

ES 500 Switch Router User Guide

Release 1.0

36-063-01 Rev. 0A



COPYRIGHT NOTICES

© 2002 by Riverstone Networks, Inc. All rights reserved.

Riverstone Networks, Inc.
5200 Great America Parkway
Santa Clara, CA 95054

Printed in the United States of America

This product includes software developed by the University of California, Berkeley, and its contributors.

© 1979 – 1994 by The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:
This product includes software developed by the University of California, Berkeley, and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Changes

Riverstone Networks, Inc., and its licensors reserve the right to make changes in specifications and other information contained in this document without prior notice. The reader should in all cases consult Riverstone Networks, Inc., to determine whether any such changes have been made.

The hardware, firmware, or software described in this manual is subject to change without notice.

Disclaimer

IN NO EVENT SHALL RIVERSTONE NETWORKS BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING BUT NOT LIMITED TO LOST PROFITS) ARISING OUT OF OR RELATED TO THIS MANUAL OR THE INFORMATION CONTAINED IN IT, EVEN IF RIVERSTONE NETWORKS HAS BEEN ADVISED, KNOWN, OR SHOULD HAVE KNOWN, OF THE POSSIBILITY OF SUCH DAMAGES.

Trademarks

Riverstone Networks, Riverstone, RS, and IA are trademarks of Riverstone Networks, Inc.

All other product names mentioned in this manual may be trademarks or registered trademarks of their respective companies.

TABLE OF CONTENTS

1	Preface	1-1
1.1	Related Documentation	1-1
2	System Upgrade Prodedures	2-1
2.1	Prom Image	2-1
2.1.1	Updating Prom Image	2-1
2.2	System Image	2-3
2.2.1	Updating System Image	2-3
3	Maintaining Configuration Files	3-1
3.1	Configuration Files	3-1
3.1.1	Changing Configuration Information	3-2
3.1.2	Displaying Configuration Information	3-3
3.1.3	Commenting out the contents of the Active via Scratchpad	3-3
3.1.4	Activating the Configuration Commands in the Scratchpad	3-3
3.1.5	Saving the Active Configuration to the Startup Configuration File	3-4
3.1.6	Viewing the Current Configuration	3-4
4	CLI Basics	4-1
4.1	Understanding CLI Command Modes	4-1
4.1.1	Accessing the command mode	4-1
4.1.2	User Mode	4-2
4.1.3	Enable Mode	4-2
4.1.4	Configure mode	4-2
4.2	Understanding CLI Commands	4-3
4.3	Getting Help with CLI Commands	4-3
4.4	Setting System Commands	4-4
4.5	Line Editing Commands	4-4
4.6	CLI and ES 500 Configuration Example	4-4
4.7	Monitoring System Settings	4-5
5	Port Management	5-1
5.1	Setting Port Modes and Parameters	5-1
5.1.1	Enabling and Disabling Ports	5-1
5.1.2	Setting Port Parameters	5-1
5.2	Monitoring Ports	5-2
6	Bridging Configuration	6-1
6.1	Spanning Tree (IEEE 802.1d)	6-1
6.2	Supported Bridging Modes	6-1
6.3	VLAN support	6-1
6.3.1	VLANs and the Device	6-2
6.3.2	Configuration Examples	6-2
6.4	Access Ports and Trunk Ports (802.1p and 802.1q Support)	6-3
6.4.1	Setting a VLAN Access port	6-3
6.4.2	Configuring VLAN Trunk Ports	6-3
6.5	Configuring Spanning Tree	6-3
6.5.1	Adjusting Spanning-Tree Parameters	6-4

6.6	Configuring a Port-Based VLAN	6-6
6.6.1	Creating a Port-Based VLAN	6-6
6.6.2	Adding Ports to a VLAN	6-6
6.7	Monitoring Bridging.....	6-6
6.8	GARP/GVRP	6-7
6.8.1	Running GVRP with STP	6-8
6.8.2	Configuring GVRP	6-8
6.8.3	Configuration Example.....	6-8
6.8.4	Monitoring GVRP status	6-10
7	SmartTRUNK Configuration	7-1
7.1	Configuring SmartTRUNKS	7-1
7.1.1	Creating a SmartTRUNK	7-1
7.1.2	Adding Physical Ports to the SmartTRUNK.....	7-2
7.1.3	Specifying Traffic Load Policy.....	7-2
7.2	Monitoring SmartTRUNK Configuration	7-3
7.3	SmartTRUNK Configuration Example	7-4
8	IP ROUTING Configuration	8-1
8.1	IP Routing Protocols	8-1
8.1.1	Unicast Routing Protocols	8-1
8.1.2	Multicast Routing Protocols	8-1
8.2	Configuring IP Interfaces and Parameters	8-1
8.2.1	Configuring and Displaying IGMP Snooping	8-2
8.2.2	Configuring IP Interfaces to Ports	8-2
8.2.3	Configuring IP Interfaces for a VLAN	8-2
8.2.4	Monitoring IP interface configuration	8-3
8.3	Configuring IP Helper	8-3
8.4	Configuration Example	8-4
8.4.1	Assigning IP Interfaces.....	8-4
9	RIP Configuration	9-1
9.1	Configuring RIP	9-1
9.1.1	Enabling and Disabling RIP	9-1
9.1.2	Configuring RIP Interfaces.....	9-1
9.2	Configuring RIP Parameters.....	9-2
9.3	Monitoring RIP	9-2
9.4	Configuration Example	9-3
10	OSPF Configuration	10-1
10.1	Configuring OSPF	10-2
10.2	Setting the router ID	10-2
10.3	Enabling OSPF	10-2
10.4	Configuring OSPF Areas.....	10-2
10.4.1	Configuring Stub Areas	10-3
10.5	Configuring OSPF interfaces.....	10-3
10.6	Configuring OSPF Interface Parameters	10-3
10.6.1	Setting the Interface State	10-4
10.7	OSPF Topology and Configuration Example	10-4
11	ACL Configuration	11-1
11.1	ACLs Basics	11-1
11.1.1	Defining Selection Criteria in ACL Rules	11-1
11.1.2	How ACL Rules are Evaluated.....	11-2
11.2	Using ACLs.....	11-2
11.2.1	Applying ACLs to Interfaces.....	11-3

11.2.2	Applying ACLs to ports.....	11-3
11.2.3	Number of ACLs	11-4
11.3	Monitoring ACLS.....	11-5
12	Security Configuration	12-1
12.1	Layer 2 Security Filters	12-1
12.1.1	Configuring Layer-2 Address Filters.....	12-1
12.1.2	Configuring Layer-2 Port-to-Address Lock Filters.....	12-1
12.1.3	Configuring Layer-2 Static Entry Filters	12-2
12.1.4	Monitoring Layer-2 Security Filters.....	12-2
13	QOS Configuration	13-1
13.1	Layer-2, Layer-3 and Layer-4 Flow Specification.....	13-2
13.2	Precedence for Layer-3 Flows.....	13-2
13.3	Queing Policies	13-2
13.4	Traffic Prioritization for Layer-2 Flows.....	13-3
13.4.1	Configuring Layer-2 QoS.....	13-3
13.4.2	802.Ip Class of Service Priority Mapping.....	13-3
13.5	Traffic Prioritization for Layer-3 and Layer-4 Flows.....	13-5
13.5.1	Configuring IP QoS Policies.....	13-5
13.6	Configuring ES 500 Queuing Policy	13-5
13.7	Monitoring QOS.....	13-6
13.8	Limiting Traffic Rate.....	13-6
13.8.1	Port Rate Limiting	13-6
13.9	Monitoring rate limiting policies.....	13-7
14	Performance Monitoring	14-1
14.1	Configuring Port Mirroring.....	14-2
14.2	Monitoring Broadcast Traffic.....	14-2
15	RMON Configuration	15-1
15.1	Configuring and Enabling RMON.....	15-1
15.1.1	RMON Groups	15-1
15.1.2	Control Tables	15-1
15.2	Using RMON	15-2
15.3	Configuring RMON Groups.....	15-2
15.3.1	Configuration Examples.....	15-2
15.4	Displaying RMON Information.....	15-3
16	Service Configuration	16-1
16.1	Configuring Rate Limit Services.....	16-1
16.2	Applying Rate Limit Services	16-1
16.2.1	Applying Aggregate and Port-Level Rate Limiting.....	16-2
16.2.2	Per-Flow Rate Limiting.....	16-4
17	Time Configuration	17-1
17.1	Setting Time and Date.....	17-1
18	SNMP Configuration.....	18-1
18.1	Configuring Access to MIB Objects.....	18-1
18.1.1	Configuring SNMP access.....	18-1
18.2	Configuring SNMP Notifications	18-2
18.2.1	Specifying the Targets.....	18-2
18.3	Monitoring SNMP Configuration.....	18-3

Appendix A Troubleshooting	A-1
A.1 Password troubleshooting.....	A-1
A.1.1 Erasing the Startup File from the FLASH	A-1
A.1.2 Overriding Forgotten Passwords.....	A-2
A.2 Configuration Troubleshooting	A-3
A.3 Erasing the Startup File	A-3
A.4 Editing the Startup File.....	A-5

LIST OF TABLES

Table 1-1	Related Documentation	1-1
Table 13-1	802.1p default priority mappings.....	13-3
Table 15-1	RMON groups	15-1

LIST OF FIGURES

Figure 2-1	The Send File window.....	2-5
Figure 2-2	The Xmodem Send File window	2-5
Figure 3-1	Commands to save configurations.....	3-2
Figure 6-1	Using GARP/GVRP on a network	6-9
Figure 7-1	SmartTRUNK configuration example	7-4
Figure 9-1	RIP Configuration Example	9-3
Figure 10-1	Example of OSPF topology	10-4
Figure 16-1	Applying aggregate and port-level rate limiting.....	16-2

1 PREFACE

This manual provides information for configuring the Riverstone ES 500 Switch Router software. It details the procedures and provides configuration examples. To install the ES 500 chassis, use the instructions in the *Riverstone ES 500 Switch Router Getting Started Guide* and perform basic setup tasks, then return to this manual for more detailed configuration information.

1.1 RELATED DOCUMENTATION

The documentation set includes the following items. Refer to these other documents to learn more about your product.

Table 1-1 Related Documentation

Guide	Contents
<i>Riverstone Networks ES 500 Switch Router Getting Started Guide</i>	Installing and setting up the device
<i>Riverstone Networks ES 500 Switch Router Command Line Reference Guide</i>	How to use CLI (Command Line Interface) commands to configure and manage the device
<i>Riverstone Networks ES 500 Switch Router Message Reference Manual</i>	SYSLOG messages and SNMP traps

2 SYSTEM UPGRADE PROCEDURES

This chapter describes updating and upgrading software. The ES 500 software includes:

- Two copies of boot prom image in **.rfb** format.
- Two copies of system image (product software) in **.arc** format.

File copies insure that if an image is corrupted in one Flash sector, the other can still be used. Both prom and system images can be downloaded through a TFTP Server and installed into the Flash memory.

2.1 PROM IMAGE

The prom image is saved in **.rfb** format. Each image is stored in a Flash sector. The prom image can be updated and used the next time the device reboots.

2.1.1 Updating Prom Image

Updating prom image using TFTP

In order to update prom image, the system must be up and running with at least one port configured with an IP interface and connected to a TFTP server.

1. Install the new prom image on the TFTP server; the file name is in the **rfb** format.
2. In Enable mode enter the command with the correct file name to be downloaded. Following is an example of the command.

```
system promimage upgrade 176.210.100.45 om3_bt_200.rfb
```

3. When the TFTP download starts, the first copy of prom image is programmed to flash and verified. If the download is successful, the second copy image is automatically programmed. The two copies are contained in the same **rfb** file and are updated in succession *without user intervention*.
4. Reset the system to see if the new prom image was programmed successfully.
5. Check that DATE AND TIME are changed. The following figure displays an example of the bootup screen.

```
BOOT Software Version Prom-1.0.0.0 Built 06-May-2002 13:52:14
Processor: MPC8245 Rev 0.12, 266 MHz (Bus: 133MHz), 64 MByte SDRAM.
I-Cache 16 KB, linesize 32.D-Cache 16 KB, linesize 32.
Cache Enabled.

rs#system promimage upgrade 16.1.1.200 om4_bt_200.rfb
```

```
09-May-2002 10:03:40 INFO    Program Download started
09-May-2002 10:03:44 INFO    Program Download completed
BOOT upgrade was successfully completed
rs#reboot
This command will reset the whole system and disconnect your telnet
session.
Do you want to continue (y/n) [n] ?y
```

Updating prom image with a corrupted file

If the new PROM image is saved to a corrupted file, the download fails. An error message informing that Boot upgrade failed displays.

```
BOOT Software Version Prom-2.00.03 Built 08-Apr-2002 16:45:25
Processor: MPC8245 Rev 0.12, 266 MHz (Bus: 133MHz), 64 MByte SDRAM.
I-Cache 16 KB, linesize 32.D-Cache 16 KB, linesize 32.
Cache Enabled.

rs#system promimage upgrade 16.1.1.200 om4_bt_200.rfb
09-May-2002 10:08:13 INFO    Program Download started
09-May-2002 10:08:15 WARNING Boot1 programming failed, it is not valid
any more
09-May-2002 10:08:15 INFO    Program Download completed
BOOT upgrade was failed
rs#
```

To recover from the corrupted file:

1. Reboot and check that the device boots with the old **.rfb** file, to ensure that the corrupted file was not downloaded.
2. Check date and time.

2.2 SYSTEM IMAGE

The system image is saved in **arc** format and is mirrored in two Flash sectors. The Flash sectors for storing the system image are called *banks* (bank 0 and bank 1). One bank called *active* stores the active copy; another bank stores a second copy. The device boots and runs from a decompressed image from the *active* bank. A user can choose manually what bank will be active after the device is rebooted.

2.2.1 Updating System Image

Updating system image

To download the new system image (arc file) through TFTP Server:

1. Place the new System image on the TFTP server. The System image file has an **.arc** extension.
2. At the RCLI prompt, enter Enable.
3. Enter **system image add** *<tftp://IP address>>/<file.arc>*. The following is an example of the syntax.

```
system image add tftp://4.1.1.2/omega4.arc
```

After the System image is downloaded, the **.arc** file is saved in either Bank 0 or Bank 1.

4. Enter **reboot**. The device reboots.
5. Enter **system image list**. The following is an example of the information that displays:

```
TWO BANKS IN FLASH

rs#system image list
Images (banks) currently available on the FLASH
bank number  bank status
bank0         -  non active
bank1         -  active    [selected for next boot]
rs#
```

Choosing the Active Bank

Two separate System image versions can be stored. To choose an active bank:

1. Enter **system image choose [bank0| bank1]** from the Enabled mode. The following is an example of the information that displays:

```
rs#system image list
Images (banks) currently available on the FLASH
bank number  bank status
bank0         -  active    [selected for next boot]
bank1         -  non active
rs#system image choose bank1
```

```

Bank1 will run next time after reboot.
rs#system image list
Images (banks) currently available on the FLASH
bank number  bank status
bank0         -  active
bank1         -  non active    [selected for next boot]
rs#

```

2. Enter reboot. ES 500 reboots and runs the latest downloaded software version.

Updating system image with wrong or corrupted arc file

If an attempt is made to download a wrong file (a file in a wrong format), the download fails with the error message informing that the file is not in the arc format. The following figure displays an example of this message.

```

Autoboot in 2 seconds - press RETURN or Esc. to abort and enter prom.

Loading Kernel.

This file is not an arc format!!

Starting from backup bank

```

No remedy is required: the system boots automatically from the backup file and the wrong file is not downloaded into the system.

If an attempt is made to download a corrupted file, the download fails with the error message informing that the file is corrupted. The following figure displays an example of this message.

```

----- Performing the Power-On Self Test (POST) -----

UART Channel Loopback Test.....PASS

Testing the System SDRAM.....PASS
EPROM Checksum Test.....PASS

Flash Image Validation Test.....PASS
Testing CPU PCI Bus Device Configuration.....PASS
The RS code is NOT valid!
Trying to use second bank!
The RS code is NOT valid!
Trying to use second bank!

```

No remedy is required: the system boots automatically from the backup file and the corrupted file is not downloaded into the system.

Updating system image with both Banks corrupted

In the event where the system image Bank 0 and Bank 1 are corrupted, perform the following:

1. Connect the device via the ASCII terminal.
2. Switch the ES500 off, wait a few seconds and switch it on. The terminal will prompt you to download software.
3. From the HyperTerminal Menu Bar, click **Transfer**.
4. From the Transfer menu, click **Send File**. The “Send File” window is displayed.




Figure 2-1 The Send File window

5. Ensure that **Xmodem** is selected in the *Protocol* field.
6. Click **Browse** and locate the software installation file, for example, eos1000_c01.arc.
7. Click **Send**. The “Xmodem File Send” window displays the progress of the downloading process.

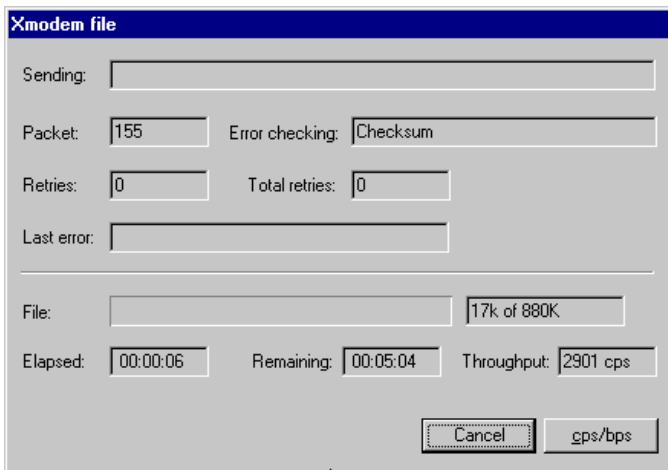


Figure 2-2. The Xmodem Send File window



Note: It is common to get “Got verify request” error message. The Xmodem protocol is designed to manage and solve this message.

