



SRstackware[®] Intelligent Network Software

Troubleshooting Guide

P/N: 6806800N83F

August 2022



Legal Disclaimer*

SMART Embedded Computing, Inc. (SMART EC), dba Penguin Solutions™, assumes no responsibility for errors or omissions in these materials. **These materials are provided "AS IS" without warranty of any kind, either expressed or implied, including but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or non-infringement.** SMART EC further does not warrant the accuracy or completeness of the information, text, graphics, links, or other items contained within these materials. SMART EC shall not be liable for any special, indirect, incidental, or consequential damages, including without limitation, lost revenues or lost profits, which may result from the use of these materials. SMART EC may make changes to these materials, or to the products described therein, at any time without notice. SMART EC makes no commitment to update the information contained within these materials.

Electronic versions of this material may be read online, downloaded for personal use, or referenced in another document as a URL to a SMART EC website. The text itself may not be published commercially in print or electronic form, edited, translated, or otherwise altered without the permission of SMART EC.

It is possible that this publication may contain reference to or information about SMART EC products, programming, or services that are not available in your country. Such references or information must not be construed to mean that SMART EC intends to announce such SMART EC products, programming, or services in your country.

Limited and Restricted Rights Legend

If the documentation contained herein is supplied, directly or indirectly, to the U.S. Government, the following notice shall apply unless otherwise agreed to in writing by SMART EC.

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data clause at DFARS 252.227-7013 (Nov. 1995) and of the Rights in Noncommercial Computer Software and Documentation clause at DFARS 252.227-7014 (Jun. 1995).

SMART Embedded Computing, Inc., dba Penguin Solutions

2900 S. Diablo Way, Suite 190

Tempe, Arizona 85282

USA

*For full legal terms and conditions, visit <https://www.penguinsolutions.com/edge/legal/>

Table of Contents

About this Manual	9
1 Introduction	13
1.1 Useful System Commands and Utilities	13
2 Debugging and Logging	17
2.1 Introduction	17
2.2 Debugging	17
2.2.1 Logging Information to a File or Syslog	17
2.2.2 Turning Off Debugging	18
3 Troubleshooting OSPF	19
3.1 Introduction	19
3.2 No OSPF Adjacency	19
3.3 Useful Show Commands	21
3.3.1 show ip ospf interface	21
3.3.2 show ip ospf neighbor	22
3.3.3 show ip ospf database	23
4 Troubleshooting RIP	25
4.1 Introduction	25
4.2 No RIP Adjacency	25
4.3 Useful Show Commands	27
4.3.1 show ip rip interface	27
4.3.2 show ip rip database	28
5 Troubleshooting VRRP	31
5.1 Introduction	31
5.2 Incorrect VRRP States	32
6 Route Selection in NSM	33
6.1 Introduction	33

Table of Contents

6.2	How Does NSM Add Routes	33
6.3	How does NSM Delete Routes	35
6.4	Show Commands	36
6.4.1	show ip route	36
6.4.2	show ip route database	38
7	Miscellaneous Issues	41
7.1	Kernel Does Not Notify the NSM about Updating the MTU/Metric	41
7.2	OSPF Adjacency Lost (System Clock)	41
7.3	Remote Devices are Unreachable	41
A	Related Documentation	43
A.1	Penguin Solutions Documentation	43

List of Figures

Figure 5-1	VRRP Topology	31
Figure 6-1	NSM Route Selection Flowchart	34
Figure 6-2	NSM Route Deletion Flowchart	35

List of Figures

List of Tables

Table 1-1	List of Commands	13
Table 3-1	OSPF Properties	19
Table 3-2	show ip ospf interface - Description of Displayed Fields	21
Table 3-3	show ip ospf neighbor - Description of Displayed Fields	23
Table 3-4	show ip ospf database - Description of Displayed Fields	24
Table 4-1	RIP Properties	25
Table 4-2	show ip rip database - Description of Display Fields	29
Table 5-1	Incorrect VRRP States	32
Table 6-1	Default Administrative Distances of Protocols	33
Table A-1	Penguin Solutions Documentation	43

List of Tables

About this Manual

Overview of Contents

This guide is intended for all network administrators and application developers who install and configure SRstackware® IP routing software. It requires that the user has a broad understanding of networking principles and network configuration. Use this information with the other technical information available with the software.

This guide contains tips for troubleshooting basic issues faced during installation, configuration and management of SRstackware IP routing software.

This manual is divided into the following chapters and appendix.

Chapter 1, Introduction on page 13

Chapter 2, Debugging and Logging on page 17

Chapter 3, Troubleshooting OSPF on page 19

Chapter 4, Troubleshooting RIP on page 25

Chapter 5, Troubleshooting VRRP on page 31

Chapter 6, Route Selection in NSM on page 33

Chapter 7, Miscellaneous Issues on page 41

Appendix A, Related Documentation on page 43

Abbreviations

This document uses the following abbreviations:


Abbreviation	Definition
BGP	Border Gateway Protocol
FIB	Forwarding Information Base
GNU	GNU's Not Unix
LDP	Label Distribution Protocol
LSA	Link-State Advertisement
MTU	Maximum Transmission Unit
NSM	Network Services Module
OSPF	Open Shortest Path First







About this Manual

Abbreviation	Definition
RIP	Routing Information Protocol
SRstackware	Switching and Routing stackware
VRRP	Virtual Router Redundancy Protocol

Conventions

The following table describes the conventions used throughout this manual.

