

DIGITAL clearVISN CoreWatch

User's Guide

Part Number: 9032685-02

December 1998

This guide provides a general overview of DIGITAL clearVISN CoreWatch.

Revision/Update Information: This is a revised document.

Notice

Cabletron Systems reserves the right to make changes in specifications and other information contained in this document without prior notice. The reader should in all cases consult Cabletron Systems to determine whether any such changes have been made.

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

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

© Cabletron Systems, Inc., 1998.

All Rights Reserved
Printed in the United States of America

Cabletron Systems and **SPECTRUM** are registered trademarks and **Cabletron**, **clearVISN**, and **GIGAswitch** are trademarks of Cabletron Systems, Inc.

DEC, DIGITAL, and the DIGITAL logo are trademarks of Digital Equipment Corporation.

Java and Solaris are trademarks of Sun Microsystems, Inc.

Netscape Navigator is a registered trademark of Netscape Communications Corp.

Pentium is a registered trademark of Intel Corp.

Windows NT is a trademark and Microsoft, Windows, and Windows 95, are registered trademarks of Microsoft Corp.

HP is a registered trademark and OpenView is a trademark of Hewlett Packard Company.

UNIX is a registered trademark of the Open Group in the US and other countries.

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

FCC Notice

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment uses, generates, and can radiate radio frequency energy and if not installed in accordance with the operator's manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause interference in which case the user will be required to correct the interference at his own expense.

WARNING: Changes or modifications made to this device which are not expressly approved by the party responsible for compliance could void the user's authority to operate the equipment.

Industry Canada Notice

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus set out in the Radio Interference Regulations of the Canadian Department of Communications.

Le présent appareil numérique n'émet pas de bruits radioélectriques dépassant les limites applicables aux appareils numériques de la class A prescrites dans le Règlement sur le brouillage radioélectrique édicté par le ministère des Communications du Canada.

VCCI Notice — Class A Computing Device

This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Cabletron Systems, Inc. Program License Agreement

IMPORTANT: Before utilizing this product, carefully read this License Agreement.

This document is an agreement between you, the end user, and Cabletron Systems, Inc. ("Cabletron") that sets forth your rights and obligations with respect to the Cabletron software program (the "Program") contained in this package. The Program may be contained in firmware, chips or other media. BY UTILIZING THE ENCLOSED PRODUCT, YOU ARE AGREEING TO BECOME BOUND BY THE TERMS OF THIS AGREEMENT, WHICH INCLUDES THE LICENSE AND THE LIMITATION OF WARRANTY AND DISCLAIMER OF LIABILITY. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, PROMPTLY RETURN THE UNUSED PRODUCT TO THE PLACE OF PURCHASE FOR A FULL REFUND.

Cabletron Software Program License

1. **LICENSE.** You have the right to use only the one (1) copy of the Program provided in this package subject to the terms and conditions of this License Agreement.

You may not copy, reproduce or transmit any part of the Program except as permitted by the Copyright Act of the United States or as authorized in writing by Cabletron.

2. **OTHER RESTRICTIONS.** You may not reverse engineer, decompile, or disassemble the Program.
3. **APPLICABLE LAW.** This License Agreement shall be interpreted and governed under the laws and in the state and federal courts of New Hampshire. You accept the personal jurisdiction and venue of the New Hampshire courts.

Exclusion of Warranty and Disclaimer of Liability

1. **EXCLUSION OF WARRANTY.** Except as may be specifically provided by Cabletron in writing, Cabletron makes no warranty, expressed or implied, concerning the Program (including its documentation and media).

CABLETRON DISCLAIMS ALL WARRANTIES, OTHER THAN THOSE SUPPLIED TO YOU BY CABLETRON IN WRITING, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WITH RESPECT TO THE PROGRAM, THE ACCOMPANYING WRITTEN MATERIALS, AND ANY ACCOMPANYING HARDWARE.

2. **NO LIABILITY FOR CONSEQUENTIAL DAMAGES.** IN NO EVENT SHALL CABLETRON OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS, PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR RELIANCE DAMAGES, OR OTHER LOSS) ARISING OUT OF THE USE OR INABILITY TO USE THIS CABLETRON PRODUCT, EVEN IF CABLETRON HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES DO NOT ALLOW THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, OR ON THE DURATION OR LIMITATION OF IMPLIED WARRANTIES, IN SOME INSTANCES THE ABOVE LIMITATIONS AND EXCLUSIONS MAY NOT APPLY TO YOU.

United States Government Restricted Rights

The enclosed product (a) was developed solely at private expense; (b) contains "restricted computer software" submitted with restricted rights in accordance with Section 52227-19 (a) through (d) of the Commercial Computer Software - Restricted Rights Clause and its successors, and (c) in all respects is proprietary data belonging to Cabletron and/or its suppliers.

For Department of Defense units, the product is licensed with "Restricted Rights" as defined in the DoD Supplement to the Federal Acquisition Regulations, Section 52.227-7013 (c) (1) (ii) and its successors, and use, duplication, disclosure by the Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at 252.227-7013. Cabletron Systems, Inc., 35 Industrial Way, Rochester, New Hampshire 03867-0505.

SAFETY INFORMATION

CLASS 1 LASER TRANSCEIVERS

The DGSRF-AA 100Base-FX Module, DGSRS-AA 1000BASE-SX Module, and DGSRL-AA 1000BASE-LX Module use Class 1 Laser transceivers. Read the following safety information before installing or operating these modules.

The Class 1 laser transceivers use an optical feedback loop to maintain Class 1 operation limits. This control loop eliminates the need for maintenance checks or adjustments. The output is factory set, and does not allow any user adjustment. Class 1 Laser transceivers comply with the following safety standards:

- 21 CFR 1040.10 and 1040.11 U.S. Department of Health and Human Services (FDA).
- IEC Publication 825 (International Electrotechnical Commission).
- CENELEC EN 60825 (European Committee for Electrotechnical Standardization).

When operating within their performance limitations, laser transceiver output meets the Class 1 accessible emission limit of all three standards. Class 1 levels of laser radiation are not considered hazardous.

Laser Radiation and Connectors

When the connector is in place, all laser radiation remains within the fiber. The maximum amount of radiant power exiting the fiber (under normal conditions) is -12.6 dBm or 55×10^{-6} watts.

Removing the optical connector from the transceiver allows laser radiation to emit directly from the optical port. The maximum radiance from the optical port (under worst case conditions) is 0.8 W cm^{-2} or $8 \times 10^3 \text{ W m}^2 \text{ sr}^{-1}$.

Do not use optical instruments to view the laser output. The use of optical instruments to view laser output increases eye hazard. When viewing the output optical port, power must be removed from the network adapter.

DECLARATION OF CONFORMITY ADDENDUM

Application of Council Directive(s):	89/336/EEC 73/23/EEC
Manufacturer's Name:	Cabletron Systems, Inc.
Manufacturer's Address:	35 Industrial Way PO Box 5005 Rochester, NH 03867
European Representative Name:	Mr. J. Solari
European Representative Address:	Cabletron Systems Limited Nexus House, Newbury Business Park London Road, Newbury Berkshire RG13 2PZ, England
Conformance to Directive(s)/Product Standards:	EC Directive 89/336/EEC EC Directive 73/23/EEC EN 55022 EN 50082-1 EN 60950
Equipment Type/Environment:	Networking Equipment, for use in a Commercial or Light Industrial Environment.

We the undersigned, hereby declare, under our sole responsibility, that the equipment packaged with this notice conforms _to the above directives.

Manufacturer	Legal Representative in Europe
Mr. Ronald Fotino	Mr. J. Solari
Full Name	Full Name
Principal Compliance Engineer	Managing Director - E.M.E.A.
Title	Title
Rochester, NH, USA	Newbury, Berkshire, England
Location	Location

Contents

Preface	xvii
About This Manual	xvii
Who Should Read This Manual?	xvii
How to Use This Manual	xvii
Related Documentation.....	xix
Chapter 1: A Look at clearVISN CoreWatch	1
What Are clearVISN CoreWatch’s Features?	2
System Requirements	2
clearVISN CoreWatch Capabilities.....	3
A Look at the Modes of clearVISN CoreWatch	3
Which MIBs Does the GSR Support?	4
Chapter 2: clearVISN CoreWatch Installation	5
Installing clearVISN CoreWatch.....	5
Installing on a Solaris System.....	5
Installing on a Windows NT, Windows 95, or Windows 98 System	6
Chapter 3: Learning clearVISN CoreWatch Basics	9
Starting clearVISN CoreWatch.....	9
Starting clearVISN CoreWatch in Solaris	10
Starting clearVISN CoreWatch in Windows NT, Windows 95, or Windows 98 ...	10
Starting clearVISN CoreWatch from within SPECTRUM Enterprise Manager	11
Starting clearVISN CoreWatch from within HP OpenView 5.x	11
A Look at the clearVISN CoreWatch Interface	12
Front Panel View	12
A Look at the Modules.....	13
Schematic View	14
Opening the Schematic View	15
Using the Schematic View	15
clearVISN CoreWatch Menus	16
clearVISN CoreWatch Toolbar.....	16
clearVISN CoreWatch Passwords.....	16
Changing the Login Password.....	16
Changing the Privileged Password	17
Accessing Help	17
Exiting clearVISN CoreWatch	18

Chapter 4: Learning Configuration Expert Basics	19
What Is Configuration Expert?	19
Starting Configuration Expert	20
Starting Configuration Expert from the Front Panel View	20
Starting Configuration Expert in Solaris.....	20
Starting Configuration Expert in Windows NT, Windows 95, or Windows 98 ...	21
A Look at the Configuration Expert Window.....	21
Configuration Tree.....	22
A Look at the Configuration Tree Icons	25
Wizards and Dialog Boxes.....	25
Copying Configuration Settings with Drag-and-Drop.....	26
Finding Objects	26
Deleting Objects.....	27
Order of Configuration Tasks.....	27
Saving and Applying Your Configuration Changes.....	28
Saving Changes to a Configuration File:	28
Loading a Configuration File into a GSR.....	28
Retrieving a Configuration File from a GSR.....	29
Exiting Configuration Expert	30
Chapter 5: Changing System Settings	31
Providing System Information	31
Configuring a GSR Chassis.....	32
Configuring Ports.....	34
Configuring Global Settings on All Ports.....	34
Configuring an Individual Port	36
Configuring the GSR for a SYSLOG Server.....	39
Configuring for DNS	41
Configuring the GSR for SNMP	42
Setting Up a Target for SNMP Traps	42
Establishing Community Strings	44
Chapter 6: Configuring GSR Bridging	47
A Look at Bridging on the GSR.....	47
Configuring the Bridging Mode of Ports.....	48
Configuring a Port for Flow-Based Bridging	49
Configuring a Port for Address-Based Bridging.....	50
Controlling the Aging State of GSR Bridging	52
Setting Up a Default Aging Timeout.....	52
Overriding the Default Timeout Interval on a Port	53
Disabling Aging on a Port.....	54
Enabling Aging on a Port.....	56
Setting Up STP on the GSR.....	58
Defining STP Settings for GSR Bridging.....	59
Defining STP Attributes on an Individual Port	60
Enabling STP on a Port.....	61
Disabling STP on a Port.....	63
Configuring SmartTRUNK Behavior on the GSR	64
Defining SmartTRUNK Settings for GSR Bridging.....	65

Chapter 7: Configuring VLANs on the GSR.....	69
A Look at VLANs on the GSR.....	69
VLAN Configuration Tips.....	70
Defining Access Ports and Trunk Ports.....	71
Creating a Protocol-Based VLAN.....	73
Creating a Port-Based VLAN.....	78
Modifying VLANs.....	82
Changing a Port-Based VLAN's Name or ID.....	83
Changing a Protocol-Based VLAN's Name, ID, or Protocol Binding.....	84
Replacing an Interface's VLAN.....	85
Changing Which Ports a VLAN Includes.....	85
Dragging Ports to Add Them to a VLAN.....	85
Adding and Removing a VLAN's Ports Through a Dialog Box.....	86
Chapter 8: Configuring IP Interfaces for the GSR.....	89
What Is IP?.....	89
A Look at IP Addresses.....	90
Creating IP Interfaces.....	92
Creating IP Interfaces Bound to a Single Port.....	92
Creating IP Interfaces Bound to a VLAN.....	98
Modifying IP Interface Definitions.....	103
Configuring the GSR for VRRP.....	106
Specifying VRRP Trace Options.....	107
Configuring a New VRRP Router.....	108
What to Do Next.....	111
Chapter 9: Configuring Unicast Routing on the GSR.....	113
Configuring Unicast Global Parameters and Static Entries.....	113
Setting Global Parameters for Unicast Routing.....	114
Defining Static ARP Entries.....	115
Defining Static Route Entries.....	117
A Look at RIP Routing in the IP Environment.....	119
What Is Preference?.....	120
Configuring the GSR for RIP.....	121
Setting RIP Global Parameters.....	121
Defining IP RIP Interfaces.....	123
Adding Trusted Gateways.....	126
Adding Source Gateways.....	127
What to Do Next.....	127
Chapter 10: Configuring Multicast Routing on the GSR.....	129
What Is DVMRP?.....	129
Configuring DVMRP Routing on the GSR.....	130
Setting DVMRP Global Parameters on the GSR.....	131
Configuring DVMRP Interfaces.....	132
Defining DVMRP Tunnels.....	133
Enabling or Disabling DVMRP on Tunnels.....	135
What Is IGMP?.....	136
Configuring IGMP Interfaces on the GSR.....	137

Setting IGMP Global Parameters	137
Enabling or Disabling IGMP on Interfaces.....	138
What to Do Next.....	140
Chapter 11: Configuring the GSR for IPX Routes.....	141
What Is IPX?	141
Creating IPX Interfaces.....	143
Creating IPX Interfaces Bound to a Single Port.....	143
Creating IPX Interfaces Bound to a VLAN.....	149
Modifying IPX Interface Definitions	154
Configuring Static IPX SAP Entries.....	157
What to Do Next.....	160
Chapter 12: Configuring QoS on the GSR	161
What Is QoS?.....	161
Establishing the GSR's Queuing Policy	162
Associating Precedences to Layer-3/Layer-4 Flows	163
Assigning IP QoS Precedence.....	163
Assigning IPX QoS Precedence.....	165
Creating QoS Profiles.....	166
Creating a QoS Profile for an IP Flow	167
Creating a QoS Profile for an IPX Flow	172
Creating a QoS Profile for a Layer-2 Flow.....	177
Modifying QoS Profiles	181
Redefining an IP Flow	181
Redefining an IPX Flow Definition.....	182
Redefining a Layer-2 Flow Definition.....	183
Changing an IP or IPX Flow's Interface List	184
Adding or Deleting a Flow's Interfaces through a Dialog Box	185
Dragging an Interface to Apply a Flow to the Interface.....	186
Changing a Layer-2 Port List.....	186
Chapter 13: Configuring Security on the GSR.....	189
A Look at ACLs	189
Setting IP Security	190
Setting IPX Security.....	194
Configuring an IPX ACL.....	194
Setting Up IPX RIP Filters.....	199
Setting Up IPX SAP Filters.....	203
Applying ACLs to IP or IPX Interfaces	206
Copying an ACL to Apply It to an Interface.....	207
Applying an ACL by Editing an Interface's Definition.....	208
Setting Layer-2 Security.....	210
Configuring Layer-2 Address Filters	211
Configuring Layer-2 Port-to-Address Lock Filters	213
Configuring Layer-2 Static-Entry Filters	216
Configuring Layer-2 Secure Port Filters	220
Modifying the GSR's Security Settings	223
Changing an ACL's Name.....	224

Adding or Modifying ACL Rules	225
Modifying Layer-2 Security Filters.....	225
Modifying a Filter’s Settings	225
Modifying a Filter’s Port Bindings	226
Chapter 14: Configuring OSPF on the GSR	229
A Look at OSPF Routing on the GSR	229
Setting OSPF Global Parameters.....	230
Configuring OSPF Area Tables.....	231
Creating OSPF Area Tables	232
Modifying Area Tables	249
Chapter 15: Configuring BGP on the GSR	251
BGP Overview	251
Configuring your GSR for BGP.....	252
Setting BGP Global Parameters	252
Configuring a New BGP Peer Group.....	253
Modifying an Existing BGP Peer Group.....	270
Chapter 16: Configuring Routing Policies on the GSR	279
Routing Policies on the GSR	279
RIP and OSPF Routing Policy Defaults	279
Setting RIP Routing Policy Defaults.....	280
Setting OSPF Routing Policy Defaults	281
A Look at the Building Blocks of Routing Policies.....	282
Export Destination Building Blocks	283
Configuring and Modifying RIP Export Destinations	283
Configuring and Modifying OSPF Export Destinations	285
Configuring and Modifying BGP Export Destinations.....	286
Aggregate Destination Building Blocks.....	288
Configuring Aggregate Destinations.....	288
Modifying Aggregate Destinations.....	290
Export Source Building Blocks.....	291
Configuring and Modifying RIP Export Sources.....	291
Configuring and Modifying OSPF Export Sources	293
Configuring and Modifying BGP Export Sources	294
Configuring and Modifying Autonomous System Path Export Sources	296
Configuring and Modifying Tag Export Sources.....	298
Configuring and Modifying Direct Export Sources	300
Configuring and Modifying Static Export Sources.....	301
Configuring and Modifying Aggregate Export Sources.....	303
Import Source Building Blocks	304
Configuring and Modifying RIP Import Sources	305
Configuring and Modifying OSPF Import Sources.....	307
Configuring and Modifying BGP Import Sources.....	309
Configuring and Modifying Aggregate Source Building Blocks.....	311
IP Route Filter Building Blocks	312
Configuring IP Route Filters	313
Modifying IP Route Filters	315

Configuring and Modifying Optional Attribute Building Blocks	316
Export Policies	317
Configuring Export Policies	318
Modifying Export Policies	323
Import Policies	324
Configuring Import Policies	324
Modifying Import Policies	328
Aggregate Policies	329
Configuring Aggregate Policies	329
Modifying Aggregate Policies	334
Redistribute Policies	335
Configuring Redistribute Policies	335
Modifying Redistribute Policies	339
Summarize Routes	340
Configuring Summarize Routes	340
Modifying Summarize Routes	344
Chapter 17: Checking System Status	347
Obtaining Chassis Information	347
Obtaining Port Information	348
Obtaining Trap Information	350
Obtaining SmartTRUNK Information	351
Chapter 18: Monitoring Real-Time Performance	353
Monitoring System Performance	353
Setting the Scaling of Dials	355
Monitoring Port Utilization	356
Obtaining Statistics About an Individual Port	358
Obtaining Packet Statistics	359
Obtaining Port Byte Statistics	360
Obtaining Port Error Statistics	361
Monitoring IP Interface Statistics	363
Obtaining IP Packet Statistics	363
Obtaining IP Reassembly Statistics	364
Obtaining IP Error Statistics	366
Monitoring IPX Interface Statistics	367
Obtaining IPX Packet Statistics	368
Obtaining IPX Error Statistics	369
Using the Graph Toolbar	371
Chapter 19: Checking the Status of Bridge Tables	373
Obtaining VLAN Information	373
Obtaining STP Port Information	374
Obtaining L2 Interface Information	376
Chapter 20: Checking the Status of Routing Tables	379
Checking IP Routing Status	379
Obtaining IP Interface Information	380

Obtaining IP Forwarding Information	381
Checking IPX Routing Status	382
Obtaining IPX Interface Information	383
Obtaining IPX Forwarding Information.....	385
Checking OSPF Routing Status.....	387
Obtaining OSPF Interface Information	387
Obtaining OSPF Area Information.....	391
Obtaining OSPF Neighbor Information	392
Obtaining OSPF Link-State Database Information	394
Obtaining OSPF Area Aggregate Information	396
Checking RIP Routing Status	398
Obtaining RIP Interface Information	398
Obtaining RIP Peer Information.....	399
Checking DVMRP Routing Status.....	401
Obtaining DVMRP Interface Information	401
Obtaining DVMRP Neighbor Information.....	403
Obtaining DVMRP Routing Information	405
Obtaining DVMRP Next Hop Information	406
Checking IGMP Status	407
Obtaining IGMP Interface Information	408
Obtaining IGMP Cache Information.....	410
Chapter 21: Checking the Status of QoS Tables	413
Obtaining Layer-2 Priority Information	413
Obtaining Flow Priority Information.....	415
Obtaining Layer-2 Switching Information	417
Obtaining Layer-3 and Layer-4 Switching Information	418
Chapter 22: Obtaining Reports	421
Saving Multiple Tables as a Report	421
Saving a Single Table as a Report	423
Appendix A: Working with Tables	425
Finding Text in a Table.....	425
Controlling the Contents of Tables.....	426
Refreshing a Table.....	427
Restoring Table Information.....	427
Obtaining Table Records.....	427
Saving a Single Table as a Report	427
Exporting Data from a Table	428
Sorting Table Information.....	428
Appendix B: clearVISN CoreWatch Menus.....	429
File Menu.....	430
Monitor Menu.....	430
Performance State Submenu	431
System State Submenu	432
Bridging State Submenu	432
Routing State Submenu.....	433

QoS State Submenu.....	435
Window Menu	436
Help Menu.....	437
Appendix C: Supported Regular Expressions.....	439
Appendix D: Error Messages	443
Missing or Invalid Field Error Messages	443
Duplicate Objects Error Messages	447
Already Exists or in Use Error Messages.....	447
Unavailable Objects Error Messages	448
Miscellaneous Error Messages	449
Glossary	455
Index	465

Preface

About This Manual

This manual provides a general overview of DIGITAL clearVISN CoreWatch and provides procedures for using that application to configure and monitor a DIGITAL GIGAswitch/Router (GSR). For product information not available in this manual, see the manuals listed in [“Related Documentation”](#) on page xix.

Who Should Read This Manual?

Read this manual if you are responsible for configuring or monitoring the GSR and you want to do so using clearVISN CoreWatch rather than using Command Line Interface (CLI) commands.

How to Use This Manual

If You Want To	See
Get an overview of clearVISN CoreWatch	Chapter 1 on page 1
Start clearVISN CoreWatch or familiarize yourself with other basic tasks and the clearVISN CoreWatch interface	Chapter 3 on page 9
Start Configuration Expert and familiarize yourself with its interface	Chapter 4 on page 19
Change system information that is needed before a GSR can be configured	Chapter 5 on page 31
Configure bridging on the GSR	Chapter 6 on page 47
Configure virtual local area networks (VLANs) on the GSR	Chapter 7 on page 69
Configure Internet Protocol (IP) interfaces that you want to use for unicast or multicast routing	Chapter 8 on page 89

If You Want To	See
Configure the GSR for the Routing Information Protocol (RIP)	Chapter 9 on page 113
Configure the GSR for the Distance Vector Multicast Routing Protocol (DVMRP) and Internet Group Management Protocol (IGMP), which IP uses to perform multicast routing	Chapter 10 on page 129
Configure Internet Packet Exchange (IPX) routes on the GSR	Chapter 11 on page 141
Configure Quality of Service (QoS) policies	Chapter 12 on page 161
Configure security on the GSR	Chapter 13 on page 189
Read an overview of OSPF routing and configure OSPF on the GSR	Chapter 14 on page 229
Configure BGP on the GSR	Chapter 15 on page 251
Configure Routing Policies on the GSR	Chapter 16 on page 279
Check the status of the GSR chassis and ports	Chapter 17 on page 347
Monitor real-time performance on the GSR	Chapter 18 on page 353
Display tables that contain bridge information and data about the GSR's VLANs	Chapter 19 on page 373
Display tables that contain information about the routing protocols you are using on the GSR	Chapter 20 on page 379
Obtain information about Layer-2, Layer-3, and Layer-4	Chapter 21 on page 413
Obtain reports that include information clearVISN CoreWatch displays in its tables	Chapter 22 on page 421
Work in clearVISN CoreWatch tables	Appendix A on page 425
Learn about the commands available on each clearVISN CoreWatch menu	Appendix B on page 429
Learn about the regular expressions clearVISN CoreWatch supports	Appendix C on page 439
Obtain information about clearVISN CoreWatch error messages	Appendix D on page 443

Related Documentation

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

For Information About	See the
The DIGITAL GIGAswitch/Router (GSR) features and the procedures for installing the GSR and setting it up for management using clearVISN CoreWatch software.	<i>DIGITAL GIGAswitch/Router Getting Started Guide</i>
How to use Command Line Interface (CLI) commands to configure and manage the GSR	<i>DIGITAL GIGAswitch/Router User Reference Manual</i>
The complete syntax for all CLI commands	<i>DIGITAL GIGAswitch/Router Command Line Interface Reference Manual</i>
SYSLOG messages and SNMP traps	<i>DIGITAL GIGAswitch/Router Error Reference Manual</i>

Correspondence

Documentation Comments

If you have comments or suggestions about this manual, send them to the DIGITAL Network Products.

Attn.: Documentation Project Manager
E-MAIL: doc_quality@lkg.mts.dec.com

Online Services

Further product information is available on the DIGITAL Network Products Business World Wide Web Sites. All sites maintain the same, rich set of up-to-date information on products, technologies and programs.

The Web Sites can be reached from the following geographic locations:

Americas: <http://www.networks.digital.com>
Europe: <http://www.networks.europe.digital.com>
Asia Pacific: <http://www.networks.digital.com.au>

Chapter 1

A Look at clearVISN CoreWatch

DIGITAL clearVISN CoreWatch is a comprehensive, easy-to-use, device management and configuration application for DIGITAL GIGAswitch/Routers (GSRs). Based on Java Programming Language, clearVISN CoreWatch provides configuration, monitoring, and reporting capabilities with the assistance of wizards, dialog boxes, and drag-and-drop operations.

clearVISN CoreWatch provides access to Configuration Expert, a utility that makes tasks such as configuring routers, virtual local area networks (VLANs), application-level Quality of Service (QoS) policies, and security filters simple and easy.

You can run clearVISN CoreWatch in the Solaris, Windows NT, Windows 95, or Windows 98 environments.

This chapter:

- summarizes the features of clearVISN CoreWatch.
- describes the system requirements of clearVISN CoreWatch.
- provides an overview of clearVISN CoreWatch capabilities.
- discusses the different modes of clearVISN CoreWatch.
- provides a list of the Management Information Bases (MIBs) clearVISN CoreWatch supports.

What Are clearVISN CoreWatch's Features?

clearVISN CoreWatch management features include the following:

- Java-based graphical user interface (GUI)
- Simplified bridging configuration
- Simplified routing configuration
- Quality of Service (QoS) policy management

QoS is a set of parameters that assign priorities to different types of traffic, define flows for Internet Protocol (IP) and Internetwork Packet Exchange (IPX) packet fields, assign a precedence to the fields of the flows you define, and establish queuing policies

- Configuration of security filters and access control lists (ACLs). An ACL is a list the GSR keeps to control access to or from various services
- Drag-and-drop VLAN setup and administration
- Detailed reporting in the hypertext markup language (HTML) format

System Requirements

clearVISN CoreWatch can run in the Solaris, Windows NT, Windows 95, and Windows 98 environments. As shown in the following table, clearVISN CoreWatch's system requirements depend upon your operating system. The table identifies which browser to use with each operating system and gives the minimum hardware requirements for each environment.

Table 1. clearVISN CoreWatch system requirements

	Solaris 2.5.1 or 2.6)	Windows NT 4.0x, Windows 95, or Windows 98
Browser	Netscape Navigator 3.0 or above	Netscape Navigator 3.0 or above, or Microsoft Internet Explorer 4.0 or above
CPU	Sparc20 or above	Pentium 133 or above
RAM	128 MB	64 MB
Disk	40 MB Free	20 MB Free

clearVISN CoreWatch Capabilities

clearVISN CoreWatch lets you perform the following operations:

- Access Configuration Expert, which is a DIGITAL utility that you use to configure your GSR as discussed later in this manual.
- Perform the following monitoring tasks on the GSR as discussed later in this manual:
 - Monitor the status of a GSR's ports, modules, power supplies, and other hardware components.
 - Check the status of each bridge table, routing table, and QoS table. These tables contain information that clearVISN CoreWatch obtains from MIBs it supports. (For a list of these MIBs, see [“Which MIBs Does the GSR Support?”](#) on page 4.)
 - Display messages stored in the GSR boot log.
 - SNMP is a protocol that provides support for monitoring and controlling network devices, collecting statistics, and managing configurations, performance, and security. SNMP is mainly used by Transmission Control Protocol/Internet Protocol (TCP/IP) networks. (TCP/IP is a suite of protocols that provide a relatively simple way to connect computers and devices from different vendors on a worldwide internetwork.)

clearVISN CoreWatch displays most monitoring information in tables and reports, but there is some data that is displayed in the form of graphs or dials.

A Look at the Modes of clearVISN CoreWatch

clearVISN CoreWatch can run in the following modes:

- User, which is the mode clearVISN CoreWatch automatically begins operating in after you log in to clearVISN CoreWatch. Use this mode to monitor the activity on the GSR or to obtain reports.
- Configure, which is the mode you use to perform any of the following tasks:
 - Change system information
 - Configure bridging and VLANs
 - Set QoS policies
 - Set security filters
- Configure multiple configuration databases as discussed later in this manual

You switch to the configure mode by starting Configuration Expert. Access to Configuration Expert is controlled by a password that your system administrator can set as discussed in [“Changing the Privileged Password”](#) on page 17.

Which MIBs Does the GSR Support?

clearVISN CoreWatch obtains information from MIBs when displaying the clearVISN CoreWatch tables discussed later in this manual. The GSR supports the following MIBs:

- IETF Standard MIBs:
 - MIB2/RFC 1213 (r/w to system group and to snmpEnableAuthTraps only)
- IETF Draft Standard MIBs:
 - OSPF-MIB/RFC 1850
 - BGP4-MIB/RFC 1657
 - RIPv2-MIB/RFC 1724
 - EtherLike-MIB/RFC 1643
 - BRIDGE-MIB/RFC 1493
- IETF Proposed Standard MIBs:
 - IF-MIB/RFC 1573
 - IP-Group IPCIDRTable-MIB/RFC 2096
- Experimental/Enterprise MIBs:
 - DOT1Q-VLAN-MIB/draft-jeya-vlan-8021q-mib-00.txt
 - IGMP/draft-ietf-idmr-igmp-mib-05.txt
 - DVMRP/draft-thaler-dvmrp-mib-04.txt
 - NOVELL RIP-SAP MIB
 - NOVELL IPX MIB
 - CTRON-YAGO-CONFIG
 - CTRON-YAGO-HARDWARE
 - CTRON-YAGO-L2
 - CTRON-YAGO-L3
 - CTRON-YAGO-SERVICE-STATUS
 - CTRON-YAGO-TRAP

Note: More information on these MIBs see the DIGITAL clearVISN CoreWatch Release Notes.

Chapter 2

clearVISN CoreWatch Installation

Installing clearVISN CoreWatch

You can install clearVISN CoreWatch on a Solaris 2.5.1 or Solaris 2.6 running CDE, Windows NT, Windows 95, or Windows 98 system. The method you use to install clearVISN CoreWatch depends on your environment. Separate discussions on installing clearVISN CoreWatch in the Solaris or Windows environments follow.

Note: clearVISN CoreWatch requires CDE to run properly on Solaris 2.5.1 and 2.6 operating systems. Ensure that your Solaris system includes CDE before attempting to run clearVISN CoreWatch.

Installing on a Solaris System

To install clearVISN CoreWatch from a CD onto a Solaris 2.5.1 or 2.6 system running CDE:

1. If you plan to integrate clearVISN CoreWatch with HP OpenView, be sure the HP OpenView daemon is running. For details, see your HP OpenView documentation.
2. Insert the clearVISN CoreWatch CD into your CD-ROM drive.
3. Log in as super user by entering the following command:

```
su - root
```

4. Ensure that you are in the appropriate subdirectory to access the CD-ROM by entering the following command:

```
cd /cdrom/cdrom0
```

5. Run the clearVISN CoreWatch installation script by entering the following command:

```
install.sh
```

clearVISN CoreWatch is installed on your system in the `/opt/CScw` directory.

6. Add `/opt/CScw/bin` to your environment path.

For details on adding items to a path, see your Solaris documentation.

Installing on a Windows NT, Windows 95, or Windows 98 System

Note: You must have Admin privileges to install clearVISN CoreWatch on a Windows NT system.

To install clearVISN CoreWatch on a Windows NT, Windows 95, or Windows 98 system:

1. If you plan to integrate clearVISN CoreWatch with HP OpenView on a Windows NT system, be sure the HP OpenView daemon is running. For details, see your HP OpenView documentation.
2. Insert the clearVISN CoreWatch CD into your CD-ROM drive.
3. After the clearVISN CoreWatch installation wizard appears, click **Next**.
4. After reviewing the license agreement, click **Yes** to accept it.
5. Enter your name and your company's name in the appropriate text boxes. Then click **Next**.
6. Specify the folder in which you want to install the software and click **Next**.
You can keep the default folder or click **Browse** and then browse to another folder.
7. Set up the type of installation by doing one of the following:
 - Choose **Typical** to install the most common options.
 - Choose **Compact** to install the minimum files needed to run clearVISN CoreWatch.
 - Choose **Custom** and click **Next** if you are an advanced user and want to specify which files to install. Options with a check mark will be installed. Click to the left of an item to select or clear its check box.
8. Click **Next**.
9. Specify a name for the clearVISN CoreWatch program group, which is DIGITAL clearVISN CoreWatch by default. Then click **Next**.

10. When the browser window reappears, close it.
11. Specify whether you want to view the clearVISN CoreWatch readme file, then click **Finish**.

Options with a check mark will be performed. Click to the left of an item to select or clear its check box.

Chapter 3

Learning clearVISN CoreWatch Basics

Before using clearVISN CoreWatch, you should be familiar with some basic clearVISN CoreWatch tasks and be familiar with the application's interface. This chapter

- discusses starting clearVISN CoreWatch.
- provides an overview of the clearVISN CoreWatch interface.
- discusses changing clearVISN CoreWatch passwords.
- discusses how to access the clearVISN CoreWatch online help.
- explains how to exit clearVISN CoreWatch.

For information on installing clearVISN CoreWatch in Solaris and Windows environments, see [Chapter 2: “clearVISN CoreWatch Installation” on page 5](#).

Starting clearVISN CoreWatch

The method you use to start clearVISN CoreWatch depends on whether you installed it in the Solaris or Windows environment. If you choose to integrate clearVISN CoreWatch with SPECTRUM or HP OpenView during installation, you can start clearVISN CoreWatch from within either system in both Solaris and Windows NT/Windows 95/Windows 98.



Caution: Before starting clearVISN CoreWatch in any environment, be sure that you can send a ping packet to the GSR and that the GSR is configured for SNMP. For details on configuring SNMP on the GSR, see the *DIGITAL GIGAswitch/Router User Reference Manual*.

Separate discussions on starting clearVISN CoreWatch in the Solaris and Windows environments and from within SPECTRUM or HP OpenView follow.

Starting clearVISN CoreWatch in Solaris

Note: clearVISN CoreWatch requires CDE to run properly on Solaris 2.5.1 and 2.6 operating systems. Ensure that your Solaris system includes CDE before attempting to run clearVISN CoreWatch.

To start clearVISN CoreWatch in the Solaris 2.5.1 or 2.6 environment:

1. Enter the following command at the Solaris prompt:

```
CoreWatch
```

The **Login Dialog** dialog box appears.

Note: If the clearVISN CoreWatch command is not found, you can locate it in `/opt/CScw/bin`.

2. Type the name or IP address and community string for the GSR. If you do not know this information, see your network administrator.
3. Click **OK**.

For details on the window that clearVISN CoreWatch opens, see [“A Look at the clearVISN CoreWatch Interface” on page 12](#).

Starting clearVISN CoreWatch in Windows NT, Windows 95, or Windows 98

To start clearVISN CoreWatch in the Windows NT, Windows 95, or Windows 98 environment:

1. Choose the **Start** menu, select **Programs**, select DIGITAL clearVISN CoreWatch, and then choose clearVISN CoreWatch. The **Login Dialog** dialog box appears.

Note: If you installed the program in a startup folder other than **Programs > DIGITAL clearVISN CoreWatch**, select that folder from the Start menu and then select clearVISN CoreWatch.

2. Type the name or IP address and community string for the GSR. If you do not know this information, see your network administrator.
3. Click **OK**.

For details on the window that clearVISN CoreWatch opens, see the [“A Look at the clearVISN CoreWatch Interface” on page 12](#).

Starting clearVISN CoreWatch from within SPECTRUM Enterprise Manager

SPECTRUM Enterprise Manager is Cabletron System's flexible and scalable network management platform based on leading-edge, object-oriented, artificial intelligence technology. SPECTRUM, which is available on Solaris and Windows NT, provides a suite of bundled applications as well as additional optional applications. The GSR is modeled in SPECTRUM using the GIGASwRtr model type. The GSR can be Auto-Discovered or manually created in a SPECTRUM Topology View and then copied to an Organization and/or Location View.

To start clearVISN CoreWatch from within SPECTRUM:

1. Start SPECTRUM.
2. If you know the topology location for your GIGASwRtr model, proceed to that location. Otherwise, open the Find View by choosing the **View** menu, selecting **New View**, and then selecting **Find**. Select **Model-Type Name** and enter the **GIGASwRtr** command to display all the GIGASwRtr models or select **Network Address** to display a particular model.
3. Bring up the menu for the GIGASwRtr model and select clearVISN CoreWatch.

This starts clearVISN CoreWatch using the GIGASwRtr model's network address and community name. For details on the window that clearVISN CoreWatch opens, see ["Front Panel View" on page 12](#).

Starting clearVISN CoreWatch from within HP OpenView 5.x

HP OpenView 5.x is network node management software for the Solaris and Windows NT environments. If HP OpenView is integrated with clearVISN CoreWatch, you may use HP OpenView to start clearVISN CoreWatch and recognize your GSRs. HP OpenView is automatically integrated with clearVISN CoreWatch when you install clearVISN CoreWatch while the HP OpenView daemon is running.

To start clearVISN CoreWatch from within HP OpenView:

1. Start HP OpenView.
2. Click a network node.
3. Select the **Misc** menu and then choose clearVISN CoreWatch. The **Login Dialog** dialog box appears.
4. Type the name or IP address and community string for the GSR. If you do not know this information, see your network administrator
5. Click **OK**.

For details on the window that clearVISN CoreWatch opens, see ["Front Panel View" on page 12](#).

A Look at the clearVISN CoreWatch Interface

clearVISN CoreWatch offers two views of the GSR and runs in different modes that you should be familiar with before using clearVISN CoreWatch. You may also find it helpful to know how to use the clearVISN CoreWatch Toolbar before using clearVISN CoreWatch. Separate discussions on each clearVISN CoreWatch view, its modes, and the clearVISN CoreWatch Toolbar follow.

Front Panel View

After you start clearVISN CoreWatch, a Front Panel view similar to the following appears:

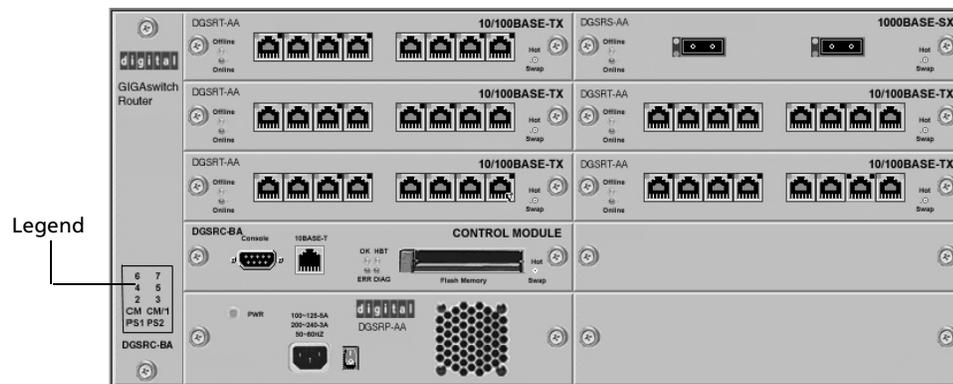


Figure 1. Front Panel view (GSR-8)

The Front Panel view is a graphical representation of a GSR-8's front-panel chassis. You can use this view to do the following:

- Obtain port statistics
- Configure ports
- Check the status of ports

The legend identified in the previous figure includes the abbreviations described in the following table:

Table 2. Legend abbreviations

Abbreviation	Description
PS1	Identifies the location of the GSR's main power supply.
PS2	Identifies the location of the GSR's redundant power supply.
CM	Identifies the location of the GSR's Control Module.
CM/1	Indicates the location of slot 1 of the GSR chassis. Slot 1 can accept either the GSR's Control Module or any module the GSR supports.
2	Indicates the location of slot 2 of the GSR chassis.
3	Indicates the location of slot 3 of the GSR chassis.
4	Indicates the location of slot 4 of the GSR chassis.
5	Indicates the location of slot 5 of the GSR chassis.
6	Indicates the location of slot 6 of the GSR chassis.
7	Indicates the location of slot 7 of the GSR chassis.

A Look at the Modules

In the Front Panel view, a GSR's modules appear similar to the following figure. This figure is for an Ethernet 10/100BASE-TX module, but the information clearVISN CoreWatch displays to represent a module depends on that module's type.



Figure 2. Front Panel view of GSR's ethernet 10/100BASE-TX module

Each port on an Ethernet 10/100BASE-TX module is represented by an object similar to the following figure. The upper-left light-emitting diode (LED) and the center LED of a port indicates whether the port is online.



Figure 3. Ethernet 10/100BASE-TX port

Each port on a Gigabit module (1000BASE-SX or 1000BASE-LX) is represented by an object similar to the following figure. The bottom LED of a port indicates whether the port is online.



Figure 4. Gigabit module (1000BASE-SX or 1000BASE-LX) port

The Front Panel view also includes online and offline LEDs similar to those on the physical chassis. Unlike their physical counterparts on the actual chassis, these LEDs do not change in the Front Panel view whenever a module goes online or offline.

Schematic View

The Schematic view, which looks similar to the following figure, is a graphical representation of a GSR's functions (such as bridging, switching, and routing services) and data objects (such as QoS flows). It also indicates which functions are active, inactive, or in error. The information in the Schematic view is updated every 30 seconds. The legend that appears at the bottom of the Schematic view indicates the scheme used to represent the various items displayed in that view.

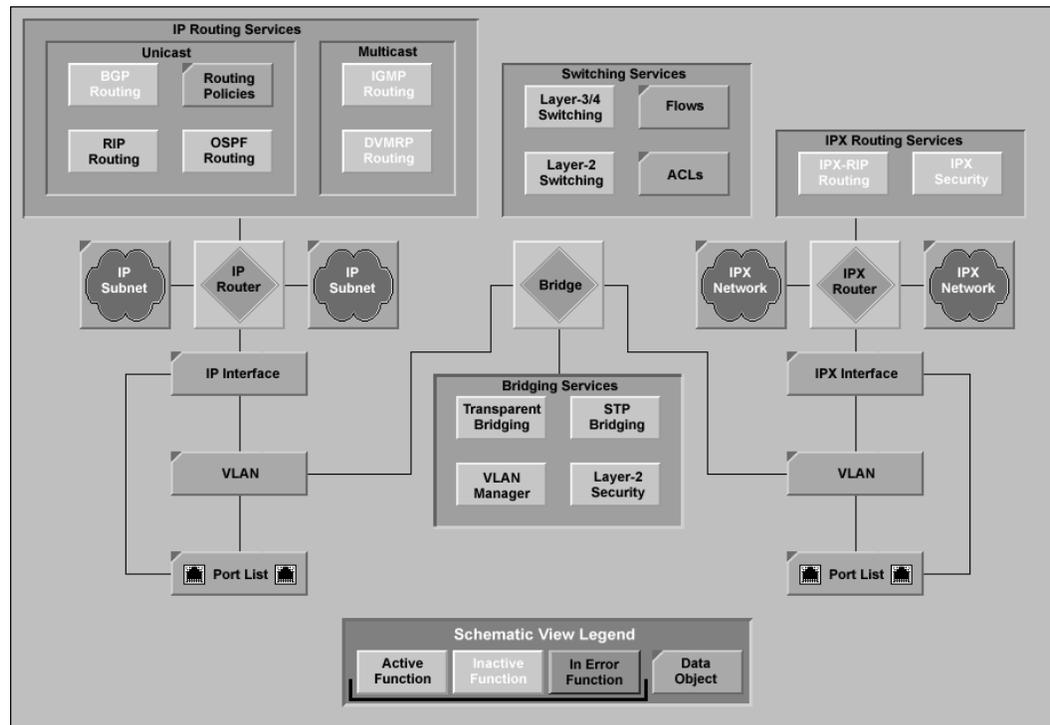


Figure 5. clearVISN CoreWatch Schematic view

Opening the Schematic View

To open the Schematic view, do one of the following:

- Select the **File** menu and then choose **Open Schematic View**.
- Click the **Open Schematic View** button  on the clearVISN CoreWatch toolbar.

Using the Schematic View

You can use the Schematic view to display the clearVISN CoreWatch tables and dials discussed later in this manual, and configure QoS flows and security filters. To do so, perform one of the following operations:

- Move the cursor to an object, click the right-mouse button, and choose a command from the menu that appears.
- Double-click a function or data object. clearVISN CoreWatch then performs the default task for that function or data object.

clearVISN CoreWatch Menus

The clearVISN CoreWatch menus are located at the top of the clearVISN CoreWatch main window. Use the commands available on these menus to perform tasks in clearVISN CoreWatch. For a description of each menu command, see [Appendix B: “clearVISN CoreWatch Menus”](#) on page 429.

clearVISN CoreWatch Toolbar

The clearVISN CoreWatch toolbar is a set of buttons located at the top of the clearVISN CoreWatch window. Clicking buttons in this toolbar performs some clearVISN CoreWatch tasks quickly. The following figure identifies the function of each clearVISN CoreWatch toolbar button. Details on the Schematic View, System Dashboard, Configuration Expert, Port Utilization Summary, reports, and online help are discussed later in this manual.

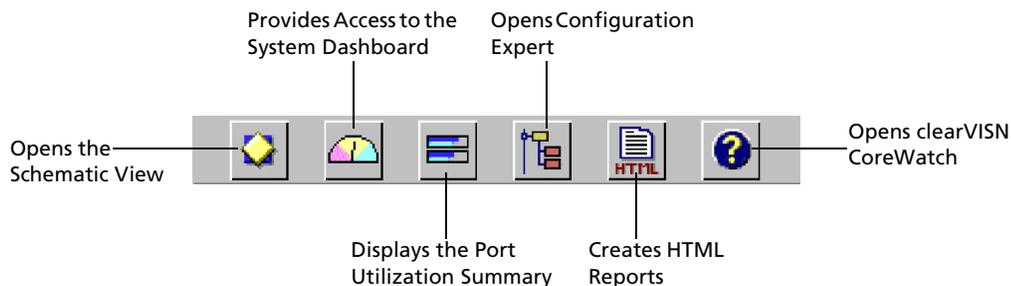


Figure 6. clearVISN CoreWatch toolbar buttons

clearVISN CoreWatch Passwords

If you can access Configuration Expert, you can change the following clearVISN CoreWatch passwords:

- Login Password, which is the password you are prompted for when you start clearVISN CoreWatch.
- Privileged Password, which is the password you are prompted for when you start Configuration Expert. This password logs you in to Configuration Expert so that you can then use that utility to configure your GSR.

Changing the Login Password

If you can access Configuration Expert, you can change your clearVISN CoreWatch Login password. To change your Login password:

1. Select the **Configure** menu and choose **Change Login Password**.

2. If you are prompted for the Privileged password, enter it to start Configuration Expert. Otherwise, skip to [step 3](#).

You are prompted for the Privileged password if you are not running clearVISN CoreWatch in Configure mode.
3. After the **Change Login Password** form appears, enter your current Login password in the **Old Password** text box.
4. Enter your new password in the **New Password** and **Re-enter New Password** text boxes.
5. Click **OK**.

Changing the Privileged Password

If you can access Configuration Expert, you can change the Privileged password. To change the Privileged password:

1. Select the **Configure** menu and choose **Change Privileged Password**.
2. If you are prompted for the Privileged password, enter it to start Configuration Expert. Otherwise, skip to [step 3](#).

You are prompted for the Privileged password if you are not running clearVISN CoreWatch in the configure mode.
3. After the **Change Privileged Password** form appears, enter your current Privileged password in the **Old Password** text box.
4. Enter your new password in the **New Password** and **Re-enter New Password** text boxes.
5. Click **OK**.

Accessing Help

When using clearVISN CoreWatch, you can access online help by choosing commands from the **Help** menu or clicking the **Help** button. If you click the **Help** button, clearVISN CoreWatch displays help specific to the form, dialog box, or other item you are currently using. The following table describes the commands you can choose from the **Help** menu.

As the table indicates, you can choose some commands directly from the **Help** menu and other commands from the **DIGITAL Web Site** submenu.

Table 3. clearVISN CoreWatch Help menu commands

	Command	Description
	Contents and Index	Opens clearVISN CoreWatch Online Help.
	Glossary	Opens the online help glossary for clearVISN CoreWatch.
	Release Note	Displays the release note(s) for your version of clearVISN CoreWatch.
DIGITAL Web Site submenu	Product News	Displays information about DIGITAL products.
	White Papers	Provides access to technical discussions of DIGITAL products.
	Frequently Asked Questions	Provides technical support for some issues or concerns that you may have while using clearVISN CoreWatch.
	Technical Support	Displays information about how to contact DIGITAL technical support.
	Send Feedback	Displays a form that you may use to let DIGITAL know what you think about its products. You can complete and send this form online.
	DIGITAL Home Page	Displays the DIGITAL home page in your Web browser.
	About clearVISN CoreWatch	Displays information about which version of clearVISN CoreWatch you are using.

Exiting clearVISN CoreWatch

To exit clearVISN CoreWatch, select the **File** menu and choose **Exit**. clearVISN CoreWatch prompts you to verify that you want to exit. Click the **Quit** button.

Chapter 4

Learning Configuration Expert Basics

You can use DIGITAL's Configuration Expert to produce configuration files for a GSR. Once you have set up a configuration file, you can load it into the GSR. This chapter:

- provides an overview of Configuration Expert.
- explains how to start Configuration Expert.
- discusses the Configuration Expert window.
- describes the different configuration files Configuration Expert uses.
- explains the purpose of Configuration Expert's wizards and dialog boxes.
- discusses finding, copying, and deleting objects.
- explains how to save and apply your changes.
- discusses exiting Configuration Expert.

If you prefer to produce a configuration file for your GSR using commands instead of Configuration Expert, see the *DIGITAL GIGAswitch/Router Command Line Interface Reference Manual*.

What Is Configuration Expert?

You can use Configuration Expert to perform the following tasks:

- Change system information

- Configure bridging polices
- Configure VLANs
- Configure IP and IPX routing
- Configure multicast routing
- Set QoS policies
- Set ACLs and security filters
- Configure multiple configuration files on the GSR

Like clearVISN CoreWatch, Configuration Expert is a Java-based GUI. This GUI offers drag-and-drop setup and administration for VLANs and ACLs. It also supports comprehensive configuration through wizards and dialog boxes.

Starting Configuration Expert

You can start Configuration Expert from the Front Panel view or independently in the Solaris or Windows operating environments. Separate discussions on starting Configuration Expert from the different environments follow.

Starting Configuration Expert from the Front Panel View

To save time navigating through your operating system, you can open Configuration Expert while in the clearVISN CoreWatch front panel view.

To open Configuration Expert from the Front Panel view:

1. Select the **File** menu and choose **Open Configuration Expert**.
2. If the **Configuration Expert Login** dialog box appears, enter the Privileged password and click **OK**. The Configuration Expert window appears.

For details on this window, see [“A Look at the Configuration Expert Window”](#) on page 21.

Starting Configuration Expert in Solaris

Note: clearVISN CoreWatch requires CDE to run properly on Solaris 2.5.1 and 2.6 operating systems. Ensure that your Solaris system includes CDE before attempting to run clearVISN CoreWatch.

To start Configuration Expert in the Solaris 2.5.1 or 2.6 environment:

- Enter the following command at the Solaris prompt:

```
ConfigExpert
```

The Configuration Expert window appears.

Note: If the ConfigExpert command is not found, you can locate it in `/opt/CScw/bin`.

For details on this window, see [“A Look at the Configuration Expert Window” on this page](#).

Starting Configuration Expert in Windows NT, Windows 95, or Windows 98

To start Configuration Expert in the Windows NT, Windows 95, or Windows 98 environment:

- Choose the **Start** menu, select **Programs**, select DIGITAL clearVISN CoreWatch, and then choose **ConfigExpert**.

Note: If you installed the program in a startup folder other than **Programs > DIGITAL clearVISN CoreWatch**, select that folder from the Start menu and then select **ConfigExpert**.

For details on this window, see [“A Look at the Configuration Expert Window” next](#).

A Look at the Configuration Expert Window

When you start Configuration Expert, a window similar to the following appears:

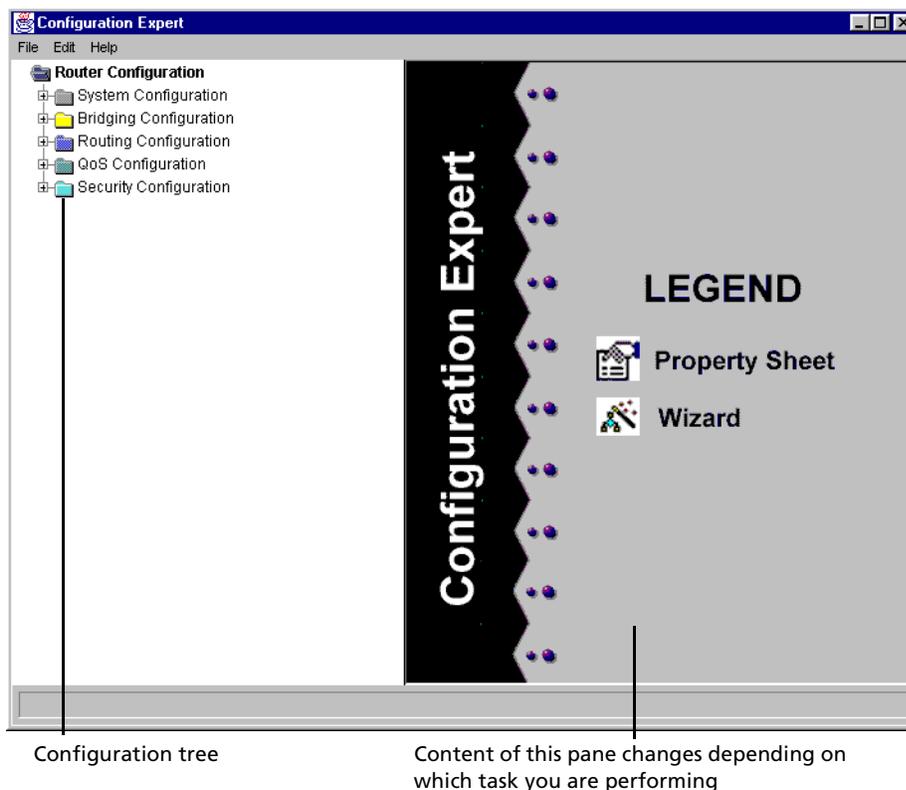


Figure 7. Configuration Expert window

The left pane of the Configuration Expert window includes the configuration tree, which you use to navigate to the objects (modules, ports, system, bridging, routing, and so on) you want to create, edit, or delete while configuring the GSR. For details on the configuration tree, see [“Configuration Tree”](#) next.

The contents of the right pane changes depending on what task you are performing in Configuration Expert. Initially the right pane contains the legend that describes some of the Configuration Expert icons. As you configure the GSR, the right pane will contain the wizard or dialog box necessary for the configuration task you are performing.

Configuration Tree

The configuration tree is a graphical representation of the GSR’s configuration. When you are configuring the GSR, you can use the configuration tree to view the contents of the configuration files so you can add, edit, or delete objects in those files.

You can view the contents of a configuration file by double-clicking its name or icon in the configuration tree. When you do so, Configuration Expert displays the file’s subtree. As shown in the following figure, a configuration file’s subtree includes system, bridging,

routing, QoS, and security configuration objects. There are separate subtrees for each configuration file.

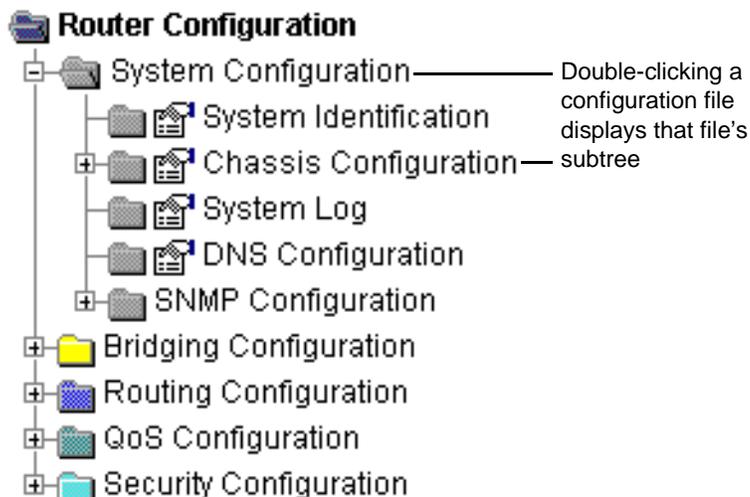


Figure 8. Configuration tree

You navigate a subtree by double-clicking its objects or clicking the plus sign (+) of objects. As you do so, Configuration Expert expands the subtree to display the object's contents. In the following figure, the System Configuration subtree has been expanded to display the contents of the Chassis Configuration object and the Module-1 object. As the figure shows, a plus sign (+) indicates an object can be expanded further and a minus sign (-) indicates an object cannot be expanded. You can close an expanded object by double-clicking it or clicking its minus sign (-).

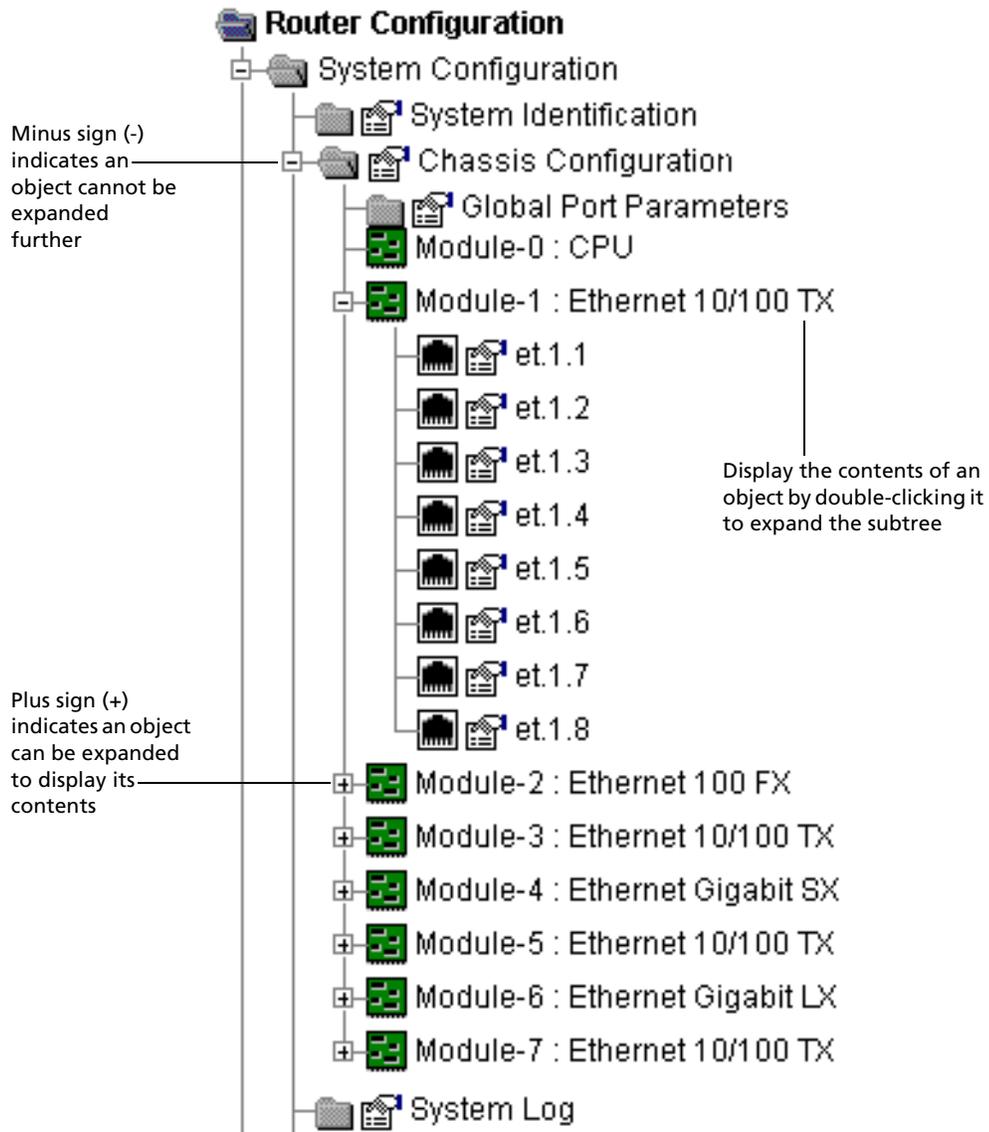


Figure 9. Expanded configuration tree

While configuring the GSR, you can continue double-clicking subtree objects until you are able to select the object you want to add, edit, or delete. After you select a configurable object, the wizard or dialog box used to configure that object appears in the right pane of the Configuration Expert window.

A Look at the Configuration Tree Icons

The configuration tree uses the icons described in the following table to represent configuration tree objects, to indicate that there is a wizard or dialog box associated with an object, and to identify configuration file changes and errors:

Table 4. Configuration tree icons and descriptions

Icon	Description
	Represents a configuration file or configurable object, such as an interface, VLAN, QoS policy, ACL, and so on. Double-clicking an object's folder icon displays the contents of that object.
	Represents a chassis module. Double-clicking a module icon displays that module's ports. However, this is not the case with the Control Module icon because there are no ports associated with the Control Module.
	Represents a port.
	Indicates that a dialog box is associated with an object. You can configure the object by using its dialog box.
	Indicates that a wizard is associated with an object. You can configure the object by using its wizard.

Wizards and Dialog Boxes

You can configure objects through wizards and dialog boxes, which are forms that prompt you for the appropriate configuration information.

Wizards take you step-by-step through the process of adding objects that are complex to configure. If you add an object using a wizard, there will be a wizard icon  next to that object.

Property sheets are used to add easily configurable objects and also to modify the configuration settings of objects once they have been added. If there is a dialog box associated with an object, there will be a dialog box icon  next to that object.

Copying Configuration Settings with Drag-and-Drop

You can drag objects to copy their configuration settings. When using this method to copy ACLs or QoS profiles, you should remember that you drag ACLs to interfaces, but you drag interfaces to QoS profiles.

Note: Dragging an object from one location to another copies the object's configuration settings. It does not move the object.

To drag objects to copy their configuration settings:

1. Navigate through the configuration tree until you locate the object possessing the configuration settings you wish to copy.
2. Navigate through the configuration tree until you locate the object to which you want to apply the configuration settings.
3. Select the object you want to copy, then drag that object to the object to which you want to apply the configuration settings.

Configuration Expert displays a folder icon while you are dragging an object. A green check mark appears whenever the cursor is on an object to which the configuration settings can be applied.

4. Release the mouse button to apply the configuration settings to the desired object.

Finding Objects

You can search for an object in the configuration tree by entering all or part of the name of an object you want to find. To find an object in Configuration Expert using a text-string search:

1. Select the **Edit** menu and choose **Find Item**. A Find Item dialog box similar to the following appears:

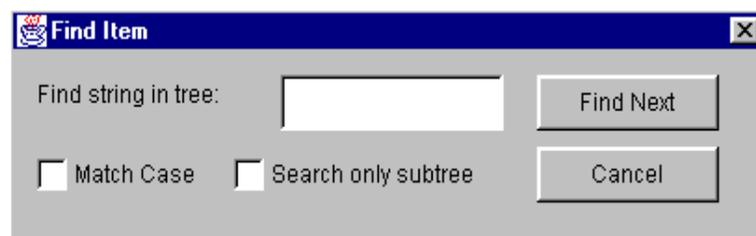


Figure 10. Find Item dialog box

2. Enter the string of text you would like to search for in the **Find String in tree** box.
3. If the text you are searching for is case sensitive, turn on the *Match Case* option.

4. If you wish to limit the range of your search to encompass only the immediate subtree, turn on the *Search only subtree* option.
5. Click the **Find Next** button. Configuration Expert searches the current configuration file (or the specified portion of it) for the text string you entered in [step 2](#).

Deleting Objects

You can delete any object you add except for VLANs. To delete an object:

1. Select the object you want to delete.
2. Select the File menu and choose **Delete Selection**.
3. Click **OK** when Configuration Expert prompts you to verify whether you want to delete the object.

Order of Configuration Tasks

Perform configuration tasks in the following order:

1. Provide system information (such as the GSR's name and location), define the GSR's ports, set up the GSR for a SYSLOG server and Domain Naming System (DNS) servers, and configure the GSR for SNMP traps. For details on performing these tasks, see [Chapter 5: "Changing System Settings" on page 31](#).
2. Configure bridging on the GSR as discussed in [Chapter 6: "Configuring GSR Bridging" on page 47](#).
3. Group physical ports on the GSR by configuring VLANs as discussed in [Chapter 7: "Configuring VLANs on the GSR" on page 69](#).
4. If you are using the GSR in an IP environment, create IP interfaces and then configure those interfaces for unicast or multicast routing.

For details on IP interfaces, see [Chapter 8: "Configuring IP Interfaces for the GSR" on page 89](#). For details on unicast routing, see [Chapter 9: "Configuring Unicast Routing on the GSR" on page 113](#). For details on multicast routing, see [Chapter 10: "Configuring Multicast Routing on the GSR" on page 129](#).

5. If you are using the GSR in an IPX environment, configure IPX interfaces as discussed in [Chapter 11: "Configuring the GSR for IPX Routes" on page 141](#).
6. If you want to control traffic on the GSR, configure the GSR's QoS policies as discussed in [Chapter 12: "Configuring QoS on the GSR" on page 161](#).
7. If you want to set security on the GSR, define ACLs and Layer-2 security filters as discussed in [Chapter 13: "Configuring Security on the GSR" on page 189](#).
8. If you want to set OSPF on the GSR, configure the GSR's OSPF behavior as discussed in [Chapter 14: "Configuring OSPF on the GSR" on page 229](#).

Saving and Applying Your Configuration Changes

When you finish working in Configuration Expert you will have to save your changes in the form of a configuration file, then load them into the GSR.

Saving Changes to a Configuration File:

To save your configuration changes:

1. Select the **File** menu and choose **Save As**. The **Save Configuration** dialog box opens.



Figure 11. Save Configuration dialog box

2. Navigate to the directory in which you wish to store your configuration file and do one of the following:
 - Select on one of the existing configuration files and click the **Save** button to overwrite it.
 - Type the name of the new configuration file you wish to add to the directory and click the **Save** button.

Loading a Configuration File into a GSR

Once you have created a new configuration file, you can use Configuration Expert to load it into your GSR as both the startup and active configuration file.

To load a configuration file into the GSR:

1. Select the **File** menu and choose **Apply Config**. An Apply Configuration dialog box similar to the following appears:

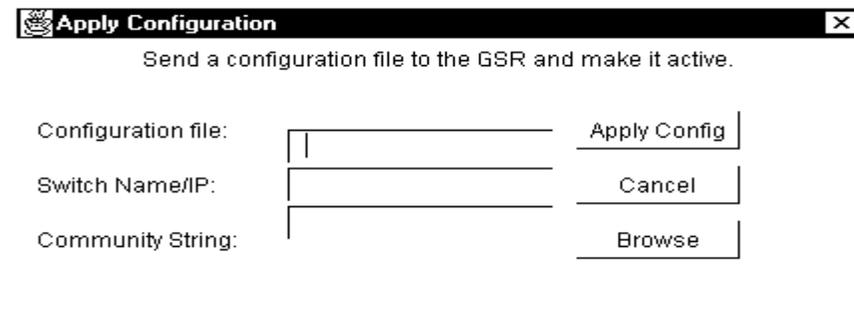


Figure 12. Apply Configuration dialog box

2. Do one of the following:
 - If you are already in the appropriate directory, enter the name of the configuration file you wish to apply to the GSR in the **Configuration file** box.
 - Choose the **Browse** button, navigate to the directory containing the configuration file you wish to apply to the GSR, and select it.
3. Enter the IP address and community string for the GSR in the **Switch Name/IP** box and **Community String** boxes, respectively.
4. Choose the **Apply Config** button. Configuration Expert loads the specified configuration file into the GSR as both the Startup and Active configuration files.

Retrieving a Configuration File from a GSR

Just as you can load configuration files into a GSR, you can retrieve the active configuration file from a GSR, as well.

To retrieve a configuration file from a GSR:

1. Select the **File** menu and choose **Retrieve Config**. A Retrieve Configuration dialog box similar to the following appears:

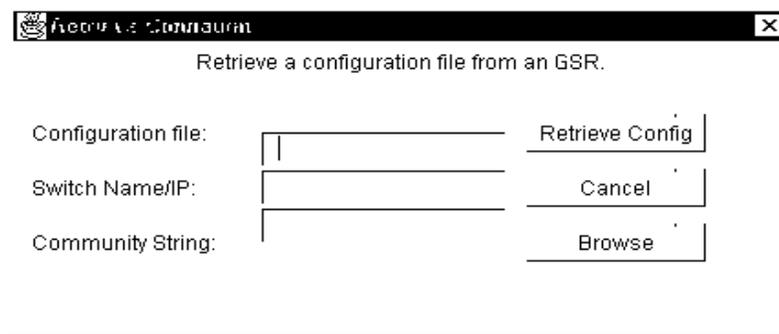


Figure 13. Retrieve Configuration dialog box

2. Enter the name of the configuration file you wish to retrieve from the GSR in the **Configuration file** box.
3. Enter the IP address and community string for the GSR in the **Switch Name/IP** box and **Community String** boxes, respectively.
4. Choose the **Retrieve Config** button. Configuration Expert retrieves the specified configuration file and opens it in the main Configuration Expert window.

Exiting Configuration Expert

To exit Configuration Expert and return to clearVISN CoreWatch, click the  button in the upper right of the Configuration Expert window.

If Configuration Expert prompts you to commit your changes because you have not already done so, click **Yes**. Then commit your changes as discussed in [“Saving and Applying Your Configuration Changes”](#) on page 28.

Chapter 5

Changing System Settings

You change system settings through the System Configuration object of a configuration file. This chapter discusses using Configuration Expert to perform the following tasks:

- Providing system information to set the GSR's name, identify who users should contact regarding the GSR, and indicate the GSR's location.
- Configuring a GSR's chassis
- Configuring a GSR's ports.
- Configuring the GSR to send system messages to a SYSLOG server.
- Configuring the GSR for Domain Naming System (DNS) servers.
- Setting up targets for SNMP traps and establishing SNMP community strings for those traps.

Providing System Information

Before setting up the GSR, you may want to set the GSR's name, provide information about who users should contact regarding the GSR, and indicate where the GSR is located. To enter this system information:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's System Configuration object.
3. Select the System Identification object.

A **System ID** dialog box similar to the following appears:

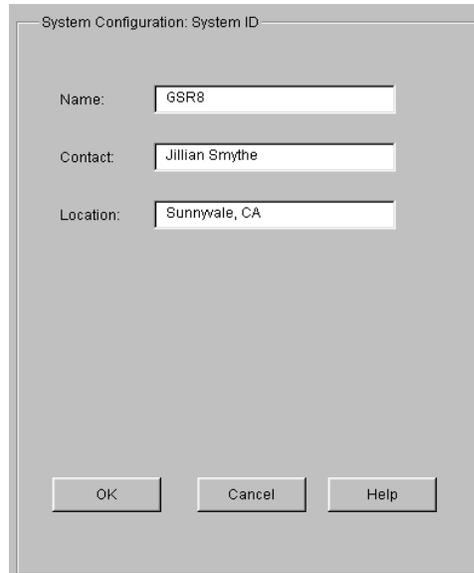


Figure 14. System ID dialog box

4. Enter the GSR's name, the name of the GSR administrator, and the location of the GSR in the appropriate text boxes.
5. Click **OK**.

Configuring a GSR Chassis

One of the first tasks you must perform when you create a new configuration file for your GSR is to properly configure the chassis, which will be the foundation of your router's configuration file. Using Configuration Expert, you can determine which modules take up which slots in your router and so on.

To configure the chassis for your GSR:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's System Configuration object.
3. Double-click the Chassis Configuration object. A **Chassis Configuration** dialog box similar to the following appears in the right hand frame of the Configuration Expert window:

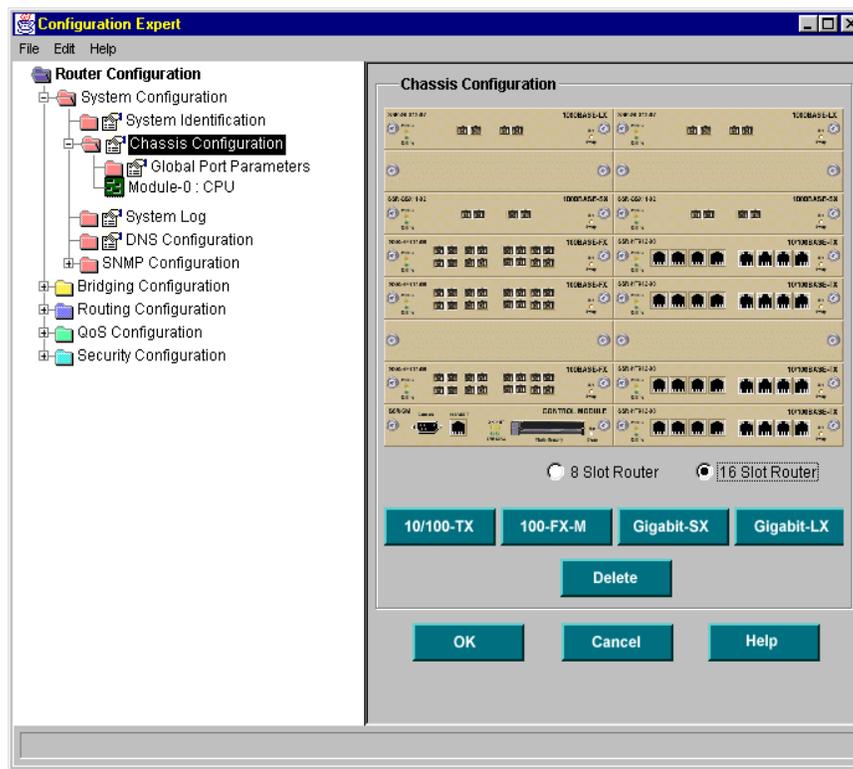


Figure 15. Chassis Configuration dialog box

4. Specify whether you are setting up a configuration file for a 8-slot, or 16-slot router by selecting the appropriate option.
5. Specify which type of module you want to configure for a given slot by taking the following steps:
 - a. Click on one of the router's empty (blank) slots to select it.
If you accidentally select a slot other than the one you intended, simply click it again to deselect it.
 - b. Click one of the four module type buttons to assign that module to the slot you selected in [step a](#). The four currently available module types are: 10/100-TX, 100-FX-M, Gigabit-SX, and Gigabit-LX.
If you accidentally assign a module other than the one you intended to a given slot, you can reverse that assignment by clicking on the slot and then choosing the Delete button.
 - c. Repeat step a and step b until you have configured all of the necessary slots for your router.
6. Click **OK**.

Configuring Ports

Configuration Expert lets you enable and disable ports as well as configure the following characteristics:

- Operating mode (half-duplex or full-duplex).
In half-duplex mode, a port can transmit data in only one direction at a time between two stations.
In full-duplex mode, a port can simultaneously send and receive data.
- Port speed (10-Mbps or 100-Mbps). This parameter applies only to ports on the 10/100 modules.
- Auto negotiation (for Gigabit Ethernet ports)
- Hash mode, which controls the distribution of flow entries in Layer-2 and Layer-3 lookup tables. (For more information on hash modes, see [“Configuring Global Settings on All Ports” on page 34.](#))

Note: DIGITAL has configured the GSR’s hash mode to its optimal setting. It is recommended you not change the hash mode unless advised to do so by DIGITAL Technical Support.

You can configure global settings for all GSR ports as well as configure individual ports. If most of a GSR’s ports are to be configured with the same or similar settings, you can first apply global settings to all ports and then modify those settings on individual ports as needed. Separate discussions on configuring default settings for all ports and configuring an individual port follow.

Configuring Global Settings on All Ports

To configure global settings on all of the GSR’s ports:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file’s System Configuration object.
3. Double-click the Chassis Configuration object.
4. Click the Global Port Parameters object.

A **Global Attributes of Ports** dialog box similar to the following appears:

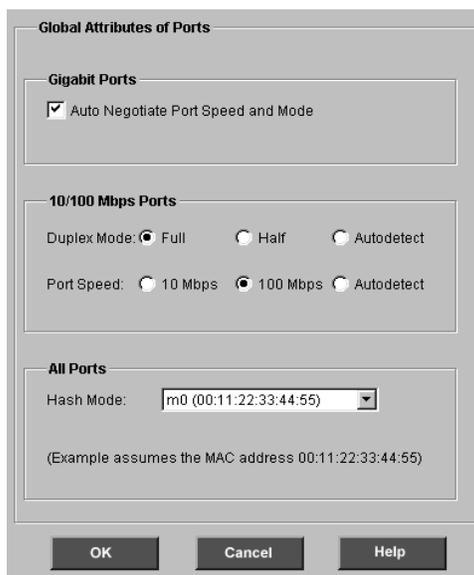


Figure 16. Global Attributes of Ports dialog box

5. If you want all Gigabit Ethernet ports to detect and then use the operating mode and speed of the network segment to which those ports are connected, select the *Auto Negotiate Port Speed and Mode* check box. Otherwise, clear the check box to disable auto negotiation on all Gigabit Ethernet ports.

All Gigabit Ethernet ports use auto negotiation. If auto negotiation is not set on a Gigabit Ethernet port, that port uses the full-duplex mode and operates at 1000 Mbps.

6. If you want to configure 10/100 Mbps Ethernet ports, do one of the following. Otherwise, skip to [step 7](#).
 - If you want all 10/100 Mbps Ethernet ports to detect and then use the operating mode or speed of the network segment to which the port is connected, select the appropriate *Autodetect* options.
 - If you want to set all 10/100 Mbps Ethernet ports' operating mode or speed to a specific setting, select the appropriate buttons.

All 10/100 Mbps Ethernet ports use autodetection by default.

The 100 FX ports do not use auto-detection. They always use full-duplex mode and operate at 100 Mbps.

Note: If you select an operating mode or speed, the setting disables autodetection for that parameter on the port. For example, if you set the speed of a segment to 10 Mbps, that segment no longer uses autodetection for the port speed and will always attempt to operate at 10 Mbps.

7. Set the Layer 2 hash mode for all ports. The hash mode controls the distribution of flow entries in Layer-2 and Layer-3 lookup tables. Assuming a MAC address of the value 0011:2233:4455, the following list describes the various hash modes.
 - M0 – 0011:2233:4455
 - M1 – 0011:2233:5544
 - M2 – 0011:3322:4455 (default hash mode)
 - M3 – 1100:2233:4455

Note: DIGITAL has configured the GSR's hash mode to its optimal setting. It is recommended you not change the hash mode unless advised to do so by DIGITAL Technical Support.
8. Click **OK**.

Configuring an Individual Port

Configure an individual port if you want to enable or disable that port or set the port's physical characteristics. To configure an individual port:

1. Start Configuration Expert if you have not already done so.

Note: If you open Configuration Expert by clicking a port in the Front Panel view, Configuration Expert assumes you want to modify that port in the Active Configuration file and automatically opens that port's dialog box. If you are modifying that port in the Active Configuration file, go to [step 6](#).
2. Open the configuration file you want to modify and then double-click that file's System Configuration object.
3. Double-click the Chassis Configuration object.

A list of modules similar to the following appears.

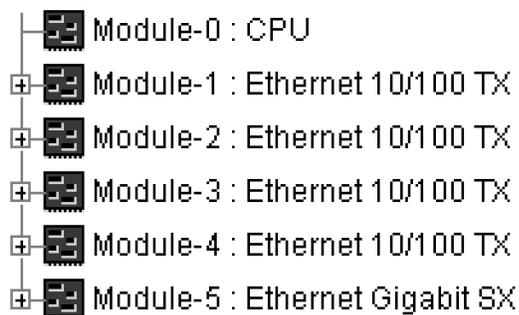


Figure 17. Sample list of modules

4. Double-click the module on which the port you want to configure is located.

The module's port list appears. The number of ports in the list depends on the module type.

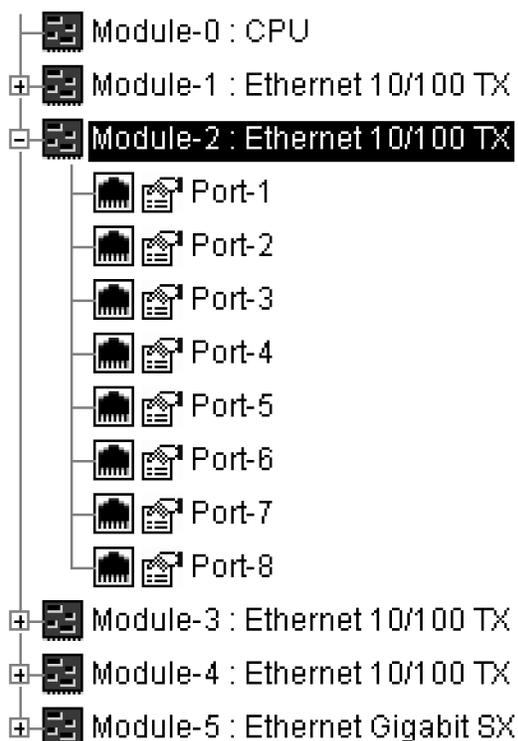


Figure 18. Sample port list

5. Select the port you want to configure.

A **Physical Attributes of Port** dialog box similar to the following appears:

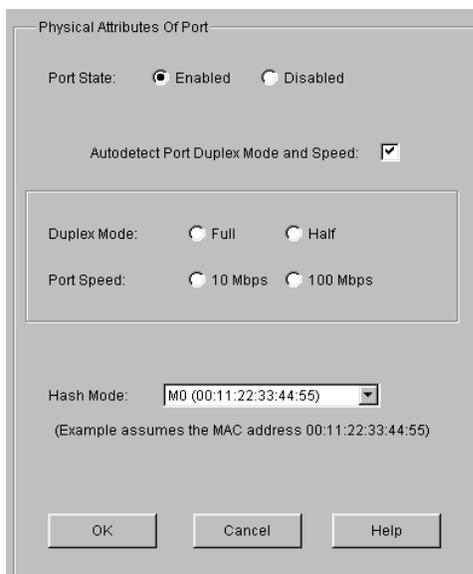


Figure 19. Physical Attributes of Port dialog box

6. Specify whether you want to enable or disable the port by selecting the appropriate option button.

Disabled ports do not send or receive any traffic. You might want to disable unused ports to prevent network users from inadvertently or unscrupulously connecting to unoccupied but enabled ports on the GSR.

7. Do one of the following:
 - If you want the port to detect and then use the operating mode and speed of the network segment to which the port is connected, select the *Autodetect Port Duplex Mode and Speed* check box. Otherwise, clear the check box to disable autodetection on the port.

All 10/100 Mbps Ethernet and Gigabit Ethernet ports use autodetection by default.

The 100 FX ports do not use auto-detection. They always use full-duplex mode and operate at 100 Mbps.

- If you want to set the port's operating mode or speed to a specific setting, select the appropriate buttons.

Note: If you select an operating mode or speed, the setting disables autodetection for that parameter on the port. For example, if you set the speed of a segment to 10-Mbps, that segment no longer uses autodetection for the port speed and will always attempt to operate at 10-Mbps.

8. Set the Layer 2 hash mode for the port. The hash mode controls the distribution of flow entries in Layer-2 and Layer-3 lookup tables. Assuming a MAC address of the value 0011:2233:4455, the following list describes the various hash modes:
 - M0 – 0011:2233:4455
 - M1 – 0011:2233:5544
 - M2 – 0011:3322:4455 (default hash mode)
 - M3 – 1100:2233:4455

Note: DIGITAL has configured the GSR's hash mode to its optimal setting. It is recommended you not change the hash mode unless advised to do so by DIGITAL Technical Support.
9. Click **OK**.

Configuring the GSR for a SYSLOG Server

You can configure the GSR to send system messages to a SYSLOG server. To do so, identify the server the GSR is to send messages to and also specify the type of messages that the GSR is to send. On the SYSLOG server, you can decide whether to discard the messages, write them to a file, or send them out to the console.

To configure the GSR for a SYSLOG server:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's System Configuration object.
3. Select the System Log object.

A **System Log** dialog box similar to the following appears:

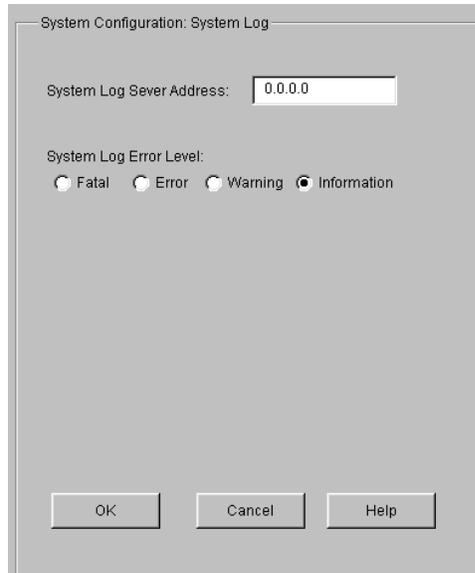


Figure 20. System Log dialog box

4. Enter the host name or IP address of the SYSLOG server.
5. Select the level of messages you want the GSR to log. You may select one of the levels described in the following table:

Table 5. SYSLOG error message levels

Level	Description
Fatal	Logs only fatal messages.
Error	Logs fatal messages and error messages.
Warning	Logs fatal messages, error messages, and warning messages. This is the default.
Information	Logs all messages, including informational messages.

6. Click **OK**.

Configuring for DNS

As an alternative to a host table on every system, some networks use a centralized Domain Naming System (DNS) server to maintain name-to-IP-address mappings. You may configure the GSR to reach up to three DNS servers. When doing so, you can also specify the domain name the GSR uses for each DNS query and the order in which the GSR searches for the specified DNS servers.

To configure the GSR for DNS servers:

1. Make sure there is a network connection to the DNS servers you want the GSR to use. You can do so by sending a ping packet to those servers.
2. Start Configuration Expert if you have not already done so.
3. Open the configuration file you want to modify and then double-click that file's System Configuration object.
4. Select the DNS Configuration object.

A **System Configuration DNS** dialog Box similar to the following appears:

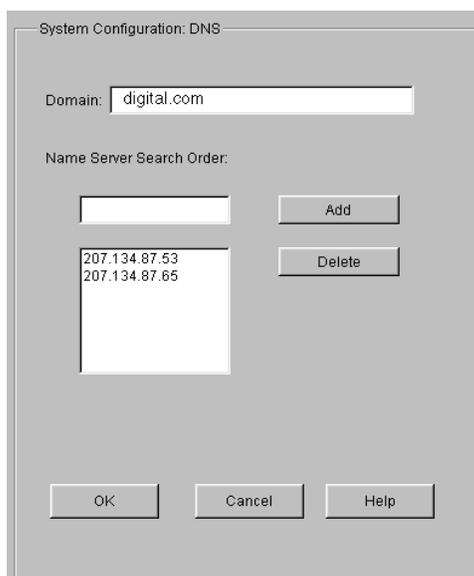


Figure 21. System Configuration DNS dialog box

5. Specify which DNS servers you want the GSR to use by taking the following steps to add those servers. When the GSR searches for DNS servers, it will do so using the order in which those servers were added.
 - a. In the *Domain* text box, enter the name of the domain on which the DNS server you want to add is an authority.

- b. In the *Name Server Search Order* text box, enter the IP address of the DNS server you want to add. Then click the **Add** button.

You can specify the address in dotted-decimal notation.

After you click the **Add** button, Configuration Expert adds the specified server to the list.

- c. If you want to configure the GSR for additional DNS servers, repeat [step a](#) and [step b](#) until you add up to three DNS servers.

If you add a server and then later want to remove it, you may do so by selecting its name in the list and then clicking the Delete button.

6. Click **OK**.

Configuring the GSR for SNMP

The Simple Network Management Protocol (SNMP) provides support for monitoring and controlling network devices, collecting statistics, and managing configurations, performance, and security. SNMP is mainly used by TCP/IP networks.

You configure the GSR for SNMP by performing the following tasks:

- Setting up targets for SNMP traps. A target is a management station to which the GSR sends SNMP traps, which are messages that describe an event (such as restarting the GSR or a link going down).
- Establishing community strings that SNMP management stations must supply to access the GSR.

Separate discussions on these SNMP tasks follow.

Setting Up a Target for SNMP Traps

Set up a target for SNMP traps to specify to which server you want the GSR to send SNMP traps. To set up an SNMP trap target on the GSR:

1. Make sure there is a network connection to the server to which you want the GSR to send SNMP traps. You can do so by sending a ping packet to the desired server.
2. Start Configuration Expert if you have not already done so.
3. Open the configuration file you want to modify and then double-click that file's System Configuration object.
4. Double-click the SNMP Configuration object.
5. Configuration Expert displays the SNMP Trap Target and SNMP Community String objects.

6. Double-click the SNMP Trap Target object.
7. Do one of the following:
 - If you are configuring a new trap target, select the Configure New Trap Target object from the list of trap targets.
 - If you are modifying an existing trap target, select it from the list that appeared after you expanded the SNMP Trap Target object.

An **SNMP Trap Target** dialog box similar to the following appears:

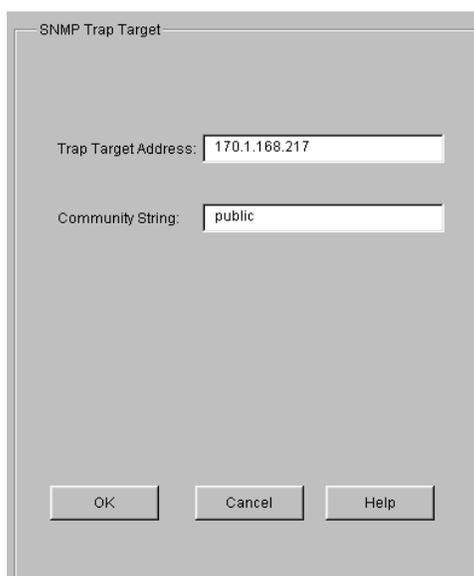


Figure 22. SNMP Trap Target dialog box

8. Enter the IP address of the management station from which you want to be able to access the traps.

Note: The target IP address should be for a management station locally attached to the GSR. Cold start traps might not reach their destination if the target requires dynamic route table entries to be forwarded correctly. The system will retry for four minutes.
9. In the *Community String* text box, enter the SNMP community for which you are setting the trap target.
10. Click **OK**.