A completion message is displayed. The ES500 will reboot and come up again with the new software image in Bank 0.

8. Ensure the new System image on the TFTP server. The System image file has an **.arc** extension.
9. At the RCLI prompt, enter **Enable**.
10. Enter **system image choose bank 1**.
11. Enter **system image add <tftp://IP address>>/<file.arc>**.
12. After the System image is downloaded enter **reboot**. The device reboots and the new system image is placed on Bank 1.

3 MAINTAINING CONFIGURATION FILES

This chapter provides information about configuration files in the Riverstone ES 500 Switch Router. It explains the different types of configuration files and procedures for changing, displaying, saving, and backing up the files.

3.1 CONFIGURATION FILES

The *Riverstone ES 500 Switch Router Getting Started Guide* introduced the following configuration files used by the ES 500:

- **Startup** – The configuration file, which the ES 500 uses to configure itself when the system is powered up. The Startup configuration remains intact even when the system is rebooted.
- **Active** – The commands from the Startup configuration file and any configuration commands that you have made active from the Scratchpad. The active configuration remains in effect until you switch off or reboot the system.



Caution

The active configuration remains in effect only during the current power cycle. If you switch off or reboot the ES 500 without saving the active changes to the Startup configuration file, these changes are lost.

- **Scratchpad** – The configuration commands you have entered during a CLI session. These commands are temporary and do not become active until you explicitly make them a part of the active configuration.

Because some commands depend on other commands for successful execution, the ES 500 Scratchpad simplifies system configuration by allowing you to enter configuration commands in any order, even when dependencies exist. When you activate the commands in the Scratchpad, the ES 500 sorts out the dependencies and executes the commands in the proper order.

The following figure illustrates the configuration files and the commands you can use to save your configuration:

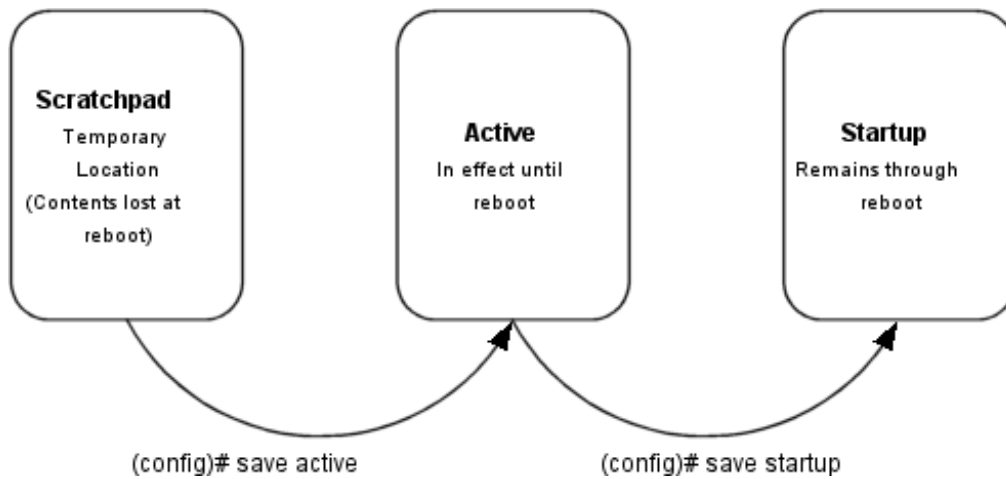


Figure 3-1 Commands to save configurations

3.1.1 Changing Configuration Information

The ES 500 provides many commands that help the user change the configuration information.

Use the **negate** command on a specific line of the active configuration to “disable” a feature or function, which has been enabled.

Commands used to change configuration information is as follows:

Enable Mode:	
Copy between Scratchpad, active configuration, startup configuration, or TFTP server (you cannot copy from Startup to Active, though)	copy <source> to <destination>
Configure Mode:	
Erase commands in the Scratchpad.	erase scratchpad
Erase the Startup configuration.	erase startup
Negate one or more commands by line numbers.	negate <line number>
Negate commands that match a specified command string.	no <string>
Save Scratchpad to active configuration.	save active
Save Active configuration to startup.	save startup

3.1.2 Displaying Configuration Information

The following table lists the commands that are useful for displaying the ES 500's configuration information.

Commands used to display configuration information is as follows:

Enable Mode:	
Show active configuration of the system.	<code>system show active-config</code>
Show the non-activated configuration changes in the Scratchpad.	<code>system show scratchpad</code>
Show the startup configuration for the next reboot.	<code>system show startup-config</code>
Configure Mode:	
Show active configuration of the system.	<code>show active</code>
Show the non-activated configuration changes in the Scratchpad.	<code>show scratchpad</code>
Show the Startup configuration for the next reboot.	<code>show startup-config</code>
Show the Active system configuration, followed by the non-activated changes in the Scratchpad.	<code>show</code>

3.1.3 Commenting out the contents of the Active via Scratchpad

When you want to comment out one or several lines in the Active file, you can use `comment out <line_number>` command in Config mode. Note, though, that if you wish this change to take effect, you need to commit these changes using `save active` command. Otherwise it does not affect the active configuration.

Example:

```
60 : snmp set target 176.240.10.33 community public status enable
61 : snmp set target 176.240.10.33 community asd status disable
rs(config)#comment out 61
rs(config)#sh
Running system configuration:
60 : snmp set target 176.240.10.33 community public status enable
61 : snmp set target 176.240.10.33 community asd status disable
***** Non-committed changes in Scratchpad *****
10: COMMENT OUT snmp set target 176.240.10.33 community asd status
disable
rs(config)#
```

3.1.4 Activating the Configuration Commands in the Scratchpad

The configuration commands you have entered are in the Scratchpad but have not yet been activated. Use the following procedure to activate the configuration commands in the Scratchpad.

1. Ensure that you are in the Enabled mode by entering the **enable** command in the CLI.
2. Ensure that you are in the Configure mode by entering the **configure** command in the CLI.
3. Enter the following command:

```
save active
```

The CLI displays the following message:

```
Do you want to make the changes Active? [y]
```

4. Type **y** to activate the changes

**Note**

When exiting the Configure mode (enter either **exit** or press **Ctrl+Z**), the CLI requests to make the Scratchpad changes active. The Scratchpad contents are added to the Active file.

**Note**

When you accept the changes to the ES 500 configuration, which are stored in the Scratchpad, and save them in the Active file (using the **save** or **copy** command, the contents of the Scratchpad are emptied.

3.1.5 Saving the Active Configuration to the Startup Configuration File

After you save the configuration commands in the Scratchpad, the control module executes the commands and makes the corresponding configuration changes to the ES 500. However, if you switch off or reboot the ES 500, the new changes are lost. Use the following procedure to save the changes in the Startup configuration file so that the ES 500 reinstates the changes when you reboot the software.

1. Ensure that you are in the Enabled mode by entering the **enable** command in the CLI.
2. Enter the following command to copy the configuration changes in the Active configuration to the Startup configuration:

```
Copy active to startup
```

3. When the CLI displays the following message, enter **yes** to save the changes.

```
Are you sure you want to overwrite the Startup configuration? [n]
```

**Note**

You also can save active changes to the Startup configuration file from within Configure mode by entering the **save startup** command.

The new configuration changes are added to the Startup configuration file stored in the control module's boot flash.

3.1.6 Viewing the Current Configuration

To view the current configuration:

1. Ensure that you are in the Enabled mode by entering the **enable** command in the CLI.
2. Enter the following command to display the status of each command line:

```
system show active-config
```


4 CLI BASICS

This chapter provides basic information about the Command Line Interface (CLI) and the ES 500. It includes the following:

- Information about CLI command modes
- Information about CLI commands
- How to get help
- How to set CLI parameters
- How to set up a basic configuration for the CLI and the ES 500

4.1 UNDERSTANDING CLI COMMAND MODES

The CLI has three separate command modes. Each command mode controls a group of related commands. This section explains the primary uses for each command mode.

4.1.1 Accessing the command mode

You can connect to ES 500 in 2 ways:

- Directly via COM port on the PC
- Via network using **telnet** application

When you connect to the ES 500 via a serial port and boot it up, you receive a command prompt:

```
rs>
```

The ES 500 doesn't require you to enter any password in order to switch between command modes. You can use **system set password** command in the Configure mode to specify passwords for each of the command modes. Refer to *Riverstone Networks ES 500 Switch Router Command Line Reference Guide* for more information.

If you want to use **telnet** to connect to the ES 500 remotely, you need to know the IP address of the ES 500 switch router. To access ES 500's command mode from a remote computer, type 'telnet <IP address>' and press Enter.

```
C:\>telnet <IP address>
```

You receive to the command prompt ">". To switch to the user mode, type 'rs' at the prompt and press Enter.

```
>rs
SYS-W-NOPASSWORD, no password for login, use 'system set password' to add one
rs>
```

4.1.2 User Mode

The initial mode on the ES 500 after booting up is the user mode. The user mode commands are a subset of the Enabled mode commands. In general, the user commands display basic information and contain basic utilities such as PING. This mode also contains the command to enter the enable mode.

The user mode command prompt consists of **rs** followed by the angle bracket (>), as shown in the following:

```
rs>
```

4.1.3 Enable Mode

The Enabled mode provides more commands than the user mode. Commands within the Enabled mode enable you to display:

- Router configuration
- Access Control Lists
- Telnet
- DHCP
- Help
- Configure (to enter the Configure mode)

To enter the Enabled mode from the user mode, type **enable** at the user command prompt. This mode also contains the command to enter the Configure mode.

The Enabled mode command prompt consists of **rs** followed by the pound sign (#):

```
rs#
```

To exit the Enabled mode and return to the user mode, either type **exit** and press the Return key, or press Ctrl+Z.

4.1.4 Configure mode

The Configure mode provides the capability of configuring and displaying all features and functions on the ES 500. The commands in this mode are *persistent* for the current session and future sessions. That is, they can be saved not only in onboard memory but also the startup configuration file. The startup configuration file is the file that the ES 500 boots from. For detailed information about the startup configuration file, see Chapter 3.

To enter the Configure mode, type **config** at the Enabled mode command prompt.

The Configure mode command prompt consists of **rs** followed by (**config**) and a pound sign (#):

```
rs(config)#
```

To exit the Configure mode and return to Enable mode, either type **exit** and press the Return key, or press Ctrl+Z.

4.2 UNDERSTANDING CLI COMMANDS

CLI commands are grouped by facilities. For example, the set of commands used to configure and display IP routing table information all start with **ip**, signifying the **ip** facility. Within the set of **ip** commands are commands such as:

- **set** – Used to configure parameters.
- **show** – Used to display.
- **start** – Used to start protocols operating.
- **stop** – Used to stop protocols from operating.
- **configure** – Used to configure parameters.

The complete set of commands for each facility is described in the *Riverstone ES 500 Switch Router Command Line Interface Reference Manual*.



Note

Some CLI modes may have commands that relate to the same facilities. For example, the Configure mode has **port** commands. These **port** commands are for making configuration changes to ports on the ES 500. The Enabled mode also has **port** commands. The **port** commands found in this mode display port statistics. They are not designed to change the way a port functions like the **port** commands in the Configure mode do.

4.3 GETTING HELP WITH CLI COMMANDS

Interactive help is available in the CLI. Invoke help by entering a question mark (?) character at any command prompt, or after a keyword in any mode. A set of facility names is displayed. These are the facilities and commands that can be used in that particular mode. Following is an example of typing ? at the enable prompt:

```
rs> ?
enable
Enable privileged user mode
exit
Exit current mode
help
Describe online help facility
l2-tables show all-macs
  -Show all MAC addresses currently in the L2 tables
l2-tables show igmp-mcast-registrations
  -Show info about multicasts registered by IGMP
l2-tables show mac
  -shows information about a particular mac address
l2-tables show mac-table-stats
  -shows statistics for the MAC address table
logout
  - logs off the system
ping
  - Ping utility
```

4.4 SETTING SYSTEM COMMANDS

The ES 500 provides commands for setting the following basic system parameters:

- System time and date
- System name
- System location
- Contact name (the person to contact regarding this device)

For more information about these commands see the *Riverstone ES 500 Switch Router Getting Started Guide*.

4.5 LINE EDITING COMMANDS

The ES 500 provides some line-editing capabilities that are similar to EMACs, a Unix text editor. For example, you can use certain line editing keystrokes to move forward or backward on a line and delete portions of a line. To use the line editing commands, you need to have a VT-100 terminal or terminal emulator. The line editing commands that you can use with CLI are detailed in the following table.

ctrl+a	Move to the beginning of line
ctrl+e	Move to the end of line
ctrl+u	Delete line from the beginning of line to cursor
ctrl+z	If inside a subsystem, it exits back to the top level. If in Enable mode, it exits back to User mode. If in Configure mode, it exits back to Enable mode.

4.6 CLI AND ES 500 CONFIGURATION EXAMPLE

The configuration demonstrates how to set:

- The system time and date
- The system name
- The system location
- The contact name (the person to contact regarding this device)

Here are the commands:

```

1. Set the date and time.
rs#system set date year 2001 month 4 day 30 hour 1 min 0 second 0
2. Change modes.
rs#configure
3. Name the device.
rs(config)#system set name "mktg-rs"
4. Set the location.
rs(config)#system set location "Santa Cruz, CA"
5. Identify the contact name.
rs(config)#system set contact "Jim Cale"

```

4.7 MONITORING SYSTEM SETTINGS

The following set of commands allows you to display in the Enabled mode information about various ES 500 parameters, settings, and resources.

Display information about system's resources	<code>system show capacity all chassis task</code>
Show contact information of the administrator of the device.	<code>system show contact</code>
Display the system date and time.	<code>system show date</code>
Display the location of the device.	<code>system show location</code>
Show system name	<code>system show name</code>
Display the time that has elapsed since the last reboot, and the system time and date when the last reboot occurred.	<code>system show uptime</code>
Display the software version running.	<code>system show version</code>

5 PORT MANAGEMENT

The following chapter provides information on commands that allow you to set and monitor different modes and parameters of the ES 500 ports.

5.1 SETTING PORT MODES AND PARAMETERS

All ports on the ES 500 are enabled on the startup and can send and receive traffic. They also all are set to auto-negotiation mode. The auto-negotiation is a process whereby ports on both sides of a connection resolve the best line speed, duplex mode and flow control scheme to communicate with each other.

5.1.1 Enabling and Disabling Ports

As it was already mentioned, all ports on ES 500 are enabled upon startup. To disable a certain port and prevent it from receiving and sending traffic, use the following command in the Configure mode:

```
rs(config)#port disable <port-list> | all-ports
```

This example shows that you have an ability to shut down a single port, a number of ports, or all ports on the device.

To activate the port and to enable it to send and receive traffic again, use the `port enable` command in the Configure mode.

```
rs(config)#port enable <port-list> | all-ports
```

5.1.2 Setting Port Parameters

You can set several port parameters, such as duplex, auto-negotiation mode, and speed, using the `port set` command in the Configure mode.

The command has the following syntax:

```
Set port mode and parameters    port set <port-list >|all-ports [auto-negotiation  
on|off] [duplex full|half] [flowctl off|enRx|enTx|both]  
[speed 10Mbps|100Mbps|1000Mbps]
```

The parameters are as follows:

- **auto-negotiation on | off** - Turns on or off auto-negotiation. Auto-negotiation is a process whereby both ports on a connection resolve the best line speed, duplex mode and flow control to communicate with each other.
- **duplex full | half** - Sets the operating mode to full or half duplex on Giga full.
- **Speed** - Sets the speed on a fast Ethernet or Gigabit Ethernet port. The speed can be set to 10Mbps, 100 Mbps, or 1000 Mbps on Giga Fixed 1000Mbps.
- **flowctl off | enRx | enTx | both** - Sets the flow control.
 - **off** - Turns off flow-control.
 - **enRx** - Acts on received PAUSE flow control frames.
 - **enTx** - PAUSE flow control frames are allowed to be transmitted.
 - **both** - Sets on both above options.

5.2 MONITORING PORTS

ES 500 features an extensive list of commands in the Enabled mode that allow you to monitor port information. The following table lists all the available commands and explains their action:

Displays auto-negotiation information. This command displays port number, administrative status, current status, and remote signaling.	<code>port show auto-negotiate</code>
Displays auto-negotiation capabilities. This command displays a list of port capabilities, advertised capabilities, and any received capabilities from another port.	<code>port show autonegotiation-capabilities</code>
Displays the user-defined description for device ports.	<code>port show description</code>
Displays current and admin. flow-control status	<code>port show flowctl</code>
Display port mirroring status information	<code>port show mirroring status</code>
Displays port status information by ports	<code>port show port-status</code>
Displays the Spanning Tree information for ports	<code>port show stp-info</code>

6 BRIDGING CONFIGURATION

The Riverstone ES 500 Switch Router provides the following bridging functions:

- Compliance with the IEEE 802.1d standard
- Compliance with the IGMP snooping standard
- Wire-speed address-based bridging
- Ability to logically segment a transparently bridged network into virtual local-area networks (VLANs), based on various criteria (e.g. physical ports)
- Frame filtering based on MAC address for bridged and Multicast traffic
- Integrated routing and bridging, which supports bridging of intra-VLAN traffic and routing of inter-VLAN traffic
- The ES 500 supports up to 8186 MACS in the bridging table.