Notation	Description
0x00000000	Typical notation for hexadecimal numbers (digits are 0 through F), for example used for addresses and offsets
0b0000	Same for binary numbers (digits are 0 and 1)
bold	Used to emphasize a word
Screen	Used for on-screen output and code related elements or commands. Sample of Programming used in a table (9pt)
Courier + Bold	Used to characterize user input and to separate it from system output
<i>Reference</i>	Used for references and for table and figure descriptions
File > Exit	Notation for selecting a submenu
<text>	Notation for variables and keys
[text]	Notation for software buttons to click on the screen and parameter description
...	Repeated item for example node 1, node 2, ..., node 12
.	Omission of information from example/command that is not necessary at the time
..	Ranges, for example: 0..4 means one of the integers 0,1,2,3, and 4 (used in registers)
	Logical OR
	Indicates a hazardous situation which, if not avoided, could result in death or serious injury

Notation	Description
	<p>Indicates a hazardous situation which, if not avoided, may result in minor or moderate injury</p>
	<p>Indicates a property damage message</p>
	<p>Indicates a hot surface that could result in moderate or serious injury</p>
	<p>Indicates an electrical situation that could result in moderate injury or death</p>
<p>Use ESD protection</p> 	<p>Indicates that when working in an ESD environment care should be taken to use proper ESD practices</p>
	<p>No danger encountered, pay attention to important information</p>

Summary of Changes

This manual has been revised and replaces all prior editions.

Part Number	Publication Date	Description
6806800N83F	August 2022	Rebrand to Penguin Solutions.
6806800N83E	March 2020	Rebrand to SMART Embedded Computing template. Updated Section 7.1.
6806800N83D	July 2017	Added registered trademark to SRstackwarea.
6806800N83C	June 2014	Rebrand to Artesyn template.
6806800N83B	October 2012	Added Note4s in the document to state.
6806800N83A	February 2012	EA Release.

Introduction

1.1 Useful System Commands and Utilities

These Unix Commands and GNU Utilities are useful in troubleshooting. Unix commands might have different syntax when used on different platforms.



Use `man <command>` to get detailed information about any Unix command.

Table 1-1 List of Commands

Command	Description
id	<p>This <code>id</code> utility displays the system identifications (ID) for a specified user. The system IDs identify users and user groups to the system. This utility displays the user name, user ID, as well as the group name and group ID of the user.</p> <p>In addition, <code>id</code> also displays the effective user and group IDs (<code>euid</code>). Install SRstackware as a root user. Use this command to verify that you are the root user.</p>
ifconfig	<p>Configures a network interface. It is used at boot time to set up interfaces or for debugging purposes. Use the <code>-a</code> flag with this command to instruct <code>ifconfig</code> to display information about all interfaces on the system.</p> <p>You can use this utility to configure the loopback address:</p> <pre>ifconfig lo 127.0.0.1</pre>
ls	Lists the contents of the current working directory.
man	<p>It provides access to online manual pages and information on how to run a specified command. For example, to learn more about <code>rm</code> (file-removal) command type:</p> <pre>man rm</pre> <p>Use option <code>-k</code> with the <code>man</code> command if you do not know which command to look for. This option directs <code>man</code> to search for manual pages containing the specified keyword.</p> <p>If the information is more than one screenful of text, the <code>man</code> command shows the first screen and prompts you with More at the bottom of the screen. Hit the spacebar when you are ready for the next screenful. Type <code>q</code> to quit.</p>

Introduction

Table 1-1 List of Commands (continued)

Command	Description
netstat	Displays different network related data structures. You can use various options to get different outputs, such as <code>-r</code> option displays routing tables. The <code>-n</code> option used along with the <code>-r</code> option displays network addresses as numbers.
ping	It is used to see if a system is operating and also to see if network connections are intact. It uses the Internet Control Message Protocol (ICMP) Echo function (detailed in RFC 792). A small packet is sent through the network to a particular IP address. The sender then listens for a return packet, if connections are good and the target system is up, a good return packet is received.
ping6	It uses the ICMPv6 Echo function (detailed in RFC 2463) to report errors encountered in processing packets and to perform diagnostics. This utility is available on Linux and FreeBSD systems. On a Solaris system the ping utility has both IPv4 and IPv6 capability.
ssh	Use Secure Shell (SSH) to log into another computer over a network, to execute commands from a remote machine, and to move files from one machine to another. SSH can be used in place of <code>telnet</code> , <code>rlogin</code> , <code>rsh</code> etc. It provides authentication and secure communications over insecure channels.
su	The <code>su</code> command changes the user ID to those of the root user or to any other specified user.
telnet	It is a user Interface to the TELNET protocol. It runs on your computer and connects it to a server on the network. You can then enter commands through the telnet program and they are executed as if you were entering them directly on the server console. This enables you to control the server and communicate with other systems on the network.
traceroute	Traces a packet from your system to a host showing the number of hops the packet requires to reach the host and how long each hop takes.
traceroute6	Displays information about the route taken by the IPv6 packets to reach the destination. This utility is available on Linux and FreeBSD systems. On a Solaris system the traceroute utility has IPv4 and IPv6 capability.
uname	Displays information about the name and version of the current operating system.

Table 1-1 List of Commands (continued)

Command	Description
useradd	Creates a new User or updates default user information. This utility works on Linux and Solaris systems. On a FreeBSD system use the pw utility to create a new user.
userdel	Deletes a user account and its related files. This utility works on Linux and Solaris systems. On a FreeBSD system use the pw utility to delete a user account.
which	Shows the complete path of the given command. If the command is missing in the path, a message is displayed stating that the command is missing. For example, using <code>which telnet</code> confirms if telnet is installed on your system and gives the path to reach it.

Debugging and Logging

2.1 Introduction

SRstackware® has a comprehensive debugging and logging facility in various protocols and components. This chapter describes how to start/stop debugging and logging using the NSM commands. For complete information about the logging commands, refer to the Layer 2 Command Reference, Layer 3 Command Reference, and the Switch Configuration Command Reference. The protocol debug commands are in corresponding command reference manuals.

