

SIEMENS

RUGGEDCOM ROX v1.16

User Guide

Preface

Introduction

1

Using ROX

2

Device Management

3

System Administration

4

Setup and Configuration

5

Upgrades

6

For RX1000, RX1000P, RX1100, RX1100P

9/2014

RC1098-EN-02

Copyright © 2014 Siemens Canada Ltd.

All rights reserved. Dissemination or reproduction of this document, or evaluation and communication of its contents, is not authorized except where expressly permitted. Violations are liable for damages. All rights reserved, particularly for the purposes of patent application or trademark registration.

This document contains proprietary information, which is protected by copyright. All rights are reserved. No part of this document may be photocopied, reproduced or translated to another language without the prior written consent of Siemens Canada Ltd..

Disclaimer Of Liability

Siemens has verified the contents of this manual against the hardware and/or software described. However, deviations between the product and the documentation may exist.

Siemens shall not be liable for any errors or omissions contained herein or for consequential damages in connection with the furnishing, performance, or use of this material.

The information given in this document is reviewed regularly and any necessary corrections will be included in subsequent editions. We appreciate any suggested improvements. We reserve the right to make technical improvements without notice.

Registered Trademarks

ROX™, Rugged Operating System On Linux™, CrossBow™ and eLAN™ are trademarks of Siemens Canada Ltd.. ROS® is a registered trademark of Siemens Canada Ltd..

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

The registered trademark Linux® is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis.

Other designations in this manual might be trademarks whose use by third parties for their own purposes would infringe the rights of the owner.

Open Source

RUGGEDCOM ROX is based on Linux®. Linux® and RUGGEDCOM ROX are made available under the terms of the [GNU General Public License Version 2.0](http://www.gnu.org/licenses/gpl-2.0.html) [http://www.gnu.org/licenses/gpl-2.0.html].

Security Information

Siemens provides products and solutions with industrial security functions that support the secure operation of plants, machines, equipment and/or networks. They are important components in a holistic industrial security concept. With this in mind, Siemens' products and solutions undergo continuous development. Siemens recommends strongly that you regularly check for product updates.

For the secure operation of Siemens products and solutions, it is necessary to take suitable preventive action (e.g. cell protection concept) and integrate each component into a holistic, state-of-the-art industrial security concept. Third-party products that may be in use should also be considered. For more information about industrial security, visit <http://www.siemens.com/industrialsecurity>.

To stay informed about product updates as they occur, sign up for a product-specific newsletter. For more information, visit <http://support.automation.siemens.com>.

Warranty

Refer to the License Agreement for the applicable warranty terms and conditions, if any.

For warranty details, visit www.siemens.com/ruggedcom or contact a Siemens customer service representative.

Contacting Siemens

Address

Siemens Canada Ltd.
Industry Sector
300 Applewood Crescent
Concord, Ontario
Canada, L4K 5C7

Telephone

Toll-free: 1 888 264 0006
Tel: +1 905 856 5288
Fax: +1 905 856 1995

E-mail

ruggedcom.info.i-ia@siemens.com

Web

www.siemens.com/ruggedcom

Table of Contents

Preface	xix
Alerts	xix
Related Documents	xix
System Requirements	xx
Accessing Documentation	xx
Training	xx
Customer Support	xx
Chapter 1	
Introduction	1
1.1 Overview	1
1.2 Security Recommendations	1
1.3 Quick Starts	3
1.3.1 Initial Configuration Before Attaching to the Network	3
1.3.2 Basic Web-Based Configuration	4
1.3.3 Physical Interface Related	4
1.3.4 Additional Configuration	5
1.4 Available Services by Port	5
Chapter 2	
Using ROX	7
2.1 The ROX Web Interface	7
2.1.1 Using a Web Browser to Access the Web Interface	7
2.1.2 The Structure of the Web Interface	8
2.2 Network Utilities	10
2.2.1 Network Utilities Main Menu	10
2.2.2 Ping Menu	11
2.2.3 Ping Check Menu	11
2.2.4 Traceroute Menu	12
2.2.5 ARP Ping Check Menu	13
2.2.6 Host Menu	14
2.2.7 Trace Menu	14
2.2.8 Tcpdump a Network Interface	15
2.2.9 Frame Relay Link Layer Trace a WAN Interface	16
2.2.10 Serial Trace a Serial Server Port	16

2.2.11 Interface Statistics Menu	17
2.2.12 Current Routing & Interface Table	17

Chapter 3

Device Management	19
3.1 Accessing the Router	19
3.1.1 Accounts and Password Management	20
3.1.2 Default Configuration	20
3.1.3 Accessing the Device Command Prompt From the Console Port	20
3.1.4 Accessing the Device Command Prompt From SSH	21
3.2 The Router Setup Shell	21
3.2.1 Configuring Passwords	21
3.2.2 Setting the Hostname and Domain	22
3.2.3 Setting the Hostname and Domain	22
3.2.4 Configuring RADIUS Authentication	22
3.2.5 Enabling and Disabling the SSH and Web Server	23
3.2.6 Enabling and Disabling the Gauntlet Security Appliance	23
3.2.7 Configuring the Date, Time and Timezone	24
3.2.8 Displaying Hardware Information	24
3.2.9 Restoring a Configuration	24
3.3 Using The LED Status Panel	25
3.4 Obtaining Chassis Information	26
3.5 Setting Up a Router Software Repository	27
3.5.1 Repository Server Requirements	27
3.5.2 Initial Repository Setup	27
3.5.3 Upgrading the Repository	28
3.5.4 Setting Up the Routers	28
3.5.4.1 An Alternate Approach	28
3.5.4.2 Upgrading Considerations	29
3.6 Reflashing the Router Software	29
3.6.1 Use Cases	29
3.6.2 Reflashing the ROX System Software	30
3.7 Maintaining the Router	31
3.7.1 Alert System	31
3.7.1.1 Alert Main Menu	32
3.7.1.2 Alert Configuration	33
3.7.1.3 Alert Filter Configuration	34
3.7.1.4 Alert Definition Configuration	35
3.7.1.5 Change Alert Definition	36
3.7.2 Backup and Restore	37
3.7.2.1 General Configuration	38

3.7.2.2	Configuration Rollback	39
3.7.2.3	Archive History	41
3.7.2.4	Archive Backup	42
3.7.2.5	Archive Restore	42
3.7.2.6	Archive Difference Tool	44
3.7.3	Decommissioning the Device	46
3.7.4	SNMP Configuration	46
3.7.4.1	SNMP Main Configuration Menu	47
3.7.4.2	System Configuration	47
3.7.4.3	Network Addressing Configuration	47
3.7.4.4	Access Control	48
3.7.4.5	Trap Configuration	50
3.7.4.6	MIB Support	51
3.7.5	Outgoing Mail	52
3.7.6	Chassis Parameters	52
3.7.7	Power over Ethernet	53
3.7.7.1	Power over Ethernet Menu	54
3.7.8	Banner Configuration	55
3.7.9	System Logs	57
3.7.9.1	Syslog Factory Defaults	58
3.7.9.2	Enabling Secure Remote Syslog	59
3.7.9.3	Remote Logging	60
3.7.10	Upgrade System	61
3.7.10.1	ROX Software Fundamentals	61
3.7.10.2	Upgrade to RX1100	62
3.7.10.3	Change Repository Server	62
3.7.10.4	Upgrading All Packages	63
3.7.10.5	Installing a New Package	64
3.7.10.6	Pre-Upgrade/Post-Upgrade Scripts	65
3.7.11	Uploading and Downloading Files	66
3.8	Configuring PPP and the Embedded Modem	66
3.8.1	PPP Interface	67
3.8.2	Authentication, Addresses and DNS Servers	67
3.8.3	When the Modem Connects	67
3.8.4	PPP Dial On Demand	68
3.8.5	LED Designations	68
3.8.6	PPP Modem Configuration	68
3.8.6.1	Modem Configuration	70
3.8.6.2	Modem PPP Client Connections	72
3.8.6.3	Dial on Demand Alternate Modem Setting	72

3.8.6.4	Modem PPP Client	73
3.8.6.5	Modem PPP Server	75
3.8.6.6	Modem Incoming Call Logs	76
3.8.6.7	Modem PPP Logs	77
3.8.6.8	Modem PPP Connection Logs	77
3.9	Configuring PPP and the Cellular Modem	78
3.9.1	PPP Interface	78
3.9.2	Authentication, IP Addressing and DNS Servers	78
3.9.3	LED Designations	79
3.9.4	PPP Cellular Modem Configuration	79
3.9.4.1	Over-The-Air Account Activation	80
3.9.4.2	Manual Account Activation	81
3.9.4.3	Cellular Modem Configuration	82
3.9.4.4	Modem PPP Client Connections	84
3.9.4.5	Modem PPP Client	85
3.9.4.6	PPP Logs, PPP Connection Logs	85
3.9.4.7	Current Route and Interfaces Table	85
3.10	Configuring Serial Protocols	86
3.10.1	Serial IP Port Features	86
3.10.1.1	LED Designations	86
3.10.2	Serial Protocols Applications	87
3.10.2.1	Character Encapsulation	87
3.10.2.2	RTU Polling	87
3.10.2.3	Broadcast RTU Polling	87
3.10.3	Serial Protocols Concepts and Issues	88
3.10.3.1	Host and Remote Roles	88
3.10.3.2	Use of Port Redirectors	88
3.10.3.3	Message Packetization	88
3.10.3.4	Use of Turnaround Delays	89
3.10.4	TcpModBus Server Application	89
3.10.4.1	Local Routing at the Server Gateway	89
3.10.4.2	MultiMaster Capability	89
3.10.5	TcpModbus Concepts and Issues	90
3.10.5.1	Host and Remote Roles	90
3.10.5.2	Port Numbers	90
3.10.5.3	Retransmissions	90
3.10.5.4	ModBus Exception Handling	90
3.10.5.5	TcpModbus Performance Determinants	91
3.10.5.6	A Worked Example	92
3.10.6	DNP (Distributed Network Protocol)	92

3.10.6.1	Address Learning for DNP	93
3.10.6.2	DNP Broadcast Messages	93
3.10.7	Serial Protocols Configuration	93
3.10.7.1	Assign Protocols Menu	94
3.10.7.2	Port Settings Menu	94
3.10.7.3	RawSocket Menu	95
3.10.7.4	TcpModBus Menu	95
3.10.7.5	DNP Menu	96
3.10.7.6	Serial Protocols Statistics Menu	98
3.10.7.7	Serial Protocols Trace Menu	99
3.10.7.8	Serial Protocols Sertrace Utility	99
3.11	Synchronous Serial Ports	100
3.11.1	Raw Socket Operation on Synchronous Ports	100
3.11.2	Synchronous Serial Port Configuration	101
3.11.2.1	Synchronous Port Settings Menu	101
3.11.2.2	Configuring Raw Socket on Synchronous Serial Ports	102
3.11.3	Synchronous Serial Diagnostics	103
3.12	Configuring SSH	103
3.12.1	Included with SSH	104
3.12.2	SSH Main Menu	104
3.12.3	Authentication	104
3.12.4	Networking	105
3.12.5	Access Control	106
3.13	Configuring the Telnet Server	106
3.13.1	Telnet Server Configuration	107
3.14	Configuring IRIG-B and IEEE1588	107
3.14.1	IEEE1588 Fundamentals	108
3.14.1.1	PTP Network Roles	108
3.14.1.2	PTP Master Election	108
3.14.1.3	Synchronizing NTP from IEEE1588	109
3.14.2	IRIG-B Fundamentals	109
3.14.2.1	IRIG-B Output Formats	109
3.14.2.2	Reference Clocks	110
3.14.2.3	How the Router Selects a Reference Clock	110
3.14.2.4	GPS Cable Compensation	110
3.14.3	IRIG-B/IEEE1588 Configuration	111
3.14.3.1	General Configuration	111
3.14.3.2	IRIG-B Configuration	112
3.14.3.3	IEEE1588 Configuration	112
3.14.3.4	IRIG-B Status	113

3.14.3.5	IEEE1588 Status	113
3.14.3.6	IRIG-B Log	114
3.15	Configuring the Intrusion Detection System	114
3.15.1	Snort Fundamentals	114
3.15.1.1	Configuring Snort	115
3.15.1.2	Which Interfaces to Monitor	116
3.15.1.3	Snort Rules	116
3.15.1.4	Alerting Methods	117
3.15.1.5	Performance and Resources	117
3.15.1.6	Troubleshooting Snort	117
3.15.2	IDS Configuration	117
3.15.2.1	Global Configuration	118
3.15.2.2	Interfaces	118
3.15.2.3	Rulesets	119
3.15.2.4	Rule Lookup by SID	120
3.15.3	Network Settings	121
3.15.4	PreProcessors	122
3.15.5	Alerts and Logging	123
3.15.6	Test Configuration	123
3.15.7	Edit Config File	124
3.16	Brute Force Attack Protection System	124

Chapter 4

System Administration	127
4.1 Webmin Configuration	127
4.1.1 IP Access Control	128
4.1.2 Ports and Addresses	129
4.1.3 Change Help Server	129
4.1.4 Logging	130
4.1.5 Authentication	131
4.1.6 Webmin Events Log	131
4.2 Configure Webmin Users	132
4.2.1 Webmin User and Group Fundamentals	132
4.2.2 RADIUS User Access Control Fundamentals	132
4.2.3 Webmin Users Menu	133
4.2.4 Edit Webmin User Menu	134
4.2.5 Current Login Sessions Menu	135
4.2.6 Password Restrictions Menu	135
4.3 Configuring the System	136
4.3.1 Bootup and Shutdown	137
4.3.2 Configuring Passwords	138

4.3.2.1	Change Password Command	138
4.3.2.2	Change Bootloader Password Command	139
4.3.3	Scheduled Commands	139
4.3.4	Scheduled Cron Jobs	140
4.3.5	System Hostname	142
4.3.6	System Time	142
4.4	Managing SSH Keys and Certificates	143
4.4.1	Uploading SSL Keys and Certificates	143
4.4.2	Regenerating SSL Keys and Certificates	144
4.4.2.1	Generating Self-Signed SSL Certificates with Scripting	145
4.4.3	Regenerating SSH Keys	146
4.5	Access Manager Secure Access Portal	146
4.5.1	What Access Manager's Secure Access Portal Protects and How	147
4.5.2	Access Manager and the Firewall	147
4.5.3	VRRP, Firewall Rules and Access Manager	150
4.5.4	Access Manager's Secure Access Portal Status Menu	150
4.5.5	Upgrading the Access Manager's Secure Access Portal	151
4.6	RADIUS Authentication	151
4.6.1	RADIUS Usage	151
4.6.2	RADIUS on ROX	152
4.6.3	RADIUS, ROX and Services	152
4.6.4	RADIUS Authentication Configuration	153
4.6.5	Edit RADIUS Server Parameters	154
4.7	RADIUS Server Configuration	154
4.7.1	Webmin Privilege Levels and FreeRADIUS	155
4.7.2	Webmin Privilege Levels and Windows IAS	156
4.7.3	PPP/CHAP and Windows IAS	159
 Chapter 5		
	Setup and Configuration	161
5.1	Configuring Networking	161
5.1.1	IPv6 Fundamentals	162
5.1.2	Network Configuration	162
5.1.2.1	Core Settings	163
5.1.2.2	Dummy Interface	164
5.1.2.3	Static Routes	164
5.1.2.4	Configuring Static Routes	165
5.1.2.5	Other Static Routes	166
5.1.2.6	Static Multicast Routing	167
5.1.2.7	DNS Client	168
5.1.2.8	Host Addresses	168

5.1.2.9	End to End Backup	169
5.1.2.10	Configuring End To End Backup	171
5.1.2.11	Current Routing and Interface Table	171
5.2	Configuring Ethernet Interfaces	171
5.2.1	LED Designations	172
5.2.2	VLAN Interface Fundamentals	172
5.2.2.1	VLAN Tag	172
5.2.2.2	ROX Functions Supporting VLANs	173
5.2.3	PPPoE On Native Ethernet Interfaces Fundamentals	173
5.2.4	IPv6 on Ethernet Fundamentals	173
5.2.5	Bridge Fundamentals	174
5.2.6	Ethernet Configuration	174
5.2.6.1	Ethernet Interfaces	175
5.2.6.2	Editing Currently Active Interfaces	176
5.2.6.3	Creating Active Virtual LAN Interfaces	176
5.2.6.4	Edit Boot Time Interfaces	177
5.2.6.5	Creating Bootup Virtual LAN Interfaces	178
5.2.6.6	Bridge Configuration	179
5.2.6.7	Bridge Filtering	180
5.2.6.8	PPPoE on Native Ethernet Interfaces	181
5.2.6.9	Edit PPPoE Interface	182
5.2.6.10	PPP Logs	182
5.2.6.11	Current Routes and Interface Table	183
5.3	Configuring Frame Relay/PPP and T1/E1	183
5.3.1	T1/E1 Fundamentals	184
5.3.1.1	Frame Relay	184
5.3.1.2	Location of Interfaces and Labelling	184
5.3.1.3	LED Designations	185
5.3.1.4	Included with T1/E1	185
5.3.2	T1/E1 Configuration	185
5.3.2.1	T1/E1 Network Interfaces	186
5.3.2.2	Strategy for Creating Interfaces	186
5.3.2.3	Naming of Logical Interfaces	187
5.3.2.4	Editing a T1/E1 Interface	188
5.3.2.5	Editing a Logical Interface (Frame Relay)	189
5.3.2.6	Frame Relay Link Parameters	189
5.3.2.7	Frame Relay DLCIs	190
5.3.2.8	Editing a Logical Interface (PPP)	190
5.3.2.9	T1/E1 Statistics	191
5.3.2.10	Link Statistics	192

5.3.2.11	Frame Relay Interface Statistics	193
5.3.2.12	PPP Interface Statistics	194
5.3.2.13	T1/E1 Loopback	194
5.3.2.14	Enabling and Disabling T1/E1 Loopback Modes	196
5.3.2.15	Upgrading Software	196
5.3.2.16	Upgrading Firmware	196
5.4	Configuring Frame Relay/PPP and T3/E3	197
5.4.1	T3/E3 Fundamentals	197
5.4.2	Location of Interfaces and Labelling	197
5.4.3	LED Designations	198
5.4.4	T3/E3 Configuration	198
5.4.4.1	T3/E3 Trunk Interfaces	198
5.4.4.2	Naming of Logical Interfaces	199
5.4.4.3	T3 Interface Parameters	199
5.4.4.4	E3 Interface Parameters	200
5.4.4.5	Editing a Logical Interface (Frame Relay)	200
5.4.4.6	Editing a Logical Interface (PPP)	201
5.4.5	T3/E3 Statistics	202
5.4.6	Current Routes and Interface Table	202
5.4.7	Upgrading Software	202
5.5	Configuring Frame Relay/PPP and DDS	202
5.5.1	Location of Interfaces and Labelling	203
5.5.2	LED Designations	203
5.5.3	DDS Configuration	203
5.5.3.1	DDS Network Interfaces	204
5.5.3.2	Editing a Logical Interface (Frame Relay)	205
5.5.3.3	Editing a Logical Interface (PPP)	206
5.5.3.4	DDS Statistics	206
5.5.3.5	DDS Loopback	207
5.5.3.6	Current Routes and Interface Table	207
5.5.3.7	Upgrading Software	208
5.6	Multilink PPP over T1/E1	208
5.6.1	Multilink PPP Fundamentals	208
5.6.2	Notes on T1/E1 Channelization	208
5.6.3	Configuring PPP Multilink over T1/E1	209
5.6.4	Multilink PPP Statistics	210
5.7	Configuring PPPoE/Bridged Mode On ADSL	210
5.7.1	PPPoE/Bridged Mode Fundamentals	211
5.7.1.1	Authentication, Addresses and DNS Servers	211
5.7.1.2	PPPoE MTU Issues	211

5.7.1.3	Bridged Mode	212
5.7.1.4	Location of Interfaces and Labelling	212
5.7.1.5	LED Designations	212
5.7.2	ADSL Configuration	213
5.7.2.1	ADSL Network Interfaces	213
5.7.2.2	Editing a Logical Interface (PPPoE)	214
5.7.2.3	Editing a Logical Interface (Bridged)	215
5.7.2.4	ADSL Statistics	216
5.7.2.5	Current Routes and Interface Table	216
5.7.2.6	Upgrading Software	216
5.8	Configuring the Firewall	216
5.8.1	Firewall Fundamentals	217
5.8.1.1	Stateless vs Stateful Firewalls	217
5.8.1.2	Linux™ netfilter, iptables and the Shoreline Firewall	217
5.8.1.3	Network Address Translation	218
5.8.1.4	Port Forwarding	218
5.8.2	Shorewall Quick Setup	219
5.8.3	ShoreWall Terminology and Concepts	219
5.8.3.1	Zones	219
5.8.3.2	Interfaces	220
5.8.3.3	Hosts	220
5.8.3.4	Policy	221
5.8.3.5	Masquerading and SNAT	221
5.8.3.6	Rules	222
5.8.4	Configuring the Firewall and VPN	224
5.8.4.1	Policy Based Virtual Private Networking	224
5.8.4.2	Virtual Private Networking to a DMZ	225
5.8.5	Firewall Configuration	225
5.8.5.1	Network Zones	227
5.8.5.2	Network Interfaces	228
5.8.5.3	Network Zone Hosts	229
5.8.5.4	Default Policies	230
5.8.5.5	Masquerading	231
5.8.5.6	Firewall Rules	232
5.8.5.7	Static NAT	234
5.8.5.8	TC (Traffic Control) Interfaces, Classes, and Rules	235
5.8.5.9	Actions When Stopped	235
5.8.5.10	Controlling the Firewall from the Command Line	235
5.9	Traffic Control	236
5.9.1	Traffic Control Example	236

5.9.2	Traffic Control Configuration	237
5.9.2.1	TC Interfaces (todevices)	238
5.9.2.2	TC Classes	239
5.9.2.3	TC Rules	241
5.10	Traffic Prioritization	244
5.10.1	Priority Queues	245
5.10.2	Filters	245
5.10.3	TOS Prioritization	245
5.10.4	Prioritization Example	246
5.10.5	Configuring Traffic Prioritization	247
5.10.6	Interface Prioritization Menu	248
5.10.6.1	Prioritization Queues	249
5.10.6.2	Prioritization Filters	249
5.10.6.3	Prioritization Transmit Queue Length	250
5.10.7	Prioritization Statistics	250
5.11	Configuring IPSec VPN	250
5.11.1	IPSec Modes	251
5.11.2	Policy-Based VPNs	251
5.11.3	Supported Encryption Protocols	252
5.11.4	Public Key and Pre-Shared Keys	252
5.11.5	X509 Certificates	252
5.11.6	NAT Traversal	253
5.11.7	Other Configuration Supporting IPSec	253
5.11.8	The Openswan Configuration Process	253
5.11.9	IPSec and Router Interfaces	253
5.11.10	L2TPD	254
5.11.11	IPSec VPN Configuration	254
5.11.12	VPN Main Menu	254
5.11.13	Server Configuration	256
5.11.14	L2TPD Configuration	257
5.11.14.1	Notes on Configuring a VPN Connection	257
5.11.15	Public Key	258
5.11.16	Pre-Shared Keys	258
5.11.17	List Certificates	259
5.11.18	VPN Connections	259
5.11.18.1	IPSec VPN Connection Details	260
5.11.18.2	Left/Right System's Settings	261
5.11.18.3	Export Configuration	262
5.11.19	Showing IPSec Status	262
5.11.20	IPSec X.509 Roaming Client Example	263

5.11.20.1	Select a Certificate Authority	264
5.11.20.2	Generate X.509 Certificates	264
5.11.20.3	VPN Networking Parameters	264
5.11.20.4	Client Configuration	265
5.11.20.5	Router IPSec Configuration	265
5.11.20.6	Firewall IPSec Configuration	266
5.11.20.7	Ethernet Port Configuration	266
5.12	Configuring Dynamic Routing	266
5.12.1	BGP Fundamentals	267
5.12.2	RIP Fundamentals	267
5.12.3	OSPF Fundamentals	268
5.12.3.1	Link State Advertisements	268
5.12.4	Key OSPF and RIP Parameters	268
5.12.4.1	Network Areas	268
5.12.4.2	Router-ID	269
5.12.4.3	Hello Interval and Dead Interval	269
5.12.4.4	Active/Passive Interface Default	269
5.12.4.5	Redistributing Routes	270
5.12.4.6	Link Detect	270
5.12.4.7	Configuring OSPF Link Costs	270
5.12.4.8	OSPF Authentication	270
5.12.4.9	RIP Authentication	270
5.12.4.10	Administrative Distances	271
5.12.5	OSPF and VRRP Example Network	271
5.12.5.1	Area and Subnets	272
5.12.5.2	VRRP Operation	272
5.12.6	Dynamic Routing Configuration	273
5.12.6.1	Enable Protocols	273
5.12.6.2	Core	274
5.12.6.3	Core Global Parameters	274
5.12.6.4	Core Interface Parameters	275
5.12.6.5	View Core Configuration	276
5.12.6.6	BGP Configuration	276
5.12.6.7	BGP Global Parameters	276
5.12.6.8	BGP Networks	278
5.12.6.9	BGP Neighbor Configuration	279
5.12.6.10	BGP Status	280
5.12.6.11	View BGP Configuration	281
5.12.6.12	OSPF	281
5.12.6.13	OSPF Global Parameters	282

5.12.6.14	OSPF Interfaces	284
5.12.6.15	OSPF Network Areas	285
5.12.6.16	OSPF Status	285
5.12.6.17	View OSPF Configuration	287
5.12.6.18	RIP	287
5.12.6.19	RIP Global Parameters	288
5.12.6.20	RIP Key Chains	289
5.12.6.21	RIP Interfaces	290
5.12.6.22	RIP Networks	291
5.12.6.23	RIP Status	291
5.12.6.24	View RIP Configuration	293
5.13	Link Backup	293
5.13.1	Path Failure Discovery	294
5.13.2	Use of Routing Protocols and the Default Route	294
5.13.3	Link Backup Configuration	295
5.13.4	Link Backup Configurations	295
5.13.5	Edit Link Backup Configuration	296
5.13.6	Link Backup Logs	297
5.13.7	Link Backup Status	297
5.13.8	Testing a Link Backup Configuration	298
5.13.9	Scheduled Link Backup Test	298
5.13.9.1	Examples	299
5.13.9.2	Logging Output	299
5.14	Configuring VRRP	300
5.14.1	The Problem with Static Routing	300
5.14.2	The VRRP Solution	300
5.14.3	VRRP Terminology	300
5.14.4	VRRP Configuration	302
5.14.5	VRRP Configuration Menu	303
5.14.6	Editing a VRRP Instance	304
5.14.7	Editing a VRRP Group	305
5.14.8	Viewing VRRP Instances Status	306
5.15	Link Layer Discovery Protocol (LLDP)	306
5.16	Configuring Generic Routing Encapsulation	307
5.16.1	GRE Configuration	308
5.16.1.1	GRE Configuration Menu	309
5.17	Configuring Layer 2 Tunnels	309
5.17.1	GOOSE Tunnel Implementation Details	310
5.17.2	Generic Layer 2 Tunnel Fundamentals	311
5.17.2.1	Generic Tunnel Implementation Details	311

5.17.3	Layer 2 Tunnels Main Menu	312
5.17.4	General Configuration Menu	312
5.17.5	GOOSE Tunnels Menu	313
5.17.6	Generic L2 Tunnels Menu	314
5.17.7	GOOSE Statistics Menu	315
5.17.8	Generic L2 Tunnel Statistics Menu	316
5.17.9	Activity Trace Menu	317
5.18	Configuring the DHCP server	317
5.18.1	DHCP Network Organizations	318
5.18.2	DHCP Client Options	318
5.18.3	Option 82 Support with Disable NAK	320
5.18.4	Example DHCP Scenarios and Configurations	320
5.18.4.1	Single Network With Dynamic IP Assignment	320
5.18.4.2	Single Network With Static IP Assignment	321
5.18.4.3	Single Network With Option82 Clients On One Switch	321
5.18.4.4	Multiple Subnets on Separate VLANs Using Option82 on One Switch	322
5.18.5	DHCP Configuration	323
5.18.5.1	DHCP Shared Network Configuration	324
5.18.5.2	DHCP Subnet Configuration	325
5.18.5.3	DHCP Group Configuration	326
5.18.5.4	DHCP Host Configuration	326
5.18.5.5	DHCP Pool Configuration	327
5.19	DHCP Relay	328
5.19.1	Configuring DHCP Relay	329
5.20	Configuring NTP Servers	330
5.20.1	The NTP Sanity Limit	330
5.20.2	NTP and the Precision Time Protocol Card	331
5.20.3	Included with NTP	331
5.20.4	NTP Configuration	331
5.20.5	Generic Options	332
5.20.6	Servers Configuration	332
5.20.7	Peers Configuration	333
5.20.8	Viewing NTP Status	333
5.20.9	Viewing The NTP Log	334
5.20.10	Viewing GPS Status	334
5.20.11	Viewing the GPS Log	335
5.21	CrossBow Station Access Controller (SAC)	335
5.21.1	Configuring CrossBow SAC	337
5.21.1.1	Configuring the SAC Connection	337
5.21.1.2	Configuring the SAC Certificates	338

5.21.1.3	Configuring SAC	338
5.21.1.4	Configuring User Access	339
5.21.1.5	Configuring Event Logs	339
5.21.2	Configuring Log Options	340
5.21.3	CrossBow Certificates	340
5.21.3.1	Generating a Certificate Signing Request (CSR)	341
5.21.3.2	Installing Certificates	342
 Chapter 6		
Upgrades		343
6.1	Installing Apache Web Server On Windows	343
6.2	Installing a Microsoft IIS Web Server	344
6.2.1	Using Microsoft Internet Information Services (IIS) Manager 6.0 or Higher as an Upgrade Repository	345
6.3	VPN/L2TP Configuration in Windows	346

Preface

This guide describes the user interface for ROX v1.16 running on the RX1000/RX1100 family of products. It contains instructions and guidelines on how to use the software, as well as some general theory.

It is intended for use by network technical support personnel who are familiar with the operation of networks. It is also recommended for use by network and system planners, system programmers, and line technicians.

Alerts

The following types of alerts are used when necessary to highlight important information.



DANGER!

DANGER alerts describe imminently hazardous situations that, if not avoided, will result in death or serious injury.



WARNING!

WARNING alerts describe hazardous situations that, if not avoided, may result in serious injury and/or equipment damage.



CAUTION!

CAUTION alerts describe hazardous situations that, if not avoided, may result in equipment damage.



IMPORTANT!

IMPORTANT alerts provide important information that should be known before performing a procedure or step, or using a feature.



NOTE

NOTE alerts provide additional information, such as facts, tips and details.

Related Documents

Other documents that may be of interest include:

- RX1000/RX1100 Installation Guide
- RUGGEDCOM Fiber Guide
- RUGGEDCOM Wireless Guide
- Industrial Defender Access Manager User Manual
- Industrial Defender Access Client User Manual
- Industrial Defender Access Manager System Installation Manual
- White Paper: Rapid Spanning Tree in Industrial Networks

System Requirements

Each workstation used to connect to the ROX Webmin user interface must meet the following system requirements:

- Must have one of the following Web browsers installed:
 - Microsoft Internet Explorer 8.0 or higher
 - Mozilla Firefox
 - Google Chrome
 - Iceweasel/IceCat (Linux Only)
- Must have a working Ethernet interface compatible with at least one of the port types on the unit
- The ability to configure an IP address and netmask on the computer's Ethernet interface
- A suitable Ethernet cable
- An SSH client application installed on a computer

Accessing Documentation

The latest Hardware Installation Guides and Software User Guides for most RUGGEDCOM products are available online at www.siemens.com/ruggedcom.

For any questions about the documentation or for assistance finding a specific document, contact a Siemens sales representative.

Training

Siemens offers a wide range of educational services ranging from in-house training of standard courses on networking, Ethernet switches and routers, to on-site customized courses tailored to the customer's needs, experience and application.

Siemens' Educational Services team thrives on providing our customers with the essential practical skills to make sure users have the right knowledge and expertise to understand the various technologies associated with critical communications network infrastructure technologies.

Siemens' unique mix of IT/Telecommunications expertise combined with domain knowledge in the utility, transportation and industrial markets, allows Siemens to provide training specific to the customer's application.

For more information about training services and course availability, visit www.siemens.com/ruggedcom or contact a Siemens sales representative.

Customer Support

Customer support is available 24 hours, 7 days a week for all Siemens customers. For technical support or general information, please contact Siemens Customer Support through any of the following methods:

- Online

Visit <http://www.siemens.com/automation/support-request> to submit a Support Request (SR) or check on the status of an existing SR.

- Telephone

Call a local hotline center to submit a Support Request (SR). To locate a local hotline center, visit <http://www.automation.siemens.com/mcms/aspa-db/en/automation-technology/Pages/default.aspx>.

- Mobile App

Install the Industry Online Support app by Siemens AG on any Android, Apple iOS or Windows mobile device and be able to:

- Access Siemens's extensive library of support documentation, including FAQs, manuals, and much more
- Submit SRs or check on the status of an existing SR
- Find and contact a local contact person
- Ask questions or share knowledge with fellow Siemens customers and the support community via the forum
- And much more...



Introduction

This chapter provides a basic overview of the ROX software. It describes the following topics:

- [Section 1.1, “Overview”](#)
- [Section 1.2, “Security Recommendations”](#)
- [Section 1.3, “Quick Starts”](#)
- [Section 1.4, “Available Services by Port”](#)

Section 1.1

Overview

Welcome to the ROX Software User Guide for the RX1000/RX1100. This Guide describes the wide array of carrier grade features made available by ROX (Rugged Operating System on Linux) software. These features include:

Routing Features

- VRRP, OSPF, BGP, RIP
- DHCP Agent (Option 82 Capable)
- Traffic prioritization, NTP Server
- IP Multicast Routing

WAN Features

- Frame Relay RFC 1490 or RFC 1294
- PPP RFC 1661, 1332, 1321, 1334
- PAP, CHAP Authentication
- PPPoE over DSL
- GOOSE messaging support

Router Software Features

- Cellular Modem Support
- Link Backup
- Precision Time Protocol (PTP) Card
- Serial IP Encapsulation
- Port Configuration and Status
- Event Logging and Alarms

Security Appliance Functions

- Integrated Router/Firewall/VPN
- Full IPSec Virtual Private Networking
- VPN with 3DES, AES128, AES256 support
- RADIUS centralized password management
- Multi-level user passwords
- SNMP v2/v3 (Simple Network Management Protocol)
- Stateful Firewall with NAT

Management Features

- Web-based, SSH, CLI management interfaces
- SNMP v2/v3
- Remote Syslog
- Rich set of diagnostics with logging and alarms
- Loopback diagnostic tests
- Raw and interpreted real time line traces
- Quick setup facility

Section 1.2

Security Recommendations

To prevent unauthorized access to the device, note the following security recommendations:

**CAUTION!**

Configuration hazard – risk of data corruption. Access to the Command Line Interface (CLI) is provided for troubleshooting purposes and should only be used by Siemens technicians. As such, this interface is not fully documented. Misuse of CLI commands can corrupt the operational state of the device and render it inaccessible.

- Replace the default passwords for all user accounts and processes (where applicable) before the device is deployed.
- Use strong passwords. Avoid weak passwords such as *password1*, *123456789*, *abcdefgh*, etc.
- Make sure passwords are protected and not shared with unauthorized personnel.
- If RADIUS authentication is being employed, configure authentication servers.
- The default SSL certificate that is provided in Webmin is a self-signed certificate, which means that the trust part of the SSL connection cannot work. Siemens highly recommends following the instructions in [Section 4.4, “Managing SSH Keys and Certificates”](#) and providing ROX with a proper X.509 certificate.
- SSL and SSH keys are accessible to users who connect to the device via the serial console. Make sure to take appropriate precautions when shipping the device beyond the boundaries of the trusted environment:
 - Replace the SSH and SSL keys with throwaway keys prior to shipping.
 - Take the existing SSH and SSL keys out of service. When the device returns, create and program new keys for the device. See [Section 4.4.3, “Regenerating SSH Keys”](#) for further information.
- Restrict physical access to the flash card to only trusted personnel. A person with malicious intent in possession of the flash card could extract critical information, such as certificates, keys, etc. (user passwords are protected by hash codes), or reprogram the card.
- Control access to the serial console to the same degree as any physical access to the device by setting a bootloader password. If a bootloader password is not set, the serial console could provide unauthorized access to BIST mode, which includes tools that may be used to gain complete access to the device.
- Restrict the IP addresses which Web management will accept connections from. See the *Webmin* menu, *IP Access Control* sub-menu. Restrict the Ethernet ports which Web management will accept connections from. See the *Webmin* menu, *Ports and Addresses* sub-menu.
- Review the IP networking settings provided in the *Network Configuration* menu, *Core Settings* sub-menu. You may wish to tighten some settings, especially *Ignore All ICMP ECHO* requests.
- Restrict the users that the SSH server will allow to connect. See the *SSH Server* menu, *Access Control* sub-menu.
- It is highly recommended that critical applications be limited to private networks, or at least be accessible only through secure services, such as IPSec. Connecting a ROX device to the Internet is possible. However, the utmost care should be taken to protect the device and the network behind it using secure means such as firewall and IPSec. For more information about configuring firewalls and IPSec, refer to [Section 5.8, “Configuring the Firewall”](#) and [Section 5.11, “Configuring IPSec VPN”](#).
- If the router is an RX1100 and you wish to use the Snort Intrusion Detection System, activate and configure it.
- If the router is an RX1100 and you wish to use the Gauntlet security appliance, activate and configure it.
- If SNMP will be used, limit the IP addresses which can connect and change the community names. Configure SNMP to raise a trap upon authentication failures.
- Only enable the services you need and expect to use.
- ROX comes with the following login banner. Replace the contents of the file */etc/issue* and */etc/issue.net* in order to change it.

WARNING: You are attempting to access a private computer system. Access to this system is restricted to authorized persons only. This system may not be used for any purpose that is unlawful or deemed inappropriate. Access and use of this system is electronically monitored and, by entering this system, you are giving your consent to be electronically monitored. We reserve the right to seek all remedies for unauthorized use, including prosecution.

- If using a firewall, configure and start the firewall before attaching the router to the public network. Configure the firewall to accept connections from a specific domain.
- Configure remote system logging to forward all logs to a central location.
- Periodically audit the device to make sure it complies with these recommendations and/or any internal security policies.

Section 1.3

Quick Starts

The following instructions are included to aid those users experienced with communications equipment that may wish to attempt to configure the router without fully reading the guide.

Section 1.3.1

Initial Configuration Before Attaching to the Network

1. Locate/mount the chassis in its final resting place and apply power.
2. The router can be configured through its web management interface, or for advanced users, through ssh. The default Ethernet addresses for ports one through four are 192.168.1.1 through 192.168.4.1. Two shell accounts, rrsetup and root, are provided. Both accounts have a default password of "admin". The web management interface uses the root account password. The rrsetup account provides a shell that configures such items as passwords, addresses, date/time and services offered by the router. The root account provides a full shell.
3. Attach a PC running terminal emulation software to the RS232 port and apply power to the chassis (default baud rate, data bits, parity - "38400 8 n 1", no hardware/software flow control). Set the terminal type to VT100. Press ENTER to obtain a login prompt.
4. Login as the rrsetup user with password "admin".
5. *Change the root and rrsetup passwords from the shell. Record the passwords in a secure manner. If RADIUS authentication will be employed, configure at least one authentication server address.*
6. Configure the router's hostname, IP address, subnet mask, and gateway addresses for the built-in Ethernet ports.
7. For an RX1100 router, the Gauntlet Security application may be configured with the passphrase allocated to the network the network address of the Command and Control Center (CCC). Note that you must also configure and activate the firewall before using the Gauntlet.
8. Ensure that the date, time and timezone fields are correctly set.
9. If Web or SSH services will not be used, these can be disabled from the setup shell.
10. All further configuration is accomplished through the web management interface. Attach the configuring host to one of the Ethernet ports configured above. Point your web browser at the address for that port, use https and specify a port number of 10000, e.g. <https://192.168.1.1:10000> (or otherwise if configured in step 4).

Login with the root user and password (configured above). If RADIUS authentication is configured and a server is available, you may also login via a RADIUS user.

Section 1.3.2

Basic Web-Based Configuration

1. Change the router password from the *System* menu, *Change Password* sub-menu.
2. If you are using the web management interface you may wish to restrict the allowed users to a specific subnet. This can be done in the *Webmin* menu, *Webmin Configuration*, *IP Access Control* sub-menu.
3. If you are planning to SSH in to the router you may wish to restrict the allowed users to a specific subnet. This can be done in the *Servers* menu, *SSH Server*, *Networking* sub-menu.
4. The router's local hostname may configured in the *System Menu*, *System Hostname* sub-menu.
5. The router may be configured to log to a remote server by the *Maintenance menu*, *System Logs* sub-menu. See [Section 3.7, "Maintaining the Router"](#) for more details.
6. The router's DNS settings may configured in the *DNS Clients* sub-menu. You may also specify the IP addresses of frequently used hosts. See [Section 5.1, "Configuring Networking"](#) for more details.

Section 1.3.3

Physical Interface Related

1. Ethernet port parameters may be changed in the *Networking* menu, *Ethernet* sub-menu. The *Ethernet Interfaces* sub-menu will configure the IP address, subnet mask, gateway address, proxy arping and media type of each interface. See [Section 5.2, "Configuring Ethernet Interfaces"](#) for more details.
2. If your router is equipped with T1/E1 WAN interfaces, the *Networking* menu, *T1/E1* sub-menu will allow you to configure them with Frame Relay or PPP connections. See [Section 5.3, "Configuring Frame Relay/PPP and T1/E1"](#) for more details.
3. If your router is equipped with T3 WAN interfaces, the *Networking* menu, *T3* sub-menu will allow you to configure them with Frame Relay or PPP connections. See [Section 5.4, "Configuring Frame Relay/PPP and T3/E3"](#) for more details.
4. If your router is equipped with DDS interfaces, the *Networking* menu, *DDS* sub-menu will allow you to configure them with Frame Relay or PPP connections. See [Section 5.5, "Configuring Frame Relay/PPP and DDS"](#) for more details.
5. If your router is equipped with ADSL interfaces, the *Networking* menu, *ADSL* sub-menu will allow you to configure them. See [Section 5.7, "Configuring PPPoE/Bridged Mode On ADSL"](#) for more details. If you wish to use PPPOE with an external ADSL modem, the *Networking* menu, *Ethernet* sub-menu will configure it.
6. If your router is equipped with an embedded modem, the *Networking* menu, *Modem* sub-menu will allow you to configure it with PPP or incoming console connections. See [Section 3.8, "Configuring PPP and the Embedded Modem"](#) for more details.
7. If your router is equipped with Serial Interfaces, the *Servers* menu, *Serial Protocols* sub-menu will allow you to configure them with an operating protocol. See [Section 3.10, "Configuring Serial Protocols"](#) for more details.
8. If your router is equipped with a Precision Time Protocol Card, the *Servers* menu, *IRIG-B* sub-menu will allow you to enable and configure its output ports. See [Section 3.14, "Configuring IRIG-B and IEEE1588"](#) for more details.

Section 1.3.4

Additional Configuration

- You may wish to configure a backup interface to use in the event of a failure of your default gateway interface. This can be done in the *Networking* menu, *Network Configuration*, *End To End Backup* sub-menu.
- If you are planning to connect your router to the Internet, configure the firewall and then activate it. This can be done in the *Networking* menu, *Shorewall Firewall* sub-menu.
- The router provides a default event logging configuration. You can modify this configuration through the *Maintenance* menu, *System Logs* sub-menu. Remote logging can be activated here.
- The routers SSH and Web Management interfaces are enabled by default. The routers DHCP server, IPSec VPN server, NTP server, OSPF/RIP protocol, VRRP protocol and firewall are disabled by default. To changes these services visit the *System* menu, *Bootup and Shutdown* sub-menu.
- You can install static IP and Multicast routings for Ethernet and WAN interfaces via the *Networking* menu, *Network Configuration*, *Routing and Default Route* and *Static Multicast Routing* sub-menus.
- You can configure the NTP server through the *Servers* menu, *NTP Server* sub-menu. See [Section 5.20, "Configuring NTP Servers"](#) for more details.
- You can configure SSH through the *Servers* menu, *SSH Server* sub-menu. SSH can be set-up to issue a login banner from this menu. See [Section 3.12, "Configuring SSH"](#) for more details.
- Traffic prioritization can be configured on the network interfaces through the *Networking* menu, *Traffic Prioritization* sub-menu.. See [Section 5.10, "Traffic Prioritization"](#) for more details.
- SNMP is disabled by default. You can configure SNMP by following the instructions in [Section 3.7.4, "SNMP Configuration"](#). You may allow read and write access, set community names, enable traps and program the router to issue traps with a specific client address.
- If your router is an RX1100, you may configure and activate the Snort Intrusion Detection System and the Gauntlet Security Appliance. If you decide to forward daily email summaries, you must configure a mail forwarder in the *Maintenance* menu *Miscellaneous* sub-menu *Outgoing Mail* sub-menu.
- When your routers configuration is stable, it is recommended that the configuration should be uploaded from the router and stored as a backup. The *Maintenance* menu *Backup and Restore* sub-menu will be useful.
- Should you need to transfer files to or from the router, the *Maintenance* menu *Upload/Download Files* sub-menu will be useful.
- Further concerns such as ensuring robustness, measuring and optimizing performance are dealt with by reading the guide fully.

Section 1.4

Available Services by Port

The following table lists the services available under ROX. This table includes the following information:

- Services
The service supported by the device.
- Port Number
The port number associated with the service.
- Port Open
The port state, whether it is always open and cannot be closed, or open only, but can be configured.



NOTE

In certain cases, the service might be disabled, but the port can still be open (e.g. TFTP).

- **Port Default**
The default state of the port (i.e. open or closed).
- **Access Authorized**
Denotes whether the ports/services are authenticated during access.

Services	Port Number	Port Open	Port Default	Access Authorized
HTTPS	TCP/443	Open (if configured with login)	Open	Yes
Telnet	TCP/23	Open (if configured with login)	Closed	Yes
NTP	UDP/123	Open (if configured)	Closed	No
SNMP	UDP/161	Open (if configured with login)	Closed	Yes
SSH	TCP/22	Open (if configured with login)	Open	Yes
TCP Modbus	TCP/502	Open (if configured)	Closed	No
TCP Modbus Gateway	TCP/502	Open (if configured)	Closed	No
IPSec IKE	UDP/500	Open (if configured)	Closed	No
IPSec NAT-T	UDP/4500	Open (if configured)	Closed	No
DNPv3	TCP/20000	Open (if configured)	Closed	No
RawSocket	TCP/configured	Open (if configured)	Closed	No
DHCP Server	UDP/67, 68	Open (if configured)	Closed	No
DHCP Server	UDP/67	Open (if configured)	Closed	No
DHCP Agent	UDP/67	Open (if configured)	Closed	No
RADIUS	UDP/1812 to send, opens random port to listen	Open (if configured)	Closed	Yes
L2TP	Random Port	Open (if configured)	Closed	Yes
PTP	IEEE 1588 UDP 319/320	Open (if configured)	Closed	No
LLDP	Open (if configured)	Open	Open	No

2 Using ROX

This chapter describes how to use the ROX interface. It describes the following tasks:

- [Section 2.1, “The ROX Web Interface”](#)
- [Section 2.2, “Network Utilities”](#)

Section 2.1

The ROX Web Interface

The ROX Web interface is provided by an enhanced version of the popular Webmin interface.

Section 2.1.1

Using a Web Browser to Access the Web Interface

Start a web browser session and open a connection to the router by entering a URL that specifies its hostname or IP address (e.g. <https://179.1.0.45:10000>). Once the router is contacted, start the login process by clicking on the *Login* link. The resulting page should be similar to that presented below.



IMPORTANT!

*Starting with ROX v1.16.1, ROX uses 1024 bit RSA certificates by default, which is required by most Web browsers. However, previous versions of ROX use 512 bit certificates and these certificates are not replaced during a simple upgrade. If upgrading from a version of ROX older than v1.16.1, replace the current certificates with new 1024 bit certificates for compatibility with most modern Web browsers. For more information, refer to the application note *Generating SSH Keys and SSL Certificates for ROS and ROX Using Windows (AN22)* available on www.siemens.com/ruggedcom.*

Enter the "root" user name and the appropriate password for that user, then click on the *Login* button. The router is shipped with a default administrator password of "admin". Once successfully logged in, the user will be presented with the main menu.

Login

You must enter a username and password to login on 192.168.0.3.

Username

Password

Login Clear

Figure 1: Signing On to the Router With a Web Browser

Section 2.1.2

The Structure of the Web Interface

The Web interface presents an web page with two frames. The leftmost or index frame selects subsystems to configure and is always displayed.

The rightmost or configuration frame presents the configuration for the currently selected subsystem, or in the case of signing-on, the home page window. The home page window presents an annotated view of the front of the chassis as well as a number of important system parameters. These parameters include:

- The router uptime and load averages for the past 1, 5 and 15 minutes. Under normal operation the load average should be less than 2.0.
- The disk usage. A disk usage higher than 92% requires attention.
- The memory usage, indicating the amount of memory used by applications. Under normal operation memory usage should be less than 60%.
- The chassis temperature.
- Any major alarms, such as the failure of hardware components.

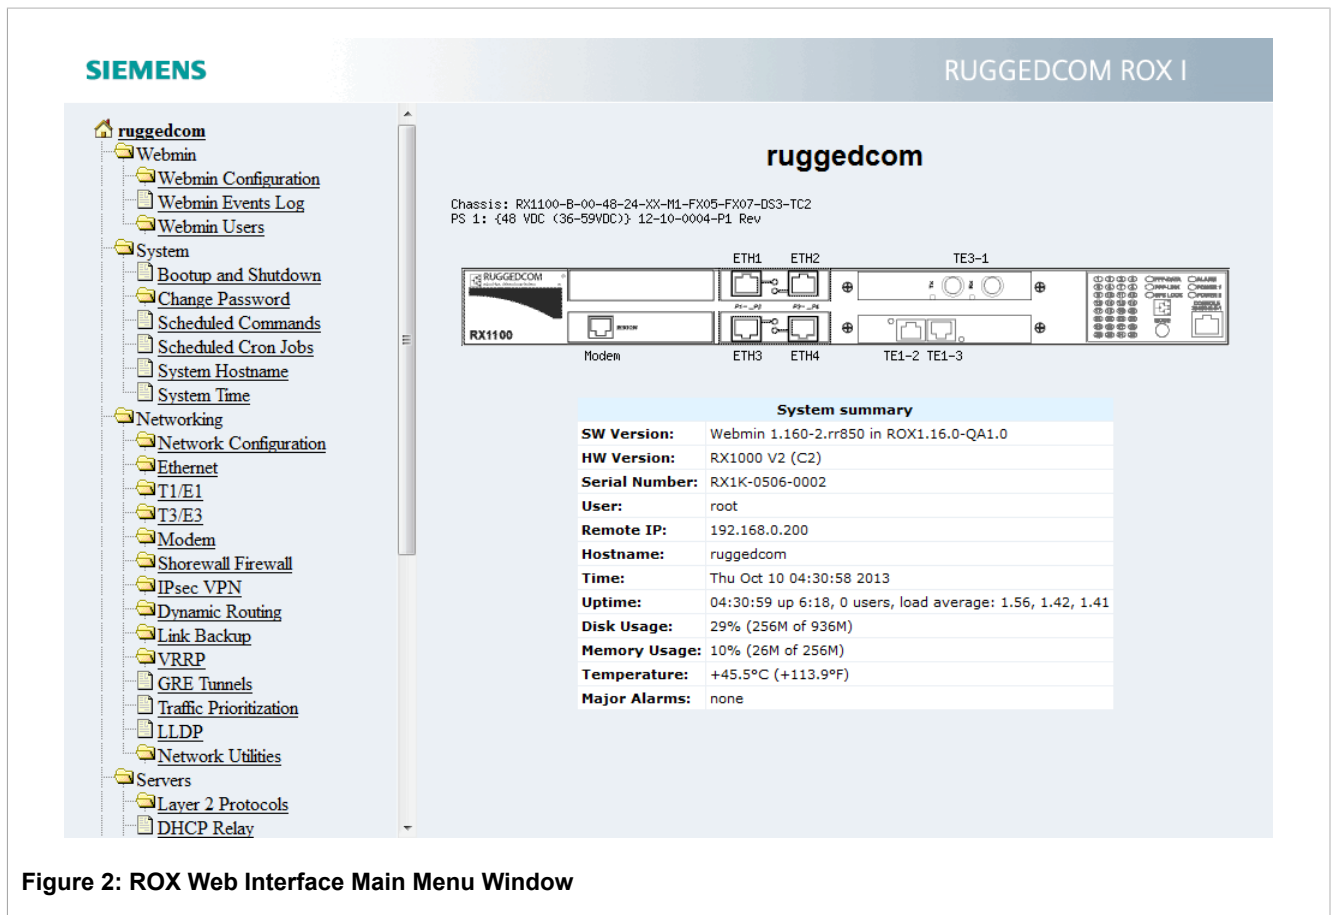






Figure 2: ROX Web Interface Main Menu Window

The index frame presents a number of entries with associated icons:

- The  icon causes home page window to be redisplayed.
- The  icon signifies that the next level contains a menu of menus.

- The  icon signifies that clicking the entry will run a single menu.
- The  icon logs out of Webmin.

The menu system entries are composed of the Webmin, System, Servers, Networking and Maintenance menus.

The Webmin Menu provides the ability to:

- Configure the sign-on password
- Specify session timeouts
- Restrict the Subnet of IP addresses that can login
- Configure and view Webmin event logs

The System Menu provides the ability to:

- Change the router password,
- Enable and disable applications from running,
- Reboot the router,
- Schedule one time and periodic tasks to run,
- Change the router's name (hostname),
- Change the time and date.

The Servers Menu provides the ability to:

- Control and configure the Serial Protocol, DHCP, NTP, IRIG-B and SSH servers.

The Networking Menu provides the ability to:

- Configure the network interfaces,
- Configure static IP and Multicast Routings and configure a default gateway,
- Select a DNS server and edit local host addresses,
- Configure End To End Backup,
- Configure DDS, T1/E1, T3 and ADSL Networking,
- Configure the embedded modem,
- Set up the firewall,
- Set up Virtual Private Networking,
- Configure Routing protocols such as OSPF and RIP,
- Configure Virtual Router Redundancy Protocol (VRRP),
- Configure Traffic Prioritization,
- Perform pings, traceroutes, host lookups and line tracing.

The Maintenance Menu provides the ability to:

- Manage the Gauntlet Security Appliance
- Backup and restore configurations,
- Configure SNMP access,
- Configure RADIUS Authentication,
- View system logs,
- Upgrade the software of the router,

- Upgrade the router type to RX1100,
- Upload/Download files to and from the router.

Section 2.2

Network Utilities

This section familiarizes the user with:

- Pinging hosts
- Running a traceroute
- Performing a host lookup
- Tracing line activity
- Showing interface statistics

Section 2.2.1

Network Utilities Main Menu

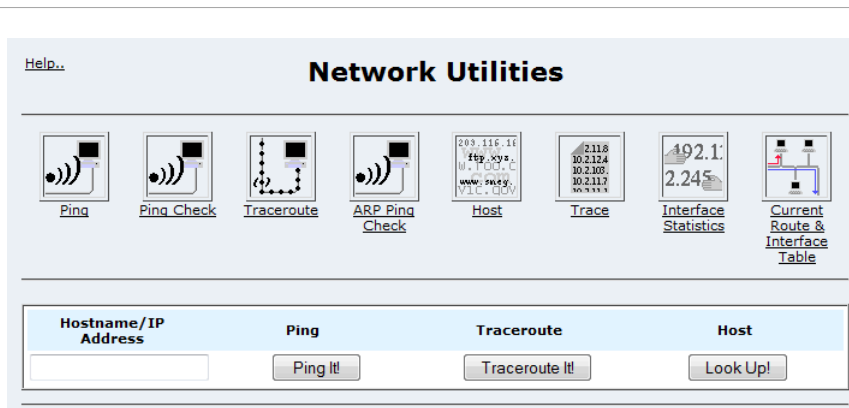


Figure 3: Network Utilities Main Menu

The lower part of the menu provides quick pinging, tracerouting and lookup of hosts.

The upper part leads to menus providing more configurable options for these commands. Additionally, Ethernet, WAN and Serial port tracing is provided. A summary of interface statistics and the current routing table is provided.

Section 2.2.2

Ping Menu

[Help..](#)

Ping

Hostname/IP Address:

☐ Verbose Output?

☐ Lookup Addresses?

How many Packets?

Packet Size?

Time between pings Seconds

Pattern(s) to send (Hex)?


 [Return to Network Utilities](#)

Figure 4: Ping Menu

- The *Hostname* field accepts the host name or IP address to ping. Note that this may be an IPv4 or and IPv6 address.
- The *Verbose Output?* field causes ping to present the maximum of output.
- The *Lookup Addresses?* field causes ping to resolve IP addresses to domain names. This can make ping behave very slowly if DNS is not properly configured.
- The *Packet Size?* field specifies the size of the data in the ping packet. The true length of the packet is 28 bytes larger due to IP/ICMP overhead.
- The *Time between pings* field limits the rate at which pings are sent.
- The *Pattern(s) to send (Hex)?* field specifies a pattern to fill the packet sent. This is useful for diagnosing data-dependent problems in a network. For example, specifying "ff" will cause the sent packet to be filled with all ones.

Section 2.2.3

Ping Check Menu

- The *Ping Check* utility is configured to ping an IP host. If the configured host fails to respond to ping checks, the utility performs a configured action in response. For example, if it is detected that a host across a PPP connection is no longer reachable, one might wish to explicitly reset the PPP connection.
- The main *Ping Check* menu displays a list of currently configured and active ping check entries:

[Help..](#)

Ping Check

Index	Interval (sec)	Ping tries	Remote IP address	Action
<input type="button" value="Add a new Ping check.."/>				


 [Return to Network Utilities](#)

Figure 5: Ping Check Menu

The main *Ping Check* menu also contains links to *Add a new Ping check...* or to *Edit* an existing entry:

Parameter	Value	Description
Interval		Waiting time between ping packet in second
Ping retries		Max ping packets to send before taking action
Remote IP address		Host to send ping packets to
Pre-defined action	None	Pre-defined action if ping check failed
Other Action		Action to be done if ping check failed

Save

Return to Ping check

Figure 6: Ping Check Edit Menu

The *Interval* field specifies the time between each successive ping request to the IP host.

The *Ping retries* field specifies the number of ping requests that are allowed to go unanswered before taking the configured action.

The *Remote IP address* field specifies the IP address of remote host to monitor.

The *Pre-defined action* drop list may be used to select a pre-defined action to be taken if the monitored IP host does not reply within the configured number of ping retries. Currently, the list contains "Restart PPP on Cellmodem" (if a cellular modem is installed) and "Restart MLPPP".

The *Other action* field accepts any valid shell command as the action to be executed if the monitored IP host does not reply within the configured number of ping retries.

Section 2.2.4

Traceroute Menu

Hostname/IP Address:

☐ Verbose Output? ☐ Lookup Addresses? ☐ Use ICMP instead of UDP?

How many Hops? 30
Packet Length? 40
Interface:

Trace It!

Return to Network Utilities

Figure 7: Traceroute Menu

The *Hostname* field accepts the host name or IP address to trace the route to. Note that this may be an IPv4 or and IPv6 address.

The *Verbose Output?* field causes ping to present the maximum of output.

The *Lookup Addresses?* field causes ping to resolve IP addresses to domain names. This can make ping behave very slowly if DNS is not properly configured.

The *Use ICMP instead of UDP?* field causes traceroute to probe with ICMP packets.

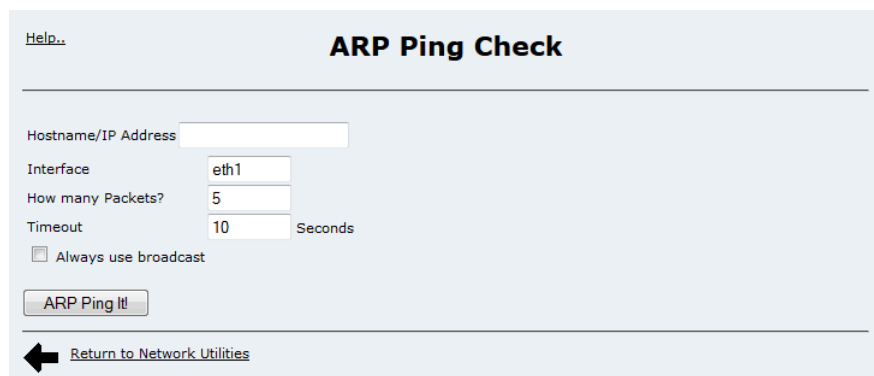
The *How many Hops?* field limits the maximum number of hops that traceroute will attempt to map.

The *Packet Length?* field specifies the size of the data in the traceroute packet.

The *Interface?* field specifies the network interface to obtain the source IP address for outgoing probe packets. Otherwise the router will manually set the address based on the actual interface taken.

Section 2.2.5

ARP Ping Check Menu



The screenshot shows the 'ARP Ping Check' menu. At the top left is a 'Help..' link. The title 'ARP Ping Check' is centered. Below the title are several input fields: 'Hostname/IP Address' (empty), 'Interface' (set to 'eth1'), 'How many Packets?' (set to '5'), and 'Timeout' (set to '10' with 'Seconds' to its right). There is a checkbox labeled 'Always use broadcast' which is currently unchecked. Below these fields is a button labeled 'ARP Ping It!'. At the bottom left is a back arrow icon, and at the bottom right is a link labeled 'Return to Network Utilities'.

Figure 8: ARP Ping Check Menu

The *Hostname/IP address* field accepts the host name or IP address to ping. Note that this may be an IPv4 or and IPv6 address.

The *Interface* field specifies the interface on the router through which the ARP ping packet will be sent.

The *How many Packets* field specifies the number of packets to be sent. The maximum is 99 packets.

The *Timeout* field specifies the number of seconds to allow for an ARP ping request. The maximum is 99 seconds.

The *Always use broadcast* field specifies whether broadcast should be used for all packets/frames. If not selected, ROX will only use broadcast for sending the first ARP packet/frame, after which, only unicast will be used.

Section 2.2.6

Host Menu

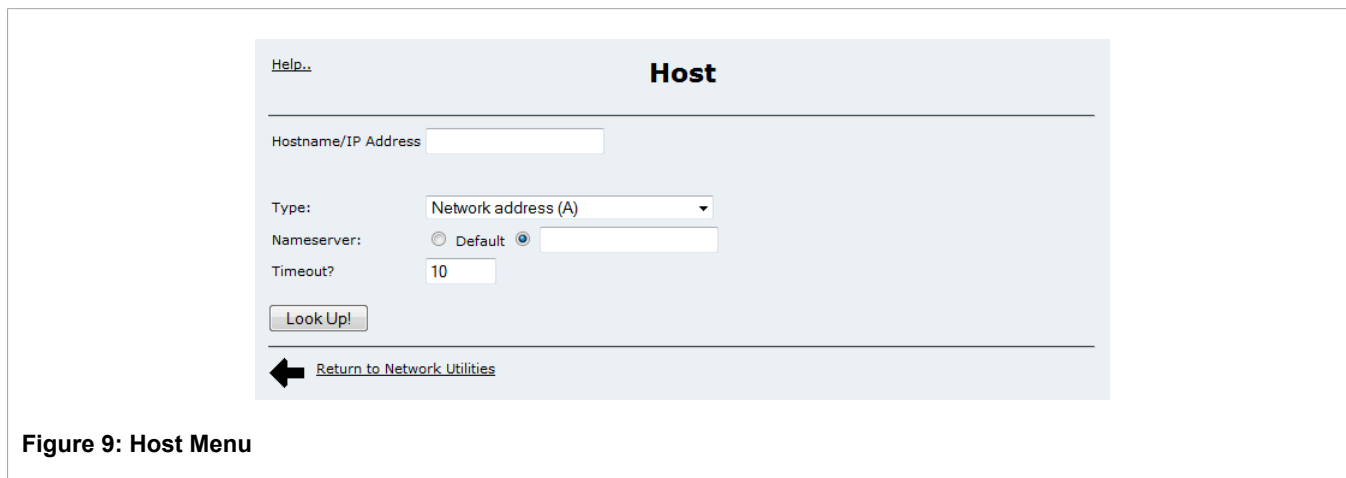


Figure 9: Host Menu

The *Hostname* field accepts the host name or IP address to ping.

The *Type* field selects the type of information to capture.

The *Nameserver* fields select the server to use to resolve with. If *Default* is left selected the DHCP, DNS or local resolv.conf setup will be used. Otherwise the address supplied will be used.

The *Timeout* field specifies the maximum time to wait before abandoning a lookup.

Section 2.2.7

Trace Menu

The *Trace Menu* contains three sections, providing the capability to trace network interfaces, frame relay interfaces, and serial server interfaces. The latter two menus will appear only if you have configured frame relay or serial server interfaces.

Section 2.2.8

Tcpdump a Network Interface

Tcpdump A Network Interface

Interface to capture on:

Maximum packets captured: (maximum 1000)

Maximum capture time: (maximum 240 sec.)

☐ Lookup addresses

☐ Display link level header

☐ Perform HEX/ASCII dump

Verbosity: ☒ Off ☐ 1 ☐ 2 ☐ 3

Ignore hostname: for protocols: ☐ SSH ☐ Webmin traffic ☐ All traffic

Ignore protocols: ☐ TCP ☐ UDP ☐ ICMP ☐ ARP ☐ VRRP ☐ IGMP ☐ OSPF ☐ ESP ☐ AH

Ports to trace:

Figure 10: Tcpdump Menu

The *Interface to capture on* field specifies the interface to show traffic on.

The *Maximum packets captured* and *Maximum capture time* fields limit the amount of traffic captured.

The *Lookup addresses* field causes IP addresses to be resolved into domain names. This can make tcpdump behave very slowly if DNS is not properly configured.

The *Display link level header* field causes this header to be displayed.

The *Perform HEX/ASCII dump* field causes the data content of the captured packets to be displayed. This option may generate a large capture data set.

The *Verbosity* fields specify the level of decoding which tcpdump supplies.

The *Ignore hostname/Only hostname* selector causes traffic to or from the specified address to be excluded or included to the exclusion of other traffic, respectively. One of the following three check boxes must be selected in order for filtering to take place. If the *SSH* box is selected, SSH traffic to or from the selected IP address will be excluded/displayed. If the *Webmin traffic* box is selected, Webmin traffic to or from the selected IP address will be excluded/displayed. If the *All traffic* box is selected, traffic to or from the selected IP address will be excluded/displayed. A typical usage of this filter is to dump all traffic on an interface and to prevent the user's own traffic from being displayed.

The *Ignore protocols/Only protocols* selector excludes or displays traffic in the protocols specified in the accompanying check boxes.

The *Ports to trace* field specifies TCP/UDP ports to trace. Enter a list of port numbers separated by spaces to trace more than a single port.

Section 2.2.9

Frame Relay Link Layer Trace a WAN Interface

Frame Relay Link Layer Trace A WAN Interface

Interface to capture on

Maximum packets captured (maximum 1000)

Maximum capture time (maximum 240 sec.)

Figure 11: Frame Relay Trace Menu

Frame Relay tracing uses the wanpipemon utility.

The *Interface to capture on* field specifies the interface to show traffic on.

The *Maximum packets captured* and *Maximum capture time* fields limits the amount of traffic captured.

Section 2.2.10

Serial Trace a Serial Server Port

Serial Trace A Serial Server Port

Trace on ports ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ All Ports ☐

Message RX/TX ☐ Incoming/Outgoing Connections ☐ Hex dump ☐

Maximum packets captured (maximum 1000)

Maximum capture time (maximum 240 sec.)

Figure 12: Serial Server Port Trace Menu

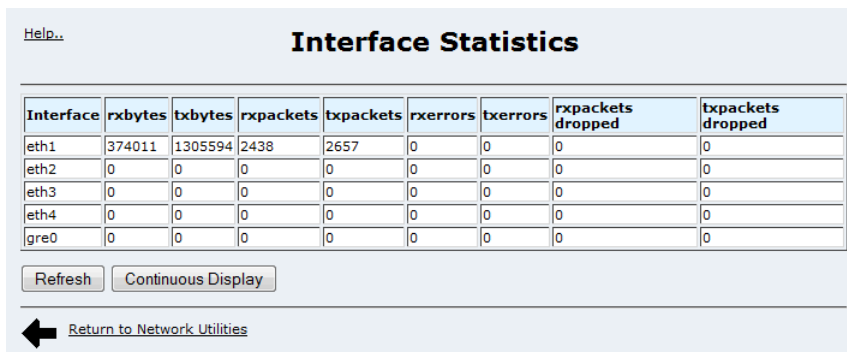
The *Trace on ports* fields specify the serial port to show traffic on.

The *Message RX/TX* and *Incoming/Outgoing Connections* fields causes data packets and *Connection activity* to be included in the trace. The *Hex dump* field causes the content of data packets to be displayed.

The *Maximum packets captured* and *Maximum capture time* fields limits the amount of traffic captured.

Section 2.2.11

Interface Statistics Menu



[Help..](#)

Interface Statistics

Interface	rxbytes	txbytes	rxpackets	txpackets	rxerrors	txerrors	rxpackets dropped	txpackets dropped
eth1	374011	1305594	2438	2657	0	0	0	0
eth2	0	0	0	0	0	0	0	0
eth3	0	0	0	0	0	0	0	0
eth4	0	0	0	0	0	0	0	0
gre0	0	0	0	0	0	0	0	0

[Refresh](#) [Continuous Display](#)

[Return to Network Utilities](#)

Figure 13: Interface Statistics Menu

This menu provides basic statistics for all network interfaces.

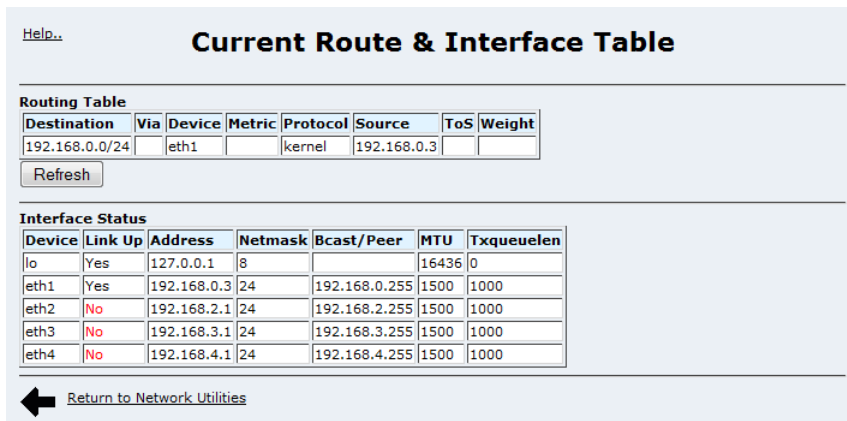
The *Refresh* button will cause the page to be reloaded.

The *Continuous Display* button will cause the browser to continuously reload the page showing the differences in statistics from the last display. *The difference is not a real time rate in bytes or packets per second.*

Note that detailed statistics for T3, T1/E1, DDS and ADSL are available within the menus that configure those interfaces.

Section 2.2.12

Current Routing & Interface Table



[Help..](#)

Current Route & Interface Table

Destination	Via	Device	Metric	Protocol	Source	ToS	Weight
192.168.0.0/24		eth1		kernel	192.168.0.3		

[Refresh](#)

Device	Link Up	Address	Netmask	Bcast/Peer	MTU	Txqueuelen
lo	Yes	127.0.0.1	8		16436	0
eth1	Yes	192.168.0.3	24	192.168.0.255	1500	1000
eth2	No	192.168.2.1	24	192.168.2.255	1500	1000
eth3	No	192.168.3.1	24	192.168.3.255	1500	1000
eth4	No	192.168.4.1	24	192.168.4.255	1500	1000

[Return to Network Utilities](#)

Figure 14: Current Routing & Interface Table

This menu displays the current routing table and the state of the router's interfaces.

Select the Refresh link in order to refresh the display.

The entries under the *Destination* field reflect the network or host which can be reached through this route. The "default" entry matches any packet which has not already matched another route.

The entries under the *Via* field reflect the address of the gateway to route packets through to reach the target network. The field is blank for non-gateway routings.

The entries under the *Device* field reflect the name of the interface this route belongs to. Packets using this route will be sent on this interface.

The entries under the *Metric* field reflect the cost of this route. The route with the lowest metric matching a destination is used.

The entries under the *Protocol* field reflect the system that created the route. It is one of "kernel" (default interface routes), "core" (dynamic routing protocol routes), "redirect" (routes added due to ICMP redirect message) or "static" (for manually added routes).

The entries under the *Source* field reflect the source address to use when originating a packet to a destination matching this route. Note that packets routed through the router have their own source address. Note that if the sending application decides to, it can manually specify the source address.

The entries under the *ToS* field reflect the ToS value a packet must match to be routed by this route.

The entries under the *Weight* field reflect the relative bandwidth or quality of this link within a multi-path route. Note that multi-path routes are shown with multiple lines for a single destination.

Interface Status

- This menu also summarizes the interface status.
- The entries under the *Device* field reflect the name of the device.
- The entries under the *Link up* field reflect the current link state of interface.
- The entries under the *Address* field reflect the local address of interface.
- The entries under the *Netmask* field reflect the netmask applied to this interface.
- The entries under the *Bcast/Peer* field reflect the broadcast address for the interface or the peer address if the interface is a point to point interface.
- The entries under the *MTU* field reflect the Maximum Transmission Unit size for the interface.
- The entries under the *Txqueuelen* field reflect the transmit queue length for the interface.

3 Device Management

This chapter describes how to configure and manage the device and its components, such as module interfaces, logs and files. It describes the following tasks:

**NOTE**

For information about how to configure the device to work with a network, refer to [Chapter 5, Setup and Configuration](#).

- [Section 3.1, “Accessing the Router”](#)
- [Section 3.2, “The Router Setup Shell”](#)
- [Section 3.3, “Using The LED Status Panel”](#)
- [Section 3.4, “Obtaining Chassis Information”](#)
- [Section 3.5, “Setting Up a Router Software Repository”](#)
- [Section 3.6, “Reflashing the Router Software”](#)
- [Section 3.7, “Maintaining the Router”](#)
- [Section 3.8, “Configuring PPP and the Embedded Modem”](#)
- [Section 3.9, “Configuring PPP and the Cellular Modem”](#)
- [Section 3.10, “Configuring Serial Protocols”](#)
- [Section 3.11, “Synchronous Serial Ports”](#)
- [Section 3.12, “Configuring SSH”](#)
- [Section 3.13, “Configuring the Telnet Server”](#)
- [Section 3.14, “Configuring IRIG-B and IEEE1588”](#)
- [Section 3.15, “Configuring the Intrusion Detection System”](#)
- [Section 3.16, “Brute Force Attack Protection System”](#)

Section 3.1

Accessing the Router

This section familiarizes the user with the ROX Serial Console interface, the ROX Setup script and signing on to the Web interface. This section describes the following procedures:

- Running the Setup Script
- Signing on the Web Interface
- Signing on to the Command Prompt
- Restoring the default configuration

You can access the router through the console, Ethernet ports, WAN ports and the modem port.

Section 3.1.1

Accounts and Password Management

The router provides an "rrsetup" account which provides a shell that quickly configures such items as passwords, addresses, date/time and services offered by the router. It is very useful to sign-in to this shell first, harden the router, and configure network addresses in order that the router be reachable from the network through Web Management.



NOTE

The rrsetup password should be changed, recorded securely and restricted to qualified personnel.

The root account provides a superuser capability for SSH shell access and the Web server.



NOTE

The root password should be changed, recorded securely and restricted to qualified personnel.

The root and rrsetup accounts may be also be managed through RADIUS authentication.

The Web management agent can be accessed through the root account. It may also be accessed through a number of RADIUS accounts via RADIUS authentication. This offers the advantage of attributing actions in logs to the specific user, as opposed to the root user.

Section 3.1.2

Default Configuration

The RX1000 is shipped from the factory with the following defaults:

- Ethernet ports are enabled and have an address of 192.168.X.1 where X is the port number,
- WAN and modem ports are disabled,
- IRIG-B output ports are disabled,
- Setup account "rrsetup", password "admin",
- Superuser account "root", password "admin",
- SSH and Web Management interfaces are enabled by default. All other services (including Serial Protocol Server, DHCP server, NTP server, End to End Backup Server, VPN Server, NFS, OSPF/RIP protocol and firewall) are disabled by default.

Section 3.1.3

Accessing the Device Command Prompt From the Console Port

Attach a terminal (or PC running terminal emulation software) to the RS232 port on the rear of the chassis. The terminal should be configured for 8 bits, no parity operation at 38.4 Kbps. Hardware and software flow control must be disabled. Select a terminal type of VT100.

Once the terminal is connected, pressing <CR> will prompt for the user to login as and that user's password. Sign-in as either the rrsetup or root user. The router is shipped with default passwords of "admin" for either of these accounts.

Section 3.1.4

Accessing the Device Command Prompt From SSH

**CAUTION!**

Configuration hazard – risk of data corruption. Access to the Command Line Interface (CLI) is provided for troubleshooting purposes and should only be used by Siemens technicians. As such, this interface is not fully documented. Misuse of CLI commands can corrupt the operational state of the device and render it inaccessible.

Use an SSH agent running the version 2 protocol. SSH to either the rrsetup or root accounts of the router at one of its IP addresses described above. The router is shipped with default passwords of "admin" for either of these accounts.

Section 3.2

The Router Setup Shell

Signing-in as the rrsetup user will automatically enter the configuration shell shown below. Quitting the shell (with cancel, or by entering escape) will cause the connection to close.

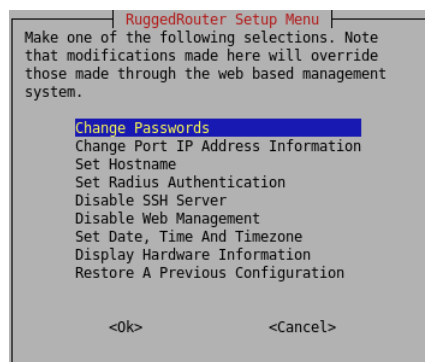


Figure 15: Router Setup Main Menu

The shell provides a number of configuration commands, described below.

Section 3.2.1

Configuring Passwords

The *Change Passwords* command changes the rrsetup and root account passwords. These passwords should be changed before installing the router on the network.




Figure 16: Router Setup Password Change Menu

Section 3.2.2

Setting the Hostname and Domain

The *Set Hostname* command sets the hostname and the domain.




Figure 17: Hostname and Domain Configuration Menu

Section 3.2.3

Setting the Hostname and Domain

The *Set Hostname* command sets the hostname and the domain.




Figure 18: Hostname and Domain Configuration Menu

Section 3.2.4

Configuring RADIUS Authentication

The *Set RADIUS Authentication* command configures the address of a RADIUS server, if one is available.




Figure 19: RADIUS Server Configuration Menu

The *Hostname/IP* field configures the RADIUS server's IP address.

The *Port Number* field sets the port number used by the RADIUS server. The default port for RADIUS is 1812.

The *Shared Secret* field configures a unique password used to authenticate communications with this server. Note that the shared secret must also be configured on the RADIUS server for the router being configured.

The *Timeout* field sets the maximum time in seconds to wait for responses from the RADIUS server before aborting a transaction.

The entry, created for both LOGIN and PPP Login, can be changed from the web interface.

Section 3.2.5

Enabling and Disabling the SSH and Web Server

By default SSH and Web Management are enabled. The *Disable SSH* and *Disable Web Management* commands allows these services to be disabled. The servers will be immediately stopped. If access to the shell has been made through ssh the session will continue, but no new sessions will be allowed.

Upon disabling the services, the titles in the main menu will change to *Enable SSH* and *Enable Web Management* to reflect the disabled state. Enabling a service automatically restarts it.

Section 3.2.6

Enabling and Disabling the Gauntlet Security Appliance

The Gauntlet security Appliance requires a pass phrase unique to your network. This menu will configure it.




Figure 20: Gauntlet Setup Menu

Section 3.2.7

Configuring the Date, Time and Timezone

The *Set The Date, Time and Timezone* command allows these parameters to be set.




Figure 21: Router Date/Time/Timezone Menu

Once set, the router will account for Daylight Savings time.

Section 3.2.8

Displaying Hardware Information

The Display Hardware Information command describes commissioned hardware.

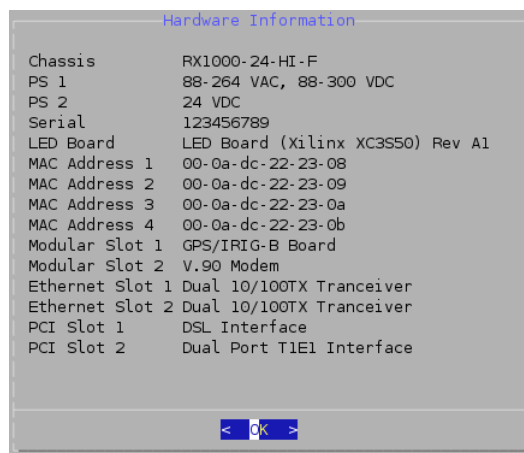


Figure 22: Router Hardware Information Menu

Section 3.2.9

Restoring a Configuration

The *Restore A Previous Configuration* command provides a means to restore a previously taken snapshot of the configuration of the router.

**NOTE**

The router will reboot immediately after restoring configuration.

The user is first prompted to select either the factory default configuration or a previously made archive.

**NOTE**

Restoring the factory defaults will reset IP addresses and may make the router impossible to reach from the network.

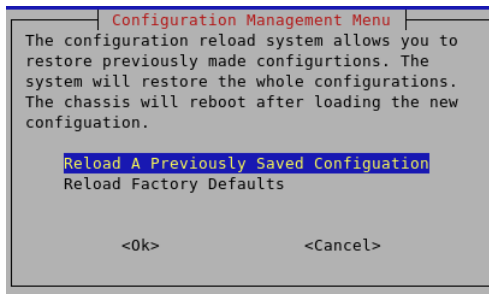


Figure 23: Selecting a configuration to reload

Initially, your router will have no previously saved configurations. The factory defaults will always be available. Once a configuration is selected the archive will be restored. After the configuration is restored, the router will reboot immediately.

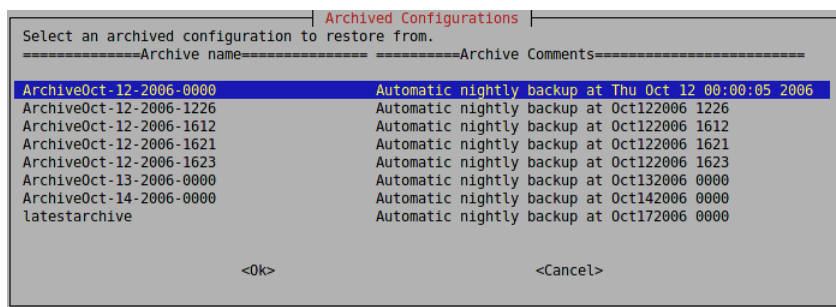



Figure 24: Selecting a previously made configuration

Section 3.3

Using The LED Status Panel

The LED status Panel provides the console port, indicates the status of hardware/software and can initiate a controlled reboot.

**Figure 25: LED Status Panel**

The LEDs are organized into three primary groups; the port group, GPS/PPP group and the Alarm/Power Supply group. The display possibilities are as follows:

Table: Meaning of LEDs

LED Name	Description
LED 1-4	Green: link activity on Ethernet port 1-4
LED 5-8	Green: link detected on Ethernet port 1-4 Red: link failure on Ethernet port 1-4
LED 9-12	Green: link activity on WAN port 1-4
LED 13-16	Green: link detected on WAN port 1-4 Red: link failure on WAN port 1-4
LED 17-20	Green: link activity on WAN port 5-8
LED 21-24	Green: link detected on WAN port 5-8 Red: link failure on WAN port 5-8
PPP-DATA	Green: link activity on PPP Modem port
PPP-LINK	green: link detected on PPP Modem port
GPS-LOCK	Green: The PTP card has acquired a GPS satellite lock
ALARM	Red: A Major Alarm exists
POWER 1	Green: Power Supply 1 is working properly Red: failure detected in Power Supply 1
POWER 2	Green: Power Supply 2 is working properly Red: failure detected in Power Supply 2

The software will cause the ALARM LED to become active for various reasons. Any condition that causes the ALARM LED to become active will activate the critical fail relay. The Web interface displays the alarms.

Pressing the pushbutton for more than five seconds will reboot the router.

Section 3.4

Obtaining Chassis Information

The chassis displays the hardware inventory at boot time. This information is captured in the `/var/log/messages` file after boot. The Web Management interface home page displays the chassis serial number.

Section 3.5

Setting Up a Router Software Repository

ROX upgrade mechanism requires a repository of software to be available. The following instructions detail:

- Requirements for a repository server,
- Initial set up of a repository,
- Upgrading the repository to the latest release,
- Maintain separate release streams for different groups of routers,
- Setting up one router to test new releases
- Configuring the network routers.

Section 3.5.1

Repository Server Requirements

In order to establish a repository you will need a host that is accessible to the routers that will be upgraded. This host must be able to act as a web server or ftp server. The host must also be able to access www.siemens.com/ruggedcom in order to download new releases of software from Siemens.

The server requirements are fairly modest. The principal requirements are for disk space, bandwidth and the ability to serve an adequate number of http sessions.

Each software release will require approximately 50 Mb of disk space. Note that this figure includes an entire software image, most upgrades will involve the transfer of only a small fraction of this amount. A large number of such releases could easily be stored on a system of only modest capabilities. In practice, only one or two releases are usually all that need be kept.

The bandwidth requirements are determined by the many factors including the number of routers, size of upgrade, when the routers upgrade, bandwidth limiting at each router and network bandwidth capability. Most web servers can serve files to the limit of the network interface bandwidth, so even a modest (e.g. 486 class machine) would prove acceptable.

The server should be able to accept at least as many http or ftp connections as there are upgradable routers in the network. In practice you will configure the routers to have staggered upgrade times in order to minimize the impact of upgrading on the network. A large upgrade (or a low bandwidth limiting value at each router) may cause all the routers to be upgrading at any one time.

Section 3.5.2

Initial Repository Setup

You must create a directory on the web server to hold the releases for the router. The directory can have any name, such as "ruggedrouter".

Some administrators like to designate one router to test the impact of new software. This will require a directory, such as "ruggedroutertest" to be created.

These directory names will be used in examples in the remainder of this section.

Ensure that the web server publishes these directories.



NOTE

If you are using Microsoft Internet Information Services (IIS) Manager 6.0 or higher as your upgrade repository, see [Section 6.2.1, "Using Microsoft Internet Information Services \(IIS\) Manager 6.0 or Higher as an Upgrade Repository"](#) for special instructions.

Section 3.5.3

Upgrading the Repository

ROX releases are obtained from www.siemens.com/ruggedcom as ZIP files. Download the ZIP file to your regular and/or test release directories and unzip them. You may delete the original ZIP file if desired.

The ZIP file name will be in the form rrX.Y.zip. The major release number X is changed when major new functionality (often hardware related) is offered. The minor release number Y is increased when minor functionality is added or bug repairs are made. The first router upgrade release is rr1.1.zip.

The zip file will extract to a directory that has the same name as the major release, e.g. "rr1". As subsequent release are made, they will also be extracted into this directory.

Section 3.5.4

Setting Up the Routers

The name of the release directory, and the major and minor release names from the zip file tells you how to set up the routers.

Suppose you have just unzipped rr1.2.zip into "ruggedroutertest" on a server available to the network at server.xyz.net. The major release is rr1 and the minor release is 2. You have chosen this directory because you want to test the release on a specific machine before propagating it to the network.

Login to the test router and visit the *Maintenance* menu, *Upgrade Software*, *Change Repository Server* sub-menu. Change the *Repository server* field to "http://server.xyz.net/ruggedroutertest" and the *Release Version* field to "rr1". You can proceed to upgrade the router manually or wait for the next nightly upgrade to take place.

After you are satisfied that the upgrade was successful you can proceed to unzip the rr1.2.zip file into your "ruggedrouter" directory (or copy the rr1/dists/rr1.2 and rr1/dists/current directories into or the "ruggedrouter" directory).

Ensure that the remainder of the routers to be upgraded have a *Repository server* field to "http://server.xyz.net/ruggedrouter" and the *Release Version* field to "rr1". They can now be upgraded.

Section 3.5.4.1

An Alternate Approach

You can eliminate the need for separate release and test directories by making your routers upgrade to a specific major and minor releases.

In this approach you will always extract releases to the same directory, e.g. "ruggedrouter".

All routers will be configured with a *Repository server* field set to "http://server.xyz.net/ruggedrouter" and the *Release Version* field initially set to "rr1.1". When you need to upgrade to rr1.2 you will visit the routers and update the *Release Version* field.

This method is simpler, but has the disadvantage that you need to visit each of the routers. This can become unwieldy when there are many routers to manage.

Section 3.5.4.2

Upgrading Considerations

The device offers you the ability to perform automatic daily upgrades, specify the download time and limit the download bandwidth. These tools automate the upgrade process and minimize the impact of upgrading on the network.

When automatic daily upgrades are used, you may wish to stagger the upgrade time of the routers. If your network has a natural "ebb flow" period of traffic activity, schedule the upgrades during this time. As an example if you have 20 routers to upgrade and they must be upgraded over an eight hour period, configure each router to start its upgrade 20 minutes after the previous router.

Be careful with limiting download bandwidth in the router. Typical upgrades will involve less than 5 MBytes of traffic. If bandwidth limiting is employed and set to 8 Kbps the upgrade will require upwards of 1.5 hours to complete.

Administrators should also be wary of routers which concentrate locally connected routers as the upgrade bandwidth consumed on the network link could reach the sum of all bandwidth limiting settings.

Routers using Frame Relay with CIR under-subscription may also encounter lengthier downloads because of retransmission.

Section 3.6

Reflashing the Router Software

ROX provides a utility to perform a complete software reinstallation. ROX persistent storage is implemented using flash memory; rewriting this memory is referred to as "reflashing".

The ROX operating system software and the complete configuration are both stored in flash memory. The reflashing process overwrites the software and reverts the configuration to its factory defaults. When reflashing to a release of ROX version 1, the reflashing utility provides an opportunity to restore a saved configuration archive.



NOTE

To restore a configuration archive using the reflashing utility, the configuration must have been made on a router using the same version of ROX as is being reflashed.

Section 3.6.1

Use Cases

The typical ROX upgrade method is through the [Section 3.7.10, "Upgrade System"](#) menu in Webmin. The the reflashing utility provides an alternative mechanism that also allows you to downgrade to an older version of ROX. The following are some examples of when to use the reflashing utility instead of the Webmin-based upgrade process:

- A network management authority may designate a particular version of ROX as the only version approved for use on the network. Newer routers added to a network may have a more recent firmware version than the approved version. Therefore, the older and approved version of ROX needs to be installed on new routers.

- A router that has become misconfigured, or whose configuration is no longer trusted for some reason, can be restored to a clean, reference state by reinstalling ROX from the factory image.
- Network staff may wish to explore how certain features operated in a previous ROX release.

Section 3.6.2

Refashing the ROX System Software

The reflashing procedure comprises the following steps:

1. Obtain the flash image of the desired ROX version from Siemens Customer Support. Make the file available to the ROX to be reflashed via a web server.
2. If necessary, back up the device's configuration, and place the configuration archive file in the same location as the image file in the previous step. Note that only a configuration archive saved using the same ROX version as the image to be reflashed will be recoverable using the reflashing utility.
3. Ensure that the ROX to be reflashed can reach the web server over the network or through a direct Ethernet connection. To verify connectivity, use the Webmin [Section 2.2.2, "Ping Menu"](#).
4. Using a serial terminal connected to the device's console port, log in and reboot the router. As the router reboots, repeatedly press the **Down** arrow key until the boot menu appears:

```
+-----+
| Debian GNU/Linux, kernel 2.6.26-2-gx1          |
| Debian GNU/Linux, kernel 2.6.26-2-gx1 (BIST mode) |
| Debian GNU/Linux, kernel 2.6.26-2-gx1 (recovery mode) |
| Software Reflash Utility                        |
|                                                 |
+-----+
```

5. Select **Software Reflash Utility** and press **Enter**.
6. Follow the prompts to provide the following information:
 - which Ethernet interface to use (1/2/3/4)
 - whether to use a static IP address or DHCP
 - when using a static IP address, provide the IP address, netmask, and gateway
 - the URL specifying the location of the reflash image file. Use a numeric IP address and not a DNS name. For example: `http://192.168.1.1/imagerr1.14.1.tar`
 - the URL specifying the location of a configuration archive file, if applicable. When reflashing to ROX2, the router does not prompt you to specify a configuration archive file.

7. After setting the reflash parameters, a confirmation prompt appears:

```
Do you want to [S]tart reflashing, change setting [1/2/3/4] or [E]xit?
```

- To change a setting, type 1, 2, 3, or 4 and press **Enter**.
 - To exit without reflashing, type **E** and press **Enter**.
 - To reflash the memory, type **s** and press **Enter**.
8. On selecting **s**, the router repartitions the flash memory and installs the specified image file, followed by the configuration archive (if one is specified). When the reflash is complete, the confirmation prompt appears again. Type **E** and press **Enter**. The router reboots in normal mode.
 9. After rebooting into the newly flashed ROX software image, you can configure the router in the usual manner through Webmin, or by logging in to a local (serial) or remote (SSH) console as the "rrsetup" user.



NOTE

The router must not lose power or be interrupted during the reflashing process. The process involves a complete rewrite of the operating software image. Interruption will require that the router be returned to the factory to have the software restored. Siemens recommends minimizing the risk by using a standalone PC as the web server and by powering both the web server and the router using an uninterruptible power supply (UPS).

Section 3.7

Maintaining the Router

This section familiarizes the user with:

- Viewing Alerts
- Backing up and restoring configurations
- Decommissioning the device
- Configuring SNMP
- Configuring RADIUS Authentication
- Configuring Outgoing Mail
- Chassis Parameters
- Power over Ethernet
- Banner Configuration
- Using System Logs
- Upgrading Software
- Using Pre-upgrade/Post-upgrade scripts
- Uploading and downloading files

Section 3.7.1

Alert System

The alert system provides the following features:

- Generates alerts, displaying them locally and/or forward them via email messages.
- Alerts are set and cleared by the daemons that own them. Active alerts are locally displayed and can be cleared manually.
- Multiple forwarders can be configured, a configurable filter level controls alert forwarding to each forwarder.
- By configuring different forwarders, low severity and high severity control centers can be set up.

Each alert is mapped to an alert definition entry, which is predefined by a daemon who owns the alert or by a user. All alert definition entries are configurable by user.

An alert filter is a user defined configuration to define the forwarders destination of active alerts. Any active alerts with Renotify Interval set to non-zero value and matches with the filter level will be forwarded to the defined forwarder destination.

Section 3.7.1.1

Alert Main Menu



Figure 26: Alert Main Menu

This menu displays active alerts and allows you to change alert system configuration and alert definitions.

Follow the *All Alerts* link to show all alerts. The view of system alerts may be limited by severity by following one of the severity links:

- Emergency
- Alert
- Critical
- Error
- Warning
- Notice
- Info
- Debug

or by category by following one of the category links:

- chassis
- performance
- interface
- daemon

Note that active alerts are volatile and will be regenerated after reboot. If you clear an alert manually, it will appear if the condition occurs again. You may disable the alert permanently by disabling the alert from its entry in the definition menu.

The *Clear Alert* link under the *Action* column allows you to clear the alert.

Clicking on the Alert Name, Specific, Severity and Date column headers will sort the alerts by those types.

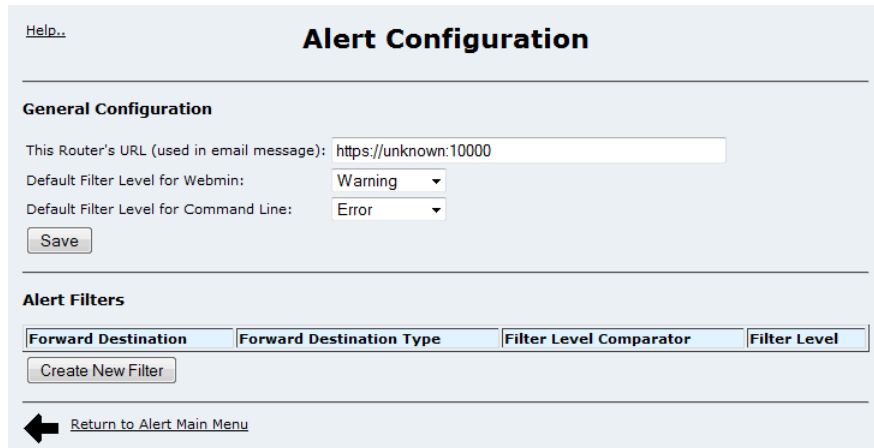
Select *Alert Configuration* to change the generic configuration and alert filter configurations.

Select *Alert Definition configuration* to change the alert definition entries.

Section 3.7.1.2

Alert Configuration

This menu configures the general information and forward filters for the alert system.



The screenshot shows the 'Alert Configuration' web interface. At the top left is a 'Help..' link. The title 'Alert Configuration' is centered. Below it is a 'General Configuration' section with three fields: 'This Router's URL (used in email message):' with the value 'https://unknown:10000', 'Default Filter Level for Webmin:' with a dropdown set to 'Warning', and 'Default Filter Level for Command Line:' with a dropdown set to 'Error'. A 'Save' button is below these fields. The next section is 'Alert Filters', which contains a table with four columns: 'Forward Destination', 'Forward Destination Type', 'Filter Level Comparator', and 'Filter Level'. Below the table is a 'Create New Filter' button. At the bottom left is a back arrow icon, and at the bottom right is a link 'Return to Alert Main Menu'.

Forward Destination	Forward Destination Type	Filter Level Comparator	Filter Level
---------------------	--------------------------	-------------------------	--------------

Figure 27: Alert Configuration Menu

The *This Router's URL* configures the link to access this router. This information will be used in the email forwarder, which user can click on the link in the email to access the router.