6.1 SPANNING TREE (IEEE 802.1D)

Spanning tree (IEEE 802.1d) allows bridges to dynamically discover a subset of the topology that is loop-free. In addition, the loop-free tree that is discovered, contains paths to every LAN segment.

6.2 SUPPORTED BRIDGING MODES

The ES 500 provides support for **Address-based bridging**.

The ES 500 performs this type of bridging by looking up the destination address in the L2 lookup table. The L2 lookup table indicates the exit port(s) for the bridged packet. If the packet is addressed to the ES 500's own MAC address, the packet is routed rather than bridged.

6.3 VLAN SUPPORT

Virtual LANs (VLANs) are a means of dividing a physical network into several logical (virtual) LANs.

VLANs are mainly used for broadcast containment. A layer-2 (L2) broadcast frame is normally transmitted all over the bridged network. By dividing the network into VLANs, the *range* of a broadcast is limited, i.e., the broadcast frame is transmitted only to the VLAN to which it belongs. This significantly reduces the broadcast traffic on a network.

The type of VLAN depends on one criterion: how a received frame is classified as the one belonging to a particular VLAN.

ES 500 in its current state supports only **Port-Based VLANs**. Ports of L2 devices (switches, bridges) are assigned to VLANs. Any traffic received by a port is classified as belonging to the VLAN to which the port itself belongs. For

example, if ports 1, 2, and 3 belong to the VLAN named “Marketing”, then a broadcast frame received by port 1 is transmitted to ports 2 and 3; it is not transmitted to any other port.

**Note**

The number of VLANs available to the user is 224. The total number of VLANs supported by the ASIC is 255, but the Op-ROS blocks 32 for internal use: 24 Fast Ethernet ports, 2 Gigabit ports, and 6 auto-created trunks.

6.3.1 VLANs and the Device

VLANs are an integral part of the ES family of switching routers. The ES 500 switching routers can function as layer-2 (L2) switches as well as fully functional layer-3 (L3) routers. Hence they can be viewed as a switch and a router in one box. To provide the maximal performance and functionality, the L2 and L3 aspects of the ES switching routers are tightly coupled.

The ES 500 can be used purely as an L2 switch. Frames arriving at any port are bridged and not routed. In this case, setting up VLANs and associating ports with VLANs is all that is required.

The ES 500 can also be used purely as a router, i.e., each physical port of the ES 500 is a separate routing interface. Packets received at any interface are routed and not bridged. In this case, no VLAN configuration is required.

6.3.2 Configuration Examples

VLANs are used to associate physical ports on the ES 500 with connected hosts that may be physically separated but need to be a part of the same broadcast domain. To associate ports to a VLAN, you must first create a VLAN and then assign ports to it. This section shows examples of creating a port based VLAN.

Creating a VLAN

In this example, servers connected to ports gi.1.(1-2) on the ES 500 need to communicate with clients connected to et.2.(1-8). You can associate all the ports containing the clients and servers to a port based VLAN called ‘BLUE’.

First, create a port-based VLAN named ‘BLUE’

```
rs(config)#vlan create BLUE port-based id 2
```

Next, assign ports to the ‘BLUE’ VLAN.

```
rs(config)#vlan add ports et.2.(1-8),gi.1.(1-2) to BLUE
```

If you want to explicitly specify that certain ports must not be added to any VLANs, you can use the `vlan forbid ports` command. This command has the following syntax:

```
rs(config)#vlan forbid ports <port-list> from <vlan-name>
```

<port-list> specifies the list of ports that should NOT be added to a VLAN called <vlan-name>

6.4 ACCESS PORTS AND TRUNK PORTS (802.1P AND 802.1Q SUPPORT)

The ports of an ES 500 can be classified in two types, based on VLAN's functionality: **access ports** and **trunk ports**. By default, a port is an access port.

Trunk ports (802.1Q) are usually used to connect one VLAN-aware switch to another. They carry the traffic, which belongs to several VLANs. For example, suppose that ES 500 A and B are both configured with VLANs V1 and V2.

Then a frame arriving at a port on the ES 500 A must be sent to the ES 500 B, if the frame belongs to VLAN V1 or to VLAN V2. Thus, the ports on the ES 500 A and B, which connect the two devices together, must belong to both VLAN V1 and VLAN V2. Also, when these ports receive a frame, they must be able to determine whether the frame belongs to V1 or to V2. This is accomplished by "tagging" the frames, i.e., by pre-pending information to the frame in order to identify the VLAN to which the frame belongs. In the ES 500 switching routers, trunk ports normally transmit and receive tagged frames only (The format of the tag is specified by the IEEE 802.1Q standard.).

6.4.1 Setting a VLAN Access port

To specify an access port for a VLAN, use the `vlan make access-port` command. Execution of this command for a specific port turns that port into a VLAN access port.

For example:

```
rs(config)#vlan make access-port et.2.4
```

6.4.2 Configuring VLAN Trunk Ports

The ES 500 supports standards-based VLAN trunking between multiple devices as defined by IEEE 802.1Q. 802.1Q adds a header to a standard Ethernet frame, which includes a unique VLAN ID per trunk between two devices. These VLAN IDs extend the VLAN broadcast domain to more than one device.

To configure a VLAN trunk, enter the following command in the Configure mode:

```
Configure 802.1Q VLAN trunks.      vlan make trunk-port <port-list>
```

6.5 CONFIGURING SPANNING TREE

By default, spanning tree is disabled on all ports of the ES 500. To enable the spanning tree on the ports of the ES 500, you need to perform the following task on the ports where you want the spanning tree enabled:

```
Enable spanning tree on one or more ports for default spanning tree.      stp enable port <port-list>
```

BPDU's are frames carrying STP configuration information that switches and bridges normally broadcast. You may wish to prevent their reception by filtering them on ports where STP is disabled.

To filter out BPDUs on the port when the STP is disabled, use the `stp filter-bpdu` command.

```
rs(config)#stp filter-bpdu <port-list>
```

6.5.1 Adjusting Spanning-Tree Parameters

You may need to adjust certain spanning-tree parameters if the default values are not suitable for your bridge configuration.

Parameters affecting the entire spanning tree are configured with variations of the bridge global configuration command. Interface-specific parameters are configured with variations of the **bridge-group interface** configuration command.

**Note**

Only network administrators with a good understanding of how bridges and the Spanning-Tree Protocol work should make adjustments to spanning-tree parameters. Poorly chosen adjustments to these parameters can have a negative impact on the performance. IEEE 802.1d specification is a good source of information on bridging.

Setting the Bridge Priority

You can globally configure the priority of an individual bridge to configure the likelihood that a bridge will be selected as a root bridge. The lower the bridge's priority, the more likely the bridge is to be selected as a root bridge. According to the adopted standards, the default priority is 32,768; however, you can change it.

To set the bridge priority, enter the following command in the Configure mode:

Set the bridge priority for the default spanning tree.	<code>stp set bridging priority <num></code>
--	--

Setting a Port Priority

You can set a priority for an interface, to determine the likelihood of a certain port to be part of the STP path. The lower the priority the more likely that the port will be in the path.

To set an interface priority, enter the following command in the Configure mode:

Establish a priority for a specified interface for the default spanning tree.	<code>stp set port <port-list> priority <num></code>
---	--

Assigning Port Costs

Each interface has a port cost associated with it. The default is determined according to the standard based on port speed/media type that is currently configured on the port. You can set different port costs.

To assign port costs, enter the following command in the Configure mode:

Set a different port cost other than the defaults for the default spanning tree.	<code>stp set port <port-list> port-cost <num></code>
--	---

**Note**

Keep in mind that STP does not have the ability to dynamically change the port cost based on the media type.

Adjusting Bridge Protocol Data Unit (BPDU) Intervals

You can adjust BPDU intervals as described in the following sections:

- Adjust the Interval between Hello BPDUs
- Define the Forward Delay Interval
- Define the Maximum Idle Interval

Adjusting the Interval between Hello Times

You can specify the interval between “Hello” times. To adjust this interval, enter the following command in the Configure mode:

Specify the interval between hello times for the default spanning tree.	<code>stp set bridging hello-time <num></code>
---	--

Defining the Forward Delay Interval

The forward delay interval is the amount of time spent listening for topology change information after an interface has been activated for bridging and before forwarding actually begins. To change the default interval setting, enter the following command in the Configure mode:

Set the forward delay interval for the default spanning tree.	<code>stp set bridging forward-delay <num></code>
---	---

Defining the Maximum Age

If a bridge does not hear BPDUs from the root bridge within a specified interval, it assumes that the network has changed and re-computes the spanning-tree topology. To change the default interval setting, enter the following command in the Configure mode:

Change the amount of time a bridge waits to hear BPDUs from the root bridge for the default spanning tree.	<code>stp set bridging max-age <num></code>
--	---

STP Dampening

STP creates a loop-free, active topology in a network by placing ports in forwarding or blocking state. When a port moves to the forwarding state, it switches from listening to learning and then to forwarding. Whenever this transition happens, there is a chance that some traffic will be lost. If this port state transition happens rarely, the traffic loss is insignificant. On the other hand, if this happens frequently, it can adversely affect the network. STP dampening addresses this issue.

When a root port stops receiving BPDUs from the root bridge, a new root port is selected. If the original root port starts receiving BPDUs from the root bridge once again and STP dampening is enabled on the port, traffic is not immediately switched back to it. Instead, the port is monitored until it satisfies a stability condition, which is that it receives a certain number of STP configuration BPDUs during a specified period of time. If the port satisfies this condition, then it is considered stable and the traffic is switched back to it. Otherwise, it remains in an unstable state and is continuously monitored until it satisfies the stability condition. This feature ensures that the ports are stable before they move to a forwarding state and the traffic is switched back to them.

Note that the dampening does not occur if the port, on which the STP dampening is enabled, goes down and comes up again. Instead, traditional STP configuration occurs.

To enable STP dampening on a port, enter the following command in the Configure mode:

Enable STP dampening on a port.	<code>stp set port <port-list> dampening enable</code>
---------------------------------	--



Note STP dampening cannot be used in conjunction with Rapid Spanning Tree Protocol.

The ES 500 has defaults for the period of time a port is monitored (10 seconds) and the number of BPDUs that need to be received (10) during this period. You can change these defaults by entering the following command in the Configure mode:

Set parameters for STP dampening.	<code>stp set bridging damp-monitor-time <seconds> damp-bpdu-count <number></code>
-----------------------------------	--

6.6 CONFIGURING A PORT-BASED VLAN

To create a port-based VLAN, perform the following steps in the Configure mode:

1. Create a port-based VLAN.
2. Add physical ports to a VLAN.

6.6.1 Creating a Port-Based VLAN

To create a VLAN, enter the following command in the Configure mode:

Create a VLAN.	<code>vlan create <vlan-name> port-based id <num></code>
----------------	--

6.6.2 Adding Ports to a VLAN

To add ports to a VLAN, enter the following command in the Configure mode:

Add ports to a VLAN.	<code>vlan add ports <port-list> to <vlan-name></code>
----------------------	--

6.7 MONITORING BRIDGING

The ES 500 displays bridging statistics and configurations.

To display bridging information, enter the following commands in the Enable mode:

Displays the MAC addresses currently in the device L2 tables.	<code>12-tables show all-macs [verbose] [vlan <VLAN-num>]</code>
---	--

Displays the Multicast MAC addresses that IGMP has	<code>12-tables show igmp-mcast-registrations [vlan <VLAN-num>]</code>
--	--

registered with the L2 tables.	
Displays the VLAN and the port number on which the specified MAC address reside.	<code>l2-tables show mac <MACaddr> vlan <VLAN-num></code>
Displays statistics for the MAC address tables on the individual.	<code>l2-tables show mac-table-stats</code>
Displays the Multicast MAC addresses that IGMP has registered with the L2 tables.	<code>l2-tables show vlan-igmp-status vlan <VLAN-num></code>
Show all MAC addresses currently in the L2 tables.	<code>l2-tables show all-macs</code>
Show information in the master MAC table.	<code>l2-tables show mac-table-stats</code>
Show information on a specific MAC address.	<code>l2-tables show mac</code>
Show all VLANs.	<code>vlan show</code>

6.8 GARP/GVRP

The Generic Attribute Registration Protocol (GARP) is a generic attribute dissemination mechanism. In the case of the GARP VLAN Registration Protocol (GVRP), the attribute is the VLAN ID (VID).

GVRP uses GARP Protocol Data Units (PDUs) to register and de-register VLAN IDs on ports. When you enable GVRP on the ES 500 and one of its ports receives a GVRP request for an existing VLAN, to which it does *not* belong, GVRP registers the VLAN ID on the port, effectively adding the port to the VLAN. For example, VLAN RED is configured on ports et.2.1 and et.2.2 of the device. Port et.2.3 receives a GVRP request from another port for VLAN RED, of which it is not a member. If GVRP is enabled on port et.2.3, it automatically joins a VLAN RED and pass the traffic for this VLAN. But if GVRP is *not* enabled on port et.2.3, VLAN registration does not occur, and the traffic for VLAN RED never reaches port et.2.3.

GVRP also provides a mechanism for dynamically creating and removing VLANs. When you turn the dynamic VLAN creation on and the ES 500 receives a request for a VLAN that does not exist on the device, GVRP dynamically creates that VLAN and adds the port that received the request.

GVRP propagates this VLAN information throughout the active topology, enabling all GVRP-aware devices to dynamically establish and update their knowledge of VLANs and their members, including the ports through which those members can be reached. (For details on GARP refer to IEEE 802.1d. For details on GVRP, refer to IEEE 802.1q.)



Note

GVRP only advertises GARP BPDUs or adds a port to a VLAN if the port is an 802.1q trunk port.

GARP/GVRP provides the following benefits:

- The administrator is not required to know ahead of time, which VLANs should be configured on the network.
- The administrator does not have to manually configure all VLANs on the network.
- It prunes unnecessary traffic if a VLAN goes down.

6.8.1 Running GVRP with STP

Anytime GARP /GVRP configures a VLAN or adds ports to a VLAN, this information needs to be propagated on all other ports that are part of the active topology. If STP is disabled, this includes all ports, except the input port. If STP is enabled, this includes all ports that are in the forwarding mode, except the input port.

6.8.2 Configuring GVRP

To configure GVRP on the ES 500, do the following:

1. Enable GVRP functionality on the ES 500 (GVRP is disabled on the ES 500 by default).
2. Enable GVRP on individual ports (GVRP is disabled on all ports on the ES 500 by default).

In order to enable the GVRP on the device, you need to use the **gvrp start** command in Configuration mode.

You can optionally set the following features by using the **garp** and **gvrp** commands described in the *Riverstone ES 500 Switch Router Command Line Interface Reference Manual*:

- Enable dynamic VLAN creation (This feature is disabled by default). When you enable this feature, VLANs are be created dynamically when there is a GVRP request for a VLAN that does not exist on the device. In addition, VLANs can be manually through the CLI.
- Set a port's registration mode to *forbidden*. Registration modes refer to whether VLAN IDs can be dynamically registered on a port. You can set a port's mode to *forbidden* to prevent it from being dynamically added to a VLAN. Setting a port to "*forbidden registration*" de-registers all dynamically registered VLANs and prevents further VLAN dynamic registration on the port.
- Set a port's status to *non-participating*. When you do so, the specified ports do not send GARP/GVRP PDUs.
- Change the default values for the following GARP/GVRP timers:
 - leave all timer default is 10,000 ms (must be higher than the leave timer)
 - leave timer default is 600 ms (When configuring the leave timer, its value should be at least three times that of the join timer)
 - join timer default is 200 ms

**Note**

For GARP/GVRP to operate properly, all layer-2 connected devices should have the same GARP/GVRP timer values.

6.8.3 Configuration Example

Consider the following configuration example:




Figure 6-1 Using GARP/GVRP on a network

Switches R4 and R5 pass the traffic between two networks. The example VLANs are as follows:

- VLAN RED on ports et.2.(1-3) on R1, and on et.2.(1-3) on R6.
- VLAN GREEN on ports et.2.(1-3) on R2, and on et.2.(1-3) on R7.
- VLAN BLUE on ports et.2.(1-3) on R3, and on et.2.(1-3) on R8.

No VLANs were configured on R4 and R5. Instead, dynamic VLAN creation is enabled. So when any of the edge routers (R1, R2, R3, R6, R7, or R8) sends a request for a VLAN to the core routers (R4 and R5), and the VLAN does not exist on the core routers, that VLAN is dynamically created on the port of the router that received the request.

For example, R4 receives a request for VLAN RED on port et.2.1. VLAN RED is created dynamically on port et.2.1 of R4. This is then propagated across the bridged LAN to all the other routers. If dynamic VLAN creation were not enabled on R4, it would have dropped the traffic for VLAN RED.