Establishing Community Strings

SNMP management stations that want to access the GSR must supply a community string that you establish on the GSR. You can establish a GSR community string by specifying the string's name and selecting the access privileges for that string. To establish community strings on the GSR:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's System Configuration object.
3. Double-click the SNMP Configuration object.

Configuration Expert displays the SNMP Trap Target and SNMP Community String objects.

4. Double-click the SNMP Community String object.
5. Do one of the following:
 - If you are configuring a new community string, select the Configure New Community String object.
 - If you are modifying an existing community string, select it from the list of community strings that appeared after you expanded the SNMP Community String object.

An **SNMP Community Strings** dialog box similar to the following appears:



Figure 23. SNMP Community Strings dialog box

6. In the *Community String* text box, enter a character string for the community string.
7. Set the level of access to the GSR by selecting one of the options described in the following table:

Table 6. Level-of-access options

Option	Description
Read-only	Allows SNMP GETs but not SNMP SETs on the SNMP management stations that access the GSR through the specified community string.
Read-write	Allows both SNMP GETs and SNMP SETs on the SNMP management stations that access the GSR through the specified community string.

8. Click **OK**.

Chapter 6

Configuring GSR Bridging

The GSR provides bridging functions. This chapter

- provides an overview of bridging on a GSR.
- discusses configuring the bridging mode of ports. A port's bridging mode determines the contents of that port's Layer-2 lookup table.
- discusses controlling the aging state of GSR bridging. The GSR's aging state determines how long the GSR stores the MAC address information.
- explains setting up the GSR for the Spanning Tree Protocol (STP), which the GSR uses to work around loops.
- explains using Configuration Expert to configure of SmartTRUNKing on the GSR.

A Look at Bridging on the GSR

The GSR uses transparent bridging to link together different segments of an Ethernet network.

In transparent bridging, the GSR operates as a learning bridge. As such, it monitors traffic on subnetworks to learn about destination locations, down bridges, traffic congestion, and other network information. This frees packets from having to carry routing information and permits the GSR to forward packets based on current network conditions.

As a learning bridge, the GSR stores the source MAC address of each packet it receives. This permits the GSR to forward future packets destined for a MAC address to be forwarded to the port on which the source address was found. It does not forward those packets to other ports, and thus reduces traffic on the network. If the GSR does not recognize a packet's MAC address, however, it forwards that packet to every port.

Configuring the Bridging Mode of Ports

You can configure ports to use either of the following bridging modes. Each port has a Layer-2 lookup table where MAC address or flows are stored. A port's bridging mode determines the contents of each Layer-2 table entry. A port can use only one type of bridging at a time.

- Flow-based bridging

If a port is configured for flow-based bridging, each Layer-2 table entry contains entries consisting of a source MAC address, destination MAC address, and VLAN ID.

- Address-based bridging

If a port is configured for address-based bridging (the default), each Layer-2 table entry contains a unique destination MAC address and VLAN ID.

Suppose that a port on the GSR is connected to a hub that is connected to three workstations, A, B, and C. If each workstation is talking to one another and sending broadcast traffic, the Layer-2 table on the GSR's port would contain the following entries for the workstations (assume that the VLAN ID is 1 for all entries):

Table 7. Bridging modes

Port Configuration	Entries
Flow-based bridging	<ul style="list-style-type: none"> • MAC addresses A->B • MAC addresses A->C • MAC addresses B->A • MAC addresses B->C • MAC addresses C->A • MAC addresses C->B • MAC addresses A->broadcast • MAC addresses B->broadcast • MAC addresses C->broadcast
Address-based bridging	<ul style="list-style-type: none"> • MAC address A • MAC address B • MAC address C • MAC broadcast address

Configuring a Port for Flow-Based Bridging

To configure a port for flow-based bridging:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the Bridging Mode object.

A **Bridging Mode** dialog box similar to the following appears:

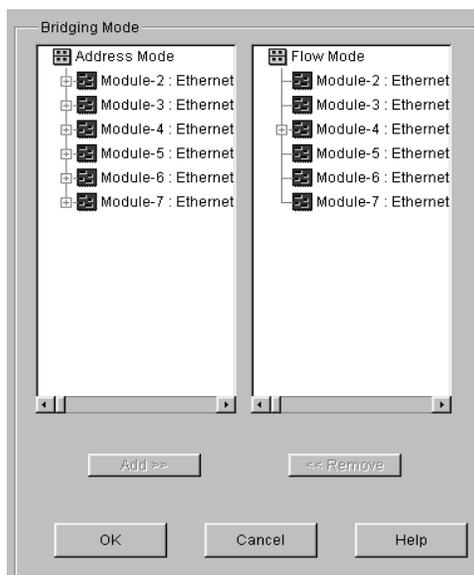
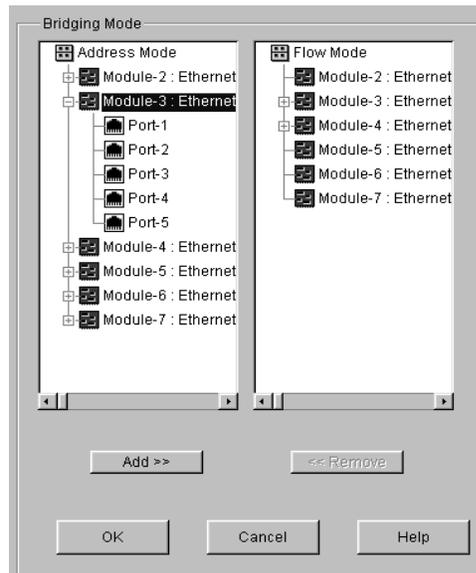


Figure 24. Bridging Mode dialog box (flow-based bridging)

4. In the **Address Mode** list, double-click the module on which the port you want to configure is located. Configuration Expert displays the module's ports that are currently using address-based bridging. From the list of ports that appears, select the port that you want to configure.



Note: Clicking a module in the **Address Mode** list rather than double-clicking it selects all of the module's ports that are currently using address-based bridging. Use this method if you are configuring all the ports on a module to use flow-based bridging.

5. Click the **Add** button.

Configuration Expert moves the selected port from the **Address Mode** list to the corresponding module in the **Flow Mode** list.

6. Click **OK**.

Configuration Expert adds the port to those found in the Flow Mode Bridging object, which is located in the Bridging Mode object.

Configuring a Port for Address-Based Bridging

Ports are configured to use address-based bridging by default. If you previously configured a port for flow-based bridging and later decide you want that port to use address-based bridging, you can change its bridging mode. To do so, take the following steps:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the Bridging Mode object.

A **Bridging Mode** dialog box similar to the following appears:

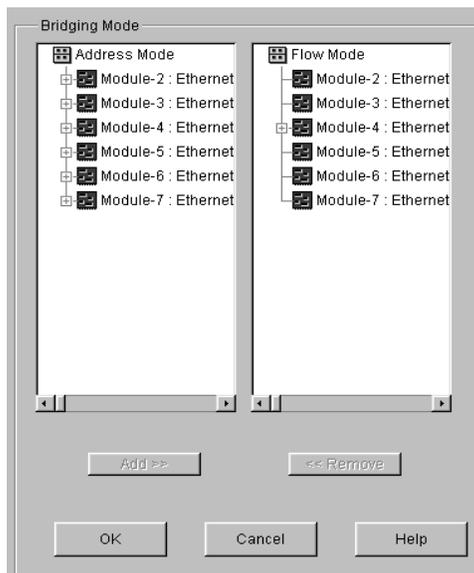
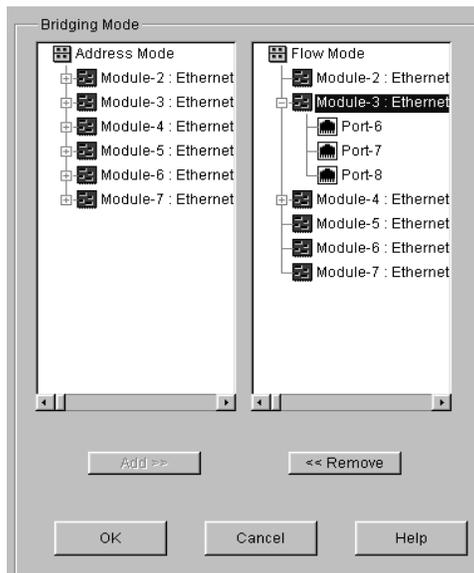


Figure 25. Bridging Mode dialog box (address-based bridging)

4. In the **Flow Mode** list, double-click the module on which the port you want to configure is located. Configuration Expert displays the module's ports that are currently using flow-based bridging. From the list of ports that appears, select the port that you want to configure.



Note: Clicking a module in the **Flow Mode** list rather than double-clicking it selects all of the module's ports that are currently using flow-based bridging. Use

this method if you are configuring all the ports on a module to use address-based bridging.

5. Click the **Remove** button.

Configuration Expert moves the selected port from the **Flow Mode** list to the corresponding module in the Address Mode list box.

6. Click **OK**.

Configuration Expert adds the port to those found in the Address Mode Bridging object, which is located in the Bridging Mode object.

Controlling the Aging State of GSR Bridging

The GSR ages learned MAC addresses in the Layer-2 lookup tables. Each port has its own Layer-2 lookup table. When a learned MAC address entry ages out, the GSR removes that entry from a port's Layer-2 lookup tables unless you disable aging on that port.

You can control the aging of learned MAC address entries in the GSR's Layer-2 lookup tables. To do so, use Configuration Expert to perform the following tasks:

- Setting up an aging timeout that ports use by default. The GSR uses the aging timeout to determine how long to keep learned MAC addresses. Aging is a regulation mechanism the GSR uses to clean up MAC address entries that have not been used for awhile.
- Overriding the default timeout interval. You do so by setting timeout intervals on any ports that are to use an interval different from the default aging timeout.
- Disabling aging on a port if you do not want the GSR to remove MAC address entries from that port's Layer-2 lookup table.

Separate discussions on each task follow.

Setting Up a Default Aging Timeout

You can set an aging time for learned MAC address entries that all ports use by default. When the aging time expires for a MAC address, the GSR removes the MAC address from the port that uses the default timeout. To set up a default aging timeout for MAC address entries:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the Aging Configuration object.
4. In the list of aging objects that appears, select the Default Aging Timeout object.

A **Default Aging Timeout** dialog box similar to the following appears:

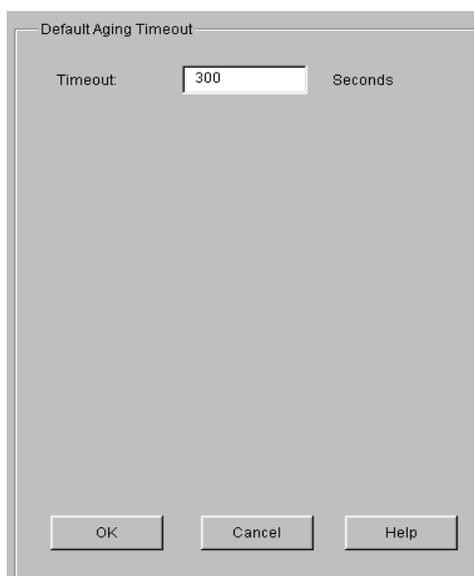


Figure 26. Default Aging Timeout dialog box

5. Enter the number of seconds that the GSR is to allow a learned MAC address to remain in the Layer-2 lookup table.

You can specify from 15 to 1,000,000 seconds. The default is 300 seconds.

6. Click **OK**.

Overriding the Default Timeout Interval on a Port

As discussed in [“Setting Up a Default Aging Timeout” on page 52](#), you can set an aging timeout for learned MAC address entries that ports use by default. You may override this default timeout if you want a port to use a different timeout interval. To override the timeout interval on a port:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file’s Bridging Configuration object.
3. Double-click the Aging Configuration object.
4. In the list of aging objects that appears, double-click the Aging Timeout Interval object. Then select the Configure Aging Timeout on a New Port object that appears.

A **Set Aging Timeout** dialog box similar to the following appears:

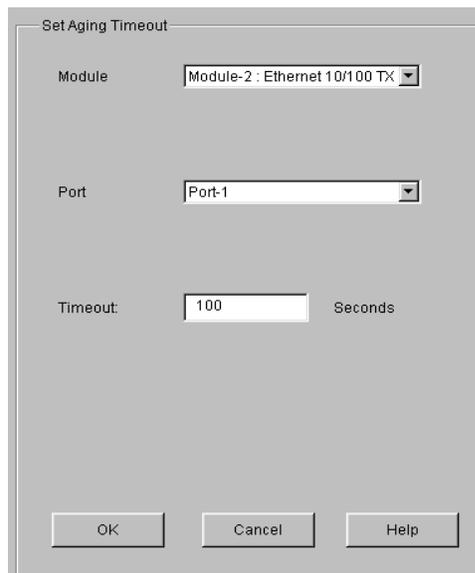


Figure 27. Set Aging Timeout dialog box

5. From the **Module** drop-down list, select the module containing the port you want to configure.
6. From the **Port** drop-down list, select the port you want to configure.
7. In the **Timeout** box, enter the number of seconds that the GSR is to allow a learned MAC address to remain in the Layer-2 lookup table for the specified port.
You can specify from 15 to 1,000,000 seconds. The default is 300 seconds.
8. Click **OK**.

Disabling Aging on a Port

Disable aging on a port if you do not want the GSR to age MAC address entries in the port's Layer-2 lookup tables.



Caution: Disabling aging on a port may eventually cause the port's Layer-2 lookup tables to become full because the GSR will not remove MAC addresses from those tables.

To disable aging on a port:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the Aging Configuration object.

- In the list of aging objects that appears, select the Aging State object.

A **Bridge Aging State** dialog box similar to the following appears:

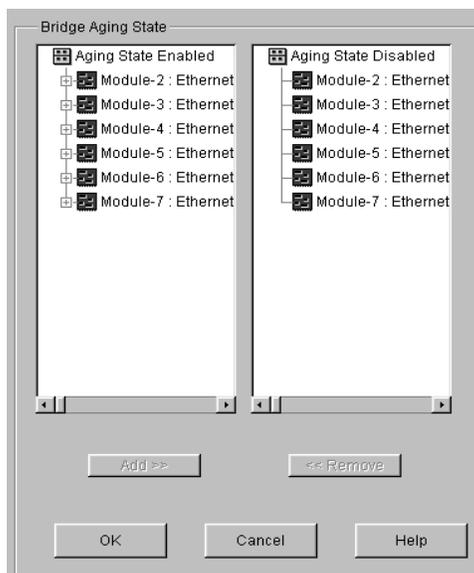
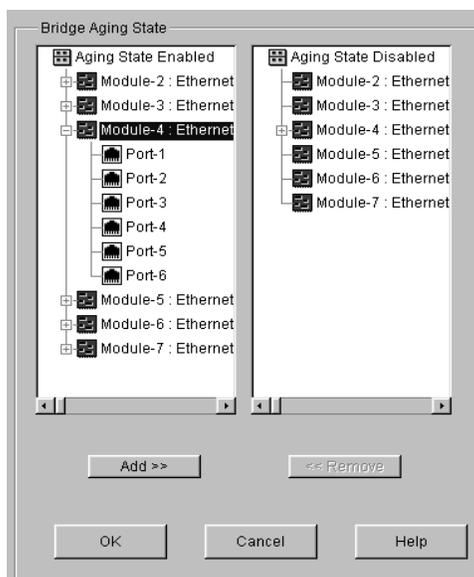


Figure 28. Bridge Aging State dialog box

- In the **Aging State Enabled** list, double-click the module on which the port you want to disable is located. Configuration Expert displays the module's ports on which aging is enabled. From the list of ports that appears, select the port on which you want to disable aging.



Note: Clicking a module in the **Aging State Enabled** list rather than double-clicking it selects all of the module's ports on which aging is currently enabled. Use this method if you want to disable aging on all of those ports.

6. Click the **Add** button.

Configuration Expert moves the selected port from the **Aging State Enabled** list to the corresponding module in the **Aging State Disabled** list.

7. Click **OK**.

Configuration Expert adds the port to those found in the Aging Disabled Ports object, which is located in the Aging State object.

Enabling Aging on a Port

The GSR removes aged MAC address entries from a port's Layer-2 lookup table if aging is enabled on that port. Aging is enabled on all ports by default. Therefore, you do not need to enable aging on a port unless you previously disabled it and then later decide you want to enable it on that port again.

To enable aging on a port:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the Aging Configuration object.
4. In the list of aging objects that appears, select the Aging State object.

A **Bridge Aging State** dialog box similar to the following appears:

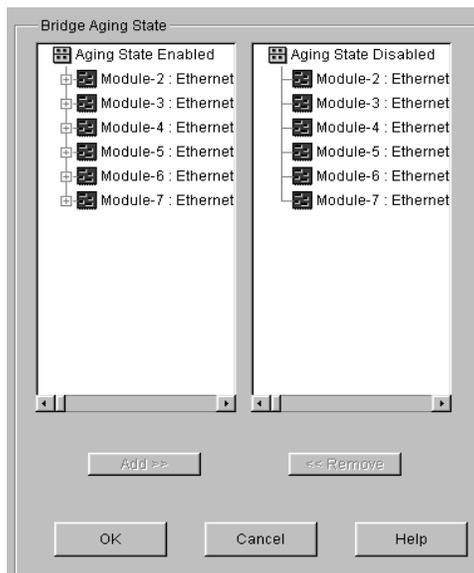
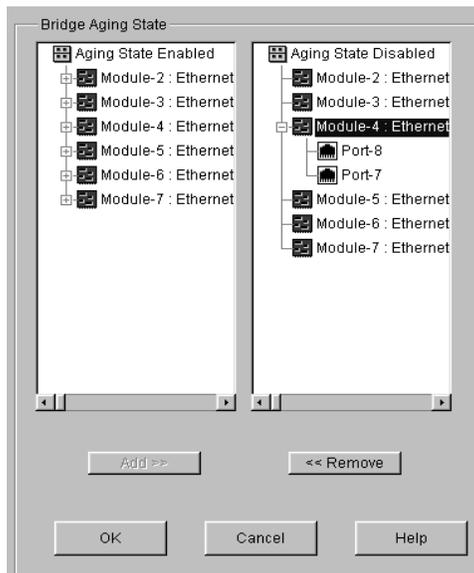


Figure 29. Bridge Aging State dialog box

5. In the **Aging State Disabled** list, double-click the module on which the port you want to enable is located. Configuration Expert displays the module's ports on which aging is disabled. From the list of ports that appears, select the port on which you want to enable aging.



Note: Clicking a module in the **Aging State Disabled** list rather than double-clicking it selects all of the module's ports on which aging is currently disabled. Use this method if you want to enable aging on all of those ports.

6. Click the **Remove** button.

Configuration Expert moves the selected port from the **Aging State Disabled** list to the corresponding module in the **Aging State Enabled** list.

7. Click **OK**.

Configuration Expert adds the port to those found in the Aging Enabled Ports object, which is located in the Aging State object.

Setting Up STP on the GSR

The GSR uses the Spanning Tree Protocol (STP) to dynamically discover a loop-free topology.

STP includes an algorithm that permits bridges to send Bridge Protocol Data Units (BPDUs). Unless a bridge is a root bridge, it uses the information in a BPDU to perform the following operations. (A root bridge is a bridge that uses STP to prevent loops by periodically exchanging topology information with other bridges.)

1. Select a root bridge.
2. Calculate the shortest path from itself to the root bridge.
3. Select a designated bridge.

A designated bridge is the bridge responsible for forwarding frames to a LAN segment.

4. Select a root port.

A root port is a port that provides the best path from the bridge to the root bridge.

5. Select which ports are included in a spanning tree.

On the GSR, you can enable STP per port. This approach offers flexibility. STP BPDUs that arrive on ports on which STP is disabled are forwarded only to ports which are part of the receiving port's VLAN.

You can set up STP on the GSR by performing the following tasks:

Note: STP is disabled on the GSR by default. If you want the GSR to use STP, you will need to define the different settings and enable STP.

- Defining STP settings for a GSR bridge.
- Defining STP settings for individual ports.
- Enabling and disabling STP on individual ports.

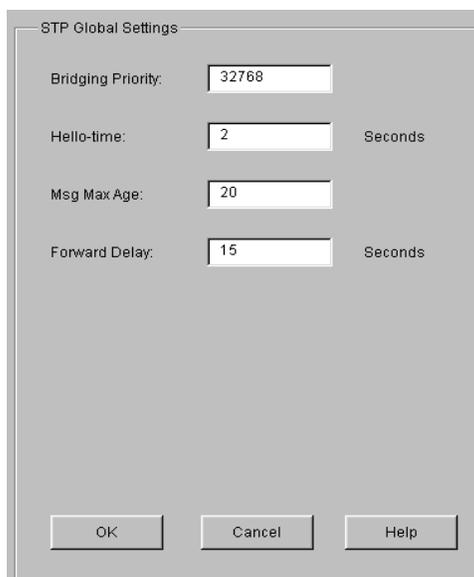
Separate discussions on these STP tasks follow.

Defining STP Settings for GSR Bridging

You can define global STP settings that the GSR uses for bridging. To define global STP settings:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the Spanning Tree Protocol object.
4. In the list of STP objects that appears, select the Global STP Settings object.

An **STP Global Settings** dialog box similar to the following appears:



The image shows a dialog box titled "STP Global Settings". It contains four input fields with the following values and units:

Field	Value	Unit
Bridging Priority	32768	
Hello-time	2	Seconds
Msg Max Age	20	
Forward Delay	15	Seconds

At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

Figure 30. STP Global Settings dialog box

5. In the **Bridging Priority** box, enter the STP bridging priority for the GSR.
The bridging priority determines which bridge functions as the root bridge. You can specify a number from 0 to 65535. The bridge assigned the lowest value is the root bridge. The default is 32768.
6. In the **Hello-time** box, enter the number of seconds that you want to elapse between the BPDUs STP sends.
Specify a number from 1 to 10. The default is 2.
7. In the **Msg Max Age** box, enter the number of seconds you want to specify for the GSR's maximum age.

The maximum age is the length of time the GSR keeps the STP-protocol information it receives. You can specify a number from 6 to 40. The default is 20.

8. In the **Forward Delay** box, enter the number of seconds you want to elapse between the transitions of the different STP states.

Specify a number from 4 to 30. The default is 15.

9. Click **OK**.

The GSR bridge will use the specified STP settings. However, individual ports also have STP settings that you can configure as discussed in [“Defining STP Attributes on an Individual Port” on page 60](#).

Defining STP Attributes on an Individual Port

In addition to the global STP settings discussed in [“Defining STP Settings for GSR Bridging” on page 59](#), the GSR uses the following STP attributes that are set on individual ports.

- Port priority

The STP algorithm uses the port priority to determine which port to use if a bridge has two ports connected in a loop.

- Port cost

This attribute specifies how much a port contributes to the total cost of the path to the root bridge when the port is the root.

A port uses the default settings for these attributes. You may, however, change these default settings on individual ports. To change a port’s priority or cost:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file’s Bridging Configuration object.
3. Double-click the Spanning Tree Protocol object.
4. In the list of STP objects that appears, double-click the Port Specific STP Settings object.

Configuration Expert displays the Configure New Port object and objects for any previously configured ports.

5. Do one of the following:
 - If you are configuring a port’s STP attributes for the first time, select the Configure STP Settings on a New Port object.
 - If you are modifying existing STP settings of a port, select that port’s object from the list that appeared after you expanded the Port Specific STP Settings object.

A **Set STP Port Specific Settings** dialog box similar to the following appears:

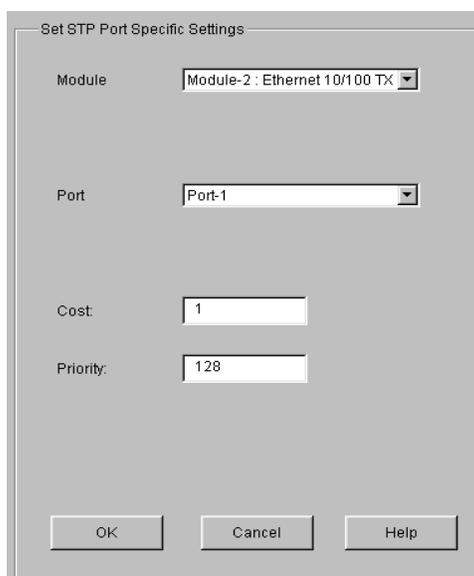


Figure 31. Set STP Port Specific Settings dialog box

6. From the **Module** drop-down list, select the module containing the port you want to configure.
7. From the **Port** drop-down list, select the port you want to configure.
8. In the **Cost** box, enter the STP cost you are assigning to the ports.
Specify a number from 1 to 65535. The default depends on the port speed: 1 for Gigabit (100-Mbps) ports, 10 for 100-Mbps ports, and 100 for 10-Mbps ports.
9. In the **Priority** box, enter the priority you want to assign to the port.
Specify a number from 0 to 255. The default is 128.
10. Click **OK**.

Enabling STP on a Port

Enable STP on a port to eliminate the possibility of there being a loop on that port. To enable STP on a port:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the Spanning Tree Protocol object.

4. In the list of STP objects that appears, select the STP Port State object.

A **Bridging STP** dialog box similar to the following appears:

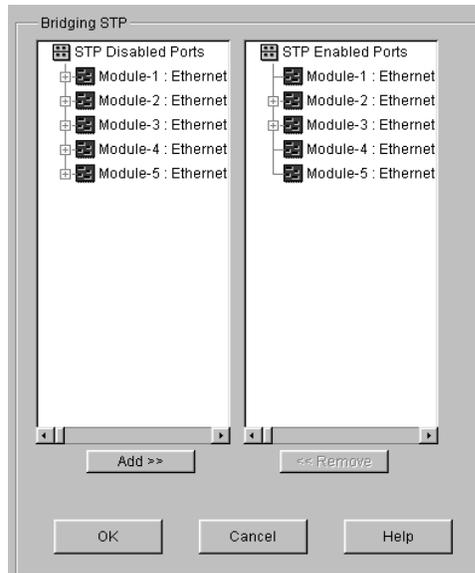
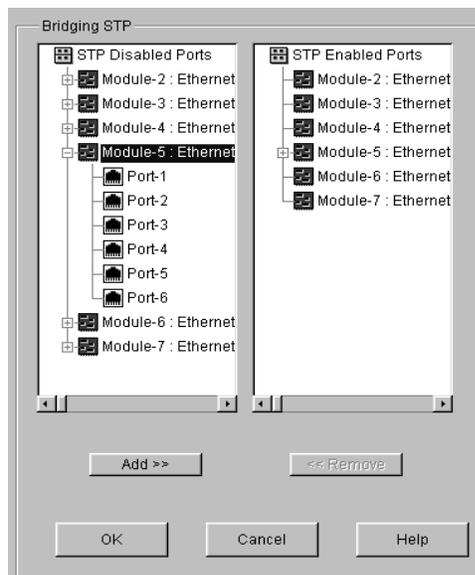


Figure 32. Bridging STP dialog box (enabling STP)

5. In the **STP Disabled Ports** list, double-click the module containing the port you want to configure. Configuration Expert displays the module's ports on which STP is disabled. From the list of ports that appears, select the port on which you want to enable STP.



Note: Clicking a module in the **STP Disabled Ports** list rather than double-clicking it selects all of the module's ports on which STP is currently disabled. Use this method if you want to enable STP on all of those ports.

- Click the **Add** button.

Configuration Expert moves the selected port from the **STP Disabled Ports** list to the corresponding module in the **STP Enabled Ports** list.

- Click **OK**.

Configuration Expert adds the port to those found in the STP Enabled Ports object, which is located in the STP Port State object.

Disabling STP on a Port

Disable STP on a port if the port does not need to participate in STP. To disable STP on a port:

- Start Configuration Expert if you have not already done so.
- Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
- Double-click the Spanning Tree Protocol object.
- In the list of STP objects that appears, select the STP Port State object.

A **Bridging STP** dialog box similar to the following appears:

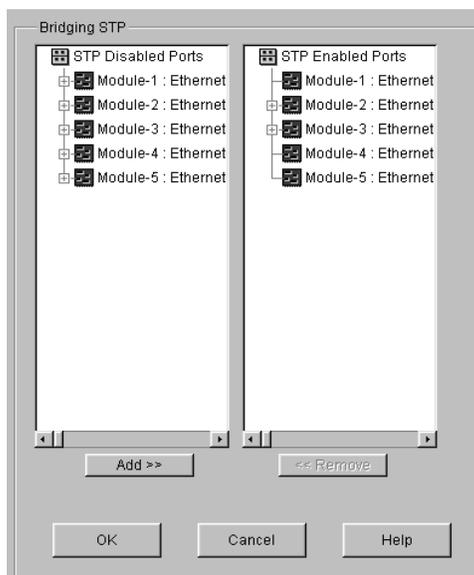
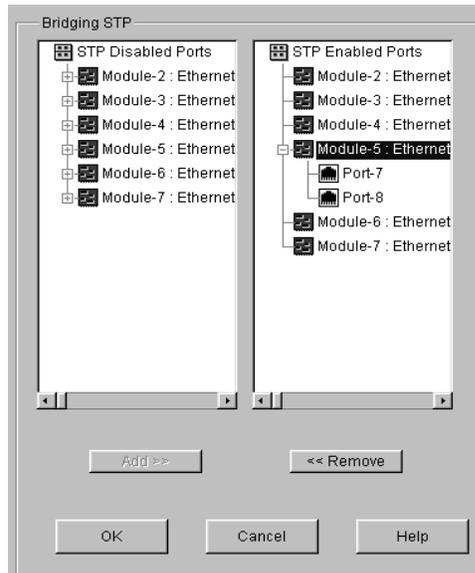


Figure 33. Bridging STP dialog box (disabling STP)

5. In the **STP Enabled Ports** list, double-click the module containing the port you want to configure. Configuration Expert displays the module's ports on which STP is enabled. From the list of ports that appears, select the port on which you want to disable STP.



Note: Clicking a module in the **STP Enabled Ports** list rather than double-clicking it selects all of the module's ports on which STP is currently enabled. Use this method if you want to disable STP on all of those ports.

6. Click the **Remove** button.

Configuration Expert moves the selected port from the **STP Enabled Ports** list to the corresponding module in the **STP Disabled Ports** list.

7. Click **OK**.

Configuration Expert adds the port to those found in the STP Disabled Ports object, which is located in the STP Port State object.

Configuring SmartTRUNK Behavior on the GSR

The GSR uses SmartTRUNKing to help maximize throughput with the help of load balancing and load sharing. A SmartTRUNK is a group of two or more ports that have been logically combined into a single port. Network traffic is divided across the ports in parallel to provide additional throughput, which provides increased throughput and link redundancy.

Defining SmartTRUNK Settings for GSR Bridging

To set up the GSR to perform SmartTRUNKing:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the SmartTRUNKing Configuration object.
4. Click the Configure New SmartTRUNKing Port object.

Configuration Expert opens the ARP wizard:

5. Click **Next**.

The SmartTRUNKing Name Entry panel appears:

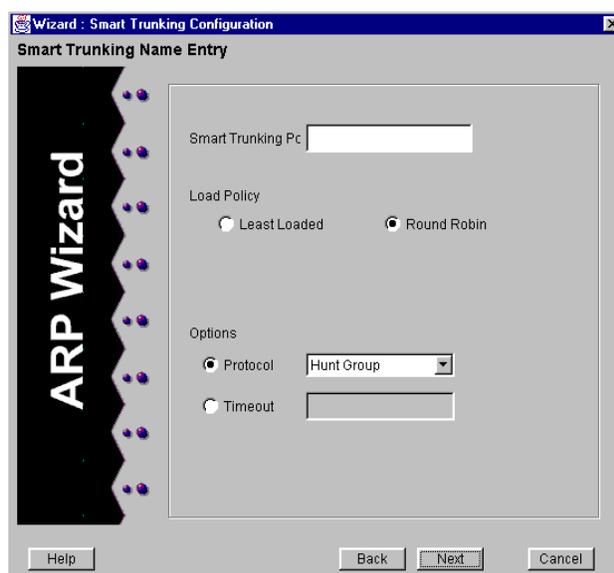


Figure 34. SmartTRUNKing Name Entry panel

6. Define the port(s) to be included in the SmartTRUNK by taking the following steps:
 - a. In the **SmartTRUNKing Port** box, enter a name for the port you wish to assign to the SmartTRUNK.

The SmartTRUNKing port name is a string up to 32 characters long. You cannot begin a SmartTRUNKing port name with an underscore (`_`) or the prefix "SYS_."

- b. Specify the *Load Policy* by selecting either the *Least Loaded* or *Round Robin* option.
 - c. Under *Options*, select either *Protocol* or *Timeout*. If you select the *Protocol* option,

you must also specify whether you wish to use the Hunt Group or no protocol by selecting the appropriate option from the drop-down list.

7. Click **Next**.

The Bound Port list panel appears:

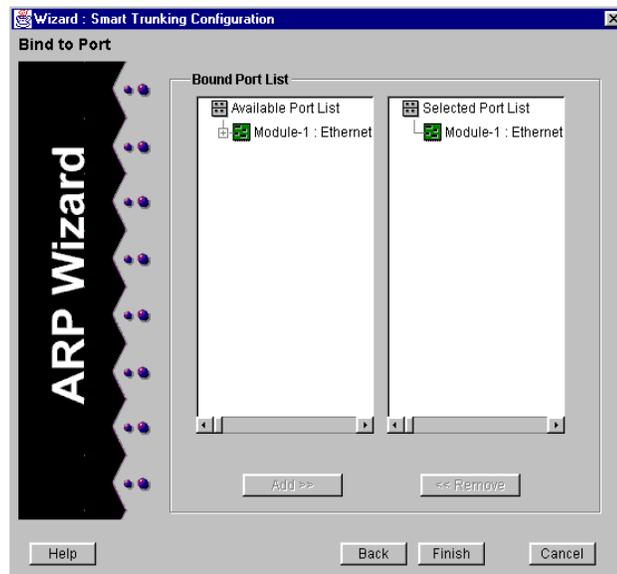


Figure 35. Bound Port list panel

8. Add a port to the SmartTRUNK by doing the following:
 - a. In the **Available Port** list, double-click the module on which the port you want to add is located. From the list of available ports that appears, select the port that you want to add.

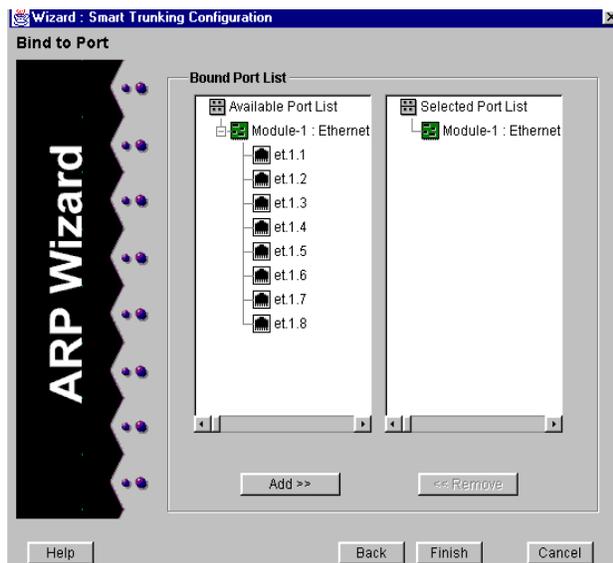


Figure 36. Expanded Bound Port list panel

- b. Click the **Add** button.

Configuration Expert moves the selected port from the **Available Port** list to the corresponding module in the **Selected Port** list.

If you accidentally add a port that you do not want to include in the SmartTRUNK, you may remove it by double-clicking that port's module in the **Selected Port** list. From the list of ports that appears, select the port you do not want in the SmartTRUNK, then click the **Remove** button.

9. Continue selecting ports and clicking the **Add** button until you have added all the ports you want the SmartTRUNK to include.

Clicking a module rather than double-clicking it in a list box selects all of the module's ports in that list box. This is a quick way to select all of a module's ports if you want to add or remove them all at the same time.

10. Click **Finish**.

Configuration Expert adds the new SmartTRUNK to the SmartTRUNKs found in the SmartTRUNKing object.

Chapter 7

Configuring VLANs on the GSR

You configure VLANs to limit the scope of traffic on the GSR. This chapter

- provides an overview of VLANs on the GSR.
- lists tips that make VLAN configuration easy.
- discusses defining ports for VLANs.
- discusses creating the different VLANs the GSR supports.
- describes copying ports to add them to a VLAN.
- discusses modifying VLANs.

A Look at VLANs on the GSR

VLANs contain Layer-2 broadcast and multicast traffic. No traffic is allowed to cross VLAN boundaries unless it passes through routers. Once connected by routers, VLANs are equivalent to subnets.

VLANs are created by grouping a set of bridged ports together as part of one bridged network. Broadcasts from one of the ports in a VLAN are received by other ports in the group but not by any ports outside of the group. Similarly, unicast traffic is bridged only between ports in a group but not to ports outside of the group. Thus, traffic is not allowed to cross the group boundary. A bridge may have multiple VLANs defined thus appearing as multiple virtual bridges on the GSR.

The GSR supports the following types of VLANs. The VLAN type determines the type of traffic the GSR will forward on the VLAN.

- Protocol-based VLAN, which divides the physical network into logical VLANs based on one or more of the following protocols:
 - IP VLAN, which is a VLAN used for IP traffic.
 - IPX VLAN, which is a VLAN used for IPX traffic.
 - Bridged-protocol VLAN, which is a VLAN used for bridged protocols (such as AppleTalk).
- Port-based VLAN, which is a VLAN that is independent of the traffic type. Port-based VLANs treat IP, IPX, and bridged protocols alike.

The ports in a VLAN can be configured as one of the following:

- Access ports

Access ports can belong to only one VLAN per protocol (IP, IPX, or a bridged protocol). This is the default for all ports.

On access ports, traffic is sent out without 802.1Q frame format.

- Trunk ports

Trunk ports can belong to any number of VLANs. Use trunk ports when you want to connect GSR routers together and send traffic for multiple VLANs on a single network segment connecting the routers.

On trunk ports, traffic is always sent out with 802.1Q frame format.

VLAN Configuration Tips

The following list includes tips that you should keep in mind while configuring VLANs on a GSR as discussed later in this chapter:

- You can quickly add ports to a VLAN, by copying ports as discussed in [“Dragging Ports to Add Them to a VLAN”](#) on page 85.
- When defining trunk ports, you can define all the ports on a module as trunk ports at one time rather than defining them individually. To do so, click rather than double-click the module in the Access Port list box of the dialog box you use to define ports. Then click the Add button.
- When defining access ports in a dialog box, you can define all the ports on a module as access ports at one time rather than defining them individually. To do so, click rather than double-click the module in the Trunk Port list box of the dialog box you use to define ports. Then click the Remove button.
- When adding ports to a VLAN, you can add all of a module’s available ports at one time rather than adding them individually. To do so, click rather than double-click the module in the Available Port List box of the dialog box you use to add ports. Then click the Add button.

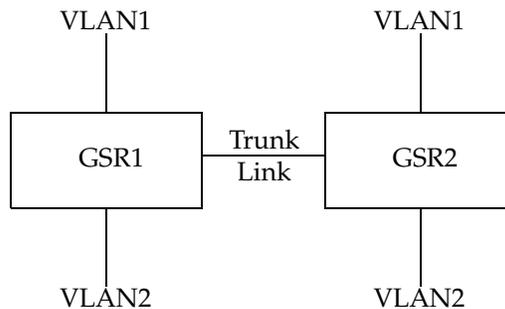
- When removing ports from a VLAN, you can remove all of a module's ports at one time rather than removing them individually. To do so, click rather than double-click the module in the Selected Port List box of the dialog box you use to add ports. Then click the Remove button.

Defining Access Ports and Trunk Ports

You can define ports as either access or trunk ports.

All ports are configured as access ports by default. Therefore, you need to define a port as an access port only if you previously defined it as a trunk port and now want to use it as an access port, instead. Access ports can be added to only one VLAN.

Define a port as a trunk port if you want to include that port in multiple VLANs. Trunk ports are useful for connecting GSRs together and for sending traffic of multiple VLANs on a single network segment connecting the routers. Suppose you have two VLANs/subnetworks of IP users on separate GSRs and those VLANs need to belong to the same layer broadcast domain. You could trunk two GSRs together as shown in the following figure:



To define access ports and trunk ports:

1. Start Configuration Expert if you have not already done so.

Note: If you start Configuration Expert from the Schematic view choosing the **Configure VLAN** command, Configuration Expert automatically expands the Active Configuration file's tree to the VLAN Configuration object. If you are adding a VLAN to the Active Configuration file, go to [step 3](#).

2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Click the Port Bridging Mode object.

A **Bridging VLAN Mode** dialog box similar to the following appears:

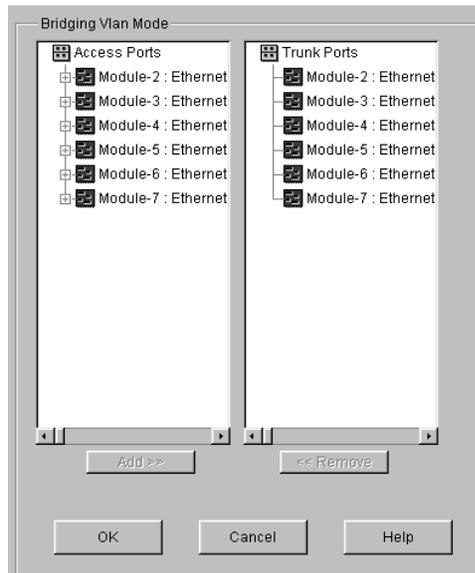
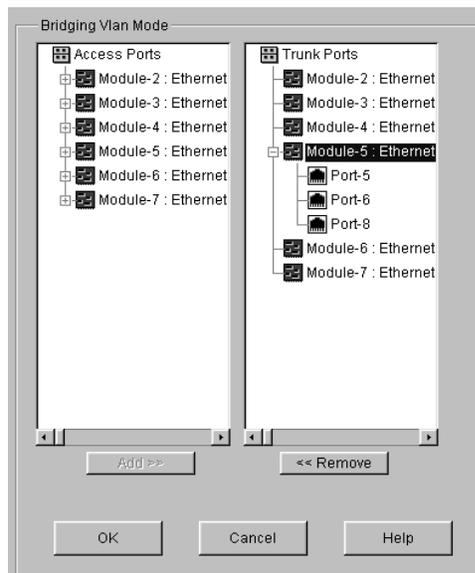


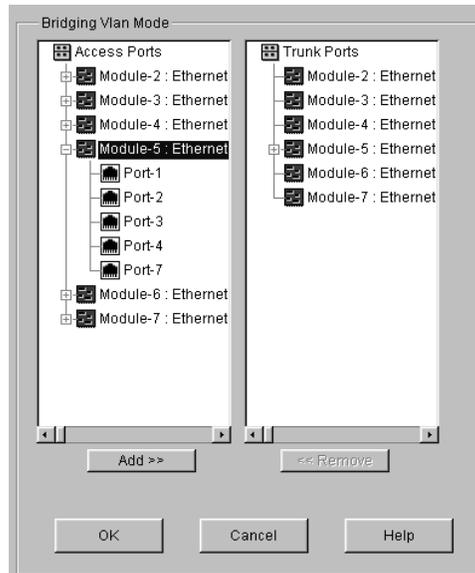
Figure 37. Bridging VLAN Mode dialog box

4. Do one of the following:
 - Define an access port by double-clicking that port's module in the **Trunk Ports** list. From the list of trunk ports that appears, select the port that you want to define as an access port. Then click the **Remove** button.



After you click the **Remove** button, Configuration Expert moves the selected port from the **Trunk Ports** list to the corresponding module in the **Access Ports** list.

- Define a trunk port by double-clicking that port's module in the **Access Ports** list. From the list of access ports that appears, select the port that you want to define as a trunk port. Then click the **Add** button.



After you click the **Add** button, Configuration Expert moves the selected port from the **Access Ports** list to the corresponding module in the **Trunk Ports** list.

5. Repeat [step 4](#) until you define all the access and trunk ports you will include in your VLANs.

Clicking a module rather than double-clicking it in a list box selects all of the module's ports in that list box. This is a quick way to select all of module's ports when you are defining those ports as all the same type.

6. Click **OK**.

Configuration Expert adds the access ports and trunk ports to those listed in the Access Ports and Trunk Ports objects, which are located in the Port Bridging Mode object.

Creating a Protocol-Based VLAN

To create a VLAN that the GSR will use for IP, IPX, or bridged-protocol traffic:

1. Start Configuration Expert if you have not already done so.

Note: If you start Configuration Expert from the Schematic view choosing the **Configure VLAN** command, Configuration Expert automatically expands the Active Configuration file's tree to the VLAN Configuration object. If you are adding a VLAN to the Active Configuration file, go to [step 3](#).

2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the VLAN Configuration object.
4. Click the Configure New VLAN object.

Configuration Expert opens the VLAN wizard.

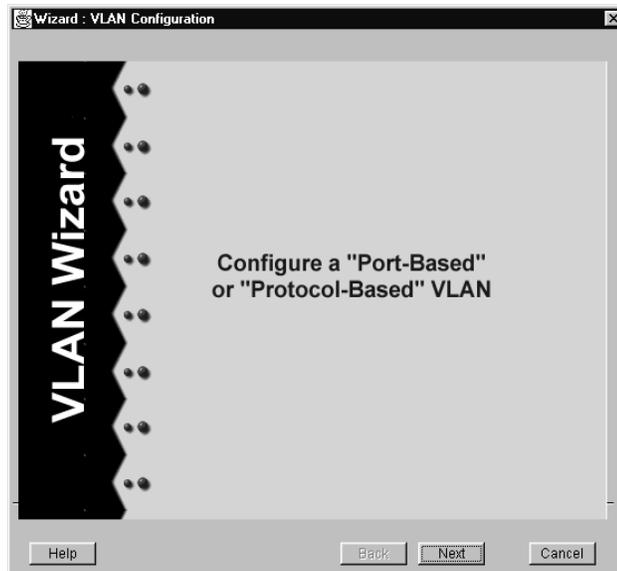


Figure 38. VLAN wizard (protocol-based)

5. Click Next.

Configuration Expert prompts you to specify which type of VLAN you want to configure.

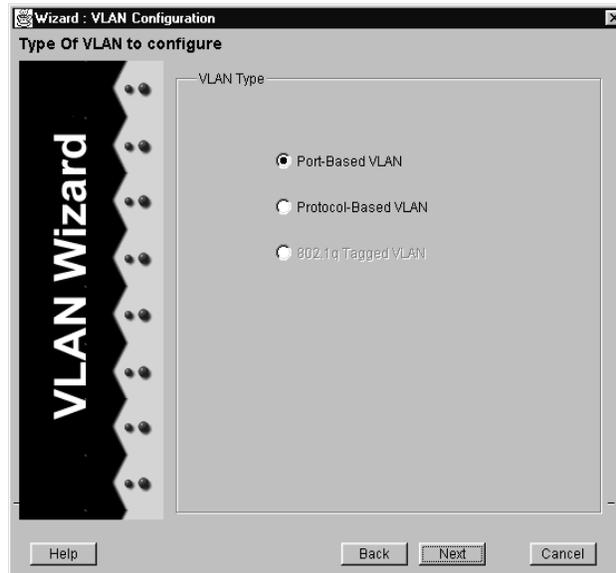


Figure 39. VLAN Type panel (protocol-based)

6. Select the *Protocol-Based VLAN* option and then click **Next**.
7. In the wizard panel that appears, define the VLAN by taking the following steps:
 - a. In the **VLAN Name** box, enter a name for the VLAN.

The VLAN name is a string up to 32 characters long. You cannot begin a VLAN name with an underscore (`_`) or the prefix "SYS_."
 - b. In the **VLAN ID** box, enter an ID number from 2 to 4093 for the VLAN.

The ID you enter must be unique for the VLAN. If you are creating a VLAN that will be used on two GSRs and you want to connect those GSRs together on the same trunk port, be sure to use the same ID when creating the VLAN on each GSR. The GSRs will use the ID to recognize that the same VLAN is configured across two GSRs.
 - c. Specify which type of traffic you want to allow on the VLAN. To do so, select one

or more of the options described in the following table:

Table 8. VLAN traffic types

Option	Description
IP	Specifies that the VLAN is for IP traffic.
IPX	Specifies that the VLAN is for IPX traffic.
Other	Specifies that the VLAN is for bridged protocols.

The following figure is an example of the information you enter to define a protocol-based VLAN:

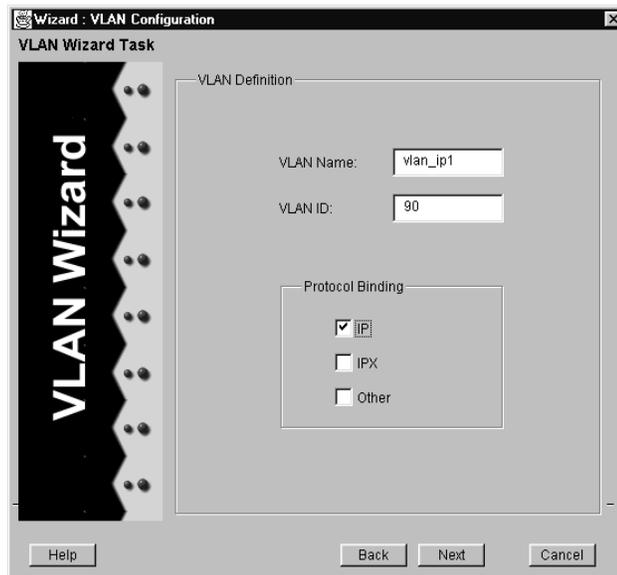


Figure 40. VLAN Definition panel (protocol-based)

8. Click **Next**.

A wizard panel similar to the following appears:

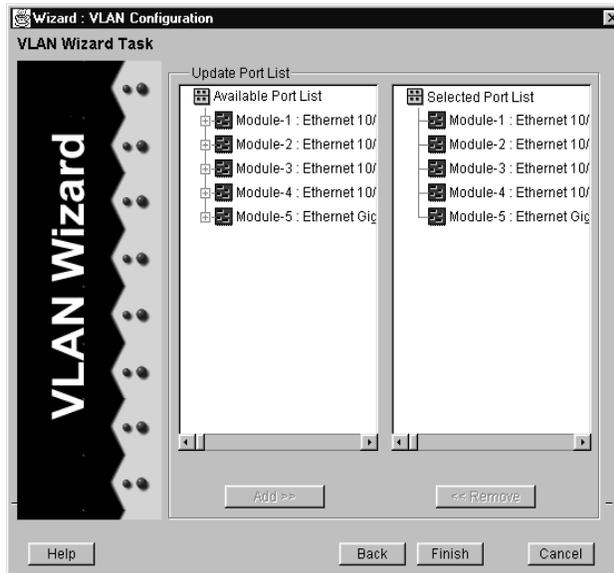


Figure 41. Update Port list panel (protocol-based)

9. Add a port to the VLAN by doing the following:
 - a. In the **Available Port** list, double-click the module on which the port you want to add is located. From the list of available ports that appears, select the port that you want to add.

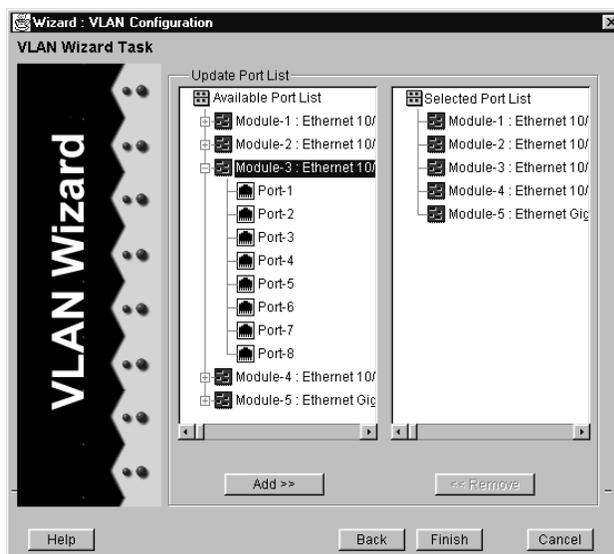


Figure 42. Expanded Update Port list panel (protocol-based)

- b. Click the **Add** button.

Configuration Expert moves the selected port from the **Available Port** list to the corresponding module in the **Selected Port** list.

If you accidentally add a port that you do not want to include in the VLAN, you may remove it by double-clicking that port's module in the **Selected Port** list. From the list of ports that appears, select the port you do not want in the VLAN. Then click the **Remove** button.

10. Continue selecting ports and clicking the **Add** button until you have added all the ports you want the VLAN to include.

Clicking a module rather than double-clicking it in a list box selects all of the module's ports in that list box. This is a quick way to select all of a module's ports if you want to add or remove them all at the same time.

11. Click **Finish**.

Configuration Expert adds the new VLAN to the VLANs found in the Protocol Based VLANs object.

Creating a Port-Based VLAN

To create a VLAN that the GSR will use for IP, IPX, and bridged-protocol traffic:

1. Start Configuration Expert if you have not already done so.

Note: If you start Configuration Expert from the Schematic view choosing the **Configure VLAN** command, Configuration Expert automatically expands the Active Configuration file's tree to the VLAN Configuration object. If you are adding a VLAN to the Active Configuration file, go to [step 3](#).

2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the VLAN Configuration object.
4. Click the Configure New VLAN object.

Configuration Expert opens the VLAN wizard.

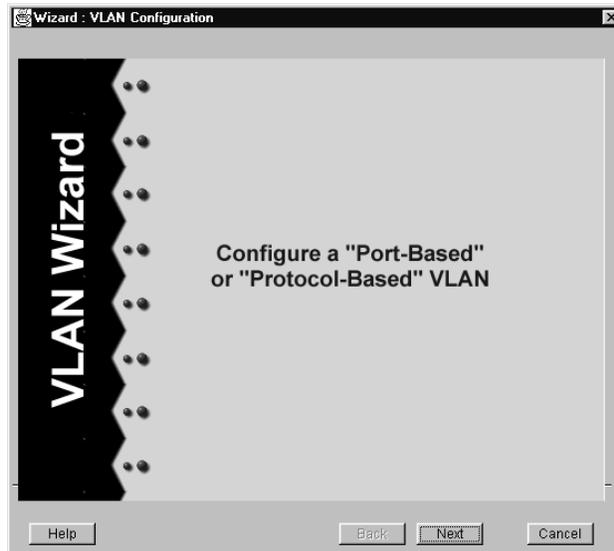


Figure 43. VLAN wizard (port-based)

5. Click **Next**.

Configuration Expert prompts you to specify which type of VLAN you want to configure.

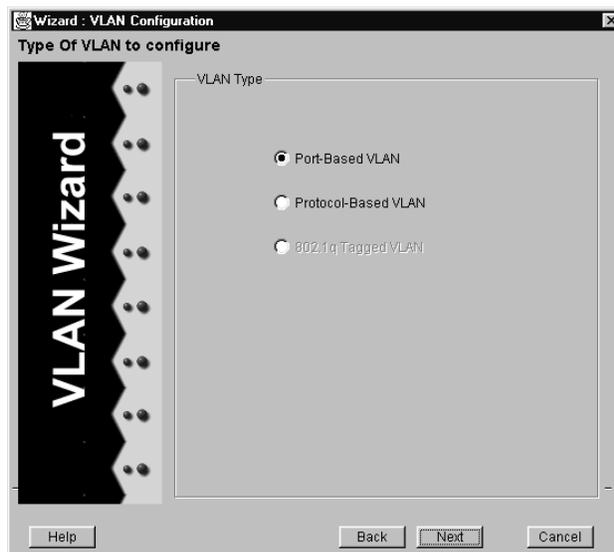


Figure 44. VLAN Type panel (port-based)

6. Select the *Port-Based VLAN* option, then click **Next**.

7. In the wizard panel that appears, define the VLAN by taking the following steps:

- a. In the **VLAN Name** box, enter a name for the VLAN.

The VLAN name is a string up to 32 characters long. You cannot begin a VLAN name with an underscore (`_`) or the prefix "SYS_."

- b. In the **VLAN ID** box, enter an ID number from 2 to 4093 for the VLAN.

The ID you enter must be unique for the VLAN. If you are creating a VLAN that will be used on two GSRs and you want to connect those GSRs together on the same trunk port, be sure to use the same ID when creating the VLAN on each GSR. The GSRs will use the ID to recognize that the same VLAN is configured across two GSRs.

The following figure is an example of the information you enter to define a port-based VLAN:

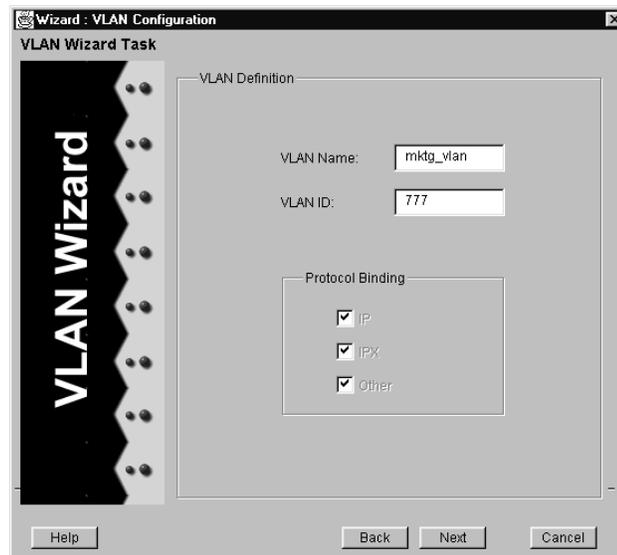


Figure 45. VLAN Definition panel (port-based)

Because port-based VLANs are used for all the different types of protocol traffic, Configuration Expert automatically selects all the Protocol Binding options described in the following table. You cannot change these selections.

Table 9. Protocol binding options

Option	Description
IP	Specifies that the VLAN is for IP traffic.
IPX	Specifies that the VLAN is for IPX traffic.
Other	Specifies that the VLAN is for bridged protocols.

8. Click **Next**.

A wizard panel similar to the following appears:

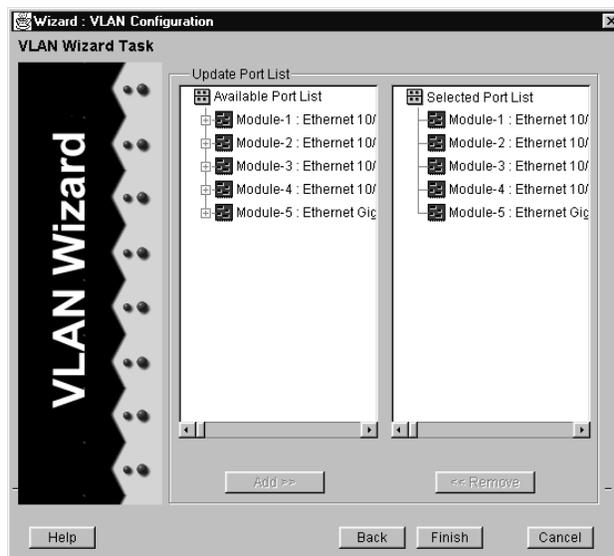


Figure 46. Update Port list panel (port-based)

9. Add a port to the VLAN by doing the following:
 - a. In the **Available Port** list, double-click the module on which the port you want to add is located. From the list of available ports that appears, select the port that you want to add.

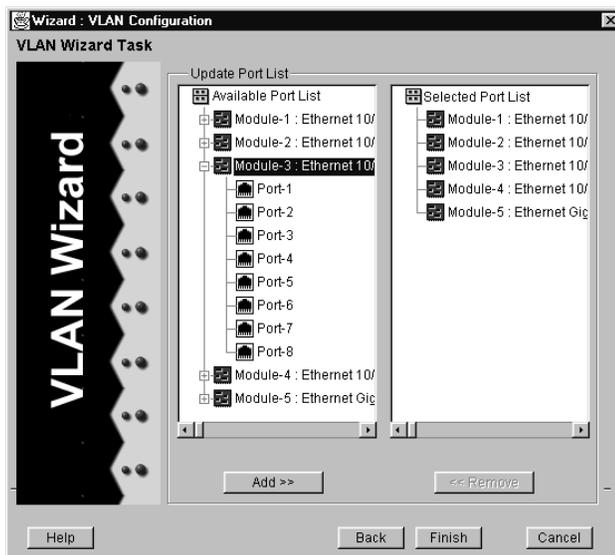


Figure 47. Expanded Update Port list panel (port-based)

- b. Click the **Add** button.

Configuration Expert moves the selected port from the **Available Port** list to the corresponding module in the **Selected Port** list.

If you accidentally add a port that you do not want to include in the VLAN, you may remove it by double-clicking that port's module in the **Selected Port** list. From the list of ports that appears, select the port you do not want in the VLAN. Then click the **Remove** button.

10. Continue selecting ports and clicking the **Add** button until you have added all the ports you want the VLAN to include.

Clicking a module rather than double-clicking it in a list box selects all of the module's ports in that list box. This is a quick way to select all of a module's ports if you want to add or remove them all at the same time.

11. Click **Finish**.

Configuration Expert adds the new VLAN to the VLANs found in the Port Based VLANs object.

Modifying VLANs

You can modify any VLAN's name, ID, and the ports included in a VLAN. On a protocol-based VLAN you can also modify the protocols bound to that VLAN. Separate

discussions on modifying the different types of VLANs and changing which ports are included in a VLAN follow.

Changing a Port-Based VLAN's Name or ID

To change the name of a port-based VLAN:

1. Start Configuration Expert if you have not already done so.

Note: If you start Configuration Expert from the Schematic view choosing the **Configure VLAN** command, Configuration Expert automatically expands the Active Configuration file's tree to the VLAN Configuration object. If you are adding a VLAN to the Active Configuration file, go to [step 3](#).

2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the VLAN Configuration object.
4. Double-click the Port-Based VLANs object. From the list of VLANs that appears, click the VLAN you want to modify.

A **VLAN Definition** dialog box similar to the following appears:

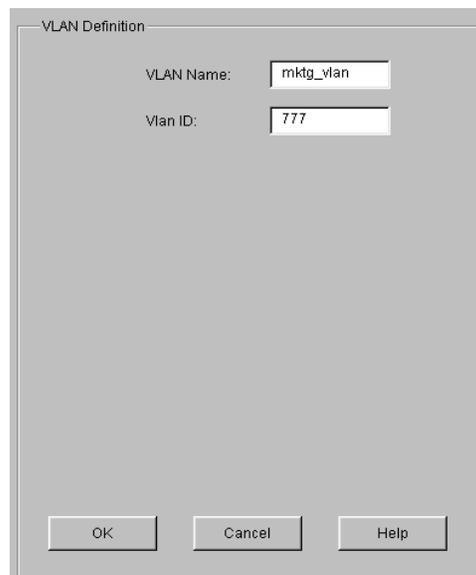


Figure 48. VLAN Definition dialog box (port-based)

5. Enter the new VLAN name or VLAN ID in the appropriate text boxes.
6. Click **OK**.

Changing a Protocol-Based VLAN's Name, ID, or Protocol Binding

To modify the name, ID, or protocol binding of a protocol-based VLAN:

1. Start Configuration Expert if you have not already done so.

Note: If you start Configuration Expert from the Schematic view choosing the **Configure VLAN** command, Configuration Expert automatically expands the Active Configuration file's tree to the VLAN Configuration object. If you are adding a VLAN to the Active Configuration file, go to [step 3](#).

2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the VLAN Configuration object.
4. Double-click the Protocol-Based VLANs object. From the list of VLANs that appears, click the VLAN you want to modify.

A **VLAN Definition** dialog box similar to the following appears:

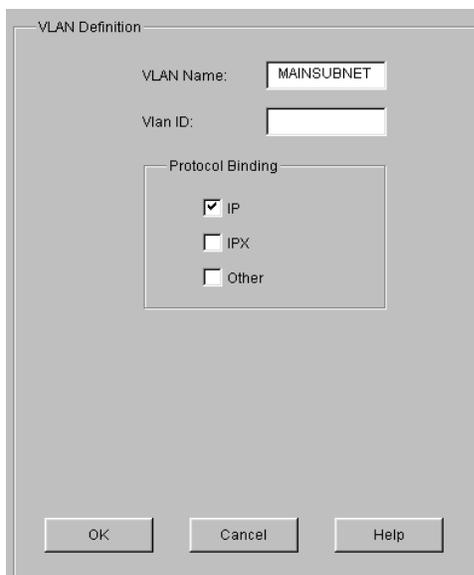


Figure 49. VLAN Definition dialog box (protocol-based)

5. If you want to change the VLAN's name or ID, edit the appropriate text boxes.
6. If you want to change which protocols the VLAN supports, select and clear the appropriate check boxes.
7. Click **OK**.

Replacing an Interface's VLAN

You can quickly replace an interface's VLAN with another VLAN. To do so, you either drag an IP VLAN to an IP interface or an IPX VLAN to an IPX interface.

To replace an interface's VLAN by dragging a VLAN to the interface:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object, then do one of the following:
 - If you are replacing the VLAN of an IP interface, double-click the IP Interface Configuration object and then double-click the IP interfaces bound to VLAN object.

Configuration Expert displays a list of IP interfaces bound to VLANs.

- If you are replacing the VLAN of an IPX interface, double-click the IPX Interface Configuration object and then double-click the IPX interfaces bound to VLAN object.

Configuration Expert displays a list of IPX interfaces bound to VLANs.

4. Double-click the Bridging Configuration object, double-click the VLAN Configuration object and then double-click the Protocol Based VLANs object.
5. Select the VLAN that will be replacing the interface's existing VLAN.
6. Drag the selected interface to the appropriate interface.

Note: You can drag only IP VLANs to IP interfaces, and you can drag only IPX VLANs to IPX interfaces.

Changing Which Ports a VLAN Includes

Configuration Expert lets you change the ports of a VLAN. You can add ports by dragging them to the VLAN. You can also add and remove ports using a dialog box. Separate discussions on each operation follow.

Dragging Ports to Add Them to a VLAN

You can quickly add ports to a VLAN by copying those ports from a port list located elsewhere in the configuration tree. To do so, take the following steps:

1. Start Configuration Expert if you have not already done so.

- Note:** If you start Configuration Expert from the Schematic view using the **Configure VLAN** command, Configuration Expert automatically expands the Active Configuration file's tree to the VLAN Configuration object. If you are adding a VLAN to the Active Configuration file, go to step 3.
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
 3. Double-click the VLAN Configuration object.
 4. Do one of the following:
 - If you want to copy ports to a port-based VLAN, double-click the Port Based VLANs object.
 - If you want to copy ports to a protocol-based VLAN, double-click the Protocol Based VLANs object.
 5. From the list of VLANs that appears, double-click the VLAN you want to modify and then double-click that VLAN's Bound Port List object.

Configuration Expert displays the modules of the ports that are currently included in the VLAN.
 6. Expand the configuration tree until the ports you want to copy appear. You can copy ports from any port list such as those found in the Chassis Configuration, STP Enabled, STP Disabled, Access Ports, or Trunk Ports objects.
 7. Click the port you want to copy. If you want to copy all of a module's ports, click that module.
 8. Drag the port to the Bound Port List object of the VLAN you want to copy the port to and release the mouse button when a green check mark appears.

Adding and Removing a VLAN's Ports Through a Dialog Box

To use a dialog box to add or remove ports of a port-based or protocol-based VLAN:

1. Start Configuration Expert if you have not already done so.

Note: If you start Configuration Expert from the Schematic view choosing the **Configure VLAN** command, Configuration Expert automatically expands the Active Configuration file's tree to the VLAN Configuration object. If you are adding a VLAN to the Active Configuration file, go to [step 3](#).
2. Open the configuration file you want to modify and then double-click that file's Bridging Configuration object.
3. Double-click the VLAN Configuration object.
4. Do one of the following:
 - If you want to change the ports of a port-based VLAN, double-click the Port Based VLANs object.

- If you want to change the ports of a protocol-based VLAN, double-click the Protocol Based VLANs object.
5. From the list of VLANs that appears, double-click the VLAN you want to modify and then click that VLAN's Bound Port List object.

A **Bound Port List** dialog box similar to the following appears:

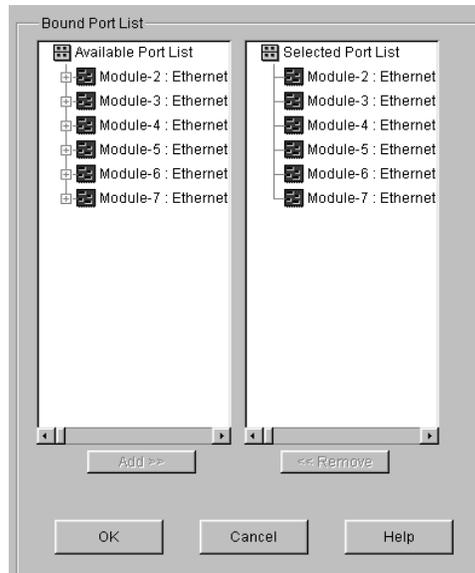


Figure 50. Bound Port list dialog box

6. Add and remove ports from the VLAN as necessary.
 - To add a port, double-click that port's module in the **Available Port** list. From the list of ports that appears, select the port you want to add and click the **Add** button. Configuration Expert moves the selected port from the **Available Port** list to the corresponding module in the **Selected Port** list.
 - To remove a port, double-click that port's module in the **Selected Port** list box. From the list of ports that appears, select the port you want to remove and click the **Remove** button. Configuration Expert moves the selected port from the **Selected Port** list to the corresponding module in the **Available Port** list box.

Clicking a module rather than double-clicking it in a list box selects all of the module's ports in that list box. This is a quick way to select all of a module's ports if you want to add or remove them all at the same time.
7. After you add and remove the desired ports, click **OK**.

Chapter 8

Configuring IP Interfaces for the GSR

Configure IP interfaces for the GSR if you want to use the GSR for IP-based unicast or multicast traffic. This chapter

- provides an overview of IP.
- describes creating and modifying IP interfaces.
- indicates what you need to do after you create IP interfaces.
- discusses configuring the GSR for VRRP.

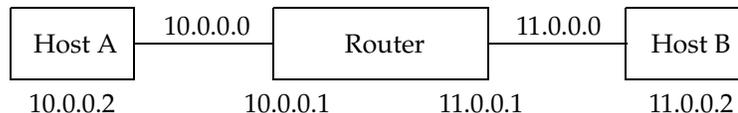
What Is IP?

The Internet Protocol (IP) is a Layer-3 (network) protocol that provides addressing and control information the GSR needs to route data packets in a network. The GSR supports IP routing as specified in RFC 1812. (IP is often referred to as TCP/IP.)

Acting as a connectionless packet-delivery service, IP does not provide a dedicated link between computers. Instead, IP provides dynamic routing by using best-effort delivery to route packets through any number of paths in the network. This dynamic routing is enabled through the routing table IP maintains. This table consists of pairs of destination addresses and next hop addresses. IP uses the table to determine that to get to a specific network address its next hop should be to a particular interface.

Suppose you have two hosts connected together with a router as shown in the following figure. Host A, which is on network 10.0.0.0, has an IP address of 10.0.0.2. Host B, which is

on network 11.0.0.0, has a network address of 11.0.0.2. The router has two interfaces and connects to both networks. (You can say the router is “well connected.”)



To communicate with one another through the router, the hosts make entries in routing tables as shown in the following table. For example, host A communicates with host B by using the route specified with the 11.0.0.0 10.0.0.1 entry in the host A routing table. That entry states that host A uses the router at 10.0.0.1 as a gateway when it wants to communicate with host B on the 11.0.0.0 network.

Table 10. Network destination and gateway routes

Host A		Router		Host B	
Destination	Gateway	Destination	Gateway	Destination	Gateway
10.0.0.0	10.0.0.2	11.0.0.0	11.0.0.1	11.0.0.0	11.0.0.2
11.0.0.0	10.0.0.1	10.0.0.0	10.0.0.1	10.0.0.0	11.0.0.1

The previous table provides information about network routes rather than host routes. If the previous table was providing information about host routes rather than network routes, each host’s routing table would list the other host’s IP address as the destination.

When a packet is sent, the entire route is not known at the source. Instead, packets travel to their destination one hop at a time. At each stop, a router calculates the next destination by matching the destination address with entries in that router’s routing table. This is similar to sending a letter through the United States postal service. The envelope has the ultimate destination address. Each post office along the way checks the ultimate destination and determines the next-hop (the next post office). When the letter gets to the post office that is serving the address on the envelope, the letter is delivered directly to the address.

In addition to providing connectionless deliveries, IP fragments and reassembles packets. IP does not guarantee that the packets will arrive in order or at all, nor does IP report packet errors back to the source. Guarantee delivery is the responsibility of TCP, and the reporting of packet errors is the responsibility of ICMP.

A Look at IP Addresses

When configuring the GSR, you will be required to enter IP address information. These 32-bit addresses, which identify a **connection to a network** rather than a particular system, include both a network ID and a host ID. If a network is part of the internet, the

network number is assigned by the Internet Network Information Center (InterNIC). Host numbers are assigned by you or another local network administrator.

The 32-bits of an IP address are grouped into four eight-bit octets, which are separated by decimal points. Each bit in an octet has a binary weight (128, 64, 32, 16, 8, 4, 2, 1). Each octet is represented in decimal format and ranges in value from 0 to 255.

There are two forms of IP addresses. One form is the dotted-quad and subnetwork mask format, which is *x.y.z.w n.n.n.x*. The other form is the Classless Inter-Domain Routing (CIDR) format, which is *x.y.z.w/xx*. Configuration Expert supports the dotted-quad and subnetwork mask format, but the CLI supports both formats. An IP address of all zeroes (0.0.0.0 0.0.0.0) means this machine on this network. An IP address of all ones (1.1.1.1 1.1.1.1) is a broadcast IP address or all hosts on all networks. Broadcast IP addresses are used to inform a router to send a packet to all hosts on a network.

Traditional IP addresses are divided into network classes, and the left-most (high-order) bit indicates the network class. The following table describes the classes used for unicast and multicast IP addresses. As the table shows, you can determine which class an address belongs to by examining the decimal value of the first octet. For example, you can determine that 129.84.6.0 is a class B IP address because the first octet of the address has a decimal value of 129.

Table 11. IP address classes

IP Address Class	Used for	Address Range	First Octet in Decimal	High-Order Bits
A	Some large organizations	1.0.0.0 to 126.0.0.0	1 to 126	0
B	Medium-size organizations	128.1.0.0 to 191.254.0.0	128 to 191	10
C	Small organizations	192.0.1.0 to 223.255.254.0	192 to 223	110
D	Multicast groups	224.0.0.0 to 239.255.255.255	224 to 239	1110

Using a netmask, the host part of the IP addresses can be divided into smaller networks. These smaller networks are called subnetworks (or subnets) and provide a more efficient way of allocating IP addresses. IP addresses broken down into subnets have the following format: | Network | Subnet | Host |. For example, 129.84.1.0, 129.84.2.0, 129.84.3.0, and 129.84.4.0 are all subnetworks on network 129.84.0.0.

Creating IP Interfaces

Create IP interfaces if you plan to configure the GSR for IP-based unicast or multicast network traffic.

When you create IP interfaces on the GSR, you provide information about the interface (such as its name, IP address, netmask, broadcast address, and so on). You also enable or disable the interface and bind the interface to a single physical port or VLAN. If you want to apply an existing ACL to an interface, Configuration Expert lets you do so either when you create the interface or afterwards.

You can bind each IP interface you create to a MAC address as discussed in the *DIGITAL GIGAswitch/Router Command Line Interface Reference Manual*. Binding an interface to a MAC address is optional. If you do not bind an interface to a MAC address, the interface uses the system address.

When creating or monitoring interfaces, you should realize that interfaces bound to a single port go down when the port goes down but interfaces bound to a VLAN remain up as long as at least one port in that VLAN remains active. The following table summarizes this:

If	Then
The port of a port-bound interface goes down	The interface bound to that port also goes down
At least one port of a VLAN remains active	The interface bound to the VLAN remains active
All the ports of a VLAN go down	The interface bound to the VLAN also goes down

The procedure for creating an IP interface depends on whether you are binding that interface to a single port or a VLAN. Separate discussions on the different procedures follow.

Creating IP Interfaces Bound to a Single Port

To create an IP interface that you want bound to a single port:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Interface Configuration object and then click the Configure New IP interface object.

Configuration Expert opens the IP Interface wizard.



Figure 51. IP Interface wizard (single port)

5. Click Next.

An IP Interface Definition panel similar to the following appears:

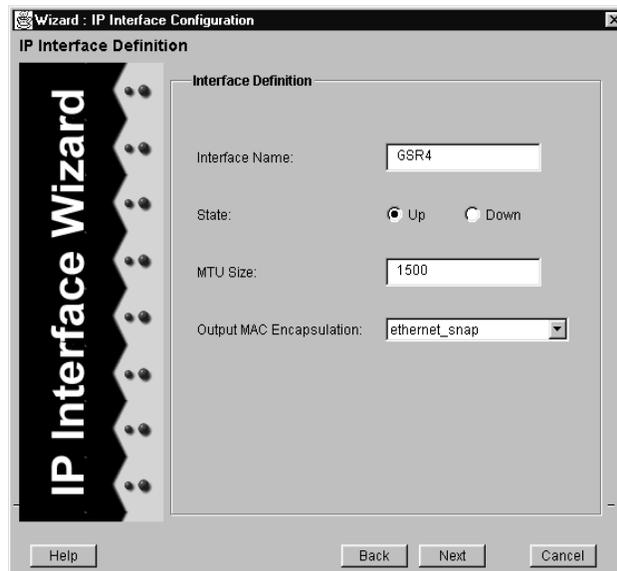


Figure 52. Interface Definition panel (single port)

6. Enter the name of the interface in the Interface **Name** box. Then either select *Up* to enable the interface or select *Down* to disable it.
7. Enter the number of bytes you want to specify for the Maximum Transmission Unit (MTU).

The MTU is the largest packet size that the GSR can transmit on the network via the interface. Any messages larger than the MTU are divided into smaller packets before being sent. You should set the MTU equal to the smallest MTU of all the networks between the GSR and a message's final destination. Otherwise, transmission speed slows down because the packets have to be fragmented.

8. Set the output MAC encapsulation you want associated with the interface by selecting one of the following from the **Output MAC Encapsulation** drop-down list:
 - *ethernet_II* (the default)
 - *ethernet_snap*
9. Click **Next**. Then enter the IP address, subnetwork mask, and broadcast address of the interface in the panel that appears.

You can specify the IP address and subnetwork mask values using the traditional format (example: 10.1.2.3/255.255.0.0).

Specifying a broadcast address is optional.

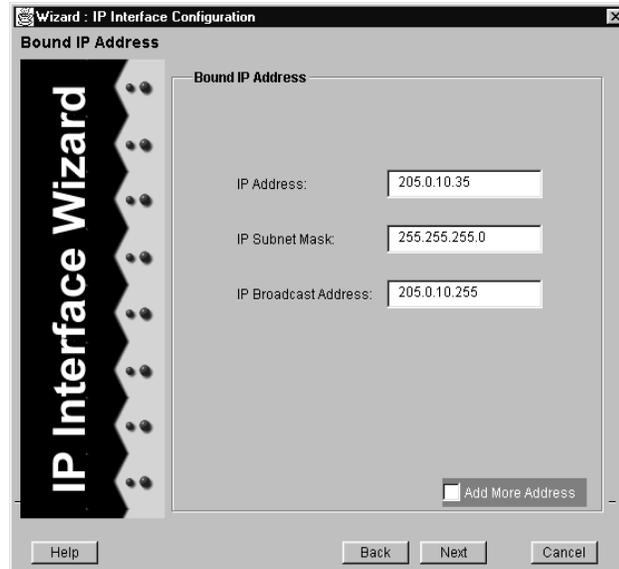


Figure 53. Bound IP Address panel (single port)

10. Click **Next** and then select the *Bind the interface to Port* option in the panel that appears.

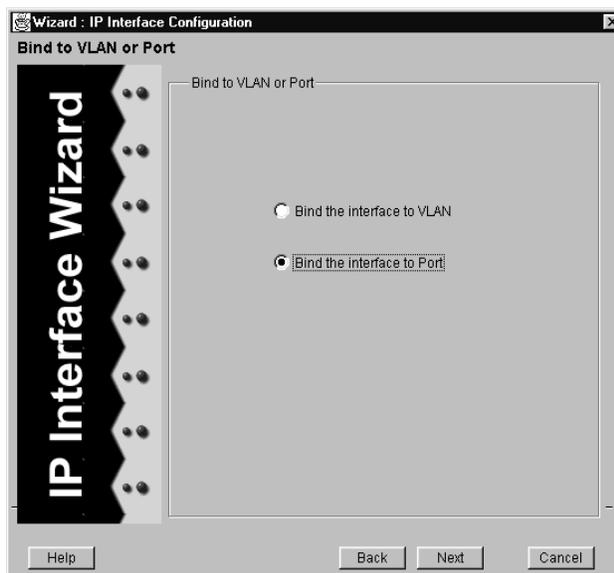


Figure 54. Bind to VLAN or Port panel (single port)

11. Click Next.

Configuration Expert displays a Bound Port List panel.

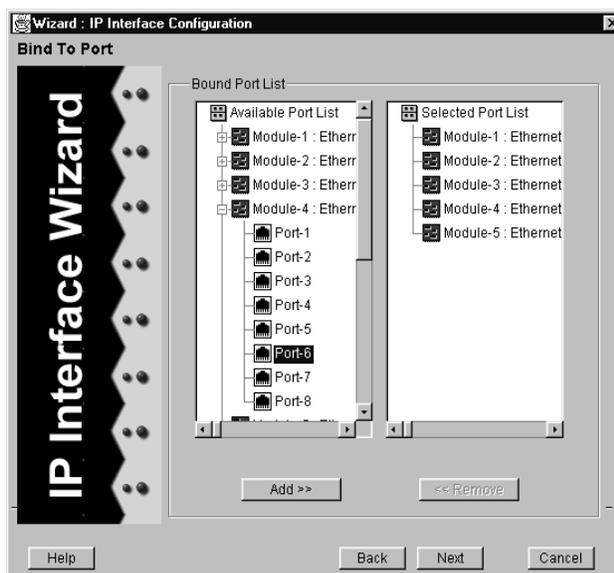


Figure 55. Bound Port List panel (single port)

12. Bind the interface to a single port by doing the following:

- a. In the **Available Port** list, double-click the module containing the desired port. From the list of available ports that appears, select the port that you want to bind to the interface.
- b. Click the **Add** button.

Note: You can bind only a single port to an interface. If you need to bind multiple ports to the interface, create a VLAN consisting of those ports and bind the interface to that VLAN.

13. Click **Next**.

The Apply ACLs panel appears.

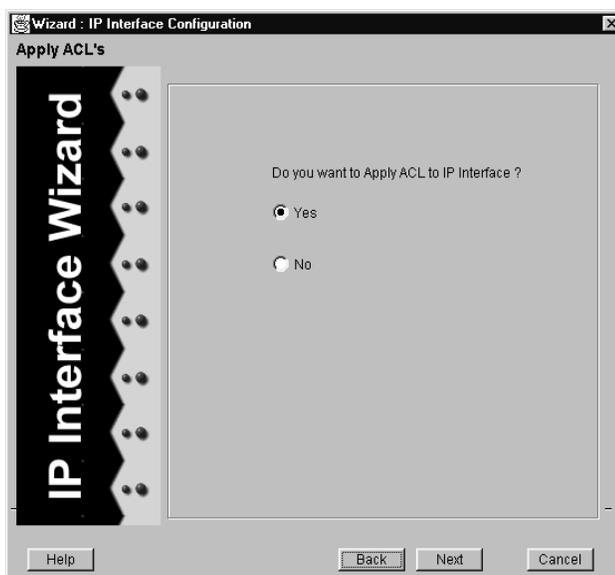


Figure 56. Apply ACLs panel (single port)

14. Specify whether you want to apply an ACL to the interface by doing one of the following:

- To not apply an ACL, select *No* and then click **Finish**. This completes the configuration of the interface.

Configuration Expert adds the new interface to those found in the IP interfaces bound to Ports object.

- To apply an ACL to the interface, select *Yes* and click **Next**.

Note: You will be able to apply an IP ACL while creating an interface only if you previously created the ACL. If you plan to apply an ACL that you have not defined yet, finish creating the interface. Then configure the

desired ACL and apply it as discussed in “Applying ACLs to IP or IPX Interfaces” on page 206.

15. If you specified you wanted to apply an ACL, use the Apply IP ACL panel that appears to apply an ACL to the interface.

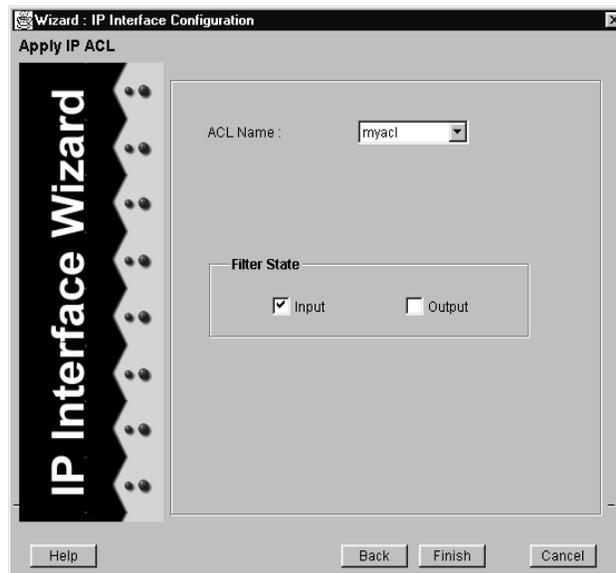


Figure 57. Apply IP ACL panel (single port)

To use the wizard to apply an IP ACL, take the following steps:

- a. Select the ACL you want to apply from the **ACL Name** drop-down list. This will apply all ACLs with that name.

Configuration Expert lets you configure multiple ACLs that use different rules but have the same name. You can only apply IP ACLs to an IP interface.

- b. Use the *Filter State* check boxes to specify whether you want to filter out inbound traffic, outbound traffic, or both inbound and outbound traffic.

Select *Input* to filter inbound traffic and select *Output* to filter outbound traffic. Selecting both check boxes filters both inbound and outbound traffic.

- c. Click **Finish**.

Configuration Expert adds the new interface to those found in the IP interfaces bound to Ports object.

Note: When you apply an ACL to an interface, the GSR appends an implicit deny rule to that ACL. The implicit deny rule denies all traffic. If you intend to allow all traffic that does not match your specified ACL rules to

go through, you must explicitly define a rule to permit all traffic. To do so, make sure the last rule of the ACL permits all traffic.

Creating IP Interfaces Bound to a VLAN

If you have created an IP VLAN, you can bind that VLAN to an IP interface while creating the interface. To create an IP interface that will be bound to an existing VLAN:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Interface Configuration object and then click the Configure New IP interface object.

Configuration Expert opens the IP Interface wizard.

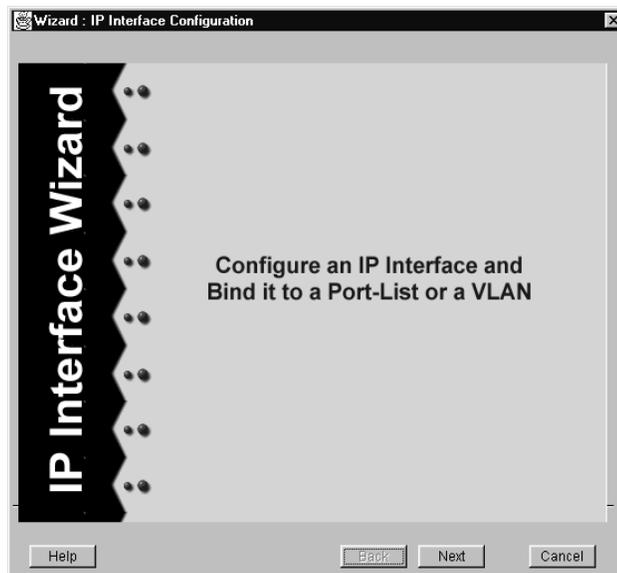


Figure 58. IP Interface wizard (VLAN)

5. Click **Next**.

An IP Interface Definition panel similar to the following appears:

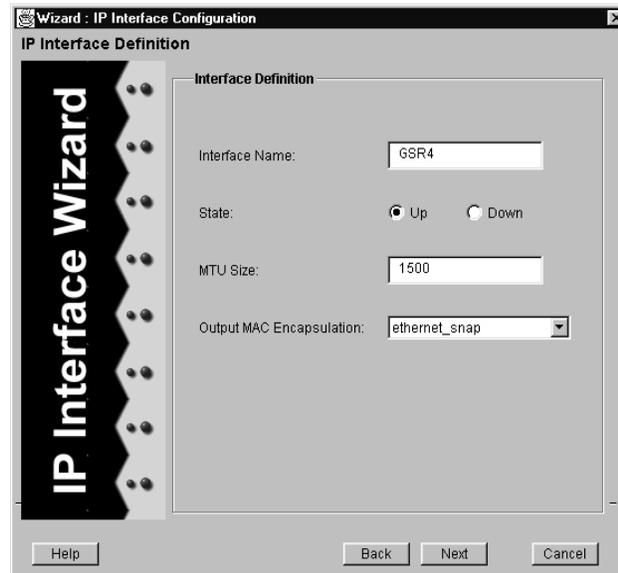


Figure 59. Interface Definition panel (VLAN)

6. Enter the name of the interface in the **Interface Name** box, then either select *Up* to enable the interface or select *Down* to disable it.
7. Enter the number of bytes you want to specify for the Maximum Transmission Unit (MTU).

The MTU is the largest packet size that the GSR can transmit on the network via the interface. Any messages larger than the MTU are divided into smaller packets before being sent. You should set the MTU equal to the smallest MTU of all the networks between the GSR and a message's final destination. Otherwise, transmission speed slows down because the packets have to be fragmented.

8. Set the output MAC encapsulation you want associated with the interface by selecting one of the following from the **Output MAC Encapsulation** drop-down list:
 - *ethernet_II* (the default)
 - *ethernet_snap*
9. Click **Next**. Then enter the IP address of the interface, subnetwork mask, and broadcast address of the interface in the panel that appears.

You can specify the IP address and subnetwork mask values using the traditional format (example: 10.1.2.3/255.255.0.0).

Specifying a broadcast address is optional.

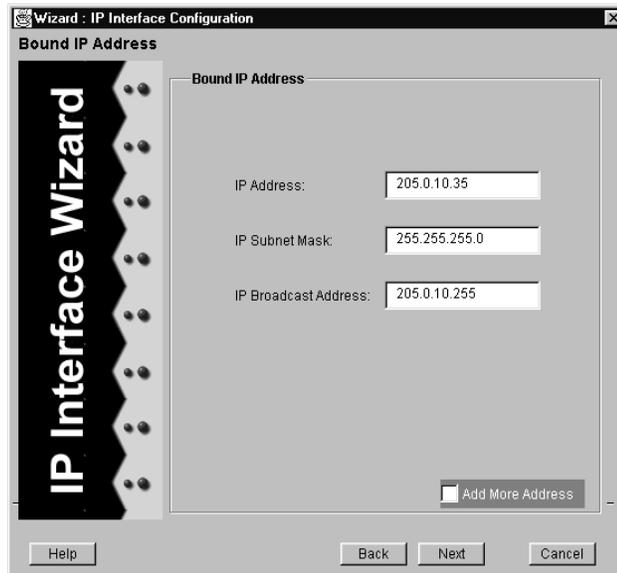


Figure 60. Bound IP Address panel (VLAN)

10. Click **Next** and then click the *Bind the interface to VLAN* option in the panel that appears. That option is available only if there are existing IP VLANs.