2.2 Debugging

In the SRstackware implementation, every protocol has debug commands. Debug commands, when used with the parameters, log parameter-specific information. For example, using the `debug ldp nsm` command, results in the router writing all messages exchanged between LDP and NSM such as: interface, bandwidth and address updates.

On using a debug command, the router continues to generate output until the `no` parameter is used with the command. The debug output and system error messages are written on the virtual terminal. Use the logging commands in the configure mode to redirect the debugging output to a file or syslog. You can set the logging levels by using parameters with these commands. Refer to the *Switch Configuration Command Reference* for details on these commands.

2.2.1 Logging Information to a File or Syslog

To send logging information to a file:

1. Use the `log file` command and specify the path and file name where the information is to be logged.
2. Turn on the debug option by using the relevant debug command.

```
# enable
# configure terminal
(config)# log file <filename>
(config)# debug <protocol> (parameter)
```
3. To log information in the system log, use the `log syslog` command:

```
(config)# log syslog
```

Debugging and Logging

The system log enables logging and analyzing configuration events and system error messages centrally. This helps in monitoring interface status, security alerts and CPU process overloads. It also allows real-time capturing of client debug output sessions. Use the **no** parameter with this command to disable system logging:

```
(config)# no log syslog
```

2.2.2 Turning Off Debugging

To turn off debugging, use the **no debug** or **undebug** command. When a protocol is specified with the **no debug** or **undebug** commands, debugging is stopped for the specified protocol. To stop all debugging, use the **all** parameter with these commands.

```
(config)# no debug bgp events
```

or

```
# undebug all
```

To turn off logging information to a file, use the **no** parameter with the **log file** command:

```
(config)# no log file (filename)
```

Troubleshooting OSPF

3.1 Introduction

In this chapter the topics are arranged sequentially. Depending on the event and time when the problem occurred, select the relevant section and follow steps sequentially. If the issue is not resolved, refer to the [Miscellaneous Issues](#) chapter in this document.

Refer to the *Product Name Short*[®] *OSPF Command Reference* for details on the commands used in this chapter.



This chapter is applicable only if LAYER3SRS is licensed.

3.2 No OSPF Adjacency

The following table lists the different properties of the commands.

Table 3-1 OSPF Properties

Property	Description
Interface Status	<p>Use the <code>show ip interface brief</code> command to make sure that the interface is not administratively shutdown. Remove this configuration setting with the <code>no shutdown</code> command, if shutdown is configured.</p> <pre># configure terminal (config)# interface eth0 (config-if)# no shutdown</pre> <p>Use the <code>show interface</code> command to make sure that the interface is up.</p>
Passive Interface	<p>Make sure that interface is not configured as a passive interface using the <code>show run</code> command:</p> <pre>! router ospf passive interface eth0 !</pre> <p>If the interface is configured as passive (as shown above), remove this configuration setting by using this command:</p> <pre>no passive interface eth0</pre>

Troubleshooting OSPF

Table 3-1 OSPF Properties (continued)

Property	Description
OSPF Enabled on the Interface	<p>Make sure that OSPF is enabled on the interface. To enable OSPF on a particular interface, use the network area command with a specified Area ID. Use the show ip ospf interface to confirm that OSPF is enabled for the interface.</p> <p>Sample output</p> <pre>eth2 is up, line protocol is up Internet Address 56.168.1.7/24, Area 0.0.0.0, MTU 1500 Router ID 7.7.7.7, Network Type BROADCAST, Cost: 10 Transmit Delay is 1 sec, State DR, Priority 1 Designated Router (ID) 7.7.7.7, Interface Address 56.168.1.7 Backup Designated Router (ID) 8.8.8.8, Interface Address 56.168.1.8 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5 Hello due in 00:00:05 Neighbor Count is 1, Adjacent neighbor count is 1 Crypt Sequence Number is 0 Hello received 625 sent 645, DD received 3 sent 4 LS-Req received 1 sent 1, LS-Upd received 5 sent 13 LS-Ack received 8 sent 5, Discarded 0</pre>
Exchange of Hello Packets	<p>Check on the interface to make sure that OSPF Hello packets are being sent and received on the interface. You can use either packet sniffer (such as, Ethereal or TCP dump) or Product Name Short log messages to verify the hello packet.</p> <p>To turn on Product Name Short logging, type:</p> <pre># configure terminal (config)# debug ospf event (config)# debug ospf packet hello</pre> <p>To display the logging message on the terminal, type:</p> <pre># terminal monitor</pre>
Mismatch between Hello Parameters	<p>It is possible that there is a mismatch between Hello parameters. Make sure that you have specified the same hello interval and dead interval values on both machines by using the show ip ospf interface command on each machine.</p>
Mismatch between MTU sizes	<p>Run show ip ospf neighbor, if you see the neighbor but the state is not full. Make sure that both routers have the same MTU size for the interfaces.</p>

3.3 Useful Show Commands

The following sections list different `show` commands.

3.3.1 `show ip ospf interface`

This command displays interface information for OSPF.

When to use this command

Use this command when you want to check the interfaces enabled for an OSPF process. It provides important information about the OSPF parameters. Confirm that the OSPF parameters match that of the neighbors. If the intended interfaces are not shown in the OSPF information, check the configuration to make sure that the IP address of the missing interface is included.

Sample output

```
eth1 is up, line protocol is up
  Internet Address 10.100.10.72/24, Area 0.0.0.0, MTU 1500
  Router ID 100.100.100.72, Network Type BROADCAST, Cost: 10, TE Metric 0
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 100.100.100.72, Interface Address 10.100.10.72
  Backup Designated Router (ID) 10.100.12.57, Interface Address
10.100.10.105
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:05
  Neighbor Count is 1, Adjacent neighbor count is 1
  Crypt Sequence Number is 0
  Hello received 19 sent 106, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 3 sent 3
  LS-Ack received 2 sent 3, Discarded 0
```

Table 3-2 `show ip ospf interface` - Description of Displayed Fields

Displayed Field	Description
Internet Address	The IP address and subnet mask of the interface.
Area	The OSPF area to which the interface belongs.
MTU	The Maximum Transmission Unit (MTU) of the interface.

