purchased, please exchange it for equipment certified for home use.

## Russia Electromagnetic Interference (EMI) Class A statement

ВНИМАНИЕ! Настоящее изделие относится к классу А.  
В жилых помещениях оно может создавать радиопомехи, для  
снижения которых необходимы дополнительные меры

## People's Republic of China Class A electronic emission statement

中华人民共和国“A类”警告声明

声 明

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

## Taiwan Class A compliance statement

警告使用者：  
這是甲類的資訊產品，在  
居住的環境中使用時，可  
能會造成射頻干擾，在這  
種情況下，使用者會被要  
求採取某些適當的對策。



---

# Index

## Symbols

[ ] 13

## Numerics

802.1Q VLAN tagging 102

## A

Access Control List (ACL) 161  
Access Control Lists. *See* ACLs.  
accessible documentation 264  
accessing the switch  
    Browser-based Interface 18, 22  
    LDAP authentication 72  
    RADIUS authentication 64  
    security 53, 63  
    TACACS+ 68  
ACL metering 162  
ACLs 83, 161  
administrator account 25, 29, 66  
anycast address, IPv6 184  
application ports 85  
assistance, getting 255, 259  
autoconfiguration  
    IPv6 185  
    link 33  
auto-negotiation  
    setup 33

## B

BBI 18  
Bridge Protocol Data Unit (BPDU) 130  
broadcast domains 99  
broadcast storm control 95  
Browser-Based Interface 18

## C

Cisco EtherChannel 117, 119  
CIST 142  
Class A electronic emission notice 264  
Class of Service queueCOS queue 170  
command conventions 13  
Command-Line Interface (CLI) 29  
Community VLANPrivate VLANs  
    Community VLAN 112  
configuration rules  
    Trunking 117  
configuring  
    port trunking 118  
    spanning tree groups 139, 144  
contamination, particulate and gaseous 263

## D

date  
    setup 31  
default password 25, 66  
Differentiated Services Code Point (DSCP) 163  
documentation format 264  
downloading software 42  
DSCP 163

## E

EAPoL 76  
electronic emission Class A notice 264  
End user access control  
    configuring 59  
EtherChannel 116  
    as used with port trunking 117, 119  
Extensible Authentication Protocol over LAN 76

## F

factory default configuration 27, 29  
failover 215  
FCC Class A notice 264  
Final Steps 38  
first-time configuration 27, 29 to ??  
flow control  
    setup 33  
frame size 100  
frame tagging. *See* VLANs tagging.

## G

gaseous contamination 263  
getting help 259

## H

hardware service and support 260  
help, getting 259  
Hot Links 213  
HP-OpenView 24, 233

## I

IBM DirectorSNMP  
    IBM Director 24, 233  
IBM support line 260  
ICMP 84  
IEEE standards  
    802.1D 128  
    802.1p 169  
    802.1Q 102  
    802.1s 142  
    802.1x 76  
IGMP 84, 191  
    Querier 195  
IGMPv3 196

- image
  - downloading 42
- Internet Group Management Protocol (IGMP) 191
- IP address 36
  - IP interface 36
- IP configuration via setup 36
- IP interfaces 36
- IP routing 36
- IP subnet mask 36
- IP subnets
  - VLANs 99
- IPv6 addressing 181, 183
- ISL Trunking 116
- Isolated VLANPrivate VLANs
  - Isolated VLAN 112

## J

- jumbo frames 100

## L

- LACP 120
- Layer 2 Failover 215
- LDAP
  - authentication 72
- Link Aggregation Control Protocol 120
- Link Layer Discovery Protocol 221
- LLDP 221
- logical segment. See IP subnets.