The *Default Filter Level for Webmin* configures the lowest alert level to show on Webmin. All active alerts higher priority than this level will be displayed on the Webmin home page.

The *Default Filter Level for Command Line* configures the lowest alert level to show when user login by console or ssh.

The *Save* button saves all changes of general configuration.

The *Create New Filter* button allows you to create a new forwarder filter for active alerts.

Section 3.7.1.3

Alert Filter Configuration

Help...

Change Filter Configuration

The filter configuration is successfully saved.

Filter Parameters

Forward Destination Type	Email
Forward Destination	admin@example.com
Filter Level Comparator	Greater Than
Filter Level	<input type="checkbox"/> Emergency <input type="checkbox"/> Alert <input type="checkbox"/> Critical <input type="checkbox"/> Error <input checked="" type="checkbox"/> Warning <input type="checkbox"/> Notice <input type="checkbox"/> Info <input type="checkbox"/> Debug

Use comma to separate multiple email addresses.

Save Delete

Return to Alert Configuration Menu

Figure 28: Alert Filter Configuration Menu

This menu configures an alert filter, which defines the forwarder destination for active alerts matching with defined filter level.

The *Forward Destination Type* configures the type of filter. Currently only type Email is supported.

The *Forward Destination* configures the destination matching with the Forwarder Destination Type. Note that multiple email addresses should be separated by comma.

The *Filter Level Comparator* configures the way to match with defined filter level.

The *Filter Level* configures what filter level is to be compared. Note that Emergency has the greatest filter level and Debug has the lowest filter level.

Section 3.7.1.4

Alert Definition Configuration

[Help..](#)

Alert Definition Configuration

View Alert Definition by Category chassis

Codepoint	Category	Name	Subsystem	Severity	Enabled	Alarmable	Renotify Interval(second)	Type
chassis:1	chassis	Inventory Problem	chassis	error	yes	yes	0	simple
chassis:2	chassis	Power Supply 1 Failure	chassis	critical	yes	yes	0	simple
chassis:3	chassis	Power Supply 2 Failure	chassis	critical	yes	yes	0	simple
chassis:4	chassis	Ledboard Push Button	chassis	error	yes	yes	0	simple
chassis:5	chassis	Ledboard	chassis	error	yes	yes	0	simple
chassis:6	chassis	Modular slot 1	chassis	error	yes	yes	0	simple
chassis:7	chassis	Modular slot 2	chassis	error	yes	yes	0	simple
chassis:8	chassis	Ethernet interface module 1	chassis	error	yes	yes	0	simple
chassis:9	chassis	Ethernet interface module 2	chassis	error	yes	yes	0	simple
chassis:10	chassis	PCI slot 1	chassis	error	yes	yes	0	simple
chassis:11	chassis	PCI slot 2	chassis	error	yes	yes	0	simple
chassis:12	chassis	Chassis	chassis	error	yes	yes	0	simple
chassis:13	chassis	Watchdog	chassis	error	yes	yes	0	simple
chassis:14	chassis	Power on Ethernet	chassis	error	yes	yes	0	simple

Create New Definition


 [Return to Alert Main Menu](#)

Figure 29: Alert Definition Configuration Menu

This menu displays matched alert definition entries. It also allows user to change an alert definition entry or create a new entry.

An alert definition entry defines an alert which will be monitored by the system.

The *View Alert Definition by Category* allows you to display alert definition entries matching with selected category.

The *Create New Definition* button allows you to create a user defined alert definition entry.

Click on one of the link under the *Codepoint* column allows you to change the configuration for that alert definition entry.

Section 3.7.1.5

Change Alert Definition

[Help...](#)

Change Alert Definition

Alert Definition Parameters

Codepoint cfgw:1	Category daemon
Name Upgrade made changes	Subsystem configwatch
Severity Warning	Alarmable <input type="checkbox"/>
Enabled <input checked="" type="checkbox"/>	Renotify Interval(second) <input checked="" type="radio"/> Disabled <input type="radio"/> (1-86400 seconds)
Type Simple	

Parameters for Shell

Sample Interval (30-86400 seconds)	Command
Comparator Greater than	Threshold
And Repeats (0-1000000 times)	And Until (0-1000000 seconds)
Not Cleared Repeats (0-1000000 times)	Not cleared Until (0-1000000 times)

Parameters for RMON

Device Name	MIB Variable
Sample Interval (30-86400 seconds)	Startup Event Rising
Rising Threshold	Falling Threshold

[Return to Alert Definition List Menu](#)

Figure 30: Change Alert Definition Menu

This menu allows you to change an existing alert definition entry.

The *Codepoint* is the key part of the alert definition entry and does not allow to be changed.

The *Category* configures which category the alert definition entry belongs to.

The *Name* configures the name of the alert definition which will be displayed by Webmin, login or email forwarder when an active alert exists.

The *Subsystem* configures which subsystem the alert definition entry belongs to.

The *Severity* configures the severity level of the alert. The severity level is sorted from highest priority to lowest priority.

The *Alarmable* configures whether the matched alert should trigger the critical relay and alarm LED on the LED panel of the router.

The *Enabled* configures whether the alert system should monitor and record matched active alert. If Enabled is not checked, matching active alert will be ignored.

The *Renotify Interval* configures how often should the matched active alert be notified according to alert filter configuration setting. If it is disabled, no notification will be forwarded.

The *Type* configures type of the alert definition entry. There are three types available: *Simple*, *Shell* and *RMON*. Currently only the first two types are supported. If users choose *Shell* type, they should complete parameters under *Parameters for Shell* table.

The *Parameters for Shell* table allows user to configure additional parameters if the alert definition entry type is *Shell*.

The *Sample Interval* configures how often should the system run configured shell command to get a sample.

The *Command* configures the shell command to run.

The *Comparator* configures how to compare with the shell command result.

The *Threshold* configures the threshold to compare with the shell command result to see whether the condition is true or false.

The *And Repeats* configures how many times the condition must be true before the alert is generated.

The *And Until* configures how many seconds the condition should be true before an alert is generated.

The *Not Cleared Repeats* configures how many times the condition must be false before the alert is cleared.

The *Not cleared Until* configures how many seconds the condition must be false before an alert is cleared.

The *Parameters for RMON table* allows user to configure additional parameters if the alert definition entry type is RMON type.

The *Device Name* configures the name of the device to be monitored.

The *MIB Variable* configures the MIB variable being monitored.

The *Sample Interval* configures how often samples should be generated.

The *Rising Threshold* configures the value that will trigger an event when the value of the variable increments past this value.

The *Falling Threshold* configures the value that will trigger an event when the value of the variable decreases past this value.

The *Startup* configures the condition that will cause the initial event.

Section 3.7.2

Backup and Restore

The Backup and Restore system provides the following features:




Figure 31: System Backup and Restore

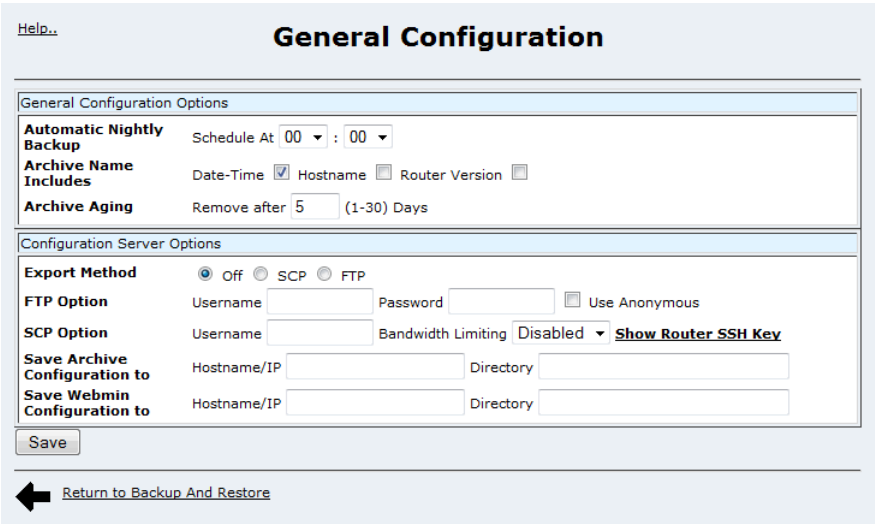
- All configuration settings are saved in a configuration archive,
- Webmin configuration settings are saved in a Webmin configuration archive,
- Archives can be used to "clone" routers, replicate a damaged resource or unwind a change,
- Archives can be created manually (including user comments) or by the Automatic nightly backup, which captures all changes over the previous 24 hours,
- The nightly backup archives can be automatically transferred via scp or ftp to a designated server,
- The nightly backup archives are kept on the router for a configurable number of days and then deleted. The most recently made archive is never destroyed. Manually created archives are never destroyed.

- If you make a configuration change you later wish to reverse you can restore a previously made archive completely. An archive difference tool is provided, showing the difference between one archive and either another archive or the current configuration. Changes in configuration can also be detected and "unwound" by applying the previous state of a router on a file by file basis.
- Archive filename is user definable and can include any of date/time, host name and/or release version,
- Archives can be uploaded to the router and restored.
- A Configuration Rollback feature that allows users to safely make modifications under a safety net.
- A factory defaults file is included.

Note the following caveats:

- Chassis specific items such as serial number, hardware inventory and MAC addresses are not saved,
- Log and history files are not saved,
- Information stored in the root and user accounts are not saved.

Section 3.7.2.1

General Configuration

The screenshot displays the 'General Configuration' web interface. At the top, there is a 'Help..' link and the title 'General Configuration'. Below this, the 'General Configuration Options' section includes: 'Automatic Nightly Backup' with a 'Schedule At' dropdown set to '00 : 00'; 'Archive Name Includes' with checkboxes for 'Date-Time' (checked), 'Hostname', and 'Router Version'; and 'Archive Aging' with a 'Remove after' field set to '5' days. The 'Configuration Server Options' section includes: 'Export Method' with radio buttons for 'Off' (selected), 'SCP', and 'FTP'; 'FTP Option' with 'Username' and 'Password' fields and a 'Use Anonymous' checkbox; 'SCP Option' with 'Username' and 'Bandwidth Limiting' (set to 'Disabled') fields, and a 'Show Router SSH Key' link; 'Save Archive Configuration to' with 'Hostname/IP' and 'Directory' fields; and 'Save Webmin Configuration to' with 'Hostname/IP' and 'Directory' fields. A 'Save' button is located at the bottom of the form. Below the form is a 'Return to Backup And Restore' link with a left-pointing arrow.

Figure 32: Backup and Restore General Configuration

This menu configures the backup system.

The *Automatic Nightly Backup* field specifies when the nightly backup is scheduled. The automatic export to a server will start (if enabled) immediately after the backup completes.

The *Archive Name Includes* field selects text fields (Date-Time, Hostname, Router Version) included in archive name.

The *Archive Aging* field specifies how long nightly backup archives are kept. Note that the most recently made nightly backup will never be deleted. Manually made archives are never aged and must be manually deleted.

The Configuration Server Options table allows user define the configuration server.

The *Export Method* field selects the method of exporting backup archives to a server. If the *Export Method* field is set to FTP, the FTP Options are used. If the *Export Method* field is set to SCP, the SCP Options are used.

The *FTP Option* field specifies FTP User name, Password or to use anonymous FTP .

The *SCP Option* field specifies SCP User name and Bandwidth Limitation when the Export Method is SCP. The *Show Router SSH Key* link will display the ssh public key for this router, which can be used in the configuration server to accept SCP from the router.

The *Save Archive Configuration to* field specifies the configuration server hostname (or IP address) and the directory in which to save configuration archives.

The *Save Webmin Configuration to* field specifies the configuration server hostname (or IP address) and the directory in which to save Webmin configuration archives.

Section 3.7.2.2

Configuration Rollback

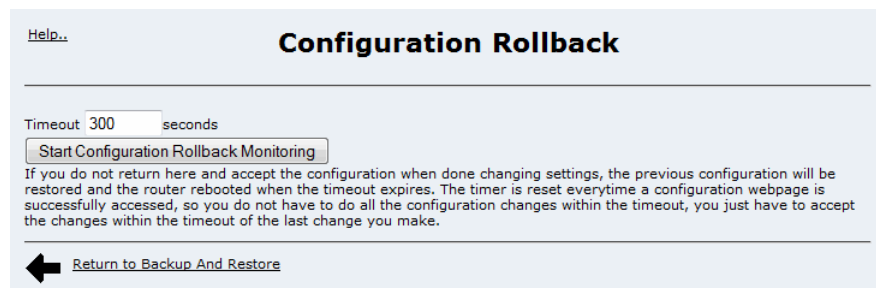


Figure 33: Configuration Rollback Menu

The Configuration Rollback menu enables the user to define a period of time in which configuration changes can be made and subsequently accepted. If the user does not explicitly accept the changes being made, then the unit will revert to a configuration snapshot that was taken when the user started the configuration rollback. In reverting to this configuration, the unit will reboot. This enables configuration changes to be made under a safety net. If a configuration change effectively blocks the user's access to the unit or has any other detrimental effect, the unit, not seeing any user acceptance, will reboot at the end of the timeout period.

When activated by pressing the *Start Configuration Rollback Monitoring* button, each subsequent configuration screen will display the line "*Configuration Rollback is active*" at the top of the page, reminding the user that there is a timeout period in effect, and an eventual reboot of the unit if changes are not accepted. For example:

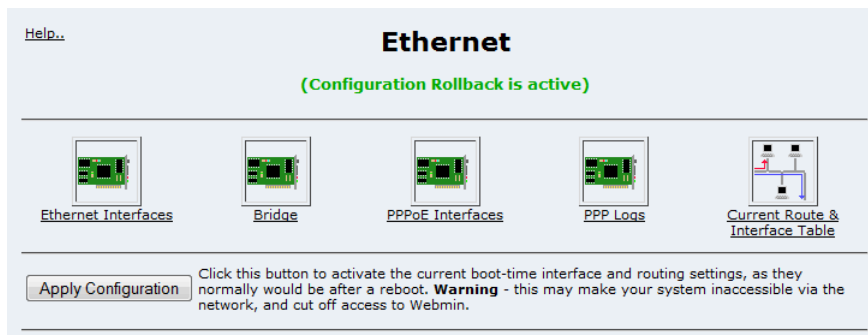


Figure 34: Ethernet main menu while Configuration Rollback is active

Please note that the timeout period is re-initialized to the value specified in the timeout entry field at each user action. Using the above timeout value the user does not have a maximum of 300 seconds. Rather, each of his actions will reset the timer to 300 seconds. The timeout mechanism will be in place from the time the user presses the *Start Configuration Rollback Monitoring* button until he eventually goes back to the Configuration Rollback menu and presses the *Accept Current Configuration* button:




Figure 35: Configuration Rollback menu ready to accept changes

Section 3.7.2.3

Archive History

The screenshot displays the 'Archive History' web interface. At the top, there is a 'Help...' link and the title 'Archive History'. Below this, the 'Configuration Archives' section shows a total size of 6749404 bytes and a list of four archives. Each archive has a checkbox, a name, a version, and a comment. A 'Remove Selected Archives' button is located below the list. The 'Webmin Archives' section shows a total size of 11796 bytes and a list of two archives, also with checkboxes, names, versions, and comments. A 'Remove Selected Archives' button is also present. Below these sections is an 'Upload configuration archives or webmin archives from your current host to this router' section. It includes an 'Archive to upload' field with a 'Browse...' button and the text 'No file selected.' Below this is an 'Upload To Router' button. At the bottom, there is a back arrow icon and a link 'Return to Backup And Restore'.

[Help...](#)

Archive History

Configuration Archives

The total size of all archived configurations is 6749404 bytes. Click on an archive to download a copy of it.

Archive Name	Version	Archive Comment
<input type="checkbox"/> Archive20131022-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
<input type="checkbox"/> latestarchive	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
<input type="checkbox"/> Archive20131019-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-19 00:00
<input type="checkbox"/> factorydefaults	rr1.16.0-QA1.1	Factory defaults

Webmin Archives

The total size of all archived configurations is 11796 bytes. Click on an archive to download a copy of it.

Archive Name	Version	Archive Comment
<input type="checkbox"/> Webmin20131022-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
<input type="checkbox"/> latestwebminarchive	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00

Upload configuration archives or webmin archives from your current host to this router

Archive to upload No file selected.

[Return to Backup And Restore](#)

Figure 36: Archive History

The Archive History menu displays current system configuration archives (including all configurations) and Webmin configuration archives (only includes Webmin configurations), sorted by date (most recent first). Following the link of an archives under the *Archive Name* field upload a copy of it.

Selecting an under the *Archive Name* field and applying the *Remove Selected Archives* button will delete the archive. Note that only manually backup archives can be deleted. Automatic nightly backup archives will automatically aged out. The latestarchive and factorydefaults archives will never be deleted.

The *Archives to upload* fields select archives to upload to the router. The *Browse...* button will allow you to select an archive. Applying the *Upload to Router* button will upload the specified archive to the router.

Section 3.7.2.4

Archive Backup

Figure 37: Archive Backup

This menu allows the user to manually create a configuration archive or Webmin archive.

The *Backup Type* field determines which type (configuration archive or Webmin archive) of archive you want to backup.

The *Archive Comment* field sets a comment which will be included in the archive file.

The *Backup archive file name* field allows you to input the candidate archive file name.

Starting the backup results in the following display.

Figure 38: Archive Backup, Complete

The archive created can be immediately uploaded if desired by following the "Upload A Copy Of This Archive.." link.



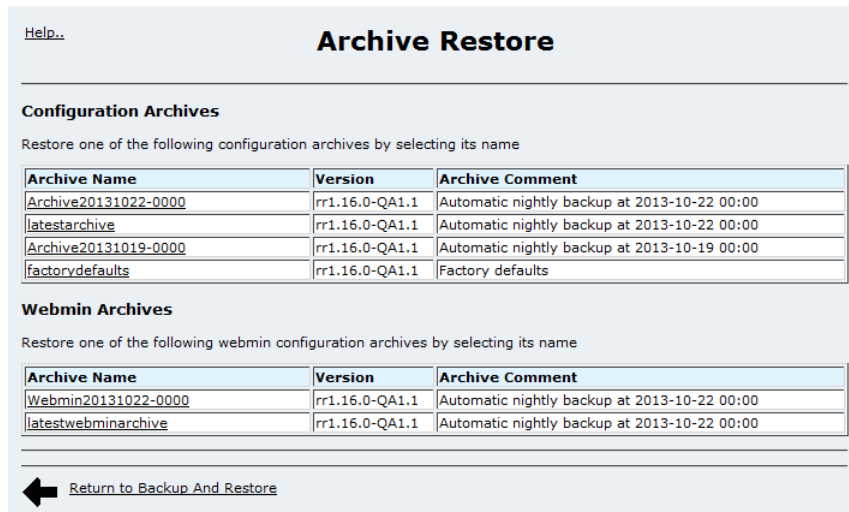
NOTE

If you use the Internet Explorer web browser, you must "Right-click" the link and save the file manually. Otherwise Internet Explorer will rename the file after uploading, preventing its use in a subsequent archive restore.

Section 3.7.2.5

Archive Restore

The restore process begins by selecting an archive to restore from. Following an archive link will restore the archive and reboot the router.



[Help..](#)

Archive Restore

Configuration Archives

Restore one of the following configuration archives by selecting its name

Archive Name	Version	Archive Comment
Archive20131022-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
latestarchive	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
Archive20131019-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-19 00:00
factorydefaults	rr1.16.0-QA1.1	Factory defaults

Webmin Archives

Restore one of the following webmin configuration archives by selecting its name

Archive Name	Version	Archive Comment
Webmin20131022-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
latestwebminarchive	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00


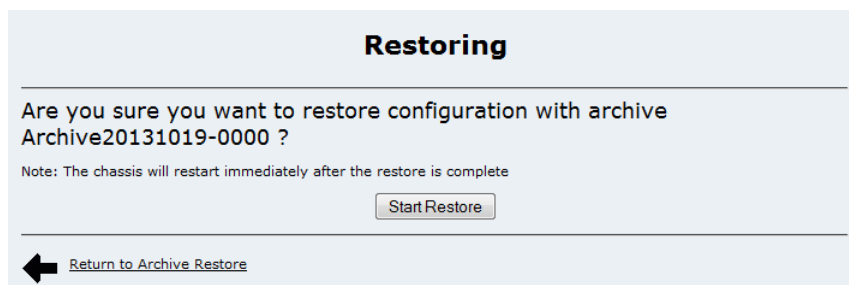
 [Return to Backup And Restore](#)

Figure 39: Archive Restore Menu

Click on one of the links under *Archive Name* to start the restore. Starting the restore results in the following display.



Restoring

Are you sure you want to restore configuration with archive
Archive20131019-0000 ?

Note: The chassis will restart immediately after the restore is complete


 [Return to Archive Restore](#)

Figure 40: Start Restore

To begin the restoring process, click the *Start Restore* button.

Section 3.7.2.6

Archive Difference Tool

[Help...](#)

Archive Differences

Select Configuration Archives to Show Differences

Archive Name	Version	Archive Comment
<input type="checkbox"/> Archive20131022-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
<input type="checkbox"/> latestarchive	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
<input type="checkbox"/> Archive20131019-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-19 00:00
<input type="checkbox"/> factorydefaults	rr1.16.0-QA1.1	Factory defaults
<input type="checkbox"/> Current Configuration	rr1.16.0-QA1.1	Current Configuration on router

Note: select two and only two targets

[Show Differences](#)

Select Webmin Configuration Archives to Show Differences

Archive Name	Version	Archive Comment
<input type="checkbox"/> Webmin20131022-0000	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
<input type="checkbox"/> latestwebminarchive	rr1.16.0-QA1.1	Automatic nightly backup at 2013-10-22 00:00
<input type="checkbox"/> Current Configuration	rr1.16.0-QA1.1	Current Configuration on router

Note: select two and only two targets

[Show Differences](#)

[← Return to Backup And Restore](#)

Figure 41: Archive Differences Menu

The Archive Difference menu shows the difference between two targets. The first target must be an archive while the second target can be either another archive or the current configuration.

Choose two and only two targets and click the *Show Differences* button.

Archive Differences List

Differences between archive Archive20131019-0000 and Current Configuration

File Name	Archive20131019-0000	Current Configuration
quagga/daemons	2013-05-30 11:36:07	2013-10-21 15:45:02
default/lldpd	2013-10-18 15:07:05	2013-10-21 15:39:14
default/service_snmpd	2013-10-08 15:52:20	2013-10-21 15:39:13
default/service_portmap	2013-10-08 15:52:17	2013-10-21 15:34:53
default/service_openl2tp	2013-10-08 15:52:17	2013-10-21 15:34:52
default/service_serserver	2013-10-08 15:52:15	2013-10-21 14:08:13
default/service_l2tunnelld	2013-10-08 15:52:15	2013-10-21 14:08:12
default/service_linkd	2013-10-08 15:52:06	2013-10-21 14:08:12
default/service_openswan	2013-10-08 15:52:26	2013-10-21 14:08:12
default/service_keepalived	2013-10-08 15:52:20	2013-10-21 14:08:11

Files only exist in archive Current Configuration

File Name	Timestamp
webmin/shorewall/version	2013-10-22 10:24:03
snmp/persistent/snmpd.conf	2013-10-21 15:39:15
openl2tp/openl2tpd.conf	2013-10-21 15:34:59
webmin/sshd/version	2013-10-21 15:11:52



[Return to Archive Differences](#)

Figure 42: Archive Differences List

The resulting menu shows the differences between the two selected targets. Files in this table are sorted by the change time (most recent changes first). Files that exist in only one of the targets are shown separately.

Following the links under *File Name* column will show a files difference between the two targets.

The difference will be shown by two methods. The difference between the two targets will be first be shown in a side by side scrollable comparison.

The difference will also be shown in a window that shows differing lines.

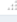


Show Difference	
Difference on 'default/lldpd' between Archive20131019-0000 and Current Configuration	
Side by Side Difference Display	
DAEMON_ARGS=""	DAEMON_ARGS="-x"
	
Differing Lines Display	
1c1 < DAEMON_ARGS="" --- > DAEMON_ARGS="-x"	
	
Note: lines beginning with '<' belong to Archive20131019-0000; lines beginning with '>' belong to Current Configuration	
Copy This File To Current Configuration	
 Return to Archive Differences List	

Figure 43: Show Difference for selected file between two targets

The *Copy This File to Current Configuration* button will be present when the destination archive is the Current Configuration. It allows user to copy the selected file from the old archive to current configuration.



NOTE

It is possible to damage your router through use of this feature! Ensure that the configuration file copied makes sense in the current version of the router.

Copying the configurations may not make any actual operating changes until the systems that own them are restarted.

If the source archive has a file that is not present in the Current Configuration, it is possible to view that file and then copy it into Current Configuration.

Section 3.7.3

Decommissioning the Device

Before taking the device out of service, either permanently or for maintenance by a third-party, make sure the device has been fully decommissioned. This includes removing any sensitive, proprietary information.

To decommission the device, do the following:

1. Obtain a copy of the ROX firmware currently installed on the device. For more information, contact Siemens Customer Support.
2. Log in to ROX. See [Section 2.1.1, "Using a Web Browser to Access the Web Interface"](#).
3. Clear the boot password. See [Section 4.3.2.2, "Change Bootloader Password Command"](#).
4. Flash the ROX firmware obtained in [Step 1](#), but do not select the option to restore the previous configuration. See [Section 3.6, "Reflashing the Router Software"](#).
5. Shut down the device. See [Section 4.3.1, "Bootup and Shutdown"](#).

Section 3.7.4

SNMP Configuration

The SNMP menus provide the following configuration features:

- System information
- agent network addresses
- Community access to the agent
- SNMP trap delivery

The SNMP (the Simple Network Management Protocol) protocol is used by network management systems and the devices they manage. SNMP is used to manage items on the device to be managed, as well as by the device itself, to report alarm conditions and other events.

The first version of SNMP, V1, provides the ability to send a notification of an event via "traps". Traps are unacknowledged UDP messages and may be lost in transit. SNMP V2 adds the ability to notify via "informs". Informs simply add acknowledgement to the trap process, resending the trap if it is not acknowledged in a timely fashion.

SNMP V1 and V2 transmit information in clear text (which may or may not be an issue depending the facilities the data is transmitted over) and are lacking in the ability to authenticate a user. SNMP V3 adds strong authentication and encryption.

Section 3.7.4.1

SNMP Main Configuration Menu




Figure 44: SNMP Main Configuration Menu

Configuring SNMP on ROX requires that the SNMP daemon be running. Enable `snmpd` (the snmp daemon) via the use the *System* folder, *Bootup and Shutdown* menu.

**NOTE**

Prior to ROX 1.10.0, SNMP was manually configured used the `com2sec`, `group`, `view` and `access` directives. If so configured, the SNMP menu will prompt you to convert the configuration to one it can manage.

Section 3.7.4.2

System Configuration

A screenshot of the 'System Configuration' menu. At the top left is a 'Help..' link. The title 'System Configuration' is centered. Below the title is a 'System Variables' section with four fields: 'System name' (empty), 'System location' (Unknown (configure /etc/snmp/snmpd.local.t), 'System contact' (Root <root@localhost> (configure /etc/snmp, and 'System description' (empty). Below these fields is a 'Save' button. At the bottom left is a back arrow icon and a 'Return to SNMP Configuration' link.

Figure 45: System Configuration Menu

The *System name*, *System location*, *System contact*, and *System description* fields configure descriptive parameters for the router.

Section 3.7.4.3

Network Addressing Configuration

For reference, the set of currently configured and active IP addresses is listed near the top of the page.

Client IP Address (Source IP):

IP Address

NOTE: If this option is not specified, the source address of SNMP packets from this host is the IP address of the interface from which the packet **exited the host**.

Figure 46: Network Addressing Configuration Menu, Client Address

The *Client address (Source IP)* field specifies the address from which snmpd will send notifications. If the field is blank, the default behaviour will be to transmit the notification from the IP address of the interface from which the message leaves the router. Snmpd will return to this behaviour if the configured address is not available when it starts.

Addresses to listen on:

Interface Name	IP Address	Listening
lo	127.0.0.1	<input checked="" type="radio"/> enabled <input type="radio"/> disabled
eth1	192.168.0.3	<input checked="" type="radio"/> enabled <input type="radio"/> disabled
eth2	192.168.2.1	<input checked="" type="radio"/> enabled <input type="radio"/> disabled
eth3	192.168.3.1	<input checked="" type="radio"/> enabled <input type="radio"/> disabled
eth4	192.168.4.1	<input checked="" type="radio"/> enabled <input type="radio"/> disabled
New		

NOTE: snmpd is currently configured to listen on all active IPV4 interfaces.

Figure 47: Network Addressing Configuration Menu, Addresses to listen on

The table of *Addresses to listen on* includes the list of currently configured and active IP addresses, and whether the address is currently listened on. The *New* field allows for the addition of other IP addresses.

Snmpd will use these addresses providing they are active at the time it starts. By default, snmpd listens on all interfaces.

Section 3.7.4.4

Access Control

SNMP V1 and V2c Communities:

No V1 or V2c communities are currently defined

Add an SNMP V1 or v2c Community Name

Community Name

Access read-only

Source IP

OID

Add

Figure 48: Access Control Menu, SNMP V1 and V2c

The first part of the Access control page allows the creation and deletion of SNMP V1 and V2c community names.

The *Community Name* field selects the name of the community. The *Access* field determines whether the community is read-only or read/write. The *Source IP* field may be used to specify an IP address or range (e.g. 10.0.0.0/24) from which access to this community name may be made. The *OID* field further restricts access to an Object Identifier (OID) tree at or below a specified OID.

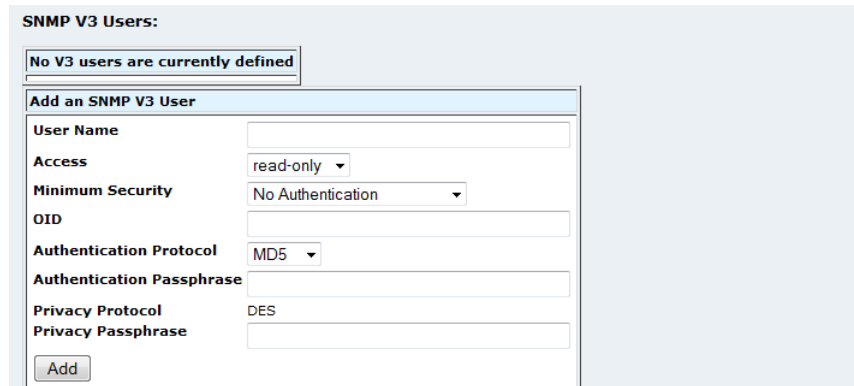
The image shows a web-based configuration window titled "SNMP V3 Users:". At the top, a message box states "No V3 users are currently defined". Below this is a section titled "Add an SNMP V3 User". This section contains several fields: "User Name" (a text input field), "Access" (a dropdown menu currently showing "read-only"), "Minimum Security" (a dropdown menu currently showing "No Authentication"), "OID" (a text input field), "Authentication Protocol" (a dropdown menu currently showing "MD5"), "Authentication Passphrase" (a text input field), "Privacy Protocol" (a dropdown menu currently showing "DES"), and "Privacy Passphrase" (a text input field). At the bottom of this section is an "Add" button.

Figure 49: Access Control Menu, SNMP V3

The second part of the Access control menu allows creation and deletion of V3 users.

The *User Name* field selects the name of the new user.

The *Access* field determines whether the community is read-only or read/write.

The *Minimum Security* field selects the level of security used by this user. It may be No Authentication (no authentication or encryption), Authentication Only (authentication by MD5 or SHA1 authentication methods, no encryption) or Authentication with Privacy (authentication by MD5 or SHA1, encryption by DES or AES ciphers).

The *OID* field further restricts access to an Object Identifier (OID) tree at or below a specified OID.

The *Authentication Protocol*, *Authentication Passphrase*, *Privacy Protocol* and *Privacy Passphrase* fields configure the protocols and passphrases used depending on the *Minimum Security* field. These settings are shared between agent and remote user.

Note that if authentication and privacy are both used, but only the authentication passphrase is provided, snmpd will use the authentication passphrase as the privacy passphrase.

Note also that if any notifications are enabled, a read-only user named *internal* will be automatically created to satisfy the requirements of the event MIB.

Section 3.7.4.5

Trap Configuration



Figure 50: Trap Configuration Menu, Trap Options

The Trap Configuration page manages SNMP trap destinations. Under *Trap Generation Options*, you may enable the generation of notifications on authentication failures or IP interface link up/down events.

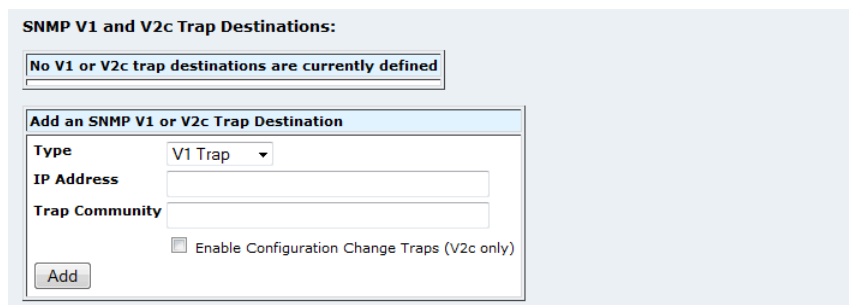


Figure 51: Trap Destinations V1 and V2c

The *SNMP V1 and V2c Trap Destinations* part of the menu allows the creation and deletion of trap destinations. The *Type* field specifies the exchange used with this destination, either V1 trap, V2c trap or V2c inform. The *IP address* and *Trap Community* fields specifies the receivers IP address and community name.

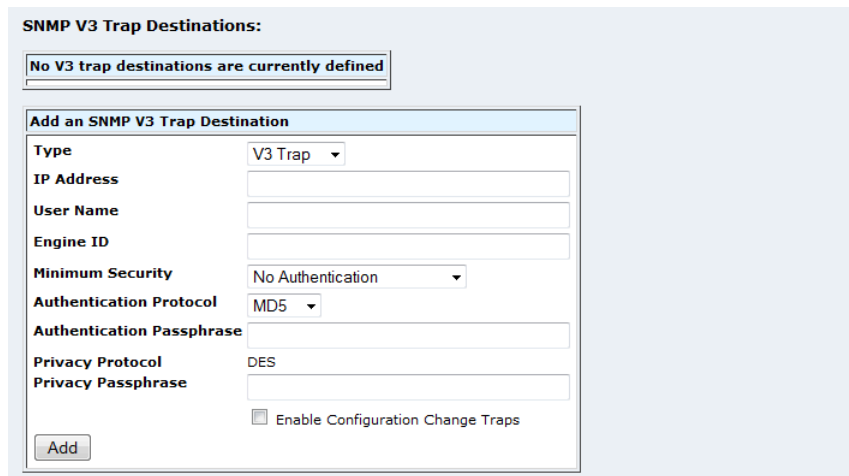


Figure 52: Trap Destinations V3

The *SNMP V3 Trap Destinations* part of the menu allows the creation and deletion of V3 trap destinations.

The *Type* field specifies the exchange used with this destination, either V3 trap or V3 inform.

The *IP address* and *User Name* fields specify the trap receiver's IP address and user name.

The *Engine ID* parameter is necessary for "inform" type notification destinations only, and must be configured by the trap receiver in order to receive these notifications. The value is considered hexadecimal ASCII. It does not require a '0x' prefix nor an 'h' suffix, and is limited to 64 characters.

The *Minimum Security*, *Authentication Protocol*, *Authentication Passphrase*, *Privacy Protocol* and *Privacy Passphrase* fields are as described in [Section 3.7.4.4, "Access Control"](#).

The *Enable Configuration Change Traps* option enables traps to be sent when the system's configuration is modified. Please note that there is only one destination allowed, for both V2c and for V3. For example, if this option is enabled in one V2c definition, then it will be the unique trap destination for configuration changes. Same with V3. If it is enabled in one V3 definition, then it is not allowed to be defined elsewhere, in V3 or V2c. Siemens SNMP MIB definitions are available from www.siemens.com/ruggedcom and from Customer Support.

Section 3.7.4.6

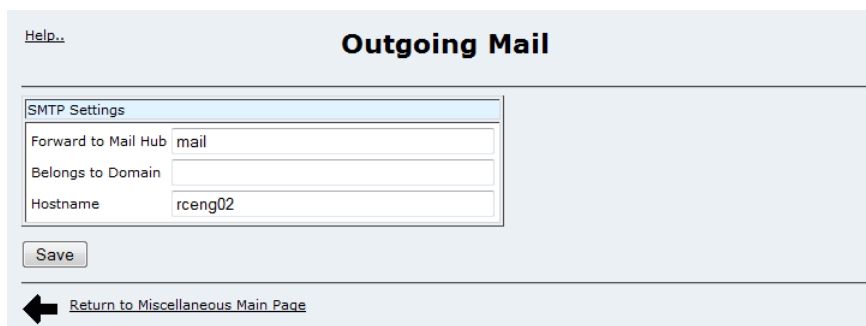
MIB Support

ROX device supports the following MIBs.

MIB Name	MIB Description
IF-MIB	The MIB module to describe generic objects for network interface sub-layers
SNMPv2-MIB	The MIB module for SNMPv2 entities
TCP-MIB	The MIB module for managing TCP implementations
IP-MIB	The MIB module for managing IP and ICMP implementations
UDP-MIB	The MIB module for managing UDP implementations
LLDP-MIB	The MIB module for managing LLDP
SNMP-VIEW-BASED-ACM-MIB	View-based Access Control Model for SNMP
SNMP-FRAMEWORK-MIB	The SNMP Management Architecture MIB
SNMP-MPD-MIB	The MIB for Message Processing and Dispatching
SNMP-USER-BASED-SM-MIB	The management information definitions for the SNMP User-based Security Model

Section 3.7.5

Outgoing Mail

**Figure 53: Outgoing Mail**

Outgoing Mail is configured from within the *Maintenance* menu *Miscellaneous* sub-menu. This menu controls where emails originated by the router are forwarded to.

The *Forward to Mail Hub* field specifies an IP address or domain name of a host that accept mail from the router.

The *Belongs to Domain* field specifies the email domain the router is part of. This information is written into the email header upon transmission.

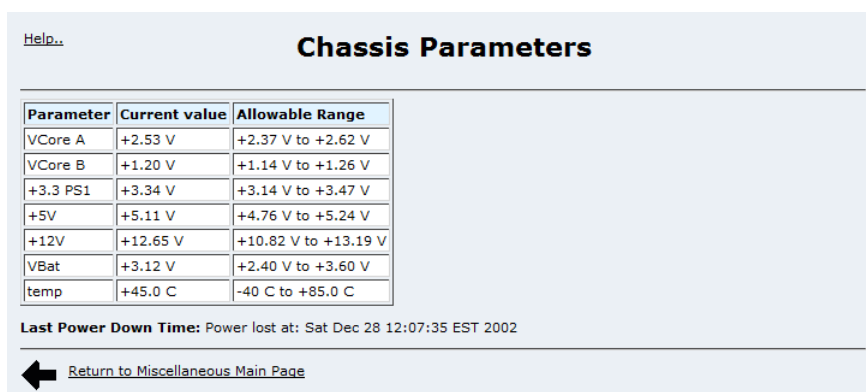
The *Hostname* field specifies the hostname to be written into the email header upon transmission.

**NOTE**

You can generate emails from scheduled commands and scripts with
"(echo "To: ops@myco"; echo -e "Subject: Hello!\n"; some-command) | sendmail -t".

Section 3.7.6

Chassis Parameters



Parameter	Current value	Allowable Range
VCore A	+2.53 V	+2.37 V to +2.62 V
VCore B	+1.20 V	+1.14 V to +1.26 V
+3.3 PS1	+3.34 V	+3.14 V to +3.47 V
+5V	+5.11 V	+4.76 V to +5.24 V
+12V	+12.65 V	+10.82 V to +13.19 V
VBat	+3.12 V	+2.40 V to +3.60 V
temp	+45.0 C	-40 C to +85.0 C

Last Power Down Time: Power lost at: Sat Dec 28 12:07:35 EST 2002

Figure 54: Chassis Parameters Menu

This menu displays the chassis temperature and, if hardware version 2, the voltage levels of chassis power supplies and a record of the last power down time. The system will highlight red any out-of-range value. The monitored values are described below:

Parameter	Description
temp	Motherboard temperature
VcoreA, VCoreB	Redundant 3.3V power supply voltages
+3.3 PS1, +3.3 PS2	Redundant 3.3V power supply voltages
+5V	5V power supply voltage
+12V	12V power supply voltage
VBat	Battery voltage

The last power down time reflects the time power was removed from the chassis as a result of a power failure, commanded reboot or an watchdog initiated reboot.

System alarms will be generated for out-of-range parameters and watchdog initiated reboots.

Section 3.7.7

Power over Ethernet

The IEEE 802.3af standard describes a method known commonly as PoE (Power over Ethernet), for providing electrical power over the twisted pair wiring most commonly used in Ethernet networks. The obvious benefit is the ability to take advantage of previously unused copper in the 10/100Base-T wiring configuration to provide power without requiring that new wiring be installed.

The device can be provisioned to supply power according to IEEE standard 802.3af. In order to provide PoE via all four router Ethernet ports, the unit must be provisioned with a PoE power supply in place of the optional redundant power supply. ROX implements PoE mode A, supplying up to 400mA at up to 48V on pins 3 and 6 (T568A pair 2) and return on pins 1 and 2 (T568A pair 3).




Figure 55: PoE Pinout on 10/100BaseT Ports

Pin	Signal	PoE Mode A
1	RX+	V -
2	RX-	V -
3	TX+	V +
4	Reserved (Do Not Connect)	
5	Reserved (Do Not Connect)	
6	TX-	V +
7	Reserved (Do Not Connect)	
8	Reserved (Do Not Connect)	
Case	Shield	

Section 3.7.7.1

Power over Ethernet Menu

Power over Ethernet

Port Name	Enabled	Power Limit(W)	Power Delivery	Class	Voltage(V)	Current(mA)	Power(W)	Status
eth1	Enabled ▾	15.0 (0-20W)	No	-	-	-	-	Ok
eth2	Enabled ▾	15.0 (0-20W)	No	-	-	-	-	Ok

Save

Refresh


 [Return to Miscellaneous Main Page](#)

Figure 56: Power over Ethernet Menu

This menu allows you to enable/disable the Power over Ethernet function and set the power limitation on available Ethernet ports. It also displays the current status on each port.

The *Port Name* column identifies the Ethernet port number.

The *Enabled* column allows you to enable or disable the Power over Ethernet function on this port.

The *Power Limit* column specifies the monitored power limitation on the port. An alarm will be generated if the power is over this limitation.

The *Power Delivery* column shows whether the power is delivered on this port.

The *Class* column shows the class of the PD device on this port.

The *Voltage*, *Current* and *Power* columns show the voltage, current and power value (in unit of Volt, mill Ampere and Walt, respectively) when there is power delivery on this port.

The *Status* column shows whether there is any error detected on this port.

Section 3.7.8

Banner Configuration

[Help..](#)

Banner Configuration

Login Banner Configuraiton - Customize console login messages

Message Before Login

WARNING: You are attempting to access a private computer system. Access to this system is restricted to authorized persons only. This system may not be used for any purpose that is unlawful or deemed inappropriate. Access and use of this system is electronically monitored and, by entering this system, you are affirming your consent to be electronically monitored. We reserve the right to

Information After Login

☒ Last Login ☒ System Information ☒ Router Status

Extra Message After Login

SSH Banner Configuration - Customize ssh login messages

Information After Login

☒ Last Login ☒ System Information ☒ Router Status

Extra Message After Login

Webmin Banner Configuration - Customize webmin login messages

Session Header

☒ Default ☐

Session Message

☒ Default ☐

Username

☒ Default ☐

Password

☒ Default ☐

Login Button

☒ Default ☐

Clear Button

☒ Default ☐

Message when Login Fails

☒ Default ☐

Logout Message

☒ Default ☐

Message for Session Timeout

☒ Default ☐

Save Configuration

Reset to Default

Figure 57: Banner Configuration Menu

This menu allows you to customize different aspects of each of the access methods to the device:

- Serial console login
- SSH login
- Webmin login

The *Login Banner Configuration* menu customizes the messages seen prior to and after login via the serial console.

The *Message Before Login* field is the banner displayed above the login prompt.

The *Information After Login* fields select from among three utilities to display to the user on a successful login:

- *Last login* causes information about the last login to be displayed: what time, from where, and on what terminal the user last logged in.
- *System Information* displays information about the running operating system kernel.
- *Router Status* displays information on the ROX release, Webmin version, serial number, uptime, temperature, disk, memory, and pending alarms.

The *Extra Message After Login* field is displayed after a successful login.

The *SSH Banner Configuration* has the same customizations as above except that no message can be displayed prior to the login prompt.

The *Webmin Banner Configuration* menu allows you to customize the web-based login box seen on connecting to the device's secure web interface:

- The *Session Header* field is the description displayed at the top of the login box.
- The *Session Message* field is the message displayed above the username and password prompts.
- The *Username* and *Password* fields are the strings displayed at the left side of the login box, describes that the input boxes are for username and password, respectively.
- The *Login Button* field is the string displayed at the left button of the login box.
- The *Clear Button* field is the string displayed at the right button of the login box.

The following image illustrates the position of each of these fields in the login box.




Figure 58: Webmin Banner Configuration Fields

1. Session Header 2. Session Message 3. Username Box 4. Password Box 5. Login Button 6. Clear Button

The *Message when Login Fails* field specifies the message displayed when the login fails.

The *Logout Message* field specifies the message displayed upon logout.

The *Message for Session Timeout* field specifies the message displayed when the session is timeout.

Section 3.7.9

System Logs

System logs are records of activities that have occurred on the router, sorted into specific categories. System logs can be invaluable when debugging configuration changes. As such, most of your use of the logs will likely consist simply of examining them.

[Help..](#)

System Logs

[Add a new system log](#)

Log destination	Active?	Messages selected	
File /var/log/auth.log	Yes	auth,authpriv.*	View..
File /var/log/syslog	Yes	*.* ; auth,authpriv.none	View..
File /var/log/cron.log	No	cron.*	
File /var/log/daemon.log	Yes	daemon.*	View..
File /var/log/kern.log	Yes	kern.*	View..
File /var/log/messages	Yes	*,=info ; *,=notice ; *,=warn ; auth,authpriv.none ; cron,daemon.none ; mail,news.none	View..
All users	Yes	*.emerg	

[Add a new system log](#)

Secure Remote Syslog Setting

Enable☐

CA Certificate

Certificate

Key

Permitted Peer Common Name

☐ Do not check ☒ Match Pattern (Multiple patterns are seperated by spaces)

Note: Secure remote syslog facility only works for TCP connection with remote syslog server. When secure remote syslog is enabled, all TCP connection will be secure connection.

Remote Syslog Source IP Bind Interface

Note: This option is used to bind the selected interface IP address as source IP for remote syslog messages. If "none" is selected, it will use the output interface IP address as the source IP address.

Click this button to make the current configuration active by killing the running `syslog` process and restarting it.

Figure 59: System Logs Menu

The *System Logs* menu screen is used to configure the system logging process. It consists primarily of a list of log destinations, each of which may be a log file on the router itself, a remote syslog server, or one of several other destinations. Please refer to [Section 3.7.9.3, "Remote Logging"](#) for more detail on log destinations. Two links, *Add a new system log* are provided above and below the list in order to configure new system logs.

Each entry in the list displays information for a particular log managed by syslog:

- *Log destination* displays the location or logging method of the log.
- *Active?* displays whether syslog is logging messages to the log.
- *Messages selected* displays the filtering criteria used to include messages in the log.

ROX implements a set of default system logs, as described in [Section 3.7.9.1, “Syslog Factory Defaults”](#).

System logs can be encrypted for secure remote logging, as described in [Section 3.7.9.2, “Enabling Secure Remote Syslog”](#).

In support of remote syslogging, the *Remote Syslog Source IP Bind Interface* field makes it possible to bind the selected network interface's IP address to syslog. Syslog messages transmitted by the router will have the selected interface's IP address as their source. If "none" is selected, the source IP address will be that of the network interface from which messages are transmitted to the remote syslog server.

Any changes made using this menu, including adding or modifying system log configurations or changing the IP Bind Interface address, require clicking *Apply Changes* in order to take effect. If the syslog daemon is not yet running, the button will instead read: *Start Syslog Server*.

Section 3.7.9.1

Syslog Factory Defaults

Although new logs can be created (and the type of information saved in existing logs changed) the factory defaults are as follows:

- *messages* - This log file catches a wide variety of generic information excluding authentication, cron and mail messages. This should be the first log you inspect when starting to debug a problem.
- *syslog* - This log file catches all information with the exception of authentications. Syslog contains all that messages contains, and more. Examine this log if you can not find relevant information in messages.
- *auth.log* - This log file catches authentication requests. View auth.log when you are trying to debug a problem in which a user is not able to sign on to a service (such as web management or ssh).
- *critical* - This log catches reports of critical failures. There should never be any messages in this log. Your Siemens Customer Support representative may ask you to inspect this file.
- *kern.log* - This log contains messages issued by the kernel (the most central part of the operating system). This log always displays messages issued at boot time, and should rarely be added to after that. Your Siemens Customer Support representative may ask you to inspect this file.
- *cron.log* (initially disabled) - This log file contains messages from the cron systems notifying of tasks started through cron. Your Siemens Customer Support representative may ask you to enable and inspect this log.
- *daemon.log* (initially disabled) - This log file contains messages from daemons (programs that run continuously in the background). Your Siemens Customer Support representative may ask you to enable and inspect this log.

Each one of the default logs above is represented in the [Figure 59](#) along with any others that may have been created.

Left unrestricted, the logging system would consume all available disk space, causing the router to fail. The router limits the memory used by the logging system by storing logs in a volatile (i.e. lost after a reboot) file system which is limited in size. Such a system will lose logging information when a power failure occurs, too much logging is generated or as the result of a user commanded reboot.

The router deals with this problem by storing compressed versions of three key files (messages, auth.log, and critical) to the permanent disk. The log files are saved every 180 seconds and upon an orderly reboot. The log files are restored during the next boot. All other files but these are cleared.

Section 3.7.9.2

Enabling Secure Remote Syslog

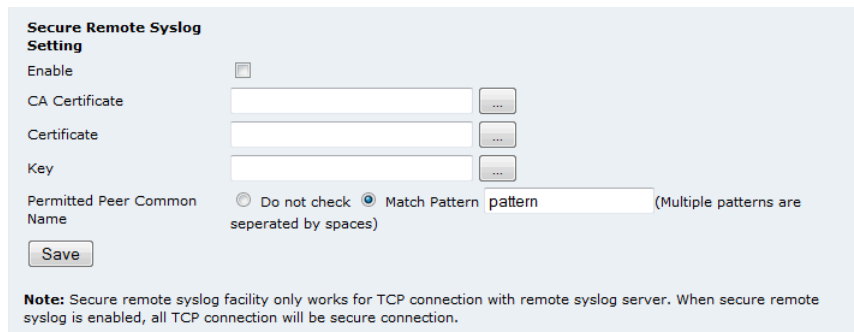


Figure 60: Secure Remote Syslog Settings

ROX supports the encryption of system logs with rsyslog.

**NOTE**

All certificates must conform to the following specifications:

- X.509 v3 digital certificate format
- PEM format
- RSA key pair, 512 to 2048 bits in length

**NOTE**

Once secure remote system logging is enabled and a remote syslog server is configured, TCP port 6514 is automatically opened.

Enable enables or disables secure remote syslog.

CA certificate specifies the path and filename of a CA (Certified Authority) certificate. The client and server certificates must be signed by the same Certified Authority (CA).

Certificate specifies the path and filename of a certificate.

Key specifies the path and filename of a key.

Permitted Peer Common Name allows you to match the common name in the certificate with one or more match patterns. Each match pattern must be separated by a space. The default value is "pattern", which can be replaced by one or more patterns, which can include wildcards (*). For example, the match pattern "*.example.com" will match "abc.example.com". Alternatively, the match pattern "a* *.example.com" will also match "abc.example.com", but it will not match "abc.example".

**CAUTION!**

Security hazard - risk of unauthorized access or exploitation. Selecting the Do not check radio button configures ROX to accept a certificate with any common name from the server. This mode is vulnerable to man-in-the-middle attacks and is not recommended.

If you do not want to match the common name, select the *Do not check* radio button.

Section 3.7.9.3

Remote Logging

Remote logging (often referred to as remote syslogging) is the process of forwarding log entries to a remote host computer. Remote logging enables central collation of logs and preserves logs in the events of security incidents. Remote logging does not require any file storage on the router and as such does not suffer from loss of information around unplanned power failures. On the other hand, remote logging cannot record events that occur before network connectivity to the logging host is established.

Remote logging can replace disk logging or can augment it.

If you wish to replace disk logging for some information type, select the appropriate link under the *System Logs* sub-menu *Log Destination* column. Enter the URL of the logging host under the *Syslog server on*.

Help..

Edit System Log

Log destination

Log to

- ☒ File: /var/log/auth.log
- ☐ Named pipe
- ☐ Syslog server on: Port Number 514
- ☐ Local users
- ☐ All logged-in users

Sync after each message? ☒

TCP ☐ UDP ☒

Logging active? ☒ Yes ☐ No

Message types to log

Facilities

- ☐ All
- ☒ Many: auth authpriv
- ☐ None

Priorities

- ☐ None
- ☒ All
- ☐ At or above..

Save View logfile Delete

Note: Secure remote syslog facility only works for TCP connection with remote syslog server.

Return to system logs

Figure 61: Changing a Syslog Entry to Log Remotely

If you wish to remote log in addition to disk log some log type, you must duplicate the log entry and then configure the logging host. Duplicate the entry by using the "Add a new system log" link on the *System Logs* sub-menu.

Finally, you may forward all information to the remote logger by creating a new system log entry and specifying "All" Facilities and all priorities, and checking the *Syslog server on* field with an appropriate address.

Section 3.7.10

Upgrade System

The screenshot shows a web interface titled "Software Upgrade System". At the top left is a "Help.." link. The interface is divided into four main sections, each with a title and a description, followed by a button:

- Upgrade to RX1100**: "Gain access to the RX1100 feature set, including Intrusion Detection Systems and Gauntlet Security." Button: "Upgrade To RX1100".
- Change Repository Server**: "The router needs to be configured with a repository server.. The router is currently operating release software rr1.16.0-QA1.1". Button: "Change Server".
- Install a New Package**: "Select the location to install a new package from..". It has three radio button options: "From local file" (selected), "From uploaded file" (with a "Browse.." button and "No file selected." text), and "From ftp or http URL". There is an "Install" button at the bottom of this section.
- Upgrade All Packages**: "Please configure a repository server..".

Figure 62: Software Upgrade System

The Software Upgrade system provides the following features:

- Upgrading from either HTTP or FTP servers (anonymous access only),
- Upgrade traffic bandwidth limiting to prevent disruption to mission critical applications,
- Manually initiated upgrades from a central server,
- Manually initiated upgrades of new versions for testing purposes,
- Manually initiated installs of new packages for testing purposes.

Section 3.7.10.1

ROX Software Fundamentals

You may be required to upgrade the router in order to take advantage of new features, security improvement and bug repairs.

Your ROX software is provided in releases of the form *rrX.Y.Z*. The platform release number X changes when new hardware platforms are released. The major release number Y is increased when important new features are added. This is called a "Major" release. The minor release number Z is increased when minor functionality is added or bug repairs are made. This is called a "Minor" release.

ROX is organized into a number of interdependent "packages". Each package contains all of the files necessary to implement a set of related commands or features, such as a firewall or ssh client. A router upgrade involves replacing some of these packages with newer versions and sometimes adding new packages. The upgrade system handles these functions automatically.

Section 3.7.10.2

Upgrade to RX1100

This menu allows you to upgrade your router. The display usefully provides a description of the current hardware in the router inventory.

[Help..](#)

Upgrade Inventory

This router has an RX1000 order code. In order to upgrade to an RX1100, contact your sales manager and provide them with following inventory record. When your salesperson returns you an updated record, overwrite the current record and press the "Upgrade Inventory" button. A reboot will then be required.

```

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

# Created by RuggedCom Inc. Final Test
# Product information
OrderCode=RX1000-F-RM-HI-00-XX-XX-TX01-TX01-TC2-S11
SerialNumber=X1K-0409-02268
MacAddressEth1=00-0a-dc-2a-3f-a0
-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.0 (MingW32)

iD8DBQFJ+KxVP2pya+G5kdYRAMTfAJ9+KEaGEFoLbTYUnDlndKXN9u1JLACggZbq
CLzIFxwT8XD8926u+aerj6I=
=Bcsc
-----END PGP SIGNATURE-----
Mainboard=12-01-0001 #{RuggedRouter MainBoard} 12-01-0001 Rev 03
ledboard=12-11-0015 #{LED Board (Xilinx XC3S50)} 12-11-0015 Rev B2
PowerSupply1=12-10-0008-P1 #{88-300VDC OR 85-264VAC} 12-10-0008-P1 Rev C3
ifboard1=12-11-0002 #{2x10/100TX RJ45} 12-11-0002 Rev C2
ifboard2=12-11-0002 #{2x10/100TX RJ45} 12-11-0002 Rev C2
pci1=13-01-0004 #{T1/E1 Channelized Dual} 13-01-0004 Rev A
pci2=13-01-0015 #{Synchronous Dual Serial Card} 13-01-0015 Rev A

```

[Upgrade Inventory](#)

[Return to Upgrade System](#)

Figure 63: Upgrade to RX1100

Section 3.7.10.3

Change Repository Server

[Help..](#)

Change Repository Server

Repository server

Release Version Use mX.Y to upgrade to that specific release or mX to upgrade to the latest release.

Bandwidth Limiting Disabled

[Save Changes](#)

[Return to Upgrade System](#)

Figure 64: Change Repository Server

**CAUTION!**

Security hazard – risk of data exposure. Make sure to establish a secure connection between the device and the repository server to prevent unauthorized users from obtaining information about the operating system.

This menu defines the server used to upgrade software. The *Repository server* field accepts a URL containing the domain name or IP address of an http or ftp server along with the directory on the server containing the upgrades.

The *release version* field accepts a software release string, such as "rr1", "rr1.7" or "rr1.7.2".

If you configure this field with only a major release number such as "rr1", the router will always pick the latest release at the server. As an example, if the router is running with release rr1.7 and release rr1.7.2 becomes available, the latter will be used.

If you configure this field with a major/minor/patch release number such as "rr1.7.2", the router will only upgrade from that release.

The *Bandwidth Limiting* selector allows you to limit the bandwidth used in the course of upgrading the system software.

Section 3.7.10.4

Upgrading All Packages

Upgrade All Packages is used for attended upgrades of the ROX system software to a newer revision. The upgrade process obtains a list of packages from the specified repository server and release version (see [Section 3.7.10.3, "Change Repository Server"](#)), automatically determines which packages need to be added or upgraded, and performing the necessary package acquisitions and installations.

Upgrade All Packages

Resynchronize package list (update) ☒ Yes ☐ No

Only show which packages would be upgraded ☒ Yes ☐ No

Upgrade Now

Figure 65: Upgrading All Packages

The *Resynchronize package list* field selects whether to obtain the full package list from the repository server. The list need only be obtained once per upgrade, so checking *No* can save time on the upgrade process if a first pass was performed with the "Only show..." option described below. This is especially true if the network link is a low-speed WAN link.

The *Only show which packages would be upgraded* field controls whether to only show the packages that will be upgraded (*Yes*) or to actually perform an upgrade (*No*).

After setting the two parameters described above, click *Upgrade Now* to begin the upgrade process.

**NOTE**

Webmin manages the upgrade of other packages. When Webmin must upgrade itself, the process requires an extra step. You will be requested to start a Webmin only upgrade. Webmin will start another program to manage the upgrade and will self-terminate. Webmin will automatically restart after the upgrade completes, after which time you may log back in.

Notes on Software Upgrade Procedures

ROX upgrades that involve a new "Major" release number generally require a router reboot after completion of the upgrade. Minor releases will never require a reboot. The release notes accompanying the upgrade will state whether a reboot will be necessary.

If a reboot is required, a notice will appear to that effect upon clicking *Upgrade Now* and before beginning the upgrade.

If the upgrade must be done in two stages, a notice to this effect will appear. The first stage will consist of an upgrade to Webmin only. After completion of the first stage of the upgrade, launch the second stage of the full upgrade process by again running *Upgrade All Packages*. Additional notifications may appear depending on the old and new ROX versions and on router hardware options.



NOTE

If the currently installed version of ROX predates release 1.14.1, the upgrade procedure must be done in two stages.

While the first stage Webmin upgrade is in progress, there will be no visual feedback from the system, since the web interface itself will be shut down. This upgrade lasts up to 5 minutes, after which time, it is recommended to click *refresh* to verify that the Webmin upgrade has completed correctly. Once the upgrade has completed, it may be necessary to log in if the session timeout has expired in the meantime.

Section 3.7.10.5

Installing a New Package

Figure 66: Installing a New Package

The *Install A New Package* feature uploads and installs packages to the router.

Select the *From local file* option if you have already moved the package to the router through http, ftp or scp. You may either enter the full path from the root directory to the package or use the file selector () to identify the package.

Select the *From uploaded file* option if you have the file locally on your workstation. You may either enter the location of the file on your local file system browse selector () to identify the package.

Select the *From ftp or http URL* if you know the network address of the package.

Complete the installation by selecting the install button.

Section 3.7.10.6

Pre-Upgrade/Post-Upgrade Scripts

The pre-upgrade and post-upgrade scripting feature allows you to execute defined scripts on the router both before and after a software system upgrade. The scripts execute before and after an actual upgrade only, and not a "dry run", during which the packages to be updated are shown and not actually updated.

In the course of an upgrade, the router will download the scripts from the same location as is configured via the *Change Repository Server* page. The script files named:

- *pre-upgrade-user* will be downloaded by the router and executed immediately prior to beginning the upgrade process. Note that in the case where a two-stage upgrade procedure is required, the pre-upgrade script will only be executed immediately prior to the second stage of the upgrade.
- *post-upgrade* will be downloaded by the router and executed immediately after the upgrade process has completed.

The scripts must start with `#!/bin/bash` or `#!/usr/bin/perl` in order to use one of either the BASH or PERL command interpreters on the router and be designed to produce consistent results in the event that they are run more than once consecutively. It is possible, for example, for an upgrade to be interrupted after the pre-upgrade script runs, and restarted at a later date.

The result of running the pre-upgrade script is included in the upgrade output.

The following post-upgrade script will send an email notification when an upgrade completes (assuming SMTP is configured properly).

```
#!/bin/bash
echo "Subject: Software upgrade for Release rrl.9.0 on `hostname` completed" > /tmp/mail
echo "To: controlcenter@ruggedcom.com" >> /tmp/mail
echo "Software upgrade for Release rrl.9.0 on `hostname` completed at `date`" >> /tmp/mail
echo >> /tmp/mail
cat /tmp/mail | ssmtp controlcenter@ruggedcom.com
rm -f /tmp/mail
```

Section 3.7.11

Uploading and Downloading Files

Figure 67: Upload/Download Menu

The Upload/Download Files menu provides a means to transfer files to and from the router.

The *Download files from the specified URLs to this router* part of the menu allows you to have the router download files from *ftp* and *http* servers. You need to specify (at least) the file URL and the directory to download it to. You may also decide to create directories cited in the download path at download time, set the user/group ownership of the file and postpone the download to a specific time.

The *Send files from your current host to the router* part of the menu allows you to send files from your host machine directly to the router. You need to specify (at least) one file to send and the directory to upload it to. Clicking on a browse button will open a file search dialog box. Select the file to upload to the router and close the dialog box. Click upon the *Send to router* button to start the transfer. You may also decide to create directories cited in the upload path at upload time, set the user/group ownership of the file and extract tar or zip files.

The *Upload a file from the router to your host* part of the menu allows you to send files from the router a specified your host machine. You need to specify the file to send. You may specify the files path directly or click on the browse button to open a file search dialog box. Select the file to upload and close the dialog box. Then click the *Upload to your host* button.

Section 3.8

Configuring PPP and the Embedded Modem

This section familiarizes the user with:

- Configuring PPP Client
- Configuring PPP Server
- Configuring Dial in console
- Viewing status

ROX may be equipped with one or two internal modems or with a serial card which allows connection to an external modem. A modem allows connections to be made over standard telephone lines. PPP (Point-to-Point Protocol) is used to establish a network connection over a modem link. As of ROX version 1.15, two internal modems can be installed in ROX.

Section 3.8.1

PPP Interface

When a PPP connection is established, a network interface is created in the system. The interface name for the first internal modem connection or an external modem connection is *ppp0*. If the router is equipped with a second internal modem, the interface name for the second modem connection is *ppp20*. Refer to these interface names when configuring firewall rules.

Section 3.8.2

Authentication, Addresses and DNS Servers

PPP authentication automatically uses either the PAP or CHAP protocols.

To create a PPP client connection in Webmin, obtain a user ID, password, and telephone number from the operator of the PPP server to which you are connecting. The operator might be an Internet Service Provider or a system administrator within your organization.

The authentication process provides a local IP address for use on the PPP interface and, optionally, the addresses of the DNS servers and default gateway. Generally, you should use these addresses unless you need to provide your own.

The PPP interface IP address obtained from the PPP server can be either a dynamic or a static IP address. Firewall configuration should be performed as is appropriate for the type of address.

In the case of a PPP server configuration, you must configure the parameters described above for incoming PPP client connections.

Section 3.8.3

When the Modem Connects

A PPP Client Connection may be configured to connect at boot time, or to dial only when there is traffic to be transmitted.

Section 3.8.4

PPP Dial On Demand

The PPP client can be configured to dial only when there is traffic to be transmitted. To do so, the the PPP interface must be configured to be the default gateway (on the device, configuring the PPP client to dial on demand, automatically enables the default gateway option). At the beginning, the PPP interface will be shown as being up but the link will not actually be connected. Only when there is traffic to be transmitted via the PPP interface will the configured number be dialed and a connection initiated.

The Dial on Demand feature requires IP addresses to be configured on for both the local and the remote ends of the PPP connection used. These addresses are required in order to instantiate PPP for the purpose of monitoring for traffic. Note that the local address, remote address, or both may be overridden by the PPP server at the time a connection is actually established.



NOTE

Do not configure the default gateway on the router if you decide to use PPP Dial on Demand.

Section 3.8.5

LED Designations

ROX provides a pair of LEDs to indicate information about the modem PPP connection.

PPP-Link will be green when the modem PPP link is established. It will flash while a connection is being established, or a console dial in session is active.

PPP-Data will flash green when there is traffic on the PPP link.

Section 3.8.6

PPP Modem Configuration

From the Configuration Main Menu, select Modem Configuration.

- If a single internal modem is installed, this page appears:

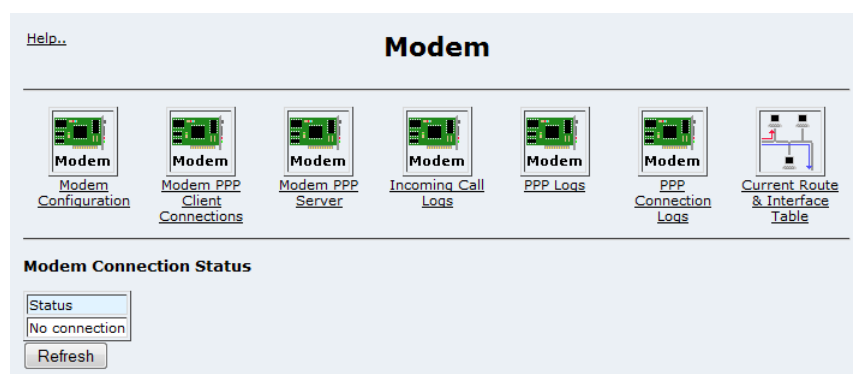


Figure 68: Modem Configuration Main Menu

On this page, review and configure the modem interface and the PPP client and server connections.

- If two internal modems are installed, this page appears:

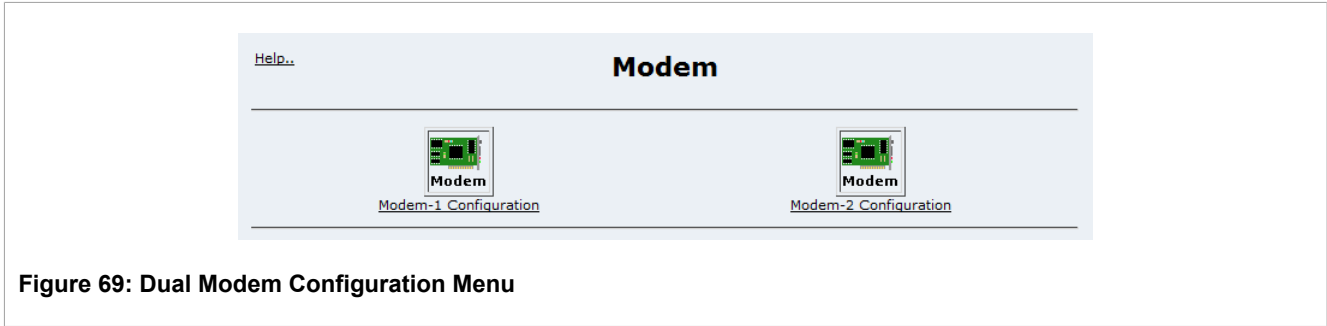


Figure 69: Dual Modem Configuration Menu

Select the modem you want to configure. The modem configuration page for the selected modem appears. In this example, the configuration page for Modem 2 is shown.

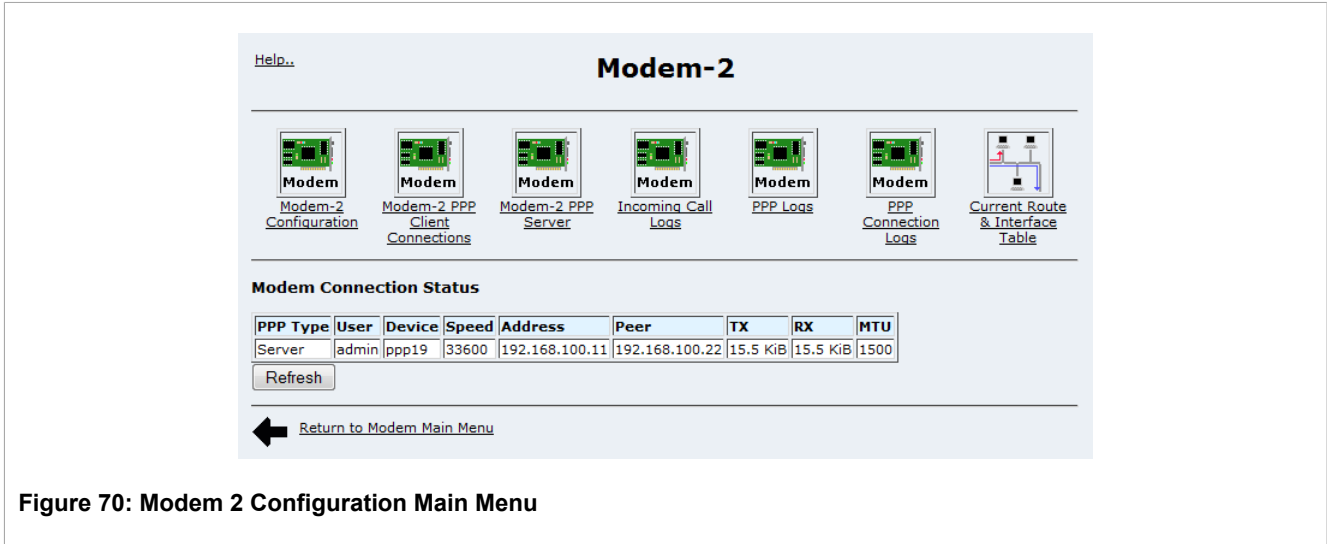


Figure 70: Modem 2 Configuration Main Menu

On this page, review and configure the modem interface and the PPP client and server connections.

Section 3.8.6.1

Modem Configuration

[Help..](#)

Modem Configuration

Parameter	Value	Description
Dial-in Console	enable <input type="checkbox"/>	Enable dial in console access
PPP Server	enable <input type="checkbox"/>	Enable incoming PPP connections
Radius Authentication	enable <input type="checkbox"/>	Radius Authenticate for incoming PPP connections
Rings before answer	1	Number of rings to wait before answering [1-10]
Additional Modem AT Init Codes	L3M0	Any extra AT codes to use when initializing the modem
Country code	Australia ▼	Set modem country code
Speaker Volume	0 ▼	Set modem speaker volume
Speaker Mode	Off ▼	Set modem speaker mode

[Save](#)

Note: Changing the country code will cause the modem to reset. Active connections will be lost.

[Return to Modem Main Menu](#)

Figure 71: Edit Internal Modem Configuration

[Help..](#)

Modem-1 Configuration

Parameter	Value	Description
Dial-in Console	enable <input type="checkbox"/>	Enable dial in console access
PPP Server	enable <input type="checkbox"/>	Enable incoming PPP connections
Radius Authentication	enable <input type="checkbox"/>	Radius Authenticate for incoming PPP connections
Rings before answer	1	Number of rings to wait before answering [1-10]
Additional Modem AT Init Codes		Any extra AT codes to use when initializing the modem

[Save](#)

[Return to Modem-1 Main Menu](#)

Figure 72: Edit External Modem Configuration

These menus allow you to configure modem settings and usage features.

The *Dial-in console* field allows the modem to answer incoming calls and present a login screen in the same way that the console serial port does. The login used for the Dial-in console is the same as that used for SSH and serial console logins.



NOTE

If RADIUS authentication is enabled, the Dial-In Console login will be in the LOGIN group and not in the PPP group. See the section:RADIUS Authentication for details.

The *PPP server* field configures the router to answer incoming modem calls and negotiate a PPP connection to the calling system to provide network access.

The *RADIUS Authentication* field will cause incoming PPP connections to be authenticated against the RADIUS servers configured in the Maintenance menu, RADIUS Authentication sub-menu.

**NOTE**

The Dial-in Console and PPP Server can be enabled at the same time. The router automatically detects whether an incoming call is PPP or console only. If PPP Client mode is active, the router attempts to maintain the PPP link at all times, and hence blocks incoming calls most of the time.

You can enable PPP Client and Dial-in Console and/or PPP Server at the same time as long as PPP Client is not configured to connect at boot-time. In this case, the modem is free to serve the incoming call or to perform the dial-out manually.

Rings before answer controls how many times to let the modem ring before answering a call, if either of Dial-in console or PPP Server is enabled.

Additional Modem AT Init Codes allows you to enter extra AT codes. The router does not verify the AT codes entered here, so take care to ensure that the codes are correct. The following lists the AT codes supported for the **SiemensInternal modem**. For external modems, refer to the external modem documentation.

- Blind Dial
 - X0 - Ignore dialtone/busy signal. Blind dial.
 - X4 - Monitor and report dialtone/busy signal. (default)
- Guard Tone Control*
 - &G0 - Disable guard tone. (default)
 - &G1 - Enable guard tone at 550Hz.
 - &G2 - Enable guard tone at 1800Hz.
- Pulse Dialing Control
 - &P0 - Make/break ratio of 39/61 at 10 pulses/second. (default)
 - &P1 - Make/break ratio of 33/67 at 10 pulses/second.
 - &P2 - Make/break ratio of 39/61 at 20 pulses/second.
 - &P3 - Make/break ratio of 33/67 at 20 pulses/second.
- Compression Control
 - %C0 - Disable data compression negotiation.
 - %C1 - Enable MNP5 compression negotiation.
 - %C2 - Enable V.42bis compression negotiation.
 - %C3 - Enable MNP5 and V.42bis compression negotiation. (default)
- Line Quality Monitoring Control
 - %E0 - Disable line quality monitor and auto-retrain.
 - %E1 - Enable line quality monitor and auto-retrain.
 - %E2 - Enable line quality monitor and fallback/fallforward. (default)
- S Registers
 - S6=X - Wait time for dialtone detection (2-255 seconds) (default=2)
 - S7=X - Wait time for carrier detection (1-255 seconds) (default=50)
 - S8=X - Pause time for comma in dial string (0-255 seconds) (default=2)
 - S9=X - Carrier detect response time (50-255 * .1 seconds) (default=6)
 - S10=X - Loss of carrier to hangup delay (50-255 * .1 seconds) (default=14)
 - S11=X - DTMF tone duration (50-255 * .01 seconds) (default=95)

S29=X - Hook flash dial modifier time (0-255 * .01 seconds) (default=70)

- *Country Code* selects which country's dialing system. If this is not set correctly, the modem might not be able to dial or connect.
- *Speaker Volume* controls the modem speaker loudness.
- *Speaker Mode* controls whether the speaker on the modem is on or off.

Section 3.8.6.2

Modem PPP Client Connections

Figure 73: Modem PPP Client Connections

To edit an existing connection, click the *Edit* link for that connection.

To create a new connection click *Add new* link.

To have the router automatically dial a connection at boot time and keep it always active, select which connection should be used from the drop down list of available connection profiles in the *Connect at boot* list.

Section 3.8.6.3

Dial on Demand Alternate Modem Setting

On systems with two modems, the *Dial On Demand* function can trigger and use an alternate modem to transmit data when the first modem is busy. For instructions on setting the *Dial on Demand* option, see [Section 3.8.6.4, "Modem PPP Client"](#).

When *Dial on demand* is selected on the *PPP Client* page, the *Dial on Demand Setting* table appears on the *PPP Client Connections* page.

[Help..](#)

Modem-1 PPP Client Connections

Connection Name	Action
Headoffice	Edit
	Add new

Parameter	Value	Description
Connect at boot	none	Which client connection to start automatically at boot

Save

Dial on Demand Setting

Target network	Alternate modem	Action
	none	Add

[Return to Modem-1 Main Menu](#)

Figure 74: PPP Client Connections page with Dial On Demand Setting table

The *Target network* is the remote network to which the alternate modem will transmit traffic. Specify the network in dotted-decimal CIDR notation (for example: 192.168.1.0/24).

The *Alternate modem* is the secondary modem to use when the primary modem is busy. Select a modem from the list. When *none* is selected, the traffic is handled by the primary modem only.



NOTE
To allow the alternate modem to handle traffic when the primary modem is busy, the Defaultroute option must be disabled for the primary modem on the PPP Client page. For example, if Modem-2 is set as the alternate modem for Modem-1, Defaultroute must not be selected for Modem-1.

Section 3.8.6.4
Modem PPP Client

[Help..](#)

Modem PPP Client

ppp0

Connection name		Disconnect on idle timeout		Seconds
PPP Username		Password		
Dial type	DTMF	Phonenumber		
Defaultroute	<input checked="" type="checkbox"/>	Use peer DNS	<input checked="" type="checkbox"/>	
Maximum Dial Attempts	0	Dial Interval (1-86400 seconds)	30	(0 means try forever)
Dial on demand	<input type="checkbox"/>	Local IP Address		Remote IP Address

Note: The Maximum Dial Attempts is the number of consecutive connection attempts the modem dial the phone number before it stops. If the Maximum Dial Attempts is 0, it will try forever. Otherwise, the reconnect button will appear in the Modem Main Menu after specified number of consecutive failed connection attempts.
The Local IP Address and Remote IP Address are only required for Dial on Demand. These IP address will be overwritten if the server side gives any of these IP addresses.

Save (Saving will reset ppp link to update settings)

delete

[Return to Modem PPP Client Connections](#)

Figure 75: Configure Modem PPP Client

The *Connection Name* field determines the name that will be used to refer to this connection when choosing which connection to dial automatically at boot, or which connection to use as a backup for another link.

The *Disconnect on idle timeout* field specifies how long an established PPP connection will wait with no data traffic before it disconnects. This option is only valid when the *Dial on demand* option is checked.

The *PPP Username* field determines the user name to use when connecting to the PPP server as specified by its operator.

The *Password* field determines the password to use when connecting to the PPP server.

The *Dial type* field determines the type of dialing system to use on the phone line. Either DTMF (Dual-Tone Multi-Frequency - commonly known as Tone dialing) or Pulse dialing. Almost all phone systems support DTMF, and DTMF is much faster at dialing. DTMF is recommended whenever possible.

The *Phonenumber* field specifies the telephone number to dial to connect to the PPP server.

The *Defaultroute* check box enables automatically setting a default route using this interface whenever it connects. If this is your primary network connection you probably want this option enabled. On systems with two modems where you want to specify a Dial On Demand alternate modem, this option must be disabled for the primary modem. For more information on Dial On Demand alternate modem settings, see [Section 3.8.6.3, "Dial on Demand Alternate Modem Setting"](#).

The *Use peer DNS* check box enables automatically setting the DNS server entries that the PPP server recommends. Enable this option unless you provide your own name servers.

The *Maximum Dial Attempts* field specifies the number of consecutive times that the modem will dial the phone number before it stops attempting to establish a connection. If the number is 0, it will never stop and dial until the connection is established. Otherwise, a Reconnect button will appear in the Modem Main Menu after specified number of consecutive failed connection attempts. For Australian A-tick compliance, a maximum number of 15 attempts is set when the country code is set to Australia when the setting is found to be either 0 or above 15.

The *Dial Interval* field determines how many seconds to wait before re-initiating the link after it terminates.

The *Dial on Demand* field specifies that this connection is designated as "dial on demand", meaning that establishment of the PPP connection is postponed until there is data to be transmitted via the interface. When *Dial on Demand* is enabled, you must specify a *Local IP Address* and a *Remote IP Address*.

On systems with two modems, *Dial On Demand* can also trigger and use the alternate modem to transmit data if the first modem is busy. When *Dial On Demand* is selected, a *Dial On Demand Setting* table appears on the *PPP Client Connections* page. For instructions on configuring the alternate modem settings, see [Section 3.8.6.3, "Dial on Demand Alternate Modem Setting"](#).

The *Local IP Address* and *Remote IP Address* fields specify the local and remote IP addresses, respectively, to use on the PPP connection. Note that one or both of these addresses may be overridden by a remote PPP server when a connection is established.

Section 3.8.6.5

Modem PPP Server

Help...

Modem PPP Server

Server Configuration

ppp0

Server IP address192.168.100.11

Client IP address192.168.100.22

Client Nameserver

Idle timeoutSeconds

Server name

System hostname (ruggedcom.localdomain)

Otherpppin

Proxy ARP

Save (Changes take effect next time user connects)

Username/Password Setting

Username	Password	Action
admin	admin	Delete
		Add

Static Route Setting

Username	Static Route	Action
		Add

Return to Modem Main Menu

Figure 76: Configure Modem PPP Server

The *Server IP address* field specifies the IP address that the router will use for the PPP interface.

The *Client IP address* field specifies the IP address to assign to an incoming PPP client connection.


The *Client Nameserver* field controls which nameserver (if any) the client should use for DNS lookups.

The *Proxy ARP* option makes the router attempt to proxy ARP the remote IP onto a local Ethernet subnet. This requires that the Client IP address be set to an IP that would be valid on one of the Ethernet subnets to which that the router is connected. If this option is selected, other machines on the Ethernet subnet will be able to communicate with the remote PPP system as though it were connected directly to the Ethernet subnet.

The *Idle timeout* field controls how many seconds to wait when there is no traffic on the PPP connection before hanging up the connection. Setting it to 0 or blank will disable the timeout.

The *Server name* fields determine the name used to authenticate the dial-in user during PPP initialization. The default value is "pppin". To use the system hostname (set with the Webmin **System Hostname** link), select the *System hostname* option. To specify a custom value, select the *Other* option and type a value in the text field.

The *Username/Password Setting* table contains a list of locally defined users that are allowed to connect to the PPP server, along with their passwords and a *Delete* button next to each entry to allow its removal. A row of empty Username and Password fields ending with an *Add* button allows the addition of new user entries.



NOTE

If RADIUS authentication is used to authenticate PPP connection requests, the Username/Password Setting table will not be displayed, and PPP user accounts must instead be configured at the RADIUS server. See [Section 4.7, "RADIUS Server Configuration"](#) for details.

Modem PPP Server

75

Each PPP user entry, whether defined on the router or on a RADIUS server, may optionally have a list of subnetworks associated with it. When the corresponding user establishes a PPP connection with the server, the configured static routes are entered into the routing table, making the listed subnetworks available via the connected PPP client.

The *Static Route Setting* table is displayed whether PPP is authenticated locally or via RADIUS. Each row of the table contains a *Username*, a list of *Static Route* entries to be entered when that user establishes a PPP connection, and a *Delete* button next to each entry to allow its removal. An empty row prompting for *Username* and ending with an *Add* button allows the addition of new Static Route entries for that PPP user, via the *Edit Routes...* menu, described below.

Edit Routes for user_one

Route (x.x.x.x/y)

Action

Add

Return to Modem PPP Server Menu

Figure 77: Add Routes for PPP User

The *Route (x.x.x.x/y)* field lists existing static routes for the listed PPP user, with a *Delete* button next to each entry to allow its removal. Static routes are added to the list by entering an IPv4 network address (x.x.x.x) followed by the number of bits in the subnet address (y), and clicking *Add*.

Section 3.8.6.6
Modem Incoming Call Logs

Help..

Incoming Call Logs

Refresh

Date	Time	Event
12/18	17:19:23	mgetty: interim release 1.1.36-Jun15
12/18	17:19:23	check for lockfiles
12/18	17:19:23	locking the line
12/18	17:19:41	mgetty: interim release 1.1.36-Jun15
12/19	07:19:51	checking if modem is still alive
12/19	07:19:51	mdm_send: 'AT' -> OK
12/19	07:19:51	waiting...
12/19	08:19:51	checking if modem is still alive
12/19	08:19:51	mdm_send: 'AT' -> OK
12/19	08:19:51	waiting...
12/19	09:19:51	checking if modem is still alive
12/19	09:19:51	mdm_send: 'AT' -> OK
12/19	09:19:52	waiting...

Refresh

Return to Modem-1 Main Menu

Figure 78: Incoming Call Logs

This page shows the latest log entries for incoming calls. This is mainly useful when trying to debug a problem with establishing incoming connections.

Section 3.8.6.7

Modem PPP Logs

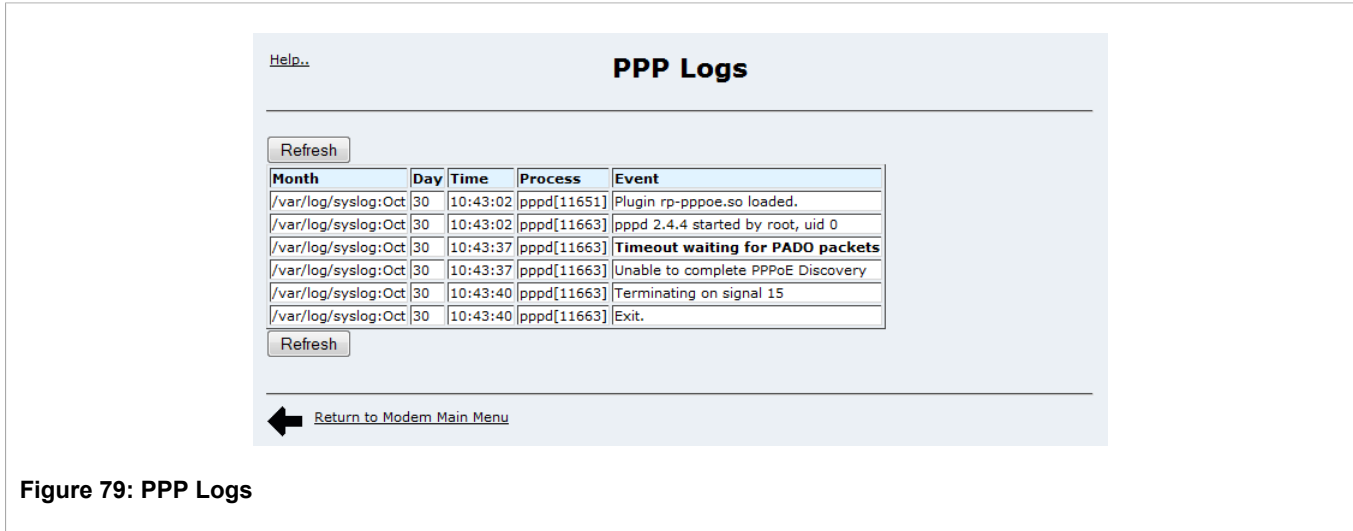


Figure 79: PPP Logs

This page shows the PPP logs. This is mainly useful when trying to debug a PPP connection problem.

Section 3.8.6.8

Modem PPP Connection Logs

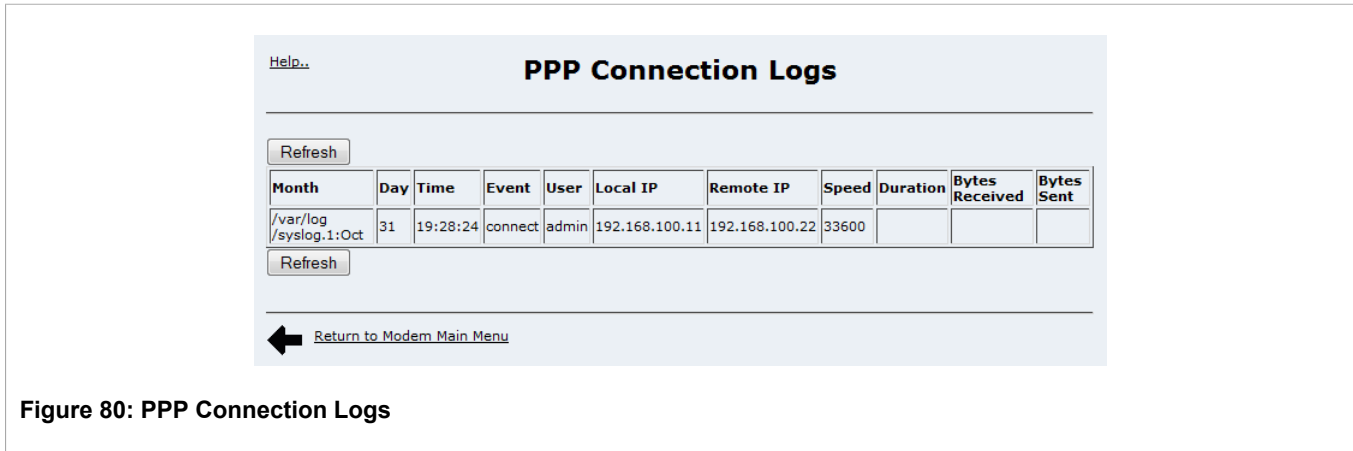


Figure 80: PPP Connection Logs

This page shows a list of PPP connections. It shows who connected, when they connected and disconnected, the connection speed, and session traffic.

Section 3.9

Configuring PPP and the Cellular Modem

This section familiarizes the user with:

- Configuring the cellular modem
- Configuring the PPP client
- Viewing connection status

The device may be equipped with an internal cellular modem instead of the land-line modem or the serial card described in the preceding section. PPP (the Point-to-Point Protocol) is used to establish an IP network connection over a cellular radio modem link.

Depending on local cellular network availability, one of three cellular modem types may be ordered:

- Edge/GPRS
- CDMA/EV-DO
- HSPA

Section 3.9.1

PPP Interface

When a PPP connection is established, a network interface is created in the system. The interface name for both internal and external modem connections is *ppp10*. Refer to this interface name when configuring firewall rules.

Section 3.9.2

Authentication, IP Addressing and DNS Servers

In contrast to the configuration for land-line modems described in the preceding section, username and password might not be required for some cellular data service providers. If username and password is not required, you can enter none in the username and password fields of the GUI, or leave them blank. If authentication is required by the cellular data service provider, again PPP authentication will automatically use PAP or CHAP. Your service provider will provide you with a username and password along with an Access Provider Name (APN), which must be entered in the GUI.

The authentication process will provide a local IP address for use on the PPP interface and optionally the addresses of the DNS servers and a default gateway address to use. You should generally use these addresses unless you need to provide your own.

The PPP interface's IP address, obtained from the PPP server, can be either a dynamic or a static IP address. Firewall configuration should be performed as is appropriate.



NOTE

A PPP Client Connection for the cellular modem may be configured to connect at boot time.

Section 3.9.3

LED Designations

ROX dedicates two LEDs to indicate cellular modem status:

The leftmost LED of the bottom row (LED #29) is the "Line" LED.

- Solid Green indicates that a PPP link has been established.
- Flashing green indicates that PPP link negotiation is in progress.
- Off indicates that the cellular modem is active but a connection to the wireless network has not yet been established.
- RED indicates that cellular modem is not currently operating.

The leftmost LED of second bottom row (LED #25) is the cellular modem's Activity LED:

- Off means that there is no data traffic on the cellular modem.
- Flashing means that there is data traffic on the cellular modem connection.

Section 3.9.4

PPP Cellular Modem Configuration

The top-level configuration interface menu for the cellular modem and PPP is accessible by clicking on the *Cellular Modem* link in the *Networking* folder. This menu allows you to display and configure the cellular modem interface.

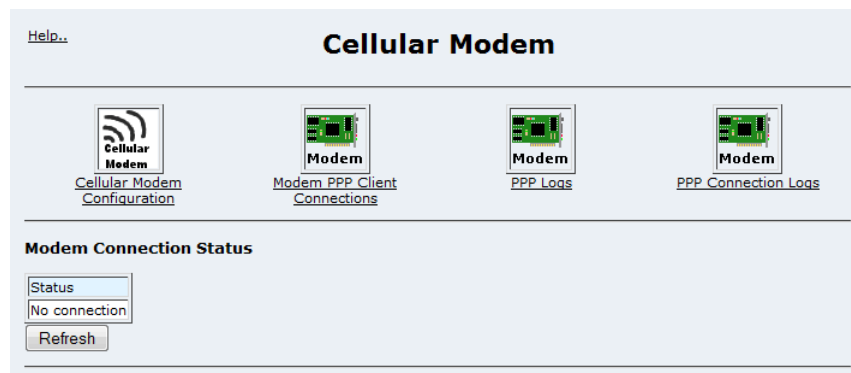


Figure 81: Cellular Modem Interface

If the installed modem is a CDMA/EV-DO type, and has not yet been activated for use on the cellular network, the following top-level menu will be seen instead, offering two different methods for activating the modem for use on the cellular network (see [Section 3.9.4.1, "Over-The-Air Account Activation"](#) and [Section 3.9.4.2, "Manual Account Activation"](#), below).




Figure 82: Cellular Modem Interface (CDMA modem not yet activated)

Prior to use, a CDMA-type cellular modem must be activated for use on a particular provider's network. Once the activation process has been completed, the modem will be able to connect to the network without further intervention. Two account activation methods are provided by ROX: "OTA (Over-the-Air)" and "Manual". Both activation methods are described in this section.

Section 3.9.4.1

Over-The-Air Account Activation

ROX supports the OTASP (Over-the-Air Service Provisioning) mechanism offered by most CDMA cellular service providers for provisioning cellular end stations for use on their networks. Using this method, the service provider, or carrier, supplies an OTASP dial string which ROX can use to contact the cellular network via the modem. During this OTASP call, the carrier authorizes and configures the modem for use on its network. Note that an OTASP dial string typically begins with "**228".

The screenshot shows a web interface titled "Over The Air Account Activation". At the top left is a "Help.." link. Below the title is a table with three columns: "Parameter", "Value", and "Description". The table has one row: "Activation Dial string" with value "*22899" and description "Dial string to automatically activate account Over The Air". Below the table is an "Activate" button. Below that is a section titled "Modem Status" containing a table with the following data:

Type of Modem	CDMA 1xRTT/EV-DO
Firmware version	p2410701,51863
Electronic Serial Number (ESN)	0x608ABD96
Received Signal Strength/EcIo	(0.0%), -125 dBm/None dB
Network Carrier ID	Verizon
Network technology currently in use	No Service
Phone number	9547890192
Activation status	Modem is activated

Below the "Modem Status" table are "Refresh" and "Reset" buttons. At the bottom left is a back arrow icon and the link "Return to Modem Main Menu".

Figure 83: Over The Air Account Activation

1. First, establish an account with the help of a service representative of the cellular network provider.
2. Enter the OTASP dial string supplied in the *Activation Dial string* field and click *Activate*.
3. The *Activation status* field will display: "Activation is in progress... Please wait." until a success or failure is detected.

4. Upon successful activation, the *Activation status* field will automatically change to display: "Activation successful". If it displays "Activation Failed", please verify the activation dial string or contact the network provider's service personnel.

Section 3.9.4.2

Manual Account Activation

If the carrier does not support Over-the-Air Service Provisioning, the cellular modem must be programmed via the *Manual Account Activation* form using settings supplied by the carrier's service personnel:

Parameter	Value	Description
Activation code		Master Subsidy Lock code provided by service provider
Phone number	9547890192	Mobile Directory Number provided by service provider
MIN	9545795619	Mobile Identification Number provided by service provider
System ID	4152	System ID provided by service provider
Network ID	65535	Network ID provided by service provider

Modem Status

Type of Modem	CDMA 1xRTT/EV-DO
Firmware version	p2410701,51863
Electronic Serial Number (ESN)	0x608ABD96
Received Signal Strength/EcIo	(0.0%), -125 dBm/None dB
Network Carrier ID	Verizon
Network technology currently in use	No Service
Phone number	9547890192
Activation status	Modem is activated


 [Return to Modem Main Menu](#)

Figure 84: Manual Account Activation

1. First, establish an account with a service representative of the cellular network provider. You will need the following settings in order to activate your modem. Note that not all of these parameters are required by all network providers:
 - *Activation code*, also known as a "subsidy lock".
 - *Phone Number*, or MDN (Mobile Directory Number).
 - *MIN* (Mobile Identification Number), often the same as the Phone Number.
 - *System ID*, or Home System ID.
 - *Network ID*
2. Click *Activate*.
3. The *Activation status* field will display: "Activation is in progress... Please wait." until a success or failure is detected.

4. Upon successful activation, the *Activation status* field will automatically change to display: "Activation successful". If it displays "Activation Failed", please verify all activation settings or contact the network provider's service personnel.

Section 3.9.4.3

Cellular Modem Configuration

The Cellular Modem Configuration menu provides information that is necessary to establish cellular network service, and allows the configuration of parameters that are necessary for the modem to access the cellular network.

Necessary parameters are configured in the top part of the screen, and modem information and status are displayed at the bottom, under "Modem Status".