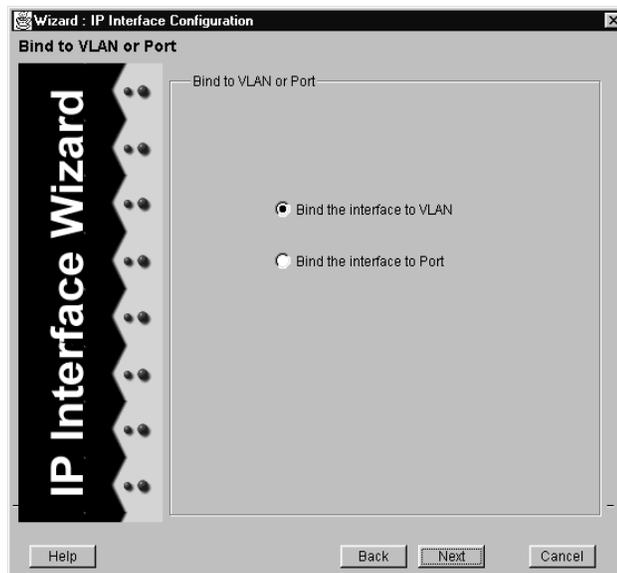


Figure 61. Bind to VLAN or Port panel (VLAN)

11. Click **Next**. In the panel that appears, select the name of the VLAN you want to bind to the interface.

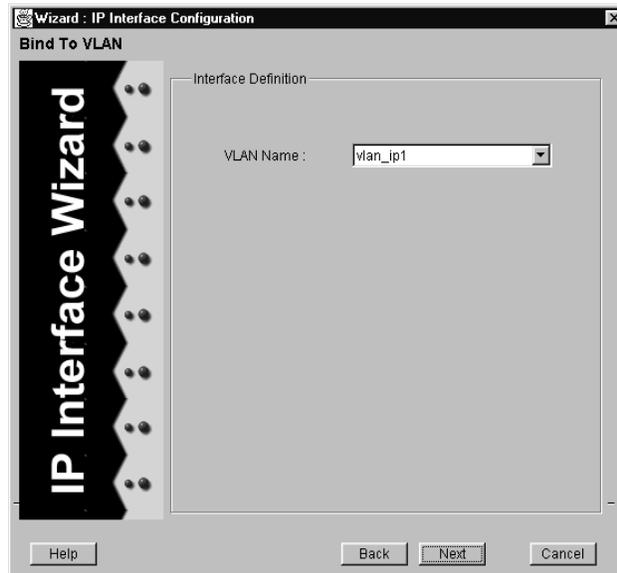


Figure 62. Interface Definition panel (VLAN)

12. Click Next.

The Apply ACLs panel appears.

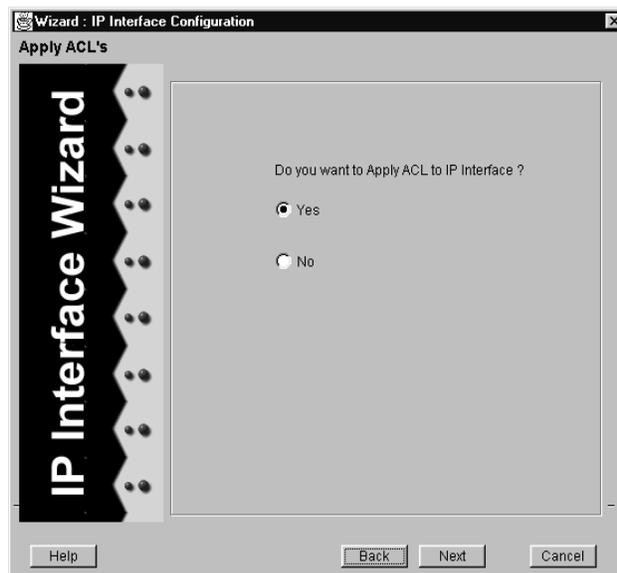


Figure 63. Apply ACLs panel (VLAN)

13. Specify whether you want to apply an ACL to the interface by doing one of the following:

- To not apply an ACL, select *No* and then click **Finish**. This completes the configuration of the interface but does not apply any ACL to that interface.

Configuration Expert adds the new interface to those found in the IP interfaces bound to VLAN object.

- To apply an ACL to the interface, select *Yes* and click **Next**.

Note: You will be able to apply an IP ACL while creating an interface only if you previously created the ACL. If you plan to apply an ACL that you have not defined yet, finish creating the interface. Then configure the desired ACL and apply it as discussed in [Chapter 13: “Configuring Security on the GSR”](#) on page 189.

14. If you specified you wanted to apply an ACL, use the Apply IP ACL panel that appears to apply an ACL to the interface.

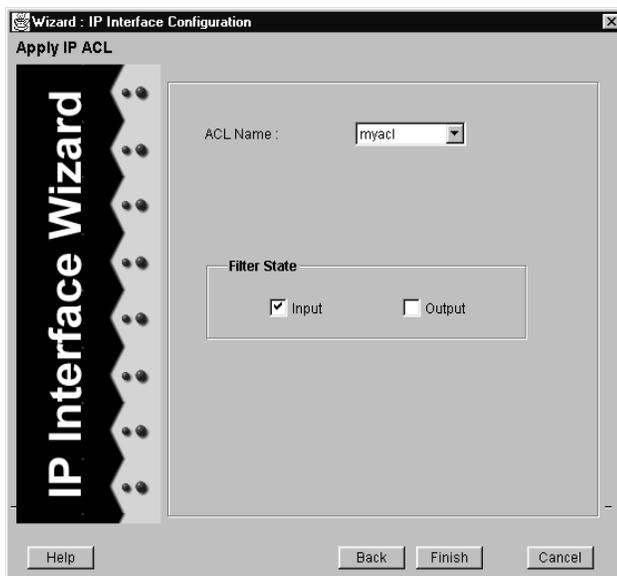


Figure 64. Apply IP ACL panel (VLAN)

To use the wizard to apply an IP ACL, take the following steps:

- a. Select the ACL you want to apply from the **ACL Name** drop-down list. This will apply all ACLs with that name.

Configuration Expert lets you configure multiple ACLs that use different rules but have the same name. You can only apply IP ACLs to IP interfaces.

- b. Use the *Filter State* check boxes to specify whether you want to filter out inbound traffic, outbound traffic, or both inbound and outbound traffic.

Inbound traffic is packets coming into the interface while outbound traffic is packets going out of that interface.

Select *Input* to filter inbound traffic and select *Output* to filter outbound traffic. Selecting both check boxes filters both inbound and outbound traffic.

- c. Click **Finish**.

Configuration Expert adds the new interface to those found in the IP interfaces bound to VLAN object.

Note: When you apply an ACL to an interface, the GSR appends an implicit deny rule to that ACL. The implicit deny rule denies all traffic. If you intend to allow all traffic that does not match your specified ACL rules to go through, you must explicitly define a rule to permit all traffic. To do so, make sure the last rule of the ACL permits all traffic.

Modifying IP Interface Definitions

Modify IP interface definitions to perform the following operations:

- Change the interface's name
- Disable or enable the interface
- Change the interface's MTU
- Change the interface's MAC encapsulation
- Change the IP address of the interface
- Bind a different port or VLAN to the interface
- Apply a different ACL to the interface
- Change which traffic an interface's ACL filters out

To modify an IP interface definition:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Interface Configuration object. Then do one of the following:
 - If the interface you want to modify is bound to a port, double-click the IP Interfaces bound to Ports object.
 - If the interface you want to modify is bound to a VLAN, double-click the IP Interfaces bound to VLAN object.

5. In the list of interfaces that appears, double-click the one you want to modify.

Configuration Expert displays the contents of the object and the **Interface Definition** dialog box of the interface.

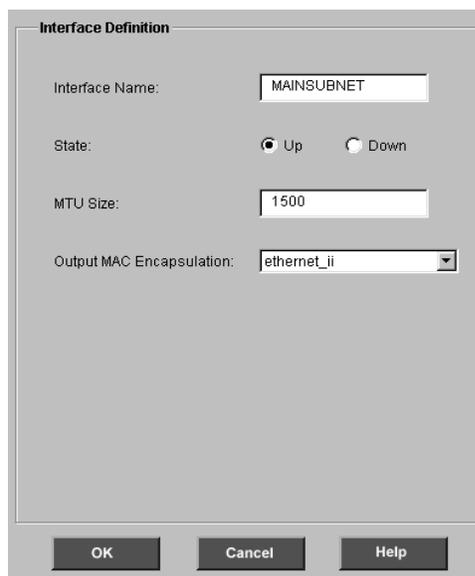


Figure 65. Interface Definition dialog box

6. If you want to edit the name, interface state, MTU, or MAC encapsulation fields, specify values as you do when creating an IP interface. Then click **OK**.
7. If you want to change the interface's IP address, IP subnet mask, or broadcast address, double-click the Bound IP Addresses object, click the interface's address object, then enter a new IP address in the Bound Port Address dialog box that appears and click **OK**.
8. If you want to bind the interface to a different port or VLAN, do one of the following:
 - If you are modifying a port-bound interface, double-click the interface's Bound Port List object. Then use the Bound Port List dialog box to remove the currently bound port and to add the new port. Click **OK**

You can add and remove ports from the dialog box's list boxes as you do when creating the interface.
 - If you are modifying a VLAN-bound interface, click the interface's VLAN object. Then select a new VLAN from the **VLAN Name** drop-down list that appears in the Interface Definition dialog box. Click **OK**.
9. If you want to change which ACLs are applied to the interface, double-click the Bound IP Security object. Use the **Update ACL List** dialog box that appears to add and remove ACLs.

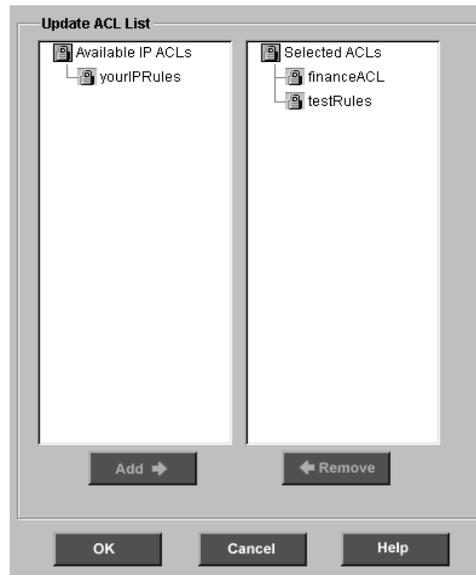


Figure 66. Update ACL List dialog box

You can add an ACL by selecting it in the **Available IP ACLs** list and then clicking **Add**. You can remove an ACL by selecting it in the **Selected ACLs** list and then clicking **Remove**.

Note: You may also apply an ACL by copying it as discussed in [“Copying an ACL to Apply It to an Interface”](#) on page 207.

10. Click **OK**.
11. If you want to change the type of traffic the ACL filters out on the interface, click the IP ACL located in the interface’s Applied IP ACL object.

The interface’s **Edit ACL** dialog box appears.

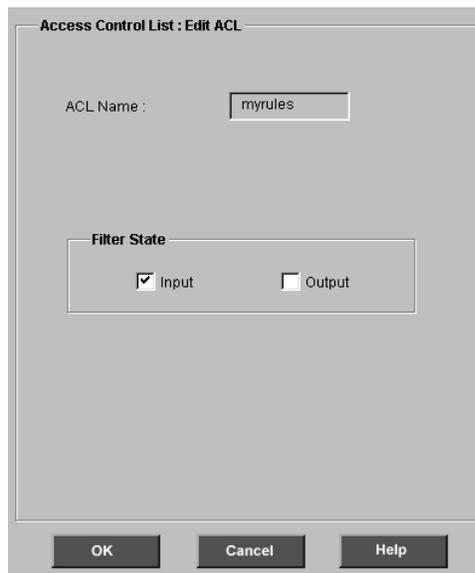


Figure 67. Edit ACL dialog box

12. Select or clear the *Filter State* check boxes to specify whether you want the ACL to filter inbound traffic (*Input*), outbound traffic (*Output*), or both input and outbound traffic. Click **OK**.



Caution: An IP interface can have up to two ACLs applied to it. If you applied two ACLs to an interface, one ACL must govern the inbound traffic and the other ACL must govern outbound traffic.

Configuring the GSR for VRRP

You can configure the GSR to operate with Virtual Router Redundancy Protocol (VRRP). End host systems on a LAN are often configured to send packets to a statically configured default router. If this default router becomes unavailable, all the hosts that use it as their first hop router become isolated, unable to send packets. VRRP provides a way to ensure that the default router on an end host is always available.

It does this by assigning IP addresses that end hosts use as their default route to a “virtual router.” A Master router is assigned to handle routes designated for the virtual router. If the Master router should become unavailable, a Backup router takes over and begins handling routes for the virtual router.

Note: You must specify at least one IP interface for your GSR before attempting to configure for VRRP.

Specifying VRRP Trace Options

To specify VRRP Trace options:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the VRRP Configuration object.
5. Click the VRRP Trace Options object.

The **VRRP Trace Options** dialog box appears:

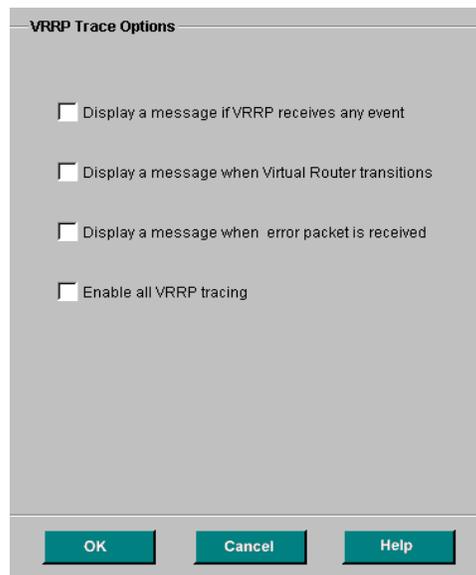


Figure 68. VRRP Trace Options dialog box

6. Select or deselect one or more of the four Trace options:
 - *Display a message if VRRP receives any event*
 - *Display a message when Virtual Router changes state*
 - *Display a message when error packet is received*
 - *Enable all VRRP tracing*

Note: Selecting the fourth option, *Enable all VRRP tracing*, automatically turns on the other three options.

7. Click **OK**.

Configuring a New VRRP Router

To set up your GSR for VRRP:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the VRRP Configuration object.
5. Double-click the VRRP Router Configuration object, then click the Create New VRRP Router object.

Configuration Expert opens the VRRP wizard:

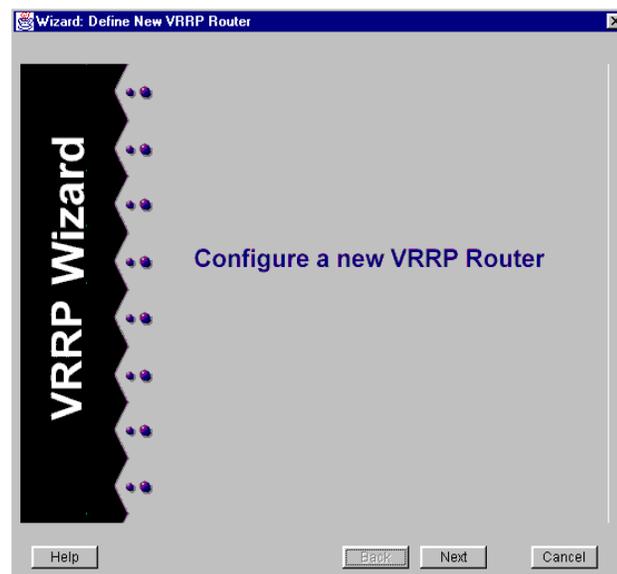


Figure 69. VRRP wizard

6. Click **Next**.

The VRRP Interface Definition panel appears:



Figure 70. VRRP Interface Definition panel

7. Select the name of your VRRP interface from the **Interface Name** drop-down list, then click **Next**.

The VRRP Router panel appears:

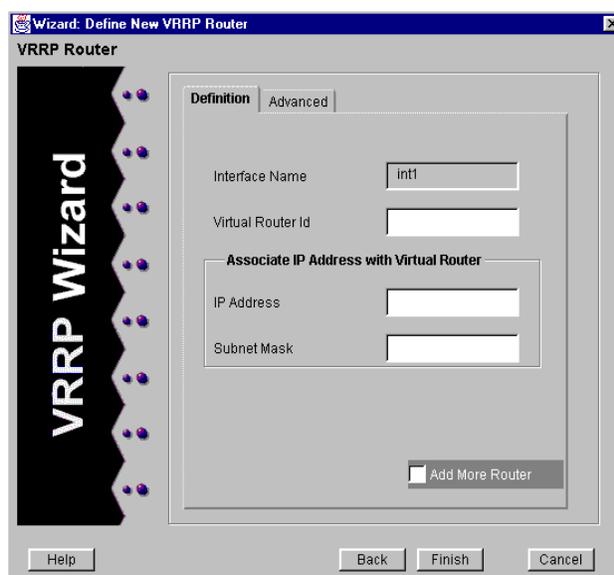


Figure 71. VRRP Router panel (Definition tab)

8. Specify a router identification number in the **Virtual Router ID** box. You *must* specify a router identification number in order to complete VRRP configuration.

Note: The identification you enter must consist of digits only. No other characters are allowed in the identification name.

9. If you wish, under *Associate IP Address with Virtual Router*, specify the IP address and subnet mask of the associated router in the appropriate text boxes.
10. Click the Advanced tab:

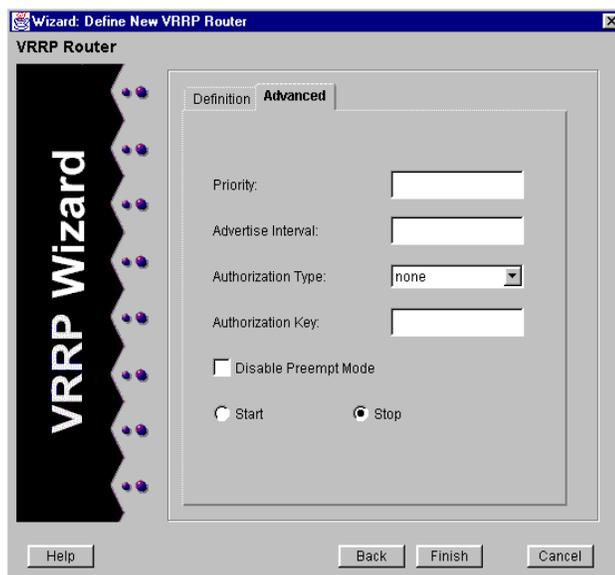


Figure 72. VRRP Router panel (Advanced tab)

11. If you wish, specify the following options on the Advanced tab of the VRRP Router panel:
 - a. Specify values in the **Priority** and **Advertise Interval** boxes. You can specify a number between 1 and 254, inclusive, for the priority, and a number between 1 and 255, inclusive, for the advertise interval.
 - b. Select an authorization type from the *Authorization Type* drop-down list.
 - c. Specify an authorization key in the **Authorization Key** box. You can specify a key string up to eight characters in length.
 - d. Specify whether you want to disable preempt mode by turning the *Disable Preempt Mode* option on or off.
 - e. Specify *Start* or *Stop* to activate or deactivate VRRP on the GSR.
12. Click **Finish**.

What to Do Next

After you define your IP interfaces, you need to configure your GSR for unicast or multicast routing. For details on unicast routing, see [Chapter 9: “Configuring Unicast Routing on the GSR”](#) on page 113. For details on multicast routing, see [Chapter 10: “Configuring Multicast Routing on the GSR”](#) on page 129.

Chapter 9

Configuring Unicast Routing on the GSR

IP can perform unicast routing using RIP. Configuration Expert lets you configure the GSR for RIP. This chapter

- discusses configuring unicast parameters and static entries that can be used by RIP.
- provides an overview of RIP.
- discusses configuring RIP global parameters, defining IP RIP interfaces, and adding RIP gateways.

In addition to configuring the GSR for unicast routing, you can configure the GSR for multicast routing and IPX. For details on multicast routing, see [Chapter 10: “Configuring Multicast Routing on the GSR” on page 129](#). For details on IPX, see [Chapter 11: “Configuring the GSR for IPX Routes” on page 141](#).

Note: Before you begin configuring your GSR for unicast routing, you should create the IP interfaces you want to use for unicast routing. For details, see [“Creating IP Interfaces” on page 92](#).

Configuring Unicast Global Parameters and Static Entries

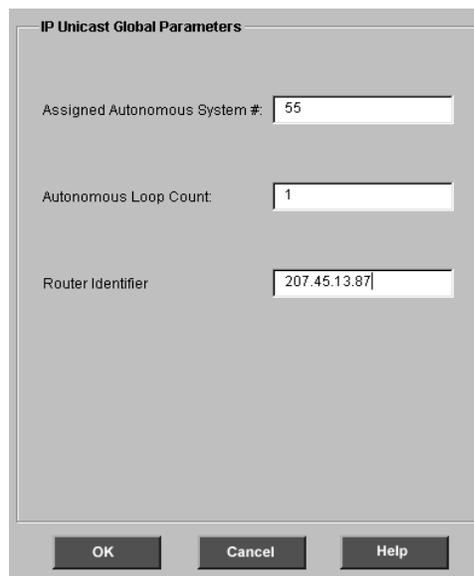
Configuration Expert lets you configure settings that will be applied to RIP. It also lets you set up static route entries and static ARP entries. Separate discussions on setting unicast global parameters and configuring the different static entries follow.

Setting Global Parameters for Unicast Routing

Configuration Expert lets you configure global unicast routing parameters for RIP. To configure these global parameters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object. Then click the Unicast Global Parameters object.

An **IP Unicast Global Parameters** dialog box similar to the following appears:



The screenshot shows a dialog box titled "IP Unicast Global Parameters". It contains three input fields with the following values: "Assigned Autonomous System #" is 55, "Autonomous Loop Count" is 1, and "Router Identifier" is 207.45.13.87. At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

Figure 73. IP Unicast Global Parameters dialog box

5. Enter a number from 1 to 65534 to specify the GSR's autonomous system number in the **Assigned Autonomous System #** box. The default is 1.
6. In the **Autonomous Loop Count** box, enter the number of times the Autonomous System may appear in the Autonomous System path. The default is 1.

This entry is only required if the router is going to run protocols that support autonomous system paths, such as the Border Gateway Protocol (BGP).

7. Enter the router ID that is for OSPF's use.

The most preferred address for the router ID is any IP address other than 127.0.0.1 on the loopback interface. If there are no secondary addresses on the loopback interface,

then the router ID is set to the address of the first interface which is in the up state that the GSR encounters. The address of a non-point-to-point interface is preferred over the local address of a point-to-point interface.

8. Click **OK**.

Defining Static ARP Entries

ARP maps network addresses to MAC addresses. Define static ARP entries if you want the GSR to use those entries rather than using ARP to automatically resolve host and MAC address entries.

To define static ARP entries:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the Static ARP Entries object. Then click the New ARP Node object.

Configuration Expert opens the ARP wizard.

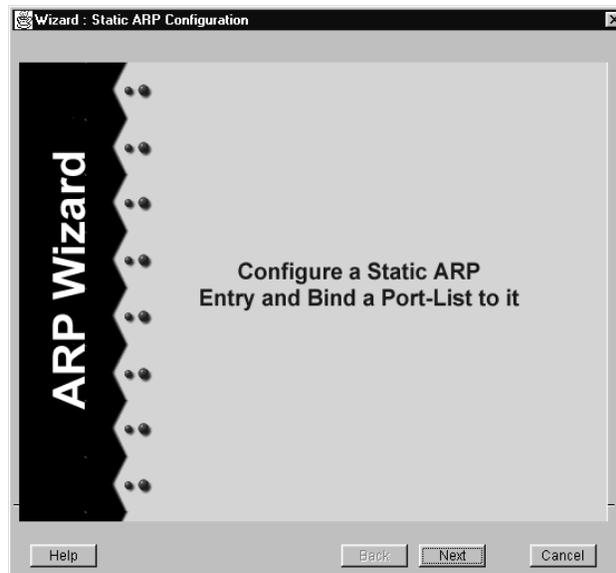


Figure 74. ARP wizard

5. Click **Next**.

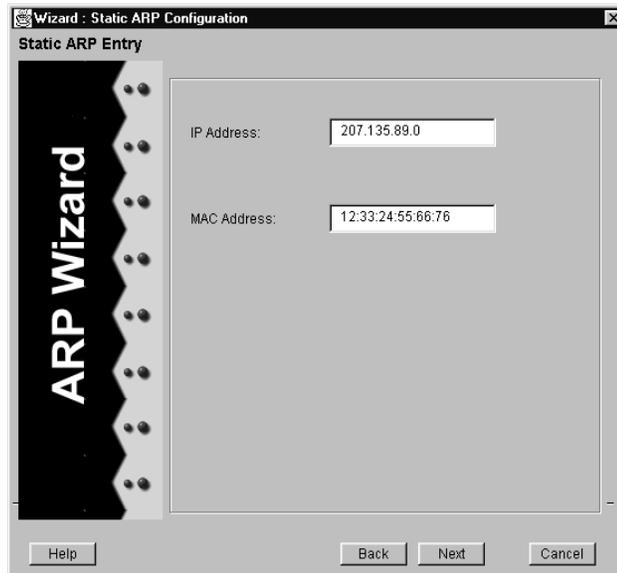


Figure 75. Static ARP Entry panel

6. Enter the IP address and MAC address of the host's ARP entry in the appropriate text boxes, then click **Next**.

Configuration Expert displays a Bound Port List panel.

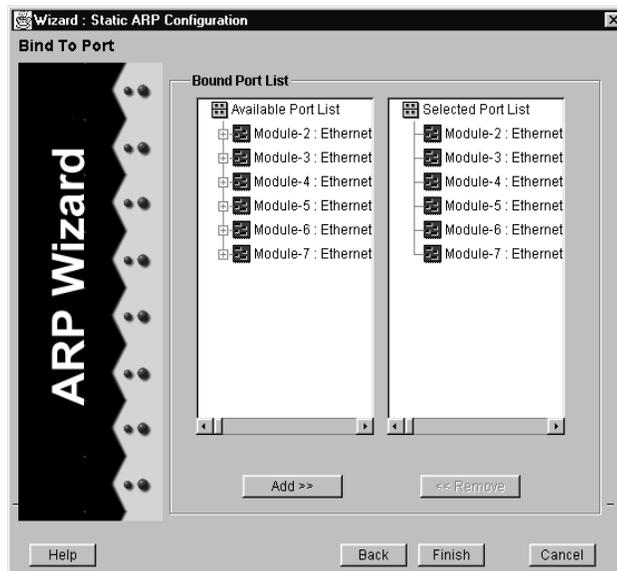


Figure 76. Bound Port list panel

7. Bind the entry to a port that the host is connected to by doing the following:
 - a. In the **Available Port** list, double-click the module containing the desired port. From the list of available ports that appears, select the port that you want to bind to the entry.
 - b. Click the **Add** button.

Note: You can bind only a single port to an entry.

If you accidentally add the wrong port, remove the port by selecting it from the **Selected Port** list and then clicking the **Remove** button.

8. Click **Finish**.

Defining Static Route Entries

Static route entries specify routes you want to explicitly configure and enter into the GSR's routing table. Define static route entries if you want to configure static routes used by the unicast routing protocol process. These routes may be overridden by routes with better preference. To define static route entries:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the Static IP Routes object. Then click the Create New IP Static Route Entry object.

An **IP Static Route Definition** dialog box similar to the following appears:

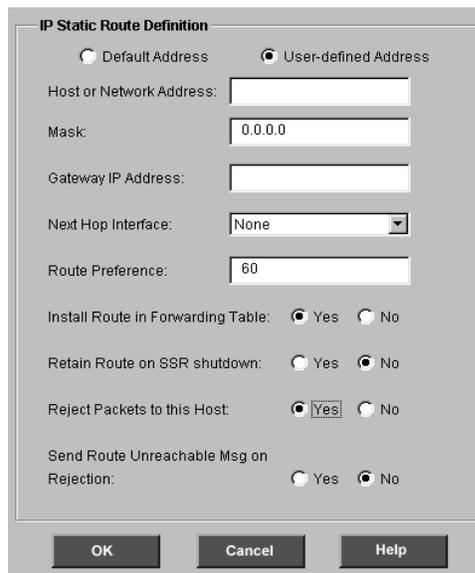


Figure 77. IP Static Route Definition dialog box

5. Enter the route's IP address and network mask in the appropriate text boxes.
6. In the **Gateway IP Address** box, enter the IP address of the next-hop gateway associated with the route.
7. In the **Next Hop Interface** box, enter the IP address of the next-hop interface associated with the route.
8. Leave the preference value set to its default (60) unless you want to assign a different preference for the static routes. To set the preference, enter a number from 0 to 255 in the **Route Preference** box.

As discussed in [“What Is Preference?” on page 120](#), you can set preferences in several places. The GSR uses preference values to determine the preference of routes from one protocol or peer over another.

Note: Do not change the default preference for static routes unless you fully understand the implications of doing so.

9. Specify whether you want to install the route in the forwarding table by selecting *Yes* or *No* for the *Install Route in Forwarding Table* option.

The GSR will export the static routes to other protocols even if it is not installing the route in the forwarding table.

10. Specify whether the route will be removed from the forwarding table during graceful shutdowns by selecting *Yes* or *No* for the *Retain Route on GSR shutdown* option.

11. Specify whether you want to cause packets to be dropped and unreachable messages to be sent to packet originators by selecting *Yes* or *No* for both the *Reject packets to this Host* and *Send Route Unreachable Msg on Rejection* options.
12. Click **OK**.

A Look at RIP Routing in the IP Environment

The Routing Information Protocol (RIP) is the most commonly used interior protocol. It is a distance-vector routing protocol for use in small networks.

The GSR supports RIP routing as specified in RFC 1721. The GSR supports both RIP v1 and RIP v2.

RIP sends update messages to RIP routers to provide information about a network's RIP routes. RIP sends these messages periodically or whenever the network's topology changes. If the GSR receives an update message that includes changes, it updates its routing table and then sends update messages to other routers on the network.

In the GSR's IP environment, each RIP interface can be configured for the following:

- To accept or not accept RIP packets from another RIP interface or gateway.
- To send or not send RIP packets to another RIP interface or gateway.

In addition to being able to specify the handling of RIP packets on individual interfaces, you can configure the GSR for the following:

- Trusted gateways

A trusted gateway is a router that is trusted to supply routing information. Configure the GSR for trusted gateways by defining a list of host names or IP addresses of the routers you want the GSR to accept messages from. The GSR will then accept update messages from those routers and reject update messages from other RIP routers.

If you do not define a list of trusted gateways, the GSR accepts update messages from all RIP routers on the network. If trusted gateways are specified, only updates from those gateways are accepted.

- Source gateways

A source gateway is a router to which the GSR will send RIP packets directly rather than sending them as multicast or broadcast packets. Configure the GSR for source gateways if you want to send routing information to specific routers. The source gateways are not affected if you configure a GSR interface not to send RIP packets as mentioned earlier.

Suppose the GSR is attached to a subnetwork that does not need to process RIP data. Also suppose the GSR is to communicate with a RIP router elsewhere on the network. You could configure the GSR to send RIP packets to the other RIP router but not send them to

the GSR's subnetwork. You could also configure the GSR to not accept RIP data from the subnetwork but to accept RIP packets from the other RIP router.

In the IP environment, RIP bases routing on a hop count. RIP only supports routes that have 0 to 15 hops. You can configure the GSR to adjust hop counts. Doing so lets you control which routes the GSR prefers to use to send or receive data when it is communicating with another RIP router.

What Is Preference?

Preference is the value the GSR uses to determine the preference of routes from one protocol or peer over another. You can set preference based on one network interface over another, from one protocol over another, or from one remote gateway over another. A default preference is assigned to each source from which the GSR receives routes. Preference values range from 0 to 255 with the lowest number indicating the most preferred route.

Preference may not be used to control the selection of routes within RIP or another Interior Gateway Protocol (IGP) because that is accomplished automatically by the protocol's metric.

Preference may be used to select routes from an Exterior Gateway Protocol (EGP), such as the Border Gateway Protocol (BGP), learned from different peers or autonomous systems.

As shown in the following table you can set preference in several places. The table identifies each item on which you can set preference and identifies the default preference value of each item. For details on setting the preference of static routes, see [“Defining Static Route Entries” on page 117](#). For details on setting the preference of RIP routes, see [“Setting RIP Global Parameters” on page 121](#). For details on setting the preference of all other items, see the *DIGITAL GIGAswitch/Router Command Line Interface Reference Manual*.

Table 12. Preferences and default values

Preference	Default
Directly connected networks	0
Open Shortest Path First (OSPF) routes	10
Static routes	60
RIP routes	100
Point-to-point interfaces	110
Routes to interfaces that are down	120

Table 12. Preferences and default values (Continued)

Preference	Default
Aggregate/generate routes	130
OSPF AS external routes	150
BGP routes	170

Note: Do not change the default preference values unless you fully understand the implications of doing so.

Even though you can set preference in several places, each route has only one preference value associated with it. The last or most specific preference value set for a route is used. The preference value is arbitrarily assigned and is used to determine the order of routes to the same destination in the GSR's routing table. The active route is chosen by the lowest preference value.

Configuring the GSR for RIP

If you are using the GSR in an IP RIP environment, you will need to configure the GSR for RIP. To do so, perform the following tasks:

1. Create IP interfaces that you want to use for RIP. For details, see [“Creating IP Interfaces” on page 92](#).
2. Enable RIP on the GSR and specify global parameters if necessary. For details on these tasks, see [“Setting RIP Global Parameters” on page 121](#).
3. Configure IP interfaces for RIP. For details on this task, see [“Defining IP RIP Interfaces” on page 123](#).
4. Configure the GSR to accept RIP updates from only specific sources. For details on this task, see [“Adding Trusted Gateways” on page 126](#).
5. Add routers that send RIP updates directly rather than using broadcast or multicast messages. For details on this task, see [“Adding Source Gateways” on page 127](#).

Setting RIP Global Parameters

Set global RIP parameters to enable or disable RIP on the GSR and to set RIP-specific parameters. To set the GSR's global RIP parameters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.

3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object. Then double-click the RIP Routing object and click the RIP Global Parameters object.

A **RIP Global Parameters** dialog box similar to the following appears:

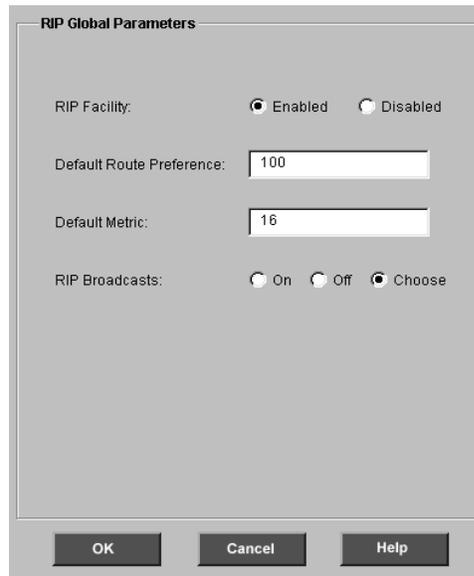


Figure 78. RIP Global Parameters dialog box

5. Specify whether you want to enable or disable RIP by selecting the appropriate option.

By default, RIP is disabled on the GSR.

6. Leave the preference value set to its default (100) unless you want to assign a different preference for the routes learned by RIP. To set the preference, enter a number from 0 to 255 in the **Default Route Preference** box.

The preference you specify applies to all IP RIP interfaces on the GSR.

Note: DIGITAL recommends against changing the default preference value for RIP routes unless you fully understand the implications of doing so.

As discussed in [“What Is Preference?” on page 120](#), you can set preferences in several places. The GSR uses preference values to determine the preference of routes from one protocol or peer over another. However, if a CLI command was used to set a preference in an import routing policy, that preference overrides the one you specify on the **RIP Global Parameters** dialog box.

7. Set the metric for routes advertised through RIP by entering a number from 1 to 16 in the **Default Metric** box.

Note: The metric 16 (the default) is equivalent in RIP to “infinite” and makes a route unreachable. You must set the default metric to a value other than 16 in order to allow the GSR to export routes from other protocols such as OSPF into RIP.

You can specify values for the default route preference and default metric in the **RIP Global Parameters** dialog box just as in the **RIP Policy Defaults** dialog box. If you change a parameter in one dialog box, that change is automatically applied to the other dialog box. See [“Setting RIP Routing Policy Defaults” on page 280](#) for more information.

8. Specify whether the GSR broadcasts RIP packets regardless of the number of interfaces present by doing one of the following:
 - Select *On* to configure the GSR to always send RIP broadcasts regardless of the number of interfaces present.
 - Select *Off* to configure the GSR to never send RIP broadcasts on attached interfaces.
 - Select *Choose* to configure the GSR to send RIP broadcasts only if more than one interface is configured on the GSR. This is the default state.
9. Click **OK**.

Defining IP RIP Interfaces

By default, RIP is disabled on all GSR IP interfaces. To enable RIP on an IP interface, you define an IP RIP interface to set the following parameters on that interface:

- Whether the interface will accept RIP updates
- Whether the interface will send RIP updates
- The RIP version (1 or 2)
- The packet type used for RIP updates (broadcast or multicast)
- The metric added to incoming RIP updates
- The metric added to outgoing RIP updates

To enable RIP on an IP interface or to modify the settings of an existing IP RIP interface:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file’s Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object and then double-click the RIP Routing object. Double-click the RIP Enabled Interfaces object.
5. Do one of the following:

- If you are modifying the settings of an interface on which RIP is enabled, select the interface you want to change.
- If you are configuring a new IP RIP interface, click the Configure New RIP Enabled Interface object.

A **RIP Interface Definition** dialog box similar to the following appears:

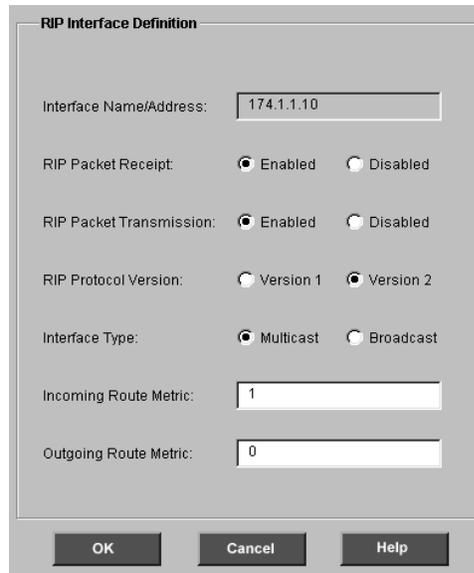


Figure 79. RIP Interface Definition dialog box

6. Select the IP interface you want to define as an IP RIP interface from the **Interface/Name Address** drop-down list.

The **Interface/Name Address** drop-down list includes existing IP interfaces. If you have not created the interface you want to configure for RIP, create one as discussed in [“Creating IP Interfaces” on page 92](#).

7. Set the RIP parameters described in the following table:

Table 13. RIP parameters

Parameter	Description
RIP Packet Receipt	<p>Specify whether the interface can receive RIP updates. Select <i>Enabled</i> if you want to receive RIP updates on the interface. Otherwise, select <i>Disabled</i>.</p> <p>This option does affect RIP updates sent from trusted gateways. If this option is disabled, the GSR will not receive any RIP updates including those sent from trusted gateways. If this option is enabled and trusted gateways have been set up, the GSR can accept updates only from those trusted gateways.</p>
RIP Packet Transmission	<p>Specify whether the interface can send RIP updates. Select <i>Enabled</i> if you want to send RIP updates from the interface. Otherwise, select <i>Disabled</i>.</p> <p>The setting of this option does not affect the sending of updates to source gateways.</p>
RIP Protocol Version	<p>Select which version of RIP is used on the interface.</p>
Interface Type	<p>Specify whether the interface sends RIP updates as multicast or broadcast messages by selecting the appropriate option. The default is multicast.</p> <p>You can set the interface type only if you select <i>Version 2</i> for the <i>RIP Protocol Version</i> option.</p> <p>Selecting broadcast specifies that RIP v1-compatible RIP v2 packets must be broadcast on the specified interface.</p>
Incoming Route Metric	<p>Enter a number from 1 to 16 to specify a metric the interface adds to incoming RIP routes before adding them to the route table.</p> <p>The default is the kernel interface metric plus 1 (the default RIP hop count). If this value is specified, it will be used as the absolute value and the kernel metric will not be used.</p> <p>Use this option to make the GSR prefer RIP routes learned from this interface less than RIP routes from other interfaces.</p>

Table 13. RIP parameters (Continued)

Parameter	Description
Outgoing Route Metric	<p>Enter the metric to be added to routes that are sent via this interface. The default is 0.</p> <p>This option is used to make other routers prefer other sources of RIP routes over the GSR.</p>

Note: Setting both the *RIP Packet Receipt* and *RIP Packet Transmission* options to *Disabled* will disable IP RIP on an interface. You may want to do so to temporarily disable RIP on an interface.

8. Click **OK**.
9. Repeat [step 5](#) through [step 8](#) until you configure all the IP RIP interfaces necessary for your network.

Configuration Expert adds the interfaces on which RIP is enabled to the list of interfaces found in the RIP Enabled IP Interfaces object. Deleting the interface from that object will disable RIP on the interface, but the interface will still be configured for IP and available as an IP interface.

Adding Trusted Gateways

A trusted gateway is a router that is trusted to supply routing information. By default, all routers on a subnet are trusted to supply routing information. If trusted gateways are specified, only updates from those gateways are accepted. Add trusted gateways if you want the GSR to accept RIP updates from only specific sources.

To add trusted gateways:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object and then double-click the RIP Routing object. Double-click the RIP Trusted Gateways object.
5. Click the Configure New RIP Trusted Gateway object.
6. In the RIP Trusted Gateway dialog box that appears, enter the IP address or host name of a source from which you want the GSR to accept RIP updates.
7. Click **OK**.

Configuration Expert adds the trusted gateway to the list of those found in the RIP Trusted Gateways object. The GSR will accept RIP updates only from the sources you include in this list.

8. Repeat [step 5](#) through [step 7](#) until you create the desired list of trusted gateways.

Adding Source Gateways

Add source gateways if you want to add routers that send RIP updates directly rather than using broadcast or multicast messages. Updates to source gateways are not affected by the setting of the interface RIP Packet Transmission option discussed in [“RIP Packet Transmission” on page 125](#). To add source gateways:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file’s Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object and then double-click the RIP Routing object. Double-click the RIP Source Gateways object.
5. Click the Configure New RIP Source Gateway object.
6. In the RIP Source Gateway dialog box that appears, enter the IP address or host name of a router that you want to send RIP updates rather than using broadcast or multicast messages.
7. Click **OK**.

Configuration Expert adds the source gateway to the list of those found in the RIP Source Gateways object.

8. Repeat [step 5](#) through [step 7](#) until you create the desired list of sourcegateways.

What to Do Next

As discussed in the following list, what you do after configuring IP interfaces for unicast routing depends on your network environment and whether you want to control traffic or set security.

- If you are planning to implement multicast routing on the GSR, configure IP interfaces for multicast routing as discussed in [Chapter 10: “Configuring Multicast Routing on the GSR” on page 129](#).
- If you are not using multicast routing but you will be using IPX, see [Chapter 11: “Configuring the GSR for IPX Routes” on page 141](#).

- If you do not need to configure the GSR for multicast routing or IPX, you can control traffic as discussed in [Chapter 12: “Configuring QoS on the GSR” on page 161](#). You can also set up security as discussed in [Chapter 13: “Configuring Security on the GSR” on page 189](#). Both tasks are optional.

Chapter 10

Configuring Multicast Routing on the GSR

Multicast routing on the GSR is supported through the Distance Vector Multicast Routing Protocol (DVMRP) and Internet Group Management Protocol (IGMP). DVMRP is used to determine forwarding of multicast traffic between GSRs. IGMP is used to determine host membership on directly attached subnets. This chapter

- provides an overview of the GSR's implementation of DVMRP, which is the routing protocol IP uses to perform multicast routing on the GSR.
- discusses configuring DVMRP routing on the GSR.
- provides an overview of the GSR's implementation of IGMP.
- discusses configuring IGMP on the GSR.

In addition to configuring the GSR for multicast routing, you can configure the GSR for unicast routing and IPX. For details on unicast routing, see [Chapter 9: "Configuring Unicast Routing on the GSR" on page 113](#). For details on IPX, see [Chapter 11: "Configuring the GSR for IPX Routes" on page 141](#).

Note: Before you begin configuring your GSR for multicast routing, you should create the IP interfaces you want to use for multicast routing. For details, see ["Creating IP Interfaces" on page 92](#).

What Is DVMRP?

DVMRP is an IP multicast routing protocol. On the GSR, DVMRP routing is implemented as specified in the draft-ietf-idmr-dvmrp-v3-06.txt file, which is an Internet Engineering

Task Force (IETF) document. The GSR's implementation of DVMRP supports the following:

- mtrace, which is a utility that tracks the multicast path from a source to a receiver.
- Generation identifiers, which are assigned to DVMRP whenever that protocol is started on a router.
- Pruning, which is an operation DVMRP routers perform to exclude interfaces not in the shortest path tree.

DVMRP uses the Reverse Path Multicasting (RPM) algorithm to perform pruning. In RPM, a source network rather than a host is paired with a multicast group. RPM permits the GSR to maintain multiple multicast groups.

On the GSR, DVMRP can be configured on a per interface basis. An interface does not have to run both DVMRP and IGMP. You can start and stop DVMRP independently from other routing protocols. IGMP starts and stops automatically with DVMRP.

To support backward compatibility on DVMRP interfaces, you can configure the router expire time and prune time on each GSR DVMRP interface. This lets it work with older versions of DVMRP.

You can use threshold values and scopes to control internetwork traffic on each DVMRP interface. Threshold values determine whether traffic is either restricted or not restricted to a subnet, site, or region. Scopes define a set of multicast addresses of devices to which the GSR can send DVMRP data. Scopes can include only addresses of devices on a company's internal network and cannot include addresses that require the GSR to send DVMRP data on the internet.

You can also configure tunnels on GSR DVMRP interfaces. A tunnel is used to send packets between routers separated by gateways that do not support multicast routing. A tunnel acts as a virtual network between two routers running DVMRP. A tunnel does not run IGMP.

Configuring DVMRP Routing on the GSR

You can configure DVMRP routing on the GSR by performing the following DVMRP-configuration tasks. You can create IP interfaces as discussed in [“Creating IP Interfaces” on page 92](#). The following DVMRP-configuration tasks are discussed in this chapter:

- Setting global parameters that will be used for all the interfaces on which DVMRP is enabled.
- Configuring DVMRP on individual interfaces. You do so by enabling and disabling DVMRP on interfaces and then setting DVMRP parameters on the interfaces on which DVMRP is disabled.
- Defining DVMRP tunnels, which IP uses to send multicast traffic between two end points.

- Enabling or disabling DVMRP on tunnels.

Setting DVMRP Global Parameters on the GSR

Set DVMRP global parameters to enable or disable multicast routing on the GSR and to specify whether or not the GSR performs pruning to exclude interfaces not in the shortest path tree.

To set the GSR's DVMRP global parameters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Multicast Routing object. Then double-click the DVMRP Routing object and click the DVMRP Global Parameters object.

A **DVMRP Global Parameters Configuration** dialog box similar to the following appears:

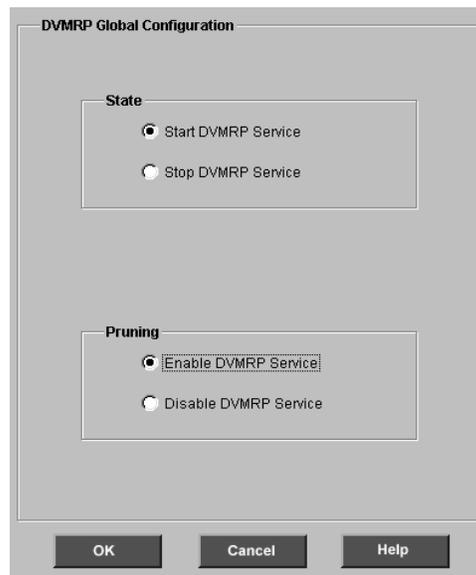


Figure 80. DVMRP Global Parameters Configuration dialog box

5. Specify whether you want to enable or disable DVMRP and pruning on the GSR by selecting the appropriate options.
6. Click **OK**.

Configuring DVMRP Interfaces

When configuring the GSR for DVMRP, you can enable or disable that protocol on IP interfaces. If you enable DVMRP on an interface, you can set DVMRP parameters on that interface. To configure DVMRP interfaces:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Multicast Routing object and then double-click the DVMRP Routing object. Double-click the DVMRP Interfaces object.
5. Do one of the following:
 - If you are modifying the settings of an interface on which DVMRP is enabled, double-click the DVMRP Enable Interfaces object. Then select the interface you want to change.
 - If you are modifying the settings of an interface on which DVMRP is disabled, double-click the DVMRP Disable Interfaces object. Then select the interface you want to change.
 - If you are configuring a new DVMRP interface, click the Configure New DVMRP Interface object.

An **Interface Bound to DVMRP** dialog box similar to the following appears:

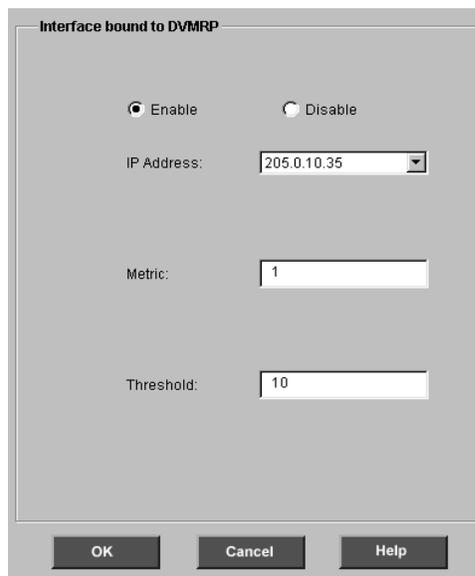


Figure 81. Interface Bound to DVMRP dialog box

6. If you are defining a new DVMRP interface, select the IP address or host name of the interface on which you are enabling or disabling DVMRP.

You will not be able to change the IP address if you are modifying an existing DVMRP interface.

7. If you enabled DVMRP on the interface, set DVMRP metric and threshold values on the interface as discussed in the following table:

Table 14. DVMRP metric and threshold values

Field	Description
Metric	Enter a number from 1 to 16 in the Metric box to specify the metric (cost) of the interface.
Threshold	Enter a number from 1 to 255 to specify the tunnel's multicast Time to Live (TTL) Threshold box. DVMRP uses the TTL to filter outbound traffic. Packets with a TTL value less than the one you specify are not sent any further.

8. Click **OK**.
9. Repeat [step 5](#) through [step 8](#) until you configure all the DVMRP interfaces necessary for your network environment.

Configuration Expert adds the interfaces on which DVMRP is enabled to the list of interfaces found in the DVMRP Enabled Interface object. Configuration Expert adds the interfaces on which DVMRP is disabled to the list of interfaces found in the DVMRP Disabled Interface object.

Defining DVMRP Tunnels

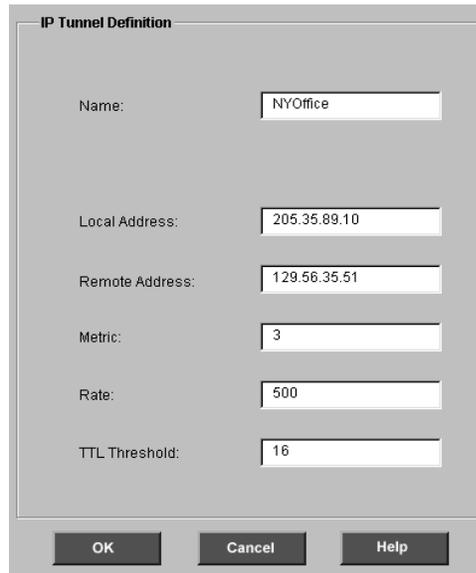
Configuration Expert lets you define DVMRP tunnels for sending multicast traffic between two end points. DVMRP treats a tunnel as another DVMRP interface. When configuring tunnels, you specify the IP addresses of the tunnel's two end points and then set DVMRP parameters as you do on interfaces.

To define a DVMRP tunnel:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Tunnel Configuration object. Then do one of the following:

- If you are modifying an existing tunnel, select that tunnel.
- If you are creating a new tunnel, select the *Configure New IP Tunnel* option.

An **IP Tunnel Definition** dialog box similar to the following appears:



The image shows a dialog box titled "IP Tunnel Definition". It contains several input fields for configuring a tunnel. The fields and their values are: Name: NYOffice; Local Address: 205.35.89.10; Remote Address: 129.56.35.51; Metric: 3; Rate: 500; TTL Threshold: 16. At the bottom of the dialog box are three buttons: OK, Cancel, and Help.

Field	Value
Name:	NYOffice
Local Address:	205.35.89.10
Remote Address:	129.56.35.51
Metric:	3
Rate:	500
TTL Threshold:	16

Figure 82. IP Tunnel Definition dialog box

5. Enter the tunnel's name.
6. Enter the IP addresses of the local and remote end points of the tunnel.

The local end point is the GSR. The remote end point is the system to which the GSR will be connected via the tunnel.

- Set the DVMRP parameters as discussed in the following table:

Table 15. DVMRP parameters

Field	Description
Metric	Enter a number from 1 to 16 in the Metric box to specify the metric (cost) of the interface.
Rate	Enter the speed at which you want outgoing traffic to pass through the interface. The field's value is specified in kilobits per second (Kbps).
TTL Threshold	Enter a number from 1 to 255 to specify the tunnel's multicast Time to Live (TTL) Threshold box. DVMRP uses the TTL to filter outbound traffic. Packets with a TTL value less than the one you specify are not sent any further.

- Click **OK**.

Configuration Expert adds the tunnel to the list of those found in the IP Tunnel Configuration object. Before the GSR can use the tunnel for multicast routing, you need to enable DVMRP on it as discussed in the next section.

Enabling or Disabling DVMRP on Tunnels

You can enable and disable DVMRP on the IP tunnels you previously created. To enable or disable DVMRP on tunnels:

- Start Configuration Expert if you have not already done so.
- Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
- Double-click the IP Routing Configuration object.
- Double-click the IP Multicast Routing object. Then double-click the DVMRP Routing object.
- Click the DVMRP Tunnels object. A **DVMRP Tunnel** dialog box similar to the following appears:



Figure 83. DVMRP Tunnel dialog box

6. Enable and disable DVMRP on tunnels.

To enable DVMRP on a tunnel, select that tunnel from the **Available Tunnels** list and click the **Add** button.

To disable DVMRP on a tunnel, select that tunnel from the **DVMRP Tunnels** list and click the **Remove** button.

7. Click **OK**.

The tunnels on which DVMRP is enabled are included in the list of tunnels found in the DVMRP Tunnels object. The tunnels on which DVMRP is disabled are included in the list of tunnels found in the IP Tunnel Configuration object.

What Is IGMP?

IGMP is a group-management protocol between hosts and routers. IP hosts use IGMP to report their host group memberships to multicast routers.

The GSR implements IGMP routing as specified in RFC 2236. Although the GSR runs IGMP v2, the GSR supports backward compatibility for hosts running IGMP v1.

An IGMP interface checks which multicast groups have members on the LAN to which the GSR is attached and to create routes for each multicast group.

On the GSR, IGMP can be configured on a per interface basis. You can configure a GSR interface to support IGMP only or both IGMP and DVMRP. If an interface is configured for both IGMP and DVMRP, IGMP starts and stops automatically with DVMRP.

On the GSR, IGMP keeps track of members on a per port basis even though an interface might contain multiple ports. Ports belonging to an interface without any IGMP memberships are not sent any multicast data traffic. This prevents ports from being sent unnecessary multicast traffic. This optimization, however, is removed if a network has multiple routers with multiple ports running a multicast protocol on a single interface. The optimization is removed to allow interaction between any host and any router within that interface. You can also remove this optimization on individual IGMP interfaces.

The GSR's implementation of IGMP allows control of the following:

- IGMP host-query intervals. Host queries are packets an IGMP router sends to hosts to learn which hosts are available. Host queries help a router determine changes to host membership.

A longer host query interval means less IGMP queries on the network. For hosts not sending explicit messages about when they leave a group, the host query interval helps the GSR determine when those hosts left that group.

- Response time, which helps in controlling the burstiness of responses from the multicasting hosts responding to IGMP queries from the GSR.

You can configure the host query interval and response time on the GSR because such configurations are supported in IGMP version 2. In IGMP version 1, however, the query interval is fixed at 125 seconds and response time is fixed at 10 seconds.

The GSR lets you control which groups are allowed or barred on a per-interface basis.

Configuring IGMP Interfaces on the GSR

You configure IGMP interfaces by

- Setting IGMP global parameters that will apply to all ports.
- Enabling or disabling IGMP on interfaces.

Setting IGMP Global Parameters

Configuration Expert lets you set some IGMP parameters that apply to all ports. These global parameters specify how often the GSR sends queries to learn which hosts are available and also how long the GSR waits for hosts to respond to such queries.

To set the global IGMP parameters that will apply to all ports:

1. Start Configuration Expert if you have not already done so.

2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Multicast Routing object. Then double-click the IGMP Protocol object and click the IGMP Global Parameters object.

An **IGMP Global Parameters Configuration** dialog box similar to the following appears:

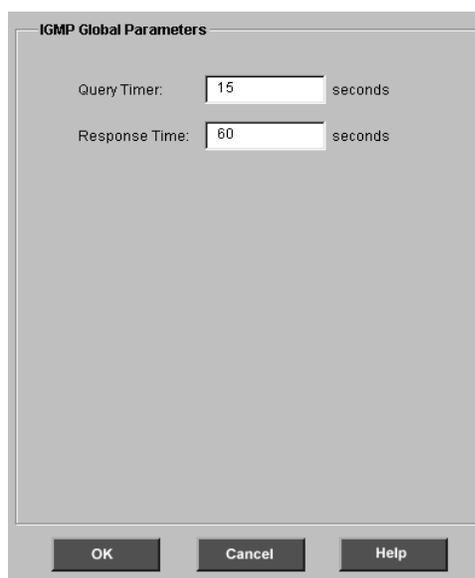


Figure 84. IGMP Global Parameters Configuration dialog box

5. In the **Query Timer** box, specify how often the GSR sends packets to learn which hosts are available.

You specify the interval in seconds and can enter values from 10 to 3600 in increments of 5 (that is, 5, 10, 15, and so on). The default is 125 seconds.

6. In the **Response Timer** box, specify how long the GSR waits for hosts to respond to the host-membership queries it sends.
7. Click **OK**.

The entries you make in the **IGMP Global Parameters** dialog box will apply to all ports.

Enabling or Disabling IGMP on Interfaces

You can enable or disable IGMP on the GSR's IP interfaces. To enable or disable IGMP interfaces:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Multicast Routing object and then double-click the IGMP Protocol object. Double-click the IGMP Interfaces object.
5. Do one of the following:
 - If you are modifying the settings of an interface on which IGMP is enabled, double-click the IGMP Enabled Interfaces object. Then select the interface you want to change.
 - If you are modifying the settings of an interface on which IGMP is disabled, double-click the IGMP Disabled Interfaces object. Then select the interface you want to change.
 - If you are configuring a new IGMP interface, click the Configure New IGMP Interface object.

An **IGMP Interface Definition** dialog box similar to the following appears:



Figure 85. IGMP Interface Definition dialog box

6. If you are defining a new IGMP interface, select the IP address or host name of the interface on which you are enabling or disabling IGMP.

You will not be able to change the IP address if you are modifying an existing IGMP interface.

7. Specify whether you want to enable or disable IGMP on the interface by selecting the appropriate options.
8. Click **OK**.
9. Repeat [step 5](#) through [step 8](#) until you configure all the IGMP interfaces necessary for your network environment.

Configuration Expert adds the interfaces on which IGMP is enabled to the list of interfaces found in the IGMP Enabled Interface object. Configuration Expert adds the interfaces on which IGMP is disabled to the list of interfaces found in the IGMP Disabled Interface object.

What to Do Next

As discussed in the following list, what you do after configuring IP interfaces for multicast routing depends on your network environment and whether you want to control traffic or set security.

- If your network environment includes IPX, see [Chapter 11: “Configuring the GSR for IPX Routes” on page 141](#).
- If you do not need to configure the GSR for IPX, you can control traffic as discussed in [Chapter 12: “Configuring QoS on the GSR” on page 161](#). You can also set up security as discussed in [Chapter 13: “Configuring Security on the GSR” on page 189](#). Both tasks are optional.

Chapter 11

Configuring the GSR for IPX Routes

The Internetwork Packet Exchange (IPX) is a datagram connectionless protocol for the Novell NetWare environment. You can configure the GSR for IPX routing and SAP. This chapter

- provides an overview of IPX on the GSR.
- discusses creating IPX interfaces, which are bound to either a port or VLAN.
- discusses configuring static IPX SAP entries.

SAP allows service-providing nodes (such as file servers and print servers) to advertise their services and addresses. You can configure static IPX SAP entries if you want to add an entry for an IPX server to the IPX SAP table.

What Is IPX?

IPX, which is a datagram connectionless protocol, performs various tasks including addressing, and routing and switching information packets from one location to another on a network.

IPX defines internetwork and intranode addressing schemes. IPX internetwork addressing is based on network numbers assigned to each network segment on a Novell NetWare internetwork. The IPX intranode address comes in the form of socket numbers. Because several processes are normally operating within a node, socket numbers provide a way for each process to distinguish itself.

The IPX packet consists of two parts: a 30-byte header and a data portion. The network node and socket addresses for both the destination and source are held within the IPX header.

The GSR uses IPX RIP to create and maintain a database of internetwork routing information. The GSR's implementation of RIP allows the following exchanges of information:

- Workstations locate the fastest route to a network number by broadcasting a route request.
- Routers request routing information from other routers to update their own internal tables by broadcasting a route request.
- Routers respond to route requests from workstations and other routers.
- Routers perform periodic broadcasts to make sure that all other routers are aware of the internetwork configuration.
- Routers perform broadcasting whenever they detect a change in the internetwork configurations.

GSR's RIP implementation follows the guidelines given in Novell's *IPX RIP and SAP Router Specification Version 1.30* document.

On the GSR, RIP automatically runs on all IPX interfaces. The GSR will keep multiple routes to the same network having the lowest ticks and hop count. Static routes can be configured on the GSR using the CLI's **ipx add route** command. Through the use of RIP filters, the GSR can control the acceptance and advertisement of networks per interface.

When creating or modifying interfaces, you should realize that interfaces bound to a single port go down when the port goes down but interfaces bound to a VLAN remain up as long as at least one port in that VLAN remains active. The following table summarizes this:

If	Then
The port of a port-bound interface goes down	The interface bound to that port also goes down
At least one port of a VLAN remains active	The interface bound to the VLAN remains active
All the ports of a VLAN go down	The interface bound to the VLAN also goes down

Each IPX interface can be bound to a MAC address as discussed in the *DIGITAL GIGAswitch/Router Command Line Interface Reference Manual*. Binding an interface to a MAC address is optional. If an interface is not bound to a MAC address, the interface uses the system address.

Creating IPX Interfaces

When you create IPX interfaces on the GSR, you provide information about the interface (such as its name, output MAC encapsulation, and IPX address). You also enable or disable the interface and bind the interface to a single port or VLAN. If you want to apply an existing ACL to an interface, Configuration Expert lets you do so either when you create the interface or afterwards.

When creating or modifying interfaces, you should realize that interfaces bound to a single port go down when the port goes down but interfaces bound to a VLAN remain up as long as at least one port in that VLAN remains active. The following table summarizes this:

If	Then
The port of a port-bound interface goes down	The interface bound to that port also goes down
At least one port of a VLAN remains active	The interface bound to the VLAN remains active
All the ports of a VLAN go down	The interface bound to the VLAN also goes down

You can bind each IPX interface you create to a MAC address as discussed in the *DIGITAL GIGAswitch/Router Command Line Interface Reference Manual*. Binding an interface to a MAC address is optional. If you do not bind an interface to a MAC address, the interface uses the system address.

The procedure for creating an IPX interface depends on whether you are binding that interface to a single port or a VLAN. Separate discussions on the different procedures follow.

Creating IPX Interfaces Bound to a Single Port

To create an IPX interface that you want bound to a single port:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IPX Routing Configuration object.
4. Double-click the IPX Interface Configuration object. Then click the Configure New IPX Interface object.

Configuration Expert opens the IPX Interface wizard.

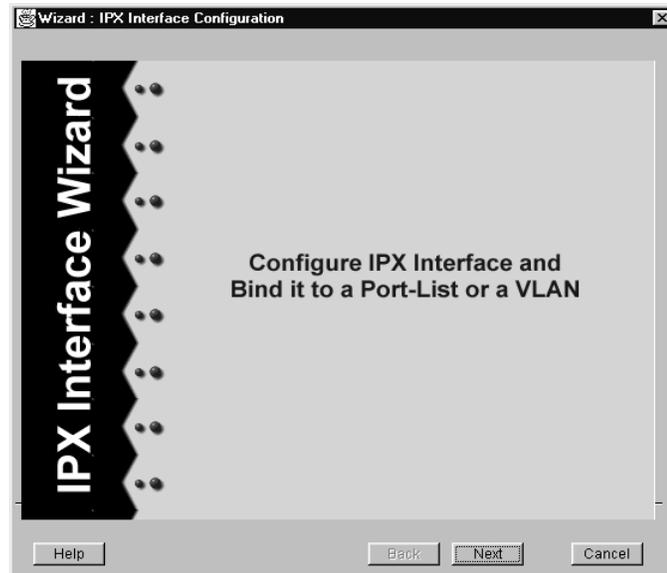


Figure 86. IPX Interface wizard (single port)

5. Click Next.

An IPX Interface Definition panel similar to the following appears:

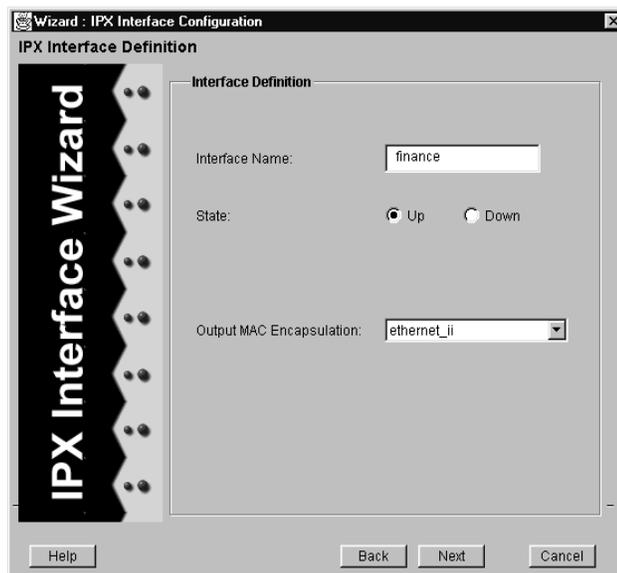


Figure 87. Interface Definition panel (single port)

6. Enter the name of the interface in the Interface **Name** box. Then either select *Up* to enable the interface or select *Down* to disable it.
7. Set the output MAC encapsulation you want associated with the interface by selecting one of the following from the Output MAC Encapsulation drop-down list:
 - *ethernet_II* (the default)
 - *ethernet_802_3*
 - *ethernet_snap*
 - *ethernet_802_2_ipx*
8. Click **Next** and then enter the IPX address of the interface in the panel that appears.

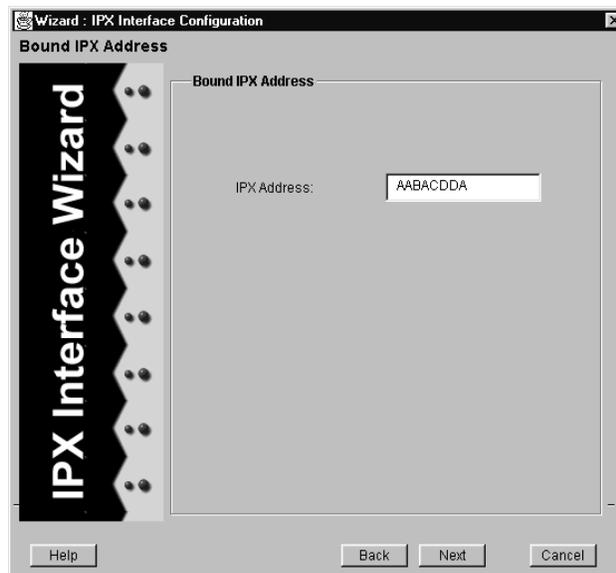


Figure 88. Bound IPX Address panel (single port)

9. Click **Next** and then select the *Bind the interface to Port* option in the panel that appears.

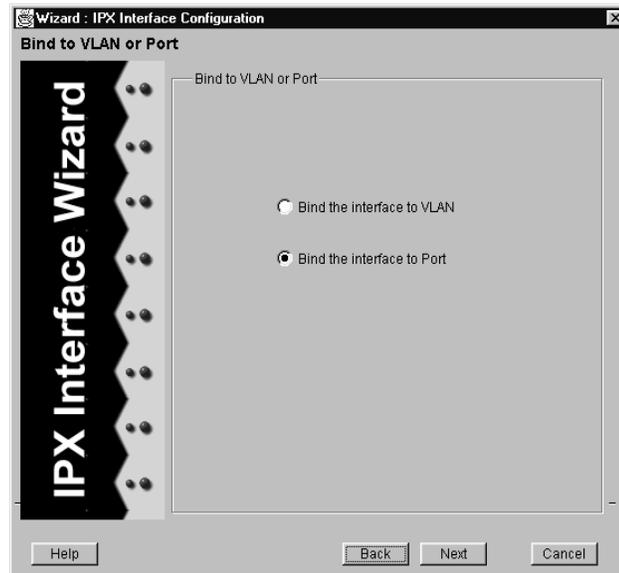


Figure 89. Bind to VLAN or Port panel (single port)

10. Click Next.

Configuration Expert displays a Bound Port List panel similar to the following:

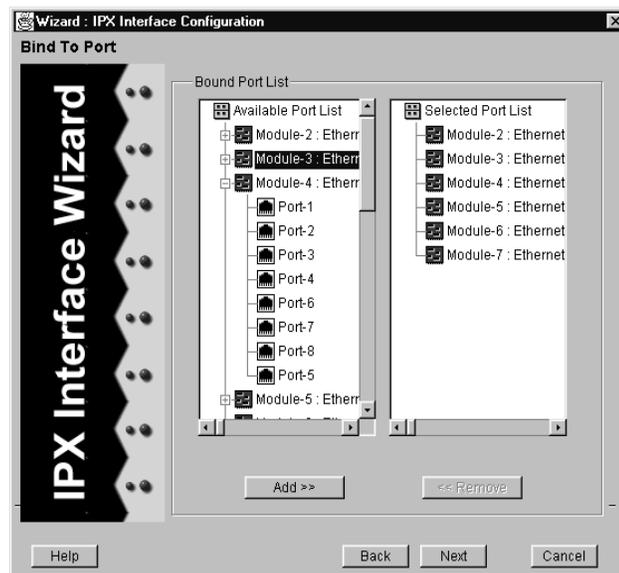


Figure 90. Bound Port List panel (single port)

11. Bind the interface to a single port by doing the following:

- a. In the **Available Port** list, double-click the module containing the desired port. From the list of available ports that appears, select the port that you want to bind to the interface.
- b. Click the **Add** button.

Note: You can bind only a single port to an interface. If you need to bind multiple ports to the interface, create a VLAN consisting of those ports and bind the interface to that VLAN.

12. Click **Next**.

The Apply ACLs panel appears.

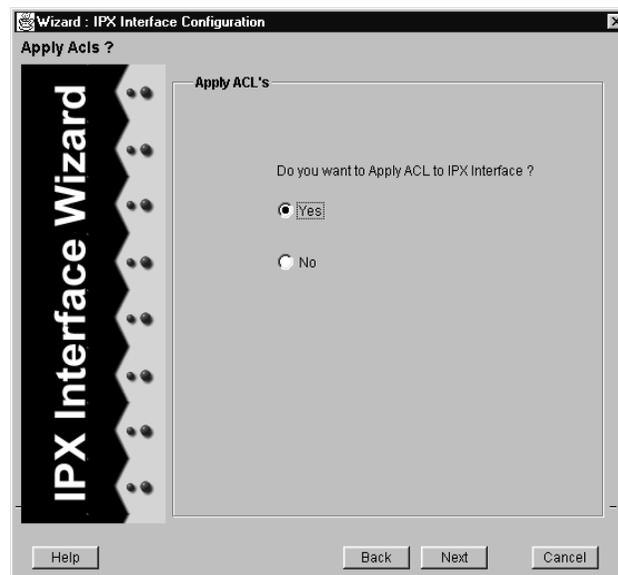


Figure 91. Apply ACLs panel (single port)

13. Specify whether you want to apply an ACL to the interface by doing one of the following.

- To not apply an ACL, select *No*. Then click **Finish**, which completes the configuration of the interface.

Configuration Expert adds the new interface to those found in the IPX interfaces bound to Ports object.

- To apply an ACL to the interface, select *Yes* and click **Next**.

Note: You will be able to apply an IPX ACL while creating an interface only if you previously created the ACL. If you plan to apply an ACL that you have not defined yet, finish creating the interface. Then configure the

desired ACL and apply it as discussed in [“Applying ACLs to IP or IPX Interfaces”](#) on page 206.

14. If you specified you wanted to apply an ACL, use the Apply IPX ACL panel that appears to apply an ACL to the interface.

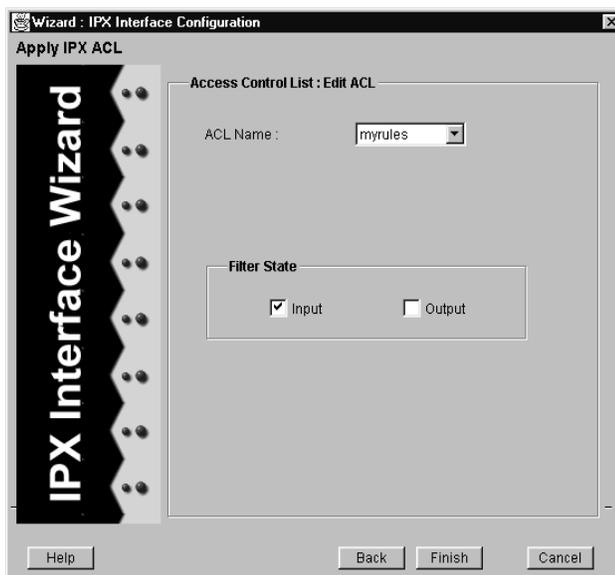


Figure 92. Access Control List: Edit ACL panel (single port)

To use the wizard to apply an IPX ACL, take the following steps:

- a. Select the ACL you want to apply from the **ACL Name** drop-down list. This will apply all ACLs with that name.

Configuration Expert lets you configure multiple ACLs that use different rules but have the same name. You can apply IPX, IPX RIP, and IPX SAP ACLs to an IPX interface.

- b. Use the *Filter State* check boxes to specify whether you want to filter out inbound traffic (*Input*), outbound traffic (*Output*), or both inbound and outbound traffic.

You can apply up to two of each of the different IPX ACLs (IPX, IPX RIP, and IPX SAP) to an IPX interface. When applying two ACLs of the same type, one ACL must govern the inbound traffic and the other must govern the outbound traffic. If you want to apply multiple ACLs to an IPX interface, finish using the wizard to apply an ACL, then apply the additional ACLs as discussed in [“Applying ACLs to IP or IPX Interfaces”](#) on page 206.

- c. Click **Finish**.

Configuration Expert adds the new interface to those found in the IPX interfaces bound to Ports object.

Note: When you apply an ACL to an interface, the GSR appends an implicit deny rule to that ACL. The implicit deny rule denies all traffic. If you intend to allow all traffic that does not match your specified ACL rules to go through, you must explicitly define a rule to permit all traffic. To do so, make sure the last rule of the ACL permits all traffic.

Creating IPX Interfaces Bound to a VLAN

If you have created an IPX VLAN, you can bind that VLAN to an IPX interface while creating the interface. To create an IPX interface that will be bound to an existing VLAN:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IPX Routing Configuration object.
4. Double-click the IPX Interface Configuration object. Then click the Configure New IPX Interface object.

Configuration Expert opens the IPX Interface wizard.

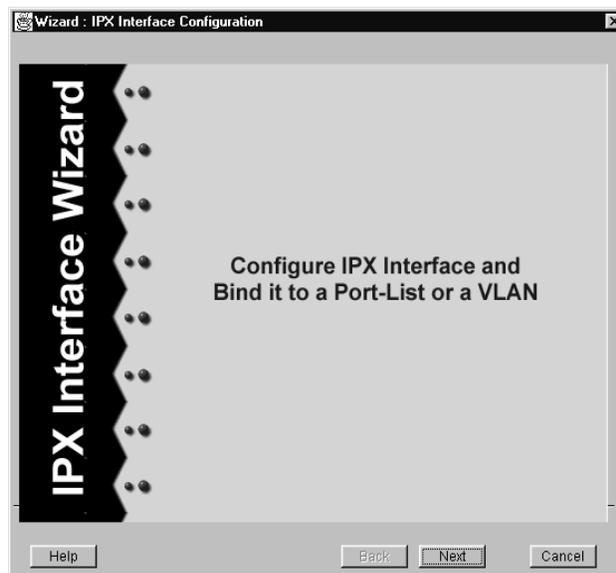


Figure 93. IPX Interface wizard (VLAN)

5. Click Next.

An IPX Interface Definition panel similar to the following appears:

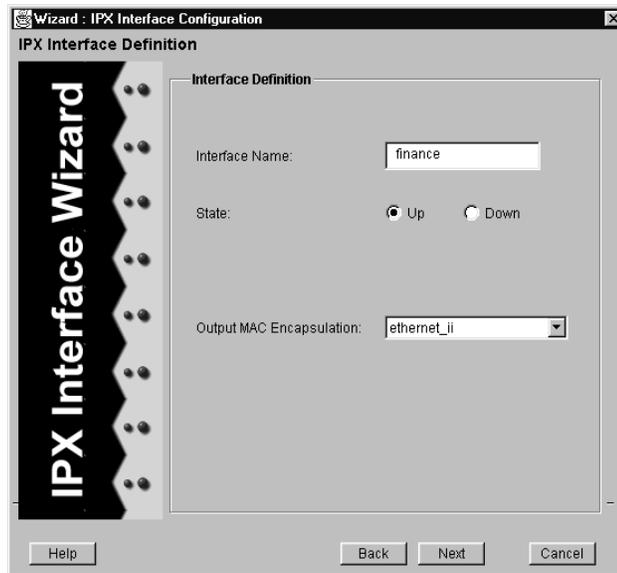


Figure 94. Interface Definition panel (VLAN)

6. Enter the name of the interface in the **Interface Name** box. Then either select *Up* to enable the interface or select *Down* to disable it.
7. Set the output MAC encapsulation you want associated with the interface by selecting one of the following from the **Output MAC Encapsulation** drop-down list:
 - *ethernet_II* (the default)
 - *ethernet_802_3*
 - *ethernet_snap*
 - *ethernet_802_2_ipx*
8. Click **Next** and then enter the IPX address of the interface in the panel that appears.

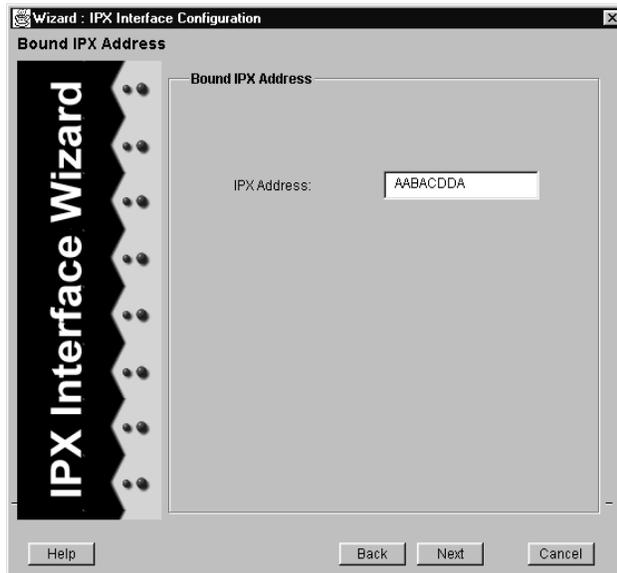


Figure 95. Bound IPX Address panel (VLAN)

9. Click **Next** and then select the *Bind the interface to VLAN* option in the panel that appears. This option is available only if there are existing IPX VLANs.

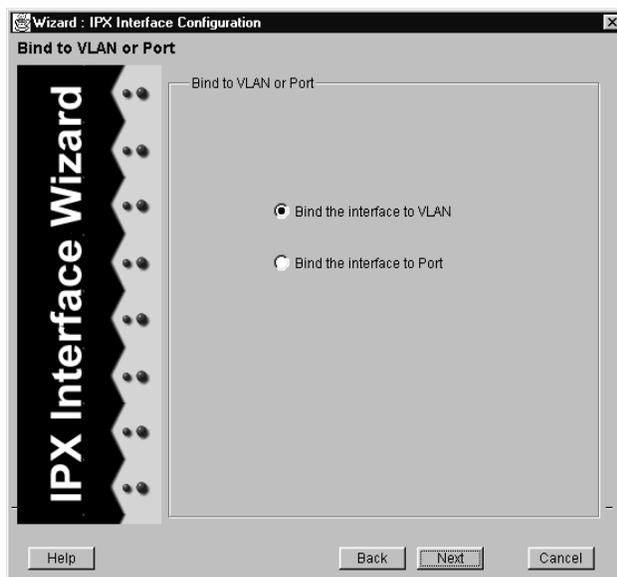


Figure 96. Bind to VLAN or Port panel (VLAN)

10. Click **Next**. In the panel that appears, select the name of the VLAN you want to bind to the interface.

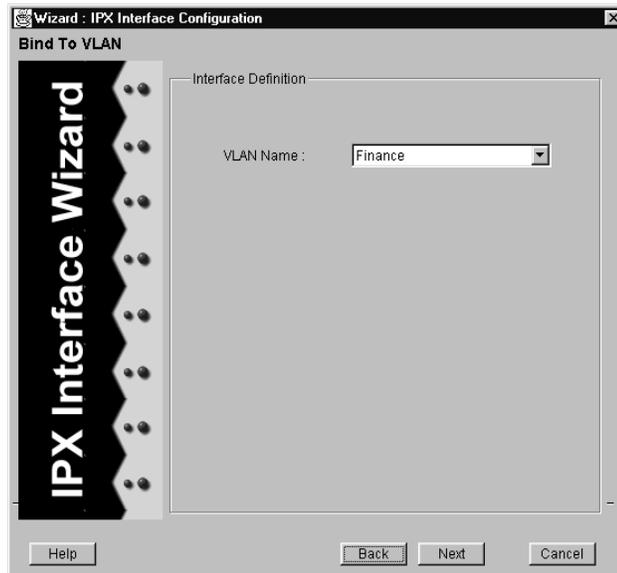


Figure 97. Interface Definition panel (VLAN)

11. Click Next.

The Apply ACLs panel appears.

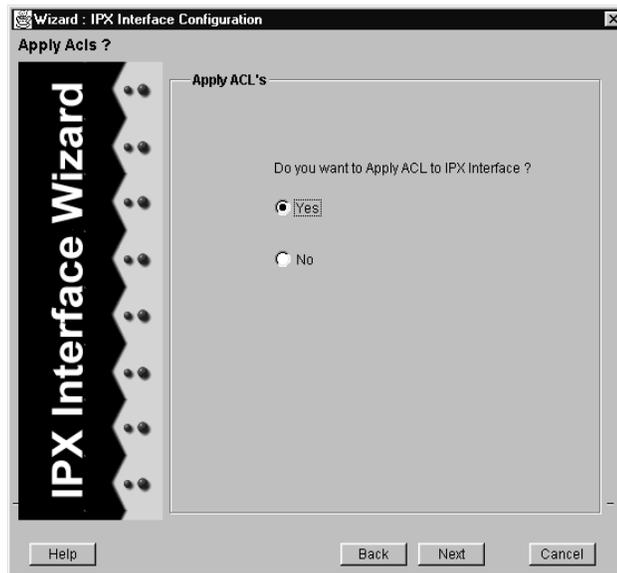


Figure 98. Apply ACLs panel

12. Specify whether you want to apply an ACL to the interface by doing one of the following.

- To not apply an ACL, select *No*. Then click **Finish**, which completes the configuration of the interface but does not apply any ACL to that interface.

Configuration Expert adds the new interface to those found in the IPX interfaces bound to VLAN object.

- To apply an ACL to the interface, select *Yes* and click **Next**.

Note: You will be able to apply an ACL while creating an IPX interface only if you previously created the ACL. If you plan to apply an ACL that you have not defined yet, finish creating the interface. Then configure the desired ACL and apply it as discussed in [“Applying ACLs to IP or IPX Interfaces”](#) on page 206.

13. If you specified you wanted to apply an ACL, use the Apply IPX ACL panel that appears to apply an ACL to the interface.

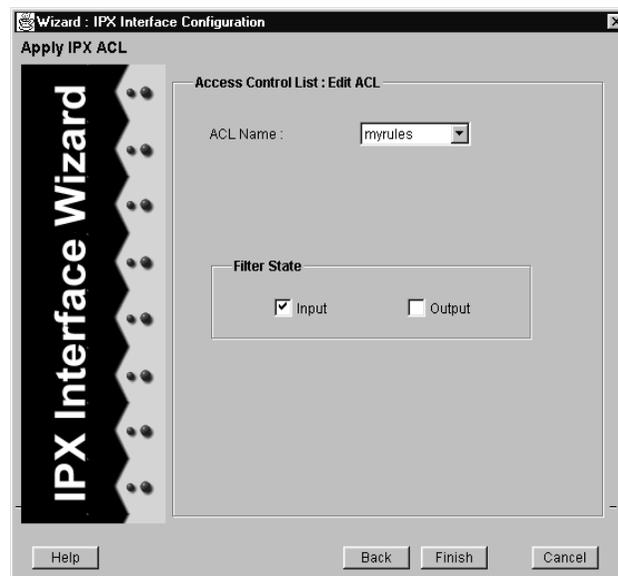


Figure 99. Access Control List: Edit ACL panel

To use the wizard to apply an IPX ACL, take the following steps:

- a. Select the ACL you want to apply from the **ACL Name** drop-down list. This will apply all ACLs with that name.

Configuration Expert lets you configure multiple ACLs that use different rules but have the same name. You can apply IPX, IPX RIP, and IPX SAP ACLs to an IPX interface.

- b. Use the *Filter State* check boxes to specify whether you want to filter out inbound

traffic, outbound traffic, or both inbound and outbound traffic.

Inbound traffic is packets coming into the interface while outbound traffic is packets going out of that interface.

Select *Input* to filter inbound traffic. Select *Output* to filter outbound traffic. Selecting both check boxes filters both inbound and outbound traffic.

You can apply up to two of each of the different IPX ACLs (IPX, IPX RIP, and IPX SAP) to an IPX interface. When applying two ACLs of the same type, one ACL must govern the inbound traffic and the other must govern the outbound traffic. If you want to apply multiple ACLs to an IPX interface, finish using the wizard to apply an ACL. Then apply the additional ACLs as discussed in [“Applying ACLs to IP or IPX Interfaces”](#) on page 206.

- c. Click **Finish**.

Configuration Expert adds the new interface to those found in the IPX interfaces bound to VLAN object.

Note: When you apply an ACL to an interface, the GSR appends an implicit deny rule to that ACL. The implicit deny rule denies all traffic. If you intend to allow all traffic that does not match your specified ACL rules to go through, you must explicitly define a rule to permit all traffic. To do so, make sure the last rule of the ACL permits all traffic.

Modifying IPX Interface Definitions

Modify IPX interface definitions to perform the following operations:

- Change the interface’s name
- Disable or enable the interface
- Change the interface’s MAC encapsulation
- Change the IPX address of the interface
- Bind a different port or VLAN to the interface
- Apply a different ACL to the interface
- Change which traffic an interface’s ACL filters out

To modify an IPX interface definition:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file’s Routing Configuration object.
3. Double-click the IPX Routing Configuration object.

4. Double-click the IPX Interface Configuration object. Then do one of the following:
 - If the interface you want to modify is bound to a port, double-click the IPX Interfaces bound to Ports object.
 - If the interface you want to modify is bound to a VLAN, double-click the IPX Interfaces bound to VLAN object.
5. In the list of interfaces that appears, double-click the one you want to modify.

Configuration Expert displays the contents of the object and the **Interface Definition** dialog box of the interface.

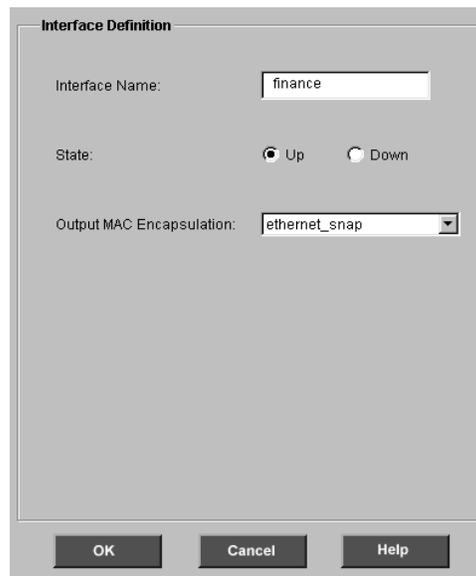


Figure 100. Interface Definition dialog box

6. If you want to edit the name, interface state, or MAC encapsulation fields, specify values as you do when creating an IPX interface. Then click **OK**.
7. If you want to change the interface's IPX address, double-click the Bound IPX Addresses object, click the interface's address object, then enter a new IPX address in the **Bound Port Address** dialog box that appears and click **OK**.
8. If you want to bind the interface to a different port or VLAN, do one of the following:
 - If you are modifying a port-bound interface, double-click the interface's Bound Port List object, use the **Bound Port List** dialog box to remove the currently bound port and to add the new port, and click **OK**.

You add and remove ports from the dialog box's list boxes as you do when creating the interface.