The following is the configuration for R1:

```

Create VLAN RED as a port-based VLAN and add ports to it.
vlan create red port-based
vlan add ports et.2.(1-3) to red
Enable GVRP
gvrp start
Enable GVRP on ports et.2.(1-3) of R1.
gvrp enable ports et.2.(1-3)

Define port et.2.1 as 802.1q trunk port (since it advertises GVRP).
vlan make trunk-port et.2.1

Ports et.2.2 and et.2.3 of R1 do not need to send GARP PDUs because they are connected to
devices that are not running GVRP. Therefore, we should set their status to non-participating.
gvrp set applicant-status non-participant ports et.2.(2-3)

```

The following is the configuration for R4:

```

Create SmartTRUNK.
smarttrunk create st.1 protocol no-protocol

Define ports et.2.(4-6) auto-negotiation off, full duplex, speed 100 Mbps
port set et.2.(4-6) auto-negotiation off, duplex full, speed 100 Mbps

```

Add ports to the SmartTRUNK.

```
smarttrunk add ports et.2.(4-6) to st.1
```

Enable GVRP

```
gvrp start
```

Define st.1 and et.2.(4-6) as 802.1q trunk port (to allow GVRP advertise and join).

```
vlan make trunk-port st.1,et.2.(4-6)
```

Enable GVRP on ports st.1, and et.2.(4-6).

```
gvrp enable ports st.1,et.2.(4-6)
```

Enable dynamic VLAN creation so when R1, R2, or R3 sends a request for a VLAN, it is dynamically be created on R4.

```
gvrp enable dynamic-vlan-creation
```

Note that because dynamic VLAN creation was enabled on R4, we did not have to manually configure any VLAN on R4.



Note

You can set ports to forbidden registration mode using the **gvrp set registration-mode** command. This mode de-registers all VLANS on the specified port and prevents any VLAN creation or registration on that port.



Note

You can set ports to forbidden registration mode using the **gvrp set registration-mode** command. This mode de-registers all VLANS on the specified port and prevents any dynamic VLAN creation or registration on that port.

6.8.4 Monitoring GVRP status

ES 500 features a set of commands that allow you to perform various monitoring of GVRP-enabled ports. The following table lists these commands and their parameters:

Display the applicant port status.	gvrp show applicant-status ports
Display the GVRP port errors.	gvrp show error-statistics
Display whether the ports are in normal registration mode, fixed registration mode, or forbidden registration mode.	gvrp show registration-mode ports
Display the port GVRP statistics.	gvrp show statistics
Display the port GVRP status.	gvrp show status

For additional information on GVRP commands, please refer to “*Riverstone Networks ES 500 Switch Router Command Line Reference Guide*”.

7 SMARTTRUNK CONFIGURATION

This chapter explains how to configure SmartTRUNKs on the ES 500. A SmartTRUNK is Riverstone's technology for load balancing and load sharing across a number of ports. SmartTRUNKs are used for building high-performance, high-bandwidth links between Riverstone's switching platforms. A SmartTRUNK is a group of two or more physical ports that have been combined into a single logical port. Multiple physical connections between devices are aggregated into a single, logical, high-speed path that acts as a single link. As flows are set up on the SmartTRUNK, traffic is balanced across all ports in the combined link, therefore balancing overall available bandwidth.

SmartTRUNKs can also interoperate with switches, routers, and servers from other vendors. SmartTRUNKs allow administrators the ability to increase bandwidth at congestion points in the network, eliminating potential traffic bottlenecks. SmartTRUNKs also provide improved data link resiliency – if one link in a SmartTRUNK fails, its flows are distributed among the remaining links.



Note

For detailed descriptions of the SmartTRUNK commands, see the “SmartTRUNK commands” section of the *Riverstone ES 500 Switch Router Command Line Interface Reference Manual*.

SmartTRUNKs are compatible with all ES 500 features, including VLANs, STP, and so on. SmartTRUNK operation is supported over different media types and a variety of technologies, including 10/100 Mbps Ethernet and Gigabit Ethernet.

7.1 CONFIGURING SMARTTRUNKS

SmartTRUNK is disabled by default upon the startup of the ES 500. The following are the steps for enabling and configuring a SmartTRUNK:

1. Make sure that all ports are of the same speed and media type, are in full duplex, and have autonegotiation disabled.
2. Create a SmartTRUNK and specify its control protocol.
3. Add physical ports to the SmartTRUNK.
4. Specify the load balancing policy. This policy determines how the flows are allocated on the SmartTRUNK's ports (this step is optional).

7.1.1 Creating a SmartTRUNK

When creating a SmartTRUNK, assign a name to the SmartTRUNK and specify its control protocol. ES 500 supports only **No Control Protocol** option at this stage.

Here is an example of creating a SmartTRUNK named **st.1**, which uses no control protocol:

```
rs(config)#smarttrunk create st.1 protocol no-protocol
```

7.1.2 Adding Physical Ports to the SmartTRUNK

You can add up to 8 of 10/100 Ethernet ports to a SmartTRUNK. If one link goes down, the traffic is redirected seamlessly to the remaining operational links.

SmartTRUNK Port Limitations

Ports added to a SmartTRUNK must meet the following criteria:

- Be regular ASIC ports
- To have no L3 interface defined on it.
- Not to belong to any VLAN.
- Not to belong to any trunk already.
- Not to be in the autoNegotiation mode.
- To be in the FULL DUPLEX mode.
- Same speed defined on all SmartTRUNK ports
- Not to be GVRP-enabled.

Here is an example of adding ports **et.2.1** through **et.2.8** to a SmartTRUNK:

```
rs(config)#smarttrunk create st.1 protocol no-protocol
rs(config)#port set et.2.(1-8) speed 100mbps
rs(config)#port set et.2.(1-8) duplex full
rs(config)#port set et.2.(1-8) auto-negotiation off
rs(config)#smarttrunk add ports et.2.(1-8) to st.1
```

7.1.3 Specifying Traffic Load Policy

The load balancing is implemented by defining per frame via which port it will be transmitted. The decision is made by calculating either the SRC MAC or the DEST MAC/IP or both.

Calculation: For each MAC/IP address you have to check only the last 3 binary digit. The options are from 0-7

Example for X ports trunk:

Mac ends with	3 ports Trunk	4 ports Trunk	7 ports trunk
0	1	1	1
1	2	2	2
2	3	3	3
3	1	4	4
4	2	1	5
5	3	2	6
6	1	3	7
7	2	4	1

Port 1 is equal to the first port you added to the trunk.

Port 2 is equal to the second port you added to the trunk

Port 3 is equal to the second port you added to the trunk.

In order to specify the load-balancing policy for the SmartTRUNK, use the following command:

```
smarttrunk set load-balancing <smarttrunk> <bridging | routing> [layer <num>] [UsedAddresses <notApplied | dstAddr | srcAddr | dstAndSrcAddr >]
```

The ES 500 performs IP load balancing (routing) based only on the destination IP.

A SmartTRUNK can be created and defined for either bridging or routing purposes. By default, load balancing is activated as follows:

- Bridging SmartTRUNKs based on Layer 2 destination address
- Routing SmartTRUNKs based on Layer 3 destination address.

Use the command above to specify the layer and address type for SmartTRUNK (bridging or routing) load balancing.



Note

- This command does not determine the type of the SmartTRUNK (bridging or routing) but rather the load balancing mechanism for each given type of SmartTRUNK.
- Currently ES 500 supports for routing SmartTrunks based only on Layer 3 destination addresses.

For additional information on using the `set load-balancing` command to specify the balancing criteria for a SmartTRUNK, please, refer to *Riverstone ES 500 Switch Router Command Line Interface Reference Manual*.

The traffic distribution policy only affects the initial assignment of L2 and L3 flows to a given port. If a link in the SmartTRUNK goes down, the flows are remapped to a different port in the same SmartTRUNK. If the flows assigned to a particular port in the SmartTRUNK exceed the bandwidth of the port, packets are dropped even if there is bandwidth available on the other ports in the SmartTRUNK.

7.2 MONITORING SMARTTRUNK CONFIGURATION

Use the following command to monitor the SmartTRUNK configuration on the ES 500:

Show information about the SmartTRUNK connection, including the MAC address of the remote switch, and the module number and port number of each remote port.	<code>smarttrunk show connections</code>
Show load-balancing criteria.	<code>smarttrunk show load-balancing</code>
Show information about the control protocol on a SmartTRUNK and the state of its ports.	<code>smarttrunk show protocol-state</code>
Show information about all SmartTRUNKs, including active and inactive ports, and the control protocol used.	<code>smarttrunk show trunks</code>

7.3 SMARTTRUNK CONFIGURATION EXAMPLE

Figure 7-1 illustrates a network design based on SmartTRUNKs. R1 is an ES 500 operating as a router, while R2 and R3 are ES 500s operating as switches.




Figure 7-1 SmartTRUNK configuration example

The following is the SmartTRUNK configuration for the ES 500 labeled 'R1' in the diagram:

```
smarttrunk create st.1,st.2,st.3 protocol no-protocol
port set et.2.(1-8) auto-negotiation off duplex full speed 100mbps
smarttrunk add ports et.2.(1-3) to st.1
smarttrunk add ports et.2.(4-6) to st.2
smarttrunk add ports et.3.(7-8) to st.3
interface create ip to-RS 38000 address-netmask 10.1.1.2/24 port st.1
interface create ip to-s1 address-netmask 11.1.1.2/24 port st.2
interface create ip to-s2 address-netmask 12.1.1.2/24 port st.3
```

The following is the SmartTRUNK configuration for the ES 500 labeled 'R3' in the diagram:

```
smarttrunk create st.2,st.4 protocol no-protocol
port set et.2.(1-8) auto-negotiation off duplex full speed 100mbps
smarttrunk add ports et.2.(1-3) to st.2
smarttrunk add ports et.2.(4-8) to st.4
```

The following is the SmartTRUNK configuration for the ES 500 labeled 'R2' in the diagram:

```
smarttrunk create st.3,st.5 protocol no-protocol
port set et.2.(1-4) auto-negotiation off duplex full speed 100mbps
smarttrunk add ports et.2.(1-2) to st.3
smarttrunk add ports et.2.(3-4) to st.5
```



Note

In the example, because R1 and R2 are operating only as switches (layer-2 traffic only), their SmartTRUNKs were not assigned to IP interfaces.

8 IP ROUTING CONFIGURATION

The ES 500 supports standards-based TCP, UDP, and IP. This chapter describes how to configure IP interfaces and general non-protocol-specific routing parameters.



Note ES 500 does not support setting of MTU.

8.1 IP ROUTING PROTOCOLS

The ES 500 supports standards-based Unicast and Multicast IP routing. Multicast routing protocols are used to determine how Multicast data is transferred in a routed environment.

8.1.1 Unicast Routing Protocols

Interior Gateway Protocols are used for routing networks that are within an “autonomous system”, a network of relatively limited size. All IP interior gateway protocols must be specified with a list of associated networks before routing activities can begin. A routing process listens to updates from other routers on these networks and broadcasts its own routing information on those same networks. The ES 500 supports the following Interior Gateway Protocols:

- Routing Information Protocol (RIP) Version 1, 2 (RFC 1058, 1723). Configuring RIP for the ES 500 is described in Chapter 9, “*RIP Configuration*”.
- Open Shortest Path First Routing (OSPF) Version 2, as defined in RFC 1583. Configuring OSPF for the ES 500 is described in Chapter 10, “*OSPF Configuration*”.

8.1.2 Multicast Routing Protocols

IP Multicasting allows a host to send traffic to a subset of all hosts. These hosts subscribe to the group membership, thus notifying the ES 500 of their participation in a Multicast transmission.

Multicast routing protocols are used to determine which routers have directly attached hosts, as specified by IGMP, that have membership to a Multicast session.

The ES 500 supports the Internet Group Management Protocol (IGMP) snooping.

8.2 CONFIGURING IP INTERFACES AND PARAMETERS

You can configure an IP interface to a single port or to a VLAN (or a Trunk). This section provides an overview of configuring IP interfaces.

Interfaces on the ES 500 are logical interfaces. Therefore, you can associate an interface with a single port or with multiple ports (or Trunks):

- To associate an interface with a single port, use the `port` option with the `interface create` command.
- To associate an interface with multiple ports, first create an IP VLAN and add ports to it, then use the `vlan` option with the `interface create` command.

The `interface create ip` command creates and configures an IP interface. Configuration of an IP interface can include information such as the interface's name, IP address, netmask, broadcast address, and so on. You can also create an interface in a disabled (down) state instead of the default enabled (up) state.



Note

You must use either the `port` option or the `vlan` option with the `interface create` command.

8.2.1 Configuring and Displaying IGMP Snooping

IGMP Snooping can be specified on a VLAN. By default, IGMP snooping is disabled on all VLANs.

The ES 500 uses IGMP to describe IGMP Snooping. When entering IGMP commands, they refer to IGMP Snooping functionality.

IGMP Snooping supports the following:

- 2K IGMP Snooping routes
- 255 IGMP Snooping sessions
- 2K IGMP Snooping FFT flows

The commands to configure IGMP Snooping is as follows:

Enable IGMP snooping on a specified VLAN.	<code>igmp enable vlan <vlan-name></code>
Set parameters for VLAN-based IGMP snooping.	<code>igmp set vlan <vlan-name> [host-timeout <num>] [querier-timeout <num>] [router-timeout <num> leave-timeout <num>] [filter-ports <port-list>] [permanent-ports <port-list>]</code>
Start IGMP snooping on enabled VLANs	<code>igmp start-snooping</code>

The command displays port, querier, Multicast group, and group membership information for each VLAN. To display IGMP Snooping configuration is as follows:

```
igmp show vlans [detail] [name <name>] [timers]
```

8.2.2 Configuring IP Interfaces to Ports

An IP interface can be directly assigned to a physical port. For example, to assign an IP interface 'RED' to the physical port et.2.4, enter the following:

```
rs(config)#interface create ip RED address-netmask 10.50.0.1/255.255.0.0 port et.2.4
```

8.2.3 Configuring IP Interfaces for a VLAN

You can configure one IP interface per VLAN. Once an IP interface has been assigned to a VLAN, you can add a secondary IP address to the VLAN. To create a VLAN called IP3, add ports et.2.1 through et.2.4 to the VLAN, and then create an IP interface on the VLAN:

```
rs(config)#vlan create IP3 port-based
rs(config)#vlan add ports et.2.(1-4) to IP3
rs(config)#interface create ip int3 address-netmask 10.20.3.42/24 vlan IP3
```

8.2.4 Monitoring IP interface configuration

To display configuration information for an IP interface use the interface show command.

```
rs(config)#interface show ip <InterfaceName> | all
```

You can view the configuration information for a specific interface or for all IP interfaces.

8.3 CONFIGURING IP HELPER

The `ip helper-address interface` command allows a user to forward the specific UDP broadcast from one interface to another. Typically, the broadcast packets from one interface are not forwarded (routed) to another interface. However, some applications use the UDP broadcast to detect the availability of a service. Other services, for example BOOTP/DHCP, require broadcast packets to be routed so that they can provide services to the clients on another subnet. An IP helper can be configured on each interface to have the UDP broadcast packets forwarded to a specific host for a specific service or forwarded to all other interfaces.

You can configure the ES 500 to forward the UDP broadcast packets received on a given interface to all other interfaces or to a specified IP address. Specify a UDP port number for which the UDP broadcast packets with that destination port number is forwarded. By default, if no UDP port number is specified, the ES 500 forwards UDP broadcast packets for the following six services:

- BOOTP/DHCP (port 67 and 68)
- DNS (port 37)
- NetBIOS Name Server (port 137)
- NetBIOS Datagram Server (port 138)
- TACACS Server (port 49)
- Time Service (port 37)

To forward the UDP broadcast packets received on the interface int1 to the host 10.1.4.5 for the six default UDP services:

```
rs(config)#ip helper-address interface int1 10.1.4.5
```

To forward the UDP broadcast packets received on the interface int2 to the host 10.2.48.8 for the packets with the destination port 111 (port mapper):

```
rs(config)#ip helper-address interface int2 10.2.48.8 111
```

To forward the UDP broadcast packets received on the interface int3 to all other interfaces:

```
rs(config)#ip helper-address interface int3 all-interfaces
```

8.4 CONFIGURATION EXAMPLE

8.4.1 Assigning IP Interfaces

To enable routing on the ES 500, you must assign an IP interface to a VLAN. To assign an IP interface named 'RED' to the 'BLUE' VLAN, enter the following command:

```
rs(config)#interface create ip RED address-netmask 10.50.0.1/255.255.0.0 vlan BLUE
```

9 RIP CONFIGURATION

This chapter describes how to configure the Routing Information Protocol (RIP) on the Riverstone ES 500 Switch Router. RIP is a distance-vector routing protocol for use in small networks. RIP is described in the RFC 1723. A router running RIP broadcasts updates at set intervals. Each update contains paired values where each pair consists of an IP network address and an integer distance to that network. RIP uses a hop count metric to measure the distance to a destination.

The Riverstone ES 500 Switch Router provides support for RIP Version 1 and 2. The ES 500 implements plain text methods for RIP Version 2.

The ES 500 supports split-horizon and poison reverse. These features are enabled by default and cannot be disabled.

9.1 CONFIGURING RIP

RIP is disabled by default on ES 500 and on each of the defined IP interfaces. To configure RIP on ES 500, follow these steps:

1. Start the RIP process by entering the `rip start` command
2. Use `rip add interface` command to inform RIP about the attached interfaces.

9.1.1 Enabling and Disabling RIP

To enable or disable RIP, enter one of the following commands in the Configure mode:

Enable RIP.	<code>rip start</code>
Disable RIP.	<code>rip stop</code>

9.1.2 Configuring RIP Interfaces

To configure RIP in the ES 500, you must first add interfaces to inform RIP about the attached interfaces.

To add RIP interfaces, enter the following command in the Configure mode:

Add interfaces to the RIP process.	<code>rip add interface <interfacename-or-IPaddr></code>
------------------------------------	--

9.2 CONFIGURING RIP PARAMETERS

No further configuration is required, and the system default parameters are used by RIP to exchange routing information. These default parameters may be modified to suit your needs by using the `rip set interface` command.

