Version 4.07

Part No. 314958-A Rev 00
June 2002

600 Technology Park Drive
Billerica, MA 01821-4130

Configuring the Contivity VPN Switch

NORTEL NETWORKS™

Copyright © 2001 Nortel Networks

All rights reserved. June 2002.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel Networks Inc.

The software described in this document is furnished under a license agreement and may be used only in accordance with the terms of that license. The software license agreement is included in this document.

Trademarks

Nortel Networks, the Nortel Networks logo, Contivity, Preside, and Optivity are trademarks of Nortel Networks.

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

AXENT and OmiGuard Defender are trademarks of AXENT Technologies, Inc.

Check Point and Firewall 1 are trademarks of Check Point Software Technologies Ltd.

Cisco and Cisco Systems are trademarks of Cisco Systems, Inc.

Entrust and Entrust Authority are trademarks of Entrust Technologies, Incorporated.

Java is a trademark of Sun Microsystems.

Linux and Linux FreeS/WAN are trademarks of Linus Torvalds.

Macintosh is a trademark of Apple Computer, Inc.

Microsoft, Windows, Windows NT, and MS-DOS are trademarks of Microsoft Corporation.

Netscape, Netscape Communicator, Netscape Navigator, and Netscape Directory Server are trademarks of Netscape Communications Corporation.

NETVIEW is a trademark of International Business Machines Corp (IBM).

NetWARE and Novel intraNetWare are trademarks of Novell, Inc.

NDS is a trademark of Novell Inc.

OPENView is a trademark of Hewlett-Packard Company.

SafeNet/Soft-PK Security Policy Database Editor is a trademark of Information Resource Engineering, Inc.

SecurID and Security Dynamics ACE Server are trademarks of RSA Security Inc.

SPECTRUM is a trademark of Cabletron Systems, Inc.

VeriSign is a trademark of VeriSign, Inc.

All other trademarks and registered trademarks are the property of their respective owners.

Restricted rights legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of conditions

In the interest of improving internal design, operational function, and/or reliability, Nortel Networks Inc. reserves the right to make changes to the products described in this document without notice.

Nortel Networks Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Portions of the code in this software product may be Copyright © 1988, Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms of such portions are permitted, provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that such portions of the software were developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from such portions of the software without specific prior written permission.

SUCH PORTIONS OF THE SOFTWARE ARE PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

In addition, the program and information contained herein are licensed only pursuant to a license agreement that contains restrictions on use and disclosure (that may incorporate by reference certain limitations and notices imposed by third parties).

Nortel Networks Inc. software license agreement

This Software License Agreement ("License Agreement") is between you, the end-user ("Customer") and Nortel Networks Corporation and its subsidiaries and affiliates ("Nortel Networks"). PLEASE READ THE FOLLOWING CAREFULLY. YOU MUST ACCEPT THESE LICENSE TERMS IN ORDER TO DOWNLOAD AND/OR USE THE SOFTWARE. USE OF THE SOFTWARE CONSTITUTES YOUR ACCEPTANCE OF THIS LICENSE AGREEMENT. If you do not accept these terms and conditions, return the Software, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

"Software" is owned or licensed by Nortel Networks, its parent or one of its subsidiaries or affiliates, and is copyrighted and licensed, not sold. Software consists of machine-readable instructions, its components, data, audio-visual content (such as images, text, recordings or pictures) and related licensed materials including all whole or partial copies. Nortel Networks grants you a license to use the Software only in the country where you acquired the Software. You obtain no rights other than those granted to you under this License Agreement. You are responsible for the selection of the Software and for the installation of, use of, and results obtained from the Software.

1. Licensed Use of Software. Nortel Networks grants Customer a nonexclusive license to use a copy of the Software on only one machine at any one time or to the extent of the activation or authorized usage level, whichever is applicable. To the extent Software is furnished for use with designated hardware or Customer furnished equipment ("CFE"), Customer is granted a nonexclusive license to use Software only on such hardware or CFE, as applicable. Software contains trade secrets and Customer agrees to treat Software as confidential information using the same care and discretion Customer uses with its own similar information that it does not wish to disclose, publish or disseminate. Customer will ensure that anyone who uses the Software does so only in compliance with the terms of this Agreement. Customer shall not a) use, copy, modify, transfer or distribute the Software except as expressly authorized; b) reverse assemble, reverse compile, reverse engineer or otherwise translate the Software; c) create derivative works or modifications unless expressly authorized; or d) sublicense, rent or lease the Software. Licensors of intellectual property to Nortel Networks are beneficiaries of this provision. Upon termination or breach of the license by Customer or in the event designated hardware or CFE is no longer in use, Customer will promptly return the Software to Nortel Networks or certify its destruction. Nortel Networks may audit by remote polling or other reasonable means to determine Customer's Software activation or usage levels. If suppliers of third party software included in Software require Nortel Networks to include additional or different terms, Customer agrees to abide by such terms provided by Nortel Networks with respect to such third party software.

2. Warranty. Except as may be otherwise expressly agreed to in writing between Nortel Networks and Customer, Software is provided "AS IS" without any warranties (conditions) of any kind. NORTEL NETWORKS DISCLAIMS ALL WARRANTIES (CONDITIONS) FOR THE SOFTWARE, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OF NON-INFRINGEMENT. Nortel Networks is not obligated to provide support of any kind for the Software. Some jurisdictions do not allow exclusion of implied warranties, and, in such event, the above exclusions may not apply.

3. Limitation of Remedies. IN NO EVENT SHALL NORTEL NETWORKS OR ITS AGENTS OR SUPPLIERS BE LIABLE FOR ANY OF THE FOLLOWING: a) DAMAGES BASED ON ANY THIRD PARTY CLAIM; b) LOSS OF, OR DAMAGE TO, CUSTOMER'S RECORDS, FILES OR DATA; OR c) DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS OR SAVINGS), WHETHER IN CONTRACT, TORT OR OTHERWISE (INCLUDING NEGLIGENCE) ARISING OUT OF YOUR USE OF THE SOFTWARE, EVEN IF NORTEL NETWORKS, ITS AGENTS OR SUPPLIERS HAVE BEEN ADVISED OF THEIR POSSIBILITY. The forgoing limitations of remedies also apply to any developer and/or supplier of the Software. Such developer and/or supplier is an intended beneficiary of this Section. Some jurisdictions do not allow these limitations or exclusions and, in such event, they may not apply.

4.General

a. If Customer is the United States Government, the following paragraph shall apply: All Nortel Networks Software available under this License Agreement is commercial computer software and commercial computer software documentation and, in the event Software is licensed for or on behalf of the United States Government, the respective rights to the software and software documentation are governed by Nortel Networks standard commercial license in accordance with U.S. Federal Regulations at 48 C.F.R. Sections 12.212 (for non-DoD entities) and 48 C.F.R. 227.7202 (for DoD entities).

b. Customer may terminate the license at any time. Nortel Networks may terminate the license if Customer fails to comply with the terms and conditions of this license. In either event, upon termination, Customer must either return the Software to Nortel Networks or certify its destruction.

c. Customer is responsible for payment of any taxes, including personal property taxes, resulting from Customer's use of the Software. Customer agrees to comply with all applicable laws including all applicable export and import laws and regulations.

d. Neither party may bring an action, regardless of form, more than two years after the cause of the action arose.

e. The terms and conditions of this License Agreement form the complete and exclusive agreement between Customer and Nortel Networks.

f. This License Agreement is governed by the laws of the country in which Customer acquires the Software. If the Software is acquired in the United States, then this License Agreement is governed by the laws of the state of New York.

Contents

Preface 23

Before you begin 23
Text conventions 23
Related publications 25
Acronyms 26
Hard-copy technical manuals 27
How to get help 27

Chapter 1
Overview 29

Management 29
Network deployment alternatives 30
Remote access versus extranet access 31
Network configurations 32
Switch behind the router, parallel with firewall 33
Switch with direct connections to the PDN 33
Switch behind a firewall and a router 34
Branch office configuration 35
Virtual private networking 36
IPSec 38
PPTP 38
L2TP 39
L2TP over IPSEC branch office support 39
L2F 40
Authentication services 40
LDAP 41
RADIUS 41
SSL and digital certificates 42
Tunnel certificates 42
Routing 43
Routing policy service (RPS) 44
Routing table 44

About RIP . . . 45
About OSPF . . . 45
About VRRP . . . 45
Private and public interfaces . . . 46
IP addressing . . . 47
Branch office connections . . . 49
Integrated firewall solutions . . . 50
Network Address Translation (NAT) . . . 50
Filters for access control . . . 51
Quality of service . . . 52
Bandwidth management . . . 52
Differentiated Services (DiffServ) . . . 52
Forwarding priority . . . 53
Call Admission priority . . . 53
RSVP . . . 53
Network Time Protocol (NTP) . . . 53
Accounting . . . 54
Local configuration and LDAP file structures . . . 54
Accounting and logging . . . 56
Command line interface . . . 56
Serial point-to-point protocol (PPP) . . . 57
Contivity 1010/1050/1100 features . . . 57
Default configuration . . . 58
DHCP Server . . . 58
PPPoE . . . 59
Branch office quick start utility . . . 59
DNS proxy . . . 60
Compact flash disk . . . 60

Chapter 2
Getting Started . . . 63

Serial Interface configuration . . . 63
Management IP address . . . 67
Safe mode configuration . . . 68
Web browser . . . 68

Preparing for configuration 69
Welcome screen 72
Setting up the Contivity 1010/1050/1100 76
Default configuration parameters 76
Branch office quick start 77
Enterprise environment 78
Service provider environment 79
Branch office quick start template 83
Connecting for internet access 83
Compact flash disk 85

Chapter 3
Configuring the System 87

System identity 87
Resetting Management IP address values 87
LAN interfaces 88
WAN interfaces 95
Hardware encryption accelerator 96
Hardware encryption accelerator and load balancing 98
Hardware encryption accelerator security aspects 98
Configuring Network Time Protocol (NTP) 98

Chapter 4
Configuring Tunnels 101

Configuring user tunnels 106
Configuring branch offices 109
Branch-to-branch with a firewall and a router 110
Branch-to-branch through a gateway switch 111
Sample branch office configuration procedure 113
Configuring the local switch 115
Configuring the remote branch office connection 119
Control tunnels 120
Control tunnel types 121
Restricted mode 123
Nailed-up control tunnels 123

Creating control tunnels 124
Creating a user control tunnel from the serial interface 127
IPSec client features 128
Split tunneling 128
Third-party IPSec clients 130
Forced logout 131
Client failover 132
Static tunnel failover 133
Client auto connect 133
Banner 134
Password storage 134
Client screen saver 135
Domain name 135
Client policy 135
Configuring L2TP over IPSec 135
Switch configuration 135
Windows 2000 configuration 140
Branch office 141
Quality of Service 143
Classifiers 143
Bandwidth management 145
Differentiated Services (DiffServ) 146
Forwarding priority 147
Call admission priority 148
RSVP 150
Proxy ARP and tunnel-to-tunnel traffic 150

Chapter 5
Configuring Routing 153

Routing overview 153
Integrated firewall and routing 155
Running dynamic routing 155
VPN routing 156
Enhanced routing 157
Static routes 158

RIP 159
Using RIP 160
Protecting against routing loops 161
Configuring RIP 161
OSPF 163
Using OSPF 164
Virtual link support 165
ECMP support 166
Configuring OSPF 166
Routing table 170
Route lookup 174
Route selection based on destination 175
Route selection based on precedence in route table 175
Routing policy service (RPS) 176
Using default routes 179
Redistribution of routes 179
Redistribution of default routes 180
Traffic patterns 180
Multicast relay 182
Equal-cost multipath (ECMP) 184
Client address redistribution 185
VRRP 189
VRRP and dynamic routing for high availability 189
VRRP and static tunnels 192
Configuring VRRP 193
Interface groups and critical interface failover 195

Chapter 6
Configuring NAT 197

NAT translation 198
Static address translation 199
Dynamic address translation 200
Dynamic port NAT 200
Dynamic pooled NAT 202
Double NAT 203

NAT usage 204
Branch office NAT 204
Interface NAT 207
NAT Traversal 209

Chapter 7
Configuring Firewalls 211

Contivity Stateful Firewall concepts 211
Stateful inspection 212
Interfaces 213
Filter rules 214
Anti-spoofing 214
Attack detection rules 215
Contivity Stateful Firewall prerequisites 216
Installing Java 2 software 217
Using Internet Explorer 217
Using Netscape 218
Using Netscape 6 220
Using Netscape on Solaris 220
Enabling firewall options 221
Configuring the Contivity Stateful Firewall 222
Selecting event log options 224
Setting up policies 225
Creating and editing firewall policies 226
Navigating rules 226
Implied rules 227
Override rules 227
Interface-specific rules 228
Default rules 230
Post-implied rules 231
Performing actions on rules 231
Header row menu 232
Row menu 232
Cell menus 233
Rule columns 234

Creating policies 241
Adding a policy 241
Deleting an existing policy 242
Copying an existing policy 242
Renaming an existing policy 242
Creating a new policy example 243
Verifying your configuration 244

Chapter 8
Configuring Servers 245

Authentication services 245
LDAP, RADIUS, and the IPSec client 249
LDAP database servers 251
LDAP server authentication 252
RADIUS authentication servers 254
RADIUS-based authentication 255
RADIUS authentication class attribute values 256
RADIUS-Assigned Framed-IP-Address attribute 259
Configuring the switch for RADIUS 259
IPSec and RADIUS 259
PPTP and RADIUS 260
AXENT software tokens and the IPSec client 260
Using a Netscape directory server 261
Quick start 262
Advanced configuration 263
Schema/attribute index cleanup 264
Schema/attribute index modifications 264
SSL 265
Exporting internal LDAP to Netscape directory server 266
RADIUS accounting configuration 268
Remote user IP address pool 269
DHCP relay 270
DHCP server 271
Configuring the switch as a RADIUS server 273

Chapter 9
Using Certificates 275

LDAP Server SSL Encryption 275
Installing LDAP certificates 276
VPN security using digital certificates 276
Public Key Infrastructure (PKI) 276
CA and X.509 certificates 277
Loading certificates on the switch 277
Installing server certificates using cut and paste #7 and #10 277
Installing server certificates using (CMP) 278
Installing trusted CA certificates 279
Server certificate settings 280
Trusted CA certificate settings 281
Group assignment by user identification 282
Allow All 282
Access control by Subject DN 283
Certificate revocation list (CRL) 283
CRL Settings 284
Enabling certificate use for tunnels 286
Identifying individual users with certificates 287
Identifying branch offices with certificates 288
IPSec authentication 288
PPTP authentication and L2TP authentication 289

Chapter 10
Switch Administration 291

Administrators 291
Backups 292
Restoring Configurations 294
Automatic Backup File Servers 295
Tools 295
Recovery 296
Upgrades 297
System configuration 298
File management 298

SNMP 298
SNMP Trap Settings 299
Shutdown 299
Switch status 300
Sessions 302
Reports 303
System 303
Health check 304
Statistics 304
Event log 304
System log 309
Switch data storage 310
RADIUS accounting 311
Accounting data 311
Accounting records 312

Chapter 11
Troubleshooting 313

Troubleshooting tools 314
Client-based tools 314
Switch-based tools 315
Other tools 316
Solving connectivity problems 316
Diagnosing client connectivity problems 317
Common client connectivity problems 318
Problems with name resolution using DNS services 321
Network browsing problems 322
Diagnosing WAN link problems 324
Hardware encryption accelerator connectivity 326
Solving performance problems 326
Eliminating modem errors 326
Hardware encryption accelerator performance 327
Additional information 336
Solving general problems 337
Web browser problems and the Contivity VPN Client Manager 337

Enabling Web browser options 338
Web browser error messages 339
Reporting a problem with a Web browser 341
Switch problems 341
Solving routing problems 343
Client address redistribution problems 343
Solving firewall problems 344

Appendix A
MIB Support 347

SNMP RFC support 347
Novell IPX MIB 347
Novell RIP-SAP MIB 347
RFC 1850 -- OSPF Version 2 Management Information Base 348
RFC 1724 -- RIP Version 2 MIB Extension 348
RFC 1213 -- Network Management of TCP/IP-Based Internets MIB 348
RFC 2667 -- IP Tunnel MIB 348
RFC 2787 -- VRRP MIB 348
RFC 2737 -- Entity MIB 349
RFC 1573 -- IanaIfType MIB 350
RFC 2233 -- If MIB 350
RFC 2571-- Snmp-Framework MIB 350
RFC2790 -- Host Resources MIB 350
RFC2495 -- DS1 MIB 351
CES MIB 351
cestraps.mib -- Nortel Networks proprietary MIB 352
newoak.mib 354
Hardware-related traps 355
Server-related traps 359
Software-related traps 361
Login-related traps 361
Intrusion-related traps 362
System-related traps 362
Information passed with every trap 363

Appendix B

Using Serial PPP 367

Establishing a serial PPP connection 367
Setting up a Dial-Up Networking connection 368
Setting up the modem 369
Setting up the switch 369
Dialing in to the switch 371
Troubleshooting Serial PPP 371
PPP option settings 373

Appendix C
System Messages 375

Certificate messages 375
ISAKMP messages 377
Branch office messages 380
SSL messages 381
Database messages 382
Security messages 383
RADIUS accounting messages 393
RADIUS authentication messages 396
Routing messages 400
Hardware messages 406

Appendix D
Configuring for interoperability 409

Configuring the Cisco 2514 router, version 11.3 409
Configuring the switch for Cisco interoperability 412
Configuring the Check Point VPN-1 and FireWall-1, Version 4.0 413
Establishing an encrypted tunnel on the switch side 414
Establishing an encrypted tunnel on the firewall side 415
Establishing a tunnel from the switch to the firewall 416
Establishing a tunnel to and from the firewall 417
Configuring the firewall's public interface 418
Configuring the switch authentication schemes 419
Configuring the firewall's encryption methods 420
Configuring the switch's public interface 422

Configuring the switch's encryption methods 423
Adding rules to the configurations 425
Editing encryption for rules 426
Configuring the switch for Check Point interoperability 429
Configuring the SafeNet/Soft-PK Security Policy Database Editor, version 1.0s 430
Connecting to IRE SafeNET/Soft-PK Security Policy Client 432
Configuring the switch for IRE interoperability 436
Third-party client installation 436
Considerations for using third-party clients 437
Configuring the switch as a branch office tunnel 439
Configuring the switch as a user tunnel 441
Configuring IPX 443
IPX client 444
Windows 95 and Windows 98 444
Windows NT 444
IPX group configuration 444
Sample IPX VPN switch topology 445
Glossary 447
Index 455

Figures

Figure 1 The Internet and an intranet form an extranet 31
Figure 2 Traditional remote access environment 32
Figure 3 Direct connections to a router 33
Figure 4 Independent WAN connections 34
Figure 5 Switch located behind a firewall 34
Figure 6 Typical branch office environment 36
Figure 7 Remote user to PDN tunnel and on to the switch 37
Figure 8 Private and public interfaces 46
Figure 9 Sample IP addressing scheme 47
Figure 10 Typical branch office environment 49
Figure 11 Switch sample NAT environment 51
Figure 12 Default configuration 58
Figure 13 Welcome screen 73
Figure 14 Navigational menu 75
Figure 15 Branch office quick start screen 82
Figure 16 LAN-to-switch connection 89
Figure 17 Switch-to-PDN configuration 95
Figure 18 Tunnel connection configuration 101
Figure 19 Base, parent, and child group associations 103
Figure 20 Typical branch office environment 110
Figure 21 Branch-to-branch with a firewall and a router 111
Figure 22 Indirectly connected branch offices 112
Figure 23 Sample branch office configuration 113
Figure 24 Setting up a branch office configuration 114
Figure 25 Branch office control tunnel 121
Figure 26 Sample control tunnel environment 122
Figure 27 Sample split tunneling environment 129
Figure 28 Traffic flow 154
Figure 29 VPN routing: allowed traffic patterns 156

Figure 30 Enhanced routing: allowed traffic patterns 157
Figure 31 Router configuration with OSPF and RIP 171
Figure 32 Grouping of route table entries 172
Figure 33 Accept and announce policies 177
Figure 34 Client address redistribution 186
Figure 35 Summarization for client address redistribution 187
Figure 36 Sample HA environment 191
Figure 37 VRRP and static tunnels 193
Figure 38 Sample NAT environment 198
Figure 39 Static address translation 200
Figure 40 Port address translation 201
Figure 41 Pooled address translation 203
Figure 42 Double NAT 204
Figure 43 Overlapping address translation 206
Figure 44 Interface NAT 207
Figure 45 Implied rules 227
Figure 46 Override rules 228
Figure 47 Interface-specific rules 229
Figure 48 Source interface rules 230
Figure 49 Destination interface rules 230
Figure 50 Default rules 231
Figure 51 Post-implied rules 231
Figure 52 Header row menu 232
Figure 53 Row menu 232
Figure 54 Cell menu (option) 233
Figure 55 Cell menu (procedure) 233
Figure 56 Rule column 234
Figure 57 Rule column for interface-specific rules 235
Figure 58 Tunnel Selection dialog box for a user tunnel 235
Figure 59 Tunnel Selection dialog box for a branch office tunnel 236
Figure 60 Network Object Selection dialog box 237
Figure 61 Network object edit dialog box 237
Figure 62 Service Object Selection dialog box 238
Figure 63 tcp object insert dialog box 239
Figure 64 Action column options 239

Figure 65 Log column options 240
Figure 66 Status column options 240
Figure 67 Firewall - Select Policy screen 241
Figure 68 Contivity VPN Switch and authentication servers 246
Figure 69 Authentication server validation flowchart 248
Figure 70 IPSec client authentication options 249
Figure 71 RADIUS authentication class attribute values 257
Figure 72 Nortel Networks logging scheme 302
Figure 73 Switch and Cisco 2514 network topology 410
Figure 74 Switch and Check Point FireWall-1 network 413
Figure 75 Switch and IRE SafeNet network topology 431
Figure 76 Split tunneling example 441
Figure 77 IPX topology 446

Tables

Table 1 Firewall filter pass-through ports . . . 35
Table 2 Sample IP addressing associations . . . 47
Table 3 Web interface configuration options . . . 70
Table 4 Configuration checklist . . . 70
Table 5 Subnet assignments . . . 79
Table 6 BOQS parameters . . . 81
Table 7 Mapping minimum data protection levels to encryption levels . . . 137
Table 8 User group properties used by IPSec transport connections . . . 139
Table 9 Bandwidth allocation per priority level . . . 148
Table 10 Call admission priority . . . 149
Table 11 Maximum connections per priority . . . 149
Table 12 RIP Statistics screen . . . 162
Table 13 RIP Database screen . . . 162
Table 14 RIP Interfaces screen . . . 163
Table 15 LSDB screen . . . 167
Table 16 OSPF Dynamic Neighbors screen . . . 168
Table 17 OSPF Interfaces screen . . . 168
Table 18 OSPF Summary screen . . . 169
Table 19 OSPF Statistics screen . . . 169
Table 20 IP Forward Table Screen . . . 173
Table 21 IP Routing Table screen . . . 173
Table 22 Redistribution rules . . . 179
Table 23 Traffic patterns . . . 181
Table 24 Multicast interface-specific rules example . . . 182
Table 25 Multicast Statistics screen . . . 183
Table 26 Multicast Interfaces screen . . . 184
Table 27 Client address redistribution statistics . . . 188
Table 28 Class attributes . . . 257
Table 29 RADIUS example details . . . 258

Table 30 Field IDs for data collection records 311
Table 31 Troubleshooting tools 316
Table 32 Trap categories 364
Table 33 DIP switches configuration 369

Preface

This guide describes the Nortel Networks* Contivity* VPN Switch. It also provides complete details to help you configure, monitor, and troubleshoot your Contivity VPN Switch.

Before you begin

This guide is for network managers who are responsible for setting up and configuring the Contivity VPN Switch. This guide assumes that you have the following background:

- Experience with windowing systems or graphical user interfaces (GUIs)
- Familiarity with the network management

Text conventions

This guide uses the following text conventions:

angle brackets (< >)	Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command. Example: If the command syntax is **ping** <*ip_address*>, you enter **ping 192.32.10.12**
bold Courier text	Indicates command names and options and text that you need to enter. Example: Use the **show health** command. Example: Enter **terminal paging** {**off** | **on**}.

braces ({ })	Indicate required elements in syntax descriptions where there is more than one option. You must choose only one of the options. Do not type the braces when entering the command. Example: If the command syntax is **`ldap-server source {external | internal}`**, you must enter either **`ldap-server source external`** or **`ldap-server source internal`**, but not both.
brackets ([ ])	Indicate optional elements in syntax descriptions. Do not type the brackets when entering the command. Example: If the command syntax is **`show ntp [associations]`**, you can enter either **`show ntp`** or **`show ntp associations`**. Example: If the command syntax is **`default rsvp [token-bucket {depth | rate}]`**, you can enter **`default rsvp`**, **`default rsvp token-bucket depth`**, or **`default rsvp token-bucket rate`**.
ellipsis points (. . .)	Indicate that you repeat the last element of the command as needed. Example: If the command syntax is **`more disk`***`n`***`:`***`<directory>`***`/...`***`<file_name>`**, you enter **`more`** and the fully qualified name of the file.
italic text	Indicates new terms, book titles, and variables in command syntax descriptions. Where a variable is two or more words, the words are connected by an underscore. Example: If the command syntax is **`ping`** *`<ip_address>`*, *`ip_address`* is one variable and you substitute one value for it.
`plain Courier text`	Indicates system output, for example, prompts and system messages. Example: `File not found.`

separator (>)	Shows menu paths. Example: Choose Status > Health Check.
vertical line (|)	Separates choices for command keywords and arguments. Enter only one of the choices. Do not type the vertical line when entering the command. Example: If the command syntax is **`terminal paging`** {**`off`** | **`on`**}, you enter either **`terminal paging off`** or **`terminal paging on`**, but not both.

Related publications

For more information about the Contivity VPN Switch, refer to the following publications:

- *Contivity 1010/1050/1100 Release Notes (V4.07)* (314963-A) provide the latest information and known problems and workarounds.
- *New Features for the Contivity 1010/1050/1100* (part number 315535-A) describes the new features for the Contivity 1010/1050/1100 series of switches.
- *Installing the Contivity 1010/1050/1100* (part number 314961-A) provides instructions on how to install the Contivity 1010, 1050, and 1100 and includes technical specifications.
- *Reference for the Contivity VPN Switch Command Line Interface* (part number 311645-C) describes the commands that you can use from the command line interface.
- *Connecting for Internet Access* (part number 314962-A) describes how to set up your Contivity 1010, 1050, or 1100 at the branch office site.

Acronyms

This guide uses the following acronyms:

ACK	acknowledgement
CA	certificate authority
CHAP	Challenge Handshake Authentication protocol
CRL	certificate revocation list
DN	distinguished name
DNS	domain name system
EAC	Extranet Access Client
FIPS	Federal Information Processing Standards
FTP	File Transfer Protocol
IP	Internet Protocol
IKE	IPSec Key Exchange
ISAKMP	Internet Security Association and Key Management Protocol
ISP	Internet service provider
L2TP	Layer2 Tunneling Protocol
LDAP	Lightweight Directory Access Protocol
LAN	local area network
MAC	media access control address
NAT	network address translation
NOC	network operations center
NTP	Network Time Protocol
OSPF	Open Shortest Path First
OSS	operations support systems
PAP	Password Authentication Protocol
PDN	public data networks
POP	point-of-presence

PPP	Point-to-Point Protocol
PPTP	Point-to-Point Tunneling Protocol
RSVP	Resource Reservation Protocol
RIP	Routing Information Protocol
SNMP	Simple Network Management Protocol
UDP	User Datagram Protocol
URL	uniform resource locator
VPN	virtual private network
VRRP	Virtual Router Redundancy Protocol
WAN	wide area network

Hard-copy technical manuals

You can print selected technical manuals and release notes free, directly from the Internet. Go to the www.nortelnetworks.com/documentation URL. Find the product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Use Adobe* Acrobat Reader* to open the manuals and release notes, search for the sections you need, and print them on most standard printers. Go to Adobe Systems at the www.adobe.com URL to download a free copy of the Adobe Acrobat Reader.

You can purchase selected documentation sets, CDs, and technical publications through the Internet at the www1.fatbrain.com/documentation/nortel/ URL.

How to get help

If you purchased a service contract for your Nortel Networks product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Nortel Networks service program, contact one of the following Nortel Networks Technical Solutions Centers:

Technical Solutions Center	Telephone
Europe, Middle East, and Africa	(33) (4) 92-966-968
North America	(800) 4NORTEL or (800) 466-7835
Asia Pacific	(61) (2) 9927-8800
China	(800) 810-5000

Additional information about the Nortel Networks Technical Solutions Centers is available from the www.nortelnetworks.com/help/contact/global URL.

An Express Routing Code (ERC) is available for many Nortel Networks products and services. When you use an ERC, your call is routed to a technical support person who specializes in supporting that product or service. To locate an ERC for your product or service, go to the http://www130.nortelnetworks.com/cgi-bin/eserv/common/essContactUs.jsp URL.

Chapter 1
Overview

This chapter introduces the Nortel Networks Contivity VPN Switch. The switch provides scalable, secure, and manageable remote access using the Internet.

Management

Management consists of initial and ongoing configuration, status and monitoring, and proactive notification of system events that indicate problems. The switch uses standard HTML Web browsers, Java* scripts, and Java applets.

Since configuration and monitoring tasks are likely to be performed by different individuals, you can assign separate administrators with view only or view and configure access rights.

The switch has a Quick Start configuration option that allows you to set up a default configuration with a single management screen in a few minutes. This allows you to quickly set up a couple of tunnels and verify their operation. The next stage of configuration is the Guided Configuration. This option leads you through an entire configuration for the switch with section introductions and online help. Once you become familiar with the navigational menu and submenu options, you can go directly to the Manage Extranet option and browse all of the management screens. You can always access context-sensitive online help by clicking on the Help button.

System monitoring provides feedback regarding the status of the system, such as the number of active sessions and the system resources in use.

The switch supports SNMP MIB II using Gets (status reports). SNMP traps allow you to react to events that need attention or that might lead to problems. Through the community name, you can define the management stations that receive the SNMP traps. The switch supports all of the SNMP management stations, including OPENVIEW*, NetView* 6000, SPECTRUM*, and Net Manager.

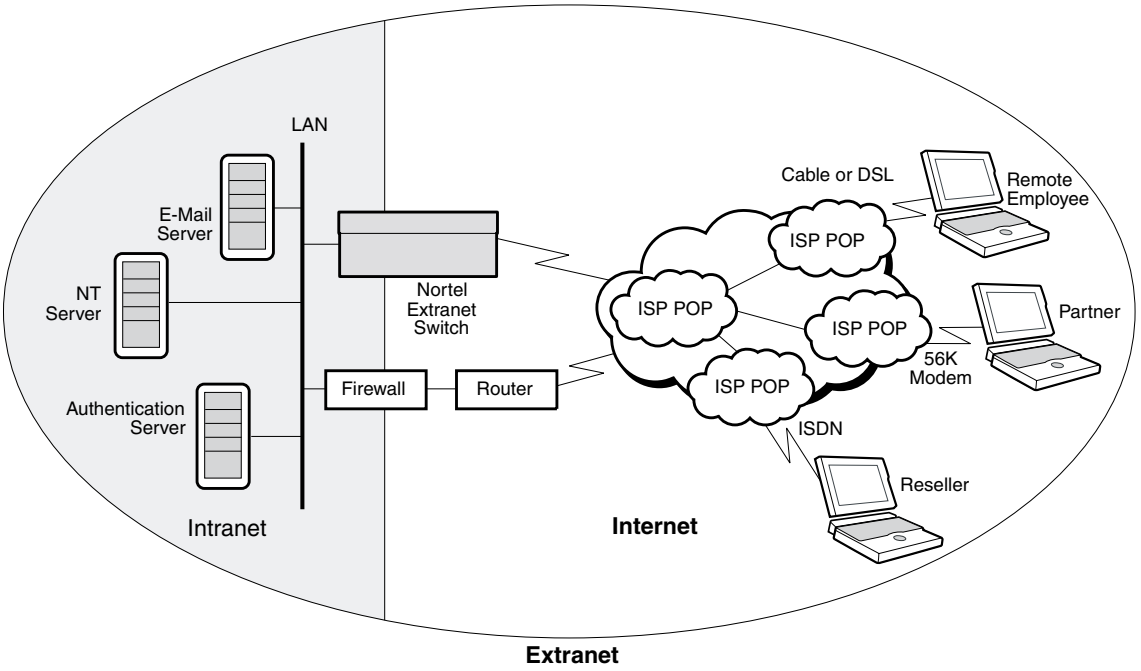
The switch allows the scripting of SNMP alerts so that a combination of system variables can signal an SNMP trap. When a trap occurs, the Nortel Networks icon on a standard management station typically turns red, signaling that there is a problem. You can double-click on the icon to open the browser to the Nortel Networks management interface.

Network deployment alternatives

With its combination of secure, manageable, and scalable features, the switch offers significant extranet access savings. You can shift information technology resources from solving the current remote user access problems to other, more proactive administrative and management areas. And you can eliminate modem-management pool problems from your organization and shift them to your extranet provider.

Extranet access allows remote users to dial in to an Internet Service Provider (ISP) anywhere in the world and reach corporate headquarters or branch offices using the ISP. The extranet provides remote users access to corporate databases, mail servers, and file servers. Figure 1 shows a typical extranet environment.

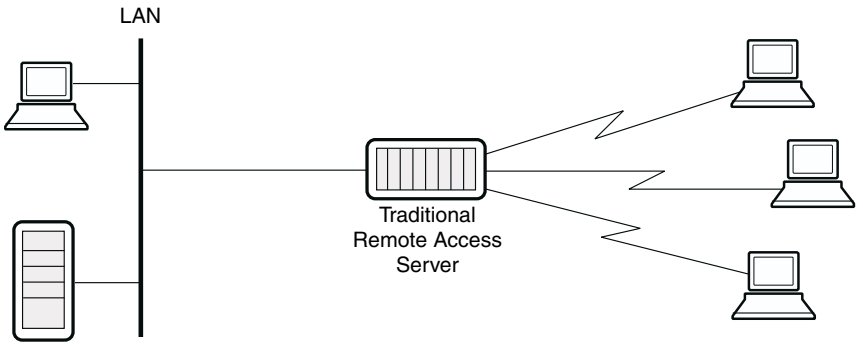
Figure 1 The Internet and an intranet form an extranet



Remote access versus extranet access

Traditional corporate remote access environments employ banks of modems to handle incoming service requests. The switch allows Internet Service Providers (ISPs) to take over the role of point-of-presence (POP) providers of modem access. This frees corporate personnel to perform other proactive duties. The switch improves performance while lowering overhead. And, the resource shift translates to significant corporate savings. Figure 2 shows a remote access environment.

Figure 2 Traditional remote access environment



You can also save by reducing the number of telephone company (Telco) T1 lines needed. Traditional T1 (channelized) lines support only 24 connections per line; a T1 connected to a switch aggregates many sessions through statistical multiplexing and can service between 100 and 200 extranet users, depending upon their usage.

Additionally, there are savings from the reduced Telco expenses for the long-distance calls from employee homes to corporate modem pool access points, as well as replacement of expensive private leased lines that link remote offices. Extranet users save by dialing into the local POP, which is usually a free call.

Network configurations

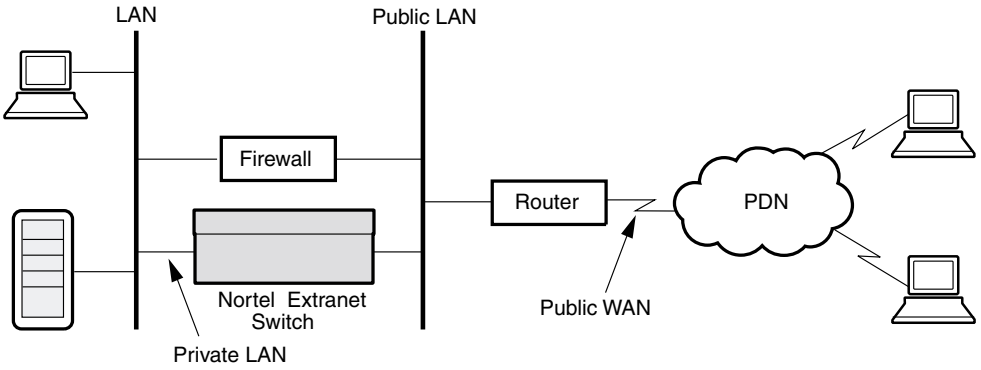
The switch can be integrated into your network in several configurations, depending on your network topology and method of Internet access. Four common topologies are described in this section:

- Switch behind the router, in parallel with a firewall
- Switch with direct connections to the PDN
- Switch behind a firewall and a router
- Branch office configuration

Switch behind the router, parallel with firewall

This configuration shows the switch connected behind a router and in parallel with a firewall. This configuration allows you to retain your preexisting configuration while still allowing secure extranet access through the switch. One of the major benefits of this configuration is that it shares an existing connection to the PDN. Figure 3 shows a firewall and the switch with connections to a router.

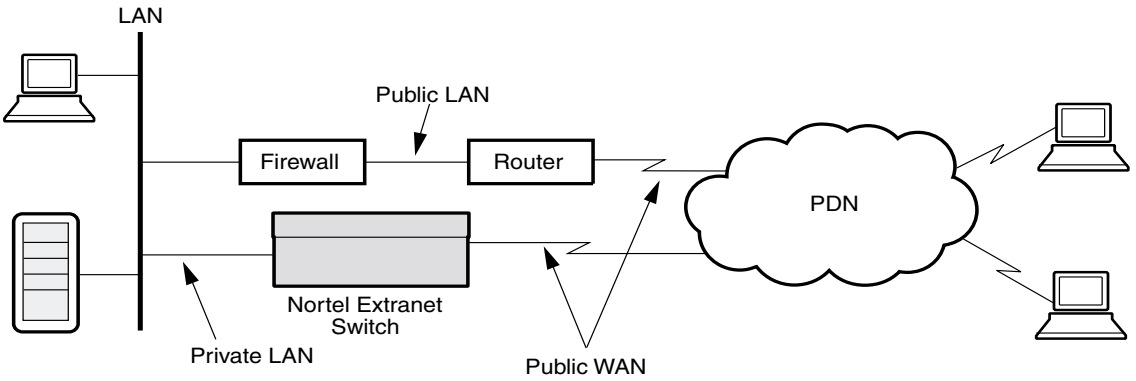
Figure 3 Direct connections to a router



Switch with direct connections to the PDN

In this configuration, the firewall and router are completely independent of the switch. Their connections and configuration remain the same. The switch is installed like a traditional remote access server except that the wide area network (WAN) connection is to the PDN rather than to a dial-up telephone network. This allows extranet bandwidth to grow without impacting other Internet traffic. Figure 4 shows the switch with independent, direct connections to a WAN.

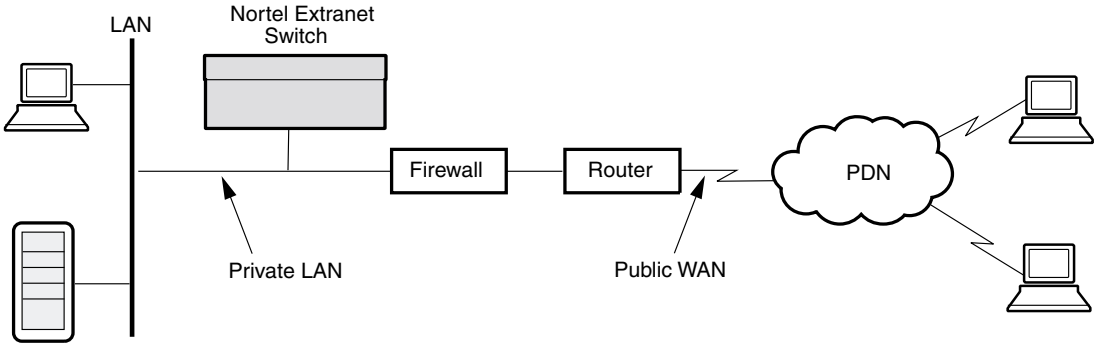
Figure 4 Independent WAN connections



Switch behind a firewall and a router

This configuration shows the switch connected behind both a firewall and a router. For this configuration to work, the router and firewall must be configured to allow tunnels to pass through to the switch. This is usually a simple matter that in most cases requires changes only to the firewall. Figure 5 shows the switch behind a firewall and a router.

Figure 5 Switch located behind a firewall



In the preceding configuration, you have to modify your firewall filters to pass PPTP, L2F, and other tunneled traffic through to the switch. Refer to Table 1 for details.

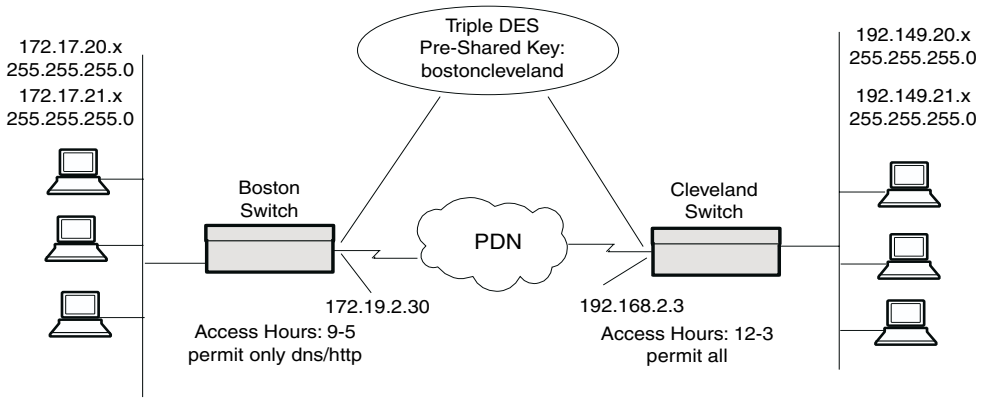
Table 1 Firewall filter pass-through ports

Service	Protocol number	Source port	Destination port
PPTP Control Connection	6 (TCP)	>1023	1723
PPTP Tunnel Encapsulation	47 (GRE)	N/A	N/A
ISAKMP/**IPSec** Key Management	17 (UDP)	500	500
IPSec Tunnel Encapsulation	50 (ESP) 51 (AH)	N/A	N/A
L2TP/L2F	17 (UDP)	>1023	1701
Reservation Control	46 (RSVP)	N/A	N/A

Branch office configuration

The branch office feature allows you to configure a tunnel connection between two private networks. Typically, one private network is behind a locally configured switch while the other is behind a remote switch. Branch office configuration allows you to configure the accessible subnetworks behind each switch. The configuration also contains the information that is necessary to set up the connection, such as the switch's IP addresses, encryption types and authentication methods. Local policy restrictions such as access hours, filter sets, and call admission priorities can be applied to limit connectivity into local subnetworks. Figure 6 shows a typical branch office environment.

Figure 6 Typical branch office environment



Virtual private networking

The switch uses the Internet and remote connectivity to create secure extranets. Remote connectivity through the PDN requires a protocol for safe transport and a connection from the remote user's PC to the PDN. The switch includes the most popular tunneling protocols, IP Security (IPSec), Point-to-Point Tunneling Protocol (PPTP), Layer 2 Forwarding Tunneling Protocol (L2TP), and Layer 2 Forwarding (L2F). IPSec uses digital certificates, password-based keys, and tokens for authentication; PPTP, L2TP, and L2F use Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP) for authentication. The PPTP implementation for the switch supports MS-CHAP authentication with 56- to 128-bit key encryption and MS PPP CHAP (MS-CHAP_V2) support.

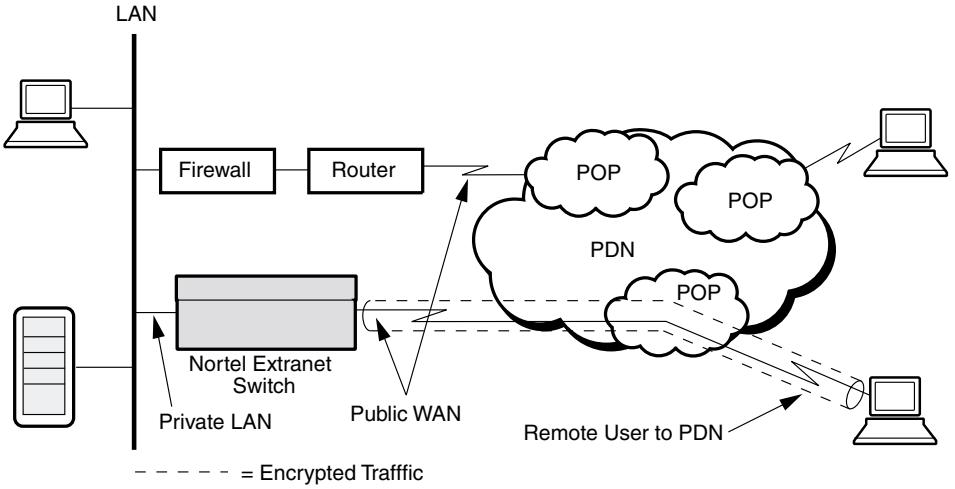
The switch provides more security than traditional remote access schemes by using a combination of authorization, authentication, privacy, and access control for each user. In addition, support for the IPSec protocol and related Internet Security Association and Key Management Protocol (ISAKMP) and the Oakley Key Establishment Protocol further enhances security.

For tunnel creation, the user establishes a connection with the PDN point-of-presence (POP), typically through an Internet Service Provider (ISP). After the Internet connection is up, the remote user launches a second dial-up connection, and this specifies a connection to a switch. Instead of a telephone number to establish the link, the second connection uses an IP address (or a name if the IP address has been entered into a Domain Name Service server).

This second connection could use either the Point-to-Point Tunneling Protocol (PPTP) or the IP Security (IPSec) tunneling protocol. Tunnels built using L2F are slightly different. The tunnel begins at a piece of networking equipment (network access server or NAS) located at the ISP instead of the remote user's PC. The user simply dials into the ISP with a telephone number that causes an L2TP or L2F tunnel to connect directly to a specific corporation. This is similar to a traditional remote-dial service except that the modems are maintained by the ISP and not the corporation.

Figure 7 shows the connection between the remote user and the ISP at the PDN and on to the switch.

Figure 7 Remote user to PDN tunnel and on to the switch



IPSec

The IPSec tunneling protocol is supported by Nortel Networks and other third-party vendors. IPSec is an emerging standard that offers a strong level of encryption. It offers DES (Data Encryption Standard), Triple DES and AES (Advanced Encryption Standard), integrity protection (MD5 and SHA), and the IETF-recommended ISAKMP and Oakley Key Determination protocols, and token codes from SecurID* and AXENT*. IPSec offers the following features:

- Client support is available from Nortel Networks and other vendors. No special ISP services are required.
- Support for IP address translation via encapsulation, packet-by-packet authentication.
- Strong encryption and token codes.

Nortel Networks provides the IPSec remote access user client software on the CD that came up with your switch. You can put the client software on a network server for your remote users to download. The client software is a Microsoft* application available for the latest releases of Windows* 95, Windows 98, Windows NT*, Windows 2000 Workstation, and Windows NT Server. The software comes with complete online help.

Nortel Networks provides two versions of the IPSec client due to export restrictions. The standard version supports DES (56-bit key) encryption, and the enhanced version supports Triple DES (3DES, 168-bit key). The self-extracting installation files for DES and Triple DES are labeled accordingly on the CD. The installation is simple; the self-extracting installation includes everything necessary to create IPSec tunnels with the switch.

PPTP

The Point-to-Point Tunneling Protocol is supported by Nortel Networks, Microsoft, and several other vendors. The Microsoft PPTP client is available for Windows 95, Windows NT Workstation (Version 4.0), Windows ME, and Windows NT Server (except Version 3.51). The Microsoft PPTP client is bundled with Windows 98 software. Network TeleSystems (www.nts.com) provides tunneling product support for Windows 3.1 and Macintosh* operating systems. You can obtain the PPTP client upgrade for Windows 95 directly from Microsoft (www.microsoft.com). Installation instructions are also available from this site.

The PPTP client is on the Contivity VPN Switch CD and is built into the Windows NT operating system. PPTP offers the following features:

- Connections can be made from a range of clients without requiring special ISP services.
- The PPTP client is available for the most common client operating systems.
- PPTP supports IP address translation using encapsulation, support for IPX tunneling, and RC4 encryption (either 56- or 128-bit, within the limits of United States export law).

L2TP

Layer 2 Tunneling Protocol is supported by Nortel Networks, Cisco Systems*, Microsoft, and other vendors. L2TP combines the best features of the L2F and PPTP tunneling protocols. L2TP tunneling allows secure remote access to corporate networks across the public Internet. L2TP tunnels are generally established between a network access server (NAS) at the ISP and the switch. L2TP allows you to specify MS-CHAP, CHAP, or PAP authentication, enable compression, and assign DNS and WINS servers to the tunnel.

You can use support for IPSec transport mode to provide security for L2TP traffic. You can use IPSec transport-protected L2TP tunneling for both remote access traffic and branch office tunnel traffic. Windows 2000 can act as a peer in a branch office connection using L2TP/IPSec or IPSec tunnel mode. Also, Windows 2000 can act as a client using L2TP/IPSec. Authentication for L2TP/IPSec tunnels can be either shared secret or digital certificate. It also provides configuration support for both voluntary and compulsory L2TP/IPSec remote access connections. (Windows 2000 authentication must be digital certificate.)

L2TP over IPSEC branch office support

To support the termination of Microsoft Windows 2000 L2TP/IPSec connections for client-to-gateway connections and gateway-to-gateway connections, the switch supports IPSec transport mode and to provide security for L2TP traffic.

L2F

The L2F (Layer 2 Forwarding) tunneling protocol is supported by Nortel Networks, Cisco Systems, Shiva*, and other vendors. L2F tunneling allows remote access to corporate networks across the public Internet. L2F tunnels are generally established between the network access server at the ISP and the switch.

There is no direct client software required for L2F beyond the PPP dialer software, such as the dial-up networking utility provided with Windows 95 and Windows 98. L2F tunnels are actually made from the ISP to the corporate switch on behalf of the user. These connections depend on the domain associated with the dial-in username. Therefore, ISPs must offer services that are based on L2F; currently, L2F is available on a limited basis. L2F provides IP address translation using encapsulation and support for IPX tunneling, but it does not perform encryption. L2F offers the following features:

- Requires special ISP services.
- No requirement for special software on the client.
- No data encryption.

Authentication services

For authentication and access control, the switch supports an internal or external Lightweight Directory Access Protocol (LDAP) server and external Remote Authentication Dial-In User Service (RADIUS) servers. External LDAP proxy server support allows authentication of users against existing LDAP databases.

The remote user attempting to dial in to the switch must be authenticated before gaining access to the corporate network. Authentication is one of the most important functions that the switch provides because it identifies users and drives many other aspects of the user-centric functionality.

Because extranet access tends to be used initially with an existing remote access system, the switch is designed to take advantage of the same authentication infrastructure that is used for traditional remote access. An external database can be used for centralized storage of user names and passwords, thereby easing administration. This database can be accessed via an intermediary protocol, such as LDAP or RADIUS.

The switch supports several authentication services and augments them with a group profile mechanism. When a user attempts to get into the network, the switch references a particular group profile to determine encryption strength, filtering profile, quality of service attributes, and more for that user.

This user- and group-specific profile allows users to be grouped together for common attributes, while preserving the flexibility to make exceptions for individual users. The product features and network access that apply to a user can be controlled by the user identity, rather than the source IP address or another mechanism. This is necessary to support mobile users and users coming from other organizations, and to create personal extranets.

Following are details of the authentication services currently supported by the switch, LDAP, or RADIUS.

LDAP

The Lightweight Directory Access Protocol (LDAP) emerged from the X.500 directory service. LDAP is gaining acceptance as the directory model for the Internet. Microsoft, Netscape*, and Novell* all support LDAP in their directory service strategies. LDAP is based on directory entries; it has an Internet person schema that defines standard attributes and can be extended to include other attributes. A directory service is a central repository of user information; for example, the switch supports the following elements using LDAP:

- Groups
- Users
- Filters
- Services

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a distributed security system that uses an authentication server to verify dial-up connection attributes and authenticate connections. RADIUS is commonly used for remote access authentication.

Many security systems can be configured with a RADIUS front end to facilitate remote access authentication. RADIUS is also the most common authentication mechanism used by ISPs. Novell NDS*, Microsoft Windows NT Domains, Security Dynamics ACE Server*, and AXENT OmniGuard Defender*, among others, all support RADIUS authentication. Windows NT Domain authentication is used to control access to NT file servers and other resources on NT networks. This is a convenient place for user passwords to be stored, as users generally remember their file server passwords. Soon, many organizations will use X.509 digital certificates as an authentication mechanism. These certificates work hand-in-hand with public key encryption to identify users to provide a level of assurance that users are who they say they are. Eventually, this type of authentication will be the most common.

SSL and digital certificates

The Secure Socket Layer protocol (SSL) can use digital certificates when establishing secure, authenticated connections between SSL clients and servers.

The switch uses a digital certificate sent from an SSL-capable LDAP server to authenticate that server. In order for digital certificate authentication to succeed, a certificate from the authority certifying the LDAP server must be imported into the switch's certificate store. This type of certificate is often referred to as a Certification Authority (CA) root certificate.

A single CA root certificate can be used to certify the authenticity of multiple LDAP servers, depending on the organization of your environment's certification hierarchy.

Tunnel certificates

The Contivity VPN Switch supports retrieval of X.509v3 certificates from Microsoft Certificate storage through the Microsoft CryptoAPI (MS CAPI). Microsoft Certificate storage also provides a mechanism to import digital certificates granted by third-party certification authorities through the use of standard messages (PKCS #12). This allows the Contivity VPN Switch and Contivity VPN Client to make use of CAs that have not been tightly integrated with the client and switch.

You can use X.509 certificates to authenticate IPSec-based tunnel connections. The switch supports RSA digital signature authentication in the IPSec ISAKMP key management protocol. Remote users can authenticate themselves to the switch using a public key pair and a certificate as credentials. In addition, the switch uses its own key pair and certificate to authenticate the switch to the user. The switch currently supports the Entrust* product suite and VeriSign* certificates.

Using certificates for tunnel connections requires the creation of a *public key infrastructure* (PKI) to issue and manage certificates for remote users and switch servers. The switch software supports the Entrust PKI, Web Connector, VPN Connector, and Entrust client software PKI components. The Entrust software issues certificates for users and the switches.

Routing

The switch's routing capabilities include allowing authorized tunneled traffic to securely flow in to and out from the corporation's private network. In addition, the switch can route traffic between two private interfaces, and between its public and private interfaces. The switch can also route traffic from its public interfaces to destinations on the Internet. As a result, you can use the switch to connect your organization to the Internet.

Once you permit traffic to flow between the public and private sides of your network, you also want to ensure that your private network is protected from unauthorized access from the public side. The switch provides a choice of two possible firewall solutions: the Contivity Stateful Firewall and Contivity Interface Filter.

With the addition of an integrated firewall, the switch can perform a variety of secure routing functions, depending upon how you set up the switch routing capabilities. For example, you can configure the switch to securely route non-tunneled traffic from its private interface, through the firewall, and out its public interface. This configuration would enable users on the private network to access the Internet without requiring a separate, dedicated router.

The integrated Contivity Tunnel Filter provides protection, using features such as packet filtering and antispoofing. The Contivity Tunnel Filter provides:

- VPN routing routes traffic to and from secure tunnels.
- Enhanced routing routes traffic between physical interfaces. This includes traffic between public and private interfaces. Enhanced routing also allows traffic to flow between a tunnel and a public interface.
- Services routing routes traffic used for the services that the switch provides. This type of routing supports tunnel protocols such as IPSec, PPTP, L2TP, and L2F. It also supports the use of HTTP and FTP protocols, which are used to manage the switch.

For OSPF to function, you must install the Advanced Routing key. The Firewall Key is optional and required only when you want to enable the redistribution capabilities of RIP and OSPF. If you configure any OSPF policy without installing the applicable key, the configured policy will not be enforced.

Routing policy service (RPS)

The IP router allows you to control the flow of routing data to and from the routing tables. The routing policy service controls this by providing IP accept and announce policies. You can enable or disable an individual route policy from the user interface.

Every IP router maintains a table of current routing information. The routing table manager receives routing updates from the network through the Internet protocols running on the router. Periodically, the routing table manager issues routing updates through the protocols.

Routing table

Like any router, the switch has a routing table that defines how traffic that comes into the switch is routed on to its destination. The routing table can contain both static and dynamic routes. Static routes are manually configured routes that do not change. Dynamic routes, however, do change, as they are learned by using the Routing Internet Protocol (RIP) or Open Shortest Path First (OSPF) from a private interface or a branch office tunnel (the switch does *not* support RIP or OSPF for public interfaces). You can show all routes or the best routes to a destination.

About RIP

The Routing Information Protocol (RIP) is a distance-vector routing protocol that allows routers to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring subnets and listen for RIP updates from the routers on those neighboring subnets. Routers use the information in the RIP updates to keep their internal routes current.

About OSPF

OSPF is a link-state routing protocol that maintains a database from which a routing table is constructed from the shortest path, using a minimum of routing protocol traffic. It provides a high functionality open protocol that allows multiple vendor networks to communicate using the TCP/IP protocol family. Some of the benefits of OSPF are:

- Fast convergence
- Variable Length Subnet Masks (VLSM)
- Hierarchical segmentation
- Area routing to provide additional routing protection and a reduction in routing protocol traffic
- Authentication

About VRRP

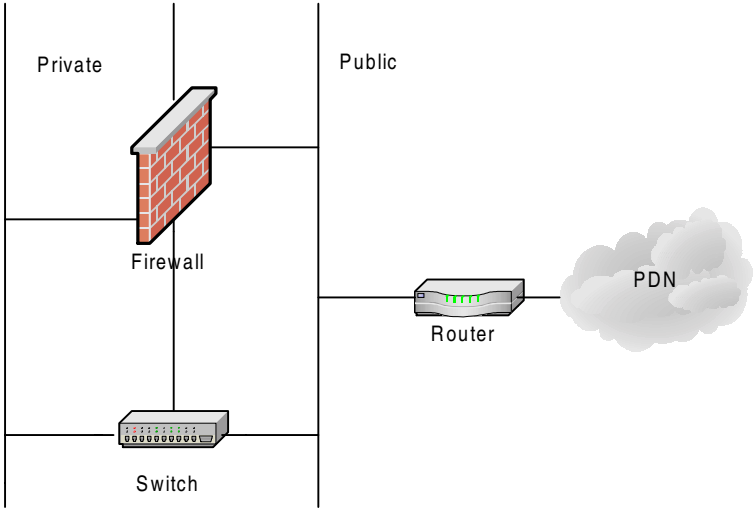
Virtual Router Redundancy Protocol (VRRP) is one method you can use to configure the switch to maintain a state of high availability. VRRP is a standard protocol that handles private interface failures. VRRP targets hosts that are configured with static next-hop routing addresses or default gateways. It provides a means of rerouting traffic in the event of a system/interface failure.

Private and public interfaces

The switch provides secure access between your local area network (LAN)-a private interface-and PDNs such as the Internet. Private refers to the LAN within your corporation, and public refers to public data networks. When a firewall is not enabled, the public interface accepts only tunneled protocols, while the private interface accepts both regular (nontunneled) and tunneled protocols. You must be careful to correctly configure each interface of the switch for proper network security.

The LAN interface on the system board is configured to be private by default. Nortel Networks recommends that its interface be connected to your corporate LAN. This interface accepts regular networking protocols such as TCP/IP, FTP, and HTTP. The private interface also accepts tunneled protocols (such as IPSec, PPTP, L2TP, and L2F) that can be used for secure management access to the switch. Figure 8 shows private and public interfaces.

Figure 8 Private and public interfaces



Additional interfaces that are inserted into the expansion slots are set to public by default. When a firewall is not enabled, the switch rejects nontunneled protocols and accepts only tunneled protocols like IPSec, PPTP, L2TP, and L2F.

The ability to ping (a program used to send ICMP echo request packets) the public interface IP address is also supported for diagnostic purposes. Due to security considerations, a public network host is given a 60-packet allotment to establish a connection. If the connection is not successfully established within this allotment, the public network host must wait approximately 60 seconds before reattempting a tunnel connection. Normally, this allotment is adequate to establish a connection. Be aware that continuous pings inhibit your ability to establish a connection.

IP addressing

Figure 9 shows sample IP address assignments in a network using a switch. Refer to Table 2 to see the IP address associations.

Figure 9 Sample IP addressing scheme

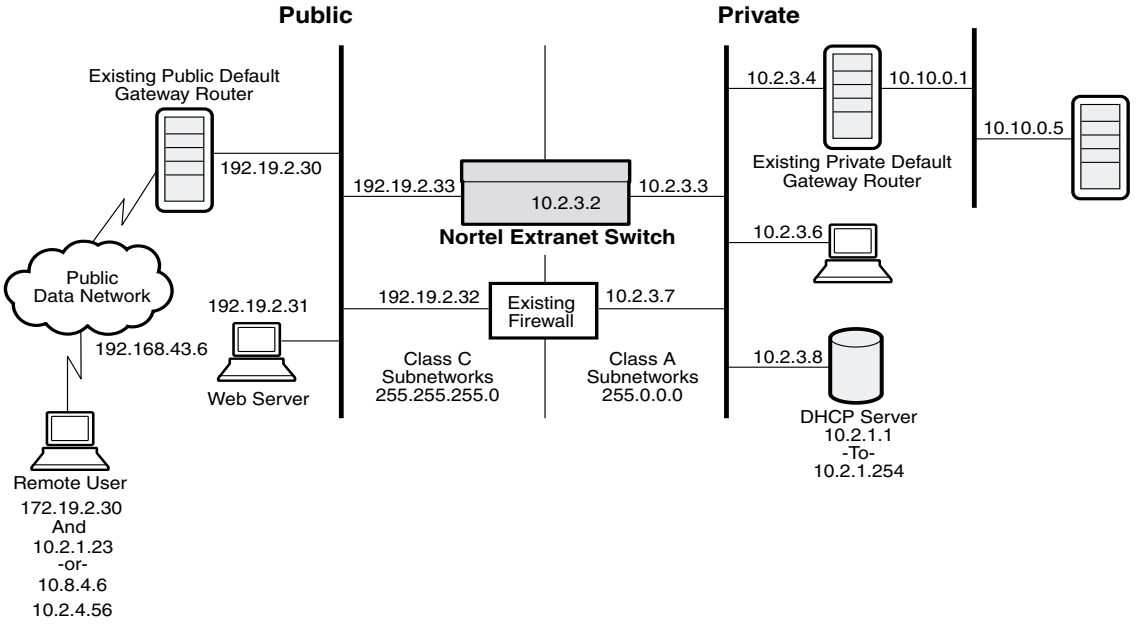


Table 2 Sample IP addressing associations

IP address	Description (when applicable, where configured)
192.168.43.6	Dial-up networking to ISP (Internet access, ISP assigned)

Table 2 Sample IP addressing associations

192.19.2.30	Public default Internet gateway router
192.19.2.33	Public LAN port IP address (remote user destination address)
192.19.2.32	Firewall public network address
10.2.3.2	Switch management IP address: System→Identity
10.2.3.3	Switch private LAN interface IP address: System→LAN Edit IP address
10.2.3.4	Private network default gateway router: System→Routing Add/Edit Default Route
10.2.3.6	Sample partners FTP server for inventory and price list
10.2.3.7	Firewall private network address
10.2.3.8	DHCP server IP address
10.2.1.1 to 10.2.1.254	Private Network Addresses Assigned to Remote Tunnel Sessions: DHCP pool: Servers→User IP Addr
172.19.2.30	ISP-assigned address
10.2.1.23	DHCP-assigned IP address for a remote user
10.8.4.6	Sample remote user static IP address: Profiles→Users Edit
10.2.4.56	Sample client-specified address: Profiles→Groups Edit IPSec/PPTP/ L2TP/L2F

The switch supports the Internetwork Packet Exchange (IPX) protocol. This allows the switch to transmit and receive IPX packets over PPTP, L2TP, and L2F tunnels, but not over IPSec.



Note: L2TP, PPTP, and L2F carry both IP and IPX traffic only for remote access connections. IPX is not supported in branch office tunnels.

The switch supports IPX by encapsulating IPX traffic within IP tunnels over PPTP, L2TP, and L2F. The private interfaces and public interfaces can carry IP and IPX traffic simultaneously. The IP addresses are not shown in the preceding illustration.

Branch office connections

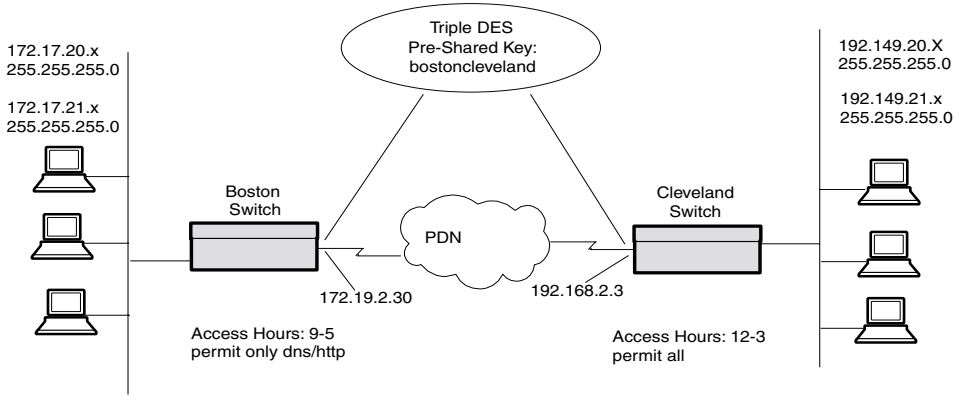
The branch office feature allows you to configure a secure tunnel connection between two private networks. Typically, one private network is behind a locally configured switch while the other is behind a remote switch. Branch office configuration allows you to configure the accessible subnetworks behind each switch. The configuration also contains the information that is necessary to set up the connection, such as the switch's IP addresses, encryption types, and authentication methods. Local policy restrictions such as access hours, filter sets, and call admission priorities can be applied to limit connectivity into local subnetworks. Figure 10 shows a typical branch office environment.

Contivity provides a keepalive feature that uses the Dead Peer Detection (DPD) technique based on an IETF draft for IPSec connections. The DPD keepalive mechanism is efficient, encrypted and scalable. Contivity automatically uses DPD keepalive if IPsec peers support the feature. Otherwise, Contivity uses a rekey keepalive approach.

The keepalive feature detects loss of connectivity; and can trigger a failover mechanism. Contivity supports keepalives for the following branch office connection types:

- Peer-to-Peer and asymmetric connections
- IPsec and L2TP/IPsec connections
- Nailed-up Connections

Figure 10 Typical branch office environment



Integrated firewall solutions

The switch includes integrated firewall solutions that are designed to meet the needs of a variety of customers. The switch provides a choice of possible firewall solutions: the Contivity Stateful Firewall and Contivity Tunnel Filter.

With the addition of the Contivity Stateful Firewall, the switch can perform a variety of secure routing functions, depending upon how you set up the switch's routing capabilities. For example, you can configure the switch to securely route non-tunneled traffic from its private interface, through the firewall, and out its public interface. This configuration would enable users on the switch's private network to access the Internet without requiring a separate, dedicated router.

By using stateful inspection, the Contivity Stateful Firewall provides a high level of security, the fastest runtime, and the flexibility to define the rules to fit your environment. The firewall delivers full firewall capabilities, assuring the highest level of network security. To do this, the firewall examines both incoming and outgoing packets running against a common security policy. All service rules are interpreted based on IP conversations (not packets) and are fully stateful. Security rules do not filter packets directly, but the firewall services determine how to process them based on the defined security policy.

The Contivity Stateful Firewall achieves optimum performance as a result of advanced memory management techniques and optimized packet inspection.

For many customers, the Contivity Tunnel Filter provides a cost-effective level of protection.

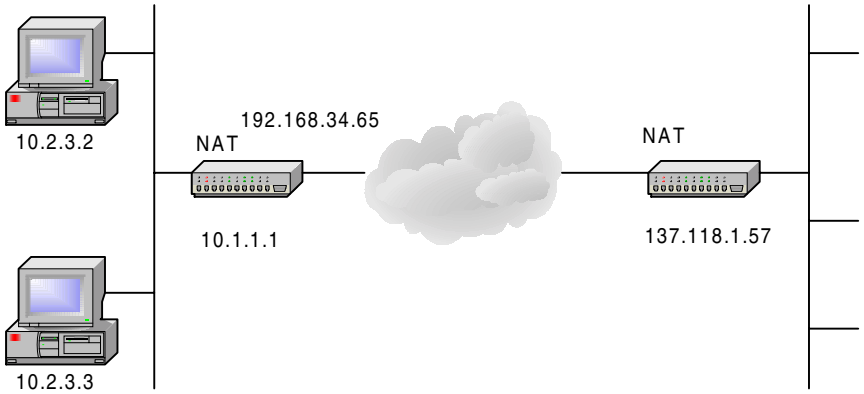
Existing customers who have Check Point* FireWall-1 V4.0 from Check Point Software Technologies, LTD. can continue to use the separately licensed integrated firewall.

Network Address Translation (NAT)

NAT is the translation of one network IP address that is used within a LAN to a different IP address that is used outside the LAN. This feature allows a system to be identified by one address on its own network, yet be identified by a totally different address to systems on a different network.

NAT allows branch office connections to eliminate problems with overlapping addresses on both sides of the connection, and it allows you to hide the LAN addresses. Figure 11 shows a sample NAT environment.

Figure 11 Switch sample NAT environment



Filters for access control

As you progressively put in place the components for your switch, access control becomes an important security mechanism. You want to have complete control over which users have access to particular servers and services.

Filtering is the mechanism that controls finely tuned access to specific hosts and services. Each user has a specific filter profile based on their group's profile that describes which resources on the network they can access. The filters are defined by:

- Protocol ID
- Direction
- Source, destination IP addresses
- Source, destination port
- TCP connection establishment

A filter profile consists of a list of rules that you create to perform precisely the action that you want. These rules are tested in order until the first match is found. Therefore, the order of the rules is very important. The filtering mechanism works such that if no rule matches then the packet is discarded (denied); this means that no traffic is transmitted or received unless it is specifically permitted.

Quality of service

With the switch fully configured and many clients dialing in to it, performance and quality of service become important. The switch supports two internal Quality of Service (QoS) mechanisms and can also participate in external network signaling to enhance performance. Forwarding priority allows for prioritized traffic, and Call admission priority allows you to reserve connection resources for high-priority users. In addition, external QoS using Resource ReSerVation Protocol (RSVP) signals the public network to reserve a portion of the network's bandwidth for a specific connection.

Bandwidth management

Bandwidth management capabilities let you manage the switch CPU and interface bandwidth resources to ensure that tunneled sessions get predictable and adequate levels of service. Bandwidth management allows you to configure the switch resources for users, branch offices, and interface-routed traffic. Bandwidth components keep track of and control the level of bandwidth being used on the physical interfaces and the tunnels.

Differentiated Services (DiffServ)

DiffServ settings classify and mark packets to receive specified per-hop forwarding behavior on each node along their path. Sophisticated classification, marking, policing, and shaping operations are implemented at network boundaries or hosts. Network resources are allocated to traffic streams by service provisioning policies that govern how traffic is marked and conditioned upon entry to a differentiated services-capable network, and how that traffic is forwarded within that network. Any DiffServ code points (DSCPs) not recognized are forwarded as if marked for the default behavior, Best Effort (BE).

Forwarding priority

Forwarding priority quality of service allows you to assign each user to one of four priority classes. Each class is guaranteed different maximum forwarding times between the interfaces of the switch. For example, high-priority traffic generated by the company CEO would be protected from high-bandwidth traffic generated by lower-priority users. Or, you might assign the sales team to Priority 1 to make sure that they could always place orders, especially during the quarter-end rush.

Call Admission priority

Call Admission priority quality of service allows you to assign each user group profile to one of four priority classes (from 1-high to 4-low) for call admission. The switch can reserve connections for each class of user, guaranteeing that a large number of low-priority users do not lock out the high-priority users. When the switch is servicing the maximum number of low-priority sessions, no further low-priority connections are accepted. Once a connection is accepted it is never dropped.

RSVP

The switch supports Resource ReSerVation Protocol (RSVP) quality of service for the Internet. Successful external network-level quality of service requires the cooperation of all the devices on the network (between the user and either the access point to the private network or the ultimate destination host). Currently, RSVP is the best-defined technology for resource reservation. However, only a few service providers offer a service that uses RSVP.

Network Time Protocol (NTP)

NTP synchronizes the clocks of various devices across networks. It also automatically adjusts the time of network devices so that they are synchronized within milliseconds. The switch receives NTP updates from an NTP time server and continuously synchronizes its clock to universal standard time. The switch supports up to eight NTP (unicast) servers and broadcast, multicast servers.

Accounting

Accounting is the final component of the switch. Detailed logs record the various activities performed by the switch, including:

- Time and amount of data transferred for each user session
- Security violations and failed authentication
- Device configuration changes, tracking who made the change and when
- System events

These logs are directly available from the management interface, and they can be exported to other applications for additional processing.

The switch supports both internal storage and backup external storage. Automatic backup and archiving of the logs ensures that the logs are available.

As part of your preventive maintenance program, periodically review the various log files for unusual events. Investigate unusual activity and, if necessary, make policy or configuration changes.

Local configuration and LDAP file structures

Configurations available for the switch are stored in two places:

- LDAP database
- Local configuration file

LDAP database information can be shared among multiple switches, while the local configuration file data is specific to a particular switch.

The LDAP database has been designed to be supported across multiple switches. The LDAP database includes:

- Profiles
- Groups
- Users
- Remote user access

- Filters
- Hours
- Tunnel authentication configuration for IPSec, PPTP, L2TP, L2F
- RADIUS authentication server
- RADIUS accounting
- Remote user IP acquisition, either DHCP or the internal address pool

Local information that is stored in the Nortel Networks configuration file is exclusive to a single switch, including:

- System identity
- Management IP address
- DNS host name
- DNS domain name
- Primary, secondary, and tertiary DNS servers
- Administrator ID and password
- LAN interfaces
- IP address
- Subnet mask
- Default gateway
- Type
- Description
- WAN interfaces
- IP address
- Description
- Remote IP address handling, either negotiated or specified
- PPP authentication settings
- PPP advanced settings
- Routing (static routes)
- LDAP
- Internal, including backup and restoration directory
- External, including the base domain name, host name or address, port, bind domain name, and bind password for the servers
- User IP address pool

- Administrator tools
- Automatic backup server configurations
- Recovery data
- FTP upgrade access information
- Saved configuration files
- Complete file system
- SNMP traps, including the enabled and disabled traps, host name, community string, trap interval
- System shutdown parameters
- Other
- Alarms

Accounting and logging

Each switch also stores its own accounting and system logging locally. The Event log contains the most recent 2000 events. The System log file is parsed into the Security log and the Configuration log.

Logging and status information are stored in the following:

- Event log
- System log
- Security log
- Configuration log

Command line interface

The command line interface allows you to make configuration changes to the switch via Telnet. You can access the command line interface by initiating a Telnet session to the switch management IP address; for example:

```
telnet 10.0.16.247
```

Then, you must log in using an administrator-privilege user account; for example:

```
Login: admin
Password: ********
CES>
```

Upon a successful login, the command line prompt appears. These prompts include CES> for the standard login prompt, CES# for enable commands, and CES(config)# for configuration mode. You can then enter any of the supported commands.

Serial point-to-point protocol (PPP)

The Serial Point-to-Point Protocol (PPP) feature allows you to manage the switch from a remote location using PPP and the serial interface. If the switch were to become unreachable over the Internet, you could still dial up and manage it through the serial interface menu.

With this feature, the serial interface becomes much like a private WAN interface. You can manage through it or even tunnel through it. When configuring Serial PPP, you can set the switch to Auto Detect, or you can specify that either PPP or the Serial menu are the options available through the serial port.

Contivity 1010/1050/1100 features

The Contivity 1010/1050/1100 family of products is a low cost, low-end IP services platform that provides Virtual Private Networking (VPN), routing, and firewall services. It is useful for carrier and enterprise customers who may want to deploy into their small business or home office locations.

This product offers the following features in addition to the existing Contivity VPN Switch features:

- Unique Default Configuration
- DHCP Server
- Point-to-Point Protocol over Ethernet (PPPoE)
- Branch office quick start utility

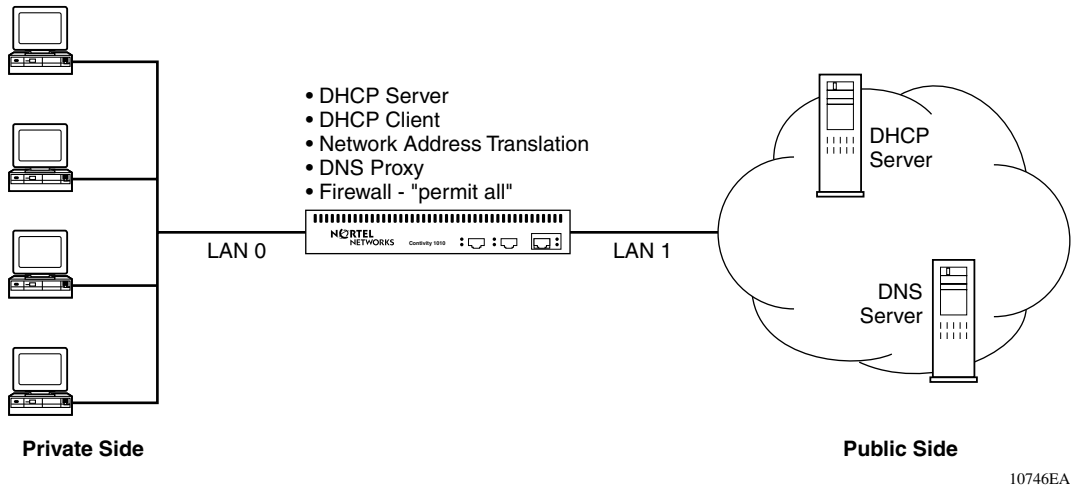
- DNS Proxy
- Compact flash disk

Default configuration

The Contivity 1010/1050/1100 series of switches uses a unique default configuration to simplify setup in the small office or home office environment.

Figure 12 show a typical default configuration.

Figure 12 Default configuration



DHCP Server

Dynamic Host Configuration Protocol (DHCP) dynamically assigns IP configuration parameters to clients and provides for centralized network administration. DHCP pushes configuration information to clients, including network address parameters and standard options. It also provides for interaction with DNS.

DHCP uses the concept of IP address leases. When a DHCP client requests an IP address, a DHCP server grants the client exclusive use of an assigned IP address for a specified period of time.

The Contivity 1010/1050/1100 series of switches include a full implementation of a DHCP server that is in compliance with RFC 2131 and RFC 2132.

PPPoE

Point-to-Point Protocol over Ethernet (PPPoE) allows PPP to run over Ethernet. Typically, PPP runs over serial interfaces and in most cases runs over phone lines connected to a server.

With DSL and cable modems, where a personal computer is connected to the Ethernet interface of the modem, ISPs cannot run PPP because PPP cannot run directly over Ethernet. ISPs often prefer to use PPP to provide features such as user authentication and bandwidth monitoring.

PPPoE is typically set up in two different configurations: PPPoE for the single user or PPPoE on a local network. For locations with single computers, the PPPoE client is typically loaded on the computer and it reaches the PPPoE server through the Ethernet connection via the DSL modem. The DSL modem then forwards the packets to the WAN interface without interpreting the PPPoE packets. The PPPoE packets reach their final destination (PPPoE server) for further handling. This is in compliance with RFC 2516.