- If you are modifying a VLAN-bound interface, click the interface's VLAN object, select a new VLAN from the **VLAN Name** drop-down list that appears in the Interface Definition dialog box, and click **OK**.
9. If you want to change which ACLs are applied to the interface, double-click the Bound IPX Security object. Use the **Update ACL List** dialog box that appears to add and remove ACLs. You can apply two of each of the different IPX ACLs (IPX, IPX RIP, and IPX SAP) to an IPX interface.

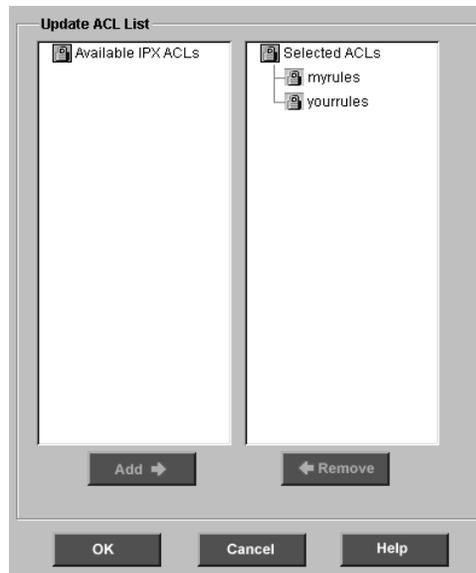


Figure 101. Update ACL List dialog box

You can add an ACL by selecting it in the **Available IPX ACLs** box and then clicking **Add**. You can remove an ACL by selecting it in the **Selected ACLs** box and then clicking **Remove**.

Note: You may also apply an IPX ACL by copying it as discussed in [“Copying an ACL to Apply It to an Interface” on page 207](#).

10. Click **OK**.
11. If you want to change the type of traffic an ACL filters out on the interface, click the IPX ACL located in the interface's Applied IPX ACL object.

The interface's **Edit ACL** dialog box appears.

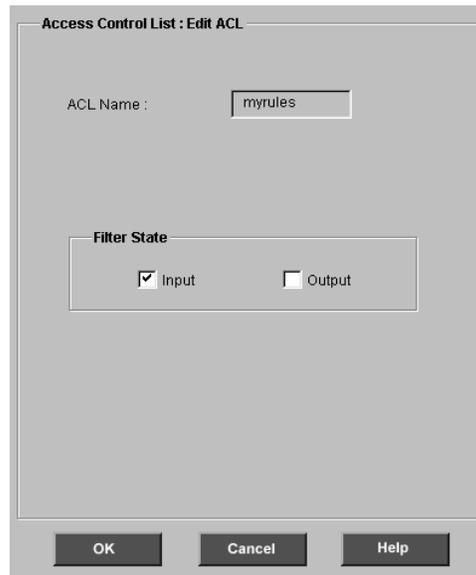


Figure 102. Edit ACL dialog box

12. Select or clear the *Filter State* check boxes to specify whether you want the ACL to filter inbound traffic (*Input*), outbound traffic (*Output*), or both input and outbound traffic. Click **OK**.



Caution: You can apply up to two of each of the different IPX ACLs (IPX, IPX RIP, and IPX SAP) to an IPX interface. When applying two ACLs of the same type, one ACL must govern the inbound traffic and the other must govern the outbound traffic.

Configuring Static IPX SAP Entries

A service access point (SAP) allows service-providing nodes (such as file servers and print servers) to advertise their services and addresses.

SAP makes the process of advertising and removing services dynamic. On the GSR, you can also add static IPX SAP entries. Through SAP, routers create and maintain a database of internetwork service information. This allows the clients on the network to determine which services are available on the network and to obtain the internetwork addresses of the nodes from which the clients can access those services. Because the SAP agent in each router keeps up-to-date information on available servers, a client wanting to locate the server can access a nearby router for the correct internetwork address.

The packet structure provided by SAP supports the following five functions:

- A workstation request for the name and address of the nearest server type

- A router request for the names and addresses of either all the servers or all the servers of a certain type on the internetwork
- A response to either a workstation or router request
- Periodic broadcast by servers and routers
- Changed server information broadcasts

You can configure a static IPX SAP entry if you want to add an entry for an IPX server to the IPX SAP table. The SAP agent broadcasts static IPX SAP entries only if an IPX interface has learned the same service dynamically.

Note: The IPX static SAP entries override dynamically learned entries, regardless of hop count. Moreover, when a dynamic route associated with the SAP entry is deleted or lost, the GSR does not advertise the IPX SAP entry until it relearns the route entry.

To configure static IPX SAP entries:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IPX Routing Configuration object.
4. Double-click the Static IPX SAP Configuration object. Then double-click the IPX SAP Entries object.
5. In the list of IPX SAP entries that appears, click one of the following:
 - An existing IPX SAP entry that you want to modify.
 - The Configure New SAP Entries object. Select this object if you want to add a new IPX SAP entry.

An **SAP Entry** dialog box similar to the following appears:

The image shows a dialog box titled "SAP Entry". It contains several input fields and a dropdown menu. The "Server Name" field is set to "ENG1". The "Service Type" dropdown is set to "Job Server (5)". The "Server Address" section is enclosed in a sub-dialog and contains three fields: "Network" set to "a1b2c3d4", "Node" set to "65:43:aa:23:11:67", and "Socket" set to "650". The "Hops" field is set to "2". At the bottom, there are three buttons: "OK", "Cancel", and "Help".

Figure 103. SAP Entry dialog box

- Configure the IPX SAP entry as discussed in the following table:

Table 16. IPX SAP fields

Field	Description
Server Name	Name of the IPX server. You can use any characters in the name except the following: " * . / : ; < = > ? [] \].
Service Type	The type of service.
Network	The IPX network address. Specify the address in the following format: a1b2c3d4f.
Node	The IPX node address. Specify the address in the following format: aa:bb:cc:dd:ee:ff.
Socket	The socket number for the IPX SAP entry.
Hops	The number of hops to the server. You can specify a number from 1 to 14.

- Click **OK**.

Configuration Expert adds the entry to those found in the IPX SAP Entries object.

What to Do Next

After configuring the GSR for IPX, you may perform the following tasks. Both tasks are optional.

- Control traffic as discussed in [Chapter 12: “Configuring QoS on the GSR” on page 161](#).
- Set up security as discussed in [Chapter 13: “Configuring Security on the GSR” on page 189](#).

Chapter 12

Configuring QoS on the GSR

After you define interfaces on the GSR, you can configure QoS policies to control traffic. This chapter

- provides an overview of QoS.
- lists the order in which you perform the various QoS-configuration tasks.
- discusses establishing the GSR's queuing policy.
- explains how to associate precedences with IP and IPX flows.
- discusses configuring QoS policies to define flows.
- provides details about modifying existing QoS profiles.

What Is QoS?

On the GSR, QoS is a set of parameters that let you do the following:

- Establish a queuing policy on the GSR to prioritize the GSR's traffic.
- Associate a precedence with Layer-3/Layer-4 flows. You do so by assigning a precedence value to each field in those type of flows. The GSR uses the precedence value assigned to the fields to break ties if packets match more than one flow.
- Classify Layer-2, Layer-3, and Layer-4 network traffic into one of the following priorities:
 - control
 - high

- medium
- low

Setting priorities for network traffic helps ensure that critical traffic will reach its destination even if the exit ports for the traffic are experiencing greater than maximum utilization.

Note: Control priority is reserved for system control traffic. Assign that priority only when necessary. Otherwise, you may hamper system operations if you assign the control priority to non-system traffic (such as HTTP, FTP, and so on).

- Define flows for Layer 2 and Layer-3/Layer-4 (IP and IPX) traffic.

Flows act as blueprints or templates for Layer-2, IP, and IPX packets. To define a flow, you specify values for the fields of Layer-2, IP, and IPX packets. The GSR then uses those field values whenever it processes such packets. If you do not specify a value for a field, the GSR assumes a wildcard value and so accepts any value for that field.

Establishing the GSR's Queuing Policy

You can establish one of the following queuing policies on the GSR to set the priority of the GSR's traffic:

- Strict priority

This policy assures the higher priorities of throughput but at the expense of lower priorities. For example, during heavy loads, low-priority traffic can be dropped to preserve throughput of control-priority traffic, and so on.

- Weighted-fair queuing

This policy distributes priority throughput among the four priorities (control, high, medium, and low) based on percentages.

The GSR can use only one queuing policy at a time. The policy you establish is used on the entire GSR. The default queuing policy is strict priority.

To establish queuing policies:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Double-click the Global Settings object.
4. Click the Queuing Discipline Configuration object.

A **Queuing Discipline Configuration** dialog box similar to the following appears:

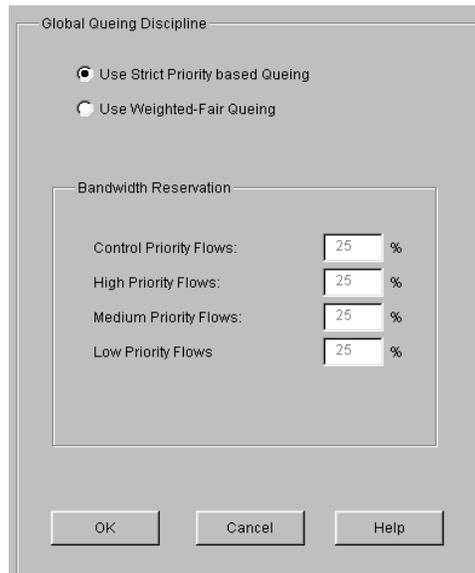


Figure 104. Queuing Discipline Configuration dialog box

5. Specify whether you want to use strict priority or weighted-fair queuing by selecting the appropriate option.
6. If you selected the *Use Weighted-Fair Queuing* option, set the amount of bandwidth you want allocated to each of the control, high, medium, and low priorities. If you selected *Use Strict Priority Based Queuing*, skip to [step 7](#).

The percentages you enter apply to all ports. Make sure the total percentages for all four priorities equals 100. You cannot set a control priority to 0%.

7. Click **OK**.

Associating Precedences to Layer-3/Layer-4 Flows

You associate precedences to Layer-3/Layer-4 flows so that the GSR will be able to break ties if packets match more than one flow. You associate a precedence with a flow by assigning a precedence value to each field of the flows.

The fields to which you assign precedence values vary for IP and IPX. Separate discussions on assigning QoS precedences for IP and IPX follow.

Assigning IP QoS Precedence

You can set the QoS precedence for the following flow fields in IP traffic:

- Destination TCP or UDP port
- Destination IP address
- Source TCP or UDP port
- Source IP address
- Type of Service (TOS) for the packet
- Incoming interface
- Protocol (TCP or UDP)

To assign the IP QoS precedence:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Double-click the Global Settings object.
4. Double-click the QoS Precedence object.
5. Click the IP QoS Precedence object.

An **IP QoS Precedence** dialog box similar to the following appears:

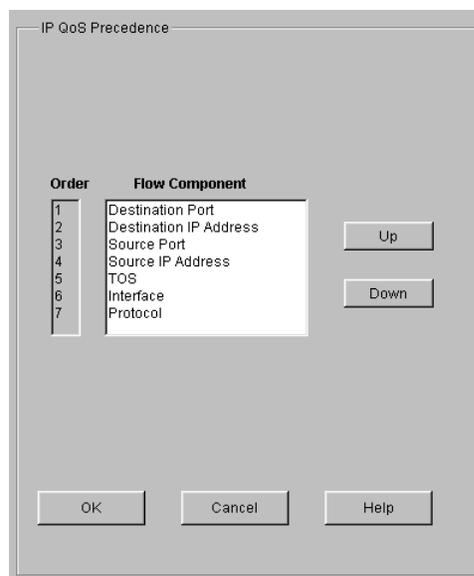


Figure 105. IP QoS Precedence dialog box

6. Assign one of the precedence values shown in the **Order** list to each of the IP-flow fields shown in the **Flow Components** list. To assign a precedence value to a field,

select that field and then click the **Up** and **Down** buttons until the field is listed next to the desired precedence value.

You may assign a precedence value from 1 to 7, where 1 is the highest precedence value. The default precedences of the fields are Destination Port (1), Destination IP Address (2), Source Port (3), Source IP Address (4), TOS (5), Interface (6), Protocol (7).

7. Click **OK**.

Assigning IPX QoS Precedence

You can set the QoS precedence for the following flow fields in IPX traffic:

- Destination network
- Source network
- Destination node
- Source node
- Destination port
- Source port
- Incoming interface

To assign the IPX QoS precedence:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Double-click the Global Settings object.
4. Double-click the QoS Precedence object.
5. Click the IPX QoS Precedence object.

An **IPX QoS Precedence** dialog box similar to the following appears:

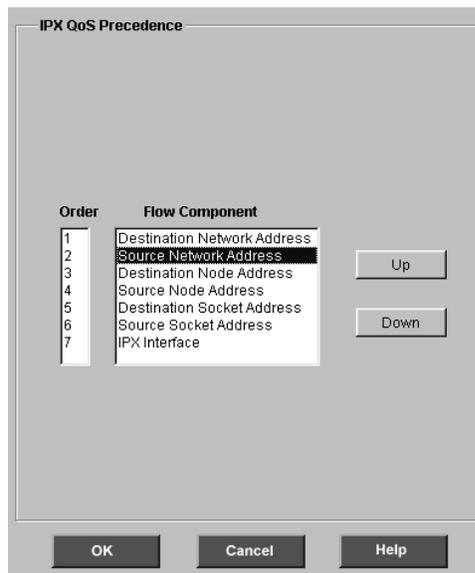


Figure 106. IPX QoS Precedence dialog box

- Assign one of the precedence values shown in the **Order** list to each of the IP-flow fields shown in the **Flow Components** list. To assign a precedence value to a field, select that field and then click the **Up** and **Down** buttons until the field is listed next to the desired precedence value.

You may assign precedence values from 1 to 7, where 1 is the highest precedence value. The default precedence values of the fields are Destination Network Address (1), Source Network Address (2), Destination Node Address (3), Source Node Address (4), Destination Socket Address (5), Source Socket Address (6), IPX Interface (7).

- Click **OK**.

Creating QoS Profiles

On the GSR, you can have multiple IP, IPX, and Layer-2 flows. You create a QoS profile for each of the flows you want to define on the GSR. When creating a QoS profile in Configuration Expert, you do so by performing the following tasks:

- Specifying the name of the flow.
- Setting a priority for flows based on the fields of the flow. Setting priorities helps ensure that critical traffic will reach its destination even if the exit ports for the traffic are experiencing greater than maximum utilization.
- Setting the values of packet fields. The GSR will use these values as blueprints or templates for the packets relevant to the flow type (IP, IPX, or Layer-2).

- Specifying to what you want the flow to apply. You apply IP and IPX flows to one or more interfaces. You apply Layer-2 flows to one or more ports.

The method you use to create a QoS profile depends on whether you are defining an IP, IPX, or Layer-2 flow. Separate discussions on creating QoS policies for the different flow types follow.

Creating a QoS Profile for an IP Flow

To create a QoS profile for an IP flow:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Click the Configure New QoS Profile object.

Configuration Expert opens the QoS wizard.



Figure 107. QoS wizard (IP flow)

4. Click **Next**.

Configuration Expert prompts you to specify which type of flow you want to define. The default is IP Flow.

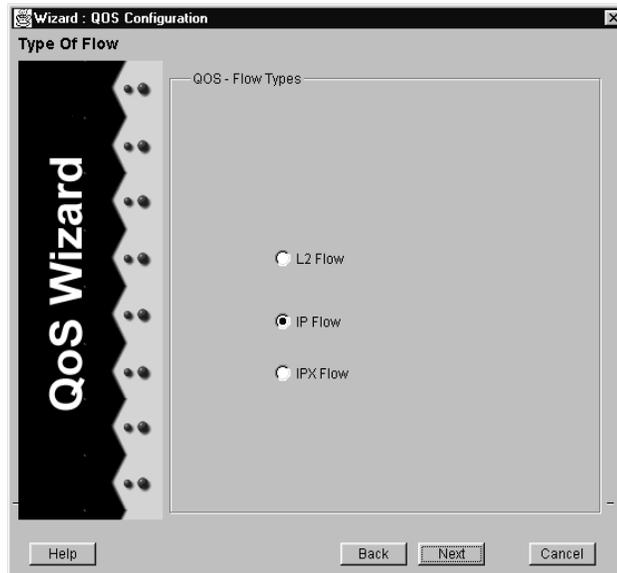


Figure 108. QoS - Flow Types panel (IP flow)

5. Click Next.

A QoS L3/L4 Flow Priority panel similar to the following appears:

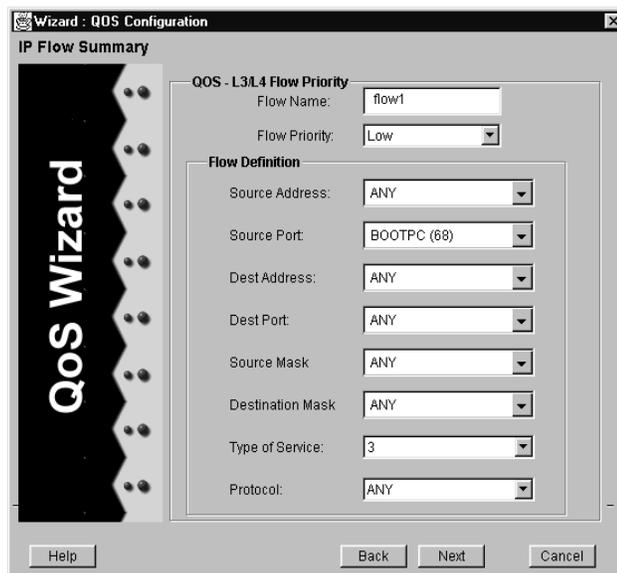


Figure 109. QoS - L3/L4 Flow Priority panel (IP flow)

6. Enter the flow's name in the **Flow Name** box.

7. From the **Flow Priority** drop-down list, select the priority you want to assign to the fields listed in the Flow Definition section of the panel.

The following table describes the priorities you may set:

Table 17. L3/L4 flow properties

Priority	Description
control	Assigns control priority to the IP flow fields you specify. This is the highest priority. Note: Control priority is reserved for system control traffic. Assign that priority only when necessary. Otherwise, you may hamper system operations if you assign the control priority to non-system traffic (such as HTTP, FTP, and so on).
high	Assigns high priority to the IP flow fields you specify.
medium	Assigns medium priority to the IP flow fields you specify.
low	Assigns low priority to the IP flow fields you specify. This is the default.

8. Define each flow field as discussed in the following table:

Table 18. L3/L4 flow fields

Field	Description
Source Address	Enter the source IP address for which you are assigning a priority. You can specify ANY to allow any value.
Source Port	Enter the source TCP or UDP port for which you are assigning a priority. Specify a port number from 1 to 65535. You can enter ANY to allow any value.
Destination Address	Enter the destination IP address for which you are assigning a priority. You can enter ANY to allow any value.
Destination Port	Enter the destination TCP or UDP port for which you are assigning a priority. Specify a port number from 1 to 65535, or enter ANY to allow any value.

Table 18. L3/L4 flow fields (Continued)

Field	Description
Source Mask	<p>Enter the network mask for which you are assigning a priority. You can specify the mask using the traditional IP address format ("255.255.0.0"). If you want to use the Classless Inter-Domain Routing (CIDR) format ("/16"), you may do so as discussed in the <i>DIGITAL GIGAswitch/Router Command Line Interface Reference Manual</i>.</p> <p>If you specify ANY instead of a network mask, the GSR assumes a wildcard "don't care" condition. If you do not specify a mask, then the GSR assumes a mask of 255.255.255.255. You cannot substitute the mask by entering ANY. The option ANY is for the entire Source Address and Source Mask pair.</p>
Destination Mask	<p>Enter the destination network mask for which you are assigning a priority. The same requirements and restrictions for Source Mask apply to Destination Mask.</p> <p>If you specify ANY instead of a network mask, the GSR assumes a wildcard "don't care" condition. If you do not specify a mask, then the GSR assumes a mask of 255.255.255.255. You cannot substitute the mask by entering ANY. The option ANY is for the entire Destination Address and Destination Mask pair.</p>
Type of Service	<p>Enter the TOS for which you are assigning a priority. Specify a number from 0 to 15, or enter ANY to allow any value.</p>
Protocol	<p>Enter the transport layer protocol (TCP or UDP) for which you are assigning priority. If you specify ANY, priority is assigned for both TCP and UDP.</p>

9. Click **Next**. In the QoS Flow Types panel that appears, specify whether you want to apply the flow to specific IP interfaces. Selecting **No** applies the flow to all IP interfaces. Selecting **Yes** allows you to apply the flow to specific IP interfaces.

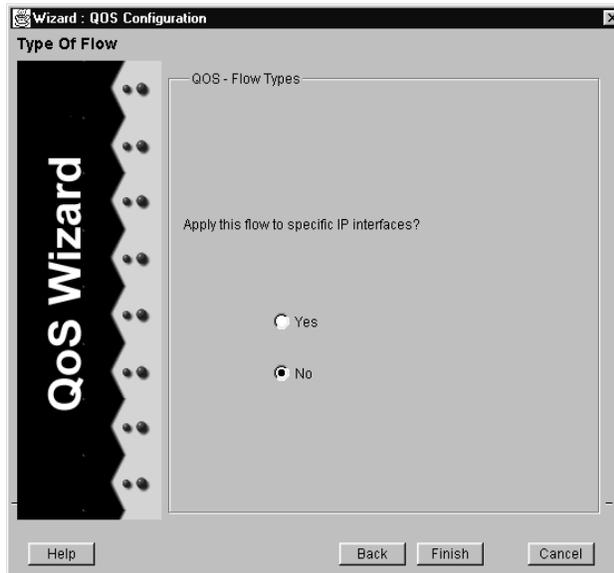


Figure 110. Apply to IP Interfaces panel (IP flow)

10. Do one of the following:

- If you selected *No*, click **Finish**.

Configuration Expert adds the QoS Profile to those included in the IP QoS Profiles object.

- If you selected *Yes*, click **Next**. In the Policy Input Interface List panel that appears, specify which interfaces you want to apply the flow to by selecting the desired interfaces from those included in the Available Interfaces list box and clicking the **Add** button. Then click the **Finish** button.

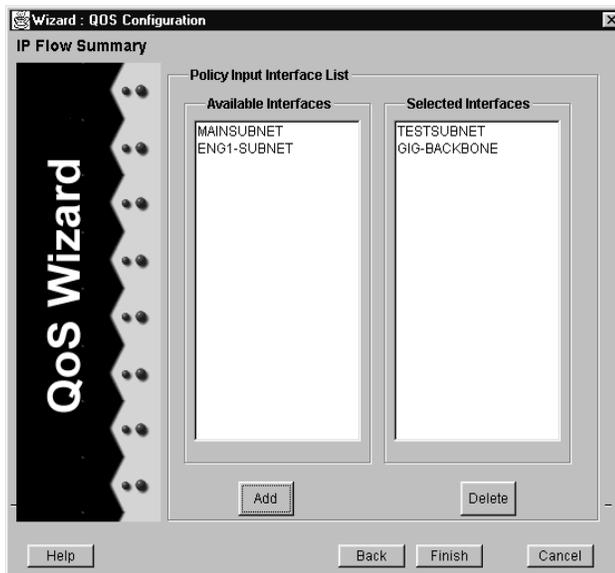


Figure 111. Policy Input Interface List panel (IP flow)

If you accidentally add a wrong interface, remove it by selecting it in the **Selected Interfaces** list and clicking the **Delete** button.

After you click **Finish**, Configuration Expert adds the QoS Profile to those included in the IP QoS Profiles object.

Creating a QoS Profile for an IPX Flow

To create a QoS profile for an IPX flow:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Click the Configure New QoS Profile object.

Configuration Expert opens the QoS wizard.



Figure 112. QoS wizard (IPX flow)

4. Click Next.

Configuration Expert prompts you to specify which type of flow you want to define.

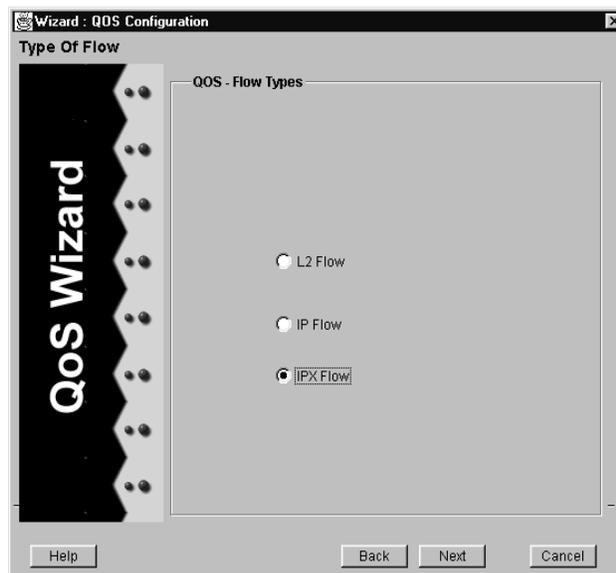


Figure 113. Qos - Flow Types panel (IPX flow)

5. Select *IPX Flow* and click Next.

An IPX Policy Definition panel similar to the following appears:

Figure 114. IPX Policy Definition panel (IPX flow)

6. Specify the flow's name in the **Name** box.
7. From the **Flow Priority** drop-down list, select the priority you want to assign to the fields listed in the Flow Definition section of the panel. The following table describes the priorities you may set:

Table 19. IPX flow priorities

Priority	Description
control	Assigns control priority to the IPX flow fields you specify. This is the highest priority. Note: Control priority is reserved for system control traffic. Assign that priority only when necessary. Otherwise, you may hamper system operations if you assign the control priority to non-system traffic (such as HTTP, FTP, and so on).
high	Assigns high priority to the IPX flow fields you specify.
medium	Assigns medium priority to the IPX flow fields you specify.
low	Assigns low priority to the IPX flow fields you specify. This is the default.

8. Define each flow field as discussed in the following table. You can enter ANY in a flow field to specify a wildcard (“don’t care”) condition.

Table 20. IPX flow fields

Field	Description
Source Network	Enter the IPX source network address. Specify it in the following format: a1b2c3d4.
Source Net Mask	Enter the IPX source network mask. Specify the mask in hexadecimal digits. If you do not specify a mask value and instead use the value ANY, the GSR internally sets the mask to FFFFFFFF.
Source MAC Addr	Enter the source node address. Specify it in the following format: aa:bb:cc:dd:ee:ff. If you specify ANY, the GSR assumes a wildcard value. All MAC addresses are then valid.
Source Port	Enter a port number from 1 to 65535, or enter ANY to allow any value.
Destination Network	Enter the IPX destination network and node address. Specify it in the following format: a1b2c3d4.
Destination Net Mask	Enter the IPX destination network mask. Specify the mask in hexadecimal digits or ANY to allow any value.
Destination MAC Addr	Enter the destination node address. The same requirements and restrictions for Source MAC Addr apply to Destination MAC Addr.
Destination Port	Enter a port number from 1 to 65535, or enter ANY to allow any value.

9. Click **Next**. In the panel that appears, specify whether you want to apply the flow to specific IPX interfaces. Selecting *No* applies the flow to all IPX interfaces. Selecting *Yes* allows you to apply the flow to specific IPX interfaces.

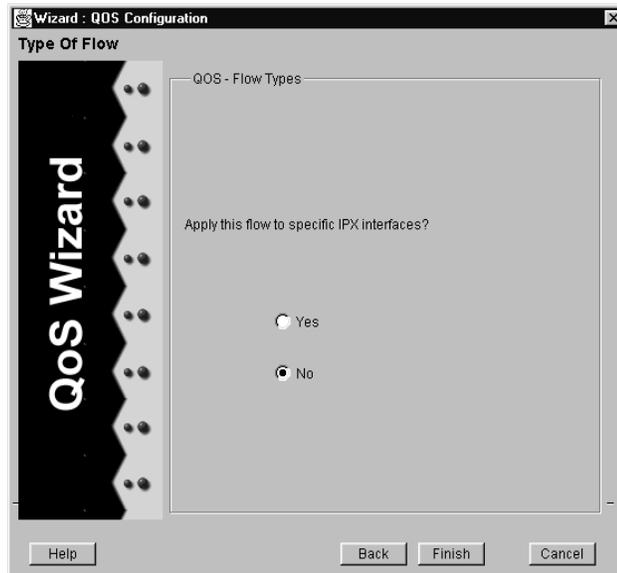


Figure 115. Apply to IPX Interfaces panel (IPX flow)

10. Do one of the following:

- If you selected *No*, click **Finish**.

Configuration Expert adds the QoS Profile to those included in the IPX QoS Profiles object.

- If you selected *Yes*, click **Next**. In the panel that appears, specify which interfaces you want to apply the flow to by selecting the desired interfaces from those included in the **Available Interfaces** list and clicking the **Add** button. Then click the **Finish** button.

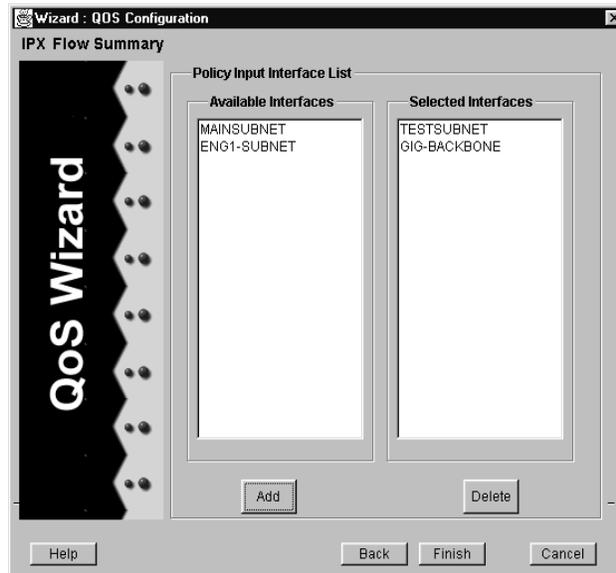


Figure 116. Policy Input Interface List panel (IPX flow)

If you accidentally add a wrong interface, remove it by selecting it in the **Selected Interfaces** list and clicking the **Delete** button.

After you click **Finish**, Configuration Expert adds the QoS Profile to those included in the IPX QoS Profiles object.

Creating a QoS Profile for a Layer-2 Flow

To create a QoS profile for a Layer-2 flow:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Click the Configure New QoS Profile object.

Configuration Expert opens the QoS wizard.



Figure 117. QoS wizard (Layer-2 flow)

4. Click Next.

Configuration Expert prompts you to specify which type of flow you want to define.

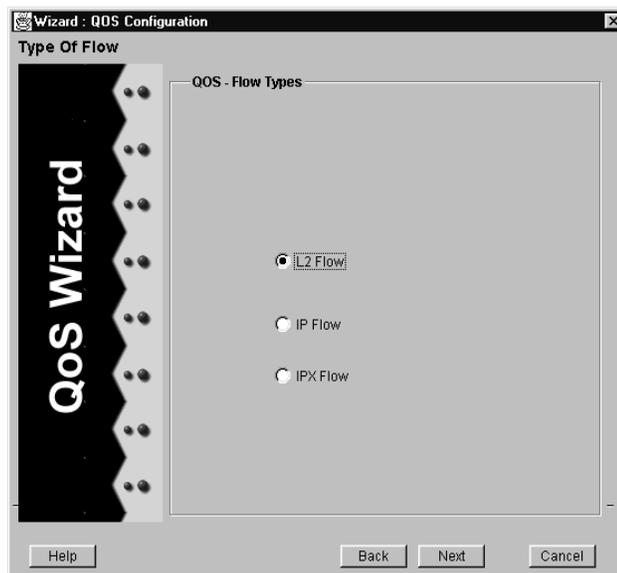


Figure 118. QoS - Flow Types panel (Layer-2 flow)

5. Select *L2 Flow* and click Next.

An L2 Flow Priority Definition panel similar to the following appears:

Figure 119. L2 Flow Priority Definition panel (Layer-2 flow)

6. Specify the flow's name in the **Name** box.
7. From the **Flow Priority** drop-down list, select the priority you want to assign to the fields listed in the Flow Definition section of the panel. The following table describes the priorities you may set:

Table 21. L2 flow priorities

Priority	Description
control	Assigns control priority to the Layer-2 flow fields you specify. This is the highest priority. Note: Control priority is reserved for system control traffic. Assign that priority only when necessary. Otherwise, you may hamper system operations if you assign the control priority to non-system traffic (such as HTTP, FTP, and so on).
high	Assigns high priority to the Layer-2 flow fields you specify.
medium	Assigns medium priority to the Layer-2 flow fields you specify.
low	Assigns low priority to the Layer-2 flow fields you specify. This is the default.

8. Define each flow field as discussed in the following table:

Table 22. L2 flow fields

Field	Description
Source MAC Addr	Enter the Layer-2 source MAC address. Specify the MAC address in the xx : xx : xx : xx : xx : xx format.
Destination MAC Addr	Enter the Layer-2 destination MAC address.
VLAN ID	Enter the ID of a VLAN.

9. Click **Next**. In the panel that appears, specify whether you want to apply the flow to specific ports. Selecting *No* applies the flow to all ports. Selecting *Yes* allows you to apply the flow to specific ports.

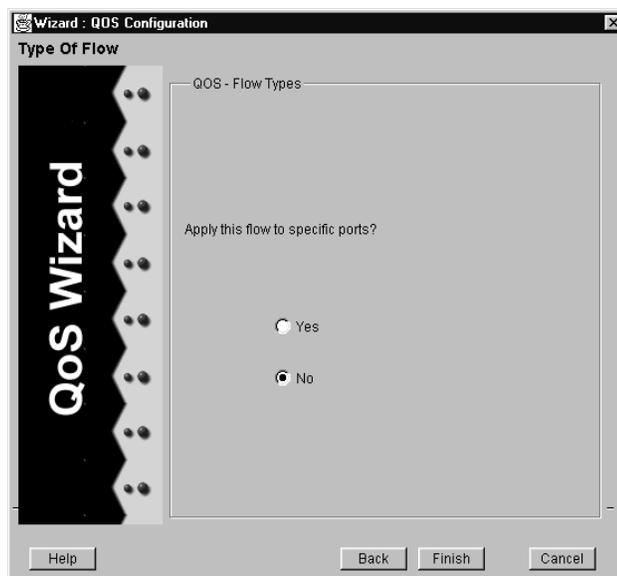


Figure 120. QoS - Flow Types panel (Layer-2 flow)

10. Do one of the following:

- If you selected *No*, click **Finish**.

Configuration Expert adds the QoS Profile to those included in the Layer-2 QoS Profiles object.

- If you selected *Yes*, click **Next**. In the Update Port List panel that appears, specify to which ports you want to apply the flow. Then click **Finish**.

To specify a port, double-click that port's module in the **Update Port** list, select the port from the port list that appears, and click the **Add** button.

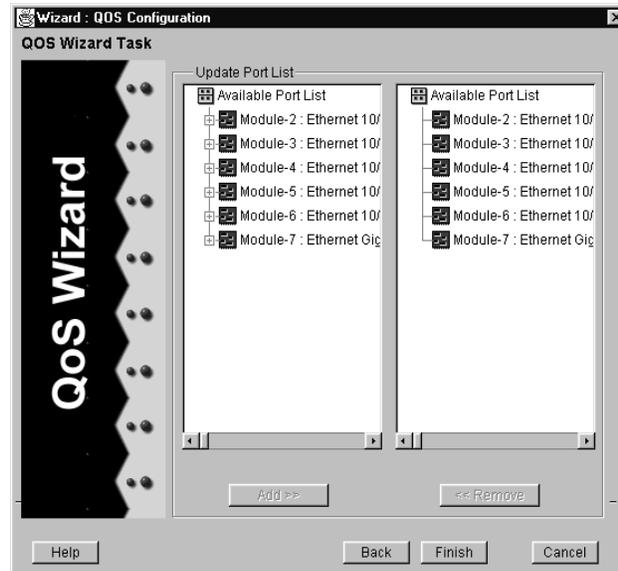


Figure 121. Update Port List panel (Layer-2 flow)

Clicking a module rather than double-clicking it selects all of that module's ports. This is a quick way to specify all of a module's ports.

If you accidentally add a wrong port, remove it by selecting it in the **Available Port** list and clicking the **Remove** button.

After you click **Finish**, Configuration Expert adds the QoS Profile to those included in the Layer-2 QoS Profiles object.

Modifying QoS Profiles

After you create a QoS profile, you can modify a flow's definition to change its name, priority, or field values. For IP or IPX flows, you can also change to which interfaces a flow applies. For Layer-2 flows, you can change which ports are bound to a flow.

Separate discussions on redefining each different flow type, changing the interface list of IP or IPX flows, and changing the ports bound to a Layer-2 flow follow.

Redefining an IP Flow

You can redefine an IP flow to change its name, priority, or fields. You can also redefine a flow to create a new one based on an existing flow. To redefine an IP flow:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Double-click the IP QoS Profiles object.
4. From the list of IP QoS profiles that appears, click the one you want to edit.

A **QoS L3/L4 Flow Priority** dialog box similar to the following appears:

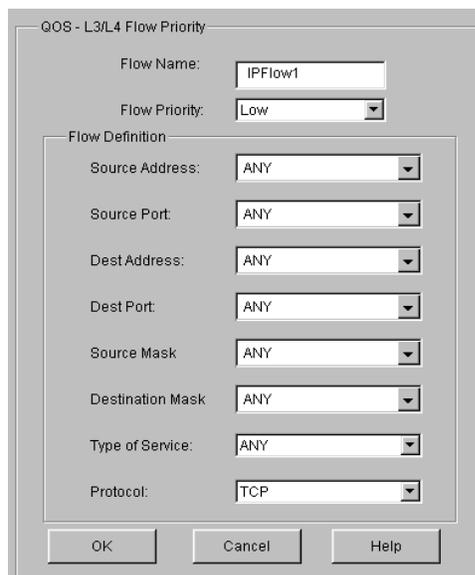


Figure 122. QoS L3/L4 Flow Priority dialog box

5. Change the flow's name, priority, or fields by editing the appropriate options.

The values for these options were specified when the QoS profile was created. For details on these options, see [“Creating a QoS Profile for an IP Flow” on page 167](#).

If you are creating a new flow based on an existing one, changing the flow name is optional. This lets you define flows that have the same name but different priority and field values.

6. Click **OK**.

Redefining an IPX Flow Definition

You can redefine an IPX flow to change its name, priority, or fields. You can also redefine a flow to create a new one based on an existing flow. To redefine an IPX flow:

1. Start Configuration Expert if you have not already done so.

2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Double-click the IPX QoS Profiles object.
4. From the list of IPX QoS profiles that appears, click the one you want to edit.

An **IPX Policy Definition** dialog box similar to the following appears:

Figure 123. IPX Policy Definition dialog box

5. Change the flow's name, priority, or fields by editing the appropriate options.

The values for these options were specified when the QoS profile was created. For details on these options, see [“Creating a QoS Profile for an IPX Flow” on page 172](#).

If you are creating a new flow based on an existing one, changing the flow name is optional. This lets you define flows that have the same name but different priority and field values.

6. Click **OK**.

Redefining a Layer-2 Flow Definition

You can redefine a Layer-2 flow to change its name, priority, or fields. You can also redefine a flow to create a new one based on an existing flow. To redefine a Layer-2 flow:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.

3. Double-click the Layer-2 QoS Profiles object.
4. From the list of Layer-2 QoS profiles that appears, click the one you want to edit.

An **L2 Flow Priority Definition** dialog box similar to the following appears:

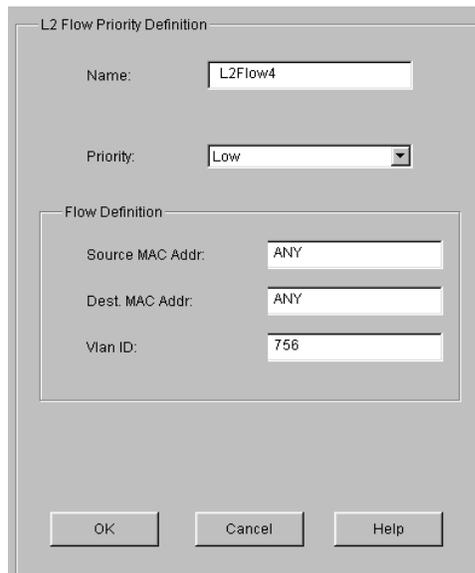


Figure 124. L2 Flow Priority Definition dialog box

5. Change the flow's name, priority, or fields by editing the appropriate options.

The values for these options were specified when the QoS profile was created. For more information on these options, see [“Creating a QoS Profile for a Layer-2 Flow” on page 177](#).

If you are creating a new flow based on an existing one, changing the flow name is optional. This lets you define flows that have the same name but different priority and field values.

6. Click **OK**.

Changing an IP or IPX Flow's Interface List

If you want to change which interfaces an existing IP or IPX flow applies to, modify the flow's list of interfaces. You can change a flow's interface list by using a dialog box to add or delete interfaces in the list, or by dragging an interface to apply a flow to that interface. Separate discussions on each task follow.

Adding or Deleting a Flow's Interfaces through a Dialog Box

To use a dialog box to add an interface to a flow's list of interfaces or delete one from the list:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and double-click the QoS Configuration object.
3. Do one of the following:
 - If you are modifying an IP flow's interface list, double-click the IP QoS Profiles object.
 - If you are modifying an IPX flow's interface list, double-click the IPX QoS Profiles object.
4. From the list of flows that appears, double-click the flow associated with the interface list you want to change.
5. Click the flow's Input Interface List object.

A **Policy Input Interface List** dialog box similar to the following appears:

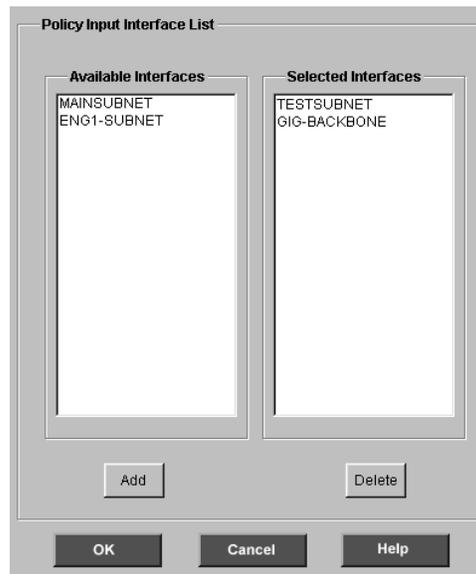


Figure 125. Policy Input Interface List dialog box

6. Specify which interface's you want the flow to apply to by adding and removing interfaces in the flow's interface list.
 - To add interfaces, select them in the **Available Interfaces** list and click the **Add** button.

- To delete interfaces, select them in the **Selected Interfaces** list and click the **Delete** button.
7. Click **OK**.

Dragging an Interface to Apply a Flow to the Interface

You can add an interface to a QoS profile's list of interfaces by dragging the interface to the flow. This will apply the flow to the interface. To apply a flow by dragging an interface to a QoS profile flow:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and double-click that file's QoS Configuration object. Then do one of the following:
 - If you want to apply a flow to an IP interface, double-click the IP QoS Profiles object.
 - If you want to apply a flow to an IPX interface, double-click the IPX QoS Profiles object.
3. Double-click the Routing Configuration object. Then do one of the following:
 - If you are applying an IP flow to an IP interface, double-click the IP Routing Configuration object. Then double-click the IP Interface Configuration object.
 - If you are applying an IPX flow to an IPX interface, double-click the IPX Routing Configuration object. Then double-click the IPX Interface Configuration object.
4. Double-click the object in which the desired interface is located.

If you are applying the flow to an IP interface, the interface is located in either the IP interfaces bound to Ports or IP interfaces bound to VLAN object.

If you are applying the flow to an IPX interface, the interface is located in either the IPX interfaces bound to Ports or IPX interfaces bound to VLAN object.
5. From the list of interfaces that appears, select the one you want to add to the flow.
6. Drag the selected interface to the desired QoS profile flow

Changing a Layer-2 Port List

If you want to change which ports an existing Layer-2 flow applies to, modify the flow's list of ports. To do so, take the following steps:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's QoS Configuration object.
3. Double-click the Layer-2 QoS Profiles object.

4. From the list of flows that appears, double-click the flow associated with the port list you want to change.
5. Click the flow's Bound Port List object.

A **Bound Port List** dialog box similar to the following appears:

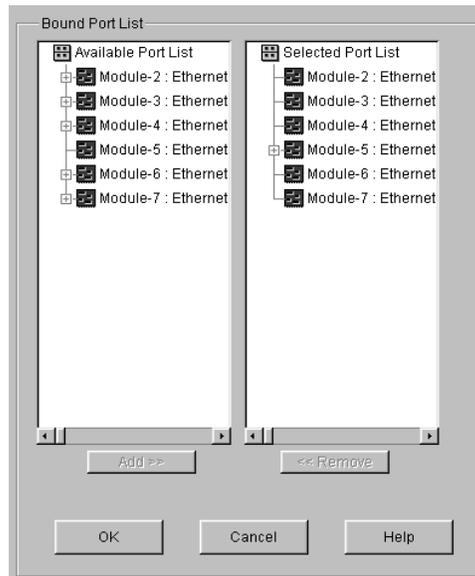


Figure 126. Bound Port List dialog box

6. Specify which ports you want the flow to apply to by adding and removing ports in the port list.
 - To add a port, double-click its module in the **Available Port** list, select the port from the list of ports that appears, and click the **Add** button.
 - To remove a port, double-click its module in the **Selected Port** list, select the port from the list of ports that appears, and click the **Remove** button.

Clicking a module rather than double-clicking it selects all of that module's ports. This is a quick way to add or remove all of a module's ports.

7. Click **OK**.

Chapter 13

Configuring Security on the GSR

You can configure security on the GSR by defining Access Control Lists (ACLs) for IP and IPX interfaces, applying those ACLs to interfaces, and setting Layer-2 filters. This chapter

- provides an overview that briefly describes ACLs.
- discusses setting security on IP networks.
- discusses setting security on IPX networks.
- describes applying an ACL to IP or IPX interfaces.
- discusses setting Layer-2 security filters.
- describes modifying the GSR's security settings.

A Look at ACLs

Access Control Lists (ACLs) allow you to restrict Layer 3 and Layer-4 traffic going through the router. Each ACL or each list consists of one or more rules describing a particular type of IP or IPX traffic. An ACL can be simple and consist of only one rule or complicated with many rules. Each rule tells the router to either permit or deny the packet that matches the rule's packet description. For more information on ACLs and rules, refer to the *DIGITAL GIGAswitch/Router User Reference Manual*.

After configuring ACLs you will need to apply them to interfaces. You may do so either while you are creating the interfaces or afterwards. For details on applying an ACL to an interface while creating the interface, see the chapter that discusses creating the type of interface (IP or IPX) to which you want to apply the ACL. If you want to apply an ACL after the interface has been created, see [“Applying ACLs to IP or IPX Interfaces” on page 206](#).

Setting IP Security

You can set security on an IP network by configuring ACLs that you will apply to IP interfaces. To set security:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the IP Security object.
4. Click the Configure New IP Security Profile object.

Configuration Expert opens the IP Security wizard.

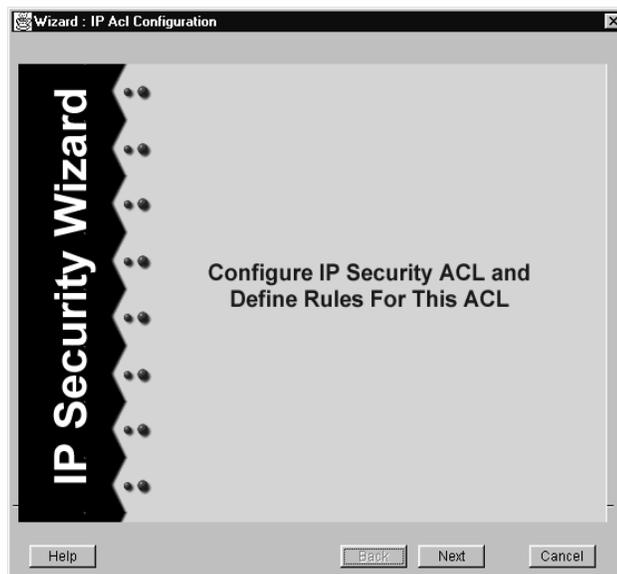


Figure 127. IP Security wizard

5. Click **Next**.

Configuration Expert prompts you for the ACL's name. You can use a string of characters or a number.

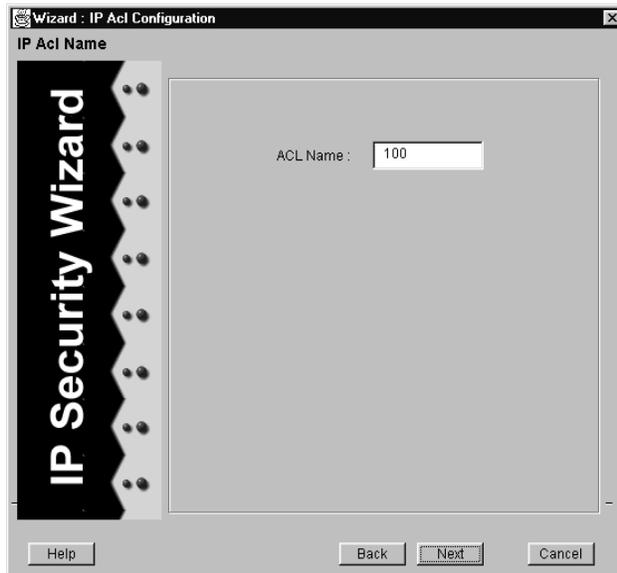


Figure 128. IP ACL Name panel

6. Enter the ACL's name in the **ACL Name** box and click **Next**.

An IP ACL Rule panel similar to the following appears:

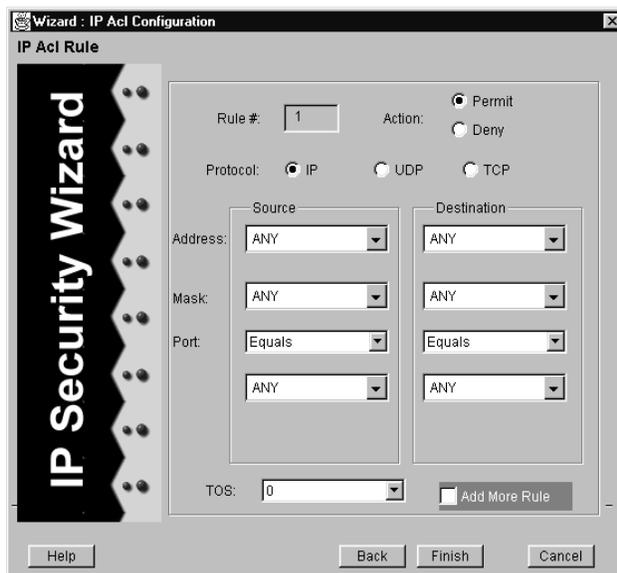


Figure 129. IP ACL Rule panel

7. If you want to permit traffic that meets the rule's criteria, select the *Permit* option. Otherwise, block such traffic by selecting the *Deny* option.
8. Specify the protocol to which the rule applies by selecting the appropriate option (*IP*, *UDP*, or *TCP*).
9. Define the rule's criteria by specifying values for the fields described in the following table. For each of the fields describing a flow, you can enter ANY to specify a wildcard ("don't care") condition. If you do not specify a value for a field, the GSR assumes that the value is a wildcard (as if you had entered ANY).

Table 23. IP/TCP/UDP ACL rule criteria fields

Field	Description
Source Address	Enter the source address of the flow.
Source Mask	<p>Enter the filtering mask of the flow. If the source address is a network or subnet address, you must supply the filtering mask. Generally, the filtering mask is the network mask of this network or subnet. If the source address is that of a host, then no mask is required.</p> <p>By default, if a mask is not supplied, the source address is treated as that of a host.</p> <p>You can specify the mask using the traditional IP address format ("255.255.0.0"). If you want to use the Classless Inter-Domain Routing (CIDR) format ("/16"), you may do so as discussed in the <i>DIGITAL GIGAswitch/Router Command Line Interface Reference Manual</i>.</p>

Table 23. IP/TCP/UDP ACL rule criteria fields (Continued)

Field	Description
Source Port	<p>For TCP or UDP, enter the number of the source TCP or UDP port. This field applies only to TCP or UDP traffic. If the incoming packet is ICMP or another non-TCP or non-UDP packet and you specified a source or destination port, the GSR does not check the port value. The GSR checks only the source and destination IP addresses in the packet.</p> <p>If a service uses the same port for both TCP and UDP, you do not need to define two separate rules. Instead, you can define one IP rule and specify the port that the service uses.</p> <p>You can specify a range of port numbers using operator symbols; for example, 10-20 (between 10 and 20 inclusive), >1024 (greater than 1024), !=1024 (not equal to 1024).</p> <p>The port numbers of some popular services are already defined as options. For example, for DNS, you can enter the port number 53 as well as select the DNS (53) option from the second Source Port field.</p>
Destination Address	Enter the destination address of the flow.
Destination Mask	<p>Enter the destination filtering mask of the flow.</p> <p>The same requirements and restrictions for Source Mask apply to Destination Mask.</p>
Destination Port	<p>For TCP or UDP, enter the number of the destination TCP or UDP port. This field applies only to incoming TCP or UDP traffic.</p> <p>The same requirements and restrictions for Source Port apply to Destination Port.</p>
TOS	Enter a TOS from 0 to 15.

10. Do one of the following:

- If you have defined all of the rules for the ACL, click **Finish**.
- If you want to define additional rules, select the *Add More Rules* check box and click **Next**.

11. If you selected the *Add More Rules* check box, define another rule in the IP ACL Rule panel that appears. To do so repeat [step 9](#) and [step 10](#) until you define all the desired rules for the ACL.

After you finish defining all of an ACL's rules, Configuration Expert adds the ACL to the IP Security object. Configuration Expert also adds a separate object for each rule and places this list of rules in the ACL object.

The rule numbers displayed in an ACL's list of rules, are automatically assigned by Configuration Expert. A rule's number is included in the Rule # box of the IP ACL Rule panel when you are defining that rule.

Setting IPX Security

You set security on an IPX network by configuring ACLs that you will apply to IPX interfaces. When setting IPX security, you can configure ACLs that perform the following functions:

- Permit or deny traffic from one computer to another. To set up this type of ACL, you configure an IPX ACL.
- Set up RIP filters, which permit or deny IPX RIP network advertisements. To set up such filters, you configure an IPX RIP ACL.
- Set up SAP filters, which permit or deny IPX SAP service advertisements. To set up such filters, you configure an IPX SAP ACL.

Separate discussions on configuring the different types of ACLs for an IPX network follow.

Configuring an IPX ACL

Configure an IPX ACL to permit or deny traffic from one computer to another. To configure an IPX ACL:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the IPX Security object.
4. Click the Configure New IPX Security Profile object.

Configuration Expert opens the IPX Security wizard.

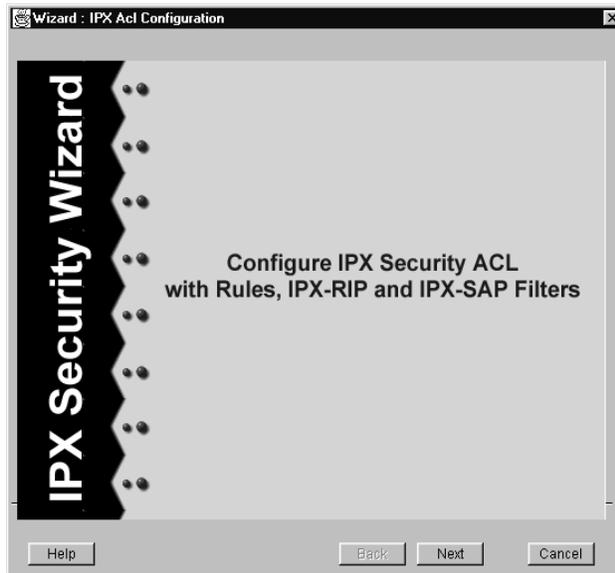


Figure 130. IPX Security wizard

5. Click Next.

Configuration Expert prompts you for the ACL's name.

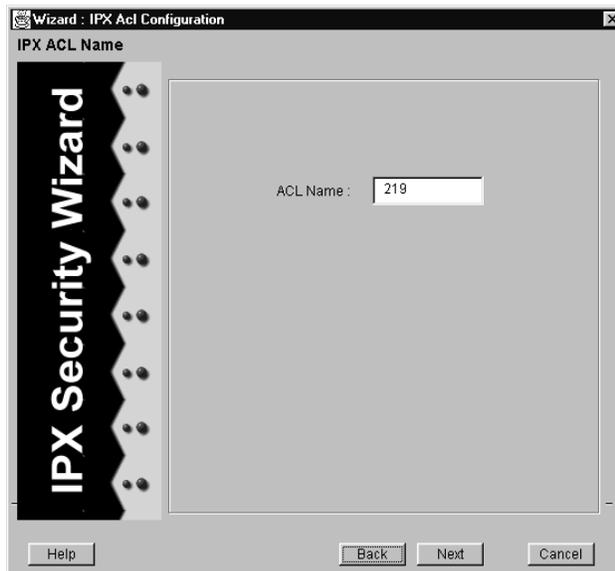


Figure 131. IPX ACL Name panel

6. Enter the ACL's name in the ACL Name box.

You can use a string of characters or a number.

7. Click **Next**.

An IPX ACL Type panel similar to the following appears:

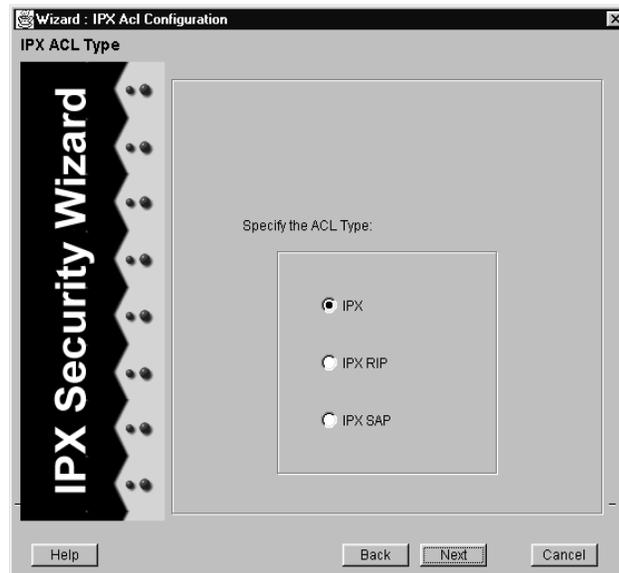


Figure 132. IPX ACL Type panel

8. Click **Next**.

An IPX ACL Rule panel similar to the following appears:

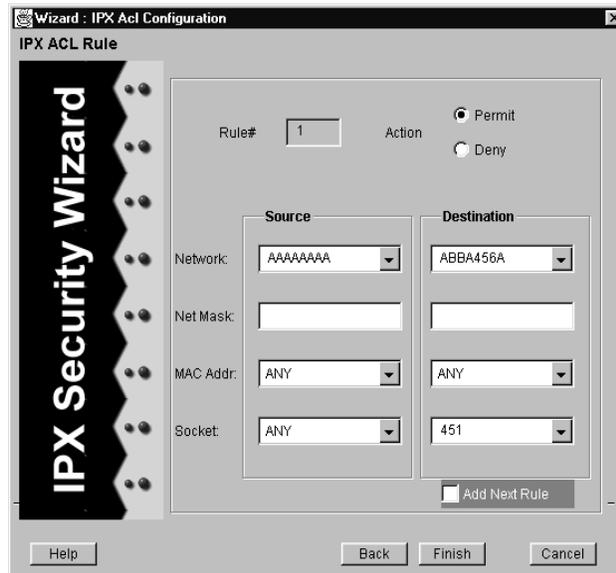


Figure 133. IPX ACL Rule panel

9. If you want to permit IPX traffic that meets the rule's criteria, select the *Permit* option. Otherwise, block such traffic by selecting the *Deny* option.
10. Define the rule's criteria by specifying values for the fields described in the following table:

Table 24. IPX ACL rule criteria fields

Field	Description
Source Network	Enter the source's network address. You can enter ANY to specify a wildcard ("don't care") condition. The GSR will interpret this number in hexadecimal format. You do not need to use a "Ox" prefix.
Source Net Mask	Enter the source's network mask. This field specifies a group of networks for which the ACL applies. This mask field is ANDed with the source's network address and the source network of the incoming packets to determine a hit. The GSR will interpret this number in hexadecimal format. You do not need to use a "Ox" prefix. Specifying this field is optional. If you do not enter a value in this field, the GSR uses the hexadecimal value FFFFFFFF.

Table 24. IPX ACL rule criteria fields (Continued)

Field	Description
Source MAC Addr	Enter the source's MAC address. You can enter ANY to specify a wildcard ("don't care") condition. The GSR will interpret this number in hexadecimal format. You do not need to use a "Ox" prefix.
Source Socket	Enter the source's IPX socket. You can enter ANY to specify a wildcard ("don't care") condition. The GSR will interpret this number in hexadecimal format. You do not need to use a "Ox" prefix.
Destination Network	Enter the destination's network address. The syntax for the destination address is the same as the syntax for the source address. You can enter ANY to specify a wildcard ("don't care") condition. The GSR will interpret this number in hexadecimal format. You do not need to use a "Ox" prefix.
Destination Net Mask	Enter the destination's network mask. This field specifies a group of networks for which the ACL applies. This mask field is ANDed with the destination's network address and the destination network of the incoming packets to determine a hit. The GSR will interpret this number in hexadecimal format. You do not need to use a "Ox" prefix. Specifying this field is optional. If you do not enter a value in this field, the GSR uses the hexadecimal value FFFFFFFF.
Destination MAC Addr	Enter the destination's MAC address. You can enter ANY to specify a wildcard ("don't care") condition. The GSR will interpret this number in hexadecimal format. You do not need to use a "Ox" prefix.
Destination Socket	Enter the source's IPX socket. You can enter ANY to specify a wildcard ("don't care") condition. The GSR will interpret this number in hexadecimal format. You do not need to use a "Ox" prefix.

11. Do one of the following:
 - If you have defined all of the rules for the ACL, click **Finish**.
 - If you want to define additional rules, select the *Add More Rules* check box and click **Next**.
12. If you selected the *Add More Rules* check box, define another rule in the IPX ACL Rule panel that appears. To do so repeat [step 10](#) and [step 11](#) until you define all the desired rules for the ACL.

After you finish defining all of an ACL's rules, Configuration Expert adds the ACL to the IPX ACLs object. Configuration Expert also adds a separate object for each rule and places this list of rules in the ACL object.

The rule numbers displayed in an ACL's list of rules, are automatically assigned by Configuration Expert. A rule's number is included in the Rule # box of the IPX ACL Rule panel when you are defining that rule.

Setting Up IPX RIP Filters

Set up IPX RIP filters to permit or deny IPX RIP network advertisements. You set up such filters by configuring an ACL for IPX RIP interfaces. To do so, take the following steps:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the IPX Security object.
4. Click the Configure New IPX Security Profile object.

Configuration Expert opens the IPX Security wizard.

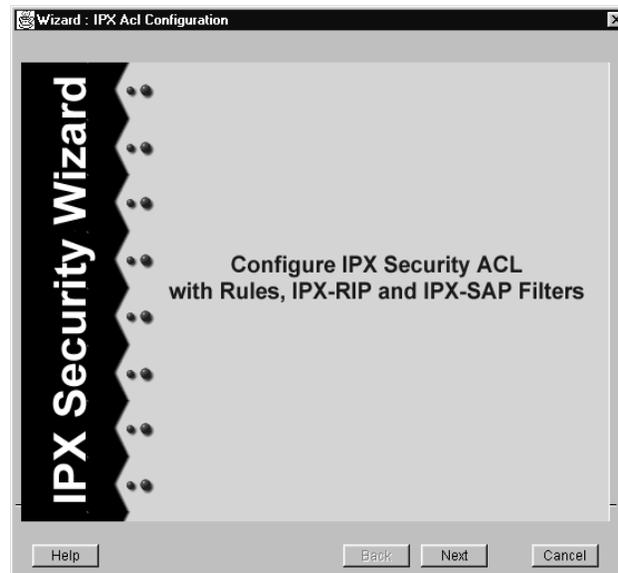


Figure 134. IPX Security wizard (RIP)

5. Click **Next**.

Configuration Expert prompts you for the ACL's name.

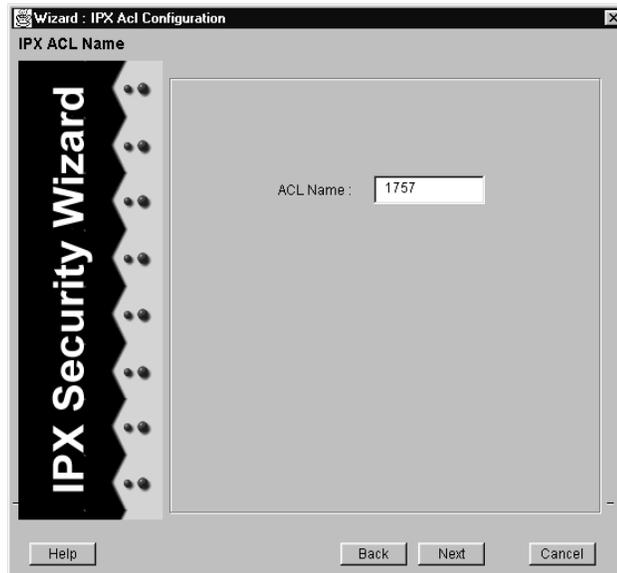


Figure 135. IPX ACL Name panel (RIP)

6. Enter the ACL's name in the **ACL Name** box.
You can use a string of characters or a number.
7. Click **Next**.

An IPX ACL Type panel similar to the following appears:

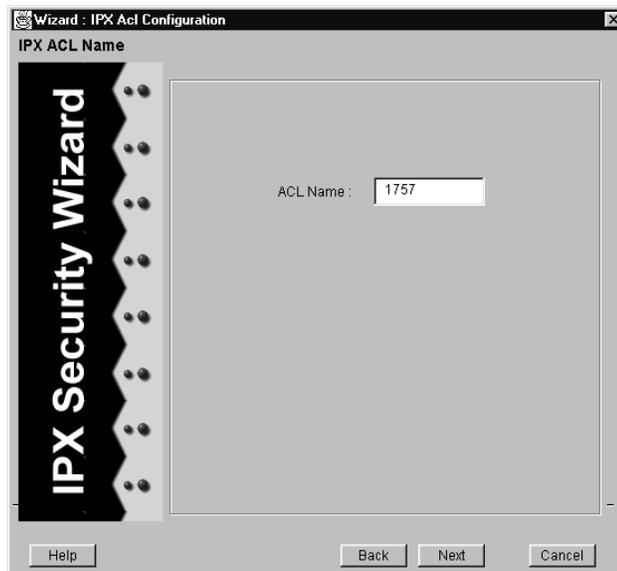


Figure 136. IPX ACL Type panel (RIP)

8. Select *IPX RIP* and click **Next**.

An IPX RIP ACL Rule panel similar to the following appears:

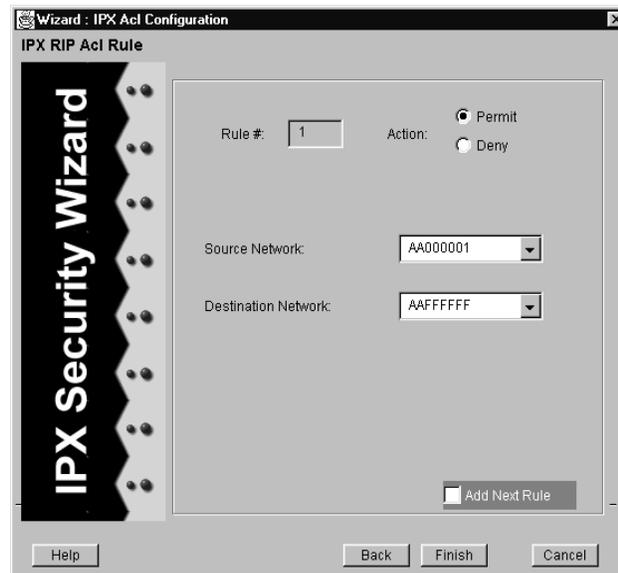


Figure 137. IPX ACL Rule panel (RIP)

9. If you want to permit IPX RIP network advertisements that meet the rule's criteria, select the *Permit* option. Otherwise, block such advertisements by selecting the *Deny* option.

10. Define the rule's criteria by specifying values for the fields described in the following table:

Table 25. IPX RIP ACL rule criteria fields

Field	Description
Source Network	Enter the source's network address. You can enter ANY to specify a wildcard ("don't care") condition. If you enter ANY, the GSR uses the value 0 for the source network address and FFFFFFFF for the destination network address.
Destination Network	Enter the destination's network address. Specifying this field is optional. If you do not enter a network address, the value that the GSR assumes depends on whether you entered ANY for the source's network address. <ul style="list-style-type: none"> • If you do not enter a value in this field and you entered ANY for the source's network address, the GSR sets the destination network address to FFFFFFFF. • If you do not enter a value in this field and did not enter ANY for the source's network address, the GSR sets destination network address to the same value you specified for the source's network address.

11. Do one of the following:
- If you have defined all of the rules for the ACL, click **Finish**.
 - If you want to define additional rules, select the *Add More Rules* check box and click **Next**.
12. If you selected the *Add More Rules* check box, define another rule in the IPX RIP ACL Rule panel that appears. To do so repeat [step 10](#) and [step 11](#) until you define all the desired rules for the ACL.

After you finish defining all of an ACL's rules, Configuration Expert adds the ACL to the IPX RIP ACLs object. Configuration Expert also adds a separate object for each rule and places this list of rules in the ACL object.

The rule numbers displayed in an ACL's list of rules, are automatically assigned by Configuration Expert. A rule's number is included in the Rule # box of the IPX ACL Rule panel when you are defining that rule.

Setting Up IPX SAP Filters

Set up IPX SAP filters to permit or deny IPX SAP service advertisements. You set up such filters by configuring an ACL for IPX SAP interfaces. To do so, take the following steps:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the IPX Security object.
4. Click the Configure New IPX Security Profile object.

Configuration Expert opens the IPX Security wizard.

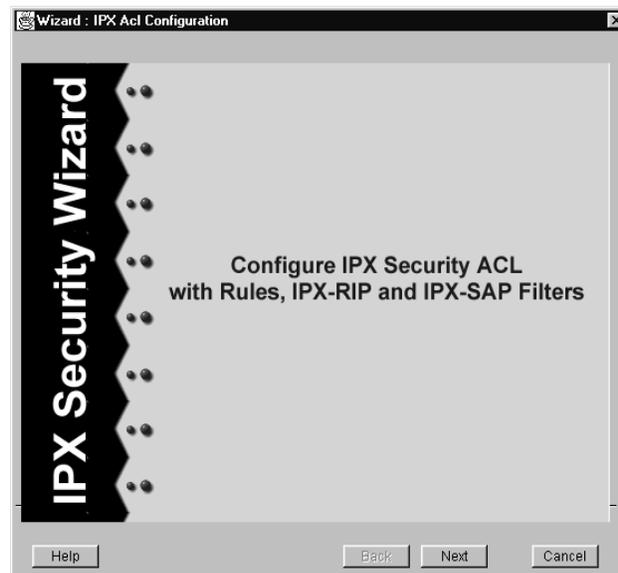


Figure 138. IPX Security wizard (SAP)

5. Click **Next**.

Configuration Expert prompts you for the ACL's name.

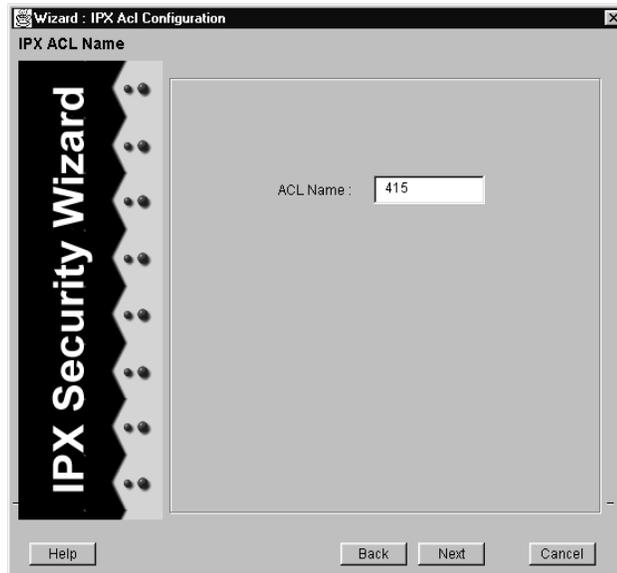


Figure 139. IPX ACL Name panel (SAP)

6. Enter the ACL's name in the **ACL Name** box.
You can use a string of characters or a number.
7. Click **Next**.

An IPX ACL Type panel similar to the following appears:

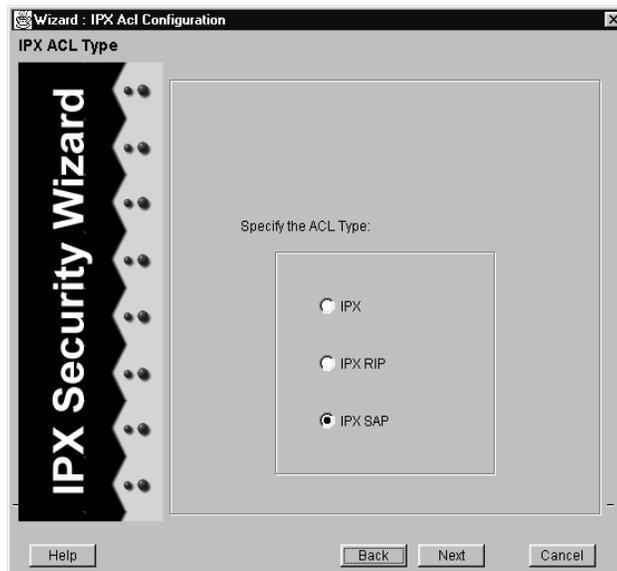


Figure 140. IPX ACL Type panel (SAP)

8. Select *IPX SAP* and click **Next**.

An IPX SAP ACL Rule panel similar to the following appears:

Figure 141. IPX ACL Rule panel (SAP)

9. If you want to permit IPX SAP service advertisements that meet the rule's criteria, select the *Permit* option. Otherwise, block such advertisements by selecting the *Deny* option.
10. Define the rule's criteria by specifying values for the fields described in the following table:

Table 26. IPX SAP ACL rule criteria fields

Field	Description
Server Name	Enter the server's name. Specifying this field is optional. You can enter ANY to specify a wildcard ("don't care") condition.

Table 26. IPX SAP ACL rule criteria fields (Continued)

Field	Description
Service Type	Enter the SAP service type. You may enter the service type as hexadecimal or select one of the choices from the Service Type drop-down list. You do not need to use a "0x" prefix. You can enter ANY to specify a wildcard ("don't care") condition.
Network Address	Enter the SAP server's network address. You can enter ANY to specify a wildcard ("don't care") condition.
Node (MAC) Address	Enter the SAP server's MAC address. You can enter ANY to specify a wildcard ("don't care") condition.

11. Do one of the following:
 - If you have defined all of the rules for the ACL, click **Finish**.
 - If you want to define additional rules, select the *Add More Rules* check box and click **Next**.
12. If you selected the *Add More Rules* check box, define another rule in the IPX SAP ACL Rule panel that appears. To do so repeat [step 10](#) and [step 11](#) until you define all the desired rules for the ACL.

After you finish defining all of an ACL's rules, Configuration Expert adds the ACL to the IPX SAP ACLs object. Configuration Expert also adds a separate object for each rule and places this list of rules in the ACL object.

The rule numbers displayed in an ACL's list of rules, are automatically assigned by Configuration Expert. A rule's number is included in the Rule # box of the IPX ACL Rule panel when you are defining that rule.

Applying ACLs to IP or IPX Interfaces

Defining an ACL specifies what sort of traffic to permit or deny. However, an ACL has no effect unless it is applied to an interface. An ACL can be applied to examine either inbound or outbound traffic. Inbound traffic is traffic coming into the router. Outbound traffic is traffic that is going out of the router. When you apply an ACL to an interface, you implicitly enable access control on that interface.

In general, you should try to apply ACLs at the inbound interfaces instead of the outbound interfaces. If a packet is to be denied, you want to drop the packet as early as possible, at the inbound interface. Otherwise, the router will have to process the packet, determine where the packet should go only to find out that the packet should be dropped at the outbound interface. In some cases, however, it may not be simple or possible for the

administrator to know ahead of time that a packet should be dropped at the inbound interface. Nonetheless, for performance reasons, whenever possible, one should create and apply an ACL to the inbound interface.

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

Note: When you apply an ACL to an interface, the GSR appends an *implicit deny rule* to that ACL. The implicit deny rule denies all traffic. If you intend to allow all traffic that does not match your specified ACL rules to go through, you must explicitly define a rule to permit all traffic. To do so, make sure the last rule of the ACL permits all traffic.

You can apply previously defined IP ACLs only to IP interfaces and previously defined IPX, IPX RIP, or IPX SAP ACLs only to IPX interfaces.



Caution: You can apply up to two IP ACLs to an IP interface, and you can apply two of each of the different IPX ACLs (IPX, IPX RIP, and IPX SAP) to an IPX interface. When applying multiple ACLs to an IP interface, one ACL must govern inbound traffic and the other ACL must govern outbound traffic. When applying multiple ACLs of the same type to an IPX interface, one ACL must govern inbound traffic and the other must govern outbound traffic.

You may apply an ACL to an interface either when you create the interface or afterwards. For details on applying an IP ACL while creating an IP interface, see [“Creating IP Interfaces” on page 92](#). For details on applying an IPX, IPX RIP, or IPX SAP ACL while creating an IPX interface, see [“Creating IPX Interfaces” on page 143](#).

You apply an IP or IPX, IPX RIP, or IPX SAP ACL to an interface after the interface is created by either copying the ACL or by editing the interface’s definition. Separate discussions on each task follow.

Copying an ACL to Apply It to an Interface

You can copy an ACL to apply it to an interface by either dragging it or using the Copy and Paste buttons. To apply an ACL by copying it to an interface:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file’s Routing Configuration object.
3. Expand the configuration tree until you locate the interface to which you want to apply the ACL. Double-click that interface’s object.

4. Do one of the following:
 - If you are applying the ACL to an IP interface, double-click the interface's Applied IP ACLs object.
 - If you are applying the ACL to an IPX interface, double-click the interface's Bound IPX Security object and then its Applied IPX ACLs object.
5. Expand the configuration tree until you locate the ACL you want to apply. Select that ACL and do one of the following:
 - If you are applying the ACL to an IP interface, drag the selected ACL to the Applied IP ACLs object of the interface to which you are applying the ACL.
 - If you are applying the ACL to an IPX interface, drag the selected ACL to the Applied IPX ACLs object of the interface to which you are applying the IPX ACL.
 - Click the Copy button and then select the Applied ACLs object of the interface to which you want to apply the ACL. Click the Paste button.
6. If you are applying multiple ACLs to an interface, configure those ACLs to govern either inbound traffic or outbound traffic. To do so, take the following steps:
 - a. Click an ACL that you want to apply to inbound traffic. In the **Access Control List: Edit ACL** dialog box that appears, ensure that there is a check mark in the *Input* check box and also make sure the *Output* check box is blank. Then click **OK**.
 - b. Click an ACL that you want to apply to outbound traffic. In the **Access Control List: Edit ACL** dialog box that appears, ensure that there is a check mark in the *Output* check box and also make sure the *Input* check box is blank. Then click **OK**.

Note: When applying multiple IP ACLs to an IP interface, one ACL must govern inbound traffic and the other ACL must govern outbound traffic. When applying multiple ACLs of the same type (IPX, IPX RIP, and IPX SAP) to an IPX interface, one ACL must govern inbound traffic and the other must govern outbound traffic.

Applying an ACL by Editing an Interface's Definition

To apply an ACL by editing an interface's definition:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Do one of the following:
 - If you are applying an IP ACL, double-click the IP Routing Configuration object. Then double-click the IP Interface Configuration object.
 - If you are applying an IPX, IPX RIP, or IPX SAP ACL, double-click the IPX Routing Configuration object. Then double-click the IPX Interface Configuration object.

4. Do one of the following:
 - If the interface you want to modify is bound to a port, double-click the bound to ports object of the interface to which you are applying the ACL.
 - If the interface you want to modify is bound to a VLAN, double-click the bound to VLAN object of the interface to which you are applying the ACL.
5. In the list of interfaces that appears, double-click the one you want to modify.
6. Do one of the following:
 - To apply an IP ACL, double-click the Applied IP ACLs object. In the **Update ACL List** dialog box that appears, remove the interface's current ACL and add the new one.
 - To apply an IPX, IPX RIP, or IPX SAP ACL, double-click the Applied IPX ACLs object. In the **Update ACL List** dialog box that appears, remove the interface's current ACL and add the new one.

The **Update ACL List** dialog box looks similar to the following figure. It includes IP ACLs when you are applying an IP ACL, and it includes IPX, IPX RIP, and IPX SAP ACLs when you are applying an IPX ACL.

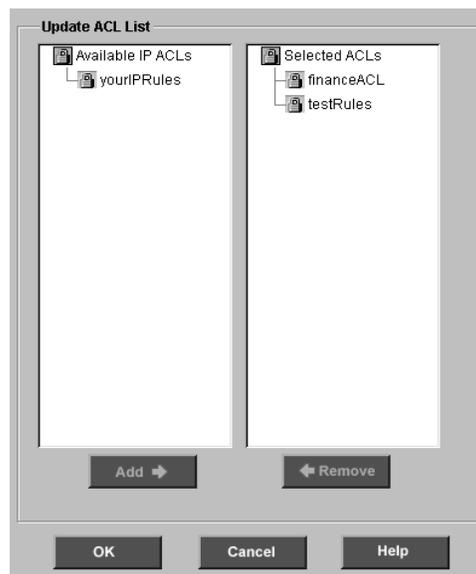


Figure 142. Update ACL List dialog box

You can add an ACL by selecting it in the list of available ACLs and then clicking the **Add** button. You can remove an ACL by selecting it in the list of selected ACLs and then clicking the **Remove** button.

7. Click **OK**.

8. If you are applying multiple ACLs to an interface, configure those ACLs to govern either inbound traffic or outbound traffic. To do so, take the following steps:
 - a. Click an ACL that you want to apply to inbound traffic. In the **Access Control List: Edit ACL** dialog box that appears, select the *Input* check box and ensure that the *Output* check box is not selected. Then click **OK**.
 - b. Click an ACL that you want to apply to outbound traffic. In the **Access Control List: Edit ACL** dialog box that appears, select the *Output* check box and ensure the *Input* check box is not selected. Then click **OK**.

Note: When applying multiple IP ACLs to an IP interface, one ACL must govern inbound traffic and the other ACL must govern outbound traffic. When applying multiple ACLs of the same type (IPX, IPX RIP, and IPX SAP) to an IPX interface, one ACL must govern inbound traffic and the other must govern outbound traffic.

Setting Layer-2 Security

Layer-2 security filters on the GSR allow you to configure ports to filter specific MAC addresses. When defining a Layer-2 security filter, you specify to which ports you want the filter to apply. You can specify the following security filters:

- Address filters

These filters block traffic based on the frame's source MAC address, destination MAC address, or both source and destination MAC addresses in flow bridging mode. Address filters are always configured and applied to the input port.
- Port-to-address lock filters

These filters prohibit a user connected to a locked port or set of ports from using another port.
- Static-entry filters

These filters allow or force traffic to go to a set of destination ports based on a frame's source MAC address, destination MAC address, or both source and destination MAC addresses in flow bridging mode. Static entries are always configured and applied at the input port.
- Secure port filters

These filters shut down access to the GSR based on MAC addresses. All packets received by a port are dropped. When combined with static entries, however, these filters can be used to drop all received traffic but allow some frames to go through.

Configuring Layer-2 Address Filters

If you want to control access to a source or destination on a per-MAC address basis, you can configure address filters. Address filters are always configured and applied to the input port. You can set address filters on the following:

- A source MAC address, which filters out any frame coming from a specific source MAC address.
- A destination MAC address, which filters out any frame destined to a specific destination MAC address.
- A flow, which filters out any frame coming from a specific source MAC address that is also destined to a specific destination MAC address.

To configure address filters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the L2 Security object.
4. Click the Configure New L2 Security Profile object.

Configuration Expert opens the L2 Security wizard.



Figure 143. L2 Security wizard (address filter)

5. Click Next.

Configuration Expert prompts you to select a filter type.

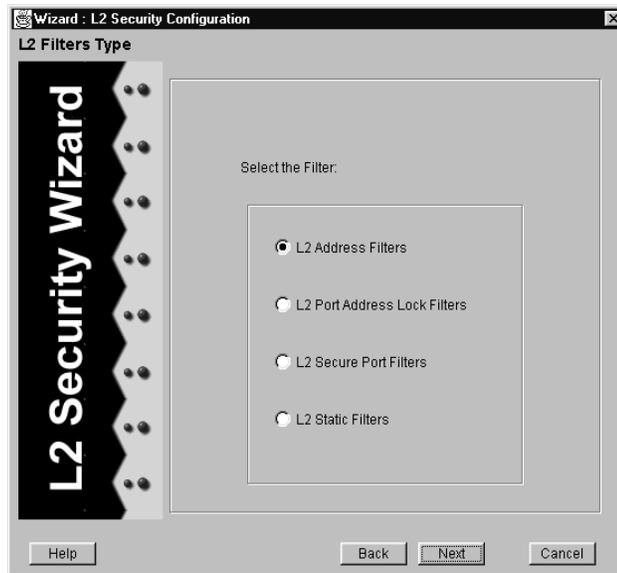


Figure 144. L2 Filter Type panel (address filter)

6. Click Next.

An L2 Address Filter panel similar to the following appears:

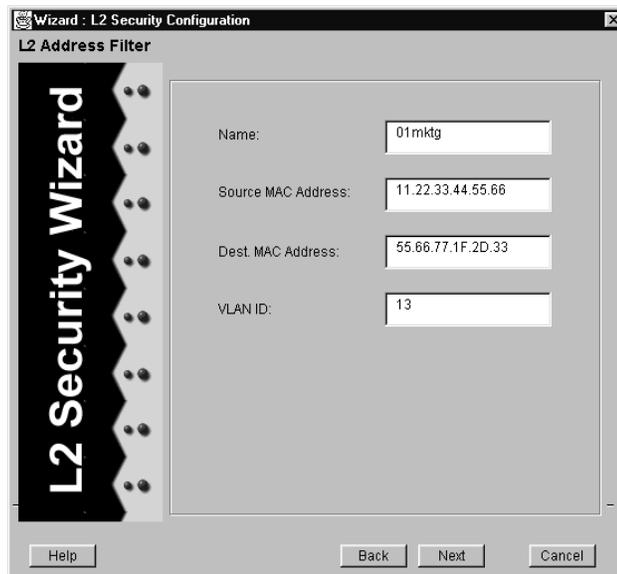


Figure 145. L2 Address Filter panel (address filter)

7. Enter the filter's name, source MAC address, destination MAC address, and VLAN ID in the appropriate text boxes.

Use the source MAC address for source or flow address filters. Use the Destination MAC Address for destination or flow static entries.

8. Click **Next**. In the Bind to Port panel that appears, specify to which ports you want to apply the filter.

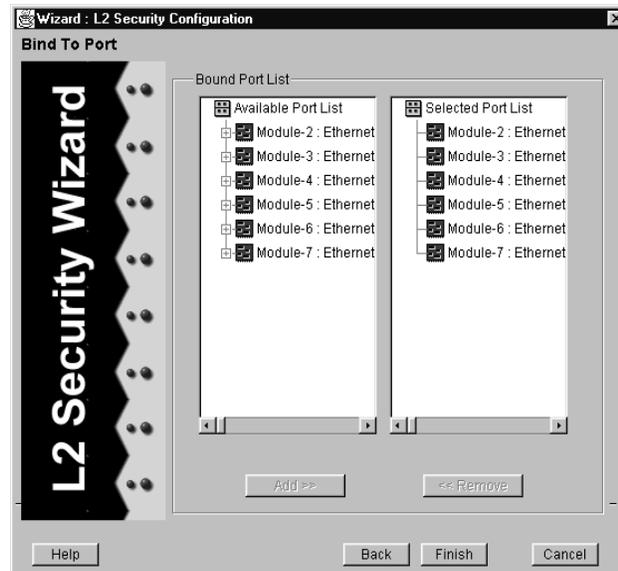


Figure 146. Bound Port List panel (address filter)

If you want to apply the filter to a port, double-click that port's module in the **Available Port** list, select the port from the port list that appears, and click the **Add** button.

Clicking a module rather than double-clicking it selects all of that module's ports. This is a quick way to apply the filter to all of a module's ports.

If you accidentally add a wrong port, remove it by selecting it in the **Selected Port** list and clicking the **Remove** button.

9. Click **Finish**.

Configuration Expert adds the filter to those found in the L2 Address Filters object. The ports to which the filter applies are included in that filter's Bound Port List object.

Configuring Layer-2 Port-to-Address Lock Filters

Port-to-address lock filters allow you to bind or "lock" specific source MAC addresses to a port or set of ports. Once a port is "locked," only the specified source MAC address is

allowed to connect to the “locked” port and the specified source MAC address is not allowed to connect to any other ports.

To configure a port-to-address lock filter:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file’s Security Configuration object.
3. Double-click the L2 Security object.
4. Click the Configure New L2 Security Profile object.

Configuration Expert opens the L2 Security wizard.



Figure 147. L2 Security wizard (lock filter)

5. Click **Next**.

Configuration Expert prompts you to select a filter type.

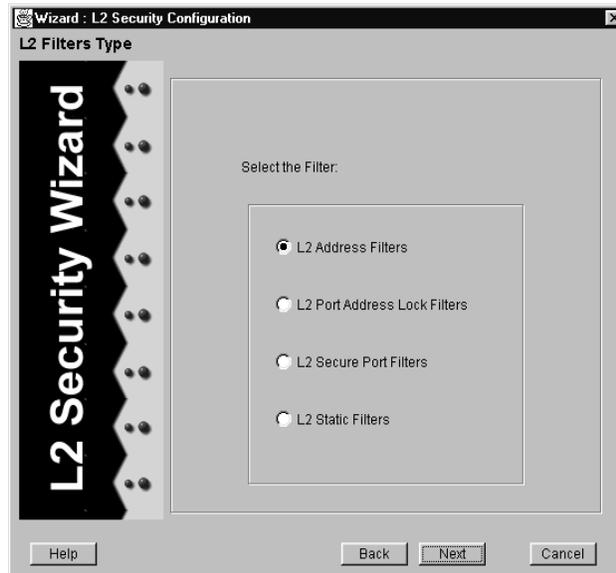


Figure 148. L2 Filter Type panel (lock filter)

6. Select *L2 Port Address Lock Filters* and click **Next**.

An L2 Port Address Lock Filter panel similar to the following appears:

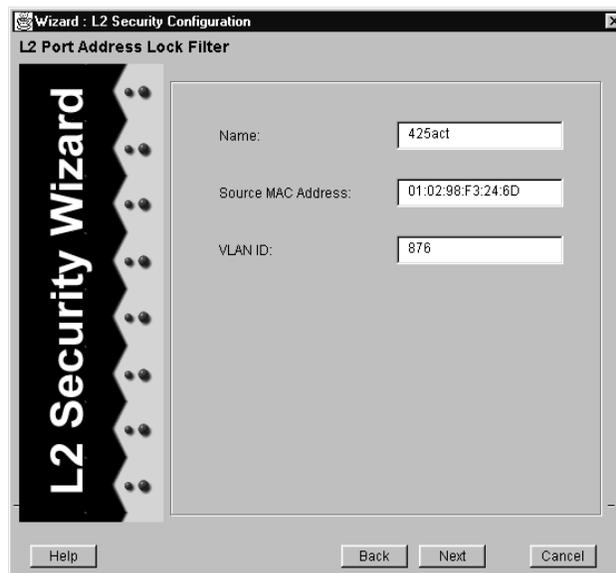


Figure 149. L2 Port Address Lock Filter panel

7. Enter the filter's name, source MAC address, and VLAN ID in the appropriate text boxes.
8. Click **Next**. In the Bind to Port panel that appears, specify to which ports you want to apply the filter. If you apply a port address lock filter to a port, you can use a static-entry filter to allow packets from a specific source to use that port even though it is locked.