## M

- manual style conventions 13
- Maximum Transmission Unit 100
- meter 89
- meter (ACL) 162
- mirroring ports 253
- monitoring ports 253
- MSTPMultiple Spanning Tree Protocol (MSTP) 142
- MTU 100
- multi-links between switches
  - using port trunking 115
- multiple spanning tree groups 134
- Multiple Spanning Tree Protocol 142

## N

- network management 18, 24, 233
- notes, important 262
- notices 261
- notices, electronic emission 264
- notices, FCC Class A 264

## O

- OSPF
  - filtering criteria 84

## P

- packet size 100
- particulate contamination 263
- password
  - administrator account 25, 66
  - default 25, 66
  - user account 25, 66
- passwords 25
- payload size 100
- Per Hop Behavior (PHB)PHB 164
- port flow control. See flow control.
- port mirroring 253
- Port Trunking 117
- port trunking
  - configuration example 118
  - description 119
  - EtherChannel 116
- ports
  - configuration 33
  - for services 85
  - monitoring 253
  - physical. See switch ports.
- Private VLANs 112
- promiscuous port 112
- protocol types 84
- PVID (port VLAN ID) 101
- PVLANprotocol-based VLAN 109

## Q

- QoS 159
- Quality of Service 159
- Querier (IGMP) 195

## R

- RADIUS
  - authentication 64
  - port 1812 and 1645 85
  - port 1813 85
  - SSH/SCP 57
- Rapid Spanning Tree Protocol (RSTP) 141
- Rapid Spanning Tree Protocol (RSTP)RSTP 141
- receive flow control 33
- re-mark 89, 162
- restarting switch setup 30
- RMON alarms 250
- RMON events 251
- RMON History 248
- RMON statistics 247
- routers
  - port trunking 116
- RSA keys 57
- RSTP 141
- rx flow control 33

## S

- SecurID 58
- security
  - LDAP authentication 72
  - port mirroring 253
  - RADIUS authentication 64
  - TACACS+ 68
  - VLANs 99
- segmentation. *See* IP subnets.
- segments. *See* IP subnets.
- service and support 260
- service ports 85
- setup facility 27, 29
  - IP configuration 36
  - IP subnet mask 36
  - port auto-negotiation mode 33
  - port configuration 33
  - port flow control 33
  - restarting 30
  - Spanning-Tree Protocol 32
  - starting 30
  - stopping 30
  - system date 31
  - system time 31
  - VLAN name 35
  - VLAN tagging 33
  - VLANs 35
- SNMP 18, 24, 233
  - HP-OpenView 24, 233
- SNMP Agent 233
- software
  - image 41
- software service and support 260
- Source-Specific MulticastSSM 196
- Spanning-Tree Protocol
  - multiple instances 134
  - setup (on/off) 32
- SSH/SCP
  - configuring 54
  - RSA host and server keys 57
- starting switch setup 30
- stopping switch setup 30
- subnet mask 36
- subnets 36
- support line 260
- support web site 260
- switch ports VLANs membership 101

## T

- TACACS+ 68
- tagging. *See* VLANs tagging.
- TCP 84
- technical assistance 259

- technical terms
  - port VLAN identifier (PVID) 102
  - tagged frame 102
  - tagged member 102
  - untagged frame 102
  - untagged member 102
  - VLAN identifier (VID) 102
- telephone assistance 260
- telephone numbers 260
- Telnet support
  - optional setup for Telnet support 39
- text conventions 13
- time
  - setup 31
- trademarks 261
- transmit flow control 33
- Trunking
  - configuration rules 117
- tx flow control 33
- typographic conventions 13

## U

- UDP 84
- upgrade, switch software 41
- USB drive 44
- user account 25, 66

## V

- Virtual Local Area Networks. *See* VLANs.
- VLAN tagging
  - setup 33
- VLANs 36
  - broadcast domains 99
  - default PVID 101
  - example showing multiple VLANs 106
  - ID numbers 100
  - interface 36
  - multiple spanning trees 129
  - multiple VLANs 102
  - name setup 35
  - port members 101
  - PVID 101
  - security 99
  - setup 35
  - Spanning-Tree Protocol 129
  - tagging 33, 101 to 107
  - topologies 106

## W

- website, publication ordering 259
- website, support 260
- website, telephone support numbers 260