The second configuration is usually seen in multi-computer locations, small offices, or branch offices where the entire LAN is connected to the Internet via DSL or cable modem. In this case, either the modem or the gateway acts as a PPPoE client.

Branch office quick start utility

The branch office quick start utility (BOQS) simplifies deployment of switches in the branch office environment. BOQS converts the Contivity 1010/1050/1100 device from an Internet Access Gateway into a Secure Access Gateway by provisioning a virtual private network (VPN) connection to a Central Office or optionally, to a Network Operation Center (NOC). BOQS adds VPN functionality to basic internet access connectivity. It allows a NOC or central office management to access the Contivity 1010, 1050, or 1100 so that network administrators can further configure the these units without going to the remote site.

DNS proxy

The Domain Name Service (DNS) is a method for mapping host names to IP addresses. It allows the Internet to provide an updated set of mappings for all Internet devices. A DNS server holds the segment of the DNS database for which it has authority. DNS clients are any TCP/IP applications that refer to hosts by host name. When an application needs to convert a host name to its IP address, it uses the client portion, which creates a DNS query specifying the host name and sends the query to a server. The server tries to find the host IP address by looking in its database or by making queries to other servers. Eventually, a DNS response is returned to the application, which contains the IP address or an error indicating that the host name is unknown.

It is common for companies to set up their own domain name system internally, and leave it to the ISP to handle all external DNS. These companies have their own DNS servers, but use the external DNS servers for non-company names. This *splits* the DNS names into two separate systems: the private, company-controlled DNS names and the Internet DNS names.

You can configure the Contivity 1010, 1050, or 1100 as a DNS proxy, which means that it can act like a DNS server for any PC on the private network. The PCs are configured to send their DNS queries to the DNS proxy, which in turn passes the query to its set of *true* DNS servers. Whether you have configured DHCP client or PPPoE determines which DNS servers will respond. When the DNS proxy receives a DNS query form a PC, it passes the query on to the DNS servers until it receives a response, which is subsequently returned to the PC.

Compact flash disk

The Contivity 1010/1050/1100 series of switches uses a compact flash disk instead of a traditional hard disk. The compact flash disk provides 32 MB of flash disk storage.

Because of the limited storage capacity, the following functionality is not provided:

- Safe mode
- Java runtime plug-in
- Graphs

- Japanese strings
- Context-sensitive help

The help files are located on the CD and on the Nortel Networks documentation Web site. When you click on the Help menu from the UI, you can enter the location of the help files on a server.

File compression is used extensively on the Contivity 1010/1050/1100 series of switches. Compressed files will retain their original names and all existing directory operations that the software performs will continue to work. The following functionality is compressed:

- VXworks image
- All Web pages
- All scripts
- Numerous text files

Chapter 2
Getting Started

This chapter describes methods for configuring and managing the Contivity VPN Switch.

Full details on hardware installation, including adding local area network (LAN) or wide area network (WAN) cards, are in the installation guide that came with the switch. You should complete the hardware installation before starting this chapter.

Serial Interface configuration

The Serial Interface allows you to give the switch a management IP address and subnet mask so that you can then use a Web browser for management. Your terminal emulator must use the following communications parameters:

- 9600 baud
- 8 data bits
- 1 stop bit
- No parity
- No flow control

The Serial Interface configuration procedure is typically only necessary in a system recovery situation.

1. Connect the serial cable (supplied with your switch) from the switch's serial port to a terminal or a communications port of a PC.
2. Power on the terminal or PC.

3 Using a terminal emulation program, such as HyperTerminal on the PC, press Enter. The Welcome screen appears and you are prompted to supply a user name and password.

```
Welcome to the Contivity VPN Switch
Copyright 1999,2000,2001 Nortel Networks

Version: V03_50.44
Creation date: Dec. 7, 2000, 20:51:06
Date: 04/27/2001
Unit Serial Number: 17563

Please enter the administrator's user name: admin

Please enter the administrator's password: setup
```

The factory default user name is *admin* and the default password is *setup*.

Before moving the switch from one network to another, change the management IP address, subnet mask, and default gateway. Otherwise, you will need to follow the Serial Interface configuration procedure to access your switch because it will not be accessible from a Web browser with an invalid address.

Before proceeding, verify that you have selected a management IP address that can be routed on the subnet attached to the private interface of the switch and have this information available.

4 After the user name and password have been entered, the following menu appears:

```
Main Menu: System is currently in NORMAL mode.

1) Interfaces
2) Administrator
3) Default Private Route Menu
4) Default Public Route Menu
5) Create A User Control Tunnel(IPsec) Profile
6) Restricted Management Mode FALSE
7) Allow HTTP Management TRUE
8) Firewall Options
9) Shutdown
B) System Boot Options
P) Configure Serial Port
C) Controlled Crash
L) Command Line Interface
R) Reset System to Factory Defaults
E) Exit, Save and Invoke Changes

Please select a menu choice (1 - 9,B,P,C,L,R,E):_
```

5 Enter **1** and press **Enter** to enter the Interface menu:

```
Interface Menu

0) Slot 0, Port 1, Private LAN
   Management IP Address = 11.0.0.100
   Subnet Mask = 255.0.0.0
   Interface IP Address = 11.0.0.10 (Subnet Mask = 255.0.0.0)
   Speed/Duplex = AutoNegotiate

1) Slot 1, Port 1, Public LAN
   IP Address = 192.168.1.20
   Subnet Mask = 255.255.255.0
   Speed/Duplex = AutoNegotiate

R) Return to the Main Menu.

Please select a menu choice:_
```

6 Select **0** and press **Enter** to enter the Slot 0, Port 1, Private LAN menu and add the management IP address.

The default settings appear, followed by the Configuration menu for the management IP address. The Old Management IP Address field is blank on a new switch.

```
Slot 0, Port 1, Private LAN
  Interface IP Address =
  Subnet Mask = 0.0.0.0
  Speed/Duplex = AutoNegotiate

  * Type 0.0.0.0 to delete.

Old Management IP Address =
New Management IP Address =
```

If you are prepared to assign a management IP address to the private interface, do so at this time at the New Interface IP Address prompt and press Enter. If you do not wish to assign a private IP address at this point, simply press Enter.

```
Old Interface IP Address =
New Interface IP Address =
```

7 The subnet mask menu will appear next. Enter the desired subnet mask and press Enter.

```
Old Subnet Mask = 0.0.0.0
New Subnet Mask =
```

The Interface Configuration menu appears. Select the menu option desired and press Enter.

8 Follow the screen prompts. The descriptions of the fields required to complete this procedure are in the "Startup configuration requirements" section at the beginning of this chapter.

9 After you complete the configuration, type E and press Enter to save the settings and exit. You can then manage the switch from a Web browser.



Note: This administrator's password is also the primary administrator's password. This password guarantees access to the switch via the serial port or a Web browser. This administrator's user ID (default=admin) and password (default=setup) combination is also called the *primary administrator*. This person always has access to all screens and controls, including the serial port and the recovery disk. There can be only one primary administrator.

Management IP address

The management IP address must reside on a defined subnet. If there are no defined subnets, the management IP address can reside on any one that you create. You cannot change your management IP address to a subnet different from an already defined subnet.

To change the management IP address when an interface IP address is defined:

1. Select Option 1 Interfaces and then select 0.
2. Select Option 3 Default Private Route.
3. Select Option 1 Interfaces.
4. Go through the menu again and change the management IP address to the address that you want.

To reset the switch IP address values using the serial interface:

- Change the management IP address, subnet mask, or gateway address to another set of values on the same logical network.
- Move the switch to a different logical or physical network.
- Completely reset the management IP addresses.
- Reset the switch so that it can be moved to a currently unknown logical network, which allows the network administrator to assign new addresses to the switch when it gets there. You might use this option if you do not know the addresses in advance.

To change IP addresses or move the switch to another network:

1. Log in to the serial interface.
2. Select Option 3 (Gateway IP Address). Set the address to 0.0.0.0.
3. Select Option 1 (Management IP Address). Set the address to 0.0.0.0.
4. Choose E.
5. Wait at least 15 seconds. Serial interface data is only written to the hard disk at 15-second intervals.

Safe mode configuration

The switch can be booted in one of the two system modes: Safe mode or Normal mode. Each mode has its own software image, configuration files, and LDAP database.

A system booted in Safe mode is only allowed to accept secured management tunnel establishment. When the secured management tunnel is established, Telnet, HTTP, and FTP traffic is allowed to come into the switch; no other VPN traffic is allowed through the secured management tunnel or the switch.

In Normal mode, the system operates with the normal software and configuration and transports both VPN traffic and management traffic.

To save your configuration into the Save Mode boot directory:

1. Select B) System Boot options.
2. Select 2) System Reset options.
3. Select 1) Reset system to Normal Mode.
4. Select 2) Reset system to Safe Mode.

Web browser

Internet Explorer 5.5 or 6.0 and Netscape 4.7.*x* or 6.2 are the supported Web browsers.

When using the Serial Interface Configuration, launch a Web browser of your choice. If you choose to use Microsoft Internet Explorer, Version 5.0 as your browser, the ie5 subdirectory on the CD contains the ie5setup.exe file. This file starts the Internet Explorer installation procedure, then links to the Microsoft Web site to download the full Internet Explorer Version 5.0 software. You must have Internet access when you run this setup program in order to complete the Internet Explorer installation.

1. Enter the management IP address to invoke the Nortel Networks Login screen. For example, if the management IP address is 10.2.3.2, then the Uniform Resource Locator (URL) is http://10.2.3.2.

2 Select an option in the navigational menu and submenu, and then you are prompted for the login and password.

3 Enter the system default login and password in lowercase characters, as follows:

Login: **admin**
Password: **setup**

At this point, follow the Quick Start Configuration procedure or the Guided Configuration procedure. Refer to Table 3 on page 70 for help in determining which procedure to use.

Preparing for configuration

To properly prepare for configuration of the Contivity VPN Switch, you should have the following items available:

- A plan to distribute IP addresses to clients when connections are requested; for example, via a DHCP server or an internal client address pool (with an address pool you need a range of IP addresses).
- An Authentication database: If you are not using internal authentication via the LDAP database, then make sure you have either the external LDAP or the RADIUS server's IP address and password or Shared Secret (password).
- An external accounting server, such as RADIUS, with its IP address and Shared Secret (password).
- Client dial-in: Prepare the clients for the type of tunneling protocol they need to use. The PPTP client application is available on the Nortel Networks CD for Windows 95, and it comes with Windows 98 and Windows NT. Nortel Networks also provides the IPSec client on the Nortel Networks CD.
- A complete network topology of the "environment" in which you are testing the switch, including the switch, the default router address, and any other IP addresses that you think might be required.

Table 3 shows the alternatives when first configuring your switch. Begin with either the Quick Start or the Guided Configuration. After you are familiar with the switch's navigational menu and capabilities, select Manage Switch.

Table 3 Web interface configuration options

Configuration type	Results
Quick Start	Configure and test a basic PPTP configuration
Guided Config	Structured switch configuration and management
Manage Switch	Comprehensive switch configuration and management

The following list provides a place for you to record the information that you need to configure basic Contivity VPN Switch parameters.

Table 4 Configuration checklist

Screen	Values required	Your Values
System→Identity	Management IP address	
System→Domain Identity	Host name Domain name	
System→DNS Service	Primary IP address Secondary IP address Tertiary IP address	
System→LAN	Private IP address Public IP address	
System→WAN (if using T1, V.35, or T3)	ISP provided information	
System→Date and Time	Manual entry of date and time or NTP configuration server broadcast or multicast IP address	
Services→Available	IPSec private address IPSec public address PPTP private address PPTP public address L2TP and L2F public address L2TP and L2F private address	

Table 4 Configuration checklist

Screen	Values required	Your Values
Services→Management	HTTP private address HTTP public address SNMP private address SNMP public address FTP private address FTP public address TELNET private address TELNET public address CRL retrieval private address CRL retrieval public address	
Routing→Static Routes Enabled/Disabled	Public gateway IP address Private gateway IP address	
Routing→OSPF Enabled/Disabled	Router ID AS boundary router (true or false)	
Routing→Rip Enabled/ Disabled	True or false	
Routing→Interfaces LAN IP Address	OSPF (enabled or disabled) RIP (enabled or disabled) VRRP (enabled or disabled)	
Servers→Radius Auth	Access (enabled or disabled Server-Supported Option (enabled or disabled) Radius Servers (enabled or disabled) Primary host name or IP addresses, public or private, Port, Shared secret/confirmed Alternate 1 host name or IP addresses, public or private, Port, Shared secret/confirmed Alternate 2 host name or IP addresses, public or private, Port, Shared secret/confirmed	

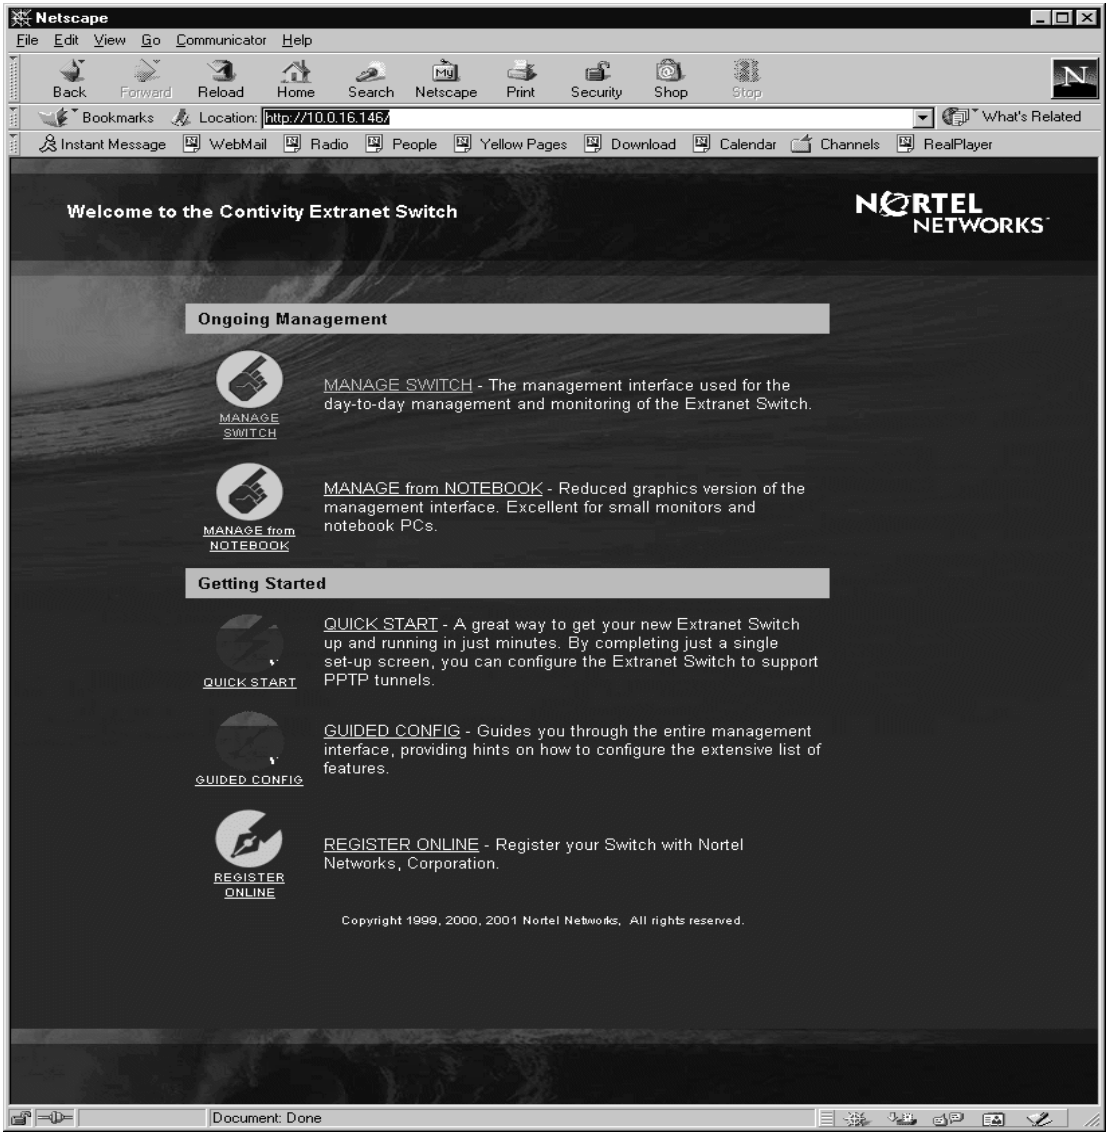
Table 4 Configuration checklist

Screen	Values required	Your Values
Servers→LDAP	Internal or external Base DN Master IP address, port or SSL Bind DN, Bind password, Confirmed Slave 1 IP address, port or SSL Bind DN, Bind password, Confirmed Slave 2 IP address, port or SSL Bind DN, Bind password, Confirmed	
Servers→User IP addr	Broadcast Any DHCP or DHCP servers: Primary IP address Secondary IP address Tertiary IP address Address pool: Pool name Start End Subnet mask	
Admin→Key Installation	Advanced routing install key Contivity Stateful Firewall install key	
Admin→Auto Backup	Automatic Backup file servers IP address of FTP servers for backup: Host Path User ID Password	

Welcome screen

The Welcome screen (Figure 13) allows access to any of the configuration areas for the Contivity VPN Switch.

Figure 13 Welcome screen



Before entering the configuration options, first register your switch to activate licenses, warranties, and services.

To start using your switch, choose from one of the following options:

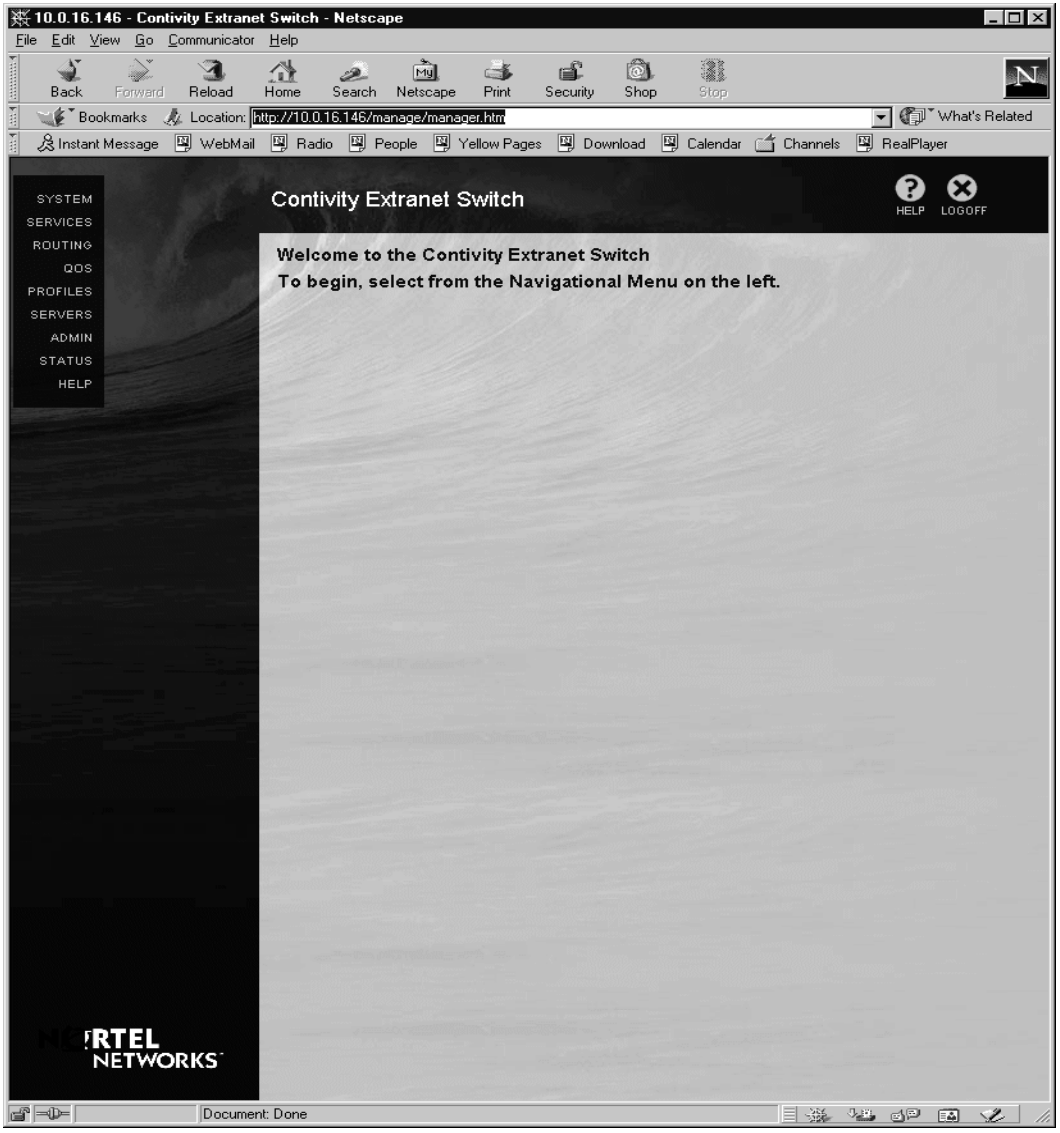
- Click on Manage Switch to begin a configuration management session. This option allows access to all Configuration Management facilities. For your first configuration, follow the Quick Start or Guided Configuration.
- Click on Manage from Notebook to run the Contivity VPN Switch Manager in notebook display mode.
- Click on Quick Start to begin the Quick Start Configuration. This option allows you to configure interfaces, set up PPTP tunnels for up to three users, and establish a connection to the switch. If you prepare for the configuration as recommended, the Quick Start can take as little as 15 minutes to complete.
- Click on Guided Config to begin the Guided Configuration. This option allows access to all Configuration Management facilities. The design and structure of the Guided Configuration, however, is such that you might want to follow the top-to-bottom layout provided. This approach walks you through the entire navigational menu from the Profiles to the Admin selections.

  Each functional area begins with a summary of the objectives of the area and then steps you through the area (for example, profiles), one subsection at a time. Context-sensitive help is available at each subsection to supplement the summary.

  Provided you have the information required to set up the switch, the Guided Configuration is estimated to take two to three hours to complete, depending on how extensive your configuration is.
- Click on Register Online to register the switch with Nortel Networks. It only takes a few minutes and it gives you access to the latest software and technical tips. Your switch requires Internet access from the private interface in order to register.

The switch's navigational menu options include the top-level configuration and monitoring areas of the switch. Each of these key areas has secondary levels, which appear once you click on an area; for example, when you click on System, the secondary level listings appear (Figure 14).

Figure 14 Navigational menu



The menu is structured to provide system configuration details, followed by profiles for groups and users. Then you can configure authentication servers, secure tunnels, administrative details, and monitor the status of the switch.

Setting up the Contivity 1010/1050/1100

The Contivity 1010/1050/1100 series of switches provides support for five (5) tunnels at introduction and 30 tunnels for licensing. The maximum tunnels include the sum of all branch office, client, and management tunnels combined. For example, if one management tunnel and two branch office tunnels are open, only two client tunnels can be connected initially (27 client tunnels with the 30 tunnel license). The license is for 25 additional tunnels. LDAP supports 150 entries.

Default configuration parameters

By default, the Contivity 1010, 1050, and 1100 are configured with the following parameters:

- The DHCP server is configured on the switch's private interface, with a default range of 192.168.1.3/24 to 192.168.1.255/24. By default, 192.168.1.1 and 192.168.1.2 are assigned to the branch office switch's private and management interfaces, respectively. The DHCP server provides its own address for the DNS server and default gateway.
- The DHCP client is configured on the switch's public interface to retrieve its IP address from the ISP's DHCP server. Other parameters retrieved from the DHCP server should include the default gateway and the DNS server.
- DNS proxy is configured to forward DNS requests to an external DNS server. The address of the DNS server is obtained during startup from the ISP's DHCP.
- Network Address Translation (NAT) translates the private IP address space (determined by default configuration of the DHCP server) into one public address assigned to the public interface by your ISP.
- Port NAT maps multiple IP addresses in the private space to a single public IP address. The default configuration only supports initiating IP sessions from the private side of the switch, which reduces security risks.
- The Contivity Interface Filter is set as the default firewall.
- The firewall setting PermitAll is the default for both the public and private interfaces. This default is different from the DenyAll default setting for other Contivity VPN switches.

Branch office quick start

Network administrators and service providers can use the branch office quick start for provisioning IP-based VPN services on a large scale. It provides VPN services using Contivity 1010, 1050, or 1100 devices as branch office VPN switches and other Contivity VPN switches as central office switches.

In addition to connectivity, the central office switch must be able to accept newly created secure connections from the Contivity 1010, 1050, or 1100. Therefore, the BOQS must be used with the knowledge and approval of a network administrator. It can only be initiated after IP addressing has been planned and the central office switch has been configured. Then you can send the provisioning parameters to the remote branch office locations.

The Contivity 1010, 1050, and 1100 must be connected to a public network and have access to the Internet before local users can use BOQS. The unique default configuration allows easy deployment of Contivity 1010, 1050, and 1100 switches in DHCP configurations (where a DHCP server is used on the public network). However, if you use static IP addressing or PPPoE on the public side, the Contivity 1010, 1050, or 1100 must be configured manually before local users can use BOQS.

All users on the private network must renew their IP addresses. For further information, see your Microsoft documentation. When the branch office tunnels are established, public access to the Internet is replaced with access to the central office.



Note: The BOQS will remain accessible after the information is entered. The network administrator must change the admin account (**username/password**) to restrict access.

After the VPN services are provisioned, branch office networks are logically connected to a central office network or to a NOC network. Branch office end users can rerun BOQS multiple times to restore the initial VPN configuration or to fix data errors.

BOQS supports two network topologies:

- Enterprise topology where the network operations center is located within the central office.
- Service Provider topology where the network operations center is independent entity from central office

Enterprise environment

Enterprises rely on the Internet for a wide range of business operations. Locating a Contivity VPN Switch at the central office and Contivity 1010/1050/1100 series of switches at the branch office locations provides a secure solution. Before you deploy the Contivity 1010, 1050, or 1100 switches at the local sites, you must configure configure routing and tunnels on the switch at the central office.

For routing, you must do the following:

- Enable global RIP service.
- Enable RIP on private interface.
- Disallow importing default routes in the group where responder tunnels are created.
- From the CLI, use the following command:

```
router rip
network a.b.c.d m.m.m.m
exit
```

For tunnels, you must do the following:

- Create one responder tunnel for each branch office Contivity 1010/1050/1100 device.
- Set the Connection Type to Responder.
- Be sure that the Control Tunnel option is NOT selected.
- Determine the connection name for the tunnel. Nortel Networks recommends that the name be the same as the initiator ID, but it could be the same as the central office tunnel name.
- Set the state to Enabled.
- Set the Local Filter to permit all.
- Set IPSEC Authentication to Text Pre-Shared Key.
- Set the Initiator ID to the same name as the central office tunnel name.

- Set the Text Pre-Shared Key to the same name as central office tunnel password.
- Set Dynamic Routing to enabled.
- Set RIP to enabled.

After the central office setup and the BOQS are complete, the Contivity 1010, 1050, or 1100 is directly accessible from the central office. This means that there is just one hop between the central office and the branch office. RIP propagates routes to this subnet across the tunnel created by BOQS.

You must have at least two more IP addresses than IP workstations on the Contivity 1010, 1050, or 1100 private network. The first address from the subnet is assigned to the private interface of the branch office switch and the second address becomes the management IP address of the switch. Each branch office must be in its own subnet.

Table 5 shows how offices with approximately 50 workstations can each have subnets assigned.

Table 5 Subnet assignments

Private Network IP address	Private Network IP Mask	Contivity 1010/ 1050/1100 Private Interface Address	Contivity 1010/1050/ 1100 Management Interface Address	BO Workstations Addresses (assigned by DHCP Server)
200.1.1.0	255.255.255.192	200.1.1.1	200.1.1.2	From 200.1.1.3 to 200.1.1.62
200.1.1.64	255.255.255.192	200.1.1.65	200.1.1.66	From 200.1.1.67 to 200.1.1.126
200.1.1.128	255.255.255.192	200.1.1.129	200.1.1.130	From 200.1.1.131 to 200.1.1.190

Service provider environment

Service providers generally have an isolated NOC from which all devices are managed. The addressing scheme could be different from a central office and require a separate designated tunnel to configure the Contivity 1010/1050/1100 series of switches.

Every Contivity 1010, 1050, and 1100 must have a distinct IP address that is visible from the NOC subnet. A NOC can assign any address reachable from a NOC network to a Contivity 1010, 1050, or 1100. BOQS configures NAT on the NOC tunnel to translate the address specified in the "Branch office switch manage NAT IP address" and "management address from branch office private subnet." If the field is empty, the NOC must use an actual management address to access the Contivity 1010, 1050, or 1100.

Because the NOC tunnel uses static routing, all Contivity 1010, 1050, and 1100 devices must be configured with a static route to the NOC private network. The NOC private address and NOC private mask fields are where a BOQS user enters this information. This information is the same for all Contivity 1010, 1050, and 1100 devices.

You must provision the NOC switch to accept control tunnel connections from the branch office. Because static routing is used in control tunnels, you do not have to enable routing protocols on the NOC switch. Use the following guidelines:

- All responder tunnels should be created in one group or in subgroups of one group for easy management. Connection Name of the tunnel should correspond to NOC tunnel name and created in an enabled state with local filter set to Permit All.
- Text Pre-Shared Key should be selected as the IPSEC authentication method, Initiator ID set to the value of Control Tunnel Name, and Text Pre-Shared Key should be equal to Control Tunnel password.
- Select Static routing. Accessible local networks should be added. All networks from which the Contivity 1010, 1050, or 1100 will be managed must be on that list.
- NAT Local option should NOT be used.
- Accessible Remote Networks should contain one address subnet (mask equal to 255.255.255.255) with Contivity 1010, 1050, or 1100 Management IP. Contivity 1010, 1050, or 1100 Management IP is either explicitly provided in the field "Branch office switch manage NAT IP address" or if this field is left empty, it is the second address from the subnet specified in the Branch Office Private IP Address and Mask fields.

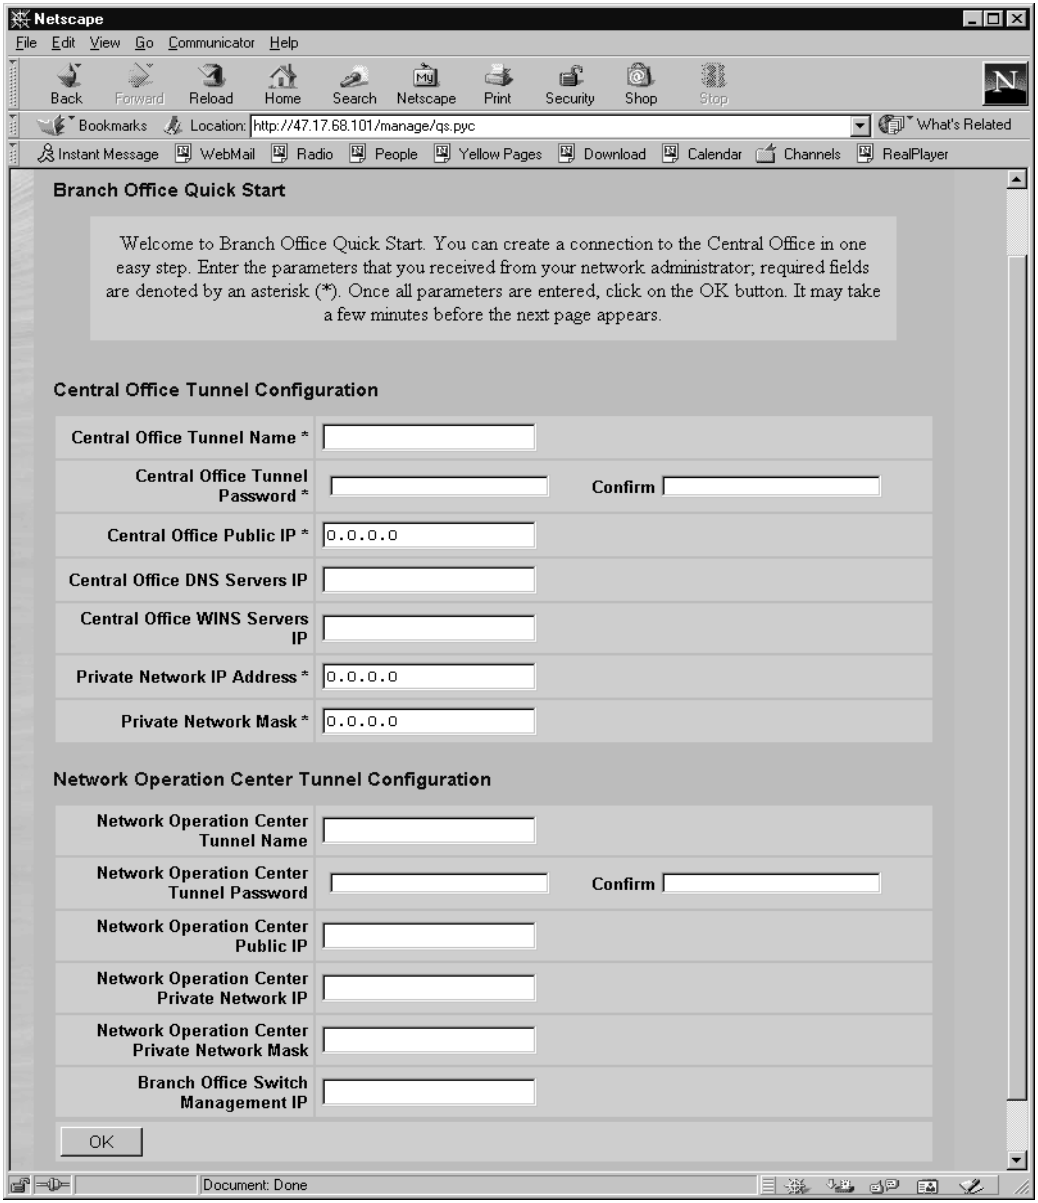
Table 6 contains the BOQS parameters.

Table 6 BOQS parameters

Central office tunnel configuration	
Central office tunnel name	Name of the branch office tunnel on the central office switch.
Central office tunnel password	Password for the branch office tunnel.
Central office public IP address	Public address of the central office switch (same for all branch offices).
Central office DNS server IP address	IP address of the DNS server in the central office. The DHCP server configured on private interface distributes this address to the branch office. You can configure multiple addresses, but you must separate them with commas. This field is optional and can be left empty.
Central office WINS sever IP address	IP address of WINS server in the central office. The DHCP server configured on private interface distributes this address to the branch office workstations. You can configure multiple addresses, but you must separate them with commas. This field is optional and can be left empty.
Private network IP address	Subnet address of the branch office network..
Private network mask	Subnet mask of the branch office network.
Network Operation Center tunnel configuration	
Network operation center tunnel name	Name of the branch office tunnel configured on NOC switch (same as initiator id on the NOC switch).
Network operation center tunnel password	Text pre-shared key used in branch office tunnel.
Network operations center public IP address	Public address of NOC switch (same for all branch offices).
Network operations center private network IP address	IP Address part of subnet address in which NOC is located (private subnet of NOC switch).
Network operations center private net mask	IP mask of subnet address in which NOC is located (private subnet of NOC switch).
Branch office switch management IP address	Address used by NOC to manage switch. Must be unique for each Contivity 1010/1050/1100 and reachable from the NOC. If left empty, can be managed with the second address of the subnet configured in branch office private network IP address/ IP mask field

Figure 15 shows the branch office quick start screen.

Figure 15 Branch office quick start screen



Branch office quick start template

The branch office quick start template provides a list of values that the local Contivity 1010, 1050 or 1100 users will need to enter on the BOQS screen. You can enter the appropriate values in the right-hand column and then fax, send, or email the template to the local user along with any other information that they may need, such as who to contact for further information or questions.

Central office tunnel configuration	**Your value**
Central office tunnel name	
Central office tunnel password	
Central office public IP address	
Central office DNS server IP address	
Central office WINS sever IP address	
Private network IP address	
Private network mask	
Network Operation Center tunnel configuration	**Your value**
Network operation center tunnel name	
Network operation center tunnel password	
Network operations center public IP address	
Network operations center private network IP address	
Network operations center private net mask	
Branch office switch management IP address	

Connecting for internet access

This section provides information on how to set up the Contivity 1010/1050/1100 series of switches for basic Internet access through a cable or DSL modem.

To set up the Contivity 1010, 1050, or 1100:

1 Plug the power cord into an AC power outlet, plug the external power supply into the port on the back labeled DC Input, and then turn on the power switch.

2. Connect the cable or DSL modem to the LAN 1 (**public**) port located on the front panel using a standard Ethernet cable (not included with the Contivity unit).

   - To connect more than one PC or other devices to the Contivity 1050 or 1100, use standard Ethernet cables to connect to the LAN 0 (**private)** ports labeled A-D located on the front panel. To connect the Contivity 1050 or 1100 directly to the PC, use a standard Ethernet cable.
   - To connect more than one PC to the Contivity 1010, connect an Ethernet switch or hub to the LAN 0 port and then connect the PCs. To connect the Contivity 1010 directly to the PC, use the cable that ships with it.

3. Because the Contivity 1010/1050/1100 has the DHCP server enabled by default on the LAN 0 port, the PC should be set up to accept IP addresses that the DHCP server will assign the LAN 0 ports. For Microsoft* Windows* operating systems:

   a. Go to the Control Panel and click on the network connections icon.

   b. Select TCP/IP protocol, click on the Properties button, and then select the "Obtain an IP address automatically" option.

   c. Reboot your PC to obtain the new IP address (range 192.168.1.3 to 192.168.1.254).

   For other operating systems, see the user documentation for those systems.

4. Determine the type of addressing to use:

   By default, the LAN 1 port acts as a DHCP client, which receives an IP address from the public side. Launch a Web browser application to verify connectivity to the Internet.

   If you use Point to Point Protocol over Ethernet (PPPoE), you must use the Web-based management utility for advanced configuration options to enable it:

   a. Open the Web browser and type http://192.168.1.2. Click on Manage Switch and enter **admin** and **setup** as the username and password.

   b. Go to the System→LAN screen, select PPPoE from the Select Protocol list and click on Apply. The Add PPPoE Interface screen appears.

   c. For Administrative State, click on Enable and for Interface Filter, select permit all.

   d. Click on OK.

e Go to an external Web site to verify Internet access.

Alternatively, you can use the following CLI commands to configure PPPoE client (with username – abc and password xyz):

```
interface fastethernet 1/1
ip address type static # is it needed
pppoe ppp username abc password xyz
pppoe enable
exit
```

If you use static IP addressing, you must use the Web-based management utility for advanced configuration options to enable it:

a Open the Web browser and type http://192.168.1.2. Click on Manage Switch and enter **admin** and **setup** as the username and password.

b Go to the System→LAN screen. select IP from the Select Protocol list and click on Apply. The Add IP Address screen appears.

c Select the Static option and enter the IP address and the subnet mask.

d Click on OK.

e Go to an external Web site to verify Internet access.

Alternatively, you can use the following CLI commands to configure the static address (for example, 47.17.2.23/24) and default gateway (for example, 47.17.2.1):

```
interface fastethernet 1/1
ip address type static
ip address 47.17.2.33 255.255.255.0
ip route 0.0.0.0 0.0.0.0 47.17.2.1 enabled
exit
```

Compact flash disk

The compact flash disk provides 32 MB of flash disk storage. Two software images can be stored on the flash disk at the same time. Operational changes for the compact flash disk are:

- The config file is saved every minute and the past three versions are kept. The config file is only written when the configuration changes.
- The on-disk system log (syslog) is not be supported. However, you can configure an external syslog server.

- No accounting information is stored on the compact flash disk. However, an external RADIUS accounting server is supported.
- The data collection log (DCLOG) is not supported, which means that the graphing capabilities of the UI are also not supported.
- The core is not saved on the compact flash disk. It is sent to an FTP server. Configuration parameters for the FTP server are stored in flash. The core file is placed on the server. To set up the FTP coredump, got to the FTP Coredump section of the Admin→Admintstrator screen, click on the Enabled checkbox and enter the appropriate FTP server information. Because many switches may be configured to coredump to the same location, the core files will have a more descriptive name: core*_date_24-hour-time_management_ip*.mem. For example, a core file generated by 10.0.8.186 on Oct.12th, 2001, at 4:46:06 PM will be named core_20011012_164606_10.0.8.186.mem.

Chapter 3
Configuring the System

This chapter describes how to configure various system-level features.Each Contivity VPN switch is uniquely identified by the system's address and domain name system (DNS) name. The DNS name can be used instead of the IP address to identify the switch and launch its management interface through a Web browser.

System identity

The System Identity screen allows you to optionally change your switch Management IP address, and provide the DNS Host Name and Domain Name. Additionally, you can assign up to three DNS addresses to resolve IP address name resolution requests.

Resetting Management IP address values

To reset the switch Management IP address values using the serial interface:

1. Enter a Management IP Address for the system. You need this address to contact all system services, such as HTTP, FTP, and SNMP. To be accessible, the Management IP Address must map to the same network as one of the private interfaces. For example, if you are planning on assigning IP address 10.2.3.3 with the subnet mask 255.255.0.0 to the private physical interface, the Management IP Address must reside in the 10.2.x.x network. If you configure the switch on one network and plan to move it to another network, change the Management IP address and private LAN interface addresses before moving the switch. Then, communicate with the switch using the new Management IP address from your browser's URL address field.
2. Enter a Name to identify the system. This should be the same name that is used by the DNS server to identify the management address of the switch that

is located on your private network. You can enter up to 64 characters maximum.

3 Enter the Name of the Internet Domain into which this system is being placed. This must be the same Internet Domain as the System Name in the Domain Name System (DNS) server. A domain is a part of the Internet naming hierarchy that refers to general groupings of networks that are based on organization-type or geography. For example, mycompany.com is the domain name for a commercial (.com) enterprise.

To configure a DNS server:

1 On the System→Identity screen, the DNS Proxy Enabled/Disabled check box allows you to select whether you want the DNS Proxy to act as a DNS server to the private side. It it enabled by default.

2 Click on the Split DNS check box if you have a split name space.

3 For Primary, enter the address of the DNS server that the DNS proxy tries to contact first.

4 For Second, enter an address for the Second Domain Name System (DNS) server. If the Primary DNS server doesn't respond in a few seconds, service is requested of the Second DNS server (if present).

5 For Third, enter an address for the Third Domain Name System (DNS) server. If the Primary and Secondary DNS servers doesn't respond, service is requested of the third DNS server (if present).

6 For Fourth, enter an address for the Fourth Domain Name System (DNS) server. If the preceding servers doesn't respond, service is requested of the fourth DNS server (if present).

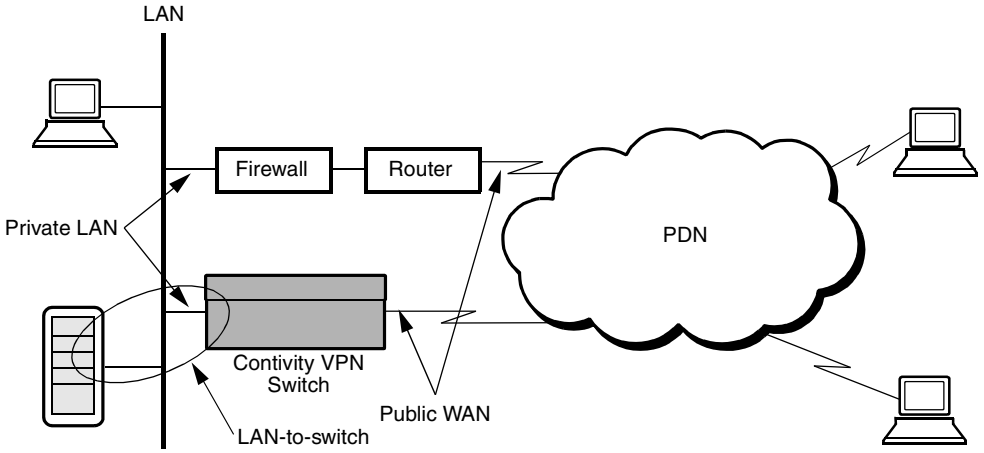
7 Click on OK. The switch checks all of the DNS addresses to see if they respond and then provides an operational or error status.

LAN interfaces

The LAN interface that is available on the system board is configured to be private by default. Connect its interface to your corporate LAN. Additional interfaces that are inserted into the expansion slots are public by default.

The private LAN interface and the management IP address must be on the same network, and the public LAN interface should be on a different network, both physically and logically. If your switch has a single network interface and you want to position the switch behind the firewall and router, then you should set the switch's interface type to private. Figure 16 shows a connection from a LAN to a switch.

Figure 16 LAN-to-switch connection



A host can send only enough packets to a public interface to establish a tunnel connection. If the tunnel is not established before a preset maximum number-of-packets-allowed counter is reached, then the packets from that host are discarded.

Public indicates that this interface is attached to a public data network like the Internet. The switch rejects nontunneled protocols and only accepts tunneled protocols like IPSec, PPTP, L2TP, L2F, and diagnostic ping on a public interface. A host can send only enough packets to a public interface to establish a tunnel connection. If the tunnel is not established before a preset maximum number-of-packets-allowed counter is reached, then the packets from that host are discarded.

When the public interface is configured to act as a DHCP client, the DHCP client needs to correspond to an external DHCP server to acquire the IP address, subnet mask and default route parameters. You can set a cost value to give preferential routing when two or more public DHCP clients are configured. In this situation, DSL and cable modem are the preferred choice for connections to the internet.

Private indicates that an interface is attached to the private network and it can accept nontunneled networking protocols such as TCP/IP, FTP, and HTTP. The Private interface also accepts tunneled protocols (for example, IPSec, PPTP, L2TP, L2F) that can be used for secure management access to the switch.



Note: The private LAN interface and the management IP address should be on the same network, and the public LAN interface should be on a different network, both physically and logically.

If you have one network only and want to position the switch behind the firewall and router, then you should use a private LAN interface only (do not use a public LAN interface).

Point to Point Protocol over Ethernet allows PPP to run over Ethernet. It is used by DSL or cable modem users and allows multiple users to have access to the Internet. PPPoE has the following usage restrictions:

- Cannot use dynamic routing on PPPoE interfaces
- Can use PPPoE on only one interface at a time
- IPX is not supported
- Must be set on a public Ethernet interface
- Must set the Administrative State to enabled (disabled by default)
- Must set the appropriate filter (deny all by default)
- PPPoE changes are dynamically applied

From the System→LAN screen, you can:

- Click on Configure to modify the interface characteristics.
- Click on Statistics to view the Link Statistics.

- From the Select Protocol list, select the tunneling protocol to use: IP is the standard Internet Protocol. and Point to Point Protocol over Ethernet (PPPoE) allows PPP to run over Ethernet.



Note: With PPPoE you cannot use dynamic routing on PPPoE interfaces. You can use PPPoE on only one interface at a time. IPX is not supported.

This screen also provides the following information about the LAN interfaces:

- IP Address shows the current IP address that is assigned to the interface.
- Subnet Mask defines which bits of the IP address represent the network the device is on and which bits represent the host's ID on the network. The device uses the Subnet Mask to determine which IP addresses are directly reachable on the network and which must be routed through a gateway. A sample IP address is 10.2.3.3 with a subnet mask of 255.255.0.0. This indicates that all hosts with addresses 10.2.*n*.*n* are directly reachable.
- Interface Filter shows whether the Contivity Stateful Firewall is in use on this LAN interface (this reflects the selection on the Services→Firewall screen). This entry also shows the interface filter that is currently being used by the Contivity Firewall. This is the interface filter that is selected on the System→LAN Interfaces→Edit IP Address screen. If no interface filter has been selected, the default of Deny All is used.

The Configure button on the System→LAN Configure screen allows you to provide optional information for the LAN interfaces. This information then appears on the System→LAN screen. Additional fields appear on the Edit LAN Interface screen for optional network cards. LAN represents the physical port interface to which you assign an IP address. Slot *n* Interface *n* represents an optional LAN card in expansion Slot *n* using Interface *n*.

1 Under the Configuration section, use the Speed/Duplex field to automatically or manually configure the LAN interface's port speed and mode.



Note: You can also use the Interface selection on the switch's Serial Port menu to set auto negotiation.

2 Select Auto-Negotiate to specify that the switch automatically set the port speed and mode to match the best service provided by the connected station, up to 100 Mbps in full-duplex mode. Auto-Negotiate is the default selection, and complies with the IEEE 802.3u auto negotiating standard.

3 Select one of the following selections to manually set the LAN interface's port speed and mode to match the speed and mode used by the connected station.

- 100Mbs/Full duplex
- 100Mbs/Half duplex
- 10Mbs/Full duplex
- 10Mbs/Half duplex



Note: You might not be able to connect to the remote system if the system is not using auto negotiation or if it uses an incompatible form of auto negotiation. If this occurs, manually set your switch's speed and mode settings to match those used by the remote system.

4 You can provide an optional description for the LAN interface. The description appears on the LAN Interfaces screen.

5 MAC Pause (Ethernet packet flow control) enables the switch to automatically adjust and control the flow of incoming and/or outgoing packets from any standard speed LAN device. Check to enable MAC Pause (Frame-based flow control) on the selected interface port. When enabled, specify the appropriate Pause parameters to be set in the hardware. Specify a value for MAC Pause Ticks. Select a value from the Free Receive FIFO Threshold list. The default is 0.

6 The State field appears on the screen for an optional LAN card in expansion Slot *n* using Interface *n*. Click to enable or disable the card.

7 The Interface Type field appears on the screen for an optional LAN card in expansion Slot *n* using Interface *n*. Click to specify whether the interface is public or private.

The Add IP address screen (System→LAN Add) allows to you assign an IP address and subnet mask to the interface. Use the Edit IP Address screen (System→LAN Edit) to modify the information.

1 Check the Static button to assign a static, unchanging IP address.

The IP address consists of 32 bits, which are written as four octets (8-bit bytes) in dotted-decimal format; for example 192.168.34.21.

The subnet mask defines how many bits of the IP address represent the network the device is on and how many bits represent the host's ID on the network.

The device uses the Subnet Mask to determine which IP addresses are directly reachable on the network and which must be routed through a gateway. A sample IP address is 10.2.3.3 with a subnet mask of 255.255.0.0. This indicates that all hosts with addresses 10.2.*n.n* are directly reachable.

2 Check the DHCP button to have the IP address assigned by a DHCP server.

3 Enter a cost. The default is 10.

4 Interface filter shows whether or not the Contivity Firewall is in use (this reflects the selection on the Services→Firewall screen).

5 This entry also shows the interface filter that is currently being used by the Contivity Firewall. Select from a list of all interface filters that have been set up on the switch (on the Profiles→Filters screen), and to select a different filter for the Contivity Firewall.



Note: If you change the interface filter setting, a message informs you that you must restart your switch before the new interface filter is used. If the Contivity Firewall is not enabled, the new selection has no effect.

6 Use the New Interface Filter link to go to the Profiles→Filters screen and create a new filter. The default Interface Filter setting is Deny All.

You can copy an existing tunnel filter for use as an interface filter or vice versa. However, when you copy a filter, the operation does not copy any components (such as a rule, port, protocol, or address) that have the same name as an existing component in the destination filter set. You should use unique names for the device in all filter rules.

To configure PPPoE:

1 Select PPPoE and click on Apply. The Add PPPoE Interface screen appears.

a Interface indicates the interface on which PPPoE is configured.

b Description shows a description of the interface, if one was entered when the interface was configured.

c Operational state shows the current operational state of the PPPoE interface.

d For Administrative State, click on Enable.

e On Demand must be enabled when running PPPoE.

f For Idle Timeout, specifiy a period of *n* minutes of inactivity, after which to tear down the connection. Enter a valid value from the range specified on screen. The Idle Timeout must be specified.

g Enter a value to indicate the cost of the connection. Enter a valid value from the range specified on screen.

h Enter a service name that describes the connection by indicating a connection location, class of service, or an ISP name, for example.

i Enter the local IP address, which is dynamically assigned to the interface.

j For Interface Filter, select permit all.

k Click OK.

2 Go back to the System→LAN screen and click on Edit.The Edit PPPoE Interface screen appears.

3 For PPP Authentication Settings, click on Configure. The PPP Local Authentication screen appears.

a Select either PAP or CHAP.

b Enter the UID and password.

c Click on OK.

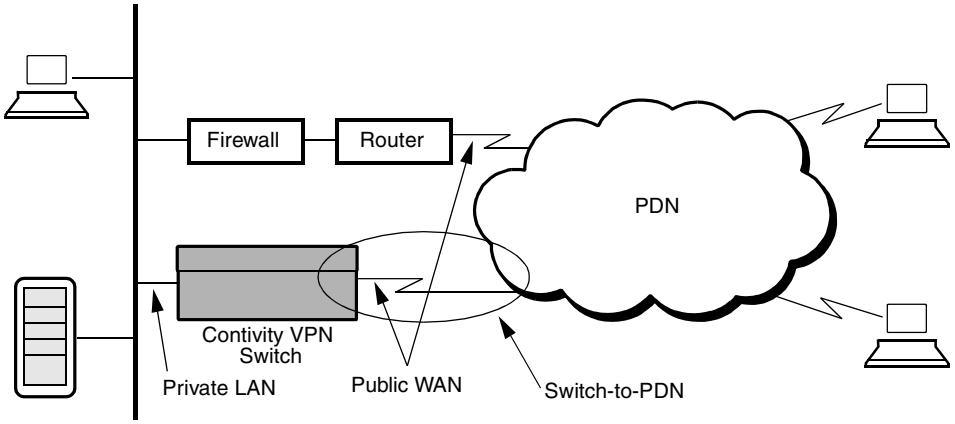
4 When you click the Configure button for PPP Advanced Settings the PPP Advanced Settings screen appears. Use this screen to configure PPP advanced settings for the PPPoE interface.

5 To verify the state of PPPoE, you can look at the Operational State in the System→LAN screen of the PPPoE. This indicates that PPPoE session is up. To verify that IP/PPP/PPPoE is established, you can either check the event log and search for PPPoE to see if IP is established or you can check the forwarding table to see if a public default route is entered through the PPPoE interface.

WAN interfaces

You assign WAN interface connections between the switch and the PDN. Figure 17 shows the connection attributes that you must configure. These attributes assign WAN interface connections between the switch and the ISP.

Figure 17 Switch-to-PDN configuration



The System→WAN Configure screen allows you to configure WAN devices with local and remote IP addresses and PPP-related settings. When you click on the PPP Authentication or Advanced Settings configuration buttons, the associated configuration screen appears. You also use this screen to specify the interface filter that is used for the option Contivity Firewall on this interface. The addresses are used by the IP Control Protocol (IPCP), which communicates IP addresses to peer connections over PPP. Many of these values are provided by your Internet (ISP).

The WAN interfaces screen (System→WAN) shows the WAN interfaces currently installed in the switch, the slot in which the cards reside, an interface description (if one has been provided), and the current state. It also indicates whether the Contivity Firewall is active and the interface filter that is in use. From this screen, you can move to another screen to configure or disable a WAN card, or view statistics.



Note: To change the IP address of a WAN link, you must disable the interface, change the address and re-enable the interface. This automatically disables static routes for the interface. If you change the IP address back to the original address, you must manually re-enable static routes.

To configure WAN interfaces:

1 Click on Configure to configure a new IP address for the associated device, or add or modify PPP advanced and authentication settings. The Configure WAN Interfaces screen appears.

2

3

Hardware encryption accelerator

The hardware encryption accelerator introduces a second-generation hardware accelerator card that off-loads bulk IPSec encryption duties from the main CPU. This card uses a single HiFn 7811 chip for encryption/compression and is PCI based. It performs triple DES or DES cryptography, LZS compression, and MD5 or SHA-1 authentication. Contact your Nortel Networks sales representative for a list of supported switch platforms. The new hardware accelerator card supports more tunnels encrypted on the card than the first-generation hardware accelerator card.



Note: Version 3.50 is the minimum software revision required for operation and the limit is one hardware encryption accelerator per Contivity unit.

The hardware encryption accelerator performs bulk encryption and compression algorithms for tunnel traffic in the switch. The encryption accelerator is an add-in device that can perform encryption, compression, or authentication on packets flowing over preestablished tunnels. The accelerator encodes and decodes packets with IPSec ESP encapsulation in hardware using triple DES and DES encryption, LZS compression, and SHA-1 or MD5 authentication. Since the hardware encryption accelerator supports these cryptographic functions, it becomes an element in the overall security envelope provided by the switch. Hardware acceleration is intended to be a background device that requires no additional handling once configured.

The encryption accelerator significantly improves the overall capacity of the switch to handle IPSec tunnel throughput. This improvement is realized when 5 or more tunnels are used, and for packet sizes greater than 512 bytes. It enables switch managers to specify (on Profiles→Groups→Connectivity) that users who require high bandwidth are given the highest priority. Then, such users are assured to run on the accelerator as long as the total number of high-priority users does not exceed the maximum sessions of the encryption accelerator card. The switch automatically performs load balancing of tunnel sessions across the accelerator and main CPUs.

The hardware accelerator screen shows the operational status that the switch reports on the hardware accelerator card. It also allows you to enable automatic recovery in case the card stops running. When the switch detects a recoverable failure, all sessions fail-over and are then handled by the software until the hardware resets and comes back on line.

The hardware accelerator UPtime field shows the elapsed time (days:hours:minutes:seconds) that the hardware accelerator has been running since it was last restarted. It might have been restarted by operator action or due to an error condition. However, if the system date and time are updated, the hardware accelerator uptime will only be correct after the next reboot.

You must have Administrator privileges to configure the card, which can be installed at any time, and you must restart the switch after configuring it.

A tunnel consists of two sessions, one each for incoming and outgoing traffic. Each session comprises a set of logical characteristics and parameters that are associated with a single communication path in a tunnel that renders a full-duplex connection. The number of sessions is double the number of tunnels, since one session is required for each direction and the tunnels are fully bidirectional.

Following is a listing of the switch's configuration, status, and monitoring paths related to the hardware accelerator:

- System→Accelerator→Hardware Accelerator: Configure
- Status→Health Check
- Status→Statistics→Hw Accel Stats
- Status→Statistics→Hw Accel Info
- Status→Event Log

The Hardware Accelerator screen shows the operational status that the switch reports for the card.

Hardware encryption accelerator and load balancing

When the hardware accelerator is installed, the switch's load balancing mechanisms account for accelerator's operation. When assigning tunnels to the hardware accelerator, the switch gives preemptive status to branch office tunnels and high priority user tunnels over non-priority user tunnels.

Hardware encryption accelerator security aspects

As part of the switch's security subsystem, it is imperative that the hardware accelerator does not compromise the overall security envelope of the switch. At startup, whenever the hardware accelerator is manually enabled, or whenever the accelerator recovers from a failure, the Power On self test diagnostics are performed to verify the integrity of the hardware. These tests include validation of the accelerator's encryption, MAC, and compression algorithms against their software counterparts. In the event POST fails, the accelerator is set offline.

Configuring Network Time Protocol (NTP)

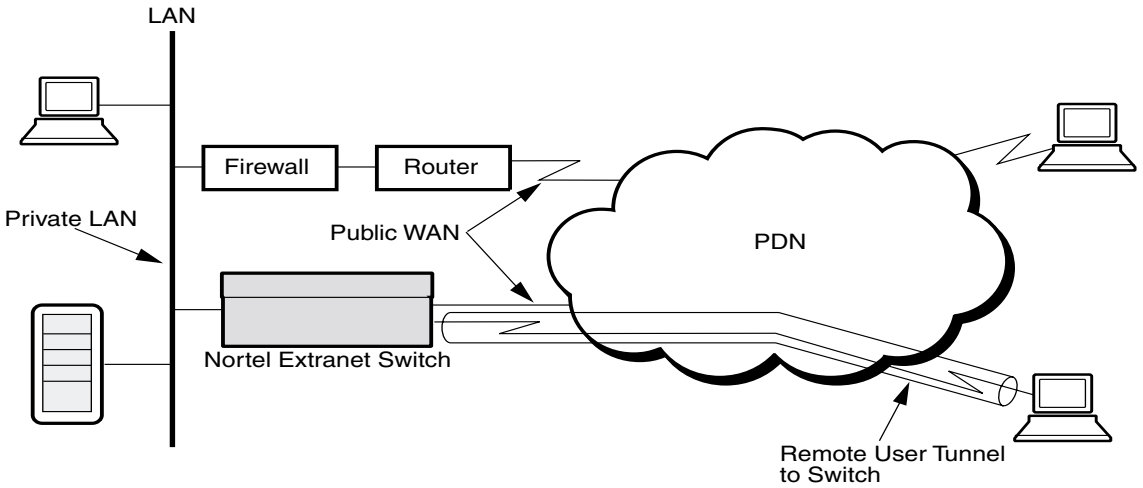
The switch provides support for the Network Time Protocol.

The System→Date and Time→Network Time Protocol screen allows you to set up NTP on the switch. NTP synchronizes the clocks of various devices across networks. It also automatically adjusts the time of network devices so that they are synchronized within milliseconds. The switch receives NTP updates from an NTP time server and continuously synchronizes its clock to universal standard time. The switch supports up to eight NTP (unicast) servers and broadcast, multicast servers.

Chapter 4
Configuring Tunnels

The switch uses the Internet and tunneling protocols to create secure extranets. The following sections describe configuring the tunnel portion of the switch. The configuration process includes setting up the authentication table and specific tunnel parameters, such as IPSec encryption, L2TP access concentrators, and L2F network access servers. Figure 18 shows a typical network illustration with the switch connected to the PDN (public data network) and to a remote user through a tunnel.

Figure 18 Tunnel connection configuration



The connection attributes that you configure in the switch enable the remote user to create a tunnel into the switch. However, you are not configuring the connection from the remote user to the Internet Service Provider (ISP) at this point. The actual connection to the switch is a tunnel that is started from the remote user's PC through its dial-up connection. That connection is to the Internet (typically using an ISP), through the Internet, and ends at the switch on the private, corporate network.

The switch associates all remote users with a group, which dictates the attributes that are assigned to a remote user session. A group can even consist of a single user, thereby creating a personal extranet.

The switch organizes groups in a hierarchical manner. At the top of the hierarchy is the base group. The base group \Base contains the default characteristics that each new group inherits. You add additional groups to the hierarchy as children of the base group.

The switch does not enforce the Maximum Number of Logins across tunnel types. If you set the number of simultaneous logins to 1, a client can still get another tunnel type connection if the client is configured to use multiple tunnel types. To limit the number of connections a client can have, configure the user for a single tunnel type.

The switch takes precautions against unauthorized users potentially hacking tunneled information when the switch is operating in Split Tunnel mode. The primary precaution is to drop packets that do not have the IP address that is assigned to the tunnel connection as its source address. For example, you establish a PPP dial-up connection to the Internet with an IP address of 192.168.21.3. When you start the tunneled connection to a switch, you are assigned a tunnel IP address of 192.192.192.192. Now, any packets that attempt to pass through the tunnel connection with a source IP address of 192.168.21.3 (or any address other than 192.192.192.192) are dropped. Furthermore, you can enable filters on the switch to limit the protocol types that can pass through a tunneled connection. To completely eliminate security risks, do not use the Split Tunneling feature.



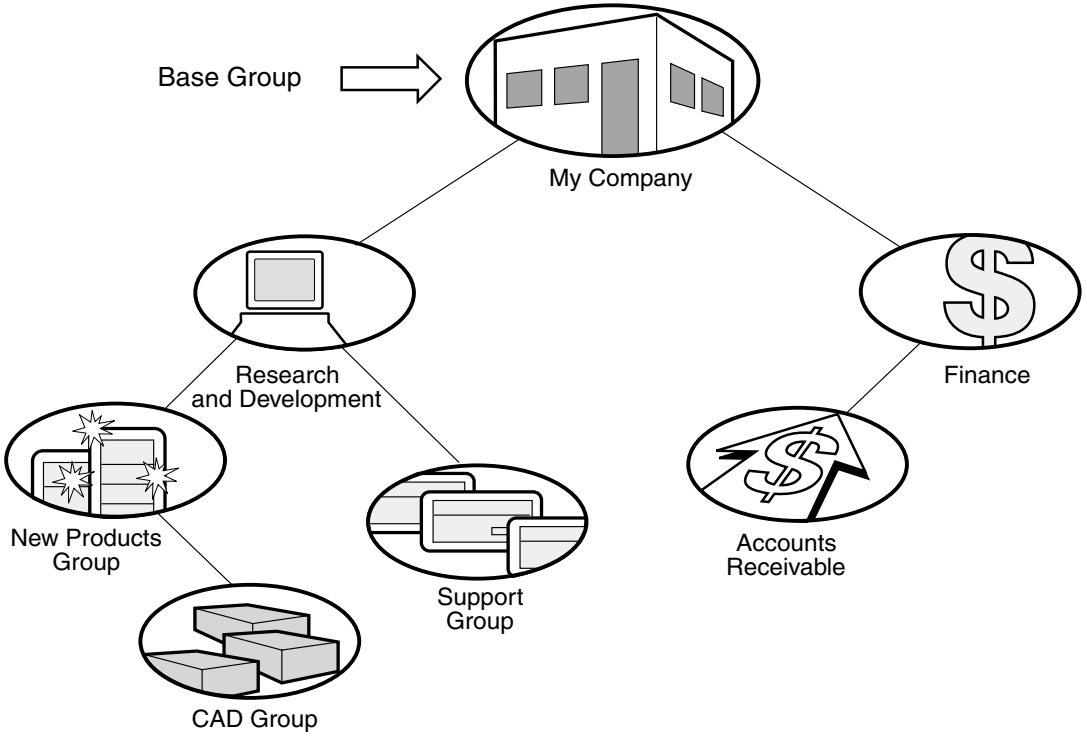
Note: PPP multilink is not supported with branch office tunnels. It is only supported with end user tunnels.

Password aging does not work for administrator accounts. Also, the following are client-specific password management symptoms:

- If you are using the IPSec client, you are warned three times that there will be an impending password expiration. You should change the password immediately. IPSec clients using versions earlier than 1.5.2 do not receive a password expiration warning.
- If you are using the PPTP client with the Connection Manager, the Connection Manager generates an impending password expiration warning.
- Other clients (L2TP and L2F) and PPTP client users who are not using the Connection Manager have no warning and no longer can log on. You must contact your system administrator if this happens. In this case, the switch is unable to notify the client because it has no actual control over the client. With PPTP, use the Connection Manager to establish a connection. With L2TP or L2F, set the Password Maximum Age to zero (never expires).

Figure 19 illustrates a group hierarchy. In this example, \Base is the base group. Research and Development and Finance are child groups of the base group, and they are parent groups to those below them.

Figure 19 Base, parent, and child group associations



Groups are collections of users with the same access attributes and rights. If all users have identical characteristics, then only one group is necessary. You create multiple groups when you need different attributes. A Lightweight Directory Access Protocol (LDAP) database stores users, groups, and their attributes. You can store this database internally (on the switch's hard disk) or externally (on a network host running LDAP server software).

The switch authenticates each user that tries to connect to the switch by checking the user ID and password against a database. The switch supports both LDAP and Remote Access Dial-In User Session (RADIUS) databases for authentication. When using LDAP for authentication, the user is always assigned to a group since LDAP also contains the user, group, and attribute information.

When authenticating a Point-to-Point Tunneling Protocol (PPTP) client against a RADIUS database, the group for a user requesting a session can be returned from the RADIUS server as a RADIUS class attribute.

When authenticating an IPSec client, the remote user is by default assigned to the group ID. If the group ID and group password are correct, the switch passes the user ID and password (or token card) to the RADIUS server for authentication.

You define a set of group attributes and give it a name. This group name is known as the Relative Distinguished Name and it is added to the LDAP database name when performing the database lookup.



Note: The group name Certificates is not allowed as a valid group name when created under the /Base directory. If you change the name to Certificate, the group is created properly. If you create the group Certificates in a /Base subdirectory, it is created properly.

In addition to assigning users to groups and providing authentication access, other group characteristics that you can configure include:

- Access hours are time ranges during which access is allowed for users in a group. These time ranges are configured from the Profiles→Hours screen.

- Call admission priority is the priority level (from low to highest) you want to permit for the group. Each level is assigned a percentage of the total number of calls allowed access to the switch. If there is a particularly high number of users logged in, new users could be denied call access, based on their call admission priority, until existing callers disconnect.
- Forwarding priority is the priority level (from low to highest) that you want to provide to sessions for users in this group. Forwarding priority assures a certain level of latency and bandwidth allocation. For example, a group with the highest forwarding priority has the highest possible bandwidth service and the lowest level of latency. Conversely, if there is a particularly high level of traffic on the line, packets for a low-priority group might be delayed or dropped. Since a low-priority group has the least amount of bandwidth and the highest level of latency, some of its packets would wait until the higher-priority-level packets have been forwarded or they would be dropped.
- Number of logins is the maximum number of simultaneous logins of a single tunnel type.
- Idle timeout is the amount of time a connection can be idle (no data has been transmitted or received through the connection for the specified amount of time). When the idle timeout expires, the session is terminated.
- User attributes include maximum password management, minimum password length, and enabling alphanumeric passwords.
- Tunnel filters are packet filters that control the type of access allowed for users in a group, based on various parameters, including Protocol ID, Direction, IP addresses, Source, Port, and TCP Connection establishment. Choose the Profiles→Filters screen to create tunnel filters.
- RSVP allows you to signal the network for required bandwidth. Only the controlled load-service is supported.
- Tunneling settings allow you to configure IPSec, PPTP, L2TP, and L2F settings.

A group inherits attributes from its parent group. For example, if the Research and Development group attributes include All Access Hours and Allow Static Addresses but deny Client-Supplied addresses, PPTP and IPSec tunneling, then the New Products (child) group would inherit these attributes.

Configuring user tunnels

To implement user tunnels, you must configure the following:

- Allowed tunnel access to the switch
- Switch tunneling protocol settings
- A user group
- Add users to the group
- A means, such as DHCP or pool, for assigning IP addresses to the client to allow user access

All tunneling protocols are enabled on the public and private networks by default. Since data in tunnels is encrypted, the default setting guarantees that all interactions with the switch are private. To prevent tunnel connections of a particular type (for all users, including administrators), you can simply disable the tunnel type.

For example, if you want to use IPSec as your only public tunneling protocol, then disable the Public selection for PPTP, L2TP, and L2F. By leaving IPSec, PPTP, L2TP, and L2F enabled on the private side, you can establish tunneled connections to the switch using any of the tunnel types from within your corporation.

To configure tunnel access to the switch:

1. Choose Services→Available.
2. Select the tunnel type.
3. Select the Management Protocol for the switch's private interface.
4. Use the RADIUS check boxes to permit RADIUS requests on the public and private interfaces of the switch. If you enable RADIUS traffic, you must also enable RADIUS on the Services→RADIUS screen.

Configuring the switch tunneling protocol settings is dependent on the tunnel type.

- For IPSec, choose Services→IPSec and select the required authentication, encryption, and authentication order.

- For PPTP, choose Services→PPTP and select the required authentication and authentication order.
- For L2TP, choose Services→L2TP and select the required authentication and authentication order and configure required L2TP access concentrators.
- For L2F, choose Services→L2F and select the required authentication and authentication order, and configure required network access servers.

To add a user group:

1. Go to Profiles→Groups and click on the Add button.
2. Enter a group name of up to 64 characters (spaces are permitted). For example, you could use Research and Development. The new group is a child of the selected parent group. Therefore, the new group initially inherits the parent group's network access attributes, including authentication, tunnel types, filtering, and priorities. When created, these inherited options can be overwritten for the new group.
3. Click on Apply and OK to add the group name.

To add a user profile in a group:

1. Select a group to which you want to add users from among those in the Group list. If you need to add a new group, select Profiles→Groups.
2. After selecting a group, you must click on Display to view the group members. This allows you to quickly change from viewing one group to another. The last names and first names of the selected group's users appear, sorted by last name.
3. Click on Add to add a user to the group; the Add User screen appears. This screen allows you to add a user profile. Only options that are enabled for the specified group appear on this screen. Also, only options that the administrator who is currently viewing the screen has rights to appear. A user profile includes:
   - User IDs
   - Passwords for the various tunneling protocols
   - Assignment of administrative rights

- An IP address that is always associated with the remote user



Note: You can assign a user to two different groups, but only if the user has two different user IDs. You cannot enter the same user ID in two different groups. A user account can have up to four user IDs, depending on the group configuration, the account. If you are creating an enterprise user ID standard, you should try to avoid schemes that might potentially create conflicts as your company grows. For example, you should not use the user's full first name and last initial.

4 Enter the first and last name of the user whose profile you want to add. This is the regular name associated with a person (for example, Mario Smith). This user can have different IDs and passwords for each tunnel type. You can move the user to a another group by selecting a different group name.

5 Enter a remote user static IP address to use in place of a pool (client-specified or DHCP) server-assigned IP address. This IP address is associated with the Static IP Address option in the Groups→Connectivity option (it is only used if the group allows it). If an IP address that is entered here is used instead of a DHCP server-assigned IP address, then only one login is allowed.

6 Enter the subnet mask. Assigning the correct subnet mask to a remote IPSec client is important when using split tunneling. When you enable split tunneling, packets destined to a host in the Split Tunnel Network list are directed into the tunnel by the IPSec client. All other traffic goes through a standard LAN or dial-up interface. This occurs on the client by adding the routes listed on the Split Tunnel Network list to the route table of the Microsoft TCP/IP stack and pointing those routes to the tunnel adapter interface. A route is also added to the route table based on the subnet mask assigned to the tunnel adapter. The IPSec Subnet Mask field allows you to specifically assign a subnet mask to a remote IPSec client that obtains an IP address either from the IP address pool, DHCP, RADIUS, or a static user configuration.



Note: If a host route for the destination address of the switch exists in the TCP/IP route table prior to launching the Contivity VPN Client, the route is deleted when the tunnel is closed.

To search within a selected group and then configure a user’s account:

1. Go to Profiles→Users→User Management.
2. Select a group from which you want to search for a particular user from the Group drop-down list box (at the top of the screen), and click on Display. The search is limited to the available groups.
3. Enter the appropriate text to search for in the input box.
4. Select one of the following as the preferred search method, then click on Search.
   - Last Name searches for a last name. You must enter the entire last name.
   - UID searches for a user ID.
   - Admin Rights searches for anyone who has View or Manage administrator privileges.
   - LDAP search allows you to enter any LDAP database attribute that is part of the person, organizational Person, or inetOrgPerson object database (for example, cn=common name or sn=surname) to generate the associated user's profile. Refer to your LDAP vendor's documentation for complete details.

Configuring branch offices

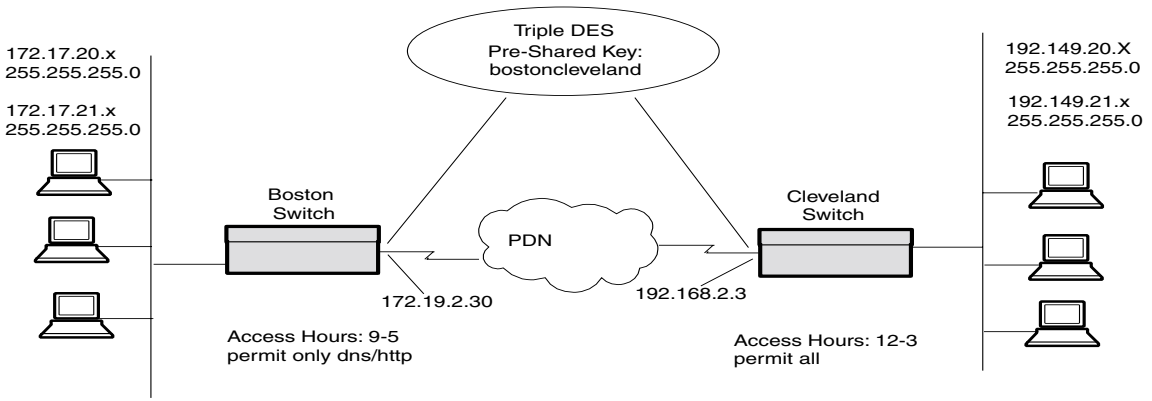
The branch office feature allows you to configure a secure tunnel connection between two private networks. Typically, one private network is behind a locally configured switch while the other is behind a remote switch. Branch office configuration allows you to configure the accessible subnetworks behind each switch. The configuration also contains the information that is necessary to set up the connection, such as the switch's IP addresses, encryption types, and authentication methods. You can apply local policy restrictions, such as access hours, filter sets, and call admission priorities, to limit connectivity into local subnetworks.

The Contivity VPN supports symmetric, or peer-to-peer branch office tunnels with fixed endpoints, and asymmetric branch office tunnels. An asymmetric branch office tunnel is a branch office tunnel where one of the endpoints does not have a fixed IP address. Such situations exist in the small branch office or SOHO

environments where the Contivity VPN Switch's public interface is behind a DSL or cable modem. The DSL or cable modem services typically do not guarantee a static IP address. Branch office tunnels in these situations are asymmetric because only one side of the tunnel can initiate a connection.

Figure 20 shows a typical branch office environment.

Figure 20 Typical branch office environment



The section "Sample branch office configuration procedure" provides sample branch office configurations for two locations, Boston and Cleveland. The initial configurations show connections established with pre-shared keys.

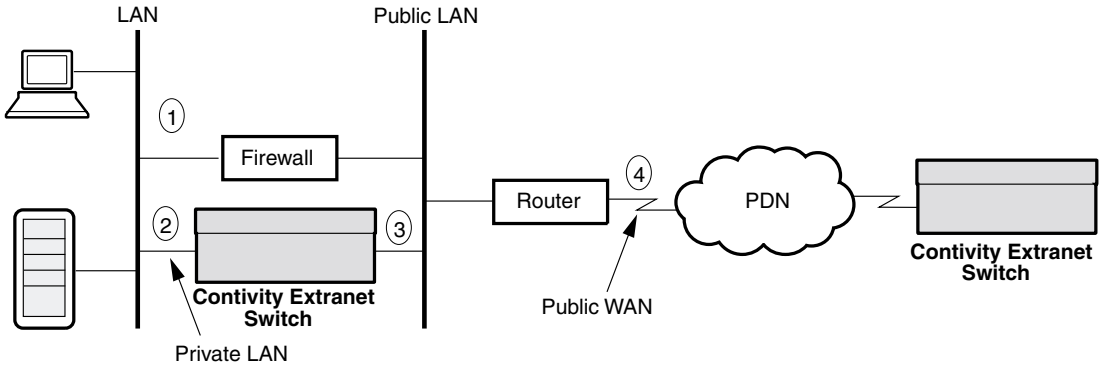
Branch-to-branch with a firewall and a router

In a mixed environment, you might want to tunnel connections to certain networks, and have all other traffic go to the Internet. You must configure the default gateway with a static route to the switch for accessible networks (refer to Profiles→Branch Office→Edit Branch Office Connection). The default private LAN router (the firewall) must redirect packets intended for remote branch office subnets.

In this case, as with any branch-to-branch configuration, you must configure each branch switch with the same encryption settings and pre-shared key (password). Of course, the accessible local and remote subnetwork addresses and subnet masks would be inverted in each switch's configuration.

Figure 21 shows a branch-to-branch configuration with a firewall and a router.

Figure 21 Branch-to-branch with a firewall and a router



In the branch-to-branch illustration, the following interactions take place with a Contivity VPN Switch:

1. The PC sends packets to the default route (the firewall).
2. The firewall redirects the packets to the local switch branch office connection.
3. The encapsulated data goes onto the public LAN.
4. The default public LAN route directs the encapsulated data to the remote switch branch office connection.