Parameter	Value	Description
Dial string	#777	Dial string to connect to Wireless network

[Save](#) [Re-activate modem](#)

Modem Status

Last updated on Fri Oct 18 16:42:24 2013

Type of Modem	CDMA 1xRTT/EV-DO
Firmware version	p2410701,51863
Electronic Serial Number (ESN)	0x608ABD96
Received Signal Strength/EcIo	((0.0%), -125 dBm/None dB)
Network Carrier ID	Verizon
Network technology currently in use	No Service
Phone number	9547890192

[Refresh](#) [Reset](#)

[Return to Modem Main Menu](#)

Figure 85: Cellular Modem Configuration (with an Edge/GPRS modem)

Modem Configuration

The *Access Point Name* (APN) is necessary only on GPRS networks (Edge or HSPA). It is the name of the cellular network access point which provides a gateway to the Internet. This information will be provided by the wireless network when you register for data service. This field is not used for CDMA modems.

The *Dial string* is special command to be sent by the cellular modem to the cellular network to establish a data connection. For example, for GSM/GPRS networks, this is typically: *99***1#. This command will depend on the wireless network. Please consult the wireless network operator for the correct dial string command for data service. A regular telephone number is usually not required to connect to a GSM/GPRS network.

Modem Status

The Modem Status section displays information about and the current status of the cellular modem installed in the ROX.



NOTE

In the case of Edge/GPRS modems, cellular modem status information is only collected if a PPP link is not currently active on the modem. If the Refresh button is clicked while a PPP link is active, the information displayed is that obtained prior to the establishment of the PPP link. In order to obtain

current modem status information, an active PPP link on the cellular modem, must be deactivated, i.e., by selecting "none" for the Modem PPP Client Connection/Connection to establish at boot time.

For EVDO/CDMA and HSPA/GPRS modems, the received signal strength can be obtained whenever the Refresh button is clicked.

The fields and format of the *Modem Status* display varies among installed cellular modem types. The fields displayed for the Edge/GPRS modem (see above) are as follows:

Edge/GPRS Modem Status

Type of Modem lists the cellular network standards supported by the modem currently installed. This information may be required by the cellular network provider in order to determine optimal compatibility between the installed cellular modem and the network.

Firmware version displays firmware revision information for the installed cellular modem.

Received Signal Strength indicates the signal level received by the cellular modem from the cell site.

Network Operator displays the identity of the wireless network provider to which the cellular modem is currently connected.

Enhanced Network Operator displays the name of Mobile Virtual Network (MVN) that the cellular modem is currently connected to.

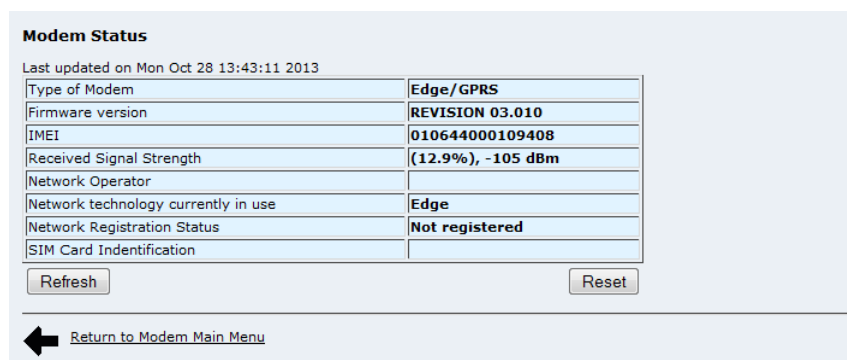
Network Registration Status displays the current registration status of the cellular modem with respect to the cellular network. Possible values are:

- Registered, home
- Registered, roaming
- Unregistered

SIM Card Identification displays the ID of the SIM card currently installed in the cellular modem.

HSPA/GPRS Modem Status

The HSPA/GPRS cellular modem lists mostly the same information as the Edge/GPRS modem, with some differences:




Last updated on Mon Oct 28 13:43:11 2013	
Type of Modem	Edge/GPRS
Firmware version	REVISION 03.010
IMEI	010644000109408
Received Signal Strength	(12.9%), -105 dBm
Network Operator	
Network technology currently in use	Edge
Network Registration Status	Not registered
SIM Card Identification	
<input type="button" value="Refresh"/> <input type="button" value="Reset"/>	
 Return to Modem Main Menu	

Figure 86: Cellular Modem Status (HSPA/GPRS)

The *IMEI* (International Mobile Equipment Identity) is a numeric identifier unique to the cellular modem card.

The *Network technology currently in use* field displays which network technology, out of the ones listed as being supported in the *Type of Modem* field, is currently in use between the modem and the network.

CDMA Modem Status

The CDMA cellular modem lists several of the same fields with some additional differences:

Modem Status
Last updated on Fri Oct 18 16:42:24 2013

Type of Modem	CDMA 1xRTT/EV-DO
Firmware version	p2410701,51863
Electronic Serial Number (ESN)	0x608ABD96
Received Signal Strength/EcIo	(0.0%), -125 dBm/None dB
Network Carrier ID	Verizon
Network technology currently in use	No Service
Phone number	9547890192

Refresh Reset

[Return to Modem Main Menu](#)

Figure 87: Cellular Modem Status (CDMA)

The *Electronic Serial Number (ESN)* is a numeric identifier unique to the cellular modem card. This corresponds to the IMEI for GSM networks.

Network Carrier ID displays the identity of the wireless network provider for which the cellular modem is currently configured.

Phone Number displays the cellular telephone number associated with the account created to provide service for the modem.

Section 3.9.4.4

Modem PPP Client Connections

[Help..](#) **Modem PPP Client Connections**

Connection Name	Action
HeadOffice	Edit
	Add new

Parameter	Value	Description
Connect at boot	HeadOffice ▼	Which dient connection to start automatically at boot

Save

[Return to Modem Main Menu](#)

Figure 88: Modem PPP Client Connections

To edit an existing connection, click the *Edit* link for that connection.

To create a new connection click *Add new* link.

To have the router automatically dial a connection at boot time and keep it always active, select which connection should be used from the drop down list of available connection profiles in the *Connect at boot* list.

Section 3.9.4.5

Modem PPP Client

[Help..](#)

Modem PPP Client

ppp10

Connection name

PPP Username Password

Defaultroute ☒ Use peer DNS ☒

Maximum Dial Attempts (0 means try forever) Dial Interval (1-86400 seconds)

Note: The Maximum Dial Attempts is the number of consecutive connection attempts the modem dial the phone number before it stops. If the Maximum Dial Attempts is 0, it will try forever. Otherwise, the reconnect button will appear in the Modem Main Menu after specified number of consecutive failed connection attempts.

(Saving will reset ppp link to update settings)


 [Return to Modem PPP Client Connections](#)

Figure 89: Configure Modem PPP Client

The *Connection Name* field determines what name will be used to refer to this connection when choosing which connection to dial automatically at boot, or which connection to use as a backup for another link.

The *PPP Username* field determines the user name to use when connecting to the PPP server as specified by your network provider. If the username is not required, you can enter "none" in this field, or leave it blank.

The *Password* field determines the password to use when connecting to the PPP server. If the password is not required, you can enter "none" in this field, or leave it blank.

The *Default Route* check box enables automatically setting a default route using this interface whenever it connects. If this is your primary connection you probably want this option enabled.

The *Use peer DNS* check box enables automatically setting the DNS server entries that the PPP server recommends. Enable this option unless you provide your own name servers.

The *Maximum Dial Attempts* field specifies number of consecutive connection attempts the modem dial the phone number before it stops. If the number is 0, it will never stop and dial until the connection is established. Otherwise, a Reconnect button will appear in the Modem Main Menu after specified number of consecutive failed connection attempts.

The *Dial Interval* field determines how many seconds to wait before re-initiating the link after it terminates.

Section 3.9.4.6

PPP Logs, PPP Connection Logs

Refer to [Section 3.8, "Configuring PPP and the Embedded Modem"](#) for information.

Section 3.9.4.7

Current Route and Interfaces Table

Refer to [Section 3.8, "Configuring PPP and the Embedded Modem"](#) for information.

Section 3.10

Configuring Serial Protocols

This section familiarizes the user with:

- RawSockets Applications
- TCP Modbus Server Applications
- DNP (Distributed Network Protocol)
- Configuring Serial ports for RawSocket
- Viewing Serial Port and TCP Connection status and statistics
- Resetting Serial ports
- Tracing Serial Port activity

Section 3.10.1

Serial IP Port Features

The ROX Serial Server provides the following features for forwarding serial traffic over IP:

- Raw Socket Protocol - a means to transport streams of characters from one serial port on the router to a specified remote IP address and TCP port
- Four independent serial ports per product
- Bit rates of 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200 or 230400 bps.
- Supports RS232, RS422 and RS485 party line operation.
- XON/XOFF flow control.
- Supports a point-to-point connection mode and a broadcast connection mode in which up to 32 remote servers may connect to a central server.
- TCP/IP incoming, outgoing or both incoming/outgoing connections mode, configurable local and remote TCP port numbers.
- Packetize and send data on a full packet, a specific character or upon a timeout.
- Supports a "turnaround" time to enforce minimum times between messages sent out the serial port.
- Debugging facilities including connection tracing and statistics

Section 3.10.1.1

LED Designations

The Quad TriplePlay Serial card includes transmit and receive LEDs. The transmit LED is leftmost when the card is in the top slot and will light while characters are being transmitted. The receive LED is rightmost when the card is in the top slot and will light while characters are being received.

Serial port numbers are as described by the "SER" labels as shown in the home page chassis diagram.

Section 3.10.2

Serial Protocols Applications

Section 3.10.2.1

Character Encapsulation

Character encapsulation is used any time a stream of characters must be reliably transported across a network.

The character streams can be created by any serial device. The baud rates supported at either server need not be the same. If configured, the router will obey XON/XOFF flow control from the end devices.

One of the routers is configured to listen to TCP connection requests on a specific TCP port number. The other server is configured to connect to its peer on the listening port number. ROX will attempt to connect periodically if the first attempt fails and after a connection is broken.

ROX can be used to connect to any device supporting TCP (e.g. a host computer's TCP stack or a serial application on a host using port redirection software).

Section 3.10.2.2

RTU Polling

The following applies to a variety of RTU protocols besides ModBus RTU, including ModBus ASCII and DNP.

The remote router communicates with host equipment through:

- native TCP connections,
- another device via a serial port or
- a port redirection package which Supports TCP.

If a ROX is used at the host end, it will wait for a request from the host, encapsulate it in a TCP message and send it to the remote side. There, the remote ROX will forward the original request to the RTU. When the RTU replies the ROX will forward the encapsulated reply back to the host end.

ModBus does not employ flow-control so XON/XOFF should not be configured.

ROX maintains configurable timers to help decide replies and requests are complete and to handle special messages such as broadcasts.

ROX will also handle the process of line-turnaround when used with RS485.

Section 3.10.2.3

Broadcast RTU Polling

Broadcast polling allows a single host connected ROX to "fan-out" a polling stream to a number of remote RTUs.

The host equipment connects via a serial port to a ROX. Up to 32 remote devices may connect to the host server via the network.

Initially, the remote servers will place connections to the host server. The host server in turn is configured to accept the required number of incoming connections.

The host will sequentially poll each RTU. Each poll received by the host server is forwarded (i.e. broadcast) to all of the remote servers. All RTUs will receive the request and the appropriate RTU will issue a reply. The reply is returned to the host server, where it is forwarded to the host.

Section 3.10.3

Serial Protocols Concepts and Issues

Section 3.10.3.1

Host and Remote Roles

ROX can either initiate or accept a TCP connection for serial encapsulation. It can establish a connection from field ("remote") equipment to the central site ("host") equipment, vice versa, or bi-directionally.

Configure ROX at the host end to connect to the remote when:

- The host end uses a port redirector that must make the connection.
- The host end is only occasionally activated and will make the connection when it becomes active.
- A host end firewall requires the connection to be made outbound.

Connect from the remote to the host if the host end accepts multiple connections from remote ends in order to implement broadcast polling.

Connect from each side to other if both sides support this functionality.

Section 3.10.3.2

Use of Port Redirectors

Port redirectors are PC packages that emulate the existence of communications ports. The redirector software creates and makes available these "virtual" COM ports, providing access to the network via a TCP connection.

When a software package uses one of the virtual COM ports, a TCP connection is placed to a remote IP address and TCP port that has been programmed into the redirector. Some redirectors also offer the ability to receive connections.

Section 3.10.3.3

Message Packetization

The server buffers received characters into packets in order to improve network efficiency and demarcate messages.

The server uses three methods to decide when to packetize and forward the buffered characters to the network:

- Packetize on Specific Character,
- Packetize on timeout and
- Packetize on full packet.

If configured to packetize on a specific character, the server will examine each received character and will packetize and forward upon receiving the specific character. The character is usually a <CR> or an <LF> character but may be any ASCII character.

If configured to packetize on a timeout, the server will wait for a configurable time after receiving a character before packetizing and forwarding. If another character arrives during the waiting interval, the timer is restarted. This method allows characters transmitted as part of an entire message to be forwarded to network in a single packet, when the timer expires after receiving the very last character of the message. This is usually the only packetizer selected when supporting ModBus communications.

Finally, the server will always packetize and forward on a full packet, i.e. when the number of characters fills its communications buffer (1024 bytes).

Section 3.10.3.4

Use of Turnaround Delays

Some RTU protocols (such as ModBus) use the concept of a turnaround delay. When the host sends a message (such as a broadcast) that does not invoke an RTU response, it waits a turnaround delay time. This delay ensures that the RTU has time to process the broadcast message before it has to receive the next poll.

When polling is performed, network delays may cause the broadcast and next poll to arrive at the remote server at the same time. Configuring a turnaround delay will enforce a minimum separation time between each message sent out the port. Note that turnaround delays do not need to be configured at the host computer side and may be disabled there.

Section 3.10.4

TcpModBus Server Application

The TcpModbus Server application is used to transport Modbus requests and responses across IP networks.

The source of the polls is a Modbus "master", a host computer that issues the polls over a serial line.

A TcpModbus Client application, such as that implemented by the RuggedServer accepts Modbus polls on a serial line from a master and determines the address of the corresponding RTU. The client then encapsulates the message in TCP and forwards the frame to a Server Gateway or native TcpModbus RTU. Returning responses are stripped of their TCP headers and issued to the master.

The TcpModbus Server application accepts TCP encapsulated modbus messages from Client Gateways and native masters. After removing the TCP headers the messages are issued to the RTU. Responses are TCP encapsulated and returned to the originator.

A "native" TcpModbus master is one that can encapsulate the Modbus polls in TCP and directly issue them to the network.

Section 3.10.4.1

Local Routing at the Server Gateway

The Server Gateway supports up to 32 RTUs on any of its four ports. When a request for a specific RTU arrives the server will route it to the correct port.

Section 3.10.4.2

MultiMaster Capability

It is possible for multiple masters to simultaneously issue requests for the same RTU. The Server Gateway will queue the requests and deliver them to the RTU in turn. This "multimaster" capability allows widely distributed masters to configure and extract information from the RTU.

Section 3.10.5

TcpModbus Concepts and Issues

Section 3.10.5.1

Host and Remote Roles

Client gateways (such as that implemented by the ROX) always make the TCP connection to the Server Gateway. The Server Gateway can only accept a connection.

Section 3.10.5.2

Port Numbers

The TCP port number dedicated to Modbus use is port 502. The Server Gateway can also be configured to accept a connection on a configurable port number. This auxiliary port can be used by masters that do not support port 502.



NOTE

The Server Gateway is capable of creating only one connection on the specified auxiliary port, whereas when Modbus is configured to use the default port, 502, it may connect to multiple RTUs.

Section 3.10.5.3

Retransmissions

The Server Gateway offers the ability to resend a request to an RTU should the RTU receive the request in error or the Server Gateway receives the RTU response in error.

The decision to use retransmissions, and the number to use depends upon factors such as:

- The probability of a line failure
- The number of RTUs and amount of traffic on the port
- The cost of retransmitting the request from the server vs. timing-out and retransmitting at the master. This cost is affected by the speed of the ports and of the network.

Section 3.10.5.4

ModBus Exception Handling

If the Server Gateway receives a request for an unconfigured RTU, it will respond to the originator with a special message called an exception (type 10). A type 11 exception is returned by the server if the RTU fails to respond to requests.

Native TcpModbus polling packages will want to receive these messages. Immediate indication of a failure can accelerate recovery sequences and reduce the need for long time-outs.

Section 3.10.5.5

TcpModbus Performance Determinants

The following description provides some insight into the possible sources of delay and error in an end-to-end TcpModbus exchange.

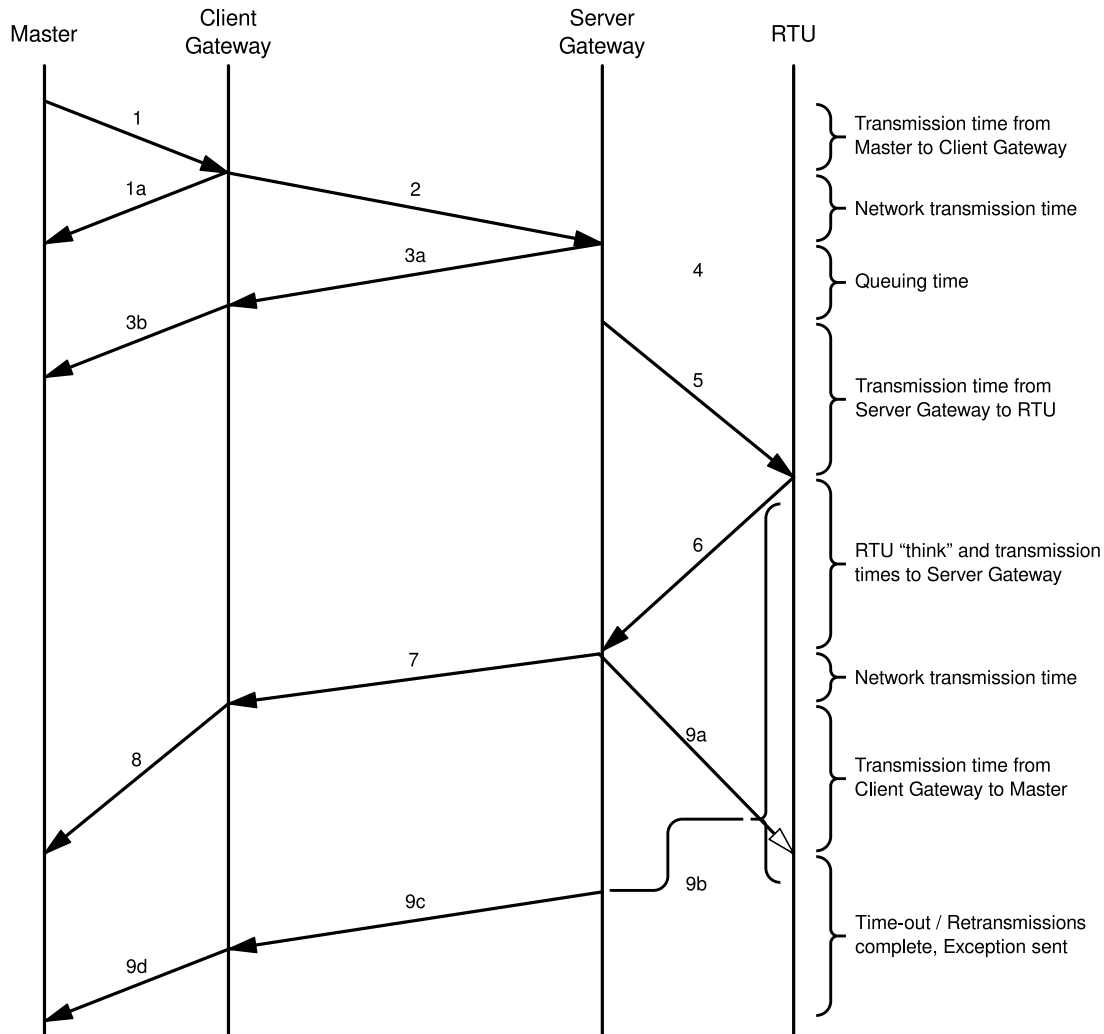


Figure 90: Sources of Delay and Error in an End to End Exchange

In step 1 the master issues a request to the Client Gateway. If the Client Gateway validates the message it will forward it to the network as step 2.

The Client Gateway can respond immediately in certain circumstances, as shown in step 1a. When the Client Gateway does not have a configuration for the specified RTU it will respond to the master with an exception using TcpModbus exception code 11 ("No Path"). When the Client Gateway has a configured RTU but the connection is not yet active it will respond to the master with an exception using TcpModbus exception code 10 ("No Response"). If the forwarding of TcpModbus exceptions is disabled, the client will not issue any responses.

Steps 3a and 3b represents the possibility that the Server Gateway does not have configuration for the specified RTU. The Server Gateway will always respond with a type 10 ("No Path") in step 3a, which the client will forward in step 3b.

Step 4 represents the possibility of queuing delay. The Server Gateway may have to queue the request while it awaits the response to a previous request. The worst case occurs when a number of requests are queued for an RTU that has gone offline, especially when the server is programmed to retry the request upon failure.

Steps 5-8 represent the case where the request is responded to by the RTU and is forwarded successfully to the master. It includes the "think time" for the RTU to process the request and build the response.

Step 9a represents the possibility that the RTU is offline, the RTU receives the request in error or that the Server Gateway receives the RTU response in error. If the Server Gateway does not retry the request, it will issue an exception to the originator.

Section 3.10.5.6

A Worked Example

A network is constructed with two Masters and 48 RTUs on four Server Gateways. Each of the Master is connected to a Client Gateway with a 115.2 Kbps line. The RTUs are restricted to 9600 bps lines. The network is Ethernet based and introduces an on average 3 ms of latency. Analysis of traces of the remote sites has determined that the min/max RTU think times were found to be 10/100 ms. What time-out should be used by the Master?

The maximum sized Modbus message is 256 bytes in length. This leads to a transmission time of about 25 ms at the Master and 250 ms at the RTU. Under ideal circumstances the maximum round trip time is given by: 25 ms (Master>client) + 3 ms (network delay) + 250 ms (server>RTU) + 100 ms (Think time) + 250 ms (RTU>server) + 3 ms (network delay) + 25 ms (client>Master). This delay totals about 650 ms.

Contrast this delay with that of a "quick" operation such as reading a single register. Both request and response are less than 10 bytes in length and complete (for this example) in 1 and 10 ms at the client and server. Assuming the RTU responds quickly, the total latency will approach 35 ms.

It is also necessary to take account such factors as the possibility of line errors and collisions between masters at the server.

The server may be configured to recover from a line error by retransmitting the request. Given a maximum frame transmission time of 250 ms and an RTU latency of 100 ms, it would be wise to budget 350 ms for each attempt to send to the RTU. Configuring a single retransmission would increase the end-to-end delay from about 650 ms to about 1000 ms.

The server can already be busy sending a request when the request of our example arrives. Using the figures from the above paragraph, the server being busy would increase the end-to-end delay from 1000 to 1350 ms.

The preceding analysis suggests that the Master should time-out at some time after 1350 ms from the start of transmission.

Section 3.10.6

DNP (Distributed Network Protocol)

ROX supports DNP 3.0, commonly used by utilities in process automation systems. DNP3 protocol messages specify source and destination addresses. A destination address specifies which device should process the data, and the source address specifies which device sent the message. Having both destination and source addresses satisfies at least one requirement for peer-to-peer communication since the receiver knows where to direct a response. Each device supporting the DNP protocol must have a unique address within the collection of devices sending and receiving DNP messages.

Section 3.10.6.1

Address Learning for DNP

ROX implements both local and remote address learning for DNP.

A local Device Address Table is populated with DNP Addresses learned for local and remote DNP devices. Each DNP address is associated with either a local serial port or a remote IP address.

When a message with an unknown DNP source address is received on a local serial port, the DNP source address and serial port number are entered into the Device Address Table. When a message with an unknown DNP source address is received from the IP network, on the IP interface that is configured as the DNP learning interface, the DNP source address and the IP address of the sender are entered into the Device Address Table.

When a message with an unknown DNP destination address is received on a local serial port, the message is sent in a UDP broadcast out the network interface configured as the DNP learning interface. When a message with an unknown DNP destination address is received from the IP network, it is sent to all local serial ports configured as DNP ports.

UDP transport is used during the DNP address learning phase.

All learned addresses will be kept in the Device Address Table, which is saved in non-volatile memory, which makes it unnecessary to repeat the DNP address learning process across a ROX reboot or accidental power loss.

An aging timer is maintained per DNP address in the table, and is reset whenever a DNP message is sent to or received for the specified address.

This learning facility makes it possible to configure the DNP3 protocol with a minimum number of parameters: a TCP/UDP port number, a learning network interface and an aging timer.

Section 3.10.6.2

DNP Broadcast Messages

DNP addresses 65521 through 65535 are reserved as DNP3 broadcast addresses. ROX supports DNP3 broadcast messages. DNP broadcast messages received on local serial ports are transmitted to all IP Addresses in the DNP Device Address Table (whether learned or statically configured). When a DNP broadcast message is received from the IP network, it is transmitted on all local serial ports configured as DNP ports.

Section 3.10.7

Serial Protocols Configuration

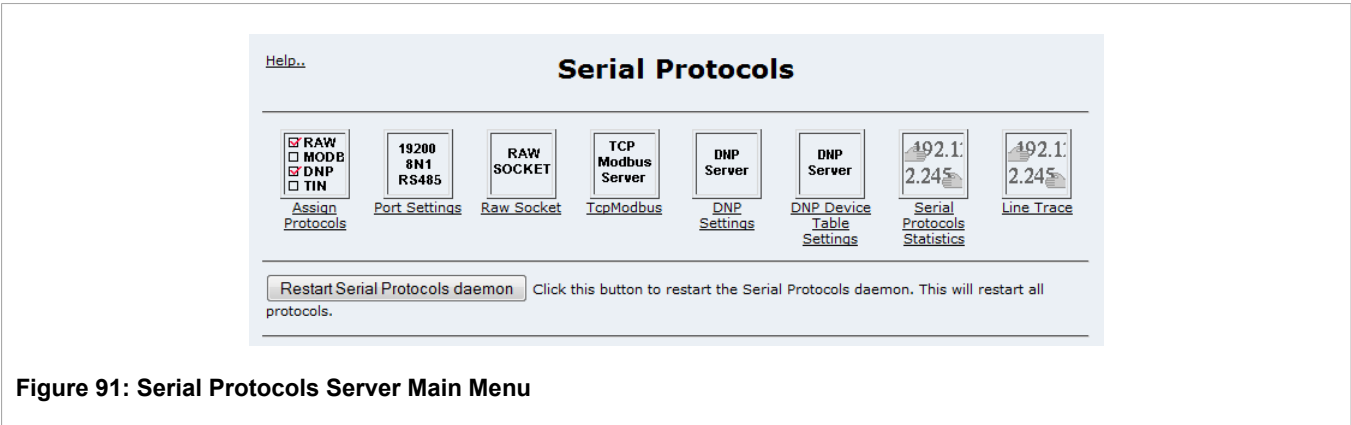


Figure 91: Serial Protocols Server Main Menu

Note that the Serial Protocols server is disabled by default and may be enabled via the *Bootup and Shutdown* menu, under the *System* folder of the main Webmin menu.

The *Assign Protocols* menu assigns a serial protocol to one of your serial ports.

The *Port Settings* menu configures the serial port and its electrical protocol.

If any of your serial ports are configured as *RawSocket* protocol, this menu will configure them.

The *Serial Protocols Statistics* menu will show you the status and statistics for any established sessions.

The *Line Trace* menu will provide a line activity trace for the serial ports.

Section 3.10.7.1

Assign Protocols Menu

[Help..](#)

Assign Protocols

Assigning a protocol to a port will make it available for configuration via a menu in the main page.

Port	Type
1	rawsocket
2	dnp
3	tcpmodbus
4	none

Save Changes

[Return to Serial Protocols](#)

Figure 92: Assign Protocols Menu

This menu associates a protocol with a serial port. Unused ports should be left associated with "none". Changing an association will immediately close the calls of the old protocol.

Section 3.10.7.2

Port Settings Menu

[Help..](#)

Port Settings

Note that all changes are made immediately.

Port	Speed	DataBits	Parity	StopBits	Flow Control	Type	Current Protocol
1	19200	8	NONE	1	NONE	RS232	rawsocket
2	9600	8	NONE	1	NONE	RS232	dnp
3	9600	8	NONE	1	NONE	RS232	tcpmodbus
4	9600	8	NONE	1	NONE	RS232	none

Save Changes

[Return to Serial Protocols](#)

Figure 93: Port Settings Menu

This menu configures the serial settings and electrical protocol associated with a serial port. Changes are made immediately.

Section 3.10.7.3

RawSocket Menu

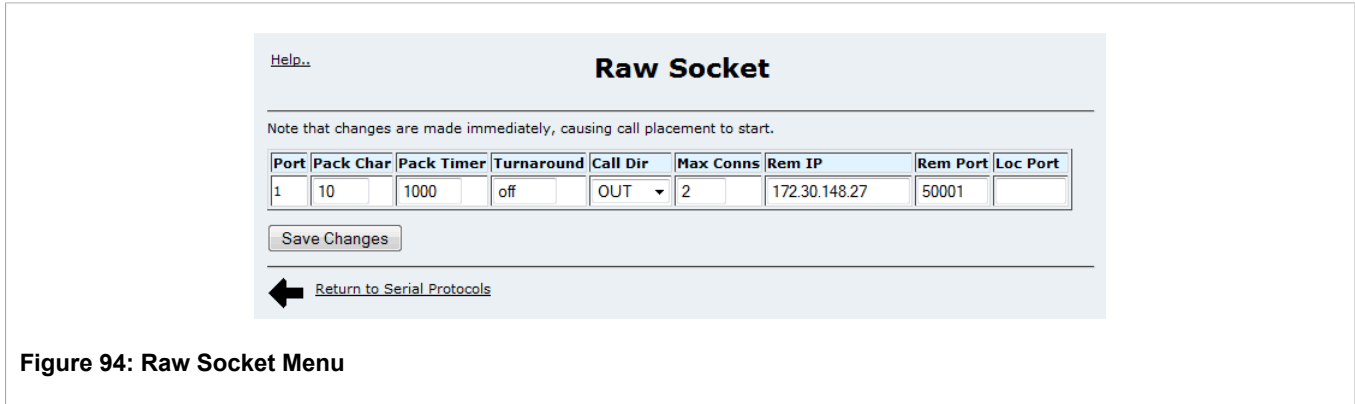


Figure 94: Raw Socket Menu

- This menu configures the Raw Socket settings for each port. Changes are made immediately.
- The *Pack Char* field configures the numeric value of the ASCII character which will force forwarding of accumulated data to the network. The Pack Char must be between 0 and 255 inclusive or the value off. If configured off, accumulated data will be forwarded based upon the packetization timeout parameter.
- The *Pack Timer* field configures the delay from the last received character until when data is forwarded. The Pack Timer must be between 5 and 1000 milliseconds inclusive.
- The *Turnaround* timer field controls the amount of delay (if any) to insert between the transmissions of individual messages out the serial port. The Pack Timer must be between 1 and 1000 milliseconds inclusive, of off.
- The *Call Dir* field configures whether to accept an incoming connection, place an outgoing connection or do both.
- The *Max Conns* field configures the maximum number of incoming connections to permit when the call direction is incoming.
- The *Remote IP* field configures the address used when placing an outgoing connection.
- The *Remote Port* field selects the TCP destination port used in outgoing connections.
- The *Local Port* field selects the local TCP port to use to accept incoming connections.

Section 3.10.7.4

TcpModBus Menu

This menu configures the TcpModbus settings for each port. Changes are made immediately.

Figure 95: TcpModbus Menu

The *Response Timer* field configures the maximum time from the last transmitted character of the outgoing poll until the first character of the response. If the RTU does not respond in this time the poll will have been considered failed. The Response Timer must be between 50 and 1000 milliseconds inclusive.

The *Pack Timer* field configures the maximum allowable time to wait for a response to a Modbus request to complete once it has started. The Pack Timer must be between 200 and 1000 milliseconds inclusive.



NOTE

The Modbus specification states the minimum time is about 640 character times at baud rates below 19200 Kbps and 256 char times + 192 ms at baud rates above 19200 Kbps. You may specify a larger value if you think your RTU will take longer to complete transmission than the calculated time.

The *Turnaround* field configures the amount of delay (if any) to insert after the transmissions of Modbus broadcast messages out the serial port. The Turnaround must be between 1 and 1000 milliseconds inclusive, or off.

The *Retransmits* field configures the number of times to retransmit the request to the RTU before giving up, should the original attempt fail.

The *Max Conns* field configures the maximum number of incoming connections.

The *Local Port* field may be used to specify an alternate local TCP port number. If this field is configured, a single connection (per serial port) may be made to this alternate port number. Note that TCP Modbus uses a default local port number of 502. There is no limit imposed on the number of connections to the default TCP port.

Section 3.10.7.5

DNP Menu

Figure 96: DNP Settings

The *Address Learning* field may be set to 'Disabled' or to an Ethernet interface name.

The *Aging Timer* sets the length of time which a learned DNP device in the Device Address Table may go without any DNP communication before it is removed from the table.

The *Max Conns* field configures the maximum number of incoming DNP connections.

The *Loc Port* field configures the TCP/UDP port number on which DNP protocol listens. DNP devices normally use port 20000.

The *Device Address field* configures a DNP device address, whether local or remote. The address may be that of a DNP device connected to a local serial port or one available via the serial port of a remote IP host. The address may be in the range 1 to 65520. Note that both local and remote serial ports must be properly configured.

The *Rem IP* field configures the IP address of the remote host that provides a connection to the DNP device with the configured address.

[Help..](#)

DNP Device Table Settings

Device Address	Rem IP	Port
<input type="text"/>	<input type="text"/>	<input type="text"/>

[Return to Serial Protocols](#)

Device Address Table

Device Address	Rem IP	Port
12	3.3.3.3	Unknown
15	0.0.0.0	2

Figure 97: DNP Device Table Settings

The *Port* field configures the serial port to which the DNP device is attached. If the entry is for a remote DNP device, i.e. the DNP device is attached to the serial port of remote IP host, the value of this parameter is 'Unknown'.

Device Address Table

This table displays all currently known active DNP devices.

Section 3.10.7.6

Serial Protocols Statistics Menu

Help...

Serial Protocols Statistics

Refresh Continuous Display

Port Statistics

Port	Protocol	Rx Chars	Tx Chars	Packet Errors	Parity Errors	Framing Errors	Overrun Errors	
1	rawsocket	1923	0	0	0	0	0	reset
2	dnp	0	0	0	0	0	0	reset
3	tcpmodbus	0	0	0	0	0	0	reset

Connection Statistics

Remote IP	Remote Port	Local Port	Rx Packets	Tx Packets	Target Serial Port(s)	Status
-----------	-------------	------------	------------	------------	-----------------------	--------

Refresh Continuous Display

← Return to Serial Protocols

Figure 98: Serial Protocols Statistics Menu

This menu presents statistics of serial port activity and established connections. The menu also allows you to reset a port, forcing call hang-up and re-establishment.

The *Port Statistics* table provides a record for each active serial port. The number of historical received and transmitted characters as well as errors will be displayed.

The *Connection Statistics* table reflects established TCP connections. Network and serial connections can be paired by examining the Target Serial Port(s) field. The *Status* field describes whether a network connection is established or in the process of being established.



NOTE

All counts are from the router's perspective. The Rx Packets count reflects packets received from the network, the contents of which are transmitted at the protocol and reflected in the Tx Chars field.

The *Refresh* button will cause the page to be reloaded.

The *Continuous Display* button will cause the browser to continuously reload the page showing the differences in statistics from the last display. *The difference is not a real time rate in bytes or packets per second.*

Protocol Specific Packet Error Statistics

The *Raw Socket* Packet Errors field reflect the number of times that a network message was received and could not be enqueued at the serial port because of output buffering constraints. This is usually symptomatic of a remote peer that uses a higher baud rate or local flow control.

Section 3.10.7.7

Serial Protocols Trace Menu

The screenshot shows a web interface titled "Line Trace". At the top left is a "Help.." link. Below the title is a warning: "Specifying large numbers of ports, entries and capture time can result in a greate deal of output..". The configuration section includes a "Port" dropdown menu, a "Trace on ports:" row with checkboxes for 1, 2, 3, 4, and "All Ports" (which is checked), a "Message RX/TX" checkbox (checked), a "Hex dump" checkbox (checked), and an "Incoming/Outgoing Connections" checkbox (checked). Below these are two input fields: "Maximum number of entries to capture" set to 20 and "Maximum time in seconds to capture over" set to 5. A "Start Trace" button is at the bottom of the configuration section. Below the button is a log of network activity showing timestamps and connection details. At the very bottom is a "Return to Serial Protocols" link with a left-pointing arrow.

Help..

Line Trace

Specifying large numbers of ports, entries and capture time can result in a greate deal of output..

Port

Trace on ports: 1 ☐ 2 ☐ 3 ☐ 4 ☐ All Ports ☒

Message RX/TX ☒ Hex dump ☒ Incoming/Outgoing Connections ☒

Maximum number of entries to capture 20 Maximum time in seconds to capture over 5

15:30:03.261 TCPCONN Opening connection to 172.30.148.37 50001:39795 for serial port 1
15:30:03.262 TCPCONN Connect failure 111 (Connection refused) 172.30.148.37 50001 for serial port 1
3 seconds to capture
2 seconds to capture
1 second to capture

Start Trace

Return to Serial Protocols

Figure 99: Serial Protocols Trace Menu

This menu displays decoded serial port and network activity.

The desired traffic sources, number of messages and length of time to capture are entered and the *Start Trace* button is pressed. The menu will display up to the provided number of messages waiting up to the specified number of seconds.

The *Trace on ports:* selections feature a list of serial ports with unused entries greyed out. The default is *All Ports*, which selects all ports.

The *Message RX/TX* field allows log entries to be printed for each received or transmitted message, and method of packetization. If the *Hex Dump* field is selected, the first 64 bytes of packet content is displayed.

The *Incoming/Outgoing Connections* field allows regular network level entries such as call connections and received/transmitted messages to be displayed. Note that some unexpected, but unusual, network messages may be displayed if they occur.



NOTE

Specifying large numbers of ports, entries and capture times can result in a great deal of output. Specifying a large capture time may require the web page to wait that interval if activity is infrequent.

Section 3.10.7.8

Serial Protocols Sertrace Utility

The command line sertrace utility offers the ability to trace the activity of serial ports in real time. A port range may be specified to limit the output to specific ports. The level of traffic to trace and the type of decoding may be specified. The tool may also be used to force the port to transmit an output test message. The following is an example of sertrace use:

```
ROX:~# sertrace -h
Trace Serial Protocol Server Activity
Usage: sertrace [-dtr] [-p portrange]
serserver -d protocol decode
serserver -t tcp level events
serserver -r raw packet display
```

```
serserver -p ports to capture (e.g 1,3,6-7)
serserver -s ports to send a test message out on (when 's' + Enter keys are pressed)
In the absence of parameters, all decoding on all ports is provided.
ROX:~# sertrace -p 1 -s 1
10:56:18.405 TCPCONN Listening on TCP Port 50002 from port 1
10:56:19.944 TCPCONN Connection opened from 10.0.10.236 4991:50002
s
10:56:47.497 RAWSOCKET Transmitting message on port 1, length 44
74 68 65 20 71 75 69 63 6b 20 62 72 6f 77 6e 20 the quick brown
66 6f 78 20 6a 75 6d 70 65 64 20 6f 76 65 72 20 fox jumped over
74 68 65 20 6c 61 7a 79 20 64 6f 67 the lazy dog
10:56:47.545 RAWSOCKET Received message on port 1, length 44 (31ms) by timer
74 68 65 20 71 75 69 63 6b 20 62 72 6f 77 6e 20 the quick brown
66 6f 78 20 6a 75 6d 70 65 64 20 6f 76 65 72 20 fox jumped over
74 68 65 20 6c 61 7a 79 20 64 6f 67 the lazy dog
10:56:47.545 TCPCONN Tx Data from port 1 44b to 10.0.10.236 4991:50002
```

Section 3.11

Synchronous Serial Ports

Section 3.11

Synchronous Serial Ports

This section familiarizes the user with configuring the device to forward data from Synchronous Serial ports over IP using the Raw Socket protocol. The ROX Serial Protocols Server (see also [Section 3.10, “Configuring Serial Protocols”](#)) that forwards data traffic from asynchronous serial ports over IP networks can also be configured to run Raw Socket connections over synchronous serial ports.

ROX Synchronous Serial ports provide the following features:

- up to four independent synchronous serial ports per router, two ports per slot
- configurable data rates from 2400 to 230400 bits per second
- internal or external clocking modes
- RTS/CTS hardware flow control
- HDLC framing
- support for the Raw Socket protocol - a means to transport data from a synchronous serial port on the router to a specified remote IP address and TCP port
- up to eight TCP Raw Socket connections per port
- debugging facilities including connection tracing and statistics

Section 3.11.1

Raw Socket Operation on Synchronous Ports

The Raw Socket protocol operates somewhat differently on synchronous ports from the way it does on asynchronous ports. One important difference is that in the case of outgoing calls only, up to eight TCP connections may be mapped to a single synchronous serial port. This means that data received on the port will be transmitted to all configured TCP endpoints, and all data received from the TCP endpoints will be transmitted on the synchronous port.

Another difference concerns framing. Forwarding to and from asynchronous serial ports involves time-outs and "pack" characters in order to aid in framing, or packetization, for transmission on an IP network. HDLC frames received on the synchronous port will be transmitted immediately on the TCP connection, and incoming data on TCP connections are written directly to the port. Note, however, that HDLC framing information is not preserved on transmission via TCP.

Section 3.11.2

Synchronous Serial Port Configuration

The *Synchronous Port Settings* menu is accessed by clicking on *Sync Serial* in the *Networking* folder of the main Webmin menu.




Figure 100: Synchronous Serial Main Menu

Note that the Serial Protocols server is disabled by default and may be enabled via the *Bootup and Shutdown* menu, under the *System* folder of the main Webmin menu. Note also that these menus configure only the physical parameters of the ports. Protocol configuration must be done via the *Serial Protocols* Configuration menu. This section describes the configuration required to run the Raw Socket protocol over the synchronous serial ports.

Section 3.11.2.1

Synchronous Port Settings Menu

This menu displays the current configuration of each synchronous serial port on the router.

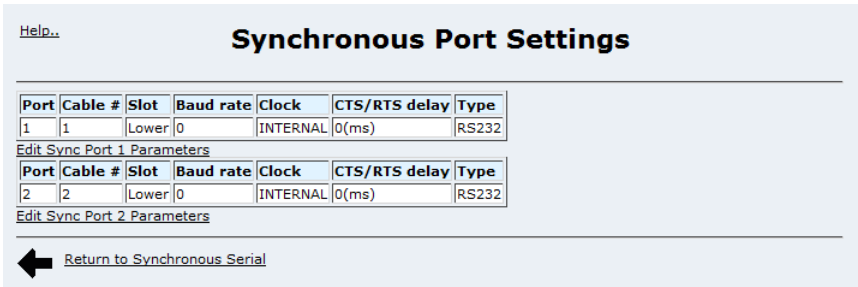


Figure 101: Synchronous Port Settings Menu

- The *Port* field denotes the system's port number for the specific synchronous serial port.
- The *Cable #* field lists the cable marking for the port.
- The *Slot* field indicates whether the port is in the router's *Upper* or *Lower* expansion slot.

The *Baud rate* field sets the bit rate in bits/s of transmitted data and the frequency in Hz of transmitted clock when the port is clocked internally (see below).

The *Clock* field selects whether the port operates from the internal clock or using a received external clock signal.

The *CTS/RTS delay* field sets the delay in milliseconds that RTS will follow CTS.

The *Type* field displays the port's interface type (depending on the type of synchronous card installed).

Underneath each port displayed in the table, a link named *Edit Sync Port X Parameters* leads to an editing menu for the corresponding port's parameters:

Port	Baud rate	Clock	CTS/RTS delay	Type
1	0	INTERNAL	0	RS232

Save Changes

← Return to Synchronous Port Settings

Figure 102: Edit Synchronous Serial Port Parameters

Section 3.11.2.2

Configuring Raw Socket on Synchronous Serial Ports

In order to enable the Raw Socket protocol the desired synchronous serial port, use the *Assign Protocols* sub-menu under the *Serial Protocols* menu (refer to [Section 3.10.7.1, "Assign Protocols Menu"](#)).

As noted above, the Raw Socket protocol operates differently for synchronous and asynchronous ports. The configuration for each also differs.

Note that changes are made immediately, causing call placement to start.

Port	Pack Char	Pack Timer	Turnaround	Call Dir	Max Conns	Rem IP	Rem Port	Loc Port
1	10	1000	off	OUT	2	172.30.148.27	50001	

Save Changes

← Return to Serial Protocols

Figure 103: Edit Synchronous Serial Raw Socket Parameters

The *Pack Char*, *Pack Timer*, and *Turnaround* timer fields are unused and are ignored in the case of synchronous serial ports.

The *Call Dir* field configures whether to accept an incoming connection (*IN*), place an outgoing connection (*OUT*) or do both (*BOTH*).

The *Max Conns* field configures the maximum number of incoming connections to permit when the call direction is incoming or the maximum number of connections to configure if the direction is outgoing.

The *Remote IP* field configures the address used when placing an outgoing connection.

The *Remote Port* field selects the TCP destination port used in outgoing connections.

The *Local Port* field selects the local TCP port to use to accept incoming connections.

When *Max Conns* is set to more than one and the call direction is outgoing, a button labelled *AddNewConns* will appear in the *Action* column of the table to allow new outbound connections to be configured.

When multiple outbound connections are configured and present in the table, a *Delete* button will appear in the *Action* column of the table to allow additional connections to be deleted.

Section 3.11.3

Synchronous Serial Diagnostics

Statistics and line tracing utilities are available for the synchronous serial ports, via the Serial Protocols menu interfaces.

- For information on port statistics, please refer to [Section 3.10.7.6, “Serial Protocols Statistics Menu”](#).
- For information on a menu for tracing Raw Socket protocol exchanges, please refer to [Section 3.10.7.7, “Serial Protocols Trace Menu”](#).
- For information on a command line utility for tracing Raw Socket protocol exchanges, please refer to [Section 3.10.7.8, “Serial Protocols Sertrace Utility”](#).

Section 3.12

Configuring SSH

This section familiarizes the user with:

- Configuring SSH Authentication
- SSH Networking and Access Control
- Setting SSH Server Options

The Secure Shell protocol provides interactive remote login service, remote command execution, and file transmission functions. It implements strong authentication and secure communications over insecure channels. The program that accepts an SSH client's connection is an SSH server. The SSH server can be programmed to enforce conditions to increase security. These conditions can be imposed upon specific hosts or upon all hosts in general.

SSH has seen two major revisions of the protocol: SSH v1 and v2. SSH v1 supported only the RSA authentication scheme, while SSH v2 supports both RSA and DSA.



CAUTION!

Avoid using SSH v1 as it is insecure and could leave your network vulnerable. Use SSH v2.

SSH provides service on TCP port 22 by default. If a firewall configured and operating on the router, it is advised to leave port 22 (or whichever port SSH has been configured to use) open in order to allow secure, authorized access from outside the firewall.

SSH also provides TCP forwarding, a means to forward otherwise insecure TCP traffic through SSH Secure Shell.

Section 3.12.1

Included with SSH

Your ROX software includes "scp", an SSH utility to perform secure copying of files and directories over the network.

If you decide to create user accounts, the ssh-keygen utility can be used to populate the account with SSH keys.

Section 3.12.2

SSH Main Menu

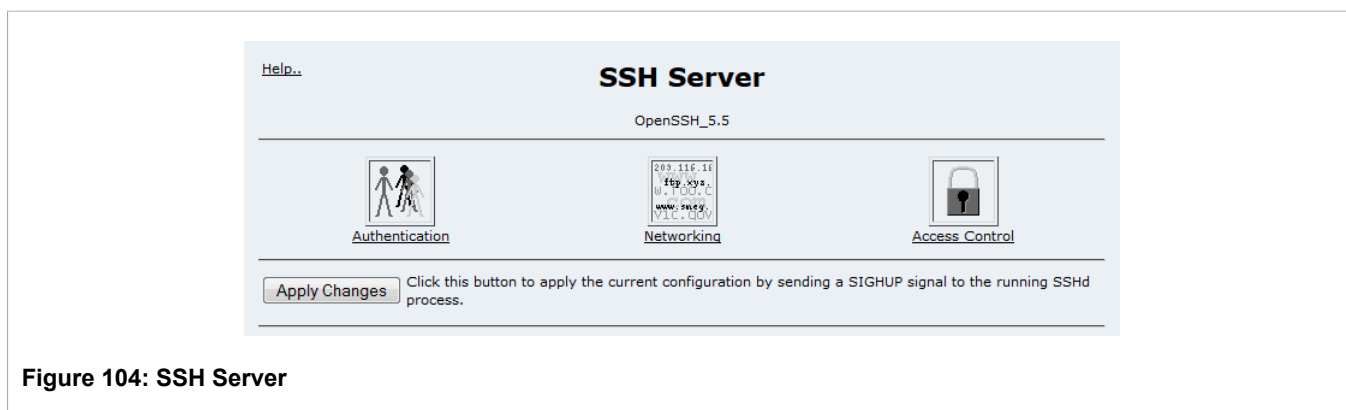


Figure 104: SSH Server

Note that the SSH server is enabled by default and may be disabled via the *System* folder, *Bootup and Shutdown* menu. When enabled, any configuration changes may be made to take effect by selecting the *Apply Changes* button. Note also that when disabling the SSH server, any currently active SSH connections to the server *will remain active*. Subsequent connection attempts will not succeed.

Section 3.12.3

Authentication

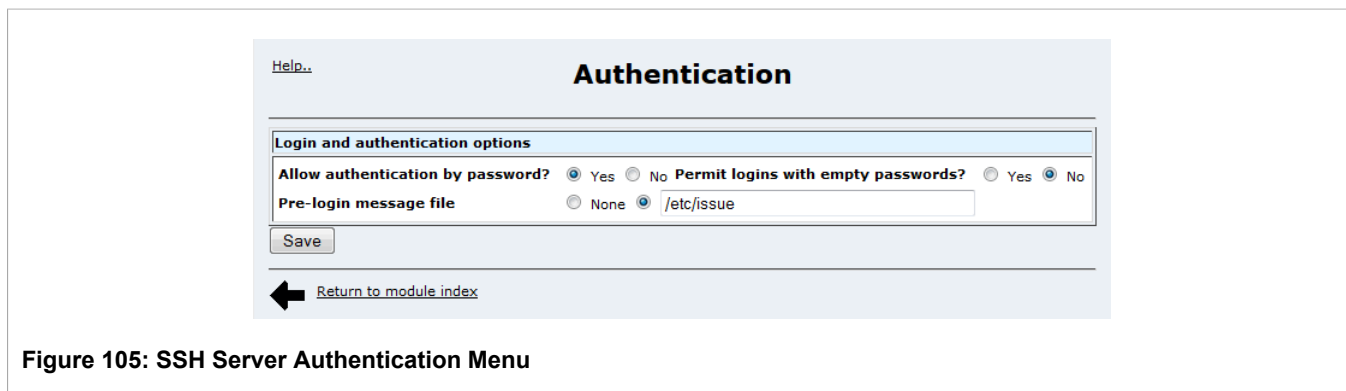


Figure 105: SSH Server Authentication Menu

Allow authentication by password? determines whether to allow clear text tunnelled passwords. If set to "Yes", the user will be allowed to enter a password for authentication if it can not be done using a public key.

Permit logins with empty passwords? (valid when authentication by password is allowed) specifies whether the server will allow login to accounts with empty passwords.

Pre-login message file specifies the name of a file that contains a message that can be displayed before the passphrase/password prompt during the login process. By default it is "/etc/issue", same as the serial console banner. If a custom message for SSH sessions is desired, specify the path to the custom file.

Section 3.12.4

Networking

The screenshot shows a web-based configuration interface titled "Networking". At the top left is a "Help.." link. Below the title is a "Networking options" section. It contains several settings: "Listen on addresses" with radio buttons for "All addresses" (selected) and "Entered below .."; a table with "Address" and "Port" headers; "Listen on port" with radio buttons for "Default (22)" and "22" (selected); "Accept protocols" with checkboxes for "SSH v1" and "SSH v2" (checked); "Disconnect if client has crashed?" with radio buttons for "Yes" (selected) and "No"; "Time to wait for login" with radio buttons for "Forever" and "120" (selected) seconds; "Allow TCP forwarding?" with radio buttons for "Yes" (selected) and "No"; and "Allow connection to forwarded ports?" with radio buttons for "Yes" and "No" (selected). A "Save" button is at the bottom left of the form. Below the form is a back arrow and a link "Return to module index".

Figure 106: SSH Server Networking

The *Listen on addresses* fields determine an IP addresses and port upon which SSH will accept a connection.

The *Listen on port* field determines the port number SSH will listen on, assuming *Listen on addresses* is set to "All addresses".

The *Accept Protocols* fields determine which versions of SSH will be allowed.

The *Disconnect if client has crashed* field determines whether the SSH server should periodically check to see if the client is still alive.

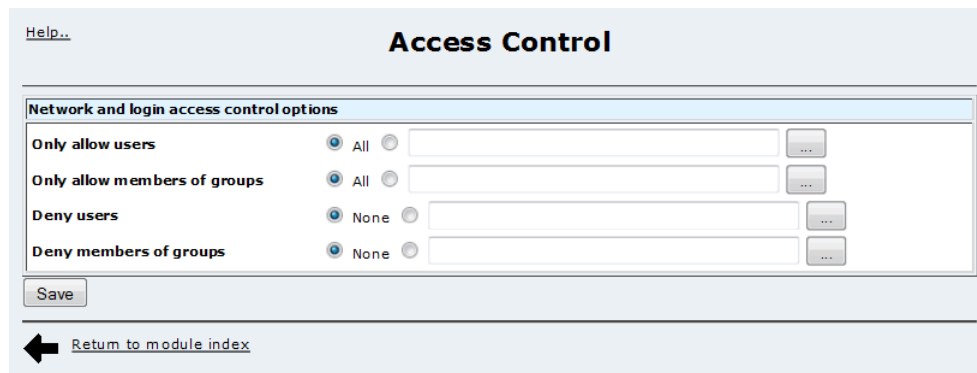
The *Time to wait for login* field determines the maximum time from a connection request until login completes, after which the client will be disconnected.

The *Allow TCP forwarding* field specifies whether TCP forwarding is permitted. If this option is set, clients on a remote network can tunnel TCP connections to machines on the device's network.

The *Allow connection to forwarded ports* field specifies whether remote hosts on the client network are allowed to connect to ports forwarded for the client.

Section 3.12.5

Access Control




The screenshot shows a web-based configuration interface titled "Access Control". At the top left is a "Help.." link. Below the title is a section labeled "Network and login access control options". This section contains four rows of configuration options, each with a radio button, a text input field, and a button with three dots (account selector):

- Only allow users:** Radio button set to "All".
- Only allow members of groups:** Radio button set to "All".
- Deny users:** Radio button set to "None".
- Deny members of groups:** Radio button set to "None".

Below these options is a "Save" button. At the bottom left is a back arrow icon, and at the bottom right is a link labeled "Return to module index".

Figure 107: SSH Server Access Control

The *Only allow users* field specifies the users allowed to connect by SSH. The specification can be a list of user name patterns, separated by spaces. Login is allowed only for user names that match one of the patterns. '*' and '?' can be used as wild cards in the patterns. Only user names are valid, a numerical user ID is not recognized. By default, login is allowed for all users. If the pattern takes the form USER@HOST then USER and HOST are separately checked, restricting logins to particular users from particular hosts.

The account selector () button can be used to build up a list of allowable users.

The *Only allow members of groups* field specifies the "group" (in the Unix sense) of users allowed to connect by SSH. The specification can be followed by a list of group name patterns, separated by spaces. If specified, login is allowed only for users whose primary group or supplementary group list matches one of the patterns. '*' and '?' can be used as wild cards in the patterns. Only group names are valid, a numerical group ID is not recognized. By default, login is allowed for all groups.

The account selector button can be used to build up a list of allowable groups.

The *Deny users* and *Deny members of groups* fields specify users and groups to deny connections to.

Section 3.13

Configuring the Telnet Server

This section familiarizes the user with configuration of the Telnet server.

Telnet is an IP network protocol (RFC854) that provides remote terminal access to a system on TCP port 23. Telnet is commonly used to refer to both the protocol and to the server and client programs that implement it.

An important consideration when using telnet is that both the authentication process and the entire session are transmitted in the clear, i.e. unencrypted. This makes it possible for an attacker to log the telnet session as it appears on the network, and acquire from it the username and password of the account used to log in via telnet, along with the entire text of a login session.

**NOTE**

Wherever possible, the use of SSH is recommended over telnet, since SSH encrypts both the authentication exchange and the session. For information on the configuring the SSH server on ROX please refer to [Section 3.12, “Configuring SSH”](#).

There are some situations in which remote network access to the router is necessary and an SSH client is not readily available, but a telnet client is. Older computer systems or network terminal devices, for example, may support telnet and not SSH, due in part to the fact that telnet itself has been in use for several decades and that it requires much fewer resources than does SSH.

One way to decrease the risk associated with running a telnet server on ROX is to configure the firewall to restrict telnet connections. It would be advisable, for example, to allow telnet connections only from the LAN side of the router, and prevent them from the WAN side.

Section 3.13.1

Telnet Server Configuration

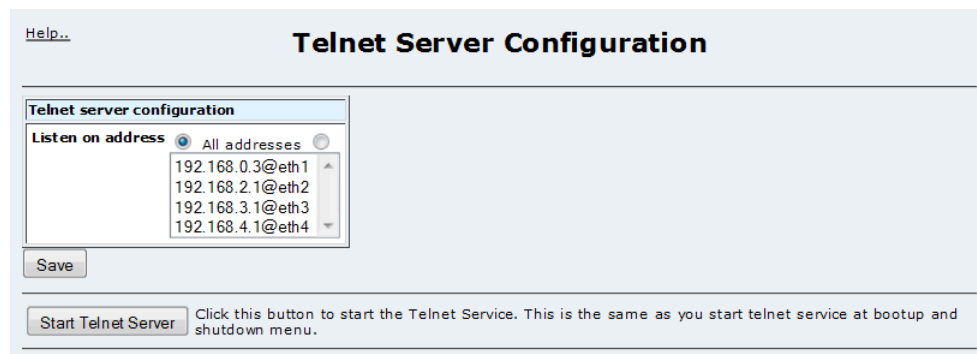


Figure 108: Telnet Server Configuration Main Menu

The *Listen on address* field and list are used to control which local router addresses the telnet server will respond from. If "All addresses" is selected, then the router's telnet server will be reachable via any configured IP address. If "All addresses" is not selected, then the telnet server will only respond from highlighted addresses in the list below.

The *Save* button will commit the configuration change. If the telnet server is not running, the configuration will be saved but will not start the server. If the telnet server is running, the configuration will take effect immediately.

The *Start Telnet Server* button allows you to start telnet server. This is the same as starting the telnet server from the *Bootup and Shutdown* menu.

Note that if the telnet server is running, the *Stop Telnet Server* button will be shown instead of the *Start Telnet Server* button. Note also that if the telnet server is stopped, currently active telnet connections will not be terminated.

Section 3.14

Configuring IRIG-B and IEEE1588

This section familiarizes the user with:

- IEEE 1588 Configuration
- IRIG-B Configuration
- Viewing IRIG-B and IEEE1588 Status

Section 3.14.1

IEEE1588 Fundamentals

The IEEE 1588 working group Precise Timing Protocol (PTP) standard details a method of synchronizing a clocks over networks, including Ethernet. ROX provides a special hardware assisted PTP capability as provided by the device's PTP card. When used in conjunction with the cards Global Positioning System (GPS) receiver, the router can provide nanosecond accuracy via IEEE1588.

Additionally, IEEE1588 may be used (in GPS failure situations) to synchronize to a remote source and provide accurate time for IRIG-B.

Section 3.14.1.1

PTP Network Roles

The IEEE 1588 standard describes regular clocks as devices having a single PTP port that can issue and receive PTP messages. PTP boundary clocks are clocks have multiple PTP ports, offering the ability to serve time to more than one subnet at a time. ROX can serve as a regular clock and communicate with boundary clocks.

The set of devices that can communicate using the PTP protocol IP multicast transmissions are said to be in the PTP subdomain. This is usually a set of PTP devices connected by a switched network or direct links. The "best" clock in the subdomain is known as the master clock. The master clock of a boundary clock is known as the grandmaster clock.

The protocol negotiates among PTP ports to identify the device with the highest quality clock source. Ports issuing messages from the master clock are said to be masters, while those that will receive the messages are slaves. When a port will not participate in the protocol its status is passive. When the network architect knows the relative quality their clock's time sources, they may configure a specific clock to be the preferred master.

Section 3.14.1.2

PTP Master Election

PTP clocks exchange SYNC messages containing information which is used by the PTP Best Master Clock (BMC) algorithm. Several factors will affect the choice of best master clock, including the preferred master clock setting, the clock identifier, grandmaster settings and clock stability.

The clock identifier is the measure of PTP clock quality and is one of the following:

PTP Identifier	Description
GPS	The PTP clock is a primary reference standard traceable to a recognized standard source of time such as GPS. The router uses this identifier when GPS is locked.
NTP	The PTP clock is a secondary reference standard clock. The router uses this identifier when it has synchronized with remote NTP server.
DFLT	After the router has power cycled but before any GPS or NTP locks have occurred.

PTP favors preferred masters over normal masters, GPS over NTP over DFLT, higher clock stability over lower clock stability.

Section 3.14.1.3

Synchronizing NTP from IEEE1588

If GPS is unavailable and PTP becomes a slave the NTP server will view the received IEEE1588 time as any other source of time. The quality (i.e. stratum) of IEEE1588 information is determined by the type of clock source at the master, the number of Boundary Clock hops and the measured network jitter.

The number of Boundary Clock hops is the number of IEEE1588 devices the original time source is relayed through (and not Ethernet hops) and is always 1 or higher.

The measured network jitter factor is 0 if jitter is higher than 1 microsecond and -1 if less than 1 microsecond.

PTP Identifier	Stratum reported to NTP
GPS	1 + Number of Hops ? 1 (if low jitter)
NTP	user configurable value (default 2) + Number of Hops ? 1 (if low jitter)
DFLT	user configurable value (default 10) + Number of Hops ? 1 (if low jitter)

The stratum number reported will be limited to a range of 1 to 16 to comply with NTP.

As an example, a directly connected PTP clock having a GPS clock source and low jitter would report a stratum of 1. With defaults a 2 hop away PTP clock having a NTP clock source and high jitter would report a stratum of 4.

Section 3.14.2

IRIG-B Fundamentals

IRIG-B outputs are provided by the Precision Time Protocol Card option.

The Inter-Range Instrumentation Group (IRIG) IRIG-B standard details the format of an output signal containing information for the current day, hour, minute and second in UTC format, broadcast at the start of each second. ROX complies to IRIG Standard 200-04 generating formats IRIG-B002 and IRIG-B003 (PWM) and IRIG-B122 and IRIG-B123 (AM).

Section 3.14.2.1

IRIG-B Output Formats

The router provides three ports by which the signal is distributed, namely:

- An Amplitude Modulated (AM) sinusoidal output port (PTP1),
- Two TTL voltage level output ports (PTP2 and PTP3) which may be configured as either pulse per second (PPS) or pulse width modulated (PWM).

The signal can be used to synchronize intelligent devices to a high quality time source, called the reference clock. The router uses a global positioning satellite (GPS) receiver, NTP or the router's local clock as the reference clock.

Section 3.14.2.2

Reference Clocks

GPS provides the highest quality reference clock. It will always be used when it is available, but may require some time after boot before becoming acquired (or "GPS locked"). Typically, GPS lock is usually acquired within five minutes of boot. When GPS is the reference clock, IRIG-B timestamps are accurate to within ns.

If GPS has not yet locked and IEEE1588 is locked, the router will use IEEE1588 server as a reference clock. When IEEE1588 is synchronized, IRIG-B timestamps are accurate to within microsecond or sub microseconds.

If GPS and IEEE1588 have not yet locked, the router will use an NTP server or peer as a reference clock. NTP typically requires less than two minutes after boot to synchronize. When NTP is the reference clock, IRIG-B timestamps can be accurate to within ms.

Before NTP is able to synchronize, the router will use the local clock to obtain the time and will emit IRIG-B timestamps on a one second basis.

Section 3.14.2.3

How the Router Selects a Reference Clock

The router can be configured to use the following as reference clocks:

- GPS, IEEE1588, NTP and the local clock,
- GPS, NTP and the local clock,
- GPS and IEEE1588,
- GPS

If the router is configured to use multiple reference clocks, it will start sending timestamps using the best ever locked reference clock (local clock is always locked). If better reference clock is locked later, the router will "step" (i.e. suddenly change) the time and use the new reference clock. If the current reference clock becomes unavailable, the router will keep running with its own high precision timing hardware. It will use this hardware until the last used reference clock is locked or a higher quality reference clock is available.

If the router is configured to use only GPS, no timestamps will be issued until GPS locks. If GPS fails, the router will keep running with its own high precision timing hardware. When GPS returns, the time will be stepped back to the GPS reference clock.

Section 3.14.2.4

GPS Cable Compensation

GPS signals received by the antenna will be delayed in time depending upon the type and length of the cable to the router. This delay will introduce inaccuracy in the calculated time and position.

ROX provides a method to account for this delay. The table below gives some examples of the delay that can be expected for a given dielectric type. Please note that cable characteristics varies from one manufacturer to the other.

Dielectric Type	Time Delay in ns/m (ns/ft)
Solid Polyethylene	4.62 (1.54)
Foam Polyethylene (FE)	3.81 (1.27)
Foam Polystyrene (FS)	3.36 (1.12)

Dielectric Type	Time Delay in ns/m (ns/ft)
Air Space Polyethylene (ASP)	3.45-3.63 (1.15-1.21)
Solid Teflon (ST)	4.38 (1.46)
Air Space Teflon (AST)	3.39-3.60 (1.13-1.20)

Section 3.14.3

IRIG-B/IEEE1588 Configuration




Figure 109: IRIG-B/1588 Main Menu

This menu allows you to configure IRIG-B and IEEE1588, display its current status and review historical changes.

Section 3.14.3.1

General Configuration




Figure 110: IRIGB/IEEE1588 General Configuration Menu

This menu allow you to configure general parameters.

The *Reference Clock Selection* field selects the order in which to prefer reference clocks.

The *Cable Compensation* field specifies the value, in nanoseconds, that will be used to compensate for the cable type and length. The compensation is done using integer nanosecond values. Fractional decimal values will be truncated.

Section 3.14.3.2

IRIG-B Configuration

Figure 111: IRIG-B Configuration Menu

This menu allow you to configure IRIG-B parameters. The save button will save the changes of configuration permanently.

The *AM Port 1 (PTP1) Output* field enables or disables the amplitude modulated output of this port.

The *TTL Port 2 (PTP2) Output* and *TTL Port 3 (PTP3) Output* fields sets the output formats of these ports to PPS, PWM and OFF.

Section 3.14.3.3

IEEE1588 Configuration

Figure 112: IEEE1588 Configuration Menu

This menu allows you to configure IEEE 1588 parameters.

The *1588 Working Mode* field allows configures whether the router will be forced to 1588 slave mode or determine its role by the BMC algorithm.

The *Preferred Master Clock* field configures the router to be preferred master clock.

The *Subdomain Name* field allows you to choose which domain you want the router to participate in. There are four domains available, each mapped to a different multicast IP address.

The *Sync Interval* field configures the rate at which SYNC messages are issued.

The router NTP daemon uses GPS as a clock source when it is available and with IEE1588 when GPS is not available.

The *Treat NTP sync'd grandmaster as stratum* field assigns the stratum number when grandmaster clock synchronized with remote NTP server but not GPS.

The *Treat Local Clock sync'd grandmaster as stratum* field assigns the stratum number when grandmaster clock synchronized with local clock but not NTP server or GPS.

Section 3.14.3.4

IRIG-B Status

Month	Day	Time	Process	Event
Oct	30	13:08:27	/usr/sbin/irigb[2657]	Force FPGA to init time
Oct	30	13:08:32	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	13:15:13	/usr/sbin/irigb[2657]	Detection of FPGA time (13/10/30 17:15:6) is out of sync with GPS (13/10/30 17:15:9), force FPGA reset clock
Oct	30	13:16:40	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	13:16:51	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	13:08:53	/usr/sbin/irigb[2657]	Detection of FPGA time (13/10/30 17:21:9) is out of sync with GPS (13/10/30 17:21:16), force FPGA reset clock
Oct	30	13:23:30	/usr/sbin/irigb[2657]	reload configuration

Figure 113: IRIG-B GPS Status

This page shows whether GPS is locked, and the source of the current reference clock.

Section 3.14.3.5

IEEE1588 Status

IEEE1588 Status
Local Clock Port IP/MAC:
IEEE1588 Status:

Figure 114: IEEE1588 Status

This page shows the historical status of IEEE1588 on the router.

The line above the table provides the local clock IP address, MAC address and the time quality information. The table will contain entries made when the clock source or status changes. The current local time on the router, the IEEE1588 status, IEEE1588 and UTC time, the offset from master in seconds, the master IP/MAC address and grandmaster MAC address are provided.

Section 3.14.3.6

IRIG-B Log

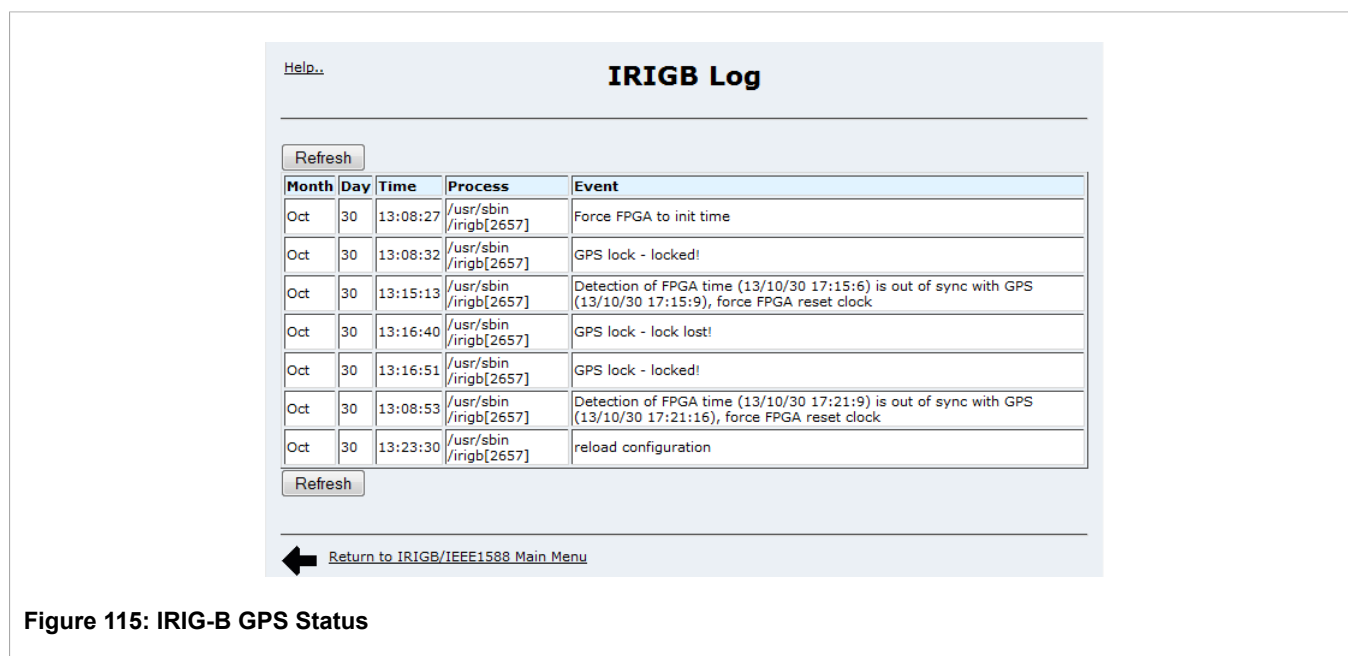


Figure 115: IRIG-B GPS Status

This page reflects reference clock changes in IRIG-B.

Section 3.15

Configuring the Intrusion Detection System

This section familiarizes the user with:

- Configuration of Snort as an Intrusion Detection System
- Generating a daily Snort analysis email

Section 3.15.1

Snort Fundamentals

The Snort Intrusion Detection System (IDS) provides a type of security management system for the router. Snort gathers and analyzes information on various network interfaces to identify possible security breaches, which include both intrusions (attacks from outside the protected network) and misuse (attacks from within the protected network). Snort examines packets received on selected interfaces, applies "rules" from its database and generates log entries to warn of "vulnerabilities".

Snort is a complex system with many capabilities and a large community of contributors and users. The interested reader is encouraged to seek more information at the project's web site: <http://snort.org>.

Section 3.15.1.1

Configuring Snort

Snort must be configured properly before it is started. In addition to the detailed Snort Manual (available at <http://snort.org>), follow these basic guidelines

- Configure the network variables for Snort according to the IP network scheme of your organization. For more information, refer to [Section 3.15.3, “Network Settings”](#).

By default, the IP address for HOME_NET and EXTERNAL_NET are set to use any IP address. If you decide to keep the default configuration for the HOME and EXTERNAL networks, make sure not to enable any rules that use “!” as a prefix to a network variable.

- Configure the preprocessor for Snort. Most of the preprocessors are enabled by default but you may want to disable the preprocessor that are not required for the rule set that you plan to enable. For more information, refer to [Section 3.15.4, “PreProcessors”](#).
- Configure the destination for Alert & Logging. The default destination for alert and logging is auth.log. You may want to change it to syslog and it is strongly recommended to setup a remote logging facility so that old log messages are not lost after the log rotation in the local logging facility. For more information, refer to [Section 3.15.1.4, “Alerting Methods”](#) and [Section 3.15.5, “Alerts and Logging”](#).
- A sample configuration for Snort is provided through Webmin. You may want to modify the configuration based on the need of your network. For more information, refer to [Section 3.15.7, “Edit Config File”](#).
- Enable Snort on the Ethernet interface(s) on which you want Snort to decode network traffic and raise alert if there is a matching condition found among the enabled rules.

**NOTE**

Due to limited CPU and memory resources on a router, it is strongly recommended to enable Snort only on one Ethernet interface which could be connected to the external network. Following is a general guideline for the maximum number of rules that could be enabled depending on the Ethernet interface(s) on which Snort is enabled.

- *One Ethernet Interface: 10,000 rules with a maximum of 400 flowbit tags used in the rules*
- *Two Ethernet Interfaces: 4,500 rules with a maximum of 200 flowbit tags used in the rules*
- *Four Ethernet Interfaces: 800 rules*

For more information, refer to [Section 3.15.1.2, “Which Interfaces to Monitor”](#) and [Section 3.15.2.2, “Interfaces”](#).

- Upload the rule sets to the “/etc/snort/rules” directory on the router using the upload/download file menu in Webmin. The rule sets can also be uploaded through WinSCP or any other method.

Include and enable the intended rule sets in the Snort.conf file. The Snort.conf file may include some pre-configured rule sets to demonstrate an example and these rule sets are not enabled by default. If the rule set that you intend to use is listed in the configuration file, remove the “#” sign from the prefix of the included rule set, otherwise the rule set needs to be included in the configuration file.

For more information about rule sets, refer to [Section 3.15.1.3, “Snort Rules”](#) and [Section 3.15.2.3, “Rulesets”](#).

- Test your configuration by using the “Test Configuration” option under Snort IDS in Webmin.

For more information, refer to [Section 3.15.6, “Test Configuration”](#).

- Enable Snort to start now and after the boot up process from the Boot Up and Shutdown menu in Webmin and verify that Snort is started and running on the intended interface. Make sure to restart Snort after making any changes under the Snort IDS menu in Webmin. Snort can be restarted by pressing the Restart Snort button under the Snort IDS menu in Webmin. Once restarted wait for the message “Snort was restarted and is running on ports: [Intended Ethernet Interface(s)]” to ensure successful operation of Snort.

Section 3.15.1.2

Which Interfaces to Monitor

Typically, the router will have an interface to an external network and interfaces comprising the local network. The firewall will cite these interfaces as belonging to the net and local zones. A key decision is whether to monitor traffic outside, or inside of the firewall.

Monitoring traffic outside the firewall (on the external network interface) has the advantage that attacks which the firewall is blocking can be seen. This method, however, will generate a large number of alerts. Additionally, firewall rules installed to eliminate vulnerabilities will not prevent future alerts since traffic is monitored before the firewall. Finally, this method will not detect misuse of the local ports.

Monitoring traffic inside the firewall (on all local interfaces) has the advantage that the number of alerts decreases as vulnerabilities are eliminated at the firewall. It is also good to monitor as much of the internal traffic as possible.

Section 3.15.1.3

Snort Rules

The Snort application in ROX does not include any rule sets. It is the responsibility of the user to download rule sets from various sources and only upload the rule sets that are intended to be used on the router running Snort. Common sources for Snort rule sets include VRT, Community and ETOpen rule sets. The VRT and Community rule sets can be downloaded from <http://snort.org> and registration is required before downloading free or subscribed rules. The ETOpen rule sets can be downloaded from <http://emergingthreats.net>, but make sure to download the nogpl version if you intend to use the VRT rule sets beside the ETOpen rule sets.

**NOTE**

Unzip the downloaded rule set files if they are zipped, select the rule sets that you plan to use on your router and copy the selected rule sets to a new folder before uploading them to the intended router. The selection of the rule sets depends on the need of your organization. You may want to enable the rules from a rule file before uploading them to a router. A rule can be enabled by removing the “#” sign from the beginning of a rule.

Each rule contains a unique Signature Identifier (SID). The SID is included in reported alerts as part of a Snort unique rule ID, a three digit number of the form [generator:SID:revision]. The "generator" field reflects the type of preprocessor or decoder used for a rule. The SID is a unique number to reflect an individual rule, while the "revision" reflects improvements to the rule.

The main Snort IDS menu provides the capability to disable individual and groups of rules.

A difference between Disabled and Enabled rule is shown in the following example:

```
#alert tcp $EXTERNAL_NET any -> $HOME_NET any (msg:"This is a test"; sid:161)
alert tcp $EXTERNAL_NET any -> $HOME_NET any (msg:"This is a test"; sid:161)
```

Section 3.15.1.4

Alerting Methods

Alerts generated by Snort are stored by one of two methods: as syslog messages or in a specified alert file.

When the local syslog method is chosen, the destination log file may be selected.

When the alert file method is chosen, a daily analysis of the file can be emailed.

The SIDs referenced in alerts can be used to quickly locate the rule via the main Snort IDS menu. The rule itself often contains HTML links to Internet resources such as <http://www.securityfocus.com/> and <http://cve.mitre.org>. These provide more in-depth descriptions of the vulnerability.

Section 3.15.1.5

Performance and Resources

The performance impact of Snort varies with the number of interfaces monitored, the number of rules enabled, the packet rate and the logging method.

Snort has been empirically determined to use about 20% of the CPU clock cycles at its maximum processing rate.

The router is capable of recording about 300 entries/second to the local syslog and 500 entries/second to the alert file. Alerts at rates exceeding the above rates will not be recorded.

Snort will require 5 MB of system memory to start with an additional 15 MB of memory for each interface monitored.

Section 3.15.1.6

Troubleshooting Snort

If Snort does not start on the intended interface(s), use the following steps to identify and solve the issue:

1. In most cases Webmin would display the error message indicating the root cause of the problem. You can also refer to the Snort start up log message, which is logged in the syslog on a router running Snort.
2. From the ROX shell, type the command `ps aux | grep snort` and find out whether Snort is running on the intended interface.
3. If Snort is enabled on one or more interface(s) and restarting Snort results in Snort not running on all intended interfaces, you may have exceeded the maximum number of rules and flowbit tags. For more information, refer to [Section 3.15.1.1, "Configuring Snort"](#).

To find out the total number of rules and flow bit tags that are currently being used by the enabled rule sets, go to the ROX shell and type either `cat /var/log/syslog | grep rules` or `cat /var/log/auth.log | grep rules` depending on your log destination. This must be done after Snort has started. If you have exceeded the maximum limit, adjust the total number of enabled rule sets and restart Snort.

Section 3.15.2

IDS Configuration

The Snort IDS menu configures Snort IDS and is composed of three sections:

- Global Snort Configuration

- Interfaces
- Rulesets



NOTE

Snort is disabled by default and may be enabled via the System folder, Bootup and Shutdown menu. If Snort is running, configuration changes must be made active by restarting it. The Restart Snort button will restart Snort and list the interfaces it is active upon.

Section 3.15.2.1

Global Configuration




Figure 116: Snort Global Configuration

The Global Configuration menu section configures parameters that apply to all interfaces.

Section 3.15.2.2

Interfaces

Interfaces								
Interface	Status	Action	Interface	Status	Action	Interface	Status	Action
eth1	✓	Disable	eth3	✓	Disable			
eth2	✓	Disable	eth4	✓	Disable			

Figure 117: Snort Interfaces

The Interfaces section selects the interfaces Snort will monitor. You must restart Snort after changing interfaces.

Section 3.15.2.3

Rulesets

Rulesets								
Rule Set	Status	Action	Rule Set	Status	Action	Rule Set	Status	Action
app-detect	✓		browser-other	✓		exploit	✓	
attack-responses	✓		browser-plugins	✓		exploit-kit	✓	
backdoor	✓		browser-webkit	✓		file-executable	✓	
bad-traffic	✓		chat	✓		file-flash	✓	
blacklist	✓		content-replace	✓		file-identify	✓	
botnet-cnc	✓		ddos	✓		file-image	✓	
browser-chrome	✓		dns	✓		file-java	✓	
browser-firefox	✓		dos	✓		file-multimedia	✓	
browser-ie	✓		experimental	✓		file-office	✓	

Figure 118: Snort Rulesets

The Rulesets section selects the rules to apply on monitored interfaces.

Each "ruleset" reflects a collection of related rules. Links in the *Action* column enable or disable all of the rules in a ruleset.

To modify individual rules in a ruleset, click the ruleset name link in the *Rule Set* column. The Edit Ruleset menu appears:

[Help..](#)

Edit Ruleset

Current Rules in app-detect.rules

Rule	Signature	Status	Action
1	alert tcp \$HOME_NET any -> \$EXTERNAL_NET \$HTTP_PORTS (msg:"APP-DETECT Absolute Software Computrace outbound connection - search.namequery.com"; flow:to_server,established; content:"Host 3A search.namequery.com 0D 0A "; fast_pattern:only; http_header; content:"TagId: "; http_header; metadata:policy security-ips drop, ruleset community, service http; reference:url,absolute.com/support/consumer/technology_computrace; reference:url,www.blackhat.com/presentations/bh-usa-09/ORTEGA/BHUSA09-Ortega-DeactivateRootkit-PAPER.pdf; classtype:misc-activity; sid:26287; rev:4;)	✓	Edit
2	alert tcp \$HOME_NET any -> \$EXTERNAL_NET \$HTTP_PORTS (msg:"APP-DETECT Absolute Software Computrace outbound connection - search.dnssearch.org"; flow:to_server,established; content:"Host 3A search.dnssearch.org 0D 0A "; fast_pattern:only; http_header; content:"TagId: "; http_header; metadata:policy security-ips drop, ruleset community, service http; reference:url,absolute.com/support/consumer/technology_computrace; reference:url,www.blackhat.com/presentations/bh-usa-09/ORTEGA/BHUSA09-Ortega-DeactivateRootkit-PAPER.pdf; classtype:misc-activity; sid:26286; rev:4;)	✓	Edit
3	alert tcp \$HOME_NET any -> \$EXTERNAL_NET \$HTTP_PORTS (msg:"APP-DETECT Chocoplayer successful installation"; flow:to_server,established; content:"/post/player.php"; http_uri; content:"type="; http_client_body; content:"mac="; distance:0; http_client_body; content:"os="; distance:0; http_client_body; metadata:policy security-ips drop, service http; reference:url,www.chocoplayer.com; classtype:misc-activity; sid:25981; rev:1;)	✓	Edit
4	alert tcp \$HOME_NET any -> \$EXTERNAL_NET 80 (msg:"APP-DETECT Ammyy remote access tool"; flow:to_server,established; content:"POST"; http_method; content:" 0A Host 3A 20 rl.ammyy.com 0D 0A "; fast_pattern:only; http_header; metadata:ruleset community, service http; reference:url,www.ammyy.com; classtype:policy-violation; sid:25947; rev:2;)	✓	Edit
5	alert tcp \$EXTERNAL_NET any -> \$HOME_NET \$HTTP_PORTS (msg:"APP-DETECT Acunetix web vulnerability scanner XSS attempt"; flow:to_server,established; content:">= 5C xa2"; fast_pattern:only; http_uri; metadata:service http; reference:url,www.acunetix.com; classtype:web-application-attack; sid:25365; rev:1;)	✓	Edit

[Return to Snort IDS Main](#)

Figure 119: Snort Ruleset Edit



NOTE

Editing rules on the device itself may conflict with updated rules that are uploaded to the device. Make sure precautions are in place to preserve any edits made on the device.

You can enable, disable and edit each rule in the ruleset.

Section 3.15.2.4

Rule Lookup by SID

The Look Up Rule button accepts an SID and displays its rule. You may elect to disable the rule or learn more information about it.

Section 3.15.3

Network Settings

[Help..](#)

Network Settings

Snort Network Settings		
Network Variable	Setting	Description
HOME_NET	any	IP Addresses in the local subnet
EXTERNAL_NET	any	IP Addresses in the external subnet
DNS_SERVERS	\$HOME_NET	Addresses of DNS servers in the local subnet
SMTP_SERVERS	\$HOME_NET	Addresses of SMTP servers in the local subnet
HTTP_SERVERS	\$HOME_NET	Addresses of HTTP servers in the local subnet
SQL_SERVERS	\$HOME_NET	Addresses of SQL servers in the local subnet
TELNET_SERVERS	\$HOME_NET	Addresses of TELNET servers in the local subnet
SSH_SERVERS	\$HOME_NET	Addresses of SSH servers in the local subnet
FTP_SERVERS	\$HOME_NET	Addresses of FTP servers in the local subnet
SIP_SERVERS	\$HOME_NET	Addresses of SIP servers in the local subnet
MODBUS_CLIENT	\$HOME_NET	Address of the Modbus client
MODBUS_SERVER	\$HOME_NET	Address of the Modbus server
AIM_SERVERS	[64.12.24.0/23,64.12.28.0/23,64.12.161.0/24,64.12.163.0]	Known AIM servers
HTTP_PORTS	[36,80,81,82,83,84,85,86,87,88,89,90,311,383,591,593,63]	Ports which serve http
SHELLCODE_PORTS	!80	Ports you want to look for SHELLCODE on
ORACLE_PORTS	1024:	Ports you want to look for ORACLE attacks on
SSH_PORTS	22	Ports used by the SSH server
FTP_PORTS	[21,2100,3535]	Ports used by the FTP server
SIP_PORTS	[5060,5061,5600]	Ports used by the SIP server
FILE_DATA_PORTS	[\$HTTP_PORTS,110,143]	List of file data ports for file inspection
GTP_PORTS	[2123,2152,3386]	Ports used for the GTP preprocessor


 [Return to Snort IDS Main](#)

Figure 120: Snort Network Settings

This menu allows you to configure the IP addresses and ports of servers in the local and external network.

The *Home Net* field defaults to "ANY" and designates the IP subnet of any local ports on the router. Configuring a specific subnet can reduce the number of alerts generated.

To specify multiple IP addresses in the *Home Net* field, enclose a comma-separated list of IP addresses in square brackets ([]). There must not be any spaces in the list of addresses. For example:

```
[10.10.10.20,192.168.1.23,172.16.30.25]
```

Section 3.15.4

PreProcessors

[Help...](#)

PreProcessors

Snort Preprocessor Settings		
Preprocessor	Options	Status
normalize_ip4		<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
normalize_tcp	ips ecn stream	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
normalize_icmp4		<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
normalize_ip6		<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
normalize_icmp6		<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
frag3_global	max_frags 65536	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
frag3_engine	policy windows detect_anomalies overlap_limit 10 mi	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
stream5_global	track_tcp yes, \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
stream5_tcp	policy windows, detect_anomalies, require_3whs 180,	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
stream5_udp	timeout 180	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
http_inspect	global iis_unicode_map unicode.map 1252 compress	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
http_inspect_server	server default \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
rpc_decode	111 32770 32771 32772 32773 32774 32775 32776 327	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
bo		<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
ftp_telnet	global inspection_type stateful encrypted_traffic no ch	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
ftp_telnet_protocol	telnet \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
ftp_telnet_protocol	ftp server default \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
ftp_telnet_protocol	ftp client default \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
smtp	ports { 25 465 587 691 } \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
ssh	server_ports { 22 } \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
dcerpc2	memcap 102400, events [co]	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
dcerpc2_server	default, policy WinXP, \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
dns	ports { 53 } enable_rdata_overflow	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
ssl	ports { 443 465 563 636 989 992 993 994 995 7801 7802	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
sensitive_data	alert_threshold 25	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
sip	max_sessions 40000, \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
imap	\	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
pop	\	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
modbus	ports { 502 }	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
dnp3	ports { 20000 } \	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
reputation	\	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled

[Return to Snort IDS Main](#)

Figure 121: Snort Preprocessors

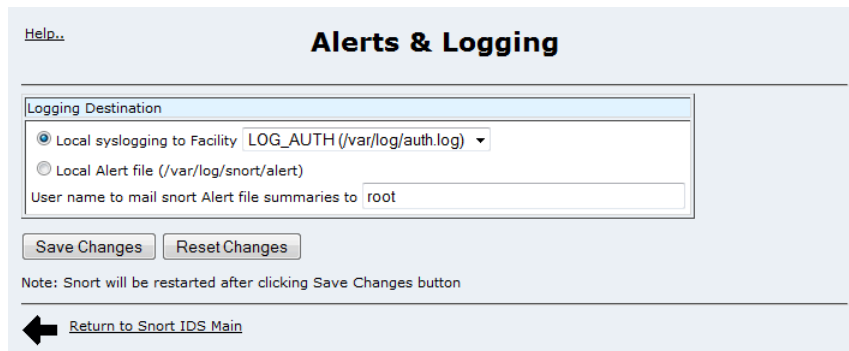
Preprocessors are plug-in modules that operate on the captured packets. Preprocessors perform a variety of transformations to make it easier for Snort to classify packets.

The configuration of preprocessors is beyond the scope of this user guide.

Section 3.15.5

Alerts and Logging

Alerts generated by Snort are stored by one of three methods: as local syslog messages, remotely syslogged messages and in an alert file.



The screenshot shows the 'Alerts & Logging' configuration page. At the top left is a 'Help..' link. The title 'Alerts & Logging' is centered. Below it is a 'Logging Destination' section with two radio buttons: 'Local syslogging to Facility' (selected) and 'Local Alert file (/var/log/snort/alert)'. The 'Local syslogging to Facility' option has a dropdown menu showing 'LOG_AUTH (/var/log/auth.log)'. Below the radio buttons is a text field for 'User name to mail snort Alert file summaries to' with the value 'root'. There are two buttons: 'Save Changes' and 'Reset Changes'. A note below the buttons states: 'Note: Snort will be restarted after clicking Save Changes button'. At the bottom left is a back arrow icon and a link 'Return to Snort IDS Main'.

Figure 122: Snort Alerts

When the *Local syslogging* method is chosen, the destination log file may be selected.

When the alert file method is chosen, a daily analysis of the file can be emailed to the user provided in the *User Name..* field. Note the you must also visit the Maintenance menu, Miscellaneous sub-menu, Outgoing Mail sub-menu in order to configure a mail forwarder.

Section 3.15.6

Test Configuration

This menu validates the Snort configuration and displays a final report. The report lists any errors found and then provides details on the current status of the configuration. Use this utility to test your Snort configuration before deployment.

Section 3.15.7

Edit Config File



Figure 123: Edit Config File

Snort is extremely flexible and not all capabilities have been described in this user guide. This menu provides the user with the ability to make raw configuration changes to the Snort configuration file from within Webmin.

**CAUTION!**

Configuration hazard – risk of data corruption. Modifications to the Snort configuration made on the device itself may conflict with rules that are later uploaded to the device from a central server.

Section 3.16

Brute Force Attack Protection System

ROX 1.16 features a Brute Force Attack (BFA) protection mechanism. This mechanism will analyze the behavior of external hosts trying to access the ssh port, specifically the number of failed logins. On the 5th or 6th failed login the IP address of the host will be blocked for 630 seconds. The range of 5 to 6 failed login exists to take into

account various methods of accessing the device, notably when same or different ports are used across a series of failed logins.

**IMPORTANT!**

The Brute Force Attack (BFA) protection system is not applicable to SNMP. Follow proper security practices for configuring SNMP. For example:

- *Do not use SNMP over the Internet*
- *Use a firewall to limit access to SNMP*
- *Do not use SNMPv1*

After the blocking time has expired, the IP will be allowed to access the device again. If the malicious behavior continues from the same IP (eg. another 6 failed login attempts), then the IP address will be blocked again, but the time blocked will increase by a factor of 1.5. This will continue as long as the IP keeps the same behavior.

**NOTE**

The failed logins must happen within 10 minutes to be considered as malicious behavior.

When BFA protection is started the following auth.log entry is displayed:

```
Jun 7 14:40:59 ruggedrouter sshguard[24720]: Started successfully [(a,p,s)=(50, 420, 600)], now ready to scan.
```

A bfaInfo SNMP v2c or V3 trap can be sent each time an IP is blocked, if SNMP traps are configured.

An auth.log entry is created when an IP address is blocked, an example is shown below:

```
Jun 7 14:43:04 ruggedrouter sshguard[24720]: Blocking 172.59.9.1:4 for >630secs: 60 danger in 5 attacks over 70 seconds (all: 60d in 1 abuses over 70s).
```

Each failed login has a value of 10. The sixth failed login represents a cumulative value of 60 and is reported as such.

Brute Force Attack protection is enabled in the "Bootup and Shutdown" page, under 'sshguard'.

**NOTE**

Enabling/disabling a firewall configuration will reset the Brute Force Attack protection mechanism.

4 System Administration

This chapter describes how to perform various administrative tasks related to device identification, user permissions, alarm configuration, certificates and keys, and more. It describes the following tasks:

- [Section 4.1, “Webmin Configuration”](#)
- [Section 4.2, “Configure Webmin Users”](#)
- [Section 4.3, “Configuring the System”](#)
- [Section 4.4, “Managing SSH Keys and Certificates”](#)
- [Section 4.5, “Access Manager Secure Access Portal”](#)
- [Section 4.6, “RADIUS Authentication”](#)
- [Section 4.7, “RADIUS Server Configuration”](#)

Section 4.1

Webmin Configuration

This section familiarizes the user with configuring the router through the Webmin menu and describes the following procedures:

- [Section 4.1.1, “IP Access Control”](#)
- [Section 4.1.2, “Ports and Addresses”](#)
- [Section 4.1.3, “Change Help Server”](#)
- [Section 4.1.4, “Logging”](#)
- [Section 4.1.5, “Authentication”](#)
- [Section 4.1.6, “Webmin Events Log”](#)




Figure 124: Webmin Configuration Menu

Section 4.1.1

IP Access Control



Figure 125: Webmin Configuration Menu, IP Access Control

Webmin uses a secure communications method called Secure Sockets Layer (SSL) to encrypt traffic with its clients. Webmin guarantees that communications with the client is kept private. But Webmin will provide access to any client that provides the correct password, rendering it vulnerable to brute force attacks. The best way of addressing this problem is to restrict access to specific IP addresses or subnets.

By default, IP access control allows all IP addresses to access Webmin.

If your router is being used on a completely private network, or IP access control is being provided by the firewall you may leave IP Access Control disabled. Select the *Allow from all addresses* field and Save.

If you wish to restrict access to a single address or subnet, select the *Only allow from listed addresses* field. Enter a single IP address or a subnetted address.

If you wish to deny access to a specific subnet, select the *Deny from listed addresses* field. Enter a single IP address or a subnetted address.

If DNS is configured you may allow and deny based upon hostname. Partially qualified domain names such as *.foo.com are acceptable.

The *Resolve hostnames on every request* field forces Webmin to perform a hostname lookup for every user access. The result of this will be that a dynamically assigned IP with a DNS entry with a Dynamic DNS registrar will be able to be checked against the IP Access Control list, just like a fixed address. This method is useful for administrators who travel or simply don't have a fixed address at their normal location.

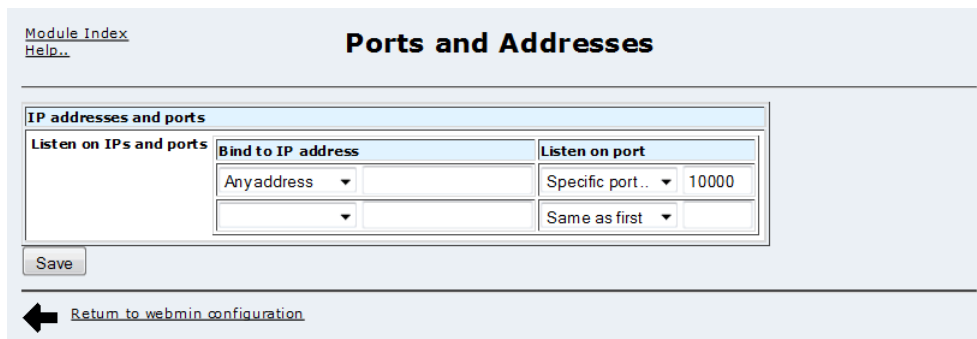


NOTE

This is not efficient if you have more than a few domain names entered in the IP Access Control list, due to the high overhead of performing a name lookup for every hostname in the list on every request.