Troubleshooting OSPF

Table 3-2 `show ip ospf interface` - Description of Displayed Fields (continued)

Displayed Field	Description
Transmit Delay	The transmit delay of the interface.
Priority	The OSPF priority of the current interface. It is used for election of Designated Router (DR) and Backup Designated Router (BDR).
Hello	The OSPF hello-interval.
Dead	The OSPF dead-interval.
Wait	The Hello wait-interval.
Retransmit	The period, in seconds, for which the router waits between retransmissions of OSPF packets that have not been acknowledged.
Hello due in	The time period for which router expects to receive hello packet.
Neighbor Count	The OSPF neighbor count.
Adjacent neighbor	The OSPF adjacent neighbor count.
Crypt Sequence Number	This is used for authentication.
Hello received 19 sent 106, DD received 4 sent 3	This line shows that this router received 19 and sent 106 Hello packets. It has received 4 and sent 3 DD packets out.
LS-Req received 1 sent 1, LS-Upd received 3 sent 3	This line indicates that this router received and sent 1 LSA request. It sent and received 3 LSA updates.
LS-Ack received 2 sent 3, Discarded 0	This line indicates that this router received 2 and sent 3 LSA acknowledgments. It discarded no LSA acknowledgment.

3.3.2 `show ip ospf neighbor`

This command displays information about OSPF neighbors.

When to use this command

Use this command to check OSPF neighbors and their states. If the expected neighbors do not show the OSPF information, make sure that the OSPF parameters match the intended neighbors and they are configured in the same area.

Sample output

```
# show ip ospf neighbor
OSPF process 100:
```

```
Neighbor ID  Pri  State  Dead Time  Address  Interface  RXmtL RqstL
DBsmL
10.100.12.57 1 Full/Backup 0:00:37 10.100.10.105 eth1:10.100.10.72 0 0 0
```

Table 3-3 *show ip ospf neighbor - Description of Displayed Fields*

Displayed Fields	Description
OSPF process	The OSPF process involved.
Neighbor ID	The OSPF Router ID of the neighbor.
Pri	The OSPF priority of the neighbor.
State	The functional state of the OSPF neighbor.
Dead Time	If a new Hello is not received within this duration, the neighbor is declared dead.
Address	The IP address of neighbor's interface attached to the network.
Interface	The interface attached to the network on which the neighbor is located.

3.3.3 show ip ospf database

This command displays information about OSPF link-state database on the router.

Sample output

```
QA72# show ip ospf database
```

```
OSPF Router process 100 with ID (100.100.100.72)
```

```
Router Link States (Area 0.0.0.0)
```

```
Link ID          ADV Router      Age   Seq#           CkSum  Link count
10.100.12.57    10.100.12.57   930  0x80000003  0x90de  2
100.100.100.72 100.100.100.72 933  0x80000004  0x7592  2
```

```
Net Link States (Area 0.0.0.0)
```

```
Link ID          ADV Router      Age   Seq#           CkSum
10.100.10.72    100.100.100.72 933  0x80000001  0x0bef
```

Troubleshooting OSPF

Summary Link States (Area 0.0.0.0)

Link ID	ADV Router	Age	Seq#	CkSum	Route
10.60.0.0	10.100.12.57	928	0x80000001	0x5108	10.60.0.0/24
71.87.120.0	10.100.12.57	928	0x80000001	0xc2c5	71.87.120.0/24
127.0.0.1	10.100.12.57	928	0x80000001	0x23fb	127.0.0.1/32

Table 3-4 *show ip ospf database - Description of Displayed Fields*

Displayed Fields	Description
Link ID	Link ID has a different meaning for different types of Link-State Advertisements (LSAs).
	Link ID for Router Link States depends on the type of network the router connects to: <ul style="list-style-type: none"> ● Point to point network - Neighbor's Router ID ● Transit network - IP address of the Designated Router's interface ● Stub network - IP network or subnet address ● Virtual link - Neighbor's Router ID
	Link ID for Net Link States: The IP address of the DR's interface.
	Link ID for Summary Link States: The IP address of the network or subnet being advertised.
ADV Router	The router ID of the router advertising the LSA.
Age	The age of the LSA.
Seq#	The sequence number of the LSA. This number increments each time a new instance of the LSA originates. This update helps other routers identify the most recent instance of the LSA.
CkSum	The Fletcher checksum of the complete LSA except the Age field.

Troubleshooting RIP

4.1 Introduction

In this chapter the topics are arranged sequentially. Depending on the event and time when the problem occurred, select the relevant section and follow steps sequentially. If the issue is not resolved, refer to the [Miscellaneous Issues](#) chapter in this document.

Refer to the *SRstackware® RIP Command Reference* for details on commands used in this chapter.



This chapter is applicable only if LAYER3SRS is licensed.

4.2 No RIP Adjacency

[Table 4-1](#) lists the different properties of the commands.