RIP Parameter	Default Value
Version number	RIP v1
Check-zero for RIP reserved parameters	Enabled
Whether RIP packets should be broadcast or Multicast	Broadcast if version is 1 and Multicast if version is 2
Preference for RIP routes	60
Metric for incoming routes	1
Metric for outgoing routes	0
Authentication	None
Update interval	30 seconds

To change RIP parameters, enter the following commands in the Configure mode:

Set RIP Version on an interface to RIP V1.	<code>rip set interface <interfacename-or-IPaddr> all version 1</code>
Set RIP Version on an interface to RIP V2.	<code>rip set interface <interfacename-or-IPaddr> all version 2</code>
Change the metric on incoming RIP routes.	<code>rip set interface <interfacename-or-IPaddr> all metric-in <num></code>

9.3 MONITORING RIP

To monitor the RIP information, enter the following commands in the Enable mode:

Show all RIP information.	<code>rip show all</code>
Show RIP global information.	<code>rip show globals</code>
Show RIP information on the specified interface.	<code>rip show interface <Name or IP-addr></code>

9.4 CONFIGURATION EXAMPLE




Figure 9-1 RIP Configuration Example

Create interface R1-if1 with ip address 1.1.1.1/16 on port et.2.1 on R-1

```
rs (config)#interface create ip R1-if1 address-netmask 1.1.1.1/16 port et.2.1
```

Configure rip on R-1

```
rs (config)#rip add interface R1-if1
```

```
rs (config)#rip set interface R1-if1 version 2
```

Change default metric-in

```
rs (config)#rip set interface R1-if1 metric-in 2
```


10 OSPF CONFIGURATION

Open Shortest Path First Routing (OSPF) is a shortest path first link-state protocol. The ES 500 supports OSPF Version 2, as defined in RFC 1583. OSPF is an interior gateway protocol that distributes routing information between routers in a independent system. OSPF chooses the least-cost path as the best path. OSPF is suitable for complex networks with a large number of routers. OSPF provides equal-cost multi-path routing enabling single destination packets to be sent via more than one interface simultaneously.

In a link-state protocol, each router maintains a database that describes the entire Autonomous System (AS) topology, which it builds out of the collected link state advertisements of all routers. Each participating router distributes its local state (i.e., the router's usable interfaces and reachable neighbors) throughout the AS by flooding. Each multi-access network that has at least two attached routers has a designated router and a backup designated router. The designated router floods a link state advertisement for the multi-access network and has other special responsibilities. The designated router concept reduces the number of adjacencies required on a multi-access network.

OSPF allows networks to be grouped into areas. Routing information passed between areas is abstracted, potentially allowing a significant reduction in routing traffic. OSPF uses four different types of routes, listed in order of preference:

- Intra-area
- Inter-area
- Type 1 ASE
- Type 2 ASE

Intra-area paths have destinations within the same area. Inter-area paths have destinations in other OSPF areas. Both types of Autonomous System External (ASE) routes are routes to destinations external to OSPF (and usually external to the AS). Routes exported into OSPF ASE as type 1 ASE routes are supposed to be from interior gateway protocols (e.g., IS-IS) whose external metrics are directly comparable to OSPF metrics. When a routing decision is being made, OSPF adds the internal cost to the AS border router to the external metric. Type 2 ASEs are used for exterior gateway protocols whose metrics are not comparable to OSPF metrics. In this case, the external cost from the AS border router to the destination is used in the routing decision.

From OSPF to RIP and visa versa, there is advertisement leaking. This to enables RIP and OSPF advertisements to transparently advertise to each other the routes learnt.

ES 500 supports the following OSPF functions:

- Definition of areas, including stub areas
- Configuration of parameters at the area, interface, or global level. Parameters that can be configured include retransmission interval, interface transmit delay, router priority, router dead and hello intervals, and authentication key.

10.1 CONFIGURING OSPF

To configure OSPF on ES 500:

- Set the router ID.
- Enable OSPF
- Create the OSPF area
- Add interfaces to the area
- Optionally, configure parameters at the global, area, and/or interface level.



Note

Important: at least one IP interface in each area must be configured before OSPF can become functional.

10.2 SETTING THE ROUTER ID

The router ID uniquely identifies the ES 500. To set the router ID to be used by OSPF, enter the following command in Configure mode:

Set the ES 500 Router ID.	<code>ip-router global set router-id <hostname-or-IPaddr></code>
---------------------------	--

If you do not explicitly specify the router ID, then an ID is chosen implicitly by the ES 500. When the router ID changes, an OSPF router has to flush all its LSAs from the routing domain.

If you explicitly specify a router ID, then it would not change, even if all interfaces were to go down.

10.3 ENABLING OSPF

OSPF is disabled by default on the ES 500. To enable or disable OSPF, enter one of the following commands in Configure mode:

Enable OSPF.	<code>ospf start</code>
Disable OSPF.	<code>ospf stop</code>
Disable OSPF.	<code>no ospf start</code>



Note

Important: all IP interfaces must be attached to each configured area before OSPF can be started.

10.4 CONFIGURING OSPF AREAS

OSPF areas are a collection of subnets that are grouped in a logical fashion. Each area maintains its own link state database. The area topology is known only within the area. A router maintains a separate link state database for each area to which it is connected.

The ES 500 supports the configuration of multiple OSPF areas, as well as three special types of areas:

- **backbone** - The backbone is responsible for distributing routing information between non-backbone areas. OSPF areas communicate with other areas via the backbone area. The OSPF area backbone contains all area border routers (ABRs).
- **stub** - A stub area is not used as a transit area. Routers within a stub area route internal traffic only.

On ES 500, you can create multiple OSPF areas, but at least one of them should be an area backbone. To configure an OSPF areas, including a stub area, enter the following command in Configure mode. To configure a backbone area, use the **backbone** keyword with the following command:

Create an OSPF area.	<code>ospf create area <area-num> backbone</code>
----------------------	---

10.4.1 Configuring Stub Areas

Information about routes, which are external to the OSPF routing domain is not sent into a stub area. Instead, if the **stub-cost** parameter is specified, the ABR generates a default external route into the stub area for destinations outside the OSPF routing domain. The **stub-cost** specifies the cost to be used to inject a default route into a stub area. If this parameter is not specified, no default route is injected into the OSPF stub area.

To define an OSPF stub area, enter the following command in Configure mode.

Specify an OSPF area to be a stub area.	<code>ospf set area <area-num> stub [stub-cost <num>]</code>
---	--



Note

The `ospf set area` command sets all the listed attributes to their values: if the attribute pertains to OSPF interface, it is applied to all the interfaces currently attached to the area specified. When a new OSPF interface is attached to an area the area setting does not affect that new interface.

10.5 CONFIGURING OSPF INTERFACES

To configure an interface for OSPF, first configure an IP interface using the **interface create** command, and then add the interface to an OSPF area. To add an IP interface to an area enter the following command in Configure mode:

Add an interface to an OSPF area.	<code>ospf add interface <name-or-IPaddr> to-area <area-addr> backbone</code>
-----------------------------------	---



Note

Important: at least one IP interface must be configured before an OSPF area can become functional.

10.6 CONFIGURING OSPF INTERFACE PARAMETERS

ES 500 provides a number of parameters that are set at the interface level. To set OSPF interface parameters, enter the following command in Configure mode:

Set OSPF interface parameters.	<code>ospf set interface <name-or-IPaddr> all [state disable enable] [cost <num>] [no-multicast] [retransmit-interval <num>] [transit delay <num>] [priority</code>
--------------------------------	--

```
<num>] [hello-interval <num>] [router-dead-interval
<num>] [poll-interval <num>]
```

**Note**

The `ospf set interface` command affects all configured OSPF interfaces in all areas, overriding previous settings.

**Note**

An OSPF command of the type "add" or "set" is merged with the previous one (if any), acting on the same object.

10.6.1 Setting the Interface State

OSPF interfaces that are added to an area are enabled by default. You can disable them by using the `state disable` option with the `ospf set interface` command.

10.7 OSPF TOPOLOGY AND CONFIGURATION EXAMPLE

The following figure shows a simple topology with 2 routers and two hubs/switches. This example will illustrate how to create interfaces, configure OSPF areas, add interfaces to OSPF areas, and start OSPF.

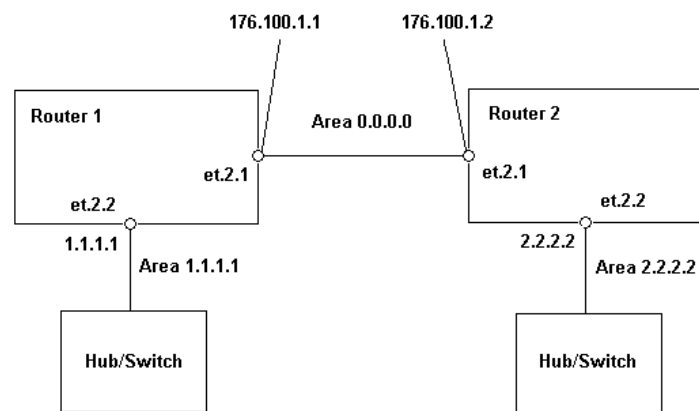


Figure 10-1 Example of OSPF topology

Here is a configuration of the first router:

```
rs (config)#interface create ip 11 address netmask 1.1.1.1 port et.2.2
rs (config)#interface create ip 12 address netmask 176.100.1.1 port et.2.1
rs (config)#ospf create area 0.0.0.0
rs (config)#ospf create area 1.1.1.1
rs (config)#ospf add interface 12 to area 0.0.0.0
rs (config)#ospf add interface 11 to area 1.1.1.1
rs (config)#ospf start
```

And this is the configuration for the second router:

```
rs (config)#interface create ip r1 address netmask 176.100.1.2 port et.2.1
rs (config)#interface create ip r2 address netmask 2.1.1.1 port et.2.2
rs (config)#ospf create area 0.0.0.0
rs (config)#ospf create area 1.1.1.1
rs (config)#ospf add interface r1 to area 0.0.0.0
rs (config)#ospf add interface r2 to area 1.1.1.1
rs (config)#ospf start
```


11 ACL CONFIGURATION

This chapter explains how to configure and use the Access Control Lists (ACLs) on the ES 500. ACLs are the lists of selection criteria for the specific types of packets. When used in conjunction with certain ES 500 functions, ACLs allow you to restrict the Layer-3/4 traffic going through the router.

This chapter contains the following sections:

- Section 10.1, "*ACL Basics*", explains how ACLs are defined and how the ES 500 evaluates them.
- Section 10.2, "*Creating and Modifying ACLs*", describes how to edit ACLs.
- Section 10.3, "*Monitoring ACLs*", lists the commands you can use to display information about ACLs active on the ES 500.

11.1 ACLS BASICS

An ACL consists of one or more *rules* describing a particular type of IP traffic. ACLs can be simple, consisting of only one rule, or complicated, with many rules. Each rule tells the ES 500 to either permit or deny the packets that match the selection criteria specified in the rule.

Each ACL is identified by a name. Such name can be a meaningful string, such as `denyftp` or `noweb` or it can be a number such as `100` or `101`.

For example, the following ACL has a rule that denies all IP packets from subnet 10.2.0.0/16 to go through the device:

```
rs (config)#acl 101 deny ip 10.2.0.0/16
```

11.1.1 Defining Selection Criteria in ACL Rules

Selection criteria in the rule are offsets inside a packet. In the example above, the selection criteria are IP packets with a Src Address = 10.2.0.0/16.

The selection criteria, which you can specify in an ACL rule depends on the type of ACL you are creating. For IP, TCP, and UDP ACLs, the following selection criteria can be specified:

- Source IP address
- Destination IP address
- Source port number
- Destination port number
- Type of Service (TOS)

These selection criteria are specified as *fields* of an ACL rule. The following syntax description shows the fields of an IP ACL rule:

```
rs (config)#acl <name> permit|deny ip <SrcAddr/Mask> <DstAddr/Mask>  
<SrcPort> <DstPort> <tos> <tos-mask> [accounting] <checkpoint interval>
```

**Note**

There are variations of the `acl permit|deny ip` command that allow you to restrict the traffic for a specific IP-based protocol (ICMP, IGMP, TCP, UDP); for example, the `acl permit|deny tcp` command allows you to restrict only the TCP traffic. These variants have the same syntax and fields as the `acl permit|deny ip` command.

Each ACL rule field is position sensitive. For example, A TCP traffic rule specifically states that the source address must be followed by the destination address, then the source socket, and then the destination socket. Not all the ACL rule fields must be specified. If a field is not specified, it is handled as a wildcard. Specified fields are matched against the packet information. Each protocol can have a number of different fields to match. For example, a rule for TCP can use socket port numbers.

Since each field is position sensitive, it may be necessary to ignore some fields in order to specify a value for another field. To ignore a field, use the keyword **any**. The following sample defines a ACL rule with following characteristics:

- `acl` — Specifies the command type.
- `error1` — Specifies the rule name.
- `deny` — Specifies the Access control action.
- `tcp` — Specifies the protocol type.
- `any` — Specifies that the field is ignored.
- `any` — Specifies that the field is ignored.
- `SMTP` — Specifies the ToS value.
- `SMTP` — Specifies the ToS mask value.

```
rs (config)#acl error11 deny tcp any any 255 255
```

Note that in the above example, the `<tos>` (Type of Service) field is not specified and is treated as a wildcard. The **any** keyword is needed only to skip a wildcard field in order to explicitly specify another field that is further down in the rule. If there are no other fields to specify, the **any** keyword is not necessary. For example, the following ACL permits all the IP traffic to go through:

```
rs (config)#acl yesip permit ip
```

11.1.2 How ACL Rules are Evaluated

For an ACL with several rules, the order of the rules is important. When the ES 500 checks a packet against an ACL, it goes through each rule in the ACL sequentially. When a match is found all subsequent rules are ignored. That is, a first-match algorithm is used. There is neither hidden or implied ordering of ACL rules, nor any precedence attached to each field. The ES 500 simply goes down the list, one rule at a time, until there is a match. Consequently, the rules that are more specific (that is, with more selection criteria) should always be listed ahead of the rules that are less specific.

11.2 USING ACLS

It is important to understand that an ACL is simply a definition of packet characteristics specified in a set of rules. An ACL must be *enabled* in one of the following ways:

- Applying an ACL to an interface, which permits or denies traffic to or from the ES 500. ACLs used in this way are known as *interface ACLs*.
- Applying an ACL to ports operating in Layer-4 bridging mode, which permits or denies bridged traffic. ACLs used in this way are known as *layer-4 Bridging ACLs*.

These uses of ACLs are described in the following sections.

11.2.1 Applying ACLs to Interfaces

An ACL can be applied to an interface to permit or deny either inbound or outbound traffic. Inbound traffic is traffic coming into the ES 500. Outbound traffic is traffic going out of the ES 500. For example, you cannot apply two or more IP ACLs to the same interface in the inbound direction. Each port can have a maximum of 224 ACLs definitions.

When a packet comes into the ES 500 at an interface where an inbound ACL is applied, the ES 500 compares the packet to the rules specified by that ACL. If it is permitted, the packet is allowed into the ES 500. If not, the packet is dropped. If that packet is to be forwarded to go out of another interface (that is, the packet is to be routed) then a second ACL check is possible. At the output interface, if an outbound ACL is applied, the packet is compared to the rules specified in this outbound ACL. Consequently, it is possible for a packet to go through two separate checks, once at the inbound interface and once more at the outbound interface.

In general, you should try to apply ACLs at the inbound interfaces instead of the outbound interfaces. If a packet is to be denied, you want to drop the packet as early as possible, at the inbound interface. Otherwise, ES 500 processes the packet, determine where the packet should go only to find out that the packet should be dropped at the outbound interface. In some cases, however, it may not be simple or possible for the administrator to know ahead of time that a packet should be dropped at the inbound interface. Nonetheless, for performance reasons, whenever possible, you should create and apply an ACL to the inbound interface.

To apply an ACL to an interface, enter the following command in Configure mode:

```
rs (config)#acl <name> apply interface <InterfaceName> input | output
```

The command creates rules in the ES 500 to filter packets according to IP flow's specifications which should be defined before by commands "`acl permit | deny`". The ES 500 filters IP routed packets on an ingress or/and egress ports defined by IP interface with name defined by "`InterfaceName`" parameter. The filtering direction is defined by keywords "input or/and output". If an IP interface is defined on a VLAN than set of ingress or/and egress ports defined by this VLAN.

11.2.2 Applying ACLs to ports

ACLs can also be created to permit or deny access to one or more ports all ACLs are for routed traffic.

Like ACLs that are applied to interfaces, ACLs that are applied to ports can be applied to either inbound or outbound traffic. To apply an ACL to a port, enter the following command in Configure mode:

```
rs (config)#acl <name> apply port <port list> input | output
```

The command applies rules in the ES 500 to filter packets according to IP flow's specifications which should be defined before by commands "`acl permit | deny`". The ES 500 filters I2 switched or IP routed packets on an ingress or/and egress ports defined by "port list" parameter. The filtering direction is defined by keywords "input or/and output". If a set of ingress or/and egress ports is defined to a VLAN, the same ACL rules apply to new port defined to the same VLAN.

The following is an example of the command's usage:

```
rs(config)#acl one permit ip 2.1.1.1
```

```
rs(config)#acl one deny ip any 2.1.1.2
rs(config)#acl one permit ip any 2.0.0.0/8
rs(config)#
rs(config)#acl one apply port et.2.3 input
```

11.2.3 Number of ACLs

The number of rules in a single ACL is limited. Each single rule is defined by one ACL entry.