Section 4.1.2

Ports and Addresses



The screenshot shows the 'Ports and Addresses' configuration page in Webmin. At the top left are links for 'Module Index' and 'Help..'. The title 'Ports and Addresses' is centered. Below it is a table with two main sections: 'Listen on IPs and ports' and 'Listen on port'. The 'Listen on IPs and ports' section has a 'Bind to IP address' dropdown menu with 'Anyaddress' selected. The 'Listen on port' section has a 'Specific port..' dropdown menu with '10000' selected and a 'Same as first' dropdown menu. A 'Save' button is at the bottom left. At the bottom right is a back arrow and a link 'Return to webmin configuration'.

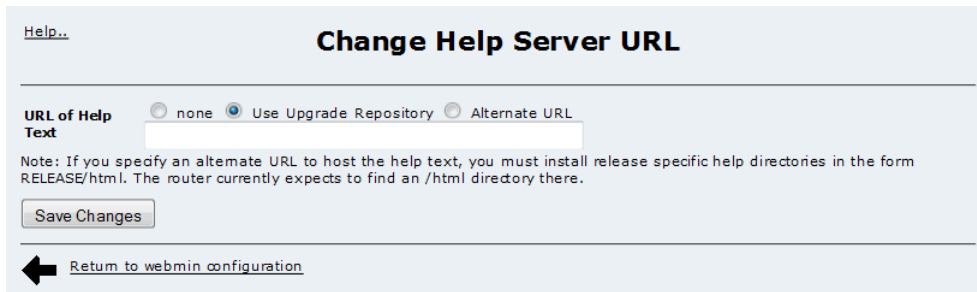
Figure 126: Webmin Configuration Menu, Ports and Addresses

This command allows you to restrict access to Webmin from one particular network interface on your server. If your Webmin server has a non-routable local address and a routable Internet address, you should decide whether anyone will ever need to be able to access the Webmin server from outside of your local network. If not, simply configure Webmin to listen on the local interface.

By default, Webmin listens on TCP port 10000 for clients. It is possible to change this default behaviour.

Section 4.1.3

Change Help Server



The screenshot shows the 'Change Help Server URL' configuration page in Webmin. At the top left is a link for 'Help..'. The title 'Change Help Server URL' is centered. Below it are three radio buttons: 'none', 'Use Upgrade Repository' (which is selected), and 'Alternate URL'. Below the radio buttons is a text input field. A note below the input field reads: 'Note: If you specify an alternate URL to host the help text, you must install release specific help directories in the form RELEASE/html. The router currently expects to find an /html directory there.' A 'Save Changes' button is at the bottom left. At the bottom right is a back arrow and a link 'Return to webmin configuration'.

Figure 127: Webmin Configuration Menu, Change Help Server

The Web management package provides context sensitive help in each of its menus. When a help link is selected the router instructs the browser to open the help text from a help server. In this way the router does not waste large amounts of disk space storing help text and network bandwidth sending large web pages. By default, the router directs the browser to the same server used to upgrade the router. This is as specified in the Maintenance menu Upgrade System sub-menu Change Repository Server command.

This command allows you to disable Web management help, use the upgrade repository server as well as specify a new server. If you specify an alternate web server to host the help text, you must install release specific help directories below the document root. The menu suggests the currently expected directory. The actual help files are provided with every release under the html directory at the repository server.

Section 4.1.4

Logging

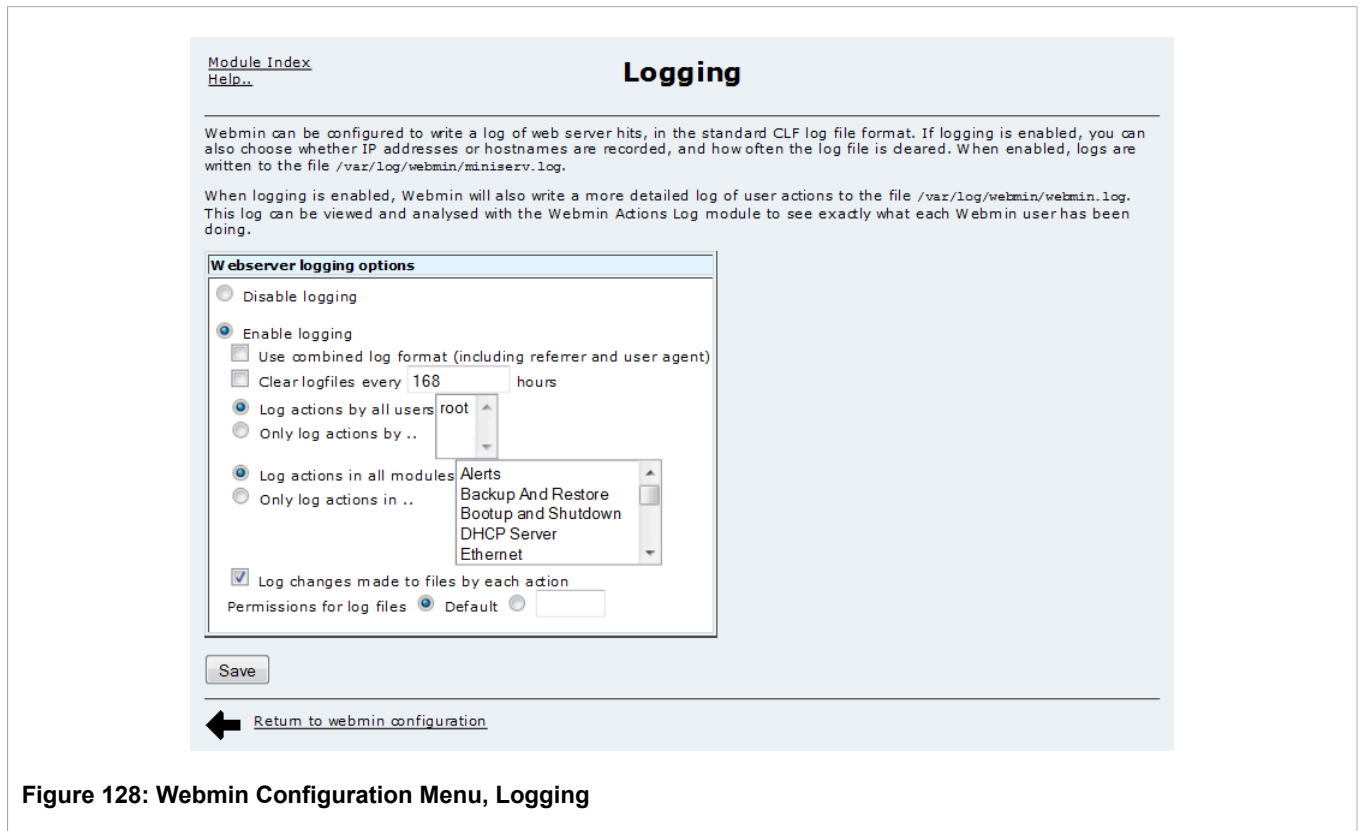


Figure 128: Webmin Configuration Menu, Logging

This menu allows you to log actions taken by Webmin administrators.

It is also possible to log actions based on the module where the actions are performed.

The *Log resolved hostnames* field will cause Webmin to provide a hostname rather than just an IP address for the client computer that performed an action.

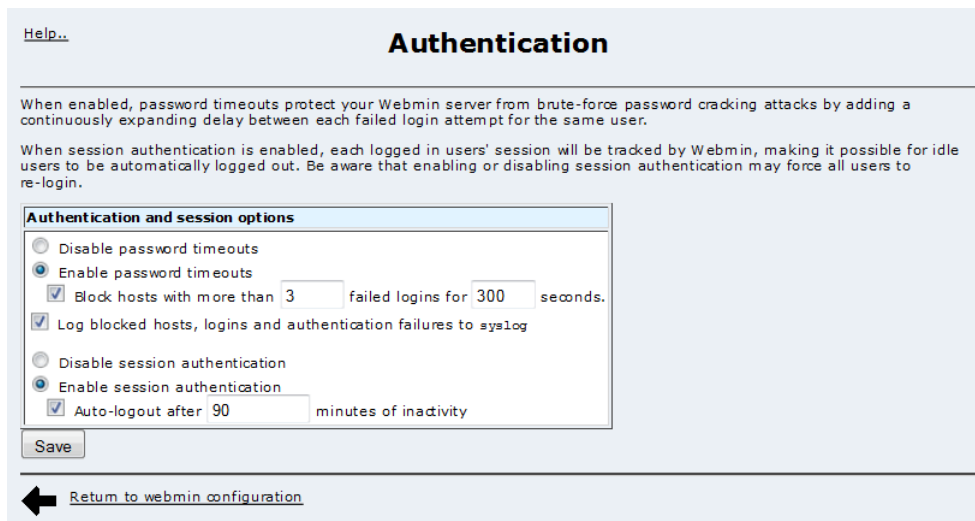
The *Clear logfiles every...hours* field causes Webmin to rotate its own logs and keep them from overfilling the disk with old logs.

Currently, the *Log actions by all users* field should be left selected.

The *Log changes made to files by each action* field causes verbose logging and should be left enabled.

Section 4.1.5

Authentication



The screenshot shows the 'Authentication' configuration page in Webmin. At the top, there is a 'Help..' link and the title 'Authentication'. Below the title, a paragraph explains that password timeouts protect the server from brute-force attacks by adding a delay between failed logins. Another paragraph explains that session authentication tracks logged-in users and can automatically log them out. The main configuration area, titled 'Authentication and session options', contains several settings: 'Disable password timeouts' (radio button), 'Enable password timeouts' (radio button, selected), 'Block hosts with more than 3 failed logins for 300 seconds' (checkbox, checked), 'Log blocked hosts, logins and authentication failures to syslog' (checkbox, checked), 'Disable session authentication' (radio button), 'Enable session authentication' (radio button, selected), and 'Auto-logout after 90 minutes of inactivity' (checkbox, checked). A 'Save' button is at the bottom left, and a 'Return to webmin configuration' link with a back arrow is at the bottom right.

Figure 129: Webmin Configuration Menu, Authentication

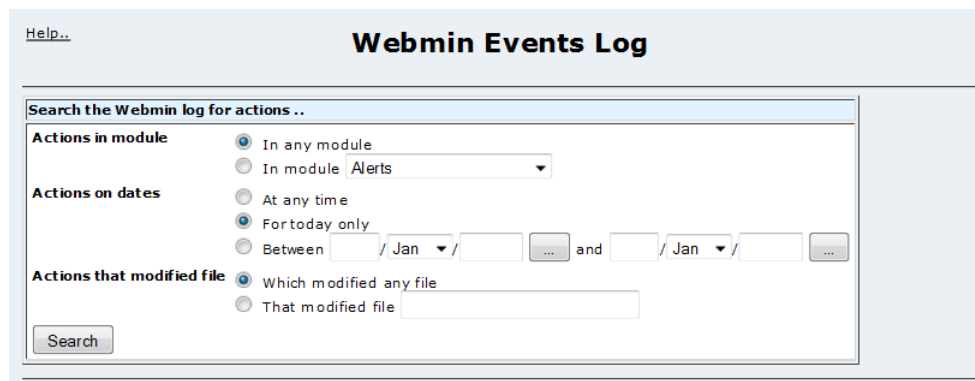
This menu allows you to configure what Webmin will do when a number of failed logins from the same IP address occur.

If the *Enable password timeouts* field is selected, the host will be blocked for the specified period of time. If the *Log blocked hosts, logins and authentication failures to syslog* field is selected, warning messages will be added to the syslog.

Enabling the *Enable session authentication* field, activating "Auto-logout after.." will cause an individual administrators session to be logged out after the specified period.

Section 4.1.6

Webmin Events Log



The screenshot shows the 'Webmin Events Log' search interface. At the top, there is a 'Help..' link and the title 'Webmin Events Log'. Below the title, a search bar is labeled 'Search the Webmin log for actions ..'. The search criteria are organized into three sections: 'Actions in module' with radio buttons for 'In any module' (selected) and 'In module Alerts' (with a dropdown menu); 'Actions on dates' with radio buttons for 'At any time', 'For today only' (selected), and 'Between' (with date pickers for month and year, and a range selector); and 'Actions that modified file' with radio buttons for 'Which modified any file' (selected) and 'That modified file' (with a text input field). A 'Search' button is located at the bottom left of the search criteria area.

Figure 130: Webmin Events Log

This menu allows you to search the Webmin log for changes made by yourself or other administrators.

Section 4.2

Configure Webmin Users

This section familiarizes the user with:

- Configuring Webmin users
- Displaying and removing existing login sessions
- Setting up password restrictions

Section 4.2.1

Webmin User and Group Fundamentals

When the Webmin package is installed for the first time, an account for the user: "root" exists on the router. Besides the root account, three groups, or privilege levels, are defined: "admin", "operator", and "guest".

- Users belonging to the "admin" group have full access to all Webmin modules.
- Users belonging to the "operator" group have full access to most Webmin modules with the following exceptions:
 - Webmin Configuration
 - Webmin Event Log
 - Webmin Users
 - Scheduled Commands
 - Scheduled Cron Jobs
 - System Hostname
 - System Time
 - SSH Server
 - Backup and Restore
 - Upgrade System
 - Upload/Download Files
- Users belonging to the "guest" group can only view configuration and statistics but can not change them. Besides this limitation, they also have no access to the modules forbidden to the "operator" group, listed above.

The "root" user must always be defined. New Webmin users can be created and deleted, and must belong to one of the three aforementioned groups. New Webmin user names must contain only the characters "a-zA-Z0-9-.@", but must not begin with "@" and must not conflict with any existing user or group name.

Section 4.2.2

RADIUS User Access Control Fundamentals

Webmin provides the ability to authenticate against a RADIUS server in order to centralize the creation and maintenance of user accounts. Multiple devices may be configured to authenticate Webmin users using a

common RADIUS server, eliminating the need to replicate the effort of configuring the same user account information on many routers.

If ROX is configured to use RADIUS to authenticate Webmin users (in the Miscellaneous module under the Maintenance category), the router will present the configured RADIUS server with the user name and password presented to a Webmin session for authentication. If the RADIUS server authenticates the user, it will return an indication of success along with the privilege level (described above) associated with the user.

A user successfully authenticated by a RADIUS server will have Webmin access corresponding to his/her privilege level, as configured for the user account on the RADIUS server. For information on how to configure user accounts on a RADIUS server, please refer to [Section 4.7, “RADIUS Server Configuration”](#).

**NOTE**

A Webmin user will only be authenticated locally if a user account of that name has already been created in Webmin.

**NOTE**

The Change Password Command can only be accessed via a locally defined user account.

Section 4.2.3

Webmin Users Menu

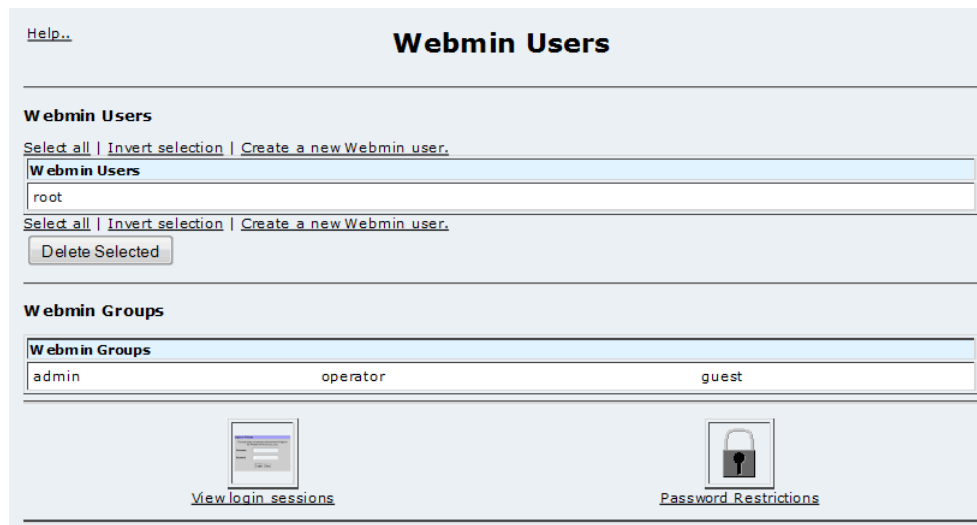


Figure 131: Webmin Users Menu

This menu allows you to create, change or delete a Webmin user, to view and remove current login sessions, and to set password restrictions.

Click the *Select all* link to select all manually created users.

Click the *Invert selection* link to deselect all manually created users.

Click the *Create a new Webmin user* link to create a new Webmin user.

Click on manually created user name to change its setting.

Click the *Delete Selected* button to delete selected users.

Click the *View login sessions* button to view all current login sessions.

Click the *Password Restrictions* button to set the password restriction rules.



NOTE

The accounts managed from this menu are local to the ROX, and are not maintained on a RADIUS server, even if one is configured.

Section 4.2.4

Edit Webmin User Menu

This menu allows you to change the user name, group membership, password, and real name for a user account.

Figure 132: Edit Webmin User Menu



IMPORTANT!

The following usernames are not permitted when creating a new user: root, admin, operator, guest, webmin, and none.

The *Username* field sets the user name for the Webmin user. This user name will be used in the login.

The *Member of group* field determines which group the user belongs to. Recall that the group is equivalent to the privilege level, which determines the user's access level for the Webmin system.

The *Password* field sets the password for the user.

The *Real name* field sets the real name for the user.

The *Save* button will save the changes permanently.

The *View Logs* button will display the action logs for this Webmin user.

The *Delete* button will delete the current user from Webmin.

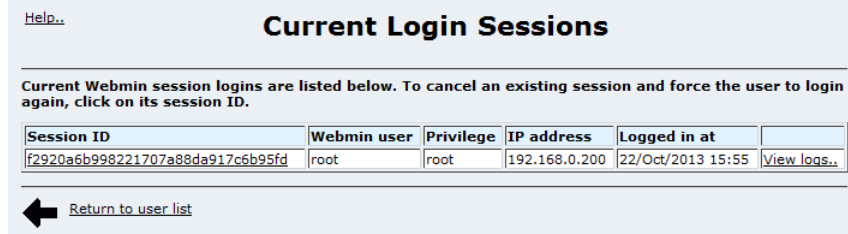


NOTE

When a Webmin user is deleted, any current session the user may have established will be terminated.

Section 4.2.5

Current Login Sessions Menu



Help..

Current Login Sessions

Current Webmin session logins are listed below. To cancel an existing session and force the user to login again, click on its session ID.

Session ID	Webmin user	Privilege	IP address	Logged in at	
f2920a6b998221707a88da917c6b95fd	root	root	192.168.0.200	22/Oct/2013 15:55	View logs..

← [Return to user list](#)

Figure 133: Current Login Sessions Menu

This menu allows you to view and delete current login sessions (delete login session will force the login user to login again).

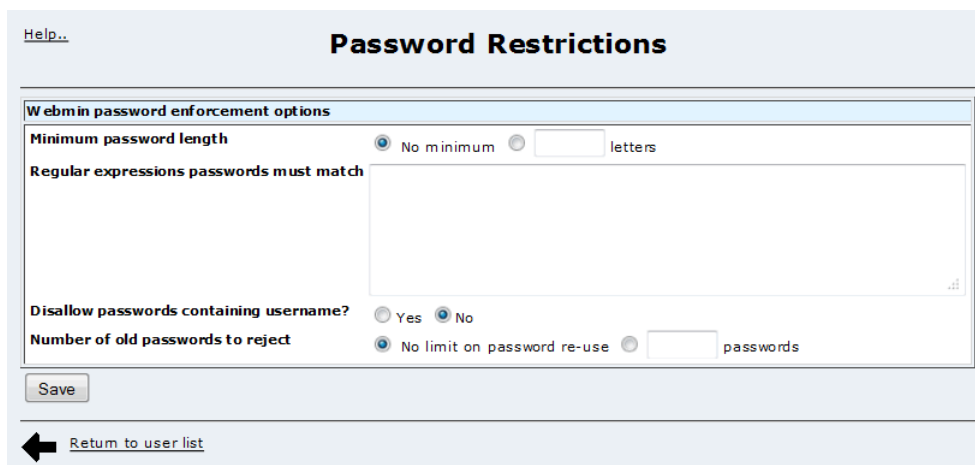
Click on *Session* link under the Session ID column to cancel a session.

Click the *Webmin user* link under the Webmin user column to display the Webmin user edit menu for that user.

Click the *View logs* link to display logs for that Webmin user.

Section 4.2.6

Password Restrictions Menu



Help..

Password Restrictions

Webmin password enforcement options

Minimum password length ☒ No minimum ☐ letters

Regular expressions passwords must match

Disallow passwords containing username? ☐ Yes ☒ No

Number of old passwords to reject ☒ No limit on password re-use ☐ passwords

← [Return to user list](#)

Figure 134: Password Restrictions Menu

This menu allows you to set restrictions for password selection in order to prevent the use of trivial, or machine-guessable passwords.

**IMPORTANT!**

Password restrictions do not apply to passwords for the root and rrsetup profiles.

The *Minimum password length* field sets the minimum length for password.

The *Regular expression passwords must match* field sets the regular expression that a new password must match. The above example restricts new passwords to begin with an alpha character followed by at least another 5 alphanumeric characters.

The *Disallow passwords containing username* field prevents new passwords from containing the user name.

The *Number of old passwords to reject* field determines after how many successful passwords settings you are allowed to reuse an old password.



NOTE

The password restriction mechanism in ROX uses PERL regular expression syntax. For the definitive reference documentation on regular expressions in PERL, refer to:

- <http://perldoc.perl.org/perlref.html>
- <http://perldoc.perl.org/perlre.html>
- <http://perldoc.perl.org/perlretut.html>
- <http://perldoc.perl.org/perlrequick.html>

If you do not have access to an Internet connection, but do have a UNIX/Linux system with PERL installed, access the local manual pages by typing on of the following commands at the command line:

- *man perlref*
- *man perlre*
- *man perlretut*
- *man perlrequick*

Root privilege is not required to access manual pages.

Section 4.3

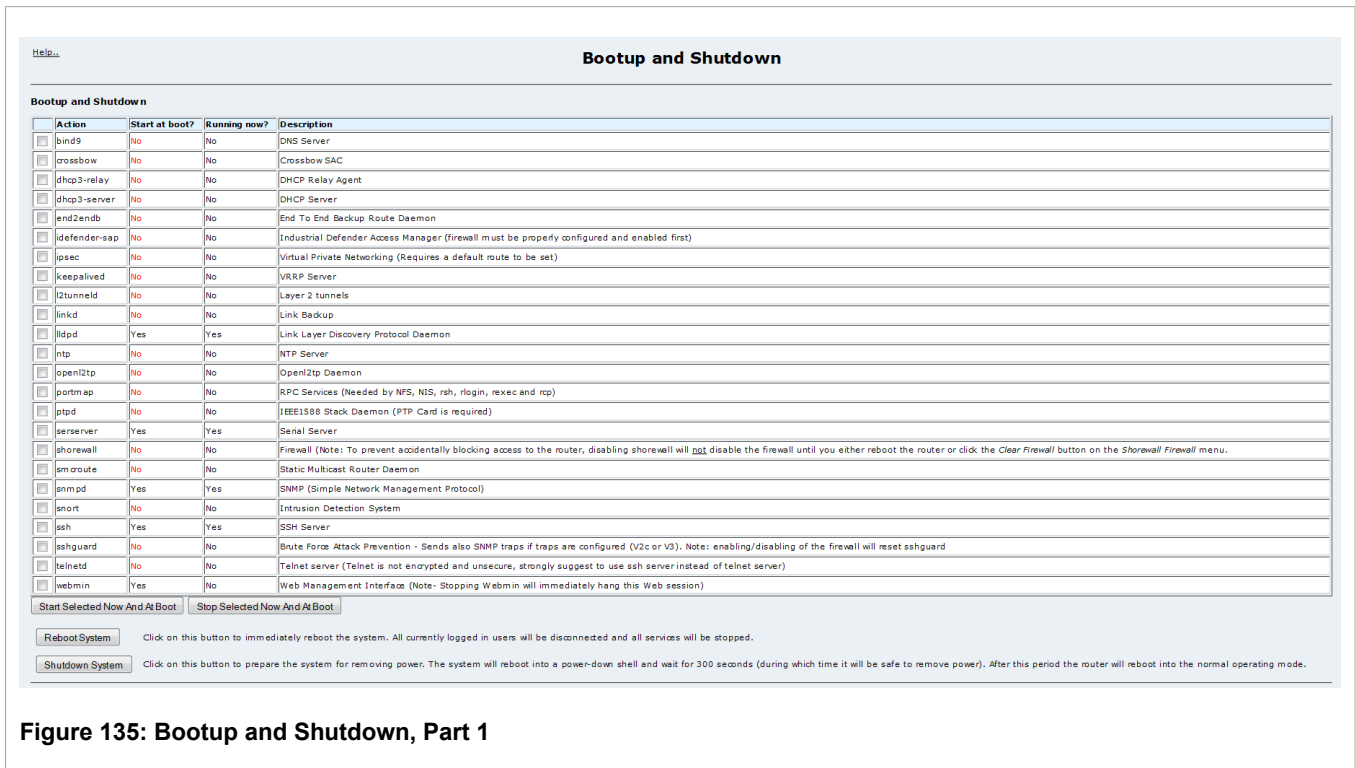
Configuring the System

This section familiarizes the user with:

- Enabling and disabling processes such as SSH and Web Management
- Changing passwords
- Shutting down and rebooting the system
- Scheduling one-off and periodic commands
- Examining system logs
- Changing the hostname
- Changing the system time and timezone

Section 4.3.1

Bootup and Shutdown

**Figure 135: Bootup and Shutdown, Part 1**

This menu allows you to enable/disable services and to perform actions at boot. The first part of the menu manages services. Check the box for the desired service and click on *Start Selected* to start the service and have it start at the next boot. Click on *Stop Selected* to stop the service and not have it start at boot.

The *Reboot System* button will cause the system to reboot.

The *Shutdown System* button shuts down the system in order to remove power.

**NOTE**

The device never enters a permanent shutdown state. If the device is instructed to shutdown, either from Webmin or from a shell command, it will reboot into a command line shell that waits five minutes before restarting.

If you really want the router to remain powered but permanently inactive, you must issue the shutdown, connect a terminal to the serial port, wait for the router to enter the shutdown shell and issue a CTRL-C.

The second part of the menu allows you to program specific actions at boot time. The script will be run after all regular boot actions have completed.

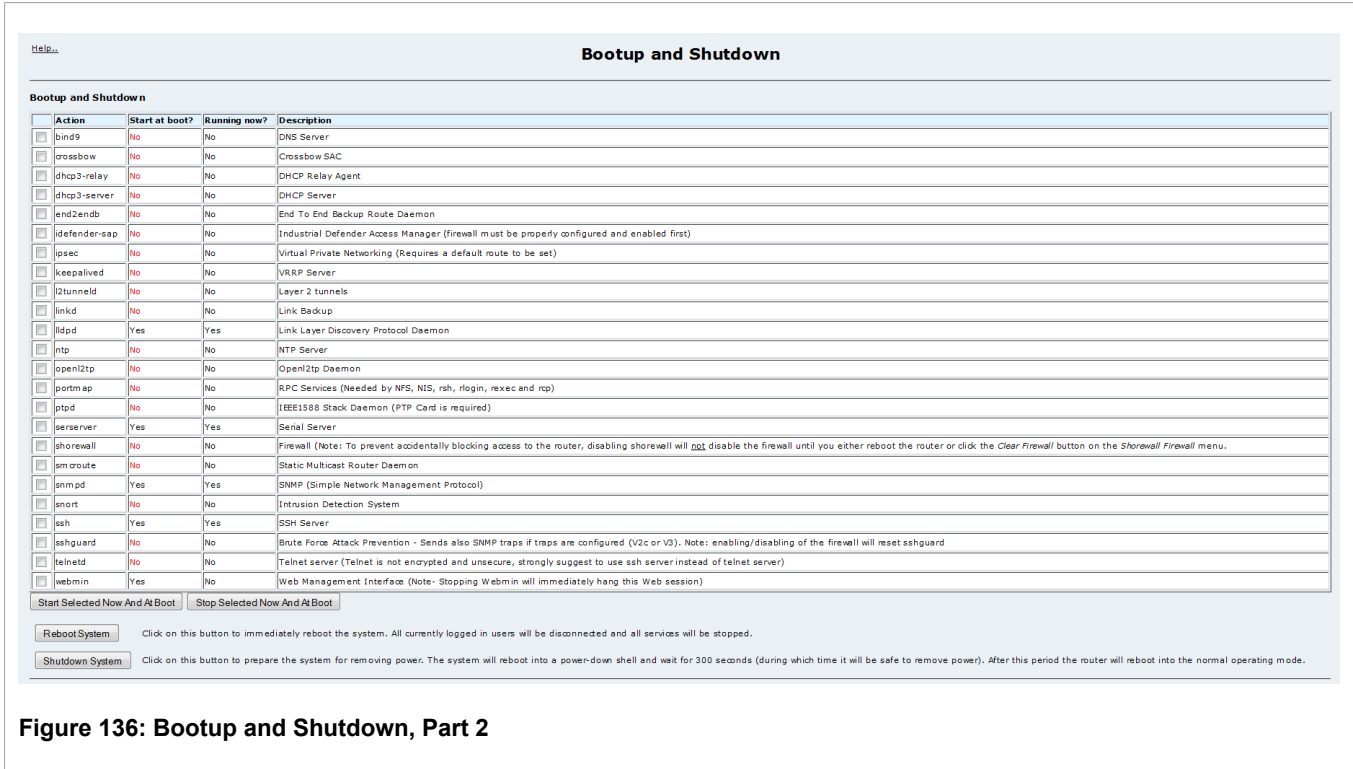


Figure 136: Bootup and Shutdown, Part 2

The actions may be a series of commands that can be executed at the command line. Each entered line is executed independently of the previous line, so change directory commands will not be effective. Always specify the absolute path of files used in commands. Selecting *Save and Run Now* will run the script and show its output, allowing you to debug it.

Section 4.3.2

Configuring Passwords

This section familiarizes the user with:

- Changing the system password
- Changing the bootloader password

Section 4.3.2.1

Change Password Command

This command changes only the root account password used to login to Webmin and the root account via the serial console or SSH.

[Help..](#)

Change Password

Change Password for user "root"

This module can be used to change the root password used to login with webmin, ssh and console..

root login password

☒ Leave unchanged ☐ Set to .. Re-enter

Save

Figure 137: System Menu Change Password Command

Section 4.3.2.2

Change Bootloader Password Command

This command changes the bootloader password used to log in to the available service modes.



CAUTION!
Security hazard – risk of unauthorized access. To prevent unauthorized personnel from accessing the available boot modes, it is strongly recommended that a bootloader password be set before the device is deployed.

[Help..](#)

Change Password

Change Password for boot loader

This module can be used to change the password (maximum 128 characters) used in boot loader

Boot loader password is not set

Boot loader password

☒ Leave unchanged ☐ Clear ☐ Set to .. Re-enter

Save

Figure 138: Bootloader Change Password Command

Section 4.3.3

Scheduled Commands

[Help..](#)

Scheduled Commands

Job ID	Run as user	Run at	Created on	Commands to execute
3	root	Fri Oct 11 01:00:00 2013	Thu Oct 10 02:50:12 2013	reboot

Figure 139: Scheduled Commands

This menu allows you to schedule a command to run in the future.

Begin by selecting the time and date you wish to run the command at using the *Run on date* and *Run at time* fields.

Use the *Run in directory* field to enter a directory to run the command in, or simply use "/".

Finally, enter the command to execute in the *Commands to execute* field.



NOTE

The command will remain scheduled after reboot. After the command is entered, the Scheduled Commands menu will display any commands and allow you cancel them.

Figure 140: Scheduled Commands Displaying a Command

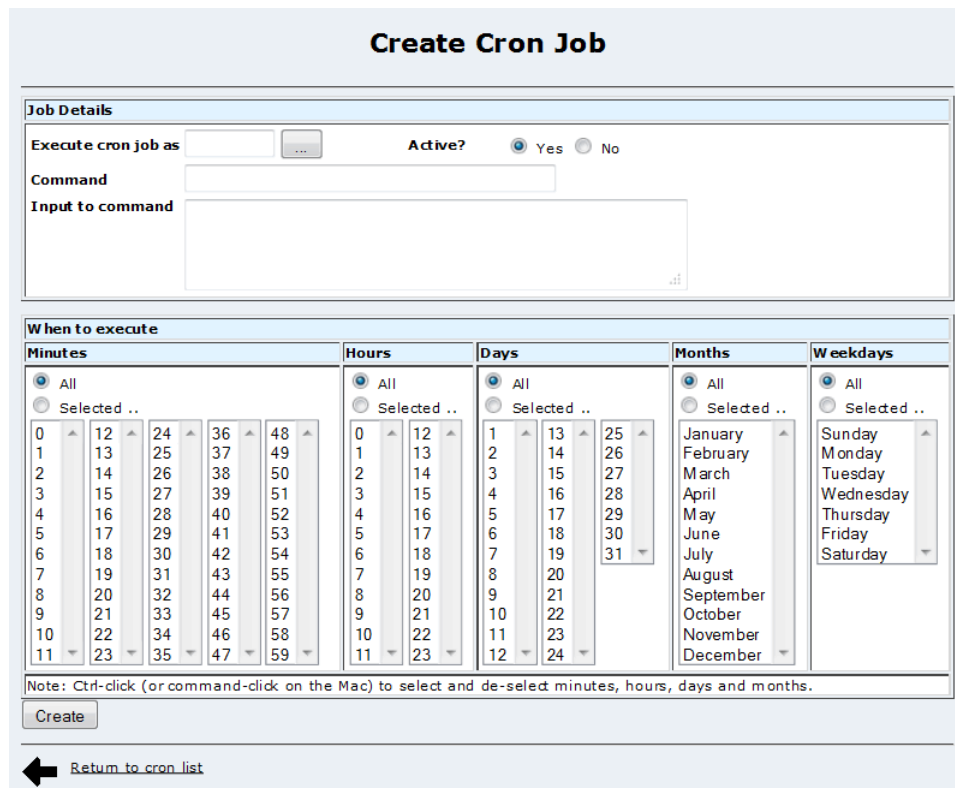
Section 4.3.4

Scheduled Cron Jobs

"Cron" is a service that allows flexible, regular scheduling of system commands. A "Cron job" is the set of a command to run and a definition of the times at which to run it. The *Scheduled Cron Jobs* menu allows you to create, edit, and delete these jobs.

Figure 141: Webmin Scheduled Cron Jobs

Initially, there will be no scheduled jobs. Follow the *create* link to create one.



Create Cron Job

Job Details

Execute cron job as: Active? ☒ Yes ☐ No

Command:

Input to command:

When to execute

Minutes	Hours	Days	Months	Weekdays
<input checked="" type="radio"/> All <input type="radio"/> Selected ..	<input checked="" type="radio"/> All <input type="radio"/> Selected ..	<input checked="" type="radio"/> All <input type="radio"/> Selected ..	<input checked="" type="radio"/> All <input type="radio"/> Selected ..	<input checked="" type="radio"/> All <input type="radio"/> Selected ..
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59	0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23	1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31	January February March April May June July August September October November December	Sunday Monday Tuesday Wednesday Thursday Friday Saturday

Note: Ctrl-click (or command-click on the Mac) to select and de-select minutes, hours, days and months.

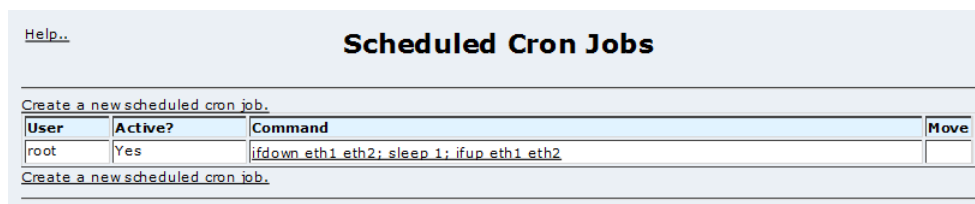
Figure 142: Creating a Cron Job

Begin the creation of a cron job specification by selecting a *user* to execute as. For most purposes, "root" will suffice. Enter the user name in the *Execute cron job as* field.

Enter the command to execute and any input to the command in the *Command* field. Select the times the script is to run from the *When to execute* table (remember to check the *selected* button above any column you edit).

The *Active* radio button at the top of the menu temporarily disables the job.

After selecting the *Create* button, the Scheduled Cron Jobs menu will display the job.



[Help..](#)

Scheduled Cron Jobs

[Create a new scheduled cron job.](#)

User	Active?	Command	Move
root	Yes	ifdown eth1 eth2; sleep 1; ifup eth1 eth2	<input type="button" value="Move"/>

[Create a new scheduled cron job.](#)

Figure 143: Scheduled Cron Jobs Menu Displaying Cron Jobs

Follow the link of a specific job in order to delete the job, edit it, or test the command part of the job by running it immediately.

If you have multiple jobs, the arrows in the *Move* column will alter the order in which they are presented.

Section 4.3.5

System Hostname

Figure 144: System Hostname

The *Hostname* field modifies the hostname as presented in the web server and shell sessions.

The *Domain* field modifies the domain as presented in the web server and shell sessions. The default is "localdomain".

Note that the new hostname and domain settings will only appear in new sessions.

Section 4.3.6

System Time

Figure 145: System Time

This menu provides a method to set the router's time and timezone.

**NOTE**

OSPF and RIP are sensitive to accurate system time. If OSPF or RIP are enabled, changing the time from this menu will cause them to be restarted.

Section 4.4

Managing SSH Keys and Certificates

The following sections describe how to manage SSH certificates and keys on the device:

- [Section 4.4.1, “Uploading SSL Keys and Certificates”](#)
- [Section 4.4.2, “Regenerating SSL Keys and Certificates”](#)
- [Section 4.4.3, “Regenerating SSH Keys”](#)

Section 4.4.1

Uploading SSL Keys and Certificates

The SSL file used in ROX I is a file that contains both the SSL private key and the SSL Certificate. Both are in PEM format. See example below.

1. Once the SSL certificate and key has been created and is available in the correct format, log in to the ROX web interface.
2. Navigate to Maintenance > Upload/Download Files.
3. Under 'Send files from your current host to router', choose the file to upload. In the 'File or directory to upload to' textbox type the following:

```
/etc/webmin/miniserv.pem
```

4. Click 'Send to router'.
5. Reboot the ROX device for Webmin to start using the new certificate.

The SSL Key/Certificate format in ROX I looks as shown in the example below.

**NOTE**

This is only an example. Do not use this key and certificate combination.

```
-----BEGIN CERTIFICATE-----
MIIC+DCCAmGgAwIBAgIJAL0J8uF/HwWXMA0GCSqGSIb3DQEBBQUAMIGuMQswCQYD
VQQGEwJDQTEQMA4GA1UECBMT250YXJpbzEQMA4GA1UEBxMHQ29uY29yZDESMBAG
A1UEChMJUnVnZ2Vky29tMRkwFwYDVQQLExBDdXN0b211ciBTdXBwb3JOMSYwJAYD
VQQDEw1WTS1WSVZFSy1URVNUUL1JVR0dFRENPTS5MT0NBTDekMCIGCSqGSIb3DQEJ
ARYVc3VwcG9ydeBydWdnZWRjb20uY29tMB4XDTEzMDUxNzIwMzkwNVoXDTE4MDUx
NjIwMzkwNVowZz4xCzAJBgNVBAYTAkNBMR4wDgYDVQQIEwdPbnRhcmlvMRAwDgYD
VQQHEwdDb25jb3JkMRIwEAYDVQQKEw1SdWdnZWRRb20xGTAxBgNVBASTEEN1c3Rv
bWVyIFN1cHBvcnQxYjAUBgNVBAMTDTE3Mi4zMCA4xNDguNDYxJDAiBgkqhkiG9w0B
CQEFVFN1cHBvcnRAcnVnZ2Vky29tLmNvbTCBnzANBgkqhkiG9w0BAQEFAAOBjQAw
gYkCgYEA3HpxLytAc/CT6gJwqWjENHUuleM11TDZt4nRcu+/JLgSssPzl9+IYYE
Gv3YoI/Ep4qqrVXwpHlOmsYlCmmaIbPFACscmLFjLzWmM/MKNDGnhSLZozArxAq
2TlsUkITmNdIdC6jKNtC8q9e674yKIS9BSMsBXa+Wi+BUdLN+K0CAwEAAAMsMCow
CQYDVR0TBAlwADAdBgNVHQ4EFgQUiijRkgmRox5jYFsRlmo7Ex8Em2swDQYJKoZI
hvcNAQEFBQADgYEAerlOmJ8YiW/OqIXh79NCaByOGuReguljLtxrwkclfTEQ626
+/kI7w90uH1TiJtcJPKdOjDqAva3eZ23VZL43t/d0jIEL8xdIBXwDtEUyqNpanPO
```

```

YVnWnMJZ2EacUlr50ONOB7K+IMjvZm0Nrrazt/9KEZ15pP/rDSXBAqnumCw=
-----END CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
MIICXQIBAAKBgQCvcenEvK0Bz8JPqAnCpaMQ0dRSV4zXVMNm3idFy778kuBKyw/O
X34hhgQa/digj8Sniqq9fCkeU6axiUJyaZohs8UBxyYsWOXNaYz8wqcMaeFItn
jMCvECrZOWxS9hOY0h18LqMo20Lyr17rvjIohL0FIywFdr5aL4FR0s34rQIDAQAB
AoGAG/YuyME9XZWMJX/1l1UpyQt4KGt1sff7cJCld7VCSiThqGEassx7YMJoMxkU
BnsDX6R1EY92+++tfHNqAcZWs/x9pUKTNekAb9CglVs/tdubOnYyIhL6TqfPNsfE
HCP58GqaBFqKYhy/1JL6N+tNknghazvdUc9wPOXn/1f4DY0CQQDZDcCt6kGeGFit
rytBBCf/bz5WzbeepqiA3oxXnMUN8uj9e1Kq3NfbhfGtJ9yK8NF1ReqLgiW0v3cA
jdq24WqbAkEAzuzigJWiBUyZMMYs4bgf9/5UdNww2fIVmwqNRMHhgOOTUCAXm/jK
JTh5XD4xslL3ttFG7EGYRR6R258yMI0aVwJBAKuY98+G98FpNhJ5/hQ0mPqUlBBE
Kvq+f1EC0K1TQ2a3uANOUBjm58qhpMnitDUUFkREthz9E5pGFGrXuyYCKCkCQQC8
r2UaulcyXdaSkyL58Fu2V0PMC7y//++ToNuQhwrzzJDXz2u33fT2W7jOVCgc42re
WZbCeE3ROT7ndRLfEsubAkAu6H27QMI73kjEaIbUqKfJXeln6Ca7Pt6Za4ShZRkW
kDHo9Q2p46T4VKXstNgk6WY36gUrb/pqXEJAXla0yqgy
-----END RSA PRIVATE KEY-----

```

Section 4.4.2

Regenerating SSL Keys and Certificates

If it is not possible to provide certificates and keys to ROX from a proper X.509 system, then Siemens recommends that the user regenerate the ROX keys and certificates on a regular basis following the steps below. These certificates are self-signed.



NOTE

For security reasons, it is highly recommended that proper X.509 certificates signed by a Certificate Authority (CA) be used. If a certificate is not signed by a CA and is self-signed, the trust portion of the certificate cannot work because the Certificate owner is essentially its own CA.

1. Copy the following text into a plain text editor.

```

cd /opt/ && openssl genrsa -out ./CA.key 2048 &&
openssl req -x509 -new -config /etc/ssl/openssl.cnf -subj "/C=CA/ST=Ontario/L=Concord/O=RuggedCom/
OU=Support/CN=hostname/" -days 1825 -key ./CA.key -out ./CA.crt &&
cat ./CA.crt ./CA.key > /etc/webmin/miniserv.pem &&
rm CA* && /etc/init.d/webmin restart

```

2. Replace the text 'CN=hostname' with the system's hostname or primary management IP address.
3. Login to the unit using SSH. See [Section 3.1.3, “Accessing the Device Command Prompt From the Console Port”](#).
4. Copy and paste the text and hit Enter, if required. The output should look like the text below:

```

RX1000:~#
RX1000:~# cd /opt/ && openssl genrsa -out ./CA.key 2048 &&
openssl req -x509 -new -config /etc/ssl/openssl.cnf -subj
"/C=CA/ST=Ontario/L=Concord/O=RuggedCom/OU=Support/CN=RX1000/" -days 1825 -key
./CA.key -out ./CA.crt &&
> cat ./CA.crt ./CA.key > /etc/webmin/miniserv.pem &&
> rm CA* && /etc/init.d/webmin restart
Generating RSA private key, 2048 bit long modulus
.....++++++
.....++++++
e is 65537 (0x10001)
Restarting webmin: webmin.
RX1000:/opt#

```

5. Open Webmin in any modern browser and verify that the certificate is a new one. The 'not valid before' value could be helpful.

Section 4.4.2.1

Generating Self-Signed SSL Certificates with Scripting

Creating a persistent script to generate self-signed SSL keys and certificates is possible with the commands referenced in [Section 4.4.2, “Regenerating SSL Keys and Certificates”](#).

1. Open a text editor such as notepad and copy the following text to it.

```
#!/bin/bash
hn=$(hostname) &&
cd /opt/ && openssl genrsa -out ./CA.key 2048 &&
openssl req -x509 -new -config /etc/ssl/openssl.cnf -subj "/C=CA/ST=Ontario/L=Concord/O=RuggedCom/OU=Support/CN=$hn/" -days 1825 -key ./CA.key -out ./CA.crt &&
cat ./CA.crt ./CA.key > /etc/webmin/miniserv.pem &&
rm CA* && /etc/init.d/webmin restart
```

2. Save the file as 'renewkey' without a filetype extension.
3. Open Webmin in a browser and navigate to the Upload/Download Files feature.

The screenshot shows the 'Upload/Download Files To The Router' interface in Webmin. It has a light blue header with a 'Help..' link. The main content area is divided into three sections, each with a light blue header. The first section, 'Download files from the specified URLs to this router', contains a large text area for 'URLs to download', a text field for 'File or directory to download to' with a browse button, and radio buttons for 'Download mode' (selected: 'Immediately, and show progress'; 'In background, at date 29 / Dec / 2002 and time 13 : 46'). A 'Download URLs to router' button is at the bottom. The second section, 'Send files from your current host to the router', contains four 'Browse...' buttons for 'Files to upload' (each with 'No file selected.' text), a 'File or directory to upload to' field with a browse button, and radio buttons for 'Extract ZIP or TAR files?' (selected: 'No'; 'Yes, then delete'; 'Yes'). A 'Send to router' button is at the bottom. The third section, 'Download a file from the router to your host', contains a file selection field with a browse button and a 'Download to your host' button.

Figure 146: Upload/Download Menu

4. Under Send files from your current host to the router, click on any of the buttons that say 'Browse' and select the file that was just created.
5. Set the option 'File or directory to upload' to '/usr/bin' and click 'Send to router'.
6. Once the upload is successful, open a SSH connection or serial console connection to ROX.
7. Type the command 'chmod 700 /usr/bin/renewkey' and hit Enter.
8. Type the command 'dos2unix /usr/bin/renewkey' if the file was created using a Windows text editor.
9. The script is now ready for use. Issuing the command 'renewkey' will regenerate a new pair of RSA keys and a certificate for Webmin. The output looks as below:

```
RX1000:~# renewkey
Generating RSA private key, 2048 bit long modulus
.....
+++
...+++
e is 65537 (0x10001)
Restarting webmin: webmin.
RX1000:~#
```

Section 4.4.3

Regenerating SSH Keys

Regenerate the SSH keys on the ROX device using the following procedure:



NOTE

Access through a physical connection to the console port is strongly recommended for this procedure. See [Section 3.1.3, “Accessing the Device Command Prompt From the Console Port”](#).

1. Delete the current SSH key by typing the following command and then pressing Enter:

```
rm /etc/ssh/ssh_host_*_key*
```

2. Generate a new SSH key and restart SSH by typing the following command and then press Enter:

```
/var/lib/dpkg/info/openssh-server.postinst configure
```

3. The following messages will appear as the script runs:

```
Creating SSH2 RSA key; this may take some time ...
Creating SSH2 DSA key; this may take some time ...
Restarting OpenBSD Secure Shell server: sshd.
```

Section 4.5

Access Manager Secure Access Portal

RX1100 owners can use Access Manager's Secure Access Portal (SAP) to restrict access to critical assets. This section details how to activate the Secure Access Portal and determine currently negotiated sessions.



NOTE

Please note that when doing firewall configuration changes while the SAP is enabled, that the SAP must first be stopped, and then explicitly restarted after the firewall configuration changes are re-enabled.



IMPORTANT!

All firewall disabling and re-enabling done using the command line must be through one of these commands:

```
/etc/init.d/shorewall clear
/etc/init.d/shorewall stop
/etc/init.d/shorewall restart
```


Details and recommendations on applying the Access Manager system to networking may be found in texts referred to in [the section called “Related Documents”](#) of the user guide. The Access Manager is configured through the Industrial Defender Main Menu, see figure below.

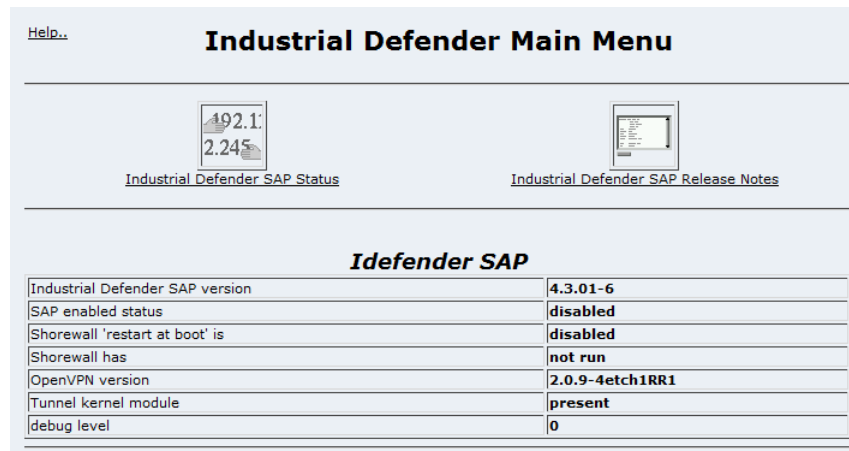


Figure 147: Access Manager Configuration Menu

Section 4.5.1

What Access Manager's Secure Access Portal Protects and How

The Secure Access Portal protects against unauthorized access to critical assets, including the router itself. The Secure Access Portal allows connection through an openVPN tunnel from known management devices to assets behind the firewall operating on known TCP/UDP port numbers.



NOTE

When restoring a previously saved configuration, the Secure Access Portal will have to be reconfigured.

Section 4.5.2

Access Manager and the Firewall

Access Manager integrates tightly with the firewall, opening it for communications between vetted clients and critical assets on a demand basis. There are four steps required to activate Access Manager's Secure Access Portal:

Step 1 of 4 : Access Manager Configuration

1. Use Access Manager to Create a Secure Access Portal (SAP). See the Industrial Defender Access Manager User Manual for details or use Help after connecting to the Access Manager.
2. Using Access Manager, authorize user(s) for defined devices behind a router (SAP).

Step 2 of 4 - Shorewall Configuration

1. Use the RX1000 Installation Guide and ROX User Guide to set up the RX1100 and gain access to rsetup via the console port. Use *Restore a Previous Configuration to Reload Factory Defaults*. Change passwords, port IP address information, set the hostname, and set the date, time and time zone. Do not use *Iddefender Setup* at this time.

From a web browser, access the ROX Webmin user interface to perform the rest of the steps:

2. Visit the *Shorewall Firewall* menu, *Network Zones* sub-menu and add the "net" and "loc" IPv4 zones. This document defines the zone for WAN interfaces as "net" and the zone for local interfaces as "loc".

Zone ID	Zone type
net	IPv4
loc	IPv4
acInt	IPv4
unusd	IPv4
fw	Firewall System

3. Visit the *Network Interfaces* sub-menu and assign interfaces to the zones. For example, eth1 = net, eth2 = loc. The exact assignment will depend upon your configuration.



NOTE

The assignment of the "acInt" zone: Industrial Defender SAP uses OpenVPN for secure communication between client and protected device. OpenVPN creates virtual "tunnel" interfaces for this purpose.

Zone ID	Interface	Address
net	eth1	detect
loc	eth2	detect
unusd	eth3	detect
unusd	eth4	detect
acInt	tun+	detect

4. Visit the *Default Policies* sub-menu and assign the following policies:

Source zone	Destination zone	Policy
fw	any	ACCEPT
loc	net	ACCEPT
acInt	any	DROP
any	any	DROP

5. Visit the *Firewall Rules* sub-menu and assign the following rules.



NOTE

Iddefender and *SAPCtl Actions* must have "log to syslog level" set to "<Don't log>".

Action	Source zone	Destination zone	Protocol	Src-Port	Dst-Port
ACCEPT	acInt	fw	any		

Action	Source zone	Destination zone	Protocol	Src-Port	Dst-Port
SAPCtl	net	fw	UDP		
Idefender	acInt	loc	any		

See also the note on VRRP, Firewall Rules, and Access Manager, below.

6. Apply the Shorewall configuration.

**NOTE**

Granting uncontrolled accesses to the router is not required in normal operation. This is a security risk and should not be done without good reason. Rules are order dependent, and so place this rule above the SAPCtl and Idefender rules.

- a. For Webmin and/or SSH access to the router you can add a rule:

ACCEPT net	fw	tcp	22,10000
------------	----	-----	----------

- b. The order of these rules is significant. Any rules added after the Idefender rules may not get processed. Rules inserted before the Idefender rules may compromise the security provided by Idefender. Contact Siemens Customer Support for assistance if you wish to add other rules.
7. Using Webmin, visit the *Bootup and Shutdown* menu and ensure that Shorewall is enabled to start at boot. Start Shorewall. Webmin access is now blocked until secure access through Access Client is opened.

Step 3 of 4 - SAP Configuration

Use *rrsetup* to define a passphrase and required setup parameters and to enable the portal.

1. Select the *required parameters* menu option and enter the "unit name" assigned to this router, and the IP address of the Access Manager which will control it.
2. Select the *SAP Passphrase* menu option and enter a valid passphrase.
3. Select the *Enable Idefender SAP* menu option to enable the software. If it is already enabled then the menu option will say "Disable Idefender SAP".

**NOTE**

Idefender-SAP can also be enabled via Bootup and the shutdown menu in the Webmin user interface.

**NOTE**

The unit name and passphrase entered at the router and the Access Manager must match or else the Access Manager will refuse to acknowledge the router. The unit name and passphrase are both case-sensitive.

Step 4 of 4 - Verification Test

1. To verify that the Access Manager, SAP and client are functioning, you can now use Access Client with a user account to connect to an authorized device.

**NOTE**

For a detailed discussion of the configuration and use of Secure Access Portals from the Access Manager's point of view, please consult the Access Manager User Manual, under Managing Secure Access Portals.

Section 4.5.3

VRRP, Firewall Rules and Access Manager

It may be necessary to specify additional firewall rules in order that Access Manager's SAP (Secure Access Portal) be able to access certain protocols, such as VRRP, on the router. If, for example, the router is configured to be a member of a VRRP Virtual Router Group, it must be able to accept VRRP communication from its peers. The following firewall rule must be added after the ACCEPT rules to UDP ports 30000 and 30001 and before the rules under Access Manager control:

Action	Source zone	Destination zone	Protocol	Source Ports
ACCEPT	net	fw	VRRP	any

The order of the firewall rules is significant. Any rules that are entered after the Access Manager rules may not be processed. Rules inserted before the Access Manager rules may compromise the security provided by Access Manager.

**NOTE**

Exposing any protocol or networked service has the potential of being a security risk and should not be done without good reason. Contact Siemens Customer Support for assistance if you wish to add other rules to the set recommended here.

Section 4.5.4

Access Manager's Secure Access Portal Status Menu

Access Manager integrates tightly with the firewall, opening it for communications between vetted clients and critical assets on a demand basis.

The status menu provides a list of validated open connections.

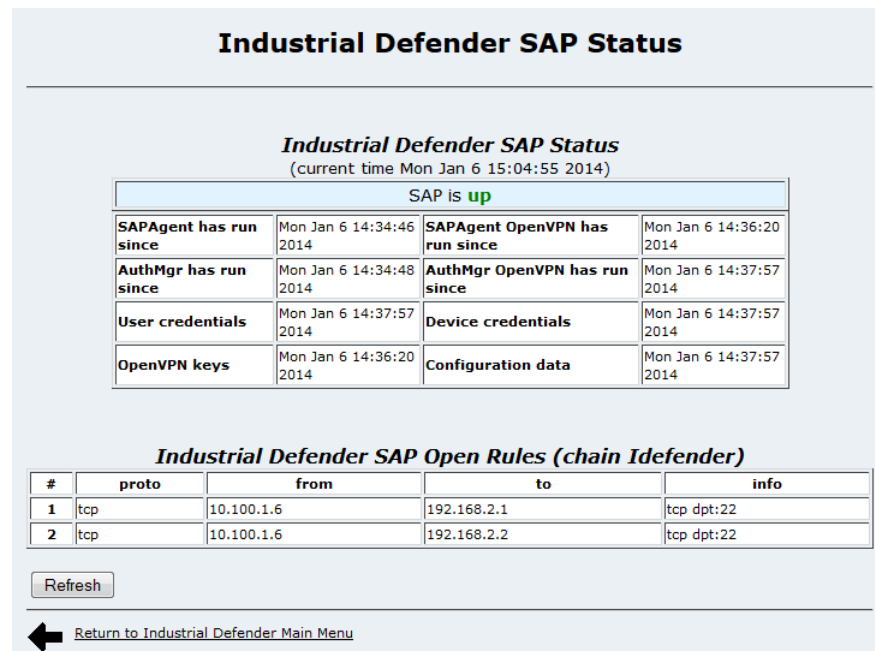


Figure 148: Access Manager's Secure Access Portal Status

Section 4.5.5

Upgrading the Access Manager's Secure Access Portal

During a ROX upgrade, all existing Access Manager Secure Access Portal protected connections will be closed.

Section 4.6

RADIUS Authentication

RADIUS (Remote Authentication Dial In User Service), described in RFC 2865, is a protocol designed to allow the centralization of authentication, authorization, and configuration of various types of services. The goal of RADIUS authentication is typically to restrict the distribution of account information and to avoid the replication of security management effort.

Section 4.6.1

RADIUS Usage

The typical mode of operation involves a Network Access Server (NAS) - in this case the ROX - and a remote RADIUS server, where account information is stored. In the course of attempting to access connection-oriented services on the NAS, a user presents credentials to the NAS for authentication. The NAS forwards these to a configured RADIUS server and accepts from it the determination of whether the user is allowed the requested access. In order to protect the security of account information and of both the NAS and the RADIUS server, transactions are encrypted and authenticated through the use of a shared secret, which is never sent in the clear.

Some administrators set the passwords of existing ROX accounts, e.g. "rrsetup" and "root", uniquely for each router, and then employ a common password per account for all routers served by RADIUS. The router-specific passwords are restricted to a very few personnel. A larger set of expert users is granted the rights to SSH login using the RADIUS root account passwords. Yet another set of users is granted access via Webmin user accounts.

Section 4.6.2

RADIUS on ROX

ROX supports RADIUS server redundancy. Multiple RADIUS servers, usually operating from a common database, may be used to authenticate a new session. If the first configured RADIUS server does not respond, subsequent servers will be tried until a positive/negative acknowledgment is received or an attempt has been made to contact all configured servers.

Each server is configured with an associated timeout which limits the time that ROX will wait for a response. An authentication request could thus require up to the sum of the timeouts of all configured servers.

RADIUS authentication activity is logged to the authorization log file, "auth.log". Details of each authentication including the time of occurrence, source and result are included.

Section 4.6.3

RADIUS, ROX and Services

RADIUS provides the means to restrict access on a per-service basis. Accounts may be configured on a RADIUS server to be allowed access only to the Webmin service, for example. ROX supports RADIUS authentication for the following services:

- *LOGIN*
- *PPP*
- *WEBMIN*

ROX provides the option of designating different servers to authenticate LOGIN, PPP or WEBMIN services separately or in combination.

The LOGIN Service

The *LOGIN* service consists of the following types of access:

- Local console logins via the serial port and modem
- Remote shell logins via SSH and Telnet
- Secure file transfers using SCP and SFTP (based on SSH)

Note that the only two accounts that typically use the LOGIN service on ROX are "root" and "rrsetup".

Authentication requests for LOGIN services will attempt to use RADIUS first and any local authentication settings will be ignored. Only when there is no response (positive or negative) from any of the configured RADIUS servers will ROX authenticate users locally.



NOTE

ROX manages both the RADIUS "login" and "ssh" services together as "LOGIN" from the Webmin interface. Please refer to [Section 4.7, "RADIUS Server Configuration"](#) for details on configuring accounts for these services at the RADIUS server.

The PPP Service

The PPP service represents incoming PPP connections via modem. Authentication requests to the PPP service use RADIUS only. In the event that no response is received from any configured RADIUS server, ROX will not complete the authentication request.

The WEBMIN Service

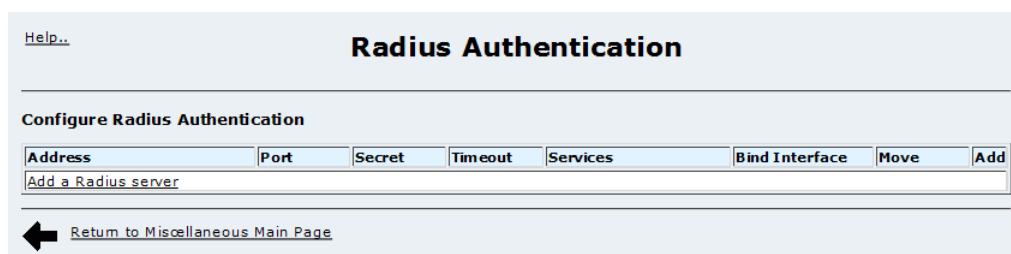
The *WEBMIN* service represents access to the Webmin user interface. Webmin accesses are authenticated first against the local user database. If the user does not exist locally, (the root account, for example, is always defined locally) then Webmin will attempt to authenticate the user via RADIUS.

The *WEBMIN* service allows the configuration of multiple operator accounts, each logged separately, and each with a different privilege level. Please refer to [Section 4.2.1, “Webmin User and Group Fundamentals”](#) for more detail on Webmin privilege levels.

The RADIUS server authenticating the WEBMIN service must also be configured to supply a "privilege-level" field which is used to allow different levels of access to different users of the web management interface. See [Section 4.7, “RADIUS Server Configuration”](#) for more information on configuring the RADIUS server.

Section 4.6.4

RADIUS Authentication Configuration



The screenshot shows a web interface titled "Radius Authentication". At the top left is a "Help.." link. Below the title is a section "Configure Radius Authentication". This section contains a table with columns: Address, Port, Secret, Timeout, Services, Bind Interface, Move, and Add. The "Add" column has a link "Add a Radius server". Below the table is a "Return to Miscellaneous Main Page" link with a left-pointing arrow.

Figure 149: RADIUS Authentication Main Menu

RADIUS Authentication is configured from within the *Maintenance* menu, *Miscellaneous* sub-menu. This menu allows you to add, delete and RADIUS servers. Add a server by clicking on the "add-above" or "add-below" arrows in the *Add* column. You may also edit a server by following its link under the *Address* field.

Reorder the servers by clicking on the arrows in the *Move* column.

Section 4.6.5

Edit RADIUS Server Parameters

[Help..](#)

Edit Radius Server Parameters

Radius Server Parameters	
Hostname/IP	186.42.4.130
Port Number	<input checked="" type="radio"/> Default <input type="radio"/> <input type="text"/>
Shared Secret	...
Timeout	10 (1-20 Seconds)
Service	WEBMIN
Bind Interface	eth1

Note: The **Bind Interface** option applies to RADIUS requests for the WEBMIN service only. The source IP address of these exchanges will be the IP address of the selected network interface.

[← Return to Radius Authentication](#)

Figure 150: RADIUS Authentication Server Parameters

This menu configures, tests and deletes RADIUS server entries.

The *Hostname/IP* field configures the RADIUS server's IP address.

The *Port Number* field sets the port number used by the RADIUS server. The default port for RADIUS is 1812.

The *Shared Secret* field configures a unique password used to authenticate communications with this server. Note that the shared secret must also be configured on the RADIUS server for the router being configured.

The *Timeout* field sets the maximum time in seconds to wait for responses from a RADIUS server before aborting the transaction with the current server. When the timeout expires, ROX will proceed to the next configured server (if one has been configured) to retry the transaction.

The *Service* field configures whether the server authenticates LOGIN, WEBMIN, PPP LOGIN or any combination of these types.

The *Bind Interface* field sets the source IP address for RADIUS requests for the WEBMIN service only. The IP address of the selected network interface will be used as the source address for these exchanges. If "none" is selected, the source IP address will be that of the network interface from which messages are transmitted to the RADIUS server. Note again that this setting applies to the WEBMIN service only.

Section 4.7

RADIUS Server Configuration

This section describes the configuration procedures for two popular RADIUS servers, "FreeRADIUS" and the Microsoft Windows "Internet Authentication Service" in order to create and manage accounts that are able to access resources on ROX. There are four RADIUS attributes required for the configuration of accounts to access services on ROX. The following table shows the RADIUS attributes required by ROX for accounts that are designated to use one or more of the "webmin", "login", "ppp", or "ssh" services:

Table: Required Attributes for various RADIUS services

RADIUS Attribute	webmin	login	ppp	ssh
User ID	required	required	required	required
Password	required	required	required	required

RADIUS Attribute	webmin	login	ppp	ssh
NAS-Identifier				
RuggedCom-Privilege-level	required			

Every account to be authenticated on behalf of the ROX must have a user ID and password. The RADIUS "NAS-Identifier" attribute may optionally be used to restrict which service an account may access:

- webmin
- login
- ppp
- ssh

Accounts that do not specify a "NAS-Identifier" attribute may access any ROX service upon authentication. Accounts may also be defined to have access to one or several services. For more information on these services on ROX, please refer to [Section 4.6.3, "RADIUS, ROX and Services"](#).

A RADIUS attribute specific to Siemens, "RuggedCom-Privilege-level", is used by Webmin to assign specific capabilities to Webmin users on a per-user basis. This attribute must be set for user accounts designated to access Webmin. Please refer to [Section 4.2.1, "Webmin User and Group Fundamentals"](#) for a complete discussion of privilege levels and their use in ROX. The following information is necessary to add support for this attribute to the vendor-specific extensions of the chosen RADIUS server:

- Siemens uses Vendor number 15004.
- "RuggedCom-Privilege-level" is attribute 2, of type "string".
- "RuggedCom-Privilege-level" must take one of the following three values:
 - "admin"
 - "operator"
 - "guest"

User accounts that require access to Webmin must be assigned a "RuggedCom-Privilege-level". Accounts that do not require Webmin access but are to be given shell login or PPP access do not require the privilege level attribute to be set.

The following two sections illustrate how to add this information to a RADIUS server configuration.

Section 4.7.1

Webmin Privilege Levels and FreeRADIUS

This section describes how to add Siemens Vendor-Specific RADIUS attributes to the FreeRADIUS "dictionary" so that they may be used in configuring accounts for ROX.

1. Locate the FreeRADIUS dictionary files (commonly in the /usr/share/freeradius/ directory on Linux systems).
2. In the dictionary directory, open the file named "dictionary", and add the line: "\$INCLUDE dictionary.ruggedcom". Note that there are typically many other vendor attribute dictionary files included in the main FreeRADIUS dictionary file.
3. Create a file named "dictionary.ruggedcom" in the dictionary directory containing the following:

```
# -*- text -*-
#
#   The RuggedCom Vendor-Specific dictionary.
#
# Version:   $Id: dictionary.RuggedCom,v 1.3.4.1 2005/11/30 22:17:24 aland Exp $
```

```
#
#   For a complete list of Private Enterprise Codes, see:
#
#   http://www.isi.edu/in-notes/iana/assignments/enterprise-numbers
#
VENDOR      RuggedCom      15004

BEGIN-VENDOR      RuggedCom

ATTRIBUTE      RuggedCom-Privilege-level      2      string

END-VENDOR RuggedCom
```

4. Create user accounts in the `/etc/freeradius/users` file. For example, in order to create a user "john" with a password "test" with "operator" access to Webmin, add the following lines to `/etc/freeradius/users`:

```
john      Auth-Type := Local
User-Password == "test",
NAS-Identifier = "webmin",
RuggedCom-Privilege-level = "operator"
```

5. Restart your freeradius server.

Section 4.7.2

Webmin Privilege Levels and Windows IAS

This section describes the steps necessary to configure Microsoft Windows IAS (Internet Authentication Service) to authenticate Webmin user accounts for ROX.

1. Create a group for each privilege level. For example, for the "operator" privilege level, create a group named `RADIUS_ROX_operator`. User accounts needing "operator" privileges would then be added to this group.
2. Use the *New Remote Access Policy Wizard* to create a custom policy with the following settings:
 - Policy conditions:
 - NAS-Identifier matches "webmin"
 - Windows-Group matches the group corresponding to the user's privilege level
 - Permission: Grant remote access permission
3. Double click the newly created policy name. In the popup window, click the *Edit Profile...* button.

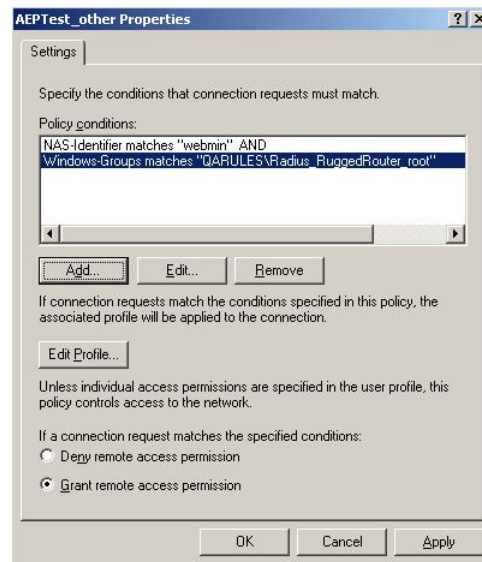


Figure 151: IAS Window - Edit Remote Access Policy

4. In the *Edit Profile* window, under the *Advanced* tab, click the *Add...* button.

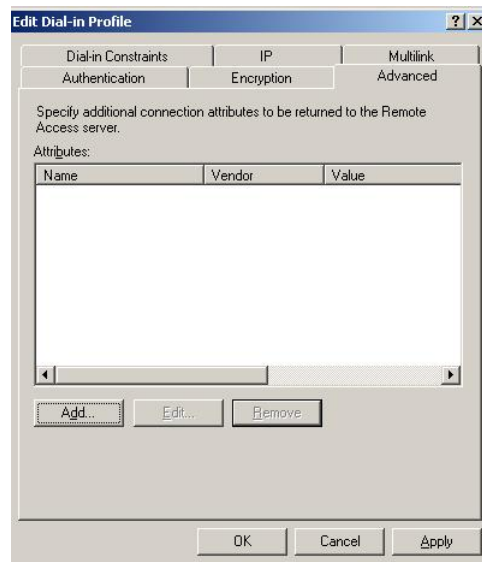


Figure 152: IAS Window - Edit Profile

5. In the *Add Attribute* window, select the *Vendor-Specific* attribute line, and click *Add*.

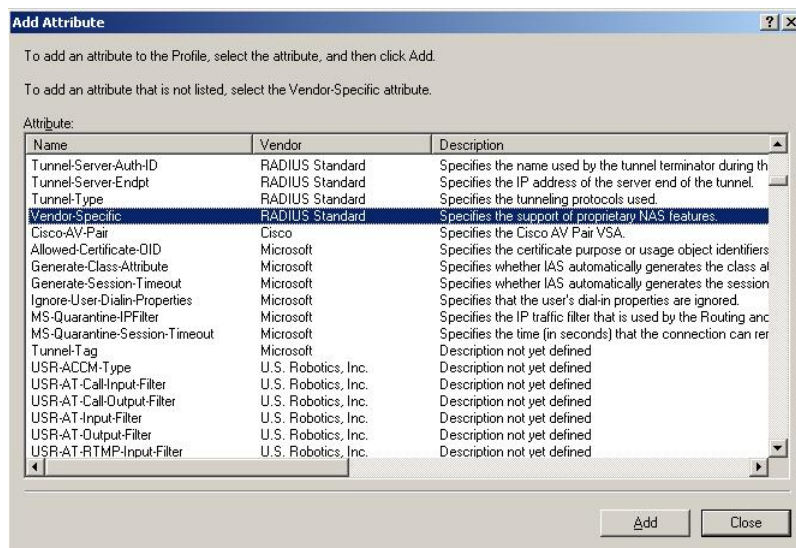


Figure 153: IAS Window - Add Attribute

6. In the *Multivalued Attribute Information* window, click the *Add* button.



Figure 154: IAS Window - Multivalued Attribute Information

7. In the *Vendor-Specific Attribute Information* window, select *Enter Vendor Code*, and enter "15004". Select *Yes. It conforms.* and click the *Configure Attribute...* button.




Figure 155: IAS Window - Vendor-Specific Attribute Information

8. In the *Configure VSA (RFC compliant)* window, in the *Vendor-assigned attribute number* edit box, enter "2". For *Attribute format*, select "string"; for *Attribute value* enter the privilege level of the group being created, e.g. "operator".




Figure 156: IAS Window - Configure VSA (RFC compliant)

Section 4.7.3

PPP/CHAP and Windows IAS

In order for Windows IAS to authenticate PPP connections that use the CHAP authentication protocol, IAS must be made to store passwords using what it calls "reversible encryption".

1. Ensure that CHAP authentication is enabled in the Remote Access Policy.
2. In the Active Directory settings for each PPP user, select "Store password using reversible encryption".

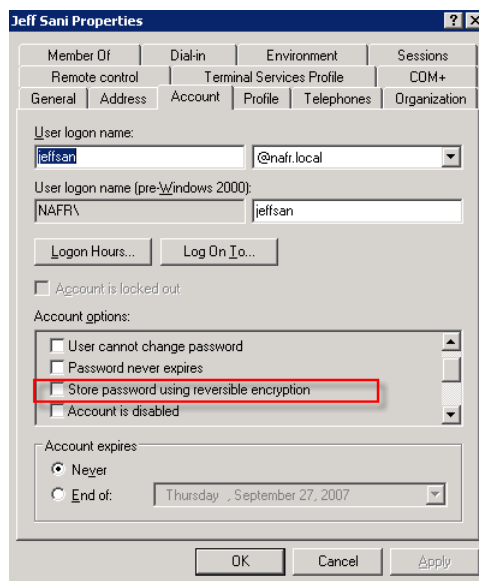


Figure 157: Active Directory - User Account Properties

5 Setup and Configuration

This chapter describes how to setup and configure the device for use on a network using the various features available in ROX. It describes the following tasks:

- [Section 5.1, “Configuring Networking”](#)
- [Section 5.2, “Configuring Ethernet Interfaces”](#)
- [Section 5.3, “Configuring Frame Relay/PPP and T1/E1”](#)
- [Section 5.4, “Configuring Frame Relay/PPP and T3/E3”](#)
- [Section 5.5, “Configuring Frame Relay/PPP and DDS”](#)
- [Section 5.6, “Multilink PPP over T1/E1”](#)
- [Section 5.7, “Configuring PPPoE/Bridged Mode On ADSL”](#)
- [Section 5.8, “Configuring the Firewall”](#)
- [Section 5.9, “Traffic Control”](#)
- [Section 5.10, “Traffic Prioritization”](#)
- [Section 5.11, “Configuring IPSec VPN”](#)
- [Section 5.12, “Configuring Dynamic Routing”](#)
- [Section 5.13, “Link Backup”](#)
- [Section 5.14, “Configuring VRRP”](#)
- [Section 5.15, “Link Layer Discovery Protocol \(LLDP\)”](#)
- [Section 5.16, “Configuring Generic Routing Encapsulation”](#)
- [Section 5.17, “Configuring Layer 2 Tunnels”](#)
- [Section 5.18, “Configuring the DHCP server”](#)
- [Section 5.19, “DHCP Relay”](#)
- [Section 5.20, “Configuring NTP Servers”](#)
- [Section 5.21, “CrossBow Station Access Controller \(SAC\)”](#)

Section 5.1

Configuring Networking

This section familiarizes the user with:

- Configuring routing and gateways
- Configuring DNS (Dynamic Name Service)
- Entering host addresses
- Configuring a pair of End To End Backup interfaces
- Viewing routing tables

Section 5.1.1

IPv6 Fundamentals

Version 6 of the Internet Protocol (IPv6, RFC 2460) has been designated to replace IPv4 throughout the Internet. Some important changes that IPv6 introduces relative to IPv4 fall into the following categories:

- Addressing

IPv6 addresses are four times the length of IPv4 addresses, at 128 bits, to be used as 64 bits of network and 64 bits of host address. The larger address space allows much greater flexibility in hierarchical network definition and routing.

- Header Format

The IPv6 packet header has been simplified relative to IPv4 in order to simplify and therefore speed the processing of packets by routing nodes. It also features more efficiently encoded options and greater flexibility in creating extensions.

- Security

Security has been designed into IPv6, rather than being treated as a component that must be added to existing IPv4 network stacks.

Section 5.1.2

Network Configuration

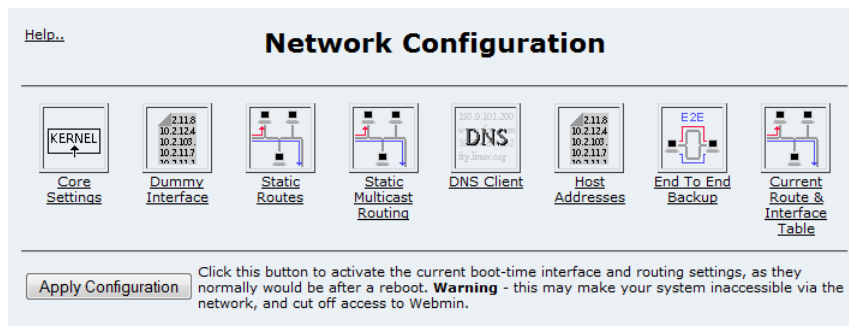


Figure 158: Network Configuration Menu

This menu allows you to configure IP networking parameters.

Select the *Core Settings* icon to configure kernel networking settings such as syncookies filtering.

Select the *Dummy Interface* in order to assign an IP Address to the router that is independent of its interfaces.

Select the *Static Route* icon to assign a gateway address.

Select the *Static Multicast Routing* icon to configure static multicast routes.

Select the *DNS Client* icon to point the router at a DNS server.

Select the *Host Addresses* icon to locally configure IP address-hostname mappings.

Select the *End To End Backup* icon to configure an end to end backup connection.

Select the *Current Routing & Interface Table* icon to view the routing table.

The *Apply Configuration* button serves to restore the permanently saved changes and restart Ethernet networking.

Section 5.1.2.1

Core Settings

Core Settings

Core Settings	
Allow IPv6 Configuration	<input type="radio"/> Yes <input checked="" type="radio"/> No
Ignore All ICMP ECHO requests	<input type="radio"/> Yes <input checked="" type="radio"/> No
Ignore ICMP Broadcasts	<input checked="" type="radio"/> Yes <input type="radio"/> No
Syncookie Protection	<input checked="" type="radio"/> Yes <input type="radio"/> No
Send ICMP Redirect	<input checked="" type="radio"/> Yes <input type="radio"/> No
Set UDP non-block in IPsec	<input type="radio"/> Yes <input checked="" type="radio"/> No

[Save and Apply](#)

[Return to network configuration](#)

Figure 159: Core Networking Settings

This menu allows you to configure core networking settings.

The *Allow IPv6 Configuration* field determines whether IPv6 may be configured via Webmin.

The *Ignore All ICMP ECHO* field corresponds to the kernel `icmp_echo_ignore_all` setting. Setting Ignore All ICMP ECHO to "yes" will cause the kernel to reject incoming ICMP ECHO request packets.

The *Ignore ICMP Broadcasts* field corresponds to the kernel `icmp_echo_ignore_broadcasts` setting. Setting Ignore ICMP Broadcasts to "yes" will cause the kernel to reject incoming ICMP ECHO request packets if their destination address is a broadcast address.

The *Syncookie Protection* field corresponds to the kernel `tcp_syncookie` setting. Setting Syncookie Protection to "yes" will cause the kernel to protect against SYN flood attacks.

The *Send ICMP Redirect* field corresponds to the kernel `send_redirect` setting. Consider a networked host H1 sending an IP datagram to a remote host H2. H1 is on the same network segment as two routers: R1 and R2. If host H2 is reachable via R2, and H1 sends an IP datagram to H2 via R1, then assuming that R1 has a route to H2 via R2, R1 will send an ICMP redirection message to H1 informing it that the route to H2 is via R2. Setting Send ICMP Redirect to "no" will cause the kernel not to send an ICMP redirect message even if one would normally be sent.

The *Set UDP non-block in IPsec* field is used to control the logging to the remote syslog server during IPsec tunnel establishment. By default, it is set to "no" which means the system will not log to the remote syslog server until the IPsec tunnel is established.

Section 5.1.2.2

Dummy Interface



Figure 160: Dummy Interface

This menu allows you to configure a dummy interface. Normally the router is reachable on any of its interface addresses, whether the interface is active or not. When OSPF and link detection is used, inactive interfaces are not advertised to the network and thus not reachable. A dummy interface is always advertised and is thus reachable. Pressing the *Save* button will save the configuration change. Pressing the *Delete* button will remove the dummy interface.

Section 5.1.2.3

Static Routes

This menu allows you to configure static routing entries, including default routes. Each static route specifies how the router can reach a remote subnet. It also allows the conversion of other static routes, obtained via DHCP for example, to permanently configured static routes.

If IPv6 support is enabled in the Core settings menu, IPv6 static routes may also be configured here.

If multiple gateways are available to route to a given remote subnet, a static route entry may be entered for each one, with the same subnet and different gateway specifications. Typically, one would also enter a different metric for each route, the lowest metric indicating the preferred route.

Multipath Routes

It is also possible to specify the same metric for each one of several alternative routes to the same remote subnet. This allows the creation of a multipath route. With such a set of redundant routes available to a remote subnet, the router will select one or another route to transmit traffic destined to the subnet.

The end result is that the aggregate of data traffic to the remote subnet is shared among the multiple routes. Note the distinction between Multipath Routing and [Section 5.6, "Multilink PPP over T1/E1"](#): whereas Multilink PPP effectively multiplies the bandwidth for all traffic by the number of links that comprise a 'bundle', Multipath Routing multiplies the capacity of the route, at link-native speeds, by the number of different routes provided.

Default Routes

A default route is a special instance of a static route. The destination network of `0.0.0.0/0` is the most general possible IPv4 network specification. Packets destined to an IPv4 subnet that is not reachable via any other routing entry in the system will be forwarded to the default gateway, i.e., the gateway for the default route. Default routes for both IPv4 and IPv6 may be configured.

Section 5.1.2.4

Configuring Static Routes

[Help..](#)

Static Routes

Configured Static Routes

Route	Network/Mask	Gateway	Interface	Metric	Comment
1	0.0.0.0/0		w2c1ppp		Inactive
2	10.2.3.0/24	192.168.100.2			Inactive
3	10.200.0.0/16	172.30.128.1			Active
4					

Save

Other Static Routes

Network/Mask	Gateway	Interface	Metric	Action
10.2.3.0/24		gre1		Save to Configured Static Routes
0.0.0.0/0		w1c1fr22		Save to Configured Static Routes

Note: This router has the following network interfaces for IPv4 routes
eth1 eth2 eth3 eth4 w1c1fr22 w2c1ppp

Note: This router is set to not support IPv6.
Note: Assign null0 to Interface field to install blockhole route.




 [Return to network configuration](#)

Figure 161: Static Routes



NOTE
Modem PPP, PPPoE on ADSL, and any interface configured to obtain IP configuration via DHCP may also negotiate default gateways independently of this configuration menu.


The *Network/Mask* field specifies the remote subnet field of a static route definition. If this field is cleared, the route will be deleted when Save is clicked. The Network is specified in dotted quad notation, and the Mask (the number of bits in the subnet mask) is an integer between 0 (for a default route) and 32 (for a host route).



NOTE
It is possible to create a route on a locally connected broadcast network (i.e. without a gateway) without also bringing up a corresponding IP address on that interface. For example, it would be possible to add 192.168.30.0/24 to eth1, which has an IP address of 10.0.1.1 but no corresponding alias address on the 192.168.30.0/24 subnet.

The *Gateway* field specifies the IP address of the 'next hop' to which to forward traffic destined to the specified subnet. If the gateway to a particular subnet is across a point-to-point link, it is not necessary to specify a gateway, but a network interface (below) must be specified.

The *Interface* field specifies the network interface to use to reach the gateway. The interface does not need to be active or even exist, but the route will not be installed until both are true. Specifying an interface is only strictly necessary when a gateway address is not specified. The menu provides a list of currently configured interfaces for quick reference.



NOTE
A "blackhole", or "null" route may be installed by entering "null0" in the interface field.

The *Metric* field specifies an integer cost metric for the route, which is used when choosing among multiple routes in the routing table that most closely match the destination address of a packet being forwarded. The route with the lowest metric is chosen as the active route.



NOTE

Multiple routes to the same destination subnet may also be specified using identical metrics in order to create multipath routes. Please refer to [the section called "Multipath Routes"](#).

The *Comment* field shows the status of the static route, and provides a basic cause when the route is not installable.

The Save button below the table will perform the following sanity checks on routing entries that have been added or modified. If the tests pass, the routes will be saved and immediately installed.

- A specified Gateway must be reachable, and if a network interface is also specified, must be reachable via that interface.
- If a specified network interface exists but is not active, the static route will be installed and marked, "Inactive (interface is not active)".
- If a specified network interface does not exist (e.g. an on-demand modem PPP connection) the static route will be installed and marked, "Inactive (interface does not exist)".

Delete routes by removing their Network/Mask addresses before saving.



NOTE

In order to redistribute static routes to other routers, the Redistribute Static option must be enabled in the corresponding OSPF, RIP, or BGP configuration's Global parameters menu in Webmin.

Section 5.1.2.5

Other Static Routes

This table will be shown if there are active static routes which were not configured manually in the *Configured Static Routes* table. The *Save to Configured Static Routes* link next to each route entry in this table will make the corresponding route permanent.



NOTE

There are situations where manually entered routes should not be converted, e.g. routes dynamically added by IPSec and GRE tunnels. Making these routes permanent may cause the daemons that add them to fail.

Section 5.1.2.6

Static Multicast Routing

[Help..](#)

Static Multicast Routing

Configured Static Multicast Routes

Route	Multicast IP Address	Input Interface	Source IP Address	Output Interfaces	Comment
1	224.10.10.10	eth1	192.168.1.254	eth4	Installed
2					Use commas to separate Output Interfaces

Note: This router has the following network interfaces
eth1 eth2 eth3 eth4 wic1 wippp


 [Return to network configuration](#)

Figure 162: Static Multicast Routing

This menu allows you to configure static multicast routing.

The *Configured Static Multicast Routes* table shows configured multicast routes.

New routings may be added by completing the bottom row of the table and selecting the **Save** button. Routings may be deleted by clearing the routings *Multicast IP Address* field and selecting the **Save** button.

The *Multicast IP Address* field specifies the multicast IP address to be forwarded.

The *Input Interface* field specifies the interface upon which the multicast packet arrives.

The *Source IP Address* specifies the multicast packet's expected source IP address.

The *Output Interface* specifies the interface to which the matched multicast packet will be forwarded.

The *Comment* field shows the current status of the routing.

The *Note* field below the table shows current active interfaces.

In order to start Multicast routing at each and every boot, you must enable it via the System folder, Bootup and Shutdown menu.

Section 5.1.2.7

DNS Client

Figure 163: DNS Client

This menu allows you to display and configure various DNS client fields.

The *DNS servers* fields allow you to specify, in order, the servers from which to request Internet Domain Name resolution.

The *Search domains* field allow you to specify the domain names of, primarily, the domain of which the router is a member, and secondarily, other domains that may be used to search for an unqualified host name (i.e. as though it were local). If a domain name is not specified here, the router will attempt to extract this information from the host addresses.

Section 5.1.2.8

Host Addresses

Figure 164: Host Addresses

This menu allows you to display and configure host addresses. Host addresses are useful when a non-changing IP address is often used or when DNS is not configured.

Click the *Add a new host address* link to add an address.

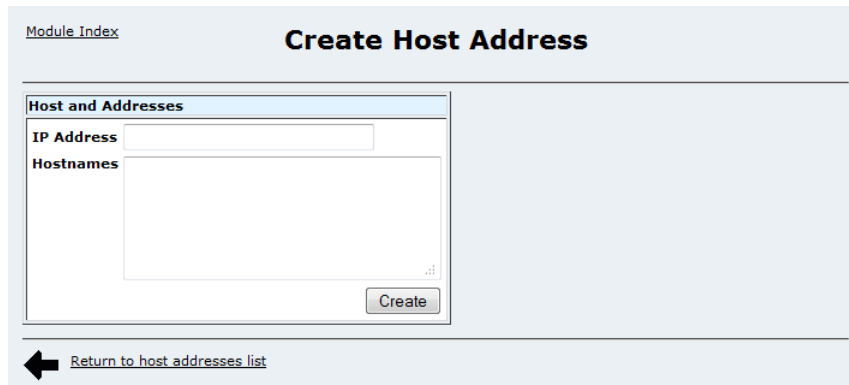
The screenshot shows a web interface titled "Create Host Address". At the top left is a link for "Module Index". The main form area is titled "Host and Addresses" and contains two input fields: "IP Address" and "Hostnames". The "Hostnames" field is a larger text area. To the right of these fields is a large, empty light blue area. At the bottom right of the form is a "Create" button. Below the form is a navigation bar with a left-pointing arrow and a link that says "Return to host addresses list".

Figure 165: Create Host Address

The *IP Address* field sets the IP address.

The *Hostnames* field sets the hostname.

The *Create* button saves the host address.

Section 5.1.2.9

End to End Backup

End to end backup is method of using two interfaces to ensure a reliable end to end connection between two routers using alternate routing, without the need to configure routing protocols.

The two interfaces are assigned as a primary:secondary backup pair. The primary interface serves as the gateway. If connectivity to the target is lost from the primary interface, traffic is migrated to the secondary interface. When connectivity is restored on the primary path, traffic will be restored to it.




Figure 166: End to End Backup Example

1. Network A 2. Network B 3. Router 1 4. Router 2

The backup is "end to end" because connectivity is determined by the availability of an interface on the target system, and not a local link. In the above figure, interface w1ppp acts as the primary interface and eth1 acts as the secondary interface. The router tests the primary path by probing 192.168.16.2 on router 2. A failure of the either w1ppp, network A or the remote link on router2 will render the primary path as "failed".

If the primary path fails, the routing table will be modified to direct packets out the secondary (eth1 in the above figure).

Presumably, the secondary is a higher cost (and perhaps lower throughput) path. In the initial deployment of this feature, the secondary path was implemented with Ethernet-CDMA modem. The modem featured a low latency connection time (initiated by the reception of packets) but had a low bandwidth capability and high monetary cost.

Note that the feature must be implemented at both routers. If the feature is only implemented at router 1, the second router's gateway will still point towards Network A after a failure of the primary path. Packets from router 1 would reach router 2 through the secondary, but the responses would disappear in the black hole of the failed path.

To configure End to End backup, see [Section 5.1.2.10, "Configuring End To End Backup"](#)

Section 5.1.2.10

Configuring End To End Backup

[Help..](#)

End To End Backup

This menu configures the end to end backup feature. This method assigns two interfaces as a primary:secondary backup pair and monitors the primary link to detect a failure. After a failure occurs, traffic is shunted to the secondary until the primary is restored. Note that in order for end to end backup to work the primary interface must act as the default interface.

End To End Routes

Primary Interface

eth1

Secondary Interface

eth2

Fail Over Time (Seconds)

1

Generate Alarms

☒ Yes ☐ No

Peer IP Address on Primary

192.168.16.2

Peer IP Address on Secondary

192.168.17.2

(0 < Fail Over Time <=60)

Save

Save and Apply


 [Return to network configuration](#)

Figure 167: End To End Backup

This menu allows you to display and configure end to end backup.

In order to start end to end backup at each and every boot, you must enable it via the System folder, Bootup and Shutdown menu. The menu will remind you if the feature is not enabled.

The *Primary Interface* field determines the primary interface. The interface selected should be configured to supply the default gateway.

The *Peer IP Address on Primary field* sets the IP address to probe for connectivity on the primary interface.

The *Secondary Interface* field determines the secondary interface.

The *Peer IP Address on Secondary* field sets the IP address to probe for connectivity on the secondary interface.

The *Fail Over Timer* field determines the amount of time the primary link must be failed before directing packets down the secondary link.

The *Generate Alarms* field determines whether alarms are generated upon configuration problems and link failures.

The *Save* button will save changes to the configuration file. The *Save and Apply* button will save changes restart the end to end backup daemon.

Section 5.1.2.11

Current Routing and Interface Table

This menu displays the current routing table and the state of the router's interfaces. Consult [Section 2.2, “Network Utilities”](#) for details of this menu.

Section 5.2

Configuring Ethernet Interfaces

This section familiarizes the user with:

- Reading the Ethernet LEDs in [Section 5.2.1, “LED Designations”](#)
- VLAN Fundamentals in [Section 5.2.2, “VLAN Interface Fundamentals”](#)
- Ethernet Bridge Fundamentals in [Section 5.2.5, “Bridge Fundamentals”](#)
- PPPoE Fundamentals in [Section 5.2.3, “PPPoE On Native Ethernet Interfaces Fundamentals”](#)
- [Section 5.2.6, “Ethernet Configuration”](#) contains instructions for the following:
 - [Section 5.2.6.2, “Editing Currently Active Interfaces”](#)
 - [Section 5.2.6.3, “Creating Active Virtual LAN Interfaces”](#)
 - [Section 5.2.6.4, “Edit Boot Time Interfaces”](#)

Siemens manufactures dual Ethernet Interface boards in a variety of formats. Some (most notably the optical interfaces) have the same outward appearance but different order numbers. A complete set of descriptions is displayed on the console during boot and can be found after boot in the file `/var/cache/ruggedrouter/inventory`.

Section 5.2.1

LED Designations

The device includes two sources of LED indicated information about Ethernet ports, the front panel LEDs and the LED Panel.

A LED is associated with each port, next to the Ethernet interface RJ45 socket. This LED is off when the link is disconnected, remains solidly on when the link is established and flashes briefly from on to off when traffic occurs.

The LED Panel also summarizes this information. LEDs 1-4 reflect traffic on Ethernet port 1-4. LEDs 5-8 reflect the link status of the same ports.

Section 5.2.2

VLAN Interface Fundamentals

A virtual LAN (VLAN) is a group of devices on one or more LAN segments that communicate as if they were attached to the same physical LAN segment. VLANs are extremely flexible because they are based on logical instead of physical connections. When VLANs are introduced, all traffic in the network must belong to one or another VLAN. Traffic on one VLAN cannot pass to another, except through an intranetwork router or layer 3 switch.

The IEEE 802.1Q protocol specifies how traffic on a single physical network can be partitioned into VLANs by "tagging" each frame or packet with extra bytes to denote which virtual network the packet belongs to.

Section 5.2.2.1

VLAN Tag

A VLAN tag is the identification information that is present in frames in order to support VLAN operation. If an Ethernet frame is VLAN tagged, the EtherType value (immediately following the Source MAC address) is set to 0x8100, denoting 802.1Q (VLAN). The next 2-bytes of the VLAN tag contain: a 3-bit User Priority Field that may be used as a priority level for Ethernet frames, a 1-bit Canonical Format Indicator (CFI) used to indicate the presence of a Routing Information Field (RIF), and finally the 12-bit VLAN Identifier (VID) which uniquely

identifies the VLAN to which the Ethernet frame belongs. These four bytes, known as the VLAN tag, are followed by the rest of the Ethernet frame, starting with the length field.

Section 5.2.2.2

ROX Functions Supporting VLANs

Functions	Support	Comments
Static Route and Default Route	Yes	
Static Multicast Routing	Yes	
End To End backup	Yes	
PPPoE	No	
Shorewall Firewall	Yes	
IPSec	Yes	
VRRP	Yes	
Traffic Prioritization	Yes	
Dynamic Routing		Both OSPF and RIP support VLAN
GRE Tunnel	Yes	
DHCP Server	Yes	

Section 5.2.3

PPPoE On Native Ethernet Interfaces Fundamentals

ROX supports PPPoE (Point-to-Point Protocol Over Ethernet) over both external modems (described here) and internal interfaces (described in [Section 5.7, "Configuring PPPoE/Bridged Mode On ADSL"](#)). The section contains more useful information on PPPOE Authentication, Addresses, DNS Servers and MTU Issues.

Only one PPPoE interface can be created on each Ethernet Interface. Each PPPoE interface name is assigned internally. The name is "pppX", where X is 10 plus the native Ethernet interface the PPPoE is created upon (e.g. a PPPoE on eth1 is ppp11).

Section 5.2.4

IPv6 on Ethernet Fundamentals

By default, IPv6 disabled on the router, in which case IPv6 addresses may not be assigned to Ethernet interfaces. IPv6 may be enabled via the *IPv6 Support* option in *Core* settings under the *Network Configuration* category.

If IPv6 is enabled on the router, and link is asserted on a given ethernet port, the system will automatically assign a "link-local" address on that port beginning with 0xfe80, for example: fe80::20a:dcff:fe1a:e401/64.

Section 5.2.5

Bridge Fundamentals

ROX supports software-based Ethernet Bridging. The bridge appears to the router as an Ethernet interface, and may be assigned an IP address statically or via DHCP. Network services, such as SSH, DHCP, NTP, VRRP, etc., may be configured to run on the bridge interface.