Table 4-1 RIP Properties

Property	Description
Interface Status	<p>Use the <code>show ip interface brief</code> command to make sure that the interface is not administratively shutdown. Remove this configuration using the <code>no shutdown</code> command, if shutdown is configured.</p> <pre># configure terminal (config)# interface eth0 (config-if)# no shutdown</pre> <p>Use the <code>show interface</code> command to make sure that the interface is up.</p>
Passive Interface	<p>Make sure that the interface is not configured as a passive interface using the <code>show run</code> command:</p> <pre>! router rip passive interface eth0 !</pre> <p>If the interface is configured as passive (as shown above), remove this configuration setting by using this command:</p> <pre>no passive interface eth0</pre>

Troubleshooting RIP

Table 4-1 RIP Properties (continued)

Property	Description
RIP enabled on Interface	<p>Confirm that RIP is enabled on the interface. To enable RIP on a particular interface, use the network command. Use the show ip rip interface to make sure that RIP is enabled for the interface.</p> <p>Sample output</p> <pre># show ip rip interface fxp0 is up, line protocol is up Routing Protocol: RIP Receive RIP packets Send RIP packets Passive interface: Disabled Split horizon: Enabled with Poisoned Reversed IP interface address: 10.15.0.60/16</pre>
Exchange of RIP Advertisements	<p>Make sure that RIP advertisements are being sent and received on the interface. You can use either a packet sniffer (such as, Ethereal or TCP dump) or the SRstackware log messages to verify the RIP advertisements.</p> <p>To turn on SRstackware logging, type:</p> <pre># configure terminal (config)# debug rip event (config)# debug rip packet detail</pre> <p>To display the logging message on the terminal, type:</p> <pre># terminal monitor</pre>
RIP Version Mismatch	<p>One router configured as RIPv1 and the other router as RIPv2 results in no RIP adjacency.</p> <p>Configure the router running RIPv2 as follows:</p> <pre>! interface eth1 ip rip send version 1-compatible ip rip receive version 1 2 !</pre>

Table 4-1 RIP Properties (continued)

Property	Description
Firewall	<p>Verify if a firewall is present. If there is a firewall, it blocks the UDP packet. You must remove the firewall if you have one. To display the existing firewall configurations, in Linux, use:</p> <pre>ipchains -L</pre> <p>Flush the existing firewall configurations by using:</p> <pre>ipchains -F</pre>

4.3 Useful Show Commands

The following sections list different **show** commands.

4.3.1 show ip rip interface

This command displays information about RIP interfaces.

When to use this command

Use this command to verify that RIP is enabled on an interface.

Sample output

```
# show ip rip interface eth1
eth1 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.10.10/24
```

Line by line description In the above output:

```
eth1 is up, line protocol is up
  Routing Protocol: RIP
```

These lines denote that the interface is UP and RIP is enabled.

```
Receive RIP packets
```

Troubleshooting RIP

Send RIP packets

These lines indicate that the interface is capable of receiving/sending both RIP version 1 and 2 packets, which is the default.

If RIP is configured to send only version 1 packets using the `ip rip send version 1` command, the output displays: `Send RIPv1 packets only`

Passive interface: Disabled

This line denotes that the specified interface is not passive and can send and receive RIP updates.

If passive interface is configured using the `passive-interface <IFNAME>` command, RIP updates are received but not sent. This configuration is required when a router does not want to advertise itself but still wants to learn RIP routes.

Split horizon: Enabled with Poisoned Reversed

This line denotes that the `split-horizon with poisoned reversed` feature is enabled on the displayed interface (`eth1`). This means that routes will not be advertised on the interface from which they are learned avoiding the problem of counting to infinity.

IP interface address:

```
10.10.10.10/24
```

These lines display the IPv4 address of the RIP enabled interface.

4.3.2 show ip rip database

This command displays the different routes learned by RIP.

When to use this command

Use this command to view details of all the routes learned by RIP.

Sample output

```
# show ip rip database
```

```
Codes: R - RIP, K - Kernel, C - Connected, S - Static, O - OSPF, I - IS-IS,  
       B - BGP
```

	Network	Next Hop	Metric From	If	Time
R	192.1.1.0/24	10.10.10.68	2 10.10.10.68	eth1	02:47
O	193.1.1.0/24	20.10.10.50	2 20.10.10.50	eth2	03:06
S	196.6.6.0/24		1	eth1	

Line by line description In the above output:

Codes: R - RIP, K - Kernel, C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP

Refer to the `show ip route` command description in [Section 6.4.1 on page 36](#) for details about these codes.

Table 4-2 `show ip rip database` - Description of Display Fields

Display Field	Description
Network	Is the network prefix
Next hop	Is the IPv4 address of the nexthop router.
Metric	Is the metric to reach the network prefix.
From	Is the IPv4 address of the neighbor's interface.
If	Is the local router's interface through which the router reaches the Network prefix.
Time	Is the duration for which the network prefix is stored in the RIP routing table.

```
R 192.1.1.0/24      10.10.10.68      2 10.10.10.68    eth1 02:47
```

This line denotes a RIP route learned from a neighbor (10.10.10.68) through interface eth1. This route belongs to the 192.1.1.0/24 network, its metric value is 2 and its nexthop is 10.10.10.68. It will remain in the RIP routing table for 2 minutes 47 seconds.

```
O 193.1.1.0/24      20.10.10.50      2 20.10.10.50    eth2 03:06
```

This line denotes an OSPF route learned from a neighbor (20.10.10.50) through interface eth2. This route belongs to the 193.1.1.0/24 network, its metric value is 2 and nexthop is 20.10.10.50. It will remain for 3 minutes and 6 seconds in the RIP routing table.

```
S 196.6.6.0/24      1                  eth1
```

This line denotes a static route connected through interface eth1. It has a metric value of 1 to reach the network prefix.

Troubleshooting VRRP

5.1 Introduction

In this chapter the topics are arranged sequentially. Depending on the event and time when the problem occurred, select the relevant section and follow steps sequentially. If the issue is not resolved, refer to the [Miscellaneous Issues](#) chapter in this document.


Refer to the *SRstackware® VRRP Command Reference* for details on commands used in this chapter.



This chapter is applicable only if LAYER3SRS is licensed.

The following figure shows the topology used for illustration.

Figure 5-1 VRRP Topology



5.2 Incorrect VRRP States

The following table lists the different properties of the commands.