There are 11 types of flows. Flow type is defined by set of packet fields, which is used for classification. For example: "acl one permit ip 1.2.3.4" and "acl two permit ip any 1.2.3.4" have two different types of flow because they use different set of packet fields for classification.

Example:

```
rs(config)#interface create ip inf1 address-netmask 1.2.3.4 port et.2.3
rs(config)#acl one permit ip 1.2.3.4
rs(config)#acl one permit ip 1.2.0.0/16
rs(config)#acl one permit ip 1.2.3.0/8
rs(config)#acl one permit ip any 1.2.3.0/8
rs(config)#acl one permit ip any 1.2.3.4
rs(config)#acl one permit ip any any http
rs(config)#acl one permit ip 1.2.3.4 any http
rs(config)#acl one permit ip 1.2.3.4 2.3.4.5 http
rs(config)#acl one permit ip 1.2.3.4 2.3.4.5 http 7777
rs(config)#acl one permit ip any 2.3.4.5 http 7777
rs(config)#acl one permit ip any any http 7777
rs(config)#acl one apply interface inf1 input
```

Rules created by these commands utilize all 11 flow types, therefore the command below, which required another flow type, results an error.

```
rs(config)#acl one permit ip any any http 7777 34
%ACL-E-NOMEM processing error - insufficient memory
```

There are 248 flows, which belong to one of 11 flow types that may be created for ports:

- 248 for et.2.1-et.2.8,
- 248 for et.2.9-et.2.16,
- 248 for et.2.17-et.2.24,
- 121 for gi.1.1
- 121 for gi.1.2
- et.2.1-et.2.8

The following illustrates an example of the syntax:

```
rs(config)#acl one permit ip 1.2.3.1
rs(config)#acl one permit ip 1.2.3.2
rs(config)#acl one permit ip 1.2.3.3
```

```
.....  
rs(config)#acl one permit ip 1.2.3.247  
rs(config)#acl one permit ip 1.2.3.248  
rs(config)#acl one apply interface inf1 input
```

Rules created by these commands utilize all 248 flows, which belong to one flow type. Therefore the command below, which require another flow result in error.

```
rs(config)#acl one permit ip 1.2.3.249  
%ACL-E-NOMEM processing error - insufficient memory
```

11.3 MONITORING ACLS

ES 500 provides a display of ACL configurations active in the system.

To display the ACL information, enter the following commands in the Enable mode:

Show all ACLs.	<code>acl show all</code>
Show a specific ACL.	<code>acl show aclname <name> all</code>

12 SECURITY CONFIGURATION

12.1 LAYER 2 SECURITY FILTERS

Filters on the ES-500 provide Layer-2 security by configuring ports to filter specific MAC addresses. A Layer-2 security filter is defined by specifying which ports to apply the filter. Filters have the effect of ensuring that specifically defined traffic is blocked or forwarded to and from the ports. The following filters are supported:

- Address filters - Block traffic based on a frame source MAC address, destination MAC address, or both. Address filters are always configured and applied on the input port.
- Port-to-address lock filters - Prohibits a user connected to a locked port or set of ports from using another port. A port or set of ports are disallowed access to other ports.
- Static entry filters - Allow traffic to go to a set of destination ports based on a frame source MAC address, destination MAC address, or both. Static entry filters are always configured and applied on the input port. Source static entry filters and destination static entry filters can be configured. Source static entry filters allow or disallow frames based on their source MAC address; destination static entry filters allow or disallow frames based on their destination MAC address.

12.1.1 Configuring Layer-2 Address Filters

To control access to a source or destination on a per-MAC address basis, configure an address filter.

Address filters are always configured and applied to the input port. You can set Address filters can be set on the following:

- A source MAC address - Any frame coming from a specific source MAC address is filtered out.
- A destination MAC address - Any frame destined to specific destination MAC address is filtered out.

To configure Layer-2 address filters, enter the following commands in Configure mode:

Configure a source MAC based address filter.	<code>filters add address-filter name <name> source-mac <MACaddr> source-mac-mask <MACadd> <VLAN-num> in-port-list <port-list></code>
Configure a destination MAC based address filter.	<code>filters add address-filter name <name> dest-mac <MACaddr> dest-mac-mask <MACadd> vlan <VLAN-num> in-port-list <port-list></code>

12.1.2 Configuring Layer-2 Port-to-Address Lock Filters

Port address lock filters to bind or “lock” specific source MAC addresses to a port or set of ports. Once a port is locked, only the specified source MAC address is allowed to connect to the locked port and the specified source MAC address is not allowed to connect to any other ports.

To configure Layer-2 port address lock filters, enter the following commands in Configure mode:

Configure a port address lock filter.	<code>filters add port-address-lock name <name> source-mac <MACaddr> vlan <VLAN-num> in-port-list <port-list></code>
---------------------------------------	--

12.1.3 Configuring Layer-2 Static Entry Filters

Static entry filters allow traffic to go to a set of destination ports based on a frame source MAC address, destination MAC address, or both source and destination MAC addresses in bridging mode. Static entries are always configured and applied at the input port. The following static entry filters can be set:

- Source static entry, specifying that any frame coming from source MAC address is allowed or disallowed.
- Destination static entry, specifying that any frame destined to a specific destination MAC address is allowed or disallowed.

To configure Layer-2 static entry filters, enter the following commands in Configure mode:

Configure a source static entry filter.	<code>filters add static-entry name <name> restriction allow disallow source-mac <MACaddr> source-mac-mask <MACaddr> vlan <VLAN-num> in-port-list <port-list> out-port-list <port-list></code>
---	--

Configure a destination static entry filter.	<code>filters add static-entry name <name> restriction allow disallow dest-mac <MACaddr> dest-mac-mask <MACaddr> vlan <VLAN-num> in-port-list <port-list> out-port-list <port-list></code>
--	--

12.1.4 Monitoring Layer-2 Security Filters

Layer-2 security filter configurations contained in the routing table can be displayed.

To display security filter information, enter the following commands in Enable mode.

Show address filters.	<code>filters show address-filter [all-source all-destination all-flow] [source-mac <MACaddr> dest-mac <MACaddr>] [ports <port-list>] [vlan <VLAN-num>]</code>
Show port address lock filters.	<code>filters show port-address-lock [ports <port-list>] [vlan <VLAN-num>] [source-mac <MACaddr>]</code>
Show static entry filters.	<code>filters show static-entry [all-source all-destination all-flow] ports <port-list> vlan <VLAN-num> [source-mac <MACaddr> dest-mac <MACaddr>]</code>

13 QOS CONFIGURATION

The ES 500 allows network managers to identify traffic and to set Quality of Service (QoS) policies without compromising the wire speed performance. The ES 500 can guarantee the bandwidth on an application-by-application basis, thus accommodating high-priority traffic even during the peak periods of usage. QoS policies can be broad enough to encompass all the applications in the network, or relate specifically to a single host-to-host application flow. The ES 500 provides four different features to satisfy the QoS requirements:

- **Traffic Prioritization** – Allows network administrators to differentiate between mission-critical network traffic and non-critical network traffic and segregate the traffic into different priority queues. Once a packet has been identified, it can be assigned to any one of the four priority queues in order to ensure delivery. The priority can be allocated based on any combination of the Layer-2, Layer-3, or Layer-4 traffic.
- **Type of Service (ToS)** – ToS rewrite provides network administrators with an access to the ToS octet in an IP packet. The ToS octet is designed to provide feedback to the upper layer application. The administrator can mark packets using the ToS rewrite feature so that the application (a routing protocol, for example) can handle the packet based on a predefined mechanism.
- **Traffic Rate Limiting** – Provides network administrators with tools to manage the bandwidth resources. The administrator can create an upper limit for a traffic profile, which is based on Layer-3 or Layer-4 information. Traffic that exceeds the upper limit of the profile can be either dropped or re-prioritized into another priority queue.

Within the ES 500, QoS policies are used to classify Layer-2, Layer-3, and Layer-4 traffic into the following priority queues (in order from highest priority to lowest):

- Control
- High
- Medium
- Low



Note

Control is for router control traffic. The remaining classes are for normal data flows.

Separate buffer space is allocated to each of these four priority queues. By default, the buffered traffic in higher priority queues is forwarded ahead of the pending traffic in lower priority queues. This is the *strict priority* queuing policy. During heavy load times, low-priority traffic can be dropped to preserve the throughput of the higher-priority traffic. This ensures that the critical traffic reaches its destination, even if the traffic exit ports are experiencing greater-than-maximum utilization. To prevent the low-priority traffic from waiting indefinitely as higher-priority traffic is sent, you can apply the Weighted Fair Queuing (WFQ) policy to set a minimum bandwidth for each class.

13.1 LAYER-2, LAYER-3 AND LAYER-4 FLOW SPECIFICATION

In the ES 500, traffic classification is accomplished by means of mapping Layer-2, -3, or -4 traffic to one of the four priorities. Each traffic classification is treated as an individual traffic in the ES 500.

For Layer-2 traffic, you can define a traffic based on the MAC packet header fields, including source MAC address, destination MAC address, and VLAN IDs. A list of incoming ports can also be specified.

For Layer-3 (IP) traffic, you can define traffic, blueprints or templates of the IP packet headers:

Ip Fields – The source IP address, destination IP address, UDP/TCP source port, UDP/TCP destination port, TOS (Type of Service), transport protocol (TCP or UDP), and a list of incoming interfaces.

For Layer-4 traffic, you can define traffic based on source/destination TCP/UDP port number in addition to the Layer-3 source/destination IP address.

If a value is not entered for a field, a wildcard value (all values acceptable) is assumed for the field.

13.2 PRECEDENCE FOR LAYER-3 FLOWS

The precedence from 1 to 7 is associated with each field in a flow. The ES 500 uses the precedence value associated with the fields to break ties if packets match more than one flow. The highest precedence is 1 and the lowest is 7. Here is the default precedence of the fields:

- IP
 - Destination port – 1
 - Destination IP address – 2
 - Source port – 3
 - Source IP address – 4
 - ToS – 5
 - Interface – 6
 - Protocol – 7

Use the `qos precedence ip` command to change the default precedence.

13.3 QUEING POLICIES

The ES 500 QoS policies apply to all ports (per device as a whole). There are two types of queuing policies you can use on the ES 500:

- **Strict priority** – Assures the higher priorities of throughput but at the expense of the lower priorities. For example, during heavy loads, the low-priority traffic can be dropped to preserve throughput of the control-priority traffic. This is the default queuing policy.
- **Weighted fair queuing** – Distributes priority throughput among the four priorities based on percentages. This queuing policy is set on a per-device basis.

13.4 TRAFFIC PRIORITIZATION FOR LAYER-2 FLOWS

13.4.1 Configuring Layer-2 QoS

When applying QoS to a layer-2 flow, priority can be assigned as follows:

- The frame gets assigned a priority within the switch. Select low, medium, high or control.
- The frame gets assigned a priority within the switch, and if the exit ports are VLAN trunk ports, the frame is marked with an 802.1Q priority. Select a number from 0 to 7.

To set a QoS priority on a layer-2 flow, enter the following command in the Configure mode:

Set a Layer-2 QoS policy.	<code>qos set l2 name <name> source-mac <MACaddr> any source-mac-mask <MACaddr> dest-mac <MACaddr> any dest-mac-mask <MACaddr> vlan <vlanID> any in-port-list <port-list> priority control high medium low <trunk-priority></code>
---------------------------	--

13.4.2 802.1p Class of Service Priority Mapping

The following table shows the default mappings of 802.1p Class of Service (CoS) values to internal priorities for frames:

Table 13-1 802.1p default priority mappings

802.1p CoS values	Internal priority queue
0,1	Low
2,3	Medium
4,5	High
6,7	Control

You can create one or more priority maps that are different from the default priority map and then apply these maps to some or all ports on the ES 500. The new priority mapping replaces the default mappings for those ports.

Creating and Applying a New Priority Map

To specify a priority map on a per-port basis, enter the following commands in the Configure mode:

Create a new priority mapping.	<code>qos create priority-map <name> <CoS number> control high medium low</code>
Apply new priority mapping to ports.	<code>qos apply priority-map <name> ports <port-list></code>

For example, the following command creates the priority map *all-low* which maps all 802.1p priorities to the low internal priority queue:

```
rs (config)#qos create priority-map all-low 0 low 1 low 2 low 3 low 4 low 5 low 6 low 7 low
```

The priority map is applied to ports as shown in the following example:

```
rs (config)#qos apply priority-map all-low ports et.2.(1-4), gi.1.(1-2)
```

**Note**

Keep in mind that in ES 500 the priority map is automatically applied to all ports, even if you specify only some of them.

You do not need to specify mappings for all 802.1p values. If you do not specify a particular mapping, the default mapping for that 802.1p priority is used. The following example creates the priority map *no-ctrl* with the same mappings as the default priority map, except that the 802.1p priority of 7 is mapped to the internal priority ‘high’ instead of ‘control’.

```
rs (config)#qos create priority-map no-ctrl 7 high
```

Removing or Disabling Per-Port Priority Map

Negating a `qos create priority-map` command removes the priority map. Before you can remove a priority map, you must negate all commands that use the priority map. Negating a `qos apply priority-map` command causes the configured ports to use the default priority mapping.

The ability to specify per-port priority maps is enabled on the ES 500 by default. You can disable the use of per-port priority maps on the ES 500. All ports on ES 500 is configured to use the default priority map only. If the commands to create and apply priority maps exist in the active configuration, they remain in the configuration but are ineffective.

To disable the use of priority maps, enter the following command in the Configure mode:

```
Disable use of per-port priority maps on the ES 500.    qos priority-map off
```

If the above command is negated, ports on the ES 500 can use per-port priority maps. If the commands to create and apply priority maps exist in the active configuration, they are reapplied.

Displaying Priority Map Information

To display priority maps and the ports on which they are applied, enter the following command in the Enable mode:

```
Display priority mapping.    qos show priority-map <name> | all
```

13.5 TRAFFIC PRIORITIZATION FOR LAYER-3 AND LAYER-4 FLOWS

QoS policies applied at Layer-3 and -4 allow you to assign priorities based on specific fields in the IP headers. You can set QoS policies for IP flows based on a source IP address, destination IP address, source TCP/UDP port, destination TCP/UDP port, type of service (TOS), and transport protocol (TCP or UDP). A QoS policy set on an IP flow allows you to classify the priority of the traffic based on:

- Layer-3 source-destination flows
- Layer-4 source-destination flows
- Layer-4 application flows

13.5.1 Configuring IP QoS Policies

To configure an IP QoS policy, perform the following tasks:

1. Identify the Layer-3 or -4 flow and set the IP QoS policy.
2. Specify the precedence for the fields within an IP flow.

Setting an IP QoS Policy

To set a QoS policy on an IP traffic flow, use the following command in the Configure mode:

Set an IP QoS policy.	<pre>qos set ip <name> <priority> <srcaddr/mask> any <dstaddr/mask> any <srcport> any <dstport> any <tos> any <port list> <interface-list> any <protocol> any <tos-mask> any <tos-precedence-rewrite> any <tos- rewrite> any</pre>
-----------------------	--

For example, the following command assigns control priority to any traffic coming from the 10.10.11.0 network:

```
rs (configure)#qos set ip xyz control 10.10.11.0/24
```

Specifying Field Precedence for an IP QoS Policy

To specify the precedence for an IP QoS policy, use the following command in the Configure mode:

Specify precedence for an IP QoS policy.	<pre>qos precedence ip [sip <num>] [dip <num>] [srcport <num>] [destport <num>] [tos <num>] [protocol <num>] [intf <num>]</pre>
--	---

13.6 CONFIGURING ES 500 QUEUING POLICY

The ES 500 default queuing policy, strict priority, can be set on a system-wide basis. To change for all the devices:

Set queuing policy to weighted-fair	<pre>qos set queuing-policy weighted-fair port all-ports</pre>
-------------------------------------	--

If you want to revert the ES 500 queuing policy from weighted-fair to strict priority (default), enter the following command in the Configure mode:

Revert the ES 500 queuing policy to strict priority	<code>negate <line within active-configuration containing qos set queuing-policy weighted-fair></code>
---	--

Allocating Weight in Percentages for a Weighted-Fair Queuing Policy

If you enable the weighted-fair queuing policy on the ES 500, you can allocate bandwidth for the queues on the ES 500. To allocate bandwidth for each queue, enter the following command in the Configure mode:

Allocate bandwidth for a weighted-fair queuing policy.	<code>qos set weighted-fair control <percentage> high <percentage> medium <percentage> low <percentage> port <port list> all-ports</code>
--	---

13.7 MONITORING QoS

The ES 500 provides the display of QoS statistics and configurations.

To display the QoS information, enter the following commands in the Enable mode:

Show all IP QoS flows.	<code>qos show ip</code>
Show QoS information for L2 flows	<code>qos show l2</code>
Show IP precedence values.	<code>qos show precedence ip</code>
Show WFQ bandwidth allocated for each port.	<code>qos show wfq [port <port-list> all-ports]</code>
Show priority mappings.	<code>qos show priority-map all</code>

13.8 LIMITING TRAFFIC RATE

Rate limiting provides the ability to control the usage of the fundamental network resource - the bandwidth. It allows you to limit the rate of traffic that flows through the specified interfaces, thus reserving bandwidth for critical applications. The ES 500 supports the following types of rate limiting:

- **Port-level Rate Limiting** – Configure policies that limit traffic coming into a particular port.