NOTE

- *When adding interfaces to the bridge, any network services running on the individual interfaces must be reconfigured to refer to the bridge interface. For example, if a DHCP server is running on eth1 and eth1 is made a member of the bridge br1, the DHCP configuration must be changed to refer to br1.*
- *The device's Ethernet bridge is implemented in software and is not as efficient as a hardware-based switch. We do not recommend connecting a switch to the Ethernet bridge interface. When connecting a switch to the ROX, configure the ROX interface as a normal Ethernet interface, rather than as an Ethernet bridge.*
- *Because the device's Ethernet bridge is implemented in software, CPU resources are required to forward broadcast, multicast, and unicast traffic on the bridge.*
- *If the router is running as a firewall, the routeback option must be enabled for the bridge interface under Firewall > Edit Network Interface.*
- *The device's Ethernet bridge turns off spanning tree by default. Do not connect the bridge interface to a ring network.*

Ethernet frames traveling over the bridge can also be filtered based their Ethernet type. When enabled, only the frames that match the allowed Ethernet type(s) are allowed to pass through. The Ethernet type is a two-octet field in the Ethernet frame used to indicate the protocol encapsulation in the payload. Bridge filtering can be used in applications that require Layer 2 packet filtering at a high data rate between Ethernet interfaces.

Section 5.2.6

Ethernet Configuration

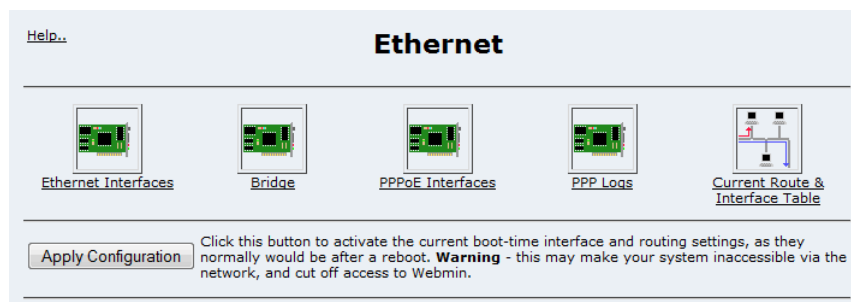


Figure 168: Ethernet Menu

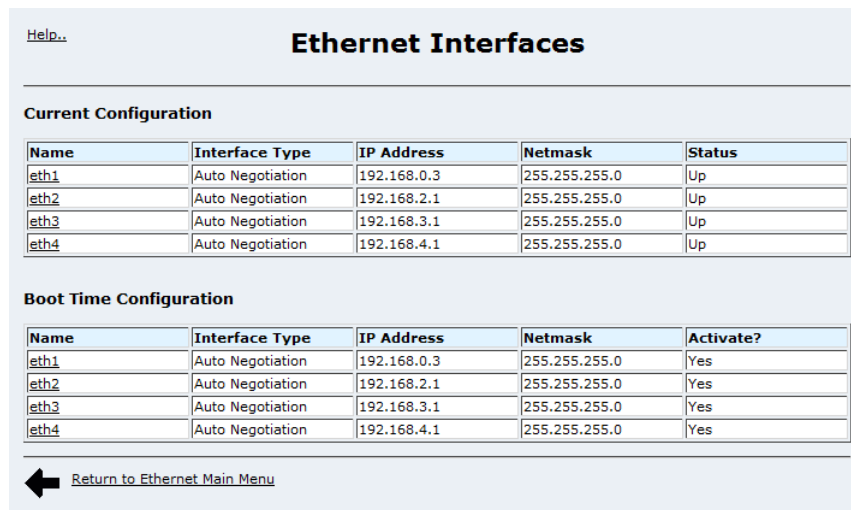
This menu allows you to configure Ethernet interface, Bridge, PPPoE and display the routes and status of all network interfaces.

Select the *Ethernet Interfaces* icon to configure Ethernet interfaces.

The Network Interfaces menu lets you edit the permanent configuration of Ethernet interfaces, or simply try out changes. The *Apply Configuration* button serves to restore the permanently saved changes and restart Ethernet networking.

Section 5.2.6.1

Ethernet Interfaces



[Help...](#)

Ethernet Interfaces

Current Configuration

Name	Interface Type	IP Address	Netmask	Status
eth1	Auto Negotiation	192.168.0.3	255.255.255.0	Up
eth2	Auto Negotiation	192.168.2.1	255.255.255.0	Up
eth3	Auto Negotiation	192.168.3.1	255.255.255.0	Up
eth4	Auto Negotiation	192.168.4.1	255.255.255.0	Up

Boot Time Configuration

Name	Interface Type	IP Address	Netmask	Activate?
eth1	Auto Negotiation	192.168.0.3	255.255.255.0	Yes
eth2	Auto Negotiation	192.168.2.1	255.255.255.0	Yes
eth3	Auto Negotiation	192.168.3.1	255.255.255.0	Yes
eth4	Auto Negotiation	192.168.4.1	255.255.255.0	Yes


 [Return to Ethernet Main Menu](#)

Figure 169: Current and Boot Time Ethernet Configuration

This menu allows you to display and configure the Ethernet interfaces in the router.

The *Current Configuration* table allows you to try out changes on the existing interfaces before making permanent changes. Any changes made take effect immediately, but will not be present after the next boot. The entries in this table can also be used to temporarily disable or re-enable an interface.

The *Boot Time Configuration* table router allows you make changes to the "permanent" configuration of any interface.

The Network Configuration menu *Apply Configuration* button applies permanent changes and restart Ethernet networking. If only temporary changes have been made, the permanent configuration will be re-applied.

In either table, edit the desired interface by clicking on its link under the *Name* column.

Section 5.2.6.2

Editing Currently Active Interfaces

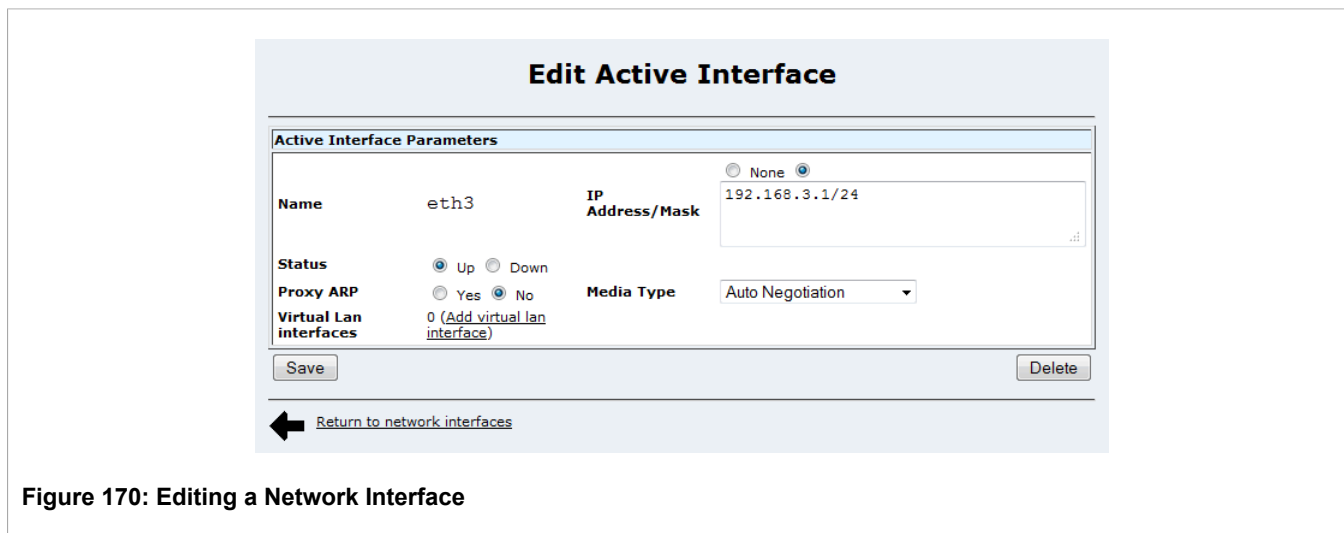


Figure 170: Editing a Network Interface

This menu allows you to make changes to the currently active interfaces. The Save button *will activate any changes, and will not affect the permanent configuration.*

The *IP Address/Mask* field sets the IP address and mask for this interface. You can assign multiple IPv4 or IPv6 addresses to the interface, one on each line. Please note that IPv6 address fe80::20a:dcff:fe0a:1540/64 in this example is the automatically assigned link-local IPv6 address.

The *Status* field provides a way to disable the interface or bring it back into service.

The Proxy ARP fields display whether the interface has proxy-arp activated.

The *Media Type* field displays the current media type. Copper interfaces may be configured to Auto-negotiable, 10 BaseT Half Duplex, 10 BaseT Full Duplex, 100 BaseT Half Duplex and 100 BaseT Full Duplex modes.

The *Virtual LAN interfaces* field displays how many VLAN interfaces are created on this interface and the link, *Add virtual lan interface* allows you to add a VLAN interface on the physical interface.

Section 5.2.6.3

Creating Active Virtual LAN Interfaces

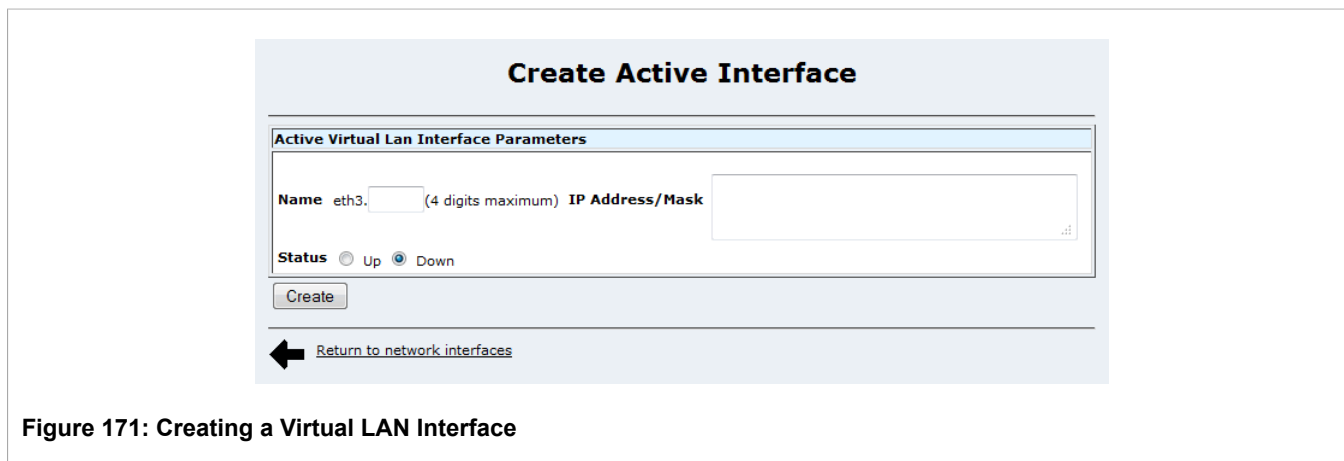


Figure 171: Creating a Virtual LAN Interface

Click the link *Add Virtual Lan Interface* in order to create a VLAN interface.

The only new parameter is the VLAN ID, which must be a numeric value between 1 and 4094. The VLAN ID will be presented automatically as 4 digits (prefixed with 0) if the input is smaller than 4 digits. For example, if the input is 2, it will be automatically changed to 0002.

Section 5.2.6.4

Edit Boot Time Interfaces

Edit Bootup Interface

Boot Time Interface Parameters

Name	eth1	IP Address/Mask	<input type="radio"/> None <input type="radio"/> From DHCP <input checked="" type="radio"/> From BOOTP 192.168.0.3/24
MTU	Automatic	Activate?	<input checked="" type="radio"/> Yes <input type="radio"/> No
Proxy ARP	<input type="radio"/> Yes <input checked="" type="radio"/> No	Media Type	Auto Negotiation
Virtual Lan interfaces	0 (Add virtual lan interface)		

Save and Apply

Return to network interfaces

Figure 172: Editing a Boot Time Interface

This menu allows you to make permanent changes to interfaces and to immediately apply those changes if desired. The *Save* button will save changes to the permanent configuration.

The *Proxy ARP*, *Media Type* and *Virtual Lan Interfaces* controls are as described above.

The *IP Address/Mask* fields allow you to manually specify one or multiple IP address/Mask for this interface, or to obtain the address from DHCP or from BOOTP. You can have both IPv4 and IPv6 (if IPv6 is enabled) addresses at the same time, one on each line.

The *Activate* fields allow you permanently disable the interface without actually deleting it.

The *Virtual LAN interfaces* field displays how many VLAN interfaces are created on this interface and the link, *Add virtual lan interface* allows you to add a VLAN interface on the boot time interface.

The *Save and Apply* button applies any changes after they have been saved.

Section 5.2.6.5

Creating Bootup Virtual LAN Interfaces

Create Bootup Interface

Boot Time Vlan Interface Parameters

Name eth1. (4 digits maximum) IP Address/Mask

MTU Automatic Activate? ☐ Yes ☒ No

VLAN QoS	Map ingress to mark	Map egress from marks
0	<input type="text"/>	default
1	<input type="text"/>	<input type="text"/>
2	<input type="text"/>	<input type="text"/>
3	<input type="text"/>	<input type="text"/>
4	<input type="text"/>	<input type="text"/>
5	<input type="text"/>	<input type="text"/>
6	<input type="text"/>	<input type="text"/>
7	<input type="text"/>	<input type="text"/>

[Return to network interfaces](#)

Figure 173: Creating a Virtual LAN Interface

Click the link *Add virtual Lan Interface* in order to create a VLAN interface.

The *Name* field is the VLAN ID, which must be a numeric value between 1 and 4094. The VLAN ID will be presented automatically as 4 digits (prefixed with 0) if the input is smaller than 4 digits. For example, if the input is 2, it will be automatically changed to 0002.

The *IP Address/Mask* field sets the IP address and mask for this interface. You can assign multiple IPv4 or IPv6 addresses to the interface, one on each line.

The *Activate* fields allow you permanently disable the interface without actually deleting it.

The *VLAN QoS* fields represent the QoS priority options. Quality of Service (QoS) mapping is used to map QoS traffic. It assigns a traffic control mark to incoming IP traffic based on the priority value of a tagged frame. The incoming traffic is then classified and placed in the priority queues according to the traffic control rules specified for the marked rule. In addition, traffic control can assign the same priority or a different priority value when a frame needs to be egressed with a VLAN tag through a traffic control interface.

The *Map ingress to mark* and *Map egress to mark* fields allow you to set ingress and egress marks. Egress markers for QoS maps are used to assign priority to traffic that shares the same mark as one of the egress marks configured for the device.

The *Create and Apply* button creates and applies the new virtual LAN interface.

Section 5.2.6.6

Bridge Configuration

The screenshot shows the 'Bridge Configuration' window. At the top is a 'Help..' link. Below it is the title 'Bridge Configuration'. The main area is titled 'Bridge Configuration Parameters'. It contains several fields: 'Enable Bridge' with radio buttons for 'Enable' (selected) and 'Disable'; 'IP Address/Mask' with radio buttons for 'None', 'From DHCP', and a selected option (likely 'Static') showing the address '192.168.40.1/24'; 'Select Bridge Devices' with a list box containing 'eth1 (Ethernet)', 'eth2 (Ethernet)', 'eth3 (Ethernet)', and 'eth4 (Ethernet)'; and 'Retain IP on Bridge Device' with a checked checkbox. At the bottom left is a 'Save' button, and at the bottom right is a back arrow and the text 'Return to Ethernet Main Menu'.

Figure 174: Creating an Ethernet Bridge

This menu allows you to configure the Ethernet bridge interface.

The *Enable Bridge* field controls whether the bridge interface is enabled. The *Bridge Filtering* icon also appears on the Ethernet menu. If the bridge interface is disabled, the other fields will be ignored.

The *IP Address/Mask* field assigns the IP address and mask on this bridge interface. The bridge interface may similarly use one or more static IPv4 or IPv6 addresses, or obtain an address via DHCP.

The *Select Bridge Devices* list is used to select which Ethernet interfaces are to be part of the bridge interface.

The *Retain IP on Bridge Device* choices enable or disable the ability to retain an Ethernet interface's IP address when it is added to the bridge. When enabled (checked), the IP address is retained and the router can be remotely accessed via the Ethernet interface. When disabled (clear), the IP address must be assigned to the bridge to remotely access the router. This should be enabled when bridge routing is enabled. For more information, see [Section 5.2.6.7, "Bridge Filtering"](#).

The Save button will save the configuration changes. Please note that the changes will be effective immediately after clicking the save button.

**NOTE**

For important information about using Ethernet bridging, see [Section 5.2.5, "Bridge Fundamentals"](#).

Section 5.2.6.7

Bridge Filtering

Help..

Bridge Filtering Configuration

Bridge Filtering Configuration

Bridge filtering ☒ Enable ☐ Disable

Bridge routing ☒ Enable ☐ Disable

Bridge Filtering Rules

Allow Ethernet Type	Action
ISO	<input type="button" value="Delete"/>
0x0800	<input type="button" value="Delete"/>
0x0806	<input type="button" value="Delete"/>

[Add new rule](#)

[Return to Ethernet Main Menu](#)

Figure 175: Configuring Bridge Filtering

This menu allows you to configure Ethernet bridge filtering. It is only available when the Ethernet bridge interface is configured.

IMPORTANT!

When bridge filtering is enabled, an IP address must be assigned to the participating physical interface to remotely access the router.

The *Bridge filtering* options enable or disable bridge filtering. When enabled, only Ethernet frames that match the allowed Ethernet Types are allowed to pass over the bridge. When disabled, Ethernet frames travel through the bridge interface as if it were a normal switch. For this option to function, the *Retain IP on Bridge Device* check box must be checked.

The *Bridge routing* options enable or disable bridge routing. When enabled, IPv4, IPv6 and ARP frames are forced up to the routing layer for processing. See [Section 5.2.6.6, “Bridge Configuration”](#).

The *Save* button will save the configuration changes. Please note that the changes will be effective immediately after clicking the save button.

The *Bridge Filtering Rules* table lists the allowed Ethernet types.

To delete an Ethernet type from the list, click *Delete*.

To add a new type, click *Add new rule*. On the *Create New Rule* page, specify the default ISO type or type a different type in the *Others* box. Click *Save* when done.

[Help..](#)

Create New Rule

Bridge Filtering Configuration

Allow Ethernet TypeNoneOthers

Save

[Return to previous page](#)

Figure 176: Creating New Rule



NOTE
For important information about using Ethernet bridging, see [Section 5.2.5, “Bridge Fundamentals”](#).

Section 5.2.6.8

PPPoE on Native Ethernet Interfaces

This menu allows you to display and configure the PPPoE interfaces on all available Ethernet ports.

[Help..](#)

PPPoE Interfaces

PPPoE Interfaces

Ethernet	Interface Name	MTU	Use Peer DNS	Default Route	Status
eth1	Add PPPoE interface..				
eth2	Add PPPoE interface..				
eth3	Add PPPoE interface..				
eth4	Add PPPoE interface..				

[Return to Ethernet Main Menu](#)

Figure 177: List PPPoE Interfaces

The PPPoE Interfaces table allows you to add a PPPoE interface on an Ethernet ports or change PPPoE interface parameters of created interfaces. Only one PPPoE interface can be created on each Ethernet port.

The *Ethernet* field shows all available Ethernet ports.

The *Interface Name* field shows created PPPoE interfaces and provides a link to edit the existing configuration or create a new one.

The *MTU*, *Use Peer DNS* and *Default Route* fields are the configured information for PPPoE interfaces.

The *Status* field shows the current PPPoE link status.

Section 5.2.6.9

Edit PPPoE Interface

This menu allows you to edit a PPPoE interface.

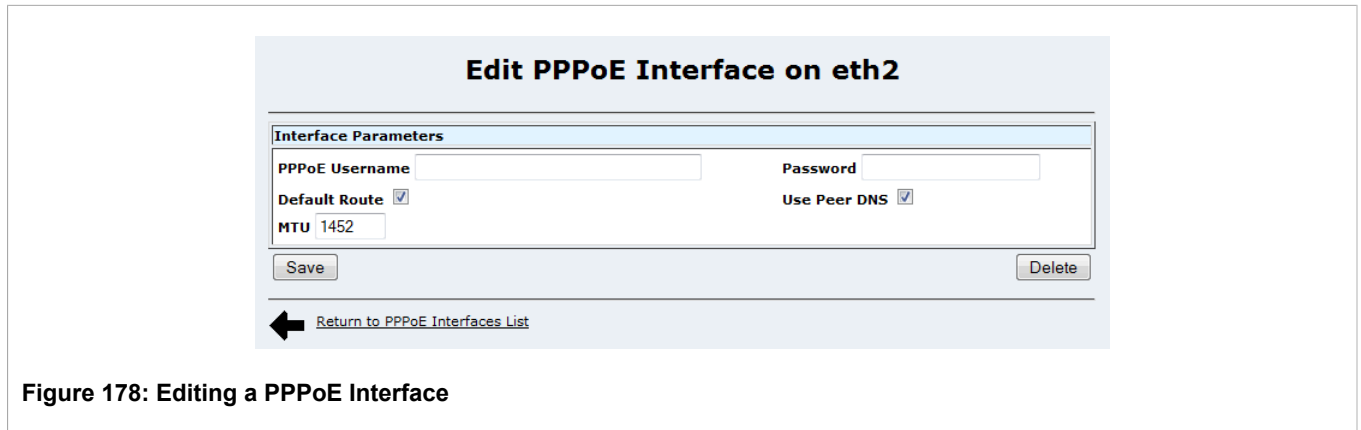


Figure 178: Editing a PPPoE Interface

The *PPPoE Username* field determines the username to use when connecting to the PPPoE server as specified by your provider.

The *Password* field determines the password provided to the PPPoE server.

The *Default Route* check box enables automatically setting a default route using this interface whenever it connects. If this is your primary connection you probably want this option enabled.

The *Use peer DNS* check box enables automatically setting the DNS server entries that the PPPoE server recommends. Enable this option unless you provide your own name servers.

The *MTU* field defines the MTU size to request when connecting to the PPPoE server. In some cases the PPPoE provider may provide a smaller MTU in which case the smaller setting will be used, or it may refuse to alter the MTU and use whatever it considers to be the default.

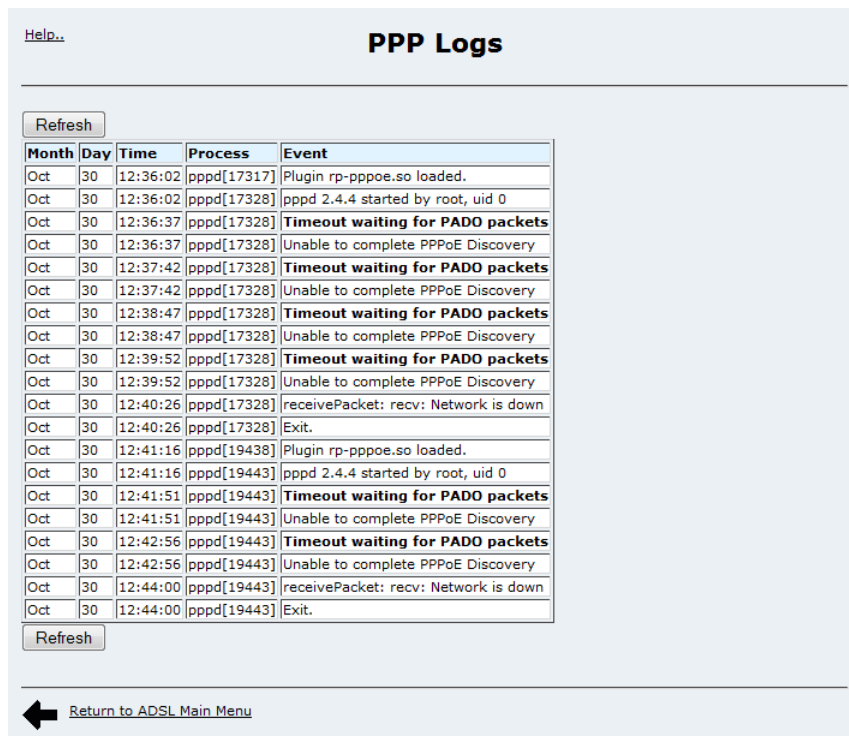
The *Save* button will update all of the changes. The current PPPoE link will be connected.

The *Delete* button will delete the PPPoE interface, closing the current PPPoE link.

Section 5.2.6.10

PPP Logs

This menu displays the native Ethernet and internal ADSL interface PPPoE connection messages. This is mainly useful when trying to debug a PPP connection problem.



Month	Day	Time	Process	Event
Oct	30	12:36:02	pppd[17317]	Plugin rp-pppoe.so loaded.
Oct	30	12:36:02	pppd[17328]	pppd 2.4.4 started by root, uid 0
Oct	30	12:36:37	pppd[17328]	Timeout waiting for PADO packets
Oct	30	12:36:37	pppd[17328]	Unable to complete PPPoE Discovery
Oct	30	12:37:42	pppd[17328]	Timeout waiting for PADO packets
Oct	30	12:37:42	pppd[17328]	Unable to complete PPPoE Discovery
Oct	30	12:38:47	pppd[17328]	Timeout waiting for PADO packets
Oct	30	12:38:47	pppd[17328]	Unable to complete PPPoE Discovery
Oct	30	12:39:52	pppd[17328]	Timeout waiting for PADO packets
Oct	30	12:39:52	pppd[17328]	Unable to complete PPPoE Discovery
Oct	30	12:40:26	pppd[17328]	receivePacket: rcv: Network is down
Oct	30	12:40:26	pppd[17328]	Exit.
Oct	30	12:41:16	pppd[19438]	Plugin rp-pppoe.so loaded.
Oct	30	12:41:16	pppd[19443]	pppd 2.4.4 started by root, uid 0
Oct	30	12:41:51	pppd[19443]	Timeout waiting for PADO packets
Oct	30	12:41:51	pppd[19443]	Unable to complete PPPoE Discovery
Oct	30	12:42:56	pppd[19443]	Timeout waiting for PADO packets
Oct	30	12:42:56	pppd[19443]	Unable to complete PPPoE Discovery
Oct	30	12:44:00	pppd[19443]	receivePacket: rcv: Network is down
Oct	30	12:44:00	pppd[19443]	Exit.

Figure 179: Display PPP Logs

Section 5.2.6.11

Current Routes and Interface Table

The table provided by this command is as described in the *Networking* menu, *Network Utilities* sub-menu. It is also provided here as a convenience.

Section 5.3

Configuring Frame Relay/PPP and T1/E1

This section familiarizes the user with:

- Frame Relay and PPP Terminology and Issues
- Configuring Frame Relay and PPP Links
- Viewing status and statistics
- Upgrading Firmware

Section 5.3.1

T1/E1 Fundamentals

A T1 is a communications circuit upon which has been imposed a digital signal 1 (DS1) signalling scheme. The scheme allows 24 "timeslots" of 64 Kbps DS0 information (as well as 8 Kbps of signalling information) to be multiplexed to a 1544 Kbps circuit.

The 24 DS0s can be used individually as standalone channels, bonded into groups of channels or can be bonded to form a single 1536 Kbps channel, referred to as a clear channel. Not all channels need be used. It is quite common to purchase N channels of 64Kbps bandwidth and leave the remainder unused, this is known as fractional T1.

The telephone network terminates the T1 line and maps each of the channels through the T1 network to a chosen T1 line. Individual and bonded DS0s from more than one remote T1 can be aggregated into a full T1 line (often referred to as central site concentration).

Whereas the T1 line itself is referred to as the physical interface, groups of DS0s form channels and the protocols that run on the channels are known as a logical interfaces. The device provides you the ability to operate Frame Relay or PPP over your logical interfaces.

An E1 is a communications circuit conforming to European standards, possessing 32 64 Kbps channels, of which one is usually reserved for signalling information.

Section 5.3.1.1

Frame Relay

Frame Relay is a packet switching protocol for use over the WAN. ROX provides the ability to construct point-to-point IP network connections over Frame Relay.

Each Frame Relay interface provides a link between a local and peer station. One of the stations must be configured as a Data Communications Equipment (DCE) device (often known as the Switch) while the peer station must be configured as a Data Terminal Equipment (DTE) device (often known as Customer Premises Equipment (CPE)). The DCE is responsible for managing the link, advertising connections to the DTE and switching packets between connections. The DTE raises individual connections and sends data on them.

When using a T1/E1 line to access a public Frame Relay provider, configure the Router as a DTE.

Unlike PPP, a Frame Relay link can provide multiple connections. Each connection is identified by a Data Link Connection Identifier (DLCI) and must match at the DCE and DTE. The use of multiple connections can support meshed network interconnections and disaster recovery.

Section 5.3.1.2

Location of Interfaces and Labelling

Unlike the Ethernet ports (which are statically located), the location of T1/E1, DDS and ADSL ports in your router depends upon the number of ports and how they were ordered. Refer to the labelled hardware image as presented in the Webmin home page.

To make labelling easy to understand, all T1/E1, T3, DDS and ADSL ports are assigned a unique port number that relates to the LEDs on the status panel.

Section 5.3.1.3

LED Designations

ROX includes two sources of LED indicated information about T1/E1 lines, the T1/E1 card itself and the LED Panel.

One LED is associated with each line, next to the interface jack. This LED is red when the link is disconnected, flashes green when the link is connecting and remains solid green when the link is established.

ROX also indicates information about T1/E1 ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port. Consult the section [Section 3.3, “Using The LED Status Panel”](#) to determine which LEDs correspond to the port.

Section 5.3.1.4

Included with T1/E1

T1/E1 includes wanpipemon, a utility that can capture traces from the T1/E1 line.

Section 5.3.2

T1/E1 Configuration




Figure 180: T1/E1 Trunks and Interfaces

This menu allows you to display and configure T1 or E1 Trunks as well as display the routes and status of the network interfaces.

Section 5.3.2.1

T1/E1 Network Interfaces

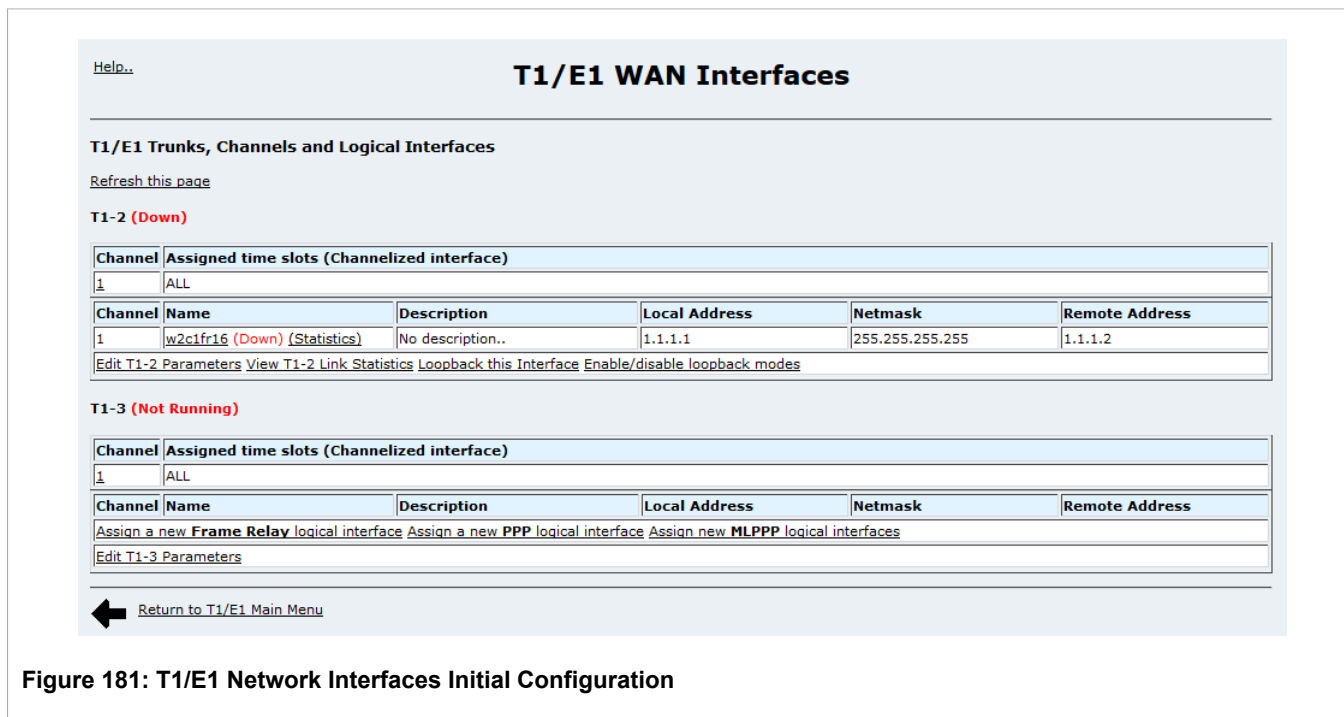


Figure 181: T1/E1 Network Interfaces Initial Configuration

This menu allows you to display and configure T1/E1 Trunk parameters, Channels and the logical interfaces that run on them. A table is presented for each interface.

Note that the interface number is the same regardless of whether it is a T1 or E1 interface. Interface numbers are as described by the "WAN" labels as shown in the home page chassis diagram.

The status of the trunks physical and logical interfaces are shown. This menu presents connection statuses but does not update them in real time. Click on the *Refresh this page* link to update to the current status.

Section 5.3.2.2

Strategy for Creating Interfaces

Initially, each interface will be configured as T1 and will have a single channel that includes all timeslots (1-24). Channelized cards can have their timeslots reassigned to make additional channels. Unchannelized cards may have timeslots removed from their single timeslot.

If the interface is to be an E1, convert it using the "Edit T1-1 Parameters" link.

If the interface is channelized and you need to have more than one channel, construct the channel groups with the desired bandwidths. This can be done by editing the single initially configured channel and removing timeslots. The unassigned timeslots will be displayed on the main menu in a link that creates channels, as shown below.

Channel	Assigned time slots (Channelized interface)				
1	1				
2	2				
Timeslots 3-24 are unused, assign a new channel					
Channel	Name	Description	Local Address	Netmask	Remote Address
Assign a new Frame Relay logical interface Assign a new PPP logical interface Assign new MLPPP logical interfaces					
Edit T1-2 Parameters					

Figure 182: T1/E1 Network Interfaces After Channel Creation

Once all timeslots have been assigned to channels, the "Timeslots.." link will no longer appear. Note that you do not have to assign all timeslots.

Assign Frame Relay or PPP to the channels by following the "Assign .. Protocol" links. The resultant menus will allow you select the desired channel.

If you are assigning multiple DLCIs, assign the first DLCI used by that interface and configure the Frame Relay Link Parameters and that DLCIs network parameters.

After assigning the first DLCI, you may revisit the interface through the link under the *Name* field and add additional DLCIs.

Once all channels have been assigned, the "Assign" links will no longer appear, as shown below. Note that any of the Frame Relay interfaces on a channel (in this case w1c4fr16 and w1c4fr17) may be used to edit the Frame Relay Link Parameters.

Help..

T1/E1 WAN Interfaces

T1/E1 Trunks, Channels and Logical Interfaces

[Refresh this page](#)

T1-1 (Up)

Channel	Assigned time slots (Channelized interface)				
1	1-24				
Channel	Name	Description	Local Address	Netmask	Remote Address
1	w1c1fr16 (Up) (Statistics)	No description..	1.1.1.1	255.255.255.255	1.1.1.2
Edit T1-1 Parameters View T1-1 Link Statistics Loopback this Interface Enable/disable loopback modes					

T1-2 (Not Running)

Channel	Assigned time slots (Channelized interface)				
1	ALL				
Channel	Name	Description	Local Address	Netmask	Remote Address
Assign a new Frame Relay logical interface Assign a new PPP logical interface Assign new MLPPP logical interfaces					
Edit T1-2 Parameters					


 [Return to T1/E1 Main Menu](#)

Figure 183: T1/E1 Network Interfaces After Interface Creation

Section 5.3.2.3

Naming of Logical Interfaces

Webmin names the logical interfaces for you (but allows you to provide a description). All interfaces start with a "w" to identify them as wan interfaces, followed by the physical interface number.

Unchannelized hardware interfaces supply only one channel (that can be composed of a varying number of timeslots) logical interface. You may configure one PPP interface or up to 992 Frame Relay DLCI interfaces. The next part of the identifier is either "ppp" or "frX" where X the frame relay channel number.

Channelized hardware allows more than one logical interface. The next part of the identifier indicates the channel the interface uses with a "c" followed by the lowest channel used. The final part of the identifier is either "ppp" or "fr" and the frame relay channel number.

**NOTE**

Once a channel is created, and an interface is constructed on it, the name of the interface will never change. This will remain true even if the number of timeslots on the channel is changed. This property is desirable since interface names used by features such as OSPF, RIP and the firewall can rely on the interface name. Channel re-assignments can, however, lead to a non-intuitive relationship between channels and timeslots.

Section 5.3.2.4

Editing a T1/E1 Interface**Figure 184: Edit T1 Interface**

This menu allows you to display and configure T1 or E1 Trunk parameters. By default the interface is set for T1 operation. The *Convert this interface to E1* link will set the interface for E1 operation and allow you to configure its settings.

If logical interfaces use a channel number larger than 24, an attempt to convert from E1 to T1 will prompt to delete the logical interface first.

T1 Settings**NOTE**

The D4 framing format option is not used. It is available solely for compatibility with legacy data models.

The *Framing* field determines the framing format used. Your line provider will indicate the correct format. Modern facilities usually employ Extended Super Frame (ESF), an enhanced T1 format that allows a line to be monitored during normal operation.

The *Line Decoding* field reflects the line encoding/decoding scheme. Almost all T1s now use B8ZS.

The *Clocking* field selects whether to accept or provide clocks. In normal use the central office provides clocks and your setting should be "Normal". You may also connect to another router by using a cross-over cable and selecting a "Master" clocking option on one of the two routers.

The *Line Build Out* field "tunes" the shape of the T1 pulses and adjusts their amplitude depending upon distances and the desired attenuation.

E1 Settings

The *Framing* and *Line Decoding* fields for E1 reflect the European variants.

The *Clocking* field performs the same function as that described for T1.

Section 5.3.2.5

Editing a Logical Interface (Frame Relay)

[Help...](#)

Edit New Logical Interface

T1-2 Channel Frame Relay Parameters

Station TypeCPE (FR DTE Interface)

Signalling typeANSI

T39110

T39216

N3916

N3926

N3934

EEK TypeOff

EEK Timer5

New Logical Interface

Channel	DLCI	Local Address	Netmask	Remote Address	Description
4			255.255.255.255		

Save

[Return to T1/E1 WAN Interfaces](#)

Figure 185: Editing a Logical Interface (Frame Relay)

This menu allows you to configure Frame Relay link and logical interface fields.

Section 5.3.2.6

Frame Relay Link Parameters

The first table presents the link parameters and applies to all logical interfaces.

The *Station Type* field determines whether the router acts as a customer premises equipment or as a frame relay switch. When a Frame Relay network provider is used, the CPE interface should be chosen. When the connection is end to end, it is typical to set the central site end to switch and the remote end to be CPE.

The *Signaling type* field reflects the Frame Relay link management protocol used, which include ANSI T1.617 Annex D, LMI and Q.933 signaling.

The *T391* (Link Integrity Verification polling) timer is valid at the CPE and indicates the number of seconds between the transmission of In-channel Signaling messages.

The *T392* (verification of polling cycle) timer is valid at the Switch and indicates the expected number of seconds between the reception of In-channel Signaling messages transmitted by the CPE.

The *N391* counter is valid at the CPE and defines the frequency of transmission of Full Status enquiry messages.

The *N392* counter is valid at both the CPE and the Switch and defines the number of errors during N393 events which cause the channel to be inactive.

The *N393* counter is valid at both the CPE and the Switch and is an event counter for measuring N392.

The *EEK Type* field controls whether End to End Keepalive messages are sent while operating as a CPE device. If this option is set to "Off", EEK is disabled. If this option is set to "Request", EEK messages are sent every *EEK Timer* x T391 seconds. This timer may be configured from 1 to 100 periods in duration.

Your network provider will inform you of what is proper for these parameters.

Section 5.3.2.7

Frame Relay DLCIs

The second table provides a listing of all DLCIs available on the channel. Only the DLCI selected from the main menu can be edited, although another DLCI can be added by following the *Add another DLCI to this channel* link.

The *DLCI Number* refers to the Data Link Connection Identifier. This number should be provided to you by your provider.

The *Local IP Address* field defines the IP address for this logical interface.

The *Netmask* field displays the network address mask. The value 255.255.255.255 indicates that the connection is point-to-point.

The *Remote IP Address* field defines the IP address for other side of this interface. As most WAN links are of point-to-point type, there is only one host connected to the other end of the link and its address is known in advance. This option is the address of the 'other end' of the link and is usually assigned by the network administrator or Internet service provider.

The *Description* field attaches a description to the logical interface viewable from the network interfaces menu.

The *Delete this logical interface* button removes the currently selected interface. Repetitive use of this button on other DLCIs assigned to the channel will free the channel up.

Section 5.3.2.8

Editing a Logical Interface (PPP)

The screenshot shows a web interface titled "Edit New Logical Interface". At the top left is a "Help.." link. Below the title is a section labeled "T1-2 Channel PPP Parameters" containing a table with the following columns: Channel, Local Address, Netmask, Remote Address, Disable Magic Number, and Description. The "Channel" column has a dropdown menu with "1" selected. The "Netmask" column contains the value "255.255.255.255". The "Disable Magic Number" column has a checkbox that is currently unchecked. Below the table is a "Save" button. At the bottom left is a back arrow icon, and at the bottom right is a link labeled "Return to T1/E1 WAN Interfaces".

Channel	Local Address	Netmask	Remote Address	Disable Magic Number	Description
1		255.255.255.255		<input type="checkbox"/>	

Figure 186: Edit Logical Interface (PPP)

The *Local Address*, *Netmask*, *Remote Address*, and *Description* fields are as described in the previous section.

Some PPP implementations exist that are unable to negotiate the LCP Magic Number feature correctly. The *Disable Magic Number* field disables PPP LCP negotiation of the Magic Number feature altogether.

Section 5.3.2.9

T1/E1 Statistics

When at least one logical interface is configured, T1/E1 Link and logical interface statistics will be available. These statistics are available from links on the T1/E1 WAN Interfaces menu.

Link Statistics are provided through the *View Link Statistics* link at the bottom of each interface table. Frame Relay and PPP statistics are available through (*Statistics*) links under the interface name column of each interface table.

The screenshot shows the 'T1/E1 WAN Interfaces' configuration page. At the top, there is a 'Help..' link and the title 'T1/E1 WAN Interfaces'. Below this is a section for 'T1/E1 Trunks, Channels and Logical Interfaces' with a 'Refresh this page' link. The main content area displays 'T1-1 (Up)' and a table for 'Assigned time slots (Channelized interface)'. This table has two columns: 'Channel' and 'Assigned time slots (Channelized interface)', with one row showing '1' and '1-24'. Below this is another table for the 'T1-1 (Up)' interface. This table has columns: 'Channel', 'Name', 'Description', 'Local Address', 'Netmask', and 'Remote Address'. It contains one row for channel '1' with the name 'w1c1fr16 (Up) (Statistics)', description 'No description..', local address '1.1.1.1', netmask '255.255.255.255', and remote address '1.1.1.2'. At the bottom of the interface table, there are several links: 'Edit T1-1 Parameters', 'View T1-1 Link Statistics', 'Loopback this Interface', and 'Enable/disable loopback modes'. A back arrow and 'Return to T1/E1 Main Menu' link are at the very bottom.

Channel	Assigned time slots (Channelized interface)
1	1-24

Channel	Name	Description	Local Address	Netmask	Remote Address
1	w1c1fr16 (Up) (Statistics)	No description..	1.1.1.1	255.255.255.255	1.1.1.2

[Edit T1-1 Parameters](#) [View T1-1 Link Statistics](#) [Loopback this Interface](#) [Enable/disable loopback modes](#)

[Return to T1/E1 Main Menu](#)

Figure 187: View T1/E1 Link Statistics

Section 5.3.2.10

Link Statistics

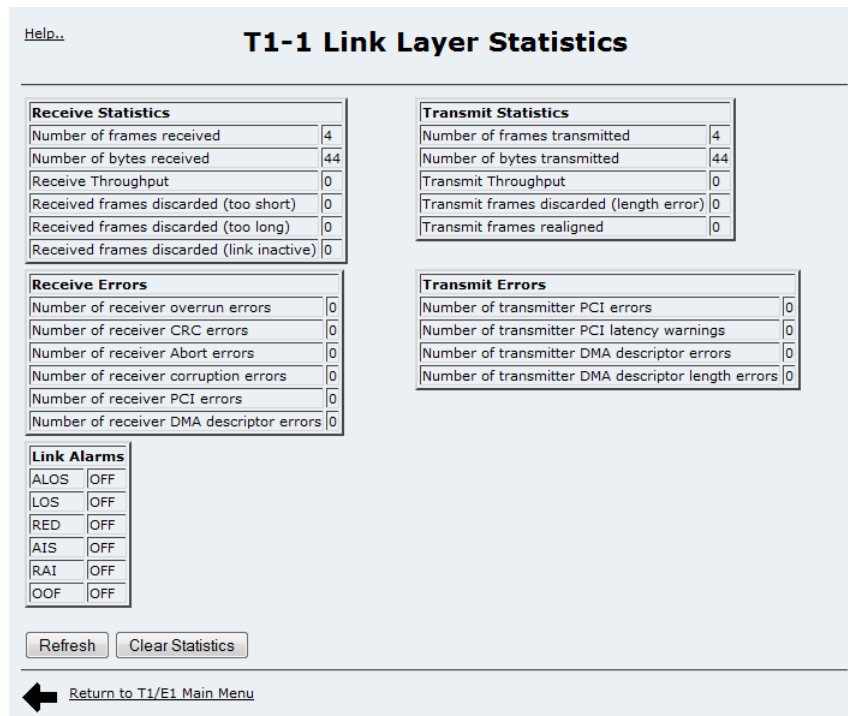


Figure 188: T1/E1 Link Statistics

The *Link Alarms* indicate ongoing problems.

ALOS/LOS (Loss of Signal) – This alarm indicates a complete absence of synchronization pulses on the line.

RED (Red Alarm) - This is a local equipment alarm. It indicates that the incoming signal has been corrupted for a number of seconds. This equipment will then begin sending a yellow alarm as its outbound signal.

AIS (Alarm Indication Signal, or BLUE alarm) - This alarm indicates the total absence of incoming signal as a series of continuous transitions (an all 1's pattern) is received.

YEL (Yellow Alarm) – This alarm is transmitted to the network and alerts it that a failure has been detected.

OOF (Out of Frame) – This alarm signifies the occurrence of a particular density of framing error events. This alarm could signify that the wrong framing mode is configured.

Section 5.3.2.11

Frame Relay Interface Statistics

[Help...](#)

w1c1fr16 Statistics

DLCI Receive Statistics		DLCI Transmit Statistics	
Information frames received	11	Information frames transmitted	9
Information bytes received	924	Information bytes transmitted	792
Received I-frames discarded due to inactive DLCI	0		
I-frames received with Discard Eligibility (DE) indicator set	0		
I-frames received with the FECN bit set	0		
I-frames received with the BECN bit set	0		

Frame Relay Trunk Statistics	
Full Status Enquiry messages received	29
Link Integrity Verification Status Enquiry msg received	143
Full Status Reply messages sent	33
Link Integrity Verification Status messages sent	139
CPE initializations	0
Current Send Sequence Number	142
Current Receive Sequence Number	143
Current N392 count	0
Current N393 count	0

Frame Relay Trunk Communications Errors	
I-frames not transmitted after a tx. int. due to excessive frame length	0
I-frames not transmitted after a tx. int. due to excessive throughput	0
Received frames discarded as they were either too short or too long	0
discarded I-frames with unconfigured DLCI	0
discarded I-frames due to a format error	0
App. didn't respond to the triggered IRQ within the given timeout period	0
discarded In-channel Signalling frames due to a format error	0
In-channel frames received with an invalid Send Seq. Numbers received	0
In-channel frames received with an invalid Receive Seq. Numbers received	0
timeouts on the T392 timer	1
consecutive timeouts on the T392 timer	0
times that N392 error threshold was reached during N393 monitored events	0

Figure 189: Frame Relay Statistics

Note that the *Frame Relay Trunk Statistics* and *Frame Relay Trunk Communications Errors* tables are common to all Frame Relay DLCIs on the trunk.

Section 5.3.2.12

PPP Interface Statistics

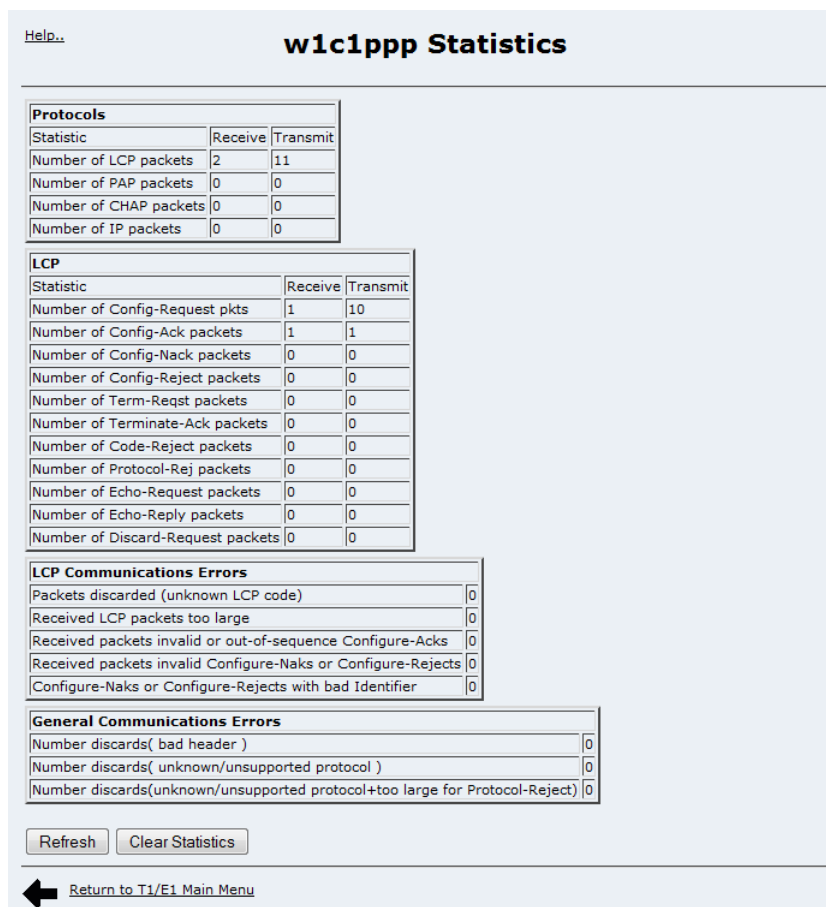


Figure 190: PPP Link Statistics

Section 5.3.2.13

T1/E1 Loopback

After configuring at least one logical interface, you can perform T1/E1 Loopback tests. After configuring an interface, a *Loopback this Interface* link appears on the Interface menu. Clicking the link displays the Loopback menu.

[Help..](#)

T1-2 Loopback

Loopback Settings

Note:
A **Digital loopback** command causes test frames to be transmitted through the digital sections of the T1/E1 interface. The frames are looped back immediately before the analog transceivers, received by the software and verified.
A **Remote loopback** command causes test frames to be transmitted through the analog transceiver to the T1/E1 line and verifies frames received from the line. You must arrange for the line to be remotely looped back (e.g. Line loopback) or employ a loopback stub with the clock mode set to "Master" for this test to succeed.
A **Line loopback** command causes frames received from the T1/E1 line to be looped back to the line. A notification is presented for each frame received during this test.

A loopback test will take down the interface, which may be undesirable when it is in use. Also note that this testing **runs independently** of the settings in 'Enable/disable loopback modes'

Select Loopback type: No loop ▾

Number of Loops: 20 (maximum 1000)

Time to run test: 20 (maximum 600 sec.)

Start Loopback

[Return to T1/E1 Main Menu](#)

Figure 191: T1/E1 Loopback Menu

The *Select Loopback Type* field selects the loopback test.


The *Number of Loops* field controls the frames sent during digital and remote loopback. This parameter is not used during line loopback.

The *Time to run test* field limits the time the sender will transmit and the router running line loopback will wait.

Running a loop test on an active interface will immediately cause it to go down. The loop test automatically initializes the trunk after completing the test.

The loopback test provides a means to test your T1/E1 digital and analog hardware and the T1/E1 line. The sender transmits a number of frames which are looped back to the sender. The router verified the returning frames for correctness.

A digital loopback starts first and verifies the digital section of the interface. If a loopback stub is inserted in the interface jack, a remote loopback verifies the interface's digital and analog sections. If the remote equipment is able to loop, the entire T1/E1 line can be verified. If the remote router is another Siemens router, starting a line loopback verifies both cards and the line. The router initiating the test displays the loopback frame count as frames arrive.

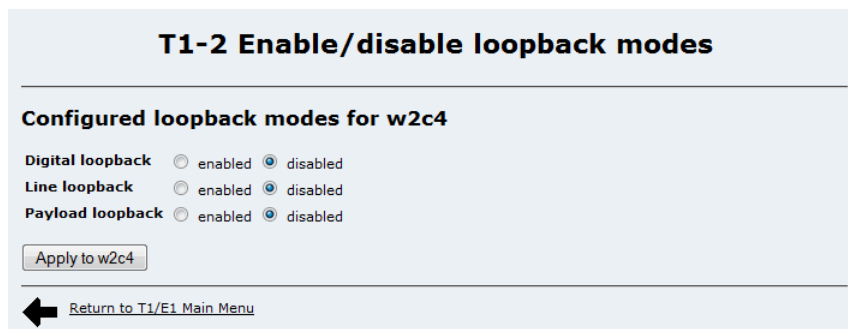
**Figure 192: T1/E1 Loopback**

Note that loopback tests are performed independently of the Enable/Disable Loopback Modes settings. Even if one or more loopback modes are disabled, the router still performs the selected loopback tests. For more information on enabling and disabling loopback modes, see [Section 5.3.2.14, "Enabling and Disabling T1/E1 Loopback Modes"](#).

Section 5.3.2.14

Enabling and Disabling T1/E1 Loopback Modes

You can enable and disable T1/E1 loopback modes from the Enable/disable loopback modes menu. After configuring an interface, an *Enable/disable loopback modes* link appears on the Interface menu. Clicking the link displays the Enable/disable Loopback Modes menu.



T1-2 Enable/disable loopback modes

Configured loopback modes for w2c4

Digital loopback ☐ enabled ☒ disabled

Line loopback ☐ enabled ☒ disabled

Payload loopback ☐ enabled ☒ disabled


 [Return to T1/E1 Main Menu](#)

Figure 193: Enable/disable Loopback Modes Menu

Select an option to enable or disable the *Digital loopback*, *Line loopback*, and *Payload loopback* modes. Click the *Apply* button. The loopback mode settings apply immediately and remain set when you reboot the appliance.

Note that enabling a loopback mode just enables the selected mode on the interface; it does not perform a loopback test. Also note that these settings are independent of the loopback tests performed from the T1/E1 Loopback menu.

Section 5.3.2.15

Upgrading Software

For some customers, access to remote sites is accomplished solely by a T1 or E1 connection. Usually a software upgrade will stop the system being upgraded, perform the upgrade, and then restart it. If T1/E1 was upgraded in this way, the upgrade would fail as the T1/E1 link was taken down. Instead, T1/E1 software upgrades modify only the software on the disk. You must schedule a reboot in order to run the new version of T1/E1 software. See [Section 3.7.10, “Upgrade System”](#) and [Section 3.7.11, “Uploading and Downloading Files”](#) for further information.

Section 5.3.2.16

Upgrading Firmware

ROX T1/E1 interfaces reside upon PCI interface cards. These cards contain FLASH memory which (from time to time) will be required to be upgraded. The upgrade process will take down the T1/E1 links, upgrade the firmware and then restart the interfaces.

**NOTE**

The upgrade process requires upwards of 15 minutes for each PCI interface card. Because of the lengthy duration required to upgrade the interfaces, ROX does not automatically perform the firmware upgrade. Instead, the scheduling of the upgrade is left to the user.

The upgrade can be performed by signing on to the platform via the console or ssh and running the command `"/usr/sbin/update-wanfirmware"`. If the ssh connection has been made over an active T1/E1 interface, the connection will fail but the upgrade will continue.

The upgrade can also be scheduled for a specific time by using the *System* menu, *Scheduled Commands* sub-menu. Set the *Commands to execute* field to `"/usr/sbin/update-wanfirmware proceed"`, set the *Run in directory* field to `"/root"` and set the *Run at time* field to the desired upgrade time.

After the upgrade completes, alarms recommending an upgrade will be cleared.

Section 5.4

Configuring Frame Relay/PPP and T3/E3

This section familiarizes the user with:

- Configuring Frame Relay and PPP Links
- Viewing status and statistics
- Upgrading Firmware

Section 5.4.1

T3/E3 Fundamentals

T3 refers to a communications link upon which has been imposed a Digital Signal 3 (DS3) signalling scheme. The scheme allows 672 time slots of 64 Kbps DS0 information to be multiplexed onto a 44.736 Mbps circuit.

E3 refers to the ITU standard corresponding to the mainly North American T3 standard. E3 calls for 512 DS0-equivalent time slots multiplexed onto a 34.368 Mbps circuit.=

ROX provides the ability to operate Frame Relay or PPP over your physical T3/E3 interfaces.



NOTE

Channel groups and fractional lines are not supported on ROX T3 and E3 interfaces.

Section 5.4.2

Location of Interfaces and Labelling

Unlike the Ethernet ports (which are statically located), the location of T1/E1, T3, DDS and ADSL ports in your router depends on the number of ports and how they are ordered. Refer to the labelled hardware image as presented in the Webmin home page.

To make labelling easy to understand, all T1/E1, T3/E3, DDS and ADSL ports are assigned a unique port number that relates to the LEDs on the status panel.

Section 5.4.3

LED Designations

ROX includes two sources of LED indicated information about T3/E3 lines, the T3/E3 card itself and the LED Panel.

One LED is associated with each line, next to the interface jack. This LED is red when the link is disconnected, flashes green when the link is connecting and remains solid green when the link is established.

ROX also indicates information about T3/E3 ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port. Consult the section [Section 3.3, "Using The LED Status Panel"](#) to determine which LEDs correspond to the port.

Section 5.4.4

T3/E3 Configuration




Figure 194: T3/E3 Trunks and Interfaces

This menu allows you to display and configure T3/E3 Trunks as well as display the routes and status of the network interfaces.

Section 5.4.4.1

T3/E3 Trunk Interfaces

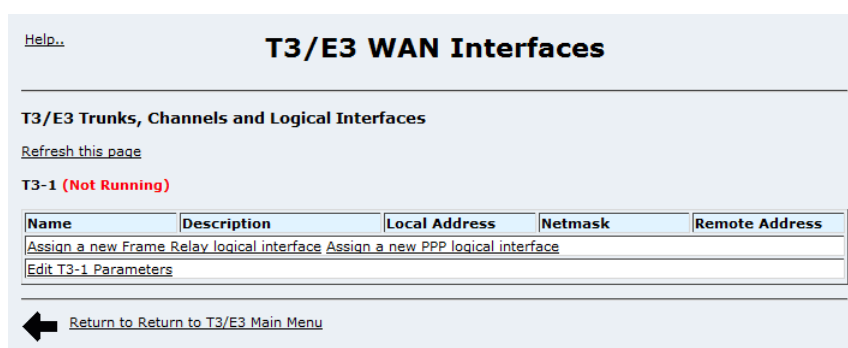


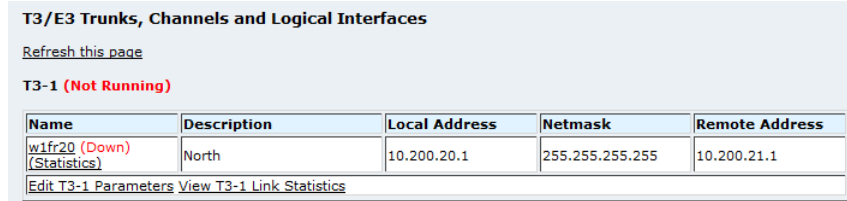
Figure 195: T3/E3 Network Interface Initial Configuration

This menu allows you to display and configure T3/E3 Trunk parameters. A table is presented for each interface. Interface numbers are as described by the "WAN" labels as shown in the home page chassis diagram.

The status of each trunk's physical and logical interface is shown. The menu presents connection status, but note that it does not update in real time. Click on the *Refresh this page* link to update the status display.

Each T3/E3 trunk may be configured as a Frame Relay link with one or more DLCIs, or as a single PPP link. Select *Assign a New Frame Relay logical interface* or *Assign a new PPP logical interface*, respectively.

The contents of the menu will change after the creation of logical interfaces, providing links to logical interface configuration and statistics and overall trunk statistics, as seen below:



T3/E3 Trunks, Channels and Logical Interfaces

[Refresh this page](#)

T3-1 (Not Running)

Name	Description	Local Address	Netmask	Remote Address
w1fr20 (Down) (Statistics)	North	10.200.20.1	255.255.255.255	10.200.21.1

[Edit T3-1 Parameters](#) [View T3-1 Link Statistics](#)

Figure 196: T3/E3 Network Interface With Logical Interfaces

Section 5.4.4.2

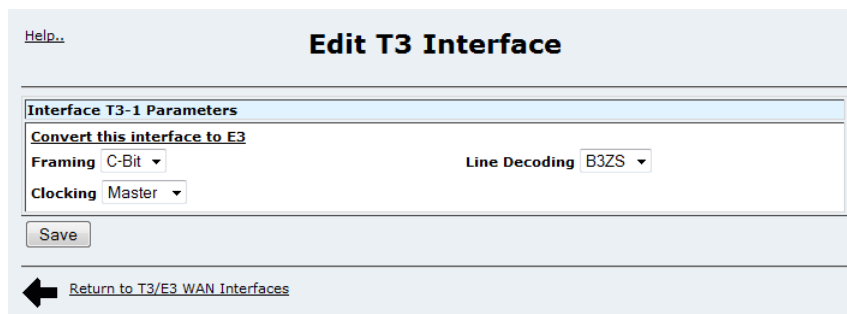
Naming of Logical Interfaces

ROX names the logical interfaces that are created for T3/E3 Trunks, but allows you to provide a description. All interfaces start with a "w" to identify them as wan interfaces, followed by the interface number. The next part of the identifier is either "ppp" or "fr" and the frame relay DLCI number.

Section 5.4.4.3

T3 Interface Parameters

The *Edit T3-X Parameters* link from the *T3/E3 WAN Interfaces* menu links to this menu, which displays and configures T3 Trunk parameters, including the option to use the interface in E3 mode.



[Help...](#)

Edit T3 Interface

Interface T3-1 Parameters

[Convert this interface to E3](#)

Framing: C-Bit Line Decoding: B3ZS

Clocking: Master

[Return to T3/E3 WAN Interfaces](#)

Figure 197: Edit T3 Interface

The *Framing* field determines the framing format used. Your line provider will indicate the correct format.

The *Line Decoding* field reflects the line encoding/decoding scheme. Almost all T3s now use B3ZS.

The *Clocking* field selects whether to accept or provide clock signal. In normal use the central office provides the clock signal in which case the setting should be "Normal". It is also possible to connect to another router, for example, by using a cross-over cable and selecting "Master" on one of the two routers to provide the clock signal.

The link: *Convert this interface to E3* reconfigures the interface for use as an E3 trunk.

Section 5.4.4.4

E3 Interface Parameters

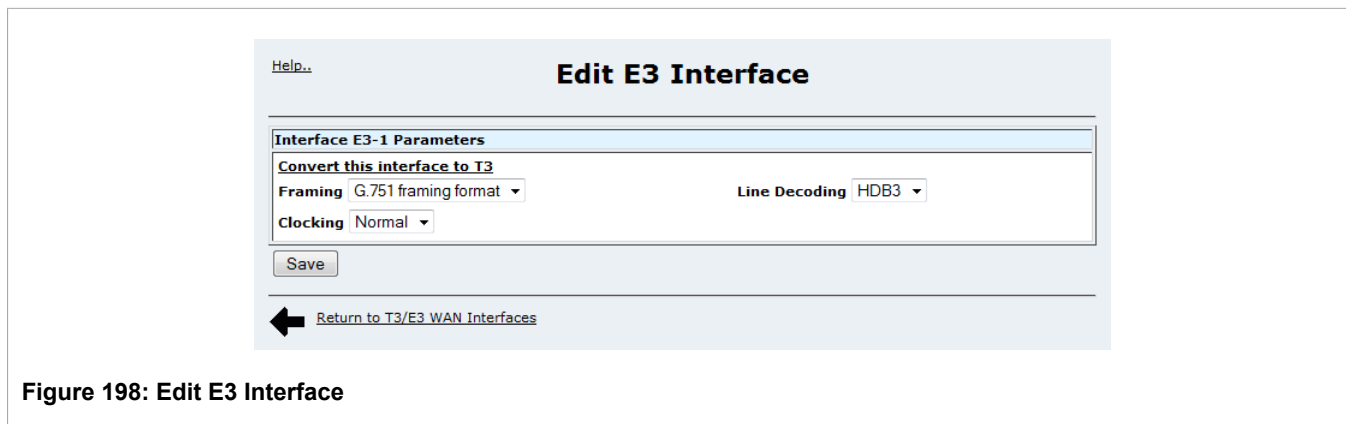


Figure 198: Edit E3 Interface

The *Framing* field determines the framing format used. Your line provider will indicate the correct format.

The *Line Decoding* field reflects the line encoding/decoding scheme.

The *Clocking* field selects whether to accept or provide clock signal. In normal use the central office provides the clock signal in which case the setting should be "Normal". It is also possible to connect to another router, for example, by using a cross-over cable and selecting "Master" on one of the two routers to provide the clock signal.

The link: *Convert this interface to T3* reconfigures the interface for use as a T3 trunk.

Section 5.4.4.5

Editing a Logical Interface (Frame Relay)

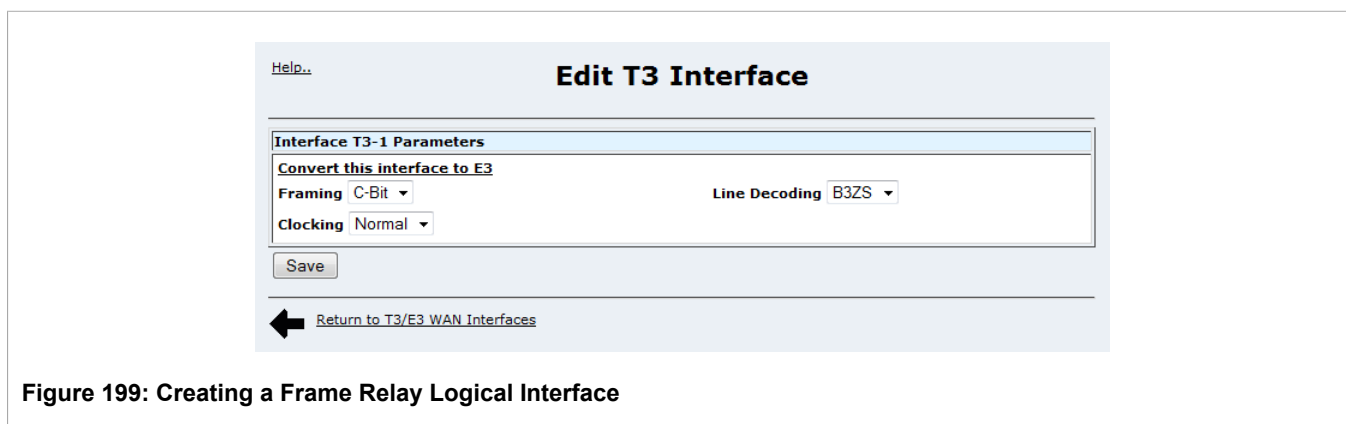


Figure 199: Creating a Frame Relay Logical Interface

This menu allows you to display and configure logical interface fields for Frame Relay. The menu is composed of two tables. The first table contains configuration parameters that apply to all DLCIs in the Frame Relay link. The second table configures network parameters of individual DLCIs.

The fields and buttons in this menu are the same as those those described in [Section 5.3.2.5, “Editing a Logical Interface \(Frame Relay\)”](#).

Once the first DLCI has been configured, revisiting the link to that DLCI from the "Trunks and Interfaces" page will display a menu that allows additional DLCIs to be configured.

[Help..](#)

Edit New Logical Interface

T3-1 Frame Relay Parameters

Station Type

CPE (FR DTE Interface)

Signalling type

ANSI

T391

10

T392

16

N391

6

N392

6

N393

4

EEK Type

Off

EEK Timer

5

New Logical Interface

DLCI	Local Address	Netmask	Remote Address	Description
		255.255.255.255		

Save

[Return to T3/E3 WAN Interfaces](#)

Figure 200: Edit Logical Interface (Frame Relay)

Section 5.4.4.6

Editing a Logical Interface (PPP)

[Help..](#)

Edit New Logical Interface

T3-1 PPP Parameters

Local Address	Netmask	Remote Address	Description
	255.255.255.255		

Save

[Return to T3/E3 WAN Interfaces](#)

Figure 201: Edit Logical Interface (PPP)

The *Local IP Address* field defines the IP address for the PPP interface.

The *Netmask* field displays the network address mask. The value 255.255.255.255 indicates that the connection is point-to-point.

The *Remote IP Address* field defines the IP address for other side of the link. This address is usually assigned by the network administrator or Internet service provider.

The *Description* field attaches a description to the logical interface viewable from the network interfaces menu.

The *Delete* button removes the currently selected interface.

Section 5.4.5

T3/E3 Statistics

When at least one logical interface is configured, T3/E3 Link and logical interface statistics will be available. These statistics are available from links on the T3/E3 WAN Interfaces menu.

Link Statistics are available via the *View T3(E3)-X Link Statistics* link at the bottom of each interface table. Frame Relay and PPP statistics are available through "(Statistics)" links under the interface name column of each interface table.

Link, Frame Relay and PPP Interface Statistics are as described in detail in the [Section 5.3.2.9, "T1/E1 Statistics"](#) section of [Section 5.3, "Configuring Frame Relay/PPP and T1/E1"](#). The differences are that the T3/E3 link reports only AIS, LOS, OOF and YEL alarms.

Section 5.4.6

Current Routes and Interface Table

The table provided by this command is the same one as described in the *Networking* menu, *Network Utilities* sub-menu. It is also provided in the T3/E3 configuration menu as a convenience.

Section 5.4.7

Upgrading Software

In some installations, the only access to a device at a remote site may be via a T3 or E3 connection. Usually a ROX system software upgrade will stop the system, perform the upgrade, and then restart it. If the T3/E3 port were to be upgraded in this way, the upgrade would fail as the T3/E3 link would be taken down. Instead, T3/E3 software upgrades modify only the software on the disk. You must schedule a reboot in order to run the new version of T3 software. See [Section 3.7.10, "Upgrade System"](#) and [Section 3.7.11, "Uploading and Downloading Files"](#) for further information.

Section 5.5

Configuring Frame Relay/PPP and DDS

This section familiarizes the user with:

- Configuring Frame Relay and PPP Links
- Viewing status and statistics
- Upgrading software

A Digital Data Services (DDS) line is a North American digital transmission method that operates at 56 Kbps synchronously over an unloaded, 4-Wire metallic-pair circuit.

The DDS line is typically a telephone grade network connection often called the "local loop". A Data Terminal Equipment (DTE) device attaches to the line and transmits data to the telephone company (TELCO), which routes the data to a remote DDS line. A short-haul, synchronous-data line driver known as a CSU/DSU terminates the line and attaches to the DTE. The DSU part of the DSU/CSU manages the format of the data signal while the CSU manages electrical levels, isolation and provides loopback to the TELCO.

The DDS port provides an integrated DTE, DSU and CSU.

Section 5.5.1

Location of Interfaces and Labelling

Unlike the Ethernet ports (which are statically located), the location of T1/E1, DDS and ADSL ports in your router depends upon the number of ports and how they were ordered. Refer to the labelled hardware image as presented in the Webmin home page.

To make labelling easy to understand, all T1/E1, T3, DDS and ADSL ports are assigned a unique port number that relates to the LEDs on the status panel.

Section 5.5.2

LED Designations

ROX indicates information about DDS ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port. Consult the section "Using The LED Status Panel" to determine which LEDs correspond to the port.

Section 5.5.3

DDS Configuration




Figure 202: DDS Trunks and Interfaces

This menu allows you to display and configure DDS Trunks. The Current Routes menu will display the routes and status of the network interfaces.

Section 5.5.3.1

DDS Network Interfaces

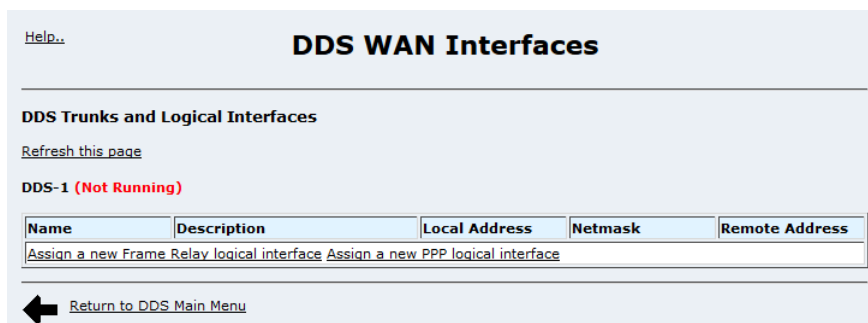


Figure 203: DDS WAN Interfaces

This menu allows you to display DDS trunks and configure the logical interfaces that run on them. A table is presented for each interface.

Interface numbers are as described by the "DDS" labels as shown in the home page chassis diagram.

The status of both the physical interface and its corresponding logical interface is shown.

If no interfaces have been configured the menu will provide links to Frame Relay and PPP configuration menus.

This menu presents connection statuses but does not update them in real time. Click on the *Refresh this page* link to update to the current status.

The menu will change after assignment of a logical interface, providing links to logical interface and link statistics.

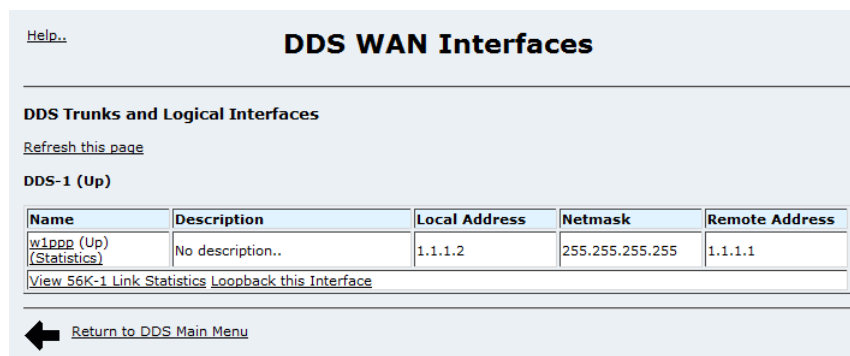


Figure 204: DDS WAN Interfaces after logical interface assignment

Webmin names the logical interfaces for you (but allows you to provide a description). All interfaces start with a "w" to identify them as wan interfaces, followed by the interface number. The next part of the identifier is either "ppp" or "fr" and the frame relay DLCI number.

Section 5.5.3.2

Editing a Logical Interface (Frame Relay)

Edit New Logical Interface

Help..

56K-1 Frame Relay Parameters

Station Type: CPE (FR DTE Interface) Signalling type: ANSI Link Failure..: Leaves IP interface up

T391: 10 T392: 16 N391: 6 N392: 6 N393: 4 EEK Type: Off EEK Timer: 5

New Logical Interface

DLCI	Local Address	Netmask	Remote Address	Description
		255.255.255.255		

Save

Return to DDS WAN Interfaces

Figure 205: Edit Logical Interface (Frame Relay), single DLCI

This menu allows you to display and configure logical interface fields for Frame Relay. The menu is composed of two tables. The first table provides link based configuration, which affect all DLCIs. The second table provides configuration parameters for individual DLCIs.

After the first DLCI has been configured, revisiting that DLCI will display a menu that allows additional DLCIs to be configured.

Edit Logical Interface w1fr17

Help..

56K-1 Frame Relay Parameters

Station Type: CPE (FR DTE Interface) Signalling type: ANSI Link Failure..: Leaves IP interface up

T391: 10 T392: 16 N391: 6 N392: 6 N393: 4 EEK Type: Off EEK Timer: 5

56K-1 Channel 1

Name	DLCI	Local Address	Netmask	Remote Address	Description
w1fr16	16	1.1.1.1	255.255.255.255	2.2.2.2	Link 123
w1fr17	17	3.3.3.3	255.255.255.255	4.4.4.4	Link 456

Add another DLCI to this channel

Save Delete this logical interface

Return to DDS WAN Interfaces

Figure 206: Edit Logical Interface (Frame Relay), multiple DLCIs

The fields and buttons in this menu are the same as those described in [Section 5.3.2.5, “Editing a Logical Interface \(Frame Relay\)”](#).

Section 5.5.3.3

Editing a Logical Interface (PPP)

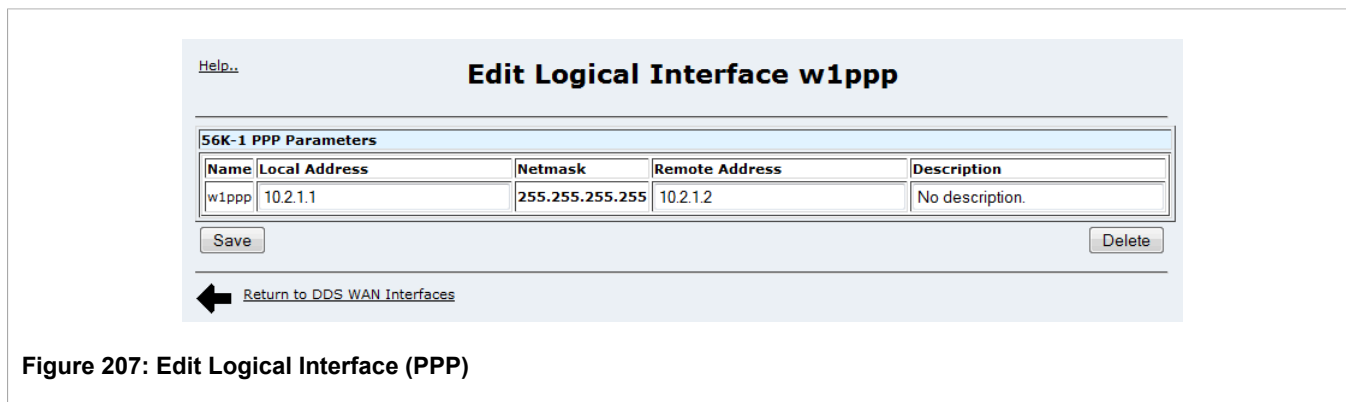


Figure 207: Edit Logical Interface (PPP)

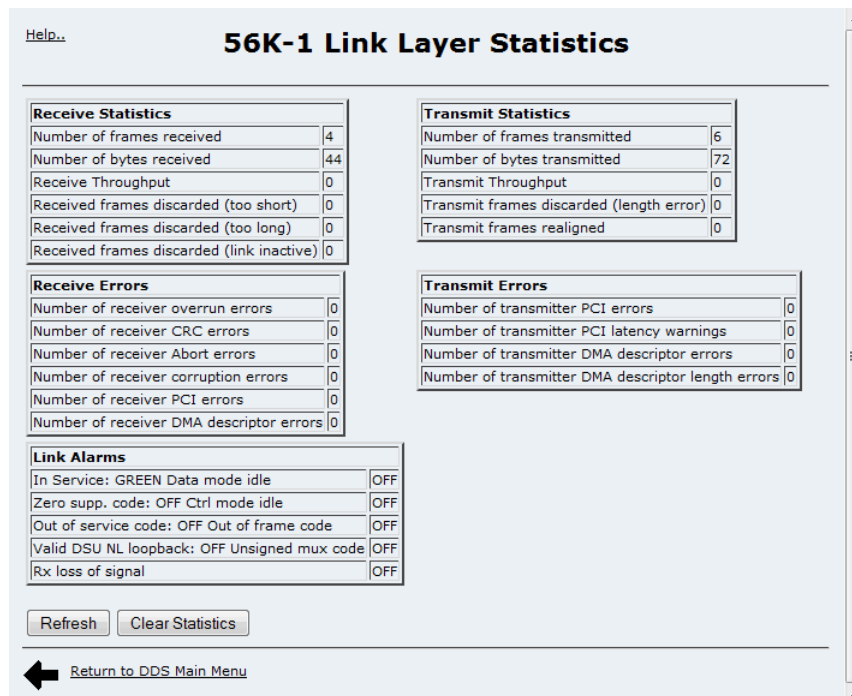
The fields and buttons in this menu are the same as those described in the *Editing A Logical Interface (PPP)* section of [Section 5.3, "Configuring Frame Relay/PPP and T1/E1"](#).

Section 5.5.3.4

DDS Statistics

When at least one logical interface is configured, DDS Link and logical interface statistics will be available. These statistics are available from links on the DDS WAN Interfaces menu.

Link Statistics are provided through the "View Link Statistics" link at the bottom of each interface table. Frame Relay and PPP statistics are available through "(Statistics)" links under the interface name column of each interface table.

**Figure 208: DDS Link Statistics**

Frame Relay and PPP Interface Statistics are as described in [Section 5.3, “Configuring Frame Relay/PPP and T1/E1”](#).

Section 5.5.3.5

DDS Loopback

When at least one logical interface is configured and that interface is active, a DDS Loopback test can be performed. This menu can be reached from a link on the DDS WAN Interfaces menu.

The remote equipment must be able to loop, allowing the entire line to be verified. If the remote equipment is another ROX, starting a line loopback will verify both cards and the line. DDS has no standard for performing digital loopback.

For more information on DDS loopback refer to [Section 5.3.2.13, “T1/E1 Loopback”](#).

Section 5.5.3.6

Current Routes and Interface Table

The table provided by this command is as described in the *Networking* menu, *Network Utilities* sub-menu. It is also provided here as a convenience.

Section 5.5.3.7

Upgrading Software

For some customers, access to remote sites is accomplished solely by a DDS connection. Usually a software upgrade will stop the system being upgraded, perform the upgrade and then restart it. If DDS port was upgraded in this way, the upgrade would fail as the DDS link was taken down. Instead, DDS software upgrades modify only the software on the disk. You must schedule a reboot in order to run the new version of DDS software. See [Section 3.7.10, "Upgrade System"](#) and [Section 3.7.11, "Uploading and Downloading Files"](#) for further information.

Section 5.6

Multilink PPP over T1/E1

This section familiarizes the user with:

- Multilink PPP in overview
- Configuring Multilink PPP
- Viewing MLPPP statistics

Section 5.6.1

Multilink PPP Fundamentals

The PPP Multilink Protocol (also known as Multilink PPP) is defined in Internet RFC 1990. Its purpose is to combine two or more PPP links into one so-called "bundle" in order to provide more bandwidth to a point to point connection.

PPP Multilink must be supported on both sides of the link, and may be used if there is more than one PPP link connecting the two endpoints. It works by multiplexing data on a per-packet basis to transmit across multiple PPP links. Sequence numbering is used to attempt to preserve the order of packets transmitted across the bundle.

ROX is capable of running PPP Multilink over two to four T1/E1 links. It is capable of defining only one MLPPP bundle.

Section 5.6.2

Notes on T1/E1 Channelization

T1/E1 lines can be configured as "channelized" or "unchannelized". A more complete discussion of this topic than the one provided below can be found in the section on [Section 5.3.2.2, "Strategy for Creating Interfaces"](#).

In *unchannelized* mode, an entire T1/E1 link is aggregated into one channel. In the *MLPPP Channel Setting* table below, unchannelized T1/E1 interfaces will be seen to have only one channel: channel 1.

In *channelized* mode, more than one channel is defined for each T1/E1 interface. The section on [Section 5.3.2.2, "Strategy for Creating Interfaces"](#) describes the process of creating multiple channels on a T1/E1 interface. Note that in order for PPP Multilink to operate optimally, it is advisable to ensure that each link in the MLPPP bundle has the same bandwidth. This means that the number of time slots, the clocking mode and rate for each T1/E1 link that is used by PPP Multilink should be the same.

Section 5.6.3

Configuring PPP Multilink over T1/E1

In order to begin creating an MLPPP bundle, click on *T1/E1* in the *Networking* folder of the main Webmin menu. *T1/E1 Trunks and Interfaces* will display the menu below:

[Help..](#)

T1/E1 WAN Interfaces

T1/E1 Trunks, Channels and Logical Interfaces

[Refresh this page](#)

T1-2 (Down)

Channel	Assigned time slots (Channelized interface)
1	ALL

Channel	Name	Description	Local Address	Netmask	Remote Address
1	w2c1fr16 (Down) (Statistics)	No description..	1.1.1.1	255.255.255.255	1.1.1.2

[Edit T1-2 Parameters](#) [View T1-2 Link Statistics](#) [Loopback this Interface](#) [Enable/disable loopback modes](#)

T1-3 (Not Running)

Channel	Assigned time slots (Channelized interface)
1	ALL

Channel	Name	Description	Local Address	Netmask	Remote Address

[Assign a new Frame Relay logical interface](#) [Assign a new PPP logical interface](#) [Assign new MLPPP logical interfaces](#)

[Edit T1-3 Parameters](#)

[Return to T1/E1 Main Menu](#)

Figure 209: T1/E1 WAN Interfaces

Click on *Assign new MLPPP logical interfaces* to specify the parameters of the MLPPP bundle. The bundle can have one or more PPP links over T1/E1.

[Help..](#)

Edit Logical Interface MLPPP

MLPPP Parameters

Local Address	Netmask	Remote Address	Description
1.1.1.1	255.255.255.255	1.1.1.10	mlppp

MLPPP Channel Setting

T1-1	T1-2
Channel 1	Channel 1

[Save](#) [Delete](#)

[Return to T1/E1 WAN Interfaces](#)

Figure 210: Edit MLPPP Logical Interface Menu

- The *Local IP address* field specifies the IP address of the MLPPP interface.
- The *Netmask* field specifies the Network Address mask.
- The *Remote Address* field specifies the IP address of the remote end of the MLPPP link.

- The *Description* field allows the administrator to store a brief description of MLPPP link.
- The *MLPPP Channel Setting* table allows one or more T1/E1 channels to be included in the MLPPP bundle.

After the fields have been entered, click the **Save** button to create the MLPPP bundle.

Section 5.6.4

Multilink PPP Statistics

Once an MLPPP interface is configured, interface statistics become available for both the T1/E1 links which comprise the MLPPP bundle and for the MLPPP interface itself.

The T1/E1 link statistics interface is described in [Section 5.3.2.10, "Link Statistics"](#).

The statistics of the PPP links comprising the MLPPP bundle can also be displayed by clicking the *(Statistics)* link below the MLPPP interface name (e.g. "w1c1mlppp") in the table for each T1/E1 interface.

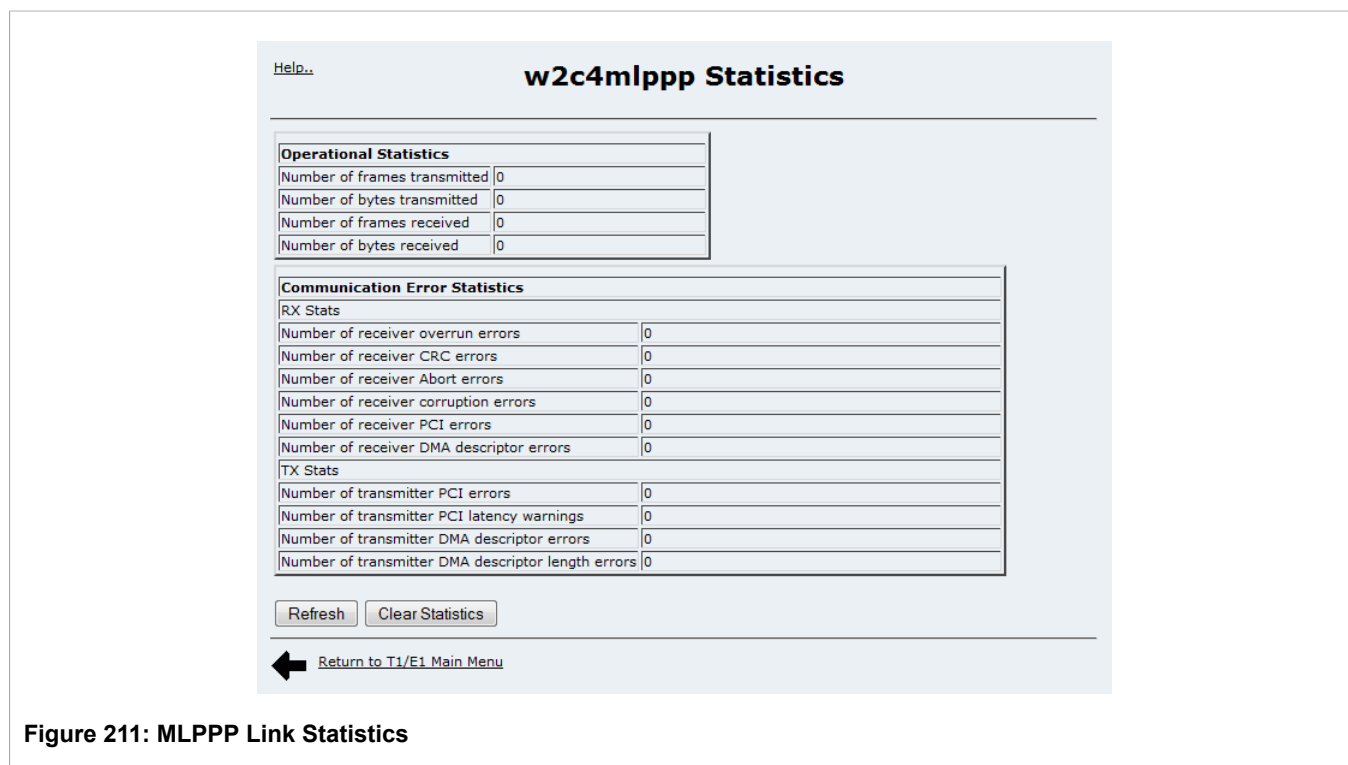


Figure 211: MLPPP Link Statistics

Section 5.7

Configuring PPPoE/Bridged Mode On ADSL

This section familiarizes the user with:

- Configuring PPPoE and Bridged Mode Links
- Viewing status

An ADSL (Asymmetric Digital Subscriber Line) line is a communications link running over regular POTS telephone service. The link is asymmetric, supporting data transfer at up to 8 Mbps from the network and up to 1 Mbps to the network. The actual bandwidth depends upon the distance between the router and telco central

office, the maximum distance of which may be up to 5480 m. An ADSL card must connect to a central ADSL DSLAM for its connection.

ADSL shares ordinary telephone lines by using frequencies above the voice band. ADSL and voice frequencies will interfere with each other. If the line will be used for both data and voice, a "splitter" should be installed to divide the line for DSL and telephone.

ADSL is almost always used to make a connection to the Internet via an ISP. There are two methods for establishing the connection, PPPoE and Bridged mode.

ADSL uses the ATM protocol to communicate with the central office DSLAM. ATM uses virtual channels to route traffic and the DSL connection needs to know which virtual channels to use. Most providers use VPI=0 and VCI=35. There are exceptions to this. Some providers that use different settings are listed in the following table.

Provider	VPI	VCI
Typical Provider	0	35
Bell South	8	35
New Edge	0	38
Sprint	8	35
US West/Qwest	0	32

Section 5.7.1

PPPoE/Bridged Mode Fundamentals

In PPPoE (Point-to-Point Protocol Over Ethernet) the PPP dial-up protocol is used with Ethernet over ADSL as the transport. PPPoE supports the protocol layers and authentication widely used in PPP and enables a point-to-point connection to be established in the normally multipoint architecture of Ethernet.

As your PPPoE connection is established a PPP interface will be created. The name will be "pppX" where X is the same as the interface number. Use this interface name in firewall rules.

Section 5.7.1.1

Authentication, Addresses and DNS Servers

PPP authentication utilizes PAP or CHAP. Your ISP will provide you with a user-ID and password which you will enter in the GUI. The authentication process will assign a local IP address and addresses of the ISPs DNS servers to the router. You should use these DNS servers unless you wish to provide your own.

You will obtain either a dynamic or static IP from your ISP. Firewall configuration should be performed as is appropriate.

Section 5.7.1.2

PPPoE MTU Issues

The use of PPPoE introduces a limitation of the maximum length of packets. The maximum Ethernet frame is 1518 bytes long. 14 bytes are consumed by the header, and 4 by the frame-check sequence, leaving 1500 bytes for the payload. For this reason, the Maximum Transmission Unit (MTU) of an Ethernet interface is usually 1500 bytes.

This is the largest IP datagram which can be transmitted over the interface without fragmentation. PPPoE adds another six bytes of overhead, and the PPP protocol field consumes two bytes, leaving 1492 bytes for the IP datagram. This reduces the MTU of PPPoE interfaces to 1492 bytes.

Packets received by hosts via Ethernet that are sized to the Ethernet MTU will be too large for the PPPoE connections MTU and will be fragmented. Large packets from hosts on the Internet will be fragmented by the ISP. The router will re-assemble these packets, but at the cost of increased latency. Configuring smaller MTUs at your hosts may reduce latency.

Section 5.7.1.3

Bridged Mode

In bridged mode, the router simply employs the ADSL interface as a carrier of Ethernet frames. The interface will be created at boot time with a 1500 byte MTU.

No authentication information is required for bridged mode.

Your ISP will provide you with one or more IP addresses and an appropriate subnet mask. Your ISP will also suggest a DNS server which you can configure via the *Networking, Network Configuration, DNS Client* menu.

Section 5.7.1.4

Location of Interfaces and Labelling

Unlike the Ethernet ports (which are statically located), the location of ADSL ports in your router depends upon the number of ports and how they were ordered. Refer to the labelled hardware image as presented in the Webmin home page.

To make labelling easy to understand, all T1E1, T3, DDS and ADSL ports are assigned a unique port number that relates to the LEDs on the status panel.

Section 5.7.1.5

LED Designations

ROX includes two sources of LED indicated information about ADSL lines, the ADSL card itself and the LED Panel.

Four LEDs are associated with the line, next to the interface jack.

Power (Green) indicates when the card is active and powered.

Link (Green) indicates when the DSL link is established.

TX (Red) indicates when data is being transmitted over DSL.

RX (Red) indicates when data is being received over DSL.

While connecting the LEDs are flashing sequentially.

ROX also indicates information about ADSL ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port. Consult the section [Section 3.3, "Using The LED Status Panel"](#) to determine which LEDs correspond to the port.

Section 5.7.2

ADSL Configuration

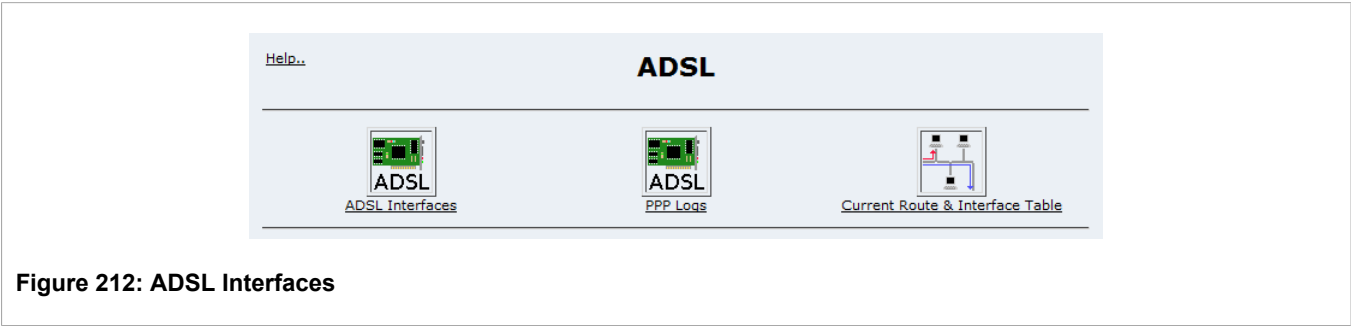


Figure 212: ADSL Interfaces

This menu allows you to display and configure ADSL interfaces. The PPP Logs menu will display a log of PPP related information. The Current Routes menu will display the routes and status of the network interfaces.

Section 5.7.2.1

ADSL Network Interfaces

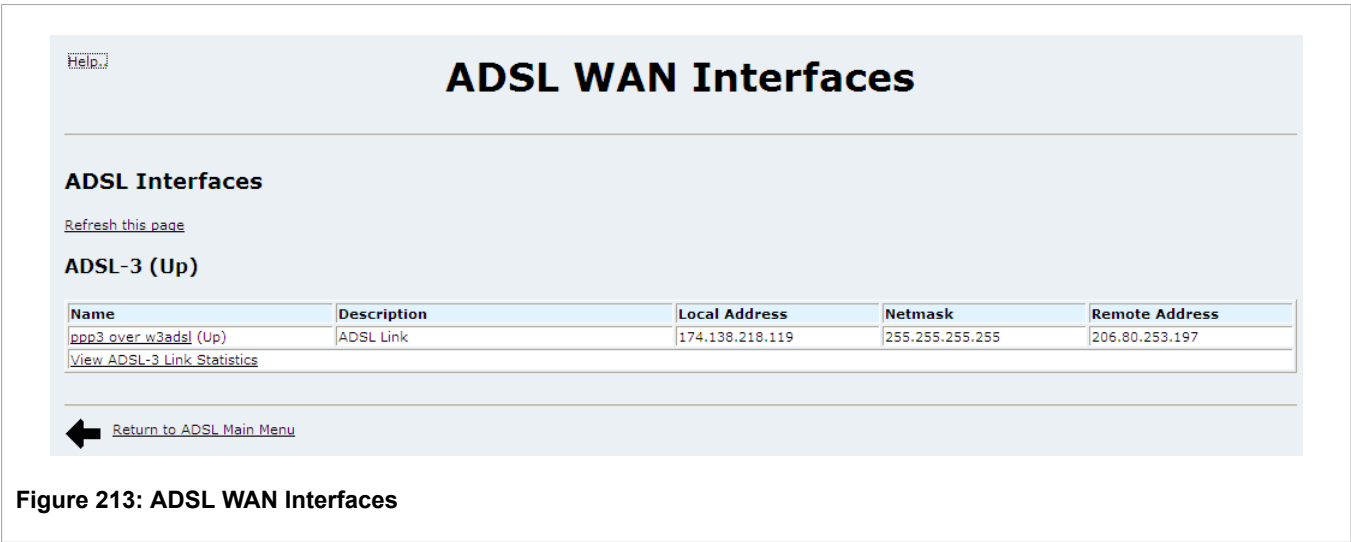


Figure 213: ADSL WAN Interfaces

This menu allows you to display and configure ADSL interfaces and the protocols that run on them. A table is presented for each interface.