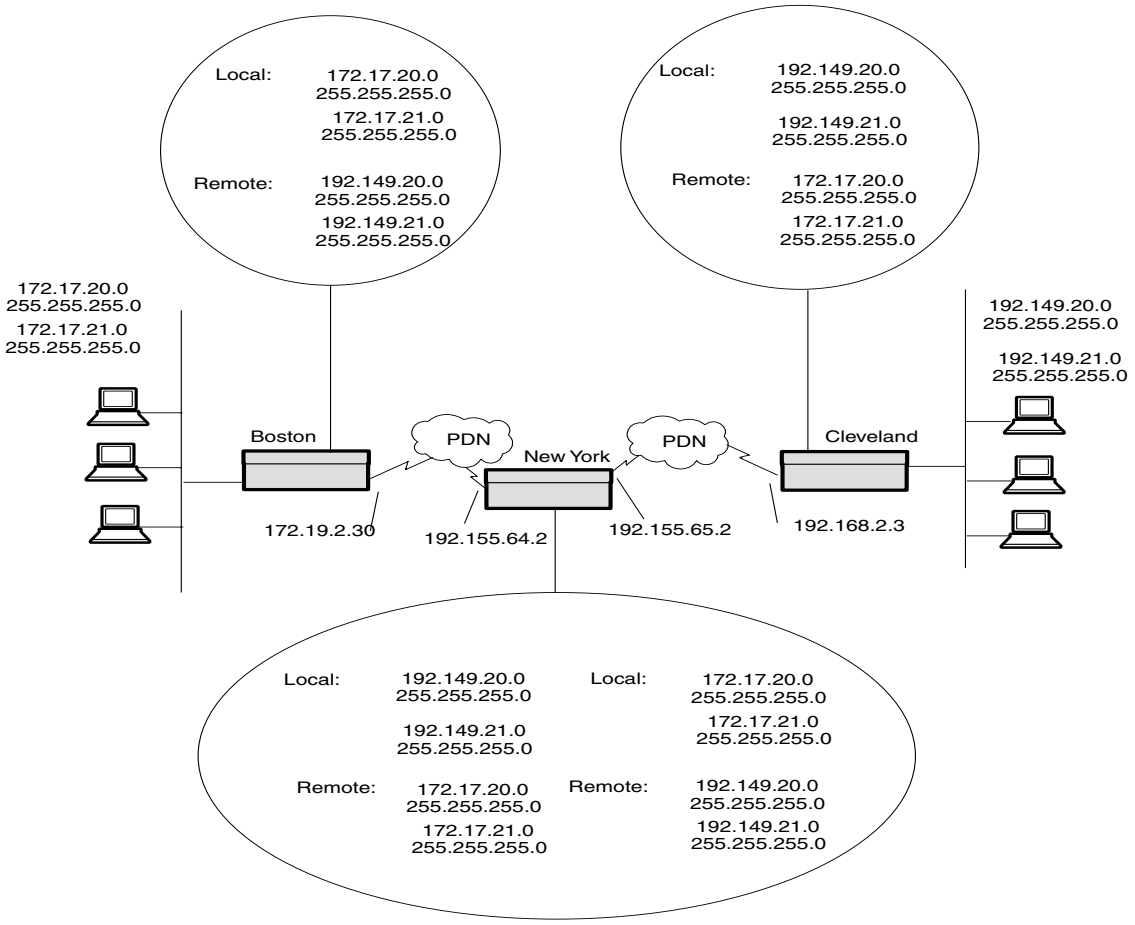
For switches that have a WAN link, actions 3 and 4 collapse together, and the encapsulated data is directed to the remote server.

Branch-to-branch through a gateway switch

In a three-switch topology, the two indirectly connected switches can create tunnels at will as long as each switch properly includes all of the local and remote subnetworks and subnetwork masks as accessible networks. Figure 22 shows the relationship between three switches and the local and remote networks that must be configured for each link to allow indirectly connected branch offices to bring up tunnels at will. The New York switch in the middle has two branch office connections configured.

All connections must have identical encryption settings. However, only adjacent connections are required to share keys. For example in the following figure, the Boston ↔ New York connection shares keys and the New York ↔ Cleveland connection shares keys. Boston and Cleveland are not required to share keys.

Figure 22 Indirectly connected branch offices



In branch offices, you might have two or more branches that use the same LAN addressing scheme. Nonetheless, users still have to communicate with one another across the branches.

NAT allows branch office connections to eliminate problems with overlapping addresses on both sides of the connection, and it allows you to hide the LAN addresses. To set up branch offices with NAT, see Chapter 6.

Sample branch office configuration procedure

This section describes an example of the procedure used to create a branch office VPN connection between two Contivity VPN Switches. In this example, the local switch is at the enterprise headquarters in Boston and the remote switch is at a sales office in Cleveland.

When you set up a branch office connection, you must perform the configuration procedure twice, once for each of the two switches that make up the connection. The branch office settings for the two switches mirror each other. For example, the local address setting that you configure on the Boston switch would be considered the remote address setting when you configure the Cleveland switch.

Figure 23 shows the configuration information and the addresses that are used in this example. It lists the procedure for setting up a branch office connection and the management Web pages that are used during the configuration process. Figure 24 also shows where the information from the figure is entered on the management Web pages.

Figure 23 Sample branch office configuration

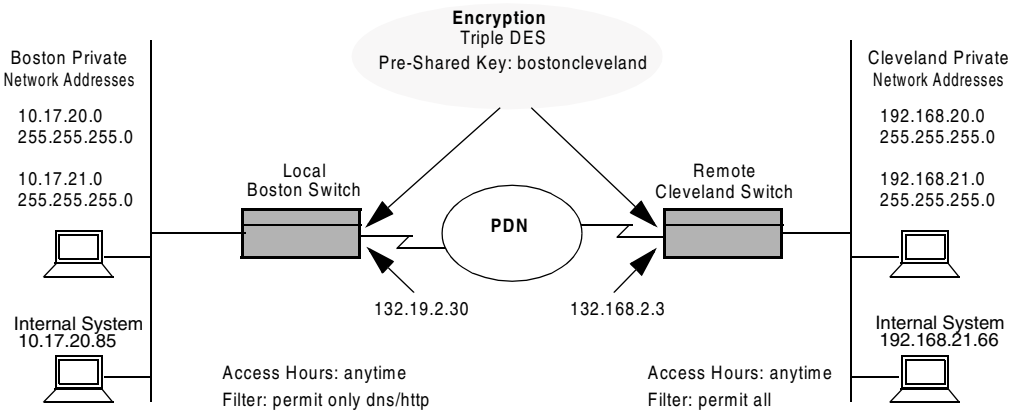
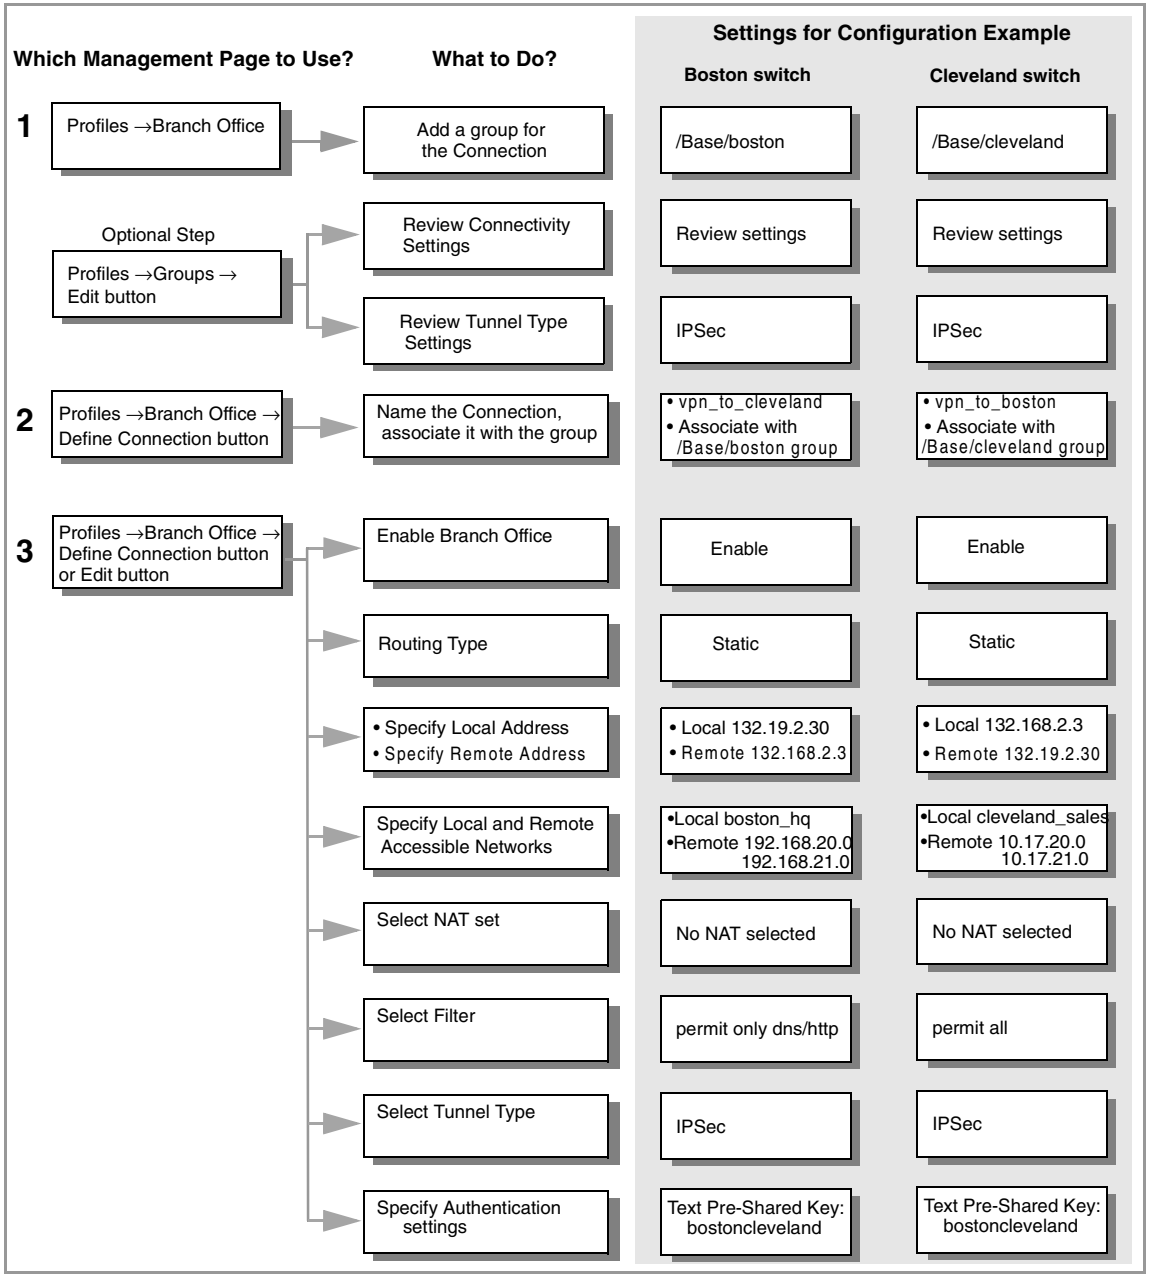


Figure 24 Setting up a branch office configuration



As the administrator of a branch office connection, you can manage the level of access that you give to users of the connection. You specify when the connection is used, what operations can be done through the connection, and which systems on the private networks can be accessed.

Before configuring a branch office connection, check that the following management Web pages are set up in accordance with your management policies and the planned usage for the connection. You use the settings on these pages when you configure the branch office connection.

- The System→Forwarding page must allow branch office-to-branch office traffic.
- The Profiles→Networks page must list the switch's private networks. In the sample configuration, the local switch's internal network name is boston_hq and the subnets are 10.17.20.0 and 10.17.21.0. The remote systems behind the remote switch can reach systems in these networks.

  The remote switch's internal network name is cleveland_sales and the subnets are 10.17.20.0 and 10.17.21.0. The remote systems behind the local switch can reach systems in these networks.
- The Profiles→Hours page must have the Hours of Access setting that you want to use. The example uses the setting of Anytime.
- The Profiles→Filters page must have the filters that you want to use for the branch office connection. For the example, the local switch uses a filter of permit only dns/http, and the remote switch uses permit all.

Configuring the local switch

You access the Branch Office screen through the Profiles→Branch Office menu selection. Use the subsequent configuration pages to set up a branch office connection. These enable you to specify the attributes of the switches that are participating in the connection and to set up network parameters, such as addresses and tunnel type, for the connection.

To create a new branch office connection, give it a name and associate it with a group. You can choose an existing group or create a new one. The branch office connection then uses that group's attributes, such as password management and encryption. You set the group's attributes on the Profiles→Groups pages.

The branch office connection then inherits the attributes of that group. You can associate multiple branch offices with the same group, thereby saving setup time and increasing management efficiency. For example, you might plan on creating several VPN connections from various remote sales offices into your enterprise headquarters. In this case, you create all of the connections in the same group so they all have the same attributes, such as hours of access, encryption method, and password management.

Use the main Branch Office screen to create new branch office connections and to edit or delete existing connections. You can also add or edit the group that is associated with your branch office connection.

To define a new branch office connection:

1. Go to the Add Group screen to create a new group. You can associate the new group with a branch office connection. You can add a new group that you want associated with the branch office connection. The new group inherits the attributes (for example, Access Hours) of its parent group, which are then used by the branch office connection.
2. The list shows all the branch office groups on the switch. Select the group whose attributes you want the new group to inherit. Refer to the Profiles→Groups→Edit→Connectivity screen for additional details on the hierarchical structure of group attributes.
3. Enter the group name that identifies the new group that you want associated with the branch office connection. The group name can be a maximum of 64 characters (spaces are permitted).

Use the initial Define Connection screen to create a new branch office connection. To define a new connection, you first enter its name, then associate the new connection with a group. The connection then inherits the attributes of the group.

To create a new branch office connection:

1. Enter the name of the new branch office connection. The name can be a maximum of 64 characters (spaces are permitted).
2. The list contains the names of all groups that have been set up on this switch. Select the group that you want to use for the branch office connection. The group is a child of its associated parent group and inherits the parent group's

network access attributes (refer to the Profiles→Groups→Edit→Connectivity screen for details). You can later modify the new group's inherited options.

3. Put a check in the Control Tunnel check box to specify that this branch office connection is for a control tunnel. When you have created and saved a new branch office connection, you cannot later change the control tunnel specification for that connection. For example, if you create a branch office connection that is not a control tunnel, you cannot later change it to be a control tunnel. In this case, you must set up a new connection with the same information as the original connection and create it as a control tunnel.

The Define Connection and the Edit Connection screens allow you to enable the Branch Office feature and to specify routing and networking information, local and remote identification, and authentication attributes for the branch office connection.

To enable a branch office:

1. Go to Profiles→Branch Office→Edit Connection. The Connection Name field displays the name of the connection.
2. Select one of the following connection types:
   - Peer-to-peer, where the connection type is the traditional branch office tunnel and either side can initiate traffic.
   - Initiator, where with Asymmetric Branch Office Tunnels (ABOT), one side must be configured as the initiator and the other as the responder. Only the Initiator can bring up the tunnel. When the connection type is set to initiator, there is no need to define a local endpoint. You should only configure an IPSec tunnel type. IPSec authentication requires an initiator ID.
   - Responder, where neither local or remote endpoints are required. You must configure IPSec authentication to specify the same initiator ID as in the associated initiator branch office tunnel.
3. The Group Name list contains the names of all groups that have been set up on this switch. Select the group that you want to use for the branch office connection. This group is a child of its associated parent group and inherits the parent group's network access attributes. You can later modify the new group's inherited options.
4. Under State, select either Disabled or Enabled.

5. Under Configuration, the Endpoint Address fields are used to specify the public interface IP addresses of the switches that form the branch office connection. The Local Endpoint address is the public interface IP address of the switch whose management interface you are using. The Remote Endpoint address is the public interface IP address of the switch that forms the opposite end of the branch office connection.
6. Click the drop-down list and choose the filter that you want this branch office connection to use. The default is permit all. You can specify one filter. Packet filtering controls the types of access allowed for users of this branch connection. Filters are based on various parameters, including protocol ID, direction, IP addresses, source, port, and TCP connection establishment. Filters are defined on the Profiles→Filters screen.
7. Under Configure Routing, click on the IP button to specify the type of routing to use for traffic going through the branch office connection.
   - If you choose static routing, you must manually specify the accessible networks (the private internal networks behind a switch that can be accessed via the branch office connection).
   - If you choose RIP, the routing protocol automatically determines the accessible networks based on information that is entered on the System→LAN Interfaces→Edit IP Address screen.
8. Under Tunnel Type, select either IPSec, PPTP, or L2TP. The default type is IPSec.



Note: If you change the tunnel type, the fields in the authentication portion of this screen change to reflect the different configuration requirements for the new tunnel type.

9. Under IPSec Authentication, the fields that appear depend on whether you are using an IPSec, PPTP, or L2TP tunnel type.



Note: If you create a branch office connection using any IPSec certificate and you choose IP address as the alternate name, you must use the IP address of the public interface that is on the branch office end of the connection.

For specific information on configuring authentication, see Chapter 9, “Using Certificates.”

Configuring the remote branch office connection

At the Define Connection screen, enter the required information for the branch office connection; for example, static routing and the IPSec tunnel type.

1. Enable the branch office connection by selecting the check box.
2. Click on the list and choose the routing type that you want to use for your branch office connection. The screen changes to show the appropriate routing fields.
3. Specify the addresses of the public interfaces of the two switches forming the connection.
4. Specify the private subnets that can be reached through the tunnel connections of this branch office connection.
   a. For the local network, click on the list and select the address. The local network is the remote switch's private network, which is specified on the Profiles→Networks screen.
   b. For the remote networks, click on the Add button to select the Add Networks screen. Then add the remote subnetworks' IP Address and subnet mask. The remote networks are the local switch's private networks. When you are done adding the networks, click on OK to return to the Define Connection configuration page.
5. Click on the list and then select the No NAT Translation option.
6. To select the filter for the connection, click on the list, then select the filter you want to use (permit all for this example).
7. Select the tunnel type (IPSec for this example). The screen changes to show IPSec fields.
8. Set up the Authentication method for the connection; for example, a text pre-shared key. Enter the key, then retype it in the Confirm Text String field to verify that you entered it correctly.
9. Click on the OK button to save the configuration settings for the remote switch.

You have completed the branch office configuration.

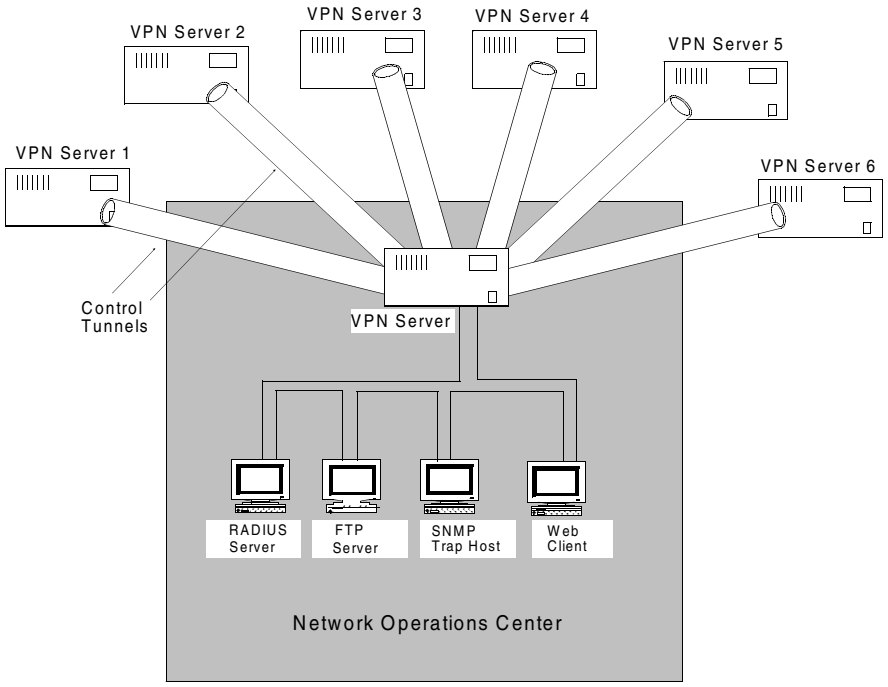
Control tunnels

Control tunnels are special tunnels that allow you to securely manage a Contivity VPN Switch over the Internet. The primary reasons for creating control tunnels are secure management and network data integrity. Control tunnels give you secure access to a customer's remote switch so that you can manage it over a network. Control tunnels also guarantee that no data from the network behind that customer's switch could ever "leak" through the control tunnel for access by the people on the network that are managing the switch.

You can allow access to FTP, DHCP, RADIUS, and DNS servers from the switch through the control tunnel. Control tunnels allow you to easily configure secure tunnels to any switch you want to manage anywhere in the world. This allows you to set up an encrypted tunnel to a customer's switch. Through that tunnel you can perform all the necessary management tasks, such as HTTP, FTP, SNMP, and Telnet.

Figure 25 shows a sample branch office control tunnel environment where a central VPN server is able to control several VPN devices and configure various services, such as RADIUS, FTP backup, SNMP Traps through Web client management, or Telnet.

Figure 25 Branch office control tunnel

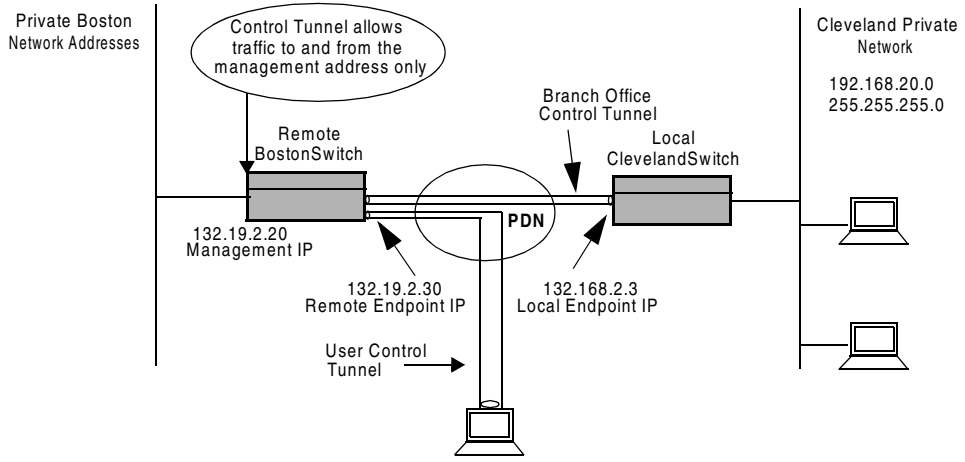


Control tunnel types

There are two types of control tunnels: a branch office control tunnel and a user control tunnel. With both tunnel types, you can establish a secure IPSec tunnel to a system that you want to manage. The traffic inside the tunnels is limited to the switch's management IP address only, which is unique to control tunnels.

Figure 26 shows a special branch office control tunnel from a network operations center in Cleveland, and also a user control tunnel.

Figure 26 Sample control tunnel environment



Branch office control tunnels allow anyone on the configured network to communicate with the switch being managed. This allows a switch to communicate with various systems within a company's network operations center or corporate headquarters (the Cleveland private network).

A user control tunnel allows a Contivity VPN Client to communicate with a switch that is being managed. This allows network management personnel from anywhere in the world access to the management tasks.

If you work at a network operations center (NOC) in Cleveland and your job is to manage a customer's switch that is located in Boston, you would want to use control tunnels. On one end of the control tunnel (the switch under management), access is always restricted to the management address only. Access to the Boston switch is limited. The Cleveland end of the tunnel could allow access to its entire private network. This allows multiple systems in your Cleveland NOC to communicate with the management address only of the Boston switch; or for the Boston switch to use remote servers (FTP, DHCP, RADIUS, and DNS servers) on the Cleveland private network.

In this environment, the remote Boston switch is a control tunnel to the local Cleveland switch. From any system on the Cleveland network, you can access the management address for the Boston switch. This allows systems on the Cleveland network to initiate management operations on the Boston switch, such as HTTP, FTP, and Telnet. Yet because it is a control tunnel, users on the Cleveland private networks cannot exchange packets with users on the private Boston Network.

Additionally, a user control tunnel is configured so that a remote user can establish a control tunnel when using the IPSec client. You create this user account with password authentication in the Control Tunnels group using the serial port.

Restricted mode

The Restricted mode feature prevents management of the switch except through a control tunnel. This limits the scope of management to someone who has the proper credentials both to set up the tunnel (if it is an end user) and to log in as an administrator (administrative access privileges). Having the proper access privileges acts as a level of security. Additionally, since in restricted mode you are forced to manage the switch through a tunnel, you are guaranteeing data protection through encryption.

You enable Restricted mode through the Serial Interface menu or the command line interface available through Telnet. In Restricted mode, you can perform the key management functions through the control tunnel, including HTTP, FTP, SNMP, and Telnet. All other attempts to perform these actions outside of the control tunnel will fail.

You cannot enter Restricted mode unless there is an active control tunnel. This ensures there is a mechanism to manage the switch in restricted mode.

Nailed-up control tunnels

You may want to have some control tunnels remain up even when there is no traffic traversing the control tunnel. This is generally the case for branch office versus end user control tunnels. You can do this by creating a script to continuously ping the switch under management. You continuously send ping packets to the management IP address of the switch on the customer premise

through the control tunnel from a host at the network operations center. The pings must occur at an interval that is less than the Idle Timeout value. These pings act as a liveliness detection and perform keepalive signals for the end connection, and report to the sender that the packet was received or that there was no response.



Note: If you change any settings to the branch office connection when using nailed up tunnels, you must bring down the tunnel for the changes to take effect.

To create a nailed-up control tunnel:

1. Go to Profiles→Branch Office screen and click on the Edit button next to the group that you want to have nailed up.
2. On the Edit Group screen, click on the Configure button under the Connectivity section.
3. On the Connectivity screen, when you click on the Configure button next to the Nailed Up field, a drop-down list gives you the option to select Enabled or Disabled.
4. Click on Enabled. Disabled is the default.
5. Click on OK.

Creating control tunnels

To create a special branch office connection, you must create a control tunnel definition on the remote customer switch. There are two methods you can use to create control tunnels:

- On the switch's GUI, use the Profiles→Branch Office screens to create the branch office connection and specify that it is for a control tunnel.
- For the command line interface, use the switch's command line interface (described in *Reference for the Contivity VPN Switch Command Line Interface*) to set up the connection as a control tunnel. This procedure is described in the following example.

To configure the local switch:

1. Initiate a Telnet session to the customer's switch.

2 Enter the appropriate control create string, following the required control create parameters already described. A sample string follows:

```
control create boston bostoncleveland 132.19.2.20 132.19.2.30
192.168.2.3 192.168.20.0 255.255.255.0
```

A special control tunnel filter is used by default with control tunnels to maximize security:

```
Management Only
```

3 To view Help, enter control help create.

These are control create parameters that you must enter:

```
CONTROL CREATE <name> <password> <MGMT/Local_P> <Local_endpoint>
<Remote_endpoint> <Remote_Subnet_Address> <Remote_Subnet_Mask>
```

If you are using the local switch's current Management IP address (132.19.2.20) rather than a substitute, then the network address translation feature is unnecessary. Otherwise, enable control on the remote switch, and enter the control address through the command line interface.

If you enter an address other than the management IP address (MGMT), NAT creates a NAT set with a static rule. The NAT set is called Control plus the name of the connection (for example, Control Boston). This also creates a network definition that is named Control plus the name of the connection (for example, Control Boston). The network definition contains the NAT management address. In this case, the branch office connection automatically fills in the correct NAT rule and accessible network. When using the control create commands, you must enter them in a complete string.

The switch that you are controlling sets a management only filter by default. This restricts access to the management IP address only. You can verify your control tunnel connection from the Profiles→Branch Office: Control Tunnels connection field.

After you configure the local switch, you must configure the switch located at the remote site. Complete the following steps to define the branch office connection for the remote switch.

To create a new group:

1 Access the Profiles→Branch Office screen.

2. Click on the Add Group button.
3. From the list, select the parent group whose attributes the new group inherits; for example, /Base.
4. Enter the name for the new group. The sample configuration uses the name Cleveland. Click on OK to save the settings and return to the Profiles→Branch Office screen. You can use the Edit button next to the group name on the Profiles→Branch Office screen to review or modify the group's attributes.

To define the branch office connection, specify a name for the connection, then associate the connection with a group:

1. On the Profiles→Branch Office screen, click on Define Branch Office Connection.
2. On the initial Define Connection screen, name the connection, for example vpn_to_boston which is specified on the Profiles→Networks screen.
3. Click on the name in the list to associate the connection with the group, for example /Base/cleveland.
4. Click on OK. The configuration screen for the new branch office connection appears.

On the Define Connection configuration screen, you enter required configuration information for the local branch office connection, for example, static routing and the IPSec tunnel type.

1. Enable the branch office connection by selecting the check box.
2. Click the list and select the routing type that you want to use for your branch office connection, for example, static routing. You *cannot* use control tunnels with RIP.
3. Specify the addresses of the public interfaces of the two switches forming the connection.
   a. For the local endpoint address, click on the list and select the address of the local switch (for example, 132.168.2.3).
   b. In the remote endpoint address field, enter the address of the remote switch (for example, 132.19.2.30) that you want to form the opposite end of the branch office connection.

4. Specify the private subnetworks that can be reached through the tunnel connections of this branch office connection.
   a. Click on the list and select the network address(es), for example, the cleveland entry previously created. The local network is the switch's private network, which is specified on the Profiles→Networks screen.
   b. Click on Add in the Remote Endpoint field, and enter the Management IP address of the local switch and its mask (for example, a host mask of 255.255.255.255).
5. Click on the list, then select the No NAT Translation option.
6. Click on the desired filter for the connection.
7. Select the IPSec tunnel type.
8. Set up the authentication method for the connection, for example, text pre-shared key. Enter the key (for example, bostoncleveland), then retype it in the Confirm Text String field.
9. Click OK to save the configuration settings.

Next, you should verify your branch office connection by sending ping packets to the management IP address of the local switch. Or, you can establish a Web connection to the local switch and attempt to configure it.

Creating a user control tunnel from the serial interface

You can create a user tunnel using the serial interface. Control tunnels allow the management of the switch without access to anything on the network other than the management IP address. This is used to force management through an encrypted tunnel and restricts access to the local resource such as outsourcing management of a switch. You create the control tunnel user in the group /Base/ Control Tunnels.

1. Open a connection to the serial interface.
2. Type **5** to Create a user management tunnel.
3. Enter a user name and password.
4. When prompted for an IP address, enter an address that would be a static IP address for the control tunnel or leave it blank to use the IP Addr Pool.

This creates a group called Control and places this user within this group.

IPSec client features

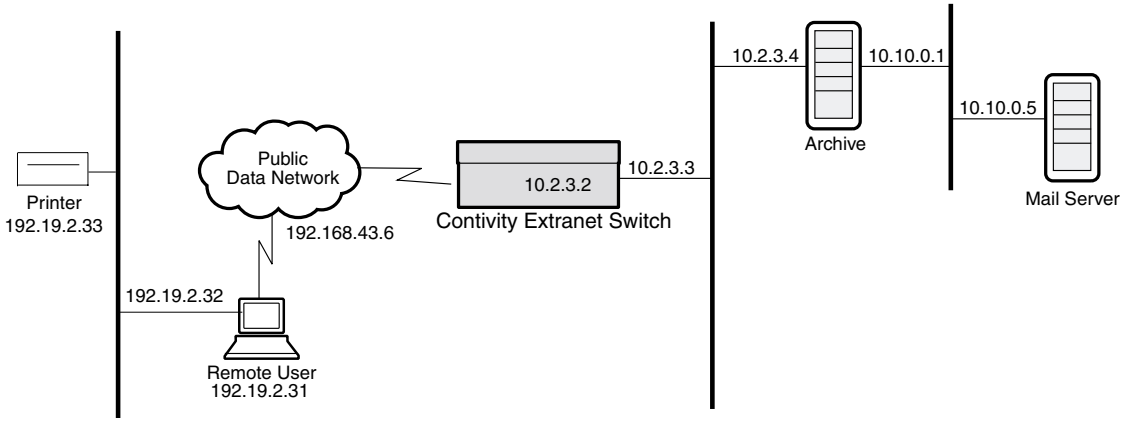
The switch supports the following IPSec client features:

- Split tunneling
- Third-party IPSec clients
- Forced logout
- Client fail-over
- Client auto connect
- Banner
- Password storage
- Client screen saver
- Domain name
- Client policy

Split tunneling

All IPSec client traffic is tunneled through the switch by default. Split tunneling allows you to configure specific network routes that are downloaded to the client. Only these network routes are then tunneled; any other traffic goes to the local PC interface. Split tunneling allows you to print locally, for example, even while you are tunneled into the switch. Figure 27 shows a sample split tunneling environment.

Figure 27 Sample split tunneling environment



The previous figure shows split tunneling enabled and split tunnel network IP addresses 10.2.3.4 and 10.10.0.5. When a client establishes an IPSec tunnel, these addresses are loaded into the client application.

The remote user, for example, then downloads his email from the mail server at 10.10.0.5, and downloads a document from the Archive at 10.2.3.4. Next, without exiting the tunnel, the remote user can print the document through the PC's local network interface 192.19.2.32 to the printer at 192.19.2.33. You can enable split tunneling through the Profiles→Groups→IPSec→Edit screen split tunneling field.

You designate which network routes to tunnel through the switch from the Profile→Networks screen. Next, you associate specific network routes to specific groups through the Profiles→Groups→IPSec→Edit screen by configuring the Split Tunnel Networks field.

The switch takes precautions against violators potentially hacking tunneled information when the switch is operating in split tunneling mode. The primary precaution is to drop packets that do not have the IP address that is assigned to the tunnel connection as its source address. For example, if you have a PPP dial-up connection to the Internet with an IP address of 192.168.21.3, and then you set up a tunneled connection to a switch and you are assigned a tunnel IP address of 192.192.192.192, then any packets that attempt to pass through the tunnel connection with a source IP address of 192.168.21.3 (or any address other than 192.192.192.192) are dropped.

Furthermore, you can enable filters on the switch to limit the protocol types that can pass through a tunneled connection.

The Client Selection feature enables you to configure your switch to accept tunnel connections from third-party clients, in addition to the Nortel Networks Contivity VPN Client.

To completely eliminate security risks, you should not use the split tunneling feature.

Third-party IPSec clients

The client selection feature enables you to configure your switch to accept tunnel connections from third-party clients, in addition to the Contivity VPN Client.

The Client Selection feature provides more flexibility and mobility than was previously available to remote users who want to connect to your switch using a client other than the Contivity VPN Client. The alternate method of connecting third-party clients requires you to set up a branch office connection and configure the remote client's IP address as the connection's remote gateway address. This branch office method binds a client machine to a fixed IP address. This can be limiting if a user needs to be able to create tunnels from multiple systems, for example, a work desktop system and a mobile laptop.

With the client selection feature, you establish an account for a remote user, rather than for a remote machine. You set up the account within the realm of remote access users, as has always been done for the Contivity VPN Client users. This gives the remote user the freedom to create tunnels to your switch from different machines, and from different locations.

When configuring for Contivity VPN Clients, the switch ignores the "Allow undefined networks for non-Contivity clients" field for clients that are not Contivity clients. The switch never allows Contivity VPN Clients to connect to undefined networks. All reachable networks must be defined on the Profiles→Networks screen.

When configuring for clients that are not Contivity VPN Clients, the fields that are preceded by an asterisk are not supported. You must select either the Split Tunneling or “Allow undefined networks field for non-Contivity clients” field for clients that are not Contivity VPN Clients. If you select both, the switch uses the Split Tunneling feature and ignores the “Allow undefined networks” selection.



Note: Nortel Networks recommends that you always specify Split Tunneling for groups used by clients other than Contivity VPN Clients. This ensures that your switch has control over the networks that the third-party client can access.

With Split Tunneling enabled, the third-party clients can only connect to networks that are listed as split tunnel networks on your switch. If Split Tunneling is disabled and “Allow undefined networks for non-Contivity VPN Clients” is enabled, the clients can connect to all internal networks.

The switch supports both preshared key and RSA digital signature authentication methods. For clients that are not Contivity VPN Clients, you must specify at least one of these authentication methods on the Services→IPSec screen.



Note: You must ensure that your remote third-party client uses the same Internet Key Exchange (IKE) Phase 1 mode that your switch uses. For Preshared Key authentication, the switch uses IKE Aggressive mode. If the client only supports IKE Main mode, it must be configured as a branch office due to the IKE restrictions. For RSA Digital Signature authentication, the switch uses IKE Main mode.

RADIUS authentication is not supported. Also, you can configure a static address for the tunnel from a client other than a Contivity VPN Client, or you can allow the client to use its own IP address as the address used within the tunnel. You can configure the client selection in the Profiles→Groups→Edit→IPSec screen.

Forced logout

For IPSec tunneling, you can specify a time after which all active users are automatically logged off. The default is 0, which means the option is turned off. The possible range is 00:00:01 to 23:59:59.

Client failover

Client fail-over uses small packets to check and maintain, or keep alive, the connection between the client and the switch.

Use the Contivity VPN Client to disable keepalives between the switch and the client. This option allows you to disable keepalives when tunneling over an ISDN link, since the link is not always active. If an idle time-out has been set on the switch, and keep alives have been disabled on the client, the client might not receive notice that the connection has been closed (due to the Idle Time-out), when the physical ISDN connection is not active.



Note: If the idle time-out on the switch logs off the client, and the client has client fail-over configured on the Services→IPSec screen, that client then fails over to the defined failover server, rather than being disconnected as desired.

When the Keep Alive parameter is disabled on the client it prevents the switch and client from exchanging keep alive messages. Therefore, if the connection is lost, the switch does not realize that the client is no longer connected until the idle time is reached. If the idle time-out can be set to Never, the resulting connection could remain established for a long time, which wastes switch resources.

If the number of Logins is set to 1, which is the default, the client cannot reconnect until the rekey happens, which by default is in 8 hours. If the user has the Disable Keep Alive parameter set on the client, and the connection goes down, the user could be prohibited from reconnecting for 8 hours or more, depending on the rekey value.

Also, do not set the idle time-out to 0. If you lose the connection in this situation, you must delete the session from the switch to reconnect.

Static tunnel failover

When a static branch office tunnel fails, all packets flowing through the tunnel are dropped. The static tunnel failover feature provides a means to detect and recover from these failures. Static tunnel failover interacts with static route API directly to remove and add static routes associated with a remote network. When a tunneling protocol detects a network failure, the static tunnel API removes static routes associated with the remote network from the route table manager.



Note: When setting up static tunnel failover, you need to configure the primary tunnel as nailed up from the Profiles→Branch Office→Edit→Connectivity screen. It should also have less cost than any secondary tunnels.

Client auto connect

The client auto connect feature enables remote clients to connect their IPSec tunnel sessions in a single step. This is similar to the way Microsoft Dial-Up Networking automatically connects to an ISP when a Web browser is launched. With auto connect, client users simply click on the desired destination, for example, a Web page on the private internal network. This first starts their dialup connection, then makes the tunnel connection to the switch, and finally makes the connection to the requested destination. What has, in the past, taken three distinct user operations is now accomplished by a single action.

The client auto connect settings specify those network connections that trigger the client's autoconnect feature. For example, you can specify that whenever a remote client attempts to connect to a site in the xyz.com domain, the client auto connect feature is started.

You must make sure that the switch is configured to allow connections to potential destinations. If the switch is not configured properly, the remote user might be able to make the connection to the switch, but cannot access the requested destination. For example, the switch's filters might be set up to deny access to finance.xyz.com, while the client auto connect is configured to start when

connections to the xyz.com domain are received. With this configuration, when a remote client tries to access finance.xyz.com, their connection to their ISP and then to the switch is automatically started. However, because of the filters, access to finance.xyz.com is denied.



Note: After you enable client auto connect, you must reboot the PC on which the client is running and manually make sure the client can connect to the switch.

When the client successfully connects to the switch, the switch downloads the list of networks and domains that trigger the autoconnect feature. This list, which is stored in the client's registry, is used to determine whether a tunnel connection should automatically be started when one is not already active.

The following client features apply only to the Contivity VPN Client:

- Password storage on client
- Client policy
- Client screen saver
- Domain name

Banner

You can customize an enterprise login banner for the Contivity VPN Client by entering text into the space provided. This banner appears at the top of the IPSec client upon login.

Password storage

You can allow client systems to save the login password in this password list, or you can require that a remote user enter the password with each request for authentication and access to an IPSec tunnel. Click on Enable to allow client systems to save the login password.



Note: When using certificates, saving the password on the client is not allowed.

Client screen saver

Setting this security feature forces the client to use a password in association with the screen saver. When it is enabled, if the user leaves the system while connected to a tunnel, the system then gets locked out of the tunnel when the screen saver kicks in.

Domain name

This setting enables you to specify the name of the domain used while an IPSec tunnel is connected. Specifying the domain name in this field ensures that domain lookup operations point to the correct domain. This is particularly important for clients that use Microsoft Outlook or Exchange, to ensure that the mail server is mapped to the correct domain.

When a tunnel is connected, the remote client's registry is updated to use the specified domain. When the client disconnects the tunnel, the remote client's original domain is again used.

Client policy

Client policy helps prevent potential security violations that could occur when you are using the split tunneling feature. Split tunneling allows client data to travel either through a tunnel to the enterprise network or directly to the Internet.

Configuring L2TP over IPSec

Windows 2000 supports only L2TP with IPSec transport mode for remote access or branch office. (L2TP cannot be used without IPSec.) It supports only RSA Digital Certificates for IPSec transport authentication with the switch. Windows 2000 Professional Server or Advanced Server can act as a Windows 2000 L2TP/IPSec client to a switch server.

Switch configuration

1 Configure an L2TP user account on the switch through the Profiles→Users page and enter an L2TP user ID and password. As with all remote access

accounts, you can also enter a static IP address on this page or let the system choose one from an address pool. As an alternative, you can configure a RADIUS server where the L2TP user account resides.

2 Before doing any per-user configuration, the switch must be issued a certificate and must have the issuer's certificate installed.

a Generate a certificate request from the System→Certificates page. This request can be transferred to a CA server that issues the certificate. The certificate can then be installed from the same page.

b You must also install the CA server's certificate on the switch through the System→Certificates page. If the Windows 2000 certificate is issued by a different CA, you must also install its certificate.

3 Configure an IPSec transport account on the switch in one of three ways:

- Configure the Subject DN or Alternate Subject Name of the Windows 2000 certificate on the same page as the L2TP user account. Pull down the Valid Issuer Certificate Authority and select the CA who issued the Windows 2000 certificate. Pull down the Server Certificate and choose the switch's certificate to be returned to Windows 2000. Windows 2000 checks the issuer's certificate also, so choose a certificate issued by a CA that Windows 2000 knows about. You may also check the Require Own IPSec Credentials now if you want to ensure that this L2TP user always uses this IPSec transport account.
- Go to the System→Certificates page and select the Enable Allow All Feature check box. For the CA that issued the Windows 2000 certificate, select the Allow All Enabled check box. Select a user group from the Default Group pull-down. Be sure the user group selected has Allow IPSec Transport enabled and configured (not inherited) in its IPSec group properties. This configuration is very useful when L2TP user accounts are in RADIUS, since no L2TP or IPSec transport information needs to be stored in the LDAP server per user.
- Create a separate user that contains the IPSec transport account. Set up the account as described in the first option. Do not check the Require Own IPSec Credentials for either this user or the L2TP user. This configuration is supported mainly for L2TP user accounts that are in RADIUS or compulsory tunneling (many L2TP users sharing an IPSec transport connection). Alternatively (for testing), you could install a single certificate on multiple Windows 2000 PCs, in which case they would be sharing a single IPSec transport account. Windows 2000 does not support compulsory tunneling.

4 Configure the L2TP profile for the user:

a At a minimum, you must set the desired minimum data protection level for the user. L2TP traffic arriving through an IPSec transport that does not meet this requirement is discarded. This is done in the L2TP properties of the group and therefore applies to all L2TP users under this group. No checking is done to determine whether the selection makes sense. For example, selecting 3DES as the minimum protection level implies that 3DES must be able to be negotiated with the Windows 2000 PC. To do this, DES must be enabled in the Services→IPSec page, must be enabled in the IPSec properties of the group containing the IPSec transport account, and must be configured on the Windows 2000 machine as an acceptable encryption type. Table 7 describes the mapping of minimum data protection levels.

Table 7 Mapping minimum data protection levels to encryption levels

Minimum data protection level	Encryption levels
Triple DES	ESP-Triple DES with SHA1 Integrity·ESP ESP-Triple DES with MD5 Integrity
56-bit DES	ESP-Triple DES with SHA1 Integrity ESP-Triple DES with MD5 Integrity ESP-56-bit DES with SHA1 Integrity ESP-56-bit DES with MD5 Integrity
40-bit DES	ESP-Triple DES with SHA1 Integrity ESP-Triple DES with MD5 Integrity ESP-56-bit DES with SHA1 Integrity ESP-56-bit DES with MD5 Integrity ESP-40-bit DES with SHA1 Integrity ESP-40-bit DES with MD5 Integrity

Minimum data protection level	Encryption levels
Authentication only	ESP-Triple DES with SHA1 Integrity ESP-Triple DES with MD5 Integrity ESP-56-bit DES with SHA1 Integrity ESP-56-bit DES with MD5 Integrity ESP-40-bit DES with MD5 Integrity ESP-40-bit DES with SHA1 Integrity ESP-NULL (Authentication Only) with SHA1 Integrity ESP-NULL (Authentication Only) with MD5 Integrity AH-Authentication Only (HMAC-SHA1) AH-Authentication Only (HMAC-MD5)
Not required	ESP-Triple DES with SHA1 Integrity ESP-Triple DES with MD5 Integrity ESP-56-bit DES with SHA1 Integrity ESP-56-bit DES with MD5 Integrity ESP-40-bit DES with SHA1 Integrity ESP-40-bit DES with MD5 Integrity ESP-NULL (Authentication Only) with SHA1 Integrity ESP-NULL (Authentication Only) with MD5 Integrity AH-Authentication Only (HMAC-SHA1) AH-Authentication Only (HMAC-MD5) Data is allowed through even if it does not come through an IPSec transport with this data protection level.

b If the Require Own IPSec Credentials check box is not selected on the L2TP user page, Require IPSec Credentials from Group must select a user group that contains a set of allowed IPSec transport accounts. These IPSec transport accounts may be contained at any level below this group. L2TP traffic that arrives through an IPSec transport not contained in this group is discarded.

c Turn on compression in the L2TP group properties if compression is desired. Compression for the PPP traffic is done if both the switch and Windows 2000 agree that compression is enabled. Windows 2000 does not support compression at the IPSec transport level.

d Authentication may be MSCHAPV1, MSCHAPV2, CHAP, or PAP. Of these, Windows 2000 prefers to perform MSCHAPV2 followed by MSCHAPV1 followed by CHAP followed by PAP. Windows 2000 does

not support L2TP encryption, so if MSCHAP V1 or V2 is enabled, make sure the Not Encrypted check box is also enabled.

5 Configure the IPSec transport profile by making sure Allow IPSec Transport is enabled in the group containing the IPSec transport account.

6 By default, Windows 2000 does not have Perfect Forward Secrecy (PFS) enabled. It is enabled by default on the switch. These two settings are not compatible and generate an appropriate error indicating such in the event log when a connection is attempted. To disable PFS on the switch, go to the IPSec properties of the IPSec transport group and disable PFS.

Table 8 User group properties used by IPSec transport connections

Property	Section	Note
Access hours	Connectivity	
Idle timeout	Connectivity	
Client selection	IPSec	Must have clients that are not Contivity VPN Clients enabled
Database authentication	IPSec	Only RSA digital signature is supported with Windows 2000
Encryption	IPSec	
IKE Encryption and Diffie-Hellman Group	IPSec	
Perfect forward secrecy	IPSec	By default, Windows 2000 has this disabled.
Forced Logoff	IPSec	
Compression	IPSec	Windows 2000 does not support IPSec compression. However, this setting does not affect connectivity.
Rekey timeout	IPSec	
Rekey data count	IPSec	
Allow IPSec transport	IPSec	

Windows 2000 configuration

Windows 2000 Professional, Server or Advanced Server may act as a Windows 2000 L2TP/IPSec client to a switch server. The steps for configuring the Windows 2000 side of this follow.

To install a certificate on the Windows 2000 PC using a Windows 2000 Microsoft CA, connect to a CA server and get a certificate. This involves pointing a browser at the CA server with the URL <*IP address*>/certsrv.

1. Choose Request a Certificate.
2. Choose Advanced request.
3. Submit a certificate request to this CA using a form.
4. On the form provide the identifying information. This becomes the subject DN in the certificate that is entered on the switch IPSec transport account.
5. Choose IPSec Certificate as the Intended Purpose.
6. Select Use local machine store under Key Options.
7. When the certificate has been issued at the CA server, return to the first page
8. Choose Check on a pending certificate.
9. Click Install this certificate. This installs the certificate in the local computer certificate store. To view this store, run the mmc command from the Start→Run prompt. Select Console→Add/Remove Snap-in. From the list of snap-ins, choose Certificates and select Computer account. At the console, expand Personal→Certificates under Certificates (Local Computer). The installed certificate should appear. Clicking on it brings up an information window that indicates its validity and that a private key exists for this certificate.

To install the CA server certificate for the Windows 2000:

1. If the switch's certificate was issued by a different CA, that server's certificate should also be installed. For the Microsoft CA, go back to the home page and select Retrieve the CA certificate or certificate revocation list.
2. Click on the Install this CA certification path. This installs the CA certificate as a trusted CA, which can be seen in mmc under Trusted Root Certification→Certificates.

To set up the dial-up networking entry to use L2TP over IPSec:

1. Click on My Computer and click on Network and Dial-up Connections. Click on Make New Connection.
2. Choose Connect to a private network through the Internet for the network connection type.
3. Enter the interface address of the switch server.
4. Edit the properties of this new connection and select the Networking tab. Change the Type of VPN server to L2TP.
5. Connect to the switch using the L2TP user ID and password entered on the switch. The certificate installed previously is automatically used to set up the IPSec transport connection.

Branch office

Windows 2000 Server or Advanced Server may act as a Windows 2000 L2TP/IPSec gateway to a switch. Both static routing and dynamic (RIP and OSPF) routing are possible through this branch connection.

To configure the switch:

1. Configure an L2TP branch connection on the switch. Go to Profiles→Branch Office.
2. Enter the IP address of the Windows 2000 server as the remote endpoint. Select L2TP as the tunnel type.
3. Choose MS-CHAPV2 unencrypted as the authentication type.
4. Enter a local UID for the switch.
5. Enter a peer UID for Windows 2000.
6. Enter a shared password.
7. Selected L2TP if you want compression. As with remote access, compression is not supported on Windows 2000 for the IPSec transport connection.
8. If you want L2TP tunnel authentication supported, you must provide an L2TP Access Concentrator definition. Windows 2000 does not support L2TP tunnel authentication.

9. Select the minimum data protection level. If you select anything other than Not Required, you must set up an IPSec account. Mappings of data protection levels to encryption levels are exactly as shown in Table 7.
10. As with remote access, the IPSec transport account must be set up. By default, Windows 2000 supports only certificate authentication, so a process exactly like that described for remote access must be performed. The CA Allow All authentication option is not available for branch office connections. The L2TP branch office must use the IPSec transport account specified in the connection if data protection is required.
11. You can set up routing as either static or dynamic.

To configure Windows 2000:

1. You must install a certificate for Windows 2000 and the CA certificates as described above.
2. Start the Routing and Remote Access administrative tool.
3. Right-click on Routing Interfaces and choose New Demand-dial Interface.
4. Choose the name of the branch connection. This name becomes the L2TP user ID of the switch. MSCHAPV2 is case sensitive for user IDs. To ensure interoperability with the switch, use lowercase user IDs.
5. Select Connect using VPN.
6. Select L2TP as the VPN type.
7. Enter the interface address of the switch.
8. Select Route IP packets on this interface and select Add a user account so a remote router can dial in.
9. Select a password for the switch L2TP user ID. If the switch initiates branch office connections to Windows 2000, this password must match that entered on the switch Branch Office Connection page. If not, then this password does not matter.
10. Choose the Windows 2000 L2TP user ID and the shared password. If the Windows 2000 initiates branch office connections to the switch, this password must match that entered on the switch Branch Office Connection page. The Domain field may be left blank.

11 The switch supports only MSCHAPV2 as a branch office L2TP authentication method, so be sure you enable this method in the properties (it is by default).

12 If you want static routes to demand dial on this connection, expand IP Routing→Static Routes and right-click on New Static Route. Select the interface just created and enter the subnet information. Be sure you enable Use this route to initiate demand-dial connections. Alternatively, you can dial the connection by right-clicking on it and selecting Connect.

Quality of Service

The switch supports two internal Quality of Service (QoS) mechanisms and can also participate in external network signaling to enhance performance. Forwarding priority allows for prioritized traffic, and Call admission priority allows you to reserve connection resources for high-priority users. In addition, external QoS using Resource ReSerVation Protocol (RSVP) signals the public network to reserve a portion of the network's bandwidth for a specific connection.

Classifiers

An MF Classifier can be defined for an interface (interface MF). The interface MF-Classifier is applied to routing traffic going through that interface.

To configure an MF classifier:

1 Go to the QoS→Classifiers screen. The Current Multi-Field (MF) Classifiers list includes all existing MF classifiers.

2 Select from the Current Multi-Field (MF) Classifiers and click Edit to edit the rules for that MF Classifier. The Edit Rule screen appears.

The Rules in Classifier list shows all of the rules that are applied to the MF Classifier.

The Available Rules list shows all of the existing rules. You can select rules from this list to move them into the Rules in Classifier list and apply them to the MF Classifier.

3 Click on a rule from the Available Rules list on the right of the screen, then click on the left arrow to add the rule. This adds the selected rule to the current

rules list. The new rule is added after the rule currently selected in the Rules in Classifier list.

4 Click on the Edit button to edit an existing rule. The Edit/Create Rules screen appears. The Classifier Rule for field shows the name of the rule.

5 Enter the source and destination addresses to limit the rule to acting on packets from and to these addresses. Source and destination are relative to the direction of the rule.

6 Click on the Modify button nest to the Source and Destination Address fields to edit either of these fields. The DiffServRules Definintion Address screen appears.

7 Select the appropriate protocol from the list. The default list of protocols include:

- ICMP - Internet Control Message Protocol is a Network protocol layer. The PING utility generates ICMP packets. PING is often used to see if a system's network is available.
- IP - Internet Protocol is a Network layer protocol in the TCP/IP stack that offers a connectionless internetwork service. IP packets that are encapsulated within other packets create “IP over IP.” Multicast IP packets (packets that have multicast destinations), carried between networks that support multicasting over intermediate networks that do not, are the most common implementation. Conferences and other services that are offered through Multicast Backbone (MBONE) are examples.
- TCP - Transmission Control Protocol is a transport layer protocol in the TCP/IP protocol stack. This is a connection-oriented protocol that provides reliable full-duplex data transmission. Web browsers using HTTP and FTP are examples.
- UDP - User Datagram Protocol is a transport layer protocol in the UDP/IP protocol stack. UDP is a connectionless service that exchanges datagrams without acknowledgment or delivery guarantees, and therefore requires that error handling and retransmissions are handled by other protocols. DNS and WINS are examples.

8 Click on the Modify button next to the Protocol filed to edit it.

9 Click on the Modify button to the right of the TCP/UDP Source and Destination Port fields to edit them. You can filter packets to or from the Source and Destination ports to permit or deny any packets transferred by the switch. The source of destination is relative to the direction of the rule.

10 Click on the Modify button to the right of the Current DSCP Value field to create and edit the DSCP value and mask. The DSCP value and mask assignments allow packets that are already marked to retain their settings or to be remarked based on their previous DSCP value.

11 Select the DSCP to be marked on the next meter, either EF (Expedited Forwarding) or an AF (Assured Forwarding) level, that this rule applies to data.

Bandwidth management

Bandwidth management capabilities let you manage the switch CPU and interface bandwidth resources to ensure that tunneled sessions get predictable and adequate levels of service. Bandwidth management allows you to configure the switch resources for users, branch offices, and interface-routed traffic. Bandwidth components keep track of and control the level of bandwidth being used on the physical interfaces and the tunnels.

Bandwidth management forces tunnels to conform to a set of rates. There are two rates (committed and excess) and excess action (mark or drop). Packets are given different drop preferences, depending on whether they are below committed rate (lowest drop preference), between committed and excess rate (higher drop preference), and above excess rate (highest drop preference if excess action is Mark). When there is congestion on the switch, packets are dropped according to their drop preference. When excess action is Drop, all the packets above excess action are dropped.

You can add admission control to guarantee that resources are available to support the committed bandwidth assigned to a user. This potentially denies a client access before the licensed limit of a switch is reached. The available bandwidth is based on the switch interface speed.

To configure bandwidth management:

1 Enable the advanced routing license on the Admin→Install Keys screen.

2 Go to the QOS→Bandwidth Rates screen to define the bandwidth settings. The maximum rate that you can create is 10 Mbps. This must be defined in bits per second (10 Mbps=10000000).

3 Go to the Profiles→Groups→Groups→Edit→Connectivity screen. In the User Bandwidth Policy section, define the committed and excess bandwidth rates.

4 Set Bandwidth Management to Enabled.

5 If desired, you can return to the Bandwidth Management to enable admission control.

6 Go to the QOS→Interfaces screen to set the over-subscription rate. Adjusting this ratio allows you to adjust for some users not using all of their allotted bandwidth simultaneously under normal circumstance. The default is 10:1.

Differentiated Services (DiffServ)

DiffServ settings classify and mark packets to receive specified per-hop forwarding behavior on each node along their path. Sophisticated classification, marking, policing, and shaping operations are implemented at network boundaries or hosts. Network resources are allocated to traffic streams by service provisioning policies that govern how traffic is marked and conditioned upon entry to a differentiated services-capable network, and how that traffic is forwarded within that network. Any DiffServ code points (DSCPs) not recognized are forwarded as if marked for the default behavior, Best Effort (BE).



Note: You can only have either DiffServ or Forwarding Priority active at any one time.

To configure DiffServ:

1 Go to QOS→Interfaces and click on the Configure button in the DiffServ Edge section.

2 In the Multi-Field Classifier State field, select Enabled or Disabled to enable or disable the application of MF Classifiers on this interface.

3 In the Ingress (inbound) field, select the MF Classifier that you want to apply when packets are coming in from this interface

4 In the Egress (Outbound) field, select from the list the MF Classifier that you want to apply when packets are going out of this interface.

5 In the Traffic Conditioning State field, select Enabled or Disabled to enable or disable traffic conditioning on this interface.

6 Enter a value, in bps, for Expedited Forwarding (EF) Shaping. Shaping is a process of delaying the packets in a stream in order to conform to a defined

traffic profile, in this case, the EF Shaping value. Nonconforming traffic is delayed, not dropped.

7 Traffic conditioning is the process of dropping and remarking a traffic stream in order to shape it into compliance with a traffic metering profile.

For Expedited Forwarding (EF) and Assured Forwarding 1, Assured Forwarding 4 (AF1-AF4), you can configure a Traffic Conditioning Meter (in bps). For EF, the rate is used as an average rate, though at times traffic can burst as much as twice the configured rate. Traffic below the rate is forwarded; traffic above the rate is dropped.

For AF1-AF4, any packets under the rate are marked as low drop precedence. Any packets under two times the configured rate are marked as medium drop precedence. Any packets above two times the configured rate are marked as high drop precedence.

To configure Egress (Outbound) Queueing Mode:

1 Got to the QOS→Interfaces screen and click on the Configure button in the DiffServ section.

2 Select the Queueing mode for this interface, either DiffServ Per-Hop behavior or Legacy Forwarding Priority. legacy Forwarding Priority provides backward compatibility for earlier versions of the switch.

Forwarding priority

Forwarding priority quality of service allows you to assign each user to one of four priority classes. Each class is guaranteed different maximum forwarding times between the interfaces of the switch. For example, high-priority traffic generated by the company CEO would be protected from high-bandwidth traffic generated by lower-priority users. Or, you might assign the sales team to Priority 1 to make sure that they could always place orders, especially during the quarter-end rush.

The technology that supports Forwarding priority is called *weighted fair queuing with random early detection* (RED). This queuing mechanism gives each of the four user classes (from 1-high to 4-low) a different weight in the amount of service time they receive by the packet-forwarding process. Each class, however, is guaranteed some level of service so that no traffic through the switch is ever

completely stalled. It is important to assign users to the four different class levels to make sure that they get the proper service and performance, especially during heavily congested times. QoS is only effective when all associated lines are capable of servicing the forwarding demands at the required speeds.

If a group profile has a Forwarding priority of 1 (highest), it has the highest possible bandwidth guarantee and the lowest level of latency. Packets sent by this group are transmitted immediately even if there is heavy traffic on the switch. Conversely, if a group profile has a Forwarding priority of 4 (lowest), it has the least amount of bandwidth allocated and possibly the highest level of latency. Therefore, fewer packets sent by this group are transmitted while there are higher-level priority packets to be sent when the switch traffic is heavy.

To illustrate how the Forwarding priority works, the example in Table 9 assumes heavy traffic and a queue of packets. Therefore, packets would be transmitted according to the approximate rates per “pass” that are cited in the table.

Table 9 Bandwidth allocation per priority level

Priority 1	Priority 2	Priority 3	Priority 4
≈60% pass	[a]25% pass	[a]10% pass	[a]5% pass

Of the total packets transmitted in a hypothetical pass, 60 percent would come from the Priority 1 queue; 25 percent from the Priority 2 queue; 10 percent from the Priority 3 queue; and 5 percent from the Priority 4 queue.

Call admission priority

Call admission priority quality of service allows you to assign each user group profile to one of four priority classes (from 1-high to 4-low) for call admission. The switch can reserve connections for each class of user, guaranteeing that a large number of low-priority users do not lock out the high-priority users. When the switch is servicing the maximum number of low-priority sessions, no further low-priority connections are accepted. Once a connection is accepted it is never dropped.

Since there is a maximum number of sessions supported on the switch, it is important to assign users to the proper Call admission priority classes. This ensures that connections are available to the appropriate users when there is heavy utilization. Although other callers may be permitted access to the switch, this access is proportional to the assigned priority level for their group.

By default, any call is admitted access for the first 50 percent of connections, regardless of the assigned Call admission priority. The next 25 percent of calls guarantee access to only Priority 1, 2, and 3 callers. The next 15 percent of calls guarantee access to only Priority 1 and 2 callers. For the final 10 percent of calls, only Priority 1 callers are guaranteed access.

For example, assuming a hypothetical maximum of 2000 sessions, Table 10 shows the connections available for each priority based on a percentage of the total capacity:

Table 10 Call admission priority

Capacity	Priority	Available connections
0 to 50%	All	1000
51 to 75%	1, 2, 3	500
76 to 90%	1, 2	300
91 to 100%	1	200

Table 11 shows the maximum number of connections available for each priority:

Table 11 Maximum connections per priority

Priority	Connections
1	2000
2	1800
3	1500
4	1000

RSVP

The switch supports Resource ReSerVation Protocol (RSVP) quality of service for the Internet. Successful external network-level quality of service requires the cooperation of all the devices on the network (between the user and either the access point to the private network or the ultimate destination host). Currently, RSVP is the best-defined technology for resource reservation. However, only a few service providers offer a service that uses RSVP.

The switch is ready to take part in the RSVP signaling that is available in some network backbones and that will increase in the future. The switch signals to the other devices on the public network and describes the level of bandwidth that is needed to ensure adequate performance. This amount of bandwidth is determined by both the data rate that the user has to the Internet, and by the data rate of the link between the Internet and the switch. This beginning stage of RSVP will be expanded over time to take advantage of advances in the technology. Meanwhile, you can build RSVP networks using the switch to gain experience in this environment.

The two key components of RSVP are:

- PATH messages, which are constant announcements by the host system or the switch that a certain amount of bandwidth must be kept available.
- RESV messages, which are responses from the client that it wants to reserve the requested bandwidth.

If the client responds to the PATH messages with RESV messages, then RSVP-ready routers attempt the resource reservation. These routers actually reserve the resources requested if they are RSVP-compliant.

Proxy ARP and tunnel-to-tunnel traffic

You can configure the Contivity VPN Switch to respond to ARP requests on any of the physical interfaces. The switch responds to the following types of routes:

- User tunnels are routes created for user tunnels. This entry is enabled by default and cannot be changed.

- Branch office tunnels are routes available through branch office connections. This option is disabled by default.
- Physical interfaces are routes available through physical interfaces. This option is disabled by default.

To configure Proxy ARP, go to the System→Forwarding screen. On this screen, you can configure the switch to enable or disable different types of tunnel to tunnel traffic, including end user-to-end user traffic, end user-to-branch office traffic, and branch office-to branch office traffic.

Chapter 5
Configuring Routing

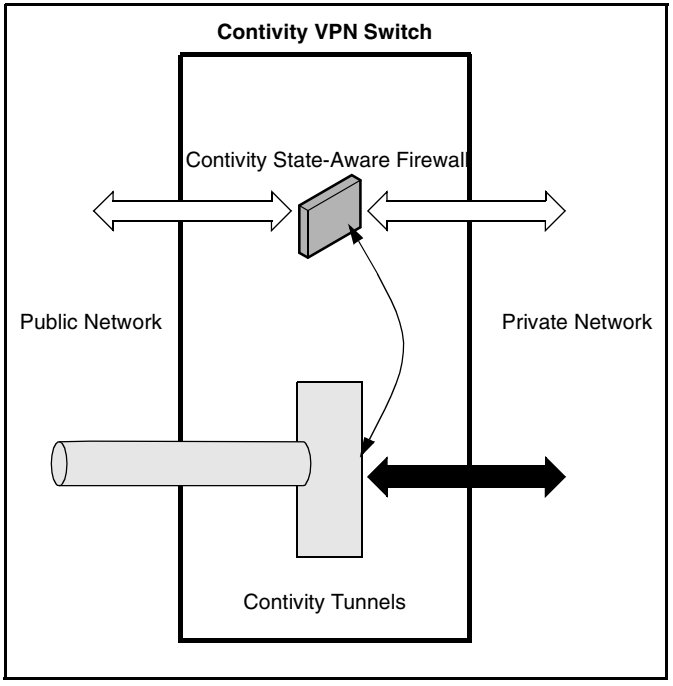
This chapter describes static and dynamic routing capabilities on the switch and the traffic patterns that are supported.

Routing overview

The switch's routing capabilities include allowing authorized tunneled traffic to securely flow in to and out from the corporation's private network. In addition, the switch can route traffic between two private interfaces, and between its public and private interfaces. The switch can also route traffic from its public interfaces to destinations on the Internet. As a result, the switch connects your organization to the Internet.

To permit traffic to flow between the public and private sides of your network, you must ensure that your private network is protected from unauthorized access from the public side. You must enable either the integrated Contivity State-Aware Firewall or the Contivity Tunnel Filter to provide this protection, using features such as packet filtering and antispoofing. The relationship between the integrated Contivity State-Aware Firewall and the switch is shown in Figure 28. The Contivity Tunnel Filter functions in a similar manner.

Figure 28 Traffic flow



The switch provides the following important routing capabilities:

- VPN routing deals with the following flow of data:
  - tunnel<-->tunnel
  - tunnel<-->private physical interfaces
- Enhanced routing deals with the following flow of traffic:
  - Private<-->private
  - Private<-->public
  - Public<-->public
  - Public<-->tunnel
- Services routing deals with data that supports the services provided by the switch.

For OSPF to function, you must install the Advanced Routing Key. The Firewall Key is optional and required only when you want to enable the redistribution capabilities of RIP and OSPF.

To install a software license key:

1. Go to the Admin→Install Keys screen.
2. Type the key that you obtained from Nortel Networks Customer Support in the box to the right of Advanced Routing.
3. Click on the Install button.

Integrated firewall and routing

The switch is a security device. Therefore, the routing configuration takes effect as it relates to the integrated firewall configuration of the switch. In all of the following sections, when there is a reference to integrated firewall, it means the Contivity Firewall option on the Services→Firewall screen. You use this option by selecting either Contivity State-Aware Firewall or Contivity Interface Filter. You must have a firewall license to use the Contivity State-Aware Firewall option; however, you do not need a firewall license if you use the Contivity Interface Filter option.

Running dynamic routing

A physical LAN interface in the switch can be configured either as a private or a public interface. At present, all WAN interfaces are configured only as public interfaces. The switch does not support dynamic routing protocols over a public interface. However, you can run dynamic routing protocols over private or tunneled interfaces.

To run dynamic routing protocols on a private LAN interface, you need to define the Contivity interface filter (with permit all) for that interface on the Systems→LAN page. Then from the Services→Firewall/Nat page, you need to enable Contivity Interface Filter. This setup does not need a Contivity State-Aware Firewall license. When you are running dynamic routing over a tunnel, the default tunnel filter, being a trusted interface, is permit all.



Note: The Advanced Routing License Key is required to enable features like OSPF and ECMP. Static routes, RIP, and Route redistribution do not need this license.

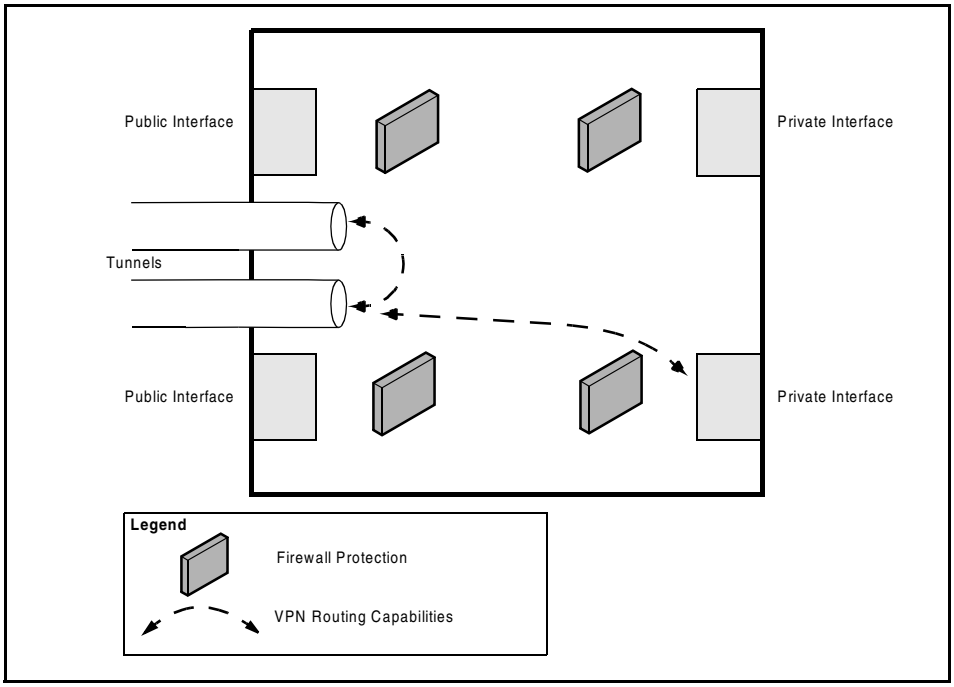
By configuring all the LAN interfaces as private interfaces, you can configure the switch simply as a router.

VPN routing

VPN routing refers to the routing that enables traffic to enter or exit the switch through a tunnel. With VPN routing, the switch supports tunnel-to-private interface traffic and tunnel-to-tunnel traffic, including branch office connections.

Tunneled users might connect to the switch using a VPN Client or through a branch office connection. This is the traditional VPN traffic pattern that has been supported by the switch since its initial release. Figure 29 illustrates the traffic that VPN routing supports.

Figure 29 VPN routing: allowed traffic patterns

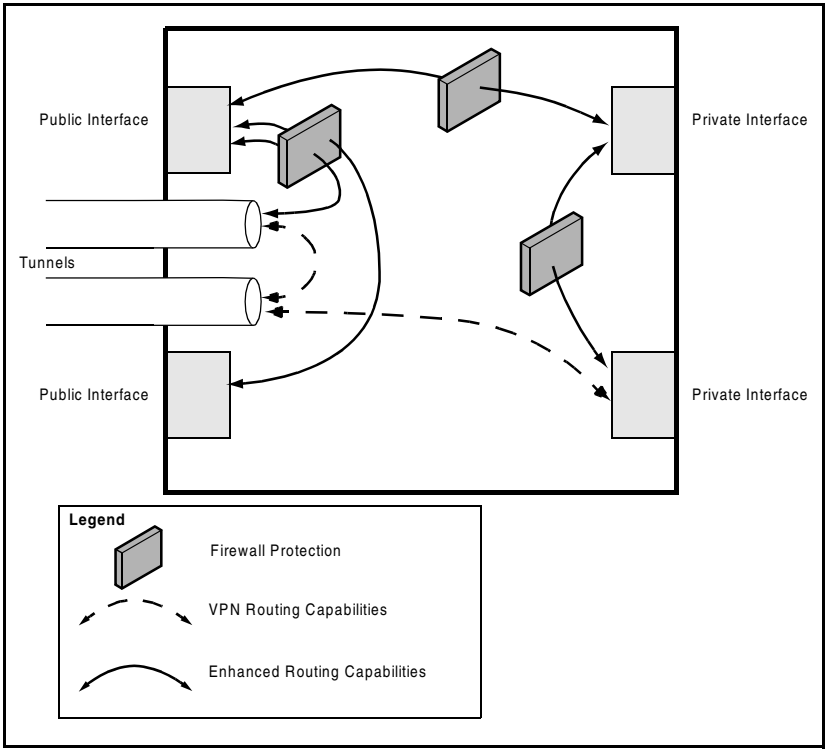


Enhanced routing

In addition to the traffic supported by the VPN routing feature, the switch's enhanced routing feature supports traffic patterns between the private network and the Internet, and enables traffic between a tunnel interface and a public interface.

Figure 30 illustrates the variety of routing possibilities supported by the switch. The new traffic type, which is provided by the enhanced routing feature, goes through the integrated firewall. Therefore, one of the switch's firewalls must be enabled to allow the switch to support this type of traffic. The integrated firewall support is an important component in enabling the switch to provide this comprehensive routing environment while maintaining the switch's emphasis on protecting the private network from unauthorized traffic.

Figure 30 Enhanced routing: allowed traffic patterns



Static routes

You can use static routes to set up routes between switches when you do not have any dynamic routing protocol, such as OSPF or RIP. Even if you do have dynamic routing protocols, you may want to use static routes because they provide stronger security. The switch supports multiple default and static routes.

You can manually configure static routes on the switch. Based on their states, they are added or removed from the Route Table Manager (RTM).

To add, edit, or delete static routes:

1. Check the Enabled box to enable static routes. When this check box is cleared, all of the static routes and default routes are disabled globally. Even if a static route is enabled, the route is not used because the static routes are globally disabled. When static routes are enabled, traffic flow depends on other configuration settings.
2. To add a public default route, click on the Add Public Route button. The Add Public Default Route screen appears and can:
   - **a** Click on enabled or disabled to select the Admin State.
   - **b** Type in the cost to select the relative cost for the switch. You would use a lower cost number, such as 1, for the least expensive route. When there are multiple default paths, the switch chooses the route with the least cost as the preferred route. The default is 10.
   - **c** Enter the IP address in the Gateway Address field.
   - **d** Click on the appropriate buttons to edit or delete default routes.
   - **e** Click on OK.
3. To add a private default route, click on the Add Private Route button. The Add Private Default Route screen appears and can:
   - **a** Click on enabled or disabled to select the Admin State.
   - **b** Type in the cost to select the relative cost for the switch. You would use a lower cost number, such as 1, for the least expensive route. When there are multiple default paths, the switch chooses the route with the least cost as the preferred route. The default is 10.
   - **c** Enter the IP address in the Gateway Address field.

d Click on OK.

4 Click on the add button to add static routes to the routing table. The Static Routes→Add screen appears. When a static route is added, the switch checks whether the next hop interface address belongs to an attached network. If it does not, the switch does not allow such static routes.

a Select enabled or disabled for the Admin state. The default is enabled.

b Select the relative cost for the switch. You would use a lower cost number (for example, 1)for the least expensive route. When there are multiple paths, the switch chooses the route with the least cost as the preferred route. The default is 10.

c Enter the IP address of the static route for the destination network.

d Enter the subnet mask for the static route for the destination network.

e Enter the gateway address where packets are routed for the destination network.

5 Click on the Show Branch Office Routes button to display the configured Branch Office tunnels that are set up as static routes. By default, a tunnel is configured as a static route between the tunnel end points.

RIP

The Routing Information Protocol (RIP) is a distance-vector routing protocol that enables routers to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring subnets and listen for RIP updates from the routers on those neighboring subnets. Routers use the information in the RIP updates to keep their internal routes current.

For RIP, the best path to a destination is the path with the fewest hops. RIP computes distance as a metric, usually the number of hops (or routers) from the origin subnet to the target subnet. RIP can handle a maximum of 15 hops.

RIP is one of the most common interior gateway protocols used in the Internet. RIP Version 2 is backward compatible with RIP Version 1 and corrected many of the RIP Version 1 shortcomings, such as subnet routing, authentication, support for multiple autonomous systems, and multicast support for route messages.

The switch supports RIP for routing traffic within the private network and between branch office connections. The switch supports both Version 1 and Version 2 of the RIP protocol. For additional information on RIP, refer to the RFCs, which are on the Internet Engineering Task Force (IETF) Web site at www.ietf.org.

- RFC 1058 – Routing Information Protocol: Describes the Routing Information Protocol (RIP), which is loosely based on the program "routed," distributed with the 4.3 Berkeley Software Distribution. The specifications in this RFC represent a combination of features taken from various implementations of this program.
- RFC 1721 – RIP Version 2 Protocol Analysis: Describes the key features of the RIP Version 2 protocol and the current implementation experience.
- RFC 1722 – RIP Version 2 Protocol Applicability Statement: Describes how RIP Version 2, which is an extension to RIP Version 1, may be useful within the Internet.
- RFC 1723 – RIP Version 2 Carrying Additional Information: Specifies an extension of the Routing Information Protocol (RIP) that expands the amount of useful information carried in RIP messages and that adds a measure of security.

To configure RIP global parameters, go to Routing→RIP. When globally enabled, RIP can be added to specific interfaces by going to the Routing→Interfaces screen.

Using RIP

The switch sends RIP broadcast/multicast messages at regular intervals. These messages contain information about routes that the switch can reach. Other routers on the network pick up these messages, and in turn update their routing tables and then send out route messages to their peer routers. The switch's RIP support allows you to:

- Enable/disable propagation of RIP messages from the switch's private and tunneled interface.
- Configure the private side of the switch to RIP passive mode or active mode. In passive mode the switch listens for RIP messages and updates its routing table, but it does not advertise RIP routes. In active mode the switch updates its routing tables and also advertises its routes to other routers.

The interface filters setting can affect the behavior of routing protocols. For example, RIP uses UDP as its transport mechanism and if the interface filters are set to deny UDP, then RIP advertisements are not sent or received.

Protecting against routing loops

A routing loop can occur when two or more routers continuously forward the same packet to each other until the hop count goes to infinity, the packets time-to-live counter expires, or the network goes down. Loops typically occur when a new router is added to the network or when a router in an existing network goes away and the remaining routers must recalculate routes. A loop detection protocol can prevent a routing loop and can speed up convergence while the situation corrects itself. The switch supports the following methods used by RIP for minimizing loops and for speeding up the convergence that is caused by the normal correction of a loop:

- Split horizon, where the switch does not send routes that it learns from a neighboring router back to that same neighbor.
- Split horizon with poison reverse, where the switch *does* send back the routes that it learns from a neighboring router, but it sets the metrics for the connection to infinity.
- Triggered updates, where an update is sent almost immediately after a routing change has been made on the switch. This is in contrast to the default RIP method, in which routes are updated at regular intervals.