Rate limiting policies work only in one direction. That is, only the traffic coming into the interface to which a policy is applied is subject to rate limiting (except for output port rate limiting policies, which are applied to egress ports). If both incoming and outgoing traffic to a network or subnet needs to be rate limited, then you should create separate policies to be applied to each interface.

13.8.1 Port Rate Limiting


Use a port rate limiting policy if incoming or outgoing traffic on a particular port needs to be rate limited. Unlike other types of rate limiting policies, you do not specify an ACL when defining this type of policy. Port rate limiting policies do not need to be applied to an interface and take effect when they are created.


To configure port rate limiting policies for input ports, you must first enable the aggregate rate limiting mode on the line card. You do not need to enable the aggregate rate limiting mode to configure a policy to limit outgoing traffic

on a port. You can configure port-level rate limiting policies on output ports in either per-flow or aggregate rate limiting mode.

To define a port rate limit policy, use the following commands in the Configure mode:

Define a port rate limit policy to limit incoming traffic on a port.	<code>rate-limit <name> port-level input port <port list> rate <num> [drop-packets no-action tos-rewrite <num></code>
Define a port rate limit policy to limit outgoing traffic on a port.	<code>rate-limit <name> port-level output port <port list> rate <rate-limit> drop-packets</code>

 **Note** The first example describes both actions permissible when Rate-limiting traffic according to the input port. The TOS rewrite action changes the DSCP value if the incoming traffic is Layer-3.

 **Note** For output port policies, the only action that you can specify if traffic exceeds the specified rate is to drop packets. The goal of the output port rate limiting is to drop outgoing traffic that exceeds the specified rate. This is achieved by the metering of frames in an egress of an ASIC. The rate limit drop action is executed for a frame, which rate is greater than the defined threshold.

In BCM5615 (5605, 5625) ASIC though, only ingress metering is available. Therefore, the output rate limiting is implemented by means of metering of aggregate frames destined for specific output ports, defined by the *port-list* parameter of the `rate-limit` command.

If you configure output port policies, all types of outgoing IP or bridged traffic are rate limited:

Specify that control traffic be not subject to output port rate limiting.	<code>rate-limit <name> port-level slot <slot-number> to limit Layer 2 high priority traffic</code>
---	---

13.9 MONITORING RATE LIMITING POLICIES

To display the information about rate limiting policies on ES 500, use the following command in the Enable mode:

```
rate-limit show [all] | [policy-type | port-level-policies | all] | [policy-name <name>] | [port-level port <port list> | all-port] | [port-level policy-name <name>]
```

This command has the following parameters:

Displays information on all rate-limiting policies.	<code>show all</code>
The type of rate limiting policy.	<code>policy-type</code>
Displays all port level policies.	<code>policy-type portlevel-policies</code>
Displays all rate limiting types.	<code>policy-type all</code>
Displays rate limiting policies by name.	<code>policy-name <name></code>
Displays all rate limiting policies by name.	<code>policy-name all</code>

Displays rate limiting policies on a port.	<code>port-level port <port list></code>
Displays rate limiting policies on all ports.	<code>port-level port all-port</code>
The name of the port-level rate limiting policy.	<code>port-level policy-name <name></code>

14 PERFORMANCE MONITORING

The ES 500 is a full wire-speed layer-2, -3 and -4 switching router. As packets enter the ES 500, layer-2, -3, and -4 flow tables are populated. The flow tables contain information on performance statistics and traffic forwarding. Thus the ES 500 provides the capability to monitor performance at Layer 2, 3, and 4.

Layer-2 performance information is accessible to SNMP through MIB-II and can be displayed by using the **l2-tables** command in the CLI. Layer-3 and 4 performance statistics are accessible to SNMP through RMON/RMON2 and can be displayed by using the **statistics show** command in the CLI. In addition to the monitoring commands listed, you can find more monitoring commands listed in each chapter of the *Riverstone ES 500 Switch Router Command Line Interface Reference Manual*.

To access statistics on the ES 500, enter the following commands in the Enable mode:

Show all TCP/UDP connections and services.	<code>ip show connections</code>
Display the IP helper addresses configuration on the system.	<code>ip show helper-address</code>
Display the IP interface configuration.	<code>ip show interfaces</code>
Show all MAC addresses currently in the L2 tables.	<code>l2-tables show all-macs [verbose] [vlan <VLAN-num>] [multicast]</code>
Show information about the master MAC table.	<code>l2-tables show mac-table-stats</code>
Show whether IGMP is on or off on a VLAN.	<code>l2-tables show vlan-igmp-status</code>
Show info about Multicasts registered by IGMP.	<code>l2-tables show igmp-mcast-registrations [vlan <VLAN-num>]</code>
Show ICMP statistics.	<code>statistics show icmp</code>
Show IP interface's statistics.	<code>statistics show ip</code>
Show port error statistics.	<code>statistics show port-errors</code>
Show TCP statistics.	<code>statistics show tcp</code>
Show UDP statistics.	<code>statistics show udp</code>
Show port packet statistics.	<code>statistics show port-packets</code>
Show normal (non-error) port statistics.	<code>statistics show port-stats</code>

Show broadcast monitoring information for ports.	<code>port show bmon</code>
--	-----------------------------

14.1 CONFIGURING PORT MIRRORING

ES 500 monitors port activity with port mirroring. Port mirroring monitors the performance and activities of ports on ES 500 or for traffic defined by an ACL through just a single, separate port. While in Configure mode, you can configure your ES 500 for port mirroring with a simple command line like the following:

Configure port mirroring.	<code>port mirroring monport <port number> targport <port number>/target-profile <acl name></code>
---------------------------	--

14.2 MONITORING BROADCAST TRAFFIC

The ES 500 allows you to monitor broadcast traffic for one or more ports, and for the control module. You can specify that a port be shut down if its broadcast traffic reaches a certain rate limit for a particular period of time. Additionally, you can configure the ES 500 to shut down for a specified period, if the packets sent to the control module reach a certain limit during a specified time interval. Packets to be monitored can be limited to broadcast packets only or all packets.

To specify the monitoring of broadcast traffic and the shut down threshold for one or more ports, enter the following command in the Configure mode:

Configure monitoring of broadcast traffic.	<code>port bmon <port list> [rate <rate>] [broadcast enable disable] [multicast enable disable] [unknown_unicast enable disable]</code>
--	---

15 RMON CONFIGURATION

You can employ Remote Network Monitoring (RMON) in your network to help monitor traffic at remote points on the network. With RMON, data collection and processing is done with a remote *probe*, namely the ES 500. The ES 500 also includes RMON *agent* software that communicates with a network management station via SNMP. Because information is only transmitted from the ES 500 to the management station when required, SNMP traffic on the network and the management station's processing load are reduced.

The ES 500 provides support for RMON MIB, as specified in RFCs 1757. While non-RMON SNMP products allow the monitoring and control of specific network *devices*, RMON returns statistics on the network *segments* at the MAC layer.

15.1 CONFIGURING AND ENABLING RMON

By default, RMON is disabled on the ES 500. It is enabled automatically the moment you insert the first entry into the RMON table.

To insert the entry, enter the following command in the Config mode:

```
rs (config)#rmon etherstats index <index-number> port <port>
```

15.1.1 RMON Groups

The RMON MIB groups are defined in RFCs 1757. The supported RMON groups are shown in the table below.

Table 15-1 RMON groups

Group	Function
EtherStats	Records Ethernet statistics (for example, packets dropped, packets sent, etc.) for specified ports.
Event	Controls event generation and the resulting action (writing a log entry or sending an SNMP trap to the network management station).
Alarm	Generates an event when specified alarm conditions are met.
History	Records statistical samples for specified ports.

15.1.2 Control Tables

Many RMON groups contain both control and data tables. Control tables specify what statistics are to be collected. For example, you can specify the port for which statistics are to be collected and the owner (name, phone, or IP address) for that port.

15.2 USING RMON

RMON on the ES 500 allows you to analyze network traffic patterns, set up alarms to detect potential problems before they turn into real congestive situations, identify heavy network users to assess their possible candidacy for moves to dedicated or higher speed ports, and analyze traffic patterns to facilitate more long-term network planning. RMON provides layer-2 information. RMON groups collect the traffic, which is flowing through the ES 500's layer-2 ASIC.

15.3 CONFIGURING RMON GROUPS

The following table shows the `rmon` command that you use to configure each RMON group:

To configure the Alarm group.	<code>rmon alarm index <index-number> variable <string> [interval <seconds>] [falling-event-index <num>] [falling-threshold <num>] [owner <string>] [rising-event-index <num>] [rising-threshold <num>] [startup rising falling both] [status enable disable] [type absolute-value delta-value]</code>
To configure the Event group.	<code>rmon event index <index-number> type none log trap both [community <string>] [description <string>] [owner <string>] [status enable disable]</code>
To configure the History group.	<code>rmon history index <index-number> port <port> [interval <seconds>] [owner <string>] [samples <num>] [status enable disable]</code>
To configure the EtherStats group.	<code>rmon etherstats index <index-number> port <port> [owner <string>] [status enable disable]</code>

15.3.1 Configuration Examples

This section shows examples of configuration commands that specify an event that generates an SNMP trap and the alarm condition that triggers the event.

The RMON Alarm group allows the ES 500 to poll itself at user-defined intervals. Alarms that constitute an event are logged into the Event table that can then be polled by the management station. The management station is able to poll more network devices this way, as it only needs to poll the RMON Event table and not the device itself. The management station can also be sent the trap information.

The following examples configure the ES 500 to create an event when a specified number of octets arrive. The managed object etherStatsOctets has an object identifier (OID) of 1.3.6.1.2.1.16.1.1.1.4 and the ES 500 polls this OID every 5 minutes (300 seconds).

The command line below is an example of an RMON Event group configuration with the following attributes:

- Index number 15 to identify this entry in the Event control table.
- The event is both logged in the Event table and an SNMP trap is generated with the community string "public".
- Event owner is "monitor".

```
rs(config)#rmon event index 15 type both community public description "Octet
count" owner "monitor"
```

The command line below is an example of an RMON Alarm group configuration with the following attributes:

- Index number 20 to identify this entry in the Alarm control table.
- The OID 1.3.6.1.2.1.16.1.1.1.4 identifies the attribute to be monitored.
- Samples taken at 300 second (5 minute) intervals.
- A “Startup” alarm generation condition instructing the ES 500 to generate an alarm if the sample is greater than or equal to the rising threshold, or less than or equal to the falling threshold.
- Compare value at time of sampling (absolute value) to the specified thresholds.
- Rising and falling threshold values are 50 and 20 respectively.
- Rising and falling event index values are 25 and 15 respectively, which trigger the previously-configured Event.

```
rs(config)#rmon alarm index 20 variable 1.3.6.1.2.1.16.1.1.1.4 interval 300
startup both type absolute-value rising-threshold 50 falling-threshold 20
rising-event-index 25 falling-event-index 15 owner "monitor"
```

15.4 DISPLAYING RMON INFORMATION

The CLI `rmon show` commands allow you to display the same RMON statistics that can be viewed from a management station. To display RMON statistics for the ES 500, use the following CLI command lines in Enable mode:

To show Ethernet statistics.	<code>rmon show etherstats <port-list> all-ports</code>
To show all events and logs.	<code>rmon show events</code>
To show all alarms.	<code>rmon show alarms</code>
To show history and logs.	<code>rmon show history <port-list> all-ports</code>

For example, the following screen shows the result of entering the `rmon show alarm` command in the Enable mode:

```
rs#rmon show alarm

RMON I Alarm Table
Index: 1, Variable: 1.3.6.1.2.1.16.1.1.1.4, Owner: monitor
-----
Rising-Event-Index   : 4
Falling-Event-Index  : 4
Rising-Threshold     : 100
Falling-Threshold    : 20
Interval              : 100
Startup Type         : both
Sample Type          : absolute-value
```

Here is a result of `rmon show etherstats` command in Enable mode:

```
rs#rmon show etherstats all-ports
```

```

RMON I Ethernet Statistics Table
Index 1 Port: et.2.1  Owner: monitor
----
RMON EtherStats          Total
-----
Octets                    0
Multicast Frames         0
Broadcast Frames         0
Collisions                0
64 Byte Frames           0
65-127 Byte Frames       0
128-255 Byte Frames      0
256-511 Byte Frames      0
512-1023 Byte Frames     0
1024-1518 Byte Frames    0
    
```

16 SERVICE CONFIGURATION

Rate limiting enables you to control the rate of traffic that flows through specified interfaces. You can configure rate limiting by using the `service rate-limit` facility.

The `service rate-limit` facility provides the following benefits:

- The `service rate-limit` commands support burst-safe rate limiting.
- You can use ACLs or a feature called the Multi-Field Classifier (MF Classifier) to define the traffic profile on which rate limiting services are applied.

The procedure for configuring rate-limiting services involves first creating the rate limit service, and then applying it to the traffic on a port or interface. This chapter describes how to use the `service rate-limit` commands to configure the rate limiting.

16.1 CONFIGURING RATE LIMIT SERVICES

The ES 500 supports the following rate limit services:

- **Aggregate Rate Limiting** – Limits an aggregation of flows to a specified rate.
- **Port-level Rate Limiting** – Limits the traffic that is coming to a particular port.
- **Per-flow Rate Limiting** – Limits individual flows to a specified rate. This is the default rate limiting mode on the ES 500.

16.2 APPLYING RATE LIMIT SERVICES

After configuring the rate limit service, you can apply it to one or more interfaces or ports, and specify the traffic profile to which it is applied. When you use the `service` commands to configure rate limit services, you can specify the traffic profile through ACLs or by using an MF Classifier.

The `service apply rate limit acl` command applies a previously-defined rate limit service to an interface or to a port and specifies the traffic profile, via the ACL, to which the rate limit service applies.

The following example shows how to apply an ACL called `acl1` to the port `et.2.8`:

```
rs(config)#service customergroup1 apply rate-limit acl1 port et.2.8
```

Use the MF Classifier to define traffic characteristics based on the fields of an IP packet. In addition, when you use the MF Classifier command, `service <name> apply rate-limit mf-classifier`, you can specify the traffic profile *and* apply the rate limit service all at once. You do not need to do this in separate commands, as shown in the following example:

```
rs(config)#service customergroup1 create rate-limit aggregate rate 10000000
rs(config)#service customergroup1 apply rate-limit mf-classifier interface
int1 source-addr-mask 10.1.1.1
```

16.2.1 Applying Aggregate and Port-Level Rate Limiting

The following example shows how you can configure an aggregate rate limiting service and a port-level rate limiting service on the ES 500. The aggregate rate limiting service restricts the traffic from 4 customers (S1, S2, S3, and S4) to 10 Mbps. The port-level rate limiting service restricts traffic from the Internet to 64 Mbps.

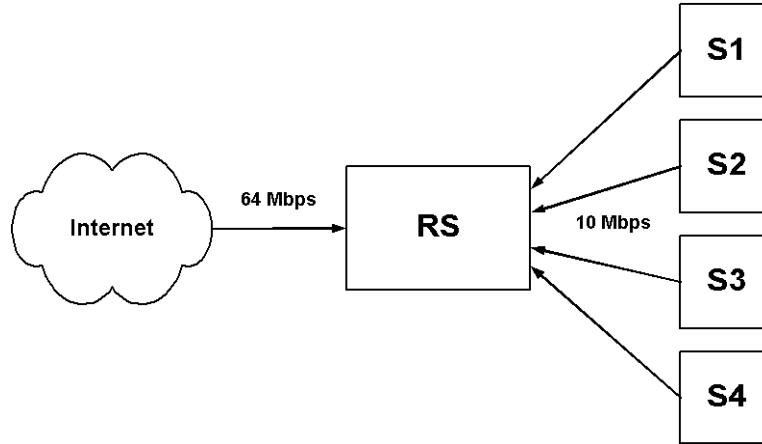


Figure 16-1 Applying aggregate and port-level rate limiting

The configuration shown in Figure 16-1 is created with the following commands. In this example, ACLs are used to define the traffic profile:

```

rs(config)#show active

Configure VLANs and interfaces

vlan create s1 ip
vlan create s2 ip
vlan create s3 ip
vlan create s4 ip

vlan add ports et.2.8 to s1
vlan add ports et.2.7 to s2
vlan add ports et.2.6 to s3
vlan add ports et.2.5 to s4

interface create ip from_s1 vlan s1 address-netmask 10.1.1.1
interface create ip from_s2 vlan s2 address-netmask 15.1.1.1
interface create ip from_s3 vlan s3 address-netmask 20.1.1.1
interface create ip from_s4 vlan s4 address-netmask 25.1.1.1

Configure the ACLs

acl s1 permit ip 10.1.1.1
acl s2 permit ip 15.1.1.1
acl s3 permit ip 20.1.1.1
acl s4 permit ip 25.1.1.1
  
```

Configure aggregate rate limiting

```
service flow1 create rate-limit aggregate rate 10000000 drop-packets
```

Apply the rate limit service to the traffic on the interfaces

```
service flow1 apply rate-limit acl s1 interface from_s1
```

```
service flow1 apply rate-limit acl s2 interface from_s2
```

```
service flow1 apply rate-limit acl s3 interface from_s3
```

```
service flow1 apply rate-limit acl s4 interface from_s4
```

Configure the port-level rate limit service

```
rs(config)#service flow2 create rate-limit input-portlevel rate  
64000000 no-action port et.2.4
```

The following example shows how you would use the MF-Classifier to define a traffic profile and apply a rate limit to it:

```
rs(config)#show active
```

Configure the VLANs and interfaces

```
vlan create s1 ip
```

```
vlan create s2 ip
```

```
vlan create s3 ip
```

```
vlan create s4 ip
```

```
vlan add ports et.2.1 to s1
```

```
vlan add ports et.2.2 to s2
```

```
vlan add ports et.2.3 to s3
```

```
vlan add ports et.2.4 to s4
```

```
interface create ip from_s1 vlan s1 address-netmask 10.1.1.1
```

```
interface create ip from_s2 vlan s2 address-netmask 15.1.1.1
```

```
interface create ip from_s3 vlan s3 address-netmask 20.1.1.1
```

```
interface create ip from_s4 vlan s4 address-netmask 25.1.1.1
```

Enable aggregate rate limiting

```
system enable aggregate-rate-limiting slot 2
```

Configure aggregate rate limiting

```
service flow1 create rate-limit aggregate rate 10000000 drop-packets
```

Apply the rate limit service to the specified traffic on the interfaces

```
service flow1 apply rate-limit mf-classifier interface from_s1
```

```
service flow1 apply rate-limit mf-classifier interface from_s2
```

```
service flow1 apply rate-limit mf-classifier interface from_s3
```

```
service flow1 apply rate-limit mf-classifier interface from_s4
```

Configure the port-level rate limit service

```
service flow2 create rate-limit input-portlevel rate 64000000 no-action  
port et.2.4
```

You can display information about the rate limit services that you configured:

Displays the specified aggregate rate limit service(s).	<code>service show rate-limit aggregate</code>
Displays all rate limit services.	<code>service show rate-limit all</code>
Displays the specified per-flow rate limit service(s).	<code>service show rate-limit per-flow</code>

Example below illustrates the effect of `service show rate-limit all` command:

```
rs#service show rate-limit all detailed
-----
Service Name : flow1          Service Type: Aggregate Rate
Limit Rate : 10000000 bps     Exceed Action : Drop Packets
s
Timeselect : 4              Credits : 204
-----
Rate Limit Service name: flow2  Type : Input Port Level
Configured on ports: et.2.1
Direction Rate   Credits  Time Interval  Exceed Action
-----
Input    64000000  1310   41.94 ms     Drop Packets
-----
```

16.2.2 Per-Flow Rate Limiting

Use a per-flow rate limiting policy if an individual traffic flow that needs to be limited to a particular rate. A single per-flow rate limiting policy can have multiple ACLs to define different traffic profiles and traffic rate limitations. When there are multiple traffic profiles, a sequence number is used to identify the order in which the profiles are applied.