Interface numbers are as described by the "ADSL" labels as shown in the home page chassis diagram.

The status of the physical interface, its corresponding logical interface and link statistics are provided.

This menu presents connection statuses but does not update them in real time. Click on the *Refresh this page* link to update to the current status.

Section 5.7.2.2

Editing a Logical Interface (PPPoE)

Figure 214: Edit Logical Interface (PPPoE)

This menu allows you to display and configure logical interface fields for PPPoE and to convert the interface to Bridged Mode.

By default, interfaces are created with PPPoE. If you want the interface to be Bridged Mode, click on the *Convert this interface to bridged* link.

The *Description* field attaches a description to the logical interface viewable from the network interfaces menu.

The *VPI* field determines the VPI number the connection uses. The default of 0 is correct for most providers. The *VCI* field determines the VCI number the connection uses. The default of 35 is correct for most providers.

The *Attempt ATM Autoconfiguration* option causes the router to attempt to automatically determine the VPI and VCI used on the connection. This does not work with all providers and may cause the connection to fail even if the link light is on. If this option is used it should only be used to find out what the correct values are if your provider isn't willing to help you, and when the correct values are found it should be disabled with the correct values entered in the VPI and VCI fields instead.

The *PPPoE Username* field determines the username to use when connecting to the PPPoE server as specified by your provider.

The *Password* field determines the password provided to the PPPoE server.

The *Default Route* check box enables automatically setting a default route using this interface whenever it connects. If this is your primary connection you probably want this option enabled.

The *Use peer DNS* check box enables automatically setting the DNS server entries that the PPPoE server recommends. Enable this option unless you provide your own name servers.

The *MTU* field defines the MTU size to request when connecting to the PPPoE server. In some cases the PPPoE provider may provide a smaller MTU in which case the smaller setting will be used, or it may refuse to alter the MTU and use whatever it considers to be the default.



NOTE

If the negotiated MTU is different from the requested MTU, a warning will be displayed on the Networking, ADSL menu.

Section 5.7.2.3

Editing a Logical Interface (Bridged)

The screenshot shows a web interface titled "Edit Logical Interface". At the top left is a "Help.." link. Below the title is a section "Interface w3adsl Parameters". Inside this section, there is a link "Convert this interface to PPPoE". The form contains several fields: "Description" (a text area), "VPI" (a text box with "0"), "Comment" (a text area), "VCI" (a text box with "35"), "Attempt ATM Autoconfiguration" (a checkbox), "Use DHCP" (a checked checkbox), "Local IP Address" (a text box with "169.254.0.1"), "Netmask" (a text box with "255.255.0.0"), "Remote IP Address" (a text box with "169.254.0.2"), and "Use as Default Route" (radio buttons for "No" and "Gateway", with "No" selected). At the bottom of the form are "Save" and "delete" buttons. Below the form is a back arrow and a link "Return to ADSL WAN Interfaces".

Figure 215: Edit Logical Interface (Bridged)

The *Description* field attaches a description to the logical interface viewable from the network interfaces menu.

The *VPI* field determines the VPI number the connection uses. The default of 0 is correct for most providers.

The *Attempt ATM Autoconfiguration* option causes the router to attempt to automatically determine the VPI and VCI used on the connection. This does not work with all providers and may cause the connection to fail even if the link light is on. If this option is used it should only be used to find out what the correct values are if your provider isn't willing to help you, and when the correct values are found it should be disabled with the correct values entered in the VPI and VCI fields instead.

The *VCI* field determines the VCI number the connection uses. The default of 35 is correct for most providers.

The *Use DHCP* field forces the router to fetch its IP address from the peer via DHCP. Note that DHCP is selected the local and remote IP addresses are immediately dummed out to 169.254.0.1 and 169.254.0.2, the netmask is set to 255.255.0.0 and default gateway option is suppressed.

The *Local IP Address* field defines the IP address for this interface.

The *Netmask* field defines the network address mask. The value 255.255.255.255 specifies a point-to-point connection which is almost always correct.

The *Remote IP Address* field defines the IP address for other side of this interface. As most WAN links are of point-to-point type, there is only one host connected to the other end of the link and its address is known in advance. This option is the address of the 'other end' of the link and is usually assigned by the network administrator or Internet service provider.

The *Gateway IP Address* field defines the IP address to use as the gateway for sending to other sites. This is usually the same as the Remote IP Address.

Section 5.7.2.4

ADSL Statistics

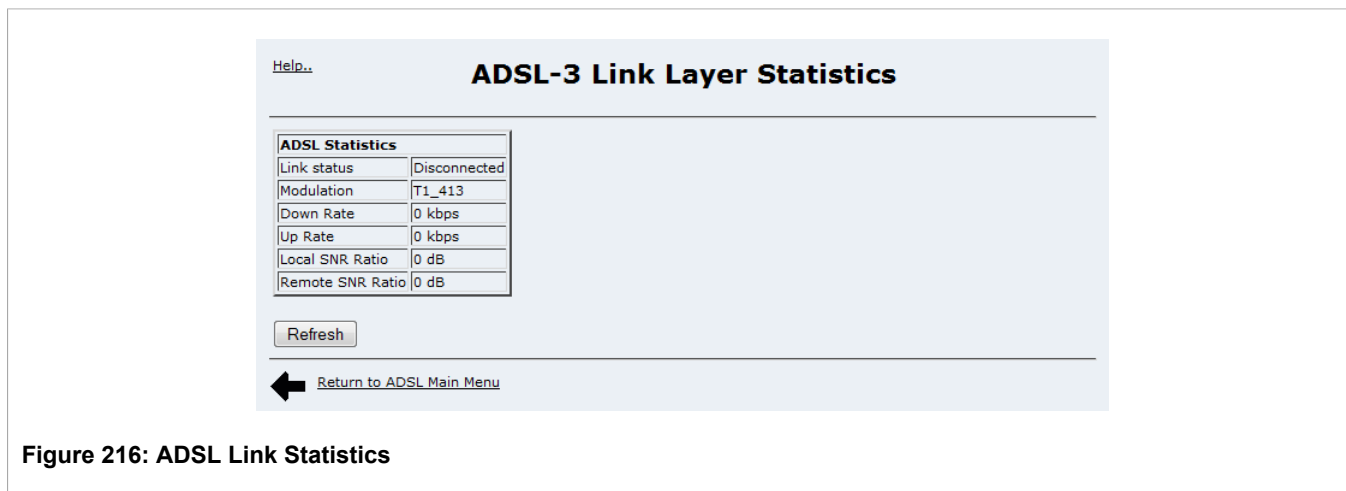


Figure 216: ADSL Link Statistics

When at least one logical interface is configured, ADSL Link statistics will be available. These statistics are available from links on the ADSL WAN Interfaces menu.

The *Local SNR Ratio* is an effective indicator of line quality. SNR values above 40 db correspond to excellent line quality, while values below 10 db result in marginal operation or failure.

Section 5.7.2.5

Current Routes and Interface Table

The table provided by this command is as described in the *Networking* menu, *Network Utilities* sub-menu. It is also provided here as a convenience.

Section 5.7.2.6

Upgrading Software

For some customers, access to remote sites is accomplished solely by an ADSL connection. Usually a software upgrade will stop the system being upgraded, perform the upgrade and then restart it. If ADSL was upgraded in this way, the upgrade would fail as the ADSL link was taken down. Instead, ADSL software upgrades modify only the software on the disk. You must schedule a reboot in order to run the new version of ADSL software. See [Section 3.7.10, "Upgrade System"](#) and [Section 3.7.11, "Uploading and Downloading Files"](#) for further information.

Section 5.8

Configuring the Firewall

This section familiarizes the user with:

- Enabling/Disabling The Firewall
- Elements of Firewall design
- How to configure the Firewall

- Checking Firewall configuration

Section 5.8.1

Firewall Fundamentals

Firewalls are software systems designed to prevent unauthorized access to or from private networks. Firewalls are most often used to prevent unauthorized Internet users from accessing private networks (intranets) connected to the Internet.

When the ROX firewall is used, the router serves a *gateway* machine through which all messages entering or leaving the intranet pass. The router examines each message and blocks those that do not meet the specified security criteria. The router also acts as a *proxy*, preventing direct communication between computers on the Internet and intranet. Proxy servers can filter the kinds of communication that are allowed between two computers and perform address translation.

Section 5.8.1.1

Stateless vs Stateful Firewalls

Firewalls fall into two broad categories: stateless and stateful (session-based).

Stateless or "static" firewalls make decisions about a traffic without regard to the history, simply opening a "hole" for the traffic's type (based upon TCP or UDP port number). Stateless firewalling is a relatively simple affair, easily handling web and email traffic. Stateless firewalls suffer from disadvantages, however. All holes opened in the firewall always open, there is no opening and closing connections based on outside criteria. Static IP filters offer no form of authentication.

Stateful firewalling adds considerable complexity the firewalling process by tracking the state of each connection.

A stateful firewall also looks at each packet and apply tests, but the tests applied or "rules" may be modified depending on packets that have already been processed. This is called "connection tracking". Stateful firewalls can also recognize that traffic on connected sets of TCP/UDP ports is from a particular protocol and manage it as a whole.

Section 5.8.1.2

Linux™ netfilter, iptables and the Shoreline Firewall

ROX employs a stateful firewall system known as *netfilter*, a set of loadable kernel modules that provides capabilities to allow session-based packet examination. The netfilter system is an interface built into the Linux kernel that allows the IP network stack to provide access to packets.

The netfilter system uses rulesets, collections of packet classification rules that determine the outcome of examination of a specific packet. The rules are defined by *iptables*, a generic table structure syntax and utility program for the configuration and control of netfilter.

In practice an iptables rule file and a script are all that are needed to load the netfilter system with rules on upon router start up. The iptables rules, however, are somewhat difficult to configure and manage.

The Shoreline Firewall, often known as shorewall, offers a more convenient approach. Shorewall is really just a front end to netfilter, maintaining the information used to generate the iptables rules in a less complicated form.

Shorewall itself does not provide a graphical front end, and instead assumes administrators will have a fair amount of familiarity with reading and editing Linux configuration files. ROX comes with a GUI front that simplifies some of the management aspects.

Section 5.8.1.3

Network Address Translation

Network Address Translation (NAT), enables a LAN to use one set of IP addresses for internal traffic and a second set for external traffic. The NAT function of netfilter makes all necessary IP address translations as traffic passes between the intranet and Internet. NAT is often referred to in Linux as IP Masquerading.

NAT itself provides a type of firewall by hiding internal IP addresses.

More importantly, NAT enables a network to use more internal IP addresses. Since they're used internally only, there's no possibility of conflict with IP addresses used by other organizations. Typically, your internal network will be setup to use one or more of the reserved address blocks described in RFC1918, namely:

10.0.0.0/8 (10.0.0.0 - 10.255.255.255)

172.16.0.0/12 (172.16.0.0 - 172.31.255.255)

192.168.0.0/16 (192.168.0.0 - 192.168.255.255)

As packets with these address reach the NAT gateway their source address and source TCP/UDP port number is recorded and the address/port number is translated to the public IP address and an unused port number on the public interface. When the Internet host replies to the internal machine's packets, they will be addressed to the NAT gateway's external IP at the translation port number. The NAT gateway will then search its tables and make the opposite changes it made to the outgoing packets and forward the reply packets on to the internal machine.

Translation of ICMP packets happens in a similar fashion but without the source port modification.

NAT can be used in static and dynamic modes. Static NAT masks the private IP addresses by translating each internal address to a unique external address. Dynamic NAT translates all internal addresses to one (or more) external address(es).

Section 5.8.1.4

Port Forwarding

Port forwarding (also known as redirection) allows traffic coming from the Internet to be sent to a host behind the NAT gateway.

Previous examples have described the NAT process when connections are made from the Intranet to the Internet. In those examples, addresses and ports were unambiguous.

When connections are attempted from the Internet to the Intranet, the NAT gateway will have multiple hosts on the Intranet that could accept the connection. It needs additional information to identify the specific host to accept the connection.

Suppose that two hosts, 192.168.1.10 and 192.168.1.20 are located behind a NAT gateway having a public interface of 213.18.101.62. When a connection request for http port 80 arrives at 213.18.101.62, the NAT gateway could forward the request to either of the hosts (or could accept it itself). Port forwarding configuration could be used to redirect the requests to port 80 to the first host.

Port forwarding can also remap port numbers. The second host may also need to answer http requests. As connections to port 80 are directed to the first host, another port number (such as 8080) can be dedicated to the second host. As requests arrive at the gateway for port 8080, the gateway remaps the port number to 80 and forwards the request to the second host.

Finally, port forwarding can take the source address into account. Another way to solve the above problem could be to dedicate two hosts 200.0.0.1 and 200.0.0.2 and have the NAT gateway forward requests on port 80 from 200.0.0.1 to 192.168.1.10 and from 200.0.0.2 to 192.168.1.20.

Section 5.8.2

Shorewall Quick Setup

For users familiar with Shorewall, the following serves as a reminder of how to build the firewall. New users may wish to read [Section 5.8.3, “ShoreWall Terminology and Concepts”](#) before continuing.

1. Logically partition your network into zones. Will you establish a DMZ? Will all Ethernet interfaces need to forward traffic to the public network? Which interfaces are to be treated in a similar fashion?
2. Assign your interfaces to the zones. If using T1/E1, have you created your T1/E1 interfaces prior to building the firewall?
3. Set the default policies for traffic from zone to zone to be as restrictive as possible. Has the local zone been blocked from connecting to the DMZ or firewall? Does the DMZ or firewall need to accept connections? Which connections should be dropped and which reset? What logs are kept?
4. How is the network interface IP assigned, i.e. dynamically or statically? Do hosts at the central site need to know the local address?
5. If your network interface IP is dynamically assigned, configure masquerading.
6. If your network interface IP is statically assigned, configure Source Network address Translation (SNAT). If a sufficient number of IP addresses are provided by the ISP, static NAT can be employed instead.
7. If your hosts must accept sessions from the Internet configure the rules file to support Destination Network address Translation (DNAT). Which hosts need to accept connections, from whom and on which ports?
8. Configure the rules file to override the default policies. Have external connections been limited to approved IP address ranges. Have all but the required protocols been blocked?
9. If you are supporting a VPN, add additional rules.
10. Check the configuration using the *Shorewall Firewall* menu, "Check Firewall" button.
11. Activate the firewall. It is usually a good idea to port scan the firewall after activation and verify that logging is functioning.

Section 5.8.3

ShoreWall Terminology and Concepts

This section provides background on various Shorewall terms and concepts. References are made to the section where configuration applies.

Section 5.8.3.1

Zones

A network zone is a collection of interfaces, for which forwarding decisions are made, for example:

Name	Description
net	The Internet
loc	Your Local Network
dmz	Demilitarized Zone
fw	The firewall itself

Name	Description
vpn1	IPSec connections on w1ppp
vpn2	IPSec connections on w2ppp

You may create new zones if you wish. For example if all of your Ethernet interfaces are part of the local network zone, disallowing traffic from the Internet zone to the local zone will disallow it to all Ethernet interfaces. If you wanted some interfaces (but not others) to access the Internet, you could create another zone.

Zones are defined in the file `/etc/shorewall/zones` and are modified from the *Network Zones* menu.

Section 5.8.3.2

Interfaces

Shorewall Interfaces are simply the Ethernet and WAN interfaces available to the router. You must place each interface into a network zone.

If an interface supports more than one subnet, place the interface in zone 'Any' and use the zone hosts setup (see below) to define a zone for each subnet on the interface.

An example follows:

Interface	Zone
eth1	loc
eth2	loc
eth3	Any
eth4	dmz
w1ppp	net

Interfaces are defined in the file `/etc/shorewall/interfaces` and are modified from the *Network Interfaces* menu.

Section 5.8.3.3

Hosts

Shorewall hosts are used to assign zones to individual hosts or subnets, on an interface which handles multiple subnets. This allows the firewall to manage traffic being forwarded back out the interface it arrived on, but destined for another subnet. This is often useful for VPN setups to handle the VPN traffic separately from the other traffic on the interface which carries the VPN traffic. An example follows:

Zone	Interface	IP Address or Network
local	eth3	10.0.0.0/8
guests	eth3	192.168.0.0/24

Interfaces are defined in the file `/etc/shorewall/hosts` and are modified from the *Network Hosts* menu.

Section 5.8.3.4

Policy

Shorewall policies are the default actions for connection establishment between different firewall zones. Each policy is of the form:

Source-zone Destination-zone Default-action

You can define a policy from each zone to each other. You may also use a wildcard zone of "all" to represent all zones.

The default action describes how to handle the connection request. There are six types of actions: ACCEPT, DROP, REJECT, QUEUE, CONTINUE and NONE. The first three are the most widely used and are described here.

When the *ACCEPT* policy is used, a connection is allowed. When the *DROP* policy is used, a request is simply ignored. No notification is made to the requesting client. When the *REJECT* policy is used, a request is rejected with an TCP RST or an ICMP destination-unreachable packet being returned to the client.

An example should illustrate the use of policies.

Source Zone	Destination Zone	Policy
loc	net	ACCEPT
net	all	DROP
all	all	REJECT

The above policies will:

- Allow connection requests *only* from your local network to the Internet. If you wanted to allow requests from a console on the ROX to Internet you would need to add a policy of ACCEPT fw zone to net zone.
- Drop (ignore) all connection requests from the Internet to your firewall or local network, and
- Reject all other connection requests.

Note that a client on the Internet that is probing the ROX's TCP/UDP ports will receive no responses and will not be able to detect the presence of the router. A host in the local network, on the other hand, will fail to connect to the router but will receive a notification.

Note that order of policies is important. If the last rule of this example were entered first then no connections at all would be allowed.

Policies are defined in the file `/etc/shorewall/policy` and are modified from the *Default Policy* menu.

Section 5.8.3.5

Masquerading and SNAT

Masquerading and Source NAT (SNAT) are forms of dynamic NAT.

Masquerading substitutes a single IP address for an entire internal network. Use masquerading when your ISP assigns you an IP address dynamically at connection time.

SNAT substitutes a single address or range of addresses that you been assigned by your ISP. Use SNAT when your ISP assigns you one or more static IP addresses that you wish to one or more internal hosts.

The masquerading/SNAT entries are defined in the file `/etc/shorewall/masq` and are modified from the *Masquerading* menu. Each entry is of the form:

Interface Subnet Address Protocol Port(s)

Interface is the outgoing (WAN or Ethernet) interface and is usually your Internet interface.

Subnet is the subnet that you wish to hide. It can be an interface name (such as eth1) or a subnetted IP address.

**NOTE**

It is always recommended to use subnetted IP address instead of interface name. When using interface name the interface has to UP and running for the firewall to start without any problems.

Address is an (optional IP) address that you wish to masquerade as.

**NOTE**

The presence of the Address field determines whether masquerading or SNAT is being used. Masquerading is used when only Interface and Subnet are present. SNAT is used when Interface, Subnet and Address are present.

Protocol (optionally) takes on the name of protocols (e.g. tcp, udp..) that you wish to masquerade.

Ports (optionally) takes on the ports to masquerade when protocol is set to tcp or udp. These can be raw port numbers or names as found in file /etc/services.

Some examples should illustrate the use of masquerading:

Rule	Interface	Subnet	Address	Protocol	Ports
1	eth1	192.168.1.0/24			
2	ppp+	192.168.0.0/24	66.11.180.161		
3	w1ppp	192.168.2.0/24	100.1.101.16		
4	w1ppp	192.168.2.0/24	100.1.101.16	tcp	smtp

1. In this masquerading rule, 192.168.1.0/24 subnet is the local network and eth1 is connected to a DSL modem. Traffic from the subnet should be translated to whatever IP is assigned to the modem. Internet clients will not be able to determine the router's public address unless some form of dynamic dns is employed.
2. In this SNAT rule a static address of 66.11.180.161 is acquired from the ISP. Traffic from the subnet 192.168.0.0/24 should be translated to 66.11.180.161 as it sent to the Internet over ppp. The + at the end of "ppp+" causes Shorewall to match any ppp interface.
3. In this SNAT rule, traffic from the subnet (192.168.2.0/24) should be translated to 100.1.101.16 as it sent to the Internet on t1/e1 port w1ppp.
4. This example is much the same as the previous one except that only SMTP from 192.168.2.0/24 subnet will be allowed.

Masquerading and SNAT rules are defined in the file /etc/shorewall/masq and are modified from the *Masquerading* menu.

Section 5.8.3.6

Rules

The default policies can completely configure traffic based upon zones. But the default policies cannot take into account criteria such as the type of protocol, IP source/destination addresses and the need to perform special actions such as port forwarding. The Shorewall rules can accomplish this.

The Shorewall rules provide exceptions to the default policies. In actuality, when a connection request arrives, the rules file is inspected first. If no match is found then the default policy is applied. Rules are of the form:

Action Source-Zone Destination-Zone Protocol Destination-Port Source-Port Original-Destination-IP Rate-Limit User-Group

Actions are ACCEPT, DROP, REJECT, DNAT, DNAT-, REDIRECT, REDIRECT-, CONTINUE, LOG and QUEUE. The DNAT-, REDIRECT-, CONTINUE, LOG and QUEUE actions are not widely used and are not described here.

**NOTE**

When applying new rules, previous traffic seen by the router might still be considered as having valid connections by the connection tracking table. For instance, a rule for the TCP and UDP protocols is applied. The router sees both TCP and UDP traffic that qualifies for NAT. The rule is then modified to allow only UDP. The router will still see TCP packets (i.e. retransmission packets).

If required, reboot the router to flush all existing connection streams.

Action	Description
ACCEPT	Allow the connection request to proceed.
DROP	The connection request is simply ignored. No notification is made to the requesting client.
REJECT	The connection request is rejected with an RST (TCP) or an ICMP destination-unreachable packet being returned to the client.
DNAT	Forward the request to another system (and optionally another port).
REDIRECT	Redirect the request to a local tcp port number on the local firewall. This is most often used to "remap" port numbers for services on the firewall itself.

The remaining fields of a rule are as described below:

Field	Description
Action	The action as described in the previous table.
Source-Zone	The zone the connection originated from.
Destination-Zone	The zone the connection is destined for.
Protocol	The tcp or udp protocol type.
Destination-Port	The tcp/udp port the connection is destined for.
Source-Port	The tcp/udp port the connection originated from.
Original-Destination-IP	The destination IP address in the connection request as it was received by the firewall.
Rate-Limit	A specification which allows the rate at which connections are made to be limited.
User-Group	A method of limiting outbound traffic from the firewall to a specific user, group of users and a specific application.

The following examples illustrate the effects of the rules file:

Rule	Action	Source-Zone	Destination-Zone	Protocol	Dest-Port	Source-Port	Original-Destination-IP
1	ACCEPT	net:204.18.45.0/24	fw				
2	DNAT	net	loc:192.168.1.3	tcp	ssh, http		
3	DNAT	net:204.18.45.0/24	loc:192.168.1.3	tcp	http	-	130.252.100.69

Rule	Action	Source-Zone	Destination-Zone	Protocol	Dest-Port	Source-Port	Original-Destination-IP
4	ACCEPT	fw	net	icmp			
5	ACCEPT	net:204.18.45.0/24	fw	icmp	8		

1. This rule accepts traffic to the firewall itself from the 204.18.45.0/24 subnet. If the default policy is to drop all requests from net to the firewall, this rule will only accept traffic from the authorized subnet.
2. This rule forwards all ssh and http connection requests from the Internet to local system 192.168.1.3.
3. This rule forwards http traffic from 204.18.45.0/24 (which was originally directed to the firewall at 130.252.100.69) to the host at 192.168.1.3 in the local zone. If the firewall supports another public IP address (e.g. 130.252.100.70), a similar rule could map requests to another host.
4. and 5) These rules allow the firewall to issue icmp requests to the Internet and to respond to icmp echo requests from the authorized subnet.

Rules are defined in the file `/etc/shorewall/rules` and are modified from the *Firewall Rules* menu.

Section 5.8.4

Configuring the Firewall and VPN

Section 5.8.4.1

Policy Based Virtual Private Networking

Begin configuration by creating local, network and vpn zones. Identify the network interface that carries the encrypted IPSec traffic and make this interface part of zone "ANY" in the interfaces menu as it will be carrying both traffic for both zones.

Visit the Zone Hosts menu and, for the network interface that carries the encrypted IPSec traffic, create a zone host with zone VPN, the correct subnet and the IPSec zone option checked. If you plan to have VPN tunnels to multiple remote sites ensure that a zone host entry exists for each (or collapse them into a single subnet). Create another zone host for the same interface with a network zone, using a wider subnet mask such as 0.0.0.0/0. It is important that the vpn zone be declared before the net zone since the more specific vpn zone subnet must be inspected first.

Host Zone	Interface	Subnet	IPSec Zone?
vpn	w1ppp	192.168.1.0/24	Yes
net	w1ppp	0.0.0.0/0	No

The IPSec protocol operates on UDP port 500 and using protocols ah (Authentication Header) and Encapsulating Security Payload (ESP) protocols. The firewall must accept this traffic in order to allow IPSec.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	net	fw	ah	
ACCEPT	net	fw	esp	
ACCEPT	net	fw	udp	500

IPSec traffic arriving at the firewall is directed to openswan, the IPSec daemon. Openswan then decrypts the traffic and forwards it back to shorewall on the same interface that originally received it. You will also need a rule to allow traffic to enter from this interface.

Action	Source Zone	Destination Zone
ACCEPT	vpn	loc

Section 5.8.4.2

Virtual Private Networking to a DMZ

If the firewall is to pass the VPN traffic through to another device (e.g. a VPN device in a DMZ) then establish a DMZ zone and install the following rules.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	net	dmz	ah	
ACCEPT	net	dmz	esp	
ACCEPT	net	dmz	udp	500
ACCEPT	dmz	net	ah	
ACCEPT	dmz	net	esp	
ACCEPT	dmz	net	udp	500

Section 5.8.5

Firewall Configuration

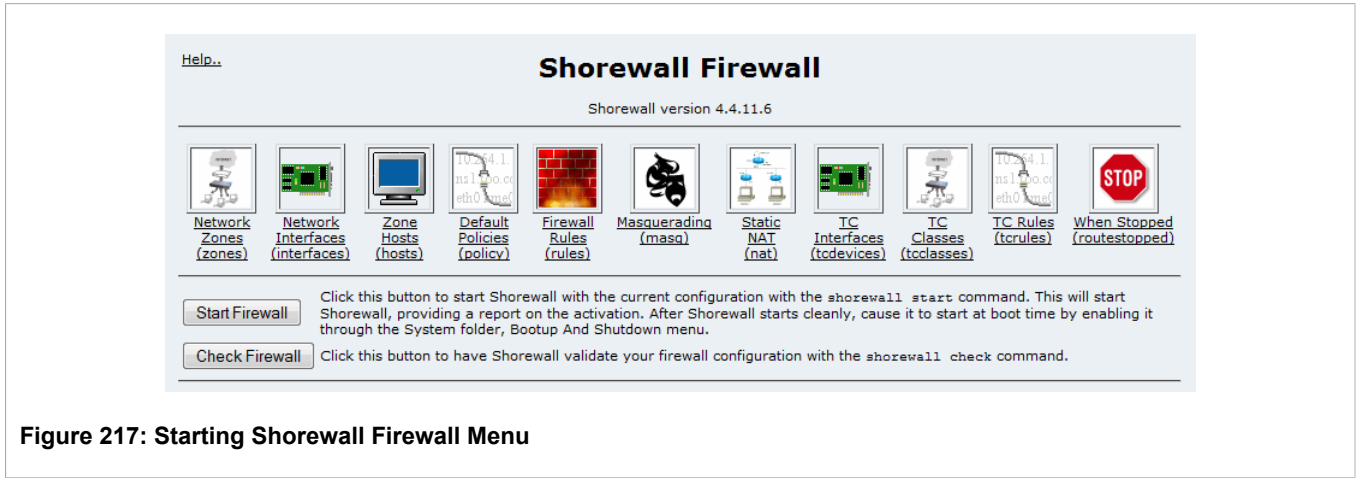


Figure 217: Starting Shorewall Firewall Menu

The above figure shows the firewall menu prior to configuration.

Configure the firewall through the provided menus. The "Check Firewall" button can be selected after each menu configuration to check the existing configuration and provide notice of items still to be configured.

When the firewall is fully configured, the "Start Firewall" button may be selected. Starting the firewall in this way will provide more detail (in the event of a problem). If the firewall starts cleanly, the menu appearance will change to that of the figure below.

In order to start the firewall at each and every boot, you must enable it via the System folder, Bootup and Shutdown menu.



Figure 218: Shorewall Firewall Menu

The "Apply Configuration" button must be used after making configuration changes. It is recommended that the "Check Firewall" button be used first to verify that any changes made are valid.

The "Refresh Configuration" button can be used to activate changes to the blacklisted host and traffic shaping configurations.

The "Clear Configuration" button will *remove the firewall rules completely and eliminate any protection they offer*. In some cases, you might wish to do this temporarily to determine if the firewall is responsible for an application problem.

The "Stop Firewall" button will stop the firewall. *Note that you should add an entry to the "When Stopped" menu to allow access from your management station while the firewall is stopped. If you do not do this, you lose web/ssh access and have to gain access via the console in order to restart the firewall.* Stopping the firewall will not disable it. Disable the firewall via the System folder, Bootup and Shutdown menu.

The "Show Status" button presents a variety of information summarizing the status of the firewall and routing system.

The "Check Firewall" button tests the current configuration to ensure it is valid.

Section 5.8.5.1

Network Zones

Help..

Network Zones

The zones listed on this page represent different networks reachable from your system, defined by name and type of zone.

[Add a new network zone.](#)

Zone ID	Zone type	Add
fw	Firewall system	

[Add a new network zone.](#)

[Manually Edit File](#) Click this button to manually edit the Shorewall file `/etc/shorewall/zones`, in which the entries above are stored.

[Return to list of tables](#)

Figure 219: Firewall Network Zones

This menu allows you to add, delete and configure zones. Add a new zone by selecting the "Add a new network zone" link or by clicking on the add-above or add-below images in the *Add* field.

The *Zone Type* field controls the type of traffic carried in the zone. The Firewall system zone type is built in to the fw zone. A zone type of IPSec is used with policy based VPNs. A zone type of IPV4 is used with normal traffic and route based VPNs.

Reorder the zones by clicking on the arrows under the *Move* field.

**NOTE**

If you define a VPN zone whose traffic is received via a network zone, it is essential that the VPN zone be declared before the network zone.

Clicking on a link under the *Zone ID* field will allow you to edit or delete the zone. Note that if you delete a zone you should remove any rules that reference it.

**NOTE**

There must be exactly one zone of type firewall. Do not delete this zone.

You may also make changes by manually editing the zone file.

Section 5.8.5.2

Network Interfaces

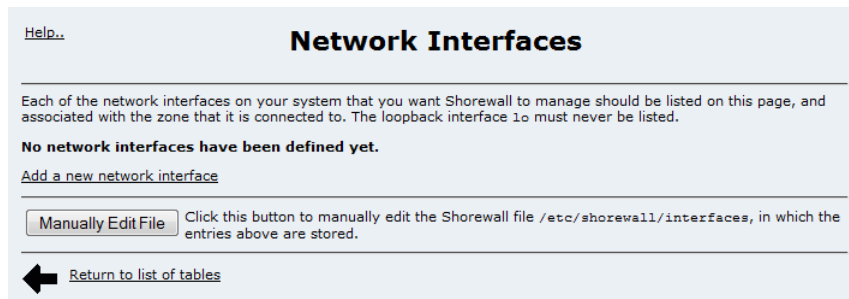


Figure 220: Firewall Network Interfaces

This menu allows you to add, delete and configure network interfaces. Add a new interface by selecting the "Add a new network interface" link or by clicking on the add-above or add-below images in the *Add* field. Reorder the interfaces by clicking on the arrows under the *Move* field.

Clicking on a link under the *Interface* field will allow you to edit or delete the interface. Note that if you delete an interface you should remove any rules that reference it.

You may also make changes by manually editing the interfaces file.



NOTE

If you use a WAN interface in the firewall, the interface will be referred to by its name. Some WAN changes (such as changing the number of channels used by a T1/E1 logical interface) will change the name. Ensure that the entries in this menu reflect the correct interface names.

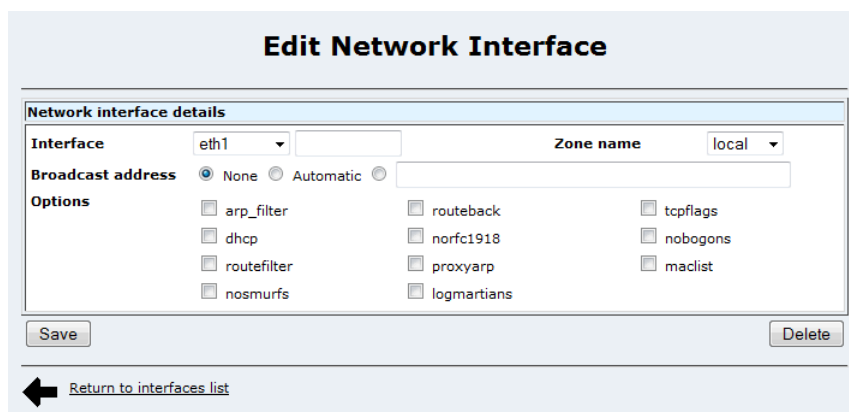


Figure 221: Editing Network Interface's Firewall Settings

The *dhcp* option should be selected if interface is assigned an IP address via DHCP or is used by a DHCP server running on the firewall. The firewall will be configured to allow DHCP traffic to and from the interface even when the firewall is stopped. You may also wish to use this option if you have a static IP but you are on a LAN segment that has a lot of laptops that use DHCP and you select the *norfc1918* option (see below).

The *arp_filter* option causes this interface to only answer ARP "who-has" requests from hosts that are routed out of that interface. Setting this option facilitates testing of your firewall where multiple firewall interfaces are

connected to the same HUB/Switch (all interfaces connected to the single HUB/Switch should have this option specified). Note that using such a configuration is strongly recommended against.

The *routeback* option causes Shorewall to set up handling for routing packets that arrive on this interface back out the same interface.

The *tcpflags* option causes Shorewall to make sanity checks on the header flags in TCP packets arriving on this interface. Checks include Null flags, SYN+FIN, SYN+RST and FIN+URG+PSH; these flag combinations are typically used for "silent" port scans. Packets failing these checks are logged according to the TCP_FLAGS_LOG_LEVEL option in /etc/shorewall/shorewall.conf and are disposed of according to the TCP_FLAGS_DISPOSITION option.

The *norfc1918* option causes packets arriving on this interface and that have a source or destination address that is reserved in RFC 1918 to be dropped after being optionally logged.

The *nobogons* option causes packets arriving on this interface that have a source address reserved by the IANA or by other RFCs (other than 1918) to be dropped after being optionally logged.

The *routefilter* option invokes the Kernel's route filtering (anti-spoofing) facility on this interface. The kernel will reject any packets incoming on this interface that have a source address that would be routed outbound through another interface on the firewall.

**NOTE**

The routefilter option should not be enabled on interfaces that are part of a multipath routing configuration.

The *proxyarp* option causes Shorewall to set proxy arp for the interface. Do not set this option if implementing Proxy ARP through entries in /etc/shorewall/proxarp.

The *maclist* option causes all connection requests received on this interface to be subject to MAC address verification. May only be specified for Ethernet interfaces.

The *nosmurf*s option causes incoming connection requests to be checked to ensure that they do not have a broadcast or multicast address as their source. Any such packets will be dropped after being optionally logged according to the setting of SMURF_LOG_LEVEL in /etc/shorewall/shorewall.conf.

The *logmartians* option causes the martian logging facility will be enabled on this interface. See also the LOG_MARTIANS option in /etc/shorewall/shorewall.conf.

Section 5.8.5.3

Network Zone Hosts

Create Zone Host

Zone host details

Zone: fw

Interface: eth1

IP address or network:

Host options: ☐ IPsec zone

Create

[Return to zone hosts list](#)

Figure 222: Firewall Zone Hosts

This menu allows you to add, delete and configure interfaces hosting multiple zones. Add a new zone host by selecting the "Add a new zone host" link or by clicking on the add-above or add-below images in the *Add* field. Reorder the hosts by clicking on the arrows under the *Move* field.

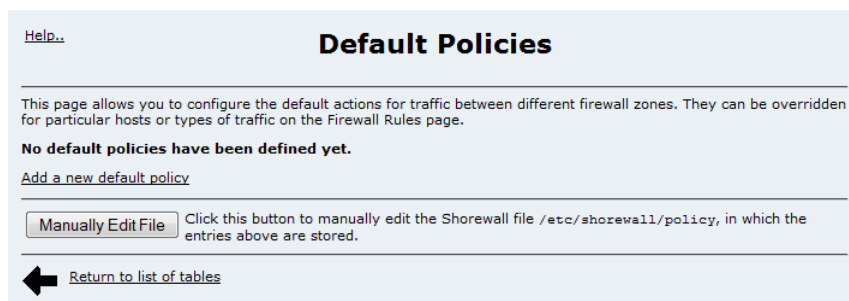
The *Zone* field selects a zone that will correspond to a subnet on the interface in question. The *Interface* field describes that interface and the *IP address or network* field describes the subnet.

Selecting the *IPSec zone Host Option* field will identify that the traffic to host in this zone is encrypted.

The *Save* and *Delete* buttons will allow you to edit or delete the zone host. You may also make changes by manually editing the policy

Section 5.8.5.4

Default Policies

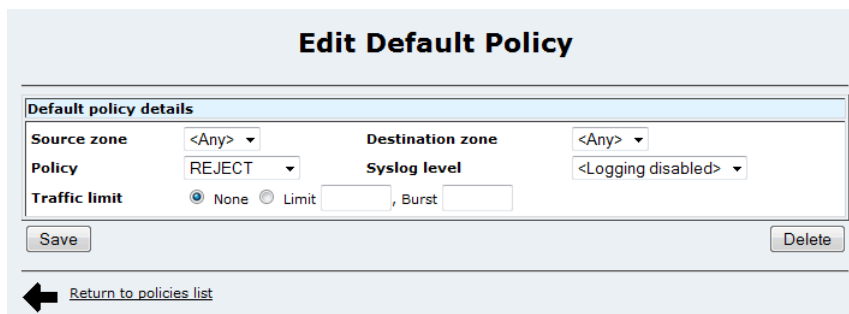


The screenshot shows the "Default Policies" configuration page. At the top, there is a "Help.." link and the title "Default Policies". Below the title, a paragraph explains that this page allows configuring default actions for traffic between different firewall zones, which can be overridden on the Firewall Rules page. It states "No default policies have been defined yet." and provides a link "Add a new default policy". A "Manually Edit File" button is present, with a tooltip that says "Click this button to manually edit the Shorewall file /etc/shorewall/policy, in which the entries above are stored." At the bottom, there is a back arrow and a link "Return to list of tables".

Figure 223: Firewall Default Policies

This menu allows you to add, delete and configure default policies. Add a new policy by selecting the "Add a new default policy" link or by clicking on the add-above or add-below images in the *Add* field. Reorder the policies by clicking on the arrows under the *Move* field.

Clicking on a link under the *Source zone* field will allow you to edit or delete the policy, as shown below. You may also make changes by manually editing the policy file.



The screenshot shows the "Edit Default Policy" configuration page. The title is "Edit Default Policy". Below it is a form titled "Default policy details". The form contains several fields: "Source zone" with a dropdown menu showing "<Any>", "Destination zone" with a dropdown menu showing "<Any>", "Policy" with a dropdown menu showing "REJECT", and "Syslog level" with a dropdown menu showing "<Logging disabled>". There is also a "Traffic limit" section with radio buttons for "None" (selected) and "Limit", followed by input fields for "Limit" and "Burst". At the bottom of the form are "Save" and "Delete" buttons. Below the form is a back arrow and a link "Return to policies list".

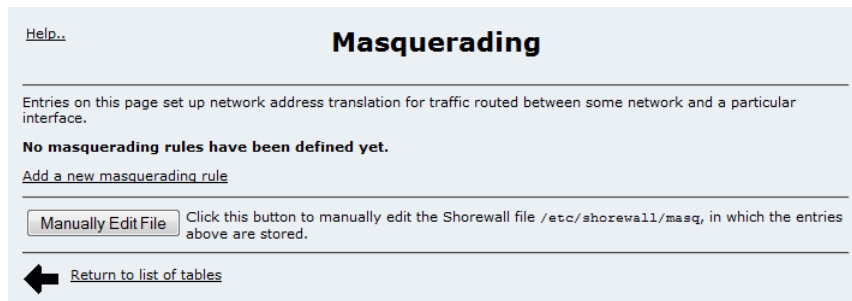
Figure 224: Editing a Firewall Default Policy

The *Syslog level* field causes a log entry to be generated every time the rule is followed.

The *Traffic limit* fields allow you to place an upper limit upon the rate at which the rule is applied. The *Limit* field is the steady state rate and is of the form "X/sec" or "X/min" where X is the number of allowed rule followings. The *Burst* field denotes the largest permissible burst and defaults to five if not configured.

Section 5.8.5.5

Masquerading

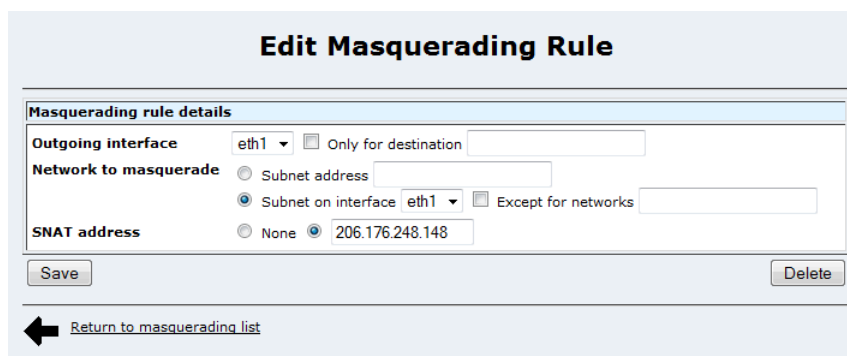


The screenshot shows a web interface titled "Masquerading". At the top left is a "Help.." link. Below the title, a paragraph states: "Entries on this page set up network address translation for traffic routed between some network and a particular interface." This is followed by the message "No masquerading rules have been defined yet." and a link "Add a new masquerading rule". A "Manually Edit File" button is present with a tooltip that reads: "Click this button to manually edit the Shorewall file /etc/shorewall/masq, in which the entries above are stored." At the bottom, there is a left-pointing arrow and a link "Return to list of tables".

Figure 225: Firewall Masquerading and SNAT

This menu allows you to add, delete and configure masquerading and SNAT rules. Add a new rule by selecting the "Add a new masquerading rule" link or by clicking on the add-above or add-below images in the *Add* field. Reorder the policies by clicking on the arrows under the *Move* field.

Clicking on a link under the *Outgoing interface* field will allow you to edit or delete the rule, as shown below. You may also make changes by manually editing the rule file.



The screenshot shows a web interface titled "Edit Masquerading Rule". It contains a form with the following fields and options:

- Masquerading rule details** (Section Header)
- Outgoing interface**: A dropdown menu showing "eth1" and a checkbox labeled "Only for destination" with an empty text field next to it.
- Network to masquerade**: Two radio buttons. The first is "Subnet address" with an empty text field. The second is "Subnet on interface" (selected) with a dropdown menu showing "eth1" and a checkbox labeled "Except for networks" with an empty text field next to it.
- SNAT address**: Two radio buttons. The first is "None". The second is selected and has a text field containing "206.176.248.148".
- At the bottom of the form are "Save" and "Delete" buttons.
- Below the form is a left-pointing arrow and a link "Return to masquerading list".

Figure 226: Editing a Masquerading Rule

The *Only for destination* field restricts the masquerading to the specified IP address.

The *Network to masquerade* fields determine the interface or subnet on the private network that you wish to masquerade. The *Except for networks* field restricts traffic from the specified subnet.

The *SNAT address* field is used to determine whether masquerading or SNAT is being performed. If checked, the entered IP address is used as a SNAT address.

Section 5.8.5.6

Firewall Rules

[Help..](#)

Firewall Rules

This table lists exceptions to the default policies for certain types of traffic, sources or destinations. The chosen action will be applied to packets matching the chosen criteria instead of the default.

[Add a new firewall rule](#)

Action	Source	Destination	Protocol	Source ports	Destination ports	Move	Add
ACCEPT	Any	Host 206.30.180.94 in zone DMZ	Any			↓	↓
DNAT	Host 66.11.180.161 in zone Internet	Any	TCP	Any	ssh	↑ ↓	↓
ACCEPT	Any	Zone DMZ	TCP	Any	ssh	↑	↓

[Add a new firewall rule](#)

[Manually Edit File](#) Click this button to manually edit the Shorewall file /etc/shorewall/rules, in which the entries above are stored.

[Return to list of tables](#)

Figure 227: Firewall Rules

This menu allows you to add, delete and configure firewall rules. These rules are inspected and applied before the default policies are used. Add a new rule by selecting the "Add a new firewall rule" link or by clicking on the add-above or add-below images in the *Add* field. Reorder the policies by clicking on the arrows under the *Move* field.

Clicking on a link under the *Action* field will allow you to edit or delete the rule, as shown below. You may also make changes by manually editing the rule file.

Edit Firewall Rule

Firewall rule details

Action: and log to syslog level:

Source zone: ☐ Only hosts in zone with addresses:

Destination zone or port: For DNAT or REDIRECT rules, this is the new destination address and/or port, normally it is the destination for the rule. ☐ Only hosts in zone with addresses:

Protocol:

Source ports: ☒ Any ☐ Ports or ranges:

Destination ports: ☒ Any ☐ Ports or ranges: For DNAT or REDIRECT rules, fill in the original destination port here.

Original destination address for DNAT or REDIRECT: ☒ None ☐

Rate limit expression: ☒ No limit ☐

Rule applies to user set: ☒ All users ☐

[Save](#) [Delete](#)

[Return to firewall rules list](#)

Figure 228: Editing a Firewall Rule

The following fields describe the information to match against an incoming connection request in order to apply this rule.

The *Action* field specifies the final action to take on incoming requests matching the rule. The *and log to syslog* field determines whether logging will take place and at which logging level.

The *Source zone* field specifies the zone from which the request originates.

**NOTE**

*When defining the destination zone for a DNAT rule, select the **Only hosts in zone with addresses** and type the IP address for the new destination in the field next to it.*

*When defining the destination zone for a REDIRECT rule, select **"Other..."** and enter the port number for the new destination.*

The *Destination zone or port* field specifies the request's destination.

Each of the Source and Destination zones may use one of the defined zone names, or one may select "Other..." and specify a zone name in the text field to the right. Both Source and Destination may be further qualified using the *Only hosts in zone with addresses* fields. Multiple comma-separated subnet, IP, or MAC addresses may be specified in the following way:

- Subnet: 192.168.1.0/24
- IP: 192.168.1.1
- IP range: 192.168.1.1-192.168.1.25
- MAC: ~00-A0-C9-15-39-78

The *Protocol* field specifies the protocol (tcp, udp or icmp) to match.

The *Source ports* and *Destination ports* fields specify TCP or UDP port numbers to match. These fields are in the form of a list of comma-separated port numbers or ranges of port numbers of the form, "first:last".

The *Original destination address* field matches the request's destination IP address.

**NOTE**

*If you use are using DNAT to implement port forwarding, enter the original destination address here and the forwarded address in the *Destination zone or port* fields **Only hosts in zone with address sub-field**.*

The *Rate limit expression* fields specify rate limit control of the form "X/sec" or "X/min" where X is the number of allowed requests in the time period. A burst limit field ":Y" where Y is the maximum consecutive number of requests and defaults to five if not configured.

The *Rule applies to user set fields* allow advanced users to match the rule against specific users and groups. This matching only takes place when the source of the traffic is the firewall itself.

Section 5.8.5.7

Static NAT

[Help...](#)

Static NAT

The static network address translation entries in this table can be used to set up a 1-1 correspondence between an external address on your firewall and an RFC1918 address of a machine behind your firewall. Static NAT is often used to allow connections to an internal server from outside your network.

[Add a new static NAT entry](#)

External address	External interface	Internal address	Move	Add
204.226.111.45	eth1	192.168.0.1	↓	↑ ↓
204.62.138.24	eth1	10.0.0.1	↑	↑ ↓

[Add a new static NAT entry](#)

[Manually Edit File](#) Click this button to manually edit the Shorewall file /etc/shorewall/nat, in which the entries above are stored.

[Return to list of tables](#)

Figure 229: Static NAT

This menu allows you to add, delete and static NAT translations. Add a new translations by selecting the "Add a new static NAT entry" link or by clicking on the add-above or add-below images in the *Add* field. Reorder the translations by clicking on the arrows under the *Move* field.

Clicking on a link under the *External Address* field will allow you to edit or delete the rule, as shown below. You may also make changes by manually editing the rule file.

Edit Static NAT

Static NAT entry details

External address	204.226.111.45	External interface	eth1 virtual
Internal address	192.168.0.1	No IP alias	<input type="checkbox"/>
Active for all hosts?	<input checked="" type="radio"/> Yes <input type="radio"/> No	Active for firewall system?	<input type="radio"/> Yes <input checked="" type="radio"/> No

[Save](#) [Delete](#)

[Return to static NAT list](#)

Figure 230: Creating a Static NAT Entry

The *External address* and *Internal address* fields specify the addresses to translate.

The *External interface* field specifies the interface to perform the translation upon.

The *No IP alias* field is used to tell the firewall not create the IP alias for the external address if it has not been created on the external interface yet.

The *Active for all hosts* field is used to specify whether access to the external IP from all firewall interfaces should undergo NAT (Yes or yes) or if only access from the interface in the INTERFACE column should undergo NAT.

The *Active for firewall system* field is used to specify whether packets originating from the firewall itself and destined for the EXTERNAL address are redirected to the internal ADDRESS.

Section 5.8.5.8

TC (Traffic Control) Interfaces, Classes, and Rules

The Traffic Control subsystem of the firewall allows sophisticated management of the amount of bandwidth that different types of traffic are permitted to use on a given interface. Please see [Section 5.9, “Traffic Control”](#) for details.

Section 5.8.5.9

Actions When Stopped

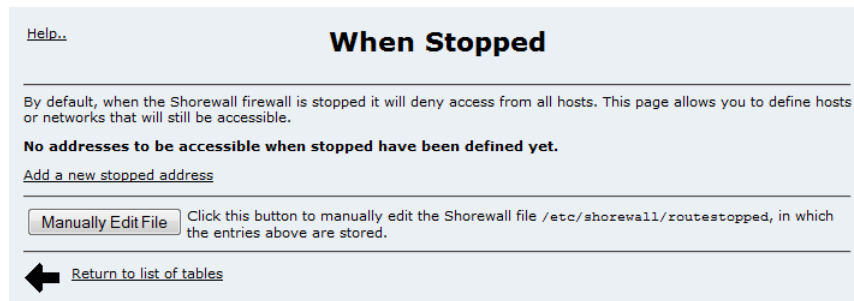


Figure 231: Actions When Stopped

This menu allows you to control which addresses the firewall will accept connections from after it has been stopped. Add a new translations by selecting the "Add a new stopped address" link or by clicking on the add-above or add-below images in the *Add* field. Reorder the translations by clicking on the arrows under the *Move* field.

Clicking on a link under the *Interface* field will allow you to edit or delete the rule, as shown below. You may also make changes by manually editing the rule file.

Section 5.8.5.10

Controlling the Firewall from the Command Line

The software provides limited control of the firewall from the command line, such as from SSH.

**IMPORTANT!**

The firewall cannot be started or stopped unless it is configured. Run `enable` first or see [Section 5.8.5, “Firewall Configuration”](#).

**CAUTION!**

Configuration hazard – risk of data corruption. Access to the Command Line Interface (CLI) is provided for troubleshooting purposes and should only be used by Siemens technicians. As such, this interface is not fully documented. Misuse of CLI commands can corrupt the operational state of the device and render it inaccessible.

The following commands run from the command line:

- To enable the firewall now or at bootup:

```
/etc/init.d/shorewall enable
```

- To disable the firewall now or at bootup:

```
/etc/init.d/shorewall disable
```

- To clear the firewall (clearing allows all connections):

```
/etc/init.d/shorewall clear
```

- To stop the firewall:

```
/etc/init.d/shorewall stop
```

- To start or restart the firewall:

```
/etc/init.d/shorewall start  
/etc/init.d/shorewall restart
```

Section 5.9

Traffic Control

Traffic Control is a subsystem of the firewall that allows management of the amount of bandwidth per network interface that different types of traffic are permitted to use.

Each interface to be managed is assigned a total bandwidth that it should allow for incoming and outgoing traffic. Classes are then defined for each interface, each with its own minimum assured bandwidth and a maximum permitted bandwidth. The combined minimum of the classes on an interface must be no more than the total outbound bandwidth specified for the interface. Each class is also assigned a priority, and any bandwidth left over after each class has received its minimum allocation (if needed) will be allocated to the lowest priority class up until it reaches its maximum bandwidth, after which the next priority is allocated more bandwidth. When the specified total bandwidth for the interface is reached, no further packets are sent, and any further packets may be dropped if the interface queues are full.

Packets are assigned to classes on the outbound interface based on either a mark assigned to the packet, or the ToS (type of service) field in the IP header. If the ToS field matches a defined class, then the packet is allocated to that class. Otherwise, it is allocated to any class that matches the mark assigned to the packet, and if no class matches the mark, then the packet is assigned to the default class.

Marks are assigned to packets either by the TC Rules based on any of a number of parameters, such as IP address, port number, protocol, packet length, and so on, or by mapping an 802.1p VLAN CoS value to a MARK in the VLAN configuration of the incoming port. Marks are also used to map back to an 802.1p CoS value on an outbound VLAN port.

Section 5.9.1

Traffic Control Example

The goal of this example is to operate Ethernet port 1 at 5Mbit/s and ensure that UDP source port 20000 traffic gets at least half the bandwidth, while ICMP and TCP ACK packets should have high priority, HTTP traffic gets at least 20% and at most 50%, and all other traffic should get what is left over but only up to 50% of the bandwidth.

The three TC menus would be configured as follows:

TC Interfaces

Interface	Inbound bandwidth	Outbound bandwidth
eth1	5000kbit	5000kbit

TC Classes

Interface	Mark	Minimum	Maximum	Priority	Options
eth1	1	full/2	full	0	
eth1	2	1kbit	full	1	tcp-ack
eth1	3	full/5	full*5/10	2	
eth1	4	1kbit	full*5/10	3	default

TC Rules

Mark	Source	Destination	Protocol	Source Port	Dest Port	Test	Length	TOS
2	Any	Any	ICMP	Any	Any	Any	Any	Any
RESTORE	Any	Any	Any	Any	Any	0	Any	Any
CONTINUE	Any	Any	Any	Any	Any	!0	Any	Any
1	Any	Any	UDP	20000	Any	Any	Any	Any
3	Any	Any	TCP	Any	80	Any	Any	Any
4	Any	Any	Any	Any	Any	0	Any	Any
SAVE	Any	Any	Any	Any	Any	!0	Any	Any

The rules first check non connection-based protocol rules (ICMP in this case) in order to assign a mark. For any packet that is still not marked, we attempt to restore a saved mark for the connection. If at this point the packet has a mark set, we stop checking rules (CONTINUE) since it is either ICMP or a packet from an existing connection which we have already assigned a mark. If still no mark is assigned, it must be a new connection so we process the packet through all the remaining rules to determine the mark it should receive. At the end we save the new mark to the connection so that any further packets for the connection do not have to go through all the rules again, in order to save processing resources. We mark all packets with no other matching rule to 4 since that represents the default class (as defined in TC Classes). This allows explicit traffic control of even unspecified network connections.

Section 5.9.2

Traffic Control Configuration



NOTE

Traffic Control is mutually exclusive of Traffic Prioritization. Do not enable both of these features at once.

Section 5.9.2.1

TC Interfaces (tcdevices)

[Help..](#)

TC Interfaces

Each of the traffic classification interfaces on your system that you want Shorewall to manage should be listed on this page, and with its maximum inbound and outbound bandwidths the link is capable of handling. The loopback interface `lo` must never be listed.

[Add a new traffic classification interface](#)

Interface	In bandwidth	Out bandwidth	Add
eth1	2000kbit	5000kbit	↑ ↓

[Add a new traffic classification interface](#)

[Manually Edit File](#) Click this button to manually edit the Shorewall file `/etc/shorewall/tcdevices`, in which the entries above are stored.

[Return to list of tables](#)

Figure 232: TC Interfaces

This Menu allows you to add, edit or remove traffic classification interfaces, and to assign the maximum inbound and outbound bandwidths that the interface can handle. Add a new traffic classification interface by selecting the *Add a new traffic classification interface* link or by clicking on the add-above or add-below images in the *Add* column.

Clicking on a link in the *Interface* column will allow you to edit or delete a traffic classification interface as shown below.

Edit TC Interface

TC interface details

Interface:

In bandwidth:

Out bandwidth:

[Save](#) [Delete](#)

[Return to tcdevices list](#)

Figure 233: Edit TC Interface

The *Interface* field specifies a network interface whose traffic will be controlled by the TC subsystem.

The *In bandwidth* field specifies the maximum inbound bandwidth that the interface can handle. If the rate exceeds this value, packets may be delayed or potentially dropped.

The *Out bandwidth* field specifies the maximum outbound bandwidth that the interface can handle. Outbound traffic above this rate is delayed or potentially dropped.

Bandwidth is specified in either kilobytes per second (kbps), or kilobits per second (kbit).



NOTE

The minimum bandwidth that may be specified to the Traffic Control subsystem for any network device is 10 kilobits per second (kbit).

The *Manually Edit File* button also allows you to make direct changes to the TC interface configuration file.

TC Interfaces

This form can be used to manually edit the Shorewall file `/etc/shorewall/tcdevices`. Be careful, as no syntax checking will be done on the changes.

```
#
# Shorewall version 4 - Tcdevices File
#
# For information about entries in this file, type "man shorewall-tcdevices"
#
# See http://shorewall.net/traffic_shaping.htm for additional information.
#
#####
#NUMBER:      IN-BANDWIDTH      OUT-BANDWIDTH  OPTIONS      REDIRECTED
#INTERFACE
eth1    2000kbit      5000kbit
#LAST LINE -- ADD YOUR ENTRIES BEFORE THIS ONE -- DO NOT REMOVE
```

Save

Undo

Return to tcdevices list

Figure 234: Edit TC Interface

Section 5.9.2.2
TC Classes

Help...

TC Classes

Classes define the bandwidth available to specific types of traffic. Exactly one class per interface **must** be flagged as default. At most one class per interface can have the `tcp-ack` flag set. Each `tos` byte match is limited to one class per interface and `tos` byte matching overrides Mark matching. A packet can be matched either by Mark or by `tos` byte.

Add a new traffic classification class

Interface	Mark to match	Minimum rate	Maximum rate	Priority	Options	Move	Add
eth1	1	full/2	full*8/10	1	tos-minimize-delay,tcp-ack	↓	↑ ↓
eth1	2	200kbit	200kbit	1	default	↑	↑ ↓

Add a new traffic classification class

Manually Edit File

Click this button to manually edit the Shorewall file `/etc/shorewall/tcclasses`, in which the entries above are stored.

Return to list of tables

Figure 235: TC Classes

This menu allows you to add, edit, or remove a traffic classification class. Please note that each class is associated with exactly one network interface. Exactly one class for each interface must be designated as default. Unmarked traffic (packets which have not been assigned a mark value in the *TC Rules* menu or via VLAN 802.1p) will be handled by the default class. Classes can match packets either by their assigned mark or by their ToS field. The ToS field takes precedence over the mark.

Add a new traffic class by selecting the *Add a new traffic classification class* link or by clicking on the add-above or add-below images in the *Add* column.

Clicking on a link in the *Interface* column will allow you to edit or delete a traffic classification class as shown below.

Edit TC Classes

TC class details

Interface: eth1

Mark to match: 1

Minimum rate: ☐ kbit ☒ full * 1 / 2

Maximum rate: ☐ kbit ☒ full * 8 / 10

Priority: 1

Options:

- ☐ default
- ☐ tos-maximize-reliability
- ☒ tcp-ack
- ☒ tos-minimize-delay
- ☐ tos-minimize-cost
- ☐ tos-maximize-throughput
- ☐ tos-normal-service

tos= (0xXX or 0xXX/0xYY where XX is two digit hex tos value and YY is two digit hex mask)

[Return to tcclases list](#)

Figure 236: Edit TC Classes

The *Interface* field specifies which network interface this TC Class applies to.

The *Mark to match* field specifies what mark value this TC Class will match. The mark may be in the range 1..255 in decimal or hex.

The *Minimum rate* field specifies the minimum bandwidth allocated to this class.

The *Maximum rate* field specifies the maximum bandwidth allocated to this class.

Bandwidth is specified in megabytes per second (mbps), megabits per second (mbit), kilobytes per second (kbps), kilobits per second (kbit), or bytes per second (bps). Alternately it can be specified as a fraction of the full port speed defined in the *TC Interfaces* menu.

The *Priority* field specifies the priority with which this class is serviced. Please note that lower value priority classes will be serviced first (and hence with lower latency). The lower priority classes are also the first to be allocated any leftover bandwidth after all classes have been provided with their minimum bandwidth. Priority may be in the range 0..255. However, a priority higher than 7 will be mapped automatically to 7.

The *default* option field sets the current class as the default class for the interface. Please note that you must define exactly one default class per interface.

The *tos-minimize-delay* option field specifies that a packet with the minimized delay ToS (ToS bit 3 set) belongs to this class.

The *tos-maximize-throughput* option field specifies that a packet with the maximize throughput ToS (ToS bit set) belongs to this class.

The *tos-maximize-reliability* option field specifies that a packet with the maximize reliability ToS (ToS bit 5 set) belongs to this class.

The *tos-minimize-cost* option field specifies that a packet with the minimize cost ToS (ToS bit 6 set) belongs to this class.

The *tcp-ack* option field specifies that a tcp ack packet with size ≤ 64 belongs to this class. This is useful for speeding up bulk downloads. Please note that the size of the ack packet is limited to 64 bytes because this option is intended to only match packets with no payload. This option is only valid for one class per interface.

The `tos=` option field allows you to define a classifier for the given value/mask combination of an IP packet's TOS byte. Value and mask are both specified in hexadecimal notation using the "0x" prefix. It is also possible to specify a diffserv marking, or DSCP (Diffserv Code Point). These are typically quoted as 6-bit values, and must be left-shifted (multiplied by 4) for use in the "tos=" field. For example, a DSCP of 0x2E (EF, or Expedited Forwarding) would be entered as 0xB8/0xFC (4 X 0x2E = 0xB8, and the two lowest order bits are masked by masking with 0xFC).

**NOTE**

ToS field matches take precedence over the assigned mark.

The **Save** button saves the class changes to the TC configuration.

The **Delete** button delete the class from the TC configuration.

The **Manually Edit File** button also allows you to make direct changes to the TC Classes configuration file.

TC Classes

This form can be used to manually edit the Shorewall file `/etc/shorewall/tcclasses`. Be careful, as no syntax checking will be done on the changes.

```
# Shorewall version 4 - Tcclasses File
#
# For information about entries in this file, type "man shorewall-tcclasses"
#
# See http://shorewall.net/traffic_shaping.htm for additional information.
#
#####
#INTERFACE:CLASS      MARK      RATE:      CEIL      PRIORITY      OPTIONS
#
#          DMAX:UMAX
eth1      1          full/2    full*8/10    1          tos-minimize-delay,tcp-ack
eth2      2          200kbit  200kbit  2          default
#LAST LINE -- ADD YOUR ENTRIES BEFORE THIS ONE -- DO NOT REMOVE
```

SaveUndo

[Return to tcclasses list](#)

Figure 237: Manually Edit TC Classes

Section 5.9.2.3

TC Rules

This menu allows you to add, edit or remove a traffic classification rule. Add a new rule by selecting the *Add a new traffic classification rule* link or by clicking on the add-above or add-below images in the *Add* column. Reorder rules by clicking on the arrows in the *Move* column.

[Help..](#)

TC Rules

Rules define which mark to assign to a packet and/or connection based on various criteria. Every rule is processed in order for each packet, and every matching rule is applied to the packet (or its connection), unless a CONTINUE rule is matched in which case no further rules are processed for the packet.

[Add a new traffic classification rule](#)

Mark	Source	Destination	Protocol	Source Ports	Destination Ports	Test	Length	TOS	Move	Add
1	Any	Any	ICMP	Any	echo-request	Any	Any	Any	↓	T D
1	Any	Any	ICMP	Any	echo-reply	Any	Any	Any	↑ ↓	T D
4	Any	Any	Any			Any	Any	Minimize-Delay	↑ ↓	T D
3	Any	Any	Any			Any	Any	Any	↑ ↓	T D
2	Any	Any	UDP	54	53	!1/C	Any	Any	↑	T D

[Add a new traffic classification rule](#)

[Manually Edit File](#) Click this button to manually edit the Shorewall file /etc/shorewall/tcrules, in which the entries above are stored.

[Return to list of tables](#)

Figure 238: TC Rules

Clicking on a link in the *Mark* column will allow you to edit or delete a traffic classification rule, as shown below.

Edit TC Rule

TC rule details

Mark

☒ Set Packet mark to 1 / in Default chain

☐ Modify packet mark by AND in Default chain

☐ SAVE connection mark with mask in Default chain

☐ CONTINUE in Default chain

Source ☒ Any

Destination ☒ Any

Protocol ICMP

Source Ports ☒ Any ☐ Ports or ranges

Destination Ports ☐ Any ☒ Ports or ranges echo-request

Test Packet mark =

Length ☒ Any ☐ Range

TOS <Any>

[Save](#) [Delete](#)

[Return to tcrules list](#)

Figure 239: Edit TC Rule

The *Mark* row determines how the mark value will be assigned for a packet or a connection:

- The *Set* field determines whether the packet or the connection is assigned the mark. The *mark to* field specifies the mark value for the rule and the */* field specifies the mask for the mark value (if the */* field is empty, the mark value will be the value set in the *mark to* field). The *in* field specifies the chain in which the rule will be processed. Mark and mask may be in the range 1..255 in decimal or hex.

- The *Modify packet mark by* field allows you to change the mark value by an AND or OR value. The *in* field specifies the chain in which the rule will be processed. The value may be in the range 1..255 in decimal or hex.
- The *SAVE/RESTORE connection mark with mask* field allows you to save or restore the connection mark value with an assigned mask value. The *in* field specifies the chain in which the rule will be processed. The mask value may be in the range 1..255 in decimal or hex.
- The *CONTINUE in* field specifies that no more TC rules be checked if the packet matches, and to forward the packet to the specified chain.

The *Source* field specifies the source IP address, subnet or MAC addresses to match. Please refer to [Section 5.8, "Configuring the Firewall"](#) for the formats of MAC and IP addresses supported by Shorewall.

The *Destination* field specifies the destination IP address, subnet or MAC addresses to match.

The *Protocol* field specifies the protocol (UDP, TCP, ICMP, etc) to match.

The *Source Ports* field specifies the source TCP or UDP port number or ICMP type to match.

The *Destination Ports* field specifies the destination TCP or UDP port number or ICMP type to match.

The *Test* field defines a test on the existing packet or connection mark. The packet or the connection mark may be checked for equality or non-equality against a reference mark. A mask may again be specified in the */field*, to apply to both marks prior to comparison. Mark and mask may be in the range 1..255 in decimal or hex.

The *Length* field specifies the packet length or length range to match.

The TOS field specifies the packet TOS value to match. A TOS value may be selected from the list, or may be specified in decimal or hex by selecting "Other". It may take on one of the following values:

- Minimize-Delay (16/0x10)
- Maximize-Throughput (8/0x08)
- Maximize-Reliability (4/0x04)
- Minimize-Cost (2/0x02)
- Normal-Service (0/0x00)

The *Manually Edit File* button also allows you to make direct changes to the TC Rules configuration file.

TC Rules

This form can be used to manually edit the Shorewall file /etc/shorewall/tcrules. Be careful, as no syntax checking will be done on the changes.

```
# Shorewall version 4 - Tcrules File
#
# For information about entries in this file, type "man shorewall-tcrules"
#
# See http://shorewall.net/traffic_shaping.htm for additional information.
# For usage in selecting among multiple ISPs, see
# http://shorewall.net/MultiISP.html
#
# See http://shorewall.net/PacketMarking.html for a detailed description of
# the Netfilter/Shorewall packet marking mechanism.
#####
#MARK  SOURCE  DEST  PROTO  DEST  SOURCE  USER  TEST
LENGTH TOS  CONNBYTES  HELPER
#
1      -      -      icmp   echo-request  PORT(S)  PORT(S)
1      -      -      icmp   echo-reply    -      -      -
4      -      -      all    -           -      -      -
Minimize-Delay
3      -      -      all    -           44
```

Save Undo

← Return to tcrules list

Figure 240: Manually Edit TC Rules

Hints on optimizing the TC Rule table

Every rule is processed in table order for every packet, unless a CONTINUE rule is matched, in which case processing stops. This can be used to improve efficiency in combination with the SAVE and RESTORE rules. For example, consider a TC Rules table organized roughly as follows (and in the same order):

- A RESTORE rule is used to restore the connection's mark to a matching, unmarked packet
- A CONTINUE if the mark is non zero
- Specific rules to check criteria to assign a mark, and finally
- A SAVE mark to connection if the mark is non zero (ie a match was found above)

Using the above structure for the TC Rules table, only the first packet of any tcp or udp connection will have to go through all the rules, while every following packet will have its mark restored by the first rule, and then CONTINUE, skipping potentially many matching rules in the remainder of the table.

Section 5.10

Traffic Prioritization

This section familiarizes the user with:

- Enabling/Disabling Traffic Prioritization
- Viewing Traffic Prioritization Statistics



NOTE

Traffic Prioritization has been retained in ROX for compatibility with older installations that may rely on it. For new configurations, please use the newer and more flexible [Section 5.9, "Traffic Control"](#) facility instead.

ROX is able to prioritize traffic transmitted on network interfaces (including Ethernet, T1E1, DSL and PPP ports), giving preferential treatment to certain classes of traffic.

It is important to note that prioritization can only be applied to outbound traffic. Inbound traffic can not be prioritized.

The two key elements of prioritization are traffic queues and filters. Each prioritized interface has its own unique set of these elements.

**NOTE**

Traffic Prioritization works most effectively with WAN interfaces. For LAN interfaces please use the [Section 5.9, "Traffic Control" interface](#).

Section 5.10.1

Priority Queues

Prioritization establishes a number of queues, each holding packets of differing priority. When the interface is ready to transmit a packet it selects a packet from the highest priority queue first.

If the interface is busy transmitting when packets arrive, they are enqueued in the appropriate queue.

If the interface is not transmitting when the frame arrives to be enqueued, the frame is immediately transmitted. Prioritization will not add additional delay to a stream of packets of differing priority. Prioritization will simply reorder the sequence of transmission of packets to send higher priority packets first.

Note that it is possible to indefinitely stall the transmission of packets from a lower priority queue if a traffic from a higher queue saturates the interface.

**NOTE**

The router mandates that you must have at least a low, normal and high priority queue. Additionally, the high queue must be of higher priority than the normal queue, which must be of higher priority than the low queue.

Section 5.10.2

Filters

For each packet to be transmitted on a prioritized interface, the packet is compared against each of the filters on that interface until a match is found. The matching filter directs the packet onto a specific queue. If no matching filter is found the packets Type of Service (TOS) bits in its IP header are examined and used.

It is possible to match on source and destination IP address/mask pairs, source and destination port numbers and protocols.


The 0.0.0.0/0 address/mask matches any IP address.

Protocols that can be matched upon include TCP, UDP, ICMP, OSPF, VRRP and IPSec.

Section 5.10.3

TOS Prioritization

The priority of an IP packet can be derived from its Type of Service field. The TOS field has the following format:



The four TOS bits (the 'TOS field') are defined as:

- MD - Minimize Delay
- MT - Maximize Throughput
- MR - Maximize Reliability
- MMC - Minimize Monetary Cost

As any (or all) of these bits may be set in a packet at a time, there are 16 possible combinations. The router maps these combinations into the high, normal and low priority queues as shown in the following table:

MD	MT	MR	MMC	Descriptions	Priority Queue
0	0	0	0	Normal Service	Normal
0	0	0	1	Minimize Monetary Cost	Low
0	0	1	0	Maximize Reliability	Normal
0	0	1	1	MR+MMC	Normal
0	1	0	0	Maximize Throughput	Low
0	1	0	1	MT+MMC	Low
0	1	1	0	MT+MR	Low
0	1	1	1	MT+MR+MMC	Low
1	0	0	0	Minimize Delay	High
1	0	0	1	MD+MMC	High
1	0	1	0	MD+MR	High
1	0	1	1	MD+MR+MMC	High
1	1	0	0	MD+MT	Normal
1	1	0	1	MD+MT+MMC	Normal
1	1	1	0	MD+MT+MR	Normal
1	1	1	1	MD+MT+MR+MMC	Normal

Section 5.10.4

Prioritization Example

A remote site router connects to a private network via a T1 line. The router uses OSPF to manage an alternate routing, but its primary purpose is to allow access to a switched network of RuggedServers implementing TcpModbus gateways (TCP/UDP port 502). The router and switches are managed through their Web interfaces, but can be managed through SSH as well. The RuggedServers are managed through Telnet. An SNMP network management polling application tracks the status of all devices.

It is generally wise to ensure that control and management capabilities are always provided. OSPF and SSH/Telnet should be assigned to the highest priority queue. OSPF packets are small and do not consume much bandwidth. SSH and Telnet are not often used but must be available when required.

TcpModbus traffic is ensured a low latency by assigning it the next lowest queue.

Web traffic will be used to manage the router and switches and should be assigned to a still lower queue.

All other traffic can be assigned to a final queue.

In all, four queues are required. The system provides three basic queues ("high", "normal" and "low") and a fourth, the "extra high" can be manually added.

Traffic filters are inspected in the order in which they are entered. To reduce load and improve performance the filters should be entered in an order which recognizes the most frequent traffic (under normal conditions). The best filter order is probably:

- match source port 502 > queue "high"
- match protocol OSPF > queue "extra high"
- match source port "snmp" > queue "extra high"
- match source port "www" > queue "normal"
- match source port "10000" > queue "normal"
- match source port "ssh" > queue "extra high"
- match source port "telnet" > queue "extra high"
- match source IP/Mask 0.0.0.0/0 > queue "low"

Note that the snmp, www, ssh and telnet keywords are defined in the file /etc/services, so we can use their mnemonics here. We could also have used the raw port numbers 161, 80, 22 and 23, respectively. The TcpModbus port number is not common, and must be explicitly entered. The Webmin port number of 10000 reflects the fact that web traffic from a router is issued on this port.

Each of the "port based" filters must match a source port. Matching is being applied to packets from the service at the well known source port to an unknown and variable destination port number.

Finally, note that the final traffic filter essentially suppresses TOS inspection by directing all unmatched traffic onto the "low" queue.

Section 5.10.5

Configuring Traffic Prioritization



NOTE

Traffic Prioritization is mutually exclusive of Traffic Control. Do not enable both of these features at once.

[Help..](#)

Traffic Prioritization

Please note that Traffic Control should be used instead

Interfaces

Interface	Prioritized?	Queues	Filters	Statistics
eth1	No	-	-	-
eth2	No	-	-	-
eth3	No	-	-	-
eth4	No	-	-	-
w2c1ppp	No	-	-	-
w2c2ppp	No	-	-	-
w2c3ppp	No	-	-	-

Figure 241: Traffic Prioritization Main Menu

This menu displays network interfaces for which prioritization may be activated. Prioritization may be configured by following the *Interface* column link. The statistics of prioritized interfaces may be viewed by following the links in the *Statistics* column.

Section 5.10.6

Interface Prioritization Menu

[Help..](#)

w2c1ppp Prioritization

Prioritization Queues

Note that you must have at least a low, normal and high priority queue. The high queue must be of higher priority than the normal queue, which must be of higher priority than the low queue. If you delete a priority queue, any filters which use that queue will be adjusted to point at the next lowest queue.

Queue Name	Move	Add
high		↑ ↓
normal		↑ ↓
low		↑ ↓

Prioritization Filters

Packets are matched against filters from the following table, in ascending order. When a match occurs the packet is entered onto the respective target queue. If no matches occur the packet's TOS bits are inspected and the packet is entered onto the low, normal or high queue.

Source IP/Netmask	Source Port	Dest IP/Netmask	Dest Port	Protocol	Target Queue	Move	Add
Add a traffic filter							

Transmit Queue Length

Packets from the above prioritization queues are collected on to a transmit queue prior to transmission. Limiting the size of this queue increases performance by preventing the buffering of a number of lower priority frames.

Length
Edit 1

[Return to Traffic Prioritization](#)

Figure 242: Interface Prioritization Menu

This menu allows you to add, delete and configure queues and filters. Add a new queue or filter by clicking on the add-above or add-below arrows in the *Add* field. You may also edit a manually created queue by following its link under the *Queue Name* column, and edit a filter by following its "Edit" link.

Reorder the queues and filters by clicking on the arrows in the *Move* field. Some restrictions apply with queues. You are not allowed to reorder queues in a way that violates the priority implicit in their name.

The Transmit Queue Length Selector allows you to make a tradeoff between latency and performance.

Remove prioritization by selecting the *Delete and Apply* button.

Section 5.10.6.1

Prioritization Queues

The screenshot shows the 'w2c1ppp Queue Configuration' window. At the top left is a 'Help...' link. The title is 'w2c1ppp Queue Configuration'. Below the title is a section 'Queue Configuration for new queue' containing a 'Queue Name' text input field. Below this is a 'Save and Apply' button. At the bottom left is a back arrow icon, and to its right is a link 'Return to Traffic Prioritization'.

Figure 243: Prioritization Queue Configuration

This menu allows you to edit the name of a priority queue and to delete the queue.

If you delete a queue referenced by filters, the filters will be adjusted to use the next lowest queue.

Section 5.10.6.2

Prioritization Filters

The screenshot shows the 'w2c1ppp Filter Configuration' window. At the top left is a 'Help...' link. The title is 'w2c1ppp Filter Configuration'. Below the title is a section 'Filter Configuration for new filter' containing several fields: 'Source IP/Netmask' and 'Source Port' (text input), 'Dest IP/Netmask' and 'Dest Port' (text input), 'Protocol' (dropdown menu), and 'Target Queue' (dropdown menu with 'extra high' selected). Below these fields is a 'Save and Apply' button. At the bottom left is a back arrow icon, and to its right is a link 'Return to Traffic Prioritization'.

Figure 244: Prioritization Filter Configuration

This menu allows you to edit and delete traffic filters.

The *Source IP/Netmask* and *Dest IP/Netmask* fields specify the IP addresses and masks used to match an outgoing packet. Use 0.0.0.0/0 to generate an "all packets" match.

The *Source Port* and *Dest Port* fields specify the port numbers used to match an outgoing packet. You may specify either a raw number or a mnemonic as specified in the */etc/services* file. This setting matches both UDP and TCP ports, unless the *Protocol* field specifies UDP or TCP.

The *Protocol* field specifies a protocol to match against, currently either TCP, UDP, ICMP, OSPF, VRRP or IPSec.

The *Target Queue* field selects one of the available priority queues.

Section 5.10.6.3

Prioritization Transmit Queue Length

The WAN protocols supplied by the ROX rely upon transmit queues to ensure their efficiency. Even as a packet is starting to be transmitted, other packets can be lining up behind it. Normally there is only one queue, the transmit queue, and packets are transmitted from it in the order in which they arrived.

The transmit queue is a means of enhancing performance.

Prioritization favors some packets over others by transmitting them with preference.

Prioritization works by establishing queues at the required priority levels filling the transmit queue with them in priority order. The aim of establishing low latency for certain traffic is foiled when transmit queue lengths are large because multiple low priority packets may have queued before a high priority packet arrives at the router.

Siemens recommends that the transmit queue length be left at its minimum default value of 1. Higher values, however, may strike a balance between latency and performance.

Section 5.10.7

Prioritization Statistics

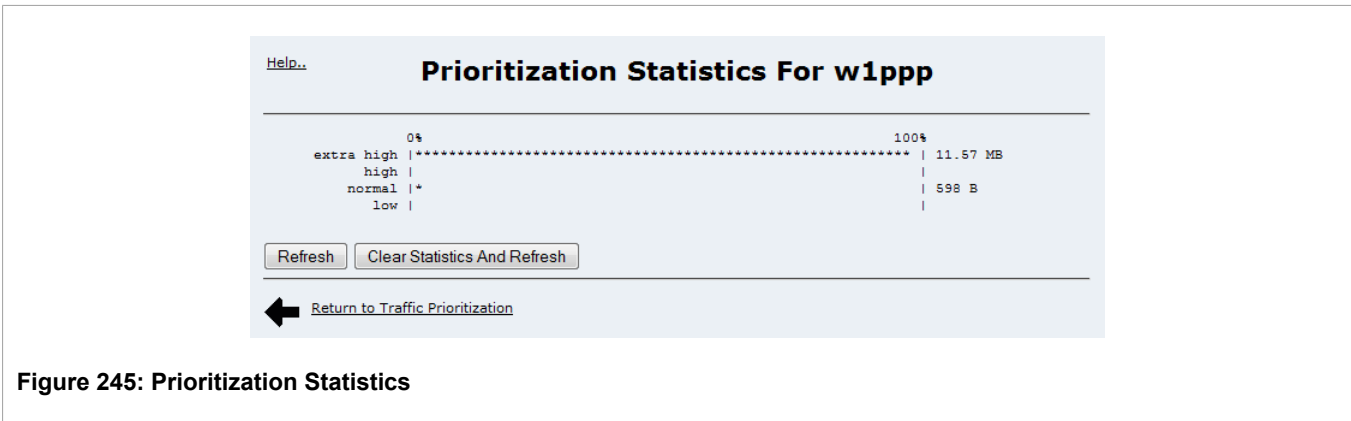


Figure 245: Prioritization Statistics

This menu displays the percentage of interface traffic that has been transmitted from each priority queue. The *Refresh* button causes the statistics to be updated. The *Clear Statistics and Refresh* button causes the statistics to be cleared and then captured after a one second interval.

Section 5.11

Configuring IPsec VPN

This section familiarizes the user with:

- Configuring IPsec VPN Global Options
- Creating VPN Connections
- Configuring L2TPD

- Enabling and Starting IPsec
- Obtaining VPN Status

IPsec (Internet Protocol SECurity) uses strong cryptography to provide both authentication and encryption services. Authentication ensures that packets are from the right sender and have not been altered in transit. Encryption prevents unauthorized reading of packet contents.

These services allow you to build *secure tunnels through untrusted networks*. Everything passing through the untrusted network is encrypted by the IPsec gateway and decrypted by the gateway at the other end. The result is a Virtual Private Network (VPN), a network which is effectively private even though it includes machines at several different sites connected by the insecure Internet.

The IPsec protocols were developed by the Internet Engineering Task Force (IETF) and are required as part of IP version 6.

Openswan is the open source implementation of IPsec used by ROX.

The protocols used by IPsec are the Encapsulating Security Payload (ESP) and Internet Key Exchange (IKE) protocols.

ESP provides encryption and authentication (ensuring that a message originated from the expected sender and has not been altered on route).

IKE negotiates connection parameters, including keys, for ESP. IKE is based on the Diffie-Hellman key exchange protocol, which allows two parties without any initial shared secret to create one in a manner immune to eavesdropping.

Section 5.11.1

IPsec Modes

IPsec has two basic modes of operation. In transport mode, IPsec headers are added as the original IP datagram is created. The resultant packet is composed of an IP header, IPsec headers and IP payload (including a transport header). Transport mode is most commonly used between IPsec end-stations, or between an end-station and a gateway.

In tunnel mode, the original IP datagram is created normally and then encapsulated into a new IP datagram. The resultant packet is composed of a new IP header, IPsec headers, old IP header and IP payload. Tunnel mode is most commonly used between gateways, the gateway acting as a proxy for the hosts behind it.

Section 5.11.2

Policy-Based VPNs

ROX supports the creation of policy-based VPNs, which may be characterized as follows:

- IPsec network interfaces are not created.
- The routing table is not involved in directing packets to the IPsec later.
- Only data traffic matching the tunnel's local and remote subnets is forwarded to the tunnel. Normal traffic is routed by one set of firewall rules and VPN traffic is routed based on separate rules.
- The firewall is configured with a VPN zone of type "IPsec".
- As IPsec packets are received, they are decoded, policy-flagged as IPsec-encoded, and presented as having arrived directly via the same network interface on which they were originally received.

- Firewall rules must be written to allow traffic to and from VPN tunnels. These are based on the normal form of source/destination IP addresses and IP protocol and port numbers. These rules, by virtue of the zones they match, use the policy flagging inserted by netkey and route matching data traffic to the proper interface.

Section 5.11.3

Supported Encryption Protocols

Openswan supports the following standard encryption protocols:

- 3DES (Triple DES) – Uses three DES encryptions on a single data block, with at least two different keys, to get higher security than is available from a single DES pass. 3DES is the most CPU intensive cipher.
- AES – The Advanced Encryption Standard protocol cipher uses a 128-bit block and 128, 192 or 256-bit keys. This is the most secure protocol in use today, and is much preferred to 3DES due to its efficiency.

Section 5.11.4

Public Key and Pre-Shared Keys

In *public key cryptography*, keys are created in matched pairs (called public and private keys). The public key is made public while the private key is kept secret. Messages can then be sent by anyone who knows the public key to the holder of the private key. Only the owner of the private key can decrypt the message.

When you want to use this form of encryption, each router configures its VPN connection to use the RSA algorithm and includes the public signature of its peer. The device's public signature is available from the output of the *Show Public Keys* menu.

In *secret key cryptography*, a single key known to both parties is used for both encryption and decryption.

When you want to use this form of encryption, each router configures its VPN connection to use a secret pre-shared key. The pre-shared key is configured through the *Pre-shared Keys* menu.



NOTE

Use of pre-shared keys require that the IP addresses of both ends of the VPN connection be statically known, so they can't be used with sites with dynamic IPs.

Section 5.11.5

X509 Certificates

When one side of the VPN connection is placed from a dynamic IP (the so-called "roaming client"), X509 Certificates may be used to authenticate the connection. Certificates are digital signatures that are produced by a trusted source, namely a Certificate Authority (CA). For each host, the CA creates an certificate that contains CA and host information and "signs" the certificate by creating a digest of all the fields in the certificate and encrypting the hash value with its *private key*. The encrypted digest is called a "digital signature". The host's certificate and the CA *public key* are installed on all gateways that the host connects to.

When the gateway receives a connection request it uses the CA *public key* to decrypt the signature back into the digest. It then recomputes its own digest from the plain text in the certificate and compares the two. If both digests match, the integrity of the certificate is verified (it was not tampered with), and the public key in the certificate is assumed to be the valid public key of the connecting host.

Section 5.11.6

NAT Traversal

Historically, IPSec has presented problems when connections must traverse a firewall providing Network Address Translation (NAT). The Internet Key Exchange (IKE) used in IPSec is not NAT-translatable. When IPSec connections must traverse a firewall IKE messages and IPSec-protected packets must be encapsulated as User Datagram Protocol (UDP) messages. The encapsulation allows the original untranslated packet to be examined by IPSec.

Section 5.11.7

Other Configuration Supporting IPSec

If the router is to support a remote IPSec client and the client will be assigned an address in a subnet of a local interface, you must activate proxy ARP for that interface. This will cause the router to respond to ARP requests on behalf of the client and direct traffic to it over its connection.

IPSec relies upon the following protocols and ports:

- protocol 51, IPSec-AH Authentication Header (RFC2402),
- protocol 50, IPSec-ESP Encapsulating Security Payload (RFC2046),
- UDP port 500.

You must configure the firewall to accept connections on these ports and protocols. See [Section 5.8.4](#), “Configuring the Firewall and VPN” for details.

Section 5.11.8

The Openswan Configuration Process

Each VPN connection has two ends, in the local router and the remote router. The Openswan developers designed the configuration in such a way that the configuration record describing a VPN connection can be used without change at either end. One side of the connection (typically the local side) is designated the "left" side and the other is designated the "right" side.

A convenient method is to configure both ends simultaneously, having two browser windows up. The relevant information is cut and pasted from window to window.

This module also includes tools to export and import the connection data. The configuration can thus be generated at one router, exported, and imported at the remote router.

Section 5.11.9

IPSec and Router Interfaces

The IPSec daemon requires router interfaces to exist before it starts. If none of the interfaces needed by IPSec exist, IPSec will check for them every minute until at least one does.

Note that in the unlikely event that IPSec uses multiple network interfaces, a stop of any of those interfaces will cause all tunnels to stop.

IPSec may have to be manually restarted after configuring network interfaces when multiple tunnels exist.

Section 5.11.10

L2TPD

L2TP stands for "Layer Two Tunneling Protocol". The main purpose of this protocol is to tunnel PPP packets through an IP network, although it is also able to tunnel other layer 2 protocols.

On ROX, L2TPd is used in conjunction with Openswan and PPP to provide support for establishing a secure, private connection with the router using the Microsoft Windows VPN/L2TP client.



NOTE

L2TPD listens on UDP port 1701. The firewall will need to be configured to allow connections to L2TPD via IPsec but to prevent connections to L2TPD directly without using IPsec.

Section 5.11.11

IPsec VPN Configuration



Figure 246: IPsec VPN Configuration Menu Before Key Generation

Upon the first entry to this menu you will be prompted to generate a VPN host key. Key generation will require about 30 seconds to complete after which the menu appearance will change.

Section 5.11.12

VPN Main Menu

The new menu appearance will resemble that of the following menu with the exception that you will be warned that VPN networking is not enabled. Enable VPN networking via the System folder, Bootup and Shutdown menu.

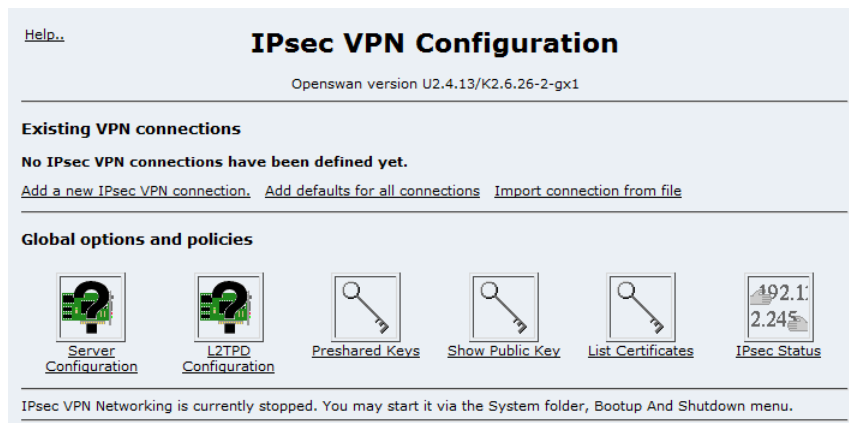


Figure 247: IPsec VPN Configuration Menu Before After Generation

After a VPN connection is created this menu will display an icon for the connection, as shown in the next view of the VPN Configuration menu.

The "Add defaults for all connections" link allows you to create a profile that will apply to all connections for items such as key type, encryption protocol and compression. These defaults can then be overridden on a per connection basis.

The "Add a new IPsec VPN connection" link creates a new connection and its icon.

The "Import connection from file" link creates new connections from imported data.

Select the *Server Configuration* icon to configure server parameters.

Select the *L2TPD Configuration* icon to configure L2TP parameters.

Select the *Pre-shared Keys* icon to create, delete and edit pre-shared keys.

Select the *Show Public Keys* icon to display the server's public key.

Select the *IPsec Status* icon to display information about the server's capabilities and any current connections.

After a VPN connection is created this menu will include a "Start Connection" button that can start or restart VPN connections. This button is shown in the next view of the VPN Configuration menu.

The "Apply Configuration" button restarts the server to activate any configuration changes that have been made, restarting VPN connections.

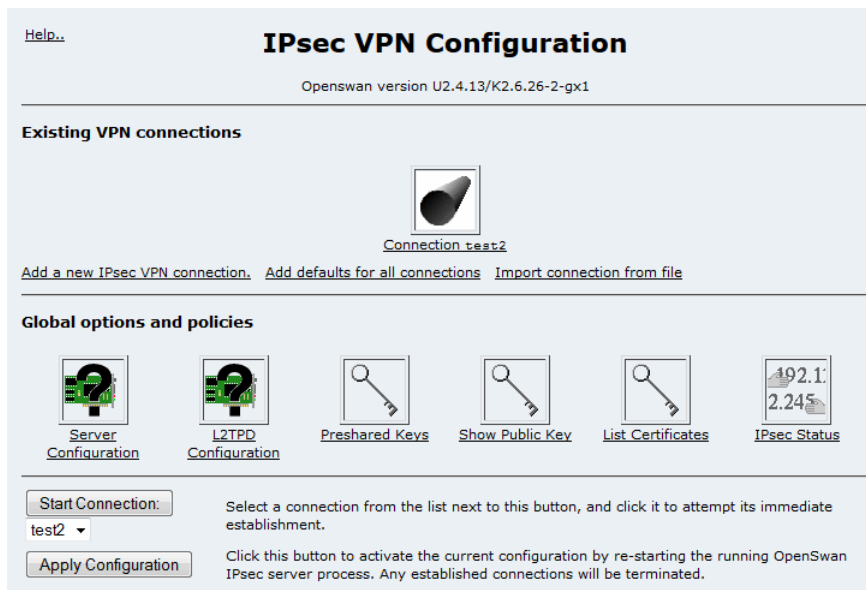


Figure 248: IPsec VPN Configuration After Connections Have Been Created

Section 5.11.13

Server Configuration

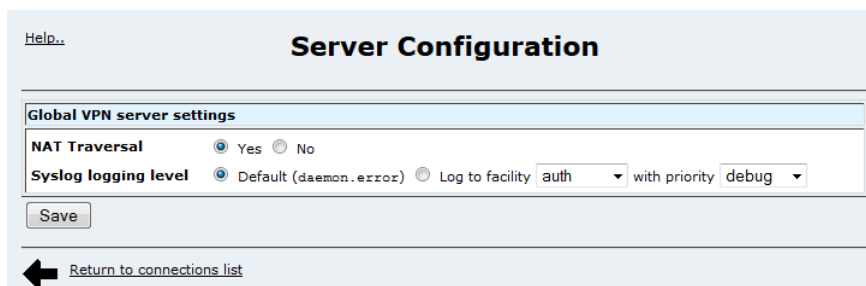


Figure 249: Server Configuration

The *NAT Traversal* fields enable and disable this feature. Enable NAT Traversal if this router originates the VPN connection and the VPN connection passes through a firewall.

The *Syslog logging level* fields determines the facility and priority of log messages generated by Openswan.

Section 5.11.14

L2TPD Configuration

L2TPD Configuration Option	
Local IP Address	<input type="text"/>
First Address in Remote IP Address Pool	<input type="text"/>
Maximum Number of Remote IP Address	<input type="text"/>
Netmask of Remote IP Address Pool	255.255.0.0
Primary DNS Server for Windows Client	<input type="text"/>
Secondary DNS Server for Windows Client	<input type="text"/>
Primary WINS Server for Windows Client	<input type="text"/>
Secondary WINS Server for Windows Client	<input type="text"/>

Figure 250: L2TPD Configuration Menu

The *Local IP Address* field sets the router's IP address for the PPP connection. Note that for all connections, the router will use the same local IP address.

The *First Address in Remote IP Address Pool* field sets the first IP address of the remote IP address pool. IP addresses from this pool are assigned to PPP clients connecting to the L2TPd.

The *Maximum Number of Remote IP address* field sets how many simultaneous connections will be allowed.

The *Netmask of Remote IP Address Pool* field is always 255.255.0.0.

The *Primary DNS Server for Windows Client* field sets the primary DNS Server IP address to be used by connecting Windows clients.

The *Secondary DNS Server for Windows Client* field sets the secondary DNS Server IP address to be used by connecting Windows clients.

The *Primary WINS Server for Windows Client* field sets the primary WINS Server IP address to be used by connecting Windows clients.

The *Secondary WINS Server for Windows Client* field sets the secondary WINS Server IP address to be used by connecting Windows clients.

Section 5.11.14.1

Notes on Configuring a VPN Connection

In addition to configuring L2TP using the menu described above, setting up a VPN connection using L2TP requires several other items to be configured.

Create a VPN connection using the *Edit Connection* menu, described below. Set the *At IPsec Startup* field to "Add Connection", set *Perfect Forwarding Secrecy* to "No", and set *L2TP* to "Yes". Set *Private subnet behind system* to "None" for both left system and right system settings.

It is recommended to configure RADIUS in order to authenticate VPN clients. Note that the ROX RADIUS configuration must be set to authenticate PPP services.

Section 6.3, “VPN/L2TP Configuration in Windows” Describes the steps required to configure the VPN client in Microsoft Windows 2000/XP.



NOTE

AES encryption was designed to be more computationally efficient than 3DES. AES256 or AES128 are therefore preferred for VPN connections, as they require less of the CPU resource. AES256 is not supported in Windows XP but is supported in Vista.

Section 5.11.15

Public Key



Figure 251: Show Public Key

This menu displays the device's public RSA key.

Section 5.11.16

Pre-Shared Keys

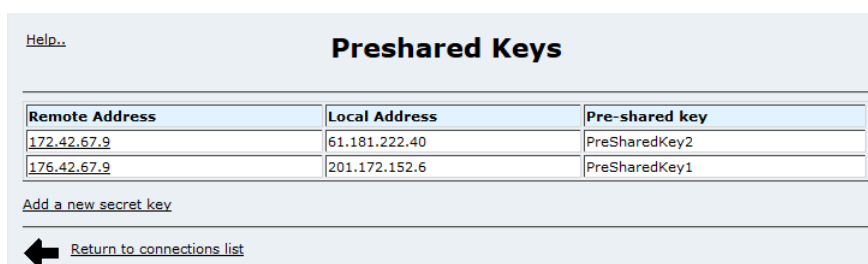


Figure 252: Pre-Shared Keys

This menu creates, deletes and edits pre-shared keys used by VPN connections using secret key encryption. Select the links under the "Remote Address" column to edit or delete a secret key.