Configuring RIP

To configure RIP:

1. Check the Enabled check box to enable RIP on the switch. By default RIP is globally disabled. If RIP is disabled on this screen, the switch does not process any RIP requests. After you enable RIP on this screen, you must also enable it on the Routing→Interfaces screen.
2. In the Update Timer edit box, enter the amount of time in seconds that you want RIP to update the routes. The default is 30 seconds. The range of values that you can specify is from 5 seconds to 65535 seconds. The hold down timer is six times the update timer. Routes are invalid after the time set on the hold down timer.

3 Configured Interfaces lists the IP address of the RIP interface and whether the the state of the RIP protocol is enabled or disabled on this interface.

4 Click on the Statistics button to display statistics about the RIP protocol in the switch. The following table describes the fields on the RIP Statistics screen.

Table 12 RIP Statistics screen

Column	Description
Global RIP Status	Enabled or disabled
Update interval	Interval in seconds
Trusted Neighbor	Enabled or disabled
Rip Domain	Set or reset
Triggered Update	Set or reset
Route Change	Number of routes changed
Query	Number of queries sent

5 Click on the Database button to display information for all of the RIP interfaces.The following table describes the fields on the RIP Database screen.

Table 13 RIP Database screen

Column	Description
Circuit	Circuit ID
Address	IP address
Mask	Network mask of IP address
Owner	Protocol
Cost	Import cost of RIP routes
Metric	Export metric of RIP routes
Gw	Gateway IP address

Click on the Interfaces button to display information for all of the RIP interfaces, including tunnels that are running RIP. The following table describes the fields on the RIP Interfaces screen.

Table 14 RIP Interfaces screen

Column	Description
Ip	RIP IP address
Subnet	Network mask of IP address
RipEnabled	Whether RIP is enabled or disabled
IntfState	Whether up or down
Auth	Authentication type
Type	Interface type
Cid	Circuit ID
RxMode	RIP receive version supported
TxMode	RIP transmit version supported
PoisonRev	Whether enabled or disabled
ImpDRoute	Whether enabled or disabled
ExpTSMetric	Disabled or metric (1-15) export tunnel static route
ExpSMetric	Disabled or metric (1-15) export static route
ExpDMetric	Disabled or metric (1-15) export default route
ExpOspfMetric	Disabled or metric (1-15 export OSPF route

OSPF

OSPF is a link-state routing protocol that maintains a database from which a routing table is constructed from the shortest path, using a minimum of routing protocol traffic. It provides a high-functionality open protocol that allows multiple-vendor networks to communicate using the TCP/IP protocol family. Some of the benefits of OSPF are:

- Fast convergence
- Variable-length subnet masks (VLSM)
- Hierarchical segmentation

- Area routing to provide additional routing protection and a reduction in routing protocol traffic
- Authentication
- Auto-virtual link support
- Equal Cost MultiPath (ECMP) support

Using OSPF

The switch sends OSPF broadcast/multicast messages at regular intervals. These messages contain information about routes that the switch can reach. Other routers on the network pick up these messages, and in turn update their routing tables and then send out route messages to their peer routers. The switch's OSPF support allows you to enable or disable OSPF protocols on the switch's private and tunneled interface.

The interface filters setting can affect the behavior of routing protocols. For example, OSPF uses IP as its transport mechanism and if the interface filters are set to deny IP, then OSPF advertisements are not sent or received.

For OSPF to function, you must install the Advanced Routing Key. The Firewall Key is optional and required only when you want to enable the redistribution capabilities of RIP and OSPF.

To install a software license key:

1. Go to the Admin→Install Keys screen.
2. Type the key that you obtained from Nortel Networks Customer Support in the box to the right of Advanced Routing.
3. Click on the Install button.

After a valid key is installed, the label Key Installed is displayed. It is only necessary to install a key once on each switch. Click on the Delete button to remove the key. A confirmation message appears and if you click Yes, the key is removed.



Note: The presence of the Advanced Routing License is checked only when OSPF is globally enabled. If you enter the Advanced Routing Key, globally enable OSPF, and then delete the Advanced Routing Key, OSPF will continue to run. However, if you then disable and re-enable OSPF, it will no longer run.

Virtual link support

The OSPF protocol can create virtual links automatically to link the segments in a backbone network. The switch must be an area border router for the automatic configuration of virtual links to operate properly.

Because version 3.x does not support OSPF vlinks, you must turn on or off a flag using the CLI Compat_3_6 enable command. However, this flag does not permit switches running version 4.0 to form a virtual link over tunnels.

If all the switches in the network are running version 4.0, then Nortel Networks recommends that you use the no Compat_3_6 enable flag on all of the switches. Even if only one switch running version 3.x participates in OSPF through tunnels, then every switch with version 4.0 needs to be set with the Compat_3_6 enable flag.

For example, if endpoint A is running version 4.0 and endpoint B is running version 3.x, end point A should have the Compat_3_6 flag enabled for OSPF to work properly.

However, if both end point A and end point B are running version 4.0 and both have enabled the Compat_3_6 enable flag, OSPF will work properly, but vlink over this tunnel may not work properly. Also, if you set the flag for endpoint A to Compat_3_6 enable and endpoint B to no Compat_3_6 enable flag, OSPF may not work properly.

If you set the flag for endpoint A to no Compat_3_6 enable and set endpoint B to no Compat_3_6 enable, OSPF and vlink will both work properly over this tunnel.

ECMP support

OSPF supports a maximum of four simultaneous equal cost multiple paths, which are used for load balancing and resource sharing. This feature can be enabled or disabled.

Configuring OSPF

To configure OSPF:

1. Click Routing→OSPF to configure OSPF global parameters. Enabled indicates that the OSPF protocol is enabled on this screen. The default setting is disabled.
2. Type in the IP address in the Router ID field. This uniquely identifies the router in an area and defaults to the lowest of the IP addresses of management/physical interfaces defined in the switch. You can change this address provided that it is unique within the area. The default is the lowest IP interface in the box.
3. An AS-Boundary-Router is a router that exchanges routing information with routers belonging to other autonomous systems and advertises AS external routing information throughout the AS. Select True from the AS-Boundary-Router drop down list. The default is false. This parameter must be set to True if you want to enable redistribution of non-OSPF routes via OSPF.
4. To automatically create virtual links to link the segments in a backbone network, select True from the drop down list. The default is False.
5. OSPF supports two types of external metrics. Type 1 external metrics are expressed in the same units as OSPF interface cost (in terms of the link state metric). Type 2 external metrics are an order of magnitude larger; any Type 2 metric is considered greater than the cost of any path internal to the AS boundary router. Use of Type 2 external metrics assumes that routing between AS boundary routers is the major cost of routing a packet, and eliminates the need for conversion of external costs to internal link state metrics. Select metric Type 1 or Type 2. Type 1 is the default.
6. Equal Cost Multipath provides load balancing of packets to a destination that is reachable over more than one physical interface. Select the maximum

number of paths (1-4). The OSPF compile time maximum equal cost paths parameter is four paths.

7 The Known OSPF Areas section displays all of the OSPF areas defined locally to the switch. The area information is not shared among switches. If you want two switches to have one of their interfaces in a common area, you must configure both switches to define the area information. Area IDs are used as representations of parts of the OSPF network. They help to manage large numbers of networks so that they can exchange information within an area. Each Area ID must be unique for OSPF. By default all switches have an area named 0.0.0.0.To add an OSPF area, click on the add button. The Routing Protocols→Add Area screen appears:

- **a** Enter the IP address in the Area ID field.
- **b** For Stub, select True of False from the drop-down list. The default is False.
- **c** For Stub Metric, enter the number of the stub metric. The default is 1.

8 Configured Physical Interfaces lists:

- **a** IP address of the configured OSPF interfaces. The default is 0.0.0.0.
- **b** Area ID of the configured OSPF interfaces. The default is 0.0.0.0.
- **c** Type is either Broadcast or Point-to-Point. The default is Broadcast.
- **d** State is either Enabled or disabled. The default is enabled.

9 You can save the LSDB table as a text file in the directory ide0/system/routing. Enter the file name and click on the Save button.

10 In the Status section, you can display. LSDB (Link State Database), Neighbor, Interfaces, or Summary. Click on the LSDB button to display the link state databases in all areas that are configured in that switch. The following table describes the information in the OSPF LSDB screen.

Table 15 LSDB screen

Column	Description
Link State ID	Link state address
Adv Router	Advertising router address
Age	Age in seconds
Seq Nbr	Sequence number

Table 15 LSDB screen (continued)

Column	Description
CheckSum	Checksum
Links	Number of links

11 Click on the Neighbor button to display a list of neighbors on all the interfaces running OSPF. The following table describes information on the OSPF Neighbors screen.

Table 16 OSPF Dynamic Neighbors screen

Column	Description
Router ID	OSPF ID of neighbor
P	Priority number
State	State of neighbor connection
Dead Time	Time until neighbor is declared dead
Address	Neighbor IP address
Interface	Local IP interface address

12 Click on the Interfaces button to display the list of interfaces that you configured for OSPF. The following table describes the fields for the OSPF Interfaces screen.

Table 17 OSPF Interfaces screen

Column	Description
IP Address	Local IP address
Area ID	OSPF area for the interface
Interface Type	Broadcast (BCAST) for Point to Point (PTPT)
Interface State	State of interface: Enabled or Disabled (physical) or Other (tunnel)
Metric Cost	Cost associated with the interface
Priority	Priority used to negotiate DR/BDR state
Designated Router	Designated router (0.0.0.0 for PTPT)

13 Click on the Summary button to display the overall summary of OSPF running on the switch.The following table describes fields on the OSPF Summary screen.

Table 18 OSPF Summary screen

Column	Description
Router ID	Unique OSPF ID of router
Router State	OSPF global configured state (up or down)
Supports TOS	Type of Service support
SPF schedule delay	Shows Shortest Path First
Hold time between two SPFs	Time between Shortest Path First calls
Minimum LSA interval	Link state advertisement interval
Minimum LSA arrival	Link state advertisement arrival minimum
Number of external LSA	Number of link state advertisements
Link State Update Interval	Time between link state updates
Link State Age Interval	Time between link state age intervals
Number of Areas in this router	Number of areas
Area	Area ID
Number of interfaces in this area	Number of interfaces in this area
SPF algorithm has executed	Number of times shortest path algorithm has been executed

14 Click on the Statistics button to display statistical information about OSPF. The following table describes the fields on the OSPF Statistics screen.

Table 19 OSPF Statistics screen

Column	Description
Interface-CID	Local IP interface address and circuit ID
Hellos	Number of “Hello” packets received and transmitted
DBs	Number of “DB” (Database Exchange) packets
LS Req	Link state requests
LS Upd	Link state updates
LS Ack	Link state acknowledgements

Routing table

Like any router, the switch has a routing table that defines how traffic that comes into the switch is forwarded on to its destination. The routing table can contain both static and dynamic routes. Static routes are manually configured routes that do not change. Dynamic routes, however, do change, as they are learned by using the Routing Internet Protocol (RIP) or Open Shortest Path First (OSPF) from a private interface or a branch office tunnel (the switch does *not* support RIP or OSPF for public interfaces).

The route table contains the best routes submitted by the routing protocols and the static route application and dynamic protocols, such as OSPF and RIP. The route table manager chooses the best route based on the following order of protocol priority: direct route, static route, OSPF route, RIP route, default route. With this and the protocol cost, the route table manager selects the best route and forwards it to the forwarding table.

To view the route table, go to the Routing→Route Table screen. You can view the best routes to a destination or all routes to a destination. You can also filter based on protocols. You can also copy all the route table entries into a file for off-line analysis.

Figure 31 shows how different routing protocols interact with the routing table manager.

Figure 31 Router configuration with OSPF and RIP

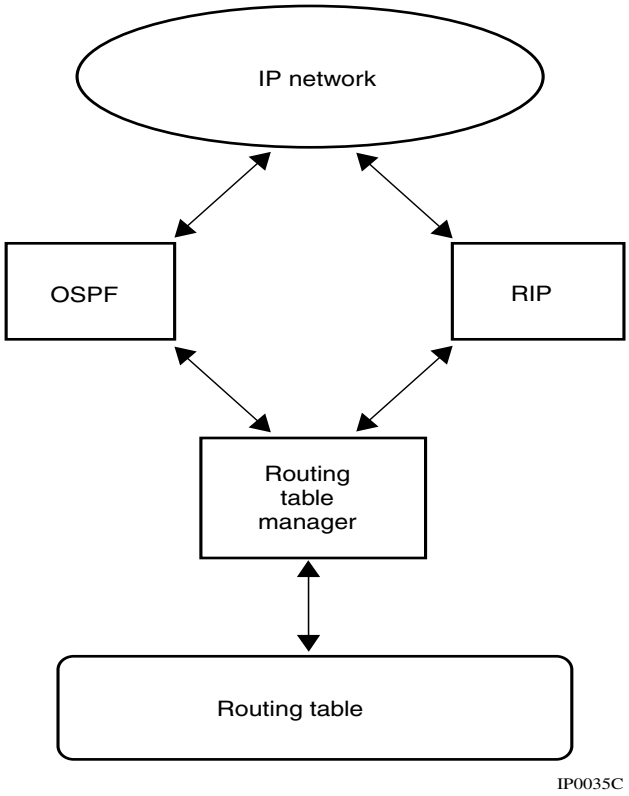
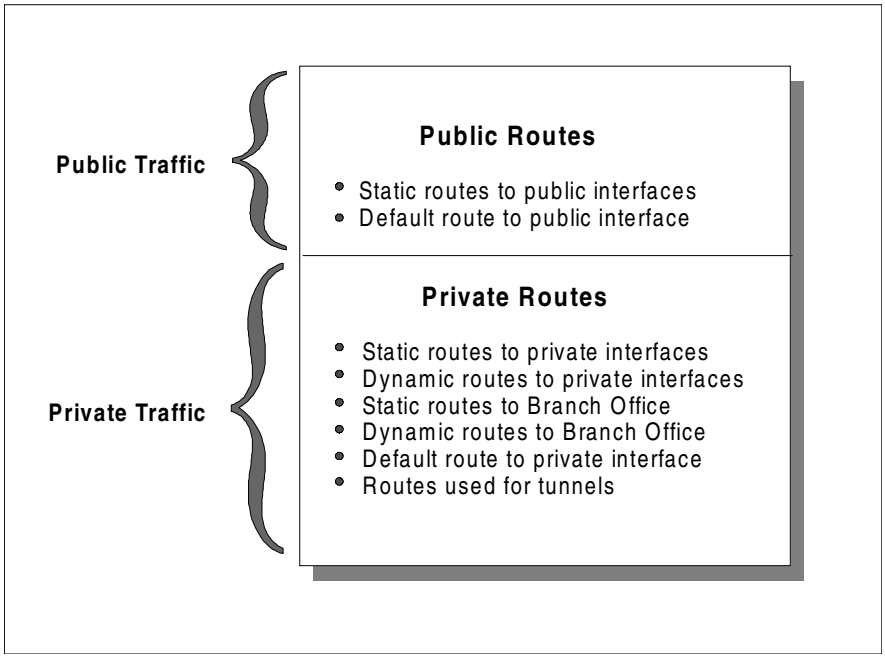


Figure 32 shows the switch's routing table entries divided into two groups: public and private. The tunnels are trusted interfaces and therefore grouped into private routes.

Figure 32 Grouping of route table entries



To search the route table:

1. In the Destination field, select the All, Host, or Network option. If you select Host, you can select whether the interface is All or choose the address from the Interface drop-down list. From the Protocol drop-down list, select All or the protocol (OSPF, RIP, Static, or Direct). You must enter the IP address in the edit box. If the destination is Network, you can select whether the interface is All or choose the address from the Interface drop-down list. From the Protocol drop-down list, select All or the protocol (OSPF, RIP, Static, or Direct). You must enter the IP address and the Network Mask in the edit boxes. You can choose the Exact or Best Match from the Search Type drop-down list.
2. In the Interface field, select All or the current switch interface.
3. In the Protocol field, select All, OSPF, RIP, Static, or Direct.
4. Click on the Search button.

To save the route table to a file:

1. Enter the file name.You can save the route table as a text file in the directory ide0/system/*xxx*/, where *xxx* is the name of the file that you specify.
2. Under Route Filter, select Best Routes to view all routes to a single or All Routes to view all destinations. The default is Best Routes.

To check on route table status:

1. Click on the IP Forwarding Table button to display the following information for the IP Route Network Table, the IP Route Host Table, and the IP Public Address Table. The following table describes fields on the IP Forward Table screen.

Table 20 IP Forward Table Screen

Column	Description
Destination/Mast	Network address and mask
Gateway	IP address of next hop gateway
Flags	Internal use flags
Refcnt	Reference count
Use	How many time used
Interface	Interface identifier
MTU	Size of packet
OuterCtxt	(For internal use only)
CircMap	(For internal use only)
RtEntryP	(For internal use only)

2. Click on the Route Table button to display the full internal routing table displays all routing information. The following table describes fields on the IP Routing Table screen.T

Table 21 IP Routing Table screen

Column	Description
Seq	Sequence number that shows the best route
Proto	Protocol
IP Address/Netmask	IP Address/NetMask

Table 21 IP Routing Table screen

Column	Description
Weight	A combination of cost and priority for the best route
NextHop	IP address of the next hop
NextHopInterface	IP address of the next hop interface
CId	Circuit ID

Route lookup

When a packet arrives, the switch performs a full lookup in its routing table. The following sequence is used to determine which route to use:

- If the switch's firewall support is enabled, all the public and private routes in the routing table are available to the traffic.
- If firewall support is not enabled, then only the private portion of the routing table is available.
- If the traffic's destination route is not found in the table, the table's public and private default routes are invoked as described in the following section.

The switch's routing table can be thought of as having two separate parts. One part contains the routes for traffic that uses the switch's public interfaces, and a second part of the table has the routes for traffic using the private interfaces. Tunnels are virtual interfaces and are treated as private interfaces.

The following list shows the types of routes that might be in the switch's routing table:

- Static routes
  - To public interfaces
  - To private interfaces
  - To a configured branch office
- Dynamic routes
  - To private interfaces
  - To a configured branch office

- Default routes
  - — To public interfaces
  - — To private interfaces
- Host routes
  - — Routes added for dialed VPN users (for example, Contivity VPN Clients or PPTP clients)
- Utunnel Routes
  - — Host/network routes for clients that log in using the client address redistribution feature

Route selection based on destination

Selection of a route to a specific destination is based on the most specific match. For example, if you have a route to the network 10.1.0.0/16 via A and another route to 10.1.2.0/8 via B, traffic to the destination 10.1.2.1 will be sent via B even though the address matches both 10.1.0.0/16 and 101.2.0/8. If B is not available, then it will flow through A.

Route selection based on precedence in route table

The routing table selects the best routes submitted by the routing protocols and submits them to the forwarding table. The order of precedence is:

1. Direct routes
2. Static routes
3. OSPF routes
4. RIP routes
5. Default - static routes (locally defined default routes)
6. Default - OSPF routes (learned from some other routers via OSPF redistribution)
7. Default - RIP routes (learned from some other routers via RIP redistribution)

Within each selection above if there are multiple paths of the same cost and if ECMP is enabled, the route table manager inserts the appropriate routes into the forwarding table.

Routing policy service (RPS)

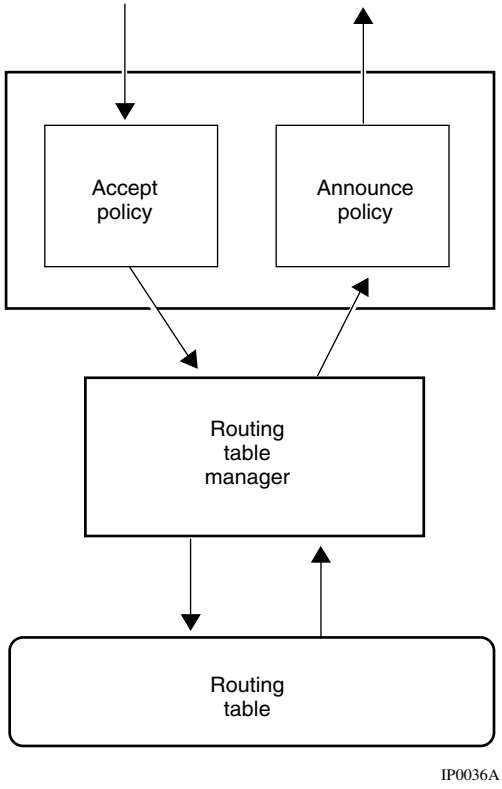
The IP routing function allows you to control the flow of routing data to and from the routing tables. The routing policy service controls this by providing IP accept and announce policies. You can enable or disable a policy as needed.

IP accept policies govern the addition of new RIP- or OSPF-derived routes to the routing tables. When RIP or OSPF receives a new routing update, it consults its accept policies to validate the information before entering the update into the routing tables. Accept policies contain search information (to match fields in incoming routing updates) and action information (to specify the action to take with matching routes).

IP announce policies govern the propagation of RIP or OSPF routing information. When preparing a routing advertisement, RIP or OSPF consults its area boundary router to determine whether the routes to specific networks are to be advertised and how they are to be propagated. Announce policies contain network numbers (to associate a policy with a specific network) and action information (to specify a route propagation procedure). For OSPF, announce policies are applied only to external routes. For RIP, announce policies apply to all routes, including external routes that are redistributed into RIP and RIP-generated routes.

Figure 33 shows the direction between the route table manager and accept/ announce policies of the routing protocol.

Figure 33 Accept and announce policies



Periodically, the routing table manager announces routes to other routers in the network. The routing table manager forwards a route for advertisement to the protocol. The protocol consults an announce policy to determine whether or not to advertise the route to the network.

OSPF link-state advertisements (LSAs) are received and placed in the link state database (LSDB) of the router. The information in the LSDB is also propagated to other routers in the OSPF routing domain. According to the OSPF standard, all routers in a given area must maintain a similar database. To maintain database integrity across the network, a router must not manipulate received LSAs before propagating them to other routers. To accomplish this, OSPF accept and announce policies act in the following manner:

OSPF accept policies control which OSPF non-self-originated external routing information is passed to the routing table manager. The accept policies control only what the local router uses; they do not affect the propagation of OSPF internal and OSPF non-self-originated external information to other routers.

OSPF announce policies control which self-originated external routing updates are placed into the LSDB for distribution according to the OSPF standard. OSPF announce policies affect what other routers learn but only with regard to the local router's self-originated information.

To configure routing policy services:

1. Go to Routing→Policy.
2. Check the Enabled box to enable RPS. The default setting is disabled.
3. Under Redistribution Table, the name of the protocol, either OSPF or RIP, is listed. Select the source of the route for each protocol, either Static, Direct, or RIP.
4. Click on the Add button to add a policy.
5. Enter the Access Name/Number. This can be any name or number that you choose. A numbered access list can use numbers between 1 and 99. A named access list can be any alphanumeric name that begins with a letter of the alphabet. You need to create an Access list before creating policy entries. To create the list, click on the New Access List link. This displays the Access Lists screen with all named items listed. You can edit or delete a selected list name or create a new one by typing the name in the edit box.
6. Enter the name of the protocol, either OSPF or RIP.
7. Enter Interface address. This is the IP address of the physical interface where you want to apply the policy. Use 0.0.0.0 if you want to apply the policy for all interfaces. Use the tunnel endpoint IP address for tunnels.
8. Enter Policy Type: accept or announce policy. You can have only one Accept or Announce policy for each protocol per interface.
9. Click on OK.

Using default routes

Default routes are used when the switch receives traffic for which no matching route exists in the routing table. The use of the default routes depends on several factors, such as whether integrated firewall support is enabled and where the traffic originated (for example, from the public or private interface). Table 22 lists how the use of default routes is determined.

The default gateways are the addresses of the next-hop router. Packets are routed through the default gateway onto the private or public network when the Routing Table does not have a specific route to the destination.

Redistribution of routes

The switch supports redistributing static, direct, and RIP routes into OSPF. Also, it supports redistributing static, direct, and OSPF routes into RIP. Such a redistribution can be further controlled on a per interface basis in RIP.

The switch is a security-aware device. Therefore, the route redistribution is also based on security configurations. Table 22 describes the rules of redistribution for RIP and OSPF with the firewall enabled or disabled.

Table 22 Redistribution rules

Redistributed Route	Firewall ON	Firewall OFF
Public Direct Route	Yes	No
Public Default Route	Yes	No
Public Static Route	Yes	No
Private Direct Route	Yes	Out physical - No; out tunnel - Yes
Private Default Route	Yes	Out physical - No; out tunnel - Yes
Private Static Route	Yes	Out physical - No; out tunnel - Yes
Tunnel Static Route	Yes	OSPF - Always Yes RIP - In general, Yes, but can be controlled on a per-interface basis
Tunnel Dynamic Route	Yes	Yes
Utunnel routes	Yes	Yes

Redistribution of default routes

With default routes (public or private), when redistributed by a dynamic routing protocol, the receiving router treats these routes as protocol-specific default routes. Therefore, any locally defined default route will have a higher precedence over any routes learned by redistribution.

Even though a public default route is represented by 0.0.0.0/32 when redistributed, it is represented as 0.0.0.0/0 to conform with the routing protocols. When static routes are redistributed into a routing protocol, default routes are also redistributed. However, if you have both private and public default routes, only one of them will be redistributed, thus reducing the number of redundant routes to the same destination through the same next-hop interface.

Traffic patterns

The switch's enhanced routing feature provides a range of routing support. Table 23 shows the traffic patterns that are supported and lists which part of the routing table the traffic uses.

This table also indicates whether the traffic goes through the switch's firewall. In this table, tunnel means either branch office or client tunnel. The private and public routes in the second column are as defined in Figure 32.

Table 23 Traffic patterns

Traffic	Which routes are used when no firewall is enabled?	Which routes are used when a firewall is used?	Does it go through the firewall when a firewall is used?
VPN routing and enhanced routing			
From tunnel: • tunnel→private • tunnel→tunnel • tunnel→public *	Private routes only	• Uses all routes. • Private default has precedence over public default.	• Yes if it is to a public interface • No for all other cases
From private interface: • private→private * • private→tunnel • private→public *	Private routes only	• Uses all routes. • Public default has precedence over private default.	• No if it is to a tunnel interface • Yes for all other cases
From public interface: • public→private * • public→tunnel *	Traffic is dropped.	• Uses all routes. • Private default has precedence over public default.	• Yes
Services routing			
Inbound destined for switch: • private→Switch • tunnel→Switch • public→Switch	Does not use routing table.	Does not use routing table.	• No if it is from a tunnel • No if it is a tunnel protocol • No if it is to the management address • Yes for all other cases
Outbound from switch private interface address (includes tunnels)	Private only	• Uses all routes. • Private default has precedence over private default.	• No if it is a tunnel protocol • No if it is from management address • Yes for all other cases
Outbound from switch public interface address	Public only	• Uses all routes. • Public default has precedence over public default.	• No if it is a tunnel protocol • Yes for all other cases

* Indicates traffic goes through the firewall.

Multicast relay

IP multicast is an extension to the standard IP network-level protocol. It provides efficient delivery of information from a single source to multiple destinations. It is useful for applications such as video conferences, shared white boards, and news feeds. IP multicast uses class D addresses, ranging from 224.0.0.0 through 239.255.255.255. Multicast routing protocols establish the distribution tree for a given multicast group.

A multicast relay listens to an incoming multicast and forwards, or relays, that broadcast to one or more destination addresses in the absence of multicast routing. The destination addresses may be unicast in the case of tunnel or multicast. Multicast relay is not supported on public interfaces.

By default, multicast relay is globally disabled. If multicast relay is disabled on this screen, the switch processes only multicast requests in the range of 224.0.1.0 to 239.255.255.255. When you enable multicast relay, received traffic is filtered according to filter checking and access lists.

You can configure the congestion threshold as a percent of system resources. The default value is 3000. If the unicast packet forwarding performance decreases due to multicast traffic, you can reduce this number.

You can forward multicast packets over a tunnel using default filter (permit all). For example, if you want to allow multicast packets received over the interface to be relayed over tunnel B01 and not over tunnel B02, you would define the interface-specific rules as shown in Table 24.

Table 24 Multicast interface-specific rules example

	type	src intf	dst intf	source	dst	service	action
receiving	SRC	LAN	ANY	S	231.0.01	voice	allow
relay	DST	ANY	BO1	S	231.0.0.1	voice	allow
relay	DST	ANY	BO2	S	231.0.0.1	voice	drop

To configure multicast relay:

1 Check the Enabled check box to enable multicast relay on the switch. By default multicast relay is globally disabled. If multicast relay is disabled on this screen, the switch processes only multicast requests in the range of 224.0.1.0 to 239.255.255.255. When you enable multicast relay, received traffic is filtered according to filter checking and access lists.



Note: Multicast requires use of the Permit All Interface filter.

2 You can configure the congestion threshold as a percent of system resources. The default value is 3000. If the unicast packet forwarding performance decreases due to multicast traffic, you can reduce this number.

3 Click on the Add button to add an interface to the multicast boundary list. The Multicast→Add screen allows you to add an interface.

- **a** Enter the Access Name/Number in the edit box.
- **b** Select the IP address for the interface.
- **c** Select the either Enabled or Disabled for the State.
- **d** Click on the New Access List link to view the existing access screen.

4 Click on the Statistics button to display the global multicast relay status and the statistics of the configured multicast interfaces, including branch office interfaces.The following table describes fields on the Multicast Statistics screen.

Table 25 Multicast Statistics screen

Column	Description
Interface	IP address of interface
CID	Circuit ID
PktsRcvd	Number of packets received
PktsSent	Number of packets sent
PktsDropped	Number of packets dropped

5 Click on the Interfaces button to display all of the configured information about enabled interfaces, including private physical and branch office tunnel

interfaces.The following table describes fields on the Multicast Interfaces screen.

Table 26 Multicast Interfaces screen

Column	Description
Interface	IP address of interface
Access-list	Name of the access list

Equal-cost multipath (ECMP)

Equal-cost multipath provides load balancing of packets to a destination that is reachable over more than one physical interface. This can increase the forwarding capacity of a switch that is media bound. Load balancing is provided on a per-packet basis or a packet stream basis. You can also use it to balance traffic over tunnels whether or not they are going out single or multiple physical interfaces.

With equal-cost multipath, the routing protocol applications can submit an equal-cost route for a destination. The routing table manager passes the set of equal-cost best paths to the forwarding table.

The switch can handle up to eight equal-cost paths. Their use depends on the support from the routing protocol application. OSPF uses a maximum of four paths, RIP uses only one, and static routes use all eight.

The load balancing and resource sharing is controlled by the following forwarding policies:

- Per-packet - packets are forwarded in a round-robin fashion. If Contivity State-Aware Firewall is enabled, this policy may cause some overhead in switching the firewall context.
- Per-destination - packets are forwarded based on source and destination address pair.
- Per-source - packets are forwarded based on source address.

You can change the forwarding algorithm to per-packet, per-destination, or per-source without affecting routing or forwarding tables.

Client address redistribution

When a client initiates a user tunnel, the switch assigns an inner address to the client. Sources for these addresses can be:

- A predefined address pool in the switch with an address range that belongs to a locally attached private network
- A predefined address pool in the switch with an address range that does not belong to any locally attached private network
- A static address configured in the switch
- A RADIUS or DHCP address
- A client-supplied private address

If addresses from any of these sources do not belong to any of the locally attached switch networks, you must enable client address redistribution to ensure that these addresses are advertised in the dynamic route updates sent out by the switch. Client address redistribution uses a route type called a Utunnel. Utunnel routes can be either host or network routes.

When client address redistribution is active, the switch creates and advertises a user tunnel host route whenever a client tunnel is created using an inner address that does not belong to a locally attached network. When the tunnel is taken down, the corresponding host route is deleted.

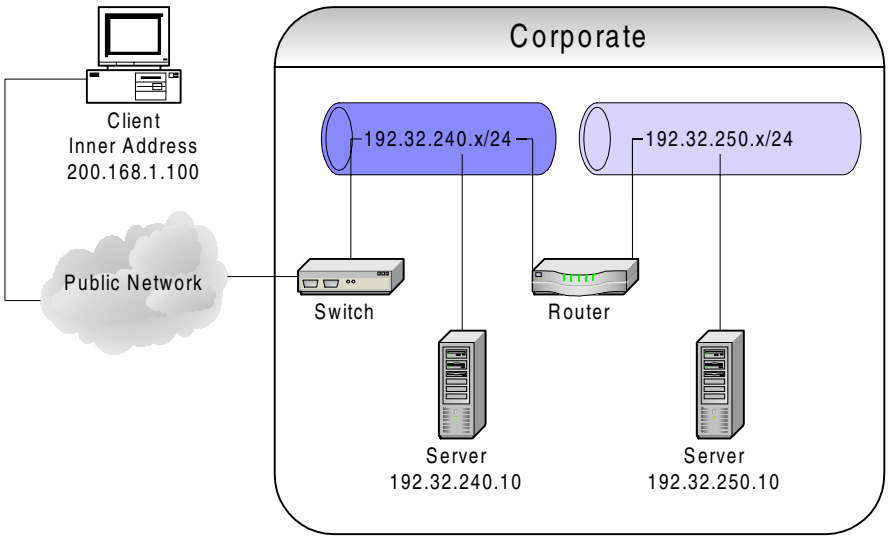
When inner addresses are being allocated from an address pool with a range that does not belong to a locally attached network, you can reduce the number of entries in the route table and the route redistribution overhead by using the summarization option. Summarization creates and advertises a single Utunnel network route covering the address pool range when a client tunnel is created using an inner address from this address pool. The network route remains in the route table until the last tunnel using an inner address from this address pool is taken down.



Note: The maximum number of Utunnel routes cannot exceed the maximum number of client tunnels supported by the corresponding hardware platform. The default value is 200.

Figure 34 shows an example of client address redistribution where the client has an inner address that is not within the local subnet of the private network. The switch creates a Utunnel route that is then propagated over the network. This allows the router on the private network to recognize the 200.168.1.100 address and route responses back to it properly.

Figure 34 Client address redistribution



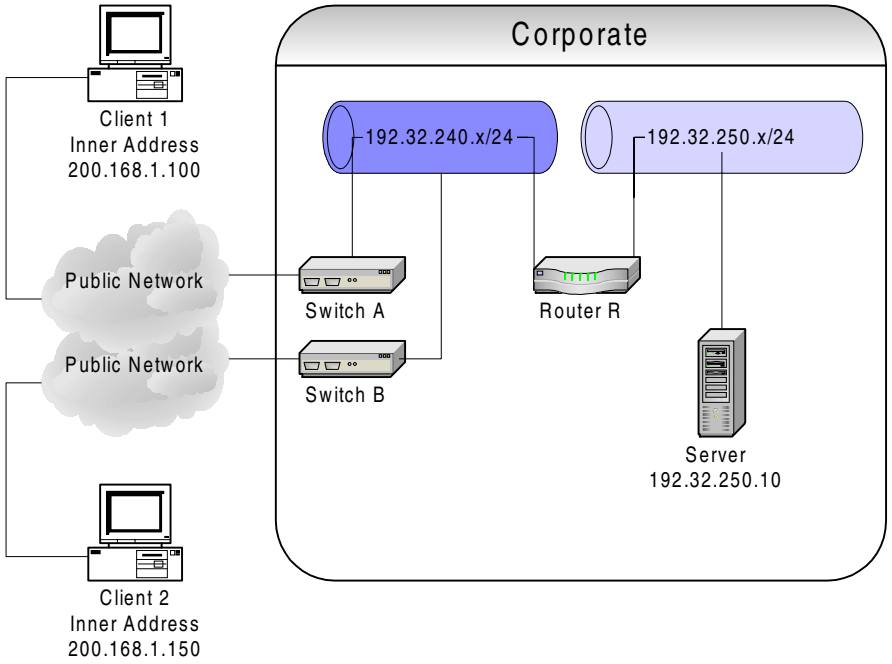
If you enable summarization, the switch identifies the subnet from the address pools where this address belongs and inserts a user tunnel network route for this subnet into the route table manager. Summarization is disabled by default.

Enabling summarization is useful for large networks where optimization reduces the number of Utunnel host entries in the RTM. However, when using summarization, you could potentially have connection problems if the subnets of the address ranges span multiple switches. If you have two switches assigning addresses that belong to the same IP subnet, do not use the summarization option.

For example, in Figure 35 Switch A has an address range of 200.168.1.100 to 200.168.1.120 and Switch B has an address range of 200.168.1.150 to 200.168.1.170. Both of these ranges are part of class C subnet 200.168.1.*x*/24. Client 1 logs in to Switch A and Client 2 logs in to Switch B. Both clients have inner addresses that are not within the local subnet of the private network, but are

in the same IP subnet. Switch A and Switch B running client address redistribution create Utunnel host routes. These routes are propagated over the network. The router on the private network recognizes addresses 200.168.1.100 and 200.168.1.150 and can route responses back to them through the designated switch.

Figure 35 Summarization for client address redistribution



If you enable summarization on both switches, both switches will advertise routes to 200.168.1.*x*/24. Router R will use one of these routes, causing either Client 1 or Client 2 to have communication problems.

The routing table manager handles Utunnel routes similarly to other route types (RIP or OSPF). You can view Utunnel routes using the Routing→Route Table Manager screen. The route policy service handles redistribution (advertisement) of Utunnel routes similarly to redistribution of other route types.

To configure client address redistribution, go to the Routing→Client-Addr-Dis screen:

1. Check the Enabled box to enable client address redistribution. If you enable client address redistribution, operations related to the Utunnel are effective. If you disable client address redistribution, the switch stops redistribution of all Utunnel routes. The routes remain in the route table, but are not advertised. Any existing user tunnels remain logged in, but may not be able to communicate with the private network and any new Utunnel routes are disallowed. The default is disabled.
2. Under Number of UTunnel Host Routes, the Configured Maximum field you can limit the maximum number of user tunnel host route entries in the system. The default value is 200.The Current field displays the current number of user tunnel hosts logged in to the system.
3. Check the Summarization box to enable summarization. If you disable summarization, the switch inserts a user tunnel host route for the client address into the route table. If you enable summarization, the switch identifies the subnet from the address pools where this address belongs and inserts a user tunnel network route for this subnet in the route table. The default is disabled.
4. Click on the Show User Tunnel Routes to display the user tunnel routes.
5. Click on the Statistics button to display the configuration of client address redistribution, including whether it is globally enabled or disabled, the utunnel limit, current utunnel count, current utunnel redistribution configuration and summarization options. The following table describes these statistics.

Table 27 Client address redistribution statistics

Column	Description
IP address	IP address
Mask	IP network mask
Next Hop	Next hop address
Interface	IP interface address
Cost	Relative cost for the switch

VRRP

Virtual Router Redundancy Protocol (VRRP) is one method you can use to configure the switch to maintain a state of high availability. VRRP is a standard protocol that handles private interface failures. VRRP targets hosts that are configured with static next-hop routing addresses or default gateways. It provides a means of rerouting traffic in the event of a system/interface failure.

VRRP provides switch-level failover in case a private physical interface fails. Currently, it does not provide such failover for public or tunneled interfaces. However, using VRRP and dynamic routing provides a high degree of failover. VRRP is managed as two separate parts. The first part handles those configuration parameters that must be the same between all switches that make up a Virtual Route (VR).

Use of an external LDAP server makes it easier to configure VRRP because it provides a common location in which information about each switch in the system can be maintained. Use of an external LDAP server enables each switch to see the settings of other switches in the system. Configuration of VRRP requires that Virtual Router IDs (VRIDs) are agreed to by all participating switches.

An external LDAP server is not a requirement. If the internal LDAP server is being used, the various switches must have these parameters configured the same and the responsibility for doing so lies with the administrator.

The second part of VRRP configuration is the information that is specific to a switch. This is information that is related to an interface and the role that the interface plays in VRRP (master or backup). This information is kept in the normal configuration file that is stored on the switch.

VRRP and dynamic routing for high availability

High Availability (HA) depends on a number of functions and their proper inter-working: the core routing features, VRRP, a dynamic routing protocol (RIP or OSPF), and consistent configuration.

The performance of HA also depends on the remote system's ability to support the standard functionality being used. For the proposed HA solutions to work, the remote systems must support the standard routing and tunnel protocol being used by the local system. This means that nothing restricts the remote system from being a non-Contivity platform.

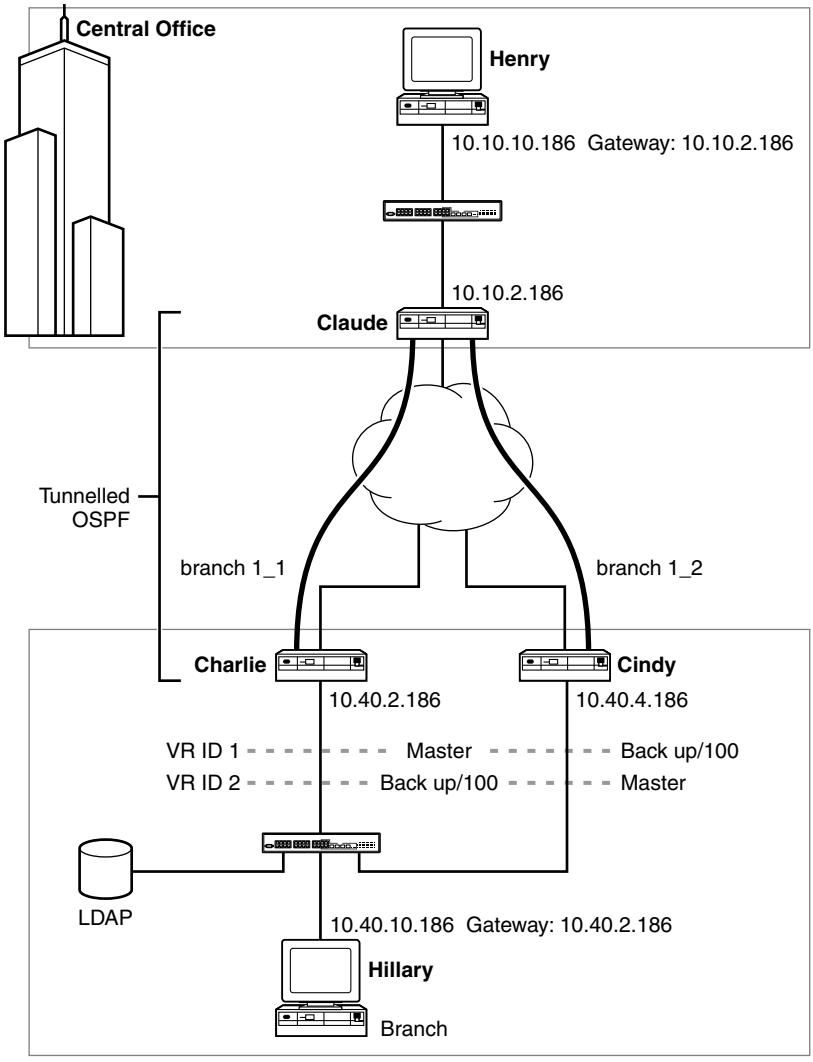
Figure 36 shows a deployment where the central office is configured with one switch, Claude, and the host Henry with the default gateway pointing to Claude. The branch is configured with Charlie and Cindy. VRRP is configured on the private side of Charlie and Cindy backing up each other's physical interface. The host, Hillary, in the branch, has the default gateway pointing to Charlie. There are also two branch office tunnels, as indicated, connecting the switches.

Consider the traffic flow between Hillary and Henry. If Charlie totally fails or the private interface of Charlie fails, Cindy will assume the mastership of the private interface of Charlie. Cindy will assume the IP address 10.2.40.186 as well as the MAC address of Charlie's interface. All the IP traffic (terminating or pass through) targeted to Charlie on the private side will now flow through Cindy. Cindy will forward all the pass through traffic, but will drop all the packets terminating at Charlie. For example, a data packet from Hillary to Henry will be forwarded, but a ping request to 10.2.40.186 will be dropped. Henry is NOT aware of such a change.

Routing configuration plays a vital role in this failover operation. Charlie and Cindy need to know that the path to Henry is via Claude; and Claude should know that there are two paths to Hillary: one via Charlie and another via Cindy. The routing information on the switches can be manually populated using static routes, but it is tedious when the network gets large. Running dynamic routing protocols such as RIP or OSPF will eliminate this tedious manual configuration.

The VRRP failover occurs within 3 seconds based on the default configuration. Use of OSPF on the tunnels guarantees a maximum failover time of 40 seconds based on the default configuration. However, by setting the appropriate value for the hello interval, this can be drastically reduced. Use of RIP takes a maximum of 2.5 minutes based on the default configuration. You can also modify the RIP parameters to reduce this time.

Figure 36 Sample HA environment



10687EA

VRRP and static tunnels

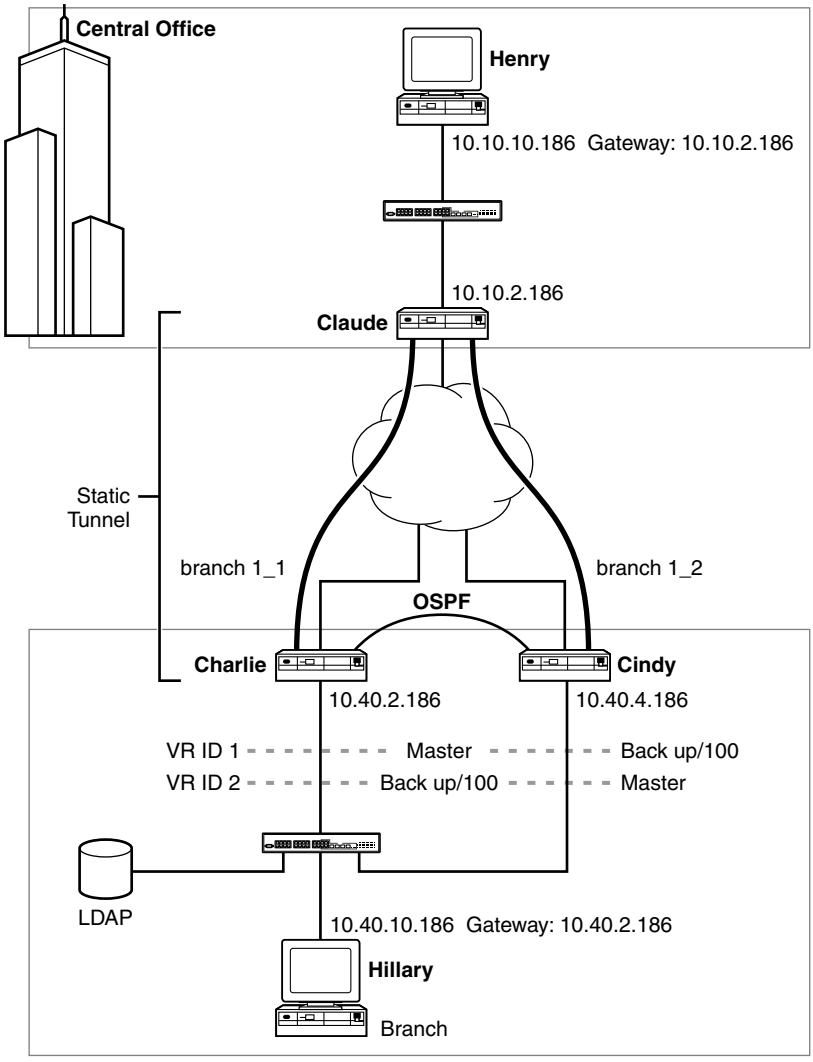
VRRP provides switch-level failover in case a private physical interface fails. This allows a VRRP critical interface to be configured for a physical interface (a tunnel or interface group). It also allows you to configure interfaces into interface groups.

An interface configured with VRRP is called a VRRP master interface. You can associate up to three interface groups with the master interface. If the master interface goes down, the switch fails over. If any of the associated interface groups goes down, the switch master interface fails, causing failover. However, when all interface groups come back up, the switch takes back master ownership. For example, if the first tunnel on a VRRP master interface goes down, failover occurs.

When a VRRP master interface fails over, you cannot manage the switch through that interface, even if the interface group caused the failover. Interfaces and interface groups can be dynamically added and deleted.

In Figure 37, if you use both tunnels as static tunnels, the fail over of the private interfaces using VRRP may not work. If Claude to Cindy has a lower cost than the cost from Claude to Charlie and Hillary has Cindy as a default gateway and a private interface, then Charlie will route the traffic from Hillary to Henry, but traffic from Henry to Hillary will be sent to Cindy and dropped. To prevent this situation, you can have an OSPF tunnel between Charles and Cindy via a public interface and then turn on redistribution of direct and static routes on that tunnel.

Figure 37 VRRP and static tunnels



Configuring VRRP

The following steps describe how to configure your switch for VRRP.

1 Go to the Routing→VRRP screen, enter the IP address, and click on Create.

2. On the VRRP→Edit VRRP IP Address screen, enter a decimal value between 1-255 for the VRID. This number must be unique to the LAN segment running VRRP.
3. In the Advertise Interval edit box, enter the interval at which the VR advertises its virtual MAC address. The range is 1-255 seconds and the default is 1 second.
4. Select None or Simple as the authentication type for this VR. None means that VRRP protocol exchanges are not authenticated and Simple means they are authenticated by a simple text password.
5. Enter up to 8 characters of text for the authentication string and confirm it.
6. You can optionally set Master Delay mode to control when a switch takes mastership of an address it owns. Normally, this occurs when the interface becomes enabled. With Master Delay mode it is possible to delay when this happens. Master Delay mode is optional. The default for a VR is that Master Delay mode is disabled (None). Master Delay mode operates in one of two possible ways: Delay or Time of Day.



Note: The Safe mode feature has the following interaction with Master Delay mode: When safe mode is enabled, a boot after an unclean failure starts the safe mode image, instead of the normal boot image. If the Safe mode image is configured with VRRP, then Master Delay mode works. However, Safe mode automatically boots the normal image after some delay. This boot looks as though it was after a clean shutdown and Master Delay mode is not invoked.

7. Go to the Routing→Interfaces screen and click on Configure on the VRRP protocol for the appropriate interface. The interfaces that correspond to addresses physically on the box are automatically displayed in the Master Status section and all others are displayed in the Backed up addresses section.
8. In the Backed up Address box, add all of the interfaces that you want to back up and click on Add.
9. In the Master Status section, enable all interfaces that you want to be Master.
10. Click on Close.

Interface groups and critical interface failover

Interface groups support the VRRP critical interface fail-over feature of the switch. An interface group is a logical grouping of interfaces (physical or tunnel) defined in a Contivity VPN Switch. The group may consist of a single physical interface, an IP address, a list of physical interfaces, a tunnel, a list of tunnels or any combination of physical interfaces and tunnels defined on the switch.



Note: Status of a Critical Interface is defined to be up when at least one of its members is up and is down when all of its members are down. The main state is mapped to the up state for the switch.

A physical interface on which a VRRP has been configured to run as master is called a VRRP master interface. With each of the VRRP master interfaces, you can associate a maximum of three interface groups. When any one these three interface groups goes down, the switch looks as if the VRRP master interface is down thereby forcing a VRRP fail over. The VRRP master interface stays in this down state until all of the associate interface groups come up and then claims the mastership.

The Configured Interface Groups section of the Routing→Interface Grp screen lists the names of configured interface groups, the number of IP interfaces included in the group and the current admin state and the operational state of the group.

To configure interface groups:

1. Click the Add button to access the Interface Group→Add screen to configure groups.
2. Enter a name for the group in the name field.
3. Select and move the available interfaces that you want to include in the group into the Interfaces in Group list.
4. To find interface groups with a given interface, enter the IP address and click on the Search button.

Chapter 6
Configuring NAT

Network address translation (NAT) enables transparent routing between address spaces. Using NAT in an extranet allows the dynamic connection of multiple private networks via secure tunnels without requiring any address space reconfiguration. Increasing use of NAT comes from two major factors:

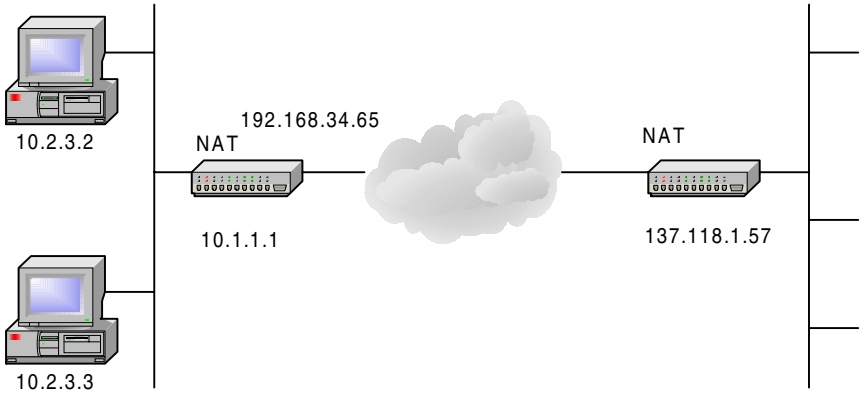
- Shortage of IP addresses. Most Internet Service Providers (ISPs) only allocate one address to a single customer and this address is assigned dynamically, so every time a client connects to the ISP a different address is provided. Because users are given only one IP address, they can have only one computer connected to the Internet at a time. With NAT running on this single computer, it is possible to share that single address between multiple local computers and connect them all at the same time. The outside world is unaware of this division and performs all communications as though only a single machine on the local network were accessible.
- Security. NAT automatically provides security without any special set-up because it only allows connections that originate on the inside network. It is still possible to make some internal servers available to the outside world by statically mapping internal addresses to externally available ones, thus making services such as FTP available in a controlled way.

In the context of virtual private networks, network address translation is necessary to allow multiple intranets with conflicting subnets to communicate. Since the configuration of branch office or partner networks may be fixed, a VPN solution must be able to securely route between these networks without requiring all the private addresses to be unique across the entire extranet.

NAT contains a pool of available global addresses that are continually reused. Internal network addresses are allocated according to internal considerations of the network. Global addresses must remain unique in order to distinguish between different hosts. When a packet is routed, NAT replaces the internal corporate address with a global address. As soon as the application session is over, the global address is returned to the pool and can be used by subsequent connections.

Figure 38 shows a sample NAT environment.

Figure 38 Sample NAT environment



NAT translation

NAT supports both static and dynamic address translations:

- Static address translation allows access to an internal IP address from the Internet, but the IP address remains unknown to the outside world. Every internal IP address mapped by static NAT is converted to a different global IP address. The assignment of NAT IP addresses to original IP addresses is unambiguous.
- Dynamic many-to-many address translation allows you to define a pool of global addresses that can be shared among local users when there are more users available than global IP addresses. Every address mapped by dynamic NAT is allowed access to the Internet/Intranet, but requests from the outside are not allowed access.

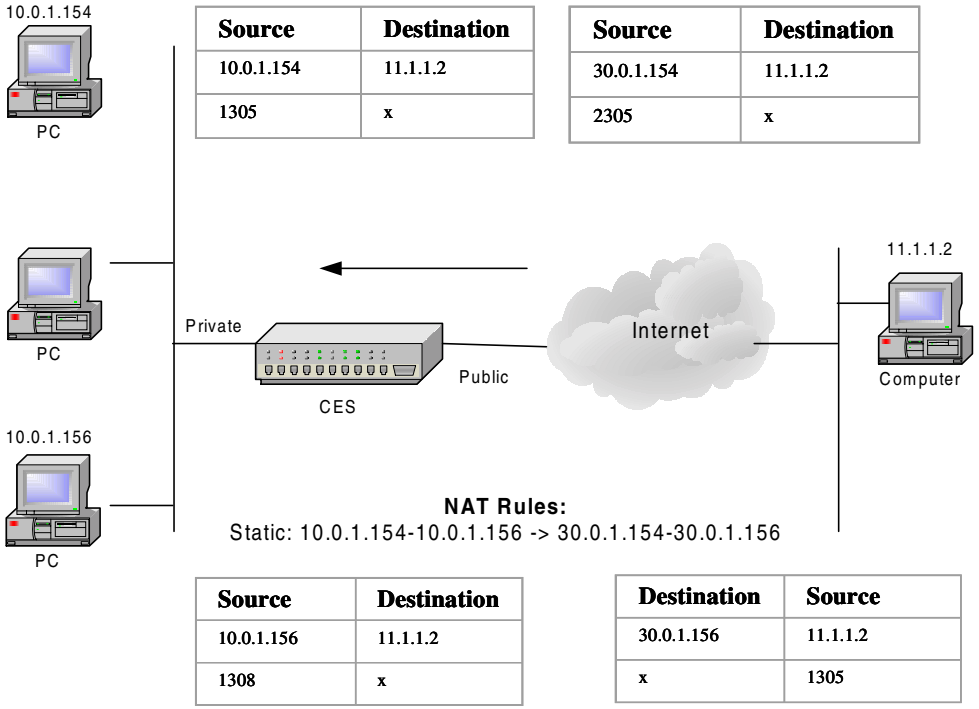
- Dynamic many-to-one translation allows all users on the private network to have one global source. Every address mapped by single NAT is allowed access to the Internet, but access is not allowed in.

Static address translation

A static address is the most specific NAT type or rule and it always overrides dynamic rules. Static addresses are one-for-one. A static address allows bidirectional access to a destination address, as long as you configure the DNS servers accordingly. Additionally, either client device can use the bidirectional address. Static address translation allocates one external host address for each internal address. This allocation is always the same.

Figure 39 shows host 10.0.1.154 on the private side statically mapped to an external address 30.0.1.154, which allows Internet host 11.1.1.2 to initiate a session using the translated external address. The host using this rule is always bound to the same external address.

Figure 39 Static address translation



Dynamic address translation

Dynamic rules include either port addressing or pooled addressing. The rules are dynamic because an address is assigned to a port or is assigned from a pool, depending on the specific situation and the conditions that currently exist for that situation.

Dynamic port NAT

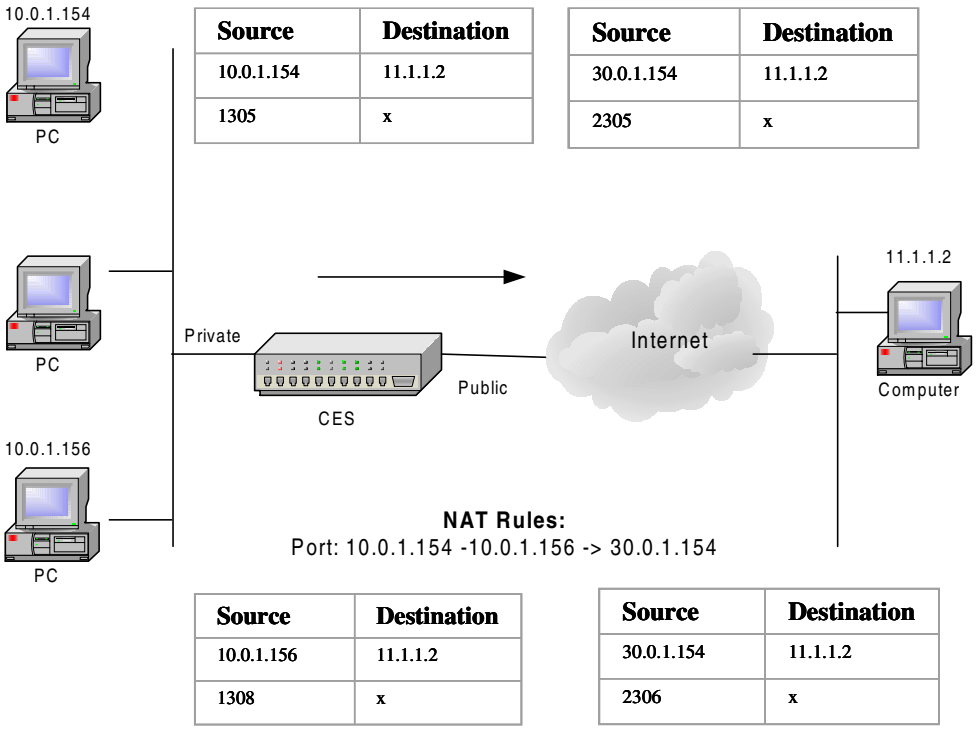
Unlike static rules, dynamic port NAT is not one-for-one. All packet transmissions must be initiated from the internal network. For dynamic port translation, the switch checks to see if the packet matches any translation table entries. If an entry exists, then it modifies the destination port and address appropriately. If there are no matching entries, the switch checks to see if the packet is initiating a connection. If so, then it allocates the next available port, adds the address and

port to the translation table, and modifies the packet accordingly. It allocates the port assignment from the range of unassigned port numbers. For an incoming packet, if there are no matching entries in the translation table, it drops the packet. Port numbers are dynamically assigned from four pools:

- Low Port List: For port numbers less than 512
- Privileged Port List: For port numbers between 512 and 1024
- User Reserved Port List: For port numbers between 1024 and 5000
- High Port List: For port numbers between 5000 and 65535

Figure 40 shows port translation where the private network 10.0.1.0 is hidden behind the public address 30.0.1.154. All requests originating from the private network (10.0.1.0) have their source IP addresses replaced with the public IP address 30.0.1.154; only the public IP address is visible from the public network. In addition, source ports are dynamically translated to unique translated ports.

Figure 40 Port address translation



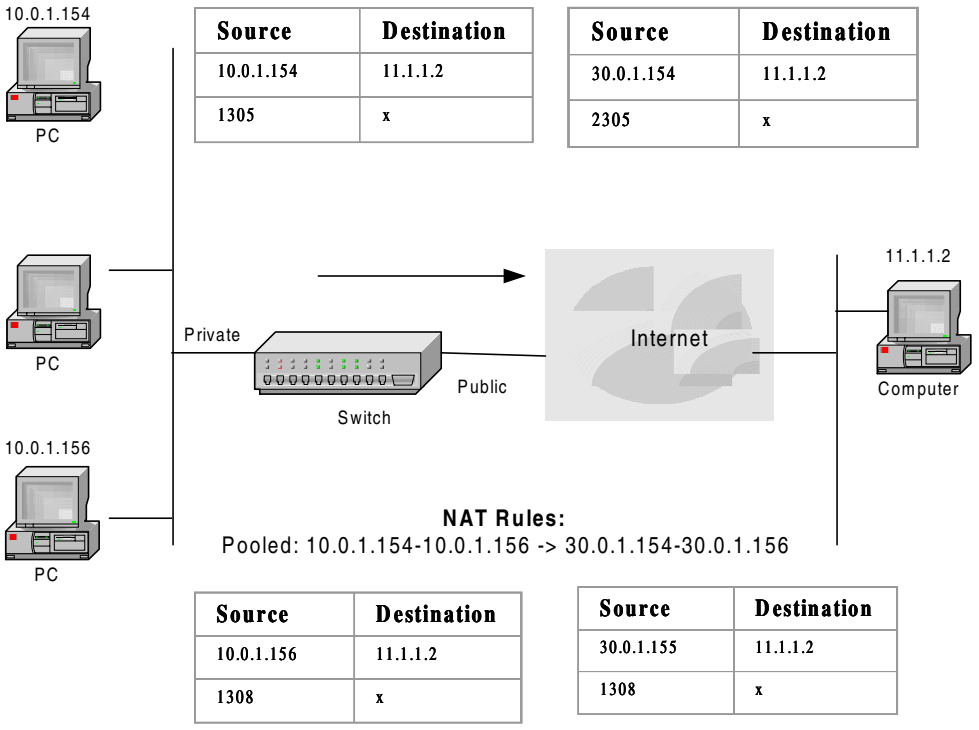
Dynamic pooled NAT

Dynamic pooled NAT is similar to dynamic port NAT. The switch checks to see if an address entry has already been allocated. If so, it updates the packet addressing and sends the packet. Otherwise, the switch attempts to allocate an address from a pool designated for this session. If an address is available, the switch adds the address pair (the original private address and the newly assigned public address) to the translation table and modifies the packet header. If there are no addresses available, it drops the packet. For an incoming packet, if there are no matching entries in the translation table, then it drops the packet.

In dynamic many-to-many NAT, only the address (not the port) is translated. Usually, the number of externally visible IP addresses is less than the number hidden behind the switch. Each time a request is made from a host on the private network, the switch chooses an external IP address that is currently unused, and then performs the translation. You can only use dynamic many-to-many for traffic initiated from an internal host.

Figure 41 shows a pooled NAT rule that converts the internal address range 10.0.1.154 - 10.0.1.164 to 30.0.1.154 - 30.0.1.154. Traffic is initiated from 10.0.1.1.54 and 10.0.156 destined to host (11.1.1.2) on the public Internet. Both addresses are dynamically translated to unique public addresses.

Figure 41 Pooled address translation

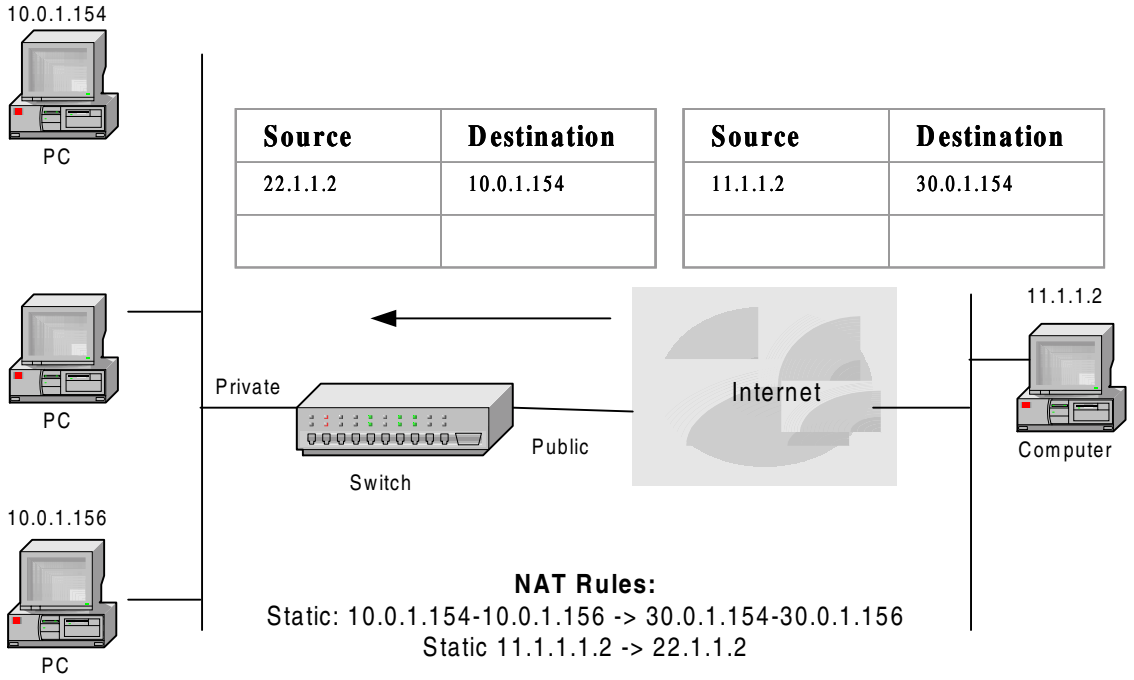


Double NAT

Double NAT allows both external and internal networks to be translated at the same time. For each packet entering and leaving the switch, both the source and destination addresses could be modified. For Contivity NAT, this is achieved using rules, one for translating the source address and one for translating the destination address. The destination address translation must use a static rule.

Figure 42 shows a host 11.1.1.2 on the Internet initiating a connection to 30.0.1.154, the translated address of the internal host. As the packet traverses NAT, both the source and destination addresses are translated.

Figure 42 Double NAT



NAT usage

NAT is applied to routed traffic passing through its physical interfaces (interface NAT) and branch office interfaces (branch office NAT) using separate NAT policies. Each branch office has one NAT policy, and there is one global NAT policy applied to non-tunneled traffic.

Branch office NAT

In branch offices, you might have two or more branches that use the same LAN addressing scheme. Nonetheless, people in the branches still have to communicate with one another.

A typical scenario might include a client on LAN 1 who tries to access the FTP server on LAN 2, and who sends a packet with a source address of 10.0.0.13 and a destination address of 10.0.0.14. Without NAT, the switch would look at the destination address and assume that the destination is on the same LAN as the source device because the addresses are both on the 10.0.0.0 network; therefore, no tunnel connection would be brought up with this packet.

To solve the problem and allow the client to access the server on the other LAN, you implement NAT on both sides of the branch office connection. In this example, switch 1 defines a remote accessible network of 12.0.0.0, and switch 2 defines a remote accessible network of 11.0.0.0. Switch 2 uses a static translation of 10.0.0.14 (server) to 12.0.0.1. Switch 1 uses a translation of 10.0.0.13 (client) to 11.0.0.1. As a result, switch 2 must define 11.0.0.0 as the remote accessible network.

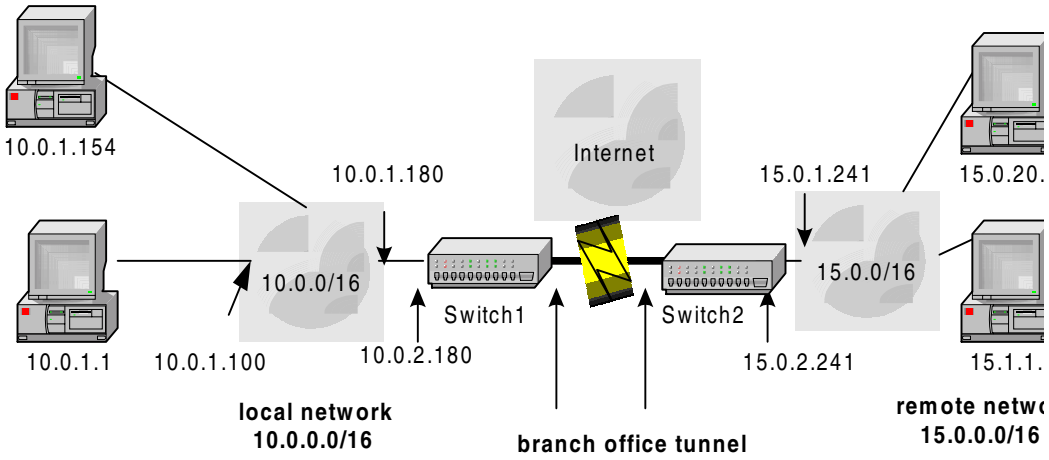
With NAT implemented on both sides of the branch office connection, the client can access the FTP server. A packet generated from the client has a source address of 10.0.0.13 and a destination address of 12.0.0.1. Switch 1 recognizes that 12.0.0.0 is the remote LAN for the branch office connection. Switch 1 translates the source address of the packet to 11.0.0.1 based on the NAT table. Switch 2 looks at the destination address of the incoming packet and translates it to 10.0.0.14, but the source address remains 11.0.0.1.



Note: When you have configured two branch offices for NAT, and assigned each office a NAT rule, you must remember that the accessible networks now use the translated addresses.

Figure 43 shows a simple branch office connection with two LANs, and a branch office tunnel across the internet. A pooled NAT rule is applied to Switch1, which connects the local network to the remote network via its branch office tunnel.

Figure 43 Overlapping address translation



Source	Destination
10.0.1.154	15.1.1.2
1305	x

Source	Destination
30.0.1.154	15.1.1.2
1305	x

NAT sets are collections of rules that make up a named set. You can create specific NAT sets for certain conditions and assign the sets as they are appropriate to the conditions. Typically, you apply NAT sets to branch offices that use either static or dynamic address schemes.

The NAT screen allows you to create NAT sets and edit or delete any currently defined NAT sets. To create new NAT sets, define a name and click on Create. When you edit an existing NAT set, the NAT Rules screen appears. The NAT Rules screen allows you to add new rules or edit existing rules. This screen lists the currently defined rules for a given set.

To create and edit NAT sets:

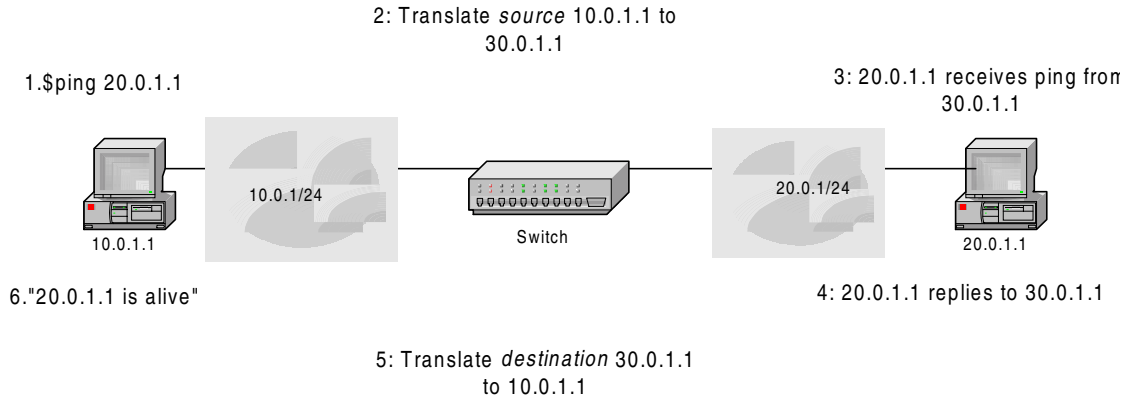
1 Click Profiles→NAT and enter the name of a new NAT set in the create box, then click on Create.

2. Click on the newly created NAT name, then click on Edit. The NAT Rules Edit screen appears.
3. Click on Add Rule. The NAT Add Rule screen appears.
4. Select a translation type, enter the internal start address in the appropriate box, enter the external address, and click on OK. The NAT Rules screen appears, showing the start and end addresses and the external address.

Interface NAT

Interface NAT is applied to IP packets going out of or coming into the switch via its physical interfaces causing either the source or destination IP address to be translated to another IP address, depending upon the NAT policy.

Figure 44 Interface NAT



NAT sets are applied to Interface NAT using the Services→Firewall screen. Interface NAT rules can be one of the following types:

- Static – for static mapping, an internal address range is mapped “one to one” to an external range.
- Port – for port mapping, the range of internal addresses is hidden behind a single external address. These external addresses are distinguished by using dynamically assigned port numbers.
- Pooled – for pooled mapping, an internal address is dynamically mapped to the next available address from the external address range.

To configure a new interface NAT policy:

1. Go to the System→Firewall NAT screen.
2. Click on the Interface NAT check box under the Enabled column. The NAT Set list box shows the currently selected NAT set and lists any additional NAT sets.
3. Click on the NAT Configuration link to go to the NAT Set page.
4. To create a new NAT set, type the name of the NAT set and click on the Create button. The NAT→Edit Set screen appears with (No rules currently defined) in the NAT Rules list box.
5. Click on the Add Rule button to add the rule. The NAT→Add Rule screen appears.
6. Under Translation Type, select Static, Pooled, or Port from the drop-down list.
7. Under Internal, enter the start (first available) and end (last available) addresses that represent the address pool that is used within the intranet.
8. Under External, enter the starting external address for the address range. If the translation type is Pooled, then you must also enter an ending external address.
9. Click on OK.
10. Click on Close to return to the NAT screen.
11. Click on the NAT Configuration link to go to the Profiles→NAT Sets screen to create a new NAT set if you do not want to use any of the existing ones.
12. You can use the Return to Firewall/NAT screen link to go back to this screen and apply the new NAT set. Selecting Interface NAT applies to non-tunneled traffic only. It does not affect the NAT sets applied to branch office tunnels. If any branch office tunnel NAT sets are assigned, they remain in effect for those branch office tunnels.
13. Reboot the switch.

To edit an existing NAT set:

1. Go to the System→Firewall NAT screen.
2. Click on the Interface NAT check box under the Enabled column. On the NAT screen, the NAT Set list box shows the currently selected NAT set and lists any additional NAT sets.

3. Select the NAT set that you want to edit and click on the Edit button. The NAT→Edit Set screen appears.
4. Select the rule within the policy that you want to edit and click on the Edit button.

To delete a NAT set:

1. Go to the System→Firewall NAT screen.
2. Click on the Interface NAT check box under the Enabled column. The NAT Set list box shows the currently selected NAT set and lists any additional NAT sets.
3. Select the NAT set that you want to delete and click on the Delete button. The NAT→Delete Set screen appears, asking if you are sure you want to delete the NAT set.
4. Click on OK. This returns you to the NAT screen.

NAT Traversal

NAT Traversal allows a number of devices on a private network to access the Internet simultaneously without each requiring its own external IP address. Most hotels and airports that provide Internet connectivity use NAT to connect to the Internet. You can enable NAT Traversal on the Services→IPSec screen. By default, NAT Traversal is disabled.

To use NAT Traversal, you must also define a UDP port that is used for all client connections to the switch. This port must be a unique and unused UDP port within the private network (supported range 1025 - 49151). By default, no UDP port is defined.

Chapter 7
Configuring Firewalls

The switch includes integrated firewall solutions that are designed to meet the needs of a variety of customers. The switch provides a choice of either the Contivity Stateful Firewall or the Contivity Interface Filter. For the Contivity Stateful Firewall to function, you must install the Firewall License Key.

With the Contivity Stateful Firewall, the switch can perform a variety of secure routing functions, depending upon how you set up the switch's routing capabilities. For example, you can configure the switch to securely route non-tunneled traffic from its private interface, through the firewall, and out its public interface. This configuration would enable users on the switch's private network to access the Internet without requiring a separate, dedicated router.

For many customers, the Contivity Interface Filter provides a cost-effective level of protection. The Contivity Interface Filter can only be disabled when Contivity Stateful Firewall is enabled. You typically enable it while you are migrating to the Contivity Stateful Firewall.

Existing customers who have Check Point FireWall-1 V4.0 from Check Point Software Technologies Ltd. can continue to use the separately licensed integrated firewall.

Contivity Stateful Firewall concepts

The Contivity Stateful Firewall provides a secure access point between an internal network and an external network, such as the Internet. The firewall allows you to protect your network and the information on it from unauthorized intrusion from external networks. The firewall provides a line of defense to allow acceptable traffic, as defined by your organization, and to drop all unacceptable traffic before it enters or leaves the network. It monitors packets and sessions to make decisions based on established rules to determine the appropriate actions to take.

By using s*tateful inspection*, the Contivity Stateful Firewall provides a high level of security, the fastest runtime, and the flexibility to define the rules to fit your environment. The firewall delivers full firewall capabilities, assuring the highest level of network security. To do this, the firewall examines both incoming and outgoing packets running against a common security policy. All service rules are interpreted on IP conversations (not packets) and are fully stateful. Service rules do not filter packets directly, but the firewall services determine how to process them based on the security policy defined. The firewall provides a user interface to help you determine the appropriate rules for your network.

The Contivity Stateful Firewall achieves optimum performance as a result of advanced memory management techniques and optimized packet inspection.

In addition, you can configure the firewall to log some or all significant events. This includes all connections over the network, such as all e-mail transactions, firewall status changes, and system failures. You can use the logged information to help enhance network security or track unauthorized use.

Stateful inspection

Some protocols are difficult to allow through a firewall securely using traditional filtering mechanisms. In FTP, for example, the control connection is typically created using a known port, but the data connection is over a random port. To allow an FTP data connection through a firewall without leaving a large number of open ports requires stateful inspection: packets are inspected at the application layer to determine which port the data connection is using. Traffic on that port can then be allowed to pass through the firewall for the duration of the FTP session.

Transport-level state inspection provides a number of ways to make TCP traffic more secure and more difficult for hackers to intercept. Stateful inspection of TCP consists of verifying the consistency of the TCP header as well as preventing some well-known TCP attacks. TCP sequence numbers are randomized to prevent sequence number guessing.