Table 5-1 Incorrect VRRP States

Property	Description
Interface running	<p>Make sure the interfaces are up and running by using the <code>show interface</code> command. In the following sample output interface eth1 is:</p> <pre># show interface eth1 Interface eth1 Hardware is Ethernet, address is 0002.b3d4.436f index 3 metric 1 mtu 1500 <UP,BROADCAST,RUNNING,MULTICAST> ... input errors 0, length 0, overrun 0, CRC 0, frame 0, fifo 0..</pre> <p>If the interface is down:</p> <p>Use <code>no shutdown</code> command, in the interface mode, to bring up the interface.</p> <p>Or</p> <p>Use the <code>ifconfig <IFNAME> up</code> command to bring up the interface.</p>
Reachability of routers	<p>Make sure that both VRRP routers can reach each other by pinging.</p> <pre># ping 10.10.11.2 PING 10.10.11.2 (10.10.11.2) 56(84) bytes of data. 64 bytes from 10.10.11.2: icmp_seq=1 ttl=255 time=0.202 ms 64 bytes from 10.10.11.2: icmp_seq=2 ttl=255 time=0.201 ms</pre> <p>If both routers cannot reach each other, check the network connections for the default Master and default Backup routers.</p>
Advertisement Intervals on both routers	<p>Check the advertisement interval on Master and Backup routers. The advertisement interval must be the same on both.</p> <p>The default advertisement interval = 1.</p> <p>Use the <code>advertisement-interval</code> command, in Router mode, to configure the advertisement interval.</p>

Route Selection in NSM

6.1 Introduction

Understanding the NSM route selection process helps in analyzing and troubleshooting route-related problems. This chapter describes the route selection process in NSM. It also describes relevant show commands and their outputs.

For every known prefix, NSM maintains a route node entry in its route table. NSM populates this table upon receiving routes from clients (BGP, OSPF, RIP), from static routes configured using CLI, from the kernel's Forwarding Information Base (FIB) or connected routes derived from interface information.

When multiple routes are available for the same prefix, NSM uses an internal route selection mechanism to select routes to be added to the FIB. The primary factor for route selection is the Administrative Distance of the protocol.



IPv4 Static routing is available in Basic SRS. Other routing protocols are supported only if LAYER3SRS is licensed.

The following table lists the default administrative distances of protocols.

Table 6-1 Default Administrative Distances of Protocols

Protocol	Administrative Distance	Preference
Connected	-	1 (highest)
Kernel	-	2
Static	1	3
eBGP	20	4
OSPF	110	5
ISIS	115	6
RIP	120	7
iBGP	200	8 (lowest)

6.2 How Does NSM Add Routes

NSM prefers routes learned from protocols with a lower administrative distance over routes learned from protocols with a higher administrative distance.

Route Selection in NSM

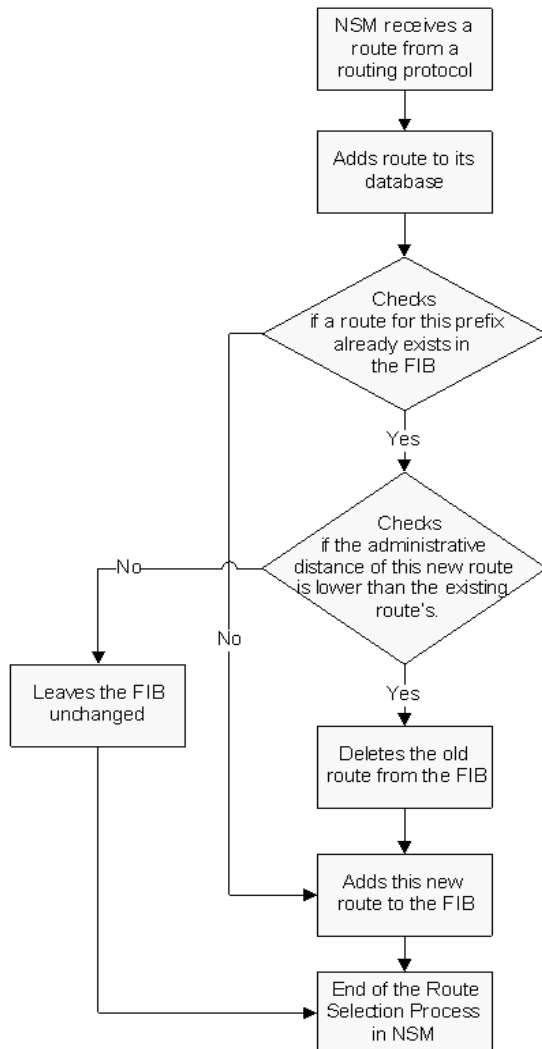
For example, the following route entries display that the Static Routes (administrative distance 1) is preferred over the OSPF Route (administrative distance 110):

```
S  *-> 10.10.34.0/24 [1/0] via 10.10.31.16, eth2
O  10.10.34.0/24 [110/31] via 10.10.31.16, eth2, 00:21:19
```

The administrative distance of routing protocols can be configured using the distance command.

The following figure displays how a route is selected in NSM.