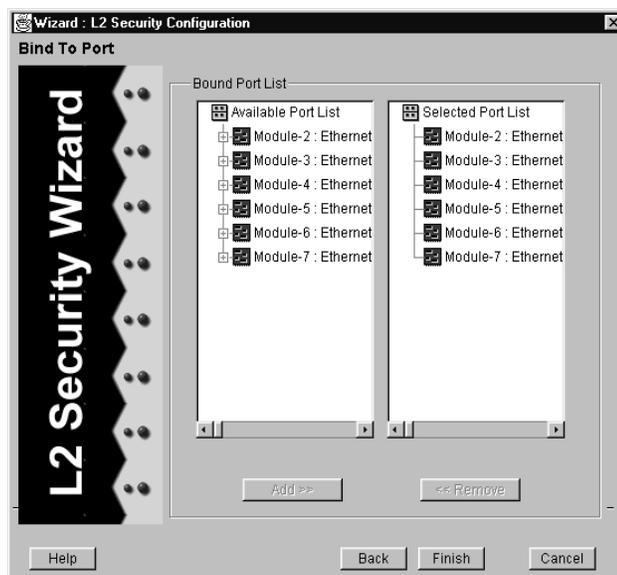


Figure 150. Bound Port List panel (lock filter)

If you want to apply the filter to a port, double-click that port's module in the **Available Port** list, select the port from the port list that appears, and click the **Add** button.

Clicking a module rather than double-clicking it selects all of that module's ports. This is a quick way to apply the filter to all of a module's ports.

If you accidentally add a wrong port, remove it by selecting it in the **Selected Port** list and clicking the **Remove** button.

9. Click **Finish**.

Configuration Expert adds the filter to those found in the L2 Port Address Lock Filters object. The ports to which the filter applies are included in that filter's Bound Port List object.

Configuring Layer-2 Static-Entry Filters

Static-entry filters allow or force traffic to go to a set of destination ports based on a frame's source MAC address, destination MAC address, or both source and destination

MAC addresses in flow bridging mode. Static entries are always configured and applied at the input port. You can set the following static-entry filters:

- Source static entry, which specifies that any frame coming from a specific source MAC address will be allowed or disallowed to go to a set of ports.
- Destination static entry, which specifies that any frame destined to a specific destination MAC address will be allowed, disallowed, or forced to go to a set of ports.
- Flow static entry, which specifies that any frame coming from a specific source MAC address that is destined to a specific destination MAC address will be allowed, disallowed, or forced to go to a set of ports.

To configure static-entry filters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the L2 Security object.
4. Click the Configure New L2 Security Profile object.

Configuration Expert opens the L2 Security wizard.

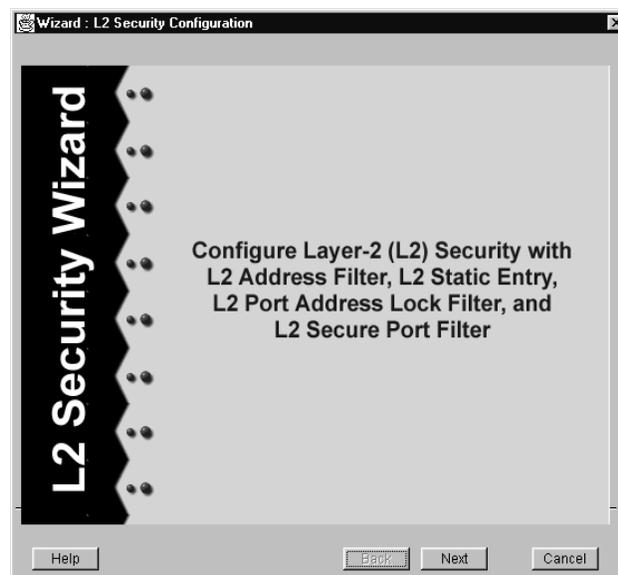


Figure 151. L2 Security wizard (static-entry filter)

5. Click **Next**.

Configuration Expert prompts you to select a filter type.

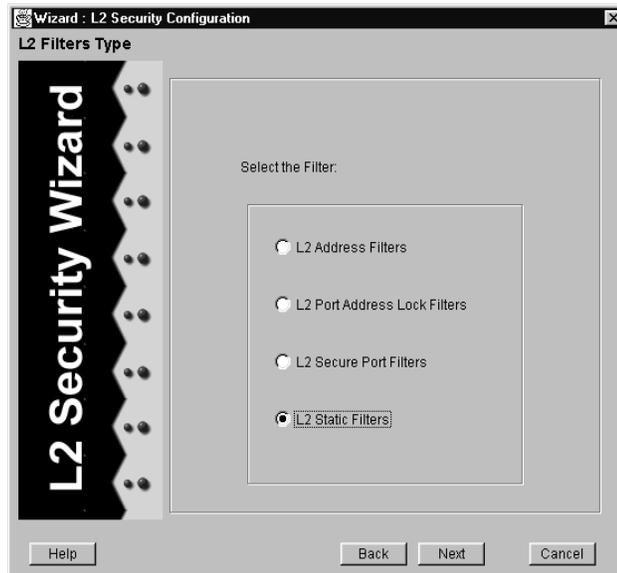


Figure 152. L2 Filter Type panel (static-entry filter)

6. Select *L2 Static Filters* and click **Next**.

An L2 Static Filter panel similar to the following appears:

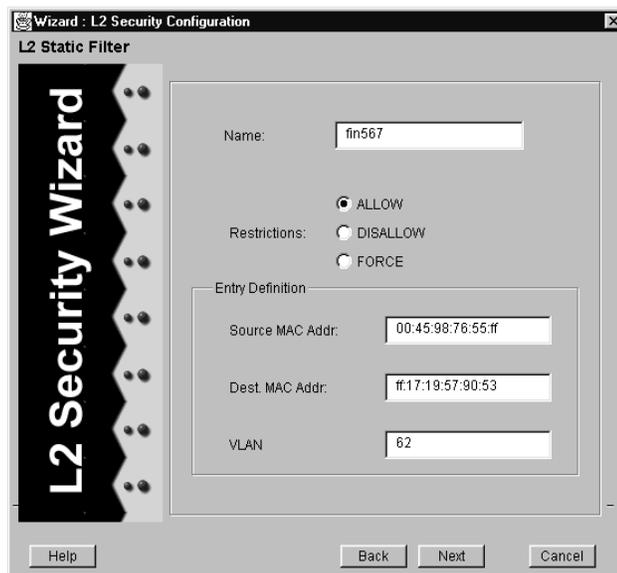


Figure 153. L2 Static Filter panel

7. Enter the filter's name in the **Name** box.

8. Specify the forwarding behavior of the static entry by doing one of the following:
 - Select *ALLOW* to allow packets to go to a specific set of ports.
 - Select *DISALLOW* to prohibit packets from going to a specific set of ports.
 - Select *FORCE* to force packets to go to a specific set of ports. The GSR will force the packets despite of any port locks in effect on the specified ports. Do not select this option if you are defining a source entry.

You will define the set of ports to which you are allowing, disallowing, or forcing packets later in the wizard.

9. Specify the source MAC address, destination MAC address, and VLAN ID in the appropriate text boxes.

Use the source MAC address for source or flow entries. Use the Destination MAC Address for destination or flow static entries.

10. Click **Next**. In the Bind to Port panel that appears, specify to which ports you want to apply the filter.

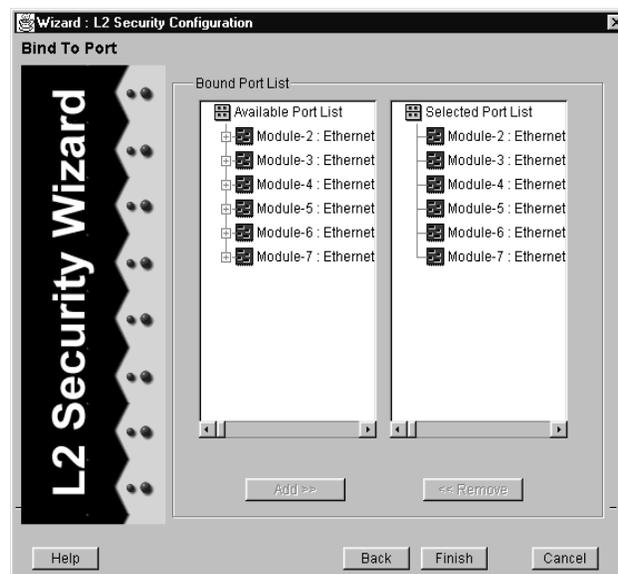


Figure 154. Bound Port List panel (static-entry filter)

To apply the filter to a port, double-click that port's module in the **Available Port** list, select the port from the port list that appears, and click the **Add** button.

Clicking a module rather than double-clicking it selects all of that module's ports. This is a quick way to apply the filter to all of a module's ports.

If you accidentally add a wrong port, remove it by selecting it in the **Selected Port** list box and clicking the **Remove** button.

11. Click **Next**. In the second Bind to Port panel that appears, specify the ports to which you are allowing, disallowing, or forcing packets. Then click **Finish**.

Configuration Expert adds the filter to those found in the L2 Static Entries object. The ports to which the filter applies are included in that filter's Bound Port List object. The ports to which you are allowing, disallowing, or forcing packets are included in the filter's Out Port List object.

Configuring Layer-2 Secure Port Filters

Secure port filters block access to a specified port. You can use a secure port filter by itself to secure unused ports. Secure port filters can be configured as source or destination port filters. A secure port filter applied to a source port forces all incoming packets to be dropped on a port. A secure port filter applied to a destination port prevents packets from going out a certain port.

You can combine this secure port filters with static entries in the following ways:

- Combine a source secure port filter with a source static entry to drop all received traffic but allow any frame coming from a specific source MAC address to go through.
- Combine a source secure port filter with a flow static entry to drop all received traffic but allow any frame coming from a specific source MAC address that is destined to a specific destination MAC address to go through.
- Combine a destination secure port with a destination static entry to drop all received traffic but allow any frame destined to a specific destination MAC address to go through.
- Combine a destination secure port filter with a flow static entry to drop all received traffic but allow any frame coming from a specific source MAC address that is destined to a specific destination MAC address to go through.

To configure secure port filters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the L2 Security object.
4. Click the Configure New L2 Security Profile object.

Configuration Expert opens the L2 Security wizard.

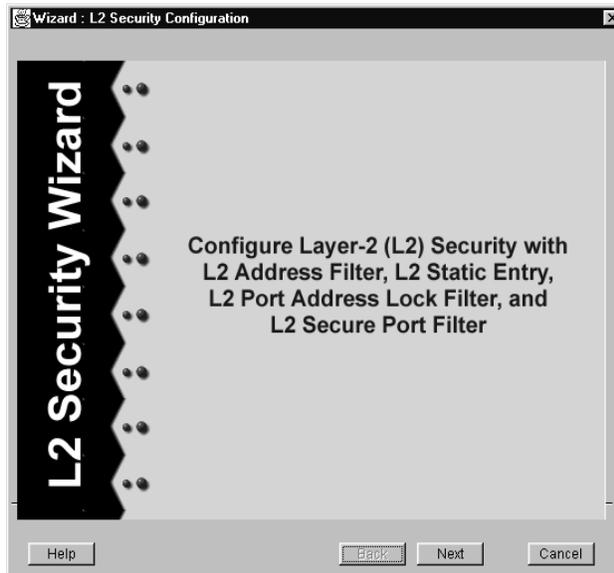


Figure 155. L2 Security wizard (secure port filter)

5. Click **Next**.

Configuration Expert prompts you to select a filter type.

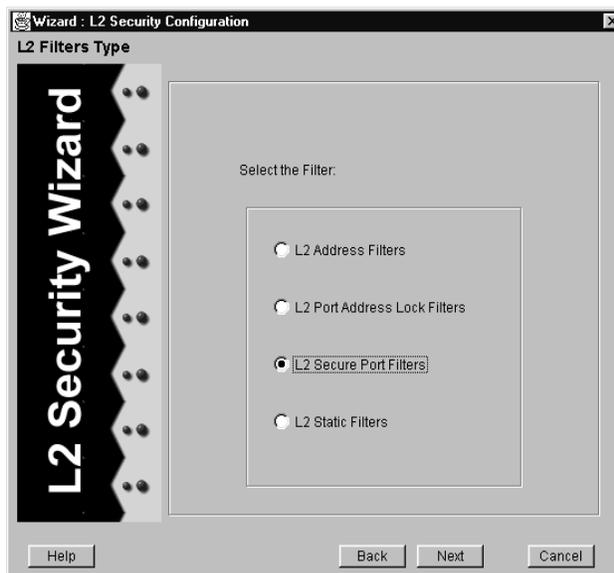


Figure 156. L2 Filter Type panel (secure port filter)

6. Select *L2 Secure Port Filters* and click **Next**.

An L2 Port Filter panel similar to the following appears:

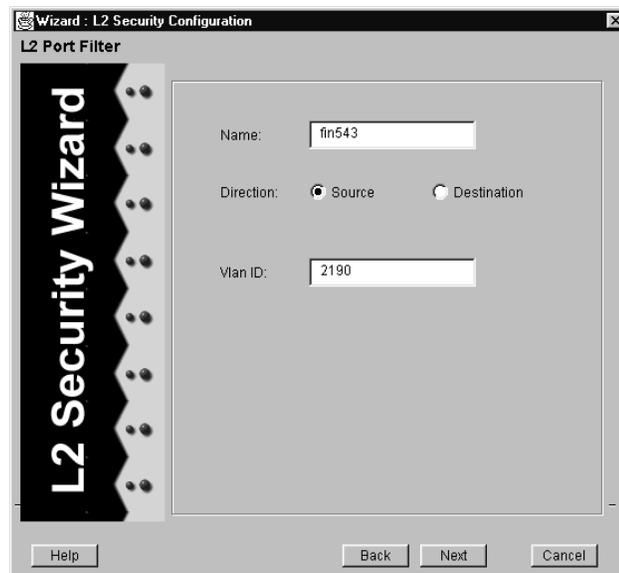


Figure 157. L2 Port Filter panel (secure port filter)

7. Enter the filter's name in the **Name** box.
8. Select either the *Source* or *Destination* option to specify whether the filter is to secure a source port or a destination port.
9. Specify the VLAN by entering its ID in the **VLAN ID** box. Then click **Next**.
10. In the Bind to Port panel that appears, specify to which ports you want to apply the filter.

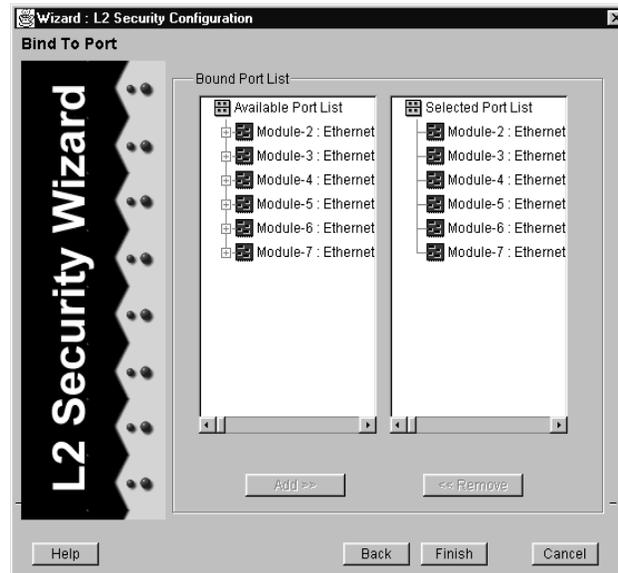


Figure 158. Bound Port List panel (secure port filter)

If you want to apply the filter to a port, double-click that port's module in the **Available Port** list, select the port from the port list that appears, and click the **Add** button.

Clicking a module rather than double-clicking it selects all of that module's ports. This is a quick way to apply the filter to all of a module's ports.

If you accidentally add a wrong port, remove it by selecting it in the **Selected Port** list box and clicking the **Remove** button.

11. Click **Finish**.

Configuration Expert adds the filter to those found in the L2 Secure Port Filters object. The ports to which the filter applies are included in that filter's Bound Port List object.

Modifying the GSR's Security Settings

You can modify a GSR's ACLs and Layer-2 security filters.

Modify an ACL if you want to change its name, add a rule to it, or edit an existing rule.

Modify a Layer-2 security filter if you want to change any of its values or redefine which ports are bound to the filter.

Changing an ACL's Name

To change an ACL's name:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Do one of the following:
 - Double-click the IP Security object and then the IP ACLs object if you are modifying an IP ACL.
 - Double-click the IPX Security object if you are modifying an IPX ACL. Then double-click the object (IPX ACLs, IPX RIP ACLs, or IPX SAP ACLs) appropriate for the type of IPX ACL you are modifying.
4. In the list of ACLs that appears, click the one you want to modify.

Configuration Expert displays the ACL's **Edit ACL** dialog box.

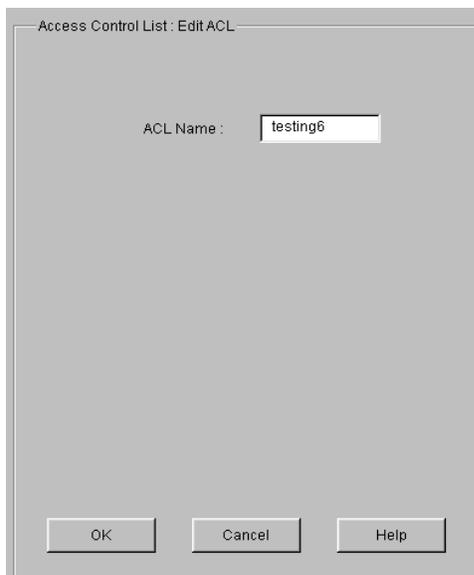


Figure 159. Edit ACL dialog box

5. Enter the new ACL name in the **ACL Name** box.
6. Click **OK**.

Adding or Modifying ACL Rules

To add a rule to an existing ACL or modify an ACL's rule:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Do one of the following:
 - Double-click the IP Security object and then the IP ACLs object if you are modifying an IP ACL.
 - Double-click the IPX Security object if you are modifying an IPX ACL. Then double-click the object (IPX ACLs, IPX RIP ACLs, or IPX SAP ACLs) appropriate for the type of IPX ACL you are modifying.
4. In the list of ACLs that appears, double-click the ACL you want to modify. Then click one of the following:
 - The rule you want to modify.
 - The Create New Rule object. Click this object if you want to add a new rule.
5. Edit the **ACL Rule Definition** dialog box that appears.

The dialog box's rule options are the same as those you specify when creating a new ACL. For details on an ACL's rule definition, see the section that discusses creating that type of ACL.

6. Click **OK**.

Modifying Layer-2 Security Filters

You can modify the settings of Layer-2 security filters and change their port bindings. Separate discussions on performing these tasks follow.

Modifying a Filter's Settings

To modify a filter's settings:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the L2 Security object.
4. Double-click the object appropriate for the type of filter you want to edit (L2 Address Filters, L2 Port Address Lock Filters, L2 Secure Port Filters, or L2 Static Entries).

5. In the list of filters that appears, click the one you want to modify.

Configuration Expert displays the selected filter's dialog box.

6. Edit the dialog box.

The options of a filter's dialog box are the same as those you specify when creating such a filter. For details on a filter's options, see the section that discusses setting that type of filter.

7. Click **OK**.

Modifying a Filter's Port Bindings

You can modify to which ports a filter applies. For Layer-2 static entries, you may also change to which ports the entry is allowing, disallowing, or forcing packets. To perform these tasks, you modify a filter's port bindings.

To modify a filter's port bindings:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Security Configuration object.
3. Double-click the L2 Security object.
4. Double-click the object appropriate for the type of filter you want to edit (L2 Address Filters, L2 Port Address Lock Filters, L2 Secure Port Filters, or L2 Static Entries).
5. In the list of filters that appears, double-click the one you want to modify.

Configuration Expert displays the filter's Bound Port List object. That object includes the ports to which the filter applies.

If you are modifying a Layer-2 static entry, Configuration Expert also displays an Out Port List object. That object includes the ports to which the entry is allowing, disallowing, or forcing packets.

6. Click the port list object you want to modify.

A **Bound Port List** dialog box similar to the following appears:

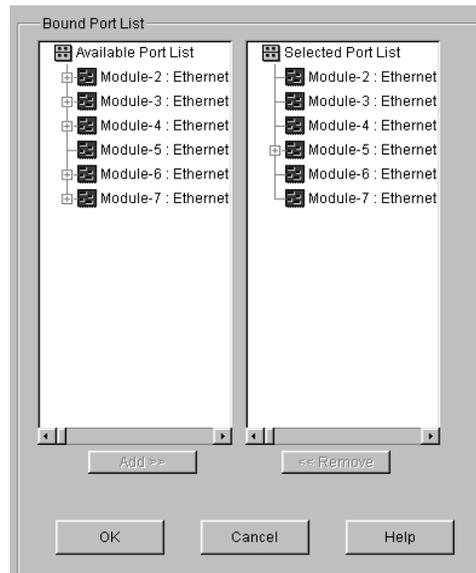


Figure 160. Bound Port List dialog box

7. Specify which port's you want to bind to the filter by adding and removing ports in the port list.
 - To add a port, double-click its module in the **Available Port** list, select the port from the list of ports that appears, and click the **Add** button.
 - To remove a port, double-click its module in the **Selected Port** list, select the port from the list of ports that appears, and click the **Remove** button.

Clicking a module rather than double-clicking it selects all of that module's ports. This is a quick way to add or remove all of a module's ports.

8. Click **OK**.

Chapter 14

Configuring OSPF on the GSR

A Look at OSPF Routing on the GSR

The Open Shortest Path Routing (OSPF) protocol is a link-state protocol. It is an Interior Gateway Protocol (IGP) that distributes routing information between routers in a single autonomous system.

OSPF chooses the least cost path as the best path. In a link-state protocol, each router maintains a database describing the entire autonomous-system topology, which it builds out of the collected link-state advertisements of all routers. Each participating router distributes its local state (such as the router's usable interfaces and reachable neighbors) throughout the autonomous system by flooding.

Each multiaccess network that has at least two attached routers has a designated router and a backup designated router. The designated router floods a link state advertisement for the multiaccess network and has other special responsibilities. The designated router concept reduces the number of adjacencies required on a multiaccess network. An adjacency is an OSPF relationship formed between selected neighboring routers for the purpose of exchanging routing information. Not every pair of neighboring routers become adjacent.

OSPF allows networks to be grouped into areas. Routing information passed between areas is abstracted, potentially allowing a significant reduction in routing traffic. OSPF uses four different types of routes, listed in order of preference: intra-area, inter-area, type 1 external, and type 2 external. Intra-area paths have destinations within the same area, inter-area paths have destinations in other OSPF areas and Autonomous System External (ASE) routes are routes to destinations external to the autonomous system.

Routes imported into OSPF as type 1 routes are supposed to be from IGPs whose external metrics are directly comparable to OSPF metrics. When a routing decision is being made,

OSPF will add the internal cost to the autonomous system border router to the external metric. Type 2 ASEs are used for Exterior Gateway Protocols (EGP) that have metrics not comparable to OSPF metrics. In this case, only the internal OSPF cost to the autonomous system border router is used in the routing decision.

From the topology database, each router constructs a tree of the shortest paths with itself as the root. This shortest-path tree gives the route to each destination in the autonomous system. Externally derived routing information appears on the tree as leaves. The link-state advertisement format distinguishes between information acquired from external sources and information acquired from internal routers, so there is no ambiguity about the source or reliability of routes. Externally derived routing information (for example, routes learned from EGP) is passed transparently through the autonomous system and is kept separate from OSPF's internally derived data. Each external route can also be tagged by the advertising router, enabling a passing of additional information between routers on the borders of the autonomous system.

OSPF intra-area and inter-area routes are always imported into the GSR routing database with a preference of 10. It would be a violation of the protocol if an OSPF router did not participate fully in the area's OSPF, so it is not possible to override this. Although it is possible to give other routes lower preference values explicitly, it is ill-advised to do so.

Hardware multicast capabilities are also used where possible to deliver link-status messages. OSPF areas are connected by the backbone area, the area with identifier 0.0.0.0. All areas must be logically contiguous and the backbone is no exception. To permit maximum flexibility, OSPF allows the configuration of virtual links to enable the backbone area to appear contiguous despite the physical reality.

All routers in an area must agree on that area's parameters. A separate copy of the link-state algorithm is run for each area. Because of this, most configuration parameters are defined on a per area basis. All routers belonging to an area must agree on that area's configuration. Erroneous configuration will lead to adjacencies not forming between neighbors, and routing information might not flow, or even loop.

Setting OSPF Global Parameters

You can use clearVISN CoreWatch to set global OSPF parameters to start or stop OSPF on the GSR and specify how often autonomous-system export link-state advertisements will be generated and flooded into OSPF.

To set the GSR's global OSPF parameters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.

4. Double-click the IP Unicast Routing object. Then double-click the OSPF object and click the OSPF Global Parameters object.

An **OSPF Global Parameters** dialog box similar to the following appears:

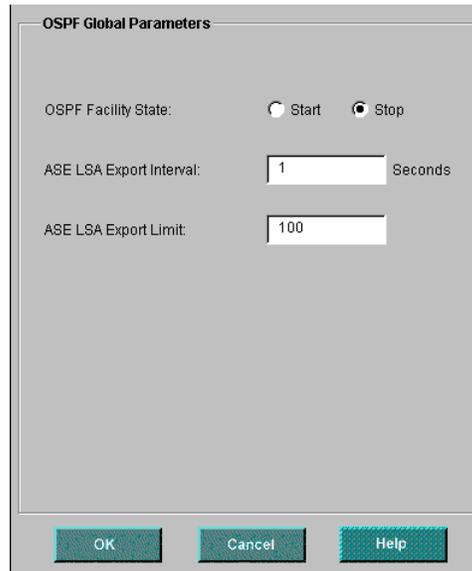


Figure 161. OSPF Global Parameters dialog box

5. Set the *OSPF Facility State* option to *Start* or *Stop*.
6. In the **ASE LSA Export Interval** box, enter the interval at which autonomous-system, link-state advertisements are generated and flooded into OSPF.
Specify an integer value equal to or greater than 1. The default is 1.
7. In the **ASE LSA Export Limit** box, specify how many autonomous systems will be generated and flooded into each batch.
Specify an integer value equal to or greater than 1. The default is 100.
8. Click **OK**.

Configuring OSPF Area Tables

You can configure an OSPF area table to create an OSPF area and add interfaces, neighbors, and stub hosts to it. The method you use to configure OSPF areas depends upon whether you are creating a new area or modifying an existing one. Separate discussions on each task follow.

Creating OSPF Area Tables

To create OSPF area tables:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object and then double-click the OSPF Area Table object.
5. Click the Configure New OSPF Area object.

Configuration Expert opens the OSPF Area wizard.

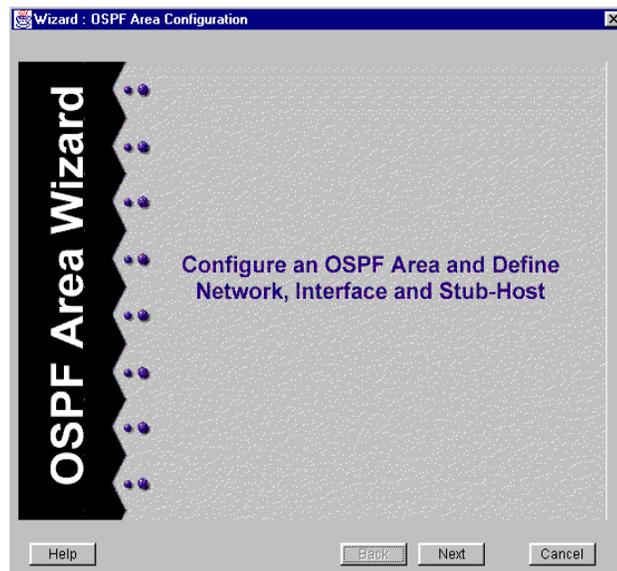


Figure 162. OSPF Area wizard

6. Click Next.

The OSPF Area Definition panel appears.

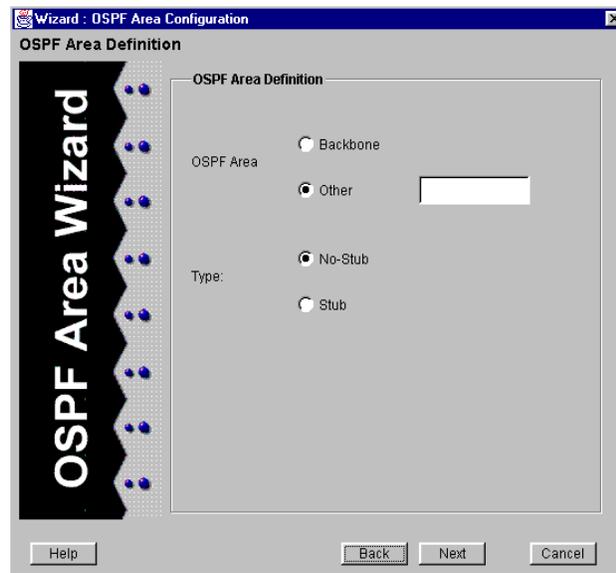


Figure 163. OSPF Area Definition panel

7. Create an area by either selecting the *Backbone* option to add the backbone area or selecting the *Other* option and then entering an ID number in the **Area ID** box.
8. Specify whether the area you are adding is a stub by selecting the appropriate option.
If you select *Stub*, specify the cost to be used to inject a default route into the area. To do so, enter a number from 0 to 65535 in the **Cost** box that appears after you select the *Stub* option.
9. Click **Next**. In the panel that appears, specify whether you want to associate a network range with the area by selecting the appropriate option. Click **Next**.

Network ranges are used for summarization. Intra-area link-state advertisements that fall within the specified ranges are not advertised into other areas as inter-area routes. Instead, the specified ranges are advertised as summary network link-state advertisements.

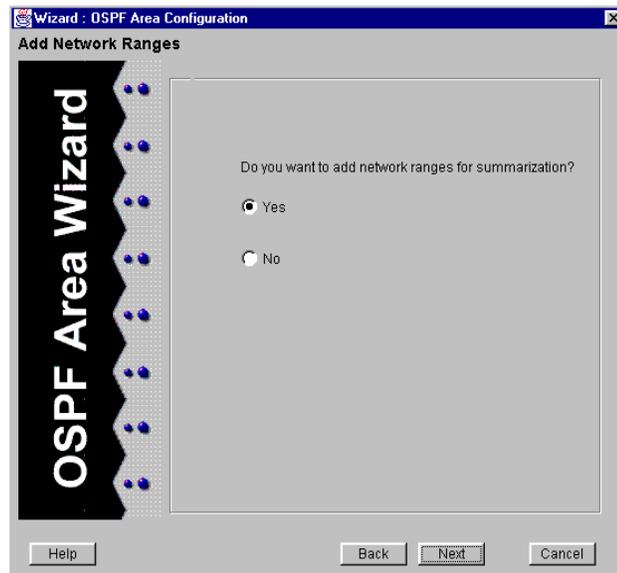


Figure 164. Add Network Ranges panel

10. If you selected *Yes* to specify you want to associate a network with the area, use the OSPF Network Definition panel that appears to specify which network you want to associate with the area. Otherwise, skip to [step 11](#).

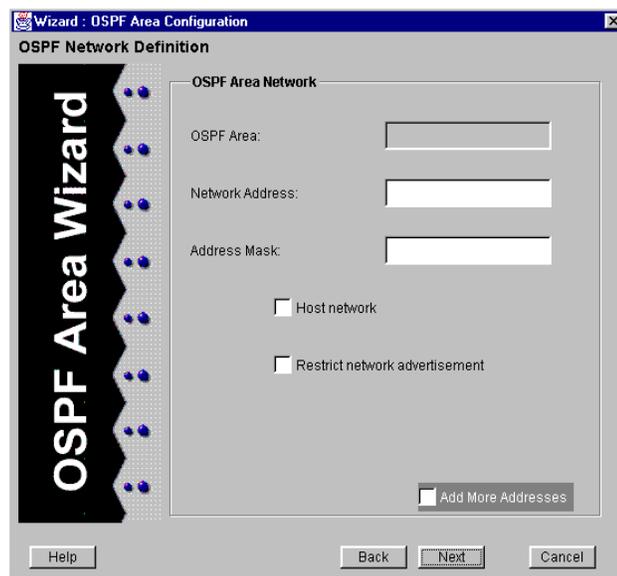


Figure 165. OSPF Area Network panel

- a. Enter the network address and address mask in the appropriate text boxes.

- b. If the specified network is a host network, then check the host-network box.
 - c. If you do not want to advertise the network or host network in the Summary Network link-state advertisements, select *Restrict Network Advertisement*.
 - d. To associate another network with the area, select the *Add More Address* checkbox, click **Next**, and repeat [step a](#) through [step c](#) until you associate all the appropriate networks.
11. Click **Next**. In the panel that appears, specify whether you want to add a broadcast interface by selecting the appropriate option. Then click **Next**.

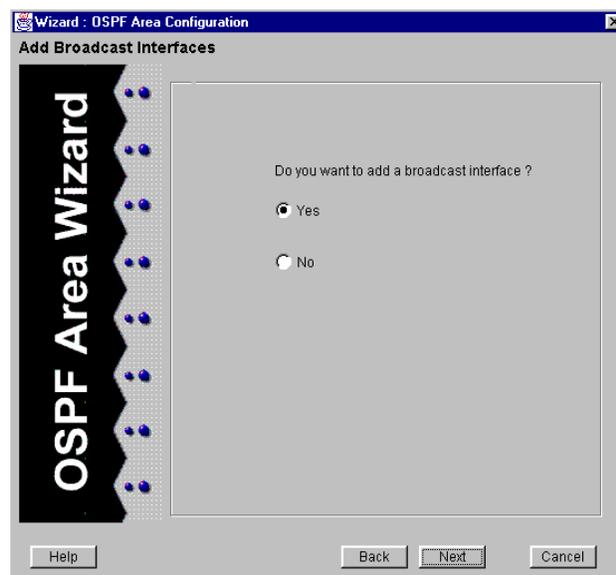


Figure 166. Add Broadcast Interface panel

12. If you selected *Yes* to specify you want to add a broadcast interface, add the interface. Otherwise, skip to [step 13](#).

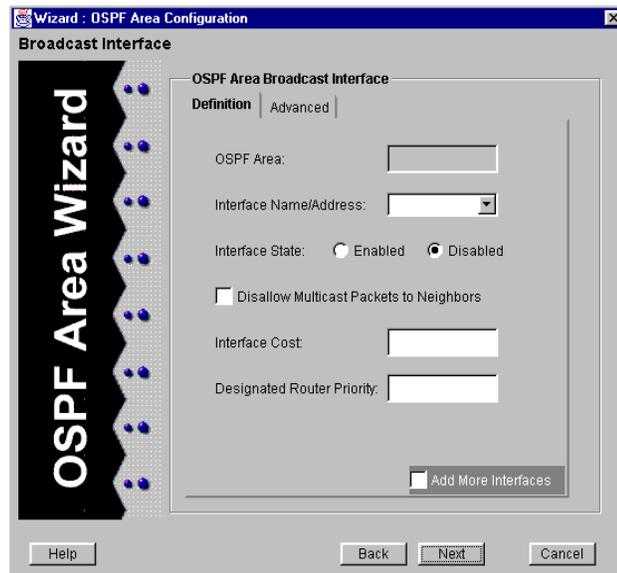


Figure 167. OSPF Area Broadcast Interface panel (Definition tab)

To add a broadcast interface, take the following steps:

- a. From the **Interface/Name Address** drop-down list, select the IP interface you want to configure as a broadcast interface.
- b. Select the appropriate *Interface State* option to enable or disable the interface.
- c. Select or clear the *Disallow Multicast Packets to Neighbors* check box to specify whether the GSR is to send multicast packets to neighbors on point-to-point interfaces.

- d. Enter the interface cost and designated router priority for the broadcast interface in the appropriate text boxes. See the following table for more detailed information:

Table 27. Broadcast interface cost and designated router priority description

Option	Description
Interface Cost	Enter the sum of all interfaces a packet from the area must cross to reach the interface.
Designated Router Priority	<p>Enter an integer value between 0 and 255 to specify the priority for becoming the designated router on the interface. The default is 0.</p> <p>When two routers attached to a network both attempt to become the designated router, the one with the highest priority wins. A router with a router priority of 0 is ineligible to become a designated router.</p>

- e. To add another interface, select the *Add More Interfaces* check box, click **Next**, and repeat [step a](#) through [step d](#) until you add all the appropriate interfaces.
- f. Click the **Advanced** tab.

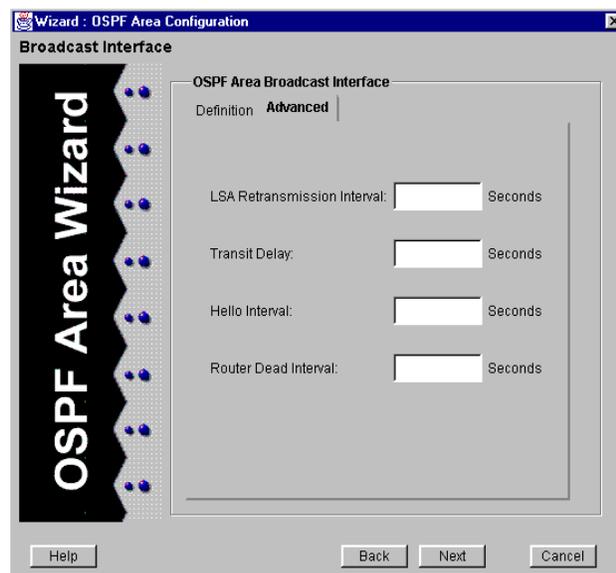


Figure 168. OSPF Area Broadcast Interface panel (Advanced tab)

- g. Set the advanced options on the interface as discussed in the following table:

Table 28. Broadcast interface options

Option	Description
LSA Retransmission Interval	Enter the number of seconds between link-state advertisement retransmissions for adjacencies belonging to the interface. Specify an integer value equal to or greater than 1. The default is 5 seconds.
Transit Delay	Enter the estimated number of seconds required to transmit a link-state update over the interface. Transit delay takes into account transmission and propagation delays and must be greater than 0. Specify an integer value equal to or greater than 1. The default is 1 second.
Hello Interval	Enter an integer value from 0 to 255 to specify the estimated number of seconds between Hello packets that the router sends on the interface. The default is 10 seconds for broadcast networks and 30 seconds for both point-to-point and non-broadcast interfaces.
Router Dead Interval	Enter an integer value from 0 to 255 to specify the number of seconds that may occur without Hello packets being heard before the router's neighbors will declare it down. By default, this option is four times the value of the hello interval.

13. Click **Next**. In the panel that appears, specify whether you want to add a non-broadcast interface by selecting the appropriate option.



Figure 169. Add Non-Broadcast Interface panel

14. If you selected *Yes* to specify you want to add a non-broadcast interface, add the interface. Otherwise, skip to [step g](#).

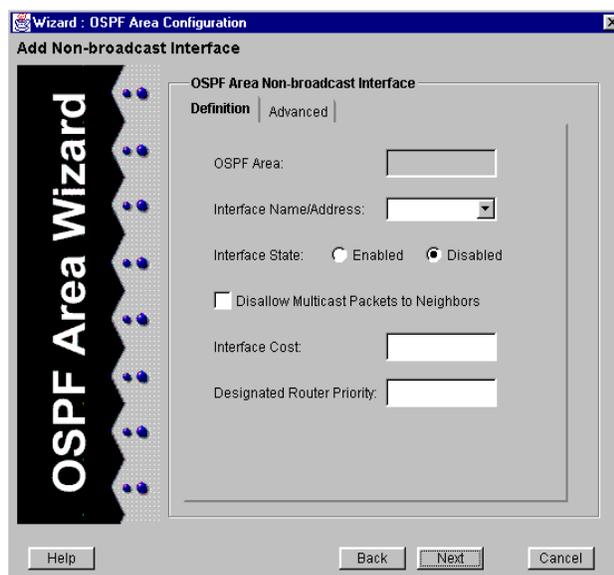


Figure 170. OSPF Area Non-Broadcast Interface panel (Definition tab)

To add a non-broadcast interface, take the following steps:

- a. From the **Interface/Name Address** drop-down list, select the IP interface you want to configure as a non-broadcast interface.
- b. Select the appropriate *Interface State* option to enable or disable the interface.
- c. Select or clear the *Disallow Multicast Packets to Neighbors* check box to specify whether the GSR is to send multicast packets to neighbors on point-to-point interfaces.
- d. Enter the interface cost and designated router priority for the non-broadcast interface in the appropriate text boxes. See the following table for more detailed information:

Table 29. Non-broadcast interface cost and designated router priority description

Option	Description
Interface Cost	Enter the sum of all interfaces a packet from the area must cross to reach the interface.
Designated Router Priority	<p>Enter an integer value between 0 and 255 to specify the priority for becoming the designated router on the interface. The default is 0.</p> <p>When two routers attached to a network both attempt to become the designated router, the one with the highest priority wins. A router with a router priority of 0 is ineligible to become a designated router.</p>

- e. Click the **Advanced** tab.

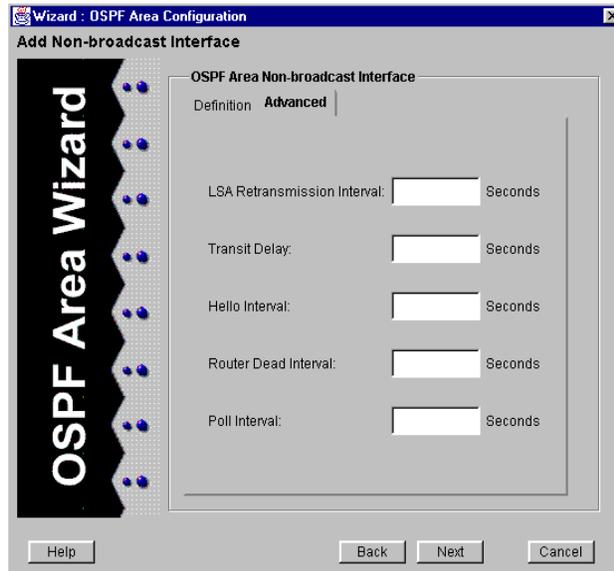


Figure 171. OSPF Area Non-Broadcast Interface panel (Advanced tab)

- f. Set the advanced options on the interface as discussed in the following table:

Table 30. Non-broadcast interface options

Option	Description
LSA Retransmission Interval	Enter the number of seconds between link-state advertisement retransmissions for adjacencies belonging to the interface. Specify an integer value equal to or greater than 1. The default is 5 seconds.
Transit Delay	Enter the estimated number of seconds required to transmit a link-state update over the interface. Transit delay takes into account transmission and propagation delays and must be greater than 0. Specify an integer value equal to or greater than 1. The default is 1 second.

Table 30. Non-broadcast interface options (Continued)

Option	Description
Hello Interval	Enter an integer value from 0 to 255 to specify the estimated number of seconds between Hello packets that the router sends on the interface. The default is 10 seconds for broadcast networks and 30 seconds for both point-to-point and non-broadcast interfaces.
Router Dead Interval	Enter an integer value from 0 to 255 to specify the number of seconds that may occur without Hello packets being heard before the router's neighbors will declare it down. By default, this option is four times the value of the hello interval.
Poll Interval	Enter an integer value from 0 to 255 to specify the estimated number of seconds between polls to the non-broadcast interface. The default value for this interval is 120 seconds.

- g. Click **Next**.

The OSPF Area Interface Neighbor panel appears.

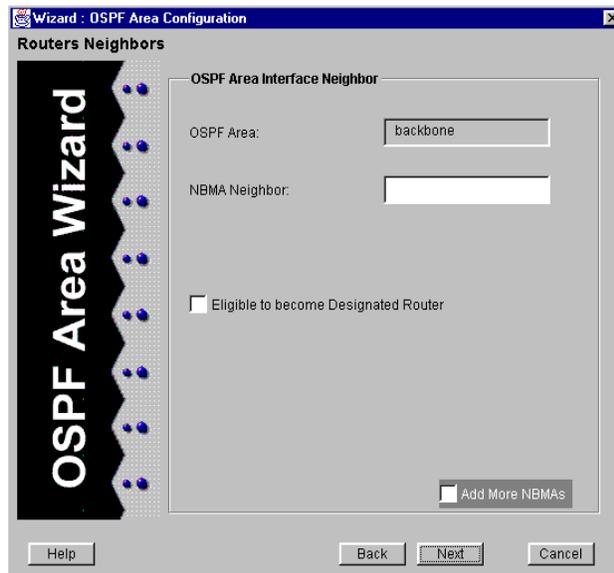


Figure 172. OSPF Area Interface Neighbor panel

- h. Add a neighbor to the interface by taking the following steps:
 - i. Specify the IP address of the NBMA neighbor you want to add to the interface.
 - ii. Select or clear the *Eligible to Become Designated Router* check box to specify whether the OSPF NBMA neighbor is eligible for becoming a designated router.
 - iii. To add another neighbor, select the *Add More NBMA's* check box, click **Next**, and repeat [step i](#) and [step ii](#) until you add all the appropriate neighbors.
- i. Click **Next**. In the panel that appears, specify whether you want to add more non-broadcast interfaces by selecting the appropriate option.

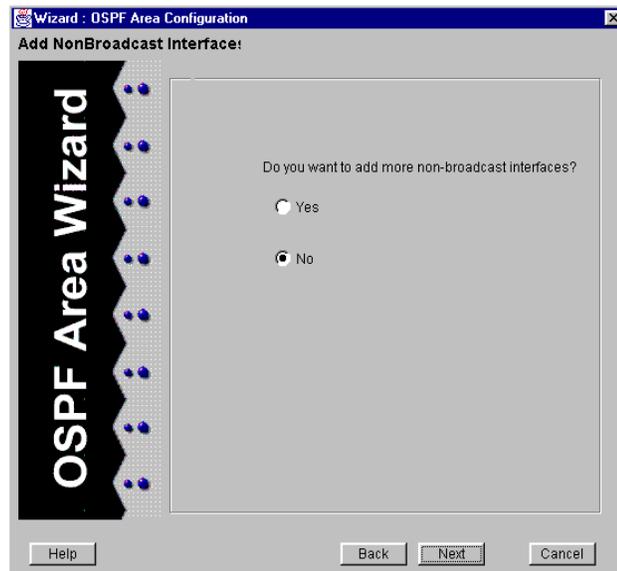


Figure 173. Add Non-Broadcast Interface panel

- j. Do one of the following:
 - If you select *Yes*, click **Next** and then add another non-broadcast interface specifying options as you did when you added the other non-broadcast interface. Then click **Next**.
 - If you select *No*, click **Next**.

Configuration Expert prompts you to specify whether you want to add a stub host.

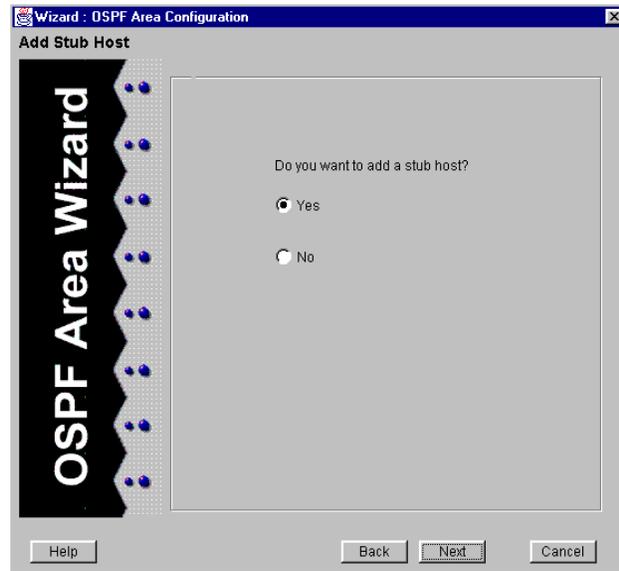


Figure 174. Add Stub Host panel

15. Do one of the following:
 - If you do not want to add a stub host, select *No*.
 - If you want to add a stub host, select *Yes* and then define the host in the OSPF Area Stub Host panel that appears.

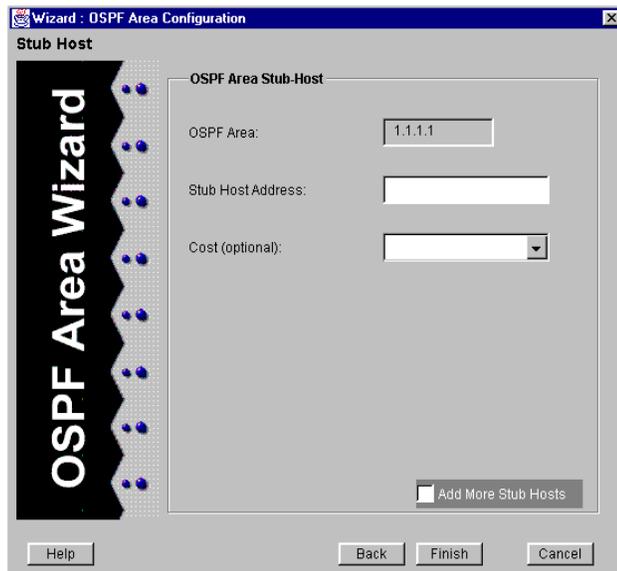


Figure 175. OSPF Area Stub Host panel

To add a stub host, take the following steps:

- a. Enter the address of the stub host.
 - b. Enter an integer value from 0 to 65535 to specify the cost that should be advertised for the directly attached stub host.
 - c. To add another stub host, select the *Add More Stub Hosts* check box, click **Next**, and then repeat [step a](#) and [step b](#) until you add all the appropriate stub hosts.
16. If you are configuring the backbone area and you need to define virtual links, do so by taking the following steps. Otherwise, skip to [step 17](#).

Note: In order to specify a Virtual Link, you must have previously configured at least one non-backbone area.

- a. Click **Next**.



Figure 176. Add Virtual Links panel

- b. Select *Yes* and then click *Next*.

A Virtual Link panel similar to the following appears:

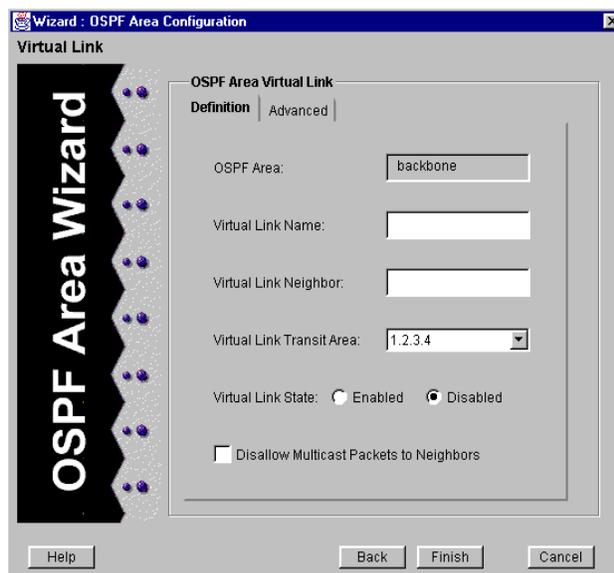


Figure 177. OSPF Area Virtual Link panel (Definition tab)

- c. Enter a virtual link name, then enter an IP address of an OSPF virtual link

- neighbor and the Area ID of the transit area in the appropriate text boxes.
- d. Enable or disable the interfaces
- e. Select or clear the *Disallow Multicast Packets to Neighbor* check box to specify whether the GSR is to send multicast packets to neighbors on point-to-point interfaces.
- f. Click the **Advanced** tab.

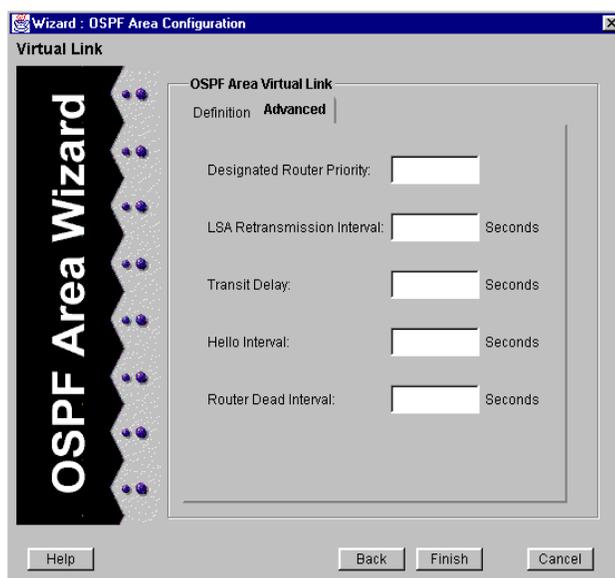


Figure 178. OSPF Area Virtual Link panel (Advanced tab)

- g. Set the options on the virtual link as discussed in the following table:

Table 31. Virtual link options

Option	Description
Designated Router Priority	<p>Enter a number between 0 and 255 to specify the priority for becoming the designated router on the virtual link.</p> <p>When two routers attached to a network both attempt to become the designated router, the one with the highest priority wins. A router with a router priority of 0 is ineligible to become a designated router.</p>

Table 31. Virtual link options (Continued)

Option	Description
LSA Retransmission Interval	Enter the number of seconds between link-state advertisement retransmissions for adjacencies belonging to the virtual link. Specify a number equal to or greater than 1.
Transit Delay	Enter the estimated number of seconds required to transmit a link-state update over the virtual link. Transit delay takes into account transmission and propagation delays and must be greater than 0. Specify a number equal to or greater than 1.
Hello Interval	Enter a number from 0 to 255 to specify the estimated number of seconds between hello packets that the router sends on the virtual link.
Router Dead Interval	Enter a number from 0 to 255 to specify the number of seconds that may occur without Hello packets being heard before the router's neighbors will declare it down.

- h. To add another link, select the *Add More Link* check box, click **Next**, and then continue repeating [step a](#) through [step g](#) until you add all the appropriate neighbors.

17. Click **Finish**.

Modifying Area Tables

To modify existing area tables:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object and then double-click the OSPF Area Table object.
5. Double-click the area table you want to modify.

Configuration Expert lists the objects of the table you are modifying.
6. Double-click the type of object (Networks Interfaces, and so on) you want to modify. Click the object you want to modify.
7. Edit the object's dialog box.

The options of the dialog box are the same as those you specified while creating the area table. For details on specifying these options, see [“Creating OSPF Area Tables” on page 232](#).

Chapter 15

Configuring BGP on the GSR

BGP Overview

The Border Gateway Protocol (BGP) is an exterior gateway protocol that allows IP routers to exchange network connectivity information. BGP became an internet standard in 1989 (RFC 1105) and the current version, BGP-4, was published in 1994 (RFC 1771). BGP is typically run between Internet Service Providers. It is also frequently used by multi-homed ISP customers, as well as large commercial networks.

Autonomous systems that wish to connect their networks together must agree on a method of exchanging routing information. Interior gateway protocols such as RIP and OSPF may be inadequate for this task since they are not designed to handle multi-AS and security issues. Similarly, using static routes may not be the best choice for exchanging AS-AS routing information because there may be a large number of routes, or the routes may change often.

Note: An *Autonomous System* (AS) is defined as a set of routers under a central technical administration that has a coherent interior routing plan and accurately portrays to other ASs what routing destinations are reachable by way of it.

In an environment where security is a concern and where using static routes is not feasible, BGP is often the best choice for an AS-AS routing protocol. BGP prevents the introduction of routing loops created by multi-homed and meshed AS topologies. BGP also provides the ability to create and enforce policies at the AS level, such as selectively determining which AS routes are to be accepted or what routes are to be advertised to BGP peers.

Configuring your GSR for BGP

The following sections describe how to use clearVISN CoreWatch to set up global BGP parameters on the GSR, configure new BGP peer groups for the GSR, and modify existing peer groups on the GSR.

Setting BGP Global Parameters

You can use clearVISN CoreWatch to set global BGP parameters to start or stop BGP on the GSR and specify BGP preference and metric values, as well as the cluster identification.

To set the GSR's global BGP parameters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the BGP object and click the BGP Global Parameters object.

A **BGP Global Parameters** dialog box similar to the following appears:

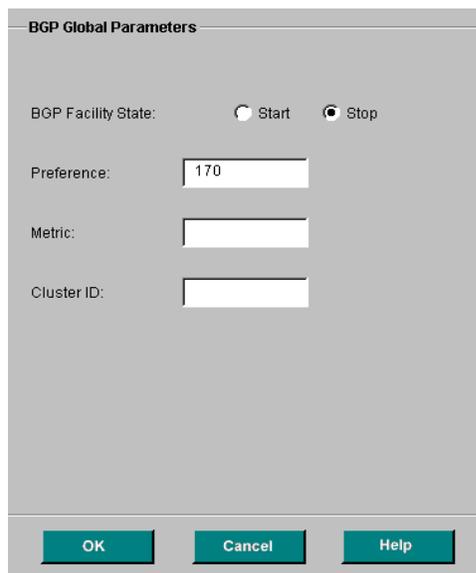


Figure 179. BGP Global Parameters dialog box

6. Set the *BGP Facility State* option to *Start* or *Stop*.
7. Enter global values for the following options:

Table 32. BGP global options

Option	Description
Preference	Defines the preference value for BGP hosts/networks. You can specify a value between 0 and 255, inclusive.
Metric	Defines the default-metric value for BGP hosts/networks. You can specify a value between 0 and 65535, inclusive.
Cluster ID	Specifies the IP address for BGP hosts/networks.

8. Click **OK**.

Configuring a New BGP Peer Group

You can configure a BGP peer group for the GSR.

To configure a BGP peer group:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the BGP Routing object.
6. Double-click the BGP Groups object and click the Create New BGP Peer Group object.

Configuration Expert opens the BGP wizard:

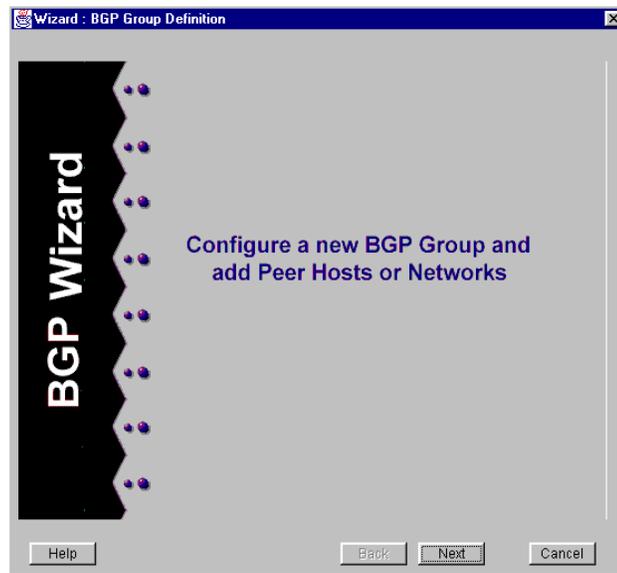


Figure 180. BGP wizard

7. Click Next.

The BGP Peer-Group Definition panel appears:

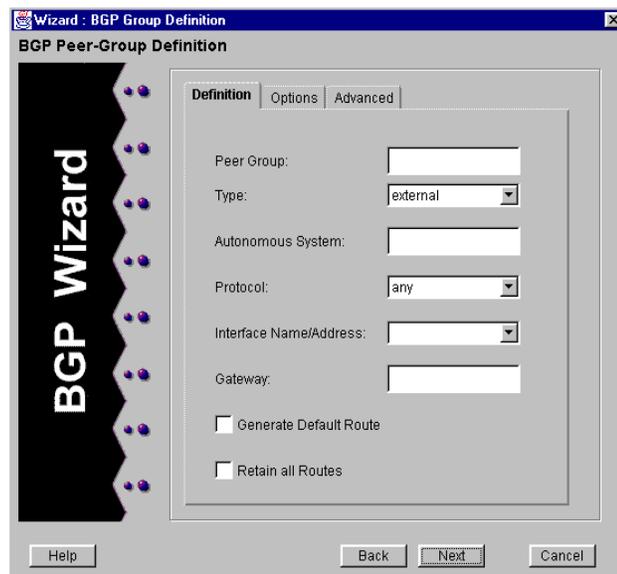


Figure 181. BGP Peer-Group Definition panel (Definition tab)

8. Specify the general BGP options for your GSR according to the following:
 - a. Specify an associated peer group for your GSR in the **Peer Group** box. You *must* specify a peer group to successfully configure BGP on your GSR.
 - b. Select a BGP type from the **Type** drop-down list.
 - c. Specify an autonomous system for your GSR in the **Autonomous System** box. This option specifies the autonomous system the router represents to this peer group.
 - d. Select the protocol you wish to use for BGP from the **Protocol** drop-down list. You can select *OSPF*, *RIP*, *Static*, or specify that the GSR is free to use any of these protocols by selecting the *any* option.
 - e. Select the interface name or address by selecting an option from the **Interface Name/Address** drop-down list.
 - f. Enter a BGP gateway for your GSR in the **Gateway** box. If a network is not part of the peer group, this option specifies a router on an associated network the next hop router for routes received from this neighbor.

Note: You must enter a host address on a locally associated network unless the host is using OSPF protocol.

- g. Select either or both of the *Generate Default Route* and *Retain All Routes* options by checking their respective boxes.

The *Generate Default Route* option specifies whether the router should generate a default route when BGP receives a valid update from its peer. The generation of default routes is enabled by default.

The *Retain All Routes* option is used to retain routes learned from a peer, even if the routes' autonomous system paths contain exported autonomous system numbers for this GSR.

9. Click the Options tab.

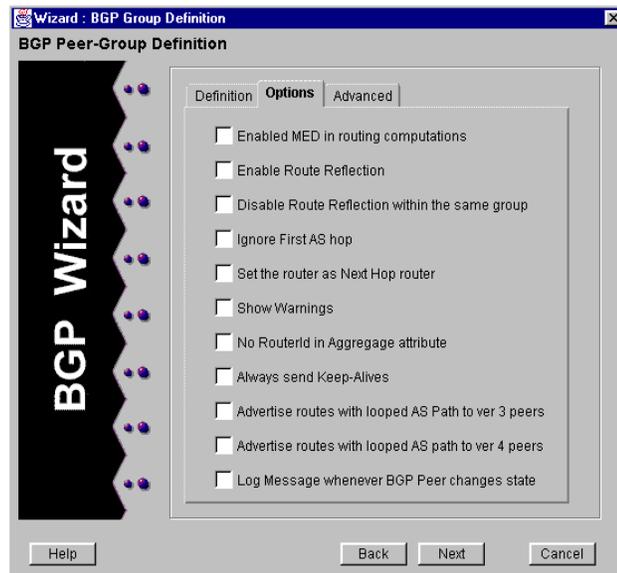


Figure 182. BGP Peer-Group Definition panel (Options tab)

10. If you wish, activate or deactivate one or more of the options on the Options tab by checking or unchecking their respective boxes. Checked options are activated; unchecked options remain inactive. Refer to the following table for a list of the available options and their descriptions:

Table 33. BGP Peer-Group options

Option	Description
Enable MED in Routing Computations	Specifies that the Multi_Exit_Disc (MED) metric is to be used in routing computations. By default, MEDs are not sent on external connections. To include MEDs for external connections, turn on this option, or activate the <i>Metric Out</i> option on the Advanced tab of the BGP Peer Group Definition panel or Add Host to Peer Group panel.
Enable Route Reflection	<p>Specifies that GateD will act as a route reflector for the peer-group.</p> <p>All routes received from any group member will be sent to all other internal neighbors, and all routes received from any other internal neighbors will be sent to the reflector clients.</p> <p>Note: This option is designed to be used with internal, routing, and IGP peer groups.</p>
Disable Route Reflection with the Same Group	<p>Specifies that routes received from reflector clients will only be sent to internal neighbors which are not in the same group as the reflector client that originated the route.</p> <p>Note: The router exports routes <i>from</i> the local autonomous system <i>back into</i> the local autonomous system when acting as a route reflector. All routes received from any group member will be sent to all other internal neighbors, and all routes received from any other internal neighbors will be sent to the reflector clients.</p>

Table 33. BGP Peer-Group options (Continued)

Option	Description
Ignore First AS Hop	<p>Specifying this option (here or on the Advanced tab of either the BGP Peer Group Definition panel or Add Host to Peer Group panel) tells GateD not to drop routes from Route Servers (servers that are able to propagate route without appending their own autonomous system to the path).</p> <p>Note: This option should only be used if it is positively known that the peer is a route server and not a normal router.</p>
Set the Router as the Next hop Router	<p>Specifies that the next hop in route advertisements for this peer or group of peers is to be this GSR, even if it would normally be possible to send a third-party next hop. This option may be useful in cases that deal with broken bridged interconnect media (in cases where the routers on the shared medium do not really have full connectivity to each other) or broken political situations.</p> <p>Note: This option may cause inefficient routes to be followed. DIGITAL recommends that you use this option only for external groups.</p>
Show Warnings	<p>This option instructs GateD to issue warning messages when receiving questionable BGP updates such as duplicate routes and/or non-existing route deletions. Normally, these events are silently ignored.</p>
No Router ID in Aggregate Attribute	<p>This option instructs GateD to specify the router identification in the aggregate attribute to be 0 (instead of its router identification). This action prevents different routers in an autonomous system from creating aggregate routes with different autonomous system paths.</p>
Always Send Keep-Alives	<p>This option instructs GateD to consistently send Keep-Alive messages, even when an update might have been a viable substitute. This allows interoperability with routers that do not completely obey the protocol specifications on this point.</p>

Table 33. BGP Peer-Group options (Continued)

Option	Description
Advertise Routes with Looped AS Path to ver 3 Peers	This option instructs GateD to advertise routes whose autonomous system paths are looped (i.e. with an autonomous system appearing more than once in the path) to version 3 external peers. Note: This option is ignored when set on internal groups or peers.
Advertise Routes with Looped AS Path to ver 4 Peers	Prevents routes with looped autonomous system paths from being advertised to version 4 external peers. This can be useful when trying to avoid advertising routes to peers that would automatically (and erroneously) forward routes on to version 3 neighbors.
Log Message Whenever BGP Peer Changes State	Instructs the GSR to log an informational message whenever an associated peer's state changes.

11. Click the Advanced tab:

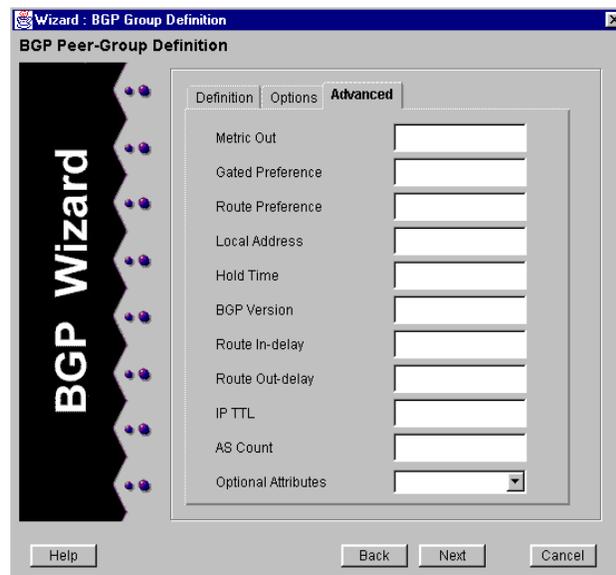


Figure 183. BGP Peer-Group Definition panel (Advanced tab)

12. If you wish, specify appropriate values for the options on the Advanced tab according to the guidelines in the following table:

Table 34. BGP Peer-Group advanced options

Option	Description
Metric Out	<p>Specifies the primary metric used on all routes sent to the specified peer group. The metric hierarchy is as follows (from highest to lowest priority/preference):</p> <ul style="list-style-type: none"> • The metric specified by export policy • Peer-level metric out • Group-level metric out • Default metric
GateD Preference	<p>Allows GateD to use the LOCAL_PREF attribute to set preference on reception, and allows the GateD preference to set the LOCAL_PREF on transmission. The set-pref metric works as a lower limit, below which the imported LOCAL_PREF may not set the GateD preference.</p> <p>Note: This option is designed to be used with internal, routing, and IGP peer groups.</p>
Route Preference	<p>Specifies the preference used for routes learned from these peers. This can differ from the default BGP preference set in the BGP set preference statement, so that GateD can prefer routes from one peer, or group of peer, over others. This preference may be explicitly overridden by import policy.</p>
Local Address	<p>Specifies the address to be used on the local end of the TCP connection with the peer or with the peer's gateway when the gateway option is used. A session with an external peer will only be opened when an interface with the appropriate local address (through which the peer or gateway address is directly reachable). In either case incoming connections will only be recognized as matching a configured peer if they are addressed to the configured local address.</p> <p>Note: This option is designed to be used with internal, routing, and IGP peer groups and should be a valid, active interface address.</p>

Table 34. BGP Peer-Group advanced options (Continued)

Option	Description
Hold Time	<p>Specifies the Hold Time value (in seconds) when negotiating peer connections. If BGP does not receive a Keep-Alive, update, or notification message from a peer within the Hold Time period specified, then the BGP connection will be closed.</p> <p>The value must either be set to 0 (no Keep-Alive messages will be sent) or a minimum value of 6.</p>
BGP Version	<p>Specifies the BGP protocol version to use. If a BGP version is specified, only the specified version will be offered during negotiation. If no version is specified, version negotiation proceeds using the highest supported version first.</p> <p>Currently, ver 2, ver 3 and ver 4 are supported.</p>
Route In-Delay	<p>Used to dampen route fluctuations. The In-Delay option specifies the amount of time (in seconds) a route learned from a BGP peer must remain stable before it will be accepted into the routing database. The default value for this option is 0, or disabled.</p>
Route Out-Delay	<p>Used to dampen route fluctuations. The Out-Delay option is the amount of time (in seconds) a route must remain included in the routing table before it is exported to BGP. The default value for this option is 0, or disabled.</p>
IP TTL	<p>By default, BGP sets the IP TTL for local peers to 1 and the TTL for non-local peers to 255. This option is provided when attempting to communicate with improperly functioning routers that ignore packets sent with a TTL of 1.</p>

Table 34. BGP Peer-Group advanced options (Continued)

Option	Description
AS Count	<p>This option determines how many times GateD will insert our own autonomous system number when we send the autonomous system path to an external neighbor.</p> <p>The default value for this option is 1. Higher values are typically used to bias upstream neighbors' route selection.</p> <p>Note: This option supersedes the <i>Advertise Routes with Looped AS Path to ver 4 Peers</i> option. Regardless of whether or not the <i>Advertise Routes with Looped AS Path to ver 4 Peers</i> option is turned on, the GSR will still send multiple copies of its own autonomous system if the this option is set to something greater than 1.</p> <p>In addition, if the value of the <i>AS Count</i> option changes and GateD is reconfigured, routes will not be sent to reflect the new setting. If this is the desired behavior, you will need to restart the peer session.</p> <p>DIGITAL recommends that you use this option only for external groups.</p>
Optional Attributes	<p>Specifies the identifier of the optional-attributes-list to be associated with this peer group.</p> <p>You can configure optional attributes using the Optional Path Attribute List dialog box as described in “Configuring and Modifying Optional Attribute Building Blocks” on page 316.</p>

13. Click **Next**.

Configuration Expert displays the Add Host or Networks panel:

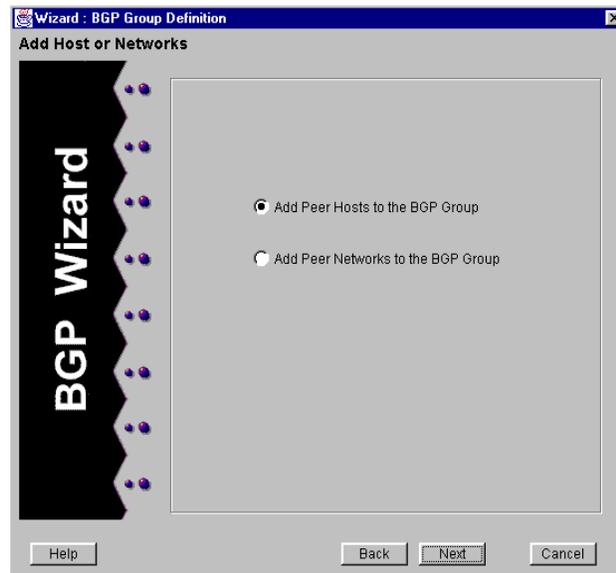


Figure 184. Add Host or Network panel

14. Select either the *Add Peer Hosts to the BGP Group* or *Add Peer Networks to the BGP Group* option and click **Next**.
15. If you specified that you want to add peer hosts to the BGP, use the Add Host to BGP Group panel that appears to add a peer host to your BGP. Otherwise, skip to [step 17](#).

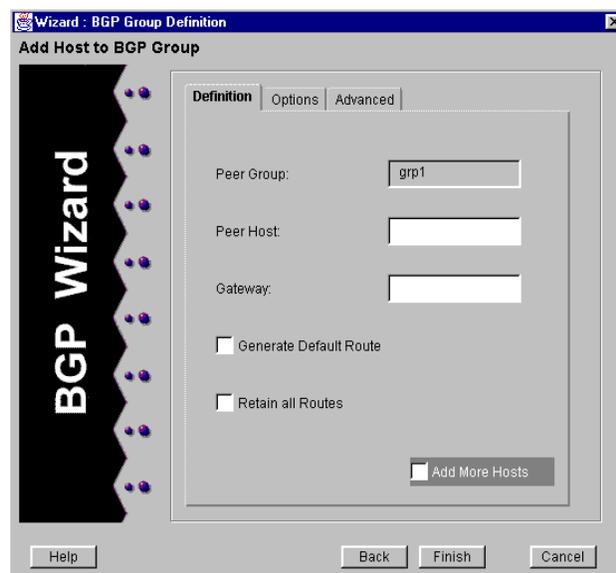


Figure 185. Add Host to BGP Group panel (Definition tab)

16. To add peer hosts to the BGP, go through the following steps:
 - a. Specify the address of the peer host you wish to add to the BGP in the **Peer Host** box.
 - b. Enter a BGP gateway for your GSR in the **Gateway** box.
 - c. Select either or both of the *Generate Default Route* and *Retain All Routes* options by checking their respective boxes.

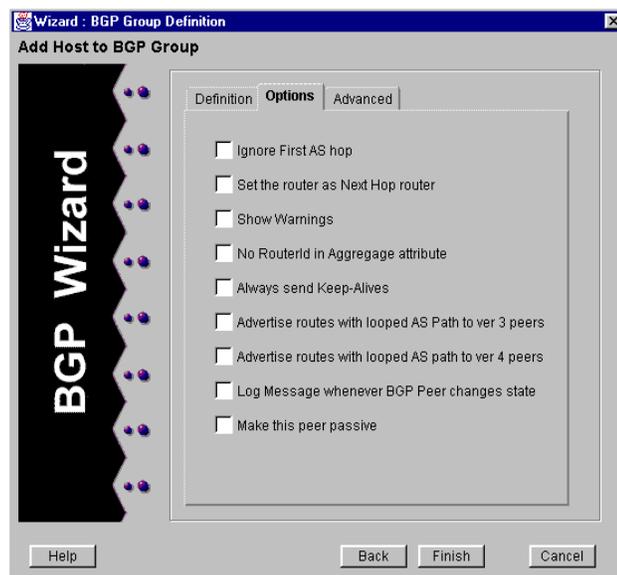


Figure 186. Add Host to BGP Group panel (Options tab)

- d. If you wish, activate or deactivate one or more of the options on the Options tab by checking or unchecking their respective boxes. Checked options are activated; unchecked options remain inactive. Refer to the following table for a list of the

available options and their descriptions:

Table 35. Add Host to BGP Group options

Option	Description
Ignore First AS Hop	<p>Specifying this option (here or on the Advanced tab of either the BGP Peer Group Definition panel or Add Host to Peer Group panel) tells GateD not to drop routes from Route Servers (servers that are able to propagate route without appending their own autonomous system to the path).</p> <p>Note: This option should only be used if it is positively known that the peer is a route server and not a normal router.</p>
Set the Router as the Next hop Router	<p>Specifies that the next hop in route advertisements for this peer or group of peers is to be this GSR, even if it would normally be possible to send a third-party next hop. This option may be useful in cases that deal with broken bridged interconnect media (in cases where the routers on the shared medium do not really have full connectivity to each other) or broken political situations.</p> <p>Note: This option may cause inefficient routes to be followed. DIGITAL recommends that you use this option only for external groups.</p>
Show Warnings	<p>This option instructs GateD to issue warning messages when receiving questionable BGP updates such as duplicate routes and/or non-existing route deletions. Normally, these events are silently ignored.</p>
No Router ID in Aggregate Attribute	<p>This option instructs GateD to specify the router identification in the aggregate attribute to be 0 (instead of its router identification). This action prevents different routers in an autonomous system from creating aggregate routes with different autonomous system paths.</p>
Always Send Keep-Alives	<p>This option instructs GateD to consistently send Keep-Alive messages, even when an update might have been a viable substitute. This allows interoperability with routers that do not completely obey the protocol specifications on this point.</p>

Table 35. Add Host to BGP Group options (Continued)

Option	Description
Advertise Routes with Looped AS Path to ver 3 Peers	<p>This option instructs GateD to advertise routes whose autonomous system paths are looped (i.e. with an autonomous system appearing more than once in the path) to version 3 external peers.</p> <p>Note: This option is ignored when set on internal groups or peers.</p>
Advertise Routes with Looped AS Path to ver 4 Peers	<p>Prevents routes with looped autonomous system paths from being advertised to version 4 external peers. This can be useful when trying to avoid advertising routes to peers that would automatically (and erroneously) forward routes on to version 3 neighbors.</p>
Log Message Whenever BGP Peer Changes State	<p>Instructs the GSR to log an informational message whenever an associated peer's state changes.</p>
Make this Peer Passive	<p>Specifies that active OPENs to this peer should not be attempted. BGP would wait for the peer to issue an OPEN. By default all explicitly configured peers (the one's configured using the add peer_host command) are active, they periodically send OPEN messages until the peer responds. Note that if it is applied to both sides of a peering session, it will prevent the session from ever being established.</p>

- e. Click the Advanced tab:

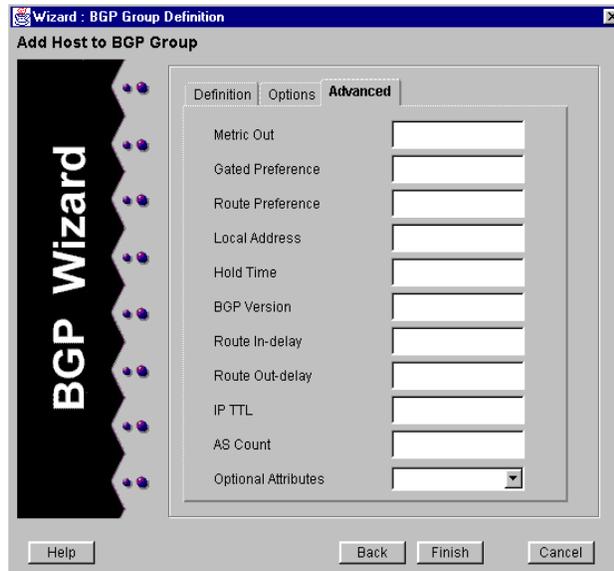


Figure 187. Add Host to BGP Group panel (Advanced tab)

- f. If you wish, specify appropriate values for the options on the Advanced tab according to the guidelines in the following table:

Table 36. Add Host to BGP Group advanced options

Option	Description
Metric Out	<p>Specifies the primary metric used on all routes sent to the specified peer group. The metric hierarchy is as follows (from highest to lowest priority/preference):</p> <ul style="list-style-type: none"> • The metric specified by export policy • Peer-level metric out • Group-level metric out • Default metric

Table 36. Add Host to BGP Group advanced options (Continued)

Option	Description
Gated Preference	<p>Allows GateD to use the LOCAL_PREF attribute to set preference on reception, and allows the GateD preference to set the LOCAL_PREF on transmission. The set-pref metric works as a lower limit, below which the imported LOCAL_PREF may not set the GateD preference.</p> <p>Note: This option is designed to be used with internal, routing, and IGP peer groups.</p>
Route Preference	<p>Specifies the preference used for routes learned from these peers. This can differ from the default BGP preference set in the BGP set preference statement, so that GateD can prefer routes from one peer, or group of peer, over others. This preference may be explicitly overridden by import policy.</p>
Local Address	<p>Specifies the address to be used on the local end of the TCP connection with the peer or with the peer's gateway when the gateway option is used. A session with an external peer will only be opened when an interface with the appropriate local address (through which the peer or gateway address is directly reachable). In either case incoming connections will only be recognized as matching a configured peer if they are addressed to the configured local address.</p> <p>Note: This option is designed to be used with internal, routing, and IGP peer groups and should be a valid, active interface address.</p>
Hold Time	<p>Specifies the Hold Time value (in seconds) when negotiating peer connections. If BGP does not receive a Keep-Alive, update, or notification message from a peer within the Hold Time period specified, then the BGP connection will be closed.</p> <p>The value must either be set to 0 (no Keep-Alive messages will be sent) or a minimum value of 6.</p>
BGP Version	<p>Specifies the BGP protocol version to use. If a BGP version is specified, only the specified version will be offered during negotiation. If no version is specified, version negotiation proceeds using the highest supported version first.</p> <p>Currently, ver 2, ver 3 and ver 4 are supported.</p>
Route In-Delay	<p>Used to dampen route fluctuations. The In-Delay option specifies the amount of time (in seconds) a route learned from a BGP peer must remain stable before it will be accepted into the routing database. The default value for this option is 0, or disabled.</p>

Table 36. Add Host to BGP Group advanced options (Continued)

Option	Description
Route Out-Delay	Used to dampen route fluctuations. The Out-Delay option is the amount of time (in seconds) a route must remain included in the routing table before it is exported to BGP. The default value for this option is 0, or disabled.
IP TTL	By default, BGP sets the IP TTL for local peers to 1 and the TTL for non-local peers to 255. This option is provided when attempting to communicate with improperly functioning routers that ignore packets sent with a TTL of 1.
AS Count	<p>This option determines how many times GateD will insert our own autonomous system number when we send the autonomous system path to an external neighbor.</p> <p>The default value for this option is 1. Higher values are typically used to bias upstream neighbors' route selection.</p> <p>Note: This option supersedes the <i>Advertise Routes with Looped AS Path to ver 4 Peers</i> option. Regardless of whether or not the <i>Advertise Routes with Looped AS Path to ver 4 Peers</i> option is turned on, the GSR will still send multiple copies of its own autonomous system if the this option is set to something greater than 1.</p> <p>In addition, if the value of the <i>AS Count</i> option changes and GateD is reconfigured, routes will not be sent to reflect the new setting. If this is the desired behavior, you will need to restart the peer session.</p> <p>DIGITAL recommends that you use this option only for external groups.</p>
Optional Attributes	<p>Specifies the identifier of the optional-attributes-list to be associated with this peer group.</p> <p>You can configure optional attributes using the Optional Path Attribute List dialog box as described in “Configuring and Modifying Optional Attribute Building Blocks” on page 316.</p>

- g. Click the Definitions tab.
- h. If you wish to add another peer host to the BGP, check the *Add More Hosts* option, click **Next**, and repeat [step a](#) through [step f](#) until you have specified all of the hosts you wish to add to the BGP. Otherwise, simply click **Next** to continue with the configuration process.

- Use the Add Networks to BGP panel that appears to add peer networks to the BGP:

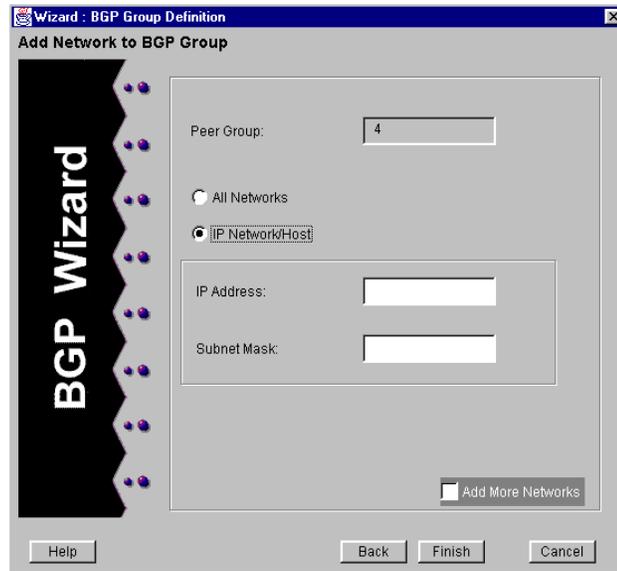


Figure 188. Add Network to BGP Group panel

- Specify whether to add all associated networks or a specific network to the BGP by selecting either the *All Networks* or *IP Network/Host* option.
- If you selected the *IP Network/Host* option in [step 18](#), enter the IP address and subnet mask for the network interface in the appropriate text boxes.
- If you wish to add another individual network to the BGP, check the *Add More Networks* option, click **Next**, and repeat [step 18](#) and [step 19](#) until you have specified all of the networks you wish to add to the BGP.
- Click **Finish**.

The BGP wizard closes and an object representing your newly-defined network information appears in the Configuration Expert configuration tree under BGP Groups.

Modifying an Existing BGP Peer Group

There will probably come a time when you find it necessary to update your peer group information to include new or relocated associated hosts or networks, and/or delete outdated hosts or networks.

To modify an existing peer group:

- Start Configuration Expert if you have not already done so.

2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the BGP Routing object.
6. Double-click the BGP Groups object and click on the object you wish to modify.

A **BGP Peer Group Definition** dialog box similar to the following appears:

The screenshot shows the 'BGP Peer Group Definition' dialog box with the 'Definition' tab selected. The fields are as follows:

- Peer Group: grp2
- Type: external
- Autonomous System: (empty)
- Protocol: any
- Interface Name/Address: (empty)
- Gateway: (empty)
- Generate Default Route
- Retain all Routes

Figure 189. BGP Peer Group Definition dialog box (Definition tab)

7. Specify the general BGP options you wish to change by going through the following procedure:
 - a. Select a BGP type from the **Type** drop-down list.
 - b. Specify an autonomous system for your GSR in the **Autonomous System** box.
 - c. Specify the protocol you wish to use for BGP from the **Protocol** drop-down list. You can select *OSPF*, *RIP*, *Static*, or specify that the GSR is free to use any of these protocols by selecting the *any* option.
 - d. Specify the interface name or address by selecting an option from the **Interface Name/Address** drop-down list.
 - e. Enter a BGP gateway for your GSR in the **Gateway** box.

- f. Select either or both of the *Generate Default Route* and *Retain All Routes* options by checking their respective boxes.
8. Click the Options tab.

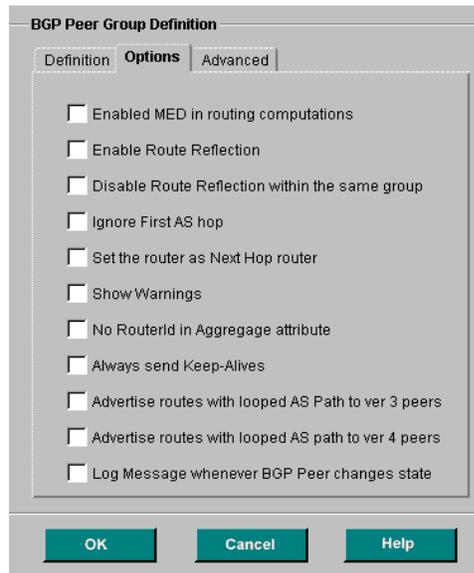


Figure 190. BGP Peer Group Definition dialog box (Options tab)

9. If you wish, activate or deactivate one or more of the options on the Options tab by checking or unchecking their respective boxes. Checked options are activated;

unchecked options remain inactive. Refer to the following table for a list of the available options and their descriptions:

Table 37. BGP Peer-Group options

Option	Description
Enable MED in Routing Computations	<p>Specifies that the Multi_Exit_Disc (MED) metric is to be used in routing computations. By default, MEDs are not sent on external connections. To include MEDs for external connections, turn on this option, or activate the <i>Metric Out</i> option on the Advanced tab of the BGP Peer Group Definition panel or Add Host to Peer Group panel.</p>
Enable Route Reflection	<p>Specifies that GateD will act as a route reflector for the peer-group.</p> <p>All routes received from any group member will be sent to all other internal neighbors, and all routes received from any other internal neighbors will be sent to the reflector clients.</p> <p>Note: This option is designed to be used with internal, routing, and IGP peer groups.</p>
Disable Route Reflection with the Same Group	<p>Specifies that routes received from reflector clients will only be sent to internal neighbors which are not in the same group as the reflector client that originated the route.</p> <p>Note: The router exports routes <i>from</i> the local autonomous system <i>back into</i> the local autonomous system when acting as a route reflector. All routes received from any group member will be sent to all other internal neighbors, and all routes received from any other internal neighbors will be sent to the reflector clients.</p>

Table 37. BGP Peer-Group options (Continued)

Option	Description
Ignore First AS Hop	<p>Specifies that the next hop in route advertisements for this peer or group of peers is to be this GSR, even if it would normally be possible to send a third-party next hop. This option may be useful in cases that deal with broken bridged interconnect media (in cases where the routers on the shared medium do not really have full connectivity to each other) or broken political situations.</p> <p>Note: This option may cause inefficient routes to be followed. DIGITAL recommends that you use this option only for external groups.</p>
Set the Router as the Next hop Router	<p>This option instructs GateD to issue warning messages when receiving questionable BGP updates such as duplicate routes and/or non-existing route deletions. Normally, these events are silently ignored.</p>
Show Warnings	<p>This option instructs GateD to specify the router identification in the aggregate attribute to be 0 (instead of its router identification). This action prevents different routers in an autonomous system from creating aggregate routes with different autonomous system paths.</p>
No Router ID in Aggregate Attribute	<p>This option instructs GateD to consistently send Keep-Alive messages, even when an update might have been a viable substitute. This allows interoperability with routers that do not completely obey the protocol specifications on this point.</p>
Always Send Keep-Alives	<p>This option instructs GateD to advertise routes whose autonomous system paths are looped (i.e. with an autonomous system appearing more than once in the path) to version 3 external peers.</p> <p>Note: This option is ignored when set on internal groups or peers.</p>

Table 37. BGP Peer-Group options (Continued)

Option	Description
Advertise Routes with Looped AS Path to ver 3 Peers	Prevents routes with looped autonomous system paths from being advertised to version 4 external peers. This can be useful when trying to avoid advertising routes to peers that would automatically (and erroneously) forward routes on to version 3 neighbors.
Advertise Routes with Looped AS Path to ver 4 Peers	Instructs the GSR to log an informational message whenever an associated peer's state changes.
Log Message Whenever BGP Peer Changes State	<p>Specifies that the next hop in route advertisements for this peer or group of peers is to be this GSR, even if it would normally be possible to send a third-party next hop. This option may be useful in cases that deal with broken bridged interconnect media (in cases where the routers on the shared medium do not really have full connectivity to each other) or broken political situations.</p> <p>Note: This option may cause inefficient routes to be followed. DIGITAL recommends that you use this option only for external groups.</p>

10. Click the Advanced tab.

The screenshot shows the 'BGP Peer Group Definition' dialog box with the 'Advanced' tab selected. The dialog contains the following fields and controls:

- Metric Out:
- Gated Preference:
- Route Preference:
- Local Address:
- Hold Time:
- BGP Version:
- Route In-delay:
- Route Out-delay:
- IP TTL:
- AS Count:
- Optional Attributes: ▼

At the bottom of the dialog are three buttons: **OK**, **Cancel**, and **Help**.