The menu will not allow more than one entry to have a specific pair of IP addresses. The menu will not allow a password shorter than eight characters in length.

Section 5.11.17

List Certificates

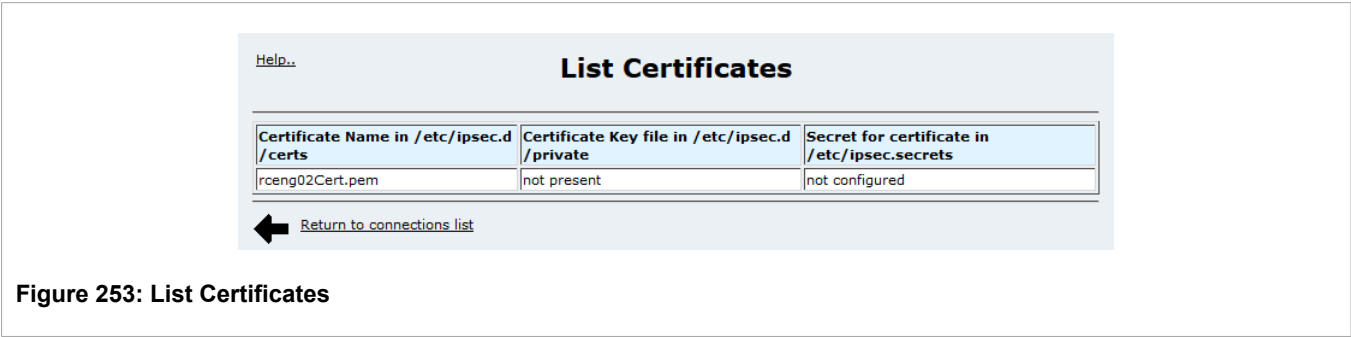


Figure 253: List Certificates

This menu lists available certificate files, their corresponding key files and details whether a public key for the certificate is configured.

Section 5.11.18

VPN Connections

The IPSec main menu "Add a new IPSec VPN connection" link leads to the "Create Connection" menu, creating a new connection and its icon. Selecting the connection's icon from the IPSec main menu displays the same menu, allowing editing and deletion.

An IPSec connection is composed of three types of information. There is information about the local host, the remote host and about the overall connection between them. The configuration data has been designed in such a way that there are identical connection specifications on both ends. Because of this, connection specifications are written in terms of "left" and "right" participants, rather than in terms of local and remote. Which participant is considered left or right is arbitrary; IPSec figures out which one it is being run on based on internal information.

The Create/Edit Connection menu is reflects this organization by being split into three sections. The first section (IPSec VPN Connection Details) describes parameters relating to the connection itself.

The next two sections (Left System's Settings, Right System's Settings) describe IP networking parameters and RSA signatures at each peer. These two sections are identical and are described once.

Section 5.11.18.1

IPSec VPN Connection Details

The screenshot shows the 'Edit Connection' window for an IPSec VPN. The 'Connection name' is 'Remote_16'. The 'At IPsec startup' dropdown is set to 'Default'. The 'Authenticate by' radio buttons are set to 'rsasig'. The 'Connection type' dropdown is set to 'Tunnel (host or network)'. The 'Phase 1 Encryption Protocols(Cipher)' and 'Phase 2 Encryption Protocols(Cipher)' checkboxes are set to 'allow only', and the 'Phase 1 Encryption Protocols(Hash)' and 'Phase 2 Encryption Protocols(Hash)' radio buttons are set to 'sha1'. The 'Compress data?' radio buttons are set to 'No'. The 'Perfect Forward Secrecy' radio buttons are set to 'Yes'. The 'Connection key lifetime' is set to 'secs'.

Figure 254: Editing a VPN Connection, Part 1

The *Connection name* field associates a name with the connection. Do not embed whitespace in the name.

The *At IPsec startup* field determines what happens to the connection after Openswan starts and includes the options "Ignore", "Add connection", "Start Connection", "Route" and "Default". A value of "Ignore" will cause the connection to be ignored. A value of "Add connection" will cause the connection to be established when explicitly started (via command line or the *IPSec VPN Configuration* menu "Start Connection" button). If "Start connection" is chosen then the connection will be authorized when Openswan is started, but not activated until an incoming request arrives. A value of "Route" will cause a route (and only the route) for packets to be established, discarding packets sent there, which may be preferable to having them sent elsewhere based on a more general route (e.g., a default route).

The *Authenticate by* fields select the authentication method. If "Default" is selected the value in the "Defaults for all connections" record is used. If "rsasig" is selected then the *System's public key* of each of the Left System's Settings and Right System's Settings sections must include an RSA signature string or an X.509 certificate must be in use. If "secret" is selected then the *Preshared key* menu must contain a key indexed by the Public IPs of the Left and Right systems.

The *Phase 1 Encryption Protocols* fields select the encryption protocols used for Phase 1 (aka ISAKMP SA). If "Default" is selected, the value in the "Defaults for all connections" record is used. If "allow only" is selected, only the selected protocols among "aes256", "aes192", "aes128" and "3des" will be included in the list of protocols to be negotiated. At connection time, the two peers will compare their capabilities and select the strongest (allowed) common protocol. In decreasing order of cryptographic strength, they are: AES256, AES192, AES128, and 3DES.

The *Phase 1 Encryption Protocols(Hash)* fields select the hash method used for Phase 1 (aka ISAKMP SA). If "Default" is selected, the value in the "Defaults for all connections" record is used. Normally, the user should select the "Default" option. However, in special cases (with some kind of VPN server, for example), you may need to clearly specify which one (sha1 or md5) you want to use.

The *Phase 2 Encryption Protocols* fields select the encryption protocols used for Phase 2 (aka IPsec SA). If "Default" is selected the value in the "Defaults for all connections" record is used. If "allow only" is selected, only the selected protocols among "aes256", "aes192", "aes128" and "3des" will be included in the list of protocols to be negotiated. At connection time the two peers will compare their capabilities and select the strongest common protocol.

The Phase 2 Encryption Protocols(Hash) fields select the hash method used for Phase 2 (aka IPSec SA). If "Default" is selected, the value in the "Defaults for all connections" record is used. Normally, the user should select the "Default" option. However, in special cases (with some kind of VPN server), you may need to clearly specify which one (sha1 or md5) you want to use.

The *Compress data?* fields will select whether data should be compressed prior to encryption. If "Default" is selected the value in the "Defaults for all connections" record is used.

The *Perfect Forward Secrecy* fields will enable PFS, causing keys to be exchanged in a manner which provides attackers that have compromised a key with no advantage in decoding previously intercepted packets or with subsequent packets. Not all clients support PFS.

The *Connection key lifetime* fields determine how long a particular instance of a connection should last, from successful negotiation to expiry. Normally, the connection is renegotiated before it expires.

The *L2TP* field determines whether this connection uses L2TP.

**NOTE**

ROX supports only DH group 1024 bits or greater for both Phase 1 and Phase 2. Please ensure that your client is configured not to use DH group sizes of less than 1024 bits.

Section 5.11.18.2

Left/Right System's Settings

Left system's settings

Public IP address ☐ From default route ☐ Automatic (%any) ☒ Address or hostname .. 206.73.193.8

System identifier ☒ Default ☐ None ☐ IP address .. ☐ Hostname ..

Private subnet behind system ☒ None ☐ behind system ☐ within system

System's public key ☐ None ☒ Entered below .. ☐ Automatic (%any) ☐ Certificate File

0sAQOb6Hs5WTVgvEvcVd1MW0NndZ3i8+xA5Hj9qGcueIWv7eRxHO6AKZq6Q8/R5IaK5dIr6tZNDkECez
dDs79455kBpKZM6KebKm9GTtn0mw915vwURXLIKCI1MBT11ZcwAjcSwojxGcgFAWIp8aJ7m58E0qkFov
ZDSe9ve7I8g1TODnNNdAp6KW3aCnFNufQO6L/ya5Foe7USi+ErOUN8v3sagCxiyJztkxv14h0ytQKG
PEAJJQLVgmTiKF7dqaumvsLzD5tB6sB7BrUmA4469/S7xSFPWvY2yFNTagvMeG2d1K8WH3qUmjKvYvF
q9Yh+T4JHdvZ1BiboC+BQGNhQ8A31YzjIAHTz51+CoCMgEF5IP

Next hop to other system ☒ Default ☐ Automatic ☐ Default route ☐ IP address ..

Figure 255: Editing a VPN Connection, Part 2

The Public IP address fields determine the IP address of the side of the connection being edited. Check the *Address or hostname..* field and provide a fixed IP address or hostname. If this side reflects a remote client whose IP address changes, select *Automatic (%any)*. Use *From default route* if the host's IP is dynamically assigned.

The *System identifier* fields provide IPSec with a way to determine which section of the connection applies to which host. Left to *Default* the parameter will use the public IP address from above. Set to *None*, the router will use an empty id. You can override these with an IP address or hostname.

The *Private subnet behind system* fields determine if this system has an internal network connected to it that the other host should be granted access to. Enter an IP network address and mask into this field. If you enter a subnet of 0.0.0.0/0 in this field, this connection will serve as a default route for all traffic.

The *System's public key* fields provide an RSA key if RSA keying is to be used. If you want to use secret keying, select *None*. When you first create a connection, this field is filled in for you with the local system's RSA key. If

you are filling in this field for the remote system, the key can be obtained from the *Show Public Key* page on that system. Select *Certificate File* and provide a certificate if using X.509 certificates.

The *Next hop to other system* fields determine the address to forward traffic to in order to reach the other system. Unless you have an unusual network setup, this field should be set to *Default route*.

**NOTE**

If you set *Next hop to other system* to "default", you must configure a default route. You can check for the existence of a default route with the *Network Configuration* menu, *Current Routing & Interface Table* icon. A default route will be indicated by a "default" in the *Destination* column.

Section 5.11.18.3

Export Configuration

Selecting the "Export Configuration" button provides a means to capture the connection specification in such a way as to be importable at the remote router.

Section 5.11.19

Showing IPSec Status

IPSec Status

```
1 interface lo/lo 127.0.0.1
2 interface eth1/eth1 10.0.0.253
3 interface eth2/eth2 204.50.190.89
4 interface wlpnp/wlpnp 206.186.238.138
5 %myid = (none)
6 debug none
7 algorithm ESP encrypt: id=2, name=ESP_DES, ivlen=8, keysize=64, keysize=64
8 algorithm ESP encrypt: id=3, name=ESP_3DES, ivlen=8, keysize=192, keysize=192
9 algorithm ESP encrypt: id=7, name=ESP_BLOWFISH, ivlen=8, keysize=40, keysize=448
10 algorithm ESP encrypt: id=11, name=ESP_NULL, ivlen=0, keysize=0, keysize=0
11 algorithm ESP encrypt: id=12, name=ESP_AES, ivlen=8, keysize=128, keysize=256
12 algorithm ESP encrypt: id=252, name=ESP_SERPENT, ivlen=8, keysize=128, keysize=256
13 algorithm ESP encrypt: id=253, name=ESP_TWOFISH, ivlen=8, keysize=128, keysize=256
14 algorithm ESP auth attr: id=1, name=AUTH_ALGORITHM_HMAC_MD5, keysize=128, keysize=128
15 algorithm ESP auth attr: id=2, name=AUTH_ALGORITHM_HMAC_SHA1, keysize=160, keysize=160
16 algorithm ESP auth attr: id=5, name=AUTH_ALGORITHM_HMAC_SHA2_256, keysize=256, keysize=256
17 algorithm ESP auth attr: id=251, name=(null), keysize=0, keysize=0
18 algorithm IKE encrypt: id=7, name=OAKLEY_AES_CBC, blocksize=16, keydeflen=128
19 algorithm IKE encrypt: id=5, name=OAKLEY_3DES_CBC, blocksize=8, keydeflen=192
20 algorithm IKE hash: id=2, name=OAKLEY_SHA, hashsize=20
21 algorithm IKE hash: id=1, name=OAKLEY_MD5, hashsize=16
22 algorithm IKE dh group: id=2, name=OAKLEY_GROUP_MODP1024, bits=1024
23 algorithm IKE dh group: id=5, name=OAKLEY_GROUP_MODP1536, bits=1536
24 algorithm IKE dh group: id=14, name=OAKLEY_GROUP_MODP2048, bits=2048
25 algorithm IKE dh group: id=15, name=OAKLEY_GROUP_MODP3072, bits=3072
26 algorithm IKE dh group: id=16, name=OAKLEY_GROUP_MODP4096, bits=4096
27 algorithm IKE dh group: id=17, name=OAKLEY_GROUP_MODP6144, bits=6144
28 algorithm IKE dh group: id=18, name=OAKLEY_GROUP_MODP8192, bits=8192
29 stats db_ops.c: {curr_cnt, total_cnt, maxsz} :context={0,6144,36} trans={0,6144,336}
   attrs={0,6144,224}
30 "openswantest": 10.0.0.0/8===204.50.190.89...204.50.190.91===192.168.1.0/24; erouted; eroute owner:
   #2997
31 "openswantest":  ike_life: 3600s; IPSec_life: 28800s; rekey_margin: 540s; rekey_fuzz: 100%;
   keyingtries: 0
32 "openswantest":  policy: PSK+ENCRYPT+TUNNEL+PFS+UP; prio: 24,8; interface: eth2;
```

```
33 "openswantest": newest ISAKMP SA: #3093; newest IPSec SA: #2997;
34 "openswantest": IKE algorithms wanted: 5_000-1-5, 5_000-1-2, 5_000-2-5, 5_000-2-2, flags=-strict
35 "openswantest": IKE algorithms found: 5_192-1_128-5, 5_192-1_128-2, 5_192-2_160-5, 5_192-2_160-2,
36 "openswantest": IKE algorithm newest: 3DES_CBC_192-MD5-MODP1536
37 "openswantest": ESP algorithms wanted: 3_000-1, 3_000-2, flags=-strict
38 "openswantest": ESP algorithms loaded: 3_000-1, 3_000-2, flags=-strict
39 "openswantest": ESP algorithm newest: AES_256-HMAC_SHA1; pfsgroup=<Phase1>
40 #3126: "openswantest" STATE_QUICK_I1 (sent QI1, expecting QR1); EVENT_RETRANSMIT in 9s
41 #3093: "openswantest" STATE_MAIN_R3 (sent MR3, ISAKMP SA established); EVENT_SA_REPLACE in 1050s;
    newest ISAKMP
42 #2997: "openswantest" STATE_QUICK_R2 (IPSec SA established); EVENT_SA_REPLACE in 19773s; newest
    IPSec; eroute owner
43 #2997: "openswantest" esp.df9839e9@204.50.190.91 esp.8e2d7255@204.50.190.89 tun.0@204.50.190.91
    tun.0@204.50.190.
```

The "IPSec Status" button produces a window of text similar to that of the above figure (except that line numbers have been inserted for purposes of illustration).

The first group (lines 1-5) describes configured interfaces.

The second group (lines 7-17) describes ESP capabilities. In this group we can see encryption capabilities (lines 7-13) and authentication capabilities (lines 14-17). At least one set of values must match between the left- and right-hand side VPN devices. This is also frequently referred to as the Phase 2 parameters, because the data encryption process is the second and final thing to occur in establishing a VPN.

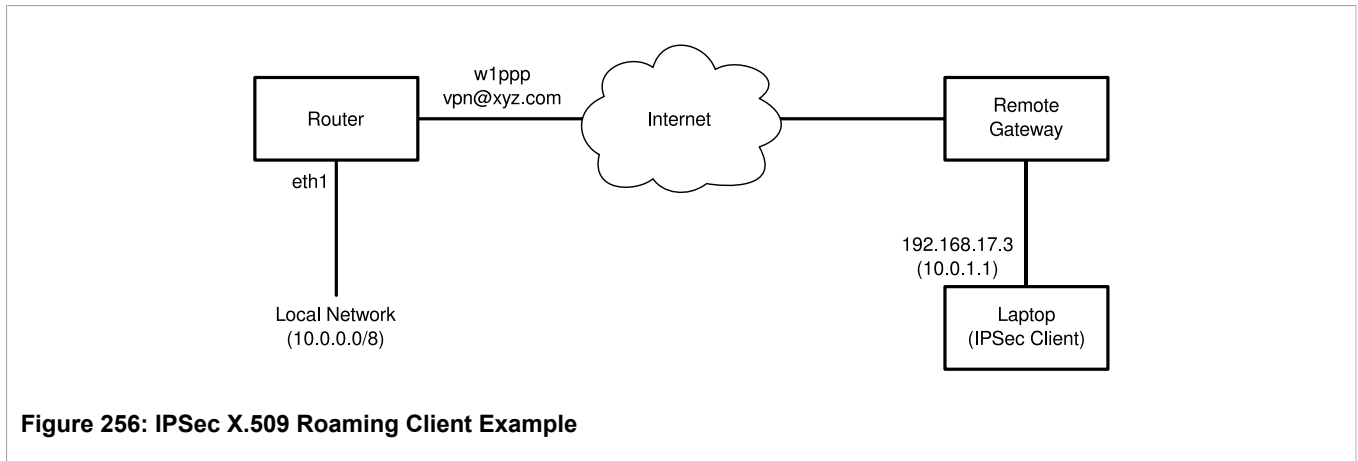
The third group (lines 18-28) describes IKE capabilities and defines the various encrypted key exchange algorithms and their parameters. At least one set of values must match between the left- and right-hand side VPN devices. This is also frequently referred to as the Phase 1 parameters, because the key exchange process is the first thing to occur in establishing a VPN.

The fourth group (lines 30-39) describe connection describe VPN connections (here "openswantest"). The first line is particularly useful since it indicates the connection addresses, subnets and that the connection is active ("erouted"). If there are no entries, then the VPN hasn't been established at all. If there are entries, but no STATE_QUICK_R2 (IPSec SA established) lines then the IPSec parameters are configured, but the tunnel hasn't been established. This can be normal, tunnels become active once the Phase 1 and Phase 2 security associations are created, and this usually only occurs after traffic is flowing. The associations then get torn down after a time-out period.

Section 5.11.20

IPSec X.509 Roaming Client Example

This example details how to set up IPSec connections using X.509 certificates on the router. The router will provide an IPSec gateway to a number of remote clients that connect via an Internet connection. Each of the clients will fetch an IP address locally from a DHCP server, and it is assumed (but not required) that network address translation will be applied at the client end. Each of the clients should "appear" on the local network on a specific IP address. In this example the clients are laptop PCs.



Section 5.11.20.1

Select a Certificate Authority

Begin by constructing the required certificates. You may construct the certificates using a ROX or a third party tool. The device that is used to build the certificates is known as the certificate authority. There are advantages and disadvantages to using the router itself as the authority. It is convenient to use if it is the only router in the network and many clients will be connecting to it. On the other hand, if the router holds the certificate authority and is compromised, all certificates must be constructed again.

Ensure that the Certificate Authority generates certificates with a reasonable life and generates keys of at least 1024 bits in length.

Section 5.11.20.2

Generate X.509 Certificates

Use the authority to produce a certificate authority public certification (cacert) and a certificate for each of the clients and a certificate for the router. The certificate authority will require some information that is shared by all certificates (e.g. a Country Name (C), a State Or Province Name (S), an Organization name (O)) and some per-client information (e.g. a Common Name (CN) and an Email address (E)). Together this information forms the Distinguished Name (DN) and is used by the router and client to validate each other.

Section 5.11.20.3

VPN Networking Parameters

The first step is to identify the key parameters required. The router public gateway (here vpn@xyz.com) and its gateway interface (w1ppp) must be known. The local network subnet (10.0.0.0/8) and each clients' internal network address (here 10.0.1.1) must be known. All client addresses should be assigned from a subnet of the local network (e.g. 10.0.1.0/24). A number of encryption parameters should be decided upon depending upon the client capabilities. Avoid selecting 3DES if possible due to its high overhead.

Section 5.11.20.4

Client Configuration

Depending upon the client, you may be required to produce the certificate in a P12 format, and may be required to include an "export" password as well. This password will be required to be known by the personnel that configure the client in order to import the certificate.

Install the client IPSec software and import the cacert and the client's own certificate and key. Configure the client with the router public gateway, the client's internal network address and the desired encryption parameters. At this point the client should be able to use its Internet connection to ping the public gateway.

Section 5.11.20.5

Router IPSec Configuration

Transfer the cacert and the router's certificate to the router. If your authority prepares a Certificate Revocation List (CRL), you will want to transfer that as well.

The cacert file should be renamed cacert.pem and installed in `/etc/IPSec.d/cacerts/`.

The CRL file should be renamed to `crl.pem` and installed in `/etc/IPSec.d/crls/`.

The router's certificate must be installed in `/etc/IPSec.d/certs/`. Its public key file (e.g. `router.key`) must be installed in `/etc/IPSec.d/private/` and a line `' : RSA router.key "Password"'` (where Password is the pass phrase that was used to generate the certificate) must be added to the end of the `/etc/IPSec.secrets` file.

**NOTE**

The Maintenance Menu, Upload/Download Files sub-menu provides a method to transfer the files directly to the indicated directories.

Enable IPSec from the *Bootup and Shutdown* menu. Visit the *IPSec VPN* menu and generate a public key.

Visit the *Server Configuration* menu and associate the IPSec0 interface with the desired interface the connection will arrive on (here `w1ppp`).

Create a connection for the clients. Set the parameters as follows:

Parameters	Value	Comments
At IPSec Startup	Add connection	We wish to add the connection when the client starts it.
Authenticate by	rsasig	X.509 certificates provide RSA
Connection Type	Tunnel	
Encryption Protocols	As desired	
Compress Data	As desired	
Perfect Forwarding Secrecy	As desired	Recommend "yes"
NAT Traversal	No	Required when the router acts as a client and is behind a NAT firewall.
Left System Settings		Router's side
Public IP Address	Address or hostname .. (IP of public gateway)	
System Identifier	Default	
Private subnet behind system	10.0.0.0/8	

Left System Settings		Router's side
System's public key	Certificate File (router.pem)	
Next hop to other system	Default	
Right System Settings		Laptop1 side
Public IP Address	Automatic	
System Identifier	Default	
Private subnet behind system	10.0.1.0/24	Assign IP based on client from within this subnet
System's public key	Entered below (%cert)	Derive identity from incoming certificate
Next hop to other system	Default	

Apply the configuration to restart the server and create an IPSec0 interface.

Section 5.11.20.6

Firewall IPSec Configuration

Create firewall Zones "vpn" and "net". Ensure that the WAN interface (here w1ppp) and IPSec0 interface are present in the Shorewall *Network Interfaces*. The WAN interfaces should be in zone "net" while IPSec0 should be in zone "vpn".

Add the following firewall rules:

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	all	fw	ah	
ACCEPT	all	fw	esp	
ACCEPT	all	fw	udp	500
ACCEPT	vpn	loc		

Restart the firewall to install the rules.

Section 5.11.20.7

Ethernet Port Configuration

Because the remote client will be assigned a local IP address but is reachable only through the IPSec connection, proxy ARP must be employed. Activate proxy ARP on the Ethernet interface that hosts the local network (here eth1) via the *Networking* Menu, *Ethernet* sub-menu *boot time entry Proxy ARP setting*. When a host on eth1 arps for the remote client address, the router will answer on behalf of the client.

Section 5.12

Configuring Dynamic Routing

This section familiarizes the user with:

- Enabling the Dynamic Routing Suite
- Enabling and starting OSPF, RIP, and BGP
- Configuring OSPF, RIP, and BGP
- Obtaining OSPF, RIP, and BGP Status
- OSPF and VRRP

Dynamic routing is provided by the Quagga suite of routing protocol daemons. Quagga provides three daemons for managing routing, the core, `ripd`, `ospfd`, and `bgpd`.

The core daemon handles interfacing with the kernel to maintain the router's routing table and to check link statuses. It tells RIP, OSPF, and BGP what state links are in, what routes are in the routing table, and some information about the interfaces.

The `ripd`, `ospfd`, and `bgpd` daemons handle communications with other routers using the RIPv2, OSPFv2, and BGP protocols respectively, and decide which routers are preferred to forward to for each network route known to the router.

In complex legacy networks, RIP, OSPF, and BGP may be active on the same router at the same time. Typically, however, one of them is employed.

Section 5.12.1

BGP Fundamentals

The Border Gateway Protocol (BGP, RFC 4271) is a robust and scalable routing protocol. BGP is designed to manage a routing table of up to 90000 routes, and is therefore used in large networks, or among groups of networks which have common administrative and routing policies. If BGP is used to exchange routing information between different networks, it is called Exterior BGP (EBGP); Interior BGP (IBGP) is used to exchange routing information between routers within the same network.

Section 5.12.2

RIP Fundamentals

The Routing Information Protocol determines the best path for routing IP traffic over a TCP/IP network based on the number of hops between any two routers. It uses the shortest route available to a given network as the route to use for sending packets to that network.

ROX's RIP daemon (`ripd`) is an RFC1058 compliant implementation of RIP support RIP version 1 and 2. RIP version 1 is limited to obsolete class based networks, while RIP version 2 supports subnet masks as well as simple authentication for controlling which routers to accept route exchanges with.

RIP uses network and neighbor entries to control which routers it will exchange routes with. A network is either a subnet or a physical (broadcast-capable) network interface. Any router that is part of that subnet or connected to that interface may exchange routes. A neighbor is a specific router to exchange routes with specified by its IP address. For point to point links (T1/E1 links for example) one must use neighbor entries to add other routers to exchange routes with. The maximum number of hops between two points on a RIP network is 15, placing a limit on network size.

Link failures will eventually be noticed although it is not unusual for RIP to take many minutes for a dead route to disappear from the whole network. Large RIP networks could take over an hour to converge when link or route changes occur. For fast convergence and recovery, OSPF is a much better choice. RIP is a fairly old routing protocol and has mostly been superseded by OSPF.

Section 5.12.3

OSPF Fundamentals

The Open Path Shortest First (OSPF) protocol routing determines the best path for routing IP traffic over a TCP/IP network based on link cost and quality. Unlike static routing, OSPF takes link failures and other network topology changes into account. Unlike the RIP routing protocol, OSPF provides less router to router update traffic.

ROX's routing protocols are supplied by the Quagga routing package.

ROX's OSPF daemon (ospfd) is an RFC 2178 compliant implementation of OSPFv2. The daemon also adheres to the RFC2370 (Opaque LSA) and RFC3509 (ABR-Types) extensions.

OSPF network design usually involves partitioning a network into a number of self contained areas. The areas are chosen to minimize intra-area router traffic, making more manageable and reducing the number of advertised routes. Area numbers are assigned to each area. All routers in the area are known as Area routers. If traffic must flow between two areas a router with links in each area is selected to be an Area Border router, and serves as a gateway.

Section 5.12.3.1

Link State Advertisements

When an OSPF configured router starts operating it issues a hello packet. Routers having the same OSPF Area, hello-interval and dead-interval timers will communicate with each others and are said to be neighbors

After discovering its neighbors, a router will exchange Link State Advertisements in order to determine the network topology.

Every 30 minutes (by default) the entire topology of the network must be sent to all routers in an area. If the link speeds are too low, the links too busy or there are too many routes, then some routes may fail to get re-announced and will be aged out.

Splitting the network into smaller areas to reduce the number of routes within an area or reducing the number of routes to be advertised may help to avoid this problem.

In shared access networks (i.e. routers connected by switches or hubs) a designated router and a backup designated are elected to receive route changes from subnets in the area. Once a designated router is picked, all routing state changes are sent to the designated router, which then sends the resulting changes to all the routers.

The election is decided based on the priority assigned to the interface of each router. The highest priority wins. If the priority is tied, the highest router-id wins.

Section 5.12.4

Key OSPF and RIP Parameters

Section 5.12.4.1

Network Areas

Network areas determine the regions within which routes are distributed to other routers. The subnets at a particular router can be added to its OSPF Area. The router will advertise these subnets to all routers in its area.

**NOTE**

OSPF areas must be designed such that no single link failure will cause the network to be split into two disjoint networks.

A router can be part of multiple areas and function as a gateway between areas. When multiple areas are used on a network, area 0 is the backbone area. All areas must have a router connecting them to area 0.

Section 5.12.4.2

Router-ID

Defines the ID of the router. By default this is the highest IP assigned to the router. It is often a good idea to configure this value manually to avoid the router-id changing if interfaces are added or deleted from the router. During elections for designated router, the router-id is one of the values used to pick the winner. Keeping the router-id fixed will avoid any unexpected changes in the election of the master router.

Section 5.12.4.3

Hello Interval and Dead Interval

The hello interval is the time between transmission of OSPF Hello packets. The dead interval is the time to wait without seeing an OSPF Hello packet before declaring a neighboring router dead and discarding its routes. It is recommended that the dead interval be at least four times the hello interval for reliable operation.

Lower values of these settings will help to speed up the change in network routes when the topology of the network changes. It will also increase the load on the router and the links, due to higher traffic caused by the increase in messages. Lower values will also put limits on the number of routes that can be distributed within an area, as will running over slower links.

**NOTE**

OSPF will not work properly if the Hello Interval and Dead Interval are not identical on every router in an area.

Section 5.12.4.4

Active/Passive Interface Default

OSPF regards router interfaces as either passive or active, sending OSPF messages on active interfaces and ignoring passive interfaces. By default, newly created interfaces are viewed as passive from OSPF until they are configured active. This is more efficient and secure for the router. The default type for new interfaces is controlled by the passive interface default option in the OSPF Global Parameters.

**NOTE**

The default setting of Passive Interface Default means that you must explicitly configure interfaces active before OSPF will attempt to use them.

Section 5.12.4.5

Redistributing Routes

Routes for subnets which are directly connected to the router but are not part of the OSPF area or RIP or BGP networks can be advertised if "redistribute connected" is enabled in the OSPF, RIP, or BGP Global Parameters. Static routes and routes handled by the kernel can also be redistributed if "redistribute static" and "redistribute kernel" are enabled, respectively.

Section 5.12.4.6

Link Detect

Link detection is enabled for active network interfaces, which ensures that the appropriate routing daemon is notified when an interface goes down, and stops advertising subnets associated with that interface. The routing daemon resumes advertising the subnet when the link is restored. This allows routing daemons to detect link failures more rapidly (as the router does not have to wait for a "dead interval" to time out). Link Detect also causes "redistributed" routes to start and stop being advertised based on the status of their interface links.

Section 5.12.4.7

Configuring OSPF Link Costs

Link cost is used when multiple links can reach a given destination, to determine which route to use. OSPF will (by default) assign the same cost to all links unless provided with extra information about the links. Each interface is assumed to be 10Mbit unless told otherwise in the Core Interface configuration.

The reference bandwidth for link cost calculations is 100Mbit by default in the OSPF Global Parameters. The reference bandwidth divided by the link bandwidth gives the default cost for a link, which by default is 10. If a specific bandwidth is assigned to each link, the costs will take this into account.

It is also possible to manually assign a cost to using a link in the OSPF Interface Configuration for each interface for cases where the speed of the link is not desired as the method for choosing the best link.

Section 5.12.4.8

OSPF Authentication

OSPF authentication is used when it is desirable to prevent unauthorized routers from joining the OSPF network. By enabling authentication and configuring a shared key on all the routers, only routers which have the same authentication key will be able to send and receive advertisements within the OSPF network. Authentication adds a small overhead due to the encryption of messages, so is not to be preferred on completely private networks with controlled access.

Section 5.12.4.9

RIP Authentication

RIP authentication is used when it is desirable to prevent unauthorized routers from joining the network. RIP authentication is supported by per-interface configuration or the use of key-chains. Separate key chains spanning different groups of interfaces and having separate lifespans are possible. By enabling authentication and configuring a shared key on all the routers, only routers which have the same authentication key will be able to send and receive advertisements within the RIP network.

Section 5.12.4.10

Administrative Distances

The router may work with different routing protocols at the same time, as well as employing local interface and statically assigned routes. An administrative distance, (from 0 to 255) is a rating of the trustworthiness of a routing information source. For a given route, the protocol having the lowest administrative distance will be chosen.

By default the distances for a connected interface is 0 and for a static route is 1. By default, OSPF will set an administrative distance of 110 and RIP will set a distance of 120.

Section 5.12.5

OSPF and VRRP Example Network

This network consists of three routers connected in a ring with T1/E1 links. Router 1 and 2 and the switched network represent a remote site in which the routers supply a redundant gateway to the hosts via VRRP and the T1/E1 links supply a redundant network connection to the rest of the network.




Figure 257: OSPF and VRRP Example

Section 5.12.5.1

Area and Subnets

As the OSPF design is simple, an area of 0 is used. The three point-to-point T1/E1 links are placed in the area by adding 1.1.1.0/24 to it. Router 1 and 2 will include their Ethernet links by adding subnet 1.1.2.0/24 to their area descriptions. Router 3 must also include 2.2.2.0/24 in its area description so that its existence is advertised.

The point-to-point T1/E1 interfaces and Ethernet interfaces on Router 1 and 2 must be made active. The Ethernet interface on Router 3 can be left passive since it does not participate in OSPF advertisements.

Router 1 and 2 must enable link-detect, to stop advertising 1.1.1.0/24 in the event of a link failure.

Section 5.12.5.2

VRRP Operation

Router 1 and 2 have VRRP setup on their Ethernet connection so that they can both function as the gateway for the clients on their network segment. Normally Router 1 is the VRRP master, and only in case of a link failure to the switch or the router failing, will Router 2 take over the virtual IP. The virtual IP used as the gateway is 1.1.2.254. Each router also has its own IP on the network so that each can be reached individually.

If Router 1 or its Ethernet link fail, VRRP will detect the link being down and remove the direct route to the 1.1.2.0/24. VRRP on Router 2 will stop seeing messages from Router 1, elect itself master and will take over the gateway for the network.

OSPF on router 1 will notice the link being down (and the route to 1.1.2.0/24 disappearing) and will use information from router 2 install a route to 1.1.2.0/24 via Router 2.

Router 3 will notice that Router 2 is now a more direct path to 1.1.2.0/24 network and start sending to Router 2 instead of Router 1.

After the failure all routers still know how to reach the entire network, and the clients on 1.1.2.0/24 can still send on the network using the same gateway address. The clients will see only a MAC address change of the gateway and experience a few seconds of network outage. When the link returns, VRRP will switch back to the master, and the routes will return to their normal state.

Note that if the Router 1 WAN link fails, Router 1 will see routes to Router 3 via the Router 1 – Router 2 WAN and Ethernet links. If the faster Router 1 – Router 2 Ethernet path fails, Router 1 will fall back to the Router 1 – Router 2 WAN link.

Note that it would not be useful to leave the Ethernet 1.1.2.0/24 subnets out of the area and turn on redistribute connected as OSPF would not use the subnets for routing.

Section 5.12.6

Dynamic Routing Configuration

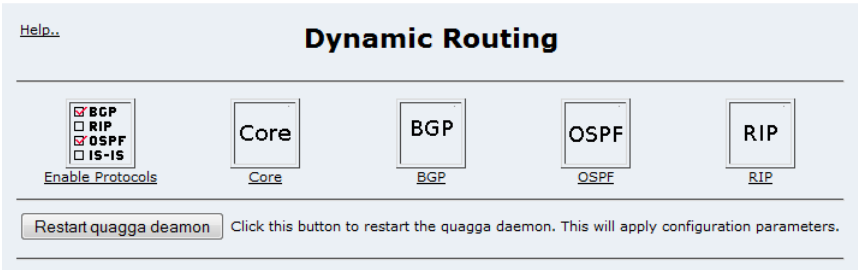


Figure 258: Dynamic Routing Main Menu

In order to begin dynamic route processing on ROX, one of the protocols, RIP, OSPF, or BGP must be enabled in the *Enable Protocols* menu under *Dynamic Routing*.

The *Core* menu configures link related items such as link-detect and link cost.

The *RIP*, *OSPF*, and *BGP* menus configure the corresponding routing protocols.

Section 5.12.6.1

Enable Protocols

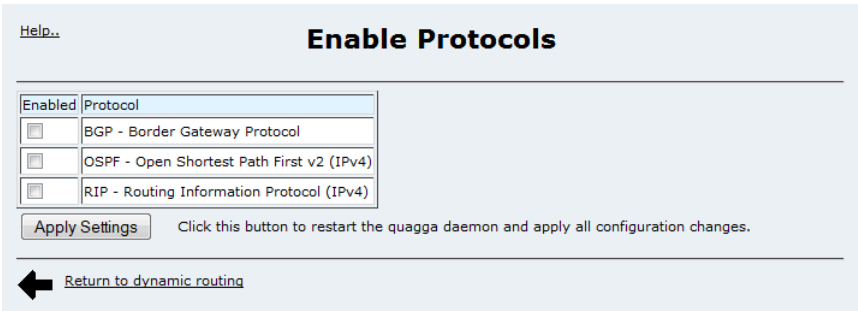



Figure 259: Dynamic Protocol Enable Menu

This menu enables RIP and OSPF for dynamic routing.

Section 5.12.6.2

Core

**Figure 260: Core Menu**

The Core routing daemon handles communications between the kernel of the router and the other dynamic routing protocols. The core handles link detection, static route monitoring, and routes for directly connected interfaces on the router. It also manages adding routes to the kernel routing table based on the routes discovered by other dynamic routing protocols. Core is always enabled whenever dynamic routing is enabled as it is required by all other dynamic routing protocols.

Section 5.12.6.3

Core Global Parameters

Parameter	Value	Description [Possible values] (default value)
Enable Password	•••••	Enable password. For configuration access. [string without spaces] (previous password)
Telnet Password	••••••••••	Telnet password. For port 2601 access. [string without spaces] (previous password)
Hostname		Identifier of router. Often the DNS name of the router. [string without spaces] (no hostname)
Router ID		Identifier of router. Often the main IP address of the router. [A.B.C.D] (highest IP of system)

Note: The router-id is automatically picked as the highest IP address assigned to the unit at startup unless the config contains a router-id to use.
In OSPF, the router-id change only updates the config and takes effect on restart. In BGP, the router-id change updates the config and takes effect immediately.

Figure 261: Core Global Parameters

The *Enable Password* field sets the password to be used for the enable command of core. This is used by the telnet interface of core to control access to the configuration.

The *Telnet Password* field sets the password to be used for telnet access to core. This is used as the login password of core when locally telnetting to port 2601 of the router.

The *Hostname* field sets the hostname for the core daemon. This value is only used as a reference for convenience. The telnet interface prompt will contain this hostname. The router's system wide hostname is used if this field is left blank.

The *Router ID* field sets the router-id to use for the core daemon. This value is used as a unique identifier for the dynamic routing protocol to identify which router sent which route advertisement. By default it uses the highest IP assigned to an interface on the router. It is recommended that this value be set to a unique fixed IP on each router. Note that this value is used by both OSPF and by BGP if not overridden under the global configuration for each protocol.

Section 5.12.6.4

Core Interface Parameters

Parameters specific to one interface are configured here.

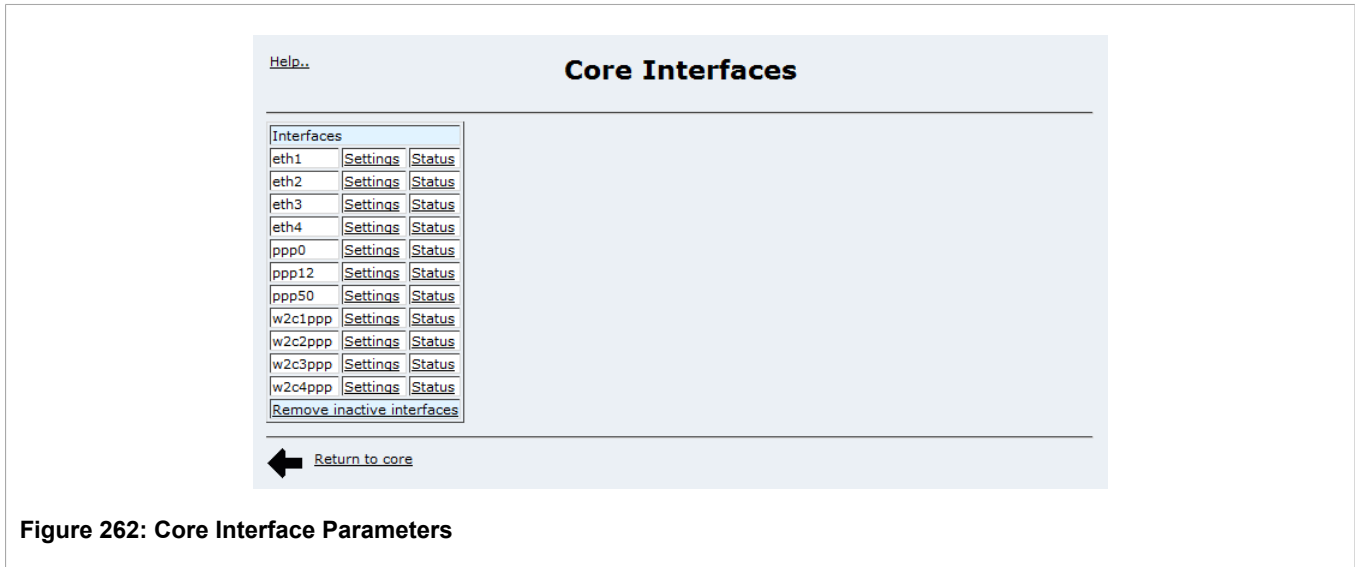


Figure 262: Core Interface Parameters

Each interface on the router is listed. Clicking on settings displays a menu of configuration options for that interface. Clicking on status displays the current status of the interface, including link state, IP address and traffic counts.

Clicking "Remove inactive interfaces" purges the list of any interfaces which are no longer configured on the router.

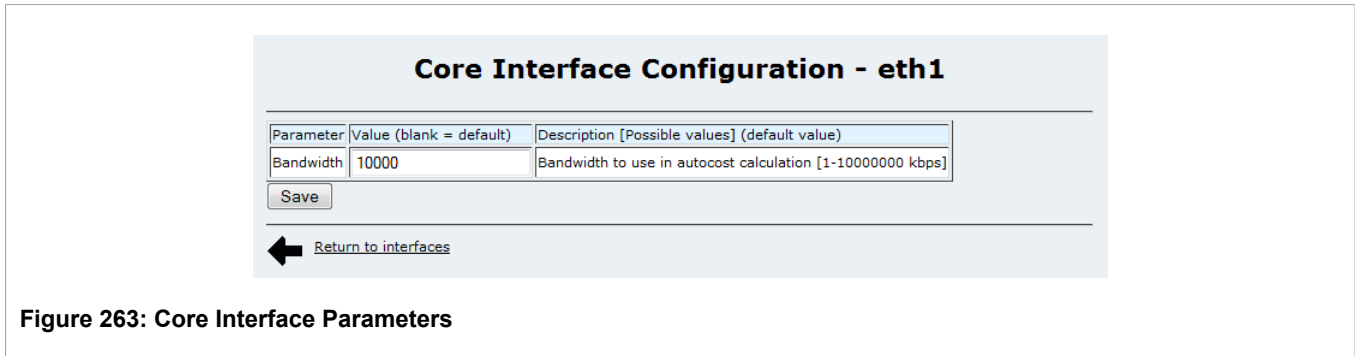


Figure 263: Core Interface Parameters

The *Bandwidth* field sets the bandwidth value to assume for the interface when automatically calculating a cost for using the link on this interface. By default all interfaces are treated as 10Mbit (10000 Kbps). OSPF by default uses an automatic cost of 10 for all links by calculating is as reference bandwidth (100Mbit) divided by the link bandwidth (10Mbit). If a manual cost is assigned to the interface in OSPF, this value is ignored. RIP does not use this parameter.

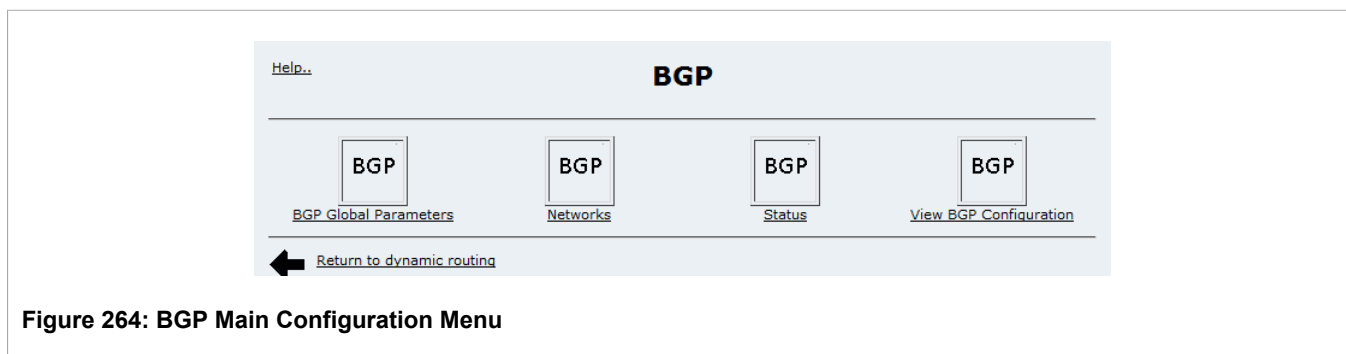
Section 5.12.6.5

View Core Configuration

This menu shows the current configuration file for the Core interfaces.

Section 5.12.6.6

BGP Configuration



This menu contains the configuration and status of BGP on the router.

The *BGP Global Parameters* and *Networks* menus configure BGP. The *Status* and *View BGP Configuration* menus display the actual status and configuration file contents of BGP

Section 5.12.6.7

BGP Global Parameters



NOTE

The AS ID, defined below, must be configured prior to any configuration of BGP networks or neighbors.

[Help..](#)

BGP Global Parameters

Parameter	Value	Description [Possible values] (default value)
Enable Password	••••••••	Enable password. For configuration access. [string without spaces] (previous password)
Telnet Password	••••••••	Telnet password. For port 2605 access. [string without spaces] (previous password)
Hostname		Identifier of router. Often the DNS name of the router. [string without spaces] (no hostname)
Distance	value: <input type="text"/> IP/Mask: <input type="text"/>	Define an administrative distance [unset,1-255] (unset=not used)
Always compare Med	enable <input type="checkbox"/>	Always comparing MED from different neighbors [enable/disable] (disable)
Default local preference	100	Local preference value (high value means preferred in IBGP) [0-4294967295] (100)
Deterministic Med	enable <input type="checkbox"/>	Pick the best-MED path among paths advertised from neighboring AS [enable/disable] (disable)
Redistribute Connected	enable <input type="checkbox"/> metric: <input type="text"/>	Redistribute routes for directly connected interfaces to BGP area routers. [enable/disable,0-16] (disabled,unset)
Redistribute Static	enable <input type="checkbox"/> metric: <input type="text"/>	Redistribute static routes to BGP area routers. [enable/disable,0-16] (disabled,unset)
Redistribute Kernel	enable <input type="checkbox"/> metric: <input type="text"/>	Redistribute kernel routes to BGP area routers. [enable/disable,0-16] (disabled,unset)
Redistribute OSPF	enable <input type="checkbox"/> metric: <input type="text"/>	Redistribute ospf routes to BGP area routers. [enable/disable,0-16] (disabled,unset)
Redistribute RIP	enable <input type="checkbox"/> metric: <input type="text"/>	Redistribute rip routes to BGP area routers. [enable/disable,0-16] (disabled,unset)
AS ID	200	Autonomous System ID. Note: When AS ID is changed, all BGP configurations related to this AS will be removed.
Router ID	192.168.10.1	Identifier of router. Often the main IP address of the router. [A.B.C.D] (highest IP of system)

Note: The router-id is automatically picked as the highest IP address assigned to the unit at startup unless the config contains a router-id to use.
The router-id change updates the config and takes effect immediately.


 [Return to bgp](#)

Figure 265: BGP Global Parameter Menu

The *Enable Password* field sets the password to be used for the bgpd *enable* command. This is used by the telnet interface to control access to the bgpd configuration.

The *Telnet Password* field sets the password to be used for telnet access to bgpd. This is used as the bgpd login password when locally telnetting to port 2605 of the router.

The *Hostname* field sets the hostname for the bgpd daemon. This value is only used as a reference for convenience. The telnet interface prompt will contain this hostname. The router's system wide hostname is used if this field is left blank.

The *Distance* field sets the administrative distance to use for all routes unless overridden by other distance settings.

The *Always Compare Med* field enables the comparison of MED for paths from neighbors in different AS (Autonomous System). By default, MED comparison is only done among paths within the same AS. Enabling this option, BGP will always perform MED comparison regardless of AS. The path with a lower MED is preferred to one with a higher value.

The *Default Local Preference* field sets the local "preference" value for the router.

The *Deterministic Med* field is used to select the best-MED path among paths advertised from neighboring AS.

The *Redistribute Connected* fields control the distribution of connected routes. When enabled, BGP will advertise routes to directly connected interfaces to other BGP routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute Static* fields control the distribution of static routes. When enabled, BGP will advertise static routes created using the Network Configuration/Routing and Default Route menu to other BGP routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute Kernel* fields control the distribution of kernel routes. When enabled, BGP will advertise routes from the kernel routing table, which includes static routes entered by the administrator, to other BGP routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute OSPF* fields control the distribution of routes learned by OSPF. When enabled, BG will advertise routes learned by OSPF.

The *Redistribute RIP* fields control the distribution of routes learned by RIP. When enabled, BGP will advertise routes learned by RIP.

The *AS ID* is the Autonomous System ID used between BGP routers. Note that for BGP routers connected to the Internet, these numbers are controlled by IANA (the Internet Assigned Numbers Authority).

The *Router ID* field sets the router-id to use for the BGP daemon. This value is used as a unique identifier for the dynamic routing protocol to identify which router sent which route advertisement. By default it uses the highest IP assigned to an interface on the router. It is recommended that this value be set to a unique fixed IP on each router

Section 5.12.6.8

BGP Networks

Help...

Networks

Neighbors Remote-AS Action

[Add a new neighbor...](#)

Networks	Action
Subnet (x.x.x.x/x) or x.x.x.x	Add

[Return to bgp](#)

Figure 266: BGP Networks Menu

Neighbors are other BGP routers with which to exchange routing information. One or more neighbors must be specified in order for BGP to operate. To add a neighbor to the BGP network, click on *Add a new neighbor* to configure its BGP attributes.

Networks may be specified in order to add BGP routers connected to the specified subnets. Note that a network specification need not be a given valid entry in the routing table. Since BGP is a border gateway protocol, one would more typically enter a more general network specification here. For example, if a routed network inside the AS comprised many different Class C subnets (/24) of the 192.168.0.0/16 range, it would be more efficient to advertise the one Class B network specification, 192.168.0.0/16, to one's BGP neighbors.



NOTE

If BGP Neighbors are specified but no Networks are specified, then the router will receive BGP routing information from its neighbors but will not advertise any routes to them.

Section 5.12.6.9

BGP Neighbor Configuration

[Help..](#)

BGP Configuration

Parameter	Value	Description [Possible values] (default value)
Neighbor IP	192.168.10.2	Neighbor IP address
Remote AS	210	Remote AS
Maximum-prefix	50	Maximum prefix accept from the neighbor

Inbound attributes for the neighbor

Metric attribute		Metric attribute for the neighbor
Local Preference attribute		Local preference attribute for the neighbor
Weight attribute		Weight attribute for the neighbor

Outbound attributes for the neighbor

Metric attribute		Metric attribute for the neighbor
Local Preference attribute		Local preference attribute for the neighbor

Save

Delete

Return to bgp networks

Figure 267: BGP Neighbor Configuration Menu

Neighbor IP is the IP address of a BGP neighbor to add. A BGP Neighbor is

Remote AS is Autonomous System ID of a BGP neighbor.

The *Metric attribute* is propagated throughout an AS. A path with a lower metric attribute is preferred over one with a higher value.

The *Local Preference attribute* is propagated throughout an AS. A higher local preference attribute is preferred over a lower one for the entire AS.

The *Weight attribute* is defined locally. If the router learns more than one route to the same destination, the one from the router with a higher weight is preferred.

Section 5.12.6.10

BGP Status

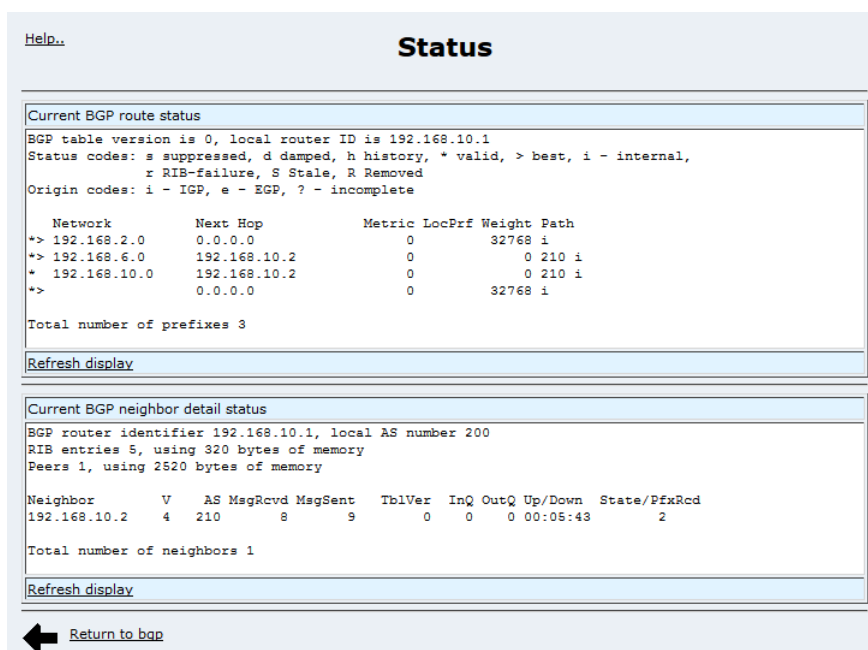


Figure 268: BGP Status Display

The BGP Status menu displays:

- A list of routes currently managed by BGP along with the status of each one
- A list of BGP neighbors along with statistics and state information for each one
- Other statistics, including known peers and memory usage

The table in the *Current BGP route status* report provides the following information:

- *Network* is the IP address for the network
- *Next Hop* is the next hop IP address
- *Metric* is the metric value
- *LocPrf* is the local preference
- *Weight* is the weight
- *Path* is the Autonomous System (AS) path

The table in the *Current BGP neighbor detail status* report provides the following information:

- *Neighbor* is the IP address for the neighbor
- *V* is the IP version
- *AS* is the Autonomous System (AS) number
- *MsgRcvd* is the number of messages received
- *MsgSent* is the number of messages sent
- *TblVer* is the table version
- *InQ* is the in queue depth

- *OutQ* is the out queue depth
- *Up/Down* is the last up/down time
- *State/PfxRcd* is either the number prefixes received (if the connection is established) or the state of the connection

Section 5.12.6.11

View BGP Configuration

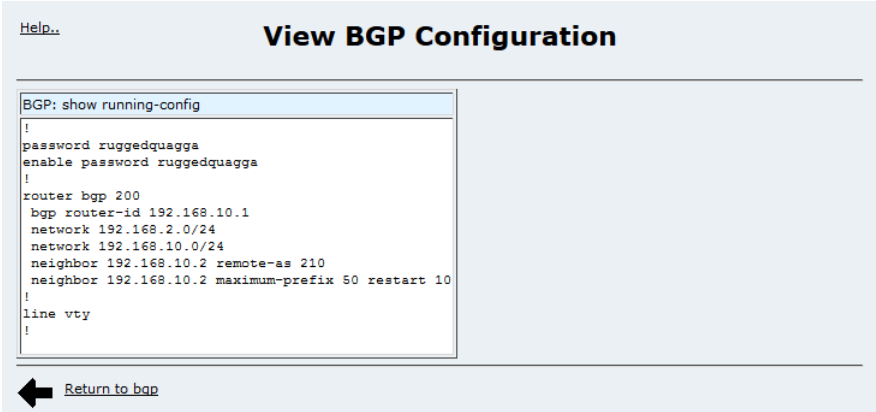


Figure 269: View BGP Configuration Menu

This menu displays the text of the active configuration file for the BGP daemon.

Section 5.12.6.12

OSPF

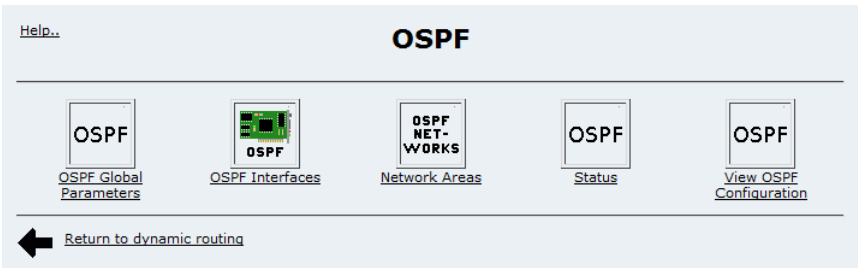


Figure 270: OSPF Menu

This menu contains the configuration and status of OSPF on the router.

The *OSPF Global Parameters*, *OSPF Interfaces* and *Network Areas* menus configure OSPF. The *Status* and *View OSPF Configuration* menu display the actual status and configuration file contents of OSPF.

Section 5.12.6.13

OSPF Global Parameters

[Help..](#)

OSPF Global Parameters

Parameter	Value	Description [Possible values] (default value)
Enable Password	••••••••••	Enable password. For configuration access. [string without spaces] (previous password)
Telnet Password	••••••••••	Telnet password. For port 2604 access. [string without spaces] (previous password)
ABR-Type	cisco ▾	Set OSPF ABR type [standard/cisco/ibm/shortcut] (cisco)
Auto Cost Reference Bandwidth	100	Calculate OSPF interface cost according to bandwidth [1-4294967 Mbps] (100)
Default-Information Originate	enable <input type="checkbox"/>	Advertise default route (disabled)
Default Metric	20	Control distribution of default information [1-16777214] (20)
Distance		Define an administrative distance [unset,1-255] (unset=not used)
Distance OSPF External		Define an administrative distance (external) [unset,1-255] (unset=use Distance)
Distance OSPF Inter-area		Define an administrative distance (inter-area) [unset,1-255] (unset=use Distance)
Distance OSPF Intra-area		Define an administrative distance (intra-area) [unset,1-255] (unset=use Distance)
Hostname		Identifier of router. Often the DNS name of the router. [string without spaces] (no hostname)
Opaque LSA	enable <input type="checkbox"/>	Enable Opaque LSA capability (disabled)
Passive Default	enable <input type="checkbox"/>	Set new interfaces passive by default (enabled)
Refresh Timer	10	Set refresh timer [10-1800 Seconds] (10)
RFC 1583 Compatibility	enable <input type="checkbox"/>	Enable compatibility with obsolete RFC1583 OSPF (current is RFC2178) (disabled)
Redistribute Connected	enable <input type="checkbox"/>	Redistribute routes for directly connected interfaces to OSPF area routers. [enable/disable,1/2,0-16777214] (disabled,2,unset)
	metric-type 2 ▾	
	metric	
Redistribute Static	enable <input type="checkbox"/>	Redistribute static routes to OSPF area routers. [enable/disable,1/2,0-16777214] (disabled,2,unset)
	metric-type 2 ▾	
	metric	
Redistribute Kernel	enable <input type="checkbox"/>	Redistribute kernel routes to OSPF area routers. [enable/disable,1/2,0-16777214] (disabled,2,unset)
	metric-type 2 ▾	
	metric	
Redistribute RIP	enable <input type="checkbox"/>	Redistribute rip routes to OSPF area routers. [enable/disable,1/2,0-16777214] (disabled,2,unset)
	metric-type 2 ▾	
	metric	
Redistribute BGP	enable <input type="checkbox"/>	Redistribute bgp routes to OSPF area routers. [enable/disable,1/2,0-16777214] (disabled,2,unset)
	metric-type 2 ▾	
	metric	
Router ID		Identifier of router. Often the main IP address of the router. [A.B.C.D] (highest IP of system)

Note: The router-id is automatically picked as the highest IP address assigned to the unit at startup unless the config contains a router-id to use.
The router-id change only updates the config and takes effect on restart.

[Return to ospf](#)

Figure 271: OSPF Global Parameters

The *Enable Password* field sets the password to be used for the enable command of ospfd. This is used by the telnet interface of ospfd to control access to the configuration.

The *Telnet Password* field sets the password to be used for telnet access to ospfd. This is used as the login password of ospfd when locally telnetting to port 2604 of the router.

The *ABR-Type* field select which method to use on area border routers to manage inter area routes. Standard follows RFC2178, Cisco and IBM follow RFC3509. Shortcut is covered by the draft-ietf-ospf-shortcut-abr-00.txt document. Standard requires all ABRs to have a backbone connection. The other three methods allow for ABRs that do not have a backbone connection.

The *Auto Cost Reference Bandwidth* field sets the reference bandwidth used to calculate auto costs for OSPF interfaces. The auto cost is the reference bandwidth divided by the interface bandwidth. By default this is 100Mbit/10Mbit = auto cost of 10. The interface cost is set in the Core Interface configuration for each interface. The cost for each interface can also be set in the OSPF Interface configuration to override the auto cost calculation.

The *Default Metric* field sets the default metric to be used for OSPF routes which don't have another metric specified.

The *Default-Information Originate* field, when enabled, causes the router to advertise its default route to the OSPF network.

The *Distance* field sets the administrative distance to use for all routes unless overridden by other distance settings.

The *Distance External* field sets the administrative distance to use for all external routes (backbone routes). The *Distance Inter-area* field sets the administrative distance to use for all routes between areas. The *Distance Intra-area* field sets the administrative distance to use for all routes within an area.

The *Hostname* field sets the hostname for the ospf daemon. This value is only used as a reference for convenience. The telnet interface prompt will contain this hostname. The router's system wide hostname is used if this field is left blank.

The *Opaque LSA* field controls the opaque LSA option. This feature is covered in RFC2370. This feature is sometimes used to distribute application specific information through a network using OSPF LSAs.

The *Passive Default* option controls the default active/passive state of new interfaces. When enabled all new interfaces will be passive by default. The passive state of individual interfaces is controlled from the OSPF Interfaces configuration.

The *Refresh Timer* field controls how frequently OSPF LSA refreshes occur.

The *RFC 1583 Compatibility* field controls support for RFC1583 compatibility. If this option is enabled OSPF will be compatible with the obsolete RFC1583 version of OSPF. By default it is compatible with RFC2178 version of OSPF only.

The *Redistribute Connected* fields control distribution of connected routes. When enabled, OSPF will advertise routes to directly connected interfaces to other OSPF routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute Static* fields control the distribution of static routes created using the Network Configuration / Routing and Default Route menu. When this parameter is enabled, OSPF will advertise these static routes to other OSPF routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute Kernel* fields control distribution of kernel routes. When enabled, OSPF will advertise routes from the kernel routing table, which includes static routes entered by the administrator, to other OSPF routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute RIP* fields control distribution of routes learned by RIP. When enabled, OSPF will advertise routes learned by RIP.

The *Redistribute BGP* fields control distribution of routes learned by BGP. When enabled, OSPF will advertise routes learned by BGP.

The *Router ID* field sets the router-id to use for the ospf daemon. This value is used as a unique identifier for the dynamic routing protocol to identify which router sent which route advertisement. If it is not set here, it defaults to using the Router ID set in [Section 5.12.6.3, “Core Global Parameters”](#). If this last is not set either, the Router ID defaults to a string containing the highest IP address assigned to an interface on the router, in dotted quad notation. It is recommended that this value be set to a unique fixed value on each router. Note that for the new router-id to take effect, the routing daemon must be restarted.

Section 5.12.6.14

OSPF Interfaces

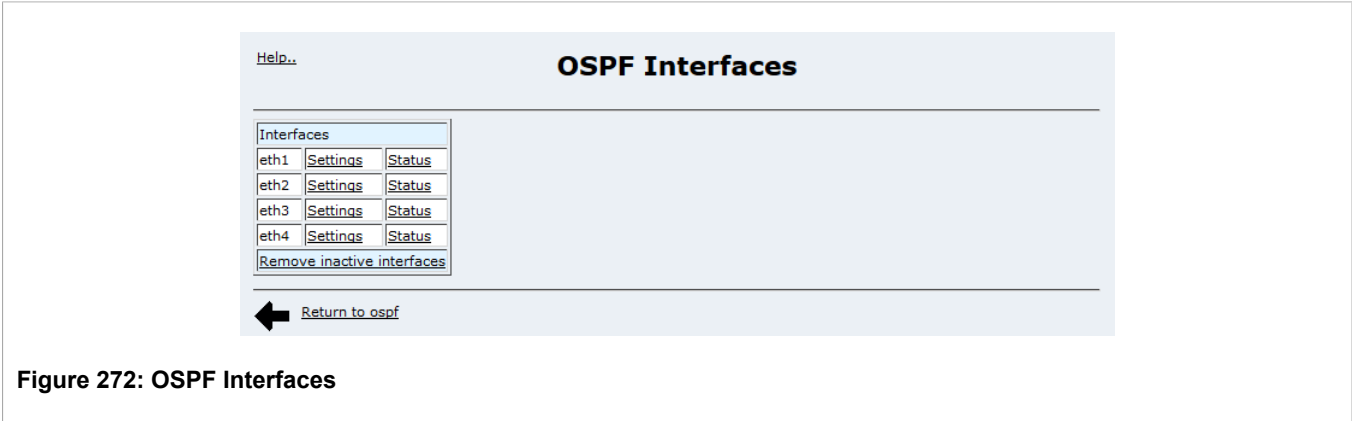


Figure 272: OSPF Interfaces

Parameters specific to one interface are configured here.

Each interface on the router is listed. Clicking on settings displays a menu of configuration options for that interface. Clicking on status displays the current status of the interface, including link state and current OSPF status on the interface. If an interface is not part of an area it will show up as OSPF not enabled on interface.

Clicking "Remove inactive interfaces" purges the list of any interfaces which are no longer configured on the router.

The *Cost* field controls the administrative cost of routing over this interface. By default the cost is auto calculated as the ospf reference bandwidth divided by the core interface bandwidth. By default this is 100Mbit/10Mbit = cost 10.

The *Priority* field controls the priority associated with this interface. By default the priority of interfaces is 1. The router with the highest priority wins elections for designated router for an area.

The *Hello Interval* field controls how often hello packets are sent to other routers in the area. This value must match on all router interfaces in an area.

The *Dead Interval* field controls how long to wait for hello packets before declaring another router dead. This should normally be set to 4 times the hello interval.

The *Retransmit Interval* field controls the delay between retransmissions.

The *Transmit Delay* field controls the estimated number of seconds to transmit a link state update packet. This should take into account transmission and propagation delays of the interface.

The *Passive Interface* option controls if an interface is active or passive. Passive interfaces do not send LSAs to other routers and are not part of an OSPF area.

The *Authentication* field controls the type of authentication to use when communicating with other routers. It can be set to "default", which uses the routing package's default value, "null" (just check for message corruption), or "message digest", which cryptographically signs each message with a shared key.

The *Message Digest Keys* table allow the addition and deletion of keys to use for areas connected to this interface when authentication is set to "message-digest".

Section 5.12.6.15

OSPF Network Areas

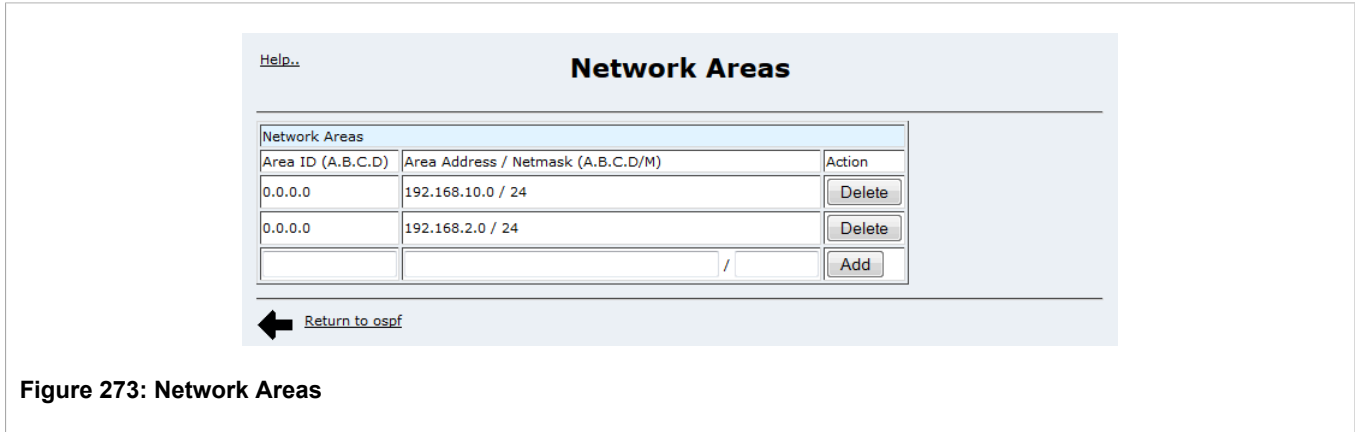


Figure 273: Network Areas

OSPF uses areas to control which routes are distributed between routers. To add a network to an area, enter the area ID and the network address and netmask and click Add. To delete an entry click the Delete button beside the entry. All networks routes that are part of the same area will be distributed to other routers in the same area.

Section 5.12.6.16

OSPF Status

This status menu shows various pieces of information about the current OSPF status. The status of each interface is shown, the current database, the current OSPF neighbors and the current OSPF routing table.

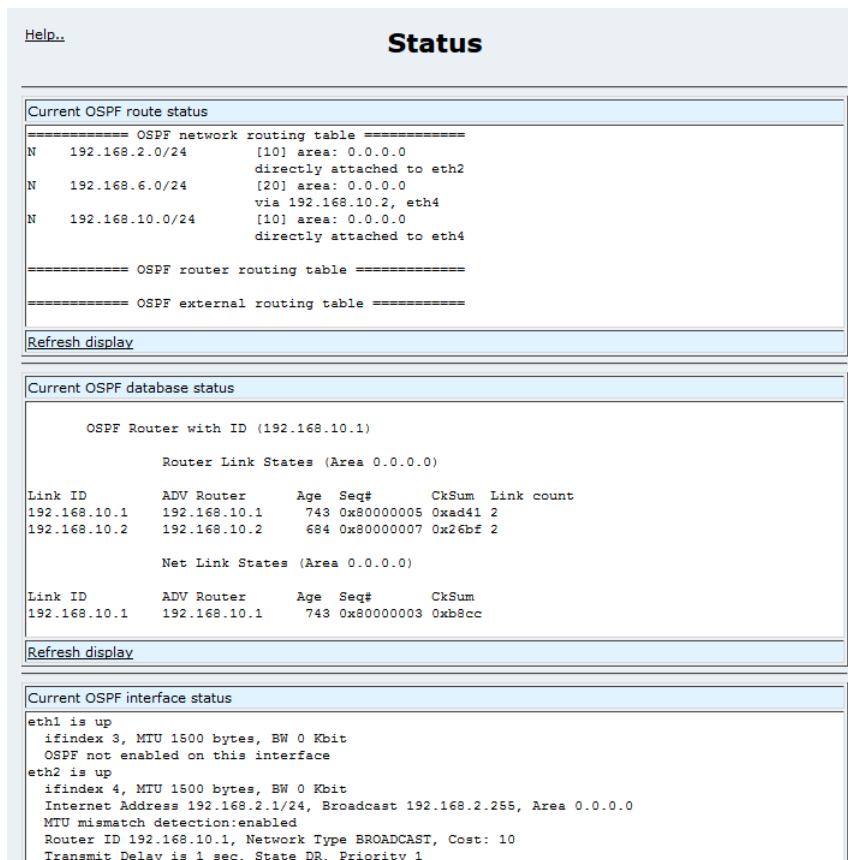


Figure 274: OSPF Status

The OSPF *Status* menu displays:

- A list of routes currently managed by OSPF along with the status of each one
- The status of the OSPF database
- The status of current OSPF interfaces

The tables in the *Current OSPF route status* report provide the following information:

- *N* is the OSPF destination network flag
- *IA* is either the inter-area flag (if displayed) or the intra-area flag (if not displayed)
- *R* is the router flag
- *E1/E2* is the external route type (E1 equals type1, E2 equals type2)
- *ChkSum* is the LS checksum

The tables in the *Current OSPF database status* report provide the following information:

- *Link ID* is the Link State (LS) ID
- *ADV Router* is the IP address for the advertising router
- *Age* is the age of the LS
- *Seq#* is the LS sequence number
- *ChkSum* is the LS checksum

- *Link Count* is the number of links
- *Route* is the network subnet

Section 5.12.6.17

View OSPF Configuration

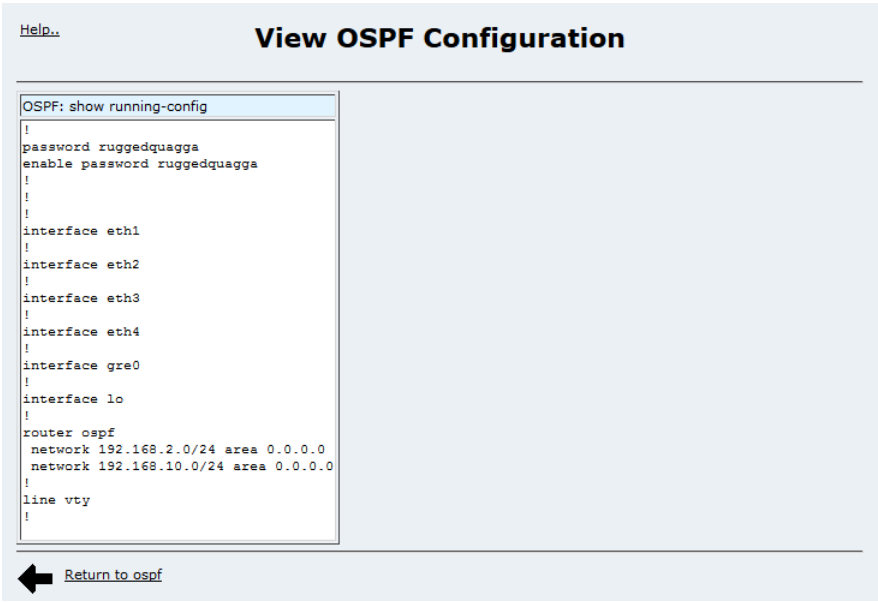


Figure 275: View OSPF Configuration

This menu displays the text of the active configuration file for the OSPF daemon.

Section 5.12.6.18

RIP

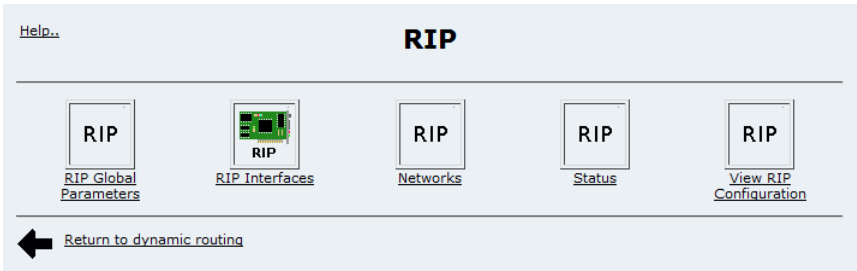


Figure 276: RIP Menu

This menu contains the configuration and status of RIP on the router.

The *RIP Global Parameters* and *RIP Interfaces* configure RIP. The *Status* and *View RIP Configuration* menu display the actual status and configuration file contents of RIP.

Section 5.12.6.19

RIP Global Parameters

[Help...](#)

RIP Global Parameters

Parameter	Value	Description [Possible values] (default value)
Enable Password	••••••••	Enable password. For configuration access. [string without spaces] (previous password)
Telnet Password	••••••••	Telnet password. For port 2602 access. [string without spaces] (previous password)
Hostname		Identifier of router. Often the DNS name of the router. [string without spaces] (no hostname)
Default-Information Originate	enable <input type="checkbox"/>	Advertise default route (disabled)
Default Metric	1	Control distribution of default information [1-16] (1)
Distance		Define an administrative distance [unset,1-255] (unset=not used)
Redistribute Connected	enable <input type="checkbox"/> metric	Redistribute routes for directly connected interfaces to RIP area routers. [enable/disable,0-16] (disabled,unset)
Redistribute Static	enable <input type="checkbox"/> metric	Redistribute static routes to RIP area routers. [enable/disable,0-16] (disabled,unset)
Redistribute Kernel	enable <input type="checkbox"/> metric	Redistribute kernel routes to RIP area routers. [enable/disable,0-16] (disabled,unset)
Redistribute OSPF	enable <input type="checkbox"/> metric	Redistribute ospf routes to RIP area routers. [enable/disable,0-16] (disabled,unset)
Redistribute BGP	enable <input type="checkbox"/> metric	Redistribute bgp routes to RIP area routers. [enable/disable,0-16] (disabled,unset)
Passive Default	enable <input type="checkbox"/>	Set new interfaces passive by default (enabled)
Update Timer	30	Routing table update timer [5-2147483647] (30)
Timeout Timer	180	Routing information timeout timer [5-2147483647] (180)
Garbage Collection Timer	120	Garbage collection timer [5-2147483647] (120)
Send Version	2	RIP version to transmit to neighbors [1, 2] (2)

Save

Key Chains

Key Chain Name	Action
	Add

[Return to rip](#)

Figure 277: RIP Global Parameters

The *Enable Password* field sets the password to be used for the enable command of ripd. This is used by the telnet interface of ripd to control access to the configuration.

The *Telnet Password* field sets the password to be used for telnet access to ripd. This is used as the login password of ripd when locally telnetting to port 2602 of the router.

The *Hostname* field sets the hostname for the rip daemon. This value is only used as a reference for convenience. The telnet interface prompt will contain this hostname. The router's system wide hostname is used if this field is left blank.

The *Default-Information Originate* field, when enabled, causes the router to advertise its default route to the RIP network.

The *Default Metric* field sets the default metric to be used for RIP routes which don't have another metric specified.

The *Distance* field sets the administrative distance to use for all routes unless overridden by other distance settings.

The *Redistribute Connected* fields control distribution of connected routes. When enabled, RIP will advertise routes to directly connected interfaces to other RIP routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute Static* fields control the distribution of static routes. When enabled, RIP will advertise static routes created using the Network Configuration/Routing and Default Route menu to other RIP routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute Kernel* fields control distribution of kernel routes. When enabled, RIP will advertise routes from the kernel routing table, which includes static routes entered by the administrator, to other RIP routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The *Redistribute OSPF* fields control distribution of routes learned by OSPF. When enabled, RIP will advertise routes learned by OSPF.

The *Redistribute BGP* fields control distribution of routes learned by BGP. When enabled, RIP will advertise routes learned by BGP.

The *Passive Default* option controls the default active/passive state of new interfaces. When enabled all new interfaces will be passive by default. The passive state of individual interfaces is controlled from the RIP Interfaces configuration.

The *Update timer* field controls how often RIP sends out routing table updates.

The *Timeout Timer* field controls how long information stays in the routing table after it is received without an update.

The *Garbage Collection Timer* field controls how long expired entries are remembered before being purged.

Section 5.12.6.20

RIP Key Chains

The Key Chains table configures authentication keys used on the interfaces. By defining the keys in a key chain, the same settings can be applied to multiple groups of interfaces. Without key chains the same settings would have to be entered for each interface separately.

Key chains also allow multiple keys to be entered in a single key chain with a start time for when that key should become valid as well as the duration the key is valid. This allows multiple keys to be set up with automatic transitions from one key to the next over time.

A key consists of a key string, which is the value used for authentication. It also has the optional lifetime to accept RIP messages with the key, and the optional lifetime to send RIP messages with that key.

Section 5.12.6.21

RIP Interfaces

[Help..](#)

RIP Interface Configuration - eth1

Parameter	Value (blank = default)	Description [Possible values] (default value)
Passive Interface	passive <input type="checkbox"/>	Control interface passive setting (not passive)
Receive Version	1 2	RIP version to accept from other routers [1, 2, 1 2 (both)] (1 2)
Send Version	2	RIP version to transmit to other routers [1, 2, 1 2 (both)] (2)
Authentication	<input checked="" type="radio"/> None <input type="radio"/> String <input type="text"/> <input type="radio"/> Key Chain <input type="text"/>	Authentication to use [None, Specified string, Specified key chain] (None)
Authentication Mode	Text	Mode of authentication to use [Plain text, MD5 RFC compliant, MD5 old ripd compatible] (Text)
Use Split Horizon	Yes	Use a split horizon [No, Yes, Yes with poisoned reverse] (No)


 [Return to interfaces](#)

Figure 278: RIP Interfaces

Parameters specific to one interface are configured here.

Each interface on the router is listed. Clicking on settings displays a menu of configuration options for that interface.

Clicking "Remove inactive interfaces" purges the list of any interfaces which are no longer configured on the router.

The *Passive Interface* option controls if an interface is active or passive. Passive interfaces do not send RIP updates to other routers.

The *Receive Version* field controls which versions of RIP messages will be accepted from. Either version 1, 2 or both versions can be accepted. By default both RIP versions are accepted.

The *Send Version* field controls which versions of RIP messages to send to other routers. Either version 1, 2 or both versions can be sent. By default only RIP version 2 messages are sent.

The *Authentication* fields choose the authentication mode this port uses. A port can either use no authentication, use a specific authentication string (used the same was as the string in a key), or use a specific key chain's settings. By default no authentication is used.

The *Authentication mode* field chooses the mode of authentication used. Options are plain text (the default), MD5 following the RIP authentication RFC, and MD5 using the method used by the old ripd implementation.

The *Use Split Horizon* field controls use of the RIP split-horizon feature (RIP v2 only). It can be disabled or enabled, and if enabled it can optionally enable the poisoned reverse feature. Split horizon controls whether routes learned through an interface should be allowed to be advertised back out that interface. By default RIP advertises all routes it knows about to everyone, which makes it take a very long time for dropped links to age out of the network. The split horizon prevents advertising those routes back out the same interface which helps to control this problem. Some network topologies with rings of routers will still have some issues with aging out dead routes even with split horizon enabled but they will still age out faster. If fast network recovery is desired, use OSPF.

Section 5.12.6.22

RIP Networks

[Help..](#)

Networks

Neighbors

Neighbor	Action
192.168.10.2	<button>Delete</button>
<input type="text"/>	<button>Add</button>

Networks

Subnet (x.x.x.x/x) or Interface	Action
192.168.10.0/24	<button>Delete</button>
192.168.2.0/24	<button>Delete</button>
<input type="text"/>	<button>Add</button>
eth1 <input type="text"/>	<button>Add</button>

[Return to rip](#)

Figure 279: RIP Networks

Neighbors are specific routers with which to exchange routes using the RIP protocol. This can be used when you want to explicitly control which routers are part of your RIP network.

Networks are used when you want to add any router that is part of a specific subnet, or connected to a specific network interface to be part of your RIP network.

Both neighbors and networks can be used at the same time.



NOTE
For point to point links (T1/E1 links for example) one must use neighbor entries to add other routers to exchange routes with. Also note that RIP v1 does not send subnet mask information in its updates. Any defined networks are restricted to the classic (in the sense of Class A, B and C) networks. RIP v2 does not have this failing.

Section 5.12.6.23

RIP Status

This status menu shows various pieces of information about the current RIP status. The status of each interface is shown, the current database, the current RIP neighbors and the current RIP routing table.

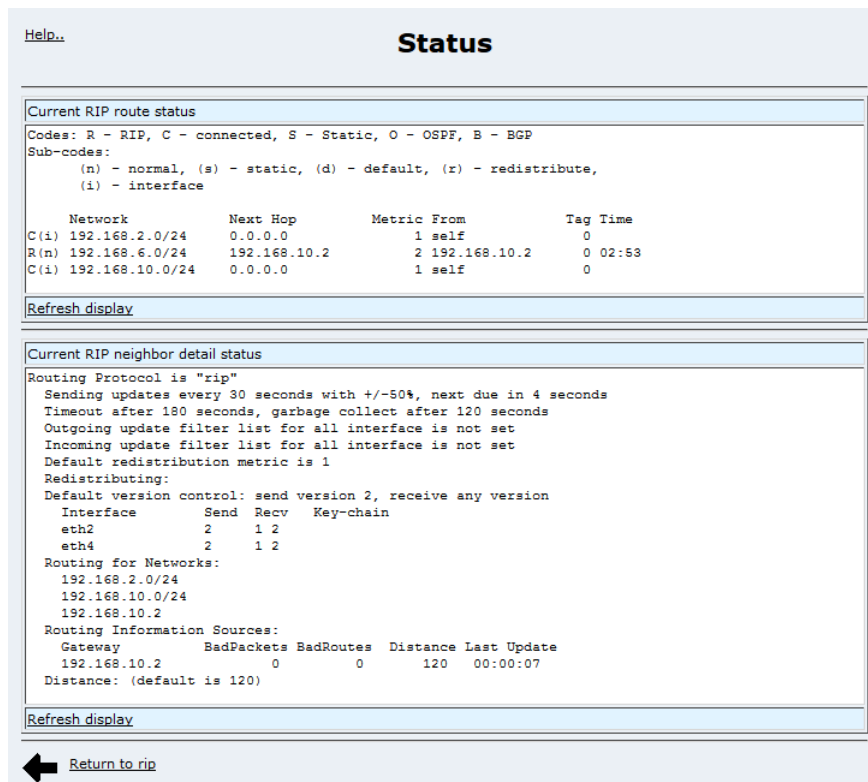


Figure 280: RIP Status

The RIP *Status* menu displays:

- A list of routes currently managed by RIP along with the status of each one
- A list of RIP neighbors along with statistics and state information for each one
- Other statistics, including known peers and memory usage

The table in the *Current RIP route status* report provides the following information:

- *Network* is the IP address for the network
- *Next Hop* is the next hop IP address
- *Metric* is the metric value
- *From* is the source IP address
- *Tag* is the tag information for the route
- *Time* is the route up remaining time

The tables in the *Current RIP neighbour detail status* report provides the following information:

- *Interface* is the interface name
- *Send* is the RIP send version
- *Recv* is the RIP receive version
- *Key-chain* is the RIP authentication key chain
- *Gateway* is the gateway address
- *BadPackets* is the number of bad packets received

- *BadRoutes* is the number of bad routes received
- *Distance* is the distance
- *Last Update* is the peer uptime since the last update

Section 5.12.6.24

View RIP Configuration

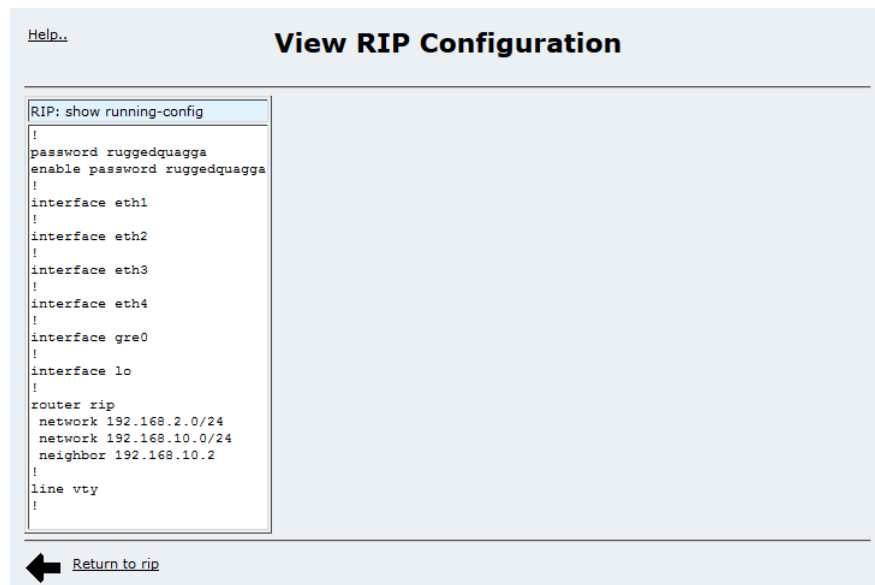


Figure 281: View RIP Configuration

This menu displays the text of the active configuration file for the RIP daemon.

Section 5.13

Link Backup

This section familiarizes the user with:

- Configuring link backup
- Obtaining system status
- Testing link backup

Link backup provides an easily configured means of raising a backup link upon the failure of a designated main link. The main and backup links can be Ethernet, CDMA or Dial Modem, TE1, DDS, ADSL or T3. The only requirement is that the main link be a "permanent" link raised at boot time.

The feature can back up to multiple remote locations, managing multiple main:backup link relationships. When the backup link is a modem, many "profiles" of dialed numbers can exist (each serving as a distinct backup link).

The feature can back up a permanent high speed WAN link to a permanent low speed WAN link. This is used when OSPF cannot be employed, such as on public links.

The feature can be used to migrate the default route from the main to the backup link.

The time after a main link failure to backup link startup and the time after a main link recovery to backup link stop are configurable.

The status of the system and a method of testing fail over is provided.

Section 5.13.1

Path Failure Discovery

In order to discover the failure of a primary path (here, through Network A) the link backup daemon will both inspect the link status of the main link and send a regular ping to a designated host. In this way, failures of network links within the cloud are discovered. It is essential that the host always respond to the ping. Another option is to configure a dummy address within the router and ping that address.




Figure 282: Link Backup Example

The daemon will construe the main link as having failed (even if its link status is "up") if the remote host fails to respond to configurable number of pings after waiting a configurable timeout for each ping.

Section 5.13.2

Use of Routing Protocols and the Default Route

If the main trunk is on a private network, employ a routing protocol to ensure that an alternate route to end network is learned after the backup trunk is raised. Ensure that OSPF/RIP are configured to operate on the secondary trunk, assigning it a higher metric cost than that of the main trunk.

If the main trunk is on a public network, employ the "transfer default route" feature.

Section 5.13.3

Link Backup Configuration




Figure 283: Link Backup Main Menu

Note that Link backup is disabled by default and may be enabled via the System folder, Bootup and Shutdown menu.

Link backup can be configured through the *Link Backup Configuration* link.

Link backup status and logs can be viewed through the *Link Backup Status* and the *Link Backup Log* link after the daemon has been started. A link backup configuration can be tested through the *Link Backup Test* link.

Section 5.13.4

Link Backup Configurations

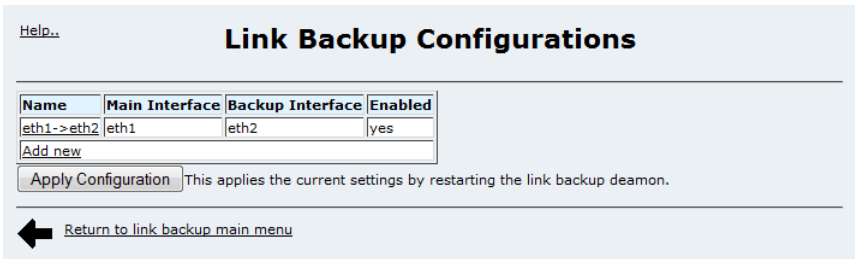


Figure 284: Link Backup Configurations

This menu displays existing main:backup link relationships. Following the links under the *Name* field to an existing pair will edit them or adds a new one.

The *Apply Configuration* button will apply changes by restarting the link backup daemon.

Section 5.13.5

Edit Link Backup Configuration

Edit Link Backup Configuration

Configure eth1 to eth2 link backup

Name

☒ Enable this configuration

☐ Transfer default gateway

Backup gateway

☐ Bring up backup link on demand

Main ping test target

Ping Interval

Ping timeout seconds

Ping retry count

Startup delay seconds

Main path down timeout seconds

Main path up timeout seconds

[Return to link backup configurations](#)

Figure 285: Edit Link Backup Configuration

Set the *Name* field to supply an identification of the pair. This field initially defaults to the "main_link_name>backup_link_name".

The *Enable this configuration* field enables this backup.

The *Transfer default gateway* field causes the gateway to be transferred to the backup link upon failure of the main link path. If the backup interface is point to point, such as PPP, the *Backup gateway* IP address can be automatically determined. Non-point to point interfaces such as Ethernet must be configured with one. The *Bring up backup link on demand* option allows protocols such as DHCP to be used to fetch an address when required.

The *Startup Delay* field configures the length of time to wait for the main link to come up at the start of day.



NOTE

*If **Startup Delay** is too low, backup may be falsely triggered at start up.*

The *Ping Interval* field configures how often pings are sent.

The *Ping timeout* field configures the duration before immediately retrying a ping.

The *Ping retry count* field configures the number of ping retries before construing a path failure.



NOTE

*The maximum time to discover a path failure is the length of the **Ping Interval** and the product of the **Number of missed pings before fail over** and the **Ping timeout**.*

The *Main path down timeout* field specifies the number of seconds the main trunk must be down before starting the backup trunk.

The *Main path up timeout* field specifies the number of seconds the main trunk must have returned to service before stopping the backup trunk.

You may delete a link backup configuration through the *Delete* button.



NOTE
If you delete a link backup configuration that has failed over (or is failing over) to its backup trunk, the link daemon will stop attempting the link backup and restore the main trunk, even if the main trunk is still down.



NOTE
When using Ethernet as on-demand backup interface, first disable the backup interface in the main Networking menu. When an interface is disabled and configured as an on-demand interface, it remains down until it is brought up when needed by link failover.

Section 5.13.6

Link Backup Logs

Help...

Link Backup Logs

Refresh

Month	Day	Time	Process	Event
/var/log/syslog:Oct	28	15:57:23	linkd[4664]	linkd initializing.
/var/log/syslog:Oct	28	15:57:23	linkd[4664]	linkd configured and started.
/var/log/syslog:Oct	28	15:57:23	webmin[4601]	root [Bootup and Shutdown] Started actions linkd

Refresh

[Return to link backup main menu](#)

Figure 286: Link Backup Log

The link backup log displays the log of recent backup events.

Section 5.13.7

Link Backup Status

Help...

Link Backup Status

Name	Main Interface		Backup Interface		Main Ping Test	Time of Last State Change	Link Backup State
	Device	Link State	Device	Link State			
eth1->eth2	eth1	N/A	eth2	N/A	N/A	2013-10-16 12:16:19	Delaying monitor start

Refresh

[Return to link backup main menu](#)

Figure 287: Link Backup Status

The link backup status menu displays the status of managed link backup sets.

Section 5.13.8

Testing a Link Backup Configuration

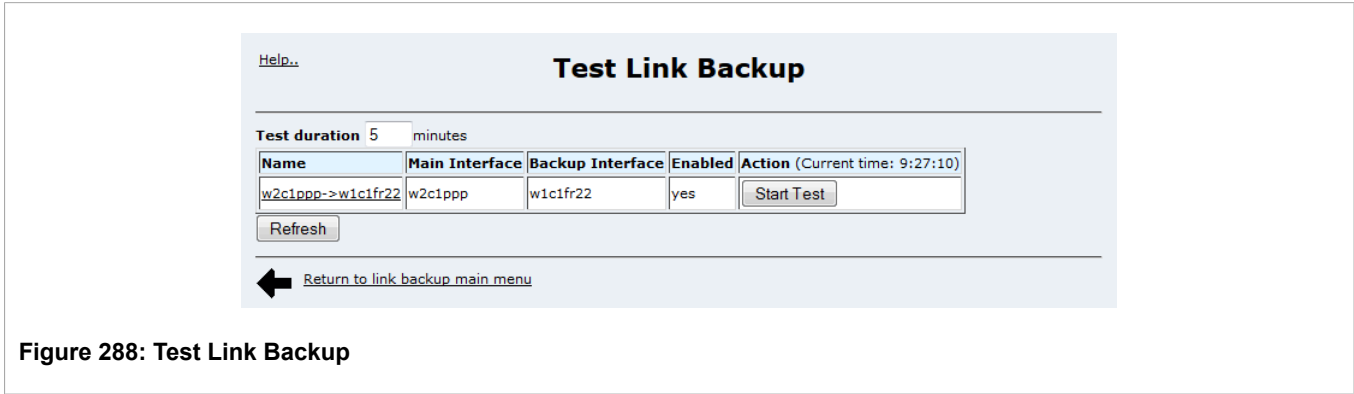


Figure 288: Test Link Backup

Clicking the *Start Test* button corresponding to a particular link backup configuration causes all data received on the "Main" interface to be discarded, in order to convince the link backup process that the main link is down. If the configuration is correct, the link backup process will then bring up the "Backup" link and the connection to the network thus protected will be re-established.

The *Test Duration* field controls the amount of time to run before restoring service to the main trunk. Please note that this duration must take into account the timing parameters of the backup configuration: The duration should comfortably exceed the *Ping Interval* plus the *Ping Timeout* multiplied by the *Ping retry count* plus the *Main path down timeout*. In the case of a dial backup configuration, also be sure to take into account the call setup and modem connection times. Add to this a time that will allow time to navigate the Webmin menus to observe that Link Backup status, link states, and routing are all as expected before, during, and after the Link Backup test.

Section 5.13.9

Scheduled Link Backup Test

The *blinktest* command-line utility provides a mechanism for scheduling tests of link backup configurations. Its function is similar to that of the Webmin-based link backup test, described above. It differs by being accessible from the command line, and by allowing a test to be scheduled to start and stop at some time in the following 24 hours.

The syntax of the *blinktest* command is as follows:

`blinktest -i <interface> -b <begin> -s <stop> <command>`
where:

- 'interface' is the main interface as configured for the link backup.
- 'begin' is the time the test is scheduled to begin, in 24-hour HH:MM format.
- 'stop' is the time the test is scheduled to stop, in 24-hour HH:MM format.
- 'command' is either 'start', in which case -i, -b and -s must be specified, or 'cancel', which only requires -i to be specified.