Figure 6-1 NSM Route Selection Flowchart

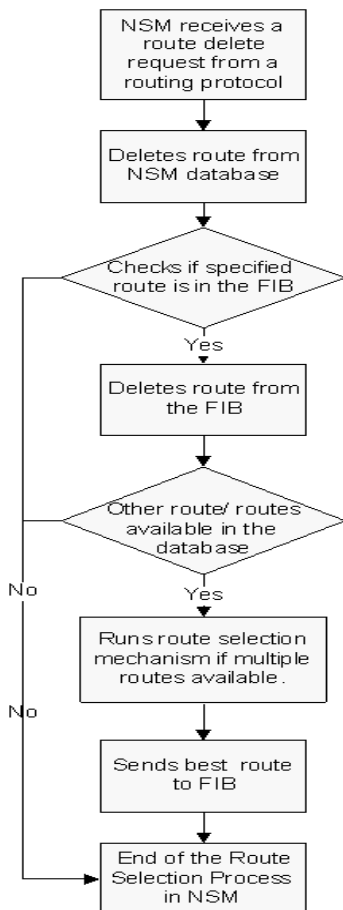


Routing protocols use different metrics to calculate the best path for a destination. The best path is sent to NSM. However, when two paths have an equal cost/metric and Equal Cost Multipath (ECMP) is enabled on a system, NSM might receive two paths from the same protocol.

6.3 How does NSM Delete Routes

When NSM receives a route delete request from a routing protocol, NSM deletes the specified route from its database. Then it checks if the specified route is in the FIB. If the route is in the FIB, NSM deletes it from the FIB and checks if another route is available in its database for the same prefix. If there is another route in the database, NSM installs this route in the FIB. When multiple such routes exist, NSM runs the route selection mechanism to choose the best route to be added to the FIB.

Figure 6-2 NSM Route Deletion Flowchart



6.4 Show Commands

The `show ip route` and the `show ip route database` commands are important tools for troubleshooting. Use these commands in conjunction to get complete information about routes received and selected by NSM. Use the `show ip route database` command to list all the routes received by NSM and use the `show ip route` command to list only routes that are selected by NSM and installed in the FIB.

6.4.1 show ip route

The `show ip route` command displays the contents of the IP routing table maintained by NSM. Routes displayed in this table are also added to the kernel routing table.

When to use this command

Use this command to view only the routing entries selected by NSM.

Sample output

```
# show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default

Gateway of last resort is 10.30.0.11 to network 0.0.0.0
K*    0.0.0.0/0 via 10.30.0.11, eth0
O     9.9.9.9/32 [110/31] via 10.10.31.16, eth2, 00:18:56
K     10.10.0.0/24 via 10.30.0.11, eth0
C     10.10.31.0/24 is directly connected, eth2
S     10.10.34.0/24 [1/0] via 10.10.31.16, eth2
O     10.10.37.0/24 [110/11] via 10.10.31.16, eth2, 00:20:54
C     10.30.0.0/24 is directly connected, eth0
S     11.22.11.0/24 [1/0] via 10.10.31.16, eth2
O E2  14.5.1.0/24 [110/20] via 10.10.31.16, eth2, 00:18:56
S     16.16.16.16/32 [1/0] via 10.10.31.16, eth2
O     17.17.17.17/32 [110/31] via 10.10.31.16, eth2, 00:20:54
C     45.45.45.45/32 is directly connected, lo
O     55.55.55.55/32 [110/21] via 10.10.31.16, eth2, 00:20:54
C     127.0.0.0/8 is directly connected, lo
```

Line by line description

Each entry in this table has a code preceding it, indicating the source of the routing entry. For example, O indicates OSPF as the origin of the route and K indicates that the route has been learned from the Kernel (most operating systems add some implicitly learnt routes when the device is started up). The first few lines of the output list the possible codes that may be seen with the route entries.

Typically, route entries are composed of the following elements:

- Code
- Network/ host ip address
- Outgoing interface name
- Administrative distance and metric
- Nexthop ip address
- Time since route entry was added
- A second Label indicating the sub type of the route.

To avoid repetition, only selected route entries comprised of different elements are described here:

O 10.10.37.0/24 [110/11] via 10.10.31.16, eth2, 00:20:54

This route entry denotes:

- This route in the network 10.10.37.0/24 was added by OSPF.
- This route has an administrative distance of 110 and metric/cost of 11.
- This route is reachable via nexthop 10.10.31.16.
- The outgoing local interface for this route is eth2.
- This route was added 20 minutes and 54 seconds ago.

C 10.10.31.0/24 is directly connected, eth2

This route entry denotes:

- Route entries for network 10.10.31.0/24 are derived from the IP address of local interface eth2.
- These routes are marked as Connected routes (C) and always preferred over routes for the same network learned from other routing protocols.
- Routes for connected networks always exist in the kernel routing table but as an exception are not marked as kernel routes because NSM always calculates entries for these routes upon learning interface information from the kernel.

Route Selection in NSM

K **10.10.0.0/24 via 10.30.0.11, eth0**

This route entry denotes:

- This route in the network 10.10.0.0/24 was learned from the kernel routing table (route was statically added using kernel commands).
- This route is reachable via nexthop 10.30.0.11.
- The outgoing local interface for this route is eth0.
- The static routes added using kernel commands and static routes added using NSM commands are different. The kernel static routes are not redistributed when the redistribute static command is used in a protocol. However, the kernel static routes can be redistributed using the redistribute kernel command.

O E2 **14.5.1.0/24 [110/20] via 10.10.31.16, eth2, 00:18:56**

This route entry denotes:

- This route is the same as the other OSPF route explained above; the only difference is that it is a Type 2 External OSPF route.

K* **0.0.0.0/0 via 10.30.0.11, eth0**

This route entry denotes:

- This is a default route and was learned from the kernel (route was statically added using kernel commands).
- This route is reachable via nexthop 10.30.0.11.
- The local interface for this route is eth0.

6.4.2 show ip route database

The show ip route database command displays all routing entries known by NSM. When multiple entries are available for the same prefix, NSM uses an internal route selection mechanism based on protocol administrative distance and metric values to choose the best route. All best routes are entered into the FIB and can be viewed using the show ip route command.

When to use this command

Use this command to view all the routing entries known by NSM.