Stateful inspection of an application is unique for each application. Any non-predicted ports used by an application are validated and allowed through the firewall using stateful inspection. The following applications are inspected:

- FTP
- TFTP

- RCMD
- SQLNET
- VDOLive
- RealAudio

A *conversation* is created for all unique end-to-end communication. For instance, an FTP session between a client and a server can consist of several streams of traffic, with both data and control packets flowing back and forth. All of this traffic is contained in the same conversation. Since the Contivity Stateful Firewall also supports tunneled traffic, each unique tunnel (end user or branch office) would be identified by its own conversation.

Interfaces

The Contivity VPN Switch can have many interfaces. Each tunnel (end user or branch office) is a virtual interface, and all switches have two or more physical interfaces. Packets can be classified by the interface on which they arrive at the switch (the source interface) or the interface on which they leave the switch (the destination interface).

The rules in a policy can be constructed to either use or ignore this classification. If the rule designates “Any” as an interface, the rule ignores this classification. If the rule designates any other name as the interface, the rule uses this classification.

The rules in any policy can use the following terms to designate an interface:

- Any – Any physical interface or tunnel.
- Trusted – Any private physical interface or tunnel.
- Untrusted – Any public physical interface.
- Tunnel:Any – Any tunnel.
- For tunnels, you can specify either a group name for user tunnels or the specific branch office tunnel for branch office tunnels.
  - — Tunnel:/base – You can specify the specific branch office tunnel. For example, /base/mktng/tony refers to branch office tony in group /base/mktng.
  - — Tunnel:user – You can specify a group name for user tunnels. For example, /base/engineering refers to all user tunnels in that group.

- Interface name – The value of the Description field assigned to the physical interface on the System→LAN (or System→WAN) screen. If the description is blank, the interface name defaults to the value of the Interface field on the same screen.

Any physical interface can be marked as Private or Public on the System→LAN→Interfaces screen. By default, the LAN interface (Slot 0) is private and all other interfaces are public.

Filter rules

Filtering uses a set of rules to determine whether a packet should be allowed through the firewall. Typical options are either accept or drop the packet, which provides a degree of security for a network. The rules determine one of the following actions:

- Accept the packet
- Drop the packet
- Reject the packet by sending a reject to the source address
- Log the packet locally, which is a modifier for the previous three actions

For further explanation of the rules and how they are applied, see "Navigating rules" on page 226.

Anti-spoofing

Anti-spoofing is a method used to prevent a packet from forging its source IP address. Typically, the source address of each packet is examined and validated. Anti-spoofing performs the following checks:

- That the source address is not equal to the destination address
- That the source address is not equal to zero
- That the source address from an external network is not one of the directly connected networks

Attack detection rules

When a common attack is launched against corporate networks, the firewall should be able to detect these attacks. It should also drop any packets resulting from the attack, preventing denial-of-service as well as non-authorized intruders. The Contivity Stateful Firewall provides a defense against denial of service attacks with well-known prevention methods.

The Contivity Stateful Firewall protects against the following:

- Jolt2 is a fragmentation attack affecting Windows PCs by sending the same fragment repetitively.
- Linux* Blind Spoof attempts to establish a spoofed connection instead of sending final ACK with correct sequence number and with no flag set. Linux does not try to verify if the ACK is not set. The firewall drops any packet that does not have the ACK set.
- SYN flood can disable your network services by flooding them with connection requests. This fills the SYN queue, which maintains a list of unestablished incoming connections, forcing it not to accept additional connections.
- UDP Bomb sends malformed UDP packets that can crash a remote system.
- Teardrop/Teardrop-2 is a fragmentation attack that sends out invalid fragmented IP packets that trigger a bug in the IP fragment reassembly code of some operating systems.
- Land attack sends a TCP packet to a running service on the target host with a source address of the same host. The TCP packet is a SYN packet used to establish a new connection and is sent from the same TCP source port as the destination port. When accepted by the target host, this packet causes a loop within the operating system, essentially locking the system.
- Ping of death sends a fragmented packet larger than 65536 bytes, causing the remote system to incorrectly process this packet. This causes the remote system to either reboot or panic during processing.
- Smurf sends a large number of ICMP echo (ping) messages to an IP broadcast address with the forged source address of the intended victim. The routing device forwarding traffic to those broadcast addresses performs the IP broadcast to layer 2 broadcast, causing most network hosts to take the ICMP echo request and issue a reply to each, multiplying the traffic by the number of hosts responding.

- Fraggle sends a large number of UDP echo messages. On a multi-access broadcast network, there could potentially be hundreds of machines to reply to each packet.
- ICMP unreachable sends ICMP unreachable packets from a spoofed address to a host, causing the host to stop all legitimate TCP connections to the host that is being spoofed in the ICMP packet.
- Data flood sends a large amount of data to a system that can be used as a denial of service attack, which exhausts available resources and stops response to other user requests.
- FTP command overflow crashes FTP servers that contain buffer overflows for commands that take arguments. This applies to the user command, which means an attacker does not need a valid account to crash the system.

Contivity Stateful Firewall prerequisites

Before you start setting up your firewall, be sure you have the following:

- The management IP address of your switch. This address is found on the switch's System→Identity screen.
- The firewall license key. Go to the Admin→Install Keys screen. Type the key that you obtained from Nortel Networks in the box to the right of Contivity Stateful Firewall and click on the Install button. It is only necessary to install a key once on each switch. Click on the Delete button to remove the key.
- The name of the firewall, which is the name that is used by the Domain Name Service (DNS) server to identify the management address of the switch. This name is entered in the DNS Host Name field of the switch's System→Identity screen.
- The names and IP addresses of your switch's interfaces. These are found on the switch's Status→Statistics: Interfaces screen.

System requirements necessary to access the Contivity Stateful Firewall Manager:

- Supported operating systems and platforms include Solaris* (OS 2.6, 7, or 8) on a x86 or SPARC* platform and Microsoft Windows 95, 98, 2000, or Windows NT 4.0.

- Required software includes Java 2 Plug-in V1.3.0, available in the Java 2 Runtime Environment V1.3.0. The J2RE is available for automatic download on a Windows platform from the switch (refer to the Java 2 Runtime Environment Installation). J2RE installation files for Windows and Solaris are also available on the Nortel Networks CD in the tools/java directory.
- Supported browsers include Internet Explorer 4 and higher and Netscape Navigator* and Netscape Communicator* 4 and higher. Netscape 6 currently comes with a version of the Java 2 Plug-in that is not supported. If you wish to use Netscape 6, refer to the Netscape section of the Java 2 Runtime Environment Installation.

Installing Java 2 software

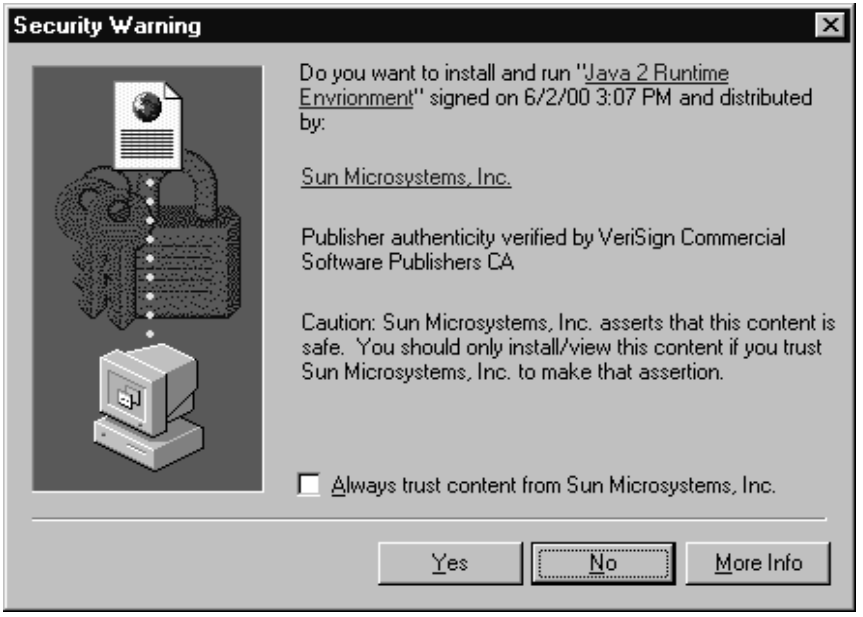
To access the Contivity Stateful Firewall Manager, the computer used to administer your switch must have the Java 2 Runtime Environment installed. There are two separate procedures that you can use to install the Java 2 software, depending on whether you use Internet Explorer or Netscape Navigator to access the switch.

Using Internet Explorer

To install the Java 2 software on Windows 9*x*, Windows 2000, or Windows NT from Internet Explorer:

1. Connect to the management IP address of the switch and log in.
2. Go to the Services→Firewall/NAT screen.
3. Click on the Manage Policies button. A popup window appears and tries to load the Contivity Stateful Firewall Manager.

4 When the Security Warning dialog box appears, click on Yes to install the Java 2 Runtime Environment.



The installation program will begin to download the software from the switch. This may take several minutes to load, depending on the speed of your connection to the switch.

5 When the installation program displays the Software Licensing Agreement, click on Yes to accept the agreement.

6 When the installation program asks for an installation location, accept the default location or choose an alternate installation location.

7 Click on Next > to finish the installation.

8 When the installation is complete, close all open Web browsers.

9 Reboot the computer for the changes to take effect.

Using Netscape

To install the Java 2 software on Windows 9*x*, Windows 2000, or Windows NT from Netscape Navigator:

1. Connect to the management IP address of the switch and log in.
2. Navigate to the Services→Firewall screen.
3. Click on the Manage Policies button. A popup window appears and tries to load the Contivity Stateful Firewall Manager. The Plug-in Not Loaded dialog box appears. (If this dialog does not appear, click on the long white or gray box that appears on the browser popup window.)
4. Click on the Get the Plug-in button to download the Java 2 Runtime Environment. The Java Plugin Download screen appears.



5. Click on the Download now link next to the Windows version of the Java Runtime Environment.
6. When the browser prompts you for a location to save the file, choose a download location and click on OK to continue. (This may take several minutes to load, depending on the speed of your connection to the switch.)
7. When the download finishes, go to the download location and double-click on the icon for the Java Runtime Environment.
8. When the installation program displays the Software Licensing Agreement, click on Yes to accept the agreement.

9 When the installation program asks for an installation location, accept the default location or choose an alternate installation location.

10 Click on Next > to finish the installation.

11 When the installation is complete, close all open Web browsers.

12 Reboot the computer for the changes to take effect.

Using Netscape 6

Netscape 6 currently includes a version of the Java 2 Plug-in that is not supported (version 1.3.01). To successfully load the Contivity Stateful Firewall Manager, you must use version 1.3.0. The following steps change the default plug-in from version 1.3.01 to version 1.3.0.

1 Install the Java 2 Runtime Environment 1.3.0 as described above and be sure to restart the computer.

2 Load the Java Plug-in Properties from Start>Settings>Control Panel>Java Plug-in.

3 Click on the Advanced tab.

4 Choose JRE 1.3.0 … from the list.

5 Click on Apply.

6 Close the window.

7 Close all open instances of Netscape.

8 Restart Netscape. The correct plug-in should be available.

Using Netscape on Solaris

The Java 2 Runtime Environment for Solaris is available on the Nortel Networks CD. The installation files and instructions are available for x86 and SPARC platforms.

To install the Java 2 software on Solaris (OS 2.6, 7, 8) from Netscape Navigator:

1 Ensure that a version of Netscape is installed on the computer.

2 Close all instances of Netscape if any are opened.

3 Go to the tools/java/solaris directory on the Nortel Networks CD.

4. Choose the subdirectory for the installed platform, either intel for x86 or sparc for SPARC.
5. Copy the binary (.bin) and the README files to the computer.
6. Follow the platform-specific installation instructions contained in the README file.
7. Set the NPX_PLUGIN_PATH environment variable to the directory containing the javaplugin.so file.

   For example, if the J2RE was installed in the /usr/j2re1.3 directory, the command to set the NPX_PLUGIN_PATH from the C shell would be:

```
setenv NPX_PLUGIN_PATH "/usr/j2re1.3/plugin/sparc on a SPARC
platform." From the C shell, it would be setenv NPX_PLUGIN_PATH
"usr/j2re1.3/plugin/i386 on a x86 platform."
```

8. Start Netscape and then close it.
9. Start Netscape again and the plug-in should be available.

Enabling firewall options

To enable your switch's firewall, select one of the following choices on the Services→Firewall/NAT screen and click on the OK button. You are then prompted to reboot your switch. You can select only one firewall choice at any one time.

- Contivity Stateful Firewall enables the switch's Contivity Stateful Firewall feature. When you enable the Contivity Stateful Firewall, you can run any combination of the following:
  - Contivity Stateful Firewall
  - Contivity Interface Filter
  - Interface NAT
  - Anti-spoofing
- Check Point FireWall-1 allows existing customers, who have Check Point FireWall-1 V4.0 from Check Point Software Technologies Ltd. to continue to use the separately licensed integrated firewall.
- No Firewall disables all firewall features on the switch. In this configuration, the switch performs VPN routing only.

After you change your Firewall selection on the Services→Firewall/NAT screen and click on the OK button, you are prompted to confirm your selection. If you selected No Firewall, click on the OK button.

If you selected either the Contivity Stateful Firewall or the Check Point FireWall-1, you must restart your switch before the firewall becomes active. After you enable the firewall support, you must configure the specified firewall.

The configuration procedures assume that you have already configured your Contivity VPN Switch (except for the firewall component) and that you have obtained the required firewall license.

Configuring the Contivity Stateful Firewall

To use the firewall on the Contivity VPN Switch, you must install the license key and enable the firewall service. Without the firewall enabled, the switch only handles traffic on the tunnel interfaces. When the firewall is enabled, the switch can also handle IP traffic between the physical interfaces and the tunnel interfaces.



Note: Before you activate the firewall on Services→Firewall/NAT screen, be sure that you shut off all traffic to the switch. You probably should do this off hours to prevent inconvenience to your users.

You must create rules for tunnel traffic before traffic on existing tunnel definitions is allowed. The Contivity Stateful Firewall uses the principle that whatever traffic is not specifically allowed is disallowed. The rule set of the active policy is applied to all traffic, including tunneled and non-tunneled traffic.Therefore, when the Contivity Stateful Firewall is first enabled, all traffic is disallowed until you configure rules specifically allowing certain types of traffic.

To get started using the firewall, follow these steps:

1 Go to the System→LAN screen. For each interface, click on Configure and enter a label in the Description field. You use this name to identify interfaces in the security policy rules. LAN represents the physical port interface to which you assign an IP address. Slot *n* Interface *n* represents an optional LAN card in expansion Slot *n* using Interface *n*.

For example, you could make Internet the description for Slot 1 Interface 1 and ServiceNet the description for Slot 2 Interface 1. The description is case sensitive and cannot be abbreviated when specifying the interface in the rules. If you do not specify a description, the default name for the interface is "Slot *n* Interface 1" (*n*=1 to 4) is case sensitive, and cannot be abbreviated.

2. Go to the Services→Firewall/NAT screen.
3. Check the box next to Contivity Stateful Firewall.
4. Click on OK. A confirmation screen appears, asking if you want to reboot.
5. Click on OK.
6. After the switch reboots, return to the Services→Firewall/NAT screen.
7. Click on Manage policies to load the Contivity Stateful Firewall Manager applet. The first time you do this on any workstation, you need to load the Java applet. You see the message "Retrieving policies."
8. Select the System Default policy, which is read-only.
9. Click on the View button to review this policy. The implied and post-implied rules are included with every new policy.
10. To exit the Contivity Stateful Firewall Manager, click on Manager→Exit. You can toggle the browser windows between the Contivity Stateful Firewall Manager applet and the Services→Firewall/NAT screen. If you use your browser to change other settings on the switch while the Contivity Stateful Firewall Manager applet is running, these changes are not reflected in the current Contivity Stateful Firewall Manager applet. Click on the Firewall icon in the Contivity Stateful Firewall Manager applet to refresh the list of policies and other switch settings. Any changes made in the Contivity Stateful Firewall Manager applet are not evident in the Services→Firewall/NAT screen until you save a policy.
11. After you exit the Contivity Stateful Firewall Manager applet, click on Refresh on the Services→Firewall/NAT screen.

You can apply (assign) a policy to the firewall in either the Contivity Stateful Firewall Manager applet or the Services→Firewall/NAT screen. The new policies you create are not automatically applied to the firewall. Only one policy at a time can be in effect on the firewall.

Note: You cannot import or export new policies. However, there are no restrictions on creating new policies.

Selecting event log options

The following options control the amount of detail recorded in the event log. None of these options apply to the system log.

- All – includes Traffic, Policy Manager, Firewall, and NAT
- Traffic – logs when flows and conversations are created or removed
- Policy Manager – logs firewall processes and when rules and policies are created
- Firewall – logs how the firewall handles packets within a flow
- NAT – logs NAT-related events
- Debug – creates special log messages intended for use only by Nortel Networks customer support

You can also select a connection number, which allows you to reserve memory for a maximum number of connections. Determining the optimum memory allocation makes it easier to tune your system for firewall traffic. Under the Connection Number section, type in any connection number. Below the box for the connection number is a suggested number range. The range displayed varies depending on the model and amount of memory for your switch. Each IPSec tunnel requires two connections. Setting the connection number to the minimum number in the number range minimizes the number of IPSec tunnels that your switch can handle and maximizes the size of dynamic memory available to the other firewall processes. Nortel Networks recommends that you set the number near the middle of the range displayed unless you have specific requirements that you need to consider.

Setting up policies

Firewall service consists of two primary components: the *service properties* and the *security policy.* The properties that define what service is being offered include a service name, the protocol (TCP, UDP, ICMP), and the port number (or range) on which the server listens.

Security policies consist of a set of rules that specify what traffic is allowed or denied. You can define custom policies when more complex security policies are needed and the standard policies are not sufficient. By customizing your policies, you can further refine the control over what traffic is allowed on your internal networks.

The firewall policies use standard actions, which represent the most commonly used policies. You can use the following actions individually or combine them in a single rule:

- Allow (accept) access if traffic arrives on a specified interface (source interface rule)
- Allow access if traffic from the firewall is to a specified interface (destination interface rule)
- Allow access if traffic is from a specific host or group of hosts
- Allow access if traffic is going to a specific host or group of hosts
- Allow access if traffic is a specific service
- Deny access (drop) or deny access and send a reject packet (reject) if specified as modifiers to the previous actions

A set of rules defines a specific security policy. A rule defines whether communication should be accepted or rejected (or logged) based on its source, destination, and service.

You must create rules for tunnel traffic before traffic on existing tunnel definitions will be allowed. The Contivity Stateful Firewall uses the principle that whatever traffic is not specifically allowed is disallowed. The rule set of the active policy is applied to all traffic, including tunneled and non-tunneled traffic.Therefore, when the Contivity Stateful Firewall is first enabled, all traffic is disallowed until you configure rules specifically allowing certain types of traffic.

Creating and editing firewall policies

Access control parameters are implemented through the graphical user interface or the command line interface (CLI). Using either interface, you can configure the following:

- Network objects
- Service objects
- Rules

See *Reference for the Contivity VPN Switch Command Line Interface* for a list of commands you can use on the CLI.

The security policy is a set of rules describing the required behavior of the service. You specify all rule fields for service policies using service objects. Each rule consists of a combination of network objects, services, actions, and logging mechanisms.

Navigating rules

The firewall policy edit screen allows you to add, delete, and modify the rules in a policy. This screen is divided into the following rule groups:

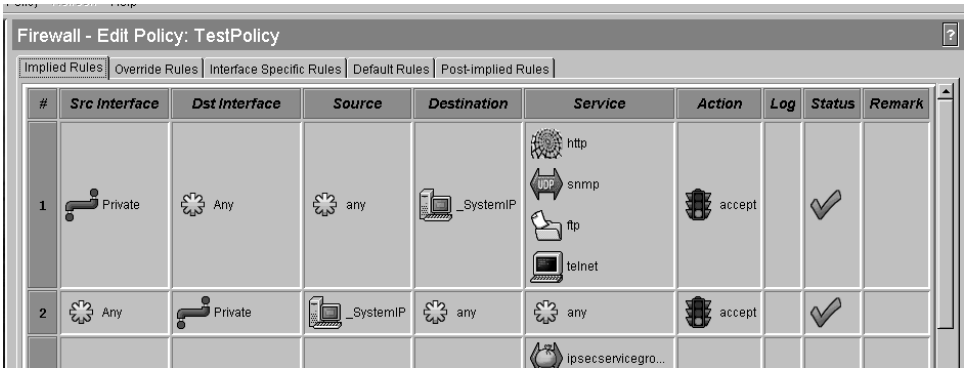
- Implied rules
- Override rules
- Interface-specific rules
- Default rules
- Post-implied rules

Rules are divided into the following types:

Implied rules

Implied rulte are processed first by the firewall. These rules permit unnel termination and access to the management interface. They are derived from the Services→Available screen. The implied (Figure 45) rules are rules that are defined by the system. You cannot modify them using the Contivity Stateful Firewall Manager. These rules are read-only and you cannot modify any part of the rule. They are generated from the Services→Available screen on the switch. These sections are for display purposes only.

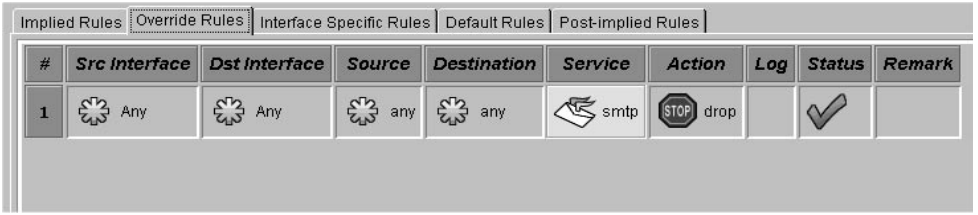
Figure 45 Implied rules



Override rules

Rules that are specified in the override rules (Figure 46) are the first set of rules specified in the policy. The purpose of these rules is to allow you to quickly override the rest of the rules in the policy, possibly for a short period while debugging a problem. These rules do not specify a specific interface in the source or destination interface column. You can only select from the interface groupings (Any, Trusted, Untrusted, Tunnel:Any).

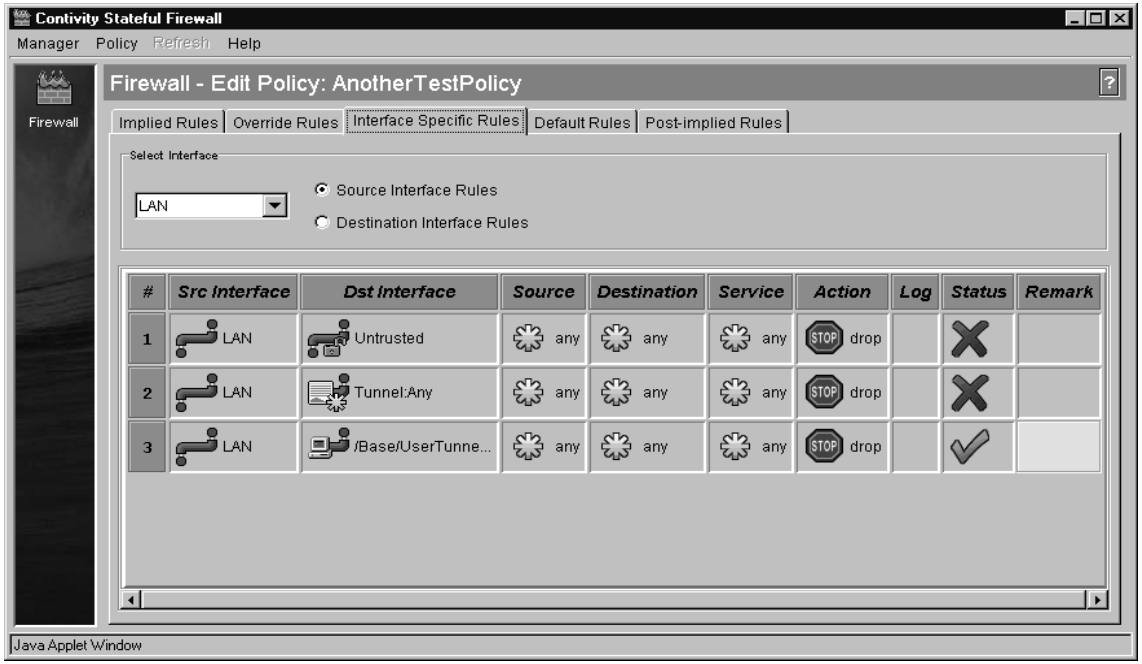
Figure 46 Override rules



Interface-specific rules

Interface-specific rules are rules that only apply to packets that enter or leave the switch through one specific interface (physical or tunnel). The interface-specific rule section (Figure 47) is divided into two types: source rules and destination rulte. Source rules are those rules that define the selected interface as their source and destination rules are those rules that define the selected interface as their destination. The interface names correspond to the names that you configured on either the System→LAN or System→WAN screen on your switch. On these screens, you can also specify whether these interfaces are private or public. For tunnels that are also interfaces, you can specify either a group name for user tunnels or the specific branch office tunnel for branch office tunnels. The interface-specific rule section only displays one interface at a time. To view all of the interface-specific rules, select All Interfaces.

Figure 47 Interface-specific rules



Selecting an interface from the selection box displays those rules that apply to the selected interface. If you select Branch Office Tunnel or User Tunnel from the primary selection box, a new selection box appears next to it that contains the branch office tunnels or user groups. This box disappears if you later select a physical tunnel.

The interface-specific section is also divided into two types: source rules and destination rules. Source rules are those rules that define the selected interface as their source and destination rules are those rules that define the selected interface as their destination. You can toggle between the two types of rules for a particular interface by clicking on the Source Rules or Destination Rules radio button.

When you add a rule to the interface-specific section, the rule is created as the same type as the one you are currently viewing.

If you are viewing the source rules (Figure 48) and you add a new rule from that screen, it is a source rule. The source interface field on a source rule is set to the current interface and is read-only. The destination interface field can contain any physical interface, tunnel, or interface grouping.

If you are viewing the destination rules (Figure 49) and you add a new rule from that screen, it is a destination rule. The destination interface field on a destination rule is set to the current interface and is read-only. The source interface field can only contain interface groupings (Any, Trusted, Untrusted, Tunnel:Any).

Figure 48 Source interface rules

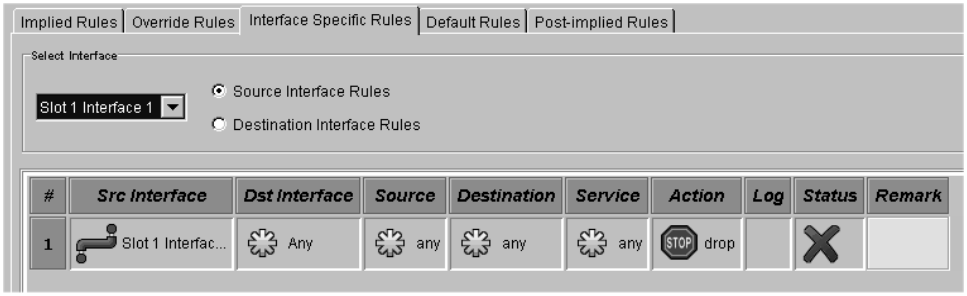
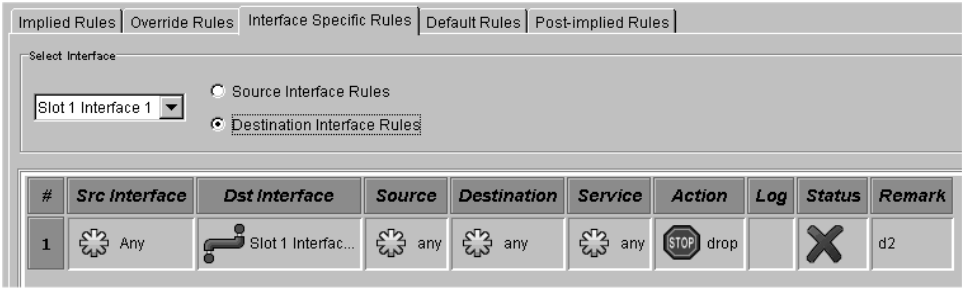


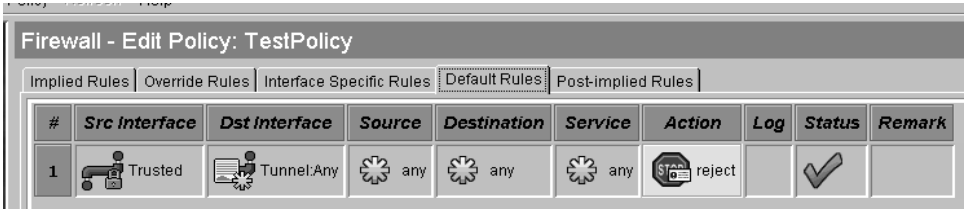
Figure 49 Destination interface rules



Default rules

A default rule (Figure 50) is a rule that is applied to all packets for which a rule was not found in the set of interface-specific rules. These rules cannot specify specific interfaces in the source or destination interface columns and can only use interface groupings (Any, Trusted, Untrusted, Tunnel:Any).

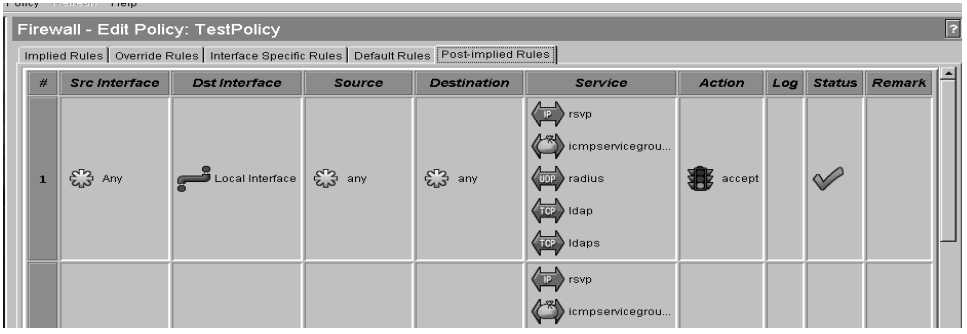
Figure 50 Default rules



Post-implied rules

Post-implied rules allow RSVP, ICMP, LDAP, and RADIUS access and access by routing protocols on trusted interfaces. You cannot configure these rules and the last post-implied rule is to drop all traffic that was not handled by a prior rule. However, you can override these rules by specifying a rule in one of the three rule groupings (override, interface-specific, or default). The post-implied (Figure 45) rules are rules that are defined by the system. You cannot modify them using the Contivity Stateful Firewall Manager. These rules are read-only and you cannot modify any part of the rule. They are generated from the Services→Available screen on the switch. These sections are for display purposes only.

Figure 51 Post-implied rules



Performing actions on rules

The actions that you can perform on rules are controlled by menus that you access by right-clicking on an option. Each of these menus controls a different aspect of the rule.

Header row menu

Right-clicking on any of the header cells brings up the Header row menu (Figure 52). This menu contains one item, Add New Rule. This menu item allows you to add a new rule to the top of the list. The new rule appears in position one and all existing rules increment by one.

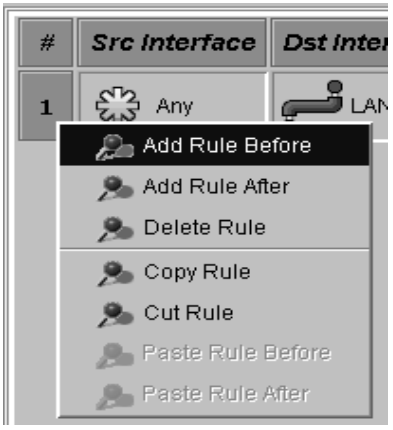
Figure 52 Header row menu



Row menu

Right-clicking on the number field in a row activates the row menu (Figure 53). This menu allows you to add a new rule at a particular location, delete the specific rule, and perform cut/copy/paste operations on a rule.

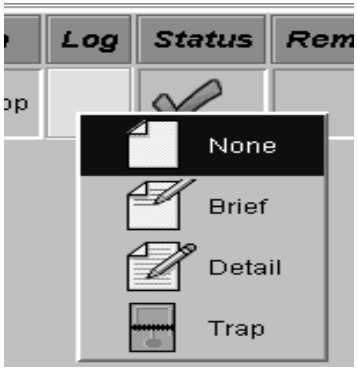
Figure 53 Row menu



Cell menus

Cell menus (Figure 54) are specific to each cell and you trigger them when you right-click on an individual cell. There are two types of cell menus: option menus and procedure menus. Option menus provide a list of possible values for the cell. These menus are similar to a drop-down list box. When you click on one of the items, the selection is displayed in the cell.

Figure 54 Cell menu (option)



Procedure menus (Figure 55) provide a list of operations that you can perform on the cell, such as Add and Edit. When you click on one of the items, either the operation is performed immediately (such as Copy) or an additional dialog box appears, prompting you for more information (such as Add).

Figure 55 Cell menu (procedure)



Rule columns

Each rule within a firewall policy has the same attributes, which are specified by the column headers. All of the columns behave the same as in all three sections, except the Src Interface and the Dst Interface columns. These two columns behave differently depending on the section where the rule appears. The following sections describe the columns within a firewall rule:

#

This column specifies the ordering of the rules within the section. The order only applies to the section in which the rule appears and does not have meaning across the entire policy. If you log a rule, this number (#) is included in the log information.

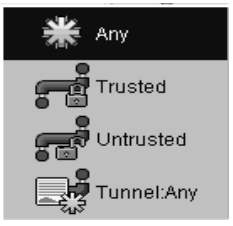
Src interface and Dst interface

These columns specify the source and destination interfaces for the rule. Right-clicking on the cell displays an option menu containing possible interfaces. What appears in this option menu depends on which section of the Firewall policy the particular column appears in. For the Override and Default rules, the interfaces may only be interface groupings. These groupings are:

- Any – Any physical interface or tunnel
- Trusted – Any private physical interface or tunnel
- Untrusted – Any public physical interface
- Tunnel:Any – Any tunnel, excluding any physical interfaces

Figure 56 shows an example of a rule column.

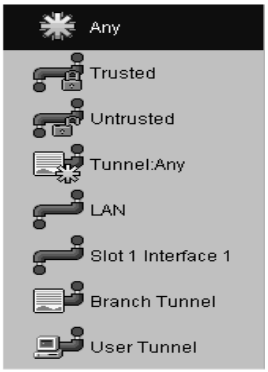
Figure 56 Rule column



For the interface-specific rules, you can specify the interfaces as either groupings or individual interfaces.

Figure 57 shows a rule column for interface-specific rules.

Figure 57 Rule column for interface-specific rules



Clicking on the user tunnel (Figure 58) or branch office (Figure 59) menu items displays the tunnel selection dialog box. This dialog box allows you to select a specific tunnel (branch office or user tunnel).

Figure 58 Tunnel Selection dialog box for a user tunnel

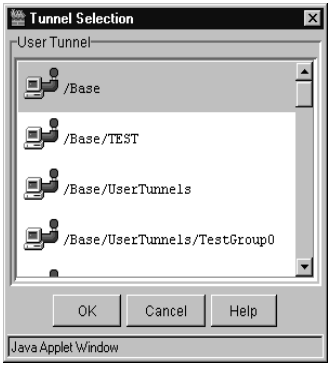
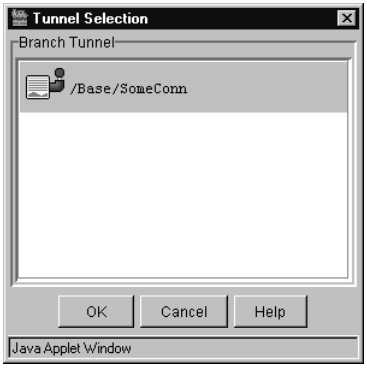


Figure 59 Tunnel Selection dialog box for a branch office tunnel



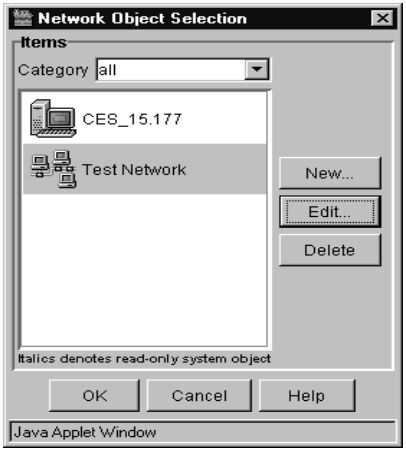
Source and Destination

These columns specify the source and destination addresses for the rule. You can modify these attributes by right-clicking on a column in the cell, which then brings up a procedure menu.

Clicking on Add displays the Network Object Selection dialog box (Figure 60). In this dialog box you define and apply a new network object. You can create the following network objects: host, network, IP range, and group (a collection of these objects).

Note: It is possible to add more than one source or destination address to a rule.

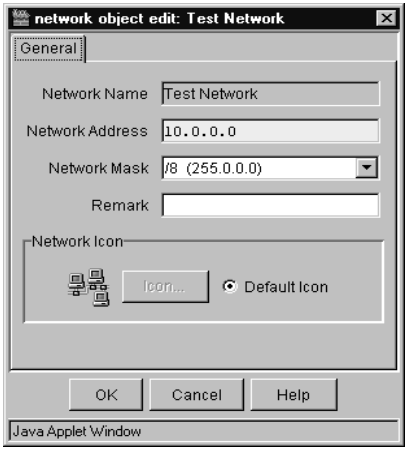
Figure 60 Network Object Selection dialog box



Italicized objects in the list are read-only and cannot be modified. The New, Edit, and Delete buttons in this dialog box allow you to create, edit and delete network objects.

Clicking on Edit displays the Network Object Edit dialog box (Figure 61). This dialog box allows you to modify the attributes for the selected network object.

Figure 61 Network object edit dialog box



Clicking on Delete removes the selected network object. If the object that you want to delete is the last object, it returns to the default value.

Clicking on Copy, Cut, or Paste performs those operations on the current network object.

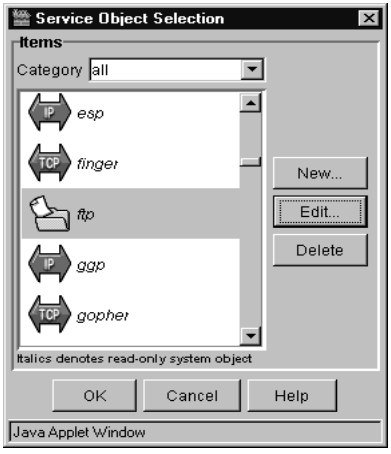
Service

This column specifies which services are handled by the selected rule. Right-clicking on the cell displays the standard procedure menu (Add or Edit).

Clicking on Add triggers the Service Object Selection dialog box (Figure 62), which allows you to define and apply a new service object. You can create the following service objects: TCP, UDP, ICMP, IP protocol, and object groups (a collection of these objects).

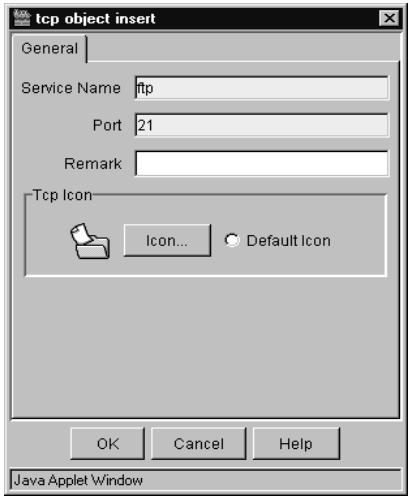
Note: It is possible to add more than one service to a rule.

Figure 62 Service Object Selection dialog box



Italicized objects in the list are read-only and cannot be modified. The New, Edit, and Delete buttons in this dialog box allow you to create, edit and delete network objects. Clicking on Edit displays the Service Object Edit dialog box (Figure 63). This dialog box allows you to modify the attributes for the selected service object.

Figure 63 tcp object insert dialog box



Clicking on Delete removes the selected service object from the cell. If the object to be deleted is the last object in the cell, the cell returns to its default value (in this case, Any).

Clicking on Copy, Cut, or Paste performs those operations on the current service object.

Action

This column (Figure 64) specifies the action that occurs when the rule is activated. Right-clicking on the cell displays an option list containing three items: Accept, Drop, and Reject. Clicking on one of these items sets the cell to the selected state.

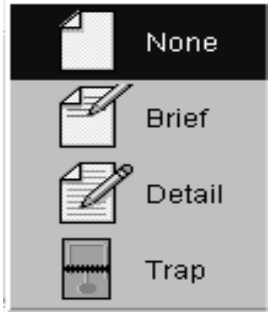
Figure 64 Action column options



Log

The Log column (Figure 65) allows you to specify the logging level for this rule. Right-clicking on this cell brings up an option list containing the following logging levels: None, Brief, Detail, and Trap.

Figure 65 Log column options



Status

This column (Figure 66) specifies the status of the particular rule. The status can be either Enabled or Disabled.

Figure 66 Status column options



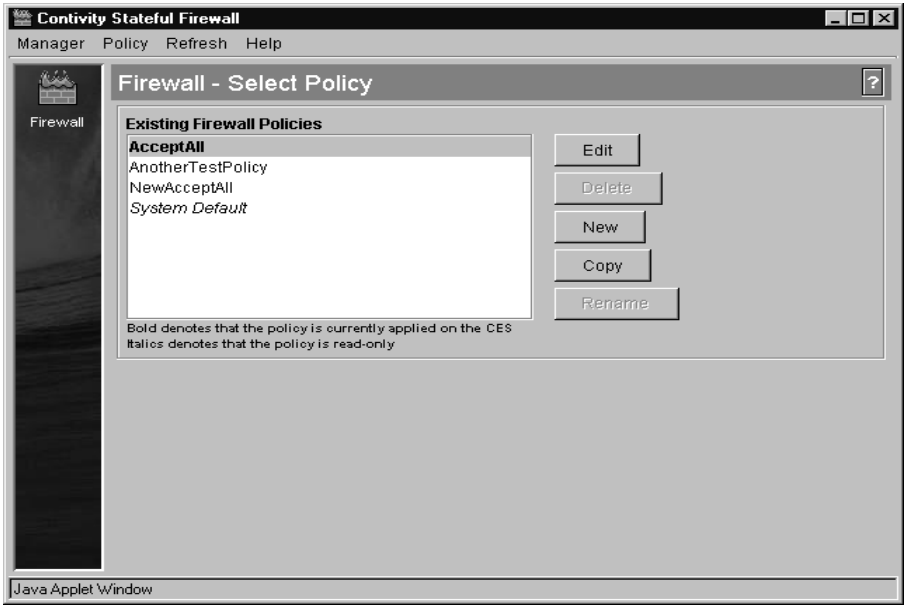
Remark

This column allows you to attach a remark to a particular rule. When you right click on Remark and choose Add or Edit remark, a dialog box appears where you can type a comment.

Creating policies

The Firewall - Select Policy screen (Figure 67) allows you to create, edit, delete, copy, or rename a firewall policy. Bold denotes the policy that is currently applied to the Contivity VPN Switch and italics denotes read-only policies. The System Default policy is always listed. This read-only policy defines the firewall behavior when no user-defined policies have been applied or when the selected policy is not available (for example, if LDAP is not available).

Figure 67 Firewall - Select Policy screen



You can add, modify or delete the rules for a policy. The screen is divided into sections that denote the rule groups. If you upgrade to a new version, the current configuration is copied to the new version. If you either upgrade or downgrade to a previous version, the configuration file and LDAP database are saved and restored.

Adding a policy

To add a new policy:

1. Click on the New button. The New Policy dialog box appears and prompts you for a name for the new policy.
2. Enter the policy name. The name must begin with a letter and may not contain the : + =] , ; " characters.
3. Click on OK to go to the Policy Edit screen, which has a blank firewall policy, or click on Cancel to return to the policy selection screen.

Refer to "Navigating rules" on page 226 to set up rules for your new policy.

Deleting an existing policy

You cannot delete a read-only policy or the policy that is currently applied to the switch. If you select one of these policies, the Delete button is not enabled. To delete an existing policy:

1. Select the policy that you want to delete and click on the Delete button. The delete policy confirmation dialog box appears.
2. Click on OK to delete the selected policy.

Copying an existing policy

To copy a firewall policy:

1. Select the policy that you want to copy.
2. Click on the Copy button. The copy dialog box appears.
3. Enter a name for the copied policy.
4. Click on OK.

The new policy appears in the list of policies in the firewall policies screen. This policy contains the same rules as the policy from which it was copied.

Renaming an existing policy

You cannot rename a read-only policy or the policy that is applied to the switch. If you select a read-only policy, the Rename button is not enabled. To rename an existing firewall policy:

1. Select the policy that you want to rename.
2. Click on the Rename button. The Rename dialog box appears.
3. Enter the new name of the policy.
4. Click on OK.

Creating a new policy example

Complete the following steps to configure your firewall policies.

1. Log in to the switch. Select Services→Firewall/NAT. The Firewall/NAT screen appears.
2. Under Configuration, click on the Enabled radio button next to Contivity Firewall.
3. Click on the Manage Policies button. The Firewall - Select Policy screen appears.
4. Click on New to create a new policy. The New Policy dialog box appears.
5. Enter the policy name and click on OK. The name must begin with a letter and may not contain the : + =] , ; " characters. The Firewall - Edit Policy: <*policyname*> screen appears with no rules defined. In this screen, you can add, delete, and modify the rules for the policy.
6. You can select the rule group as follows:
   - Implied rules (view only)
   - Override rules
   - Interface-specific rules
   - Default rules
   - Post-implied rules (view only)
7. Select the Interface Specific Rules tab.
8. Select the interface from the drop-down list.
9. Select either Source Interface Rules or Destination Interface Rules.
10. Right-click on the appropriate cell to add a new rule.
11. Repeat these steps to add more rules.
12. Select Policy and click on Save Policy to save your changes.

13 When the policies are saved, go to the Manage menu and click on Close Manager.

Successful completion of these steps indicates that your switch's firewall is functioning and that the switch's routing patterns are available.

Verifying your configuration

When you complete the configuration tasks for the firewall, you should check the switch's routing patterns. To verify that the firewall functions properly, you can use a procedure similar to the following:

1 Make sure the firewall is using a security policy that allows the type of traffic you use for the test (or you can use an Accept All policy for the testing).

2 Verify public-to-private traffic. Perform an FTP operation from a host on the public side of the switch to a host on the private side.

3 Verify private-to-public traffic. Perform an FTP operation from a host on the private side of the switch to a host on the public side.

4 Verify tunnel-to-internal network traffic. Connect a remote Contivity VPN Switch system to the local switch. From the client, access a Web page on the internal network.

5 Verify tunnel-to-Internet traffic. Connect a remote Contivity VPN Client system to the switch. From the client, access a Web page on the Internet.

Chapter 8
Configuring Servers

This chapter describes how to configure the following authentication servers for users who are tunneling into the switch:

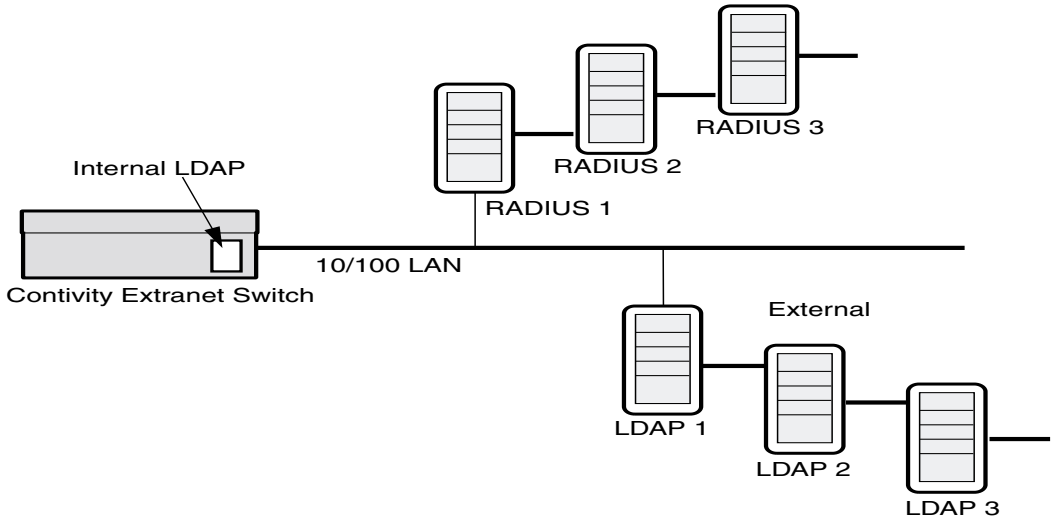
- Internal LDAP stores the group and user profiles on the internal server of the switch. External LDAP is a standard protocol for Internet directory services that is based on directory entries. A directory service is a central repository of user information.
- LDAP proxy server uses an existing LDAP database of users on an LDAP server for authentication.
- External RADIUS is a distributed security system that uses an authentication server to verify dial-up connection attributes and authenticate connections.
- RADIUS accounting logs user sessions with RADIUS-style records containing detailed connection statistics.
- The switch can function as a simple RADIUS server.

Authentication services

The switch supports LDAP and RADIUS authentication servers. The switch always attempts to authenticate a remote user against the LDAP internal or external profiles. If a user ID and password are found, the switch uses the attributes that are defined for that user's group. The switch can also authenticate against an LDAP proxy server or RADIUS database. When using either of these for authentication, you can create groups to take advantage of different profiles, or you can simply assign all LDAP proxy server or RADIUS server users into a single default group.

Figure 68 shows a switch and authentication servers.

Figure 68 Contivity VPN Switch and authentication servers



The user ID (UID) is checked against the LDAP profile database. If the UID is found in the LDAP database, the user is assigned to a group and acquires that group's attributes. Next, the password is checked, and if it is correct, the switch allows a tunnel to be formed.

If the UID is not in the profile LDAP (internal or external) database and if you specified RADIUS as the next server to check, the UID and password will be checked against the RADIUS database. If the UID and password are correct, the switch checks to see if the RADIUS server returned a class attribute. The RADIUS Class Attribute is treated as an LDAP group name. If a RADIUS class attribute is returned, and it names an existing LDAP group, the switch applies the attributes of this group to this user's session, and forms a tunnel. If the group name does not exist, the user is given the RADIUS default group's attributes. If the UID and password are incorrect, the switch rejects the user request.

IPSec behaves the same as a PPTP session; the RADIUS server defines the group for that user after authentication using the class attribute group identifier. The only difference between IPSec and PPTP is that in the event the RADIUS server does not return a class attribute, the group associated with the IPSec group ID is used

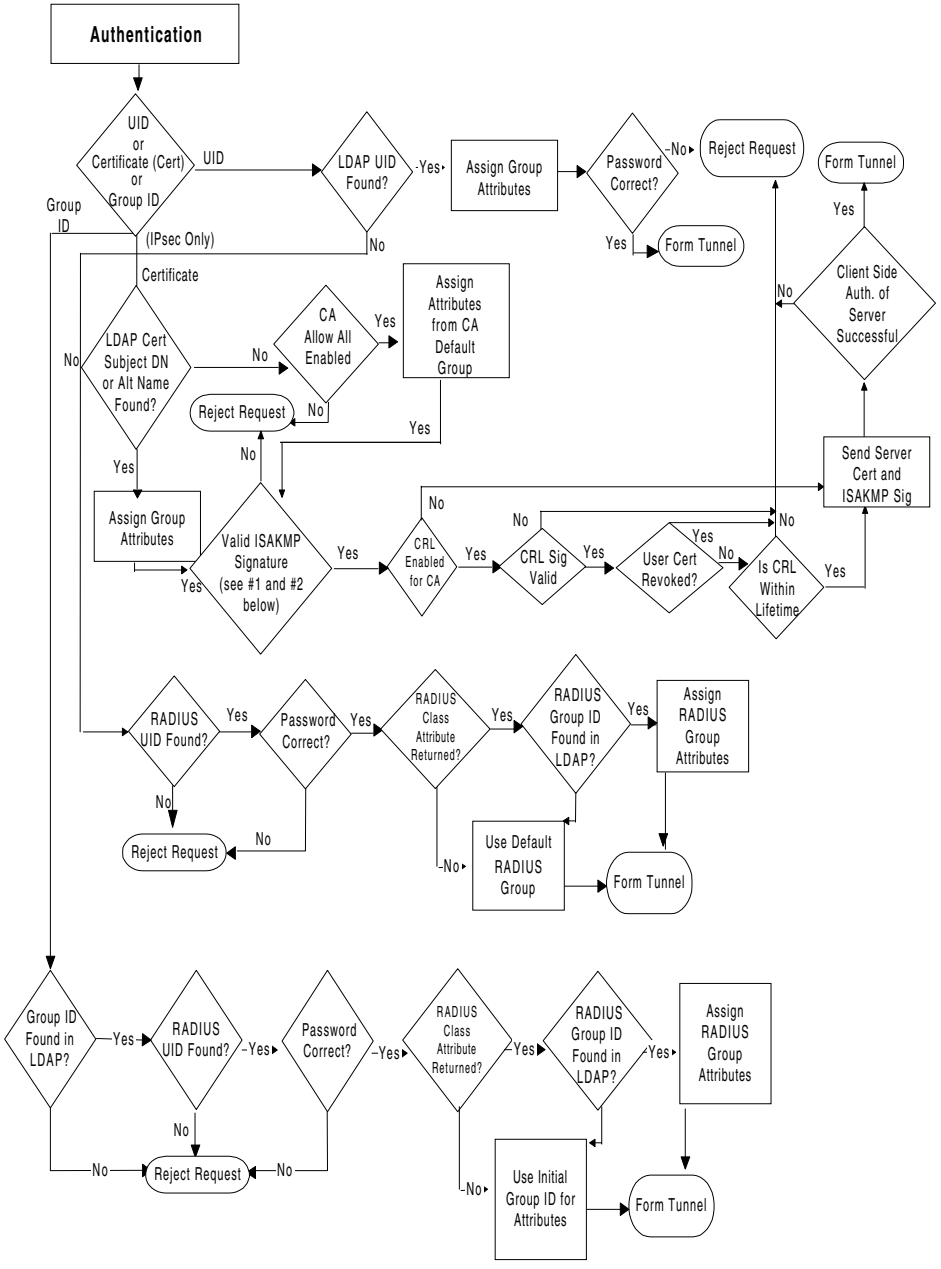
instead of the RADIUS default group. You configure the IPSec Group ID in the Authentication section of the Profiles→Groups→Edit→Configure IPSec screen. You configure the PPTP default group on the Servers→RADIUS Auth screen, RADIUS Users Obtain Default Settings from the Group option.

The group that the user is bound to must allow the authentication method that is used when the session is started.

If the UID is not in the profile LDAP (internal or external) database and if you specified LDAP authentication as the next server to check, the UID and password will be checked against the LDAP database.

Figure 69 illustrates the steps in the Nortel Networks user validation.

Figure 69 Authentication server validation flowchart



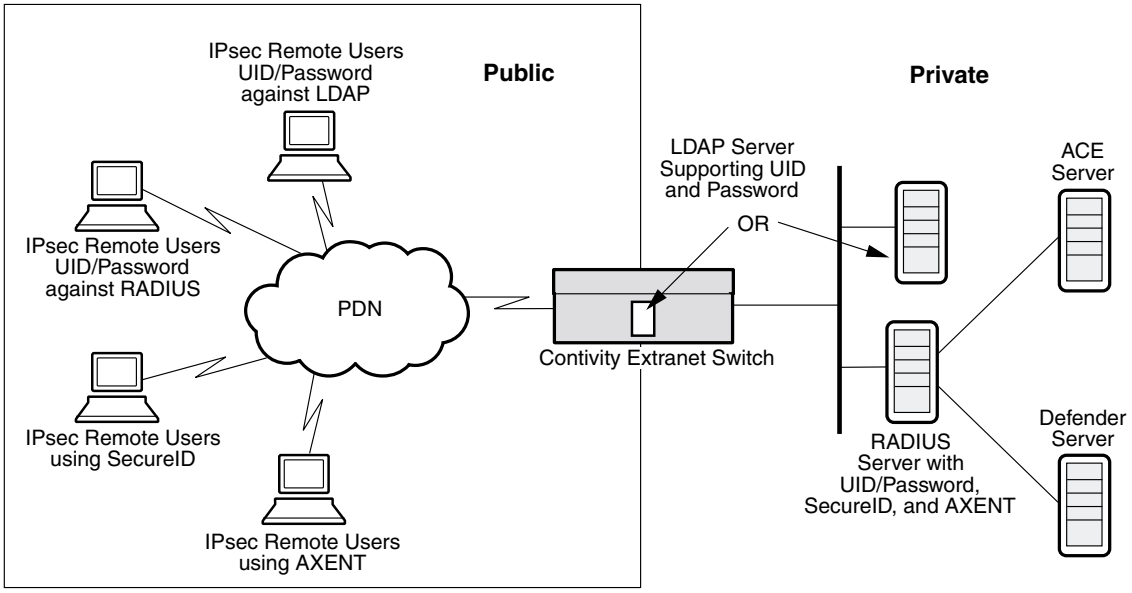
LDAP, RADIUS, and the IPSec client

Authentication using the Nortel Networks IPSec client provides several options for remote users connecting through a switch. These include:

- UID and password authentication against an LDAP database
- Group password authentication using a RADIUS server
- Token Security methods (RSA SecurID* and AXENT Omniguard/Defender*)

Figure 70 shows IPSec client options.

Figure 70 IPSec client authentication options



Each authentication option has the following positive attributes:

- Diffie-Hellman key exchange (ISAKMP/Oakley Aggressive Mode) to build the Security Association (SA).
- The user name and the password are never transmitted in the clear; a cryptographic hash function (SHA-1) is used to protect the user's identity.

- Mutual authentication between the client and the switch using a keyed hash algorithm (HMAC).
- Protection against authentication replay attacks through the use of session "cookies."

The IPSec client and the switch support authentication against an LDAP database.



Note: The IKE encryption algorithm does not affect the choice of the encryption algorithm used to encrypt data in IPSec. For example, you can use DES to encrypt the IKE exchanges, and then negotiate Triple DES for use in IPSec.

The type of encryption used to encrypt the data passed through an extranet connection can affect performance. For IPSec, Triple DES is a stronger encryption method than DES and so the computational requirements of encrypting data with Triple DES are higher. Only the network administrator can make this change.

You can set the global IKE encryption and Diffie-Helman group for IPSec on the Services→IPSec screen. The following lists the choices for use with IPSec:

- Both 56-bit DES with Group 1 and Triple DES with Group 2
- Triple DES with Group 2 (1024-bit prime)
- 56-bit DES with Group 7 (768-bit prime)

This database can reside internally in the switch or it can reside on a server. You can configure the switch to authenticate the remote access user with a simple ID and password. In this mode, the user ID is protected from network snooping by hashing the user name with a standard one way function (SHA-1). Mutual authentication is achieved by each knowing the password hash and the use of data "cookies" by each side to prevent replay attacks.

To enable IPSec authentication against the LDAP database:

1. Select this option in at least one group profile from the Profiles→Groups→Edit→IPsec screen.
2. Select User name and Password as one of the Authentication options.
3. Add users to the group for authentication from the Profiles→Users screen; select Add User for the appropriate group.

LDAP database servers

The Lightweight Directory Access Protocol (LDAP) is a standard protocol for Internet directory services that is based on directory entries. A directory service is a central repository of user information; for example:

- Groups
- Users
- Filters
- Services

An entry is a collection of attributes that has a distinguished name (DN), which refers to the entry unambiguously. Each entry attribute has a type and one or more values. Types are typically mnemonic strings; for example, **cn** represents common name and **mail** represents e-mail address. The values depend on the attribute type. For example, a mail attribute value might resemble jchirac@elysee.france.gov.

LDAP directory entries are arranged in a hierarchical tree-like structure that reflects political, geographic, and organizational boundaries. Country entries appear at the tree top. The next entries represent states or national organizations. The third-branch entries represent people, organizations, servers, files, or any other readable database entry. LDAP allows you to read, search, add, and remove information from the centralized database.

The switch uses an LDAP server to centrally store remote access profiles and corporate networking details such as the addressing mechanism; for example, group attributes including hours of access, filters, and authentication servers. The switch queries the LDAP server for access information when a user establishes a tunnel connection. The LDAP query can be serviced locally by the internal LDAP server; or it can be redirected to an external LDAP server, such as the Netscape Directory Server or Novell NetWare* Directory Services.

The switch's internal LDAP server is designed to operate only as an internal server and does not respond to external queries. Therefore, two or more switches cannot share the same internal LDAP database. To allow sharing between switches, and to take full advantage of LDAP-based directory service replication and centralization, you should use a dedicated directory services offering from one of the above vendors.

The switch synchronizes its cache every 15 minutes. For example, if you delete a user from an external LDAP database it can take up to 15 minutes before all of the switches recognize the change. Additionally, the LDAP server's status is recorded in the Event Log every 15 minutes.

To configure internal LDAP:

1. Go to the Servers→LDAP screen. The internal LDAP server is internal to the switch. If you are using more than one Contivity VPN Switch or if you are using LDAP authentication for other network services, you should consider using an external LDAP server.
2. Click to enable access to the external LDAP server. The internal server is disabled if you enable an external LDAP server.
3. Under General Configuration, click to remove the user's fully qualified ID suffix from the UID before sending it to the RADIUS server. A user ID and suffix, where Rcole is the UID and acme.com is the suffix is cole@acme.com. Specify the character that separates the suffix from the UID.
4. Click on Stop Server or Start Server, as appropriate, when you intend to back up or restore a configuration, or after you have completed the restoration of a configuration. The LDAP server must be stopped before you can perform the backup and restore procedures.

LDAP server authentication

The switch supports authentication against an existing LDAP database rather than creating a second user database for use with the switch. The server can reside on either a private or public network that is connected to the switch. If the user cannot be authenticated through the existing methods, the LDAP server authentication database is used. LDAP server authentication also allows you to configure the type of authentication methods that are allowed to access the LDAP server. There are five options available: PAP (Bind authentication), PAP, CHAP, MS-CHAP, MS-CHAP V2 (Bindname authentication).

To configure LDAP server authentication on the switch:

1. Go to the Servers→LDAP Auth(entication) screen and select Enable Access to LDAP Authentication.

2. Under LDAP proxy servers, enter a base distinguished name (DN) for the servers. This is usually in the form of: ou=organizational unit, o=organization, c=country.
3. For the remote LDAP server, enter the Master, Slave 1, and Slave 2 LDAP server host names or IP addresses. Should the master server become unavailable, the switch attempts to initiate a connection with the slave servers.
4. In the Connection section, enter the port number (default 389) and the associated SSL port number (default 636) that your LDAP server listens to queries on.
5. Enter the bind distinguished name (DN), which is the LDAP equivalent of a user ID and is required to access the base DN and its subentries. Leave this field blank if your LDAP server allows anonymous access.
6. Enter the bind password, which can consist of up to 32 characters. This password allows the switch to prove its identity (the bind DN) to the LDAP server.
7. In the Username/Password Attributes section, you can have case-insensitive character strings that are allowable in LDAP search filters. The default value for each field is blank. If you do not supply a specified value, the authentication will fail. Enter the user name and password. To enable additional policy checking specify an LDAP search filter, which can be a case-insensitive character string that forms an allowable LDAP search filter (default is blank).
8. In the User Policy Attributes section, specify attributes used to store the Contivity group, static IP address/netmask, and customized user filter. These fields can hold case-insensitive character strings that are allowable in LDAP search filters. The default value for each field is blank. Without a specified attribute name, the LDAP proxy server will not attempt to extract this information.
9. Click on the SSL Encryption hyperlink to go to the LDAP server SSL encryption screen. This allows you to select the encryption types the switch uses during negotiation with the external LDAP server. If this server does not support one of the selected encryption types, then the connection is not established.
10. Go to the Services→IPSec screen and use the Swap Server Order 2 and 3 to change the order in which the switch applies authentication.

11 Go to the Profiles→Groups screen and either add or select the group that you want to be the default group for LDAP users (this is the group a user is assigned to if the LDAP server does not send back a class attribute).

12 Go to the Profiles→Groups→IPSec Configure screen. In the Authentication area, click on the Configure button. The Authentication method for the group is already selected. Select the default server certificate from the list, which allows all users with certificates from this CA. Enter the Group ID and Group Password. Consider using the LDAP group name as the default group, because you must remember a default name once you enter it.

To configure IPSec and LDAP on the client:

1 In the Contivity VPN Client UI, go to Options→Authentication Options, and select Group Security Authentication.

2 Enter the group ID and group password.

3 Select one of the group authentication options.

4 Click on OK.

RADIUS authentication servers

RADIUS is a distributed security system that verifies connection attributes and authenticates connections.

The RADIUS Authentication Servers screen allows you to configure up to three servers for remote authentication. It is imperative that the RADIUS servers contain the same user data. The alternative RADIUS servers are used only when no response is received from the primary RADIUS server.

Most RADIUS servers support CHAP and PAP authentication, and some can now support MS-CHAP (Funk, for example).



Note: If you require PPTP-encrypted tunnels and RADIUS authentication, then you must use a RADIUS server that supports MS-CHAP. The alternative is to use an LDAP server for PPTP authentication.

RADIUS-based authentication

The switch supports authentication against a RADIUS server. This server can reside on either a private or public network that is connected to the switch. To enable RADIUS authentication, you must configure the switch with the RADIUS server host name, port number (typically 1645, but port 1812 is the RFC standard), and a shared secret. Access the switch management screen from the Servers→RADIUS Authentication screen.

The RADIUS Authentication screen also allows you to configure the type of authentication methods that are allowed to access the RADIUS server. There are five options, of which only four are IPSec-related: AXENT, SecurID, CHAP, and PAP. MS-CHAP is available for PPTP tunnel users only (it is not applicable to IPSec tunneling applications).

If you are using token cards for authentication, you must select the appropriate technologies (AXENT and/or SecurID).The RADIUS Authentication screen also allows you to configure the type of authentication methods that are allowed to access the RADIUS server.



Note: Under no circumstances is the UID or password passed in the clear for IPSec client from the remote client or from the switch communicating with the RADIUS server. If you use PAP authentication for a PPTP session, both the username and the password are passed in the clear to the switch over the Internet.

From a security point of view, there is no significant difference between using CHAP or PAP, since the connection between the switch and the RADIUS server is protected by encryption. PAP authentication consumes fewer instructions during the authentication process, which is a minor consideration.

When you are using RADIUS-based authentication, the IPSec client and the switch require a second set of credentials that are used for mutual authentication. These credentials are referred to as the group ID and group password.

The remote access client information is documented in the Contivity VPN Client online help. On the IPSec client side, the remote user must:

1 Select Options→Authentication Options

2 Click on User Group Security Authentication.

3 Enter the group ID and group password that you provide.

4 Select one of these options:

- Challenge Response Token
- Response Only Token
- Group Password Authentication

To complete the RADIUS setup, you must configure at least one group profile for RADIUS users. In this profile, you need to enter the group ID, password, and the allowed group authentication options. You can configure the group profile from the Profiles→Groups→IPSec: Edit screen.

1 Set up and test the operation of the RADIUS server with ACE and/or Defender servers, depending on the type of token security you want. You should do this before attempting authentication by an IPSec client to verify that everything on this side of the network is operating properly.

2 Identify and create the groups for authenticating token users, and supply the group ID and password to all users doing either token card or group password authentication. AXENT and SecurID users are created and maintained in their respective servers, not in the switch. Add the groups in the Profiles→Groups→IPSec screen.

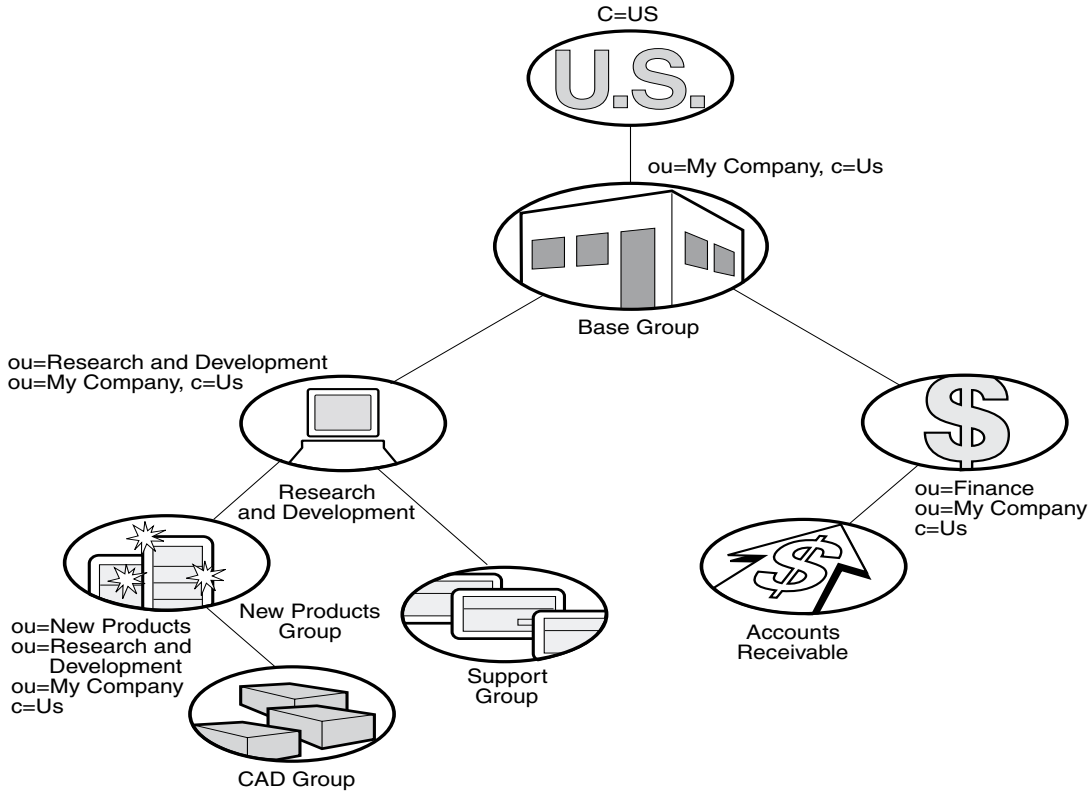
3 Define the RADIUS server configuration settings for token security.

4 Define the Tunnels settings for IPSec. Add a RADIUS server, if necessary.

RADIUS authentication class attribute values

Figure 71 shows the relationship between RADIUS authentication class attribute values for switch users. C is the class attribute for country, and OU is the class attribute for organizational unit.

Figure 71 RADIUS authentication class attribute values



The switch supports RADIUS-supplied attributes, such as IP address and MPPE key and additional specific attributes, if they are returned from a RADIUS server; any other returned attributes are ignored. The specific attributes are detailed in Microsoft* documentation and defined in RFC 2548. This data overrides the corresponding data stored in LDAP, if any. Table 28 shows common examples of class attributes.

Table 28 Class attributes

Name	Value Format	Function
Class	ou=groupname	The user is assigned to the group ‘groupname’, if it exists.
Framed-IP-Address	dotted decimal	If static addresses are allowed, this IP address is used for the tunnel session.

Table 28 Class attributes

Name	Value Format	Function
Framed-IP-NetMask	dotted decimal	Subnet mask to be used with above IP address.
Filter-ID	filter name	If defined, this filter name is applied to the tunnel session.
DNS	domain server names	If used, the domain server name system.
NBNS	protocol name	NetBIOS protocol. If used, provides a translation of the NetBIOS Windows domain name to the IP address; an internet naming service.

Table 29 shows sample details that you enter into your RADIUS server.

Table 29 RADIUS example details

User ID	Class attribute value	Assigned group
Lee Madison	ou=New Products, ou=Research and Development	New Products
Julie Lane	None	Default
Bill Sullivan	ou=Staff	Default (since ou=Staff does not exist)

The RADIUS server uses the class attribute value to associate the user ID with a group in the LDAP database.

RADIUS-Assigned Framed-IP-Address attribute

You must configure a RADIUS-Assigned Framed-IP-Address attribute on the RADIUS server for the ID being authenticated by the switch. If the option Allow Static Addresses (Profiles→Groups→Edit→Connectivity screen) is enabled for the assigned group, then the returned IP address is used for the tunnel session. Otherwise, an IP pool address is assigned.



Note: Only a single IP address is returned by the RADIUS server; therefore, only one active tunnel connection is permitted per user ID.

Configuring the switch for RADIUS

The following procedures describe how to configure the switch to interoperate with a RADIUS server while using either IPSec or PPTP.

IPSec and RADIUS

To configure IPSec and RADIUS on the switch:

1. Go to the Servers→ Radius Auth(entication) screen and select Enable Access to RADIUS Authentication.
2. Enable an authentication method from the Server Supported Authentication Options.
3. Enable a server and enter the server's host name or IP address, the interface type, port number (1645), and secret. Click on OK.
4. Go to the Services→IPSec screen and use the Add RADIUS button to add a RADIUS server to the Authentication Order table.
5. Go to the Profiles→Groups screen and either add or select the group that you want to be the default group for RADIUS users (this is the group a user is assigned to if the RADIUS server does not send back a class attribute).
6. Next, go to the Profiles→Groups→IPSec Configure screen. In the Authentication area, click on the Configure button.
7. On the next screen, the Authentication method for the group is already selected. Enter the group ID and group password. Consider using the LDAP group name as the default group, because you must remember a default name

once you enter it. If your RADIUS server returns a class attribute, ensure that the authentication method is enabled for that group. However, you do not need a group ID and group password for the group that is being returned as a class attribute.

To configure IPSec and RADIUS on the client:

1. In the Contivity VPN Client GUI, go to Options→Authentication Options, select Group Security Authentication, and enter the group ID and group password.
2. Select one of the group authentication options.
3. Click on OK.

PPTP and RADIUS

To configure PPTP and RADIUS on the switch:

1. Go to the Servers→Radius Auth(entication) screen and select Enable Access to RADIUS Authentication.
2. Enable an authentication method.

If a valid class attribute is not returned, then PPTP users are placed in the default group as configured on the Servers→RADIUS Auth(entication) screen.



Note: Everything about the authentication type must match; for example, if you send an encrypted password, then MS-CHAP must be enabled on the RADIUS authentication screen and the RADIUS server must support MS-CHAP.

AXENT software tokens and the IPSec client

The following steps enable the switch to work with AXENT software tokens and AXENT's RADIUS server using the IPSec client. You want the switch to look and appear like a network access server (NAS) to the AXENT RADIUS server. You must have AXENT RADIUS Server version 3.0 or later.

1. Go to RADIUS→Authorization Profiles on the AXENT Defender user interface.
2. Create a profile for switch-authenticated users that has a service type of Authenticate Only and a framed protocol of PPP.
3. Associate this profile with all switch users.
4. Make sure you have the proper class configuration for dial-in security for each user profile. That is, make sure the associated class allows tokens (for example, DHT Required), and that RADIUS is an allowed authorization type.
5. Enable the switch as a valid agent and set the agent key properly on the DMS Remote Manager screen.
6. Perform a DB Sync or RADIUS Sync from the DSS Remote Manager screen after making these changes.

The AXENT RADIUS server can operate with different levels of logging as configured from the command line. This logging is a useful debugging tool. Generally, problems occur when the RADIUS server cannot communicate correctly with the DSS server. This is apparent in the log file.

Using a Netscape directory server

Nortel Networks is currently providing Netscape Directory Services Version 4.0 on the Contivity VPN Switch CD. You cannot currently perform remote LDAP authentication over a Branch Office connection. However, you must have a local LDAP server (either an internal server or an external server on the private network) to authenticate the remote site.

If you are using an external LDAP server, it must be on the private network. To install the Netscape Directory Server Software:

1. Place the Contivity VPN Switch CD in the CD-ROM drive of the server configured with Windows LDAP. Locate the LDAP/Netscape folder on the CD.
2. Unzip the Netscape file, and double-click on the Setup icon to start the Netscape Directory Server installation.
3. Follow the installation instructions, selecting the defaults. The Netscape server software takes approximately 50 MB of disk space after the installation

is complete, plus space for the user database itself. Supply an administrator user name and password when you are prompted to do so. This is the user name and password that is required to run the Netscape Console administration tool on the LDAP server.

At the end of the installation, the installation procedure prompts you to launch the Netscape Console administration tool. This is a Java-based management tool.

Quick start

The following instructions describe setting up a Netscape Directory Server (NSDS, Versions 3.x and 4.0) to function as an external configuration server for your switch. Generally, these steps pertain to all users. Some of the steps pertain only to users who have installed the Netscape Directory Server from a source other than the Nortel Networks CD.

The default configuration directory location of a Netscape Directory Server 3.x is /netscape/suitespot/slapd-<*server_name*>/config. The default configuration directory location of a Netscape Directory Server 4.0 is /netscape/server4/slapd-<*server_name*>/config. Unless otherwise stated, all files to be edited are in these default configuration directories. The main configuration file in this directory is slapd.conf.

The following steps describe how to connect your switch to a newly installed Netscape directory server. These instructions assume that Secure Socket Layer (SSL) is not being used.

For Version 3.1:

1. Stop the Netscape Directory Server. Edit slapd.dynamic_ldbm.conf as the configuration directory in the following instructions.
2. Using FTP, transfer the file /system/slapd/config/bnes.idx from the switch into the default configuration directory. Add the following line to the end of the Netscape file:

   `include "<config_directory>/bnes.idx"`

   where <*config_directory*> is the full path to the configuration directory on the Netscape Directory Server disk.

For all others:

1. Start the Netscape directory server.
2. Go to the Servers→LDAP screen and choose: switch to External Server.
3. Enter the value of the first Netscape 3.x slapd.conf file suffix or the Netscape 4.0 slapd.ldbm.conf file suffix attribute in the Base DN field.
4. Enter the value of the Netscape slapd.conf file rootdn attribute in the Master server Bind DN field. This is the same DN that you entered as the *Unrestricted User* when you created your Netscape Directory Server.
5. Enter the same bind password and confirm password that you entered for the *Unrestricted User* when you created the Netscape directory server.
6. Click on OK and go to the Profiles→Users screen. After a brief delay you see a screen with the group /Base.

You are now connected to the external LDAP server and your skeletal database has been created. You can add users, groups, filters, and hours of access records through the Profiles screens. To connect other switches to the same Netscape LDAP server, repeat steps 2 through 6.

Advanced configuration

This section contains details for modifying the configuration of your Netscape Directory Server. This section details the procedure for modifying the server's schema attribute index list and SSL configuration, and then configuring the switch to connect to it. Netscape Directory Server 4.0 for Windows NT is provided on the Nortel Networks CD.

You must first disable the UID uniqueness plug-in through the Netscape console. This is located in the Netscape Directory Server Console Configuration tab. And you must restart the Netscape directory server after disabling the plug-in.

For an NSDS 4.0 server, no other configuration changes are required as long as the switch binds to the server as an appropriately privileged user. The switch attempts to modify the schema and attribute indices of the server after connecting to it. To do this, the switch must have write access to the schema entry (cn=schema) and the database configuration entry (cn=index, cn=config, cn=ldbm). If you are using the root DN as your bind DN, the switch already has the required access privileges to write to the schema and database configuration entries.

Schema/attribute index cleanup

If you have NSDS 3.x or earlier installed, its switch schema modifications must be updated manually. All attributes, objectclasses and attribute indexes prefixed with newoak must be removed. These attributes and objectclasses are contained in the files config/slapd.at.conf and config/slapd.oc.conf. The indexed attributes are contained in config/slapd.dynamic_ldbm.conf. You must remove all lines that contain an attribute or objectclass that is prefixed by newoak from these files.