Section 5.13.9.1

Examples

In order to schedule a link backup test on eth1 at 10:30AM that will last for 10 minutes:

```
blinktest -i eth1 -b 10:30 -s 10:40 start
```

In order to cancel this test, or to terminate it before the 'stop' time:

```
blinktest -i eth1 cancel
```

Section 5.13.9.2

Logging Output

blinktest logs activity and any error conditions to the system log file: `/var/log/messages` and to the serial console. The following items are logged:

- scheduled begin and stop times of a test and the main interface name of the link backup configuration under test
- actual 'begin' time of a test
- actual 'stop' time of a test
- cancellation of a test
- any errors

Log messages have the following format:

- date and time
- router host name
- application name (blinktest)
- blinktest version
- status (OK or ERROR)
- message string

Some example log messages follow:

```
Mar 29 14:57:43 brouter blinktest 1.0: OK link eth3 backup test  STOP scheduled for: Mar 29
14:59:00 2010
Mar 29 14:57:43 brouter blinktest 1.0: OK link eth3 backup test BEGIN scheduled for: Mar 29
14:58:00 2010
Please note that the scheduled date/time shown in the above log messages is the actual
scheduled date/time as reported by the 'at' utility.
When the tests actually begins or stop the following is shown:
Mar 29 12:40:00 brouter blinktest: backup link test BEGIN
Mar 29 12:41:00 brouter blinktest: backup link test STOP
```

Some example error messages:

```
Mar 27 12:43:11 brouter blinktest 1.0: ERROR Begin: Invalid interface: 0
Mar 27 12:43:45 brouter blinktest 1.0: ERROR Invalid start time:
Mar 29 12:49:09 brouter blinktest 1.0: OK link backup test for eth3 canceled at: Mon Mar
29 12:49:09 EDT 2010
```

Section 5.14

Configuring VRRP

This section familiarizes the user with:

- Configuring VRRP
- Enabling and Starting VRRP
- Obtaining VRRP Status

The Virtual Router Redundancy Protocol (VRRP) eliminates a single point of failure associated with statically routed networks by providing automatic failover using alternate routers. The ROX VRRP daemon (keepalived) is an RFC 2338 version 2 compliant implementation of VRRP.

Section 5.14.1

The Problem with Static Routing

Many network designs employ a statically configured default route in the network hosts. A static default route is simple to configure, requires little if any overhead to run and is supported by virtually every IP implementation. When dynamic host configuration protocol (DHCP) is employed, hosts may accept configuration for only a single default gateway.

Unfortunately, this approach creates a single point of failure. Loss of the router supplying the default route or the router's WAN connection results in isolating the hosts relying upon the default route.

There are a number of ways that may be used to provide redundant connections to the host. Some hosts can configure alternate gateways while others are intelligent enough to participate in dynamic routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First routing protocol (OSPF). Even when available, these approaches are not always practical due to administrative and operation overhead.

Section 5.14.2

The VRRP Solution

VRRP solves the problem by allowing the establishment of a "virtual router group", composed of a number of routers that provide a specific default route. VRRP uses an election protocol to dynamically assign responsibility for the "virtual" router to one of the routers in the group. This router is called the VRRP Master. If the Master (or optionally its WAN connection) fails, the alternate (i.e. backup) routers in the group elect a new Master. The new master provides the virtual IP address and issues a gratuitous ARP to inform the network of where the gateway can be reached.

Because the host's default route does not change and MAC address is updated, packet loss at the hosts is limited to the amount of time required to elect a new router.

Section 5.14.3

VRRP Terminology

Each physical router running VRRP is known as a VRRP Router. Two or more VRRP Routers can be configured to form a "Virtual Router". Each VRRP Router may participate in one or more Virtual Routers.

Each Virtual Router has a user-configured *Virtual Router Identifier* (VRID) and an *Virtual IP address* or set of IP addresses on the shared LAN. Hosts on the shared LAN are configured to use these addresses as the default gateway.

One router in the Virtual Router Group will be elected as the *Master*, all other routers in the group will be *Backups*.

Each router in the group will run at a specific *Priority*. The router with the highest priority is elected Master. The value of Priority varies from 1 to 255.

VRRP can also monitor a specified interface and give up control of a VRIP if that interface goes down.

In the following network, host 1 uses a gateway of 1.1.1.253 and host 2 uses a gateway of 1.1.1.252. The 1.1.1.253 gateway is provided by VRID 10. In normal practice router 1 will provide this virtual IP as its priority for VRID 10 is higher than that of router 2. If router 1 becomes inoperative or if its w1ppp link fails, it will relinquish control of VRIP 1.1.1.253 to router 2.

In a similar fashion host 2 can use the VRID 11 gateway address of 1.1.1.252 which will normally be supplied by router 2.




Figure 289: VRRP Example

1. Network 2. Remote Router 1 3. Remote Router 2 4. Switch 5. Host 1 6. Host 2

In this example traffic from host 1 will be sent through router 1 and traffic from host2 through router 2. A failure of either router (or its wan link) will be recovered by the other router.

Note that both routers can always be reached by the hosts at their "real" IP addresses.

Two or more VRRP instances can be assigned to be in the same *VRRP Group*, in which case, they can fail over together.

In the following network, both host 1 and host 2 use a gateway of 192.168.3.10. The external side can access the internal side by gateway 192.168.2.10. The VRID_20 and VRID_21 are grouped together. Normally the Router 1 will provide both internal and external access gateway as its priority is higher than those on Router 2. When either internal or external side of Router 1 becomes inoperative, it will remove the control of both VRIP 192.168.2.10 and 192.168.3.10 and gives the control to Router 2.




Figure 290: VRRP Group Example

1. Network 2. Remote Router 1 3. Remote Router 2 4. Switch 5. Host 1 6. Host 2

Other VRRP parameters are the *Advertisement Interval* and *Gratuitous ARP Delay*.

The advertisement interval is the time between which advertisements are sent. A backup router will assume mastership *four advertisement intervals* after the master fails, so the minimum fail-over time is four seconds. If a monitored interface goes down, a master router will immediately signal an election and allow a backup router to assume mastership.

The router issues a set of gratuitous ARPs when moving between master and backup state. These unsolicited ARPs teach the hosts and switches in the network of the current MAC address and port associated with the VRIP. The router will issue a second set of ARPs after the time specified by the Gratuitous ARP delay.

Section 5.14.4

VRRP Configuration

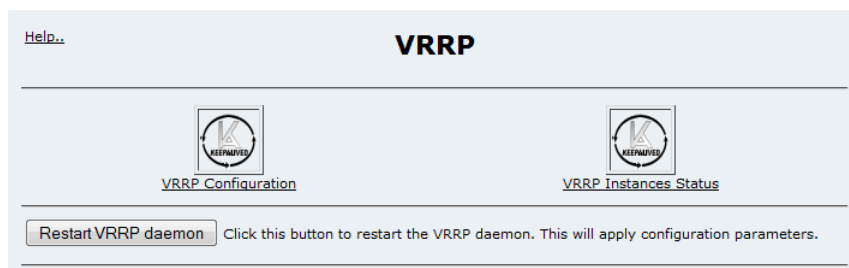


Figure 291: VRRP Main Menu

Note that VRRP is disabled by default and may be enabled via the System folder, Bootup and Shutdown menu.

VRRP can be configured through the *VRRP Configuration* link before the daemon is started.

When enabled, any configuration changes may be made to take effect by selecting the *Restart VRRP daemon* button.

The *VRRP Instances Status* link presents the status VRRP instances existing as of the last restart of keepalived.

Section 5.14.5

VRRP Configuration Menu

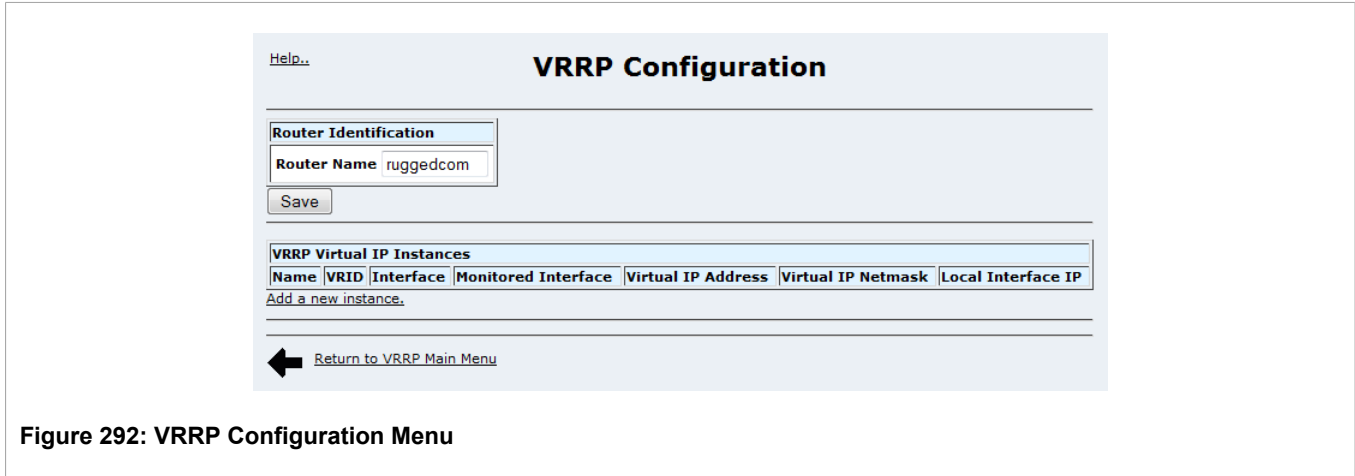


Figure 292: VRRP Configuration Menu

Set the *Router Name* field to supply an identification of the router for VRRP logs. This field initially defaults to the current hostname.

The VRRP instances under the *Name* column define virtual IP groups. Clicking on a link will allow you to edit that instance.

The VRRP groups under the *Group Name* column define virtual IP groups. Clicking on a link will allow you to add members to that group.

Section 5.14.6

Editing a VRRP Instance

Edit VRRP Instance

Virtual IP Instance Parameters

Name	VRID_20	Interface	eth2	Virtual Router ID	20
Priority	100	Advert Interval	1	Gratuitous ARP Delay	5
Extra Interface to Monitor	none	Use Virtual MAC	<input type="checkbox"/>		
Virtual IP Address	192.168.2.10	Virtual IP Netmask	24		

Save Delete

Return to VRRP Configuration

Figure 293: VRRP Instance

The *Name* field is purely for informational purposes.

The *Interface* field configures the interface that VRRP packets are sent upon.

The *Virtual Router ID* field determines the VRID number. Ensure that all routers supplying the same VRIP have the same VRID. The value of the VRID varies from 1 to 255.

The *Advert Interval* field configures the time between VRRP advertisements. Ensure that all routers supplying the same VRID have the same interval.



NOTE

The VRRP advertisement interval must be configured to the same value in the master and backup routers.

The *Gratuitous ARP Delay* field controls the number of seconds after the router changes between master and backup state that a second set of gratuitous ARPs are sent. This mechanism offers a second chance to teach the switching fabric and hosts of the new provider of a gateway address.

The *Extra Interface to Monitor* field causes VRRP to release control of the VRIP if the specified interface stops running. This prevents the situation in which a host forwards information to a gateway router that itself has no way to forward the traffic.



NOTE

*The **Extra Interface to Monitor** field allows monitoring of both logical and physical network interfaces. Examples of a physical interface include:*

- a WAN port, "w1"
- a channelized interface on a WAN port, "w1c1"

Logical interfaces ultimately provide transport for IP on top of physical interfaces. Examples of logical interfaces implemented on top of physical interfaces might include:

- a PPP interface on a channelized WAN port, "w1c1ppp"
- a Frame Relay interface on an unchannelized WAN port, "w1fr16"

Generally, one will need to monitor logical network interfaces, as they participate directly in the IP network. For the purposes of VRRP, one generally wants to monitor status at the highest network layer that is practical (e.g. IP, layer 3).



NOTE

Monitoring a physical interface does not provide an aggregation of the status information of logical interfaces configured on top of it. For example, it could be that all the logical interfaces on a given physical interface are not running, but the physical interface itself is still up and running.

The *Use Virtual MAC* option determines whether or not to use a virtual MAC address for the virtual IP address. By default, ROX does not use a virtual MAC address for VRRP; the virtual IP address is bound to the physical interface on which VRRP works. When *Use Virtual MAC* is selected, ROX creates a virtual interface with the name *VRRP.vrid*, where *vrid* is the virtual router ID set for this VRRP instance. When it becomes the Master, the virtual IP address is bound to this interface and the virtual MAC address will be *00:00:5e:00:01:vrid*, where *vrid* is the hex value for the virtual router ID set for this VRRP instance.



NOTE

When Use Virtual MAC is selected, a virtual interface is created and the virtual IP address will be bound to this interface. If you have the firewall configured, you must set it to allow VRRP messages to be sent and received on this virtual interface.

When Use Virtual MAC is not selected, the virtual IP address is bound to the physical interface. If you have the firewall configured, you must set it to allow VRRP messages to be sent and received on the physical interface.

The *Virtual IP address* field configures the VRIP IP addresses associated with this VRID. Multiple virtual IP addresses (maximum 200) can be configured in this field. Each IP address is separated by a return key.

The *Virtual IP Netmask* field configures the Virtual IP address mask.

The *Save* button saves the virtual instance.

The *Delete* button deletes the virtual instance. After you save or delete an instance you must restart the daemon to action your change.

Section 5.14.7

Editing a VRRP Group

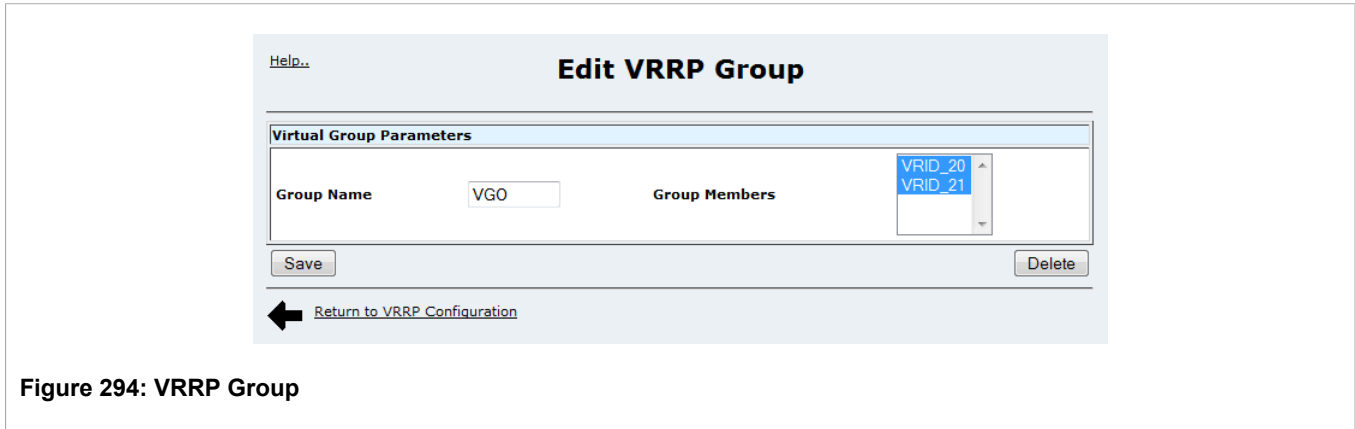


Figure 294: VRRP Group

The *Group Name* field is only for information purpose.

The *Group Members* field determines the group members in this VRRP group. At least two members are needed in order to establish a group.

Section 5.14.8

Viewing VRRP Instances Status

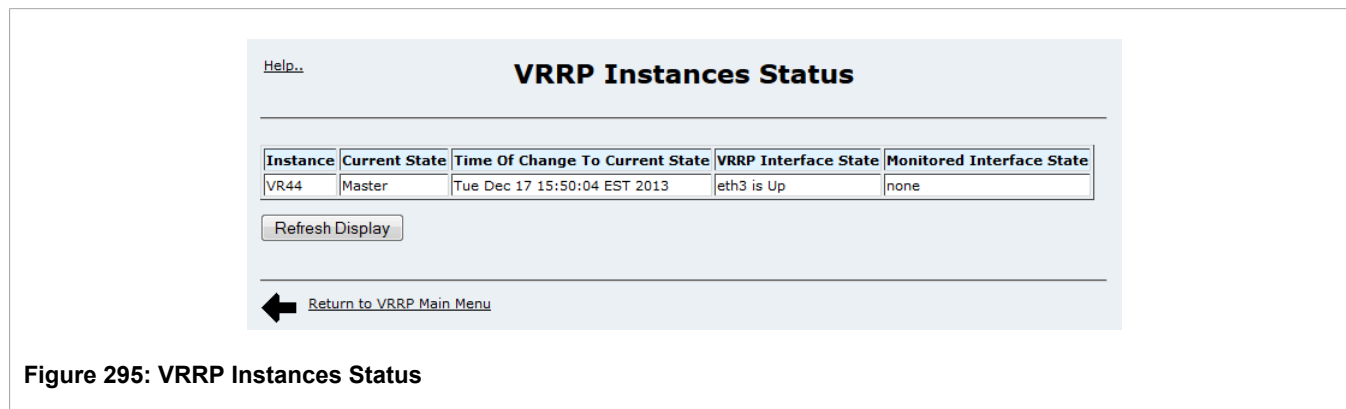


Figure 295: VRRP Instances Status

The VRRP Instances Status menu displays the current status of VRRP instances. This menu does not update status in real time. Click on the *Refresh Display* button to update to the current status.

The entries under the *Instance* column reflect the name of VRRP instances existing as of the last restart of kepalived.

The entries under the *Current State* column reflect the state VRRP instances. An instance can be in one of Master (master for the VRIP), Backup (backup for the VRIP) or Fault (VRRP interface or Monitored interface) is down.

The entries under the *Time Of Change To Current State* column reflect when the current state was entered.

The entries under the *VRRP Interface State* column reflect the link state of the interface that the instance runs upon.

The entries under the *Monitored Interface State* column reflect the link state of the monitored interface or "none" if an interface is not configured.

Section 5.15

Link Layer Discovery Protocol (LLDP)

The IEEE standard, 802.1AB Link Layer Discovery Protocol (LLDP), promises to simplify troubleshooting of enterprise networks and enhance the ability of network management tools to discover and maintain accurate network topologies in multi-vendor environments. LLDP data are made available to NMS (Network Management Systems) via SNMP.

The LLDP service is enabled by default, but may be disabled via [Section 4.3.1, "Bootup and Shutdown"](#).

**CAUTION!**

LLDP is not secure by definition. Avoid enabling LLDP on devices connected to external networks. Siemens recommends using LLDP only in secure environments operating within a security perimeter.

**NOTE**

In order to make LLDP information available via SNMP, the SNMP service must be configured, running, and accessible.

LLDP information can be polled using the standard LLDP-MIB.

LLDP status can also be seen directly using Webmin. Please note that there is a 30 second delay between updates. Pressing the *Refresh* button will query LLDP for current information.

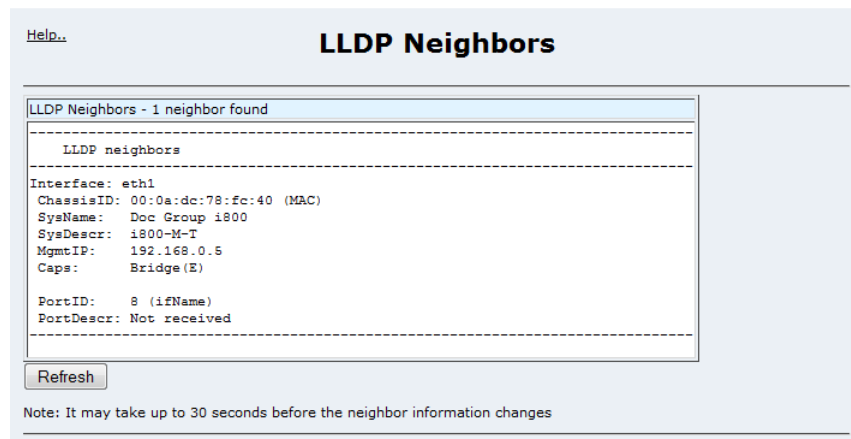


Figure 296: LLDP Summary Display

Section 5.16

Configuring Generic Routing Encapsulation

This section familiarizes the user with:

- Enabling/Disabling GRE
- Viewing GRE Status

ROX is able to encapsulate multicast traffic and IPv6 packets and transport them through an IPv4 network tunnel.

A GRE tunnel can transport traffic through any number of intermediate networks. The key parameters for GRE in each router are the tunnel name, local router address, remote router address and remote subnet.



NOTE

A GRE tunnel is a virtual interface, but it requires a physical network interface to reach the tunnel endpoint. If that physical interface is down, the GRE interface and, subsequently, the routes configured for the interface remain up. The GRE interface does not monitor the status of the network interface to which it sends traffic.




Figure 297: Example – GRE Tunnel Configuration

1. Router 1 2. Router 2

In the above example, Router 1 will use a GRE tunnel with a local router address of 172.16.17.18, a remote router address of 172.19.20.21, and a remote subnet of 192.168.2.0/24.



NOTE

If you are connecting to a CISCO router (in place of Router 1 in the example above), the local router address corresponds to the CISCO IOS "source" address and the remote router address corresponds to the "destination" address.

You may also set a cost for the tunnel. If another method of routing between Router1 and Router2 becomes available, the tunneled packets will flow through the lowest cost route. You can optionally restrict the packets by specifying the local egress device (in the case of router1, w1ppp).

Section 5.16.1

GRE Configuration

This menu displays configured GRE tunnels. The tunnel status will read *active* if the tunnel was created successfully.

Generic Routing Encapsulation Tunnels									
Tunnels									
Tunnel Name	Tunnel IP address	Tunnel MTU	Remote Net	Local IP	Remote IP	Local Egress Port	Multicast	Cost	Tunnel Status
Add a new GRE tunnel..									

Figure 298: GRE Main Menu

Section 5.16.1.1

GRE Configuration Menu

Help...

New Tunnel Configuration

new tunnel

This menu will prefix "gre" to the tunnel name upon saving, legal tunnel names are 12 characters or less in length and contain only a-z or 0-9.

Tunnel Name Tunnel IP address Tunnel MTU

Remote Net Multicast ☐

Local IP Remote IP

Cost Local Egress Port

Save and Apply

Return to Generic Routing Encapsulation Tunnels

Figure 299: GRE Tunnel Configuration Menu

This menu allows you to add or edit a tunnel.

The *Tunnel Name* field will be presented if the tunnel is being created. The tunnel name is purely for informational purposes. A network interface device with this name will be created. In order that the name not collide with those used by other interfaces, it will be prefixed with "gre".

The *Tunnel IP address* field (optionally) configures an IP address on the "gre..." network interface.

The *Remote Net* field configures the target network, at the ingress/egress at the remote end of the tunnel, whose traffic is forwarded through the tunnel. It may be an individual IP address or an IP subnet address, e.g. 192.168.0.0/24. A given Remote Net must not be used by another tunnel.

Setting the *Multicast* option enables multicast traffic on the tunnel interface.

The *Local IP* field configures the IP address of the local end of the tunnel.

The *Remote IP* field configures the IP address of the local remote of the tunnel.



NOTE

Each tunnel must have a unique combination of local and remote addresses, or it will not be activated.

The *Cost* field configures the routing cost associated with networking routing that directs traffic through the tunnel. The cost will default to zero if left unset.

The *Local Egress Port* configures a port to bind the tunnel to. If set, tunneled packets will only be routed via this port and will not be able to escape to another device when the route to the endpoint changes.

The *Tunnel MTU* field allows you to configure the MTU (Maximum Transmission Unit) value. This is useful in preventing the fragmentation of GRE packets.

Section 5.17

Configuring Layer 2 Tunnels

ROX is capable of extending the range of services that communicate solely via Layer 2 protocols (i.e. at the level of Ethernet) by tunneling them over routed IP networks. The Layer 2 Tunnel Daemon supports the IEC61850 GOOSE protocol as well as a generic mechanism for tunneling by Ethernet type.

This section familiarizes the user with:

- Configuring GOOSE tunnels
- Configuring generic Layer 2 tunnels
- Viewing tunnel status and statistics
- Tracing tunnel activity

IEC61850 is an international standard for substation automation. It is a part of the International Electrotechnical Commission's (IEC) Technical Committee 57 (TC57) architecture for electric power systems. An important feature of IEC61850 is the fast transfer of event data. Transfers of Generic Substation Events (GSEs) are accomplished through the GOOSE (Generic Object Oriented Substation Event) protocol.

IEC61850 uses Layer 2 multicast frames to distribute its messages and hence is incapable of operating outside of a switched Ethernet Network. The GOOSE tunnel feature provides a capability to bridge GOOSE frames over a WAN.

GOOSE tunnels provide the following features:

- GOOSE traffic is bridged over the WAN via UDP/IP.
- One GOOSE traffic source can be mapped to multiple remote router Ethernet interfaces in mesh fashion.
- To reduce bandwidth consumption, GOOSE daemons may be located at each of the "legs" and at the center of a star network. The centrally located daemon will accept GOOSE packets and re-distribute them.
- Statistics reports availability of remote GOOSE daemons, packet counts and Round Trip Time (RTT) for each remote daemon.
- When Virtual Router Redundancy Protocol (VRRP) is employed, GOOSE transport is improved by sending redundant GOOSE packets from each VRRP gateway.
- You can enable GOOSE forwarding by configuring a generic Layer 2 tunnel. When configured, ROX listens for GOOSE packets on one VLAN and forwards them to another VLAN.

Section 5.17.1

GOOSE Tunnel Implementation Details

The GOOSE protocol is supported by the Layer 2 Tunnel Daemon. The daemon listens to configured Ethernet interfaces and to the network itself (i.e. for tunnel connections from other daemon instances) on a configurable UDP port.

The Media Access Control (MAC) destination address of frames received from Ethernet is inspected in order to determine which GOOSE group they are in. The frames are then encapsulated in network headers and forwarded (with MAC source and destination addresses intact) to the network as GOOSE packets.

IEC61850 recommends that the MAC destination address should be in the range 01:0c:cd:01:00:00 to 01:0c:cd:01:01:ff.

GOOSE Packets received from the network are stripped of their network headers and forwarded to Ethernet ports configured for the same multicast address. The forwarded frames contain the MAC source address or the originating device, and not that of the transmitting interface. The VLAN used will be that programmed locally for the interface and may differ from the original VLAN. The frame will be transmitted with the highest 802.1p priority level (p4).

Packets received from the network will also be forwarded to any other remote daemons included in the group.

To enable forwarding for GOOSE packets, configure a generic Layer 2 tunnel to listen for GOOSE packets on one VLAN and forward them to a second VLAN. To configure the generic Layer 2 tunnel for this operation, set the following for the tunnel:

- **Ethernet Interface:** select the VLAN on which the GOOSE packets originate.
- **Ethernet Type:** set as 0x88b8. 0x88b8
- **Remote Daemon:** select the VLAN to which to forward the GOOSE packets.

Section 5.17.2

Generic Layer 2 Tunnel Fundamentals

The Layer 2 Tunnel Daemon also supports a generic mode of operation based on the Ethernet type of Layer 2 data traffic seen by the router. Multiple tunnels may be configured, each one with:

- Ethernet type
- Tunnel ingress (Ethernet interface, VLAN interface)
- Tunnel egress (either another locally connected Ethernet interface, VLAN interface, or the remote IP address of another Layer 2 Tunnel daemon instance running on another ROX)

Section 5.17.2.1

Generic Tunnel Implementation Details

For each tunnel configured, the daemon monitors the specified Ethernet or VLAN interface for Ethernet (Layer 2) frames of the specified type. If the configured egress is another local Ethernet port or VLAN interface, frames are simply forwarded on that port, unmodified.

If the configured tunnel egress is a remote IP address, the daemon encapsulates the frames and forwards them to that address, where a corresponding Layer 2 Tunnel Daemon must be configured to receive tunneled frames for local retransmission. Encapsulation headers are stripped in order that the retransmitted frames are identical to those received at the tunnel ingress.

Other notes:

- Source and destination Ethernet MAC addresses are preserved, whether they are forwarded locally or remotely.
- Packets received from the network will also be forwarded to any other remote daemons included in the group.
- The UDP port number for inter-daemon communication must be the same throughout the network
- Enabling Generic L2 Tunneling on an Ethernet interface does not interfere with other (Layer 3) networking configuration on that interface, e.g. firewall rules, IP routing, etc.

**NOTE**

Avoid network configurations where the daemons can form a traffic loop. The simplest such configuration is a triangle network where each daemon forwards to two other routers. Frames arriving at one router will start cycling in clockwise and counter-clockwise directions.

To avoid such "packet storms", frames forwarded to the network are tagged with an initial time to live count. The count is decremented at each relay to the network and prevents the frame from being relayed indefinitely.

Section 5.17.3

Layer 2 Tunnels Main Menu

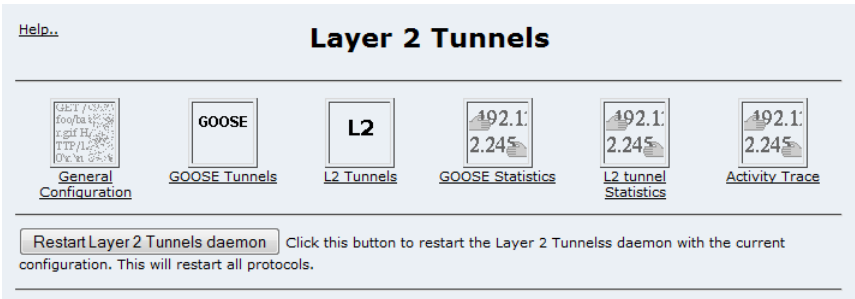


Figure 300: Layer 2 Tunnels Main Menu


Note that the Layer 2 Tunnel daemon is disabled by default and may be enabled via the System folder, Bootup and Shutdown menu.

The *General Configuration* menu changes parameters that apply to all protocols.

The *GOOSE Tunnels* and *GOOSE Statistics* menu configures and display statistics for these tunnels.

The *L2 Tunnels* and *L2 Statistics* menu configures and display statistics for these tunnels.

The *Activity Trace* menu provides a capture and trace facility fdoor.



NOTE

When enabled, any configuration changes may be made to take effect by selecting the *Restart Layer 2 Tunnels daemon* button.

Section 5.17.4

General Configuration Menu

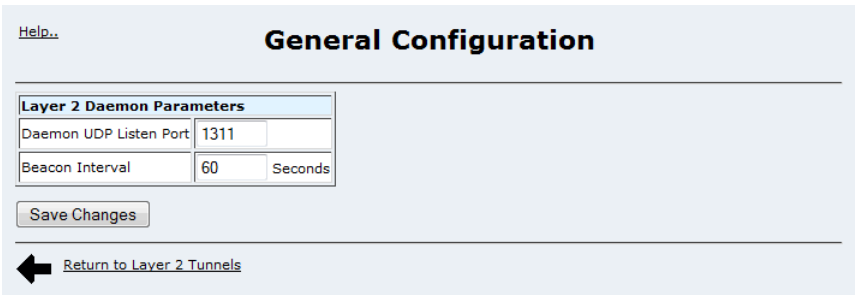


Figure 301: General Configuration Menu

This menu configures general settings for the daemon that apply to all supported tunnel configurations.

The *Daemon UDP Listen Port* field configures port used by the daemon to communicate with other daemons.



NOTE
All Layer 2 Tunnel daemons in the network must use the same port number. If the router employs a firewall, ensure that a hole is opened for each of the remote daemons using this port number.

The Beacon Interval field configures how often a Round Trip Time (RTT) measurement message is sent to each remote peer. The interval takes the values "Off" to disable RTT measurement or a time of 10 – 3600 seconds.

Section 5.17.5

GOOSE Tunnels Menu

Figure 302: GOOSE Tunnels Menu

This menu displays configured GOOSE tunnels. Edit the configuration of an existing tunnel by following the link under the *Ethernet Interface* field, or create a new tunnel by clicking *Add a new GOOSE tunnel*.

Figure 303: Edit GOOSE Tunnel Menu

This menu configures a GOOSE tunnel.

The *Ethernet Interface* field configures suitable (i.e. VLAN eligible) interfaces to listen on for GOOSE frames. You may set this field to "none" if the intent is simply to relay encapsulated traffic between remote tunnel endpoints.

The *Multicast Address* field configures the address to listen for.

The *Remote Daemon* and *Add a new Daemon* fields specify the IP addresses of remote daemons.

Section 5.17.6

Generic L2 Tunnels Menu

The screenshot shows a web interface titled "L2 Tunnels". At the top left is a "Help.." link. Below it is a table with three columns: "Ethernet Interface", "Ethernet Type", and "Remote Tunnel Addresses". Under the "Ethernet Interface" column, there is a link "Add a new L2 tunnel..". At the bottom of the interface is a button with a left-pointing arrow and the text "Return to Layer 2 Tunnels".

Figure 304: Generic L2 Tunnels Menu

This menu displays configured GOOSE tunnels. Edit the configuration of an existing tunnel by following the link under the *Ethernet Interface* field, or create a new tunnel by clicking *Add a new L2 tunnel*.

The screenshot shows a web interface titled "Create L2 Tunnel". At the top left is a "Help.." link. Below it is a form titled "L2 Tunnel". The form contains several fields: "Ethernet Interface" with a dropdown menu set to "None"; "Add Ethernet Type" with a dropdown menu set to "None" and an "Others" field; "Add a new daemon" with a dropdown menu set to "None" and a "Remote IP" field; and a checkbox for "Replace Sender's MAC address". At the bottom of the form are "Save" and "Delete" buttons. Below the form is a button with a left-pointing arrow and the text "Return to L2 Tunnels".

Figure 305: Create an L2 Tunnel

This menu configures a generic layer 2 tunnel.

The *Ethernet Interface* field configures suitable Ethernet interfaces *including VLAN interfaces* to listen on for L2 Ethernet frames of the specified type. You may set this field to "none" if the intent is simply to relay encapsulated traffic between remote tunnel endpoints.

The *Add Ethernet Type* field configures the predefined Ethernet type code that specifies the layer 2 traffic that is to be tunneled. If the desired Ethernet type is not in the predefined list, select *Other*, and enter the type code in the *Others* field:

The *Others* field configures any Ethernet type code that is not predefined in the list box described above. Note that the Ethernet type code must be in hexadecimal format. (e.g. 0x8037 for Novell Netware IPX, 0x88b8 for GOOSE etc.)

Add a new Daemon fields specify the tunnel egress, which may be a local Ethernet interface or the IP addresses of a remote daemon. Several endpoints may be added with these fields using successive edits of the tunnel configuration.

The *Replace Sender's MAC address* option signals the daemon at the tunnel egress to replace the sender's Ethernet MAC address in tunneled frames with the MAC address of output Ethernet interface. By default, the source MAC address of all tunneled frames remain intact.

Help..

Edit L2 Tunnel

L2 Tunnel

Ethernet Interfaceeth1

Ethernet TypeISO

Add Ethernet TypeNoneOthers

Remote Daemoneth4

Add a new deamonNoneRemote IP

Replace Sender's MAC address

SaveDelete

Return to L2 Tunnels

Figure 306: Edit Generic L2 Tunnel

The menu to edit a Generic L2 Tunnel configuration adds only two fields to the creation menu described above. The existing list of configured *Ethernet Types* for the tunnel and a list of *Remote Daemon* items. Note that in this context a *Remote Daemon* may be either the IP address of a remote daemon or the name of a local Ethernet interface.

Section 5.17.7

GOOSE Statistics Menu

Help..

GOOSE Statistics

RefreshContinuous Display

Ethernet Statistics

Interface	L2 MAC Address	Rx Frames	Tx Frames	Rx Chars	Tx Chars	Errors
eth3.0030	01:0c:cd:01:00:33	167302	0	9368912	0	0

Connection Statistics

Interface	L2 MAC Address	Remote IP	Rx Packets	Tx Packets	Rx Chars	Tx Chars	Errors
eth3.0030	01:0c:cd:01:00:33	192.168.2.20	167302	0	10038120	0	0

Round Trip Times

Remote IP	Transmitted	Received	Minimum RTT	Average RTT	Maimum RTT	Std. Deviation
192.168.2.20	2843	2843	0.051 ms	0.867 ms	27.948 ms	0.768 ms

RefreshContinuous DisplayClear RTT Statistics

Return to Layer 2 Tunnels

Figure 307: GOOSE Statistics Menu

This menu presents statistics of GOOSE activity at the Ethernet and Network Layers.

The *Ethernet Statistics* table provides a record for each GOOSE tunnel. The number of historical received and transmitted characters as well as errors will be displayed.

The *Connection Statistics* table reflects UDP connections. Network and Ethernet connections can be paired by examining the L2 MAC Address field.

**NOTE**

All counts are from the router's perspective. The Rx Packets count reflects packets received from the network, the contents of which are transmitted at the protocol and reflected in the Tx Chars field.

The *Round Trip Times* table reflects the measured RTT to each remote daemon. The minimum, average, maximum and standard deviation of times is presented. Wntries with a large difference between the *Transmitted* and *Received* fields indicate potential problems.

The *Refresh* button will cause the page to be reloaded.

The *Continuous Display* button will cause the browser to continuously reload the page showing the differences in statistics from the last display. *The difference is not a real time rate in bytes or packets per second.*

Section 5.17.8

Generic L2 Tunnel Statistics Menu

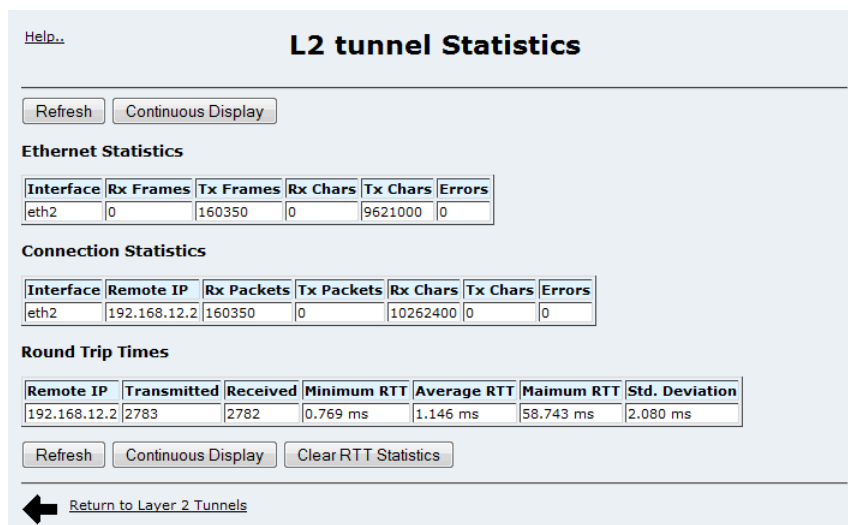


Figure 308: Generic L2 Statistics Menu

This menu presents statistics of tunneled L2 traffic.

The *Ethernet Statistics* table provides a record local Ethernet interface that is part of a tunnel configuration. The number of historical received and transmitted characters as well as errors will be displayed.

The *Connection Statistics* table provides a record for each L2 tunnel that terminates on a local Ethernet interface or on a remote UDP/IP connection.

The *Round Trip Times* table reflects the measured RTT to each remote daemon. The minimum, average, maximum and standard deviation of times is presented. Entries with a large difference between the *Transmitted* and *Received* fields indicate potential problems.

The *Refresh* button will cause the page to be reloaded.

The *Continuous Display* button will cause the browser to continuously reload the page showing the differences in statistics from the last display. Note that the difference is not a real time rate in bytes or packets per second.

Section 5.17.9

Activity Trace Menu

[Help..](#)

Activity Trace

Specifying large numbers of protocols, entries and capture time can result in a greate deal of output..

Trace Layer 2 Tunnels	
Trace on protocols: GOOSE <input type="checkbox"/> L2GEN <input type="checkbox"/> All Protocols <input checked="" type="checkbox"/>	
Message Decode <input checked="" type="checkbox"/> Hex dump <input type="checkbox"/> Packets <input type="checkbox"/> RTT Measurement Messages <input type="checkbox"/>	
Maximum number of entries to capture	20
Maximum time in seconds to capture over	10

```
15:20:27.011 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
15:20:28.012 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
15:20:29.013 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
15:20:30.014 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
15:20:31.017 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
15:20:32.020 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
15:20:33.021 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
15:20:34.033 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
15:20:35.039 L2_GEN Received message from eth3, length 60
DST MAC 00:0a:dc:0f:4f:40 SRC MAC 00:0a:dc:0f:4f:41 APP ID 65278 (0xfefe)
1 second to capture
```

Start Trace

[Return to Layer 2 Tunnels](#)

Figure 309: Activity Trace Menu

This menu displays captured and decoded network activity on configured layer 2 tunnels.


The desired traffic sources, number of messages and length of time to capture are entered and the *Start Trace* button is pressed. The menu will display up to the provided number of messages waiting up to the specified number of seconds.

The *Trace on protocols:* selections feature a (all to short) list of protocols with unused entries greyed out. The default is *All Protocols*.

The *Message Decode* field causes received/transmitted frame entries to include protocol specific information. If the *Hex Dump* field is selected, the first 64 bytes of packet content is displayed.

The *Packets* field causes received/transmitted packet entries to be displayed.

The *RTT Measurement* field displays Beacon messages used for RTT measurement.



NOTE

Specifying large numbers of ports, entries and capture times can result in a great deal of output. Specifying a large capture time may require the web page to wait that interval if activity is infrequent.

Section 5.18

Configuring the DHCP server

This section familiarizes the user with:

- DHCP Server Configuration
- Use of Option 82

Dynamic Host Configuration Protocol (DHCP) is a method for centrally and consistently managing IP addresses and settings for clients, offering a variety of assignment methods. IP addresses can be assigned based on the Ethernet MAC address of a client, sequentially, or by using port identification provided by a DHCP relay agent device.

Section 5.18.1

DHCP Network Organizations

The information to assign addresses in DHCP is organized to deal with clients at the host, group, subnet, pool and shared network level.

Hosts entries assign specific settings to a client based on its Ethernet MAC address.

Groups allow identical settings to be created for a group of hosts, making it simpler to manage changes to the settings for all the hosts contained within the group. Groups contain hosts.

Pools contain ranges of IP addresses to hand out to clients with access rules to determine which clients should receive addresses from that pool.

Subnets control settings for each subnet that DHCP serves. A subnet can include a range of IP address to hand out to clients. Only one subnet can contain dynamic IP address ranges without any access restrictions on any given physical port since DHCP doesn't know which subnet a client should belong to when the request is received. Subnets contain groups, pools and hosts.

Shared networks are used when multiple subnets should be served by a single physical port. This applies both when using a DHCP relay agent connected to the port with additional subnets behind the relay agent, or when multiple virtual networks exist on one physical interface. Each subnet then gets its own subnet definition inside the shared network rather than at the top level. Shared networks contain subnets, groups and hosts.

Section 5.18.2

DHCP Client Options

The following options apply to single hosts, subnets of hosts, pools (potentially discontinuous ranges of addresses), shared networks (a single physical networks for which distinct subnets of hosts coexist and request addresses) and groups. The meaning of each option is the same in each case, while the type of target determines which clients it applies to.

In DHCP, settings at a more specific level overrides higher levels. For example you can configure a DNS server for all clients, then create a group that overrides the setting. This allows defaults to be set at a high level to apply to most clients, while exceptions can be places just where they are needed. Many settings are only supported by certain specific types of clients, and are ignored by the majority of clients.

Basic options you should pay attention to include:

- Address ranges: The range of addresses to use for dynamic IP clients.
- Default lease time: The default length of leases assigned to clients, if the client doesn't request a lease length.
- Maximum lease time: The maximum length of leases allowed to clients. If a client requests a higher value it will be refused.
- Client hostname: The hostname the client should use.
- Default routers: The default gateway the client should use.

- Domain name: The DNS domain name the client should use.
- DNS servers: The IPs of the DNS servers the client should use.
- NTP servers: The IPs of the NTP servers the client should use.
- Static routes: Static routes the client should use.
- Time servers: The IPs of the time servers the client should use.

Lesser used client options include:

- Subnet mask: The subnet mask the client should use. Rarely needed.
- Broadcast address: The broadcast address the client should use. Rarely needed.
- Log servers: The IPs of the LOG servers the client should use.
- Swap server: The IP of the swap server the client should use. Normally only used for diskless network booted clients.
- Root disk path: The path the client should use for its root device. Normally only used for diskless network booted clients.
- NIS domain: The NIS domain the client should use.
- NIS servers: The IPs of the NIS server the client should use.
- Font servers: The IPs of the font servers the client should use. Normally only used for X terminals.
- XDM servers: The IPs of the XDM servers the client should use. Normally only used for X terminals.
- NetBIOS name servers: The IPs of the Netbios name servers the client should use.
- NetBIOS node type: The NetBIOS name resolution method the client should use.
- NetBIOS scope: The NetBIOS scope the client should use.
- Time offset: The offset from a time server the client should be using.
- Custom options allows you to add additional DHCP options required by a client.

BOOTP and Dynamic DNS related options include:

- Boot filename: The filename the client should request from a tftp server to boot from. This only applies to network booted clients.
- Boot file server: The IP address of the tftp server to boot from. This only applies to network booted clients.
- Server name: The hostname of the boot server. This only applies to network booted clients.
- Lease length for BOOTP clients: How long the IP assigned to a BOOTP client should be considered valid.
- Lease end for BOOTP clients: Cut off date for all BOOTP client leases.
- Dynamic DNS enabled: Should DNS information be updated on the DNS server when a client receives an IP address.
- Dynamic DNS domain name: The domain name to update dynamic DNS information in.
- Dynamic DNS hostname: Use the specified hostname for clients, or use the hostname supplied by the client.
- Dynamic DNS reverse domain: The reverser DNS domain to update dynamic information in for the reverse DNS entry.
- Dynamic DNS reverse domain: The reverser DNS domain to update dynamic information in for the reverse DNS entry.

Lesser used DHCP server configurations include

- Allow unknown clients: Should DHCP accept requests from clients it has never seen before or only from clients that have already received leases in the past.

- Server is authoritative: If the server is authoritative, it will send deny messages to any client which tries to renew a lease which the server knows the client shouldn't have.
- Option 82 Support.

Section 5.18.3

Option 82 Support with Disable NAK

If DHCP relay clients (option 82 clients) are used on the same subnet as the DHCP server, some clients will immediately try to renew a lease right after receiving it by requesting a renewal directly from the DHCP server. Since the DHCP server is only configured to provide that lease through a relay agent with the right option 82 fields added, the server will send the client a NAK to disallow use of the lease. Enabling this option disables this reject message, so that the renewal request that the DHCP relay agent sends a moment later (which the DHCP server accepts since it has the right option 82 fields added) will be the only message for which the client receives a reply. If the DHCP server and clients are not on the same subnet, this option is not required. The meaning of the value of many fields depends on the client's interpretation of the field, so the actual meaning of a field is determined by the client. See the documentation of the client to determine what values are required by the client for special options.

Section 5.18.4

Example DHCP Scenarios and Configurations

This section contains the following scenario examples:

- [Section 5.18.4.1, "Single Network With Dynamic IP Assignment"](#)
- [Section 5.18.4.2, "Single Network With Static IP Assignment"](#)
- [Section 5.18.4.3, "Single Network With Option82 Clients On One Switch"](#)
- [Section 5.18.4.4, "Multiple Subnets on Separate VLANs Using Option82 on One Switch"](#)

Section 5.18.4.1

Single Network With Dynamic IP Assignment

In this example the eth1 interface is provided with IP address 192.168.1.1/24 while addresses 192.168.1.101 through 192.168.1.200 are assigned to the clients. The router serves as the default gateway.

1. Enable eth1 in the 'Edit Network Interfaces' menu.
2. Click 'add a subnet', and configure it for network address 192.168.1.0 with netmask 255.255.255.0.
3. Set the assigned address range to 192.168.1.101 - 192.168.1.200.
4. Click 'Create' then edit the subnet just created and click 'Edit Client Options'.
5. Set default routers to 192.168.1.1 and save.
6. Restart the DHCP server or apply changes.

Section 5.18.4.2

Single Network With Static IP Assignment

In this example the eth1 interface is provided with IP address 192.168.1.1/24.

Assign address 192.168.1.101 to a DHCP client with MAC 00:11:22:33:44:01.

Assign address 192.168.1.102 to a DHCP client with MAC 00:11:22:33:44:02.

Assign address 192.168.1.103 to a DHCP client with MAC 00:11:22:33:44:03.

The router serves as the default gateway.

1. Enable eth1 in the 'Edit Network Interfaces' menu.
2. Click 'add a subnet', and configure it for network address 192.168.1.0 with netmask 255.255.255.0.
3. Click 'Create' then edit the subnet just created and click 'Edit Client Options'.
4. Set default routers to 192.168.1.1 and save it.
5. Click 'add a new host'.
6. Set the hardware address to Ethernet 00:11:22:33:44:01 and the fixed IP to 192.168.1.101. Assign the client a hostname as well.
7. Click 'Create'.
8. Repeat steps 5) through 7) for the other hosts with the appropriate address, MAC and hostname for each client.
9. Restart the DHCP server or apply changes.

Section 5.18.4.3

Single Network With Option82 Clients On One Switch

In this example the eth1 interface is provided with IP address 192.168.1.1/24

A switch connected to eth1 and uses address 192.168.1.2/24.

The switch port 1 is connected to the router while its ports 2 through 8 provide DHCP relay support. The switch has its DHCP relay server address set to router's address 192.168.1.1. The switch has all ports in VLAN 1. The switch base MAC address is 00:0A:DC:11:22:00.

Assign a client at switch port 2 address 192.168.1.102.

Assign a client at switch port 3 address 192.168.1.103.

Assign multiple clients at switch port 4 dynamic addresses 192.168.1.151 through 192.168.1.200.

The router serves as the default gateway.

1. Enable eth1 in the 'Edit Network Interfaces' menu.
2. Add a new subnet, and configure it for network address 192.168.1.0 with netmask 255.255.255.0.
3. Enable the 'Disable NAK of option82 clients for this subnet?' option to prevent confusing some DHCP clients due to the client being on the same network as the DHCP server and the DHCP relay agent (the switch).
4. Save it then edit the subnet just created and click 'Edit Client Options'.
5. Set default routers to 192.168.1.1 and save it.
6. Click 'add an address pool' to the subnet.
7. Set the address range to 192.168.1.102 to 192.168.1.102.

8. Click 'Create'.
9. Edit the pool by clicking on the link for the pool with address range 192.168.1.102 - 192.168.1.102.
10. Click 'add an option82 client'.
11. Give the client a unique alpha numeric name (for example client0102).
12. Set the remote id to the switch MAC address (00:0A:DC:11:22:00 in this case).
13. Set the circuit id to the switches circuit id identifier to the port (00:01:00:02 for VLAN 1 port 2 on a SiemensRUGGEDCOM switch).
14. Click 'Create'.
15. Click 'Save'.
16. Repeat steps 6) through 15) for clients 192.168.1.103 changing the pool address range and circuit id.
17. Repeat steps 6) through 15) for port 4 using the address range 192.168.1.151 to 192.168.1.200 and the circuit id for port 4.
18. Restart the DHCP server or apply changes.

Section 5.18.4.4

Multiple Subnets on Separate VLANs Using Option82 on One Switch

In this example the eth1 interface is provided with IP address 192.168.1.1/24

A switch connected to eth1 and using address 192.168.1.2/24

The switch port 1 is connected to the router while its ports 2 through 8 provide DHCP relay support. The switch has its DHCP relay server address set to router's address 192.168.1.1. The switch has all ports in VLAN 1. The switch base MAC address is 00:0A:DC:11:22:00.

The switch port 2 is on vlan2 using subnet 192.168.2.0/24 and should assign addresses 192.168.2.101 to 192.168.2.200 and default gateway 192.168.2.1.

The switch port 3 is on vlan3 using subnet 192.168.3.0/24 and should assign addresses 192.168.3.101 to 192.168.3.200 and default gateway 192.168.3.1.

The switch port 4 is on vlan4 using subnet 192.168.4.0/24 and should assign addresses 192.168.4.101 to 192.168.4.200 and default gateway 192.168.4.1.

1. Enable eth1 in the 'Edit Network Interfaces' menu.
2. Add a new subnet, and configure it for network address 192.168.1.0 with netmask 255.255.255.0.
3. Save it.
4. Add a new shared network.
5. Name the shared network (for example "eth1") and select the subnet 192.168.1.0 to be included in the shared network.
6. Save it.
7. Edit the shared network again.
8. Add a new subnet, and configure it for network address 192.168.2.0 with netmask 255.255.255.0
9. Save the new subnet and then save the shared network settings.
10. Edit the subnet just created and click 'Edit Client Options'.
11. Set default routers to 192.168.2.1 and save it.
12. Click 'add an address pool' to the subnet.

13. Set the address range to 192.168.2.101 to 192.168.2.200.
14. Click 'Create'.
15. Edit the pool by clicking on the link for the pool with address range 192.168.2.101 - 192.168.2.200.
16. Click 'add an option82 client'.
17. Give the client a unique alpha numeric name (for example subnet0102).
18. Set the remote id to the switch MAC address (00:0A:DC:11:22:00 in this case).
19. Set the circuit id to the switches circuit id identifier to the port (00:02:00:02 for VLAN 2 port 2 on a SiemensRUGGEDCOM switch).
20. Click 'Create'.
21. Click 'Save'.
22. Repeat steps 8) through 20) for vlan3 through vlan4 changing the subnet, default routers, pool address range and circuit id for each vlan.
23. Restart the DHCP server or apply changes.

Section 5.18.5

DHCP Configuration

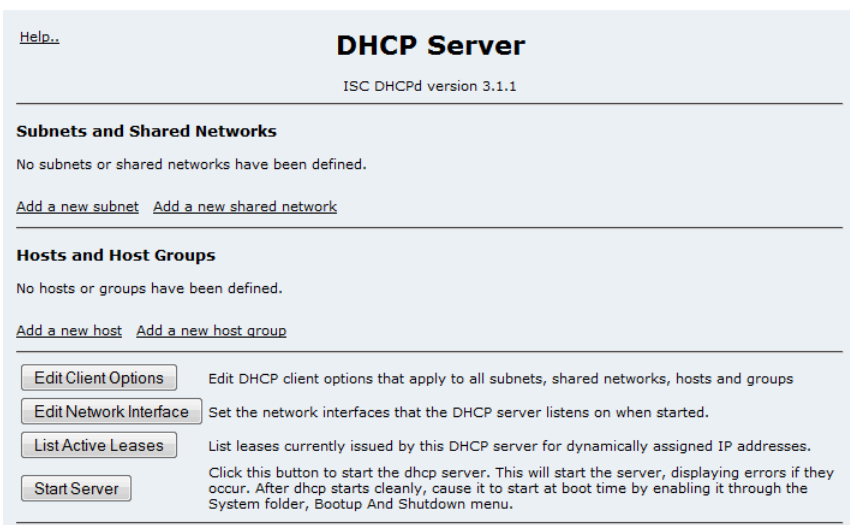


Figure 310: DHCP Server Menu

The DHCP Server main menu shows the subnets configured for DHCP, as well as any groups and hosts. New subnets, groups and hosts can be added, and existing entries can be edited or deleted.

The *Edit Client Options* button allows you to set global client settings for the DHCP server. Settings made here apply to all clients unless overridden at a lower level in the configuration.

The *Edit Network Interface* button allows you to select which interfaces DHCP should listen for DHCP requests on. Note that you must also have a subnet matching the IP address of the selected interface configured in DHCP in order to actually have DHCP listen for requests on a port.

The *List Active Leases* button displays IP leases that are currently assigned to clients from the dynamic IP address pool. Note that static MAC address to IP address assignments handled by DHCP are not displayed in

this list. From this list, there is a button labelled *List all active and expired leases*, which will additionally display leases that have been granted that are no longer active.

The *Start Server* button starts the server to check the configuration. To permanently enable DHCP you should enable it in the bootup and shutdown menu.

The *Apply Changes* button applies new settings to the running DHCP server. Use this after making any changes to the configuration.

Section 5.18.5.1

DHCP Shared Network Configuration



NOTE

The menu interfaces for creating a DHCP Shared Network Configuration and for editing an existing one are the same - only the title differs (*Create* versus *Edit*).

Figure 311: DHCP Shared Network Configuration

The settings specific to the *Create/Edit Shared network* menu are:

- The *Shared network description* field is used to describe the shared network as desired.
- The *Network name* field is a unique name to assign to the shared network. It could be the name of the interface the shared network is on, for example.
- Within a shared network you can create subnets, hosts, and groups of hosts.

Section 5.18.5.2

DHCP Subnet Configuration

**NOTE**

The menu interfaces for creating a DHCP Subnet Configuration and for editing an existing one are the same - only the title differs (Create versus Edit).

[Help...](#)

Edit Subnet

Subnet Details

Subnet description

Local Network

Network address

192.168.2.0

Netmask

255.255.255.0

Address ranges

192.168.2.0 - 192.168.2.200

Shared network

<None>

Boot filename

☒ None ☐

Boot file server

☒ This server ☐

Lease length for BOOTP clients

☒ Forever ☐ secs

Dynamic DNS enabled?

☐ Yes ☐ No ☒ Default

Dynamic DNS reverse domain

☒ Default ☐

Allow unknown clients?

☐ Allow ☐ Deny ☐ Ignore ☒ Default

Server is authoritative for this subnet?

☐ Yes ☒ Default (No)

Hosts directly in this subnet

Default lease time

☒ Default ☐ secs

Maximum lease time

☒ Default ☐ secs

Server name

☒ Default ☐

Lease end for BOOTP clients

☒ Never ☐

Dynamic DNS domain name

☒ Default ☐

Dynamic DNS hostname

☒ From client ☐

Disable NAK of option82 clients for this subnet?

☐ Yes ☒ Default (No)

Groups directly in this subnet

Save

Edit Client Options

List Leases

Delete

Add a new host

Add a new host group

Address Pools for Subnet

Pool	Address Ranges	Option 82 Clients (clientname = remote-id / circuit-id)
Add an address pool		

←

[Return to subnet list](#)

Figure 312: DHCP Subnet Configuration

The settings specific to the *Create/Edit Subnet* menu are:

- The *Subnet description* field is used to describe the subnet as desired.
- The *Network address* and *Netmask* fields define a subnet containing a span of addresses to assign.
- Within a subnet you can create hosts, groups of hosts, and address pools.

Section 5.18.5.3

DHCP Group Configuration



NOTE

The menu interfaces for creating a DHCP Group Configuration and for editing an existing one are the same - only the title differs (Create versus Edit).

Figure 313: DHCP Group Configuration

The settings specific to the *Create/Edit Host Group* menu are:

- The *Group description* field is used to describe the group as desired.
- The *Use name as client hostname* field determines whether host entries should use the hosts entry name as the client hostname to provide to the client.
- Within a group you can create hosts.

Section 5.18.5.4

DHCP Host Configuration



NOTE

The menu interfaces for creating a DHCP Host Configuration and for editing an existing one are the same - only the title differs (Create versus Edit).

[Help..](#)

Create Host

Host Details

Host description

Host name

Hardware Address

ethernet

Fixed IP address

Boot filename

☒ None

☐

Boot file server

☒ This server

☐

Lease length for BOOTP clients

☒ Forever

☐ secs

Dynamic DNS enabled?

☐ Yes

☐ No

☒ Default

Dynamic DNS reverse domain

☒ Default

☐

Allow unknown clients?

☐ Allow

☐ Deny

☐ Ignore

☒ Default

Host assigned to

Toplevel

Default lease time

☒ Default

☐ secs

Maximum lease time

☒ Default

☐ secs

Server name

☒ Default

☐

Lease end for BOOTP clients

☒ Never

☐

Dynamic DNS domain name

☒ Default

☐

Dynamic DNS hostname

☒ From client

☐

Create

Return to host list

Figure 314: DHCP Host Configuration

The *Host description* field is used to describe the host as desired.

The *Host name* field is the unique name to refer to the host within the DHCP configuration.

The *Hardware address* field is the Ethernet MAC of the client associated with the host entry.

The *Fixed IP address* field is the IP to assign to the matching client.

Section 5.18.5.5

DHCP Pool Configuration



NOTE
The menu interfaces for creating a DHCP Address Pool Configuration and for editing an existing one are the same - only the title differs (Create versus Edit).

[Help..](#)

Edit Address Pool

In subnet 192.168.2.0/255.255.255.0

Address pool options

Address ranges 192.168.2.11 - 192.168.2.11 ☐ Dynamic BOOTP ?
☐ Dynamic BOOTP ?

Failover Peer ☒ None ☐

Clients to allow **Clients to deny**

Default lease time ☒ Default ☐ secs

Boot filename ☒ None ☐ **Maximum lease time** ☒ Default ☐ secs

Boot file server ☒ This server ☐ **Server name** ☒ Default ☐

Lease length for BOOTP clients ☒ Forever ☐ secs **Lease end for BOOTP clients** ☒ Never ☐

Dynamic DNS enabled? ☐ Yes ☐ No ☒ Default **Dynamic DNS domain name** ☒ Default ☐

Dynamic DNS reverse domain ☒ Default ☐ **Dynamic DNS hostname** ☒ From client ☐

Allow unknown clients? ☐ Allow ☐ Deny ☐ Ignore ☒ Default

Option 82 clients

Client Name	Remote ID	Circuit ID
Add an option82 client		

[Return to subnet](#)

Figure 315: DHCP Pool Configuration

The settings specific to the *Create/Edit Address Pool* menu are:

- The *Failover peer* field is the IP address of a DHCP peer server if a fail over pool is created.
- The *Clients to allow/deny* field can be used to control which clients can get IP address from the pool. See documentation for `dhcpd3` for syntax and allowed values. Very rarely needed. The *Allow unknown clients* setting already handles the most common use of this option.

Section 5.19

DHCP Relay

This section familiarizes the user with the use and configuration of the device's DHCP Relay feature.

The device can be configured to act as a DHCP Relay Agent. A DHCP Relay Agent forwards DHCP and BOOTP requests from clients on one layer 2 network to one or more configured DHCP servers on other networks. This allows one to implement some measure of isolation between DHCP clients and servers.

The DHCP Relay Agent is configured to listen for DHCP and BOOTP requests on particular Ethernet and VLAN network interfaces, and to relay to a list of one or more DHCP servers. When a request is received from client, it forwards the request to each of the configured DHCP servers. When a reply is received from a server, it forwards the reply back to the originating client.

**NOTE**

The current release of DHCP Relay Agent on ROX only supports operation on Ethernet interfaces. It also does not support the "Circuit ID" and "Remote ID" Information sub-options.

**NOTE**

While DHCP Relay and DHCP Server may both be configured to run concurrently, they may not be configured to run on the same network interface.

Section 5.19.1

Configuring DHCP Relay

The DHCP Relay is disabled by default and may be enabled via the *Bootup and Shutdown* menu under the *System* folder.

Help...

DHCP Relay Configuration

DHCP Relay Configuration

Servers 192.168.54.3

Interfaces eth1, eth2, eth3, eth4

Save

Apply Configuration Click this button to apply configuration. This will restart the DHCP Relay Agent

Figure 316: DHCP Relay Configuration

This menu allows you to configure DHCP Relay Agent.

The *Servers* field configures the list of DHCP servers to which DHCP/BOOTP requests will be forwarded.

**NOTE**

In general, the DHCP servers configured here will themselves need to be configured to serve the subnets from which DHCP/BOOTP client requests will be forwarded. Refer to [Section 5.18.1, "DHCP Network Organizations"](#) for a brief discussion of the DHCP server configuration required to support Relay Agents.

The *Interfaces* field selects the network interfaces on which the relay agent will listen for DHCP/BOOTP requests. Select both the network interface to which clients and servers are attached.

**NOTE**

ROX allows the use of DHCP relay over GRE tunnels and PPP/Frame Relay interfaces for cases in which the DHCP server resides behind the far endpoint of the GRE tunnel or WAN link. In case of GRE, it is mandatory that the GRE interface have the 'Tunnel IP address' parameter configured. Not setting this parameter will result in a functional GRE configuration, but DHCP relay will not work over the tunnel.

The *Save* button will save the configuration permanently.

The *Apply Configuration* button will restart the DHCP Relay Agent with the saved configuration.

Section 5.20

Configuring NTP Servers

This section familiarizes the user with:

- Enabling/Disabling NTP
- Setting servers and peers
- Setting generic NTP options
- NTP Tools

NTP (Network Time Protocol) is an Internet protocol used to synchronize the clocks of computers to some time reference. Variants of NTP such as SNTP (Simple NTP, a reduced functionality NTP) and XNTP (Experimental NTP) exist. NTP itself is available in versions 3 and 4 (ROX includes version 4).

NTP is a fault-tolerant protocol that allows an NTP daemon program to automatically select the best of several available time sources, or reference clocks, to synchronize to. Multiple candidates can be combined to minimize the accumulated error. Temporarily or permanently wrong time sources are detected and avoided.

The NTP daemon achieves synchronization by making small and frequent changes to the router hardware clock.

The NTP daemon operates in a *client-server mode*, both synchronizing from *servers* and providing synchronization to *peers*.

If NTP has a number of servers to choose from, it will synchronize with the lowest stratum server. The stratum is a measure of the number of servers to the (most highly accurate) reference clock. A reference clock itself appears at stratum 0. A server synchronized to a stratum n server will be running at stratum $n + 1$.

You will generally configure lower stratum NTP hosts as servers and other NTP hosts at the same stratum as peers. If all your configured servers fail, a configured peer will help in providing the NTP time. It is generally a good idea to configure one at least one server and peer.

The NTP daemon will know about the NTP servers and peers to use in three ways.

- It can be configured manually with a list of servers to poll,
- It can be configured manually with a list of peers to send to,
- It can look at advertisements issued by other servers on multicast or broadcast addresses.

Note that if multicasting or broadcasting is used, it is strongly recommended to enable authentication unless you trust all hosts on the network.

NTP uses UDP/IP packets for data transfer because of the fast connection setup and response times UDP offers. The NTP protocol uses port UDP port 123. Note that if your router employs a firewall and acts as a client it must open UDP port 123. Additionally, if the router acts as a server the firewall must allow connection requests on port 123 as well.

Section 5.20.1

The NTP Sanity Limit

The NTP daemon corrects the system time through two means, "stepping" and "slewing". If the difference between the local clock and the reference chosen by NTP (the "offset") is more than 128ms for a period of more than 900 seconds, NTP will "step", or instantaneously correct, the time. If the time difference is less than 128ms, NTP will "slew" the time by no more than 500 microseconds every second toward the correct time, in such a way that to an application on the system, the time never appears to be flowing backwards.

NTP will step the system time when it starts up. This is almost always at boot time. Stepping the time afterwards can cause protocols (such as OSPF) that rely upon accurate real time to fail. The router deals with this problem by restarting these protocols if they are running when NTP restarts.

After booting, NTP uses slewing to achieve synchronization by making small and frequent changes to router hardware clock. If the reference server's clock differs from the local clock by more than 1000 seconds, the NTP daemon decides that a major problem has occurred and terminates.

Usually, NTP will succeed in synchronizing the clock at boot time. If it fails to synchronize the clock (perhaps due to a downed WAN link), the NTP daemon may terminate. The router, however, will note the termination and will automatically restart the NTP daemon.

Section 5.20.2

NTP and the Precision Time Protocol Card

If the router is equipped with a Precision Time Protocol card, NTP will treat the Global Positioning System signals received from the card (when GPS locks) as a stratum 0 reference clock. The router will always preferentially use this reference above all others.

Section 5.20.3

Included with NTP

Your ROX software includes the `ntpq`, `ntpd`, `ntptrace` and `ntp-keygen` command line utilities. The `ntpq` utility program can be used to monitor the NTP daemon operations and determine how well it is running. The `ntpd` utility program is used to query the NTP daemon about its current state and to request changes in that state. The `ntptrace` utility is a utility trace a chain of NTP servers back to the primary source.

The `ntp-keygen` utility can be used to generate secure public keys for authentication.

Section 5.20.4

NTP Configuration

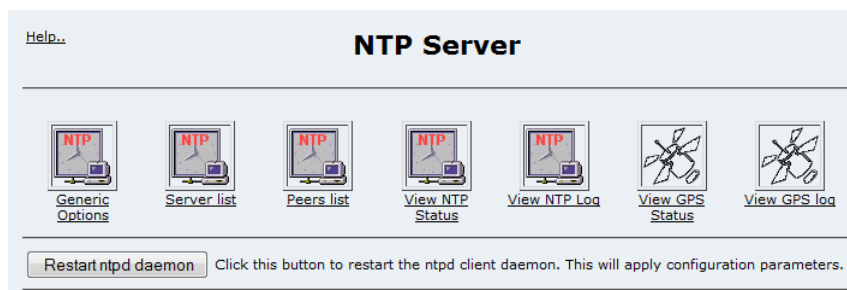


Figure 317: NTP Server

Note that the NTP server is disabled by default and may be enabled via the System folder, Bootup and Shutdown menu. When enabled, any configuration changes may be made to take effect by selecting the *Restart ntpd daemon* button. The *View GPS Status* and *View GPS log* sub-menus appear if the router is equipped with a Precision Time Protocol card.

Section 5.20.5

Generic Options

[Help..](#)

NTP Generic Options

NTP OPTIONS

Broadcast Client: No

Multicast address>: ☒ Default ☐ Custom

Custom address:

Bind Interface (for sending NTP packets): none

Save options

Note: The Bind Interface option allows ntpd always use IP address of the selected network interface as source IP address of outbounding ntp packets.

[Return to NTP Main Menu](#)

Figure 318: NTP Generic Options

Set the *Broadcast Client* option to "Yes" if you wish to act on NTP broadcast messages.

The default multicast address used for NTP is 224.0.1.1. Select a custom multicast address with the *Custom address* field if you wish to use a different address.

The Bind interface option allows you to select an existing interface so that ntp will use the selected interface IP address as the source IP address for outbounding NTP packets.

Section 5.20.6

Servers Configuration

[Help..](#)

NTP Server List

IP ADDRESS	VERSION	KEY	PREFERRED	CHECK
pool.ntp.org	Default (4)	None	No	Contact

Create new

[Return to NTP Main Menu](#)

Figure 319: NTP Server List

The servers under the *IP address* column are used as primary synchronization devices. Clicking on a link will allow you to edit that server.

By default the router includes the links 0.debian.pool.ntp.org. The 0.debian.pool.ntp.org address selects a random low stratum server from a pool of ntp servers on the Internet.

If you are operating in a private network you will want to delete both of these addresses and substitute that of a locally known low stratum server.

The *Version* field indicates the version of the NTP protocol used to communicate with this host. Change this only if it is known that the host requires a version other than 4.

The *Key* field provides an authentication key ID associated with this host. The creation of keys is not supported by the ROX Web interface.

The *Preferred* field determines whether this host is preferred over other hosts in the list.

The *Check* field link leads to a page that displays the result of an NTP query to this host. Use this feature to determine if the configured host is active.

Section 5.20.7

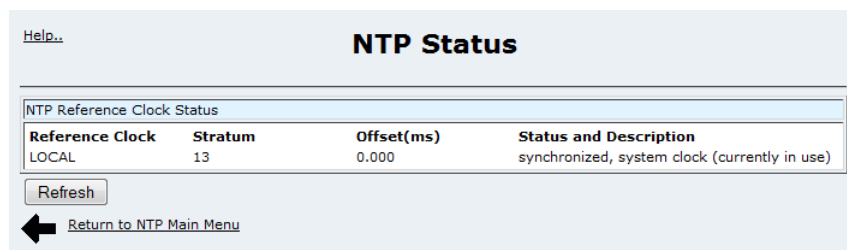
Peers Configuration

This menu allows you to enter and edit peers. Peers are NTP servers of the same stratum as the router, and are useful when contact is lost with the hosts in the NTP servers menu.

The per-peer configuration information is as described in the previous menu.

Section 5.20.8

Viewing NTP Status



NTP Reference Clock Status			
Reference Clock	Stratum	Offset(ms)	Status and Description
LOCAL	13	0.000	synchronized, system clock (currently in use)

Figure 320: NTP Status

The NTP Status menu displays possible sources and currently used reference clocks

Section 5.20.9

Viewing The NTP Log

Help...

NTP Log

Refresh

Month	Day	Time	Process	Event
/var/log	24	18:25:23	ntpd[6059]	ntpd 4.2.4p4@1.1520-o Sat Feb 5 16:17:18 UTC 2011 (1)
/syslog:Oct	24	18:25:23	ntpd[6059]	precision = 3.000 usec
/var/log	24	18:25:23	ntpd[6059]	Listening on interface #0 wildcard, 0.0.0.0#123 Disabled
/syslog:Oct	24	18:25:23	ntpd[6059]	Listening on interface #1 wildcard, ::#123 Disabled
/var/log	24	18:25:23	ntpd[6059]	Listening on interface #2 lo, ::1#123 Enabled
/syslog:Oct	24	18:25:23	ntpd[6059]	Listening on interface #3 eth3.0030, fe80::20a:dcff:fe08:a01e#123 Enabled
/var/log	24	18:25:23	ntpd[6059]	Listening on interface #4 eth2, fe80::20a:dcff:fe08:a01d#123 Enabled
/syslog:Oct	24	18:25:23	ntpd[6059]	Listening on interface #5 eth1, fe80::20a:dcff:fe08:a01c#123 Enabled
/var/log	24	18:25:23	ntpd[6059]	Listening on interface #6 lo, 127.0.0.1#123 Enabled
/syslog:Oct	24	18:29:13	ntpd[6285]	kernel time sync status change 0001

Refresh

Figure 321: NTP Log

The NTP Log menu displays the log of recent NTP events.

Section 5.20.10

Viewing GPS Status

Help...

GPS Status

GPS Status

Latitude	Longitude	GPS Lock	Number of Satellites Tracked
43° 48.4722'[N]	79° 32.4232'[W]	locked	10

Tracked Satellite Status

Satellite ID	Satellite Strength (dBHz)
18	42
22	47
24	41
21	47
14	48
15	50
27	40
19	39
06	37
03	35

refresh

← Return to NTP Main Menu

Figure 322: GPS Status

If the router is equipped with a Precision Time Protocol card, this page will shows the status of the GPS module.

The *Latitude* and *Longitude* fields show the current position of the GPS antenna.

The *GPS Lock* field show the GPS lock status.

The *Number of Satellites* shows how many satellites are currently being tracked by the GPS module.

The *Tracked Satellite Status* table shows the ID and signal strength of tracked satellites.

Section 5.20.11

Viewing the GPS Log

[Help...](#)

GPS Log

Refresh

Month	Day	Time	Process	Event
Oct	30	13:08:32	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	13:16:40	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	13:16:51	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	13:26:04	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	13:26:16	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	13:29:05	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	13:29:16	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	13:41:35	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	13:41:47	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	13:51:05	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	13:51:13	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	13:54:05	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	13:54:18	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	14:03:35	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	14:03:46	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	14:06:36	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	14:06:44	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	14:16:05	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	14:16:16	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	14:19:06	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	14:19:20	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	14:19:35	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	14:19:43	/usr/sbin/irigb[2657]	GPS lock - locked!
Oct	30	14:24:08	/usr/sbin/irigb[2657]	GPS lock - lock lost!
Oct	30	14:24:17	/usr/sbin/irigb[2657]	GPS lock - locked!

Refresh

[Return to NTP Main Menu](#)

Figure 323: GPS Log

The GPS Log menu displays the log of recent GPS events.

Section 5.21

CrossBow Station Access Controller (SAC)

There may be times when it is not possible or practical to access a facility’s devices via the network connection to the CrossBow server. To address this possibility, Siemens has developed the Station Access Controller (SAC). The Station Access Controller acts as a local version of the CrossBow server. The SAC is installed on a device

that is physically located within the facility. During normal operation, communications occur as usual between the remote CrossBow server (the enterprise server) and the devices within the facility.

However, if network connectivity is lost, or if network speed makes it impractical for an on-site operator to use the enterprise server connection, a CrossBow client launched from within the facility's network can access the facility's devices via the SAC.

On-site operators can access the local facility devices using their usual CrossBow interface. Operations initiated via the SAC are logged, and can be uploaded to the enterprise server database once the network connection is restored.

The Station Access Controller appears as a device in the Device View in the main CrossBow database.



NOTE

CrossBow SAC is disabled by default and may be enabled via the Bootup and Shutdown menu under the System folder.



NOTE

For more information about the Station Access Controller (SAC), refer to the RUGGEDCOM CrossBow User Guide.

This section familiarizes the user with:

- Configuring CrossBow SAC
- Configuring the log level
- Managing certificates

Section 5.21.1

Configuring CrossBow SAC

Module Index

CrossBow SAC Config Options

Connection Configuration

Server Address : 172.30.151.151Server Port : 21000

Client Connection Timeout : 15 minutes (set to 0 for no timeout)Device Session Timeout : 15 minutes (set to 0 for no timeout)

Server Certificate Configuration

CA Cert File Path : /etc/crossbow/cxb_test_ca_cert.pem

Cert File Path : /etc/crossbow/sac_151_certificate.pem

Private Key File Path : /etc/crossbow/sac_151_privatekey.pem

Private Key Phrase :

Station Access Controller Configuration

SAM-I Common Name : crossbowserverSAM-I Host Address : 10.200.22.232SAM-I Host Port : 21000

SAM-II Common Name : crossbowsamSAM-II Host Address : 10.200.20.172SAM-II Host Port : 21000

User Access Configuration

Max Login Attempts : 3

Event Log Message Configuration

Event Log Message Limit : 100

Save options

Return to CrossBow Main Menu

Figure 324: CrossBow SAC Config Options

This menu allows you to configure CrossBow SAC, including SAC, the SAC connection, certificates/keys, user access, and event logs.

Section 5.21.1.1

Configuring the SAC Connection

Connection Configuration

Server Address : Server Port : 21000

Client Connection Timeout : 15 minutes (set to 0 for no timeout)Device Session Timeout : 15 minutes (set to 0 for no timeout)

Figure 325: Connection Configuration

This menu allows you to configure the CrossBow SAC connection.

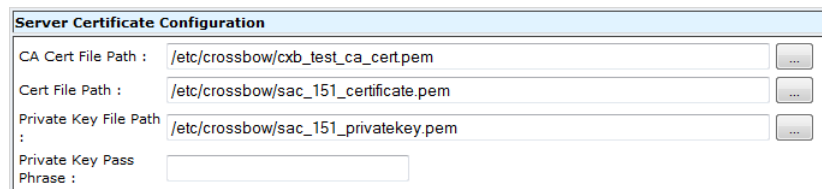
The *Server Address* field configures the IP address for the Station Access Controller (SAC) to which a client will connect.

The *Server Port* field configures the TCP port for the Station Access Controller (SAC) to which a client will connect. The default is 21000.

The *Client Connection Timeout* and *Device Connection Timeout* fields configure the time-out period. If the client or device does not reply before the time-out period ends, the connection is terminated.

Section 5.21.1.2

Configuring the SAC Certificates



Server Certificate Configuration	
CA Cert File Path :	/etc/crossbow/cxb_test_ca_cert.pem
Cert File Path :	/etc/crossbow/sac_151_certificate.pem
Private Key File Path :	/etc/crossbow/sac_151_privatekey.pem
Private Key Pass Phrase :	

Figure 326: Server Certificate Configuration

This menu allows you to configure certificates and private keys for CrossBow SAC.

The *CA Cert File Path* field defines the path to the CA certificate. Use the *Browse* button to navigate to the folder in ROX that contains the certificate.

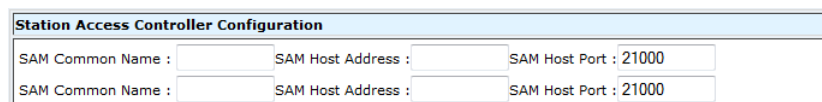
The *Cert File Path* field defines the path to the self-signed certificate. Use the *Browse* button to navigate to the folder in ROX that contains the certificate.

The *Private Key File Path* field defines the path to the private key. Use the *Browse* button to navigate to the folder in ROX that contains the key.

The *Private Key Pass Phrase* field configures the passphrase for the private key.

Section 5.21.1.3

Configuring SAC



Station Access Controller Configuration		
SAM Common Name :	SAM Host Address :	SAM Host Port : 21000
SAM Common Name :	SAM Host Address :	SAM Host Port : 21000

Figure 327: Station Access Controller Configuration

This menu allows you to configure the Station Access Controller (SAC). ROX supports up to two Secure Access Managers (SAMs).



NOTE

The common name, host address and host port must be defined for one or more SAMs.

The *SAM Common Name* field defines the common name in the certificate that the SAM, the parent of the SAC, will present when mutually authenticating with the SAC. The common name must not contain spaces.

The *SAM Host Address* field defines the IP address for the SAM.

The *SAM Host Port* field defines the TCP port for the SAM. The default is 21000.

Section 5.21.1.4

Configuring User Access


A screenshot of the 'User Access Configuration' dialog box. It has a title bar with the text 'User Access Configuration'. Below the title bar, there is a label 'Max Login Attempts' followed by a dropdown menu showing the value '3'.

Figure 328: User Access Configuration

This menu allows you to configure the maximum number of times a user can attempt to log in to the Station Access Controller (SAC).

The *Max Login Attempts* field defines the maximum number, from 0 (no limit) to 10.

Section 5.21.1.5

Configuring Event Logs

A screenshot of the 'Event Log Message Configuration' dialog box. It has a title bar with the text 'Event Log Message Configuration'. Below the title bar, there is a label 'Event Log Message Limit' followed by a dropdown menu showing the value '100'.

Figure 329: Event Log Message Configuration

This menu allows you to configure the maximum number of event messages logged by the Station Access Controller (SAC). When the limit is reached, the oldest message will be removed for each new message added to the log.

The *Event Log Message Limit* field defines the maximum number, from 0 (no limit) to 5000. The default is 100.

Section 5.21.2

Configuring Log Options

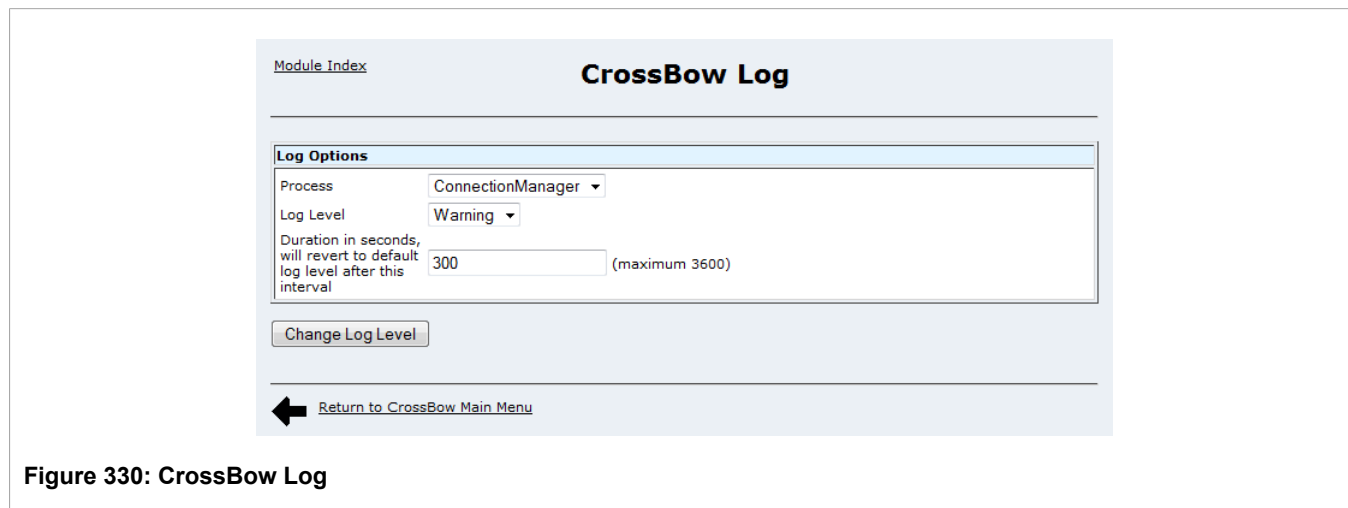


Figure 330: CrossBow Log

This menu allows you to configure the log level for a specific CrossBow process. CrossBow SAC must be enabled to configure this option.

The *Process* field selects the process.

The *Log level* field selects the log level (e.g. Warning).

The *Duration in seconds* field defines how long the log level will last.

Section 5.21.3

CrossBow Certificates

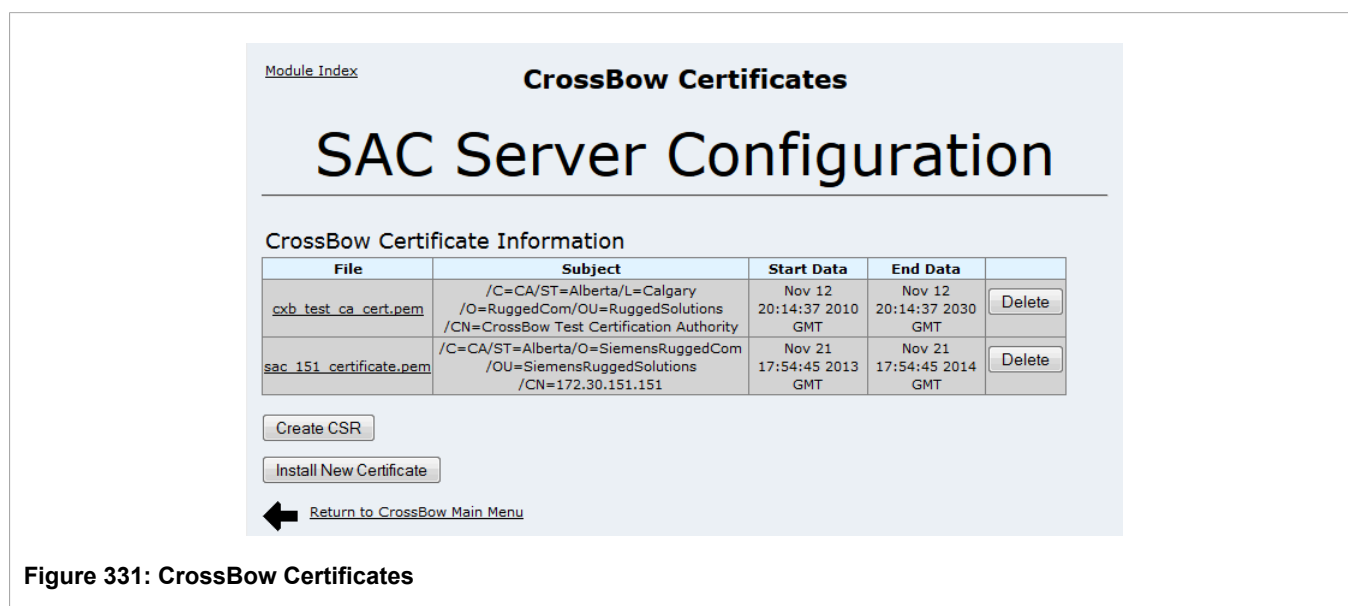


Figure 331: CrossBow Certificates

This menu allows you to review the certificates that have been installed, generate a Certificate Signing Request (CSR), and upload certificates.

For each certificate listed in the table, click the name of the certificate to view more details.

Click the *Delete* button to delete a certificate.

Click the *Create CSR* button to create a Certificate Signing Request (CSR). See [Section 5.21.3.1, “Generating a Certificate Signing Request \(CSR\)”](#).

Click the *Install New Certificate* button to upload and install new certificates. See [Section 5.21.3.2, “Installing Certificates”](#).

Section 5.21.3.1

Generating a Certificate Signing Request (CSR)



The screenshot shows a web interface titled "Certificate Signing Request". At the top left is a link for "Module Index". The main form is titled "Generate CSR" and contains several input fields: "Common Name" (with the value "ruggedcom" and an asterisk), "Pass Phrase" (with an asterisk), "Department", "Organization", "Locality(Eg: City)", "State/Province", and "Country Code". Below these fields are radio buttons for "RSA key size" with options "1024" (selected) and "2048". A "Generate CSR" button is located below the form. At the bottom left, there is a back arrow icon and a link "Return to CrossBow Certificates".

Figure 332: Certificate Signing Request

The *Common Name* field specifies the common name in the certificate.

The *Passphrase* field sets a passphrase for the CSR.

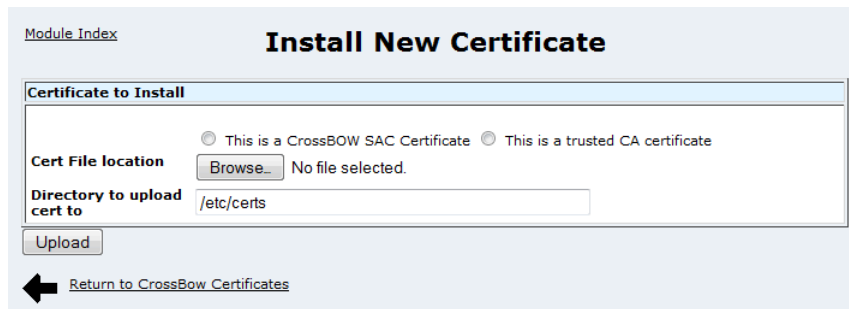
The *Department*, *Organization*, *Locality*, *State/Province* and *Country Code* fields define the address of the Certification Authority (CA).

The *RSA key size* options specify the size of the keys: 1024 or 2048 bits.

The *Generate CSR* button generates the final CSR.

Section 5.21.3.2

Installing Certificates



The screenshot shows a web interface titled "Install New Certificate". At the top left is a link for "Module Index". The main form is titled "Certificate to Install" and contains two radio buttons: "This is a CrossBOW SAC Certificate" (selected) and "This is a trusted CA certificate". Below these are two fields: "Cert File location" with a "Browse..." button and the text "No file selected.", and "Directory to upload cert to" with a text input field containing "/etc/certs". At the bottom of the form is an "Upload" button. Below the form is a back arrow icon and a link "Return to CrossBow Certificates".

Figure 333: Install New Certificate

This menu allows you to upload new certificates for CrossBow SAC.

The *Cert file location* fields specify the file, its location and what type of certificate it is. Use the *Browse* button to navigate to and select the file from your local PC or network.

The *Directory to upload cert to* field specifies where to save the certificate on the device.

The *Upload* button uploads the selected certificate to the device.

6 Upgrades

This chapter details how to install optional software to facilitate upgrading the ROX software across a network. It describes the following tasks:

- [Section 6.1, “Installing Apache Web Server On Windows”](#)
- [Section 6.2, “Installing a Microsoft IIS Web Server”](#)
- [Section 6.3, “VPN/L2TP Configuration in Windows”](#)

Section 6.1

Installing Apache Web Server On Windows

A number of customers have asked for advice and instructions on setting up a web server on Windows. Siemens recommends the Apache web server, because it is secure, robust, easy to install and configure as well as being able to be installed on a wide variety of Windows platforms.

Begin by identifying a host computer and its physical and logical location on the network. The *Repository Server Requirements* of the appendix "Setting Up A Repository" provide some guidance on host requirements. The Apache installation process will prompt you for an IP address and domain name with which to serve the web pages. Later in the install, you will also need to provide the directory where the ROX releases will be kept. Ensure that a web servers is not already installed.

Obtain Apache by visiting the web page of www.apache.org [http://www.apache.org/]. Visit the "HTTP Server" portion of the web site and click on the "Downloads" page. Identify the latest version of Apache and find its Win32 version, usually under "httpd/binaries/win32/". You should be able to find a Microsoft System Installer Version (e.g. apache_2.0.55-win32-x86-no_ssl.msi), as well as platform specific notes. Download and install this version.

Verify the web server by opening a web browser on another host on the network and entering the URL http:// followed by the IP address Apache was installed with. Note that you may also verify Apache from a browser on the web server itself by browsing <http://localhost>. [http://localhost/] If properly set-up, the Apache default web page will be shown.

If you can see this, it means that the installation of the [Apache web server](#) software on this system was successful. You may now add content to this directory and replace this page.

Seeing this instead of the website you expected?

This page is here because the site administrator has changed the configuration of this web server. Please **contact the person responsible for maintaining this server with questions**. The Apache Software Foundation, which wrote the web server software this site administrator is using, has nothing to do with maintaining this site and cannot help resolve configuration issues.

The Apache [documentation](#) has been included with this distribution.

You are free to use the image below on an Apache-powered web server. Thanks for using Apache!



Figure 334: Apache Default Web Page

Apache serves the web pages contained in the directory known as the "DocumentRoot". You must change the document root by, from the desktop, clicking Start > All programs > Apache HTTP Server > Configure Apache Server > Edit the Apache httpd.conf file. Search the file for the DocumentRoot variable and change it to the directory where your ROX releases are kept. Restart Apache by clicking Start > All programs > Apache HTTP Server > Control Apache Web Server > Restart.

Return to the web browser used earlier to verify Apache and refresh the screen. It should now reflect the contents of your ROX release directory. You should now be able to perform an upgrade from a router.

Section 6.2

Installing a Microsoft IIS Web Server

This section provides general advice and instructions on setting up a Microsoft IIS web server on Windows. For complete installation details, refer to the official Microsoft documentation.

Begin by identifying a host computer that has IIS and its physical and logical location on the network. For some guidance on host requirements, see [Section 3.5.1, "Repository Server Requirements"](#).

Start to install IIS by clicking on *Start menu, Control Panel, Add or Remove Programs, Add/remove Windows Components*. In the resultant menu check the *Internet Information Services(IIS)* box and select next.

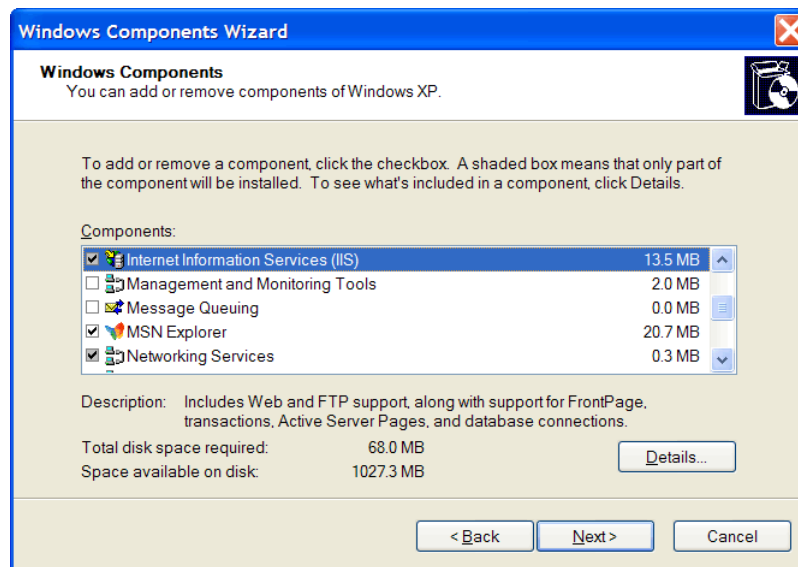


Figure 335: Installing IIS

Download the desired release (e.g. rr1.9.0.zip) from www.siemens.com/ruggedcom. Create the directory ruggedcom under the IIS root directory `C:\inetpub\wwwroot\`. Unzip the rr1.9.0.zip file within `C:\inetpub\wwwroot\ruggedcom`.

Start to enable IIS by clicking on *Start menu, Control Panel, Administrative Tools, Internet Information Services*. Right click on *Internet Information Services, Connect* and enter the host computer's IP address, e.g. 192.168.0.1.

Verify the IIS web server by opening a web browser on another host on the network and entering the URL `http://` followed by the IP address IIS was installed with, followed by `/ruggedcom`, e.g. `http://192.168.0.1/ruggedrouter`.

Visit the router you wish to upgrade and visit the *Maintenance* menu, *Upgrade System* sub-menu. Click on the *Change Server* button and set the Repository Server field (e.g. `http:// 192.168.0.1/ruggedcom`). Set the Release Version field to rr1. Save the configuration and return to the *Maintenance* menu. Set the *Only show which packages would be upgraded* radio button to *No* and click on the *Upgrade Now* button to start the upgrade.

Section 6.2.1

Using Microsoft Internet Information Services (IIS) Manager 6.0 or Higher as an Upgrade Repository

When using Microsoft Internet Information Services (IIS) Manager 6.0 or higher as your ROX upgrade repository, you must add a new application/octet-stream MIME type named * to the IIS properties. The new MIME type is required for IIS to consider ROX upgrade packets as an application/octet stream. If the new MIME type is not added, ROX upgrades will fail.

To add the new MIME type, do the following on your IIS server.

1. In the Windows **Start** menu, right-click on **My Computer** and select **Manage**. The **Computer Management** dialog appears.
2. Under **Services and Applications**, locate the **Internet Information Services (IIS) Manager** node. Right-click on your ROX upgrade repository website and select **Properties**. The **Properties** dialog appears.
3. Select the **HTTP Headers** tab and click **MIME Types**. The **MIME Types** dialog appears.

4. Click **New**. The **MIME Type** dialog appears.
5. In the **Extension** field, type *****.
In the **MIME type** field, type **application/octet-stream**.
6. Click **OK** on the **MIME Type**, **MIME Types**, and **Properties** dialog boxes.

Section 6.3

VPN/L2TP Configuration in Windows

This section describes how to set up a VPN/L2TP connection on Windows XP/2000. There are two ways to establish a connection in Windows: using a pre-shared key (in the case of Windows XP) or using a certificate (for either Windows XP or Windows 2000).

Here are the steps to establish a connection with a pre-shared key:

1. Start the "New Connection Wizard" (accessed via the Start > All Programs > Accessories > Communications menus).
2. On the "Network Connection Type" page, select the option "Connect to the network at my workplace".
3. On the "Network Connection" page, select the option "Virtual Private Network Connection".
4. On the "Connection Name" page, enter a name for the new connection.
5. On the "Public Network" page, select either "Do not dial the initial connection" or "Automatically dial this initial connection" according to your requirements.
6. On the "VPN Server Selection" page, enter the IP address of your device.
7. After the connection has been created, double click the connection. A "Connect" window will appear - select "Properties".
8. Select the "Security" tab. You will need to disable L2TP/PPP encryption (unless you want double encryption) using either of the following methods:
 - a. In the "Security" tab, click the "IPSec Settings" button, enable the "Use pre-shared key for authentication" check box and enter the pre-shared key (If you decide to use a certificate, disable the "Use pre-shared key for authentication" check box).
 - b. Select the "Networking" tab; select "L2TP IPSec VPN" for "Type of VPN".
 - c. Click the "OK" button to save the Properties settings.
 - d. Now you are back to "Connect" window; enter your user name and password to begin the connection.

More information about how to import a certificate in Windows XP/2000 can be found at the link:

<http://www.jacco2.dds.nl/networking/openswan-l2tp.html#Certificates>