Figure 191. BGP Peer Group Definition dialog box (Advanced tab)

11. If you wish, specify appropriate values for the options on the Advanced tab according to the guidelines in the following table:

Table 38. BGP Peer-Group advanced options

Option	Description
Metric Out	<p>Specifies the primary metric used on all routes sent to the specified peer group. The metric hierarchy is as follows (from highest to lowest priority/preference):</p> <ul style="list-style-type: none"> • The metric specified by export policy • Peer-level metric out • Group-level metric out • Default metric
Gated Preference	<p>Allows GateD to use the LOCAL_PREF attribute to set preference on reception, and allows the GateD preference to set the LOCAL_PREF on transmission. The set-pref metric works as a lower limit, below which the imported LOCAL_PREF may not set the GateD preference.</p> <p>Note: This option is designed to be used with internal, routing, and IGP peer groups.</p>
Route Preference	<p>Specifies the preference used for routes learned from these peers. This can differ from the default BGP preference set in the BGP set preference statement, so that GateD can prefer routes from one peer, or group of peer, over others. This preference may be explicitly overridden by import policy.</p>
Local Address	<p>Specifies the address to be used on the local end of the TCP connection with the peer or with the peer's gateway when the gateway option is used. A session with an external peer will only be opened when an interface with the appropriate local address (through which the peer or gateway address is directly reachable). In either case incoming connections will only be recognized as matching a configured peer if they are addressed to the configured local address.</p> <p>Note: This option is designed to be used with internal, routing, and IGP peer groups and should be a valid, active interface address.</p>

Table 38. BGP Peer-Group advanced options (Continued)

Option	Description
Hold Time	<p>Specifies the Hold Time value (in seconds) when negotiating peer connections. If BGP does not receive a Keep-Alive, update, or notification message from a peer within the Hold Time period specified, then the BGP connection will be closed.</p> <p>The value must either be set to 0 (no Keep-Alive messages will be sent) or a minimum value of 6.</p>
BGP Version	<p>Specifies the BGP protocol version to use. If a BGP version is specified, only the specified version will be offered during negotiation. If no version is specified, version negotiation proceeds using the highest supported version first.</p> <p>Currently, ver 2, ver 3 and ver 4 are supported.</p>
Route In-Delay	<p>Used to dampen route fluctuations. The In-Delay option specifies the amount of time (in seconds) a route learned from a BGP peer must remain stable before it will be accepted into the routing database. The default value for this option is 0, or disabled.</p>
Route Out-Delay	<p>Used to dampen route fluctuations. The Out-Delay option is the amount of time (in seconds) a route must remain included in the routing table before it is exported to BGP. The default value for this option is 0, or disabled.</p>

Table 38. BGP Peer-Group advanced options (Continued)

Option	Description
IP TTL	By default, BGP sets the IP TTL for local peers to 1 and the TTL for non-local peers to 255. This option is provided when attempting to communicate with improperly functioning routers that ignore packets sent with a TTL of 1.
AS Count	<p>This option determines how many times GateD will insert our own autonomous system number when we send the autonomous system path to an external neighbor.</p> <p>The default value for this option is 1. Higher values are typically used to bias upstream neighbors' route selection.</p> <p>Note: This option supersedes the <i>Advertise Routes with Looped AS Path to ver 4 Peers</i> option. Regardless of whether or not the <i>Advertise Routes with Looped AS Path to ver 4 Peers</i> option is turned on, the GSR will still send multiple copies of its own autonomous system if the this option is set to something greater than 1.</p> <p>In addition, if the value of the <i>AS Count</i> option changes and GateD is reconfigured, routes will not be sent to reflect the new setting. If this is the desired behavior, you will need to restart the peer session.</p> <p>DIGITAL recommends that you use this option only for external groups.</p>
Optional Attributes	<p>Specifies the identifier of the optional-attributes-list to be associated with this peer group.</p> <p>You can configure optional attributes using the Optional Path Attribute List dialog box as described in “Configuring and Modifying Optional Attribute Building Blocks” on page 316.</p>

12. Click **OK**. Configuration Expert updates the peer group information for your GSR in the Configuration Tree.

Chapter 16

Configuring Routing Policies on the GSR

Routing Policies on the GSR

You can use clearVISN CoreWatch to configure the following types of routing policies on the GSR:

- Export Policies
- Import Policies
- Aggregate Policies
- Redistribute Policies
- Summarize Routes

This chapter discusses the necessary procedures involved in both configuring and modifying these routing policies for the GSR.

RIP and OSPF Routing Policy Defaults

Before setting out to configure any definitive routing policies, you can first use Configuration Expert to configure default (or global) routing policy attributes for all RIP and OSPF routes on the GSR. For example, you can define such attributes as route metrics, route preferences, and/or interface types for RIP and OSPF routes. Defining default characteristics for these RIP and OSPF routes gives you an advantage when it comes time

to configure particular types of routing policies, in that you can simply direct the GSR to use the default attributes you will have already configured.

Setting RIP Routing Policy Defaults

To set default values for RIP's metric and preference that will be used for RIP across all routing policies:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Global Default Parameters object and click the RIP Defaults object.

A **RIP Policy Defaults** dialog box similar to the following appears:

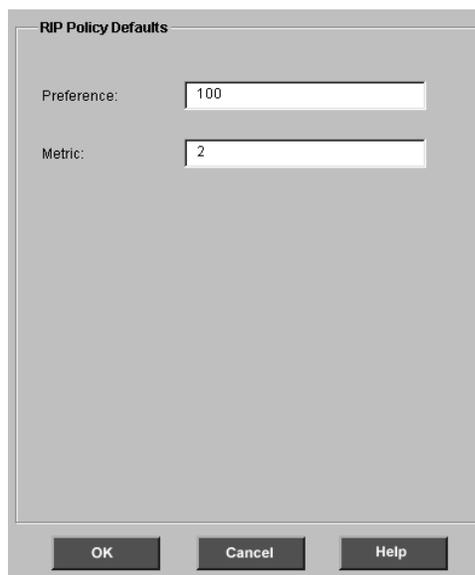


Figure 192. RIP Policy Defaults dialog box

Note: You can specify values for the default route preference and default metric in the **RIP Policy Defaults** dialog box just as in the **RIP Global Parameters** dialog box. If you change a parameter in one dialog box, that change is automatically applied to the other dialog box. See [“Setting RIP Global Parameters” on page 121](#) for more information.

7. Assign a preference for the routes learned by RIP. To do so, enter a number from 0 to 255 in the **Default Route Preference** box.

The preference you specify applies to all IP RIP interfaces on the GSR. The default value for route preference is 100.

8. Set the metric for routes advertised through RIP by entering a number from 1 to 16 in the **Default Metric** box.

Note: The metric 16 (the default) is equivalent in RIP to infinite and makes a route unreachable. You must set the default metric to a value other than 16 in order to allow the GSR to export routes from other protocols such as OSPF into RIP.

9. Click **OK**.

Setting OSPF Routing Policy Defaults

To set default values for OSPF's metric and preference that will be used for OSPF across all routing policies:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Global Default Parameters object and click the OSPF Defaults object.

A **OSPF Policy ASE Defaults** dialog box similar to the following appears:

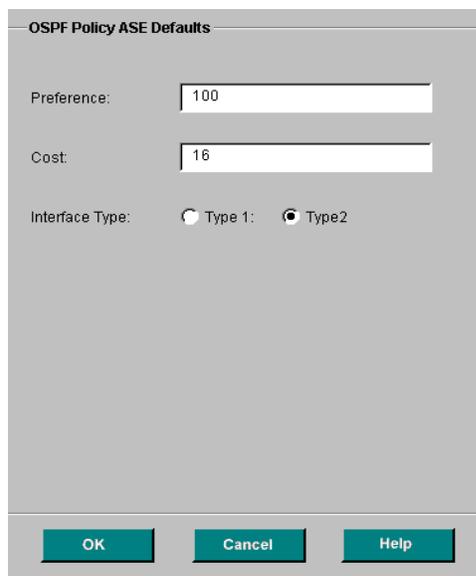


Figure 193. OSPF Policy ASE Defaults dialog box

7. Assign a preference for the routes learned by OSPF. To do so, enter a number from 0 to 255 in the **Preference** box.

The preference you specify applies to all OSPF interfaces on the GSR. The default value for route preference is 100.
8. Set the cost for routes advertised through OSPF by entering a number from 0 to 65535 in the **Cost** box.
9. Specify the OSPF interface type by selecting the appropriate option (either *Type 1* or *Type 2*) under *Interface Type*.
10. Click **OK**.

A Look at the Building Blocks of Routing Policies

Configuration Expert's building blocks form the basis of the routing policies you create.

You will need to create building blocks depending upon your network environment, the routing operations you want the GSR to perform on that network, and the routing policies you wish to create. You can define the following types of building blocks and then create a routing policy by including one or more of these building blocks in the policy:

- Export destinations and export sources, which you will include in your export policies. A destination can be a particular interface or gateway. A source can be a particular interface, gateway, autonomous system, or aggregate route.
- Import sources that you will include in your import policies.
- Filters, which are used to match a certain set of routes by destination or by destination and mask.
- Aggregate destinations and aggregate sources for aggregate routes.
- Optional Path Attribute Lists to include with your export destinations and BGP input sources.

Export Destination Building Blocks

A GSR's export routing policy can include multiple export destination building blocks that specify an export destination for RIP, OSPF, or BGP routes. Each routing policy can include export destinations for RIP, OSPF, BGP or any combination of the three.

You can use Configuration Expert to create RIP, OSPF, and BGP export destination building blocks. The method you use to define these building blocks depends on which protocol the building block is for. Separate discussions on defining RIP, OSPF, and BGP export destination building blocks follow.

Configuring and Modifying RIP Export Destinations

A RIP export destination building block specifies an export destination for RIP routes.

To define RIP export destination building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Destination object.
8. Double-click the RIP object, and then do one of the following:
 - If you are creating a new export destination, click the Configure New RIP Export Destination object.
 - If you are modifying an existing export destination, select that export destination.

A **RIP Export Destination** dialog box similar to the following appears:

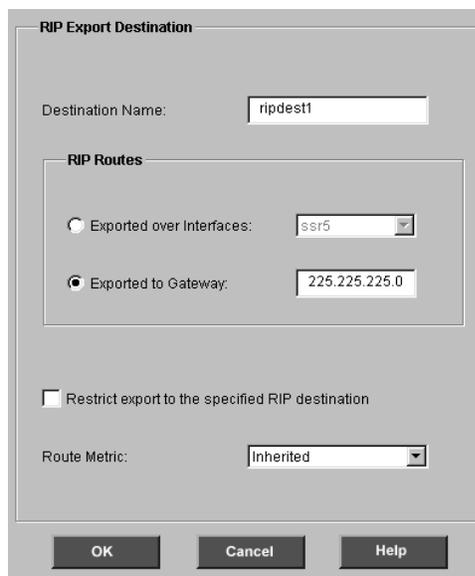


Figure 194. RIP Export Destination dialog box

9. Specify a name for the export destination in the **Destination Name** box.
10. Under *RIP Routes*, specify whether you want to export the routes to a specific interface or a gateway by doing one of the following:
 - To export routes to an interface, select the *Exported Over Interfaces* option and then select an IP interface from the associated drop-down list.
 - To export routes to a gateway, select the *Exported to Gateway* option and then specify IP address for the given gateway in the associated text box.
11. Specify whether the GSR restricts the export route or associates a metric with the exported route by doing one of the following:
 - If you do not want the GSR to export any routes to the destination, select the *Restrict export to the specified RIP destination* option.
 - If you want the GSR to inform the destination about how far away the GSR is from the export route, select the route metric to be associated with exported routes. Select a number from 1 to 16 from the **Default Metric** drop-down list or select *Inherited* if you want the GSR to use the default metric from the routing policy.
12. Click **OK**.
13. Repeat [step 8](#) through [step 12](#) until you create all the RIP export destinations you plan to include in your export policies.

Configuring and Modifying OSPF Export Destinations

An OSPF export destination building block specifies an export destination for OSPF routes.

Note: You can export OSPF routes only into OSPF ASE routes.

To define OSPF export destination building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Destination object.
8. Double-click the OSPF object, and then do one of the following:
 - If you are creating a new export destination, click the Configure New OSPF Export Destination object.
 - If you are modifying an existing export destination, select that export destination.

An **OSPF Export Destination** dialog box similar to the following appears:



The image shows a dialog box titled "OSPF Export Destination". It contains the following fields and controls:

- Destination Name:** A text input field.
- Tag:** A text input field.
- ASE Route Type:** Two radio buttons labeled "Type 1" and "Type 2". "Type 2" is selected.
- Restrict export to the specified OSPF destination:** An unchecked checkbox.
- Route Metric:** A dropdown menu with "Inherited" selected.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom.

Figure 195. OSPF Export Destination dialog box

9. Specify a name for the new export destination in the **Destination Name** box.
10. Specify the tag you wish to associate with OSPF routes in the **Tag** box.
11. Specify whether the routes are to be exported as type 1 or type 2 autonomous system external (ASE) routes by selecting the appropriate option.
12. Specify whether the GSR restricts the export route or associates a metric with the exported route by doing one of the following:
 - If you do not want the GSR to export any routes to the destination, select the *Restrict export to the specified OSPF destination* option.
 - If you want the GSR to inform the destination about how far away the GSR is from the export route, select the route metric to be associated with exported routes. Select a number from 1 to 16 from the **Default Metric** drop-down list or select *Inherited* if you want the GSR to use the default metric from the routing policy.
13. Click **OK**.
14. Repeat [step 8](#) through [step 13](#) until you create all the OSPF export destinations you plan to include in your export policies.

Configuring and Modifying BGP Export Destinations

A BGP export destination building block specifies an export destination for BGP routes.

To define BGP export destination building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Destination object.
8. Double-click the BGP object, and then do one of the following:
 - If you are creating a new export destination, click the Configure New BGP Export Destination object.
 - If you are modifying an existing export destination, select that export destination.

A **BGP Export Destination** dialog box similar to the following appears:

Figure 196. BGP Export Destination dialog box

9. Specify a name for the new export destination in the **Destination Name** box.
10. Specify the autonomous system for the new export destination in the **Autonomous System** box.
11. If you wish, specify an optional attribute for the export destination by selecting a pre-defined attribute from the **Optional Attributes** drop-down list. The available optional attributes in this drop-down list are a collection of pre-defined optional attributes you may have previously specified, as described in [“Configuring and Modifying Optional Attribute Building Blocks”](#) on page 316.
12. If you wish, specify a sequence number for the export destination in the **Sequence Number** box.
13. Specify whether the GSR restricts the export route or associates a metric with the exported route by doing one of the following:
 - If you do not want the GSR to export any routes to the destination, select the *Restrict export to the specified BGP destination* option.
 - If you want the GSR to inform the destination about how far away the GSR is from the export route, select *Inherited* from the **Default Metric** drop-down list to specify that you want the GSR to use the default metric from the routing policy.
14. Click **OK**.
15. Repeat [step 8](#) through [step 13](#) until you create all the BGP export destinations you plan to include in your export policies.

Aggregate Destination Building Blocks

Define an aggregate destination to create or modify an aggregate generation destination and control the export of routes from a source with routes to directly attached interfaces.

The procedure for defining an aggregate destination depends on whether you are creating a new destination or modifying an existing one. Separate discussions on these tasks follow.

Configuring Aggregate Destinations

To create aggregate destination building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Aggregate Destinations object and click the Configure New Aggregate Destination object.

Configuration Expert opens the Aggregate Destination wizard.

8. Click Next.

An Aggregate Destination panel similar to the following appears:

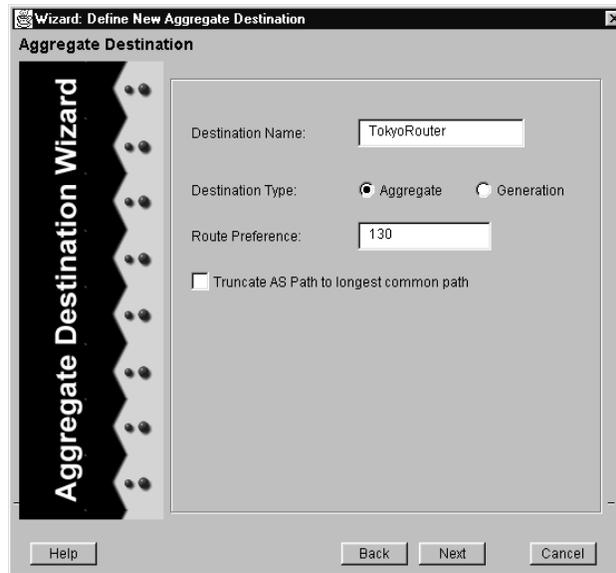


Figure 197. Aggregate Destination panel

9. Specify a name for the aggregate destination in the **Destination Name** box.
10. Under *Destination Type*, specify whether the destination is an aggregate or a generate by selecting the appropriate option.
11. In the **Route Preference** box, specify the preference you want to assign to the resulting aggregate route. You can specify a value between 0 and 64,000. The default value is 130.
12. If you wish, specify that the autonomous-system path should be truncated to the longest common autonomous-system path by turning on the *Truncate AS Path to longest common path* option. Otherwise, simply leave this option unchecked, specifying that the GSR is to build an autonomous-system path consisting of sets and sequences from all contributing autonomous-system paths.
13. Click **Next**.

An Aggregate Network panel similar to the following appears:

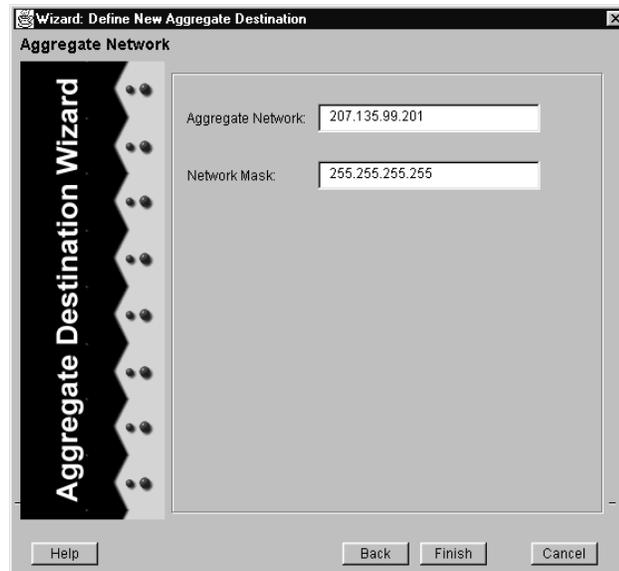


Figure 198. Aggregate Network panel

14. Specify the aggregate network by entering the appropriate IP address and subnetwork mask in the **Aggregate Network** and **Network Mask** boxes, respectively.
15. Click **Finish**.

Modifying Aggregate Destinations

You can modify aggregate destinations to change their name, type, route preference, or path name format.

To create aggregate destination building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Aggregate Destinations object and select the aggregate destination you want to modify.

8. Modify the destination by editing the fields of the **Aggregate Destination Definition** dialog box that appears.

The dialog box includes some of the same fields you specified when creating the destination. For more information on these fields, see [“Configuring Aggregate Destinations”](#) on page 288.

9. Click **OK**.

Export Source Building Blocks

You can define export source building blocks for the following types of interface:

- RIP
- OSPF
- BGP
- Autonomous System Path
- Tag
- Direct
- Static
- Aggregate

Once you have configured these export sources, you can use them to create or modify a number of routing policies using a variety of source protocols.

The procedure for defining an export source depends on whether you are configuring a new source or modifying an existing one. Separate discussions on these tasks for the various protocol types follow.

Configuring and Modifying RIP Export Sources

A RIP export source building block specifies a source for exporting RIP routes into other protocols.

To define RIP export source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.

5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Source object.
8. Double-click the RIP object, and then do one of the following:
 - If you are creating a new export source, click the Configure New RIP Export Source object.
 - If you are modifying an existing export source, select that export source.

A **RIP Export Source** dialog box similar to the following appears:

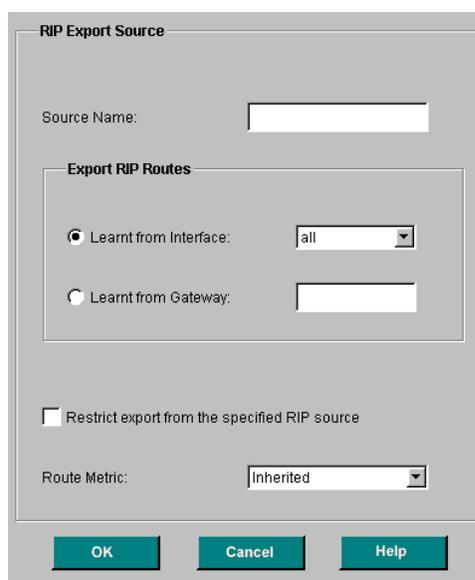


Figure 199. RIP Export Source dialog box

9. Specify the name you wish to assign to the RIP export source in the **Source Name** box.
10. Under *Export RIP Routes*, specify whether the GSR restricts the export or associates a metric with the exported route by doing one of the following:
 - If you do not want the GSR to export any routes from the source, select the *Restrict export from the specified RIP source* option.
 - If you want the GSR to associate a metric with the exported route, select a metric from the **Route Metric** drop-down list. Select *Inherited* if you want the route to use the policy's default metric. Otherwise, select a number from 1 to 16.

If you do not enter a metric value in this dialog box, the GSR inherits the metric value from the export destination. If a metric value is not specified for the export destination, the GSR inherits the default metric specified at the protocol level.

11. Click **OK**.
12. Repeat [step 8](#) through [step 11](#) until you create all of the RIP export sources you plan to include in your export policies.

Configuring and Modifying OSPF Export Sources

An OSPF export source building block specifies a source for exporting OSPF routes into other protocols.

To define OSPF export source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Source object.
8. Double-click the OSPF object, and then do one of the following:
 - If you are creating a new export source, click the Configure New OSPF Export Source object.
 - If you are modifying an existing export source, select that export source.

An **OSPF Export Source** dialog box similar to the following appears:

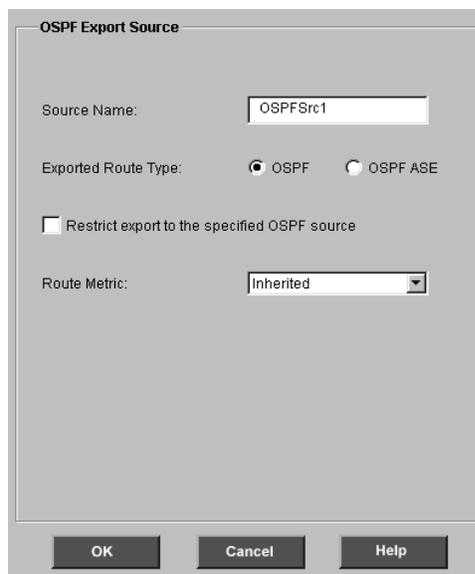


Figure 200. OSPF Export Source dialog box

9. Specify the name you wish to assign to the export source in the **Source Name** box.
10. Under *Exported Route Type*, specify whether the exported routes are OSPF routes or OSPF autonomous system external (ASE) routes by selecting the appropriate option.
11. Specify whether the GSR restricts the export or associates a metric with the exported route by doing one of the following:
 - If you do not want the GSR to export any routes from the source, select the *Restrict export to the specified OSPF source* option.
 - If you want the GSR to associate a metric with the exported route, select a metric from the **Route Metric** drop-down list. Select *Inherited* if you want the route to use the policy's default metric. Otherwise, select a number from 1 to 16.
12. Click **OK**.
13. Repeat [step 8](#) through [step 12](#) until you create all of the OSPF export sources you plan to include in your export policies.

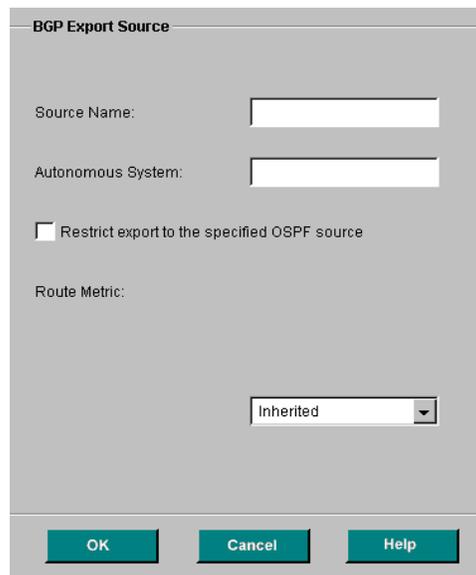
Configuring and Modifying BGP Export Sources

A BGP export source building block specifies a source for exporting BGP routes into other protocols.

To define BGP export source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Source object.
8. Double-click the BGP object, and then do one of the following:
 - If you are creating a new export source, click the Configure New BGP Export Source object.
 - If you are modifying an existing export source, select that export source.

A **BGP Export Source** dialog box similar to the following appears:



The image shows a dialog box titled "BGP Export Source". It contains the following fields and controls:

- Source Name:** A text input field.
- Autonomous System:** A text input field.
- Restrict export to the specified OSPF source**
- Route Metric:** A dropdown menu with "Inherited" selected.
- Buttons: **OK**, **Cancel**, and **Help**.

Figure 201. BGP Export Source dialog box

9. Specify the name you wish to assign to the BGP export source in the **Source Name** box.
10. Specify the autonomous system for the BGP export source in the **Autonomous System** box.

11. Specify whether the GSR restricts the export or associates a metric with the exported route by doing one of the following:
 - If you do not want the GSR to export any routes from the source, select the *Restrict export to the specified BGP source* option.
 - If you want the GSR to associate a metric with the exported route, select a metric from the **Route Metric** drop-down list. Select *Inherited* if you want the route to use the policy's default metric. Otherwise, select a number from 1 to 16.
12. Click **OK**.
13. Repeat [step 9](#) through [step 12](#) until you create all of the BGP export sources you plan to include in your export policies.

Configuring and Modifying Autonomous System Path Export Sources

An autonomous system path export source building block specifies a source for exporting AS path routes into other protocols.

To define AS path export source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Source object.
8. Double-click the Autonomous System Path object, and then do one of the following:
 - If you are creating a new export source, click the Configure New AS Export Source object.
 - If you are modifying an existing export source, select that export source.

An **Autonomous System Path Export Source** dialog box similar to the following appears:

Figure 202. Autonomous System Path Export Source dialog box

9. Specify the name you wish to assign to the AS path export source in the **Source Name** box.
10. Under *Source Definition*, specify the following source criteria:
 - a. Specify the autonomous system path regular expression in the **AS Path Regular Expression** box.
 - b. Specify a source protocol by selecting one of the following seven options from the **Protocol** drop-down list:
 - *all*
 - *static*
 - *direct*
 - *aggregate*
 - *rip*
 - *ospf*
 - *bgp*

Note: If you select the *all* option, the GSR will accept any of the six protocol types as the source protocol for the route(s).

- c. Specify the origin for your AS export source by selecting the appropriate type from the **Origin** drop-down list. You can select one of the following four options:
 - *any*
 - *igp*
 - *egp*
 - *incomplete*
11. Specify whether the GSR restricts the AS path export or associates a metric with the exported route by doing one of the following:
 - If you do not want the GSR to export any routes from the source, select the *Restrict export to the specified AS Path source* option.
 - If you want the GSR to associate a metric with the exported route, select a metric from the **Route Metric** drop-down list. Select *Inherited* if you want the route to use the policy's default metric. Otherwise, select a number from 1 to 16.
12. Click **OK**.
13. Repeat [step 9](#) through [step 12](#) until you create all of the AS Path export sources you plan to include in your export policies.

Configuring and Modifying Tag Export Sources

A tag export source building block specifies a source for exporting routes associated with a particular tag type into other protocols.

To define tag export source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Source object.
8. Double-click the Tag object, and then do one of the following:
 - If you are creating a new export source, click the Configure New Tag Export Source object.
 - If you are modifying an existing export source, select that export source.

A **Tag Export Source** dialog box similar to the following appears:

Figure 203. Tag Export Source dialog box

9. Specify the name you wish to assign to the tag export source in the **Source Name** box.
10. Under *Source Definition*, specify the following source criteria:
 - a. Specify a source protocol by selecting one of the following seven options from the **Protocol** drop-down list:
 - *all*
 - *static*
 - *direct*
 - *aggregate*
 - *rip*
 - *ospf*
 - *bgp*

Note: If you select the *all* option, the GSR will accept any of the six protocol types as the source protocol for the route(s).
 - b. Specify the tag you wish to associate with your export routes in the **Tag** box.
11. Specify whether the GSR restricts the tag export or associates a metric with the exported route by doing one of the following:

- If you do not want the GSR to export any routes from the source, select the *Restrict export to the specified Tag source* option.
 - If you want the GSR to associate a metric with the exported route, select a metric from the **Route Metric** drop-down list. Select *Inherited* if you want the route to use the policy's default metric. Otherwise, select a number from 1 to 16.
12. Click **OK**.
 13. Repeat [step 9](#) through [step 12](#) until you create all of the tag export sources you plan to include in your export policies.

Configuring and Modifying Direct Export Sources

A direct export source building block controls the export of routes from a source with routes directly attached to interfaces. To define direct export source building blocks that can be used for either RIP or OSPF:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Source object.
8. Double-click the Direct object, and then do one of the following:
 - If you are creating a new export source, click the Configure New Direct Export Source object.
 - If you are modifying an existing export source, select that export source.

A **Direct Export Source** dialog box similar to the following appears:

Figure 204. Direct Export Source dialog box

9. Specify the name you wish to assign to the direct export source in the **Destination Name** box.
10. Select the interface associated with the direct routes from the **Interface Name/Address** drop-down list.
11. Specify whether the GSR restricts the export or associates a metric with the exported routes by doing one of the following:
 - If you do not want the GSR to export any routes from the source, select the *Restrict export from the specified Direct Source* option.
 - If you want the GSR to associate a metric with the exported routes, select a metric from the **Route Metric** drop-down list. Select *Inherited* if you want the route to use the policy's default metric. Otherwise, select a number from 1 to 16.
12. Click **OK**.
13. Repeat [step 9](#) through [step 12](#) until you create all the direct export sources you will want to include in your export policies.

Configuring and Modifying Static Export Sources

A static export source building block specifies a source for exporting static routes into other protocols.

To define static export source building blocks that can be used for either RIP or OSPF:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Source object.
8. Double-click the Static object, and then do one of the following:
 - If you are creating a new export source, click the Configure New Static Export Source object.
 - If you are modifying an existing export source, select that export source.

A **Static Export Source** dialog box similar to the following appears:

The image shows a dialog box titled "Static Export Source". It contains the following fields and controls:

- Destination Name:** A text input field.
- Interface Name/Address:** A dropdown menu with "all" selected.
- Restrict Route Export from this Source:** An unchecked checkbox.
- Route Metric:** A dropdown menu with "Inherited" selected.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom.

Figure 205. Static Export Source dialog box

9. Specify the name you wish to assign to the export source in the **Destination Name** box.
10. Select the interface associated with the static routes from the **Interface Name/Address** drop-down list.

11. Specify whether the GSR restricts the export or associates a metric with the exported routes by doing one of the following:
 - If you do not want the GSR to export any routes from the source, select the *Restrict export from the specified Static Source* option.
 - If you want the GSR to associate a metric with the exported routes, select a metric from the **Route Metric** drop-down list. Select *Inherited* if you want the route to use the policy's default metric. Otherwise, select a number from 1 to 16.
12. Click **OK**.
13. Repeat [step 9](#) through [step 12](#) until you create all the static export sources you will want to include in your export policies.

Configuring and Modifying Aggregate Export Sources

An aggregate export source allows you to export multiple routes over various protocols while simultaneously associating an identifier with the routes.

Note: In order to define an aggregate export source building block, you must have previously configured at least one aggregate destination building block, as described in [“Aggregate Destination Building Blocks” on page 288](#).

To define aggregate destination building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Export Source object.
8. Double-click the Aggregates object, and then do one of the following:
 - If you are creating a new export source, click the Configure New Aggregate Export Source object.
 - If you are modifying an existing export source, select that export source.

An **Aggregate Export Source** dialog box similar to the following appears:

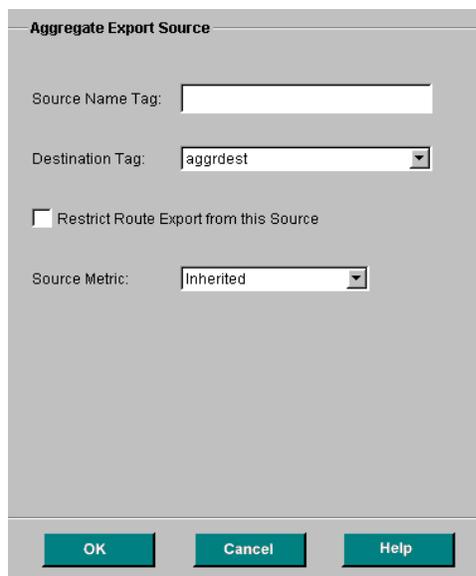


Figure 206. Aggregate Export Source dialog box

9. Specify the tag name you wish to assign to the export source in the **Source Name Tag** box.
10. Specify the tag you wish to assign to this destination by selecting it from the **Destination Tag** drop-down list.
11. Specify whether the GSR restricts the export or associates a metric with the exported routes by doing one of the following:
 - If you do not want the GSR to export any routes from the source, select the *Restrict export from this Source* option.
 - If you want the GSR to associate a metric with the exported routes, select a metric from the **Source Metric** drop-down list. Select *Inherited* if you want the route to use the policy's default metric. Otherwise, select a number from 1 to 16.
12. Click **OK**.
13. Repeat [step 9](#) through [step 12](#) until you create all the aggregate export sources you wish to include in your export policies.

Import Source Building Blocks

You can define import source building blocks to specify the interfaces or gateways from which the GSR learns routes. You can include your import source building blocks in a

GSR's import routing policies to control which routes are added to the GSR's routing table and to specify the GSR's preference of routes from one protocol or peer over another.

The method you use to create import source building blocks depends on whether you are defining one for RIP, OSPF, or BGP. Separate discussions on creating import source building blocks for the different protocols follow.

Configuring and Modifying RIP Import Sources

To define RIP import source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Import Source object.
8. Double-click the RIP object, and then do one of the following:
 - If you are creating a new import source, click the Configure New RIP Import Source object.
 - If you are modifying an existing import source, select that import source.

A **RIP Import Source** dialog box similar to the following appears:

Figure 207. RIP Import Source dialog box

9. Specify the name you wish to assign to the import source in the **Source Name** box.
10. Specify whether you want to control the importing of routes from an interface or gateway by doing one of the following:
 - If you want to control the importing of routes from an interface, select the *Imported over Interface* option and then select the desired IP interface from the associated drop-down list.
 - If you want to control the importing of routes from a gateway, select the *Imported from Gateway* option and specify that gateway's IP address in the associated text box.
11. Specify whether the GSR restricts the import or associates a preference with the imported routes by doing one of the following:
 - If you do not want the GSR to import any routes from the source, select the *Restrict Route Import from the specified RIP Source* option.
 - If you want the GSR to associate a preference with the imported routes, enter a number from 0 to 255 in the **Route Preference** box.

Configuration Expert lets you set preferences when you configure RIP, OSPF, and BGP routing policies. The GSR uses preference values to determine the preference of routes from one protocol or peer over another.

Note: The preference value you enter in the **RIP Import Source** dialog box overrides the preference value set in the **RIP Global Parameters** dialog box.

12. Click **OK**.
13. Repeat [step 8](#) through [step 12](#) until you create all of the RIP import sources you plan to include in your import policies.

Configuring and Modifying OSPF Import Sources

To define OSPF import source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Import Source object.
8. Double-click the OSPF object, and then do one of the following:
 - If you are creating a new import source, click the Configure New OSPF Import Source object.
 - If you are modifying an existing import source, select that import source.

An **OSPF Import Source** dialog box similar to the following appears:

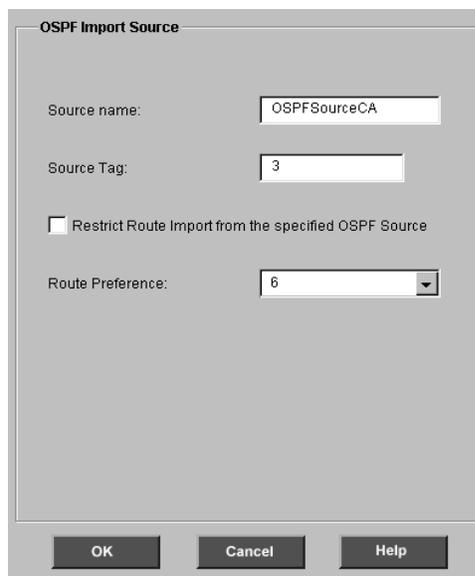


Figure 208. OSPF Import Source dialog box

9. Specify the name you wish to assign to the import source in the **Source Name** box.
10. Specify the tag number of the source tag you want associated with the imported routes in the **Source Tag** box.
11. Specify whether the GSR restricts the import or associates a preference with the imported routes by doing one of the following:
 - If you do not want the GSR to import any routes from the source, select the *Restrict Route Import from the specified OSPF Source* option.
 - If you want the GSR to associate a preference with the imported routes, enter a number from 0 to 255 in the **Route Preference** box.

Configuration Expert lets you set preferences when you configure RIP, OSPF, and BGP routing policies. The GSR uses preference values to determine the preference of routes from one protocol or peer over another.

Note: The preference value you enter in the **OSPF Import Source** dialog box overrides the preference value set in the **OSPF Global Parameters** dialog box.

12. Click **OK**.
13. Repeat [step 8](#) through [step 12](#) until you create all of the OSPF import sources you plan to include in your import policies.

Configuring and Modifying BGP Import Sources

To define BGP import source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Import Source object, click the BGP object, and then do one of the following:
 - If you are creating a new BGP import source, click the Configure New BGP Import Source object.
 - If you are modifying an existing BGP import source, select that source from the list of BGP import sources Configuration Expert displays.

A **BGP Import Source** dialog box similar to the following appears:

Figure 209. BGP Import Source Definition dialog box

8. Specify a name for your BGP import source in the **Source Name** box.

9. Under *Import BGP Routes based on*, specify the source of your BGP routes by doing one of the following:
 - To specify the source of your BGP routes as a particular autonomous system, select the *Autonomous System* option and specify the autonomous system identity in the associated text box.
 - To specify the source of your BGP routes as a regular expression from a particular autonomous system, select the *AS Path Regular Expression* option and specify the expression in the associated text box.
10. Specify the origin for your BGP import source by selecting the appropriate type from the **Origin** drop-down list. You can select one of the following four options:
 - *any*
 - *igp*
 - *egp*
 - *incomplete*
11. If you wish, specify an optional attribute for the BGP import source by selecting one from the **Optional Attribute** drop-down list. The available optional attributes in this drop-down list are a collection of pre-defined optional attributes you may have previously specified, as described in [“Configuring and Modifying Optional Attribute Building Blocks”](#) on page 316.
12. If you wish, specify a sequence number for the export destination in the **Sequence Number** box.
13. Specify whether the GSR restricts the import or associates a preference with the imported routes by doing one of the following:
 - If you do not want the GSR to import any routes from the source, select the *Restrict Route Import from the specified BGP Source* option.
 - If you want the GSR to associate a preference with the imported routes, enter a number from 0 to 255 in the **Route Preference** box.

Configuration Expert lets you set preferences when you configure RIP, OSPF, BGP and routing policies. The GSR uses preference values to determine the preference of routes from one protocol or peer over another.
14. Click OK.
15. Repeat [step 8](#) through [step 14](#) until you create all of the BGP import sources you plan to include in your import policies.

Configuring and Modifying Aggregate Source Building Blocks

You can define an aggregate source to specify the source's protocol, restrict the source's routes, specify whether the routes are considered contributors, and set the route's preference.

To create aggregate source building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Aggregate Sources object, and then do one of the following:
 - If you are creating a new aggregate source, click the Configure New Aggregate Source object.
 - If you are modifying an existing aggregate source, select that source from the list of aggregate sources Configuration Expert displays.

An **Aggregate Source Definition** dialog box similar to the following appears:

Figure 210. Aggregate Source Definition dialog box

8. Specify a name for your aggregate source in the **Aggregate Source** box.
9. From the **Route Protocol** drop-down list, select the protocol of the contributing aggregate source.
10. Do one of the following to specify which method of route restriction you wish to use:
 - To restrict selection of routes to those learned from the specified autonomous system, select the *Learned from AS* option and then enter a number from 1 to 65534 in the associated text box.

This selection may also be carried out by using route filters to explicitly list the set of routes to be accepted.
 - To restrict the selection of routes to those identified by a tag, select the *Having Tag* option and specify the desired tag in the associated text box.
 - To restrict selection of routes to those specified by the autonomous-system path, select the *AS Path Matching* option and specify the appropriate path in the associated text box.
 - To allow unrestricted route selection, select the *none* option.

This selection may also be carried out by using route filters.
11. Specify whether the routes are to be considered as contributors to the aggregate source by turning the *Resulting routes are not part of this aggregate* option on or off.
12. If the routes are to be considered as contributors of the aggregate source, specify the preference to assign to the contributing routes by entering a number from 0 to 255 in the **Route Preference** box.
13. Click **OK**.
14. Repeat [step 8](#) through [step 13](#) until you create all of the aggregate sources you plan to include in your aggregate policies.

IP Route Filter Building Blocks

The GSR uses IP route filters to match a certain set of routes by destination or by destination and mask. You define IP route filters by specifying a destination, mask, and modifiers that specify which routes the filter is for.

Note: A route will match the most specific filter that applies. Specifying more than one filter with the same destination, mask, and modifiers will generate an error.

The method you use to define an IP route filter depends on whether you are creating a new filter or modifying an existing one. Separate discussions on these tasks follow.

Configuring IP Route Filters

The GSR uses IP route filters to match a certain set of routes by destination or by destination and mask.

Note: A route will match the most specific filter that applies. Specifying more than one filter with the same destination, mask, and modifiers will generate an error.

To create IP route filters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Filters object, and click the Configure New Filter object.

Configuration Expert opens the Route Filter wizard.

8. Click **Next**.

A Filter Name panel similar to the following appears:

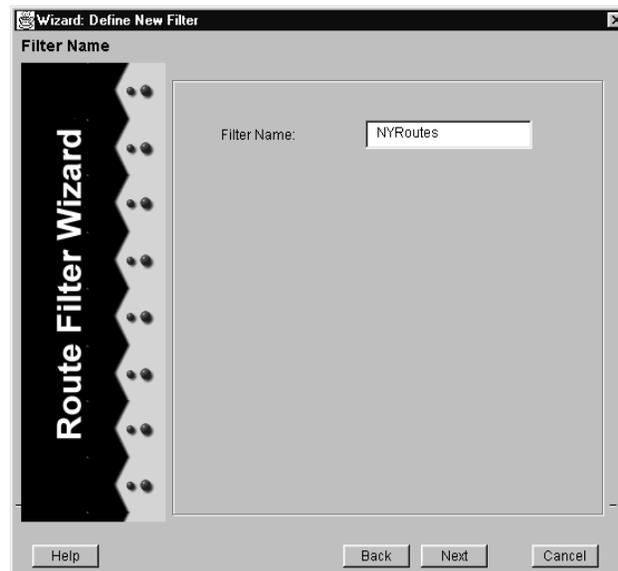


Figure 211. Filter Name panel

9. Specify the name you wish to assign to the filter in the Filter Name box.
10. Click **Next**.

A Filter Network Specification panel similar to the following appears:

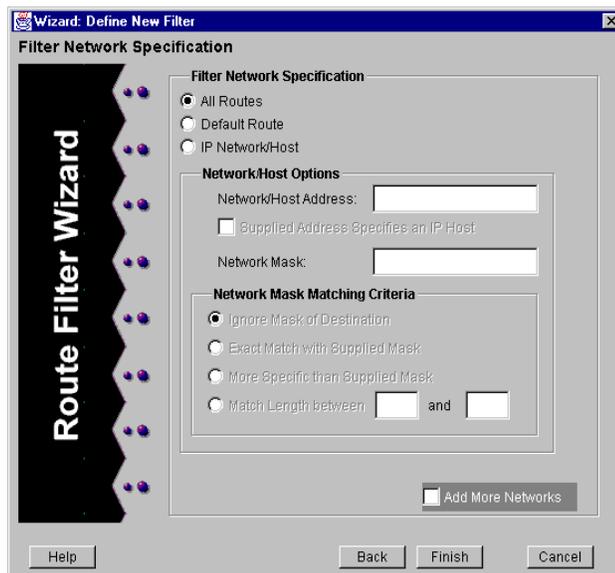


Figure 212. Filter Network Specification panel

11. Specify the level of route access by selecting the *All Routes* option, the *Default Route* option, or the *IP Network/Host* option.

If you select the *IP Network/Host* option, you can then proceed with the following step. Otherwise, skip to [step 13](#).

12. Under *Network/Host Options*, define the following characteristics:
 - a. Specify the network or host IP address and the network mask for your filter network in the **Network/Host Address** and **Network Mask** boxes, respectively.
 - b. If the IP address you specify in the Network/Host Address box points to an IP host, activate the *Supplied Address Specifies an IP Host* option and skip to [step 13](#). Otherwise, continue to the next step.
 - c. Under *Network Mask Matching Criteria*, select the appropriate option to fulfill your mask-matching requirements.

You can select one of the following four options:

- *Ignore Mask of Destination*
- *Exact Match with Supplied Mask*

- *More Specific than Supplied Mask*
- *Match Length Between <number> and <number>*

Note: If you select the *Match Length Between <number> and <number>* option, you must specify values defining the acceptable range to justify a match to your network mask. You can specify values between 0 and 32 for the acceptable length.

13. If you wish to specify additional filter networks, be sure to activate the *Add more filters* option, click **Next**, and repeat [step 11](#) and [step 12](#). Otherwise, click **Finish**.

Modifying IP Route Filters

Modify an IP route filter if you want to change its name, set different filter options, or add another network to the filter.

To modify IP route filters:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Filters object and then do one of the following:
 - If you want to change the name of the filter, select the filter you wish to change from the list of filters displayed in the Configuration Expert configuration tree.
 - If you want to change the filter settings of a network or host, select that network or host from those listed for the filter you are modifying. In the **Filter Network Specification** dialog box that appears, define the filter's settings.
 - If you want to add a new network or host to the filter, select the Add New Network object. In the **Filter Network Specification** dialog box that appears, define the filter's settings.

For details on the fields of the **Filter Network Specification** dialog box that appears when you are changing filter settings or adding a new network or host to the filter, see [“Configuring IP Route Filters” on page 313](#).

8. After making all your changes to the filter, click **OK**.

Configuring and Modifying Optional Attribute Building Blocks

You can configure optional attributes to include with your export destinations and BGP input sources.

To define optional path attribute building blocks:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Building Blocks object.
7. Double-click the Optional Attributes object, and then do one of the following:
 - If you are creating a new optional attribute, click the Configure New Optional Attributes List object.
 - If you are modifying an existing optional attribute, select the tag corresponding to that attribute from the list of optional attributes Configuration Expert displays.

An **Optional Path Attribute List** dialog box similar to the following appears:

The dialog box is titled "Optional Path Attribute List". It features a "Tag:" label followed by a text input field. Below this is a radio button labeled "Community Attributes" which is selected. Underneath is a bordered box containing "Community Id:" and "Autonomous System:" labels, each with a text input field. Below this box are two radio buttons: "Well Known Community" with a dropdown menu showing "no-export", and "Reserved Community" with a text input field. At the bottom are three buttons: "OK", "Cancel", and "Help".

Figure 213. Optional Path Attribute List dialog box

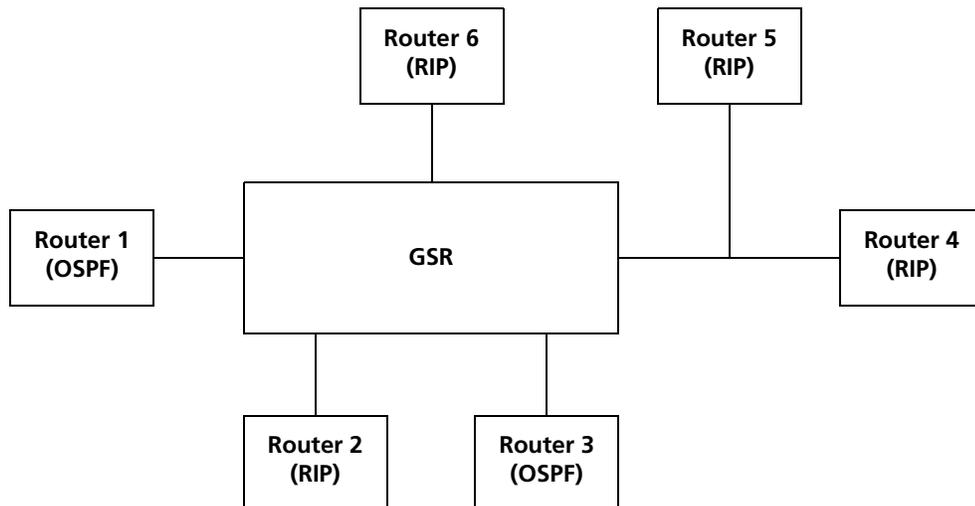
8. Specify a tag for your optional path attribute in the **Tag** box.
9. Specify the community type by doing one of the following:
 - To define specific community attributes for your optional path, select the *Community Attributes* option and specify the relative community identification and autonomous system in the **Community ID** and **Autonomous System** boxes, respectively.
 - To specify that you wish to tag routes with a well-known community identification, select the *Well Known Community* option and specify its relative behavior by selecting one of the four following options from the associated drop-down list:
 - *no-export*
 - *no-advertise*
 - *no-export-subconfed*
 - *none*
 - To specify that you wish to tag routes with a particular reserved community identification, select the *Reserved Community* option and specify the community identification in the associated text box.
10. Click **OK**.

Export Policies

A GSR's export routing policies control which routes the GSR advertises to other systems. Route exportation is controlled by both destination and source. Export policies specify the destinations to which the router will export, the sources from which the router will export, and the types of routes that will be exported, which can be either specific routes or all the routes a router has learned.

Suppose the GSR is attached to RIP and OSPF routers as shown in the following figure. Also suppose that RIP router 5 has learned all of the routes attached to RIP router 4. You could configure an export policy that directs the GSR do the following:

- Export all the routes of OSPF router 3 to RIP router 6.
- Export to RIP router 2, all of the routes directly attached to RIP router 5 but only some of the routes RIP router 5 has learned from RIP router 4.



Configuring Export Policies

After you create export destinations and sources and IP route filters as discussed earlier in this chapter, you can configure export routing policies by going through the following procedure.

To configure an export policy:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Export Policies object and click the Configure New Export Policy object.

Configuration Expert opens the Export Policy wizard.

7. Click **Next**.

A Routing Policy: Destination panel similar to the following appears:

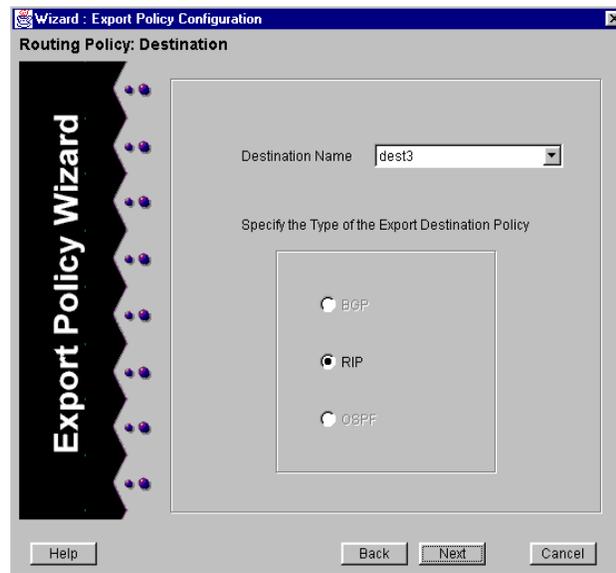


Figure 214. Routing Policy: Destination panel

8. Specify a destination name for your export policy by selecting it from the **Destination Name** drop-down list. The available destinations in this drop-down list are a collection of any pre-defined export destinations you may have previously configured, as described in [“Export Destination Building Blocks” on page 283](#).
9. Specify the export destination type by selecting the *BGP*, *RIP*, or *OSPF* option under *Specify the Type of the Export Destination Policy*.

Note: If you have not previously defined a particular type of export destination building block, then the option(s) corresponding to that type will remain inactive (dimmed) in this panel.

For example, suppose you have only defined one OSPF export destination building block and one RIP export destination building block. The *BGP* option under *Specify the Type of the Export Destination Policy* would remain inactive (dimmed).

10. Click **Next**.

An Export Policy Destination Source panel similar to the following appears:

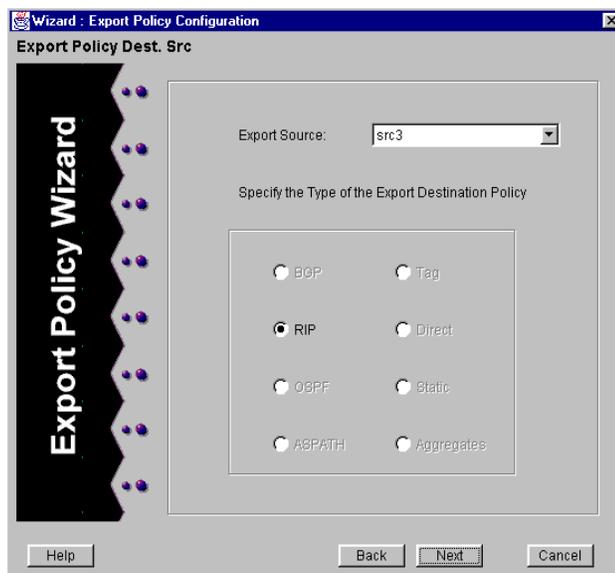


Figure 215. Export Policy Destination Source panel

11. Specify the source for your export policy by selecting it from the **Export Source** drop-down list. The available sources in this drop-down list are a collection of your pre-defined export sources.
12. Specify the export destination type by selecting the following eight options under *Specify the Type of the Export Destination Policy*:
 - *BGP*
 - *RIP*
 - *OSPF*
 - *ASPATH*
 - *Tag*
 - *Direct*
 - *Static*
 - *Aggregates*

Note: If you have not previously defined a particular type of export destination building block, then the option(s) corresponding to that type will remain inactive (dimmed) in this panel.

For example, suppose you have only defined one OSPF export destination building block and one RIP export destination building block. All of the other

options under *Specify the Type of the Export Destination Policy* would remain inactive (dimmed).

13. Click **Next**.

A Filter Specification panel similar to the following appears:

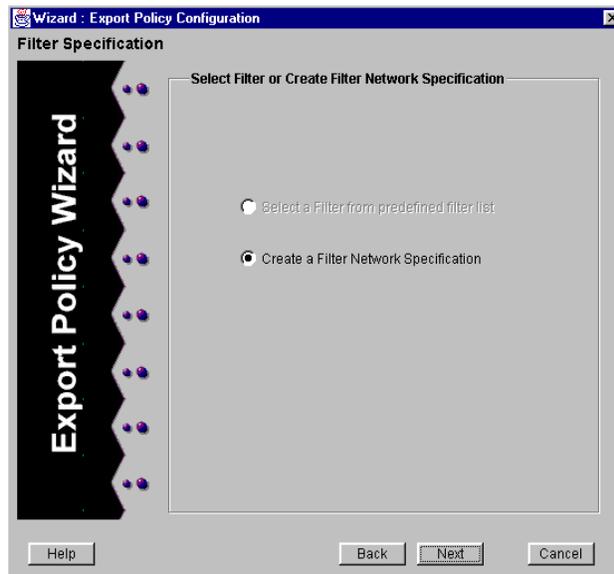


Figure 216. Filter Specification panel

14. Specify whether you wish to use an existing filter or create a new filter for your export policy by selecting the appropriate option.

Note: The *Select a filter from predefined filter list* option will remain inactive (dimmed) unless you have previously defined one or more filter building blocks as described in [“IP Route Filter Building Blocks” on page 312](#).

15. Click **Next**.

A Filter Network Specification panel similar to the following appears:

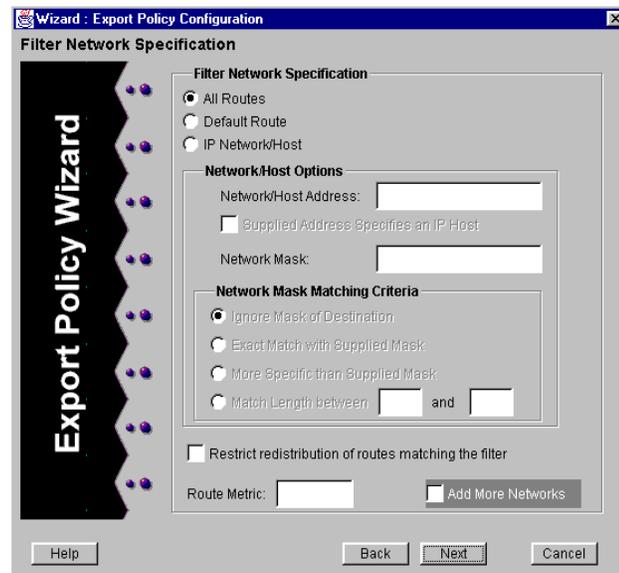


Figure 217. Filter Network Specification panel

16. Specify the level of route access by selecting the *All Routes* option, the *Default Route* option, or the *IP Network/Host* option.

If you select the *IP Network/Host* option, you can then proceed with the following step. Otherwise, skip to [step 18](#).

17. Under *Network/Host Options*, define the following characteristics:
 - a. Specify the network or host IP address and the network mask for your export policy network in the **Network/Host Address** and **Network Mask** boxes, respectively.
 - b. If the IP address you specify in the **Network/Host Address** box points to an IP host, activate the *Supplied Address Specifies an IP Host* option and skip to [step 18](#). Otherwise, continue to the next step.
 - c. Under *Network Mask Matching Criteria*, select the appropriate option to fulfill your mask-matching requirements.

You can select one of the following four options:

- *Ignore Mask of Destination*
- *Exact Match with Supplied Mask*
- *More Specific than Supplied Mask*
- *Match Length Between <number> and <number>*

Note: If you select the *Match Length Between <number> and <number>* option, you must specify values defining the acceptable range to justify a match to your network mask. You can specify values between 0 and 32 for the acceptable length.

18. Specify whether or not you wish to restrict redistribution of routes that match your defined filter by turning the *Restrict redistribution of routes matching the filter* option on or off, accordingly.
19. If you wish, you can specify a route metric for the filter network by entering any numerical value in the **Route Metric** box.
20. If you wish to specify additional filter networks, be sure to activate the *Add more filters* option, click **Next**, and repeat [step 16](#) through [step 19](#). Otherwise, click **Finish**.

Modifying Export Policies

You can modify export policies to change the filter and/or network specifications for one or more export sources associated to a particular export destination.

To modify an existing export policy:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Export Policies object.
7. Double-click the name of the existing export policy you wish to modify.

Once the list of existing export sources for this policy is displayed, double-click the export source object you wish to modify.

8. Do one of the following:
 - Double-click the Filters object and then select the name of the export source filter you wish to modify.
 - Double-click the Network Specification object and then select the name of the export source network you wish to modify.
9. Modify the attributes of the export source filter and/or network by editing the fields of the **Export Policy Filter** dialog box or **Filter Network Specification** dialog box that appears.

These dialog boxes include some of the same fields you specified when configuring the export policy. For more information on these fields, see [“Export Policies” on page 317](#).

10. Click **OK**.

Import Policies

A GSR's import routing policies control which routes are added to the GSR's routing table. Import routing policies also determine the GSR's preference of routes from one protocol or peer over another.

Configuring Import Policies

After you create import destinations and sources and IP route filters as discussed earlier in this chapter, you can configure import routing policies by going through the following procedure.

Note: Before you attempt to configure an import policy, you must ensure that you have previously configured at least one import source building block, as described in [“Import Source Building Blocks” on page 304](#).

To configure an import policy:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Import Policies object and click the Configure New Import Policy object.
Configuration Expert opens the Import Policy wizard.
7. Click **Next**.

An Import Source panel similar to the following appears:

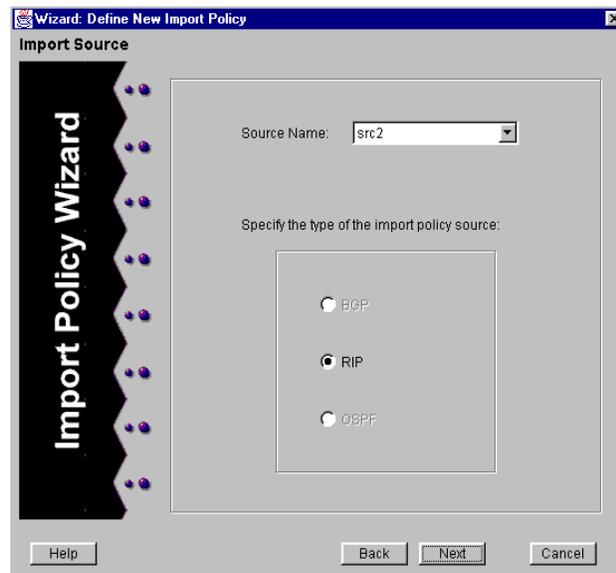


Figure 218. Import Source panel

8. Specify a source name for your import policy by selecting it from the **Source Name** drop-down list. The available sources in this drop-down list are a collection of any pre-defined import sources you may have previously configured, as described in [“Import Source Building Blocks” on page 304](#).
9. Specify the import source type by selecting the *BGP*, *RIP*, or *OSPF* option under *Specify the type of the import policy source*.

Note: If you have not previously defined a particular type of import source building block, then the option(s) corresponding to that type will remain inactive (dimmed) in this panel.

For example, suppose you have only defined one OSPF import source building block and one RIP import source building block. The *BGP* option under *Specify the type of the import policy source* would remain inactive (dimmed).

10. Click **Next**.

A Filter Specification panel similar to the following appears:

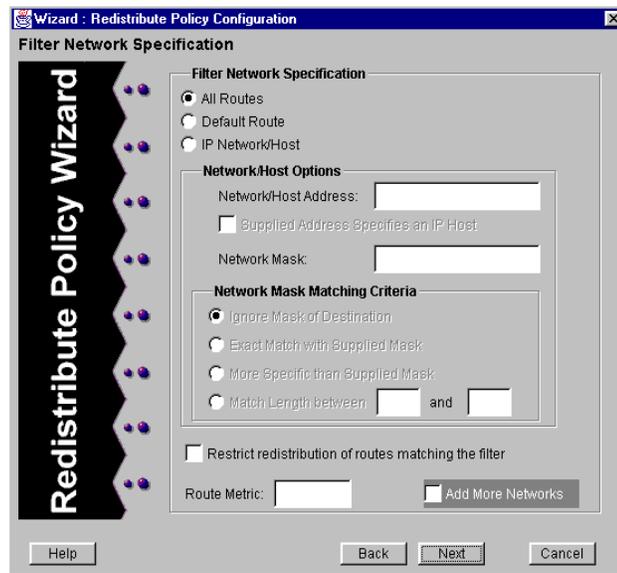


Figure 219. Filter Specification panel

11. Specify whether you wish to use an existing filter or create a new filter for your import policy by selecting the appropriate option.

Note: The *Select a filter from predefined filter list* option will remain inactive (dimmed) unless you have previously defined one or more filter building blocks as described in [“IP Route Filter Building Blocks”](#) on page 312.

12. Click **Next**.

A Filter Network Specification panel similar to the following appears:

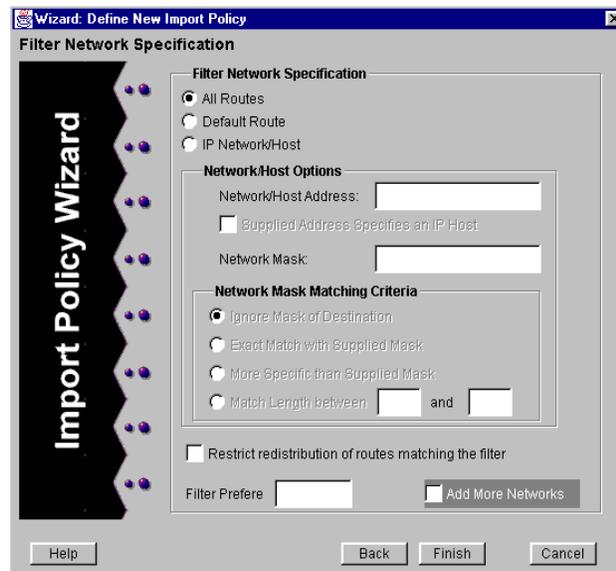


Figure 220. Filter Network Specification panel

13. Specify the level of route access by selecting the *All Routes* option, the *Default Route* option, or the *IP Network/Host* option.

If you select the *IP Network/Host* option, you can then proceed with the following step. Otherwise, skip to [step 15](#).

14. Under *Network/Host Options*, define the following characteristics:
 - a. Specify the network or host IP address and the network mask for your import policy network in the **Network/Host Address** and **Network Mask** boxes, respectively.
 - b. If the IP address you specify in the **Network/Host Address** box points to an IP host, activate the *Supplied Address Specifies an IP Host* option and skip to [step 15](#). Otherwise, continue to the next step.
 - c. Under *Network Mask Matching Criteria*, select the appropriate option to fulfill your mask-matching requirements.

You can select one of the following four options:

- *Ignore Mask of Destination*
- *Exact Match with Supplied Mask*
- *More Specific than Supplied Mask*
- *Match Length Between <number> and <number>*

Note: If you select the *Match Length Between <number> and <number>* option, you must specify values defining the acceptable range to justify a match to your network mask. You can specify values between 0 and 32 for the acceptable length.

15. Specify whether or not you wish to restrict redistribution of routes that match your defined filter by turning the *Restrict redistribution of routes matching the filter* option on or off, accordingly.
16. If you wish, you can specify a preference for the filter network by entering a value between 0 and 65,536 in the **Filter Preference** box.
17. If you wish to specify additional filter networks, be sure to activate the *Add more filters* option, click **Next**, and repeat [step 13](#) through [step 16](#). Otherwise, click **Finish**.

Modifying Import Policies

You can modify import policies to change the filter and/or network specifications for one or more import sources.

To modify an existing import policy:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Import Policies object.
7. Double-click the name of the existing import source you wish to modify.
8. Do one of the following:
 - Double-click the Filters object and then select the name of the import source filter you wish to modify.
 - Double-click the Network Specification object and then select the name of the import source network you wish to modify.
9. Modify the attributes of the import source filter and/or network by editing the fields of the **Import Policy Filter** dialog box or **Filter Network Specification** dialog box that appears.

These dialog boxes include some of the same fields you specified when configuring the import policy. For more information on these fields, see [“Import Policies” on page 324](#).
10. Click **OK**.

Aggregate Policies

Configuration Expert lets you create aggregates to restrict the selection of routes, to specify whether the routes are considered to be contributors, and if the routes are considered to be contributors to specify route preference.

An aggregate route is a general route composed of two or more contributing routes. In the following figure, California is an aggregate route the GSR has learned. That aggregate route consists of the Santa Barbara, Monterey, and Eureka contributing routes.



The GSR can export the California aggregate route to the New York router as long as one of the contributing routes remain active. The New Yorker router learns the California aggregate route but it does not learn the individual contributing routes.

Configuring Aggregate Policies

After you create aggregate export sources, aggregate destinations, and aggregate sources and IP route filters as discussed earlier in this chapter, you can configure aggregate routing policies by going through the following procedure.

Note: Before you attempt to configure an aggregate policy, you must ensure that you have previously configured at least one unused Aggregate Destination and one unused Aggregate Source building block, as described in [“Export Policies” on page 317](#) and [“Export Policies” on page 317](#).

To configure an aggregate policy:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Aggregate Policies object and click the Configure New Aggregate Policy object.

Configuration Expert opens the Aggregate Policy wizard.

7. Click **Next**.

An Aggregate Generate panel similar to the following appears:

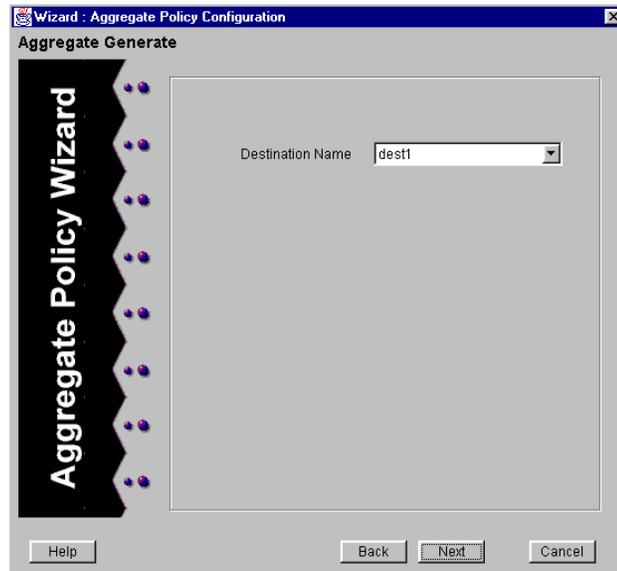


Figure 221. Aggregate Generate panel

8. Specify the destination name for your aggregate policy by selecting it from the **Destination Name** drop-down list. The available destinations in this drop-down list are a collection of any pre-defined aggregate destinations you may have previously configured, as described in [“Aggregate Destination Building Blocks”](#) on page 288.
9. Click **Next**.

An Aggregate Policy Source panel similar to the following appears:

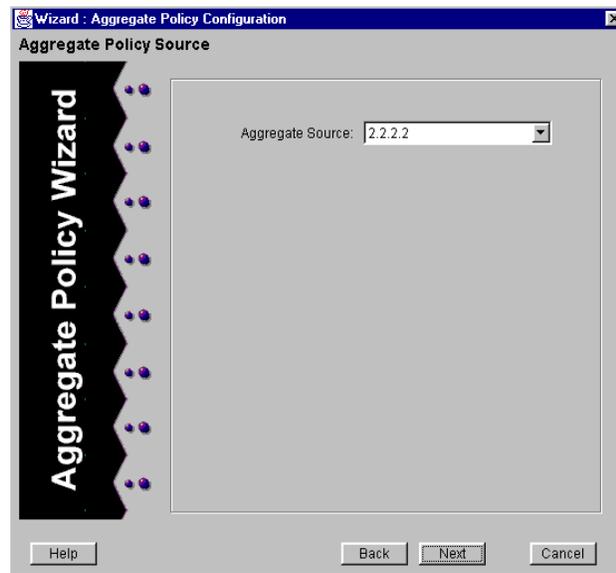


Figure 222. Aggregate Policy Source panel

10. Specify the source for your aggregate policy by selecting it from the **Aggregate Source** drop-down list. The available sources in this drop-down list are a collection of any pre-defined aggregate sources you may have previously configured, as described in ["Configuring and Modifying Aggregate Source Building Blocks"](#) on page 311.
11. Click **Next**.

A Filter Specification panel similar to the following appears:

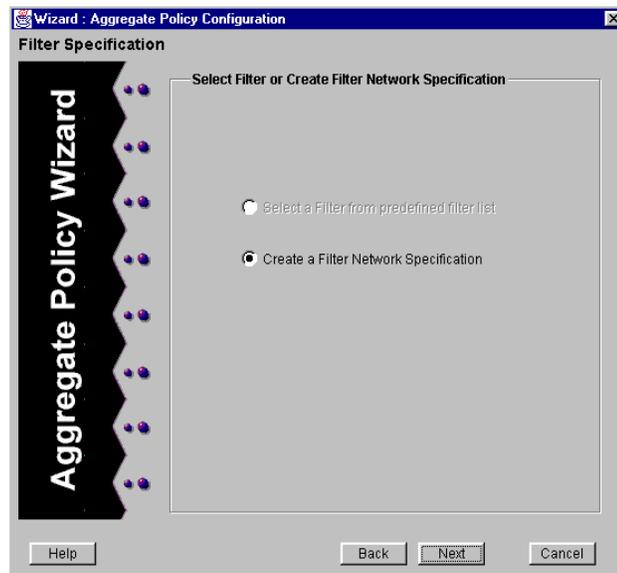


Figure 223. Filter Specification panel

12. Specify whether you wish to use an existing filter or create a new filter for your aggregate policy by selecting the appropriate option.

Note: The *Select a filter from predefined filter list* option will be inactive (dimmed) unless you have previously defined one or more filter building blocks as described in [“IP Route Filter Building Blocks”](#) on page 312.

13. Click **Next**.

A Filter Network Specification panel similar to the following appears:

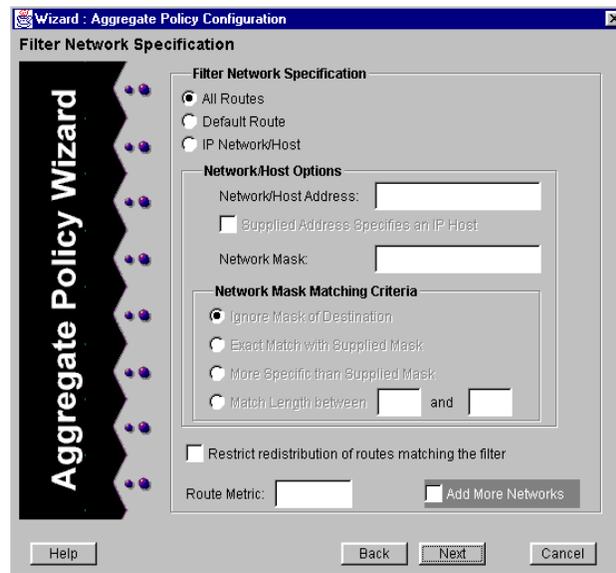


Figure 224. Filter Network Specification panel

14. Specify the level of route access by selecting the *All Routes* option, the *Default Route* option, or the *IP Network/Host* option.

If you select the *IP Network/Host* option, you can then proceed with the following step. Otherwise, skip to [step 16](#).

15. Under *Network/Host Options*, define the following characteristics:
 - a. Specify the network or host IP address and the network mask for your aggregate policy network in the **Network/Host Address** and **Network Mask** boxes, respectively.
 - b. If the IP address you specify in the **Network/Host Address** box points to an IP host, activate the *Supplied Address Specifies an IP Host* option and skip to [step 16](#). Otherwise, continue to the next step.
 - c. Under *Network Mask Matching Criteria*, select the appropriate option to fulfill your mask-matching requirements.

You can select one of the following four options:

- *Ignore Mask of Destination*
- *Exact Match with Supplied Mask*
- *More Specific than Supplied Mask*
- *Match Length Between <number> and <number>*

Note: If you select the *Match Length Between <number> and <number>* option, you must specify values defining the acceptable range to justify a match to your network mask. You can specify values between 0 and 32 for the acceptable length.

16. Specify whether or not you wish to restrict redistribution of routes that match your defined filter by turning the *Restrict redistribution of routes matching the filter* option on or off, accordingly.
17. If you wish, you can specify a route metric for the filter network by entering a value between 0 and 255 in the **Route Metric** box.
18. If you wish to specify additional filter networks, be sure to activate the *Add more filters* option, click **Next**, and repeat [step 14](#) through [step 17](#). Otherwise, click **Finish**.

Modifying Aggregate Policies

You can modify aggregate policies to change the filter and/or network specifications for one or more aggregate sources associated to a particular aggregate destination.

To modify an existing aggregate policy:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Aggregate Policies object.
7. Double-click the name of the existing aggregate policy you wish to modify.

Once the list of existing aggregate sources for this policy is displayed, double-click the aggregate source object you wish to modify.

8. Do one of the following:
 - Double-click the Filters object and then select the name of the aggregate source filter you wish to modify.
 - Double-click the Network Specification object and then select the name of the aggregate source network you wish to modify.
9. Modify the attributes of the aggregate source filter and/or network by editing the fields of the **Aggregate Policy Filter** dialog box or **Filter Network Specification** dialog box that appears.

These dialog boxes include some of the same fields you specified when configuring the aggregate policy. For more information on these fields, see [“Aggregate Policies” on page 329](#).

10. Click **OK**.

Redistribute Policies

You can use Configuration Expert to create redistribute policies for the GSR. Redistribute policies allow you to redistribute learned routes from one routing protocol to another. Using redistribute policies is a simple way to configure export policies based primarily on routing protocols.

Configuring Redistribute Policies

To configure a new redistribute policy for the GSR:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Redistribute Policies object and click the Configure New Redistribute Policy object.

Configuration Expert opens the Redistribute Policy wizard.

7. Click **Next**.

A Source Protocol panel similar to the following appears:

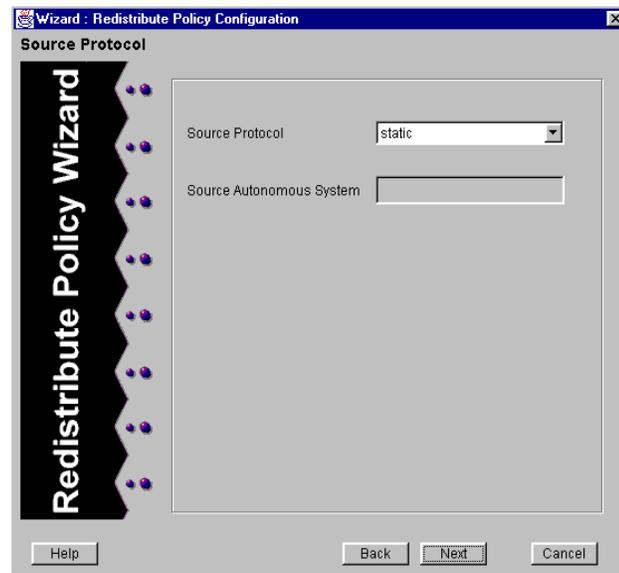


Figure 225. Source Protocol panel

8. Specify the source protocol for your redistribute policy by selecting one of the available options from the **Source Protocol** drop-down list.

You can select from the following seven options:

- *static*
- *direct*
- *aggregate*
- *rip*
- *ospf*
- *ospf-ase*
- *bgp*

Note: If you select the *bgp* option, you can also choose to specify the autonomous system for your BGP source in the **Source Autonomous System** box.

9. Click **Next**.

A Target Protocol panel similar to the following appears:

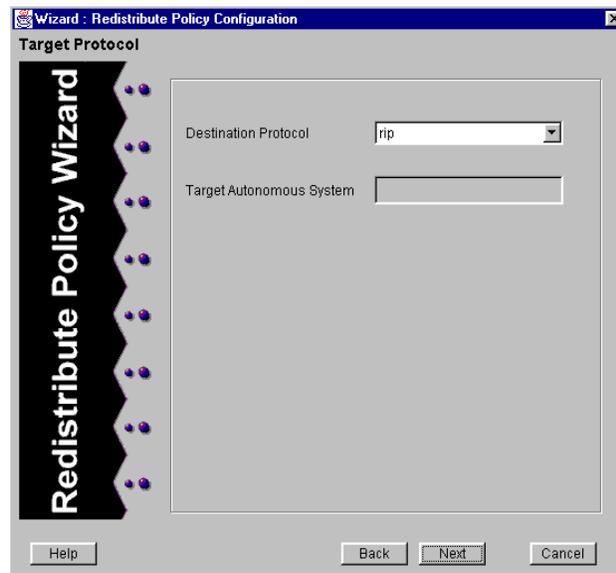


Figure 226. Target Protocol panel

10. Specify the destination protocol for your redistribute policy by selecting one of the available options from the **Destination Protocol** drop-down list.

You can select from the following three options:

- *rip*
- *ospf*
- *bgp*

Note: If you select the *bgp* option, you can also choose to specify the autonomous system for your BGP target in the **Target Autonomous System** box.

11. Click **Next**.

A Filter Network Specification panel similar to the following appears:

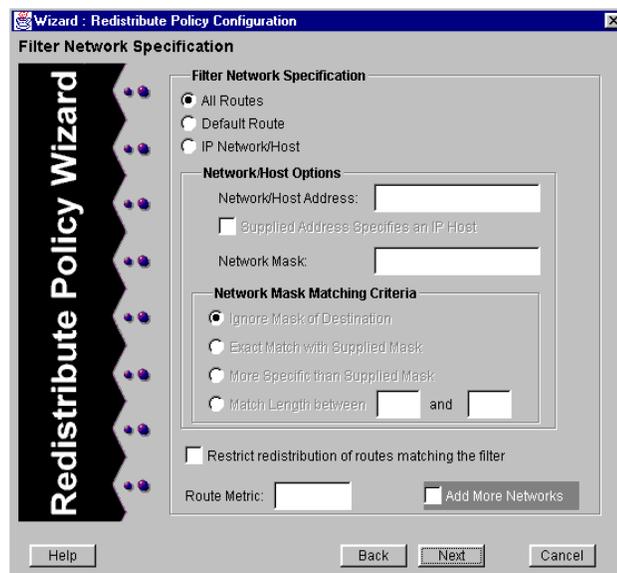


Figure 227. Filter Network Specification panel

12. Specify the level of route access by selecting the *All Routes* option, the *Default Route* option, or the *IP Network/Host* option.

If you select the *IP Network/Host* option, you can then proceed with the following step. Otherwise, skip to [step 14](#).

13. Under *Network/Host Options*, define the following characteristics:
 - a. Specify the network or host IP address and the network mask for your redistribute policy network in the **Network/Host Address** and **Network Mask** boxes, respectively.
 - b. If the IP address you specify in the **Network/Host Address** box points to an IP host, activate the *Supplied Address Specifies an IP Host* option and skip to [step 14](#). Otherwise, continue to the next step.
 - c. Under *Network Mask Matching Criteria*, select the appropriate option to fulfill your mask-matching requirements.

You can select one of the following four options:

- *Ignore Mask of Destination*
- *Exact Match with Supplied Mask*
- *More Specific than Supplied Mask*
- *Match Length Between <number> and <number>*

Note: If you select the *Match Length Between <number> and <number>* option, you must specify values defining the acceptable range to justify a match to your network mask. You can specify values between 0 and 32 for the acceptable length.

14. Specify whether or not you wish to restrict redistribution of routes that match your defined filter by turning the *Restrict redistribution of routes matching the filter* option on or off, accordingly.
15. If you wish, you can specify a route metric for the filter network by entering a value between 0 and 255 in the **Route Metric** box.
16. If you wish to specify additional filter networks, be sure to activate the *Add more filters* option, click **Next**, and repeat [step 12](#) through [step 15](#). Otherwise, click **Finish**.

Modifying Redistribute Policies

You can modify redistribute policies to change the network specifications for one or more redistribute destination protocol types associated with a particular redistribute source protocol type.

To modify an existing redistribute policy:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Redistribute Policies object.
7. Double-click the existing source protocol type you wish to modify.

Once the list of existing destination protocol types associated with this source protocol type is displayed, double-click the destination protocol type you wish to modify.

8. Select the Network Specification object.
9. Modify the attributes of the redistribute source network by editing the fields of the **Filter Network Specification** dialog box that appears.

This dialog box includes some of the same fields you specified when configuring the redistribute policy. For more information on these fields, see [“Redistribute Policies” on page 335](#).

10. Click **OK**.

Summarize Routes

You can use Configuration Expert to configure summarize routes for the GSR. Summarize routes are a simple form of aggregate configuration. When you define aggregates, you must expressly specify which routes will become continuing routes. However, summarize routes simplify the process by treating all routes as continuing routes.

Configuring Summarize Routes

To configure a new Summarize Route for the GSR:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Summarize Routes object and click the Configure New Summarize Route object.

Configuration Expert opens the Summarize Route wizard.

7. Click **Next**.

A Summarize Network panel similar to the following appears:

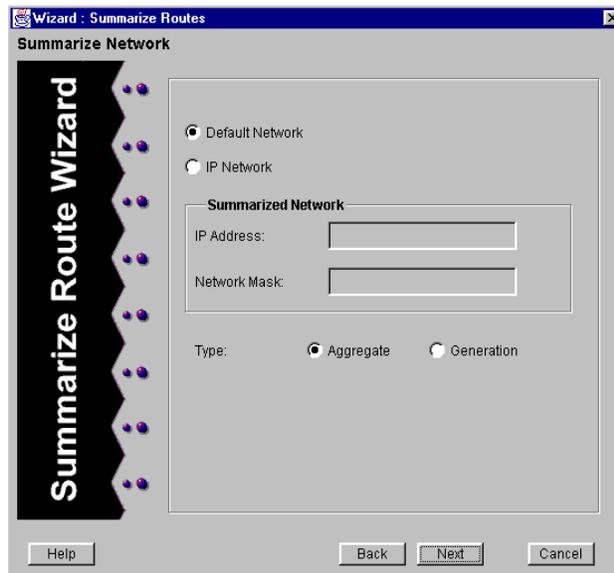


Figure 228. Summarize Network panel

8. Specify whether the new summarize network is to be a default network or an associated IP network by selecting the appropriate option.

If you selected the *IP Network* option, proceed with the following step. If you selected the *Default Network* option, skip to [step 10](#).

9. Under *Summarized Network*, specify the IP address and network mask for the associated IP network in the **IP Address** and **Network Mask** boxes, respectively.
10. Specify the network type by selecting either the *Aggregate* or *Generation* option.
11. Click **Next**.

A Source Protocol panel similar to the following appears:

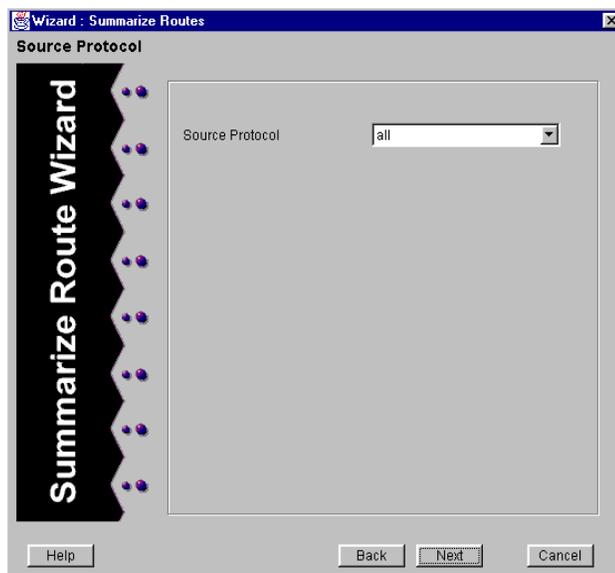


Figure 229. Source Protocol panel

12. Specify the source protocol for the summarize network by selecting one of the six available options from the **Source Protocol** drop-down list.

You can select from the following six options:

- *all*
- *rip*
- *ospf*
- *bgp*
- *direct*
- *static*

Note: Selecting the *all* option allows you to use all of the protocols listed as your source protocol.

13. Click **Next**.

A Filter Network Specification panel similar to the following appears:

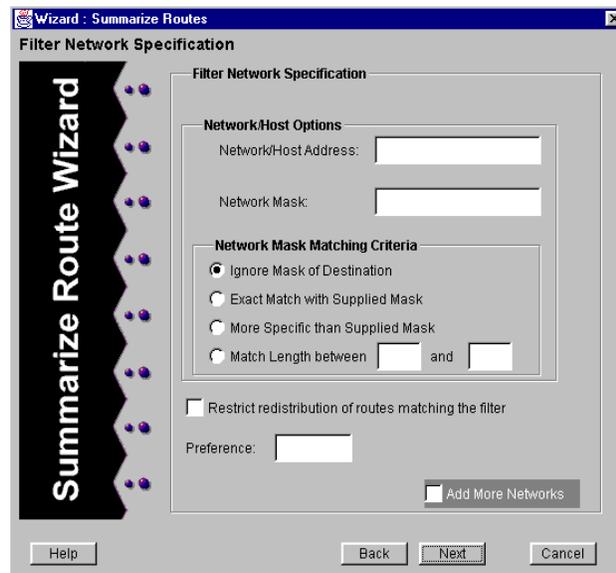


Figure 230. Filter Network Specification panel

14. Under *Network/Host Options*, define the following characteristics:
 - a. Specify the network or host IP address and the network mask for your summarize network in the **Network/Host Address** and **Network Mask** boxes, respectively.
 - b. Under *Network Mask Matching Criteria*, select the appropriate option to fulfill your mask-matching requirements.

You can select one of the following four options:

- *Ignore Mask of Destination*
- *Exact Match with Supplied Mask*
- *More Specific than Supplied Mask*
- *Match Length Between <number> and <number>*

Note: If you select the *Match Length Between <number> and <number>* option, you must specify values defining the acceptable range to justify a match to your network mask. You can specify values between 0 and 32 for the acceptable length.

15. Specify whether or not you wish to restrict redistribution of routes that match your defined filter by turning the *Restrict redistribution of routes matching the filter* option on or off, accordingly.

16. If you wish, you can specify a level of preference for the filter network by entering any numerical value in the **Preference** box.
17. If you wish to specify additional filter networks, be sure to activate the *Add more filters* option, click **Next**, and repeat [step 14](#) through [step 16](#). Otherwise, click **Finish**.

Modifying Summarize Routes

You can modify summarize routes to change their IP address and/or network specifications.

To modify an existing summarize route:

1. Start Configuration Expert if you have not already done so.
2. Open the configuration file you want to modify and then double-click that file's Routing Configuration object.
3. Double-click the IP Routing Configuration object.
4. Double-click the IP Unicast Routing object.
5. Double-click the Routing Policy Configuration object.
6. Double-click the Summarize Routes object.
7. If you wish to modify the route definition for a particular summarize route, select the IP address for that route from the list of existing summarize routes and modify the attributes of the route definition by editing the fields of the **Summarized Route Definition** dialog box that appears and click **OK**.