Note Per-flow rate limiting is enabled on the ES 500 by default.



Note Non-IP ACLs cannot be use for per-flow rate limit policies.

The command to define a per-flow rate limit policy command in Configure mode:

```
service <name> create rate-limit per-flow rate <rate> [exceed-action  
<action>]
```

The command to displays the specified per-flow rate limit service is as follows:

```
service show rate-limit per-flow <name> | all [show-applied]
```


17 TIME CONFIGURATION

This chapter discusses how to set time and date on the ES 500.

17.1 SETTING TIME AND DATE

To set the date and time on the ES 500, use the `system set date` command in Enable mode. You can set any or all components of the date and time, including the year, month, day, hour, minute, or second. After entering the command, you see a confirmation of the time change.

For example, the following command sets the date to February 6, 2002 and the time to 10:10:40 in the morning:

```
rs#system set date year 2002 month 2 day 6 hour 10 min 10 second 40
Time changed to: 2002-02-06 10:10:40
```

You can also use the `system set date` command to set a single component of the date and time. For example, the following command sets the minutes to 30:

```
rs#system set date min 30
Time changed to: 2002-02-06 10:30:09
```


18 SNMP CONFIGURATION

The Simple Network Management Protocol (SNMP) is an application layer protocol used to monitor and manage TCP/IP-based networks. It provides for the storage and exchange of management information. The ES 500 supports two SNMP versions:

- SNMP Version 1 (SNMPv1) (RFC 1157)
- SNMP Version 2c (SNMPv2c) (RFC 1901, RFC 1905, and RFC 1906)

SNMPv1 and SNMPv2c can coexist in the same managed network (RFC 2576). You can run any or all of the SNMP versions on the ES 500, depending on the one used by the SNMP management stations. (For additional information on the different SNMP versions, refer to the RFCs for each version.)

18.1 CONFIGURING ACCESS TO MIB OBJECTS

Riverstone supports most of the standard networking SNMP MIB modules, as well as proprietary MIB modules. Each MIB module is a collection of managed objects, which can be accessed by the SNMP management stations.

SNMP management stations send SNMP SET and GET requests for the management objects stored in the MIB modules. ES 500 runs an SNMP agent that listens for these SNMP messages on UDP port 161. In SNMPv1 and v2c, the SNMP managers provide a community string (or password) when they send their requests. If the ES 500 recognizes the community string, it processes the request. If the string is not recognized a message stating an authorization violation has occurred is displayed and if the system is configured to send traps, a trap is sent.

18.1.1 Configuring SNMP access

Following are the tasks for configuring SNMP access if you are running SNMPv1 and v2c:

- Configure a community string. This is required.
- Configure the agent's identity.

Each of these tasks is discussed in the sections that follow.

Configuring Community Strings

To run SNMPv1 and v2c on the ES 500, you must define at least one community string. The ES 500 has no default community strings. When you define an SNMP community string, you also need to specify its access level, which is either read-only (allows only SNMP GETs), or read-write (allows SNMP SETs and GETs). In the following example, separate community strings are defined for read-only access and for read-write access:

```
rs(config)#snmp set community public privilege read
rs(config)#snmp set community private privilege read-write
```

An SNMP manager that sends a GET request for a MIB object on ES 500 provides the community string *public* or *private*; and an SNMP manager that sends a SET request should provide the community string *private*.

Configuring the SNMP Agent's Identity

You can use the CLI to set certain MIB objects, such as those that describe the agent's identity, as shown in the following example:

```
rs(config)#system set name RS8-1
rs(config)#system set contact "IT dept"
rs(config)#system set location "building 1 closet"
```

The example sets the MIB objects sysName to *RS8-1*, sysContact to *IT dept*, and sysLocation to *building 1 closet*.

18.2 CONFIGURING SNMP NOTIFICATIONS

ES 500 sends notifications to pre-defined targets. The targets are the SNMP management stations that receive the notifications. Notifications inform the SNMP managers about conditions on the network, such as an error condition or an authentication failure.

SNMPv1 defined only one type of notification; these were called traps. SNMP agents sent traps to alert SNMP managers about conditions on the network. Traps did not require acknowledgements from the receivers. Therefore, the SNMP agent never knew whether a trap was received.

To configure SNMP notifications you need to specify the targets and the community.

18.2.1 Specifying the Targets

To send SNMP notifications, you need to specify the following:

- the targets that receive the notifications.
- a community string.

Targets are defined by their IP addresses. Each target that is defined receives a copy of the notifications generated and sent by the ES 500 agent.

In addition, you need to specify a community string for the notifications. For security reasons, the community strings in notifications should be different from the read/write community strings. So when the ES 500 sends notifications, unauthorized users that capture the notifications are not be able to use the community string to change information in the MIB modules.

In the following example, Inform notifications are sent to the target with address 10.10.10.1.

```
rs(config)#snmp set community community1 privilege read
rs(config)#snmp set target 10.10.10.1 community community1
```



Note

Command "**snmp set target <IP_addr> community public status enable**" opens a community as read-only for IP. If this command is defined it is possible to execute GET SNMP MIB on the device without "**snmp set community public privilege read**" command.

**Note**

If the IP address of the target is more than one hop away from the ES 500, configure the ES 500 with a static route to the target. If the ES 500 is rebooted, the static route allows a cold-start notification to be sent to the target. Without a static route, the cold-start notification is lost while the routing protocols are converging.

18.3 MONITORING SNMP CONFIGURATION

Using the `snmp show` command you can view the following information:

- Community strings set on the device
- IP address of SNMP trap target server

The following table explains the purpose of different command parameters:

Show all SNMP information (equivalent to specifying all the other keywords)	<code>snmp show all</code>
Display the device community string	<code>snmp show community</code>
Display the IP address of the trap target server	<code>snmp show trap</code>

APPENDIX A TROUBLESHOOTING

When working with the ES 500, there are two operational areas that can be corrupted due to hardware failures or configuration errors. This troubleshooting appendix describes how to manage these and other problems. The following sections are described in this appendix:

- Password troubleshooting
- Configuration Troubleshooting
- System image troubleshooting

A.1 PASSWORD TROUBLESHOOTING

When a password is active in a specified mode, and the password is forgotten, the password-protected mode cannot be entered. There are two methods of retrieving forgotten password:

- Erasing the Startup File from the FLASH.
- Overriding Forgotten Passwords.

A.1.1 Erasing the Startup File from the FLASH

To regain access ES 500 if a password has been forgotten perform the following:

1. Erase the Startup File from the FLASH, see *Erasing the Configuration* in the *Getting Started Guide*.



Note

A terminal or PC running terminal emulation software connected directly to the ES 500 via the DB-9 console port is needed to complete this process.

2. After the Startup file is erased from the FLASH, new passwords for the User, Enable, and Configure modes can be assigned.

A.1.2 Overriding Forgotten Passwords

Forgotten passwords can be overridden without deleting the configuration. To override a forgotten password:

3. Reboot the device. The following message is displayed:

```
----- Performing the Power-On Self Test (POST) -----
UART Channel Loopback Test.....PASS
Testing the System SDRAM.....PASS
EPROM Checksum Test.....PASS
Flash Image Validation Test.....PASS
Testing CPU PCI Bus Device Configuration.....PASS

BOOT Software Version Prom-2.00.00 Built 22-Mar-2002 12:57:00
Processor: MPC8245 Rev 0.12, 266 MHz (Bus: 133MHz), 64 MByte SDRAM.
I-Cache 16 KB, linesize 32.D-Cache 16 KB, linesize 32.
Cache Enabled.
Autoboot in 2 seconds - press RETURN or Esc. to abort and enter prom.
```

4. Press Enter. The following boot screen displays.

```
*****
River Stone Networks TM
*****

Startup menu
[1] Download sw
[2] Erase from Flash
[3] Erase Flash
[4] Erase Nvram

Enter your choice:.
```

5. Within 2 seconds, select option “2) Erase file from Flash”. The following message is displayed.

```
Warning! About to erase the file from flash
Are you sure (Y/N)?
```

6. Confirm by pressing “Y”. The following message is displayed.

```
? Flash file name (8 characters, Enter for none.)
```

7. Enter **config** and press <Enter>. The **config** file is erased and the device reboots without the startup file commands. The defined passwords can now be bypassed.

8. In Enabled mode enter **copy startup to scratchpad**. The startup file is copied to the scratchpad.

9. In Config mode enter **save active**. The following message is displayed:

```
Do you want to make the changes active? [y]
```

10. Enter **y**. The changes are made active.

11. In Config mode enter **show active**. The message below is an example of information that is displayed.

```
rs#system show active
Running system configuration:
!
! Last modified from Console on Mon Mar 27 12:12:19 2002
!
1 : system set name "rs"
2 : system set location "Houston, TX"
3 : system set contact "John Smith"
4 : system set hashed-password login jNIssH
c976b667e681d03ccd5fc527f219351a
5 : system set hashed-password enable zcGzbO
5d1f73d2d478ceaa062a0b5e0168f46a
6 : system set hashed-password diag jdfbyp
67e681d3d2d478cf21935a0b5e016f2193
```

12. Negate the command containing the forgotten password.

13. In Config mode enter **save active**. The changes are saved and the command containing the forgotten password is deleted.

14. In Config mode enter **save to startup**. The startup file is overridden with identical startup file, however, the command containing the forgotten password was deleted.

15. Reboot the device. The device is rebooted with the new configuration.

A.2 CONFIGURATION TROUBLESHOOTING

This section describes troubleshooting issues during device configuration, and is specific to startup file corruption problems that occur during the bootup. File corruption may be caused by a bad command configuration. If a device configuration problem occurs due to an incorrectly entered or corrupted command, the following methods resolve the problem:

- Erasing the Startup File
- Editing the Startup File



Note

If the fatal error is caused by a hardware problem, contact Riverstone Networks immediately.

A.3 ERASING THE STARTUP FILE



Note

This process deletes all configuration commands.

The Startup file can be erased from the FLASH. To erase the Startup file from the FLASH perform the following.

1. Reboot the device. The following message is displayed.

```
----- Performing the Power-On Self Test (POST) -----
UART Channel Loopback Test.....PASS
Testing the System SDRAM.....PASS
EPROM Checksum Test.....PASS
Flash Image Validation Test.....PASS
Testing CPU PCI Bus Device Configuration.....PASS

BOOT Software Version Prom-2.00.00 Built 22-Mar-2002 12:57:00
Processor: MPC8245 Rev 0.12, 266 MHz (Bus: 133MHz), 64 MByte SDRAM.
I-Cache 16 KB, linesize 32.D-Cache 16 KB, linesize 32.
Cache Enabled.
Autoboot in 2 seconds - press RETURN or Esc. to abort and enter prom.
```



Note

If no note is displayed, check the terminal; baud rate, and communication settings.

2. Press Enter. The following message is displayed.

```
*****
River Stone Networks TM
*****

Startup menu

[1] Download sw
[2] Erase from Flash
[3] Erase Flash
[4] Erase Nvram

Enter your choice:
```

3. Within 2 seconds, select option “[2] Erase file from flash”. The following message is displayed.

```
Warning! About to erase the file from flash
Are you sure (Y/N)? Flash file name (Up to 8 characters, Enter
for none.
```

4. Enter **startup**. The startup file is erased, and the device reboots.

**Note**

If the incorrect file name is entered, the startup file is not erased and the device automatically reboots with the current startup file.

The following message is displayed.

```
S> en
%SYS-W-NOPASSWD, no password for enable, use 'system set
password' in Config mod
e
rs#system show startup-config
!
! Startup configuration for the next system reboot
vfs_process_cat: Cannot open /int-flash/cfg/startup
rs#system show active-config
%CONFIG-I-NOCONFIG, the running system has no configuration
rs#system show scratchpad
%CONFIG-I-NOCHANGES, there are no non-committed changes
rs#
```

5. Reconfigure the device as required.

A.4 EDITING THE STARTUP FILE

The Startup file can be edited. The Startup file is edited by copying the startup file to the Scratchpad file to view the problem. The file is copied to the Active file for editing, and then copied to the Startup file for booting. To edit the Startup file perform the following:

1. Reboot the device. The following message is displayed.

```
----- Performing the Power-On Self Test (POST) -----
UART Channel Loopback Test.....PASS
Testing the System SDRAM.....PASS
EPROM Checksum Test.....PASS
Flash Image Validation Test.....PASS
Testing CPU PCI Bus Device Configuration.....PASS

BOOT Software Version Prom-2.00.00 Built 22-Mar-2002 12:57:00
Processor: MPC8245 Rev 0.12, 266 MHz (Bus: 133MHz), 64 MByte SDRAM.
I-Cache 16 KB, linesize 32.D-Cache 16 KB, linesize 32.
Cache Enabled.
Autoboot in 2 seconds - press RETURN or Esc. to abort and enter prom.
```



Note

If no note is displayed, check the terminal, baud rate, and communication settings.

2. Press Enter. The following message is displayed.

```
*****
      River Stone Networks TM
*****

Startup menu

[1] Download sw
[2] Erase from Flash
[3] Erase Flash
[4] Erase Nvram

Enter your choice:
```

3. Within 2 seconds, select option “[2] Erase file from flash”. The following message is displayed.

```
Warning! About to erase the file from flash
Are you sure (Y/N)? Flash file name (Up to 8 characters, Enter
or none.
```

4. Enter **config**.



Note

If the incorrect file name is entered, the config file is not erased and the device automatically reboots with the current startup file.

The config file is erased, and the device reboots. The Startup file is rebooted intact, while both the Configuration file and the Scratchpad are empty. The following message is displayed.

```
rs#system show active-config
%CONFIG-I-NOCONFIG, the running system has no configuration
rs#system show active-config
%CONFIG-I-NOCONFIG, the running system has no configuration
rs#system show startup-config
! Startup configuration for the next system reboot
rs#
rs#system show scratchpad
%CONFIG-I-NOCHANGES, there are no non-committed changes
rs>
```

5. Enter **enabled** and press Enter. The prompt is displayed.

```
rs#
```

6. Enter **copy startup to scratchpad** and press Enter.
7. Enter **configure** and press Enter. The prompt is displayed.

```
rs(config)#
```

8. Enter **show**. The following message is displayed.

```
%CONFIG-I-NOCONFIG, the running system has no configuration
***** Non-committed changes in Scratchpad *****
 1*: snmp set target 176.240.10.33 community public status enable
 2*: snmp set target 176.240.10.33 community asd status disable
rs(config)#
rs(config)#
Do you want to make the changes Active [yes]?
```

**Note**

The non-committed changes shown above are examples only. The non-committed changes that display, represent the actual problem and may not be represented here.

9. Enter **configure** and press Enter.
10. Enter **show**. The following message is displayed.

```
Running system configuration:
 1 : snmp set target 176.240.10.33 community public status
enable
 2 : snmp set target 176.240.10.33 community asd status disable
rs(config)#
```

11. Enter **negate 1** to remove the first corrupt command in the list and press Enter. Continue to negate other corrupt commands as required. The following message is displayed.

```
Do you want to make the changes Active [yes]?
```

12. Enter **yes**. A prompt is displayed.
13. Enter **copy active to startup** and press enter.
14. Enter **reboot** and press Enter.

**Note**

If the changes are not activated, the device reboots with its self-configuration.