Schema/attribute index modifications

You must manually modify the NSDS 4.0 schema only if the switch bind DN (cited earlier) that is used to access the server does not have privileges to modify the schema or database index entries.

To manually modify the schema or database index entries:

1. FTP the following files from the /system/slapd/ldif directory on the switch to any directory on your NSDS file system:
   - bnes.idx
   - bnes.at
   - bnes.oc
2. Add the following line to the end of the config/slapd.dynamic_ldbm.conf (NSDS 3.x) or config/slapd.ldbm.conf file (NSDS 4.0), where <*config_directory*> is the full directory path used in step 1:

   ```
   include "<config_directory>/bnes.idx"
   ```

3. Add the following line to the end of the config/slapd.user_at.conf file, where <*config_directory*> is the full directory path used in step 1:

   ```
   include "<config_directory>/bnes.at"
   ```

4. Add the following line to the end of the config/slapd.user_oc.conf file, where <*config_directory*> is the full directory path used in step 1:

   ```
   include "<config_directory>/bnes.oc"
   ```

5. Restart the NSDS server.

SSL

You must configure the Netscape directory server with compatible encryption types and you must have a certificate signed by a Certificate Authority that is trusted by the switch to use Secure Socket Layer (SSL). You enable SSL and configure encryption types on the Netscape Server Administration management screens. You install certificates and keys through the Netscape General Administration screens.

To configure the switch:

1. Start the Netscape directory server.
2. Go to the Servers→LDAP screen and select: switch to External Server.
3. Enter the Netscape directory server DNS host name or IP address in the Host Name field.
4. Change the port or SSL port number, as necessary, to the port selected when the Netscape directory server was configured. The SSL port must be selected to have the switch use SSL to connect to the Netscape directory server. Also, the Netscape directory server's certificate must be signed by a trusted Certificate Authority and must support at least one compatible encryption type to make a successful connection.
5. Enter the distinguished name, in the Base DN field, of the entry in your Netscape directory server that is serving as the base entry for the switch. The parent entry must always exist. For example, if the base DN is:

   ```
   ou=Remote Access Users, ou=East, o=MyCompany, c=US
   ```

   The organizational unit, Remote Access Users, can exist but does not have to exist, but its parent organizational unit, East, must exist.

   For a new Netscape directory server 3.x, you must specify the following, since no entries yet exist in the database, as the base DN (assuming this is also the value of the suffix attribute):

   ```
   o=MyCompany, c=US
   ```

   An entry similar to the preceding one appears in the 4.x database automatically. This causes the database and base DN of o=MyCompany, c=US to be created on the Netscape directory server. The base DN is used as the base group for the switch.

6 In the Bind DN field for the Master server, enter the DN of a user that has read and write access to the base DN and its subentries. For the slave servers, this DN only needs read access to the base DN and its subentries.

7 In the bind password and confirm password fields, enter the password that corresponds to the base DN.

8 Click on OK. After a short delay you see either operational or error in the server status column. If the status is error, something prevented the switch from successfully converting to the server. Possibilities are:

- The IP address, port bind DN, or password is incorrect.
- The schema on the server either could not be updated (NSDS 4.0) or is incorrect.
- The UID uniqueness plug-in is enabled (NSDS 4.0).
- The base DN is not an organizational unit (ou) or an organization (o).

If the status is operational, you are now connected to the external LDAP server and a skeletal database (contained in the switch's file /system/slapd/ldif/template.ldf) has been created. You can now add users, groups, filters, and hours of access records through the Profiles screens.

Exporting internal LDAP to Netscape directory server

This procedure describes exporting a database from the switch's internal LDAP server to a Netscape directory server. The default LDIF directory location of a Netscape directory server 3.x is:

/netscape/suitespot/slapd-<*server_name*>/ldif

The default LDIF directory location of a Netscape directory server 4.0 is:

/netscape/server4/slapd-<*server_name*>/ldif

The Netscape directory server must first be prepared for additional information. To export a database from the switch's Internal LDAP server to a Netscape directory server:

1 Click on Servers→LDAP.

2 Stop the internal LDAP.

3. Enter a name in the Backup to File field.
4. Click Backup Now. Wait for the backup to complete (this can take several minutes if the database is large). You can monitor the backup status on the Status→Health Check screen. The status of the internal LDAP server is "backup in progress" during the backup and reverts to "Internal Server is down" when the backup is complete.
5. Using FTP, transfer the following file from the switch to the Netscape LDIF directory:

   switch/system/slapd/ldif/<*backup_file_name*>

6. Edit the file to change all occurrences of the string o=Nortel Networks, c=US to the value of your directory suffix (for example, o=MyCompany, c=US). This might be a subentry of your directory suffix as long as the parent of this entry exists. For instance, if your database contains the entry:

   ```
   ou=East, ou=Sales, o=MyCompany, c=US
   ```

   You might want to replace o=Nortel Networks with:

   ```
   ou=Remote Access Users, ou=East, ou=Sales, o=MyCompany, c=US
   ```

   In this case, you also need to change two other lines in the base entry of the exported LDIF file. The base entry is the entry containing the line:

   ```
   o:Nortel Networks
   ```

   Change this line to:

   ```
   ou:Remote Access Users
   ```

   In the same entry, change the line:

   ```
   objectclass: organization
   ```

   to:

   ```
   objectclass: organizationalUnit
   ```

7. Restart the Netscape directory server (if necessary).
8. To add all entries in the LDIF file to the Netscape directory server database, go to the LDIF directory and enter ../../bin/slapd/server/ldapmodify -D<*bind_DN*> -w<*password*> -a -f <*ldif_file_name*> or use the graphical interface provided by Netscape to load the LDIF file.

9 Change the switch's LDAP server from Internal to External and configure it to use the Netscape directory server. The certificate validity check consists of the following:

- Are the certificate start and end dates valid?
- Is the basic constraint extension present. If so, does it indicate an end entity?
- Does the certificate's key usage indicate digital signature?
- Is the certificate signature by the CA valid?

10 Use the certificate to verify the signature on the ISAKMP packet.

RADIUS accounting configuration

The RADIUS Accounting configuration screen allows you to specify how your switch saves RADIUS accounting results. By default, the results are stored locally. You can optionally also save the RADIUS accounting information to a remote RADIUS Server.



Note: If you set the date ahead and then set it back, external RADIUS accounting no longer works.

To configure RADIUS accounting:

1 Go to the Servers→Radius Acct screen.

2 Click to enable or disable internal RADIUS accounting. Internal RADIUS accounting is enabled by default.

3 Enter an interval when a snapshot of the current active tunnel sessions is recorded in a journal file. Use the format, hh:mm:ss, for the interval. The journal file stores the session information until the user logs out of the tunnel session, after which the session stop record is saved on the local disk. In the event of a system crash, upon reinitialization the switch translates the journal file into a series of stop records on a per-session basis. This minimizes accounting data loss. A low interval creates system overhead and requires additional processing. The default interval is 00:10:00 (10 minutes).

4 Click to enable or disable the Interim RADIUS Accounting Record feature. This selection is enabled by default.

5. Enter the interval at which time interim RADIUS records are sent to the specified external RADIUS Server. Use the format hh:mm:ss for the interval. A frequent interval creates system overhead which requires additional processing. The default interval is 00:10:00 (10 minutes).
6. Click on Enable to specify that the switch send its accounting records to the external RADIUS accounting server.
7. Enter the external RADIUS server's host name or IP address. If you enter a host name, use a fully qualified domain name; for example Finance.mycompany.com.
8. Enter the server port number that you want the RADIUS accounting requests to use. The default is Port 1646.
9. Enter the external RADIUS server's required secret (password).
10. Reenter the remote server's secret (password) to verify that you typed the password correctly.
11. Click on the Test Server button to verify the connectivity from your switch to the external RADIUS server. A message at the top of the screen shows the results of the test.

Remote user IP address pool

Remote access users who are using tunneling protocols require two IP addresses to form packets. The addresses are normally referred to as *outer* and *inner* addresses. The outer address, or public address, is visible when packets are traveling through the public data networks (PDNs). This address is negotiated between the client and the ISP to which it is connected. The switch does not have control of this address.

The inner IP address is the one that eventually appears on the private network when the outer layers of the packet are removed. Therefore, this address must lie within the private network address space. The switch provides the remote user with the inner IP address during tunnel setup. This address can come from an internal address pool, from an external DHCP server(s), or from a RADIUS server.

The switch assigns the inner IP address from one of several sources, using the following order:

1. User-specified (excluding IPSec)
2. Static address, either the switch's LDAP database or the RADIUS server
3. Local address pool, either the switch's internal address pool or the DHCP-acquired address pool

To configure the user IP address pool, go to Server→User IP address pool.

DHCP relay

The DHCP relay agent on a switch forwards DHCP and BOOTP messages between a server and a client on different subnets. When a locally attached host issues a DHCP or BOOTP request as a broadcast message, the switch will relay the message to a specified DHCP or BOOTP server. The DHCP relay agent also forwards DHCP replies from server to client.



Note: The DHCP relay agent can run only on all the private physical interfaces and tunnels.

To configure DHCP relay server:

1. Go to the Servers→DHCP Relay screen.
2. Click on the Add button to configure a new DHCP Relay server. The DHCP RelayAdd screen appears.
3. Under the Physical Interface (private) section, select the Physical Interface from the list of existing private interfaces.
4. Under State, enable or disable the DHCP Relay agent by choosing from the list.
5. Under DHCP Servers, enter the Server 1 address. The DHCP relay agent will unicast DHCP packets only to the specified Helper servers (up to 3). Server 1 address is required. Server 2 and Server 3 addresses are optional. Additionally, you can enable and disable each DHCP server by checking or unchecking the Enable checkbox.

You can enable or disable DHCP relay for each interface and specify the DHCP servers for each interface. when DHCP relay is enabled on an interface, the switch forwards DHCP requests from the interface to the DHCP server configured for the same interface.

The DHCP relay agent will unicast DHCP packets only to the specified Helper servers (up to 3). Server 1 address is required. Server 2 and Server 3 addresses are optional. Additionally, you can enable and disable each DHCP server by checking or unchecking the Enable checkbox.

DHCP server

DHCP pushes configuration information to clients and provides for interaction with DNS.The following restrictions apply to the DHCP server:

- DHCP server is enabled by default on the private (trusted) interface
- DHCP Relay and the DHCP Server are mutually exclusive on a physical port.

Because the Contivity 1010, 1050, and 1100 have DHCP server enabled by default, branch office users should set up their PC to accept the IP addresses that the DHCP server will provide the LAN 0 (**private**) ports.

To configure the DHCP server:

1. Go to the Servers→DHCP screen.
2. Click on the Enable/Disable Server button to select the state of the DHCP server.
3. In the Default Options section, specify the lease time in the ddd:hh:mm:ss format or select Infinite to indicate an unspecified period of time.
4. Click Add in the Standard Options section to access the Add Option screen. The standard options section shows the current status of any added options and lets you add new options:
   - Select the desired options from the drop-down list.
   - Select the desired Type from the drop-down list.
   - Enter the appropriate value.

5. In the Pool section, click on the Add button to add a pool. The Add Pool screen appears:
   a. Enter the base IP address for the pool.
   b. Enter the subnet mask for the pool.
   c. Enter a description of the pool.
   d. Click on OK.
6. Select Pool and click on the Configure button to return to the Pool screen.
7. The Inclusion Range section allows you to add blocks of IP addresses that you can then give out.
   a. Under Inclusion Range, click on the Add button. The Pool Inclusion screen appears.
   b. Enter the base IP address for the Start Address.
   c. Enter the End IP address.
   d. Click on OK.
8. Optionally, you can select an Exclusion Range for further control of the IP addresses that you give out. For example, if you have a pool with the range 2.0.1.1 to 2.0.1.255 and want to exclude 2.0.1.50, you would specify 2.0.1.50 as both the start and end address.
   a. Under Exclusion Range, click on the Add button. The Pool Exclusion screen appears.
   b. Enter the Start Address for the range.
   c. Enter the End Address for the range.
   d. Click on OK.
9. Optionally, you can force the DHCP server to assign a fixed IP address to a host every time it logs in. You can do this with host reservations under the Host section.
   a. Click on the Add button. The Host screen appears.
   b. Enter the host name that is registered with DNS.
   c. Enter the IP address that you always want to reserve.
   d. Enter the Ethernet (MAC) address.
   e. Click on OK.

10 The server does not implement configuration changes until it is restarted. Return to the Server→DHCP screen and select the Restart Server option to restart the DHCP server.

11 To verify the configuration changes, go to the Status→Health Check screen or click on the DHCP Stats button on the Status→Statistics screen.

Configuring the switch as a RADIUS server

The RADIUS Service feature allows the switch to function as a simple RADIUS server. When the feature is enabled, the switch accepts RADIUS authentication requests from an external RADIUS client, then searches its LDAP database for the requested user. If the user is listed in the database and has the proper credentials, a positive response is returned to the requesting system. Otherwise the authentication request is rejected. Here are some key points about RADIUS Service that you should consider:

- RADIUS Service accepts standard RADIUS Access-Request packets. It then returns either a standard RADIUS Access-Accept or Access-Reject packet, based on the authentication result. The packet contains only a Reply-Message attribute. No other attributes are returned.
- RADIUS Service supports PAP, CHAP, and MS-CHAP authentication methods.

For users with multiple user accounts, RADIUS Service attempts to authenticate against each account type. If the given username/password matches any of the user's accounts, the authentication succeeds. The authentication is done in this order: PPTP, IPSec, L2F, L2TP.

Chapter 9
Using Certificates

Digital certificates are a means of binding an entity's public encryption or signing key to its identity, and having that identity verified and vouched for by a trusted third party (the certification authority). Digital certificates are used for authenticating both LDAP and VPN connections.

LDAP Server SSL Encryption

Secure Socket Layer (SSL) provides Internet security and privacy and ensures privacy between the switch and the external LDAP server. The SSL protocol negotiates encryption keys and authenticates the server before any data is exchanged. SSL maintains the transmission channels' security and integrity through encryption, authentication, and message authentication codes. The switch implementation of SSL supports the following encryption methods:

- RC4 128-bit MD5 allows clients to request RC4 128-bit MD5 encryption. This is the most secure method. The longer the encryption key, the more secure the encryption. US export law controls the export of 128-bit encryption keys.
- DES 56-bit SHA allows clients to request DES 56-bit SHA encryption. This is the mid-level encryption method, less secure than RC4-128 yet more secure than RC4-40.
- RC4 40-bit MD5 allows clients to request RC4 40-bit MD5 encryption. This is the least secure method of encryption.

You see the link to configuring SSL parameters when you switch from internal to external LDAP servers.

Installing LDAP certificates

The LDAP connection between the switch and the directory server is authenticated asymmetrically. Initially a one-way authenticated SSL connection is established when the directory server passes its certificate to the switch. After SSL authentication is established, the switch authenticates itself to the directory server by presenting its LDAP bind DN and password.

For the SSL connection to be successful, the switch must trust the issuer of the certificate being presented by the directory server during the initial SSL authentication. To import SSL certificates, select the Import→SSL Certificate on System→Certificates. Paste the PKCS#7 formatted CA certificate into the input box and click on OK

VPN security using digital certificates

You can use X.509 certificates to authenticate IPSec and tunnelled IPSec/ L2TP-based connections. The switch supports RSA digital signature authentication in the IPSec IKE key management protocol. Remote users can authenticate themselves to the switch using a public key pair and a certificate as credentials. In addition, the switch uses its own key pair and certificate to authenticate the switch to the user. In addition to requesting and installing certificates for its own use, the switch must explicitly import and mark as trusted the CA certificate that is issuing certificates for the tunnel initiator.

Public Key Infrastructure (PKI)

A PKI is used to issue certificates to both switches and end users. An important decision you must make about the design of a PKI is how you provide CA services. You can use commercially available products from a vendor such as Entrust, where the CA would reside in your facility and be controlled and operated by you. Or, you can subscribe to a CA provider, such as the VeriSign OnSite service, where the CA would actually be operated by VeriSign from a remote location.

CA and X.509 certificates

The CA creates and revokes certificates. By publishing a certificate in its public repository, the CA certifies that it has issued the certificate to the named subscriber. The CA operates on information submitted to it by switch administrators and users of certificate client software, and it stores the results of its work in its LDAP-accessible certificate repository.

Loading certificates on the switch

Two types of certificates must be installed in the switch: server certificates and trusted CA certificates. Server certificates are certificates that the switch requests for itself, and uses to prove its identity to connecting tunnels. Trusted CA certificates are certificates that are issuing end user or branch office tunnel certificates, and are imported by the switch to establish a common trust point.

Server certificates can be requested either manually (using cut and paste #7 and #10) or automatically with the Certificate Management Protocol (CMP) support that is available in 4.0.

Installing server certificates using cut and paste #7 and #10

To install server certificates using PKCS #7 and #10:

1. Go to System→Certificates: Generate Certificate Request screen.
2. Click on PKCS #10 Certificate Request.
3. Fill out the required information.
4. Click OK.
5. Copy and paste or save your CA information to a file.
6. Follow the instructions from your CA provider on how to obtain a certificate.
7. Submit the request to the applicable CA by pasting the encoding into the CA's request screen, following the instructions provided by the CA.
8. Click on Server Certificate to indicate that you are importing a server certificate. Importing a server certificate must correspond to a previous server certificate request.

9 Click on Trusted CA Certificate to indicate that you are importing a trusted CA certificate.

10 Copy the Public Key Cryptography Standard (PKCS) #7 ASCII text that is returned by the Create New Key and Certificate Request into the dialog box. Typically, a Certificate Authority provides this text to you.

11 Click on OK to generate a certificate and store it on the switch.



Note: When you are using Entrust CA, this request must have a subject distinguished name with a common name that is equal to the Entrust Reference Number that is used to preauthorize the certificate issuance.

Installing server certificates using (CMP)

The Certificate Management Protocol (CMP) allows you to create a CMP compliant certificate request. CMP targets management functions for the entire certificate/key life. It defines message formats and includes its own message protection. To initialize the switch for initial certificate enrollment with CMP, you must have:

- Issuer name - CA distinguished name
- Subject name - EE distinguished name (common name, organization, organizational unit)
- Reference number - used to identify the secret value
- Transaction ID or authorization code - initial secret value
- Enrollment URL or destination (host name or IP address) and optional port number
- Imported root CA certificate

To enter this information:

1 Go to the System→Certificates screen and click on the PKCS#10 Certificate Request button.

2 Click on the Certificate Management Protocol (CMP) link to got to the Certification Request - CMP screen. This screen displays the status of any outstanding requests and the fields to fill out for a new request.

3 Enter the reference value.

4. Enter the authentication key.
5. Enter the common name (for example, the Entrust Reference number).
6. Enter the organizational unit.
7. Click on the drop-down list to select one of the following exportable Public Key Sizes in bits (generally, larger keys are more secure):
   - 512
   - 768
   - 1024
   - 2048 (US only)
8. Enter the registration address or URL with an optional port number.
9. Copy the contents of the certificate request into your Web browser's copy buffer.
10. Submit the request to the applicable CA by pasting the encoding into the CA's request screen, following the instructions provided by the CA.
11. Click on Server Certificate to indicate that you are importing a server certificate. Importing a server certificate must correspond to a previous server certificate request.
12. Click on Trusted CA Certificate to indicate that you are importing a trusted CA certificate.
13. Copy the Public Key Cryptography Standard (PKCS) #7 ASCII text that is returned by the Create New Key and Certificate Request into the dialog box. Typically, a Certificate Authority provides this text to you.
14. Click on OK to generate a certificate and store it on the switch.

Installing trusted CA certificates

The trusted CA certificate is the issuer of the certificate that the remote user or branch office tunnel is using to authenticate, and it must be loaded and marked as trusted in the switch.

To import trusted CA certificates in PKCS #10 format:

1. Go to the Contivity VPN Switch URL.

2 Go to the System→Certificates screen.

3 Select Import Tunnel Certificate.

4 Select Trusted CA Certificate (default).

5 Paste the certificate into the paste box.

6 Click on OK. The Installed Tunnel Certificates table displays the certificate entry.

7 Enable Allow All if desired.

8 Click on OK. You have now obtained the CA certificate against which remote users can authenticate. Repeat this operation if multiple CAs will be issuing user certificates.

   Optionally, you can configure a CRL distribution point to enable revocation checking of client certificates. Click on the System→Certificates: Installed Tunnel Certificates: CA Details button, enter the appropriate CRL Information, and click on OK.

The Enabled check box enables CRL checking of certificates for the particular CA. The Search Base, Host, Connection, and Update frequency values must be set for proper access to the CRL LDAP directory store.

Server certificate settings

You can set the following from the System→Certificates→Certificate Configuration screen:

1 Under Certificate Signature Requirements, select Key Usage Extension Required: to require the Key Usage V3 extension to be present in all certificates presented as part of a tunnel initiation (user and branch office).

2 Under Installed Tunnel and Transport Certificates, select Enable Allow All Feature: to enable this feature for all CA certificates. This allows all tunnel requests authenticated by a particular CA to be allowed in, providing a significant configuration savings because individual users are not required to be provisioned into the switch.

3 Select Trusted to indicate that the certificate is trusted. For CA certificates, this indicates that tunnel requests presenting this issuer as the signer of their

certificate are trusted. For server certificates, this is a method of turning off the certificate, without having to delete it.

The System→Certificate Details screen provides the following certificate details:

- This Certificate Belongs To shows the certificate owner's X.500 distinguished name.
- This Certificate Was Issued By shows the issuer of the certificate (the Certificate Authority). In addition to the main attributes, this field also shows the issuer's certificate serial number.
- Validity Dates show the starting and ending dates through which the certificate is valid (for example, 01/29/98 to 01/29/99).
- Certificate Fingerprint shows the unique identifier that is derived from MD5 hashing the certificates. The identifier should be compared with the fingerprint supplied directly from the certificate's issuer (for example, a CA). If the fingerprints do not match exactly, the certificate has been forged or modified.

You must configure a group that is using certificate-based authentication to present a server certificate to remote parties that are initiating tunnel requests. The Default Server Certificate is the Subject DN of the certificate that you want to use as the identity of the switch when initiating or responding to a connection request associated with that group. Tunnel requests are bound to a particular group by the CA certificate that the remote party is presenting as the signer of their certificate. You can set up the local identity for the group on the Profiles→Groups→Edit screen.

Trusted CA certificate settings

Default Group provides a drop-down list of the existing groups for that particular switch. Every CA certificate must be associated with a group to be used to authenticate incoming tunnel requests. The group assignment of incoming tunnel requests is accomplished by either finding the user provisioned in the switches directory (internal or external), or by allowing all users issued by a particular CA to gain access. If all users issued by a particular CA are allowed, there are two ways of determining the group that an initiator gets assigned to (direct assignment into the group assigned to that CA, or access control by subject DN).

Group assignment by user identification

If the subject DN of the certificate presented by the remote initiator of the tunnel is a user that is located on that switch, the group that the user is bound to is the one indicated in that user's configuration.

Allow All

Using Allow All, the switch trusts the CA to establish the true identity of a user. If the user's certificate is within the certificate validity period, the certificate's signature can be verified using the CA certificate, and the user's certificate is not on the CA's CRL, the tunnel connection is permitted. Using the Allow All policy means that once users are certified by the CA, they can create a tunnel connection as long as their certificate is in good standing.

You can allow all users with certificates issued by this CA to authenticate with the switch, regardless of whether they have a user entry in the switch's LDAP database. By default, the CA certificate does not allow all users authentication. Only users with their subject distinguished names (DNs) entered into the Profiles→Users screen are able to authenticate using certificates issued by this CA. If you enable Allow All users to authenticate, you must also select a group for these users from the Default Group drop-down list box. If you want only specific instances of users to authenticate with the CA authority, you must configure each of these users from the Profiles→Users→Edit screen, and disable Allow All authentication for this CA. Only these users can then perform IPSec RSA Digital Signature Authentication using a certificate issued by this particular CA.

The Allow All feature must be enabled for each CA certificate against which you want to permit authentication without an explicit user entry. This allows anyone with a valid certificate from the particular CA to establish a tunnel connection.

Also, you must associate a default group with that certificate. The client authenticating with the Allow All feature then uses the attributes associated with that group.



Note: Branch Office connections do not support the CA Certificate Allow All feature. Therefore, you must configure an explicit Branch Office connection.

Access control by Subject DN

This form of mapping incoming requests to groups allows the subject DN of incoming certificates to be parsed to a configured depth and associated with a corresponding group. During the client authentication process, we will try to find a match of the client's certificate subject DN from all the associations of the CA. The match could be a partial match or an exact match. In case of a partial match, we will use the longest match from the root of DN. After a match is found, the client is assigned to the corresponding group. If no match is found, the client is assigned to the default group of the CA.

A DN has multiple components (RDN). The most common ones are Common Name (CN), Country Name (C), Locality Name (L), State/Province Name (S), Organization (O) and Organizational Unit (OU). The order of the RDN does not matter unless multiple OUs are present, but ordering the DN in the following sequence avoids ambiguity: C, S, L, O, OU, and CN.

The following example shows group mappings:

```
ou=Contivity, o=Nortel, c=US/base/contivity
ou=Engineering, ou=Contivity, o=Nortel, c=US/base/contivity/
Engineering
ou=Marketing, ou=Contivity, o=Nortel, c=US/base/contivity/
Marketing
ou=Engineering, o=Bay Networks, L=Boston, S=MA, c=us/base/bay
```

Certificate revocation list (CRL)

A CA can revoke user and server certificates whenever the associated key pair is no longer valid, the key pair has been compromised, the user has left the organization, or a server has been retired, among other reasons. When a certificate is revoked, the CA updates an associated revocation list with the revoked certificate's serial number. This list is referred to as a certificate revocation list (CRL). A CA can have one or more associated CRLs.

CRLs are published by the CA in an associated LDAP-accessible directory service. The publication frequency is set by the CA administrator. In an Entrust environment, a new CRL can be automatically published at a set time, at any time manually set by an administrator, or whenever a certificate is revoked. In a VeriSign OnSite environment, new CRLs are published at a fixed interval, typically 24 hours.

The switch can optionally use CRLs to verify the revocation status of user certificates. If enabled on the switch, CRLs are periodically retrieved from the CA's LDAP directory store and cached into the switch's associated LDAP database. This allows for rapid verification of user certificates during IPSec tunnel establishment. You can configure the frequency with which the switch checks for a new CRL.

Because a CRL is signed using the CA's private key, it is protected against tampering. The switch verifies the CRL signature each time it is used. A CRL server must be configured for each trusted CA certificate that is imported into the switch.



Note: The LDAP server that contains CRLs for the CA certificates on the switch must be reachable from the public or private interface.

CRL Settings

The following list provides explanations for CRL settings that you can consider setting:

- CRL Checking Enabled shows CRL usage enabled on the switch on a per CA basis. To enable the use of CRLs for a CA, select the Details button on the main System→Certificates screen. The section labeled Certificate Revocation List Information is used to configure the necessary information. The Enabled check box turns on CRL checking of certificates for the particular CA. The Search Base, Host, Connection, and Update frequency values must be set for proper access to the CRL LDAP directory store.
- CRL Checking Mandatory determines if a CRL must be present when an IPSec tunnel is established to a particular CA. If checked, when the connection is made, a CRL for that CA must be present. If no CRL is present, the tunnel is not established. This setting has no effect unless CRL Enabled is checked.

- CRL Update Frequency allows you to enter a value in minutes that represents the frequency with which the switch should query the CA's LDAP server for a newly published CRL. The default value of 0 indicates that this switch does not update any CRLs. This is useful when many switches share an LDAP database, but you want only one switch to actually perform the update operation. To minimize the load on an external LDAP server, it is important to make sure only 1 or 2 switches are updating a shared CLR entry in a multiple-switch, shared external LDAP environment.
- CRL System Status is read-only and is automatically updated by the switch to reflect the CRL updating activity.

To configure CRL servers:

1. From the System→Certificates screen, click on the Manage CRL Servers button to access the Manage CRL Servers screen. Use this screen to configure and manage CRL servers. Shown at the top of this screen is a list of currently configured CRL servers for the CA that you can edit or delete. The New CRL Server section allows you to configure and add a new CRL server.
2. In the Search Base field, enter the portion of the X.500 directory where the CA stores certificate revocation lists. Following is a sample search base entry:

   ```
   ou=Engineering, o=Nortel Networks, c=US
   ```

3. In the host field, enter the host name or IP address of the LDAP-accessible directory server that is storing the published CRLs. This host must be reachable via one of the switch's private interfaces, and if a host name is used in place of an IP address, then one or more DNS servers must be configured on the switch's System→Identity screen.
4. In the Connection field, enter the port number that is associated with the LDAP server. Optionally, enable the use of the Secure Socket Layer (SSL) to secure the connection with the LDAP server. SSL is not required in general for handling CRLs since a CRL is signed and is therefore protected against modification and spoofing.
5. Select the Enabled or Disabled state from the list box to enable or disable the CRL server.

Enabling certificate use for tunnels

For IPSec, you must enable RSA digital signature support for groups. Specifically, this includes any default groups that are associated with CAs, and the groups containing any specific instances of users who are doing certificate-based authentication.

1. From the Profiles→Groups→Edit→IPSec→Configure: RSA Digital Signature field, click on the RSA Digital Signature check box to enable.
2. Select the appropriate Default Server Certificate from the drop-down list box. This is the certificate that is sent to clients to authenticate the switch's identity. This server certificate should be issued from the same Entrust PKI that has issued the remote access clients' certificates.
3. Click on OK.

For L2TP/IPSec authentication:

1. Click on the drop-down list and select the authentication method that you want to use for the branch office connection.



Note: When you change the authentication type, the screen immediately changes to reflect the requirements of the new authentication method. Any changes that you might have made on the Authentication part of the previous screen are lost.

2. Enter the local UID, which is the user ID of the local switch that you are configuring.
3. Enter the peer UID, which is the user ID of the remote switch that you are configuring.
4. Enter the password for the UID, then confirm the password to verify that you entered it correctly. If you selected a variation of MS-CHAP V2 authentication, no password is required for the Local UID.

Identifying individual users with certificates

As an alternative to allowing all users issued by a particular CA to gain access to the switch, users can be identified explicitly by certificate attributes.

To create IPSec certificate credentials:

1. Select a valid issuer Certificate Authority from the drop-down list. These Certificate Authorities are configured from the System→Certificates: Generate Certificate Request screen.
2. You can use either the relative distinguished name or the full distinguished name. The relative distinguished name is a collection of the following components that uniquely identify the remote peer in an IPSec certificate environment.
3. Enter the organization with which the user is associated.
4. Enter the organizational unit with which the user is associated.
5. Enter the common name with which the user is associated.
6. Enter the country in which the user resides.
7. Enter the state or province in which the user resides.
8. Enter the locality in which the user resides.
9. You can directly enter the full distinguished name (FDN) in this field, rather than entering the individual components in the previously described relative distinguished name fields. A sample entry follows:

   ```
   CN=MySwitch, O=MyCompany, C=US
   ```

10. You can optionally use a subject alternative name in place of a subject DN, and specify the format of the name. The following formats are acceptable:
    - E-mail name (for example, net_admin@company.com)
    - DNS name (for example, gateway.cleveland.company.com)
    - IP address (for example, 192.168.34.21)

Identifying branch offices with certificates

The Authentication section of the Profiles→Branch Office→Edit Connection screen allows you to configure the authentication that is used between the local and remote branch office switches. The fields that appear in this screen depend on whether you are using an IPSec, PPTP, or L2TP tunnel type.

Click on the drop-down list and select the authentication method that you want to use for the branch office connection.

Note: When you change the authentication type, the screen immediately changes to reflect the requirements of the new authentication method. Any changes that you might have made on the Authentication part of the previous screen are lost.

IPSec authentication

In the Authentication section of the screen, fill out the following information:

1. Enter the pre-shared key in either text or hex string. This is an alphanumeric text or hexadecimal string that is used between the local and remote branches for authentication. In order for authentication to occur, you must use the same pre-shared string on both the local and remote branch offices.
2. Certificates are associated with each endpoint gateway and allow for mutual authentication between two connections. The certificate portion of the screen includes information about the remote branch office system, the authority that issued the certificate, and the certificate identification.
3. Remote Identity is the name of the remote peer initiating the tunnel connection. You can use either a subject distinguished name (subject DN) or a subject alternative name to uniquely identify the remote branch office system. Specifying both a full subject DN and a subject alternative name on this screen allows the remote peer to use either identity form when making a connection.
4. Select a valid issuer Certificate Authority from the drop-down list box. This CA is the issuer of the remote peer's certificate or a higher level CA in the remote peer's certificate hierarchy. The CA must have the trusted flag set via the Certificates screen. If a CA hierarchy is being used, all intermediary CAs

below the trusted CA must have been imported to the switch. These Certificate Authorities are configured from the System→Certificates: Generate Certificate Request screen.

5 If you are using a distinguished name to identify the remote branch office site, you can choose to enter the DN as either a relative distinguished name or a full distinguished name. The DN entered here must exactly match the DN in the remote peer's certificate.

6 The relative distinguished name has the following supported components:



Note: Do not include the attribute type as part of your entries in the Relative section. For example, for a name of CN=MySwitch, your entry would be MySwitch (without the CN attribute type).

- Common Name -- Enter the common name with which the server is associated.
- Org Unit -- Enter the organizational unit with which the server is associated.
- Organization -- Enter the organization with which the server is associated.
- Locality -- Enter the locality in which the server resides.
- State/Province -- Enter the state or province in which the server resides.
- Country -- Enter the country in which the user resides.

PPTP authentication and L2TP authentication

Under the authentication section of the screen, fill out the following information:

1 Under Local UID, enter the user ID of the local switch that you are configuring.

2 Under Peer UID, enter the user ID of the remote switch that you are configuring.

3 Enter the password for the UID, then confirm the password to verify that you entered it correctly. If you selected a variation of MS-CHAP V2 authentication, no password is required for the Local UID.

4 Click on Enable or Disable to enable or disable compression.

5. Click to enable or disable the Compression/Encryption Stateless Mode selection.This selection is not used if encryption and compression are both disabled.
6. The L2TP Access Concentrator (for L2TP authentication only) field appears if you have selected L2TP as the preferred tunnel type for the branch office connection. Use this entry to specify the L2TP Access Concentrator that you want to perform authentication between the switch and the NAS.

Chapter 10
Switch Administration

This chapter describes administration tasks that help you operate the switch. These tasks provide details on scheduling backups, upgrading the software image, saving configuration files, performing file maintenance, creating recovery diskettes, and system shutdown.

Administrators

The switch supports multiple administrators. You can assign different rights to allow or prevent administrative users from managing or viewing switch and user configuration information. You assign administrative privileges and rights on the Profiles→User→Edit screen. The switch also supports a primary administrator.

The Manage Switch and Manage Users access right settings can be assigned one of the following privilege levels:

- None - This user does not have administrator rights to manage the switch or manage users; the user cannot view or manage Switch Configuration or User settings.
- View - This user has administrator rights to view (monitor) switch configuration or user rights settings; however, the user cannot manage (change) them. This is the lowest level of administrator rights.
- Manage - This user has administrator rights to view (monitor) and manage (configure) other switch configuration or user rights settings. This is the highest level of administrative rights.
- Add Subgroups is a check box that lets you give the user the authority to add and delete subgroups under the given directory when the user only has View authority with Manage Switch access rights.

The Administrator Settings screen allows you to change the primary administrator user ID and password. It also controls the Administrator Idle Timeout Setting for all administrators, the default language, and serial port settings.

There can be only one primary administrator. The primary administrator user ID and password combination provides the user with this information access to all screens and control settings. The primary administrator user ID and password are also used to access the serial port and the recovery disk.



Note: The primary administrator user ID and password are only saved during a system shutdown. Therefore, once you set these parameters, you must implement an Admin→Shutdown to save the new settings. Doing a reset (using the Reset button on the back of the switch) does not store the parameters.



Note: Do not lose or forget the password once the switch has been configured. Losing or forgetting your password would require you to return the switch to Nortel Networks for reconfiguration to default settings. All settings and backups would be lost. There is no way to access the system without the primary administrator password.

You can change the primary administrator user ID and password on the Admin→Administrator screen.

Backups

The switch checks at regular intervals to see whether system file changes have been made. When there are changes, they are written to each of the backup servers. All of the system files are backed up the first time; thereafter, only the files that have changed are backed up.

Backups are performed using the File Transfer Protocol (FTP). Files are transferred from the switch's hard disk to the designated path under the default FTP directory on the backup file server.

You must provide the backup host server's IP address, path, backup interval, user ID, password, and password confirmation on the Automatic Backup screen. This information is stored in flash memory and appears on the Automatic Backup screen when you access it later. Additionally, the primary backup data is stored during the recovery procedure (refer to the section "Recovery"), in the unlikely event that it is needed.



Note: The User ID and Password fields are required, and may not be left blank or be defined as spaces. Note that the user interface screen does not enforce this and report an error. However, when the actual backup process begins, the error in the configuration is detected and the backup does not occur. After you configure the backup, you should perform a manual backup (by clicking on the Backup button) to ensure that all parameters are correctly configured.

Some FTP servers do not use standard FTP port numbers, and they cannot be used as download FTP servers for Nortel Networks software. To enable automatic backups, go to the Admin→Auto Backup screen.



Note: Any changes to backup parameters made while a backup is in process do not take effect until the currently running backup completes.

If you want to save a certain configuration for a later date, you must understand that there are two components that define a given configuration: the configuration file and the LDAP database.

Saving a configuration by using the Save Current Configurations option on the Admin→Configs screen saves only the operational parameters in the configuration file, such as interface IP addresses and subnet masks, backup host IP addresses, and DNS names.

To completely save the Contivity VPN Switch configuration on the internal LAN server, you must also save the LDAP database, which contains the group and user profiles, filters, backup file names, and so forth. Go to the Servers→LDAP screen and click on Stop Server. Next, enter a file name in the Backup/Restore LDAP Database field. You should conform to the 8.3 MS-DOS naming convention and append the file name with .ldf; for example, LDAPOne.ldf.

You can use the same backup server for multiple switches. Each switch creates a unique directory based on its serial number. Assuming that the home directory for the two FTP servers is C:/software and assuming that the switch has serial number of 01001, then the backups would be performed to the two hosts, as follows:

```
192.164.34.2 in C:/software/backup/v101/SN01001
172.19.2.30 in C:/software/backup/SN01001
```

The switch allows you to configure regular intervals when your system files are saved to designated host backup file servers. You can designate up to three backup file servers.

Configure Automatic Backups immediately so that you do not lose system or configuration information in case of problems.You configure the Automatic Backup servers from the Admin→Automatic Backup screen.

The switch does not begin a backup for at least 5 minutes after rebooting. This time period is to allow all resources to start operating. This delay occurs even if you go into the Admin→Auto Backup screen and request that a backup be started immediately.



Note: After entering the Automatic Backup File Servers information, click on the screen and press [Alt] [Print Screen] to save the screen image to a buffer. Next, paste the image into a text file and keep it as a record of the backup file servers that you are using.

Restoring Configurations

If you want to save a certain configuration for a later date, you must realize that there are two components that define a given configuration: the configuration file and the LDAP database.

Saving a configuration from the Admin→Configs screen Save Current Configurations option saves only the operational parameters in the configuration file, such as interface IP addresses and subnet masks, backup host IP addresses, and DNS names.

To completely save the Contivity VPN Switch configuration on the internal LAN server, you must also save the LDAP database, which contains the group and user profiles, filters, backup file names, and more. Go to the Servers→LDAP screen and click Stop Server. Next, enter a file name in the Backup/Restore LDAP Database field. Note that you should conform to the eight-dot-three MS-DOS® naming convention and append the file name with .ldf; for example, LDAPOne.ldf.

Automatic Backup File Servers

To enable automatic backup:

1. Click on the Enable button to enable the associated Host Backup File Server.
2. Enter the Backup File Server Host name or IP address.
3. Enter the Backup File Server Path, for example, Building3/Switch_backups.
4. Select the Specific Time option to execute the backup at a specific time. Enter the time at which you want the backup to occur.
5. Select the Interval option to execute the backup at certain intervals of time. Specify in hours the time period after which the system automatically backs up changed files to the backup file server. The minimum interval is 1 hour, and the maximum is 8064 (336 days); the default is 5 hours.
6. Enter the User ID that is required for FTP login to the backup file server.
7. Enter the Password that is required for FTP login to the backup file server.
8. Reenter the Password that is required for FTP login to the backup file server.
9. Click on the Backup button to execute a backup to each enabled server now. This action also synchronizes the hard disk drives when there is more than one in a device. Otherwise, the hard disks synchronize automatically every 60 minutes.

Tools

The switch supports standard IP tools such as `ping`, `Traceroute`, and ARP `show` and `delete`. You can access these tools through the Admin→Tools screen.

The `ping` command generates an ICMP echo-request message, which is sent by any host to test node reachability across a network. The ICMP echo-reply message indicates that the node can be successfully reached.

The Traceroute tool is used for measuring a network round-trip delay. Messages are sent per hop and the wait occurs between each message. If the address is unreachable, it uses this formula to determine how long it takes for the Traceroute to time out.

```
maximum hops (30) x the wait timeout (5) x 3 seconds
```

The Address Resolution Protocol (ARP) dynamically discovers the low-level physical network hardware address that corresponds to the high-level IP address for a host. ARP is limited to physical network systems that support broadcast packets that can be heard by all hosts on the network.

Recovery

The Recovery screen allows you to configure a recovery diskette for restoring the software image and file system to the hard drive of the switch in the unlikely event that there is a hard disk crash. The recovery diskette is included with your switch. You can also use this screen to create additional copies of the recovery diskette, as well as to reformat a diskette.



Note: The Contivity 1000 does not have a floppy drive in the unit. The Contivity 600 and Contivity 1010/1050/1100 series also do not have floppy drives, but the recovery image is stored in a PROM; you can invoke it by pressing a switch on the back of the unit.

Remove the switch's front cover to gain access to the diskette drive. Refer to *Getting Started* for details. Booting the switch using the recovery diskette, utilities accomplish restoration of the switch’s hard disk, including:

- Hard disk reformatting
- FTP access to the hard disk
- Restoration of the previously backed-up software image and file system from a backup host to the hard disk

- Downloading a new factory default software image and file system from a file server to the hard disk

These utilities are accessed via HTTP management of the switch after it has been booted off of the recovery diskette.

Upgrades

You can download the latest Nortel Networks software for the Contivity VPN Switch using File Transfer Protocol (FTP). In addition to retrieving software, you can select which version of software to run.

The Host, Path, User ID, and Password fields are already filled in. You only have to select the version that you want to download. Click on the View link to go to the Nortel Networks Web site where you can select the latest software version.

Additionally, you can enter the necessary information to download software from your own FTP site. This situation assumes that you have previously downloaded the software from Nortel Networks to your organization's FTP site (or loaded it from the Nortel Networks CD). In this case, simply write over the listed Nortel Networks details.

Nortel Networks recommends that you keep a maximum of four versions on the Contivity VPN Switch system disk. If four versions exist on the Admin→Upgrade screen, you must delete one using the Admin→File System before you FTP another release.

You cannot upgrade switch software through a branch office tunnel that is translating the management address with dynamic NAT.



Note: Some FTP servers do not use standard FTP port numbers, and they cannot be used as download FTP servers for Nortel Networks software. For more information, contact Nortel Networks Customer support.

To initiate an upgrade, go to the Admin→Upgrade screen. After the new version has loaded, select Click here to complete the upgrade. If the link reappears off the main welcome page, it means there is a browser caching problem and you should reload the page.

System configuration

You can save the current or delete existing system configuration files through the Admin→Config screen. Additionally, you can select one of the previously named configurations and restore it as the current configuration.

File management

You can navigate through the switch file system through the Admin→File System→File System Maintenance screen. It lists the devices (drives) and directories. This provides flexibility in viewing details of a file or directory, and it allows you to delete unnecessary files. For example, if you had problems performing an FTP transfer with a specific file, you could view the file details to learn its file size and when it was last modified for troubleshooting purposes. Additionally, you can toggle between hard drives when a backup drive is available.

SNMP

The SNMP screen (Admin→SNMP) allows you to configure the switch to generate SNMP Version 1 Traps, based on MIB II. You can do the following:

- Designate the remote SNMP management stations that are authorized to send SNMP Gets to the switch.
- Designate the trap hosts to which the traps can be sent.
- Configure the traps.

The SNMP counters measure packet attributes that are based on the outer IP header. In the tunneled environment there is also an inner IP header, but this IP header does not contribute to the SNMP MIB counters. For example, the outer packet header might be a good packet header and counted, but the inner packet header might be corrupted and would not contribute to the drop counter.

You can view the Health Check screen for the results of SNMP traps.



Note: A Nortel Networks proprietary MIB is included on the Nortel Networks CD. Click on the CesTraps.mib file to load the MIB. See Appendix A, "MIB Support," for a description of the CesTraps.mib.

SNMP Trap Settings

The Admin→SNMP→Trap Settings screen allows you to specify the level of severity that is reported for the trap. You can also specify that the trap be sent only once.



Note: The results of many of the selections you make on the Trap Settings screen are reported on the Health Check screen. Refer to *Reference for the Contivity VPN Switch* for additional information.

Shutdown

The Shutdown options allow you to shut down immediately, to wait until current users are logged off, or to wait until a designated time. A normal shutdown safely terminates connections so that no data is lost, compared with a spontaneous loss of power.

Additionally, you can select whether to power off or restart after shutdown, and also choose the configuration file to use upon restarting. To allow you to conduct an orderly shutdown, you can disable new logins, and you can disable logins after the shutdown to perform system maintenance.

Always use the System Shutdown screen to shut down the system rather than the Power or Reset buttons on the back of the switch. This ensures the integrity of your file system.



Note: After performing a system shutdown, click on the Reload/Refresh button to see the latest switch information.

Go to the Admin→Shutdown screen to turn off the switch.

You can boot either in normal mode or safe mode. Each mode has its own system image and configuration information. A system booted in safe mode can only accept secured management tunnel connections. Once this tunnel is established, Telnet, HTTP, and FTP traffic is allowed, but no other VPN traffic is allowed through the secured management tunnel or the switch. In normal mode, the system operates with the normal software and configuration and transports both VPN traffic and management traffic.

The system can be booted in safe mode when it is requested at the system shutdown or when the system is not shut down normally. Once in safe mode, the system stays in that mode for a configurable amount of time (default is 5 minutes). If there is no administrative user logged into the switch over a secured management connection or the serial console within this time period, the system reverts back to normal mode.

You can configure safe mode settings on the System→Settings screen.

Switch status

The System Status screens allow you to see from the Web interface which users are logged on, their traffic demands, and a summary of your switch's hardware configuration, including available memory and disk space.

The status screens include:

- Sessions
- Reports
- System Status

- Health check
- Statistics
- Firewall

The switch has several logs that provide different levels of information, including:

- Event log
- System log

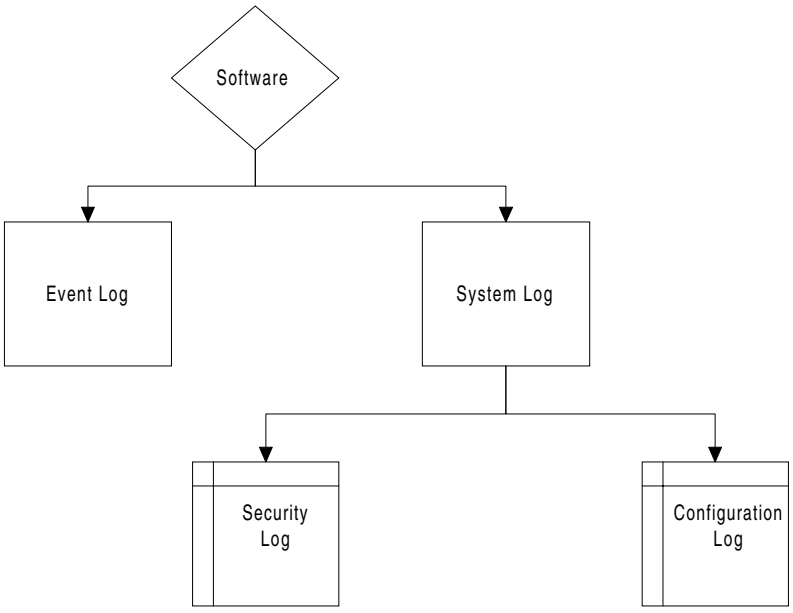
The logs are stored in text files on disk and they indicate what happened, when, and to which user (IP address and user ID).

The Event log captures real-time logging over a relatively short period of time (for example, the Event log could wrap its 2000 possible entries in minutes). The System log captures data over a longer period of time, up to 61 days.

Most events are sent to the Event log first. Significant events from the Event log are sent to the System log. Not all data that is saved by the System log comes from the Event log, but that is generally the case. The switch filters from the System log security entries for the Security log and configuration entries for the Configuration log.

Figure 72 shows a Nortel Networks logging scheme.

Figure 72 Nortel Networks logging scheme



The different log options allow you to write specific event levels to the log files and view them, including:

- Normal
- Urgent
- Detailed
- All

Sessions

You can monitor which users are tunneled into the switch, when they logged in, and the number of bytes and packets they have transmitted or received. Additionally, you can see selected session details, and you can even log off users.

Once a session becomes connected, detailed information about the connection is available from the Status→Sessions screen. This screen lists all connected sessions, including administrative sessions. Click on the appropriate buttons beside each session to either log out the session or view detailed information about it. As well as statistics, this information contains what encryption was negotiated, the SOIs of the security associations and more.

Reports

The Reports feature allows you to generate comprehensive reports of users and related information. You generate reports in an on screen tabular format, and you can import them into a spreadsheet or database through the comma-delimited format.

At midnight (12:00 a.m.), the data collection task performs summary calculations and rewrites history files, along with other management and cleanup functions. You should leave the switch running overnight to perform this task. The switch must be running at midnight to generate a historical graph for the day.

The Status Reports screen allows you to view system and performance data in text or graphical format. You can generate current or historical graphs of valuable system data. The Reports feature provides a comprehensive display or down-loadable reports on user activity.

If you have multiple switches throughout the world, you might want to use the Greenwich Mean Time (GMT) standard. This synchronizes the various log files so that the timestamps are directly comparable.

System

The System Status screen shows the switch's up time, software and hardware configurations, and the current status of key devices. When there is a pending shutdown or an IPX Public Network Address change that requires a reboot, such events are listed at the top of this screen.

Health check

The Health Check screen provides an overall summary of the current state of the switch's hardware and software components at a glance. It lists all aspects of unit operation, with the most critical information to check at the top of the page. By clicking on the link on the right side of the screen, you can go directly to the screen for configuration of that feature.

Statistics

The Statistics screen provides many subscreens with a wealth of general and diagnostic information about the switch hardware, software, and connections. Much of the information is specifically designed for Nortel Networks Customer Support personnel to assist them in diagnosing problems. Some screens, however, such as the LAN Counters, Interfaces, and the WAN Status, might provide you with some traffic information. The Status→Statistics screen allows you to see text displays of system-level statistics that can help resolve lower-level problems with connections. These displays are similar to command-line output from the operating system.

In normal operation, and routine troubleshooting, many of these screens will never need to be examined. Some of the information, such as routing information, is also available through other screens, such as System→Routing.

Event log

The Event log is a detailed recording of all events that take place on the system. These entries are not necessarily written to disk, as with the System log. The Event log retains all system activity in-memory but only the significant entries are saved in the System log (on disk). The Event log includes information on tunneling, security, backups, debugging, hardware, security, daemon processes, software drivers, and interface card driver events.

The Event log maintains the most recent 2000 events in memory and is overwritten when it becomes full. Some of the events are written to the disk log files (the Config log, Security log, and System log). These are maintained on the hard drive for up to 60 days. They are available for review for any given day.

The following is the general format of the log entries:

- Date
- Time
- Task
- Log
- Message

The first digit of the [*xx*] entry is 1 if the information is written to the system log and 0 if it is a memory-only entry.

The second digit is:

- 1 - Debug
- 2 - Low Priority
- 3 - Medium Priority
- 4 - High Priority

The following sample Event log is followed by the descriptions of the line entries.

1 07/03/2000 22:46:11 0 Security [11] Session: IPSEC[jdoe] attempting login

2 07/03/2000 22:46:11 0 Security [01] Session: IPSEC[jdoe] has no active sessions

3 07/03/2000 22:46:11 0 Security [01] Session: IPSEC[jdoe] Jay Doe has no active accounts

4 07/03/2000 22:46:11 0 ISAKMP [02] Oakley Aggressive Mode proposal accepted from jdoe (134.177.58.151)

5 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 SHARED-SECRET authenticate attempt...

6 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 attempting authentication using LOCAL

7 07/03/2000 22:46:12 0 Security [11] Session: IPSEC[jdoe]:6454 authenticated using LOCAL

8 07/03/2000 22:46:12 0 Security [11] Session: IPSEC[jdoe]:6454 bound to group /Base/Pittsburgh/Jay Doe

9 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 using group filter permit all

10 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 IN FILTER 1 permit IP any any

11 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 OUT FILTER 1 permit IP any any

12 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 deny TCP any GT 1023 any EQ 80

13 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 deny TCP any GT 1023 any EQ 21

14 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 deny TCP any GT 1023 any EQ 20

15 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 deny UDP any any EQ 161

16 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 deny TCP any GT 1023 any EQ 23

17 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 deny TCP any GT 1023 any EQ 256

18 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 deny TCP any EQ 256 any GT 1023

19 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 deny TCP any EQ 257 any GT 1023

20 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 RESTRICTED FILTER 1 permit IP any any

21 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 OUT FILTER 1 permit IP 47.184.90.26 0.0.0.0 any

22 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 LOCAL IN FILTER 1 permit TCP any GT 1023 any EQ 80

23 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 LOCAL IN FILTER 1 permit UDP any any EQ 161

24 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 LOCAL IN FILTER 1 permit ICMP any any 8

25 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 LOCAL IN FILTER 1 permit ICMP any any 0

26 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 LOCAL IN FILTER 1 permit ICMP any any 3

27 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 LOCAL IN FILTER 1 permit UDP any EQ 68 any EQ 67

28 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 LOCAL IN FILTER 1 permit TCP any GT 1023 any EQ 17

29 07/03/2000 22:46:12 0 Security [01] Session: IPSEC[jdoe]:6454 LOCAL IN FILTER 1 permit TCP any GT 1023 any EQ 586

30 07/03/2000 22:46:12 0 Security [11] Session: IPSEC[jdoe]:6454 authorized

31 07/03/2000 22:46:12 0 Security [12] Session: IPSEC[jdoe]:6454 physical addresses: remote 134.177.58.151 local 47.184.90.27

32 07/03/2000 22:46:13 0 Security [12] Session: IPSEC[jdoe]:6454 assigned IP address 47.184.91.56, mask 255.255.254.0

33 07/03/2000 22:46:13 0 ISAKMP [02] ISAKMP SA established with jdoe (134.177.58.151)

34 07/03/2000 22:46:13 0 Security [12] Session: IPSEC[jdoe]:6454 physical addresses: remote 134.177.58.151 local 47.184.90.27

35 07/03/2000 22:46:14 0 Security [12] Session: IPSEC[jdoe]:6454 physical addresses: remote 134.177.58.151 local 47.184.90.27

36 07/03/2000 22:46:14 0 Outbound ESP from 47.184.90.27 to 134.177.58.151 SPI 0x001d7dbd [03] ESP encap session SPI 0xbd7d1d00 bound to cpu 1

37 07/03/2000 22:46:14 0 Inbound ESP from 134.177.58.151 to 47.184.90.27 SPI 0x0000c959 [03] ESP decap session SPI 0x59c90000 bound to cpu 1

38 07/03/2000 22:46:14 0 IPvfy.03d9aad4{Tun} [01] SetExpectedSrcAddress: 0x385bb82f, Bcast 0xff5bb82f

39 07/03/2000 22:46:14 0 IP Redirector [10] TunnelAddrReg(47.184.91.56 0x3c88fc4 65535)

40 07/03/2000 22:46:14 0 ISAKMP [03] Established IPsec SAs with jdoe (134.177.58.151):

41 07/03/2000 22:46:14 0 ISAKMP [03] ESP 3DES-CBC-HMAC-MD5 outbound SPI 0x1d7dbd

42 07/03/2000 22:46:14 0 ISAKMP [03] IPcomp LZS outbound CPI 0x3471

43 07/03/2000 22:46:14 0 ISAKMP [03] ESP 3DES-CBC-HMAC-MD5 inbound SPI 0xc959

44 07/03/2000 22:46:14 0 ISAKMP [03] IPcomp LZS inbound CPI 0x907b

45 07/03/2000 22:46:15 0 Address Pool [11] Dhcp: address bound to 47.184.91.66-47.184.91.66 -- renewal in 21600 seconds.

Lines 1-4 are a result of receiving the first IKE (UDP port 500) packet and are normal for a client login. The characters in brackets, e.g., [jdoe], indicate the username that is attempting login for these entries. The fourth line, (the ISAKMP entry), indicates the public IP address from which the session is being initiated.

Lines 5-8 show that the user is attempting to log in with a shared secret (username and password) from the internal LDAP. The third line indicates that the authentication attempt was successful. The fourth line shows the group profile and user entry that was associated with the user. Authentication is now complete. The next section covers authorization (the privileges the newly authorized user has).

Line 9 shows the filters that were applied to the user. This line tells what filter name, permit all in this case, was assigned to the user in the group profile.

Lines 10-11 show no restrictions on IP traffic between the client and the internal networks.

Lines 12-20 show restricted filter definitions.

Line 21 allows any traffic from the management interface outbound.

Lines 22-29 show the LOCAL filter indicates what traffic is allowed from the client to the switch itself. In order, the filters allow for HTTP management from the client; SNMP queries from the client; ICMP echo replies, ICMP echoes, ICMP destination unreachable, BOOTP client requests to a BOOTP server port, Quote-of-the-day messages, and the Password server. The BOOTP messages are used by the client to ask the Contivity for its DNS and WINS information. The Quote-of-the-day server is used to provide the banner messages, if any, for the IPsec client.

Line 30 shows authorizations are complete.

Line 31 indicates the tunnel endpoints. These are the address of the interface on the switch that accepted the connection and the address of the computer that initiated the connection (most often the ISP assigned address of the PC).

Line 32 indicates the address and netmask that have been assigned to the client.

Lines 33-35 show phase I negotiations are complete and the ISKMP SA has been established. The tunnel endpoints are listed again.

Lines 36-37 indicate the IPsec Security Associations are established and the Security Parameter Indices (SPI) that have been assigned for the outbound (encapsulation) and inbound (decapsulation) sessions. (The CPU that the session was assigned to is also noted, which is only of importance to a 4x00 series unit, which has multiple processors.)

Lines 38-39 show the assigned address and associated netmask are registered with the switch. (The line with the IPvfy message lists the addresses in hex and the octets in reverse order.)

Line 40-44 indicate that the IPSec SAs have been established, as well as the fact that compression has been negotiated.

Line 45 indicates that the switch is assigning IP addresses to clients that it received from a DHCP server. After handing an address to this client, it requests another address from the DHCP server to keep the requested number of addresses in its cache to hand out to clients.

System log

The System log contains all system events that are considered significant enough to be written to disk, including those displayed in the Configuration and Security logs. Events that would appear in the System Log include:

- LDAP activity
- Configuration activity
- Server authentication and authorization requests

The following is the general format of the log entries:

- Time stamp
- Task that issued the event (tEvtLgMgr, tObjMgr, tHttpdTask)
- Number that indicates the CPU that issued the event (0=CPU(0), 1=CPU(1)
- Software module that issued the event
- First number indicates whether it is a persistent event (0=non-persistent, 1=persistent)
- Second number indicates severity level (0=Debug, 1=Low, 2=Medium, 3=High)
- Indicates that the packet matched the rule in the listed section
- Indicates the matching packet source, destination, protocol, and action configured for that rule

The following example shows a System log:

11:29:31 tEvtLgMgr 0 : CSFW [12] Rule[OVERRIDE 1]Firewall: [192.32.250.204:1024-10.0.18.12:2048, icmp], action: Allow

Switch data storage

Accounting records are detailed logs that record the various activities performed by the switch. The logs are directly available from the management interface and can be exported to other applications for additional processing. The Contivity VPN Switch gathers and stores data about the current state of the switch and the connections. The data is stored in files on the switch's hard drive.

- Session Status: RADIUS Accounting–the switch stores copies of RADIUS accounting records.These records, which can be retrieved via FTP or sent to a RADIUS server, contain information about each VPN session initiated to the switch.
- System Data: Data Collection Task–The data collection task runs on the Contivity VPN Switch and gathers data about the system's status. Each minute, data is captured by the task and written to a data file. The information captured by this task is used to create the graphs and reports available from the Status→Reports page.

RADIUS accounting

The Contivity VPN Switch stores copies of RADIUS accounting records. These records are normally sent to a standard RADIUS Accounting server. To configure a RADIUS accounting server, use the Servers→RADIUS Acct screen.

You can view the information in these records on the Status→Accounting page. These are standard RADIUS accounting records. the switch creates a file for each day and keeps the most recent 60 days of data, which are stored in the SYSTEM/ACCTLOG directory.

Accounting data

A data collection record consists of 16 pairs of entries for each data collection object currently being collected. Each value pair consists of a Field ID and an integer value. The following is an sample data collection record:

0-930057960,1-3,2-3,3-0,4-0,5-0,6-0,7-0,8-0,9-0,10-56,11-76,12-1,13-11021,14-40,15-38,16-0

The following field Ids are currently implemented:

Table 30 Field IDs for data collection records

Field identification	Collected field value	Description
0	TIMESTAMP	Seconds since Jan 1, 1970 - 00:00:00 Hours
1	TOTALSESSIONS	Summary of all sessions
2	ADMINSESSIONS	Number of Admin sessions
3	PPTPSESSIONS	Number of PPTP sessions
4	IPSECSESSIONS	Number of IPSEC sessions
5	L2FSESSIONS	Number of L2F sessions
6	L2TPSESSIONS	Number of L2TP sessions
7	IPADDRESSUSE	Number of IP Addresses in use
8	CPUUSE	Un-filtered CPU usage measurement {integer representing a percent between 0 and 100}

Table 30 Field IDs for data collection records

Field identification	Collected field value	Description
9	CPUSMOOTH	Filtered CPU usage measurement {integer representing a percent between 0 and 100}
10	MEMUSE	Filtered memory usage measurement {integer representing a percent between 0 and 100}
11	BOXPACKETSIN	Number of Inbound Packets
12	BOXPACKETSOUT	Number of Outbound Packets
13	BOXBYTESIN	Number of Inbound bytes
14	BOXBYTESOUT	Number of Outbound bytes
15	BOXDROPPEDPACKETS	Number of discarded packets
16	FAILEDAUTHATTEMPTS	Number of failed authentication attempts
17	LASTFIELDID (this field is never written to data record)	

Accounting records

The data collection system stores records in text-based files stored in the system/dclog subdirectory. The system stores the most recent 60 days of data. The system stores daily files, summary files, and summary history files. Ongoing administration tasks include monitoring the configuration files, backing up and restoring the switch or the LDAP database, and upgrading images and clients.

Chapter 11
Troubleshooting

This chapter introduces the concepts and practices of advanced network configuration and troubleshooting for the Contivity VPN Switch. Its purpose is two-fold: to provide configuration details for you to consult when setting up or modifying your extranet, and to serve as a resource when diagnosing client and network problems.

Typically, there are three types of problems to address when managing an extranet:

- Connectivity
- Performance
- General

As a network administrator, your primary concern is maintaining connectivity. For extranet access, this means maintaining the secure connections between your remote users and the private intranet serviced by the switch. Performance is another area of concern. Paying attention to performance can help you address issues before they become problems.

Connectivity problems occur when the remote user cannot establish a connection with areas of their private corporate network. There are several points of failure to consider when diagnosing connectivity problems. Problems can range from something as simple as a modem configuration error on the client workstation to a complex HDLC protocol error on the T-1 WAN interface.

Troubleshooting remote access problems typically starts at the client end when the remote user cannot establish a connection, loses a connection, or has difficulty browsing the network or printing. When connectivity problems occur and the source of the problem is unknown, it is usually best to generally follow the OSI network architecture layers. Therefore, start diagnosing the physical environment, the modem, and cables before moving up to the network and application layers (for example, pinging a host and Web browsing).

As with connectivity, there are many places in the extranet network where network performance can be affected. By regularly checking your network statistics, logs, and health check information, and by informing users of good network practices, you can often avoid problems and enhance the productivity of your extranet.

General problems are categorized here as problems other than those related to connectivity or network performance. For the latest release-specific problems, check the release notes.

Troubleshooting tools

For the Contivity VPN Switch administrator, a robust troubleshooting “toolbox” is filled with both standard and special tools for diagnosing network problems. Standard tools like Telnet, PING, Trace Route (tracert.exe), sniffers, and analyzers are a basic necessity. To this collection, some special tools are added to the Contivity VPN Switch manager and remote client applications. These special tools include client- and switch-based tools.

Client-based tools

IPSec Contivity VPN Client Monitor provides network statistics on device, connection, and network errors that are helpful for monitoring traffic flow and assessing IPSec connection performance. Statistic counters are updated once a second. For more information on the IPSec Contivity VPN Client Monitor, refer to the Contivity VPN Client online Help.

Microsoft PPTP Dial-Up Networking Monitor provides network statistics on device, connection, and network protocols that are helpful for monitoring traffic flow and assessing PPTP connection performance. For more information on the PPTP Dial-Up Networking Monitor, refer to the PPTP help or your Microsoft PPTP client documentation.

Switch-based tools

The Manager Status→Health Check screen provides a detailed picture of how the switch is performing. View colored status indicators to evaluate individual component status, and click on associated hyperlinks to jump to manager screens for corrective action.

The Manager Status→Statistics screen **a**llows you to delve into the inner workings of the switch where you can view detailed system and network statistics.

The Manager Status→Security, Config, System, and Event Log screens allow you to view various logs recording system and network events that help you trace problems and determine their origins. For more information on the Logs screens, refer to *Reference for the Contivity VPN Switch*.

Other tools

Table 31 lists the tools that are helpful for diagnosing connectivity problems from Windows 95, Windows 98, and NT workstations:

Table 31 Troubleshooting tools

Windows 95/Windows 98	Windows NT	Use for...
Winipcfg command	Ipconfig command	Obtaining IP address, DNS, WINS information
Netstats command	Netstats command	Viewing statistics from Microsoft TCP/IP stack
Ping and tracert commands	Ping and tracert commands	Testing connectivity, name resolution, route tracing
Dial-Up Monitor status	Dial-Up Monitor status	Viewing modem settings, throughput and errors

Solving connectivity problems

This section lists many of the common connectivity problems that can occur and the recommended solutions. Problems, and some typical client user responses that can help with diagnosis, are categorized as follows:

Modem and dial-up problems

“I cannot browse the Web or check my e-mail over my dial-up connection.”

“I cannot ping my ISP site.”

Extranet connection problems

“I can browse the Web over my dial-up connection, but I cannot log in to my network over the extranet connection.”

Problems with name resolution using DNS services

“I logged in to my corporate network, but I get messages saying the host is unknown.”

"I can ping the host using its IP address, but not using its host name."

Network browsing problems

"I cannot browse the corporate network."

"I cannot print."

"I cannot access the Internet over my extranet connection."

Diagnosing client connectivity problems

A connection can fail at varying points in an extranet. If remote users have a problem accessing their corporate network, and the source of the problem is unknown, Nortel Networks recommends that they follow these steps to first determine whether the problem is with their modem, PPP dial-up, or with the extranet connection:

1. Confirm that the remote user's modem is attached and working properly by having them run a terminal emulation program at their remote workstation, such as, Hyperterminal*, and issuing the AT command. The response should be AT OK if the modem is operating correctly.
2. Verify that the remote user has a PPP dial-up connection over the internet. To do this, before they try to establish an extranet access or PPTP connection, have them try Web browsing www.nortelnetworks.com or another Web site that they commonly access. If the remote user can access the Web site, their PPP dial-up connection is working properly. Refer to the section "Common client connectivity problems" to further troubleshoot the connection problem. If the remote user still cannot verify that their dial-up connection is working properly, continue with step 3.
3. Ask the remote user to check that their modem type and settings are configured properly. To do this, they need to right-click on the Dial-Up Networking connection icon (the icon they click on to dial their connection) on their desktop to view its properties. Verify that these settings are correct for their modem configuration.
4. If the remote user is connected, but unable to access any resources or servers, have them check their system's connection information by going to the Start menu, selecting Run, and typing winipcfg in the text box (or `ipconfig` if

using Windows NT). Ask them to view the statistics for their PPP adapter and confirm that the entries match those provided by their ISP.

5 If the remote user is still unable to view resources or servers over their PPP dial-up connection, contact their ISP to see if they have logged any connection attempts from the user, and for additional troubleshooting assistance.

Common client connectivity problems

Extranet connection problems

If the client is successfully connecting to their ISP, but is having problems accessing their intranet over their PPTP or IPSec Contivity VPN Client connection, have them check the following areas to further troubleshoot their connection problem.

The following messages and their associated cause and action statements are directed to the IPSec Contivity VPN Client user at the remote workstation. This information is also available in the Contivity VPN Client online Help.

Remote host not responding

Cause: This indicates that the switch never responded to the IPSec connection attempt or that UDP port 500 is blocked.

Action: Verify that the switch is accessible by pinging the host name or IP address that you filled into the destination field. To `ping` a host called extranet.corp.com, for example, open an MS-DOS* command prompt and type `ping extranet.corp.com`. If you receive a reply message, it indicates that the switch is accessible but it is not responding. If you received a message that says Request Timed Out from the `ping` command, it means that the switch is inaccessible. You may be able to further diagnose the problem using the MS-DOS Trace Route command (`tracert.exe`) on Windows systems.

The switch only allows a certain number of PING packets from another Internet host before requiring a tunnel connection to be established.

Maximum number of sessions reached

Cause: This indicates that the maximum number of users for the account you are using are currently logged in.

Action: If you are the only user with access to your account, it is possible to get this error if you restarted an IPSec connection immediately after losing the dial-up connection to your ISP. This is because the switch takes up to one minute to determine that your connection has been dropped and log you off from your account. Simply wait a minute and retry your connection.

Login not allowed at this time

Cause: This indicates that your account has been limited to specific hours of access and you are trying to connect outside of the allowed time.

Action: Contact your network administrator if you are unsure of your specific hours of access.

Authentication failed

Cause: The IPSec user name is incorrect or the password is invalid for the user name entered.

Action: Verify that the user name you entered is correct and retype the password before trying the connection again.

No proposal chosen

Cause: The switch you are connecting to is not configured to handle the authentication method configured under the current connection profile.

Action: Verify that you are using the correct IPSec parameters, such as a choice of ESP-3DES with SHA1. Make sure it matches what the client (for example an International client) is capable of doing.

Other IPSec errors

Cause: Typically other error messages indicate an error in configuration on the switch that must be corrected by the network administrator.

Action: Contact your Network Administrator with the specific error message.

Extranet connection lost

If the PPTP or IPSec Contivity VPN Client connection was initially established and then fails, one of two error messages appear: “The physical connection has been lost” or “The secure extranet connection has been lost.”

The physical connection has been lost

Cause: The PPP connection to your ISP was disconnected.

Action: Re-establish the PPP dial-up connection to your ISP before you re-establish the extranet connection to the remote network.

The secure extranet connection has been lost

Cause: For IPSec only, the switch that you are connected to has either logged your connection off or the switch is no longer responding.

Action: Re-establish the extranet connection by clicking the Connect button. If this works, the connection was probably lost due to the Idle Timeout configured on the switch. If no data is transferred through the extranet connection for a long period of time, normally 15 minutes or more, the switch automatically disconnects the connection.

If you were unable to successfully re-establish the extranet connection, the dial-up connection may be preventing data from traveling between the Extranet VPN Client and the switch. Hang up the dial-up connection and reconnect before you try to re-establish a connection to the switch. If you are still unable to connect to the switch, open an MS-DOS Command Prompt and try pinging the switch using the host name or address that you specified in the Destination field. If you receive a Destination Unreachable error message there is a routing problem at the ISP. If you receive a Request Timed Out error message, the switch is probably not available, and you should contact your network administrator.

Auto disconnect closes the dial-up connection during data transfer activity

Cause: In Windows 95 only, The Microsoft Auto Disconnect feature does not recognize data activity unless it passes through Internet Explorer. Microsoft has documented this as a known problem in Windows 95.

Action: At the remote workstation, disable the Auto Disconnect feature if you are not using Internet Explorer to access data on the remote network. To do this, open the Control Panel and choose the icon labeled Internet. Select the Connection property sheet and deselect the "Disconnect if idle for" box.

Problems with name resolution using DNS services

When the client can ping a host using an IP address, but not with its host name, or receives messages that the host name cannot be resolved, DNS misconfiguration is usually the problem.

Cause: A DNS server may not be configured for PPTP or IPSec connections on the switch.

Action: Validate that the Contivity VPN Client has been configured with a DNS entry. For Windows NT 4.0, open a command prompt and enter `ipconfig/all`. Verify that a DNS server entry is listed. For Windows 95, from the Start menu on the task bar, select Run and enter `winipcfg`. Select Nortel Networks Extranet Switch Extranet Access Adapter from the list of adapters and click on More Info. Record the information displayed under the DNS Server entry and verify it with the network administrator.

Cause: The hostname being resolved has both a public and a private IP address. This is commonly referred to as a split-horizon DNS.

Action: Open a command prompt and ping the host you are trying to reach with a fully qualified host name (for example, www.nortelnetworks.com). If you receive a response, verify that the IP address returned on the first line (for example, www.nortelnetworks.com [207.87.31.127]) is an IP address from the remote corporate network. If it is not, notify your network administrator that the internal hostname should be modified so that it is not the same as the external hostname.

Cause: The retail release of Windows 95 contained a bug that prevented use of more than one DNS server. This problem was fixed in OS Release 2.

Action: If you are using a release earlier than OS Release 2 of Windows 95, a patch is available from Microsoft to upgrade the winsock.dll. This patch is downloadable from www.microsoft.com.

Network browsing problems

Cannot browse the network (with NetBEUI)

Cause: For both PPTP and IPSec, the switch does not currently support the NetBEUI protocol.

Action: To be able to browse resources on a remote domain through a connection to a switch, it is necessary to remove the NetBEUI protocol and have a WINS server configured. By removing NetBEUI, the Microsoft Client uses NetBIOS over TCP/IP to browse network resources. This applies to both the PPTP dial-up client provided by Microsoft and the Contivity VPN Client provided by Nortel Networks.

Cannot access Web servers on the Internet after establishing an Contivity VPN Client connection

Cause: For both PPTP and IPSec, this condition occurs as a result of having all network traffic passed through the corporate network. Typically, firewalls and other security measures on the corporate network limit your access to the Internet.

Action: The switch administrator can set up a default route on the switch to forward traffic to the Internet. If this default route is not configured, you need to disconnect the extranet connection to Web browse the Internet through your ISP connection.

Alternatively, if you are using a proxy-based firewall, you must set your Web browser to use the firewall to proxy for HTTP traffic when your tunnel connection is in use.

Cannot access network shares after establishing an extranet access connection

Cause: A Windows Internet Name Service (WINS) server may not be configured for PPTP or IPSec connections on the switch.

Action: Validate that the Contivity VPN Client has been configured with a WINS server. Follow the steps outlined above under "Problems with name resolution using DNS services" to run `ipconfig` at a command prompt on Windows NT 4.0 or to run `winipcfg` on Windows 95. Verify that a primary WINS server is listed under the section for the adapter named IPSecShm on Windows NT 4.0, and on Windows 95 verify that a primary WINS server is listed in `winipcfg` for the Contivity VPN Client adapter. If there is no primary WINS server listed, notify the network administrator that the switch may not be properly configured.

Cause: Your system may be set up for a different domain other than the one on the remote network.

Action: Skip the initial domain login when Windows 95 starts and choose Log on to the Remote Domain under the Options menu of the Contivity VPN Client dialog box. You are then prompted to log in to the domain of the remote network after the extranet connection is made. This is the recommended method for users with docking station configurations.

Alternatively, on NT 4.0, Windows 98, and Windows 95, change your workstation to be a member of a workgroup instead of a domain:

1. From the Start menu, select Settings→Control Panel. In the Control Panel, double-click on the icon labeled Network. The Network Control Panel applet opens.
2. Select the Identification tab. In Windows 95, the entries on the Identification tab can be modified directly; on NT 4.0 it is necessary to click the Change button to change the entries.
3. Change to use a Workgroup and verify that the computer name does not match the entry on the remote network. The name for the workgroup is not important; you can enter anything.
4. Click on OK to save the changes and reboot the machine.
5. When accessing a resource on the remote domain, if you are prompted for a user name and password, you need to have a domain name precede your user

ID. For example, if your user ID is JSmith and you are accessing a machine on the remote domain named CORP, you would enter your user name as CORP\JSmith.

Diagnosing WAN link problems

WAN link problems can occur between the switch and the Public Data Network (PDN) at three levels:

1. T-1/V.35 interface
2. HDLC framing
3. PPP layer

If a connectivity problem occurs with the WAN link, there are two approaches you can use to diagnose and correct the problem.

- Start from the bottom to verify that physical connectivity exists, then make sure that the HDLC link is up, and finally examine the PPP status to see if it is passing IP packets back and forth.
- Start from the top down to go in the opposite direction, looking at PPP first and working down to the physical connection. An important point to remember when taking this approach is that at the higher protocol layers, there are more options to misconfigure, but changing them is easier and generally involves less effort.

A key point to remember when diagnosing WAN link problems is that the T-1 service provider should usually be involved in the troubleshooting effort. This is not only because they can help diagnose the problem, but also because an ISP can bring down a link if it detects errors on the line. You should notify the ISP administrator if you are planning on working on the link.

Check the T1/V.35 interface

To diagnose a problem at the WAN physical layer, use the following steps to verify that the T-1/V.35 interface to the Public Data Network (PDN) is operating correctly, and that the T-1 line is properly connected:

1. Have your ISP run a loopback test from their end to the CSU/DSU to verify that the external line is working correctly.

2 Check the connections between the switch and the CSU/DSU. The V.35 cable should be a straight-through cable and firmly seated. Also the CSU/DSU should be configured to use internal clocking, and NRZ encoding with CCITT CRC for the checksum.

3 Make sure that all the control signals are asserted (CTS, DCD, DSR, RTS, and DTR). You can check these signals on the switch from the Manager WAN Statistics screen. If any of these signals are incorrect, you can try disabling or enabling the link from the Manager WAN Interfaces screen, or unplugging and plugging in the link. If these steps do not resolve the problem, try switching ports on the same card, switching cables, or switching to a new card, if available.

4 If the previous steps fail to resolve the problem, and you still suspect a problem with the physical connection, try rebooting the switch to reinitialize the WAN interface.

Check the HDLC framing

Assuming that the T-1/V.35 interface is operating correctly, use the following steps to determine whether the HDLC layer is up and running properly, and to provide information for Nortel Networks Customer Support for further diagnosis:

1 Check to make sure that there are no input or output errors reported on the Manager WAN statistics screen. Also look to see if the input and output counters are incrementing at all. If the input/output counters are not incrementing, or are incrementing by huge amounts, then there are probably framing or timing errors on the link. Also, a large percentage of input errors may indicate a problem with the FCS (Frame Check Sequence) calculation.

2 Examine the Manager Statistics event log with debugging enabled. Any WAN-related log messages probably indicate some sort of error.

3 Report any of the preceding errors and messages to Nortel Networks Customer Support for assistance in diagnosing the HDLC framing problem.

Check the PPP layer

If the WAN link appears to be passing frames back and forth, yet IP packets are not flowing, the problem may be with how PPP is configured.