This dialog box includes some of the same fields you specified when configuring the summarize route. For more information on these fields, see ["Summarize Routes" on page 340](#).

8. If you wish to modify the filter network for a particular summarize route, go through the following procedure:
 - a. Double-click the IP address for that summarize route. Once the list of existing source protocol types for this IP address is displayed, double-click the source protocol type you wish to modify.
 - b. Select the IP address for the filter network you wish to modify from the list of existing filter networks.

- c. Modify the attributes of the filter network by editing the fields of the **Filter Network Specification** dialog box that appears.

This dialog box includes some of the same fields you specified when configuring the summarize route. For more information on these fields, see [“Summarize Routes” on page 340](#).

9. Click **OK**.

Chapter 17

Checking System Status

clearVISN CoreWatch can display the following system information:

- Details about which modules are installed in the GSR chassis and the number of ports available on those modules.
- Details about individual ports. This includes data about which module each port is on, a port's bridging status, information about the VLAN associated with each port, and IP address information of a port.

Obtaining Chassis Information

Obtain chassis information if you want to know which modules are installed in the GSR, the slot number of each module, or the number of ports on each module. To access such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **System State**, and then choose **Chassis Table**.

A **Chassis Info** table similar to the following appears:

The screenshot shows a window titled "207.135.89.55 : Chassis Info Table". It contains a table with the following data:

Slot Number	Module Type	Module Description	Number of Ports	Module Version
1	cpuModule	Control Module	0	0
3	bt100	10/100-TX	8	1
4	bt100	10/100-TX	8	0
5	bt100	10/100-TX	8	1
6	bt100	10/100-TX	8	0
7	bt100	10/100-TX	8	1

Below the table is a "Selection Details" section showing the details for the selected item (Slot Number 1):

Slot Number	1
Module Type	cpuModule
Module Description	Control Module
Number of Ports	0
Module Version	0

At the bottom left, it says "7 rows displayed".

Figure 231. Chassis Info table

The following table describes the fields of the **Chassis Info** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 39. Chassis Info table fields

Field	Description
Slot Number	Identifies the slot in which the module is installed.
Module Type	Identifies whether the module is a Control Module or a 10/100-TX, 100-FX, Gigabit-LX, or Gigabit-SX Ethernet, or WAN-Serial or WAN-HSSI module.
Module Description	Provides a brief description of the module. This description is defined when the GSR is configured.
Number of Ports	Indicates how many ports are on the module.
Module Version	Indicates the module's version.

Obtaining Port Information

Obtain port information if you want to know which modules are installed in the GSR, the slot number of a module, or the number of ports on each module. To access such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **System State**, and then choose **Port Table**.

A **Port** table similar to the following appears:

Port ID	Module ID	Bridge Port Status	VLAN Name	Bound IP Interface	IP
2	7	disabled	VLAN BR-GIG-BACKB...	None	N/A
2	7	disabled	VLAN GIG-BACKBONE	GIG-BACKBONE	207.1
1	7	disabled	VLAN BR-MAINSUBNET	None	N/A
1	7	disabled	VLAN MAINSUBNET	MAINSUBNET	207.1
8	6	disabled	VLAN BR-MAINSUBNET	None	N/A

Selection Details	
Port ID	2
Module ID	7
Bridge Port Status	disabled
VLAN Name	VLAN BR-GIG-BACKBONE
Bound IP Interface	None
IP Address	N/A

84 rows displayed

Figure 232. Port table

The following table describes the fields of the **Port** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 40. Port table fields

Field	Description
Port ID	Identifies the number of the port.
Module ID	Indicates the port's slot number.
Bridge Port Status	Indicates whether the port is configured for bridging.
VLAN Name	Identifies the VLAN associated with the port.
Bound IP Interface	Identifies the interface to which the port is bound.
IP Address	Identifies the port's IP address.
IP Subnet Mask	Identifies the port's IP subnet mask.
Broadcast IP Address	Identifies the broadcast IP address used for the port.

Obtaining Trap Information

Obtain trap information if you want to know what traps are currently configured for the GSR, to which community they are associated, and to which IP address(es) they will be sent. To access such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **System State**, and then choose **Trap Table**.

A **Trap** table similar to the following appears:

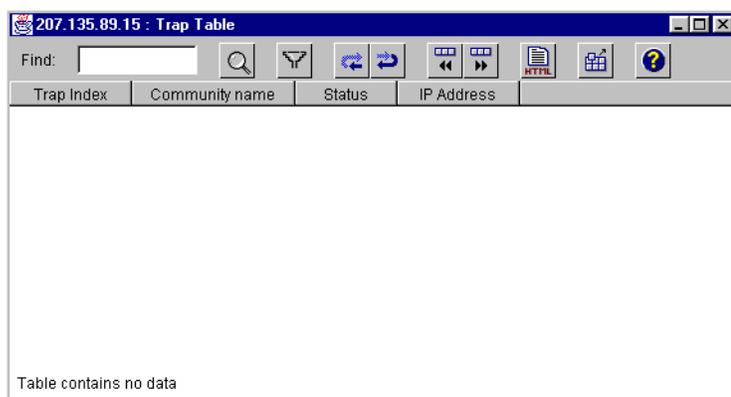


Figure 233. Trap table

The following table describes the fields of the **Trap** table's upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 41. Trap table fields

Field	Description
Trap Index	Displays all of the currently configured traps for the GSR. This identifier is simply a number between 1 and 32.
Community Name	SNMP with string sent to management for purpose of query.
Status	Displays the current state of the trap, either "enabled" or "disabled". When a trap is disabled, no trap is sent to the associated IP address.
IP Address	Displays the IP address to which the destination trap will be sent.

Obtaining SmartTRUNK Information

Obtain SmartTRUNK information if you want to know which SmartTRUNKs have been configured for the GSR, what type of ports they are, and at what locations. To access such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **System State**, and then choose **SmartTRUNK Table**.

A SmartTRUNK table similar to the following appears:

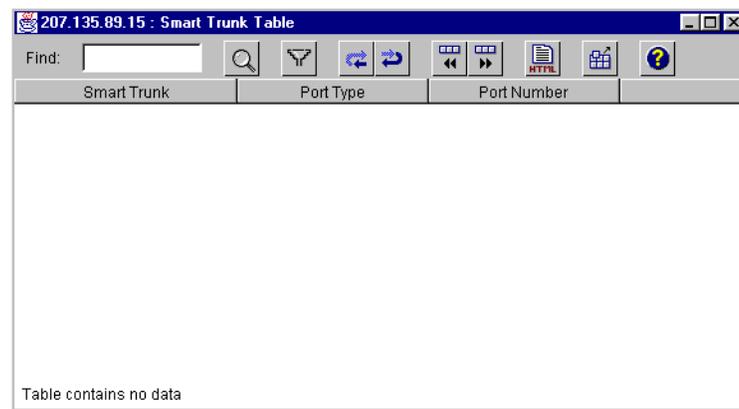


Figure 234. SmartTRUNK table

The following table describes the fields of the SmartTRUNK table's upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 42. SmartTRUNK table fields

Field	Description
SmartTRUNK	Displays the identifier (name) for the given SmartTRUNK
Port Type	Displays the type of port that has been designated as a SmartTRUNK. For example, "ethernet" or "gigether".
Port Number	Displays the true port identifier (number) of the SmartTRUNK.

Chapter 18

Monitoring Real-Time Performance

You may obtain current statistics about the following:

- System performance, which indicates the speed at which the GSR is transmitting and receiving bytes and packets.
- Overall use of all the ports of a GSR.
- Packets, bytes, and errors on specific ports.
- IP traffic of a GSR including details about various IP packets on the router, reassembly and fragmentation activities, and errors related to the IP data the GSR receives or sends.
- IPX traffic of a GSR including details about incoming and outgoing packets, and errors related to the IPX data the GSR receives or sends.

Separate discussions on the different statistics you can obtain follow.

Monitoring System Performance

Monitor a GSR's system performance by displaying the System Dashboard, which is a set of dials that indicate the following information. You may select which of these dials you want clearVISN CoreWatch to display.

- Bits per second at which the GSR is transmitting and receiving
- Packets per second at which the GSR is transmitting and receiving

To monitor system performance:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Do one of the following:
 - Select the **Monitor** menu, choose **Performance**, and then choose **System Dashboard**.
 - Click the **System Dashboard** button  on the clearVISN CoreWatch toolbar.

A **Select Dials to be Displayed** dialog box similar to the following appears:

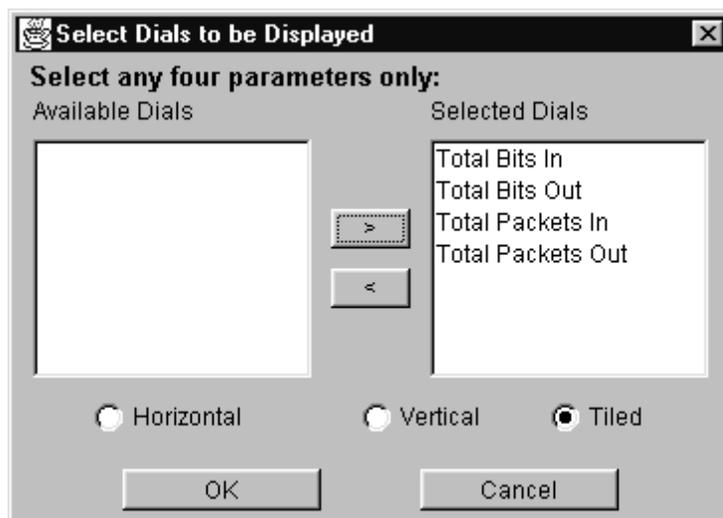


Figure 235. Select Dials to be Displayed dialog box

3. Specify which dials you want to display by selecting the desired dial names from the list box on the left. The following table identifies which dial each of the options displays. You can select up to four of these dials.

Select	To display a dial that indicates
Total Bits In	Number of bits per second the GSR is receiving.
Total Bits Out	Number of bits per second the GSR is sending.
Total Packets In	Number of packets per second the GSR is receiving.
Total Packets Out	Number of packets per second the GSR is sending.

To select more than one dial name in Windows 95 or Windows NT, hold down the Ctrl key while making your selections.

4. Click the > button.

clearVISN CoreWatch moves the selected dial names to the list box on the right. If you accidentally selected a name of a dial that you do not want to display, move it back to the left list box by selecting it and then clicking the < button.

- If you selected multiple dials, specify how you want clearVISN CoreWatch to display the dials by selecting one of the options described in the following table:

Select	To display the dials
Horizontal	In a single row.
Vertical	In a single column.
Tiled	In multiple rows and columns. The number of rows and columns will depend on how many dials the GSR is displaying.

- Click **OK**.

A **System Dashboard** dialog box similar to the following figure appears. The dials clearVISN CoreWatch displays depends on which dials you selected in the **Select Dials to be Displayed** dialog box.

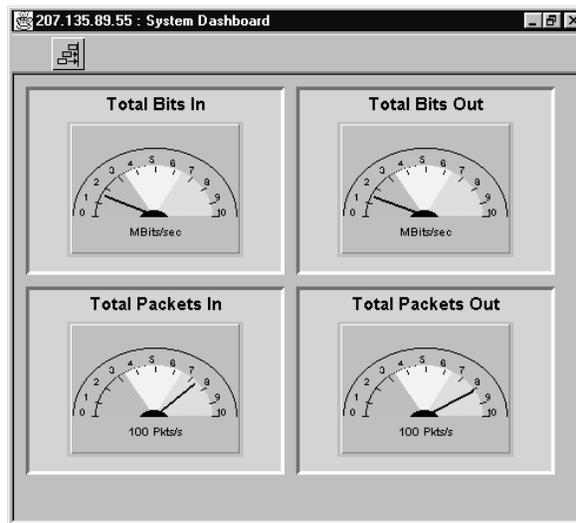


Figure 236. System Dashboard dialog box

Setting the Scaling of Dials

clearVISN CoreWatch automatically adjusts the scales of the System Dashboard dials. For example, clearVISN CoreWatch will automatically change the scaling of the Bits Out dial from KBits/Sec to 10 KBits/Sec whenever the GSR is sending more than 10 Kilobits per second. If you want to manually set the scaling of dials rather than having clearVISN CoreWatch automatically adjust it, take the following steps:

1. Click the **Options**  button, which is located at the top of the **System Dashboard** dialog box.
2. In the **Dial Options** dialog box that appears, select the check boxes of each dial you want to scale and also select the desired scale for those dials.

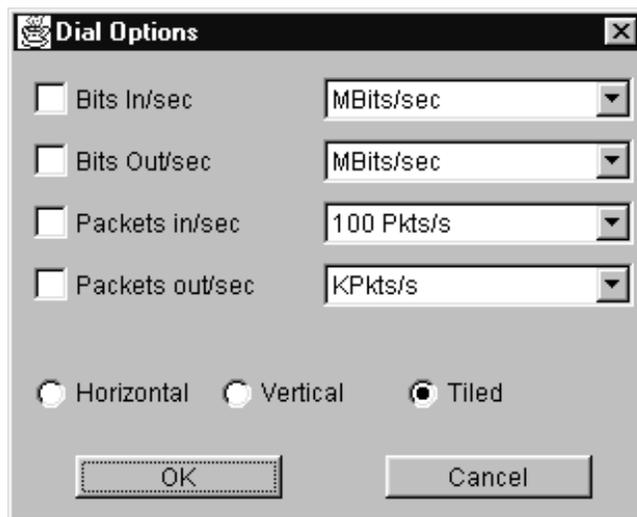


Figure 237. Dial Options dialog box

3. If you selected multiple dials, specify the manner in which you want clearVISN CoreWatch to display the dials by selecting one of the options described in the following table:

Select	To display the dials
Horizontal	In a single row.
Vertical	In a single column.
Tiled	In multiple rows and columns. The number of rows and columns will depend on how many dials the GSR is displaying.

4. Click **OK**.

Monitoring Port Utilization

clearVISN CoreWatch lets you obtain a Port Utilization Summary that provides the status of each GSR port and identifies the percentage of traffic being transmitted and received on each of those ports. To display such a summary:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.

Do one of the following:

- Click the Port Utilization Summary button  on the clearVISN CoreWatch toolbar.
- Select the **Monitor** menu, choose **Performance**, and then choose **Port Utilization Summary**.

A **Port Utilization Summary** dialog box similar to the one described in the following figure appears:

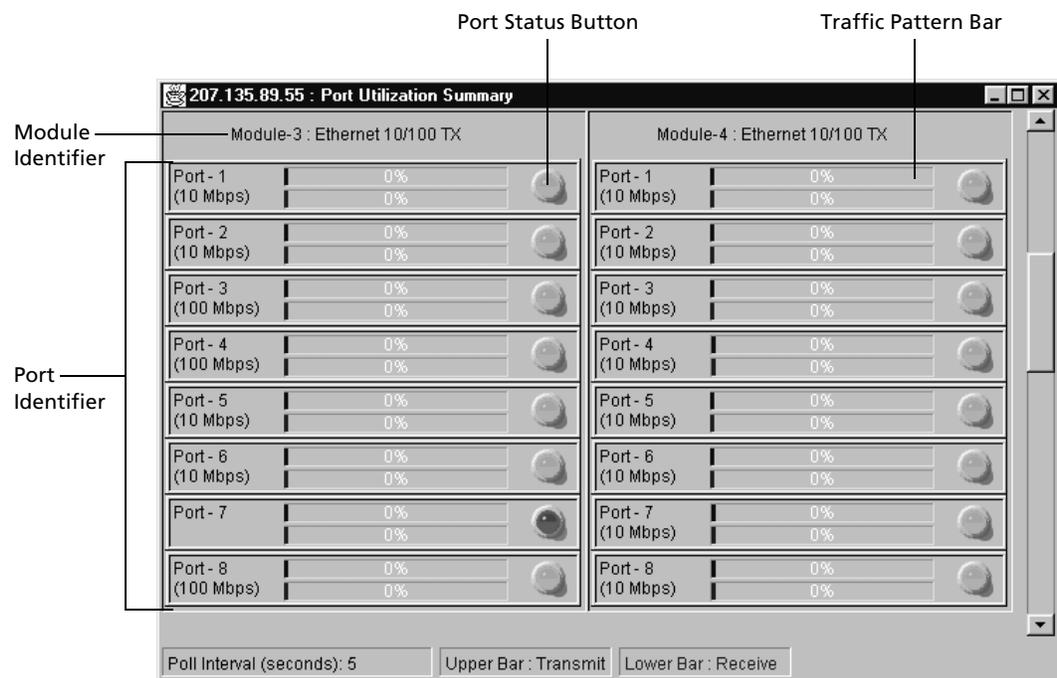


Figure 238. Port Utilization Summary dialog box

The following table describes the different items of the **Port Utilization Summary** dialog box:

Table 43. Port Utilization Summary items

Item	Description
Module Identifier	Indicates which of the following Ethernet modules are installed in a GSR slot: <ul style="list-style-type: none"> • 10/100-TX • 100-FX • Gigabit-LX • Gigabit-SX
Port Identifier	Identifies each port's number and indicates the speed of those ports.
Port Status Button	Indicates whether a port is online or offline. Bright green indicates a port is online, and dark green indicates a port is offline. There is a separate button for each port.
Traffic Pattern Bar	Identifies the percentage of traffic being transmitted and received on that port. The upper bar represents how much of a port's traffic is outgoing. The lower bar represents how much of a port's traffic is incoming. A port's high watermark indicates the maximum amount of outgoing and incoming traffic attained on that port.

Obtaining Statistics About an Individual Port

You can obtain statistics about the number of

- unicast, multicast, and broadcast packets a port is sending or receiving.
- bytes a port is sending and receiving.
- erroneous packets that are sent or received on a port as well as how many of those erroneous packets the GSR discarded.

clearVISN CoreWatch displays an individual port statistics in graphs. Details on obtaining these statistics and the different port-statistics graphs follow.

Obtaining Packet Statistics

You can obtain statistics about a port's incoming and outgoing unicast, multicast, and broadcast packets. To display such information:

1. In the Front Panel view, click the port that you want to monitor.
2. Do one of the following:
 - Double-click the port that you want to monitor.
 - Click the port that you want to monitor. Then select the **Monitor** menu, choose **Performance**, choose **Port**, and then choose **Packet Statistics**.
 - Click the port that you want to monitor. Then click the right mouse button and choose **Packet Statistics** from the pop-up menu that appears.

A **Port Packet Statistics** graph similar to the following appears. By examining the graph, you can determine how many of the different packets are being sent or received on a port and the time at which those packets were sent or received.

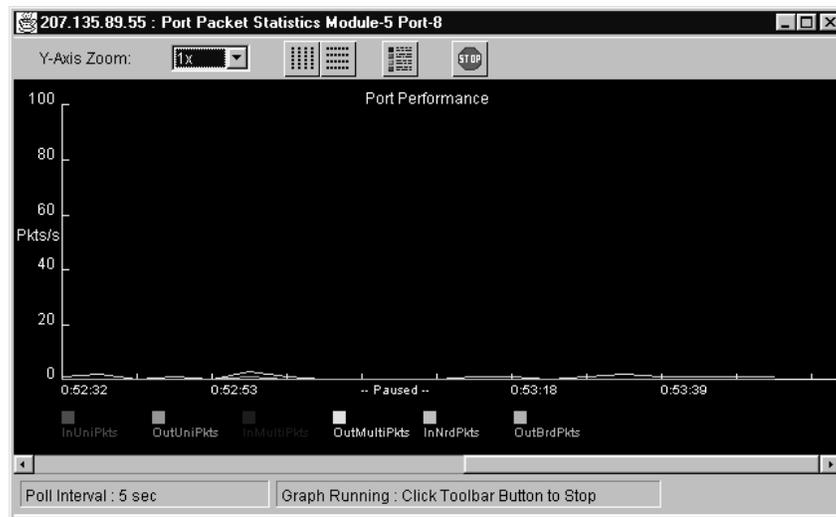


Figure 239. Port Packet Statistics graph

The following table describes the abbreviations used in the legend located at the bottom of the graph. You can control the graph's appearance and pause or resume the

gathering of statistics by using the Graph toolbar as discussed in “Using the Graph Toolbar” on page 371.

Table 44. Port Packet Statistics graph abbreviations

Abbreviation	Description
InUniPkts	Incoming unicast packets.
OutUniPkts	Outgoing unicast packets.
InMultiPkts	Incoming multicast packets.
OutMultiPkts	Outgoing multicast packets.
InBrdPkts	Incoming broadcast packets.
OutBrdPkts	Outgoing broadcast packets.

Obtaining Port Byte Statistics

You can obtain statistics about how many bytes a port is sending and receiving. To display such information:

1. In the Front Panel view, click the port that you want to monitor.
2. Do one of the following:
 - Select the **Monitor** menu, choose **Performance**, choose **Port**, and then choose **Byte Statistics**.
 - Click the right mouse button and choose **Byte Statistics** from the pop-up menu that appears.

A **Port Byte Statistics** graph similar to the following appears. By examining the graph, you can determine how many bytes are being sent or received on a port and the time at which they were sent or received.

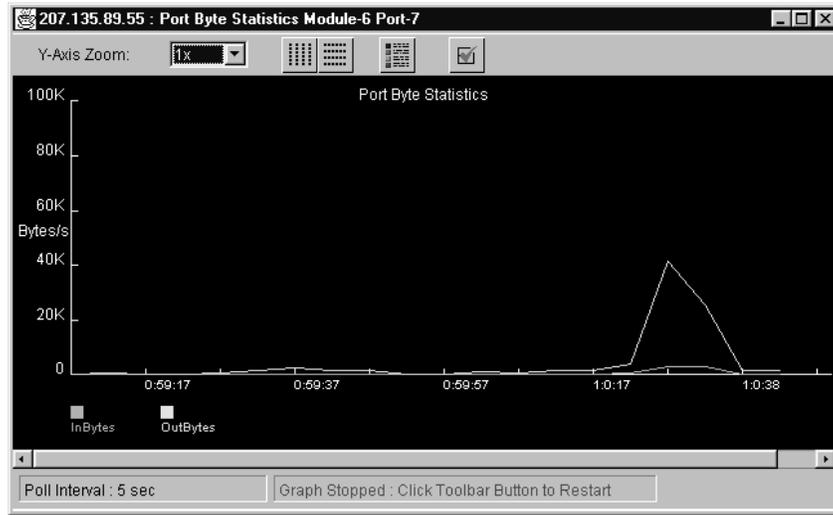


Figure 240. Port Byte Statistics graph

The following table describes the abbreviations used in the legend located at the bottom of the graph. You can control the graph's appearance and pause or resume the gathering of statistics by using the Graph toolbar as discussed in [“Using the Graph Toolbar” on page 371](#).

Table 45. Port Byte Statistics graph abbreviations

Abbreviation	Description
InBytes	Bytes received on the port.
OutBytes	Bytes sent out on the port.

Obtaining Port Error Statistics

You can obtain statistics about the erroneous packets that are sent or received on a port as well as how many of those erroneous packets the GSR discarded. To display such information:

1. In the Front Panel view, click the port that you want to monitor.
2. Do one of the following:
 - Select the **Monitor** menu, choose **Performance**, choose **Port**, and then choose **Error Statistics**.
 - Click the right mouse button and choose **Error Statistics** from the pop-up menu that appears.

A **Port Error Statistics** graph similar to the following appears. By examining the graph, you can determine how many of the different errors a port is experiencing and the time at which those errors occurred.

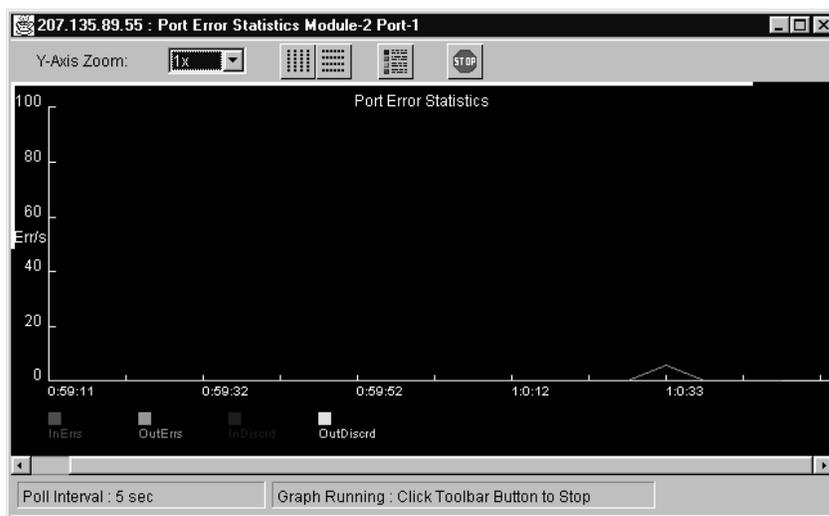


Figure 241. Port Error Statistics graph

The following table describes the abbreviations used in the legend located at the bottom of the graph. You can control the graph's appearance and pause or resume the gathering of statistics by using the Graph toolbar as discussed in [“Using the Graph Toolbar”](#) on page 371.

Table 46. Port Error Statistics graph abbreviations

Abbreviation	Description
InErrs	Erroneous packets received on the port.
OutErrs	Erroneous packets sent out from the port.
InDiscrd	Incoming packets the GSR discarded although no errors had been detected to prevent the GSR from delivering them to a higher-layer protocol. Freeing buffer space is one possible reason for discarding such packets.
OutDiscrd	Outgoing packets the GSR discarded although no errors had been detected to prevent the GSR from sending them elsewhere. Freeing buffer space is one possible reason for discarding such packets.

Monitoring IP Interface Statistics

clearVISN CoreWatch lets you obtain current statistics for

- the number of unicast, multicast, and broadcast packets being sent and received on the IP interfaces of a GSR.
- reassembly and fragmentation activities on the IP interfaces of a GSR.
- errors related to the IP data the GSR receives or sends.

clearVISN CoreWatch displays an IP interface statistics in graphs. Details on obtaining these statistics and the different IP interface graphs follow.

Obtaining IP Packet Statistics

Obtain IP packet statistics if you want to determine how many IP packets the GSR sent, received, or forwarded. To display such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Performance**, choose **IP**, and then choose **Packet Statistics**.

An **IP Packet Statistics** graph similar to the following appears. By examining the graph, you can determine how many of the different types of packets were sent, received, or forwarded on a GSR and the time at which they were sent, received, or forwarded.

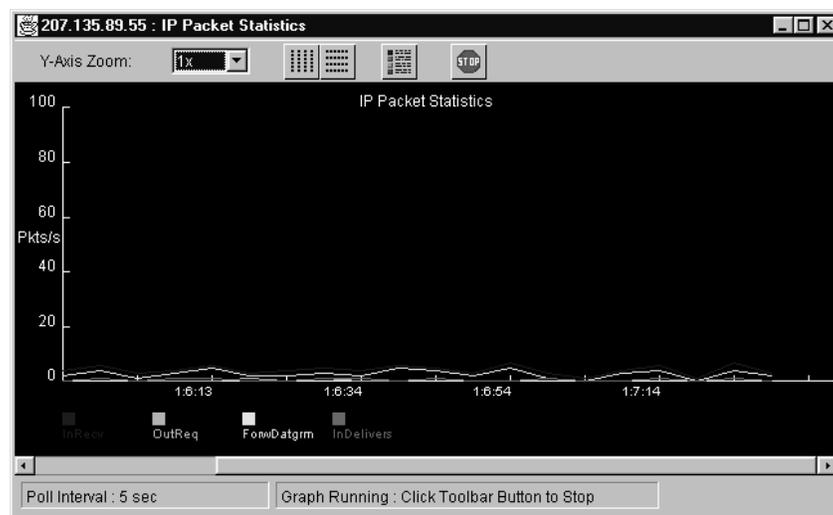


Figure 242. IP Packet Statistics graph

The following table describes the abbreviations used in the legend located at the bottom of the graph. You can control the graph's appearance and pause or resume the gathering of statistics by using the Graph toolbar as discussed in [“Using the Graph Toolbar” on page 371](#).

Table 47. IP Packet Statistics graph abbreviations

Abbreviation	Description
InRecv	IP packets the GSR received, including those that have errors.
OutReq	Packets that various IP protocols (including ICMP) supplied to IP for the GSR to send out. The statistic for these packets does not include any packets the GSR forwards.
ForwDatgrm	Incoming IP packets that the GSR received and then forwarded because the GSR was not their final destination.
InDelivers	Incoming IP packets the GSR successfully delivered to IP user-protocols (including ICMP).

Obtaining IP Reassembly Statistics

Data packets too large to be transmitted are broken down into fragments and then reassembled once they arrive at their destination. Obtain IP reassembly statistics if you want to examine the reassembling and fragmenting activities a GSR is performing on IP packets. To display such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Performance**, choose **IP**, and then choose **Reassembly Statistics**.

An **IP Reassembly Statistics** graph similar to the following appears. By examining the graph, you can determine how well the GSR is reassembling and fragmenting information and when the reassembly or fragmenting activities occurred.

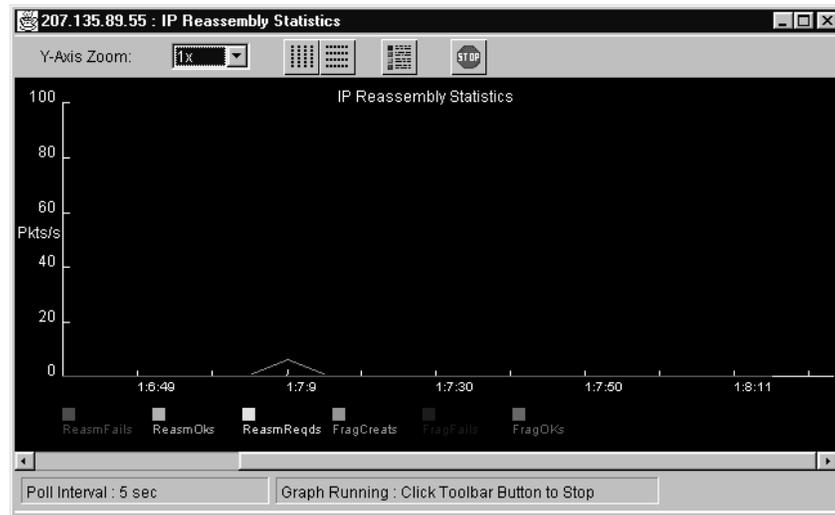


Figure 243. IP Reassembly Statistics graph

The following table describes the abbreviations used in the legend located at the bottom of the graph. You can control the graph's appearance and pause or resume the gathering of statistics by using the Graph toolbar as discussed in [“Using the Graph Toolbar” on page 371](#).

Table 48. IP Reassembly Statistics graph abbreviations

Abbreviation	Description
ReasmFails	Reassembles that failed because of time outs, errors, or other problems.
ReasmOKs	IP datagrams the GSR successfully reassembled.
ReasmReqds	IP fragments the GSR received and needs to be reassembled.
FragCreates	IP fragments the GSR created while breaking down datagrams.
FragFails	IP datagrams that needed to be fragmented but were discarded by the GSR because the Don't Fragment flag was set on those datagrams.
FragOKs	IP datagrams the GSR successfully fragmented.

Obtaining IP Error Statistics

Obtain IP error statistics for a GSR if you want to determine how many datagrams that GSR discarded because there was not enough buffer space or there were problems that prevented the processing of datagrams. To display such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Performance**, choose **IP**, and then choose **Error Statistics**.

An **IP Error Statistics** graph similar to the following appears. By examining the graph, you can determine how many of the different IP errors a GSR is experiencing and the time at which they occurred.

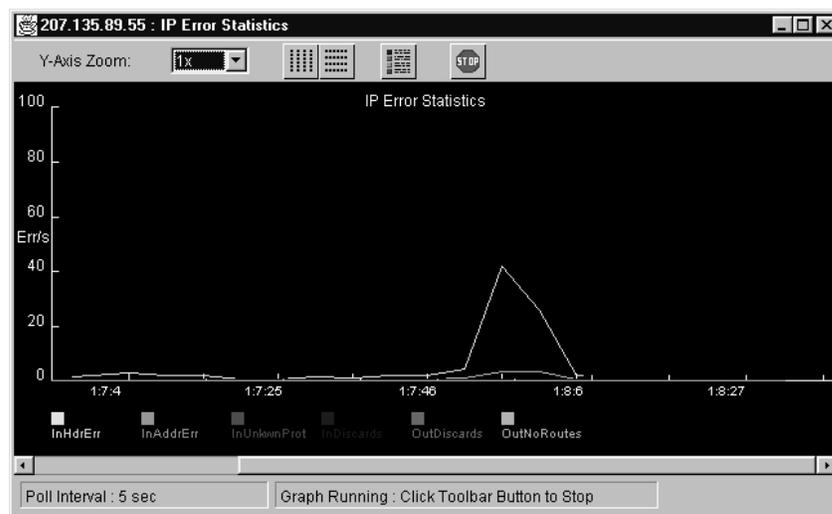


Figure 244. IP Error Statistics graph

The following table describes the abbreviations used in the legend located at the bottom of the graph. You can control the graph's appearance and pause or resume the

gathering of statistics by using the Graph toolbar as discussed in [“Using the Graph Toolbar” on page 371](#).

Table 49. IP Error Statistics graph abbreviations

Abbreviation	Description
InHdrErr	Incoming IP datagrams the GSR discarded because of problems in their headers (such as bad checksums, version number mismatching and other formatting errors, and problems in processing IP options).
InAddrErr	Incoming IP datagrams the GSR discarded because the IP header's destination field included an invalid IP address (such as 0.0.0.0) or an IP address of an unsupported class (such as Class E).
InUnkownProt	Incoming IP datagrams addressed to the GSR that were received successfully but were discarded because of an unknown or unsupported protocol.
InDiscards	Incoming IP datagrams the GSR discarded although no errors had been detected to prevent the GSR from processing them. Freeing buffer space is one possible reason for discarding such datagrams. This statistic does not include any datagrams discarded while awaiting reassembly.
OutDiscards	Outgoing IP datagrams the GSR discarded because of lack of buffer space.
OutNoRoutes	Outgoing or forwarded IP datagrams discarded because no route could be found to transmit them to their destination. The statistics shown for this item also include any datagrams the GSR could not route because its default gateways are down.

Monitoring IPX Interface Statistics

clearVISN CoreWatch lets you obtain current statistics for

- the number of incoming and outgoing IPX packets.
- the number of IPX packets that had bad header information, incorrect checksums, invalid structures, or other problems.

Obtaining IPX Packet Statistics

Obtain IPX packet statistics if you want to determine the following:

- How many IPX packets the GSR received and how many of those packets were delivered on the GSR.
- How many times the GSR was requested to send IPX information and how many IPX packets the GSR actually sent.

To display IPX packet statistics:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Performance**, choose **IPX**, and then choose **Packet Statistics**.

An **IPX Packet Statistics** graph similar to the following appears. By examining the graph, you can determine how many of the different types of packets the GSR received or delivered and the time at which those packets were received or delivered. You can also determine how many times IPX data was requested and how many packets the GSR successfully sent.

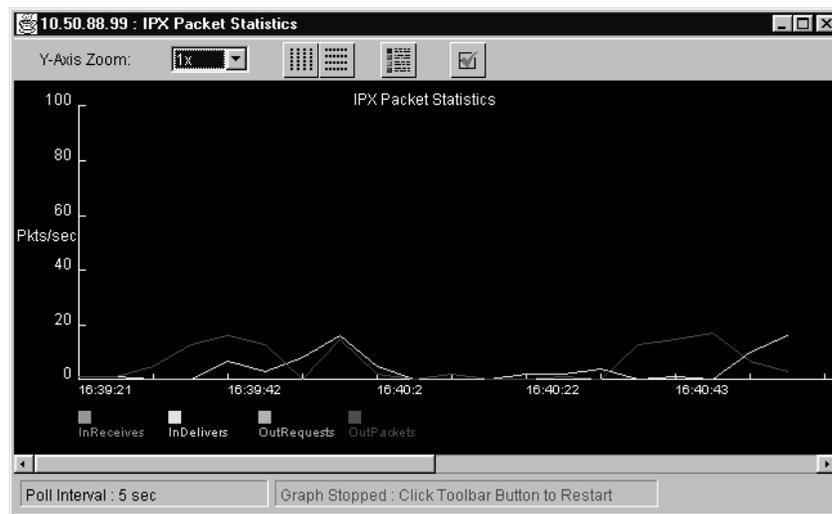


Figure 245. IPX Packet Statistics graph

The following table describes the abbreviations used in the legend located at the bottom of the graph. You can control the graph's appearance and pause or resume the

gathering of statistics by using the Graph toolbar as discussed in [“Using the Graph Toolbar” on page 371](#).

Table 50. IPX Packet Statistics graph abbreviations

Abbreviation	Description
InReceives	Total of all IPX packets the GSR received. This includes packets that have errors.
InDelivers	IPX packets that were destined for the GSR and were successfully delivered.
OutRequests	Number of requests made to the GSR to send out IPX data. This statistic does not include any packets the GSR forwards.
OutPackets	Number of IPX packets that were successfully sent from the GSR.

Obtaining IPX Error Statistics

Obtain IPX error statistics for a GSR if you want to determine how many IPX packets had bad header information, incorrect checksums, invalid structures, or other problems. The IPX error statistics also indicates how many times no route destination was found for IPX information. To display IPX error statistics:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Performance**, choose **IPX**, and then choose **Error Statistics**.

An **IPX Error Statistics** graph similar to the following appears. By examining the graph, you can determine how many of the different IPX errors the GSR is experiencing and the time at which they occurred.

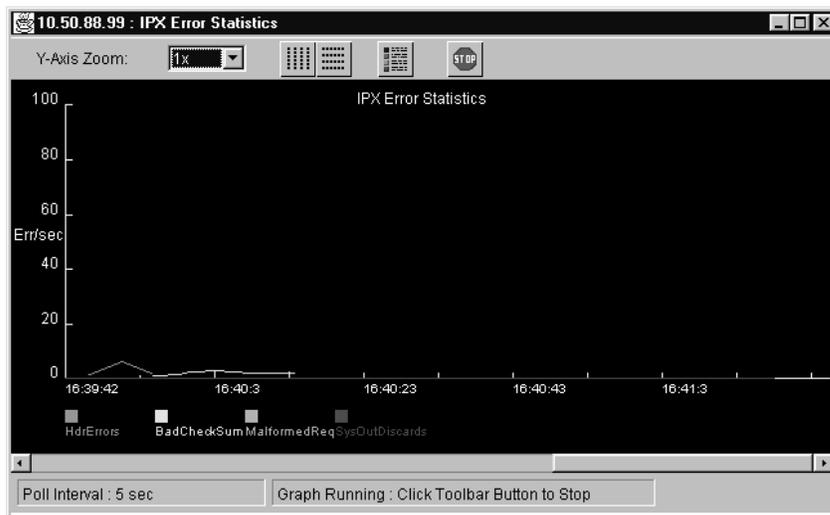


Figure 246. IPX Error Statistics graph

The following table describes the abbreviations used in the legend located at the bottom of the graph. You can control the graph's appearance and pause or resume the gathering of statistics by using the Graph toolbar as discussed in [“Using the Graph Toolbar” on page 371](#).

Table 51. IPX Error Statistics graph abbreviations

Abbreviation	Description
HdrErrors	Incoming IPX packets the GSR discarded because those packets had problems in their headers. This includes any IPX packets with a size less than the 30-byte minimum.
BadChecksum	Incoming IPX packets the GSR received with incorrect checksums.
MalformedReq	Outgoing IPX packets supplied locally (by the GSR) that contain errors in their structure.
SysOutDiscards	Outgoing IPX packets the GSR discarded because of errors. This does not include those packets accounted for by the OutMalformed statistic.

Using the Graph Toolbar

A toolbar is located at the top of each port-statistics graph. Use this toolbar to control a graph's appearance and stop or start the gathering of statistics as summarized in the following figure:

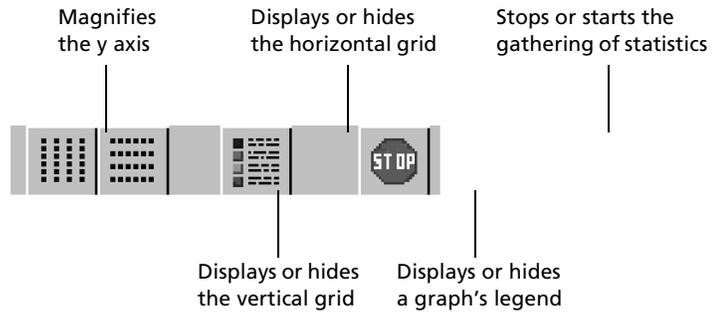


Figure 247. Graph toolbar

Chapter 19

Checking the Status of Bridge Tables

clearVISN CoreWatch lets you obtain tables that contain the following:

- Information about which ports are associated with which VLANs and the type of module on which those ports are located.
- Information about the ports on which STP is enabled.

Obtaining VLAN Information

Obtain VLAN information to display information about which ports and modules are associated with the VLANs configured on a GSR. To access such information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Bridging State**, and then choose **VLAN Table**.
- In the Schematic view, double-click the VLAN object.

A **VLAN** table similar to the following appears:

VLAN Name	Port ID	Module ID
VLAN CONTROL	2	7
VLAN BR-GIG-BACKBONE	2	7
VLAN GIG-BACKBONE	2	7
VLAN CONTROL	1	7
VLAN BR-MAINSUBNET	1	7
VLAN MAINSUBNET	1	7

Selection Details	
VLAN N...	VLAN CONTROL
Port ID	2
Module ID	7

126 rows displayed

Figure 248. VLAN table

The following table describes the fields of the **VLAN** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 52. VLAN table fields

Field	Description
VLAN Name	Name of the VLAN for which clearVISN CoreWatch is displaying information.
Port ID	Identifies the number of the port associated with the VLAN.
Module ID	Identifies the port's slot number.

Obtaining STP Port Information

Obtain STP Port information if you want to examine STP information about the ports on which STP is enabled. To access such information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Bridging State**, and then choose **STP Table**.
- In the Schematic view, double-click the STP Bridging function.

An **STP Port** table similar to the following appears:

The screenshot shows a window titled "207.135.89.55 : STP Port Table". It contains a table with the following columns: Port Number, Port Priority, State, Path Cost, Designated Root, Designated Cost, and Description. The table lists ports 9 through 13, all of which are disabled. Below the table is a "Selection Details" pane showing the properties of the selected port (Port Number: 9, Port Priority: 0, State: disabled, Path Cost: 1, Designated Root: 00 00 00 00 00 00 00, Designated Cost: 0).

Port Number	Port Priority	State	Path Cost	Designated Root	Designated Cost	Description
9	0	disabled	1	00 00 00 00 00 00...	0	00 00
10	0	disabled	1	00 00 00 00 00 00...	0	00 00
11	0	disabled	1	00 00 00 00 00 00...	0	00 00
12	0	disabled	1	00 00 00 00 00 00...	0	00 00
13	0	disabled	1	00 00 00 00 00 00...	0	00 00

Selection Details

Port Number	9
Port Priority	0
State	disabled
Path Cost	1
Designated Root	00 00 00 00 00 00 00
Designated Cost	0

42 rows displayed

Figure 249. STP Port table

The following table describes the fields of the **STP Port** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 53. STP Port table fields

Field	Description
Port Number	Identifies the port for which the table is providing STP information.
Port Priority	Indicates the priority assigned to the port. The value of this field is a number from 0 (highest priority) to 255 (lowest priority). The STP algorithm uses the port priority to determine which port to use if a bridge has two ports connected to a loop.

Table 53. STP Port table fields

Field	Description
State	<p>Indicates whether or not the port is functioning. If the port is functioning, this field identifies the port's state as one of the following:</p> <ul style="list-style-type: none"> • Blocking – signifies the GSR is not accepting any incoming frames except STP BPDUs. • Learning – signifies the GSR is learning the MAC addresses of packets it receives on the port, but is not forwarding those packets. • Listening – signifies the GSR is receiving frames but is not learning them. • Forwarding – signifies the GSR is both learning and forwarding packets it receives on the port. • Broken – signifies the port is broken. • Disabled – signifies the port is disabled. The GSR drops all frames including BPDUs.
Path Cost	Indicates the cost assigned to the port.
Designated Root	Identifies the root bridge.
Designated Cost	Indicates the path cost of the designated port.
Designated Bridge	Identifies which bridge is the designated bridge on the LAN.
Designated Port	Identifies which port on the designated bridge is used for forwarding data.
Forward Transitions	Indicates how many times the port has changed from the learning state to the forwarding state.

Obtaining L2 Interface Information

Obtain L2 Interface information if you want to examine which ports are currently serving as L2 interfaces, to which VLAN those ports belong, and the port identification (number). To access such information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Bridging State**, and then choose **L2 Interface**.

An **L2 Interface** table similar to the following appears:

VLAN Name	Port Type	Port Number
DEFAULT	gigether	gi.7.2
DEFAULT	gigether	gi.7.1
DEFAULT	ethernet	et.6.8
DEFAULT	ethernet	et.6.7
DEFAULT	ethernet	et.6.6
DEFAULT	ethernet	et.6.5
DEFAULT	ethernet	et.6.4
DEFAULT	ethernet	et.6.3
DEFAULT	ethernet	et.6.2
DEFAULT	ethernet	et.6.1
ACTONET	ethernet	et.1.7

15 rows displayed

Figure 250. L2 Interface table

The following table describes the fields of the **L2 Interface** table's upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 54. L2 Interface table fields

Field	Description
VLAN Name	Displays the name of the VLAN to which the L2 interface in question belongs.
Port Type	Displays the type of L2 interface in question. For example, "ethernet" or "gigether".
Port Number	Identifies the ports currently serving as L2 interfaces.

Chapter 20

Checking the Status of Routing Tables

clearVISN CoreWatch can display tables that include the following routing data:

- Details about IP, IPX, OSPF, RIP, DVMRP, and IGMP interfaces.
- Information about the routing of IP and IPX packets the GSR forwarded.
- Information about OSPF areas, neighboring routers, and link-state advertisements.
- Information about RIP peers.
- Information about DVMRP neighbors, routes, and hops.
- Details about IGMP caching.

Checking IP Routing Status

clearVISN CoreWatch can display the following IP information:

- The ports and addresses of each IP interface of a GSR.
- Details about an IP route's destination, the next hop on the route, whether a route is valid, whether a valid route is to a local or remote device, which protocol was used to send IP the route information, and other data useful for resolving routing problems.

Obtaining IP Interface Information

Obtain IP interface information to display details about the ports and addresses of each IP interface of a GSR. To access such information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Routing State**, choose **IP State**, and then choose **IP Interface Table**.
- In the Schematic view, double-click the IP Interface object.

An **IP Interface** table similar to the following appears:

Interface Name	Bound IP Address	IP Subnet Mask	Broadcast IP Address	Port ID	Module ID
GIG-BACKBONE	207.135.89.145	255.255.255.240	207.135.89.159	2	7
MAINSUBNET	207.135.89.55	255.255.255.128	207.135.89.127	1	7
MAINSUBNET	207.135.89.55	255.255.255.128	207.135.89.127	8	6
MAINSUBNET	207.135.89.55	255.255.255.128	207.135.89.127	7	6
MAINSUBNET	207.135.89.55	255.255.255.128	207.135.89.127	6	6

Selection Details	
Interface Name	GIG-BACKBONE
Bound IP Address	207.135.89.145
IP Subnet Mask	255.255.255.240
Broadcast IP Address	207.135.89.159
Port ID	2
Module ID	7

42 rows displayed

Figure 251. IP Interface table

The following table describes the fields of the **IP Interface** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 55. IP Interface table fields

Field	Description
Interface Name	Identifies the name of the interface.
Bound IP Interface	Indicates the IP address of the interface agent to which this interface is bound.
IP Subnet Mask	Identifies the subnetwork mask associated with the interface's IP address.

Table 55. IP Interface table fields (Continued)

Field	Description
Broadcast IP Address	Identifies the destination address the interface uses when sending broadcast packets.
Port ID	Identifies the number of the port the interface uses.
Module ID	Identifies the slot number of the interface's port.
VLAN Name	Indicates which VLAN is associated with the interface.

Obtaining IP Forwarding Information

Obtain IP forwarding information to display information about the routes used by the IP interfaces configured on a GSR. The table lists each IP route's destination and provides details about the next hop on a route, whether a route is valid, whether a valid route is to a local or remote device, which protocol was used to send IP the route information, and other data useful for resolving routing problems.

To access IP forwarding information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Routing State**, choose **IP State**, and then choose **IP Forwarding Table**.
- In the Schematic view, double-click the IP Router function.

An **IP Forwarding** table similar to the following appears:

Dest. IP Address	Dest. IP Mask	Route TOS	Next Hop IP Address	Local IP Interface	Route Type
0.0.0.0	0.0.0.0	0	207.135.89.135	4346	remote
127.0.0.0	255.0.0.0	0	0.0.0.0	0	other
127.0.0.1	255.255.255.255	0	127.0.0.1	4344	local
207.135.64.0	255.255.255.0	0	207.135.89.135	4346	remote
207.135.64.82	255.255.255.255	0	207.135.89.135	4346	remote

Selection Details	
Dest. IP Address	0.0.0.0
Dest. IP Mask	0.0.0.0
Route TOS	0
Next Hop IP Address	207.135.89.135
Local IP Interface	4346
Route Type	remote

9 rows displayed

Figure 252. IP Forwarding table

The following table describes the fields of the **IP Forwarding** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 56. IP Forwarding table fields

Field	Description
Dest IP Address	Identifies the destination IP address of the route. An entry with a value of 0.0.0.0 is considered a default route.
Dest IP Mask	Identifies the subnet mask associated with the destination IP address.
Route TOS	Indicates the route's TOS.
Next Hop IP Address	Identifies the IP address of the next hop of this route. (If the route is bound to an interface that is accessed through a broadcast medium, the value of this field is the agent's IP address on that interface.)
Local IP Interface	Indicates the local interface through which the next hop of the route should be reached.
Route Type	Indicates the type of route. <ul style="list-style-type: none"> • Remote—signifies an indirect path going through a non-local host, network, or subnetwork. • Local—signifies a path connected directly to a network or subnetwork. • Invalid—signifies an invalid path. • Other—signifies a path that is not remote, local, or invalid.
Routing Protocol	Indicates the routing protocol through which the route was learned.
Route Age	Indicates how many seconds have elapsed since this route was last updated or otherwise determined to be correct.
Next Hop AS	Identifies the IP address of the next hop of this route. (In the case of a route bound to an interface which is realized through a broadcast media, the value of this field is the agent's IP address on that interface.)

Checking IPX Routing Status

clearVISN CoreWatch can display the following IPX information:

- Details about the packets sent from each IPX interface including information which WAN router an interface uses.
- Routing information of a GSR's IPX interfaces. This includes a list of each IPX destination and provides details about the routes IPX interfaces use to reach those destinations.

Obtaining IPX Interface Information

Obtain IPX Interface information to display details about each IPX interface configured on the GSR. This information identifies each IPX interface, provides details about the packets sent from each of those interfaces, and indicates which WAN router an interface may use. To access such information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Routing State**, choose **IPX State**, and then choose **IPX Interface Table**.
- In the Schematic view, double-click the IPX Interface object.

An **IPX Interface** table similar to the following appears:

IPX Identifier	Index	Exist State	Operational State	Interface Index	Name
1	7	on	up	7	ipx-100...
1	8	on	down	8	ipx-backup
1	9	on	up	9	ipx-100...
1	10	on	up	10	ipx1001
1	11	on	down	11	ipx1002

Selection Details	
IPX Identifier	1
Index	7
Exist State	on
Operational State	up
Interface Index	7
Name	ipx.1000L.0000

5 rows displayed

Figure 253. IPX Interface table

The following table describes the fields of the **IPX Interface** table's upper frame. The Selection Details frame displays information about the item currently selected in the

upper frame. For details on using the Table toolbar, see [Appendix A: “Working with Tables” on page 425](#).

Table 57. IPX Interface table fields

Field	Description
IPX Identifier	Indicates which instance of IPX the interface is using.
Index	Indicates the number IPX uses to identify the interface.
Exist State	Indicates whether the interface is valid.
Operational State	Indicates whether the interface is up, down, or sleeping. If the interface is up, it may be used to send IPX data. If the interface is down, it cannot be used to send IPX data. If the interface is sleeping, it is waiting for data.
Interface Index	Indicates the interface's index value.
Name	Identifies the interface's name.
Type	Indicates the interface's type as one of the following: broadcast, point-to-point, WAN RIP, unnumbered RIP, dynamic, WAN WS, or other.
Dialing Name	Indicates the symbolic name IPX uses to reference the dialing information used to create this circuit. This field is blank for LAN interfaces.
Max Packet Size	Indicates the maximum number of bytes that the GSR supports locally for packets (including headers) on this circuit.
Compression State	Indicates whether data compression is enabled on the interface.
Compression Slots	Identifies how many compression slots area available on the interface.
Static Status	Indicates whether the information about static routes and services reached through this interface matches that saved in permanent storage (current).
Compressed Tx	Indicates how many compressed packets were sent on the interface.
Compressed Init Tx	Indicates how many compressed packets were requested to be sent from the interface.
Compressed Rejected Tx	Indicates how many compressed packets sent on the interface were rejected.

Table 57. IPX Interface table fields (Continued)

Field	Description
Uncompressed Tx	Indicates how many packets were sent without being compressed even though compression was turned on for the interface.
Compressed Rx	Indicates how many compressed packets were received.
Compressed Init Rx	Indicates how many compression initialization packets were received.
Compressed Rejected Rx	Indicates how many compressed packets received on the interface were rejected.
Uncompressed Rx	Indicates how many packets were received without being compressed even though data compression was enabled on the interface.
Media Type	Indicates that the interface is for an Ethernet network
IPX Network number	Identifies the IPX network number of the interface.
No of State Changes	Indicates how many times the interface has changed state.
Initialization failed times	Indicates how many times the interface failed to start.
Delay (mS)	Indicates how many milliseconds it takes to transmit one byte of data (excluding protocol headers) to a destination on the other end of the interface if there is no other traffic on that interface.
Throughput	Indicates how much data (in bits per second) may flow through the interface if there is no other traffic on that interface.
Neighbor Name	Identifies the name of the neighboring router located on a WAN.
Neighbor Network Number	Identifies the internal network number of the neighboring router located on a WAN.

Obtaining IPX Forwarding Information

Obtain IPX Forwarding information to display details about the routes of a GSR's IPX interfaces. To access such information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Routing State**, choose **IPX State**, and then choose **IPX Forwarding Table**.
- In the Schematic view, double-click the IPX Router function.

An **IPX Forwarding** table similar to the following appears:

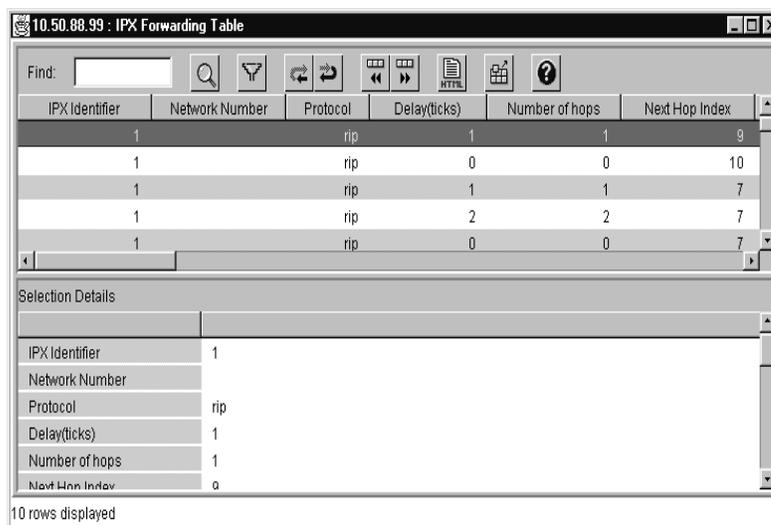


Figure 254. IPX Forwarding table

The following table describes the fields of the **IPX Forwarding** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 58. IPX Forwarding table fields

Field	Description
IPX Identifier	Indicates which instance of IPX the interface is using.
Network Number	Identifies the IPX network number of the destination.
Protocol	Indicates the routing protocol through which the route was learned.
Delay (ticks)	Indicates how many ticks it takes to reach the destination.
Number of hops	Indicates the number of hops necessary to reach the destination.
Next Hop Index	Indicates the unique identifier of the route used to reach the next hop.
Next Hop NIC Address	Identifies the NIC address of the next hop.
Next Hop Network Number	Identifies the IPX network address of the next hop.

Checking OSPF Routing Status

clearVISN CoreWatch can display the following information for OSPF routes that can be configured through CLI commands:

- The OSPF interfaces configured on a GSR.
- The areas that the GSR can communicate with.
- The neighboring routers of a GSR.
- The link-state advertisements (LSAs) of the areas to which the GSR is attached.

An LSA is a packet that provides information about neighboring routers and adjacencies to OSPF. The collected link-state advertisements of all routers and networks forms the protocol's topological database. Routers use the information included in link-state advertisements to update their routing tables.

- The IP address and TOS data that OSPF uses to determine the destination of packets.

Note: For details on OSPF, see [“A Look at OSPF Routing on the GSR” on page 229](#).

Obtaining OSPF Interface Information

Obtain OSPF interface information if you want details about a GSR's OSPF interfaces. To access this information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Routing State**, choose **OSPF State**, and then choose **OSPF Interface Table**.
- In the Schematic view, double-click the OSPF Routing function.

An **OSPF Interface** table similar to the following appears:

OSPF Interface Address	OSPF Address-less Interface	OSPF Interface Area ID	Type	Admin. Status
207.135.89.145	0	ANY	broadcast	enabled

OSPF Interface Address	207.135.89.145
OSPF Address-less Interface	0
OSPF Interface Area ID	ANY
Type	broadcast
Admin. Status	enabled
Print	A

1 rows displayed

Figure 255. OSPF Interface table

The following table describes the fields of the **OSPF Interface** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 59. OSPF Interface table fields

Field	Description
OSPF Interface Address	Identifies the interface's IP address.
OSPF Addressless Interface	Contains the interface index for those interfaces that do not have an IP address. A value of 0 is used in this field to identify those interfaces that have an IP address.
OSPF Interface Area ID	Identifies the area to which the interface is connected. An area ID of 0.0.0.0 is used for the OSPF backbone.
Type	Indicates whether the interface is attached to a broadcast, NBMA, or point-to-point network.
Admin Status	Indicates the interface's administrative status. If the interface is external to OSPF, the interface's status is listed as disabled.

Table 59. OSPF Interface table fields (Continued)

Field	Description
Priority	<p>Indicates the interface's priority. In multi-access networks, this field is used in the designated router election algorithm.</p> <p>The value 0 signifies that the router is not eligible to become the designated router on this particular network. In the event of a tie in this value, routers use their router ID as a tie breaker.</p>
Transit Delay	Indicates how many seconds it takes to transmit a link-state update packet over the interface.
Retransmission Interval	Indicates how many seconds elapse between link-state advertisement retransmissions for adjacencies belonging to this interface. The value in this field is also used when retransmitting database description and link-state request packets.
Hello Interval	Indicates how often the GSR sends hello packets on the interface. The value is in seconds and must be the same for all routes attached to the network.
Router-Dead Interval	<p>Indicates how many seconds that a router's Hello packets have not been seen before its neighboring routers declare the router down.</p> <p>The value of this field is a multiple of the Hello Interval field's value, and it must be the same for all routers attached to a common network.</p>
Poll Interval	<p>Indicates how often Hello packets are sent after a neighboring router becomes inactive because no Hello packets have been seen and the time specified in the Router-Dead Interval field has elapsed. If a neighboring router has become inactive, it may still be necessary to send Hello Packets to the dead neighbor.</p> <p>The value of this field should be much larger than that specified in the Hello Interval. For example, 2 minutes can be specified for an X.25 Packet Data Network (PDN).</p>
State	Indicates the OSPF state of the interface. The value of this field will be one of the following: Down, Loopback, Waiting, Point-to-Point, DR Other, Backup, or DR.
Designated Router	Lists the IP address of the designated router.

Table 59. OSPF Interface table fields (Continued)

Field	Description
Backup Designated Router	Lists the IP address of the backup designated router.
Number of Events	Indicates how many times the SPF interface has changed its state or that an error has occurred.
Authentication Key	<p>Identifies the interface's authentication key. If the value of the Authentication Type field is simple, the key size is restricted to 8 bytes.</p> <p>When read, this field's value always returns an octet string of length zero.</p>
Status	Indicates the status of the interface.
Multicast Forwarding	<p>Indicates the way multicast packets are forwarded on the interface. Possible values indicate that such packets are blocked (not forwarded), forwarded as data-link multicasts, or forwarded as data-link unicasts.</p> <p>Data link multicasting is not meaningful on point-to-point and NBMA interfaces, and setting this field to 0 effectively disables all multicast forwarding.</p>
Truth Value	Indicates whether demand OSPF procedures (hello suppression to full neighbors and setting the DoNotAge flag on propagated link-state advertisements) should be performed on the interface.
Authentication Type	<p>Indicates the type of authentication that is configured for the interface. The authentication type can be one of the following:</p> <ul style="list-style-type: none"> • Simple – signifies that the authentication method the interface uses is a simple password in which an authentication key of up to 8 characters is included in the packet. If this authentication method was chosen when the OSPF interface was created, a key-chain identifier should also have been specified. • MD5 – signifies that the interface uses the MD5 algorithm to create a crypto-checksum of an OSPF packet and an authentication key of up to 16 characters. • None – signifies no authentication method is associated with the interface.

Obtaining OSPF Area Information

Obtain OSPF area information if you want details about the configuration and cumulative statistics of the GSR's attached areas. To access such information:

1. If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Routing State**, choose **OSPF State**, and then choose **OSPF Area Table**.

An **OSPF Area** table similar to the following appears:

Area ID	Import As Extern	SPF Runs	Area Border Routers	Autonomous Sys. Bdr Routers
2.0.0.0	importExternal	2	0	2

Selection Details	
Area ID	2.0.0.0
Import As Extern	importExternal
SPF Runs	2
Area Border Routers	0
Autonomous Sys. Bdr Routers	2
I & Count	5

1 rows displayed

Figure 256. OSPF Area table

The following table describes the fields of the **OSPF Area** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 60. OSPF Area table fields

Field	Description
Area ID	Identifies the 32-bit integer that uniquely identifies the area. The OSPF backbone has an area ID of 0.0.0.0.
Import as Extern	Indicates whether the area supports importing autonomous system external link-state advertisements.
SPF Runs	Indicates how many times the intra-area route table has been calculated using the area's link-state database.

Table 60. OSPF Area table fields (Continued)

Field	Description
Area Border Routers	Indicates how many area border routers are reachable within the area. The value of this field is initially zero (0), and is calculated in each SPF run.
Autonomous Sys. Bdr Routers	Indicates how many autonomous system border routers are reachable within the area. The value of this field is initially zero (0), and is calculated in each SPF run.
LSA Count	Indicates how many link-state advertisements are in the area's link-state database. The value in this field does not include link-state advertisements of external autonomous systems.
LSA Checksums	Identifies the 32-bit unsigned sum of the link-state advertisements' link-state checksums contained in the area's link-state database. The value in this field excludes external (LS type 5) link-state advertisements. You may use this figure to determine if there has been a change in a router's link state databases, and to compare the link-state database of two routers.
Area Summary	Indicates whether the GSR imports summary link-state advertisements into stub areas. This field has no effect on other areas. If the value of this field is noAreaSummary, the GSR will neither originate nor propagate summary link-state advertisements into the stub area. It will rely entirely on its default route. If it is sendAreaSummary, the GSR will both summarize and propagate summary link-state advertisements.
Area Status	Indicates the status of the entry. This field's value is always valid.

Obtaining OSPF Neighbor Information

Obtain OSPF neighbor information if you want details about the neighboring routers of a GSR. To access this information:

1. If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Routing State**, choose **OSPF State**, and then choose **OSPF Neighbor Table**.

An **OSPF Neighbor** table similar to the following appears:

IP Address	Addressless Index	Router ID	Options	Designated Router Priority	State
207.135.8...	0	207.135...	2	4	full

Selection Details	
IP Address	207.135.89.146
Addressless Index	0
Router ID	207.135.89.146
Options	2
Designated Router Priority	4
State	full

1 rows displayed

Figure 257. OSPF Neighbor table

The following table describes the fields of the **OSPF Neighbor** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 61. OSPF Neighbor table fields

Field	Description
IP Address	Identifies the IP address the neighbor is using. On addressless links, the value of this field will be the address of another of the neighbor's interfaces.
Addressless Index	Displays the corresponding value of the interface index in the Internet Standard MIB if the interface is an addressless interface. If an interface has an IP address, the value of this field is zero (0).
Router ID	Identifies the IP address that uniquely identifies the neighboring router in the autonomous system.
Options	Indicates the optional capabilities supported by the neighbor.
Designated Router Priority	Indicates the priority of the neighbor in the designated router election algorithm. The value zero (0) signifies that the neighbor is not eligible to become the designated router on this particular network.

Table 61. OSPF Neighbor table fields (Continued)

Field	Description
State	Indicates the state of the relationship with the neighbor.
Nbr. Events	Indicates how many times the neighbor relationship has changed state or an error has occurred.
Retransmission Q Length	Indicates the current length of the retransmission queue.
Status	Indicates the status of the entry.
Nbr. Permanence	Indicates how the neighbor became known. This field shows whether the neighbor was dynamically learned (dynamic) or was statically configured (permanent).
Hello Suppressed	Indicates whether hello packets are being suppressed to the neighbor.

Obtaining OSPF Link-State Database Information

Obtain OSPF link-state database information if you want details about the link-state advertisements of the areas to which the GSR is attached. To access this information:

1. If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Routing State**, choose **OSPF State**, and then choose **OSPF Link State DB Table**.

An **OSPF Link State DB** table similar to the following appears:

Area ID	Area Type	Link-State ID	Router ID	Sequence	LS Age	Checks...
ANY	routerLink	207.135.89.55	207.135...	ANY	1332	24356
ANY	routerLink	207.135.89.146	207.135...	ANY	881	45846
ANY	networkLink	207.135.89.146	207.135...	ANY	881	29950

Selection Details	
Area ID	ANY
Area Type	routerLink
Link-State ID	207.135.89.55
Router ID	207.135.89.55
Sequence	ANY
LS Age	1332

3 rows displayed

Figure 258. OSPF Link State DB table

The following table describes the fields of the **OSPF Link State DB table's** upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 62. OSPF Link State DB table fields

Field	Description
Area ID	Indicates the 32-bit identifier of the area from which the link-state advertisement was received.
LS DB Type	Indicates the link-state advertisement's type. On the GSR, the possible values in this field are routerLink, networkLink, summaryLink, and asSummaryLink (represents an autonomous system summary link).
Link-State ID	Identifies the piece of the routing domain the link-state advertisement is describing. The information in this field depends on the link-state type and it is either a router ID or an IP address.
Router ID	Identifies which router sent the link-state advertisement.

Table 62. OSPF Link State DB table fields

Field	Description
Sequence	Indicates the sequence number of the link-state advertisement. This field is used to detect old and duplicate link state advertisements. The larger the number the more recent the link-state advertisement.
LS Age	Indicates how old the link-state advertisement is. The value in this field is given in seconds.
Checksum	Indicates the checksum of the complete contents of the link-state advertisement, except for the age field. The age field is excluded so that a link-state advertisement's age can be incremented without updating the checksum.
LS DB Advertisement	Contains the entire link-state advertisement including its header.

Obtaining OSPF Area Aggregate Information

Obtain OSPF area aggregate information if you want details about the IP address and TOS data that OSPF uses to determine the destination of packets. To access this information:

1. If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Routing State**, choose **OSPF State**, and then choose **OSPF Area Aggregate Table**.

An **OSPF Area Aggregate** table similar to the following figure appears:

Area ID	LS DB Type	Net/Subnet Address	Net/Subnet Mask	Status	Effect
2.0.0.0	ANY	ANY	207.135.89.145	active	advertiseMatching

Selection Details	
Area ID	2.0.0.0
LS DB Type	ANY
Net/Subnet Address	ANY
Net/Subnet Mask	207.135.89.145
Status	active
Effect	advertiseMatching

1 rows displayed

Figure 259. OSPF Area Aggregate table

The following table describes the fields of the **OSPF Area Aggregate** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 63. OSPF Area Aggregate table fields

Field	Description
Area ID	Identifies the area the address aggregate is to be found within.
LS DB Type	Indicates the address aggregate's type, which specifies the link-state database type that the address aggregate applies to.
Net/Subnet Address	Identifies the IP address of the network or subnetwork indicated by the range.
Net/Subnet Mask	Identifies the subnet mask that pertains to the network or subnetwork.
Status	Indicates the status of the area network range.
Effect	Indicates whether subnetworks contained by ranges trigger the advertisement of the indicated aggregate or result in the subnetworks not being advertised at all outside the area.

Checking RIP Routing Status

clearVISN CoreWatch can display information about the following:

- RIP interfaces configured on a GSR. This includes details about how many packets were discarded on each interface, how many route entries of valid RIP packets were ignored on each interface, how many RIP updates were sent on the interface, and whether a RIP interface is functioning.
- Active RIP peers. This includes details about the last time the GSR was updated with a peer's routing information, which version of RIP a peer is running, how many packets were discarded on each peer, and how many route entries from each peer were ignored.

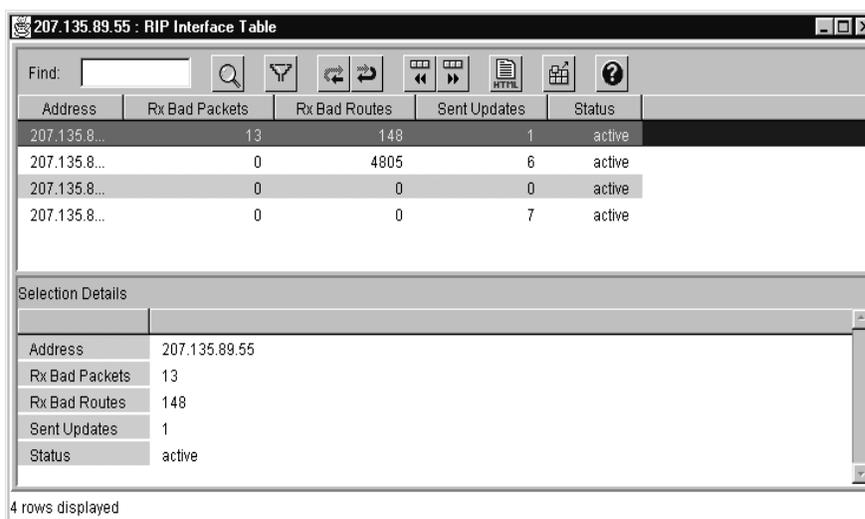
Obtaining RIP Interface Information

Obtain RIP interface information if you want details about how many packets on a RIP interface the GSR discarded, how many route entries of valid RIP packets the GSR ignored, how many RIP updates the GSR has sent on the interface, and whether a RIP interface is functioning.

To obtain RIP interface information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Routing State**, choose **RIP State**, and then choose **RIP Interface Table**.
- In the Schematic view, double-click either the RIP Routing function or the IPX-RIP Routing function.

A **RIP Interface** table similar to the following appears:



Address	Rx Bad Packets	Rx Bad Routes	Sent Updates	Status
207.135.8...	13	148	1	active
207.135.8...	0	4805	6	active
207.135.8...	0	0	0	active
207.135.8...	0	0	7	active

Selection Details	
Address	207.135.89.55
Rx Bad Packets	13
Rx Bad Routes	148
Sent Updates	1
Status	active

4 rows displayed

Figure 260. RIP Interface table

The following table describes the fields of the **RIP Interface** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 64. RIP Interface table fields

Field	Description
Address	Identifies the network IP address of the interface. Each interface has a unique index that the GSR uses to identify interfaces without an IP address. The GSR does so by appending the index number to an IP address of 0.0.0. Suppose an unnumbered interface's index is 55, then the GSR would use 0.0.0.55 as the interface's IP address.
Rx Bad Packets	Indicates how many of the interface's response packets the GSR received and then discarded because those packets were empty, included an invalid command, or had some other error.
Rx Bad Routes	Indicates how many of the interface's routes, included in valid RIP packets, the GSR ignored. The GSR ignores a route entry if it includes an unknown address family, an invalid hop number, or some other error.
Sent Updates	Indicates how often the GSR was updated with the interface's routing information. The figure in this field includes only updates sent because of route changes. This figure does not include updates sent containing new information.
Status	Indicates whether the interface is functioning (1) or not functioning (2).

Obtaining RIP Peer Information

Obtain RIP Peer information if you want details about the last time the GSR was updated with a peer's routing information, which version of RIP a peer is running, how many packets the GSR discarded from a peer, and how many of a peer's route entries the GSR ignores. You may find such information useful when resolving problems between the GSR and its RIP peers.

To obtain RIP peer information, do one of the following:

- If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
- Select the **Monitor** menu, choose **Routing State**, choose **RIP State**, and then choose **RIP Peer Table**.

A RIP Peer table similar to the following appears:

Peer Address	Domain	Last Update	Version	Rx Bad Packets	Rx Bad Routes
207.135.89.135		1 day 15 ho...	2	0	4807

Selection Details	
Peer Address	207.135.89.135
Domain	
Last Update	1 day 15 hours 55 minutes
Version	2
Rx Bad Packets	0
Rx Bad Routes	4807

1 rows displayed

Figure 261. RIP Peer table

The following table describes the fields of the **RIP Peer** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 65. RIP Peer table fields

Field	Description
Peer Address	Identifies the IP address the peer is using as its source address.
Domain	Includes the value found in the Routing Domain field of the RIP packets the GSR received from the peer. RIP 2 does not use this field so the value in this field is always 0.
Last Update	Indicates how long ago the GSR was updated with routing information from the peer.
Version	Indicates which RIP version number was included in the header of the last packet the GSR received from the peer.
Rx Bad Packets	Indicates how many RIP response packets from the peer the GSR discarded as invalid.
Rx Bad Routes	Indicates how many of the routes from the peer the GSR ignored because the route-entry format was invalid.

Checking DVMRP Routing Status

clearVISN CoreWatch can display information about the following:

- The DVMRP interfaces configured on a GSR. This includes details about the configuration of DVMRP interfaces, whether a DVMRP interface is functioning, how many packets were discarded on each interface, and how many route entries of valid DVMRP packets were ignored.
- The DVMRP neighboring routers. This includes details about the length of time those routers have been neighbors to the GSR, their capabilities, and DVMRP traffic.
- The multicast routes DVMRP uses instead of unicast routes. This includes details about the GSR's multicast routes such as their sources, upstream neighbors, and hop counts.
- The next hops of the DVMRP interfaces to which the GSR is sending IP multicast packets.

Obtaining DVMRP Interface Information

Obtain DVMRP Interface information if you want to examine details about DVMRP interfaces. This includes information about the configuration of DVMRP interfaces, whether a DVMRP interface is functioning, how many packets were discarded on each interface, and how many route entries of valid DVMRP packets were ignored.

To access such information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Routing State**, choose **DVMRP State**, and then choose **DVMRP Interface Table**.
- In the Schematic view, double-click the DVMRP Routing function.

A **DVMRP Interface** table similar to the following appears:

Local Address	Distance Metric	Status	Rx Bad Packets	Rx Bad Routes
172.1.1.10	1	active	0	0
207.135.89.64	1	active	0	0
10.135.89.10	1	active	0	0
190.1.0.1	1	notInSe...	0	0
207.135.122.11	1	active	0	0
10.10.1.10	1	active	0	0

Selection Details	
Local Address	172.1.1.10
Distance Metric	1
Status	active
Rx Bad Packets	0
Rx Bad Routes	0

7 rows displayed

Figure 262. DVMRP Interface table

The following table describes the fields of the **DVMRP Interface** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 66. DVMRP Interface table fields

Field	Description
Local Address	Identifies the interface's IP address.
Distance Metric	Identifies the cost assigned to the interface. The cost is a number from 1 to 31 that the GSR system administrator sets. This metric is similar to the one RIP uses.
Status	Indicates whether DVMRP is enabled (Up) or disabled (Down) on the interface.
Rx Bad Packets	Indicates how many invalid packets were received on the interface. of the interface's packets the GSR received and then discarded as invalid. The GSR discards packets because they include a bad report, have a bad packet format, or are from an unrecognized neighbor.
Rx Bad Routes	Indicates how many invalid route entries were received on the interface.

Obtaining DVMRP Neighbor Information

Obtain DVMRP Neighbor information if you want to examine details about DVMRP neighboring routers. This includes information about the length of time those routers have been neighbors to the GSR, the capabilities of those routers, and the traffic the GSR receives from those neighbors.

To access such information:

1. If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Routing State**, choose **DVMRP State**, and then choose **DVMRP Neighbor Table**.

A **DVMRP Neighbor** table similar to the following appears:

Index	Address	Up Time	Expiry Time	Generation ID	Major Version	Minor Version	
2	10.135...			1253429	255	0	IF
4	207.135...			210	255	0	IF
5	10.40.1.1			267	255	0	IF

Selection Details	
Index	2
Address	10.135.89.67
Up Time	
Expiry Time	
Generation ID	1253429
Major Version	255

3 rows displayed

Figure 263. DVMRP Neighbor table

The following table describes the fields of the **DVMRP Neighbor** table's upper frame. The Selection Details frame displays information about the item currently selected in

the upper frame. For details on using the Table toolbar, see [Appendix A: “Working with Tables”](#) on page 425.

Table 67. DVMRP Neighbor table fields

Field	Description
Index	Identifies the value of the virtual interface index used to reach the neighbor.
Address	Identifies the IP address of the neighbor.
Up Time	Indicates how long ago the DVMRP neighbor last became a neighbor to the GSR.
Expiry Time	Indicates the minimum amount of time remaining before the neighbor's entry will be deleted from the DVMRP Neighbor table.
Generation ID	Identifies the neighboring router's generation identifier.
Major Version	Indicates the neighboring router's major DVMRP version number. The GSR's major version number is 3.
Minor Version	Indicates the neighboring router's minor DVMRP version number. The GSR's minor version number is 255.
Capabilities	<p>Describes the neighboring router's capabilities. These capabilities are:</p> <ul style="list-style-type: none"> • Leaf – signifies that the router has only one interface with neighbors. • Prune – signifies that the router supports pruning. • GenerationID – signifies that the neighbor recognizes generation identifiers and sends that information in Probe messages. • mtrace – signifies that the neighbor can handle mtrace requests. The mtrace command tracks the multicast path from a source to a receiver.
Rx Routes	<p>Indicates how many routes the GSR received in valid DVMRP packets received from the neighbor.</p> <p>This figure is useful in diagnosing problems as well as indicating the level of DVMRP traffic.</p>
Rx Bad Packets	Indicates how many invalid packets the GSR received from the neighbor.
Rx Bad Routes	Indicates how many of the neighbor's routes the GSR ignored. The GSR ignores invalid route entries.

Obtaining DVMRP Routing Information

Obtain DVMRP routing information if you want details about the multicast routes DVMRP uses instead of unicast routes. This routing information includes such things as the identity of a route's source and upstream neighbor, the route's hop count, and an indication of how long ago the GSR learned of the route.

To access DVMRP routing information:

1. If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Routing State**, choose **DVMRP State**, and then choose **DVMRP Route Table**.

A **DVMRP Routing** table similar to the following appears:

Source IP Address	Source SubnetMask	Upstream Neighbour	IP Interface	Route Metric
4.0.20.16	255.255.255.240	207.135.122.10	5	32
4.0.35.16	255.255.255.240	207.135.122.10	5	32
4.0.36.0	255.255.255.240	207.135.122.10	5	32
4.0.48.16	255.255.255.240	207.135.122.10	5	32
9.2.27.0	255.255.255.128	207.135.122.10	5	10

Selection Details	
Source IP Address	4.0.20.16
Source SubnetMask	255.255.255.240
Upstream Neighbour	207.135.122.10
IP Interface	5
Route Metric	32
Expiration Time	2 minutes 35 seconds

102 rows displayed

Figure 264. DVMRP Routing table

The following table describes the fields of the **DVMRP Routing** table's upper frame. The Selection Details frame displays information about the item currently selected in

the upper frame. For details on using the Table toolbar, see [Appendix A: “Working with Tables”](#) on page 425.

Table 68. DVMRP Routing table fields

Field	Description
Source IP Address	Indicates the network address of the source for which the table entry contains multicast routing information. This address is combined with a source's network mask to identify that source.
Source Subnet Mask	Indicates the network mask of the source for which the table entry contains multicast routing information. This mask is combined with a source's network address to identify that source.
Upstream Neighbor	Identifies the address of the upstream neighbor from which the source's IP packets are received. The upstream neighbor is a Reverse Path Forwarding (RPF) neighbor. RPF is a routing technique that finds the next neighbor by determining which router send the GSR the packet.
IP Interface	Identifies the index of the interface on which IP datagrams sent by the source are received.
Route Metric	Indicates the hop count to the source's subnet.
Expiration Time	Indicates the minimum amount of time remaining before the route's entry will be deleted from the DVMRP Routing table.
Up Time	Indicates how long ago the GSR learned of the route.

Obtaining DVMRP Next Hop Information

Obtain DVMRP Next Hop information if you want to display information about the next hop to which the GSR is sending IP multicast packets. This information includes details about a hop's source and indicates whether the hop has downstream neighbors. To access such information:

1. If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Routing State**, choose **DVMRP State**, and then choose **DVMRP Next Hop Table**.

A **DVMRP Next Hop** table similar to the following appears:

Source Address	Source Mask	If Index	Type
4.0.20.16	255.255.255.240	1	leaf
4.0.20.16	255.255.255.240	2	leaf
4.0.20.16	255.255.255.240	3	branch
4.0.20.16	255.255.255.240	6	branch
4.0.35.16	255.255.255.240	1	leaf
4.0.35.16	255.255.255.240	2	leaf

Selection Details	
Source Address	4.0.20.16
Source Mask	255.255.255.240
If Index	1
Type	leaf

100 rows displayed

Figure 265. DVMRP Next Hop table

The following table describes the fields of the **DVMRP Next Hop** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables"](#) on page 425.

Table 69. DVMRP Next Hop table fields

Field	Description
Source Address	Identifies the source network address.
Source Mask	Identifies the source network mask.
If Index	Indicates the index value of the interface used as the next hop
Type	Indicates whether the next hop is a leaf or branch.

Checking IGMP Status

clearVISN CoreWatch can obtain the following:

- Information about which interfaces IGMP is enabled on and details about the IGMP configuration of those interfaces.
- Information about the multicast groups of IGMP interfaces. This includes details about which interfaces the multicast groups are on, whether the GSR is a member of those groups, and other IGMP-configuration data.

Obtaining IGMP Interface Information

Obtain IGMP Interface information if you want to learn on which interfaces IGMP is enabled or examine details about the IGMP configuration of those interfaces. To access such information, do one of the following:

- In the Front Panel view, select the **Monitor** menu, choose **Routing State**, choose **IGMP State**, and then choose **IGMP Interface Table**.
- In the Schematic view, double-click the IGMP Routing function.

An **IGMP Interface** table similar to the following appears:

The screenshot shows a window titled "10.50.88.162 : IGMP Interface Table". It contains a table with the following data:

Index	Query Interval	Status	Version	If Querier	Max Response Time	Ver. 1 Quer
1	125	active	version2	10.100.10.2	10	
2	125	active	version2	10.100.20.2	10	
3	125	active	version2	ANY	10	
4	125	active	version2	11.11.11.8	10	
5	125	active	version2	10.100.30.2	10	

Below the table is a "Selection Details" pane showing the following information for the selected item (Index 1):

Index	1
Query Interval	125
Status	active
Version	version2
If Querier	10.100.10.2
Max Response Time	10

5 rows displayed

Figure 266. IGMP Interface table

The following table describes the fields of the **IGMP Interface** table's upper frame. The Selection Details frame displays information about the item currently selected in the

upper frame. For details on using the Table toolbar, see [Appendix A: “Working with Tables” on page 425](#).

Table 70. IGMP Interface table fields

Field	Description
Index	Identifies an IP interface on which IGMP is enabled.
Query Interval	Indicates how often the GSR sends IGMP Host-Query packets on the interface. An IGMP router sends hosts a Host-Query packet to learn which hosts are available. The specified interval applies to all GSR ports. The time given is in seconds. The default is 125 seconds.
Status	Indicates whether IGMP is currently enabled on the interface.
Version	Indicates which version of IGMP is running on the interface. All IGMP routers on a LAN must be running the same version of IGMP on that LAN. The GSR runs version 2, but can communicate with hosts running version 1.
If Querier	Identifies the IP address of the IGMP Querier on the IP subnetwork to which the interface is attached. The IGMP Querier looks for the hosts of an IGMP router.
Max Response Time	Indicates how long the GSR waits for IGMP Host Membership responses from the hosts. The wait time you set applies to all ports on the GSR.
Ver. 1 Querier Time	Indicates how much longer the GSR will assume that there are no IGMP version 1 routers running on the interface. As long as this time is greater than 0, the GSR continues to reply to all queries with version 1 membership reports.
Wrong Ver. Queries	Indicates how many queries the GSR received from hosts running a version of IGMP different from the one running on the GSR. A value greater than 0 indicates a configuration problem because IGMP requires that all routers on a LAN be configured to run the same version of IGMP.

Table 70. IGMP Interface table fields (Continued)

Field	Description
Joins	Indicates how many times a group membership has been added on the interface. By examining this field, you can determine the amount of IGMP activity over time.
Groups	Indicates how many groups are on the interface.
Robustness	Indicates the expected packet loss on a subnet. Robustness is a measurement of a lossy network condition. The GSR assumes a robustness of 2.

Obtaining IGMP Cache Information

Obtain IGMP Cache information if you want to learn about the multicast groups of IGMP interfaces, identify whether the GSR is a member of those groups, and examine other IGMP configuration information about those groups. To access such information:

1. If you are not in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **Routing State**, choose **IGMP State**, and then choose **IGMP Cache Table**.

An IGMP Cache table similar to the following appears:

IP Multicast Address	IGMP Interface	Cache Status	Local Membership	Member Since
224.0.1.24	1	active	false	45 seconds
224.0.1.32	4	active	false	54 seconds
224.1.127.255	4	active	false	1 minute 10 seco...
224.2.3.34	4	active	false	40 seconds
224.2.127.253	4	active	false	1 minute 9 seconds

Selection Details	
IP Multicast Address	224.0.1.24
IGMP Interface	1
Cache Status	active
Local Membership	false
Member Since	45 seconds
Local Queried Member	207.135.89.10

8 rows displayed

Figure 267. IGMP Cache table

The following table describes the fields of the **IGMP Cache** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables"](#) on page 425.

Table 71. IGMP Cache table fields

Field	Description
IP Multicast Address	Identifies the IP address of the multicast group.
IGMP Interface	Indicates the interface on which the IP multicast group is configured.
Cache Status	Indicates whether the multicast group on the interface is active, not in service, not ready, create and go, create and wait, or destroyed.
Local Membership	Indicates whether a router is a member of the group. The GSR is never a member of a group.
Member Since	Indicates how long ago the GSR joined the multicast group address. A 0 signifies that the GSR is not a member of the group.
Last Reported Member	Identifies the IP address of the source that sent the GSR the last membership report for the multicast group on the interface identified in the IGMP Interface field. An IP address of 0.0.0.0 signifies that the GSR has not received a membership report on the multicast group's interface.
Cache Expiry Time	Indicates the minimum amount of time remaining before the entry will be deleted from the IGMP Cache table.
Time Before Expiry	Indicates how much longer the GSR will assume that there are no IGMP version 1 members on the IP subnetwork attached to the multicast group's interface. If this time is greater than 0, any version 2 IGMP messages the GSR receives on the interface are ignored. GSR resets this field whenever it receives a version 1 IGMP membership report on the interface.

Chapter 21

Checking the Status of QoS Tables

Examining Quality of Service (QoS) information in tables, you can identify historical trends. You can obtain the following QoS information in clearVISN CoreWatch:

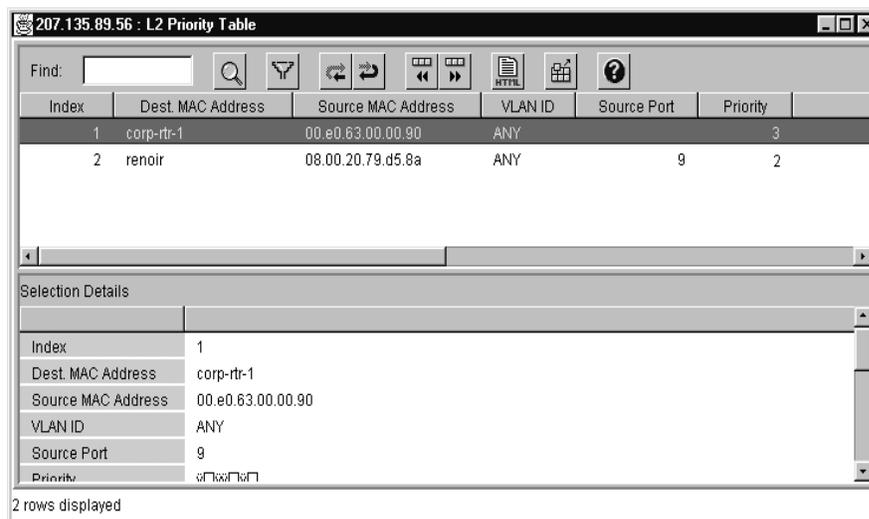
- Details about QoS priorities of Layer-3/ Layer-4 flows.
- Information about the routing of any Layer-2, Layer-3, or Layer-4 data sent to a port rather than being sent to the Control Module for further processing.

Obtaining Layer-2 Priority Information

Obtain Layer-2 priority information if you want to examine the QoS priorities of Layer-2 flows. clearVISN CoreWatch identifies a flow's priority and provides information about the fields for which that priority applies. To access such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **QoS State**, and then choose **L2 QoS**.

An **L2 Priority** table similar to the following appears:



Index	Dest. MAC Address	Source MAC Address	VLAN ID	Source Port	Priority
1	corp-rtr-1	00.e0.63.00.00.90	ANY		3
2	renoir	08.00.20.79.d5.8a	ANY	9	2

Selection Details	
Index	1
Dest. MAC Address	corp-rtr-1
Source MAC Address	00.e0.63.00.00.90
VLAN ID	ANY
Source Port	9
Priority	3

2 rows displayed

Figure 268. L2 Priority table

The following table describes the fields of the **L2 Priority** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 72. L2 Priority table fields

Field	Description
Index	Identifies the entry's table index.
Name	Identifies the name of the flow.
Dest. MAC Address	Identifies the Layer-2 destination's MAC address.
Source MAC Address	Identifies the Layer-2 source's MAC address.
VLAN ID	Identifies the VLAN to which the packet belongs.
Source Ports	Indicates which GSR ports have priority when L2 packets enter the GSR. The priority does not apply to exit ports.
Priority	Indicates the traffic priority assigned to the flow. The value of this field will be control, high, medium, or low. It specifies how much priority is given to data that matches the criteria set through the flow's parameters. Packets assigned the control priority are transmitted first.