Sample output

```
# show ip route database
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
```

```
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
> - selected route, * - FIB route, p - stale info
K *> 0.0.0.0/0 via 10.30.0.11, eth0
O *> 9.9.9.9/32 [110/31] via 10.10.31.16, eth2, 00:19:21
K *> 10.10.0.0/24 via 10.30.0.11, eth0
O 10.10.31.0/24 [110/1] is directly connected, eth2, 00:28:20
C *> 10.10.31.0/24 is directly connected, eth2
S *> 10.10.34.0/24 [1/0] via 10.10.31.16, eth2
O 10.10.34.0/24 [110/31] via 10.10.31.16, eth2, 00:21:19
O *> 10.10.37.0/24 [110/11] via 10.10.31.16, eth2, 00:21:19
K * 10.30.0.0/24 is directly connected, eth0
C *> 10.30.0.0/24 is directly connected, eth0
S *> 11.22.11.0/24 [1/0] via 10.10.31.16, eth2
O E2 *> 14.5.1.0/24 [110/20] via 10.10.31.16, eth2, 00:19:21
O 16.16.16.16/32 [110/11] via 10.10.31.16, eth2, 00:21:19
S *> 16.16.16.16/32 [1/0] via 10.10.31.16, eth2
O *> 17.17.17.17/32 [110/31] via 10.10.31.16, eth2, 00:21:19
C *> 45.45.45.45/32 is directly connected, lo
O *> 55.55.55.55/32 [110/21] via 10.10.31.16, eth2, 00:21:19
K * 127.0.0.0/8 is directly connected, lo
C *> 127.0.0.0/8 is directly connected, lo
```

Line by line description

The routes added to the FIB are marked with a *. When multiple routes are available for the same prefix, the best route is indicated with the > symbol. All unselected routes have neither the * nor the > symbol.

In the case of Connected routes, 2 entries exist in the route database; one learned from the kernel and the other derived from interface information.

```
K * 10.30.0.0/24 is directly connected, eth0
C *> 10.30.0.0/24 is directly connected, eth0
```

Route Selection in NSM

These route entries denote:

- Both these routes are in the same network 10.30.0.0/24.
- The first route has originated from the kernel. The * indicates that it has been added to the FIB (Forwarding Information Base).
- The second route is derived from the IP address of local interface eth0. It is marked as a Connected route (C). Since a Connected route has the lowest administrative distance, it is the selected route.

```
S    *> 10.10.34.0/24 [1/0] via 10.10.31.16, eth2
O      10.10.34.0/24 [110/31] via 10.10.31.16, eth2, 00:21:19
```

These route entries denote:

- The same prefix was learned from OSPF and from static route configuration.
- Since Static routes are preferred over OSPF routes, the static route is selected and installed in the FIB.



When the static route becomes unavailable, NSM automatically selects the OSPF route and installs it in the FIB.

Miscellaneous Issues

7.1 Kernel Does Not Notify the NSM about Updating the MTU/Metric

Live MTU/metric updates are sent by NSM to the protocols on Linux. To trigger this information from NSM to protocols, you need to administratively bring the interface DOWN, modify the MTU/metric and then bring the interface UP. This sends the new MTU/metric update to the protocols.

7.2 OSPF Adjacency Lost (System Clock)

When changing the system clock (moving it backward or forward) the OSPF daemon on Solaris loses adjacency. OSPF adjacency is lost and stuck in "Init" state. This is a known Solaris issue.

When changing the system time using any mechanism, users need to shutdown the system and bring it up. So when changing system time on Solaris, shutdown the system and restart SRstackware protocols.

7.3 Remote Devices are Unreachable

If you cannot reach remote devices when restarting NSM, make sure that there is at least one static IP address configured from the primary interface of the OS. Failure to do so can result in the device becoming unreachable from the outside. The connectivity established through SRstackware is lost when SRstackware is killed and the device requires manual intervention.

To configure a static address from the primary interface, use the `ifconfig` command or edit a file in the `network-scripts`:

1. Edit file `ifcfg-eth<x>` in `/etc/sysconfig/network-scripts` with this minimum configuration:

```
DEVICE = eth<x>
ONBOOT = yes
PADDR = 10.10.10.222
NETMASK = 255.255.255.0
BROADCAST = 10.10.10.255
```
2. Make sure `/etc/rc.d/init.d/network` does exist to configure a network interface with a static IP address at boot time.

Related Documentation

A.1 Penguin Solutions Documentation

Technical documentation can be found by using the Documentation Search at <https://www.penguinsolutions.com/edge/support/> or you can obtain electronic copies of documentation by contacting your local sales representative.

Table A-1 Penguin Solutions Documentation

Document Title and Source	Document Number
SRstackware Intelligent Network Software Layer 2 Command Reference	6806800N88
SRstackware Intelligent Network Software VRRP Command Reference	6806800N84
SRstackware Intelligent Network Software RIP Command Reference	6806800N85
SRstackware Intelligent Network Software Layer 2 Configuration Guide	6806800N86
SRstackware Intelligent Network Software OSPF Command Reference	6806800N87
SRstackware Application Programming Interface Developer Guide	6806800N90
SRstackware Intelligent Network Software Layer 3 Configuration Guide	6806800N89
SRstackware Intelligent Network Software Switch Configuration Command Reference	6806800N92
SRstackware Intelligent Network Software Layer 3 Command Reference	6806800N93
SRstackware Intelligent Network Software Protocol Demo Guide	6806800N07
SRstackware FAQ	6806800N91

Related Documentation

PENGUINTM

SOLUTIONS 

Penguin Solutions is a trade name used by SMART Embedded Computing, Inc., a wholly owned subsidiary of SMART Global Holdings, Inc. Penguin Edge is a trademark owned by Penguin Computing, Inc., a wholly owned subsidiary of SMART Global Holdings, Inc. All other logos, trade names, and trademarks are the property of their respective owners. ©2022 SMART Embedded Computing, Inc.