Obtaining Flow Priority Information

Obtain flow priority information if you want to examine the QoS priorities of Layer-3 and Layer-4 flows. clearVISN CoreWatch indicates a flow's priority and provides information about the fields for which that priority applies. To access such information:

1. If you are not currently in the Front Panel view, switch to it by clicking the clearVISN CoreWatch main window.
2. Select the **Monitor** menu, choose **QoS State**, and then choose **L3/L4 QoS**.

A **Flow Priority** table similar to the following appears:

Flow Name	Priority	Interface	Protocol	Source Address	Source Port	Destination Add
xyz	High	ANY	ANY	1.1.1.1	45	10.1.1.1
ipsev	High	00 00	ANY	ANY	ANY	120.2.34.1

Selection Details

Flow Name	
Priority	
Interface	
Protocol	
Source Address	
Source Port	

2 rows displayed

Figure 269. Flow Priority table

The following table describes the fields of the **Flow Priority** table's upper frame. The Selection Details frame displays information about the item currently selected in the

upper frame. For details on using the Table toolbar, see [Appendix A: “Working with Tables” on page 425](#).

Table 73. Flow Priority table fields

Field	Description
Flow Name	Identifies which flow is assigned the traffic priority shown in the IP flow's Priority field.
Priority	Indicates the traffic priority assigned to the flow. The value of this field will be control, high, medium, or low. It specifies how much priority is given to data that matches the criteria set through the flow's parameters. Packets assigned the control priority are transmitted first.
Interface	Indicates which IP interfaces are given priority in the flow. The value ANY indicates that all interfaces are accepted for this field.
Protocol	Indicates which transport layer protocol (TCP or UDP) is given priority in the flow. When configuring the GSR, the network administrator can set this field to <i>any</i> for packets with the protocol set to TCP or UDP if the other fields match.
Source Address	Indicates the source IP addresses and network masks given priority in the flow. The value ANY indicates that all addresses and masks are accepted for this field. The value 255.255.255.255 signifies that no mask entry was specified when this field was set.
Source Port	Indicates the source TCP or UDP port given priority in the flow. The value ANY indicates a wildcard.
Destination Address	Indicates the destination IP address and network mask given priority in the flow. The value ANY indicates that any address and mask are accepted for this field. The value 255.255.255.255 signifies that no mask entry was specified when this field's priority was set.
Destination Port	Indicates the destination TCP or UDP port given priority in the flow. The value ANY indicates that all such ports have priority.
TOS	Indicates the TOS given priority in the flow.

Obtaining Layer-2 Switching Information

Obtain Layer-2 switching information to examine details about the routing of Layer-2 data sent directly to a port rather than being sent to the Control Module for further processing. To access this information:

1. In the Front Panel view, select the **Monitor** menu, choose **QoS State**, and then choose **L2 Flows**.
2. In the Schematic view, double-click the Layer 2 Switching function.

An **L2 Forward** table similar to the following appears:

Index	Dest. MAC Address	Source MAC Address	VLAN ID	Dest. Port	Source Port
0			255	0	00 00 04 00 ...
0			255	0	00 00 00 88 ...
0			255	0	00 00 40 00 ...
0			255	0	00 00 40 00 ...
0			255	0	00 89 00 00 ...
n			0	74	00 00 00 00 ...

Selection Details	
Index	0
Dest. MAC Address	
Source MAC Address	
VLAN ID	255
Dest. Port	0
Source Port	00 00 04 00 00 00 00

100 rows displayed

Figure 270. L2 Forward table

The following table describes the fields of the **L2 Forward** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 74. L2 Forward table fields

Field	Description
Index	Indicates in which row the MAC address entry appears in the L2 Forward table.
Dest. MAC Address	Identifies the Layer 2 destination's MAC address.
Source MAC Address	Identifies the Layer 2 source's MAC address.

Table 74. L2 Forward table fields (Continued)

Field	Description
VLAN ID	Identifies the VLAN that is combined with a MAC address to uniquely identify the entry. The same MAC address may be learned on different VLANs. MAC addresses are combined with VLAN names to create unique identifiers for each entry.
Dest. Port	Indicates to which destination port the GSR forwarded the data. Traffic destined to the entry's MAC address is forwarded to the destination port.
Source Port	Indicates on which port the GSR received the Layer-2 packets.

Obtaining Layer-3 and Layer-4 Switching Information

Obtain Layer-3 and Layer-4 switching information to examine details about the routing of Layer-3 or Layer-4 data sent directly to a port rather than being sent to the Control Module for further processing. To access this information, take the following steps:

1. Do one of the following:
 - In the Front Panel view, select the **Monitor** menu, choose **QoS State**, and then choose **L3/L4 Flows**.
 - In the Schematic view, double-click either the Layer 3/4 Switching function or the Flows object.

The **Flow Table Filter** dialog box appears.

The image shows a dialog box titled "Flow Table Filter Form". It contains several text input fields for filtering: "Port Of Entry", "Protocol" (with "TCP" entered), "Source Address", "Source Port", "Destination Address", "Destination Port", "TOS", and "Total Packets". At the bottom right, there is a radio button labeled "And" which is selected. An "OK" button is located at the bottom center.

Figure 271. Flow Table Filter dialog box

2. Do the following to restrict the amount of data clearVISN CoreWatch receives from Layer 3 and Layer 4:
 - To include an item and not limit which values are accepted for that item, leave its box blank. This is the default for each item.
 - To include an item but limit its contents, enter the desired value in the appropriate box.
3. Do one of the following:
 - Leave the **And** button selected if you want the **Flow** table to include items matching all the selection criteria you specified in the form's text boxes.
 - Click the **And** button to clear it so that the **Flow** table will include items matching any of the specified selection criteria.
4. Click **OK**.

A **Flow** table similar to the following appears:

Port Of Entry	Protocol	Source Address	Source Port	Destination Address	Destination Port
Port et.2.2	TCP	207.135.89.183	64561	207.135.89.44	7000
Port et.2.3	UDP	207.135.89.180	47810	207.135.89.44	PORTMAPPER
Port et.2.6	UDP	207.135.89.179	NFS (2049)	207.135.89.29	63082
Port et.2.1	UDP	207.135.89.178	59351	207.135.89.44	732
Port et.2.5	UDP	207.135.89.182	37173	207.135.89.33	724

Selection Details	
Port Of Entry	Port et.2.3
Protocol	UDP
Source Address	207.135.89.180
Source Port	47810
Destination Address	207.135.89.44
Destination Port	PORTMAPPER (411)

102 rows displayed

Figure 272. Flow table

The following table describes the fields of the **Flow** table's upper frame. The Selection Details frame displays information about the item currently selected in the upper frame. For details on using the Table toolbar, see [Appendix A: "Working with Tables" on page 425](#).

Table 75. Flow table fields

Field	Description
Port Of Entry	Indicates on which port the GSR received data for the flow.
Protocol	Indicates the transport layer protocol of the flow (such as TCP or UDP).
Source Address	Indicates the network layer address of the source that originally sent the packet.
Source Port	Indicates the Layer-4 port number corresponding to the source of the data.
Destination Address	Indicates the network layer address of the destination to which the data was sent.
Destination Port	Indicates the Layer-4 port number corresponding to the destination of the data.
TOS	Indicates the value of the TOS field of packets in the flow.
Total Packets	Indicates how many packets for the flow the GSR sent directly to the destination port.

Chapter 22

Obtaining Reports

While monitoring a GSR, you may want to keep a record of the information found in the GSR boot log or any clearVISN CoreWatch table. You may keep such records by saving clearVISN CoreWatch table information to a clearVISN CoreWatch report, which is an HTML file.

This chapter discusses obtaining reports that include boot log information or data from multiple clearVISN CoreWatch tables. This chapter also discusses saving a single table as a report.

You can view a clearVISN CoreWatch report in your Web browser either immediately after the report is generated or at a later time.

Saving Multiple Tables as a Report

To obtain a report that includes data from more than one clearVISN CoreWatch table:

1. Do one of the following:
 - Select the **File** menu and then choose **Reports**.
 - Click the **Reports** button  on the clearVISN CoreWatch toolbar.

The **Report Selection** dialog box appears.

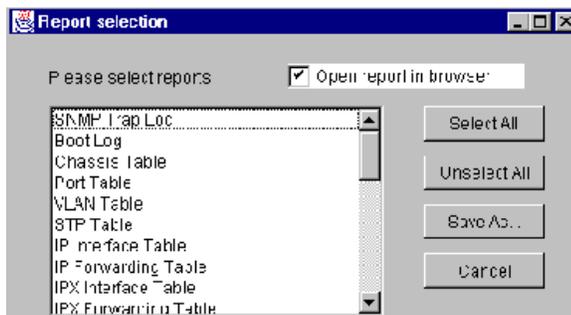


Figure 273. Report Selection dialog box

2. In the **Please select reports** list, select one or more items that you want the report to include.

If you want the report to include information from all clearVISN CoreWatch tables, click the **Select All** button.

If you make selections and then decide you do not want to select any of those tables, click the **Unselect All** button.
3. Specify whether you want to view the report immediately after clearVISN CoreWatch generates it. To do so, do one of the following:
 - If you want to view the report immediately after clearVISN CoreWatch generates it as well as be able to examine it later, leave the *Open report in browser* option selected.
 - If you prefer to look at the report later and do not want to examine it immediately after clearVISN CoreWatch generates it, clear the *Open report in browser* option.
4. Click the **Save As** button.

The **Save As** dialog box appears.
5. Enter a name for the report in the **File Name** box. If necessary, browse to the folder in which you want to save the report, then click the **Open** button.

If the specified file already exists, clearVISN CoreWatch prompts you to verify whether you want to overwrite the existing file. Clicking **OK** overwrites the file. Clicking **Cancel**, entering a different filename, and then clicking the **Open** button saves the information to another file.

clearVISN CoreWatch saves the report as an HTML file. clearVISN CoreWatch displays this file in your Web browser if the *Open report in browser* option in the **Report Selection** dialog box was selected. Otherwise, use your Web browser to open the file.

You may print the report using your Web browser's print feature.

Saving a Single Table as a Report

While monitoring a GSR, you may want to keep a record of the information found in a clearVISN CoreWatch table. You may do so by saving the table's data to a clearVISN CoreWatch report, which is an HTML file. To save the information of one clearVISN CoreWatch table to a report:

1. Open the table that you want to include in the report.
2. Click the **Report** button  on the Table toolbar.

The **Save As** dialog box appears.

3. Enter a name for the report in the **File Name** box. If necessary, browse to the folder in which you want to save the report, then click the **Open** button.

If the specified file already exists, clearVISN CoreWatch prompts you to verify whether you want to overwrite the existing file. Clicking **OK** overwrites the file. Clicking **Cancel**, entering a different filename, and then clicking the **Open** button saves the information to another file.

clearVISN CoreWatch saves the report as an HTML file and displays that file in your Web browser. You may print the report using your Web browser's print feature.

Appendix A

Working with Tables

You can perform the following operations in any clearVISN CoreWatch table:

- Find text in a table
- Control the contents of tables
- Refresh table information
- Restore table information
- Obtain additional records
- Save a table as a report
- Export table information to a file
- Sort data

Most of these tasks you perform using the Table toolbar. Separate discussions on these tasks, finding text in a table, and sorting data follow.

Finding Text in a Table

To find text in a clearVISN CoreWatch table:

1. Enter the desired text in the **Find:** box on the Table toolbar.
2. Click the **Find** button  on the Table toolbar.

clearVISN CoreWatch finds the first instance of the specified text. To find the next instance of that text, click the **Find** button again.

Controlling the Contents of Tables

You can control which fields are included in a clearVISN CoreWatch table and also limit which values are displayed in each of the table's fields.

To control the contents of a clearVISN CoreWatch table:

1. Click the **Filter** button  on the Table toolbar.
The table's filter form appears. The fields in this form depends on which table you currently have open.
2. Do one of the following to control which fields the table includes or limit the values of those fields. For details on a table's fields, see the topic that discusses that table.
 - To not include a field, select that field's check box and leave its text box blank.
 - To include a field and not limit which values are displayed in that field, leave that field's check box and text box blank. This is the default for each field.
 - To include a field but limit the contents of those fields, do one of the following:
 - To include a field only if it has a specific value, enter the desired value in the field's text box. Leave the field's check box blank.

Suppose the table includes a Port field and you enter 2 in that field. The table will then include only those items relevant to port 2 of each module.

 - To include a field only if it is not equal to a specific value, enter that value in the field's text box and select the field's check box.

Suppose the table includes a *Port* field. If you enter 2 in that field and also select that field's check box, the table will include only those items relevant to any port except for port 2 of each module.

3. If you entered values in one or more text boxes, select one or more of the options described in the following table:

Select	To match the
<i>Case Sensitive</i>	Case (upper or lower) of text-box entries
<i>Regular Expression</i>	Text box entries using wildcards, which let you search for character patterns and for text appearing at the beginning or end of a line. clearVISN CoreWatch supports the Perl5 regular expressions, which are described in Appendix C: "Supported Regular Expressions" on page 439.
<i>Match Whole Field</i>	Text box entries as whole words.

4. Select one of the options discussed in the following table:

Select	If the table is to include items matching
<i>And</i>	All the selection criteria specified in each of the form's text boxes and check boxes.
<i>Or</i>	Any of the selection criteria specified in each of the form's text boxes and check boxes.

5. Click OK.

Refreshing a Table

After you open a table, you may want to refresh it so that it contains the most up-to-date information. To do so, click the **Refresh** button  on the Table toolbar.

Restoring Table Information

If you control a table's content and then decide you no longer want to filter the data in that table, click the **Restore** button  on the Table toolbar.

Obtaining Table Records

If the information clearVISN CoreWatch obtains for a table is more than the table can include at one time, you can update the table with additional records. To do so, do one of the following:

- To obtain the next set of records clearVISN CoreWatch can display at the same time, click the **Next Rows** button on the Table toolbar.
- To obtain the previous set of records clearVISN CoreWatch displayed at the same time, click the **Previous Rows** button on the Table toolbar.

Saving a Single Table as a Report

While monitoring a GSR, you may want to keep a record of the information found in a clearVISN CoreWatch table. You may do so by saving the table's data to a clearVISN CoreWatch report, which is an HTML file. This section discusses obtaining a report for one table, but you can obtain reports that include data from multiple tables and the boot log as discussed in [“Saving Multiple Tables as a Report” on page 421](#).

To save the information of one clearVISN CoreWatch table to a report:

1. Open the table that you want to include in the report.
2. Click the **Report** button  on the Table toolbar.
The **Save As** dialog box appears.
3. Enter a name for the report in the **File Name** box. If necessary, browse to the folder in which you want to save the report. Then click the **Open** button.

If the specified file already exists, clearVISN CoreWatch prompts you to verify whether you want to overwrite the existing file. Clicking **OK** overwrites the file. Clicking **Cancel**, entering a different filename, and then clicking the **Open** button saves the information to another file.

clearVISN CoreWatch saves the report as an HTML file and displays that file in your Web browser. You may print the report using your Web browser's print feature.

Exporting Data from a Table

To export data from a clearVISN CoreWatch table to another application, such as Lotus 1-2-3 or Microsoft Excel, save it to an ASCII text file. To do so, take the following steps:

1. Open the table that contains the information you want to export.
2. Click the **Export** button  on the Table toolbar.
The **Save As** dialog box appears.
3. Enter a name for the report in the **File Name** box. If necessary, browse to the folder in which you want to save the report. Then click the **Open** button.

If the specified file already exists, clearVISN CoreWatch prompts you to verify whether you want to overwrite the existing file. Clicking **OK** overwrites the file. Clicking **Cancel**, entering a different filename, and then clicking the **Open** button saves the information to another file.

clearVISN CoreWatch saves the data to an ASCII text file that you can open in other applications.

Sorting Table Information

To sort information in a table, double-click the column heading you want to sort by. To switch between ascending and descending sort order within the column, double-click that column heading again.

Appendix B

clearVISN CoreWatch Menus

This appendix describes the following clearVISN CoreWatch menus that are located at the top of the clearVISN CoreWatch main window. Use the commands available on these menus to perform tasks in clearVISN CoreWatch.

- File
- Monitor
- Window
- Help

File Menu

The clearVISN CoreWatch File menu includes the commands described in the following table:

Table 76. File menu commands

Command	Description
Open Schematic View	Opens the Schematic view.
Reports	Lets you select which reports you want to generate.
Properties	Lets you change clearVISN CoreWatch properties. DIGITAL has configured the clearVISN CoreWatch properties to their optimal settings. Changing some of these properties may affect system performance.
GSR name or IP address	Opens the clearVISN CoreWatch main window for the GSR represented by the name or IP address you select.
Exit	Closes clearVISN CoreWatch.

Monitor Menu

Monitor menu provides access to the submenus described in the following table. When you select a submenu, clearVISN CoreWatch displays a list of commands or additional submenus. Separate discussions on each submenu follow.

Table 77. Monitor menu commands

Submenu	Description
Performance	Provides access to the System Dashboard dials, the Port Utilization Summary, and the clearVISN CoreWatch graphs that provide statistical information about the GSR traffic.
System State	Lets you display information about a GSR chassis and individual ports
Bridging State	Lets you display information about VLANs or STP bridges.
Routing State	Lets you display routing information for the following protocols: IP, IPX, OSPF, RIP, DVMRP, and IGMP.
QoS State	Lets you display information about Layer-2, Layer-3, and Layer-4 switching as well as QoS policies.

Performance State Submenu

If you select the Monitor menu and then choose Performance, the Performance submenu appears. That submenu offers the commands and additional submenus discussed in the following table. As the table indicates, you choose the System Dashboard and Port Utilization Summary commands directly from the Performance submenu, but you choose the other commands from another submenu.

Table 78. Performance State submenu commands

Submenu	Command	Description
	System Dashboard	Lets you select and then display the dials that permit you to monitor incoming and outgoing data.
	Port Utilization Summary	Displays information that indicates the status of each GSR port and identifies the percentage of traffic being transmitted and received on each of those ports.
Port	Packet Statistics	Displays the graph that lets you monitor a port's multicast, unicast, and broadcast packets.
	Byte Statistics	Displays the graph that lets you monitor a port's incoming and outgoing bytes.
	Error Statistics	Displays the graph that lets you monitor erroneous packets on a port. This graph also indicates how many packets the GSR discarded although no errors had been detected that would have prevented the GSR from delivering them to a higher-layer protocol.
IP	Packet Statistics	Displays the graph that lets you monitor IP packets the GSR sent, received, forwarded, or delivered.
	Reassembly Statistics	Displays the graph that lets you monitor reassembly and fragmentation activities on a GSR.
	Error Statistics	Displays the graph that lets you monitor IP errors on a GSR.
IPX	Packet Statistics	Displays the graph that lets you monitor IPX packets the GSR received or delivered. The graph also lets you monitor the number of requests the GSR received to send IPX data and the number of packets the GSR sent successfully.
	Error Statistics	Displays the graph that lets you monitor IPX errors on a GSR.

System State Submenu

If you select the Monitor menu and then choose System State, a submenu that includes the following commands appears:

Table 79. System State submenu commands

Command	Description
Chassis Table	Displays information about which modules are installed in the GSR, the slot number of each module, and the number of ports on each module.
Port Table	Displays information about which modules are installed in the GSR, the slot number of each module, or the number of ports on each module.
Trap Table	Displays information about traps that are currently configured for the GSR.
SmartTRUNK Table	Displays information about any ports that have been designated as SmartTRUNKs.

Bridging State Submenu

If you select the Monitor menu and then choose Bridging State, a submenu that includes the following commands appears:

Table 80. Bridging State submenu commands

Command	Description
VLAN Table	Displays information about the ports and modules associated with the VLANs of a GSR.
STP Table	Displays information about the ports associated with the STP bridges of a GSR.
L2 Interface Table	Displays information about ports that are currently serving as L2 interfaces on the GSR.

Routing State Submenu

If you select the Monitor menu and then choose Routing State, the Routing State submenu appears. That submenu provides access to the additional submenus listed in the following table. The table describes the commands available on these additional submenus.

Table 81. Routing State submenu commands

Submenu	Command	Description
IP State	IP Interface Table	Displays information about which port each IP interface uses and provides address information about those interfaces.
	IP Forwarding Table	Displays information about the routes used by the IP interfaces on a GSR.
IPX State	IPX Interface Table	Displays information about packets sent from a GSR's IPX interfaces, and indicates which WAN router each of those interfaces may use.
	IPX Forwarding Table	Displays information about the routes used by the IPX interfaces on a GSR.
OSPF State	OSPF Interface Table	Displays information about the routes, status, authentication, and other details about each OSPF interface on a GSR.
	OSPF Area Table	Displays information about the different areas defined for an OSPF autonomous system.
	OSPF Neighbor Table	Displays information about the neighboring routers of a GSR.
	OSPF Link State DB Table	Displays information about the link-state advertisements of the areas to which the GSR is attached.
	OSPF Area Aggregate Table	Displays information about the IP addresses and TOS that are used by OSPF to determine the destination of packets.
RIP State	RIP Interface Table	Displays information about how many packets on a RIP interface the GSR discarded, how many route entries of valid RIP packets the GSR ignored, how many RIP updates the GSR has sent on the interface, and whether a RIP interface is functioning.

Table 81. Routing State submenu commands (Continued)

Submenu	Command	Description
	RIP Peer Table	Displays information about RIP peers.
DVMRP State	DVMRP Interface Table	Displays information about the configuration of a GSR's DVMRP interfaces, whether those interfaces are functioning, how many packets were discarded on those interfaces, and how many route entries of valid DVMRP packets the GSR ignored.
	DVMRP Neighbor Table	Displays information about a GSR's DVMRP neighboring routers.
	DVMRP Route Table	Displays information about the multicast routes DVMRP uses instead of unicast routes.
	DVMRP Next Hop Table	Displays information about the next hop to which the GSR is sending IP multicast packets. This information includes details about a hop's source and indicates whether the hop has downstream neighbors.

Table 81. Routing State submenu commands (Continued)

Submenu	Command	Description
IGMP State	IGMP Interface Table	Displays information about the configuration of the interfaces on which IGMP is enabled.
	IGMP Cache Table	Displays information about the multicast groups of IGMP interfaces.

QoS State Submenu

If you select the Monitor menu and then choose QoS State, a submenu that includes the following commands appears:

Table 82. QoS State submenu commands

Command	Description
L2 QoS	Displays information about the QoS policies of a GSR.
L3/L4 QoS	Displays information about Layer-3 and Layer-4 QoS policies of a GSR.
L2 Flows	Displays information about the forwarding of Layer-2 data sent directly to a port rather than being sent to the Control Module for further processing.
L3/L4 Flows	Opens the Flow Table Filter form, which lets you control the contents of the Flow table. That table provides information about the data coming in to clearVISN CoreWatch from Layer 3 and Layer 4.

Window Menu

The clearVISN CoreWatch Window menu includes the commands described in the following table:

Table 83. Window menu commands

Command	Description
Tile Horizontally	Arranges all open windows in rows so that you can view them all at the same time.
Tile Vertically	Arranges all open windows in columns so that you can view them all at the same time.
Cascade	Arranges all open windows one in front of another. The title bar of each window remains visible, which can help you identify which windows are open.
Close All	Closes all the clearVISN CoreWatch windows open at the time you choose the command. This does not, however, close the clearVISN CoreWatch main window.

Help Menu

The clearVISN CoreWatch Help menu includes the commands and submenu discussed in the following table. As the table indicates, you choose some commands directly from the Help menu and you choose some commands from the DIGITAL Web Site submenu.

Table 84. Help menu commands

Submenu	Command	Description
	Contents and Index	Opens the clearVISN CoreWatch online help.
	Glossary	Opens the online help glossary for clearVISN CoreWatch.
	Release Note	Displays any release note information included with your version of clearVISN CoreWatch.
DIGITAL Web Site	Product News	Displays information about DIGITAL products.
	White Papers	Provides access to technical discussions of DIGITAL products.
	Frequently Asked Questions	Provides technical support for some issues or concerns that you may have while using clearVISN CoreWatch.
	Technical Support	Displays information about how to contact DIGITAL technical support.
	Send Feedback	Displays a form that you may use to let DIGITAL know what you think about its products. You can complete and send this form online.
	DIGITAL Home Page	Displays the DIGITAL home page in your Web browser.
	About clearVISN CoreWatch	Displays information about which version of clearVISN CoreWatch you are using.

Appendix C

Supported Regular Expressions

When controlling the contents of a clearVISN CoreWatch table, you may find it useful to enter a regular expression that clearVISN CoreWatch will use as a wildcard. clearVISN CoreWatch will then filter entries in the table based on the specified regular expression.

clearVISN CoreWatch supports the following Perl5 regular expressions:

- Alternatives separated by |
- The quantified atoms described in the following table:

Table 85. Supported Perl5 quantified atoms

Quantified Atom	Description
{n,m}	Match at least n but not more than m times.
{n,}	Match at least n times.
{n}	Match exactly n times.
*	Match 0 or more times.
+	Match 1 or more times.
?	Match 0 or 1 time.

- The following atoms:
 - Regular expression within parentheses
 - A . matches everything except \n
 - Character classes [such as (abcd) and ranges (such as (a-z)]

You may include special backslashed characters within a character class (except for back-references and boundaries).

To represent a backspace in a character class enter \b

 - a ^ is a null token that matches the beginning of a string or line
 - a \$ is a null token that matches the end of a string or line
 - The special backslashed characters described in the following table. Any other backslashed character matches itself.

Table 86. Supported special (backslashed) characters

Character	Description
\b	Null token that matches a word boundary (\w on one side and \W on the other)
\B	Null token that matches a boundary that is not a word boundary
\A	Match only at the beginning of string
\Z	Match only at the end of string or before a new line at the end
\n	New line
\r	Carriage return
\t	Tab
\f	Formfeed
\d	Digits from 0-9
\D	Non-digits
\w	Alphanumeric characters (0-9, a-z, and A-Z)
\W	Nonalphanumeric characters (0-9, a-z, and A-Z)
\s	The following white space characters: tab, new line, carriage return, and formfeed
\S	Characters other than the following white space characters: tab, new line, carriage return, and formfeed
\xnn	Hexadecimal representation of character

Table 86. Supported special (backslashed) characters (Continued)

Character	Description
\cD	Matches the corresponding control character
\nn or \nnn	Octal representation of character unless a back-reference.
\1, \2, \3, and so on	A back-reference, which matches whatever the first, second, third, and so on parenthesized group matched. If there is no corresponding group, the number is interpreted as an octal representation of a character.
\0	Matches the null character

- clearVISN CoreWatch matches expressions within parentheses as subpattern groups and uses certain methods to save those matches.

By default, a quantified subpattern matches as many times as possible without causing the rest of the pattern not to match. If you want the quantifiers to match the minimum number of times possible, without causing the rest of the pattern not to match, insert a question mark (?) after the quantifier as described in the following table:

Table 87. Supported quantifier subpattern matching

Expression	Description
*?	Match 0 or more times
+?	Match 1 or more times
??	Match 0 or 1 time
{n}?	Match exactly n times
{n,}?	Match at least n times
{n,m}?	Match at least n times but not more than m times

- clearVISN CoreWatch also supports all of the Perl5 extended regular expressions, which are described in the following table:

Table 88. Supported Perl5 extended regular expressions

Expression	Description
(?#text)	An embedded comment that you enter if you want text to be ignored.
(?:regexp)	Groups items such as “()” without causing the group match to be saved.
(?=regexp)	A zero-width positive lookahead assertion. Suppose you enter <code>\w+(?=\s)</code> . clearVISN CoreWatch then matches a word followed by white space, without including white space in the resulting match.
(?!regexp)	A zero-width negative lookahead assertion. Suppose you enter <code>port(?!7)</code> matches any occurrence of “port” that is not followed by “7”. Because this is a zero-width assertion, entering <code>a(?!b)d</code> will cause clearVISN CoreWatch to match <code>ad</code> because <code>a</code> is followed by a character that is not <code>b</code> (the <code>d</code>) and <code>a d</code> follows the zero-width assertion.
(?imsx)	Embedded pattern-match modifiers. You may enter one or more of these modifiers. <code>i</code> enables case insensitivity, <code>m</code> enables multiline treatment of the input, <code>s</code> enables single line treatment of the input, and <code>x</code> enables extended white space comments.

Appendix D

Error Messages

This appendix describes error messages you may encounter while using clearVISN CoreWatch and Configuration Expert. This appendix also includes possible solutions to the errors. The error messages are presented alphabetically within the following categories:

- Missing or invalid field error messages
- Duplicate objects error messages
- Already exists or in use error messages
- Unavailable objects error messages
- Miscellaneous error messages

Missing or Invalid Field Error Messages

The following error messages are generated when configuring new objects or modifying existing objects through configuration wizards or property sheets to indicate that a particular field has a missing or invalid value.

Solution:

To resolve any of the following error messages, enter a valid value in the appropriate field to continue creating or changing the configuration of the corresponding object.

'ACL Name' is missing or invalid. Specify a valid 'ACL Name' to continue.

A 'Port' is missing. Add a port to continue.

VLAN ID is missing or invalid. Specify VLAN ID as an integer value to continue.

'Autonomous System Patch Matching' field is missing or invalid. Specify a regular expression for the 'Autonomous System Path Matching' to continue.

'Broadcast Address' is missing or invalid. Specify a valid 'Broadcast Address' to continue.

'Broadcast Interface Address' is missing or invalid. Specify a valid 'Broadcast Interface Address' to continue.

'Community String' is missing or invalid. Specify a valid 'Community String' to continue.

'Cost' is missing or invalid. Specify a valid 'Cost' value to continue.

'Destination MAC Address' is missing or invalid. Specify a valid 'Destination MAC Address' (example: 00:20:3F:09:3C:F0) to continue.

'Destination Mask' field is missing or invalid. Specify a valid 'Destination Mask' as a IP subnet mask (example: 255.255.255.0) to continue.

'Destination' name is missing or invalid. Specify a valid 'Destination' name to continue.

'Export Destination' is missing or invalid. Specify a valid 'Export Destination' to continue.

'Export Destination' is missing or invalid. Specify a valid 'Export Destination' to continue.

'Export Source' is missing or invalid. Specify a valid 'Export Source' to continue.

'Export Source' is missing or invalid. Specify a valid 'Export Source' to continue.

'Filter' name is missing or invalid. Specify a valid 'Filter' name to continue.

'Flow Name' is missing or invalid. Specify a valid 'Flow Name' to continue.

'Gateway IP Address' is missing or invalid. Specify a valid 'Gateway IP Address' to continue.

'Gateway IP Address' is missing or invalid. Specify a valid 'Gateway IP Address' to continue.

'Gateway IP Address' of 0.0.0.0 is invalid. Specify a valid Gateway IP Address to continue.

'Having Tag' field is missing or invalid. Specify a valid value for 'Having Tag' field to continue.

'Interface Name' is missing or invalid. Specify a valid 'Interface Name' to continue.

IP Address entered is missing or invalid. Specify a valid IP address to continue.

IP Address of 0.0.0.0 is invalid. Specify a valid IP address to continue.

'IP Interface Name' is missing or invalid. Specify a valid 'IP Interface Name' to continue.

'IP Network/Host Address' is missing or invalid. Specify a valid IP Network or Host address to continue.

'IP Subnet Mask' is missing or invalid. Specify a valid 'IP Subnet Mask' (example: 255.255.255.0) to continue.

'IP Tunnel' is missing or invalid. Specify the 'IP Tunnel' name to continue.

'Learned from Autonomous System' field is missing or invalid. Specify a valid value for the 'Learned from Autonomous System' to continue.

'Local Address' and 'Remote Address' fields are missing or invalid. Specify a valid IP address for each to continue.

'MAC Address' is missing or invalid. Specify a valid 'MAC Address' (example: 00:20:3F:09:3C:F0) to continue.

'Match Length' is missing or invalid. Specify valid values for both the 'Match Length' fields to continue.

'Metric' value is missing or invalid. Specify a valid 'Metric' value to continue.

'NBMA Neighbor' is missing or invalid. Specify a valid 'NBMA Neighbor' to continue.

'Network Destination' address is missing or invalid. Specify a valid 'Network Destination' address to continue.

'Network Source' address is missing or invalid. Specify a valid 'Network Source' address to continue.

'Non Broadcast Interface Address' is missing or invalid. Specify a valid 'Non Broadcast Interface Address' to continue.

'Priority' value is missing or invalid. Specify a valid 'Priority' value to continue.

'Protocol' name is missing or invalid. Specify a valid 'Protocol' name to continue.

Range value entered is incorrect. The first value in the range should be less than the second value.

'Rate' is missing or invalid. Enter a valid 'Rate' value to continue.

'Route Export Condition' is missing or invalid. Specify a valid 'Route Export Condition' value as a regular expression to continue.

SNMP 'Trap Target' IP address is missing or invalid. Specify a valid SNMP 'Trap Target' IP address to continue.

'Source MAC Address' is missing or invalid. Specify a valid 'Source MAC Address' (example: 00:20:3F:09:3C:F0) to continue.

'Source Mask' field is missing or invalid. Specify a valid 'Source Mask' as a IP subnet mask (example: 255.255.255.0) to continue.

'Source Tag' is missing or invalid. Specify a valid 'Source Tag' value to continue.

'Source' name is missing or invalid. Specify a valid 'Source' name to continue.

'Stub Cost' is missing or invalid. Specify a valid 'Stub Cost' to continue.

'Tag' is missing or invalid. Specify a valid 'Tag' to continue.

'Threshold' value is missing or invalid. Specify a valid 'Threshold' value to continue.

'Timeout' value is missing or invalid. Specify a valid 'Timeout' value to continue.

'Virtual Link Neighbor' field is missing or invalid. Specify a valid 'Virtual Link Neighbor' IP address to continue.

'Virtual Link Transit Area' field is missing or invalid. Specify a valid 'Virtual Link Transit Area' IP address to continue.

'Virtual Link' name is missing or invalid. Specify a valid 'Virtual Link' name to continue.

'VLAN ID' is missing or invalid. Specify a valid 'VLAN ID' to continue.

'VLAN Name' is missing or invalid. Specify a valid 'VLAN Name' to continue.

Duplicate Objects Error Messages

The following error messages are generated when you attempt to create or modify names of objects that conflict with another existing object.

Solution:

To resolve any of the following error messages, provide a unique name of that object to continue.

Duplicate 'Destination' name error. Specify a unique 'Destination' name to continue.

Duplicate 'Filter' name error. Specify a unique 'Filter' name to continue.

Duplicate 'Network Address' error in the selected Aggregate Destination. Specify a new 'Network Address' to continue.

Duplicate 'Network Address' error in the selected filter. Specify a unique 'Network Address' to continue.

Duplicate 'RIP Source Gateway' IP address error. Specify an unused 'RIP Source Gateway' IP address to continue.

Duplicate 'RIP Trusted Gateway' IP address error. Specify an unused 'RIP Trusted Gateway' IP address to continue.

Duplicate 'Source' name error. Specify a unique 'Source' name to continue.

Duplicate 'Virtual Link' name error. Specify a unique 'Virtual Link' name to continue.

Duplicate 'VLAN ID' error. Specify an unused 'VLAN ID' to continue.

Already Exists or in Use Error Messages

The following error messages indicate that an object is either in use or already exists.

Solution:

To resolve any of the following error messages, either use another object to continue or re-examine the existing configuration to ensure consistency.

A 'Filter' by this name already exists. Specify a unique 'Filter' name to continue.

A 'L2 Flow' by this name already exists. Specify a unique 'L2 Flow' name to continue.

An ACL by this name already exists. Select a unique 'ACL Name' to continue.

An 'IP Flow' by this name already exists. Specify a unique 'IP Flow' name to continue.

An 'IPX Flow' by this name already exists. Specify a unique 'IPX Flow' name to continue.

An 'IPX Interface' by this name already exists. Specify a unique 'IPX Interface' name to continue.

Specified 'OSPF Area' already exists. Specify a different 'OSPF Area' to continue.

This IP Address has already been used. Specify another IP Address to continue.

This IP Address is already in use. Specify a different IP Address to continue.

This Network Address has already been used. Specify another Network Address to continue.

This 'IP Interface Name' already exists. Specify a unique 'IP Interface Name' to continue.

This 'IP Tunnel' already exists. Retry with another 'IP Tunnel' name to continue.

This 'Remote Address' is in use. Select another 'Remote Address' to continue.

This 'VLAN' already exists. Specify another VLAN to continue.

'Non Broadcast Interface Address' already exists. Specify a different 'Non Broadcast Interface Address' to continue.

'Stub Host' already exists. Specify a different 'Stub Host' to continue.

'Virtual Link' already exists. Specify a different 'Virtual Link' to continue.

Unavailable Objects Error Messages

No more IP addresses available to configure a new DVMRP Interface.

Existing DVMRP interfaces may have used up existing IP addresses.

Solution:

Try to assign a new IP address to configure a new DVMRP interface.

No more IP addresses available to configure new IGMP Interface.

Existing IGMP interfaces may have used up all existing IP addresses.

Solution:

Try to assign a new IP address to configure a new DVMRP interface.

No Access Control List (ACL) available, ... Define one or more Access Control Lists (ACLs) first and then retry this operation.

No ACL is available which is required for this operation to complete.

Solution:

Create one or more ACLs. (If you are trying to apply an ACL while creating an interface and that ACL is not currently available, you may finish creating the interface and then apply the ACL later.)

No IPX ACL available.

This configuration requires IPX ACLs which are not available.

Solution:

Create one or more IPX ACLs. (If you are trying to apply an ACL while creating an interface and that ACL is not currently available, you may finish creating the interface and then apply the ACL later.)

Miscellaneous Error Messages

All rows satisfy filter parameters: {0} rows displayed

Indicates that clearVISN CoreWatch is not including anything in a table because all of the table's data was filtered out. This message appears in a clearVISN CoreWatch table's status bar if the filter criteria you enter in the table's filter form was broad enough to filter out all entries in that table.

Solution:

Enter less specific filter criteria in the table's filter form.

Already Logged in

Indicates you tried to log in after already having done so.

This error message is generally indicative of some network problem or other error that puts clearVISN CoreWatch in the wrong state. clearVISN CoreWatch sometimes displays this message if you try to log in and the GSR did not clean up an earlier session because it was busy. Normally, however, this should not occur.

Solution:

If the operation you are attempting needs a login, restart clearVISN CoreWatch.

Another user at <IP address> is currently configuring the Switch. If you terminate that session the other user will lose all changes which have not been committed. Do you wish to end the other users session?

Indicates that you tried to log in to Configuration Expert while another user is already running clearVISN CoreWatch in the Privileged mode to configure the GSR.

Solution:

Specify whether you want to terminate the other user's session by clicking either the Yes or No button. If you click No, the other user's session will not be terminated and Configuration Expert is not started. Clicking Cancel closes the error message.

At least one 'Port' must be added to continue.

You have not added a port to an object that requires at least one port.

Solution:

Add a port to the object.

Cannot connect to switch

Indicates that the GSR could not be reached when you tried to log in. This could be because of bad connection, the GSR is not on, or another such problem.

Solution:

Check whether the machine is reachable by sending a ping packet. Also verify that you have entered the correct community string and log in password. Then try logging in again.

Cannot display an empty report

Indicates that clearVISN CoreWatch cannot generate a report because the table associated with that report is empty.

Solution:

Generate the report after the table contains some data.

Cannot Delete Selected Object. This object is being referenced by at least one other configuration object. To delete this object, all associated references to this object must be deleted first.

The selected object is being referenced by other objects and requires you to first delete these references before attempting to delete the object to ensure consistency.

Solution:

Eliminate the references. Otherwise, the configuration will be inconsistent and result in unknown behavior.

Cannot open file {0}. Privileges to write data not available

Indicates you do not have permission to write to a file.

Solution:

Check permissions and write to a writable file or directory.

Configure Password is invalid

Indicates that the wrong Privileged password was entered when you tried to start Configuration Expert.

Solution:

Start Configuration Expert again being sure to enter the correct Privileged password.

Could not Change password on Switch

Indicates that the clearVISN CoreWatch Login password or the Change Privileged Password could not be changed on the GSR. This could be because the new one was not valid or the switch could not be reached or some unknown error.

Solution:

Try changing the password again.

Export Data: Error in opening file {0} to export data

clearVISN CoreWatch could not open the specified export file.

Solution:

Try exporting the table again. If that attempt fails, try to export to a different file.

Export Data Error: Table not instantiated

Indicates that clearVISN CoreWatch failed to obtain table data and so the export operation was stopped.

Solution:

Try exporting the table data again.

Incorrect Total Percentage - the sum of percentage values of all fields should be equal to 100%.

In configuring QoS Global Queuing Discipline, which uses the weighted fair queuing method, the total bandwidth reservation for each of the priority levels (control, high, medium and low) should add up to 100 percent.

Solution:

Adjust the priority levels so they add up to 100 percent.

Monitor Password is invalid

Indicates the wrong password was entered when you started clearVISN CoreWatch.

Solution:

Enter the correct clearVISN CoreWatch Login password.

No IPX ACL available for this configuration. Create IPX ACLs first and then retry this operation.

During creation of IPX interfaces, you may optionally bind IPX ACLs. An attempt was made to bind an IPX ACL to this IPX interface when no IPX ACLs exist.

Solution:

Finish creating the IPX interface, create the IPX ACLs, and then apply those IPX ACLs to the IPX interface.

No more data available

Indicates there are no more rows in the table. This message appears in a table's status bar if you choose the Get next rows button and there are no more rows in the table.

Solution:

Do not press the Get next rows button again unless there are additional entries in the table.

Old Password is not correct

The wrong password was entered in the Old Password text box of the Change Login Password form or Change Privileged Password form.

Solution:

Enter the correct password. If you are changing the clearVISN CoreWatch Login password, enter the one you were prompted for when you last started clearVISN

CoreWatch. If you are changing the Privileged password, enter the one that currently provides access to Configuration Expert.

Password Re-entered and New Passwords do not match

Different passwords were entered in the New Password and the New Password Re-entry text boxes of the Change Login Password form or Change Privileged Password form.

Solution:

Enter the same password in both the New Password and the New Password Re-entry text boxes

Request denied by switch

Indicates that the GSR denied a connection when you tried to log in.

Solution:

Try logging in again. If that attempt fails, contact the administrator responsible for the GSR.

Static ARP Entry port binding error. Select only one 'Port' per Static ARP Entry to continue.

Static ARP Entry requires only one port to be bound to it. An attempt was made to bind more than one port.

Solution:

Remove all but one port.

The table contains no data. Cannot create report

Indicates that clearVISN CoreWatch cannot generate a report because the table associated with that report is empty.

Solution:

Generate the report after the table contains some data.

This Configuration Wizard requires at least one 'Interface' to continue.

The configuration wizard requires you to select one interface to continue.

Solution:

Select an interface.

WARNING – REVIEW BEFORE YOU DELETE: The object you are about to delete could be referenced by other configuration objects. To ensure consistency, remove 'all'

references to this object before you delete it. Refer to clearVISN CoreWatch documentation for more information.

The selected object could be referenced by other configuration objects and requires you to first delete all the references before attempting to delete the object to ensure consistency.

Solution:

Eliminate the references. Otherwise, the configuration will be inconsistent and result in unknown behavior. Suppose you want to delete an IP ACL. Before doing so, check whether the IP ACL is applied to an IP interface. If it is, delete that reference before you delete the IP ACL.

Glossary

Access Control List (ACL)

List the GSR keeps to restrict Layer 3/4 traffic going through the router. Each ACL or each list consists of one or more rules describing a particular type of IP or IPX traffic. An ACL can be simple and consist of only one rule or complicated with many rules. Each rule tells the router to either permit or deny the packet that matches the rule's packet description. For more information on ACLs and rules, refer to the *DIGITAL GIGAswitch/Router User Reference Manual*.

Adjacency

An OSPF relationship formed between selected neighboring routers for the purpose of exchanging routing information. Not every pair of neighboring routers become adjacent.

Aggregation

The combining of the characteristics of several different routes in such a way that a single route can be advertised. Aggregation reduces the amount of information that the routers must store and exchange.

Area

A set of networks grouped together but located in the same autonomous system. Setting up areas reduces OSPF routing traffic because each area's topology is invisible to the rest of the autonomous system.

AS Path

Attribute composed of a sequence of autonomous system path segments. The paths can be defined in ordered or unordered sets. The sets include the autonomous systems a route has traversed. The sets are used in aggregation and reduce routing information by listing each system only once regardless of how many times those systems are included in different paths.

Authentication Key

Parameter that guarantees that routing information is only imported from trusted routers. Many protocols, such as RIP version 2 and OSPF, provide mechanisms for authenticating protocol exchanges. An authentication key permits generation and verification of the authentication field in protocol packets.

Autonomous System

A set of routers under a single technical administration, using an interior gateway protocol and common metrics to route packets within the autonomous system, and using an exterior gateway protocol to route packets to other autonomous systems.

Since this classic definition was developed, it has become common for a single autonomous system to use several interior gateway protocols and sometimes several sets of metrics within an autonomous system. The use of the term “autonomous system” stresses that even when multiple IGPs and metrics are used, the administration of an autonomous system appears to other autonomous systems to have a single coherent interior routing plan and to present a consistent picture of what networks are reachable through it.

Backup

Interface state that indicates a router is the backup designated router on the attached network.

Backup Designated Router

OSPF router that will take over the functions of the designated router if that one fails. The backup designated router establishes adjacencies to all other routers.

Boot Log

File containing the messages the GSR sends when starting up. You may want to examine these messages to look at the configuration of the GSR and determine which functions were online when you started the GSR.

Boundary Router

Router that exchanges routing information with routers belonging to other autonomous systems. Such a router has autonomous system external routes that are advertised throughout the autonomous system. The path to each autonomous system boundary router is known by every router in the autonomous system.

This classification is completely independent of the other classifications; autonomous system boundary routers may be internal or area border routers and may or may not participate in the backbone.

Branch

Interface for which downstream neighbors exist.

Bridge Protocol Data Unit (BPDU)

Hello packet STP periodically sends to exchange information among bridges on a network.

Broadcast

Packets sent to all of the nodes of a network.

Broadcast Network

Network supporting two or more attached routers and is capable of addressing a single physical message to all of those routers.

Neighboring routers are discovered dynamically on broadcast networks using OSPF's Hello Protocol. The Hello Protocol takes advantage of the broadcast capability. The protocol also makes use of existing multicast capabilities.

Ethernet is an example of a broadcast network.

Command Line Interface (CLI)

Set of commands that let you monitor and configure GSRs.

Configuration Expert

A utility that lets you perform the following tasks:

- Change system information
- Configure bridging polices
- Configure VLANs
- Configure IP and IPX routing
- Configure multicast routing
- Set QoS policies
- Set ACLs and security filters
- Configure multiple configuration files on the GSR

Designated Bridge

Bridge responsible for forwarding frames to the LAN segment.

Designated Router

OSPF router that generates the network's link-state advertisement, establishes adjacencies, and has other special responsibilities in the running of the protocol. Designated routers reduce the number of adjacencies required on a multi-access network. This in turn reduces the amount of routing protocol traffic and the size of the topological database.

Each multiaccess network that has two or more attached routers has a designated router.

DIGITAL GIGAswitch/Router (GSR)

DIGITAL product that is capable of switching traffic at Layer-2, Layer-3, and Layer-4. The GSRs provide full-function routing at Gigabit speeds, pinpoint control over application usage, and can handle enterprise and ISP backbone traffic.

Distance Vector Multicast Routing Protocol (DVMRP)

An IP multicast routing protocol. It is a distance-vector based protocol (similar to RIP) that keeps its own unicast routing information.

Down

Interface state that indicates an interface is unusable. The interface cannot send or receive traffic.

DR

OSPF interface state that indicates the router itself is the designated router on the network to which the router is attached.

DR Other

OSPF interface state that indicates that the interface's router is neither the designated router nor the backup designated router. The interface's router forms adjacencies to both the designated router and the backup designated router (if applicable).

Exterior Gateway Protocol (EGP)

Internet protocol used to exchange of routing information between autonomous systems.

Flow

Template for specifying the characteristics of network traffic.

Front Panel View

Graphical representation of a GSR's front-panel chassis. For an example Front Panel view, see ["Front Panel View" on page 12](#).

Generation Identifier

Instance identifier assigned to DVMRP whenever that protocol is started on a router.

Hello Packets

Packets sent to acquire neighbors, which are routers on the same network as the router sending the packet.

Hop Count

Routing metric that measures the distance between a source and destination. A hop is the moving of packets through a router interface. Data transmitted directly from one router to another router counts as one hop. Suppose data is being transmitted between router A and router D. If the data must go through routers B and C before it gets to router D, the hop count for the route is three.

Protocols that base routing on the hop count will use the route with the least amount of hops.

Hypertext Markup Language (HTML)

Language that describes the structure of a Web document's contents and defines some of that document's behavior. HTML permits a file to be linked to other resources, such as text files, graphic files, sound files, and so on.

Interior Gateway Protocol (IGP)

Internet protocol that supports the exchange of routing information within an autonomous system.

Internet Control Message Protocol (ICMP)

Protocol that reports IP packet errors and provides other information about the processing of IP packets.

Internet Group Management Protocol (IGMP)

Protocol IP hosts use to report their host group memberships to any multicast routers to which the IP hosts are connected.

Internet Group Management Protocol (IGMP) Host-Query Packet

Packet an IGMP router sends to hosts to learn which hosts are available. Host queries help a router determine changes to host membership.

Internetwork Packet Exchange (IPX)

A datagram connectionless protocol.

Internet Protocol (IP)

Layer-3 (network) protocol that allows connectionless internetworking. IP uses best-effort delivery to route packets through any number of paths. It fragments and reassembles packets, but does not guarantee delivery of those packets. TCP is responsible for guaranteeing delivery.

Internet Service Provider (ISP)

Company that provides access to the Internet.

Leaf

Interface for which no downstream dependent neighbors exist.

Link-State Advertisement (LSA)

Packet that describes the local state of a router or network. This includes the state of the router's interfaces and adjacencies. Each link-state advertisement is flooded throughout the routing domain. The collected link-state advertisements of all routers and networks form the protocol's topological database.

Local Area Network (LAN)

Network that connects workstations, terminals, printers, and other devices located in a small geographical area (such as a single building).

Login Password

Password you are prompted for when you start clearVISN CoreWatch.

Loopback

Interface state that indicates the router's interface to the network is looped back. Although this makes the router unavailable for regular traffic, you can still send ICMP pings to the router or perform error-bit tests.

Lossy

Trait of a network likely to lose data when it becomes heavily loaded.

MAC Address

Address of a port or computer that other devices use to locate those ports or computers. These addresses are also used to create and update routing tables.

Management Information Base (MIB)

Network management information stored in a database. SNMP uses and maintains this information.

mtrace

Command that tracks the multicast path from a source to a receiver.

Multi-Access Network

Network to which multiple routers are attached.

Multicast Group

Set of IP hosts that listen to an IP multicast address. It is possible for members of the group to be spread across separate physical networks.

In IP multicasting, data is sent to this single IP address rather than being sent to each host's individual IP address. These addresses are in the range of 224.0.0.0 to 239.255.255.255.

Multicast Packets

Individual packets sent to a network address that represents a single multicast group.

Multicast Routing

Method of routing in which individual packets are sent to many destinations.

In IP multicast routing, data is sent from one or more routers to a multicast group. The data originates as a single packet destined to a multicast group. This differs from unicast routing in which a router sends multiple unicast packets so that there is one packet for each destination.

Neighboring Routers

Two routers that have interfaces to a common network.

NetWare

Network operating system developed by Novell that provides various network services, such as transparent remote file access.

Network Interface Card (NIC)

Adapter that provides a computer's networking capability.

Nonbroadcast Multiaccess (NBMA) Networks

Network that supports two or more attached routers but does not have broadcast capability.

Open Shortest Path First (OSPF)

OSPF is a link-state routing protocol that supports the distribution of routing information between routers belonging to a single autonomous system.

Path Cost

A number from 1 to 65535. This number indicates how much a port is contributing to the total cost of the path to the root bridge when the port is the root port.

Point-to-Point

OSPF interface state that indicates the interface is operational and is connected to either a physical point-to-point network or to a virtual link. Upon entering this state, the router attempts to form an adjacency with its neighboring router. The router sends hello packets.

Point-to-Point Network

A network that joins a single pair of routers.

Port Utilization Summary

Use the Port Utilization Summary to identify the traffic patterns of GSR ports. This summary indicates the percentage of traffic that is being transmitted and received on each port. For a detailed description of an example summary, see [“Monitoring Port Utilization” on page 356](#).

Precedence Value

QoS parameter associated with the IP or IPX packet fields that GSR uses to break ties if packets match more than one flow. You assign a precedence from 1 to 7 with each field. The highest precedence value is 1 and the lowest is 7.

Privileged Password

Password you are prompted for when you start Configuration Expert. This password logs you in to Configuration Expert so that you can then use that utility to configure your GSR.

Probe Messages

Messages sent in a reverse path from the receiver back to the source. As a probe message passes from hop to hop, it collects information such as interface address and packet counts from each router.

Pruning

Operation DVMRP routers perform to exclude interfaces not in the shortest path tree.

Quality of Service (QoS)

Set of parameters that do the following:

- Assign priorities to different types of traffic.
- Define flows that act as templates for some IP and IPX packet fields.
- Assign a precedence to the fields of the flows you define.
- Establish queuing policies to specify how the GSR handles the different traffic priorities.

Requests for Comments (RFCs)

Notes about the Internet that discuss a variety of topics regarding computing and computer communication. RFCs typically focus on networking protocols, concepts, procedures, and programs.

Reverse Path Forwarding (RPF)

Routing technique that finds the next neighbor by determining which router sent the GSR the packet.

RIP Peer

Routers from which the GSR has received a valid RIP update within the last 180 seconds.

Root Bridge

Bridge that forms the root of a bridged LAN. Root bridges use STP to prevent loops by periodically exchanging topology information with other bridges.

Routing Information Protocol (RIP)

Distance-vector routing protocol that selects the route with the lowest metric as the best route. The metric is a hop count representing the number of gateways through which data must pass in order to reach its destination. The longest path that RIP accepts is 15 hops.

Schematic View

The Schematic view is a graphical representation of a GSR's functions (such as bridging, switching, and routing services) and data objects (such as QoS flows). It also indicates which functions are active, inactive, or in error. For details on opening this view and using it to perform many clearVISN CoreWatch tasks, see [“Schematic View” on page 14](#).

Simple Network Management Protocol (SNMP)

Protocol that provides support for monitoring and controlling network devices, collecting statistics, and managing configurations, performance, and security. SNMP is mainly used by TCP/IP networks.

Simple Network Management Protocol (SNMP) Community

Logical group of network devices that are on the same domain.

Simple Network Management Protocol (SNMP) Trap Log

A file containing the SNMP traps the GSR sends. An SNMP trap is a message describing an event (such as restarting the GSR or a link going down) that was detected by a GSR function (such as a MIB or OSPF).

Spanning Tree Protocol (STP)

Protocol that enables a bridge to create a spanning tree so that the bridge can dynamically work around loops. If a bridge detects a loop, it removes that loop by shutting down bridge interfaces.

System Dashboard

Set of dials that indicate the following information:

- Bits per second at which the GSR is sending and receiving
- Packets per second at which the GSR is sending and receiving

clearVISN CoreWatch lets you select which of these dials you want to display.

Tick

Routing metric that is approximately 1/18th of a second.

In IPX, the distance to a destination is determined by how long it takes to arrive at that destination. This delivery time is usually one second on Ethernet LANs.

Transmission Control Protocol (TCP)

The Internet transport layer protocol that provides reliable communication over packet-switched networks.

Transmission Control Protocol/Internet Protocol (TCP/IP)

Suite of protocols that provide a relatively simple way to connect computers and devices from different vendors on a worldwide internetwork.

Trap

Message the GSR sends that describes an event (such as restarting the GSR or a link going down). These events are detected by different GSR functions (such as a MIB or OSPF) of the GSR. The GSR maintains an SNMP trap log if your system administrator has configured it to do so.

Tunnel

Used to send packets between routers separated by gateways that do not support multicast routing. A tunnel acts as a virtual network between two routers running DVMRP. A tunnel does not run IGMP.

Type of Service (TOS)

Used by some protocols, such as OSPF, to calculate separate routes for each type of service. There can be multiple routes to a given destination, one for each type of service (such as, low delay or high throughput). When routing a packet, a router running OSPF

uses both the destination address and type of service fields in an IP header to choose a route.

Unicast Packet

Packet sent to a single destination. The packet is going from one point to another point.

Unicast Routing

Routing method in which a packet is sent to a single host. This differs from multicast routing in which individual packets are sent to many destinations.

Update Message

Message a router sends to inform its peers of which routes are currently available and which ones are unfeasible.

User Mode

Mode in which you can monitor the GSR. This mode does not permit you to make configuration changes (such as changing passwords, configuring bridges or routers, or setting security filters).

If you plan to make configuration changes while in clearVISN CoreWatch, you may do so by starting Configuration Expert from within clearVISN CoreWatch. You will be prompted for the Privileged password.

User Datagram Protocol (UDP)

Connectionless protocol in the TCP/IP protocol stack. It performs the same delivery role as TCP, but does not guarantee delivery.

UDP is useful for sending small amounts of data in environments in which you can easily resend data if delivery fails.

Virtual Local Area Network (VLAN)

Devices on different segments of a LAN that are configured to communicate as though they are on the same wire. VLANs are very flexible because they are logical connections rather than physical connections.

Waiting

OSPF interface state that indicates the router is trying to determine the identity of the Backup Designated Router for the network. The router cannot elect neither a Backup Designated Router nor a Designated Router until it transitions out of Waiting state. This prevents unnecessary changes of a (Backup) Designated Router.

Wide Area Network (WAN)

A network extending over broad distances.

Index

Symbols

- ...
- ... obtaining information 351
- ... table 351
- ... Table command 432
- ...s
 - adding ports 66
 - creating 65–67
 - removing ports 67

A

- About clearVISN CoreWatch command 437
- Access Control Lists. *See* ACLs
- access ports
 - defining 71–73
 - in VLANs 70
- accessing online help 17
- ACL Rule Definition dialog box 225
- ACLs
 - adding rules 225
 - applying 96, 101, 147, 152
 - applying by editing definitions 208
 - applying multiple 207
 - changing names 224
 - copying to apply 207
 - implicit deny rule 207
 - IP 190–194
 - IPX 194–199
 - IPX RIP 199–202
 - IPX SAP 203–206
 - modifying rules 225
 - overview 189
- adding
 - ACL rules 225
 - interfaces to flows 185
 - ports to ...s 66
 - ports to VLANs 77, 81
 - source gateways 127
 - trusted gateways 126–127
- address filters

- configuring 211–213
- defined 210
- Address Resolution Protocol. *See* ARP
- address-based bridging 48
- Aggregate Export Source dialog box 304
- Aggregate Source Definition dialog box 311
- aging state
 - default 52–53
 - disabling on ports 54–56
 - enabling on ports 56–58
 - of GSR bridging 52–58
 - overriding default 53–54
- applying
 - ACLs 96, 101, 147, 152, 206–210
 - flows to interfaces 186
- area aggregate information, OSPF 396–397
- area information, OSPF 391–392
- ARP entries, defining 115–117
- ARP wizard 115
- associating precedences 163–166
- auto negotiation 34, 35
- Autonomous System Path Export Source dialog box 297

B

- BGP Export Destination dialog box 287
- BGP Export Source dialog box 295
- BGP Import Source Definition dialog box 309
- BGP Peer Group Definition dialog box 271
- Bound Port List dialog box 87, 187, 226
- bridge tables, checking status 373–377
- bridging
 - address-based 50–52
 - and STP 58
 - configuring 47–64
 - flow-based 49–50
 - overview 47
- Bridging Mode dialog box 49, 50
- bridging mode, ports 48–50
- Bridging State submenu 432

Bridging STP dialog box 62, 63
Bridging VLAN Mode dialog box 71
browser requirements 2
Byte Statistics command 431
byte statistics, obtaining 360

C

Cascade command 436
changes, committing 28
changing

- ACL name 224
- interface list of flows 184–186
- port list of flows 186
- QoS profiles 181–187
- system settings 31–45

Chassis Info table 347
Chassis Table command 432
chassis, obtaining information 347
checking

- DVMRP routing status 401–407
- IGMP routing status 407–411
- IP routing status 379–382
- IPX routing status 382–386
- OSPF routing status 387–397
- RIP routing status 398–400
- system status 347–351

checking the status

- of bridge tables 373–377
- of QoS tables 413–420
- of routing tables 379–411

clearVISN CoreWatch

- access modes 3
- basics 9–18
- commands 429–437
- exiting 18
- features 2
- installing
 - in Solaris 5–6
 - in Windows 6
- interface 12–16
- linking with HP OpenView 11
- linking with SPECTRUM Enterprise Manager 11
- menus 429–437
- online help 17
- overview 1–4
- requirements 2
- starting
 - in Solaris 10
 - in Windows 10

toolbar 16
views 12–15
Close All command 436
commands

- ... Table 432
- About clearVISN CoreWatch 437
- Byte Statistics 431
- Cascade 436
- Chassis Table 432
- Close All 436
- Contents and Index 437
- DIGITAL Home Page 437
- DVMRP Interface Table 434
- DVMRP Neighbor Table 434
- DVMRP Next Hop Table 434
- DVMRP Route Table 434
- Error Statistics 431
- Exit 430
- Frequently Asked Questions 437
- Glossary 437
- IGMP Cache Table 435
- IGMP Interface Table 435
- IP Forwarding Table 433
- IP Interface Table 433
- IPX Forwarding Table 433
- IPX Interface Table 433
- L2 Flows 435
- L2 Interface Table 432
- L2 QoS 435
- L3/L4 Flows 435
- L3/L4 QoS 435
- Open Schematic view 430
- OSPF Area Aggregate 433
- OSPF Area Table 433
- OSPF Interface Table 433
- OSPF Link State DB 433
- OSPF Neighbor Table 433
- Packet Statistics 431
- Port Table 432
- Port Utilization Summary 431
- Product News 437
- Properties 430
- Reassembly Statistics 431
- Release Note 437
- Reports 430
- RIP Interface Table 433
- RIP Peer Table 434
- Send Feedback 437
- STP Table 432
- System Dashboard 431

- Technical Support 437
 - Tile Horizontally 436
 - Tile Vertically 436
 - Trap Table 432
 - VLAN Table 432
 - White Papers 437
 - committing configuration file changes 28
 - community strings
 - establishing 44–45
 - SNMP 43
 - Configuration Expert
 - basics 19–30
 - copying objects 26
 - deleting objects 27
 - dialog boxes 25
 - dragging objects 26
 - exiting 30
 - finding objects 26
 - icons 25
 - overview 19
 - starting
 - Front Panel view 20
 - in Solaris 21
 - in Windows 21
 - tasks 27
 - window 21
 - wizards 25
 - configuration files
 - committing changes 28
 - icons 25
 - loading 28
 - retrieving 29
 - subtree 22
 - viewing contents 22
 - configuration tasks, order 27
 - configuration tree
 - copying objects 26
 - deleting objects 27
 - description 22–24
 - dragging objects 26
 - finding objects 26
 - icons 25
 - configuring
 - address filters 211–213
 - address-based bridging 50–52
 - DVMRP interfaces 132–133
 - flow-based bridging 49–50
 - for a SYSLOG server 39–40
 - for DNS 41
 - for DVMRP 130–136
 - for SNMP 42–45
 - GSR bridging 47–64
 - IGMP interfaces 137–140
 - IP interfaces 89–111
 - IPX interfaces 141–157
 - ports 34–39
 - port-to-address lock filters 213–216
 - QoS 161–187
 - RIP 121–127
 - secure port filters 220–223
 - static IPX SAP entries 157–159
 - static-entry filters 216–220
 - the GSR 27
 - VLANs 69–87
 - VRRP 108–110
 - Contents and Index command 437
 - contents, controlling in tables 426
 - controlling contents of tables 439–442
 - copying
 - ACLs to interfaces 207
 - objects 26
 - CPU requirements 2
 - creating
 - ...s 65–67
 - IP interfaces 92–103
 - IP QoS profiles 167–172
 - IPX QoS profiles 172–177
 - Layer-2 QoS profiles 177–181
 - QoS profiles 166–181
 - VLANs
 - port-based 78–82
 - protocol-based 73–78
- ## D
- Default Aging Timeout dialog box 53
 - defining
 - access ports 71–73
 - ARP entries 115–117
 - DVMRP tunnels 133–135
 - IP RIP interfaces 123–126
 - static route entries 117–119
 - STP settings 59–60
 - trunk ports 71–73
 - deleting
 - interfaces from flows 185
 - objects 27
 - Dial Options dialog box 356
 - dialog boxes
 - ACL Rule Definition 225
 - Aggregate Export Source 304

- Aggregate Source Definition 311
- Autonomous System Path Export Source 297
- BGP Export Destination 287
- BGP Export Source 295
- BGP Import Source Definition 309
- BGP Peer Group Definition 271
- Bound Port List 87, 187, 226
- Bridging Mode 49, 50
- Bridging STP 62, 63
- Bridging VLAN Mode 71
- Default Aging Timeout 53
- Dial Options 356
- Direct Export Source 300
- DVMRP Global Parameters Configuration 131
- DVMRP Tunnel 135
- Edit ACL 105, 156, 224
 - filters 226
- Flow Table Filter 418
- Global Attributes of Ports 34
 - icon 25
- IGMP Global Parameters Configuration 138
- IGMP Interface Definition 139
- Interface Bound to DVMRP 132
- Interface Definition 104, 155
- IP QoS Precedence 164
- IP Static Route Definition 117
- IP Tunnel Definition 134
- IP Unicast Global Parameters 114
- IPX Policy Definition 183
- IPX QoS Precedence 165
- L2 Flow Priority Definition 184
- Login Dialog 10, 11
- Optional Path Attribute List 316
- OSPF Export Destination 285
- OSPF Export Source 294
- OSPF Import Source 308
- OSPF Policy ASE Defaults 282
- Physical Attributes of Port 38
- Policy Input Interface List 185
- Port Utilization 357
 - purpose 25
- QoS L3/L4 Flow Priority 182
- Queuing Discipline Configuration 162
- Report Selection 421
- RIP Export Destination 284
- RIP Export Source 292
- RIP Global Parameters 122
- RIP Import Source 306
- RIP Interface Definition 124
- RIP Policy Defaults 280
- SAP Entry 158
- Save As 422
- Select Dials to be Displayed 354
- Set Aging Timeout 53
- Set STP Port Specific Settings 61
- SNMP Community Strings 44
- SNMP Trap Target 43
- Static Export Source 302
- STP Global Settings 59
- System Configuration DNS 41
- System ID 32
- System Log 40
- Tag Export Source 299
- Update ACL List 104, 156
- VLAN Definition 83, 84
- VRRP Trace Options 107
- dials
 - selecting 354
 - System Dashboard 355
 - Total Bits In 354
 - Total Bits Out 354
 - Total Packets In 354
 - Total Packets Out 354
- DIGITAL Home Page command 437
- DIGITAL Web Site submenu 437
- Direct Export Source dialog box 300
- disabling
 - DVMRP on tunnels 135
 - IGMP on interfaces 138–140
 - ports 38
 - RIP 122
- disk requirements 2
- displaying
 - ... information 351
 - chassis information 347
 - DVMRP interface information 401
 - DVMRP neighbor information 403–404
 - DVMRP next hop information 406–407
 - DVMRP routing information 405–406
 - IGMP cache information 410–411
 - IGMP interface information 408–410
 - IP forwarding information 381
 - IP interface information 380
 - IPX forwarding information 385
 - IPX interface information 383
 - L2 interface information 376–377
 - Layer-2 priority information 413
 - Layer-2 switching information 417–418
 - Layer-3/4 flow priority information 415–416

- Layer-3/4 switching information 418–420
 - OSPF area aggregate information 396–397
 - OSPF area information 391–392
 - OSPF interface information 387
 - OSPF link-state database information 394–396
 - OSPF neighbor information 392–394
 - port information 348
 - RIP interface information 398
 - RIP peer information 399–400
 - STP Port information 374–376
 - trap information 350
 - VLAN information 373
 - Distance Vector Multicast Routing Protocol. *See* DVMRP
 - DNS. *See* Domain Naming System
 - Domain Naming System, configuring 41
 - dragging
 - interfaces to apply flows 186
 - ports to VLANs 85
 - to copy objects 26
 - VLANs 85
 - DVMRP
 - configuring 130–136
 - configuring interfaces 132–133
 - global parameters 131
 - neighbor information 403–404
 - next hop information 406–407
 - overview 129
 - routing information 405–406
 - routing status 401–407
 - tunnels 133–136
 - DVMRP Global Parameters Configuration dialog box 131
 - DVMRP Interface table 401
 - DVMRP Interface Table command 434
 - DVMRP interfaces, obtaining information 401
 - DVMRP Neighbor table 403
 - DVMRP Neighbor Table command 434
 - DVMRP Next Hop table 406
 - DVMRP Next Hop Table command 434
 - DVMRP Route Table command 434
 - DVMRP Routing table 405
 - DVMRP State submenu 434
 - DVMRP Tunnel dialog box 135
- E**
- Edit ACL dialog box 105, 156, 224
 - editing interface definitions 208
 - enabling
 - DVMRP on tunnels 135
 - IGMP on interfaces 138–140
 - ports 38
 - RIP 122
 - error messages 443–454
 - Error Statistics command 431
 - error statistics, obtaining 361, 366, 369
 - establishing
 - community strings 44–45
 - GSR queuing policy 162
 - Exit command 430
 - exiting
 - clearVISN CoreWatch 18
 - Configuration Expert 30
 - exporting
 - table data 428
- F**
- features of clearVISN CoreWatch 2
 - filters
 - configuring address 211–213
 - configuring port-to-address lock 213–216
 - configuring secure port 220–223
 - configuring static-entry 216–220
 - Layer-2 security 210–227
 - modifying ports 226
 - RIP 194
 - SAP 194
 - setting on Layer-2 210–227
 - setting up for RIP 199–202
 - setting up for SAP 203–206
 - finding objects 26
 - flow
 - priorities of IPX 174
 - priorities of Layer-2 179
 - priorities of Layer-3/4 168
 - Flow Priority table 415
 - flow priority, obtaining information 415–416
 - Flow table 419
 - Flow Table Filter dialog box 418
 - flow-based bridging 48
 - flows
 - adding interfaces to 185
 - deleting interfaces from 185
 - dragging interfaces to 186
 - field of IPX 175
 - fields of IP 169
 - fields of Layer-2 180
 - modifying list of interfaces 184–186
 - modifying port list of 186

- redefining IP 181
- redefining IPX 182
- redefining Layer-2 183
- folder icon 25
- forwarding information
 - IP 381
 - IPX 385
- Frequently Asked Questions command 437
- Front Panel view
 - description 12–14
 - legend 13
 - modules 13–14
 - starting Configuration Expert 20
- full-duplex 34

G

- gathering
 - statistics 371
- GigaSwitch/Router. *See* GSR
- Global Attributes of Ports dialog box 34
- global parameters
 - DVMRP 131
 - IGMP 137
 - RIP 121, 123
 - unicast routing 114–115
- global settings of ports 34–36
- Glossary command 437
- graphs
 - IP Error Statistics 366
 - IP Packet Statistics 363
 - IP Reassembly Statistics 364
 - IPX Error Statistics 369
 - IPX Packet Statistics 368
 - Port Byte Statistics 360
 - Port Error Statistics 362
 - Port Packet Statistics 359
 - toolbar 371

GSR

- administrator 32
- aging state 52–58
- and STP 58–64
- and SYSLOG server 39–40
- bridging 47–64
- community strings 44–45
- configuring for DVMRP 130–136
- configuring for RIP 121–127
- location 32
- name 32
- queuing policy 162
- security 189–227

- supported MIBs 4
- VLANs 69–87

H

- half-duplex 34
- hash mode 34, 36, 39
- Help menu 18, 430, 437
- help, accessing 17

I

- icons
 - configuration file 25
 - dialog box 25
 - module 25
 - object 25
 - port 25
 - wizard 25
- IGMP
 - configuring 137–140
 - global parameters 137
 - obtaining cache information 410–411
 - overview 136
 - routing status 407–411
- IGMP Cache table 410
- IGMP Cache Table command 435
- IGMP Global Parameters Configuration dialog box 138
- IGMP Interface Definition dialog box 139
- IGMP Interface table 408
- IGMP Interface Table command 435
- IGMP interfaces
 - obtaining information 408–410
- IGMP State submenu 435
- implicit deny rule 97, 207
- installing clearVISN CoreWatch
 - in Solaris 5–6
 - in Windows 6
- Interface Bound to DVMRP dialog box 132
- Interface Definition dialog box 104, 155
- interfaces
 - defining RIP 123–126
 - modifying list of flows 184–186
 - obtaining information 380, 383, 387, 398, 401, 408–410
 - replacing VLANs 85
- Internet Group Management Protocol. *See* IGMP
- Internet Protocol. *See* IP
- IP
 - ACLs 190–194

- addresses 90
 - creating QoS profiles 167–172
 - forwarding information 381
 - obtaining error statistics 366
 - obtaining packet statistics 363
 - obtaining reassembly statistics 364
 - overview 89–91
 - QoS precedences 163
 - routing status 379–382
 - setting security 190–194
 - IP address, community string 10, 11
 - IP addresses 90
 - IP Error Statistics graph 366
 - IP Forwarding table 381
 - IP Forwarding Table command 433
 - IP Interface table 380
 - IP Interface Table command 433
 - IP Interface wizard 93, 98
 - IP interfaces
 - applying ACLs 206–210
 - bound to port 92–98
 - bound to VLAN 98–103
 - configuring 89–111
 - creating 92–103
 - disabling 138–140
 - enabling IGMP 138–140
 - modifying 103–106
 - monitoring 363–367
 - obtaining information 380
 - IP Packet Statistics graph 363
 - IP QoS Precedence dialog box 164
 - IP Reassembly Statistics graph 364
 - IP Security wizard 190
 - IP State submenu 433
 - IP Static Route Definition dialog box 117
 - IP submenu 431
 - IP Tunnel Definition dialog box 134
 - IP Unicast Global Parameters dialog box 114
 - IPX
 - ACLs 194–199
 - creating QoS profiles 172–177
 - forwarding information 385
 - obtaining error statistics 369
 - obtaining packet statistics 368
 - overview 141
 - QoS precedences 165
 - routing status 382–386
 - setting security 194–199
 - IPX Error Statistics graph 369
 - IPX Forwarding table 386
 - IPX Forwarding Table command 433
 - IPX Interface table 383
 - IPX Interface Table command 433
 - IPX Interface wizard 143, 149
 - IPX interfaces
 - applying ACLs 206–210
 - bound to port 143–149
 - bound to VLAN 149–154
 - configuring 141–157
 - modifying 154–157
 - monitoring 367–370
 - obtaining information 383
 - IPX Packet Statistics graph 368
 - IPX Policy Definition dialog box 183
 - IPX QoS Precedence dialog box 165
 - IPX RIP
 - ACLs 199–202
 - setting security 199–202
 - IPX SAP
 - ACLs 203–206
 - configuring static entries 157–159
 - setting security 203–206
 - IPX Security wizard 194–206
 - IPX State submenu 433
 - IPX submenu 431
- L**
- L2 Flow Priority Definition dialog box 184
 - L2 Flows command 435
 - L2 Forward table 417
 - L2 interface
 - information 376–377
 - L2 Interface table 377
 - L2 Interface Table command 432
 - L2 Priority table 413
 - L2 QoS command 435
 - L2 Security wizard 211–223
 - L3/L4 Flows command 435
 - L3/L4 QoS command 435
 - Layer-2
 - changing flow port list 186
 - creating QoS profiles 177–181
 - modifying filters 225
 - obtaining priority information 413
 - obtaining switching information 417–418
 - setting filters 210–227
 - Layer-3/4
 - obtaining flow priority information 415–416
 - obtaining switching information 418–420
 - Layer-3/4 flows, associating precedences 163–

- 166
- legend, Front Panel view 13
- linking
 - HP OpenView 11
 - SPECTRUM Enterprise Manager 11
- link-state information, OSPF 394–396
- link-state protocol 229
- loading configuration files 28
- Login Dialog dialog box 10, 11
- Login passwords 10, 11

M

- MAC addresses, aging 52–58
- MAC encapsulation 94, 99, 104, 150
- Management Information Bases (MIBs) 4
- Maximum Transmission Unit. *See* MTU
- menus
 - commands 429–437
 - File menu 430
 - Help 18, 430, 437
 - Monitor 430
 - Window 436
- metric 122, 133, 135
- minus sign 23
- modifying
 - ACL name 224
 - ACL rules 225
 - filter port bindings 226
 - interface list of flows 184–186
 - IP interfaces 103–106
 - IPX interfaces 154–157
 - port list of flows 186
 - QoS profiles 181–187
 - security 223–227
 - VLANs 82–84
- modules
 - icon 25
 - in Front Panel 13–14
- Monitor menu 430
- monitoring
 - IP interface statistics 363–367
 - IPX interface statistics 367–370
 - port utilization 356–358
 - real-time performance 353–371
 - system performance 353–356
- MTU 94, 99, 104
- multicast routing 129–140

N

- neighbor information
 - DVMRP 403–404
 - OSPF 392–394
- next hop 118

O

- objects
 - copying 26
 - dragging 26
- obtaining
 - ... information 351
 - chassis information 347
 - DVMRP interface information 401
 - DVMRP neighbor information 403–404
 - DVMRP next hop information 406–407
 - DVMRP routing information 405–406
 - IGMP cache information 410–411
 - IGMP interface information 408–410
 - IP error statistics 366
 - IP forwarding information 381
 - IP interface information 380
 - IP packet 363
 - IP reassembly statistics 364
 - IPX error statistics 369
 - IPX forwarding information 385
 - IPX interface information 383
 - IPX packet statistics 368
 - L2 interface information 376–377
 - Layer-2 priority information 413
 - Layer-2 switching information 417–418
 - Layer-3/4 flow priority information 415–416
 - Layer-3/4 switching information 418–420
 - OSPF area aggregate information 396–397
 - OSPF area information 391–392
 - OSPF interface information 387
 - OSPF link-state database information 394–396
 - OSPF neighbor information 392–394
 - port byte statistics 360
 - port error statistics 361
 - port information 348
 - port packet statistics 359
 - port statistics 358–362
 - reports 421–423
 - RIP interface information 398
 - RIP peer information 399–400
 - STP port information 374–376
 - table records 427

- trap information 350
 - VLAN information 373
 - online help, accessing 17
 - Open Schematic View command 430
 - opening the Schematic view 15
 - OpenView
 - linking with clearVISN CoreWatch 11
 - linking with SPECTRUM 11
 - Operating mode 35
 - operating mode 34, 38
 - Optional Path Attribute List dialog box 316
 - OSPF
 - area aggregate information 396–397
 - area information 391–392
 - link-state database information 394–396
 - neighbor information 392–394
 - overview 229
 - routing status 387–397
 - OSPF Area Aggregate table 396
 - OSPF Area Aggregate Table command 433
 - OSPF Area table 391
 - OSPF Area Table command 433
 - OSPF Export Destination dialog box 285
 - OSPF Export Source dialog box 294
 - OSPF Import Source dialog box 308
 - OSPF Interface table 387
 - OSPF Interface Table command 433
 - OSPF interfaces, obtaining information 387
 - OSPF Link State DB table 394
 - OSPF Link State DB Table command 433
 - OSPF Neighbor table 393
 - OSPF Neighbor Table command 433
 - OSPF Policy ASE Defaults dialog box 282
 - OSPF State submenu 433
 - overriding default aging timeout 53–54
 - overview
 - of ACLs 189
 - of IPX 141
 - of OSPF 229
 - of QoS 161
- P**
- Packet Statistics command 431
 - packet statistics, obtaining 359, 363, 368
 - password, Login 10, 11
 - peer information, RIP 399–400
 - Performance submenu 431
 - Physical Attributes of Port dialog box 38
 - plus sign 23
 - Policy Input Interface List dialog box 185
 - Port Byte Statistics graph 360
 - Port Error Statistics graph 362
 - port list 24, 37
 - Port Packet Statistics graph 359
 - port speed 34, 35, 38
 - port statistics, obtaining 358–362
 - Port submenu 431
 - Port table 349
 - Port Table command 432
 - Port Utilization dialog box 357
 - Port Utilization Summary
 - example 357
 - monitoring 356–358
 - Port Utilization Summary command 431
 - port-based VLANs 70, 78–82, 83
 - ports
 - access 70
 - adding to ...s 66
 - adding to VLANs 77, 81, 85–87
 - address-based 50–52
 - blocking access 220–223
 - bound to IP interfaces 92–98
 - bound to IPX interfaces 143–149
 - bridging mode 48–50
 - configuring
 - global settings 34–36
 - individual 36–39
 - cost 60
 - defining STP attributes 60–61
 - disabling 38
 - disabling aging 54–56
 - disabling STP 63–64
 - enabling 38
 - enabling aging 56–58
 - enabling STP 61–63
 - flow-based 49–50
 - icon 25
 - modifying in filters 226
 - obtaining byte statistics 360
 - obtaining error statistics 361
 - obtaining information 348, 374–376, 376–377
 - obtaining packet statistics 359
 - overriding default aging 53–54
 - priority 60
 - removing from ...s 67
 - removing from VLANs 78, 82, 85, 86, 87
 - trunk 70
 - port-to-address lock filters
 - configuring 213–216
 - defined 210

- precedences, associating 163–166
- preference 118, 120, 122
- priorities
 - flow for IPX 174
 - flow for Layer-2 179
 - flow for Layer-3/4 168
 - strict 162
 - weighted-fair queuing 162
- priority, obtaining information 413
- Product News command 437
- Properties command 430
- protocol, link-state 229
- protocol-based VLANs 70, 73–78, 84
- providing system information 31–32

Q

- QoS
 - configuring 161–187
 - flow priority 415
 - IP precedences 163
 - IPX precedences 165
 - Layer-2 priority information 413
 - Layer-2 switching information 417
 - Layer-3/4 switching 418
 - overview 161
- QoS L3/L4 Flow Priority dialog box 182
- QoS profiles
 - creating 166–181
 - for IP 167–172
 - for IPX 172–177
 - for Layer-2 177–181
 - modifying 181–187
- QoS State submenu 435
- QoS tables, checking status 413–420
- QoS wizard 167, 177
- Quality of Service. *See* QoS
- Queuing Discipline Configuration dialog box 162
- queuing policy, establishing 162
- quitting
 - clearVISN CoreWatch 18
 - Configuration Expert 30

R

- RAM requirements 2
- rate 135
- real-time performance, monitoring 353–371
- Reassembly Statistics command 431
- reassembly statistics, obtaining 364
- records in tables 427

- redefining
 - IP flows 181
 - IPX flows 182
 - Layer-2 flows 183
- refreshing table information 427
- regular expressions 439–442
- Release Note command 437
- removing ports from ...s 67
- removing ports from VLANs 78, 82, 85, 86, 87
- replacing VLANs 85
- Report Selection dialog box 421
- reports
 - containing multiple tables 421
 - containing single table 423
 - containing the boot log 421
 - obtaining 421–423
 - saving a single table 423
 - viewing 421
- Reports command 430
- requirements
 - browser 2
 - clearVISN CoreWatch 2
 - CPU 2
 - hardware 2
- restoring filtered data 427
- retrieving configuration files 29
- RIP
 - configuring 121–127
 - defining interfaces 123–126
 - disabling 122
 - enabling 122
 - filters 194
 - global parameters 121, 123
 - overview 119–121
 - peer information 399–400
 - routing status 398–400
 - setting up filters 199–202
- RIP Export Destination dialog box 284
- RIP Export Source dialog box 292
- RIP Global Parameters dialog box 122
- RIP Import Source dialog box 306
- RIP Interface Definition dialog box 124
- RIP Interface table 398
- RIP Interface Table command 433
- RIP interfaces, obtaining information 398
- RIP Peer table 400
- RIP Peer Table command 434
- RIP Policy Defaults dialog box 280
- RIP State submenu 433
- routing

- multicast 129–140
 - unicast 113–128
 - Routing Information Protocol. *See* RIP
 - Routing State submenu 433
 - routing tables, checking status 379–411
 - rules
 - adding 225
 - modifying 225
 - of IP ACLs 191–194
 - of IPX ACLs 196–199
 - of IPX RIP ACLs 201–202
 - of IPX SAP ACLs 205–206
- S**
- Safety information
 - laser 6
 - SAP
 - filters 194
 - setting up filters 203–206
 - SAP Entry dialog box 158
 - Save As dialog box 422
 - saving
 - configuration file changes 28
 - multiple tables 421
 - single table 423
 - Schematic view
 - description 14–15
 - opening 15
 - using 15
 - secure port filters
 - combining with static entries 220
 - configuring 220–223
 - defined 210
 - security
 - configuring 189–227
 - modifying 223–227
 - setting on IP networks 190–194
 - setting on IPX networks 194–206
 - setting on Layer-2 210–227
 - Select Dials to be Displayed dialog box 354
 - selecting dials 354
 - Send Feedback command 437
 - Set Aging Timeout dialog box 53
 - Set STP Port Specific Settings dialog box 61
 - setting
 - DVMRP global parameters 131
 - global unicast parameters 114–115
 - IGMP global parameters 137
 - RIP global parameters 121, 123
 - setting up
 - default aging timeout 52–53
 - RIP filters 199–202
 - SAP filters 203–206
 - STP 58–64
 - targets for SNMP traps 42–43
 - Simple Network Management Protocol. *See* SNMP
 - SNMP
 - configuring for 42–45
 - setting up trap targets 42–43
 - SNMP Community Strings dialog box 44
 - SNMP Trap Target dialog box 43
 - Solaris
 - clearVISN CoreWatch requirements 2
 - installing clearVISN CoreWatch 5–6
 - starting clearVISN CoreWatch 10
 - starting Configuration Expert 21
 - sorting table information 428
 - source gateways, adding 127
 - Spanning Tree Protocol. *See* STP
 - specifying
 - VRRP Trace options 107
 - starting
 - clearVISN CoreWatch
 - in Solaris 10
 - in Windows 10
 - Configuration Expert
 - Front Panel view 20
 - in Solaris 21
 - in Windows 21
 - static entries, IPX SAP 157–159
 - Static Export Source dialog box 302
 - static route entries, defining 117–119
 - static-entry filters
 - configuring 216–220
 - defined 210
 - statistics
 - gathering 371
 - IP error 366
 - IP packet 363
 - IP reassembly 364
 - IPX error 369
 - IPX packet 368
 - port bytes 360
 - port error 361
 - port packet 359
 - stopping
 - clearVISN CoreWatch 18
 - Configuration Expert 30
 - STP

- defining port attributes 60–61
- disabling on ports 63–64
- enabling on ports 61–63
- port information 374–376
- setting up 58–64
- STP Global Settings dialog box 59
- STP Port table 374
- STP settings, defining 59–60
- STP Table command 432
- strict priority 162
- submenus
 - Bridging State 432
 - DIGITAL Web Site 437
 - DVMRP State 434
 - IGMP State 435
 - IP 431
 - IP State 433
 - IPX 431
 - IPX State 433
 - OSPF State 433
 - Performance 431
 - Port 431
 - QoS State 435
 - RIP State 433
 - Routing State 433
 - System State 432
- subtree
 - configuration files 22
 - navigating 23
- switching information
 - obtaining for Layer-2 417–418
 - obtaining for Layer-3/4 418–420
- SYSLOG server 39–40
- System Configuration DNS dialog box 41
- System Dashboard 355
- System Dashboard command 431
- System ID dialog box 32
- System Log dialog box 40
- system performance, monitoring 353–356
- system settings, changing 31–45
- System State submenu 432
- system status, checking 347–351

T

tables

- ... 351
- Chassis Info 347
- controlling contents 426, 439–442
- DVMRP Interface 401
- DVMRP Neighbor 403

- DVMRP Next Hop 406
- DVMRP Routing 405
- exporting data 428
- finding text 425
- Flow 419
- Flow Priority 415
- IGMP Cache 410
- IGMP Interface 408
- IP Forwarding 381
- IP Interface 380
- IPX Forwarding 386
- IPX Interface 383
- L2 Forward 417
- L2 Interface 377
- L2 Priority 413
- obtaining records 427
- OSPF Area 391
- OSPF Area Aggregate 396
- OSPF Interface 387
- OSPF Link State DB 394
- OSPF Neighbor 393
- Port 349
- refreshing 427
- restoring contents 427
- RIP Interface 398
- RIP Peer 400
- saving as reports 421–423
- sorting information 428
- STP Port 374
- Trap 350
- VLAN 373
- working with 425–428
- Tag Export Source dialog box 299
- targets, SNMP trap 42–43
- Technical Support command 437
- text, finding in tables 425
- Tile Horizontally command 436
- Tile Vertically command 436
- Time to Live threshold 133, 135
- toolbars
 - clearVISN CoreWatch 16
 - graph 371
- traffic priority 162
- traffic rules 189
- Trap table 350
- Trap Table command 432
- traps
 - obtaining information 350
- traps, setting targets 42–43
- trunk ports

- defining 71–73
- description 70
- trusted gateways, adding 126–127
- tunnels
 - defining DVMRP 133–135
 - disabling DVMRP 135
 - DVMRP 133–136
 - enabling DVMRP 135

U

- unicast routing
 - description 113–128
 - global parameters 114–115
- Update ACL List dialog box 104, 156

V

- views
 - Front Panel 12–14
 - Schematic 14–15
- Virtual Local Area Network. *See* VLAN
- VLAN Definition dialog box 83, 84
- VLAN table 373
- VLAN Table command 432
- VLAN wizard 74, 78
- VLANs
 - adding ports 77, 81, 85–87
 - bound to IP interfaces 98–103
 - bound to IPX interfaces 149–154
 - configuration tips 70
 - configuring 69–87
 - going down 92, 142
 - modifying 82–84
 - obtaining information 373
 - overview 69
 - port-based
 - creating 78–82
 - description 70
 - modifying 83
 - protocol-based
 - creating 73–78
 - description 70
 - modifying 84
 - removing ports 78, 82, 87
 - replacing 85
- VRP wizard 108
- VRRP Trace Options dialog box 107
- VRRP Trace options, specifying 107
- VRRP, configuring 108–110

W

- weighted-fair queuing 162
- White Papers command 437
- Window menu 436
- Windows 95/98
 - clearVISN CoreWatch requirements 2
 - installing clearVISN CoreWatch 6
 - starting clearVISN CoreWatch 10
 - starting Configuration Expert 21
- Windows NT
 - clearVISN CoreWatch requirements 2
 - installing clearVISN CoreWatch 6
 - starting clearVISN CoreWatch 10
 - starting Configuration Expert 21
- wizards
 - ARP 115
 - icon 25
 - IP Interface 93, 98
 - IP Security 190
 - IPX Interface 143, 149
 - IPX Security 194–206
 - L2 Security 211–223
 - purpose 25
 - QoS 167, 177
 - VLAN 74, 78
 - VRRP 108