To examine the state of the PPP connection, and to provide information for Nortel Networks Customer Support for further diagnosis:

1. Check to see whether the state of the PPP connection is changing at all by periodically clicking on the Refresh button while viewing the WAN statistics screen. If the state is always Down, PPP might not know that the link is up. If the state toggles between Dead and LCP Negotiating, PPP is trying to come up but cannot. This is probably due to a problem with the underlying layers, although it could also be a bad configuration of the LCP options.
2. If the connection fails during authentication, then try disabling the PPP Authentication settings. A problem during Network Negotiating is usually due to misconfigured IPCP options.
3. Verify that all the authentication settings match the ISP-recommended router configuration.
4. If the PPP layer still does not come up, enable the interface debugger to generate large amounts of packet traces in the event log. Report this information to Nortel Networks Customer Support for further diagnosis.

Hardware encryption accelerator connectivity

If the hardware encryption accelerator fails, all sessions are automatically moved over to be handled by the software.

Solving performance problems

This section describes ways for improving the performance of the remote workstation connection to the corporate network through a switch. It also includes Microsoft networking and client setup and operation tips.

Eliminating modem errors

Modem hardware errors can impact performance when connecting to your corporate network over a dial-up connection. If modem hardware errors are occurring, try the following techniques to correct these errors and improve performance:

- Adjust the modem speed –I f the speed of the modem is set too high it can cause hardware overruns. Reset the modem speed to match the real speed of the modem.
- Disable hardware compression – The data passed through the extranet connection is encrypted, and encrypted data is typically not compressible. Depending on the algorithm the modem is using to compress the encrypted (non-compressible) data, the data may expand in size and overrun the modem's buffers.

Hardware encryption accelerator performance

Expanded bytes and packet counts reflect packets that expanded when compressed by the LZS algorithm. These counters are important because they indicate a heavier load on the accelerator, since packets that expand must be sent a second time to the accelerator with compression disabled. If there are many sessions transporting incompressible traffic (such as a video stream), the overall performance of the switch degrades relative to its performance when all sessions carry compressible data (such as FTP of text files).

Performance tips for configuring Microsoft networking

For Microsoft networking to work as designed over the extranet, each of the following components, if configured, must be working together:

- DHCP Server assigns IP addresses to clients.
- WINS Server provides a translation of the NetBIOS domain name to the IP address.
- DNS Server provides a translation of the IP Host name to the IP address.
- Master Browser is an elected host that maintains lists of all NetBIOS resources.
- Domain Controller maintains a list of all clients in the NetBIOS domain and manages administrative requests such as logins.
- Contivity VPN Switch terminates tunnels and routes Microsoft networking requests.

The following questions and answers are particularly directed toward the WINS server and browsing issues that can help you verify whether you have correctly set up these components.

What needs to be configured on the switch for network browsing?

In the group profiles, set the values of the DNS server and the WINS server. Remember that these are inherited values, so that if all subgroups of a given group use the same servers, it is sufficient to configure them in the parent group.

If these servers are not on a directly reachable subnet from the switch, or accessible through a default gateway, a static route must be configured on the switch in order to reach them.

What should be configured on the PPTP or IPSec client?

The client should have the protocols for NetBIOS and TCP/IP configured. NetBEUI should not normally be configured.

A Windows 95 or Windows 98 client should be configured to be in the correct workgroup for the NT domains it is trying to reach. For example, if there are domains named Engineering and Admin, and the client is to use the Engineering domain, then it must be configured that way in its own configuration.

For PPTP only, you must also select the option Log onto Network under My Computer→Dial Up Networking→Connection_Name properties.

The client system's NetBIOS name must be unique in the private network to which the client is connecting. Do not use the same name as your office desktop machine or something like "my computer." Uniqueness is required.

What is the preferred way to access neighbors on the network?

Microsoft recommends against browsing the Network Neighborhood when tunneling. Another way to access a network resource is through the `run` command. For example, to access shared folders on the machine HotDog, choose Start→Run and type in \\HotDog. If you experience delays using Network Neighborhood, you may want to try this method instead.

Why should WINS settings be different for extranet access?

WINS servers cache a correspondence between IP addresses and NetBIOS names. These cached values are only invalidated by a timer and not by network activity. Therefore, a WINS server that is used heavily by clients should have its expiration timeouts set low.

In a static environment, where names and addresses correspond forever, this is not an issue. But in the extranet environment, clients are assigned new IP addresses whenever they form a tunnel. Therefore, the correspondence is transitory.

Microsoft default values for the timeouts are enormous (for example, 3 weeks). These need to be reduced for an extranet environment.

What WINS settings are recommended?

The WINS settings are available on the WINS server through the Start menu→Programs→Administrator Tools. Nortel Networks is currently experimenting successfully with the following values for a WINS server:

- Server Configuration
- Renewal Interval: 41 minutes
- Extinction Interval: 41 minutes
- Extinction Timeout: 24 hours
- Verify Interval: 576 hours

The renewal interval governs how often a client must reregister its name with the WINS server. It begins trying at one-half of the renewal interval. The extinction interval governs how long it is between the time a client name is released until it becomes extinct. These intervals are the most important to control when using dynamic addresses.

There is a trade-off in setting these intervals. If they are set too small there is too much additional client registration network activity. If they are set too large, transient client entries do not time out soon enough. If you also have secondary WINS servers, the renewal interval should be the same on the secondary servers as on the primary server.

For additional information on setting interval values for a WINS configuration, refer to the Microsoft Knowledge Base article “Min. and Max. Interval Values for WINS Configuration” available at http://support.microsoft.com/support. A WINS server that has a heavy CPU load or network load will not perform well. To help performance:

- Do not run other intensive tasks on the WINS server.
- In the WINS configuration, disable detailed logging.
- If you have primary and secondary WINS servers, try to assign them for a balanced load.

Hosts that never change IP addresses can be given static entries in the WINS database. For example, you could configure the address of the Primary Domain Controller as static. To do this, you also have to have a statically reserved DHCP address for the primary domain controller.

What can you try on the WINS server when it is not working?

You can request that the WINS server clean up its database. You can do this by going into the Mappings menu item and selecting Initiate Scavenging.

If the database becomes very large, it can be compacted using the jetpack.exe program in \winnt\system32. Please consult the WINS Help before doing this because the server must be shut down.

In the WINS mappings entry, enter a `show database` command. Note the entry for -__MSBROWSE__. This is the machine that is actually the elected master browser, and it changes frequently. If this entry is pointing to an invalid machine, this can cause problems.

Can I control which machine is the master browser?

When you start a computer running Windows NT Workstation or Windows NT Server, the browser service looks in the registry for the configuration parameter MaintainServerList to determine whether a computer will become a browser. This parameter is under:

```
\HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Browser\
Parameters
```

For Windows 95, this parameter is under:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\VNETSUP\
MaintainServerList
```

MaintainServerList parameter values are:

- No–This computer will never participate as a browser.
- Yes–This computer will become a browser.
- Auto–This computer, referred to as a potential browser, may or may not become a browser, depending on the number of currently active browsers.

The registry parameter IsDomainMasterBrowser impacts which servers become master browsers and backup browsers. The registry path for this parameter is:

```
\HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Browser\
Parameters.
```

Setting the IsDomainMasterBrowser parameter entry to True or Yes makes the computer a preferred master browser. Whenever a preferred master browser starts, it forces a browser election.

When the browser service is started on the preferred master browser computer, the browser service forces an election. Preferred master browsers are given priority in elections, which means that if no other condition prevents it, the preferred master browser always wins the election. This gives an administrator the ability to configure a specific computer as the master browser.

To specify a computer as the preferred master browser, set the parameter for IsDomainMasterBrowser to True or Yes in the following registry path:

```
\HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Browser\
Parameters
```

Unless the computer is configured as the preferred master browser, the parameter entry is always False or No. There is no user interface for making these changes; you must modify the registry.

Why are subnet masks important?

If a client does not have a WINS server or is unable to contact it, it must broadcast a query to try to locate a host. Unfortunately, Windows 95, Windows 98, and Windows NT clients do not always use the correct broadcast address when tunneling.

The following example helps explain this problem. Suppose that you are using a private net 10 address space. Assume further that you have a client with IP address 10.1.2.3 and subnet mask 255.255.0.0. This means that the net 10 space is being used like a class B address space, which is perfectly legal. The correct broadcast for this client is 10.1.255.255. However, Microsoft clients may broadcast to 10.255.255.255, using the natural class A for net 10, in spite of their configuration.

If all hosts that the client is trying to reach lie on the same physical segment, this probably will work. This is because every host on the physical network receives the all subnets broadcast and probably responds, if appropriate.

All hosts on the segment receive the broadcast to 10.255.255.255, even if they are on different subnets (10.1.x.x. and 10.2.x.x). However, in a routed environment the situation changes. In this case, a broadcast from 10.1.2.3 to 10.255.255.255 is not forwarded to the other 10.2 subnet.

In the extranet environment you should make the remote client appear as much as possible to be on the local LAN. If the extranet host is assigned address 10.1.2.3, it should behave as if it is on the 10.1 LAN.

When 10.1.2.3 broadcasts to find a network neighbor, it (incorrectly) sends to 10.255.255.255. Normal routing functionality would not forward such a packet. The switch finds the best match among its physical interfaces (10.1 in this case) and modifies the broadcast to be correct for that interface (10.1.255.255 here).

In this example, if the switch's 10.1 interface had been configured with any subnet mask other than 255.255.0.0, the broadcast would not have been converted as desired.

What should I do about subnets?

Every private interface on the switch should be configured to have the same subnet mask as all of the clients residing on that subnet.

Why is there a delay in discovering the Network Neighborhood (with tunnels)?

NetBIOS treats the modem interface as if it is two different interfaces: the original modem and the tunnel. It designates the original modem as the primary interface. (You can observe this by typing `route print` in a DOS command shell.) If you tunnel over a LAN instead of a modem, the LAN adapter is designated as the primary interface.

When first instructed to seek the Network Neighborhood, NetBIOS always tries the primary interface first. This is always the wrong choice because NetBIOS always tries to send using the IP address assigned by the ISP (or possibly the address of another adapter) instead of the address assigned to the tunnel by the switch.

The outcome is somewhat different for IPSec and PPTP. For IPSec, the switch client recognizes this incorrect behavior and refuses to even send the packets. You can see a counter of the number of invalid packets of this type on the client under the status Invalid IP address.

With PPTP, the client does send the packets, but they are rejected at the switch as invalid tunneled packets because the source address does not match the switch-assigned address. If you inspect the event log, there are messages of the form Bad source address in tunnel and the session/details counter for source address drops increases.

After about 10 to 15 seconds, NetBIOS gives up on the primary interface, switches to the correct tunnel interface, and starts to browse the Network Neighborhood.

Why can't I browse another client in a different tunnel?

Cause: If you are not using a WINS server, this is not possible because network browsing requires broadcasts from one tunnel to another.

Action: Use a WINS server to browse another client in a different tunnel. When the clients tunnel in, they should register with the WINS server. Be sure that the client you want to browse has the Log onto Network setting enabled under My Computer→Dial Up Networking→Connection_Name properties.

Where can I get more information on troubleshooting dial-up connections?

The Microsoft Knowledge Base article "Dial-Up Networking 1.2 Dun12.doc" file, available from http://support.microsoft.com/support, contains help for resolving common dial-up problems.

Depending on the service provider, some POPs may not support LCP options. If your connection constantly gets declined after the modems synchronize, and you know your password is correct, try disabling this option. The Microsoft Knowledge Base "Service Pack 2 May Cause Loss of Connectivity in Remote Access" article contains more details.

Where can I get more information on configuring PPTP on my client?

There are many articles in the Microsoft Knowledge Base on configuring PPTP for Windows NT, Windows 98, and Windows 95. Refer to the section "Additional information" for a partial list. In addition, Microsoft has the following white papers available at http://support.microsoft.com/support that contain helpful information:

- Microsoft Windows 95/Windows NT White Paper, "Installing, Configuring, and Using PPTP with Microsoft Clients and Servers"
- Microsoft Windows NT Server White Paper, "Understanding PPTP"

You must create a connection definition for your initial Internet link via your service provider. A separate connection definition is needed for creating the PPTP tunnel. A common configuration problem experienced during initial PPTP setup is the failure to select the PPTP VPN adapter (instead of the modem) on the PPTP connection definition in Dialup Networking.

What DNS and WINS servers should I set for the dial-up connection?

There should be no need to set these servers statically on your dial-up client because information is dynamically downloaded from the switch for PPTP, IPSec, and L2F tunnels at connect time.

Why does DNS resolve hosts to different addresses when a tunnel connection is active?

Cause: When a tunnel connection is activated, additional DNS servers are downloaded from the extranet device (for example, the switch) to your client. In the case of Microsoft Windows 95, Windows 98, and Windows NT operating systems, the new DNS servers are added to the list of DNS servers that were assigned by your ISP. This applies to PPTP as well as IPSec tunnels. In general, the DNS servers downloaded by the extranet device provide host-name-to-address translation for hosts within a private network while the ISP-based DNS servers can be expected to translate public host names.

For Windows 95/98 and Windows NT, when a host name must be translated to an IP address (for example to browse the Web or get e-mail), all DNS servers are queried in a shotgun style. The first server to respond with an IP address wins. This can produce some interesting behavior if a host name resolves to one address on the private network and another on the public Internet. For example, host mail.mycompany.com could Internally resolve to 10.0.0.282 and externally to 146.113.64.231.

Action: To avoid problems when using a mixture of internal and external DNS services, it is essential to avoid using names that can resolve to different addresses. In the preceding example, host 10.0.0.282 should probably be renamed pop.mycompany.com. Then users can be informed to use the hostname pop.mycompany.com to retrieve electronic mail, whether in the office or connected via a tunnel link.The original retail release of Windows 95 requires the Winsock DNS Update (wsockupd) to properly function with multiple DNS servers.

My downloaded DNS servers for my tunnel connection do not work

Cause: The Microsoft Windows 95/98 and Windows NT operating systems attempt to ping new DNS servers before adding them to the current list of servers.

Action: As a quick test, try to ping (with the tunnel connection active) the DNS servers that the extranet device is downloading at tunnel startup. If you cannot ping the servers, a basic connectivity problem using the tunnel connection exists.

To view the current list of DNS servers at any time use the MS-DOS command `ipconfig/all` on Windows NT or `winipcfg` on Windows 95 or Windows 98.

Why, after disconnecting a PPTP tunnel, do I get an immediate error reconnecting?

Cause: After you disconnect a PPTP tunnel, then immediately try to reconnect, the PPTP client indicates that the connection is busy or otherwise unavailable. On Windows 95 this is caused by the PPTP control channel socket being improperly shut down by the client.

Action: You can wait for the socket to time out, but it is often more expedient to reboot. On Windows NT a similar problem can be encountered, but caused by a TCP checksum error generated by the Microsoft IP stack. The only current resolution for the Windows NT error condition is to reboot.

Additional information

Below is a list of some of the Microsoft Knowledge Base topics you can browse for information related to dial-up and tunnel configuration. To view these topics, go to http://support.microsoft.com/support. Use the Search Support Online feature to search on the title you want:

- Troubleshooting Internet Service Provider Login Problems
- Service Pack 2 May Cause Loss of Connectivity in Remote Access
- Troubleshooting Modem Problems Under Windows NT 4.0
- Dial-Up Networking 1.2 Dun12.doc File (Windows 95 PPTP Troubleshooting)

- How to Troubleshoot TCP/IP Connectivity with Windows NT
- Remote Access Service (RAS) Error Code List for Windows NT 4.0
- RAS Error 720 When Dialing Out
- Troubleshooting PPTP Connectivity Issues in Windows NT 4.0
- PPTP Registry Entries
- Connecting to Network Resources from Multihomed Computer
- How to Force 128-bit Data Encryption for RAS
- Login Validation Fails Using Domain Name Server

Solving general problems

This section contains general recommendations and explains some common problems that can occur with common Web browsers, the Nortel Networks Contivity VPN Client Manager, and the switch.

Web browser problems and the Contivity VPN Client Manager

If you have a problem browsing the Nortel Networks Contivity VPN Client Manager, start by checking the following recommendations to ensure that you are using the correct Web browser version and settings. For additional troubleshooting, check the described Web browser problems and solutions, error messages, and tips described later in this section.

Nortel Networks Contivity VPN Client Manager uses Java* and HTML features. For the management interface to function properly, verify that your Web browser meets the following minimum requirements:

- Platforms supported include Windows 95, Windows 98, Windows NT, or Macintosh*.
- Display setting of 256 colors or greater.
- Browser versions supported include Microsoft Internet Explorer, Version 4.0 or later and Netscape Communicator,Version 4.0 or later. Not using a recent version of Internet Explorer causes the upper-left corners of the management screens to remain gray rather than displaying the navigational menu and the current menu selection, respectively.

- For ActiveX Scripts, Java, and JavaScript, you must Enable both ActiveX and Java programs in Internet Explorer, and enable both Java and JavaScript in Netscape Communicator for proper switch Web management screens. These options are enabled by default on both Web browsers.

Enabling Web browser options

To make sure these options are enabled in Internet Explorer, from the Internet Explorer menu bar, choose View→Options→Security and choose:

- Run ActiveX scripts–If this option is disabled, navigational titles are not updated, and the Logoff and Help buttons do not work.
- Enable Java programs–If this option is disabled, navigational menus do not appear.

To make sure these options are enabled in Netscape*, from the Netscape menu bar, choose Edit→Preferences→Advanced, and choose:

- Enable Java – If this option is disabled, navigational menus do not appear.
- Enable JavaScript – If this option is disabled, navigational titles are not updated, and the Logoff and Help buttons do not work.

Long delays when Web browsing

Cause: HTTP. Sometimes while using the HTTP Web interface you can experience long delays (greater than five minutes).

Action: Wait until the requested screen is fully delivered before clicking on a new screen request.

Improving performance with Internet Explorer 4.0

Nortel Networks recommends that you create a DNS server entry for your management IP address. This alleviates a noticeable delay in loading the initial Main menu and navigational screens.

Clearing your Web browser cache when upgrading

To avoid problems when upgrading software revision levels (for example, moving from Version 01_01.16 or V01_05_01.28 to V01_00.33), Nortel Networks recommends that you clear your browser cache and exit the browser and all associated windows (such as mail and news readers). Refer to the following section for browser cache clearing instructions.

Clearing cache

A browser caches pages to improve performance when the same page is requested again. The switch's HTTP server allows the browsers to cache the Java class files, and all image files, but not to cache the body pages that contain the dynamically generated information. Both Internet Explorer and Netscape allow you to clear the browser cache which causes all pages to be re-requested the next time they are required. To manually clear the browser cache in Internet Explorer V4.x, select View→Internet Options, and click on Delete Files. To manually clear the browser cache in Netscape V4.x, go to Edit→Preferences →Advanced→Cache and select Clear disk and memory cache.

Web browser error messages

No data in post message

Cause: This message often appears on the main body page if you use the browser's back arrow to revisit a previously displayed page. The browser displays this message when it knows you are revisiting a dynamically generated page.

Action: To see the page, you must use the left navigational area to select it.

Internal error message

Cause: The HTTP server was unable to allocate memory. This indicates that the switch is very low on memory.

Action: Terminate any unnecessary tasks to free up memory. It may be necessary to reboot the switch. If this condition recurs, there may be a serious problem. Contact Nortel Networks Customer Support.

Document not found message

Cause: This message is returned when the HTTP server cannot find the requested page on the switch. This could happen because the Java navigation index file is out of synch with the rest of the system. A corrupted or incorrectly cached index file could also cause this problem.

Action: Clear your browser cache or restart your browser to correct this problem.

New administrator login ignored

Cause: Internet Explorer saves your user ID and password in its cache and automatically resends those values on subsequent login attempts. Therefore, when prompted after an idle timeout, the user ID and password value you enter are ignored, and Internet Explorer sends the original user ID and password. For example, if you log in as administrator with password abc123De, and you subsequently log out. If you log in again, this time as DottieDoe with password FGh45678, in spite of the different login and password, Internet Explorer sends Administrator with passwordabc123De.

Action: When you log off the switch, close out of the Web browser completely (shut down the browser). This clears the cache and the next time that you log in you are starting fresh.

Excess resource consumption using Internet Explorer

Cause: Internet Explorer has a known problem with excessive memory consumption using Java applets. Over time, this problem can cause serious overall system performance degradation.

Action: If you notice that your system's performance seems to slow down for no reason, close and restart Internet Explorer. This releases unused memory and should improve system performance. Go to http://premium.microsoft.com/support/kb/articles/q173/1/45.asp for details.

Internet Explorer 4.0 multiple help windows

Cause: In Internet Explorer 4.0, if you select context-sensitive help and do not close the help window after viewing, you might end up with multiple help windows open.

Action: Close help windows after viewing them.

Distorted background images

Cause: In Netscape versions prior to 4.0 where you configure your Windows 95, Windows 98, or Windows NT system for 8-bit color (256 colors or less), images might appear distorted in the navigational area.

Action: To avoid this situation, increase the color display setting to 256 or greater. You should check with your video card manufacturer's documentation to confirm that your video card supports 256 colors or greater.

Reporting a problem with a Web browser

When reporting a problem with a browser to Nortel Networks, include the following information:

- Workstation operating system and version
- Browser vendor and version (major and minor version)
- Cache setting (size in Netscape, percent of drive for Internet Explorer)
- Verify document setting (every time or once per session)

Switch problems

Excessive active sessions logged

Cause: The number of active sessions can reach more than 4 billion. This is an erroneous number that results from a negative number of sessions.

Action: Restart the system.

Power failure

Cause: The power supplies can become unseated during shipping. When this problem occurs, you might not be able to start the switch, or a warning might be posted to the Status→Health Check screen indicating a potential problem.

Action: If necessary, remove the front bezel as directed in the *Getting Started Guide*, then push the bottom of the power supply in to reseat it.

Cannot convert from an internal address pool to an external DHCP server

Cause: You cannot convert IP address distribution from an internal address pool to an external DHCP server while sessions are active.

Action: Go to the Admin→Shutdown screen, and select Disable Logins after Restart. After everyone has logged off of the switch, then you can convert from an internal address pool to an external DHCP server.

Group and user profile settings not saved

Cause: When you use the Save Current Configurations option on the Admin→Configs screen, it saves only the operational parameters in the configuration file, such as interface IP addresses and subnet masks, backup host IP addresses, DNS names.

Action: To completely back up the switch's configuration, you must also back up the LDAP database, which contains the group and user profiles, filters, and backup file names.

1. Go to the Servers→LDAP screen and click on Stop Server.
2. Enter a file name in the Backup/Restore LDAP Database field. The name should conform to the MS-DOS naming conventions and append the filename with LDF (for example, ldapone.ldf). The restore process can take anywhere from five minutes for a very small LDAP database to several hours for a very large database.
3. You can view the progress of the restoration from the Admin→Health Check screen.

Restart fails after using recovery and reformatting the hard disk

Cause: When you are using the recovery disk and reformatting the hard disk, sometimes the system will not restart.

Action: Power-cycle the system using the green power button on the back of the switch.

Solving routing problems

The following sections describe routing problems.

Client address redistribution problems

The number of current Utunnel host users may display more than the configured maximum.

Cause: This is not an error and is the running state of the system. For example, if you configured a maximum of 200 and have 150 logins, the screen will display the maximum as 200 and the current as 150. If you then modify the maximum to 100, the screen will display the maximum as 100 and the current as 150. As users log out, the current number will eventually be no greater than the maximum.

Action: No action.

Client address redistribution is enabled and the client is logged in, but the client is not communicating with the private network.

Cause: Client address redistribution is not enabled.

Action: Have the client log in again. Client address redistribution only takes effect if the client logs in when it is enabled.

1. Check the Routing→Policy screen to be sure Utunnel routes are enabled.
2. Check to be sure that the routing protocols OSPF and RIP are properly set up.
3. Check to be sure you have the correct address ranges if you configured summarization.
4. Be sure you have an Advanced Routing license if you are using OSPF for client address redistribution.

Solving firewall problems

An error occurred while parsing the policy

Description: The policy that you are attempting to view or edit cannot be opened because it does not conform to the required format. This may be caused by an error in the LDAP database or a problem with the connection to the switch.

Action:

1. Close the Contivity Stateful Firewall Manager.
2. Close all instances of the browser used to load the Contivity Stateful Firewall Manager.
3. Be sure that the connection to the switch is established.
4. Be sure that the LDAP server containing the policy is properly configured and is active.
5. Restart the browser and navigate to the System→Firewall screen.
6. Reload the Contivity Stateful Firewall Manager.

An error occurred while communicating with the switch

Description: The Contivity Stateful Firewall Manager encountered an error while retrieving the data from the switch. This may have been caused by a network error or the switch may have stopped responding.

Action:

1. Close the Contivity Stateful Firewall Manager.
2. Close all instances of the browser used to load the Contivity Stateful Firewall Manager.
3. Be sure that the connection to the switch is established.
4. Restart the browser and navigate to the System→Firewall screen.
5. Reload the Contivity Stateful Firewall Manager.

Authorization failed. Please try again.

Description: This error occurs when the wrong authentication credentials are entered. The user is re-prompted for credentials until they are either correct or the user clicks on Cancel.

Action: No action required.

Unable to communicate with the switch

Description: The Contivity Stateful Firewall Manager cannot establish a connection to the switch. This may have been caused by a network error, or the switch may not be responding to requests.

Action:

1. Close the Contivity Stateful Firewall Manager.
2. Close all instances of the browser used to load the Contivity Stateful Firewall Manager.
3. Be sure that the connection to the switch is established.
4. Restart the browser and navigate to the System→Firewall screen.
5. Reload the Contivity Stateful Firewall Manager.

The contents of the database may have changed...

Description: This error occurred because the LDAP database has changed in such a way that the current data in the Contivity Stateful Firewall Manager might not be valid. This error is encountered when the:

- Internal LDAP server has been shut down and restarted.
- External LDAP server in use is switched to the internal LDAP server.
- Internal LDAP server in use is switched to an external LDAP server.
- External LDAP server's port or IP address changes.

Action:

To ensure that the most current data is loaded:

1 Close the current policy, if opened. Saving is not permitted until this error is remedied.

2 From the policy selection screen, select All from the Refresh menu.

System files were not loaded properly

Description: This error occurred because the files necessary to load the Contivity Stateful Firewall Manager were either not downloaded from the switch properly or were not initialized properly.

Action:

If this error is encountered:

1 Close the Contivity Stateful Firewall Manager.

2 Close all instances of the browser used to load the Contivity Stateful Firewall Manager.

3 Restart the browser and navigate to the System→Firewall screen.

4 Reload the Contivity Stateful Firewall Manager.

If the error continues to occur or if the Contivity Stateful Firewall Manager is being accessed through a user tunnel:

1 Open the Java Plug-in Properties.

2 On Windows systems, navigate to Start→Settings→Control Panel→Java Plug-in. For all other systems, refer to the Java Plug-in documentation.

3 Be sure that the check box for Cache JARs in Memory is deselected.

4 Click on Apply and close the Java Plug-in Properties window.

5 Close the Contivity Stateful Firewall Manager.

6 Close all instances of the browser used to load the Contivity Stateful Firewall Manager.

7 Restart the browser and navigate to the System→Firewall screen.

8 Reload the Contivity Stateful Firewall Manager.

Appendix A
MIB Support

The Contivity VPN Switch supports the management information base (MIB) for use with network management protocols in TCP/IP-based Internets and TCP/IPX-based networks. The switch only supports SNMP Gets. It does not support SNMP Sets.

Nortel Networks also provides proprietary MIBs for the switch's SNMP trap support. The MIBs, cestraps.mib and newoak.mib, are available on the Contivity VPN Switch distribution CD in the Doc directory.

SNMP RFC support

This section discusses the SNMP-related RFCs that the Contivity VPN Switch supports.

Novell IPX MIB

The switch supports the IPX MIB that is distributed by Novell, Inc.

Novell RIP-SAP MIB

The switch supports the IPX RIP-SAP MIB that is distributed by Novell, Inc.

RFC 1850 -- OSPF Version 2 Management Information Base

The switch supports RFC 1850, *OSPF Version 2 Management Information Base*. As stated in the introduction to the RFC, the RFC "defines a portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing the Open Shortest Path First Routing Protocol."

RFC 1724 -- RIP Version 2 MIB Extension

The switch supports RFC 1724, *RIP Version 2 MIB Extension*. As stated in the introduction to the RFC, the RFC "defines a portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it defines the objects for managing RIP Version 2."

RFC 1213 -- Network Management of TCP/IP-Based Internets MIB

The switch supports RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB II*. This RFC provides the architecture and system for managing TCP/IP-based internets. With the exception of the EGP Group (Section 6.10) and the Transmission Group (Section 6.11), the switch provides full support for the RFC.

RFC 2667 -- IP Tunnel MIB

The switch supports RFC 2667, *IP Tunnel MIB*. As stated in the introduction to the RFC, it "describes a Management Information Base (MIB) used for managing tunnels of any type over IPv4 networks, including GRE [16,17], IP-in-IP [18], Minimal Encapsulation [19], L2TP [20], PPTP [21], L2F [25], UDP (e.g., [26]),ATMP [22], and IPv6-in-IPv4 [27] tunnels."

RFC 2787 -- VRRP MIB

The switch supports RFC *2787, Definitions of Managed Objects for the Virtual*

Router Redundancy Protocol. As stated in the introduction to the RFC, it "defines an extension to the Management Information Base (MIB) for use with SNMP-based network management. In particular, it defines objects for configuring, monitoring, and controlling routers that employ the Virtual Router Redundancy Protocol (VRRP)."

RFC 2737 -- Entity MIB

This MIB contains five tables, we have partially implemented two of these tables.

```
*entPhysicalTable
entLogicalTable
entLPMappingTable
*entAliasMappingTable
entPhysicalContainsTable
```

The entPhysicalTable provides a listing of the hardware elements that are present in the system. For example each slot is listed, and if there is a card in the slot, then the card and any ports on the card. The exception to this is the hardware accelerator which does not appear in the table). The listing shows element relationships via the columns entPhysicalContainedIn and entPhysicalParentRelPos. The only columns that have been implemented are:

```
entPhysicalIndex
entPhysicalDescr (although the value is not strictly what the MIB
specifies)
entPhysicalContainedIn
entPhysicalClass
entPhysicalParentRelPos
entPhysicalName
entPhysicalIsFRU
```

All other columns will return an appropriate default value for the object.

The entAliasMappingTable provides a mapping from entPhysicalIndex to ifTable.ifIndex. Hence by walking this table, a management station can deterring the ifIndex associated with a physical port.

RFC 1573 -- IanaIfType MIB

This MIB contains the enumerations for rfc2233 ifTable.ifType. These enumerations describe the various types of interfaces that ifTable can support.

RFC 2233 -- If MIB

This MIB is the latest evolution of rfc1213 Interfaces group, plus several new objects.

RFC 2571-- Snmp-Framework MIB

This MIB provides textual conventions and object definitions used in the SNMP agent architecture.

RFC2790 -- Host Resources MIB

The Host Resources MIB defines a uniform set of objects for the managing host computers. Host computers are independent of the operating system, network services, or any software application.The Host Resources MIB defines objects that are common across many computer system architectures.

The Contivity VPN Switch does not support the following groups/objects:

- hrSystem Group
  - hrSystemInitialLoadDevice
  - hrSystemInitialLoadParameters
  - hrSystemNumUsers
  - hrSystemProcesses
  - hrSystemMaxProcesses
- hrStorage Group

  hrStorageAllocationFailures
- hrDevice Group
  - hrDevice Table

    hrDeviceErrors

  - — hrNetworkTable
  - — hrPrinterTable
  - — hrDiskStorageTable

    hrDiskStorageCapacity
  - — hrPartitionTable

    hrPartitionSize
  - — hrFSTable

    hrFSLastFullBackupDate

    hrFSLastPartialBackupDate
- hrSWRun Group

  hrSWRun
- hrSWRunPerf Group

  hrSWRunPerf

RFC2495 -- DS1 MIB

These objects are used when you use a DS1/E1/DS2/E2 interface. At present, this applies to the ifType variable in the Internet-standard MIB ds1 (18).

This MIB provides an alternative reporting method for monitoring line status on a T1 line. ANSI reporting is still supported, but the reporting method is either ANSI or DS1 MIB.

CES MIB

This MIB contains CES proprietary MIB data. For instance the “ping MIB” is contained in this file. The ping MIB, via an SNMP GET REQUEST, causes the CES to ping another device and get statistics based on the results of the ping. For instance sending a PDU specifying pingAverageTime.192.32.250.248.4.4076, sends four pings, of 4076 bytes, to address 192.32.250.248. (It actually sends five

because one ping is sent by itself so that if the device being pinged is the other end of a Branch Office tunnel, it ensures that the tunnel is brought up before trying to send pings through the tunnel. This ping is not counted in the statistics.) The object returns the values of:

```
-2 Invalid parameter(indices).
-1 No reply.
0 Less than 16ms average time.
>0 The average time.
```

The objects and their parameters(indices) are:

```
pingAverageTime - returns the average ping time for the set of
specified pings.
pingPercentLoss - returns the percentage of loss.
```

The first index is the IP address to ping. The second index is the number of pings, if this is not specified or is an invalid value it defaults to 3. The third index is the size of the ping request. If it is not specified or is an invalid value then it defaults to 1024.

CES MIB has been expanded to provide trap acknowledgement.

cestraps.mib -- Nortel Networks proprietary MIB

This section lists the contents of the cestraps.mib, the Nortel Networks MIB for the Contivity VPN Switch.

```
-- Each Trap contains the Trap OID as well as the following OIDs:
--     SeverityLevel
--     System Name
--     System Date
--     System Time
--     System Uptime
--
       NEWOAKTRAP DEFINITIONS ::= BEGIN

       IMPORTS
          enterprises                FROM RFC1155-SMI
          DisplayString              FROM RFC1213-MIB
          OBJECT-TYPE                FROM RFC-1212
          TRAP-TYPE                  FROM RFC-1215;

   -- This MIB module uses the extended OBJECT-TYPE macro as
   -- defined in [9], and the TRAP-TYPE macro as defined in [10].

   contivity         OBJECT IDENTIFIER ::= { enterprises 2505 }


   ContivitySnmpTraps OBJECT-TYPE
      SYNTAX  DisplayString
      ACCESS  read-only
      STATUS  mandatory
      DESCRIPTION
      "Nortel Networks Inc's Enterprise trap."
   ::= {contivity 1}
```

newoak.mib

This section provides the contents of the newoak.mib, which defines the "newoak" enterprise ID, the "contivity" object identifier, and the sysObjectIDs for each of the switch models.

```
--  This MIB module uses the extended OBJECT-TYPE macro as
        --  defined in [9], and the TRAP-TYPE macro as defined in
[10].

        newoak    OBJECT IDENTIFIER ::= { enterprises 2505 }

-- The following MODULE-IDENTITY definition can be commented out if
the MIB parser
-- you are using has trouble parsing it. If you do comment it out,
then uncomment
-- the following object identifier defintion.
--    contivity OBJECT IDENTIFIER ::= {newoak 1}
--
     contivity  MODULE-IDENTITY
          LAST-UPDATED "0004252130Z"  -- April 25, 2000 7:30pm EST
          ORGANIZATION "Nortel Networks,Inc."
          CONTACT-INFO
            "support@nortelnetworks.com
            Postal: Nortel Networks,Inc.
                    80 Central St.
                    Boxboro, MA 01719
            Tel:    +1 978 264 7100
            E-Mail: support@nortelnetworks.com"

          DESCRIPTION
              "This MIB defines the sysObjectIDs for different
               variations ofthe Convitiy Extranet Switch."
          ::= { newoak 1 }
-- IDENTIFIER ::= {newoak 1}
     contivityExtranetSwitch2000 OBJECT IDENTIFIER ::= {newoak 2}
     contivityExtranetSwitch1000 OBJECT IDENTIFIER ::= {newoak 3}
     contivityExtranetSwitch4500 OBJECT IDENTIFIER ::= {newoak 4}
     contivityExtranetSwitch15XX OBJECT IDENTIFIER ::= {newoak 5}
     contivityExtranetSwitch2500 OBJECT IDENTIFIER ::= {newoak 6}
     contivityExtranetSwitch2600 OBJECT IDENTIFIER ::= {newoak 7}
     contivityExtranetSwitch1600 OBJECT IDENTIFIER ::= {newoak 8}
     contivityExtranetSwitch4600 OBJECT IDENTIFIER ::= {newoak 9}

      END
```

Hardware-related traps

```
hardwareTrapInfo OBJECT IDENTIFIER
    ::= {ContivitySnmpTraps 1}

-- Trap #1001
hardDisk1Status OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS mandatory
    DESCRIPTION "Hard Disk Number 1 Status."
    ::= {hardwareTrapInfo 1}

-- Trap #1002
hardDisk0Status OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Hard Disk Number 0 Status."
    ::= {hardwareTrapInfo 2}

-- Trap #1003
memoryUsage OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Memory Usage Status."
    ::= {hardwareTrapInfo 3}

-- Trap #1004
LANcardStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of any LAN cards on the system."
    ::= {hardwareTrapInfo 4}

-- Trap #1005
CPUtwoStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of second CPU."
    ::= {hardwareTrapInfo 5}

-- Trap #1006
fanOneStatus OBJECT-TYPE
    SYNTAX  DisplayString
```

```
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of the first CPU fan."
    ::= {hardwareTrapInfo 6}

-- Trap #1007
fanTwoStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of the second CPU fan."
    ::= {hardwareTrapInfo 7}

-- Trap #1008
chassisFanStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of the Chassis fan."
    ::= {hardwareTrapInfo 8}

-- Trap #1009
fiveVoltsPositive OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of +5 Volt power."
    ::= {hardwareTrapInfo 9}

-- Trap #10010
fiveVoltsMinus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of -5 Volt power."
    ::= {hardwareTrapInfo 10}

-- Trap #10011
threeVoltsPositive OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of +3 Volt power."
    ::= {hardwareTrapInfo 11}

-- Trap #10012
twoDotFiveVA OBJECT-TYPE
    SYNTAX  DisplayString
```

```
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of 2.5VA power."
    ::= {hardwareTrapInfo 12}

-- Trap #10013
twoDotFiveVB OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of 2.5VB power."
    ::= {hardwareTrapInfo 13}

-- Trap #10014
twelveVoltsPositive OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of +12 Volt power."
    ::= {hardwareTrapInfo 14}

-- Trap #10015
twelveVoltsMinus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of -12 Volt power."
    ::= {hardwareTrapInfo 15}

-- Trap #10016
normalTemperature OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of normal temperature reading."
    ::= {hardwareTrapInfo 16}

-- Trap #10017
criticalTemperature OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of critical temperature reading."
    ::= {hardwareTrapInfo 17}

-- Trap #10018
chassisIntrusion OBJECT-TYPE
    SYNTAX  DisplayString
```

```
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "The chassis intrusion sensor indicates that
                 the unit has been opened."
    ::= {hardwareTrapInfo 18}

-- Trap #10019
dualPowerSupply OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of the redundant power supplies."
    ::= {hardwareTrapInfo 19}

-- Trap #10020
t1WANStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of T1 WAN card(s)."
    ::= {hardwareTrapInfo 20}

-- Trap #10021
t3WANStatus OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of T3 WAN card(s)."
    ::= {hardwareTrapInfo 21}
```

Server-related traps

```
serverTrapInfo OBJECT IDENTIFIER
    ::= {ContivitySnmpTraps 2}

-- Trap #3001
radiusAcctServer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of External Radius Accounting Server."
    ::= {serverTrapInfo 1}

-- Trap #3002
backupServer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of External Disk Backup Server."
    ::= {serverTrapInfo 2}

-- Trap #3003
diskRedundency OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of Local Disk Redundancy."
    ::= {serverTrapInfo 3}

-- Trap #3004
IntLDAPServer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of Internal LDAP Server."
    ::= {serverTrapInfo 4}

-- Trap #3005
LoadBalancingServer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of Load Balancing Server."
    ::= {serverTrapInfo 5}

-- Trap #3006
DNSServer OBJECT-TYPE
    SYNTAX  DisplayString
```

```
      ACCESS  read-only
      STATUS  mandatory
      DESCRIPTION "Status of DNS Server."
      ::= {serverTrapInfo 6}

   -- Trap #3007
   SNMPServer OBJECT-TYPE
      SYNTAX  DisplayString
      ACCESS  read-only
      STATUS  mandatory
      DESCRIPTION "Status of SNMP Server."
      ::= {serverTrapInfo 7}

   -- Trap #3008
   IPAddressPool OBJECT-TYPE
      SYNTAX  DisplayString
      ACCESS  read-only
      STATUS  mandatory
      DESCRIPTION "Status of the IP address pool."
      ::= {serverTrapInfo 8}

   -- Trap #3009
   ExtLDAPServer OBJECT-TYPE
      SYNTAX  DisplayString
      ACCESS  read-only
      STATUS  mandatory
      DESCRIPTION "Status of External LDAP Server."
      ::= {serverTrapInfo 9}

   -- Trap #30010
   radiusAuthServer OBJECT-TYPE
      SYNTAX  DisplayString
      ACCESS  read-only
      STATUS  mandatory
      DESCRIPTION "Status of Radius Authentication Server."
      ::= {serverTrapInfo 10}

-- Trap #30011
     certificateServer OBJECT-TYPE
         SYNTAX  DisplayString
         ACCESS  read-only
         STATUS  mandatory
         DESCRIPTION "Status of Certificates Validity."
         ::= {serverCESTrapInfo 11}
```

Software-related traps

```
softwareTrapInfo OBJECT IDENTIFIER
    ::= {ContivitySnmpTraps 3}

-- Trap #5001
NetBuffers OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Network buffer usage."
    ::= {softwareTrapInfo 1}

-- Trap #5002
fireWall OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Status of internal firewall."
    ::= {softwareTrapInfo 2}
```

Login-related traps

```
loginTrapInfo OBJECT IDENTIFIER
    ::= {ContivitySnmpTraps 4}

    -- Trap #101
    failedLogin OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Failed Login Attempt."
    ::= {loginTrapInfo 1}
```

Intrusion-related traps

```
intrusionTrapInfo OBJECT IDENTIFIER
  ::= {ContivitySnmpTraps 5}

-- Trap #201
securityIntrusion OBJECT-TYPE
  SYNTAX  DisplayString
  ACCESS  read-only
  STATUS  mandatory
  DESCRIPTION "Login Security Intrusion."
  ::= {intrusionTrapInfo 1}
```

System-related traps

```
-- Trap #401
powerUpTrap OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Power Up."
    ::= {ContivitySnmpTraps 6}

-- Trap #601
periodicHeartbeat OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "Periodic Heartbeat."
    ::= {ContivitySnmpTraps 12}
```

Information passed with every trap

```
SeverityLevel OBJECT-TYPE
    SYNTAX INTEGER
    {
    fatal(1),
    major(2),
    minor(3),
    informational(4),
    insignificant(5),
    reversal(6)
    }
        ACCESS read-only
        STATUS mandatory
        DESCRIPTION "Severity of specific trap."
        ::= {ContivitySnmpTraps 7}
systemName OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "System Name."
    ::= {ContivitySnmpTraps 8}

systemDate OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "System Date."
    ::= {ContivitySnmpTraps 9}

systemTime OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "System Time."
    ::= {ContivitySnmpTraps 10}

systemUpTime OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION "System Up Time."
    ::= {ContivitySnmpTraps 11}
```

The following table provides explanations for the traps.

Table 32 Trap categories

Hardware	
1.3.6.1.4.1.2505.1.1.0.1001	hardDisk1StatusTrap
1.3.6.1.4.1.2505.1.1.0.1002	hardDisk0StatusTrap
1.3.6.1.4.1.2505.1.1.0.1003	memoryUsageTrap
1.3.6.1.4.1.2505.1.1.0.1004	lanCardStatusTrap
1.3.6.1.4.1.2505.1.1.0.1005	cpuTwoStatusTrap
1.3.6.1.4.1.2505.1.1.0.1006	fanOneStatusTrap
1.3.6.1.4.1.2505.1.1.0.1007	fanTwoStatusTrap
1.3.6.1.4.1.2505.1.1.0.1008	chassisFanStatusTrap
1.3.6.1.4.1.2505.1.1.0.1009	fiveVoltsPosStatusTrap
1.3.6.1.4.1.2505.1.1.0.10010	fiveVoltsMinusTrap
1.3.6.1.4.1.2505.1.1.0.10011	threeVoltsPositiveTrap
1.3.6.1.4.1.2505.1.1.0.10012	twoDotFiveVATrap
1.3.6.1.4.1.2505.1.1.0.10013	twoDotFiveVBTrap
1.3.6.1.4.1.2505.1.1.0.10014	twelveVoltsPositveTrap
1.3.6.1.4.1.2505.1.1.0.10015	twelveVoltsMinsTrap
1.3.6.1.4.1.2505.1.1.0.10016	normalTemperatureTrap
1.3.6.1.4.1.2505.1.1.0.10017	criticalTemperatureTrap
1.3.6.1.4.1.2505.1.1.0.10018	chassisIntrusionTrap
1.3.6.1.4.1.2505.1.1.0.10019	dualPowerSupplyTrap
1.3.6.1.4.1.2505.1.1.0.10020	t1WANStatusTrap
1.3.6.1.4.1.2505.1.1.0.10021	t3WANStatusTrap
1.3.6.1.4.1.2505.1.1.0.10022	hwAccelTrap
Server	
1.3.6.1.4.1.2505.1.2.0.3001	radiusAcctServerTrap
1.3.6.1.4.1.2505.1.2.0.3002	backupServerTrap
1.3.6.1.4.1.2505.1.2.0.3003	diskRedundencyTrap
1.3.6.1.4.1.2505.1.2.0.3004	intLDAPServerTrap
1.3.6.1.4.1.2505.1.2.0.3005	loadBalancingServerTrap
1.3.6.1.4.1.2505.1.2.0.3006	dnsServerTrap

Table 32 Trap categories

Server	
1.3.6.1.4.1.2505.1.2.0.3007	snmpServerTrap
1.3.6.1.4.1.2505.1.2.0.3008	ipAddressPoolTrap
1.3.6.1.4.1.2505.1.2.0.3009	extLDAPServerTrap
1.3.6.1.4.1.2505.1.2.0.30010	radiusAuthServerTrap
1.3.6.1.4.1.2505.1.2.0.30011	certificateServerTrap
Software	
1.3.6.1.4.1.2505.1.3.0.5001	netBuffersTrap
1.3.6.1.4.1.2505.1.3.0.5002	FireWallTrap
1.3.6.1.4.1.2505.1.3.0.5003	FipsStatusTrap
Failed Login	
1.3.6.1.4.1.2505.1.4.0.101	FailedLoginTrap
Intrusion	
1.3.6.1.4.1.2505.1.5.0.201	SecurityIntrusionTrap
Presence	
1.3.6.1.4.1.2505.1.0.401	PowerUpTrapEntry
1.3.6.1.4.1.2505.1.0.601	PeriodicHeartbeatTrap

Appendix B
Using Serial PPP

The Serial Point-to-Point Protocol (PPP) feature allows you to manage the switch from a remote location using PPP and the serial interface. If the switch were to become unreachable over the Internet, you could still dial up and manage it through the serial interface menu.

With this feature, the serial interface becomes much like a private WAN interface. You can manage through it or even tunnel through it. You can enable Serial PPP support on the switch using the Web interface (System→Settings). When configuring Serial PPP, you can set the switch to Auto Detect, or you can specify that either PPP or the Serial Menu are the options available through the serial port.

Serial PPP authentication is performed by the Password Authentication Protocol (PAP), which uses a standard user ID and a password that is sent in the clear. When authenticated, the serial interface acts like a private WAN interface.

Establishing a serial PPP connection

To enable Serial PPP:

1. Set up a Dial-Up Networking connection.
2. Set up the modem.
3. Set up the switch.
4. Dial into the switch using the Primary Administrator's user name and password.

Setting up a Dial-Up Networking connection

To establish a Serial PPP connection using a Microsoft Dial-Up Networking connection from the client system:

1. Double-click on My Computer.
2. Double-click on the Microsoft Dial-Up Networking icon.
3. Set the COM port baud rate on the client system to be compatible with the switch's baud rate. The rates should be the same to establish a connection. Possible rates are:
   - 9600 (default)
   - 19200
   - 38400
   - 56000
4. Go to Server Types, and under Type of Dial-Up Server, select PPP: Internet, Windows NT Server, Windows 95. Make sure that none of the Advanced options are set.
5. Go to Allowed network protocols, and select TCP/IP.
6. Go to TCP/IP Settings, and specify your IP address. This is the Management IP address that the switch uses to communicate with the client that is dialing in through the modem.
7. Click Server Assigned name server addresses.
8. Unclick IP header compression.
9. Click Use default gateway on remote network.
10. Do not configure Scripting and Multilink.
11. Click Configure the client modem, and use the following settings:
    - 8 data bits
    - 1 stop bit
    - No parity
    - Hardware flow control

    Do not choose Log On to Network if the selection appears.

Setting up the modem

The following procedure assumes you are using a 3Com/US Robotics 56K x2 modem. It describes how to set up a modem to communicate with the switch using a Dial-Up Networking connection.

Table 33 DIP switches configuration

Parameter	Setting
Data Terminal Ready	On
Verbal Result Codes	On
Suppress Result Codes	On
Echo Offline Commands	Off
Auto Answer (must be set)	On
Carrier Detect Normal	On
Load NVRAM Defaults	On
Dumb Mode	Off

Setting up the switch

To set up the switch's parameters through the Web interface (System→Settings):

1 Go to the System→System Settings screen and select one of the following modes of operation under the Serial Port option:

- Serial Menu (default)--leaves the switch's serial interface in the traditional serial menu mode. In this mode, no serial PPP is supported. When connecting a program such as Hyper Terminal to the interface, the standard serial interface menu appears. In Auto Detect Mode, if you are using a terminal emulator, such as Hyper Terminal, you must press Enter several times to get the logon and password prompt. Also, you can ignore the modem initialization string (which might not be in use) that is displayed on the Hyper Terminal screen.
- PPP--you can set up the switch to use the Point-to-Point Protocol (PPP) protocol over the serial port. This feature allows you to manage the switch from a remote location using PPP and the serial interface. If the switch were to become unreachable over the internet, you could still dial up and manage it through the serial interface menu. This feature allows you to

access all of the management services (HTTP, Telnet, FTP, SNMP) through the Web interface. Once a session is established through PPP, the serial interface acts as a private WAN interface with an internal IP address (0.0.1.35).

- Auto detect--automatically detects whether the connected device is using PPP or serial menu mode at startup. The switch cannot determine the device's baud rate, nor can it determine a change from PPP to serial menu mode, except upon startup. Auto Detect checks the mode each time the switch is restarted. When performing its Auto Detect check, the switch sends out AT command set characters to configure a modem if one is attached.

  When the switch is in Auto Detect mode, and if a terminal session is connected and the terminal baud rate is the same as the switch's, the terminal displays the AT command sets on the screen. Simply press Enter several times until a serial menu session starts. You should use the Auto Detect Mode rather than PPP Mode. Using PPP mode could leave the switch in a state such that you could never manage it from the serial interface menu directly. If this happened, you would still have to manage the switch through a PPP application (such as Dial-Up Networking). Directly connecting a serial cable and running Hyper Terminal would not work because the interface would only recognize PPP.

2 Select one of the following Baud Rates to match the baud rate of your terminal. After you select the baud rate, you must click the Reset button to change the port to the selected baud rate. This option is necessary for PPP if a modem initialization string specifies a fixed baud rate.

- 57600
- 38400
- 19200
- 9600 (default)

3 Enter the modem initialization string. Refer to the manufacturer's documentation to learn the vendor-specific character initialization string. Preconfiguring the modem and using the switch's default initialization string (ATZ) provides the best results.

A sample 3Com/US Robotics 56K modem initialization string that instructs the external modem to connect at 19,200 Kb/s is ATZAT&B1AT&N10.

4 Click the Reset button to reset the port to the selected baud rate and apply any other modem changes.

Dialing in to the switch

Use the standard dial-up networking procedure to connect to the switch. After connecting, you can then manage the switch using either Telnet (for the command line interface) or the browser-based GUI. Use the switch's management IP address for the telnet session or the browser's destination URL.

Troubleshooting Serial PPP

When the serial port is set up for PPP only, you can still do in-band Web management.

Cause:

I have a modem connected, but I cannot get a PPP connection.

Actions:

- Verify that the modem supports the switch's selected baud rate. Most connection problems occur because the modem is not operating at the same baud rate as the switch. For example, a 3Com/US Robotics 56 Kb/s modem's default baud rate when attempting to establish a connection to the switch is 38400, but the switch's default baud rate is 9600.
- Verify that the switch is set up do PPP over the serial port. You can verify this by checking the settings in the Web interface (System→Settings).
- Verify that you clicked Reset from the Web interface when making changes to the screen (System→System Settings). This guarantees the serial port resets and initializes the modem. This is especially true with a modem connected to a switch that was restarted.
- Check the event log for failures.
- Make sure you have the correct dial-up networking settings. Refer to the section, "Setting up a Dial-Up Networking connection."
- Make sure you have the remote modem set to auto answer and that it is in smart mode so that it can respond to the AT command set.
- Verify that the auto detection did not fail, and that the switch is in serial menu mode.

Cause:

You were dialed in and managing the switch remotely using PPP and you changed the baud rate and applied it, but now you cannot manage the switch.

Action:

To manage the switch, disconnect the dial-up connection and attempt to reestablish it. This gives the modem a chance to renegotiate the baud rate with the switch.

Cause:

You are set up to use PPP but want to use the serial port for the serial menu.

Action:

Choose the serial port mode Serial Menu. Press OK using the Web management interface (System→System Settings) and restart the switch. A serial cable must be installed in place of the modem in order to use the Serial Menu. Remember to power off the switch when plugging in and unplugging the serial port connection; otherwise, you might damage system components.

Cause:

You are set up to use the Serial Menu but want to use the port for PPP.

Action:

You can change the serial port settings (System→System Settings) or the Serial Menu itself. For these changes to take effect, restart the switch. For the best results, connect the modem while the switch is turned off.

Cause:

You are using a dial-up serial PPP connection and you encounter repeated CRC errors.

Action:

Make sure that the modem that is connected to the switch has hardware flow control enabled.

PPP option settings

These settings describe the switch's behavior when negotiating Serial PPP.

For IP:

- IP Address negotiation is enabled.
- The switch needs the peer's IP address to make a connection.
- The peer should not suggest an IP address for the switch. The switch uses its management IP address.
- The switch rejects VJ compression.
- The switch rejects VJ connection ID compression.

For LCP:

- The switch does not initiate a connection.
- The switch accepts magic number negotiation.
- The switch rejects address control field compression.
- The switch rejects protocol field compression.
- The switch does not allow asynchronous character map to be negotiated.
- The switch accepts Maximum Receive Unit (MRU) requests.

For authentication:

- The switch does not authenticate itself to a peer with PAP upon request.
- The switch requires that peers perform PAP authentication using the administrator's login and password.
- The switch does not authenticate itself to a peer with CHAP upon request.
- The switch does not require that the peer authenticate itself with CHAP.

Appendix C
System Messages

System forwarding (syslog) enables you to forward information from your switch's system log to different host machines using the system logging daemon (syslogd).

This appendix provides a listing of possible syslog messages that the Contivity VPN Switch might write to a remote system. Each message is followed by a description and the recommended corrective action, if any.

Certificate messages

Error removing CA certificate file: xxx

Description: The switch can be manufactured with a trusted CA certificate for use by SSL. The temporary manufacturing file containing the certificate is removed the first time you boot the switch. This error message indicates that the switch is unable to remove the temporary certificate file. The error might be caused by a general problem with the local file system.

Action: Manually delete all files in the /system/cert/ca directory.

Installed new CA certificate from file: *xxx*

Description: The switch can be manufactured with trusted CA certificates for use by SSL. This informational message indicates a trusted SSL CA certificate was installed when the switch was manufactured.

Action: No action required.

tCert: Shutdown complete

Description: This informational message indicates that the task responsible for certificate maintenance has shut down. This is usually part of the normal system shutdown.

Action: No action required.

tCert: task creation failed

Description: The task responsible for X.509 certificate maintenance on the switch failed to start properly. This most likely indicates severe resource exhaustion on the switch.

Action: Reboot the switch. If the reboot does not fix the problem, contact Nortel Networks Technical Support.

tCert: X.509 certificates disabled in flash memory

Description: This is an informational message that indicates the use of X.509 certificates by the switch has been totally disabled.

Action: No action required.

Warning: System CA certificates may have been tampered with, please reinstall!

Description: The switch performs a periodic integrity check of the SSL-related X.509 certificates that are stored on the switch's local file system. This message signals a failure during the integrity check. This indicates that one or more of the SSL-related certificates might have been tampered with, or that a certificate has been corrupted.

Action:

1 Delete, then reinstall any SSL-related certificates. It is not necessary to delete and reinstall the tunnel-related certificates since they are stored in the LDAP database and not on the local file system.

2 Manually verify the tunnel-related certificate fingerprints. You should perform this procedure any time you suspect tampering.

ISAKMP messages

ISAKMP [13] No proposal chosen in message from *xxx* (a.b.c.d)

In many cases, a Session:IPsec message precedes the ISAKMP message. If the Session:IPsec message indicates an error, then the Session message describes the cause and required action. If there is no Session:IPsec error message, refer to the following list of causes and solutions for explanations.

Description: The encryption types proposed by branch office *xxx* do not match the encryption types configured locally.

Action: Check the encryption types on both sides to make sure they match. If necessary, reconfigure the encryption on one system.

Description: The requested authentication method (for example, RSA Digital Signature) is not enabled.

Action: Enable all required authentication types. Make sure the unneeded types are disabled.

Description: One side of the connection is configured to support dynamic routing while the other side is configured for static routing, where branch office is *xxx*.

Action: Configure both sides to use the same routing type.

Description: Both sides are configured to support static routing. However, the local and remote network definitions of the two sides do not match, where branch office is *xxx*.

Action: Configure both sides to have matching local and remote network definitions.

Description: The Perfect Forward Secrecy (PFS) setting of the two sides do not match. Branch office *xxx* does not have PFS enabled, while PFS is required by the local settings.

Action: Make sure the PFS settings on both sides match. Either enable PFS on the remote side, or disable PFS locally.

ISAKMP [13] Error notification (No proposal chosen) received from *xxx* (a.b.c.d)

Description: The proposal made by the local switch has been rejected by a Contivity VPN Client. This usually indicates that the client is using an international version (56-bit) while the switch has stronger encryption enabled.

Action: The encryption methods used by the client and the switch must match. Either provide the user with a Contivity VPN Client version that supports the stronger encryption method used by the switch, or enable 56-bit encryption on the switch.

Description: The proposal made by the local switch has been rejected by a remote branch office switch, or by an IPsec implementation from another vendor.

Action: Check with the administrator of the remote system to determine the cause of the problem. If the remote system is another switch, the cause is noted in that system's log.

ISAKMP [13] Authentication failure in message from *xxx* (a.b.c.d)

In many cases, a Session:IPsec message precedes the ISAKMP message. If the Session:IPsec message indicates an error, the Session message describes the cause and required action. If there is no Session:IPsec error message, refer to the following list of causes and solutions for explanations.

Description: No encryption types are enabled for the account in question.

Action: Enable the desired encryption types.

Description: The requested authentication method (for example, RSA Digital Signature) is not enabled.

Action: Enable all required authentication types. Make sure the unneeded types are disabled.

ISAKMP [13] Error notification (Authentication failure) received from *xxx* (a.b.c.d)

Description: A Contivity VPN Client attempted to connect, but the user supplied the wrong password.

Action: Make sure that the user and the switch have the same password.

Description: A remote branch office switch rejected your switch's attempt to authenticate.

Action: Contact the administrator of the remote system. If the remote system is another Contivity VPN Switch, the cause is noted in that system's log.

No response from client - logging out

Description: Your switch has lost network connectivity with the remote side.

Action: Verify the network connectivity between your switch and the remote side.

Description: A remote branch office using pre-shared key authentication is using a key that is different from what is configured on the local switch. Because the two sides are using a different encryption key, your switch cannot decrypt the encrypted messages from the other side, and therefore drops the messages.

Action: Make sure that both systems are using the same pre-shared key.

ISAKMP [13] *xxx* (a.b.c.d) has exceeded idle timeout - logging out

Description: The remote system has been idle (meaning that no traffic has been sent) for the amount of time configured in the Idle Timeout parameter (Profiles→Groups→Connectivity).

Action: If the Idle Timeout value is too low, increase it. To disable idle timeouts entirely, set the Idle Timeout value to 00:00:00.

ISAKMP [13] Invalid ID information in message from *xxx* (a.b.c.d)

Description: One side of the connection is configured to support dynamic routing while the other side is configured for static routing. Branch office is *xxx*.

Action: Configure both sides to use the same routing type.

Description: Both sides are configured to support static routing, however the local and remote network definitions of the two sides do not match. Branch office is *xxx*.

Action: Configure both sides to have matching local and remote network definitions.

ISAKMP [13] Error notification (Invalid ID information) received from *xxx* (a.b.c.d)

Description: One side of the connection is configured to support dynamic routing while the other side is configured for static routing. Branch office is *xxx*.

Action: Configure both sides to use the same routing type.

Description: Both sides are configured to support static routing. However, the local and remote network definitions of the two sides do not match. Branch office is *xxx*.

Action: Configure both sides to have matching local and remote network definitions.

Branch office messages

Couldn't install route for *remxxx@xxx*

Description: The switch was unable to install the route for the remote network (indicated by *remxxx@xxx)*. This might result when the route collides with an existing static route.

Action: Remove the existing static route or change the route for the remote network to be a subset or superset of the static route.

SSL messages

Checking chain: invalid parent cert, *xxx*

Description: The given certificate in the chain is not valid. This might indicate that the certificate installed at the external LDAP server has expired or is invalid in some other way.

Action: Verify that the certificate is valid or use a certificate that you know is valid.

Checking chain: invalid child cert, *xxx*

Description: The given certificate in the chain is not valid. This might indicate that the certificate installed at the external LDAP server has expired or is invalid in some other way.

Action: Verify that the certificate is valid or use a certificate that you know is valid.

Child cert *[xxx]* not valid signature by *[xxx]* - *xxx*

Description: The given certificate in the chain is not properly signed. This error can indicate that the certificate was incorrectly installed at the external LDAP server.

Action: Reinstall the certificate at the external LDAP server.

Invalid root cert, *xxx*

Description: One of the root certificates passed to the switch during SSL negotiations was invalid.

Action: Configure the remote side to pass a valid chain of certificates to the switch.

No matching trusted CA certs

Description: None of the certificates in the chain are trusted CA certificates. This message can result if the CA certificate has not been installed or has not been marked as trusted on the switch.

Action: Make sure the CA certificate has been installed and that the certificate is marked as trusted on your switch.

Database messages

Configuration file: *xxx* does not exist

Description: The slapd.cnf file does not exist on the disk, therefore the internal LDAP server could not start. This error can occur if the switch disk has been modified.

Action: Reinstall the switch software.

Failed to start

Description: The internal LDAP server did not start. This can be can caused by a missing configuration file.

Action: Reinstall the switch software.

Index file for attribute *xxx* from file *xxx* could not be created

Description: The given attribute index file for the internal LDAP server could not be created. This might indicate that the switch disk is full or that the database index files are corrupt.

Action: Restore the switch software from an FTP backup or reimport the database from the LDIF file.

LDIF file: *xxx* could not back up

Description: The internal LDAP server database could not be backed up to the specified LDIF file. This can result if the name of the LDIF file is not in 8.3 format.

Action: Make sure the backup file has an 8.3 file name.

LDIF file: could not restore *xxx*

Description: The internal LDAP server database could not be restored from the specified LDIF file. This might indicate that the LDIF file does not exist.

Action: Choose an LDIF file that currently resides on the switch disk.

Security messages

Account: *xxx[xxx]* uid *xxx* not found in account

Description: A UID of the remote entity was not found in the account used to initiate a branch office connection (the UID entry in the message is a UID for PPTP or L2TP, and a remote gateway address for IPsec). This error can result if the credentials given by the remote side of the branch office connection do not match the local configuration.

Action: Make sure the Remote Identity information of the IPsec Authentication Certificates section (Profiles→Branch Office→Edit Connection) is configured properly.

AuthServer: ldap inconsistent; no server type in entry *xxx*

Description: An LDAP entry for an authentication server does not contain a server type. This can indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

CaAuthServer: failed remove - *xxx*

Description: An LDAP entry for a CA authentication server was not fully created and then could not be removed. This can result if the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

CaAuthServerCollection: authenticate *xxx* cert *[xxx]* invalid signature by *[xxx]* - *xxx*

Description: The certificate passed in with the authentication request does not have a valid signature, based on the CA certificate configured on the switch. This can indicate either an incorrect certificate at the remote side (either a client or branch office), or an incorrect CA certificate was installed on the switch.

Action: Make sure that both sides have the correct certificates installed.

CaAuthServerCollection: authenticate *xxx[xxx]*:*xxx* bad certificate - *xxx*

Description: The certificate passed in with the authentication request is not a valid X.509 certificate. This error can result if the certificate configured either at the client or the other side of the Branch Office is incorrect.

Action: Install the correct certificates.

Security: store new system IP address *xxx* failed - *xxx*

Description: The system IP address could not be stored in the switch configuration LDAP entry. Possible cause: the LDAP server is not accessible.

Action: Start the LDAP server or change the external LDAP server configuration to make it accessible.

Security: store new system name *xxx* failed - *xxx*

Description: The system name could not be stored in the switch configuration LDAP entry. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Security: store new system subnet mask *xxx* failed - *xxx*

Description: The system subnet mask could not be stored in the switch configuration LDAP entry. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Entry is referenced *[xxx]* - *xxx*

Description: The LDAP entry is being referenced by another LDAP entry (for example, a filter set being referenced by a User Group or Branch Office Connection).

Action: Remove all references to the LDAP entry in question, then delete the entry.

Error copying entry *[xxx]* to *[xxx]* - *xxx*

Description: An error occurred while copying an LDAP entry.

Action: Delete the new copy that caused the error and retry the rename operation.

Error copying subentries of *[xxx]* to *[xxx]* - *xxx*

Description: An error occurred while copying a set of LDAP entries. This can be caused by an unreachable LDAP server.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Error copying tree *[xxx]* to *[xxx]* - *xxx*

Description: An error occurred while copying a tree of LDAP entries. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Error deleting entry *[xxx]* - *xxx*

Description: An error occurred while deleting an LDAP entry. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

Error deleting tree *[xxx]* - *xxx*

Description: An error occurred while deleting a tree of LDAP entries. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

LocalAuthServer: failed remove - *xxx*

Description: An LDAP entry for an LDAP authentication server was not fully created and then could not be removed. This might indicate that the LDAP server is not accessible.

Action: Start the LDAP server, or change the external LDAP server configuration to make it accessible.

SchemaCls: Database schema not available

Description: The external LDAP server does not support a schema entry so it is not possible to update its schema over the network. This error can occur if the external LDAP server is not a Netscape server and does not support the cn=schema entry.

Action: Update the external LDAP server schema manually, then reconnect to it.

xxx xxx being referenced by *xxx*

Description: The LDAP entry is being referenced by another LDAP entry (for example, a filter set being referenced by a User Group or Branch Office Connection).

Action: Remove all references to the LDAP entry in question, then delete the entry.

Session: *xxx* uid invalid - authentication failed

Description: The given IPsec hashed UID was not found in the LDAP database. This can be caused if the UID typed in at the client was invalid or the account no longer exists.

Action: Make sure the correct UID was typed at the client and make sure the account is valid.

Session: *xxx[xxx]* invalid uid - authentication failed

Description: The given group UID was not found in the LDAP database, or the UID was found under a group account and this was not a group login. This error might result if the UID was mistyped at the client or the account no longer exists.

Action: Make sure the correct UID was typed at the client and make sure the account is valid.

Session: *xxx[xxx]* session rejected - system is initializing

Description: The switch had to reject the incoming request since it is still initializing.

Action: Wait a short time to make sure that the switch has initialized, then try again.

Session: *xxx[xxx]* session rejected - system is shutting down

Description: The switch rejected the incoming request since it is shutting down.

Action: Wait for the switch to restart, then try again.

Session: *xxx[xxx]:xxx xxx* auth method not allowed

Description: The authentication method of the incoming request is not allowed in the group that the session is bound to. The session is bound to a group by one of the following:

- The group that the user's account is in (in LDAP)
- The RADIUS default group
- The RADIUS class attribute
- The CA authentication server's default group

Action: Enable the authentication method for the bound group.

Session: *xxx[xxx]:xxx* - authentication failed using all authservers

Description: The incoming request could not be authenticated by any configured authentication servers (LDAP, RADIUS, or CA).

Action: Provide the correct credentials. (For example, create a new user account.)

Session: *xxx[xxx]:xxx* AddLink failed *[xxx]* current links *xxx*

Description: The multilink session could not be created. This can be caused by any of the following:

- New logins are disabled.
- The max sessions on the switch has been reached.
- There is not enough heap on the switch.
- The call admission priority slot is full.
- The call admission priority slot is outside of access hours.
- The max links configured for the group has been reached.

Action: Verify the correct settings for each of the possible causes.

Session: *xxx[xxx]:xxx* IP address assignment failed

Description: An address could not be assigned to the session. This can result if the static address for the session is in use or if the address pool is exhausted.

Action: Expand the number of addresses in the pool, or change the static address on the account.

Session: *xxx[xxx]:xxx* L2TP host *[xxx]* account misconfigured

Description: The L2TP Access Concentrator on the Branch Office Connection does not exist or does not have a LAC or switch UID.

Action: Recreate the L2TP Access Concentrator entry and make sure this entry is linked to the Branch Office Connection.

Session: *xxx[xxx]:xxx* account has max links *(xxx)*

Description: The maximum number of multilink sessions has been reached.

Action: Increase the maximum number of allowed PPP links on the Profiles→Groups→Edit→Connectivity screen.

Session: *xxx[xxx]:xxx* account has max sessions *(xxx)*

Description: The maximum number of sessions for the given account has been reached.

Action: Increase the number of logins on the Profiles→Groups→Edit→Connectivity screen.

Session: *xxx[xxx]:xxx* account is disabled

Description: The account is not currently enabled. This error can occur if the Branch Office Connection request is a different tunnel type than the local switch.

Action: Make sure that both sides are configured to support the same tunnel type.

Session: *xxx[xxx]:xxx* account not allowed now

Description: The session request is outside the permitted hours of access.

Action: Change the Access Hours setting assigned to the group on the Profiles→Groups→Edit→Connectivity screen.

Session: *xxx[xxx]:xxx* authentication failed using *xxx*

Description: The credentials for the session could not be validated by any of the authentication servers.

Action:

1. Make sure you are using the correct credentials.
2. Expand the capability of the RADIUS authentication server to handle the authentication method.
3. Add a new account with the given credentials.

Session: *xxx[xxx]:xxx* client assigned address *[xxx]* already in use

Description: The address given by the tunnel client is currently is use. This might indicate that the address is either being used in a static or dynamic route, or that the address is assigned to an active tunnel.

Action: Configure the client to use a different address.

Session: *xxx[xxx]:xxx* connect Qos level *xxx* full

Description: The switch does not have any more slots for the session's call admission priority. This can indicate that the configured Call Admission Priority for the group that the request is assigned to is too low.

Action: Increase the Call Admission Priority on the Profiles→Groups→Edit→Connectivity screen.

Session: *xxx[xxx]:xxx* invalid password - master admin authentication failed

Description: The primary administrator password was invalid. This can result from using the wrong password or from making a mistake while typing the password.

Action: Make sure you are using the correct password, and make sure you typed it correctly.

Session: *xxx[xxx]:xxx* login rejected - new logins disabled

Description: New logins are currently disabled. This can result if the switch was shut down with one of the following settings enabled on the Admin→Shutdown screen:

- The **Disable new logins** checkbox is selected, or
- The **Disable logins after restart** checkbox is selected.

Action: Enable new logins by deselecting the disable login settings on the Admin→Shutdown screen and then restart the switch.

Session: *xxx[xxx]:xxx* no memory free: *xxx* threshold: *xxx*

Description: There is not enough heap memory available to establish the session. This can result if the switch has consumed a large amount of memory while processing management requests.

Action: Increase the amount of physical memory on the switch, or wait until the management requests are complete.

Session: *xxx[xxx]:xxx* only one session/static address allowed

Description: An address can be used by only one session. This error occurs if the switch receives a second login to an account that has a static address configured.

Action: Change the account to use dynamic addresses from either a static address pool or DHCP.

Session: *xxx[xxx]:xxx* pool address *[xxx]* already in use

Description: The returned static pool address is currently is use. This error can occur if another tunnel is using this address via a static address configuration or another address pool. The error also occurs if a static host route using this address has been added.

Action: No action is necessary. The switch attempts to allocate a different address.

Session: *xxx[xxx]:xxx* session directed to use server *xxx*

Description: This is an informational message indicating that load balancing is enabled and the session is being redirected to another switch. This occurs when the switch is either more heavily CPU-loaded or session-loaded than the other switch.

Action: No action is necessary.

Session: *xxx[xxx]:xxx* static address *[xxx]* already in use

Description: The static address assigned to the account is in use by another tunnel or via a static host route.

Action: Change the static address.

Session: *xxx[xxx]:xxx* system has max sessions *(xxx)*

Description: The switch has reached its maximum number of sessions. This occurs when the switch reaches the maximum number of tunnels that can be configured.

Action: Use load balancing with another switch (if you are using IPsec clients), or upgrade the switch to the next higher model.

RADIUS accounting messages

RADIUS: Cannot send accounting request to <*server-name*>, possibly due to DNS translation failure

Description: This message indicates a connection failure. While sending a request, an error occurred due to a socket creation problem. This usually indicates a DNS resolution problem.

Action: Verify the following:

- DNS host name is correct.
- DNS server is configured properly.
- DNS server is available.

RADIUS: no reply from server <*server-name*>(<*port number*>)

Description: This message indicates a connection failure. The connection timed out while waiting for a response.

Action: Verify the following:

- RADIUS server's IP address and port number are correct.
- RADIUS server is available.
- Shared secret is correct.

RADIUS: <*server-name*> server timed out

Description: This message indicates a connection failure. The connection timed out while waiting for a response.

Action: Verify the following:

- RADIUS server's IP address and port number are correct.
- RADIUS server is available.
- Shared secret is correct.

RADIUS: network socket failure with <*server-name*>, recvfrom error: <*error*>

Description: This message indicates a connection failure. An error occurred while receiving the response.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

RADIUS: <*server-name*> server failed

Description: This message indicates a connection failure. An error occurred while receiving the response.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Indicated packet length too large

Description: This message indicates that an invalid response was received. The length of the response packet is not equal to the number of bytes received.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

RADIUS: failure sending <*user-name*> accounting record to <*server-name*>

Description: This message indicates that an invalid response was received. The length of the response packet is not equal to the number of bytes received.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Non-matching ID in server response

Description: This message indicates that an invalid response was received. The Transaction ID in the response packet is not the expected value.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Unsupported response type (<*number*>) received from server

Description: This message indicates that an invalid response was received. The response packet type is not one of the expected types: Access-Accept, Access-Reject, or Access-Challenge.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Received bad attribute type from server

Description: This message indicates that an invalid response was received. The RADIUS Attribute value is incorrect.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Response OK

Description: This message indicates that a valid response was received.

Action: No action necessary.

RADIUS: <*user-name*> accounting record sent to <*server-name*> OK

Description: This message indicates that a valid response was received.

Action: No action necessary.

RADIUS authentication messages

RADIUS: Cannot send request to <*server-name*>, possibly due to DNS translation failure

Description: This message indicates a connection failure. While sending a request, an error occurred due to a socket creation problem. This usually indicates a DNS resolution problem.

Action: Verify the following:

- DNS host name is correct.
- DNS server is configured properly.
- DNS server is available.

Login failure due to: Server network connection failure

Description: This message is received by the Contivity VPN Client, and indicates a connection failure. While sending a request, an error occurred due to a socket creation problem. This usually indicates a DNS resolution problem.

Action: Verify the following:

- DNS host name is correct.
- DNS server is configured properly.
- DNS server is available.

RADIUS: no reply from RADIUS server <*server-name*>(<*port number*>)

Description: This message indicates a connection failure. The connection timed out while waiting for a response.

Action: Verify the following:

- RADIUS server's IP address and port number are correct.
- RADIUS server is available.
- Shared secret is correct.

RADIUS: <*server-name*> server timed out authenticating <*user-name*>

Description: This message indicates a connection failure. The connection timed out while waiting for a response.

Action: Verify the following:

- RADIUS server's IP address and port number are correct.
- RADIUS server is available.
- Shared secret is correct.

RADIUS: network socket failure with <*server-name*>, recvfrom error: <*error*>

Description: This message indicates a connection failure. An error occurred while receiving the response.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

RADIUS: <*server-name*> server error while authenticating <*user-name*>

Description: This message indicates a connection failure. An error occurred while receiving the response.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Indicated packet length too large

Description: This message indicates that an invalid response was received. The length of the response packet is not equal to the number of bytes received.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

RADIUS: <*server-name*> sent invalid response packet for <*user-name*>

Description: This message indicates that an invalid response was received. The length of the response packet is not equal to the number of bytes received.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Non-matching id in server response

Description: This message indicates that an invalid response was received. The Transaction ID in the response packet is not the expected value.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Unsupported response type (<*number*>) received from server

Description: This message indicates that an invalid response was received. The response packet type is not one of the expected types: Access-Accept, Access-Reject, or Access-Challenge.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Received bad attribute type from server

Description: This message indicates that an invalid response was received. The RADIUS Attribute value is incorrect.

Action: Retry authentication attempt and verify that RADIUS server packets are properly formed.

Invalid reply digest from server, possible shared secret mismatch

Description: This message indicates that an invalid response was received. The computed authenticator does not match the value in the packet.

Action: Verify that the shared secrets match.

RADIUS: <*server-name*> sent packet with invalid response authenticator for <*user-name*>

Description: This message indicates that an invalid response was received. The computed authenticator does not match the value in the packet.

Action: Verify that the shared secrets match.

RADIUS server returned access challenge

Description: This message indicates that a valid access-challenge response was received.

Action: No action required.

RADIUS: <*server-name*> sent challenge for <*user-name*>

A valid access-challenge response was received.

Action: No action required.

RADIUS access challenge received

Description: This message is received by the Contivity VPN Client. A valid access-challenge response was received.

Action: No action required.

RADIUS server rejected access

Description: This message indicates that a valid access-reject response was received.

Action: No action required.

RADIUS: <*user-name*> access DENIED by server <*server-name*>

Description: This message indicates that a valid access-reject response was received.

Action: No action required.

Response OK

Description: This message indicates that a valid access-accept response was received.

Action: No action required.

RADIUS: <*user-name*> access OK by server <*server-name*>

Description: This message indicates that a valid access-accept response was received.

Action: No action required.

Routing messages

Unable to create *xxx* for OSPF

Description: The switch could not create the necessary components to initialize OSPF. This could happen if the switch runs out of free memory.

Action: Disable and enable OSPF globally in Routing→OSPF screen. If this does not work, disable OSPF, boot the switch and enable OSPF in Routing→OSPF screen.

OSPF Disabled

Description: The administrator disabled OSPF from the Routing→OSPF screen.

Action: No action required.

Closing OSPF-RTM connection

Description: OSPF closed the RTM connection, which occurs if the administrator disables OSPF from Routing→OSPF screen.

Action: No action required.

Ospf_Global.State changed from ENABLED to DISABLED by user 'admin' @ *x.x.x.x*

Description: The administrator disabled OSPF from the Routing→OSPF screen.

Action: No action required.

Opened OSPF-RTM connection

Description: The administrator enabled OSPF from the Routing→OSPF screen and successfully registered with RTM.

Action: No action required.

OSPF Enabled

Description: The administrator enabled OSPF from the Routing→OSPF screen.

Action: No action required.

Ospf_Global.State changed from DISABLED to Enabled by user 'admin' @ *x.x.x.x*

Description: The administrator disabled OSPF from the Routing→OSPF screen.

Action: No action required.

Can not accept x.x.x.x as router id

Description: OSPF can not accept the given router ID in the Routing→OSPF screen.

Action: You must change router ID in the Routing→OSPF screen. Invalid router IDs are 127.0.0.1 and 0.0.0.0.

LoadOspfAreas Failed

Description: OSPF failed to load all areas of information from the config file. This could happen if the config file is damaged.

Action: Delete all OSPF areas, re-create them from the Routing→OSPF screen, and reboot the switch.

LoadOspfIntf Failed

Description: OSPF failed to load information for all interfaces from the config file. This could happen if the config file is damaged.

Action: Delete all OSPF interfaces, re-create them from the Routing→Interface screen, and reboot the switch.

VR *xxx*: Starting *xxx* as Master for *xxx*

Description: Logged when VRRP is starting as a master for an address. The parameters are:

- The VRID of this VR
- The reason for starting, either because it was enabled or the interface went up
- The IP address

Action: No action required.

VR *xxx*: Starting *xxx* as Backup for *xxx*

Description: Logged when starting as a backup for an address. The parameters are:

- The VRID of this VR
- The reason for starting, either because it was enabled or the interface went up
- The IP address

Action: No action required.

VR *xxx*: Starting *xxx* as master delayed Backup for *xxx*

Description: Logged when master delay mode is in effect. The parameters are:

- The VRID of this VR
- The reason for starting, either because it was enabled or the interface came up
- The IP address

Action: No action required.

VR *xxx*: Shutting down *xxx* on *xxx*

Description: Logged when VRRP is stopping. The parameters are:

- The VRID of this VR
- The reason for stopping, either because it was disabled or the circuit went down
- The IP address

Action: No action required.

Unable to get configuration for VR *xxx*

Description: This is an error event. It is logged when VRRP is enabled but the common configuration parameters are missing. These are the items set under the Routing→VRRP screen. The parameter is the IP address that is missing information.

Action: No action required.

RIP *xxx*: RIP Enabled

Description: Logged when RIP is globally enabled.

Action: No action required.

RIP *xxx*: RIP Disabled

Description: Logged when RIP is globally disabled.

Action: No action required.

RIP *xxx*: Can't alloc main node

Description: Logged when there is not enough memory to allocate RIP parameters.

Action: No action required.

RIP *xxx*: Circuit *xxx* created

Description: Logged when the RIP circuit was created. The parameter stands for circuit ID.

Action: No action required.

RIP *xxx*: Circuit *xxx* deleted

Description: Logged when the RIP circuit was deleted. The parameter stands for circuit ID.

Action: No action required.

RIP *xxx*: Unable to register with UDP

Description: Logged when you cannot register with UDP protocol.

Action: No action required.

RIP *xxx*: setsockopt RIP socket *xxx* SO_RCVBUF *xxx* failed

Description: Logged when RIP receive buffers are not large enough. This can happen when a large numbers of RIP neighbors send their RIP updates simultaneously. The first parameter is the socket number and the second parameter is the maximum receive buffer size.

Action: No action required.

RIP *xxx*: bind RIP socket *xxx* failed

Description: Logged when RIP failed to bind the socket.

Action: No action required.

RIP *xxx*: Unable to spawn Dispatcher task *xxx* for RIP

Description: Logged when RIP failed to spawn the main task responsible for receiving RIP packets. The parameter stands for the name of the task.

Action: No action required.

RIP *xxx*: Unable to spawn timer task *xxx* for RIP

Description: Logged when RIP failed to spawn the timer task. The parameter stands for the name of the task.

Action: No action required.

RIP *xxx*: cid xxx mismatched auth password from *xxx*

Description: Logged when RIP authentication failed while receiving RIP packets. The first parameter is the circuit ID on which it was receiving RIP packets and the second parameter is the IP address from which it received RIP packets.

Action: No action required.

Hardware messages

The switch software provides informational messages when cards are removed and replaced. When you exchange two cards with each other, the switch considers this two simultaneous replacements.

Interface [*nnn*] not present, deleting from config

Description: This indicates that the configuration file contains an interface [*nnn*] entry, but there is no card in the slot. The interface [*nnn*] entry is deleted from the configuration.

Action: No action required.

Interface [*nnn*] replaced, resetting config

Description: This indicates the card type specified in the configuration file does not match the card type currently in the slot. The configuration information is reset to defaults then initialized with the current hardware.

Action: No action required.

Interface [*nnn*] replaced, deleting from config

Description: This indicates the card type specified in the configuration file does not match the card currently in the slot. The interface is deleted from the configuration. This applies when the replaced card has more ports than the current card.

Action: No action required.

HWAccel [*nnn*] not present, deleting from config

Description: This indicates the configuration file contains a HWAccel [*nnn*] entry, but there is no hardware accelerator in the slot. The HWAccel [*nnn*] entry is deleted from the configuration.

Action: No action required.

Appendix D
Configuring for interoperability

This chapter explains the requirements and procedures for setting up different vendor hardware or software to interoperate with the Contivity VPN Switch. These instructions enable you to establish encrypted tunnels to and from the switch with the noted vendors. These requirements and procedures are subject to change based on hardware and software changes by the vendors.

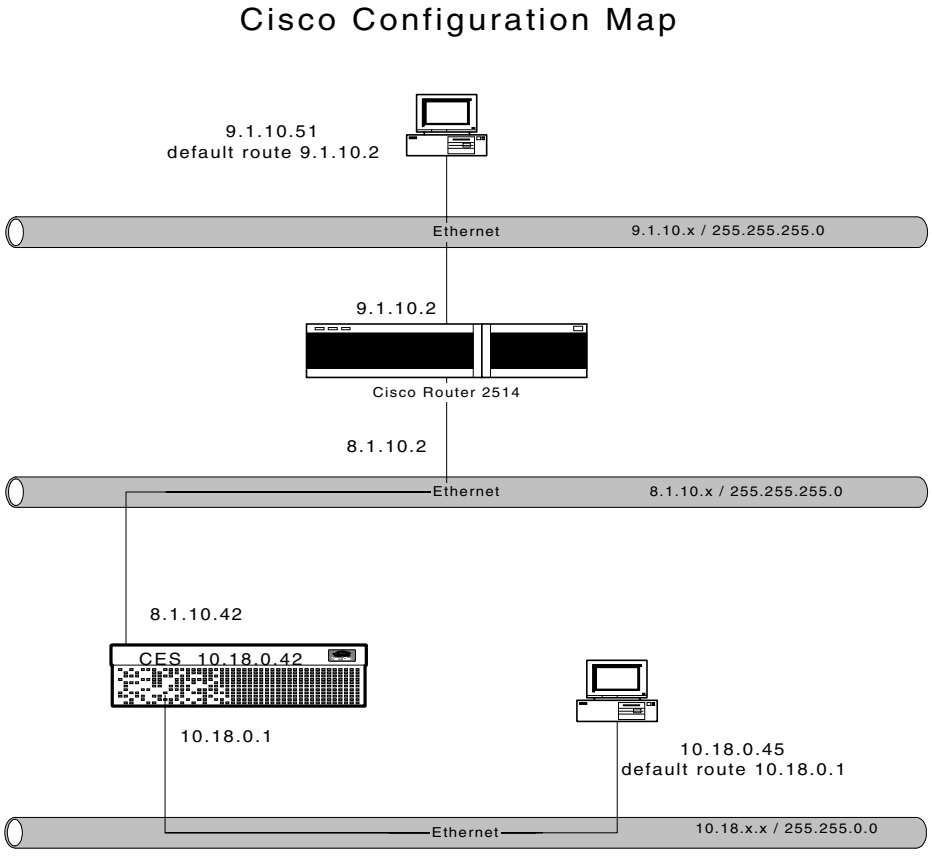
Procedures are available for the following products:

- Cisco 2514 Router*, Version 11.3
- Check Point, Version 4.0 with VPN-1 and FireWall-1
- Information Resource Engineering, Inc. (IRE), SafeNet/Soft-PK Security Policy Database Editor*, Version 1.0s

Configuring the Cisco 2514 router, version 11.3

To set up the switch to establish encrypted tunnel connections with the Cisco 2514 router, as shown in Figure 73, you should configure the Cisco 2514 with the Show Configuration command.

Figure 73 Switch and Cisco 2514 network topology



The following is a `show config` command:

```
Cisco2514# show config
Using 1088 out of 32762 bytes
version 11.3
no service password-encryption
hostname Cisco2514
enable secret 5 $1$aSJB$Xz/o4I4IqCY.FT2RH372/1
enable password password
!
crypto isakmp policy 1
 hash md5
 authentication pre-share
 lifetime 3000
crypto isakmp key test address 8.1.10.42
!
crypto ipsec transform-set esp1 esp-des esp-md5-hmac
!
crypto map bay 11 ipsec-isakmp
 set peer 8.1.10.42
 set session-key lifetime seconds 3000
 set transform-set esp1
 match address 132
!
!
interface Ethernet0
 ip address 9.1.10.2 255.255.255.0
 no mop enabled
!
interface Ethernet1
 ip address 8.1.10.2 255.255.255.0
 no mop enabled
 crypto map bay
!
interface Serial0
 no ip address
 no ip mroute-cache
 shutdown
!
interface Serial1
 no ip address
 shutdown
!
ip classless
ip route 10.18.0.45 255.255.255.255 8.1.10.42
access-list 132 permit ip host 9.1.10.51 host 10.18.0.45
access-list 132 permit ip host 10.18.0.45 host 9.1.10.51
dialer-list 1 protocol ip permit
```

```
dialer-list 1 protocol ipx permit
snmp-server community public RO
line con 0
line aux 0
line vty 0 4
password terminal
login
end
```

Configuring the switch for Cisco interoperability

To configure the switch for Cisco* interoperability, go to the Profiles→Networks screen and click on Edit.

Create any local accessible networks that you want available.

1. Enter the IP address for the new subnet; for example, 10.18.0.45.
2. Enter the subnet mask for the new network.
3. Click on Add. The Networks Edit screen appears and shows the newly created subnet in the Current Subnets list for the named network.
4. Add each local subnet for which you want tunneled connections coming to or going from the switch to a Network profile.
5. Verify that your settings are synchronized with the Cisco router on the Profiles→Branch Office: Edit GROUP screen.

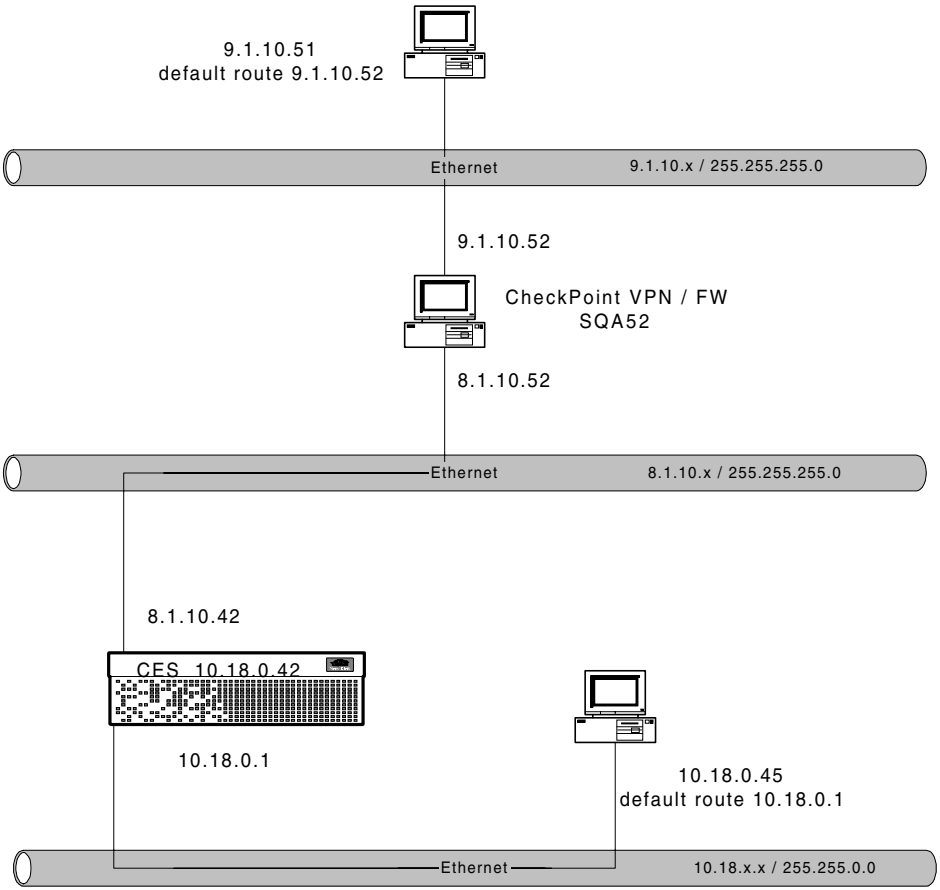
   For Cisco, turn off Vendor ID and Perfect Forward Secrecy (PFS). Go to the Profiles→Groups→IPSec: Configure screen.

6. Create and configure the IPSec branch office connection on the switch, using the network profile you just created for the local accessible network.
7. On the Profiles→Branch Office screen, you must enable IPSec Authentication: Text Pre-Shared Key.

Configuring the Check Point VPN-1 and FireWall-1, Version 4.0

To set up the switch to establish encrypted tunnel connections with the Check Point FireWall-1* (Version 4.0) as shown in Figure 74, you must create certain objects on the firewall.

Figure 74 Switch and Check Point FireWall-1 network



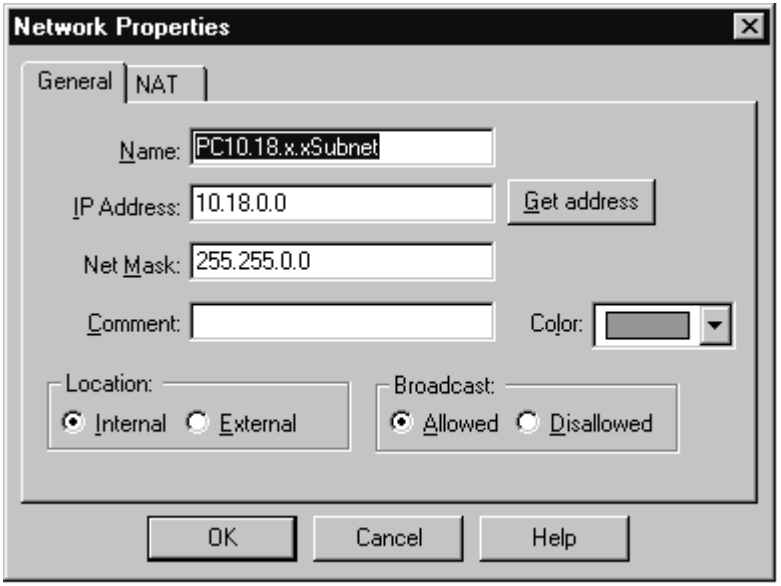
These objects are required for the two devices to interoperate according to the preceding network configuration example:

- Manage→Network Objects→New→Network - Subnets on Check Point.
- Manage→Network Objects→New→Network - Subnet on the switch.
- Manage→Network Objects→ New→ Workstation - End point of Check Point.
- Manage→Network Objects→ New→ Workstation - End point of the switch.
- Create the Encryption rules in the firewall policy.

Establishing an encrypted tunnel on the switch side

To establish an encrypted tunnel to and from a specific network on the switch side (10.18.0.0) of the tunnel:

1 Select Manage→Network Objects→New: Network. The Network Properties: General dialog box opens.



2 Enter this general information; for example:

- Name of the switch's subnet: PC10.18.xx.xSubnet

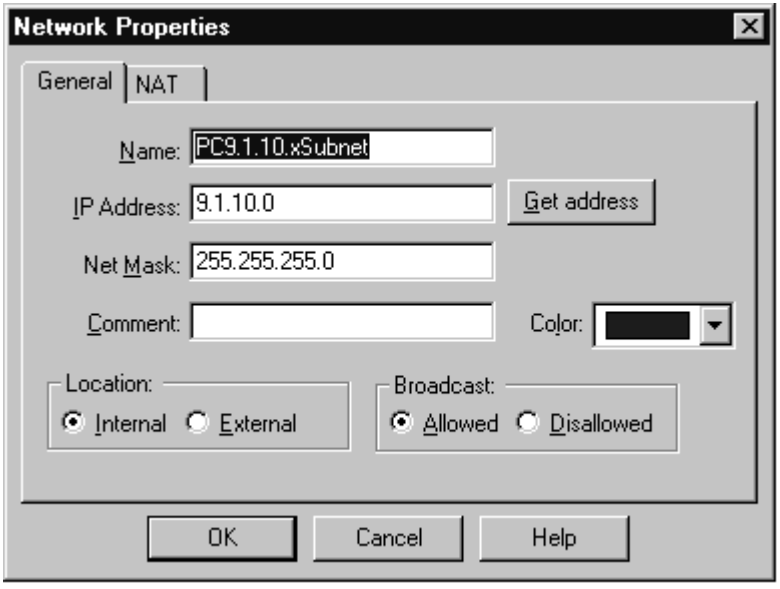
- The switch's private subnet IP address: 10.18.0.0
- The switch's private subnet Mask: 255.255.0.0
- Select location: Internal
- Select broadcast type: Allowed

3 Click on OK. The Network Objects dialog box opens.

Establishing an encrypted tunnel on the firewall side

To establish an encrypted tunnel to and from a specific network on the firewall side (9.1.10.0) of the tunnel:

1 Select Manage→Network Objects→New: Network. The Network Properties: General dialog box opens.



2 Enter the following information; for example:

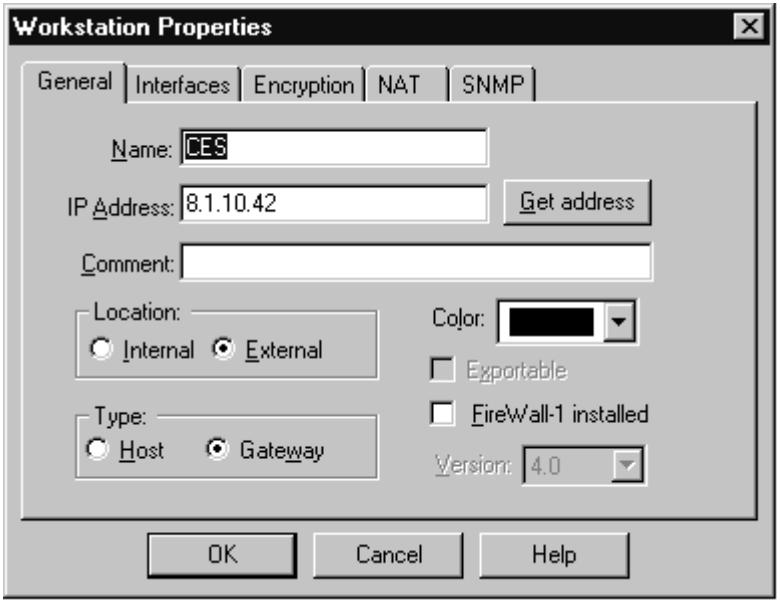
- Name of the firewall's subnet: PC9.1.10.xx.xSubnet
- The firewall's private subnet IP address: 9.1.10.0
- The firewall subnet mask: 255.255.255.0
- Select location: Internal
- Select broadcast type: Allowed

3 Click on OK. The Network Objects dialog box opens.

Establishing a tunnel from the switch to the firewall

To establish an encrypted tunnel to and from the switch to the firewall:

1 Select Manage→Network Objects→New: Workstation. The Workstation Properties dialog box opens.



2 Enter this general information; for example:

- Name of the switch's Endpoint: CES
- The switch's Public IP Address: 8.1.10.42
- Select Location: External
- Type: Gateway

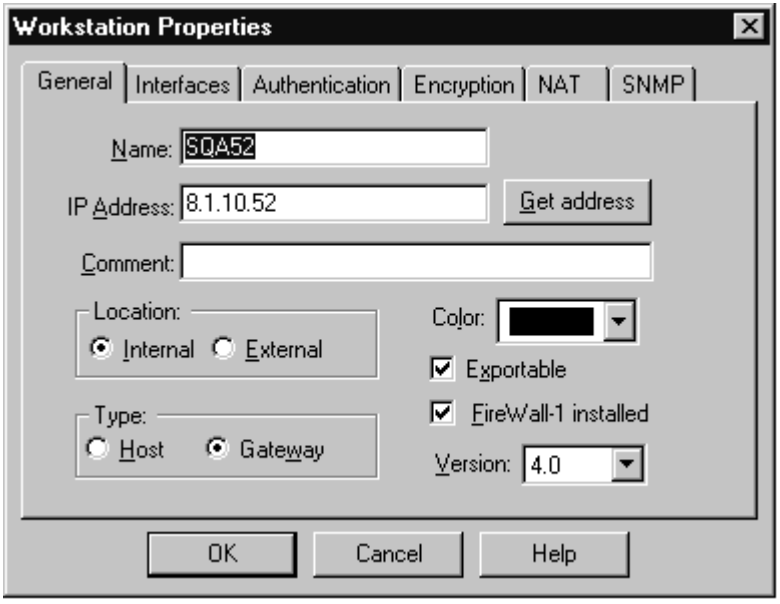
3 Click OK.

The configuration of this workstation object should be done after the firewall's workstation object has been created.

Establishing a tunnel to and from the firewall

To establish an encrypted tunnel to and from the firewall to the switch:

1 Select Manage→Network Objects→New: Workstation. The Workstation Properties dialog box appears.



2 Enter this general information; for example:

- Name of the firewall Endpoint: SQA52

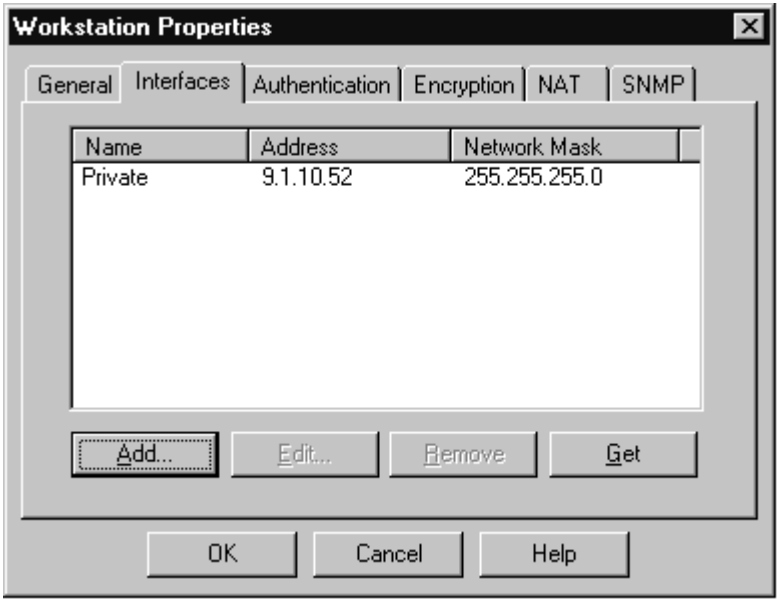
- The firewall's IP Address: 8.1.10.52
- Select Location: Internal
- Select Type: Gateway
- Click on Exportable
- Click on FireWall-1 installed

3 Click OK.

Configuring the firewall's public interface

To configure the firewall's public interface:

1 Click on the Workstation Properties: Interfaces tab. The Workstation Properties: Interfaces dialog box opens.



2 Click on Add.

3 Enter the interfaces information:

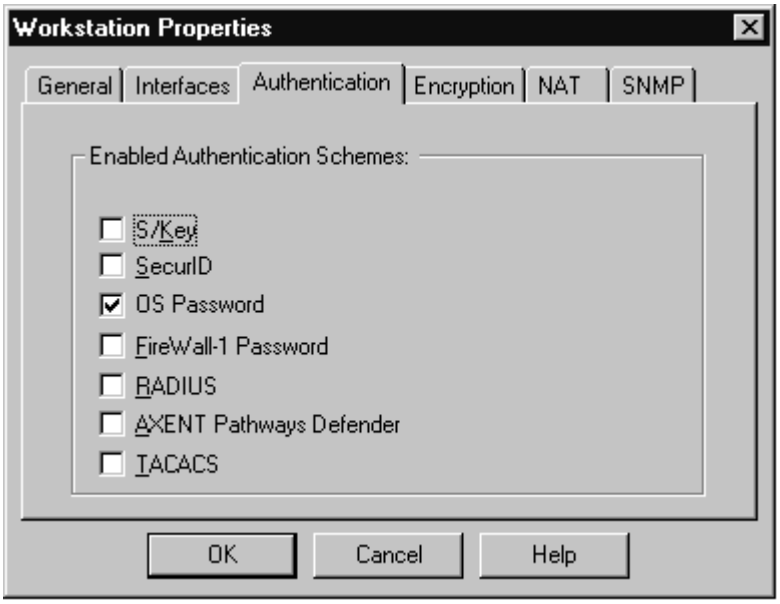
- Name: Private
- Net Address: 9.1.10.52
- Net Mask: 255255.255.0

4 Click on OK. The Workstation Properties: Interfaces dialog box opens with the newly created interface listed.

Configuring the switch authentication schemes

To configure the switch's authentication schemes:

1 Click on the Workstation Properties: Authentication tab.



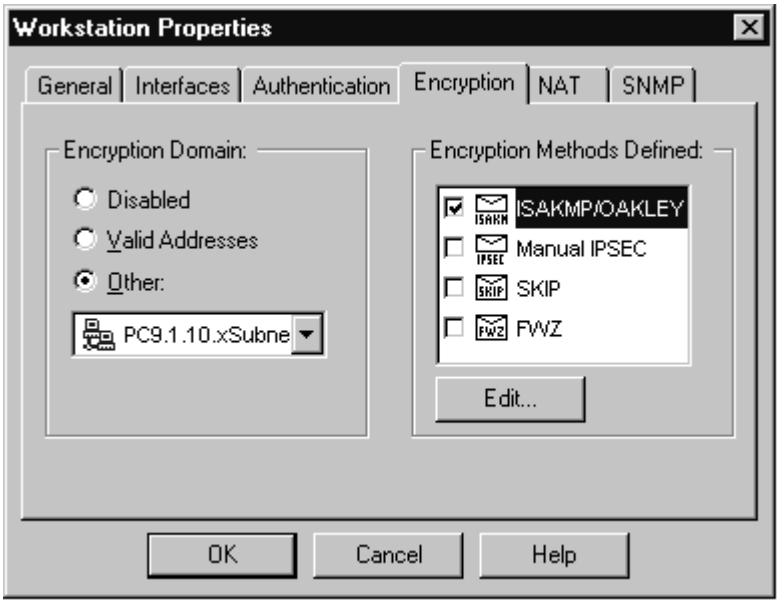
2 Click on OS Password

3 Click on OK.

Configuring the firewall's encryption methods

To configure the firewall's encryption methods:

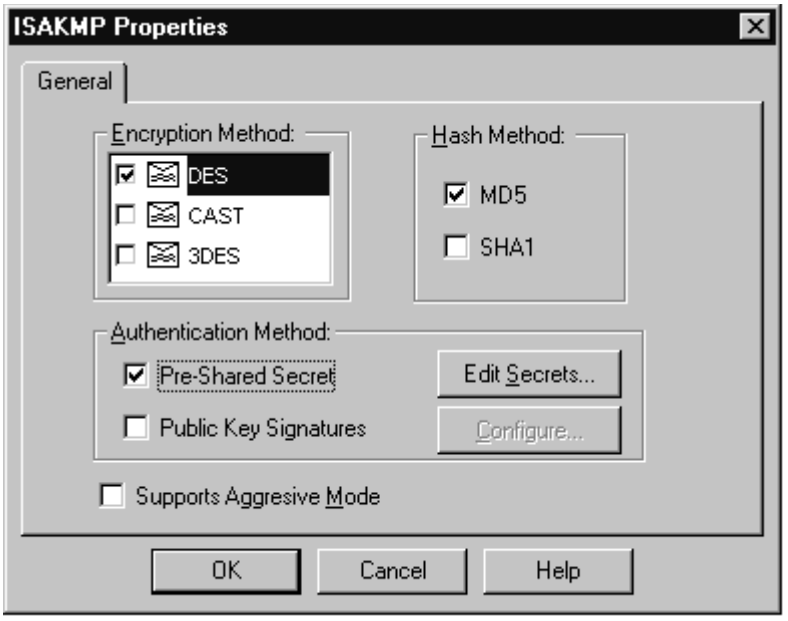
1 Click on the Workstation Properties: Encryption tab. This is where you assign the firewall responsibility for the 9.1.10.x subnet.



2 Fill out this tab as follows:

- In the Encryption Domain section, select Other and choose PC9.1.10.xSubnetwork from the drop down list.
- In the Encryption Methods Defined section, select ISAKMP/OAKLEY.

3 Click on Edit. The ISAKMP Properties: General dialog box opens.



4 On the ISAKMP Properties→General tab, select:

- Encryption Method: DES
- Authentication Method: Pre-Shared Secret
- Hash Method: MD5

5 Click on Edit Secrets. The Shared Secrets dialog box opens.



6 In the Shared Secrets List, double-click on the appropriate shared secret; for example, CES (the switch workstation object). This creates a shared secret between the switch and the firewall.

7 Finish configuring the switch's workstation object.

Configuring the switch's public interface

To configure the switch's public interface:

1 Select Manage→Network Objects→Select CES Workstation Object.

2 Click on the Workstation Properties: Interfaces tab. The Workstation Properties: Interfaces dialog box opens.

3 Click on Add.

4 Enter the following information:

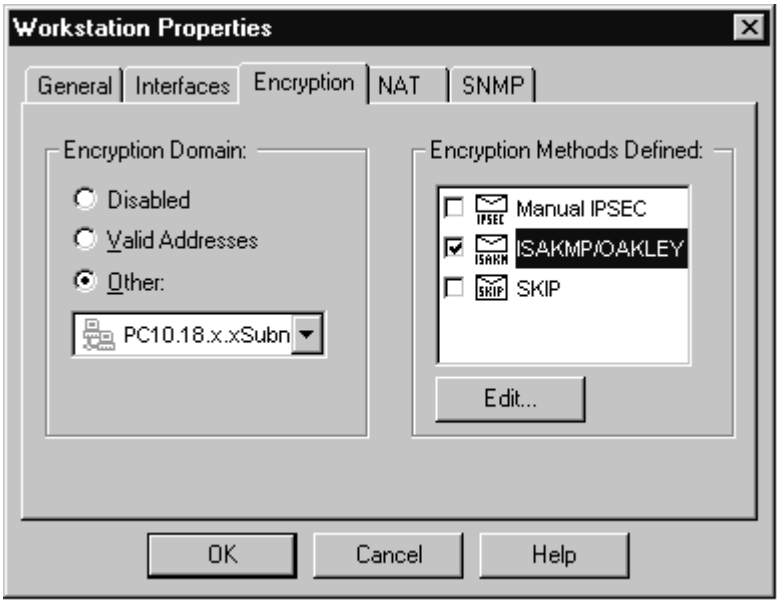
- Name: CES Public
- Net Address: 8.1.10.42
- Net Mask: 255.255.255.0

5 Click on OK. The Workstation Properties: Interfaces dialog box opens with the newly created interface listed.

Configuring the switch's encryption methods

To configure the switch's encryption methods:

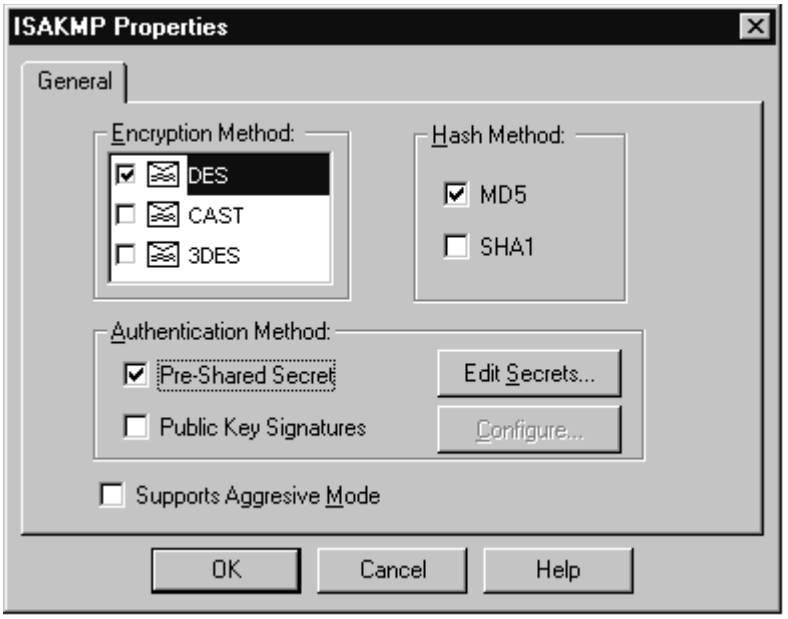
1 Click on the Workstation Properties: Encryption tab. The Workstation Properties: Encryption dialog box opens.



2 Select the following:

- In the Encryption Domain section, select Other and choose PC9.1.10.xSubnetwork from the drop down list.
- In the Encryption Methods Defined section, select ISAKMP/OAKLEY.

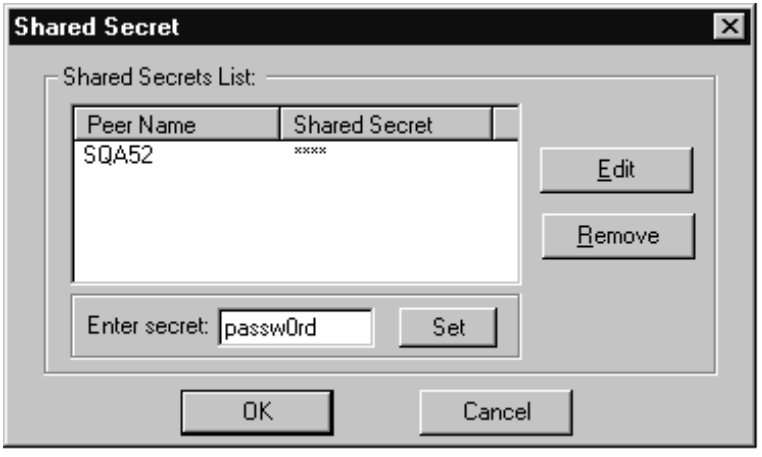
The ISAKMP Properties: General dialog box opens.



3 On the ISAKMP Properties→General screen, select:

- Encryption Method: DES
- Authentication Method: Pre-Shared Secret
- Hash Method: MD5

4 Click on Edit Secrets. The Shared Secrets List dialog box opens.



5 In the Shared Secrets List, double-click on the appropriate shared secret; for example, SQA52 (the FireWall workstation object). This creates a shared secret between the switch and the firewall.

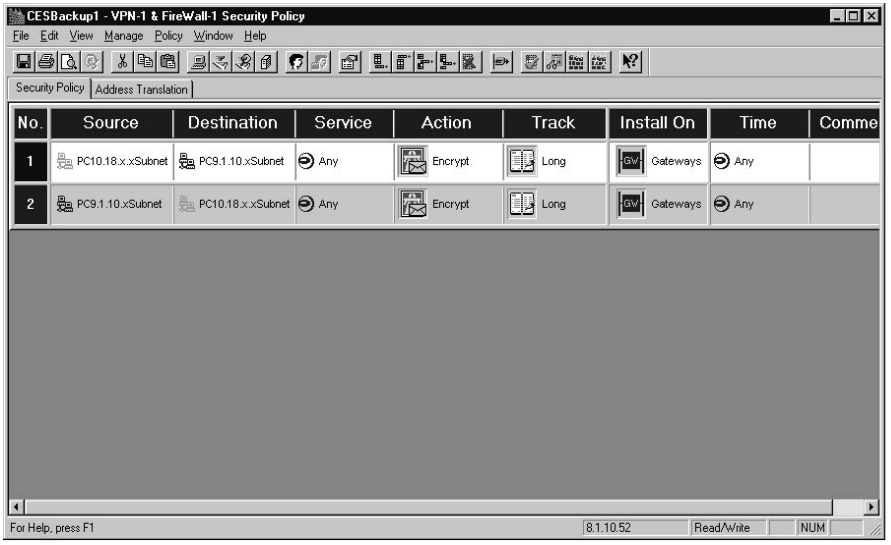
6 Enter the secret (for example, passwOrd), click on set and then on OK. The Shared Secret List now shows asterisks for the shared secret.

7 Click on OK. The ISAKMP Properties dialog box opens.

8 Click on OK. The Workstation Properties dialog box opens.

Adding rules to the configurations

You must add rules to enable the configurations that you have just created. After you add the rules, the VPN-1 & FireWall-1 Security Policy screen should look like this:



1 Use the Edit: Add Rule menu choice. From the Rule editor, right-click on the Action column.

2 Select Source, and right-click on Any.

3 Click on Add.

4 Select the 10.18.x.x subnet (source).

5 Select Destination, and right-click on Any.

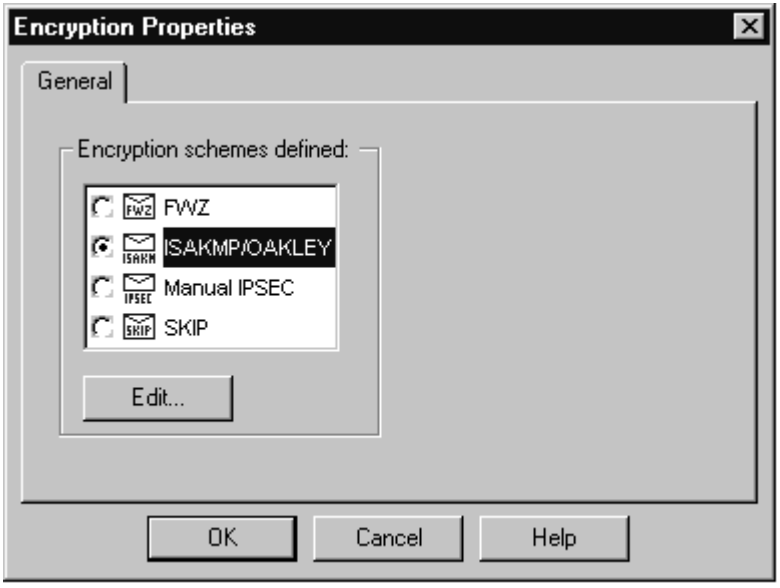
6 Click on Add.

7 Select the 9.1.10.x subnet (destination).

8 Select Edit Properties→General.

9 Select Action.

10 Click on Drop and right-click on Encryption.

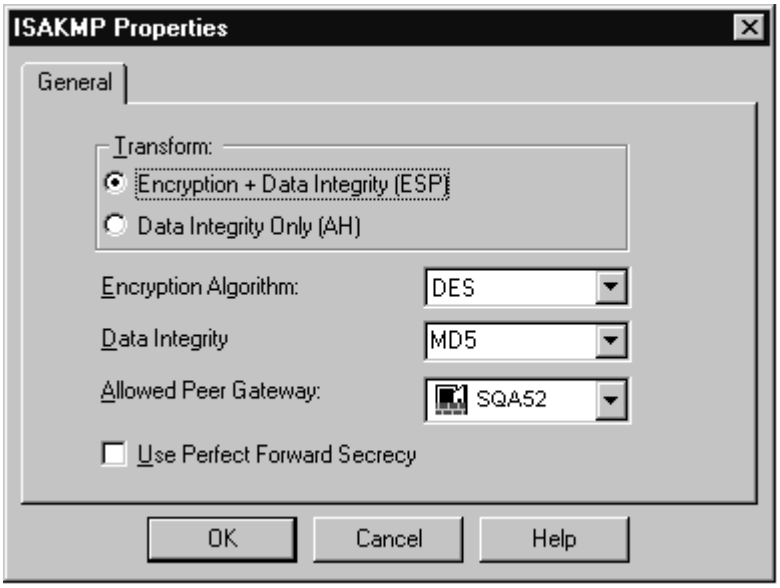


11 Right-click and select Edit Properties.

Editing encryption for rules

Edit the encryption for each rule; for example:

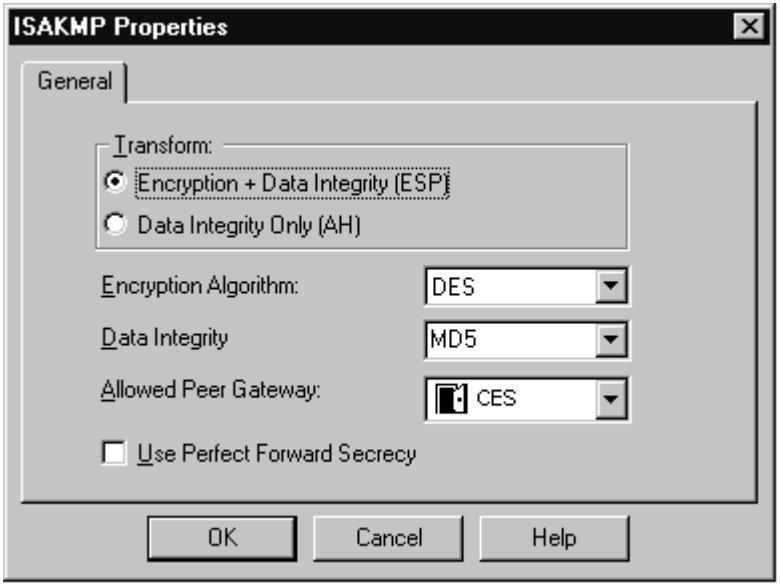
1 If the encryption schemes are defined, click on ISAKMP/OAKLEY. The ISAKMP Properties: General dialog box opens.



2 On the ISAKMP Properties→General screen, click:

- Transform: Encryption + Data Integrity (ESP)
- Encryption Algorithm: DES
- Data Integrity: MD5
- Allowed Peer Gateway: SQA52.

3 If the encryption schemes are defined, select ISAKMP/OAKLEY. The ISAKMP Properties: General dialog box opens.

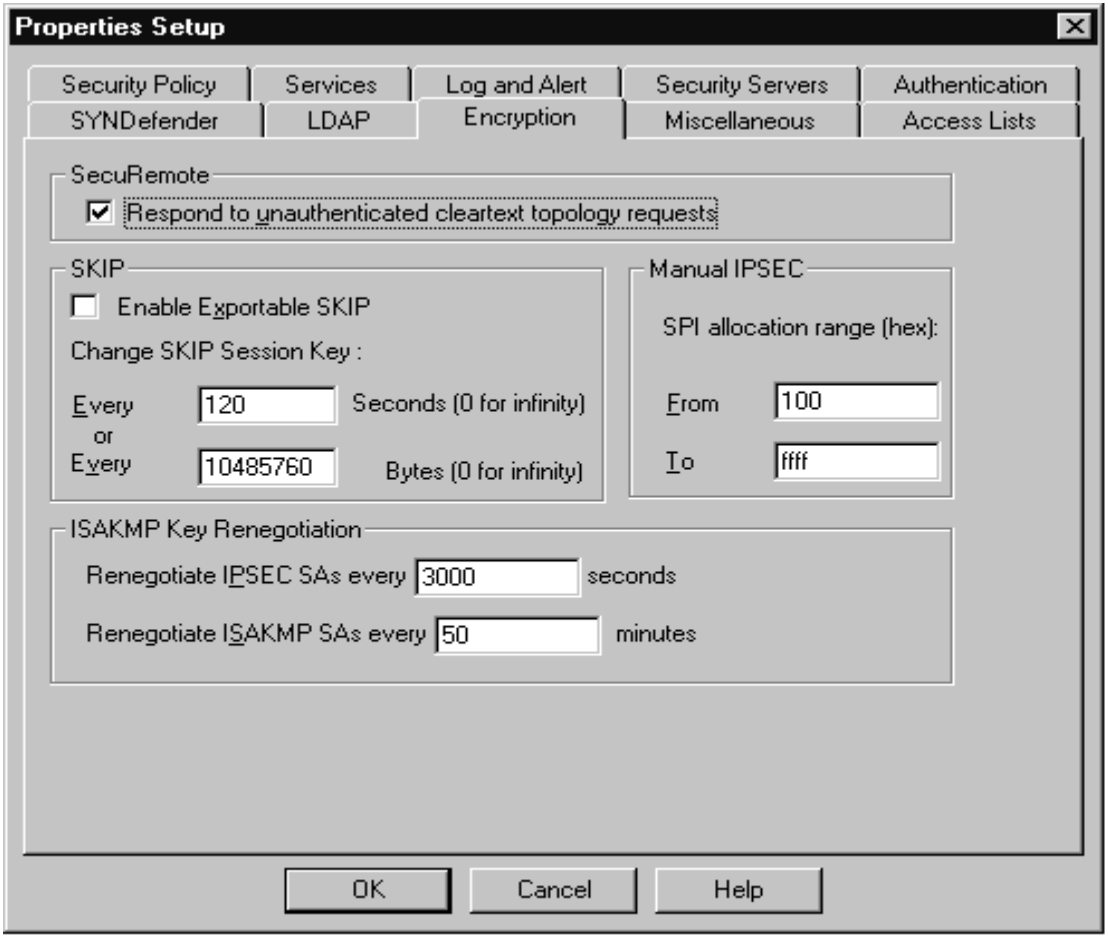


4 On the ISAKMP Properties→General dialog box, click:

- Transform: Encryption + Data Integrity (ESP)
- Encryption Algorithm: DES
- Data Integrity: MD5
- Allowed Peer Gateway: CES

On the Profiles→Branch Office: Connectivity screen, the Forced Logoff attribute is equivalent to the ISAKMP SA time field on the Check Point screen (Policy→ Properties→Encryption).

On the Profiles→Branch Office: Connectivity IPSec screen, the Rekey time-out is equivalent to the IPSEC SA field found on the Check Point screen (Policy→Properties→ Encryption). A sample Check Point Policy→Properties→Encryption screen follows.



Configuring the switch for Check Point interoperability

To configure the switch for Check Point interoperability, go to the Profiles→Networks screen and click on Edit. Create the network object used for local accessible networks; for example:

1. Enter the IP address for the new subnet: 10.18.0.45.
2. Enter the subnet mask for the new network.

3 Click on Add. The Networks Edit dialog box opens and shows the newly created subnet in the Current Subnets list for the named network.

4 Add each local subnet for which you want tunneled connections coming to or going from the switch to a Network profile. Refer to Chapter 5 for additional information.

5 Verify that your settings are synchronized with the firewall (Profiles→Branch Office: Edit GROUP).

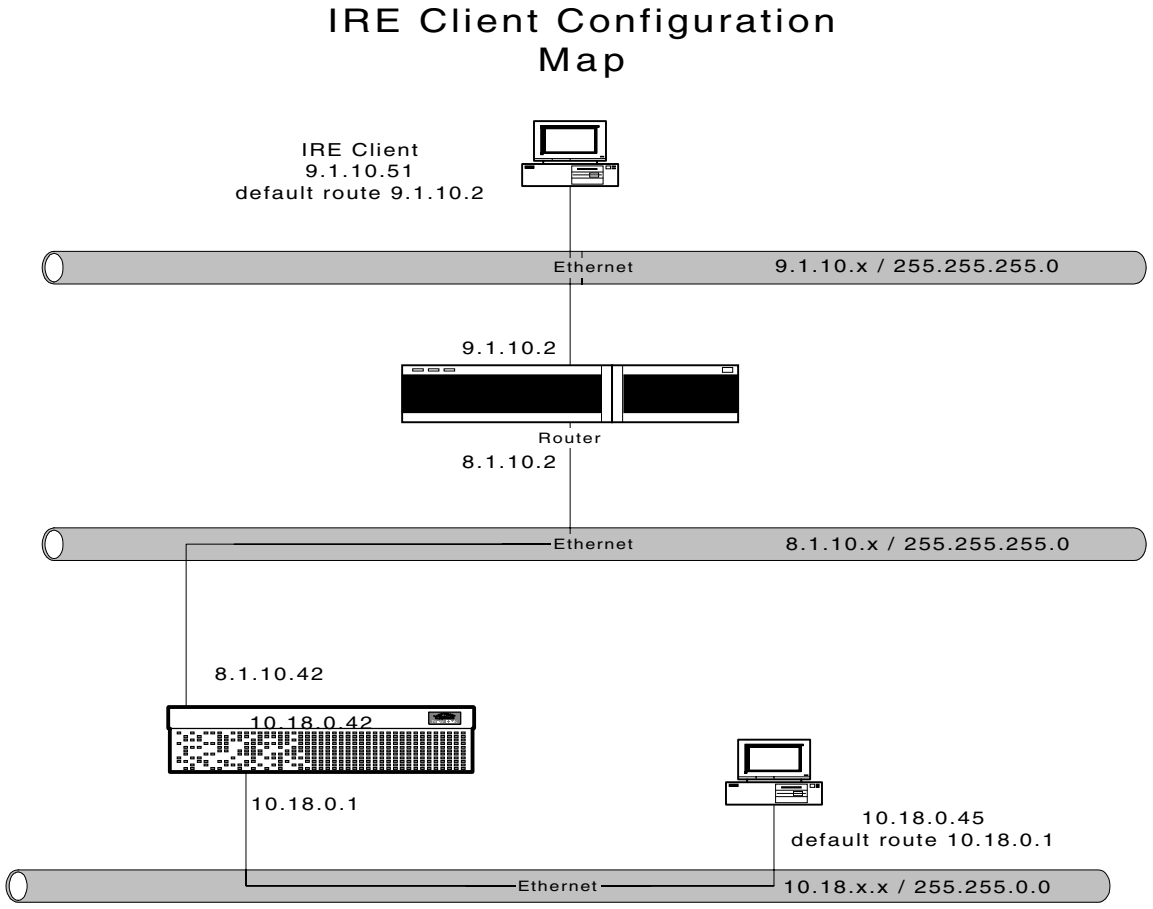
For some vendors, you want to turn off Vendor ID and Perfect Forward Secrecy (PFS). Go to the Profiles→Groups→IPSec: Configure screen to do this.

Create and configure the Branch Office connection on the switch, using the network profile you just created for the local accessible network. On the Profiles→Branch Office screen, you must enable the Authentication: Text Pre-Shared Key option.

Configuring the SafeNet/Soft-PK Security Policy Database Editor, version 1.0s

To set up the switch to establish encrypted tunnel connections with the IRE Soft-PK Security Policy Client as illustrated in Figure 75, you should configure the screens as described on following pages.

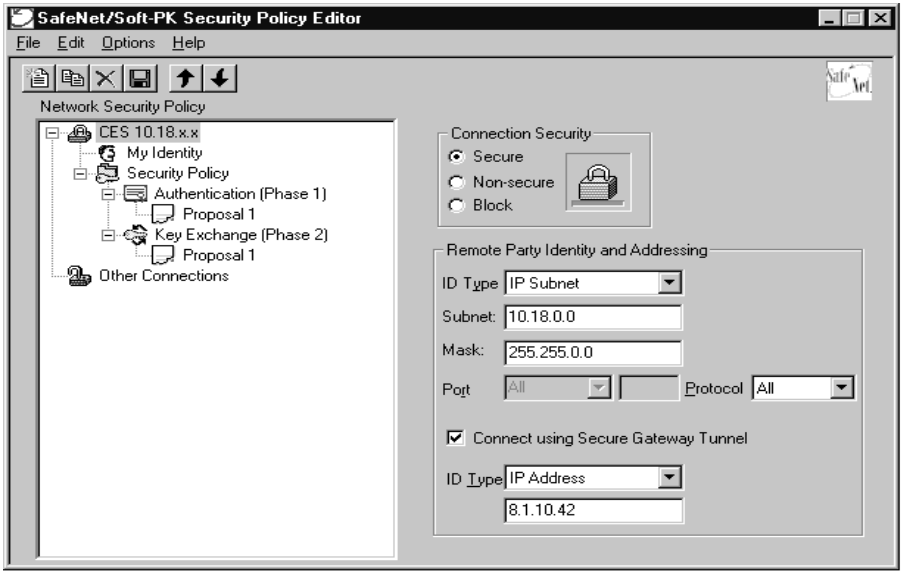
Figure 75 Switch and IRE SafeNet network topology



Connecting to IRE SafeNET/Soft-PK Security Policy Client

To set up the switch to establish encrypted tunnel connections with the IRE SafeNet/Soft-PK Security Policy Client, follow these instructions:

1 Open the SafeNet/Soft-PK Security Policy Client, and click on File: New. The following screen configures the network so that any packets going to the 10.18.0.0 subnet goes through the switch's 8.1.10.42 interface to establish a tunnel.



2 Click on the switch: CES 10.18.x.x.

3 Click on Connection Security: Secure.

4 Under Remote Party Identity and Addressing, select the following:

- ID Type: IP Subnet
- Subnet: 10.18.0.0.
- Mask: 255.255.0.0
- Protocol: All

5 Under Connect using Secure Gateway Tunnel, select the following:

- ID Type: IP Address

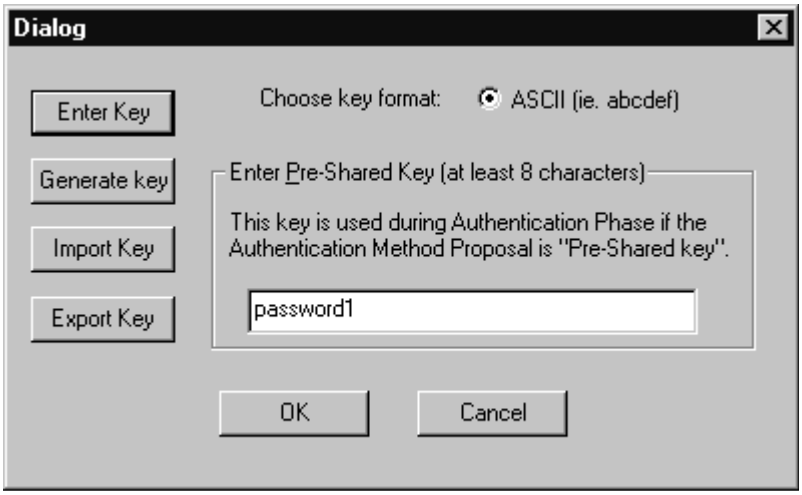
- 8.1.10.42

The SafeNet/Soft PX Security Policy Editor dialog box opens.

6 Click on My Identity to configure the SafeNet client, and select the following:

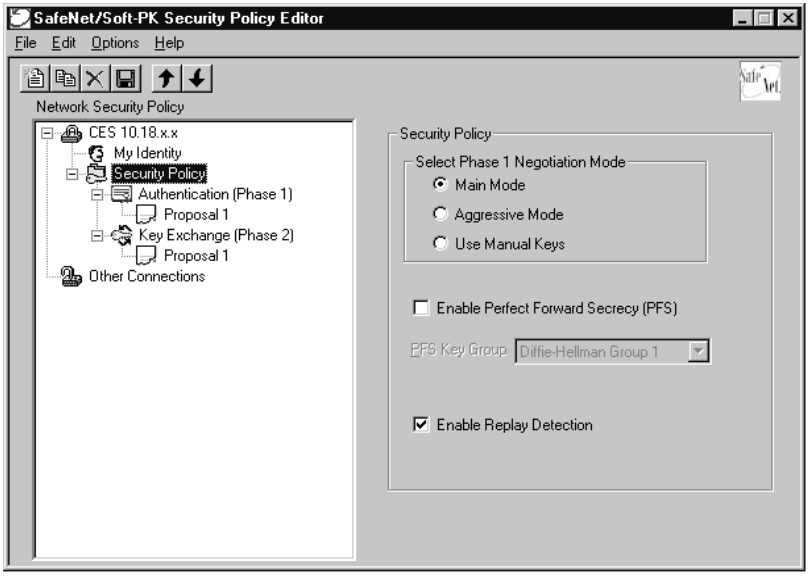
- Select Certificate: None
- ID Type: IP Address
- Port: All

7 Click on Pre-Shared Key. The Pre-Shared Key dialog box appears.



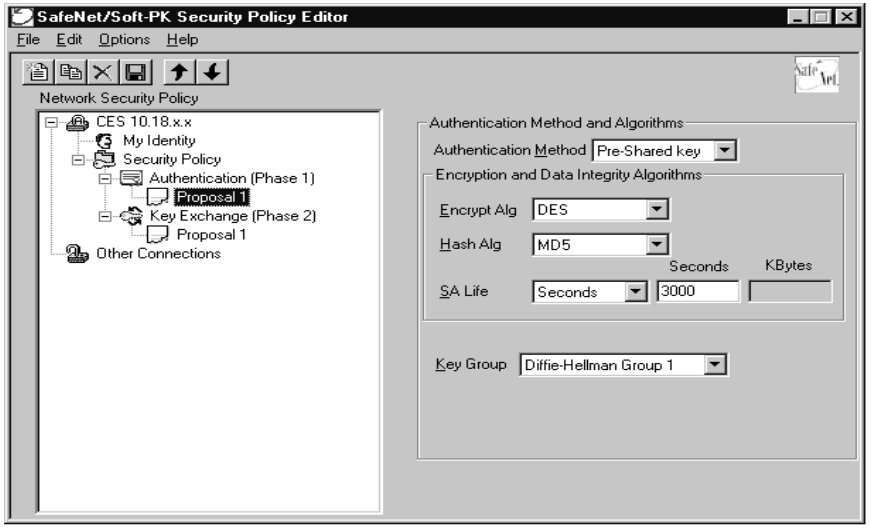
8 In the Pre-Shared Key dialog box, click on Enter Key, then enter the pre-shared key.

9 Click on OK. The SafeNet/Soft-PK Security Policy Editor dialog box appears.



10 Click on Security Policy: Select Phase 1 Negotiation Mode: Main Mode.

11 Click on Enable Replay Detection.



12 On the Authentication (Phase 1), Proposal 1, Authentication screen, enable the following:

- Authentication Method: Pre-Shared key
- Encrypt Alg: DES
- Hash Alg: MD5
- SA Life: Seconds and 3000 (Seconds)
- Key Group: Diffie-Hellman Group 1

13 On the Key Exchange (Phase 2), Proposal 1 screen, enable the following:

- Encapsulation Protocol (ESP)
- Encrypt Alg: DES
- Hash Alg: MD5
- Encapsulation: Tunnel
- SA Life: Seconds and 3000 (Seconds)

Configuring the switch for IRE interoperability

To configure the switch for IRE interoperability, go to the Profiles→Networks screen and click on Edit.

Create the network object used for local accessible networks:

1. In the Networks Edit screen, enter the IP address for the new subnet; for example, 10.18.0.45.
2. Enter the subnet mask for the new network: 255.255.0.0.
3. Click on Add. The Networks Edit screen reappears and shows the newly created subnet in the Current Subnets list for the named network.
4. Add each local subnet for which you want tunneled connections coming to or going from the switch to a network profile.
5. Verify that your settings are synchronized with the SafeNet client (Profiles→Branch Office: Edit GROUP).
6. Create and configure the IPSec Branch Office connection on the switch, using the network profile you just created for the local accessible network. On the Profiles→Branch Office screen, you must enable the IPSec Authentication: Text Pre-Shared Key option.
7. For some vendors, you want to turn off Vendor ID and/or Perfect Forward Secrecy (PFS). Go to the Profiles→Groups→IPSec: Configure screen to do this.

Third-party client installation

The switch supports third-party IPSec clients and includes support for the following:

- Authentication using either pre-shared authentication (using IKE Aggressive mode) or digital signature certificate authentication (using IKE Main mode) into a switch remote access user's IPSec account for third-party IPSec clients.

- Client address assignment to be used within the IPSec tunnel formed as a result of the Quick Mode negotiation. The client's external IP address or a pre-arranged internal IP address can be used as the address that is negotiated during the IKE Quick Mode exchange.
- Split tunneling with third-party IPSec clients, such that if split tunneling is enabled on the switch, then the subnet that the client specifies as the switch's identity within the tunnel during IKE Quick Mode must be listed as one of the split tunnel networks for the Quick Mode proposal to be accepted. If split tunneling is not enabled on the switch, then the switch identity that the client specifies for Quick Mode can be any value that the client chooses.

Depending on the third-party client that you are using, you need to configure either a branch office tunnel or a user tunnel. For example, the switch has been configured and tested with the Linux FreeS/WAN* client. If you are using the FreeS/WAN Linux* client, you must configure your user and the switch as a branch office tunnel. If you are using another client that supports IPSec Aggressive mode, you can configure your switch as a user tunnel.

Considerations for using third-party clients

There are several considerations regarding the use of third-party clients with Contivity:

- Client Dynamic Addressing -- Many third-party clients now support the Aggressive mode method of establishing a security association. The advantage of Aggressive mode for remote user access is that, unlike Main mode, the VPN server does not authenticate the security association based on prior knowledge of the IP address of the user. Therefore, the remote user can be dynamically assigned an address by their ISP.
- Client Address Advertisement -- When connecting to the Contivity VPN client, the switch assigns the client-side inner address of the IPSec tunnel from the enterprise address space. This is the address that devices on the private network send data to in response to requests from the client. The switch captures packets destined for those addresses and sends them through the public interface encapsulated within IPSec, addressed to the ISP-assigned outer address of the client.

In the case of third-party clients, the switch does not have a mechanism to assign the inner address of the client. The inner address of the client tunnel is normally set the same as the ISP-assigned outer address. Servers in the enterprise need to find a route back to these clients. You must configure the switch as the default gateway on the network. The switch can then forward tunneled traffic to served clients and forward other traffic to the Internet or other default gateways. This option is not always desirable because of the impact on the customer network infrastructure.

- Authentication -- Various authentication services supported with the Contivity VPN Client are not supported with third-party clients. Radius, SecureID, Axent, and other Radius-based services will not work with the switch, even if the third-party client has the support available. LDAP with pre-shared key and unmanaged certificates are the only authentication services supported by the switch with third-party clients.
- Client Customization -- This capability allows a service provider to customize the look of the client with their branding. In addition, it allows the service provider to preconfigure the service profiles (switch destination and authentication options) and lock down the client configuration for the end-user so that they cannot modify or change these attributes.
- Load Balancing -- Traditional load balancers do not often work with the IPSec protocol because of the security features on individual packets and separate key management and data channels. The switch has built-in load balancing features for IPSec client terminations that allow two switches to load balance and failover connections. This feature works with third-party clients.
- QoS -- The Contivity VPN Client is subject to manager-defined QoS policies. Connection slots can be reserved for different classes of user, and they can be assigned differing forwarding priorities for their traffic. The switch preserves Diff-Serv markings for dial tunnels, copying the Diff-Serv Code Point from the inside packet to the tunnel header.
- Advanced attribute definition from the server -- On a group-by-group basis, you can load the client with its tunneled IP address and subnet mask, a Microsoft domain name, both WINS and DNS servers, a message of the day and the Contivity banner. The network manager can also determine access days and hours, crypto strength, how often the client will re-key, and whether the client can store a password for the group. It can initiate a password-protected screen saver if the user leaves the PC, and can log off idle connections. You can filter traffic in the tunnel based on IP address and/or port number and can configure to close the tunnel if certain network

applications are run. You can set the tunnel to automatically start when predefined applications or destinations are accessed, and close when these application are completed. These features are not available with third-party clients.

- Address Assignment -- Client-tunneled IP addresses can be assigned via a DHCP server, on a per-group basis from a named pool, via RADIUS attribute, or statically. The client receives the inner IP address from the enterprise address space. Third-party remote access clients get their inner address assigned the same as the outer, which is normally what the ISP assigns, and is not part of the enterprise address space.
- Split Tunneling -- On a group-by-group basis, a service provider can determine which IP addresses will go into the tunnel and which will use the local adapter (for general Internet access, or local printing/server usage). With third-party clients, you should enable split tunneling. If disabled, the client must be put into a group configured to allow undefined networks.
- Advanced Security features -- The Contivity VPN Client tunnel only accepts packets originating from the machine on which it is loaded. If attempts are made to route packets through a Contivity VPN Client, the tunnel is closed. When non-split tunneling is enabled, only packets that have passed through the VPN (have been correctly decrypted, and authenticated) will be accepted; other packets are dropped. If any attempt is made to change the station address of the client, the tunnel is automatically closed. Third-party clients do not necessarily have this security.
- Tight integration with MS-DUN and IPASS -- This allows one-click access that dials and authorizes the ISP connection and then creates the VPN connection automatically. This makes it significantly easier for the end user. Third-party clients typically do not have this ease-of-use feature.
- High end PKI integration -- The switch integrates software from the leading certificate vendors, for a high-end managed PKI implementation. Managed PKI features like automated enrollment and automatic renewal are critical for large-scale rollouts. Other clients have loose or no integration for managed PKI and rely on the features of a browser or simple cut-and-paste methods. This is not available with third-party clients when used with the switch, even if the client has the support built in.

Configuring the switch as a branch office tunnel

To configure the switch as a branch office tunnel:

1. Go to the Profiles→Branch Office screen and click on the Define Branch Office Connection button. The Branch Office→Define Connection screen appears.
2. Specify the addresses of the public interfaces of the two switches forming the connection.
   a. For the local endpoint address, select the address of the local switch from the drop down list.
   b. For the remote endpoint address, enter the address of the remote switch that will form the opposite end of the branch office connection.
3. Set the tunnel type to IPSec.
4. Depending on what your third-party clients support, you can use either pre-shared key or digital certificate authentication. Click to enable the user name and password to authenticate user identity. The user name is the user's IP address and the password can be any password. Match the pre-shared secret with the client shared secret.
5. Click on RSA Digital Signature to enable certificate authentication if your third-party client supports RSA Digital Signature authentication. You must then select a default server certificate from the drop down list. You configure servers from the System→Certificates screen.
6. Go the Profiles→Branch Office screen and click on the Edit button, scroll down to the IPSec section and click on the Configure button. The Branch Office screen appears.
7. Select the encryption type supported by your third-party client.
8. Select Enable or Disable for the VendorID.
9. Set Perfect Forward Secrecy (PFS) to match the client side.
10. In the Rekey Time-out section, enter the amount of time to which you want to limit the lifetime of a single key used to encrypt data. The default is 08:00:00 (8 hours).
11. In the Rekey Data Count section, you can choose to set a rekey data count depending on how much data you expect to transmit through the tunnel with a single key. The default is 0 KB; a setting of 0 disables this count.

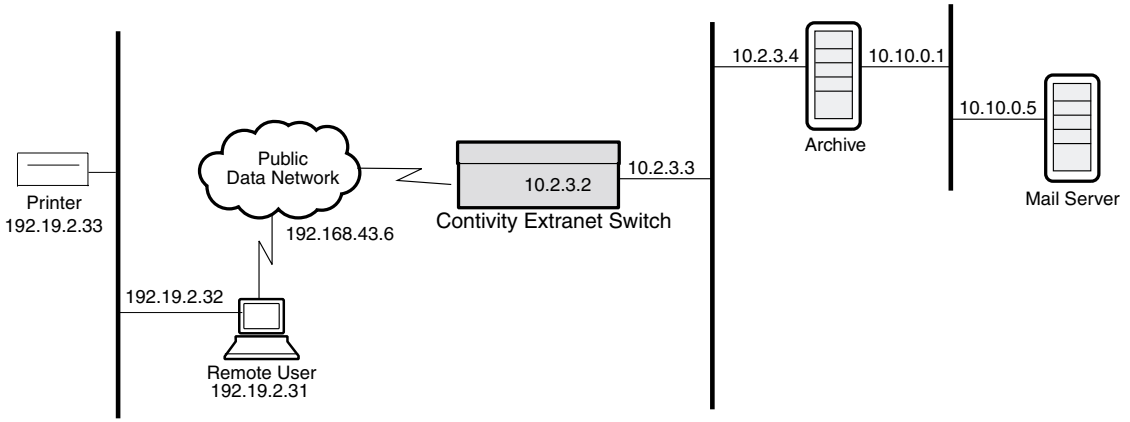
Configuring the switch as a user tunnel

If you have third-party client software that supports Aggressive mode IPSec, you can configure the switch as a user tunnel. You must use either the LDAP database or the certificate authentication. The switch supports both pre-shared key and RSA digital signature authentication methods and you must specify at least one of these methods.

Nortel Networks recommends enabling split tunnels for all groups that support third-party clients. If split tunneling is disabled, third-party clients can only connect if the group is configured to allow undefined networks. This means that the client can establish IPSec security associations for all networks. If you do not enable split tunneling, you must enable the Allow undefined networks option.

Figure 76 shows a network with a split tunneling environment.

Figure 76 Split tunneling example



To configure the switch as a user tunnel:

1. Go the Profiles→Groups screen and click on the Add button. Enter a group name of up to 64 characters (spaces are permitted); for example, Research and Development.
2. Click on the Edit button next to the name of the new group, scroll down to the IPSec section, and click on the Configure button. The IPSec Edit screen appears.

3. Set Split tunneling to Enabled if you want your switch to have control over the networks that the third-party client can access. If split tunneling is disabled and Allow undefined networks for non-Contivity VPN Clients is enabled, the clients can connect to all internal networks. If you select both Split Tunneling and Allow undefined networks for non-Contivity VPN Clients, the switch uses the split tunneling feature and ignores the Allow undefined networks selection.
4. Under Client Selection, select Non-Contivity VPN Clients (Linux*) or Both Contivity and Non-Contivity VPN Clients from the drop-down list.
5. Third-party clients can use either pre-shared key or digital certificate authentication. Click on the check box to enable the user name and password to authenticate user identity. If you are using Main mode, the user name is the user's IP address and the password can be any password.

   Click on RSA Digital Signature to enable certificate authentication if your client supports this. You must then select a default server certificate from the drop down list. You configure servers from the System→Certificates screen.
6. Selections in the Encryption fields are dependent on the type of encryption that your third-party client supports.
7. Click on the checkbox to enable Perfect Forward Secrecy (PFS). PFS ensures that if one key is compromised, subsequent keys are not compromised.
8. In the Forced Logoff edit box, specify a time after which all active users are automatically logged off. The default is 0, which means the option is turned off. The possible range is 00:00:01 to 23:59:59.
9. Click on the checkbox to enable compression for IPsec tunneling.
10. In the Rekey Time-out section, enter the time to which you want to limit the lifetime of a single key used to encrypt data. The default is 08:00:00 (8 hours).
11. In the Rekey Data Count section, you can choose to set a rekey data count depending on how much data you expect to transmit through the tunnel with a single key. The default is 0 KB; a setting of 0 disables this count.
12. Select Enable or Disable, depending on whether you want to allow IPsec Data Protection.

Configuring IPX

The Internetwork Packet Exchange (IPX) protocol is the Novell adaptation of the Xerox Networking System (XNS) protocol. IPX has the following characteristics:

- It is a connectionless datagram delivery protocol. A datagram is a unit of data that contains all of the addressing information required for it to be delivered to its destination.
- It does not guarantee the delivery of packets. Higher-level protocols assume the responsibility for reliability.

The switch supports IPX by encapsulating IPX traffic within PPTP client connections. Note that the switch's IPX support is not available for the IPSec tunneling protocol.

IPX is the network-layer routing protocol used in the Novell NetWare environment. The primary tasks of IPX are addressing, routing, and switching information packets from one location to another on a network. In a LAN-based client the network interface card (NIC) provides network node addressing; in a tunneled environment, the switch provides the network node addressing.

Network addresses form the basis of the IPX internetwork addressing scheme for sending packets between network segments. Every network segment of an internetwork is assigned a unique network address by which routers forward packets to their final destination network. On the switch, all public interfaces are treated as a single network segment with a unique network address. A network address in the NetWare environment consists of eight hexadecimal characters. In the following example, 0x indicates that this is a hexadecimal number, and *n* is any hexadecimal character.

`0x`*`nnnnnnnn`*

Socket numbers are the basis for an IPX *intranode address* (the address of an individual entity within a node). They allow a process (for example, IPX Routing Information Protocol [RIP] and Service Access Points [SAP]) to distinguish itself to IPX. To be able to communicate on the network, the process must request a socket number. Any packets IPX receives addressed to that socket are then passed on to the process within the node.

The switch uses IPX RIP and SAP to dynamically learn and advertise IPX routes and services. The switch assigns IPX addresses to tunneled clients; remote users cannot configure the IPX tunnel address for their systems.

The switch does not forward IPX packets from a private nontunneled LAN to another private nontunneled LAN, nor does it propagate routing or server tables from a private nontunneled LAN to another private nontunneled LAN.

IPX client

On the PPTP client (for example, Microsoft Dial-Up Networking), you must enable the dial-up networking IPX option. Enabling the IPX option allows you to tunnel using IPX, IP, or IPX and IP according to the dial-up networking selections.

Windows 95 and Windows 98

When running Windows 95 or Windows 98, load the intraNetWare client, which is available from the Novell* Web site:

```
http://www.novell.com
```



Note: The NetWare* client for Windows 95 and Windows 98 does not function properly; therefore, you must use the Novell intraNetWare client when using IPX with PPTP.

Windows NT

The NetWare client is already on Windows NT systems. You can use that or the Novell intraNetWare client, which you can access from the Novell Web site at http://www.novell.com.

IPX group configuration

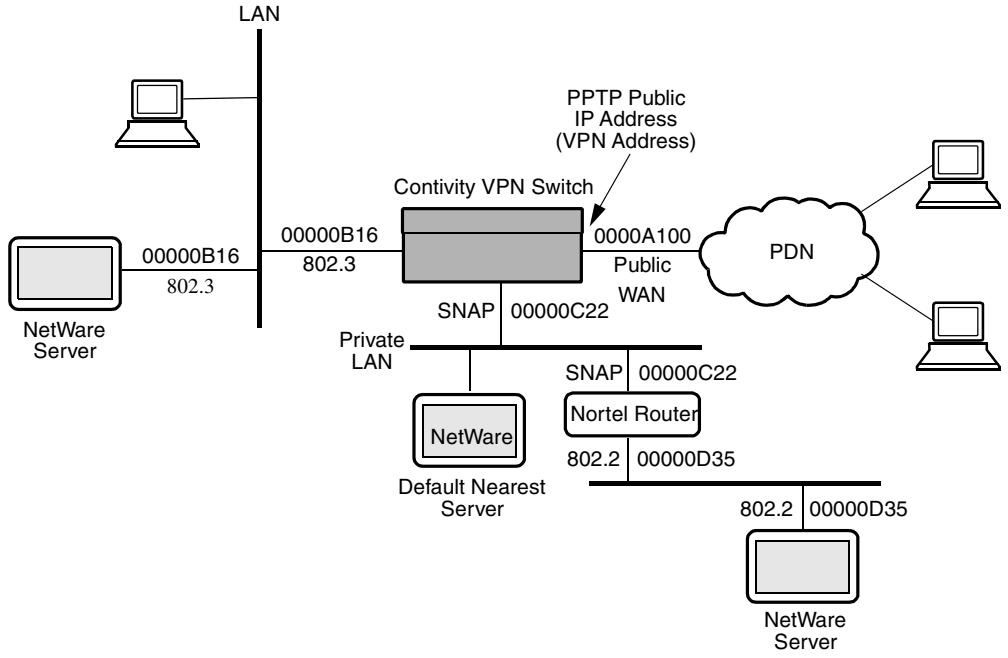
IPX is disabled on a per-group basis by default. Therefore, you must enable IPX for group users to access IPX. Enable IPX for group users from the Profiles→Groups→Edit→Connectivity screen.

Sample IPX VPN switch topology

Regardless of the number of IPX public interfaces that are configured on the switch, they all use the same IPX network address. You must enable the private interfaces that you want to use for IPX traffic, and for each private interface you must configure the IPX network address and IPX frame type. The IPX network address that you configure must match the IPX network address for that LAN, and the IPX frame type must match the IPX frame type for that LAN. In the following figure, the public interface IPX network address that the switch provides is 0000A100.

In Figure 77, the private interface network address to the NetWare server is 00000B16 and the Frame Type is 802.3; similarly, the private interface network address to the Nortel Networks Router is 00000C22 and the Frame Type is SNAP.

Figure 77 IPX topology



Note: The Private LAN can also carry IP and IPX traffic simultaneously. The IP addresses are not shown in this figure.

Glossary

acknowledgement (ACK)

A type of message sent to indicate that a block of data arrived at its destination without error.

address masks

IP addresses used to represent a series or range of IP addresses.

authentication

A security procedure where a user verifies his identity before accessing networks protected by a firewall.

bandwidth

The difference between the highest and lowest frequencies of a transmission channel; amount of data that can be sent through a given communications circuit.

certification authority (CA)

An authority that issues digital certificates and manages the life cycle of certificates.

Challenge Handshake Authentication Protocol (CHAP)

A peer entity authentication method for PPP, using a randomly-generated challenge and requiring a matching response that depends on a cryptographic hash of the challenge and a secret key.

client

A system or process that requests a service of another system or process.

default route

A route that is used when the switch receives traffic for which no matching route is in the routing table.

Diffie-Helman

A key agreement algorithm that does key establishment, not encryption. However, the key it produces may be used for encryption, for further key management, or any other cryptography.

digital certificate

A certificate document in the form of a digital data object to which is appended a computed digital signature value that depends on the data object.

distinguished name (DN)

An identifier that uniquely represents an object in the X.500 Directory Information Tree. An X.509 public-key certificate or CRL contains a DN that identifies its issuer, and attribute certificate identifies its subject.

Domain Name System (DNS)

A general purpose distributed, replicated, data query service used to look up host IP addresses based on host names. DNS applications can perform name-to-address and address-to-name translations.

dynamic routes

Routes that are learned via the switch's RIP support, and are used for branch office connections and the private interface.

encryption

The manipulation of a packet’s data to prevent any but the intended recipient from reading the data.

encryption certificate

A public-key certificate that contains a public key that is intended to be used for encrypting data, rather than for verifying digital signatures or performing other cryptographic functions.

Federal Information Processing Standards (FIPS)

Technical guidelines for U.S. Government procurements of information processing system equipment and services.

File Transfer Protocol (FTP)

A TCP-based application layer Internet Standard protocol that transfers files to and from a remote host.

firewall

A collection of hardware and software components that controls communication between two networks, such as a private network and the Internet. All information passed between the two networks must pass through the firewall. The firewall allows only authorized traffic to pass between the networks.

gateway

A communications device or program that passes data between networks having similar functions but dissimilar implementations.

integrated firewall

The switch supports two integrated firewalls: the Contivity Firewall and the FireWall-1 Inspection Module from Check Point Technologies, LTD. Also referred to as an embedded firewall.

interface

The connection between a router and one of its attached networks. An interface to a network has a single IP address and mask associated with it.

Internet

The single, interconnected, worldwide system of commercial, government, educational, and other computer networks that share a set of protocols.

Internet Protocol (IP)

The transport layer protocol used by the Internet Protocol family for transporting information among computers.

Internet Security Association and Key Management Protocol (ISAKMP)

Defines how encryption keys for sessions are initiated and updated.

intranet

A computer network, especially one based on Internet technology, that an organization uses for its own internal, and usually private, purposes and that is closed to outsiders.

IP address

The identifiers used by the protocols that govern Internet information exchange. The Internet Network Information Center assigns these numbers to uniquely identify different machines on the Internet.

IPsec

A tunneling protocol that offers a strong level of encryption, integrity protection. It uses digital certificates, password-based keys, and tokens for authentication.

IPsec Key Exchange (IKE)

An Internet IPsec key-establishment protocol that puts in place authenticated keying information for use with ISAKMP and for other security associations.

key agreement

A method for negotiating a key value without transferring the key, even in an encrypted form, such as Diffie-Helman.

Layer2 Tunneling Protocol (L2TP)

Tunneling protocol that enables secure remote access to enterprise networks across the Internet.

Lightweight Directory Access Protocol (LDAP)

Protocol based on directory entries that provide access for management and browser applications that provide read/write interactive access to the X.500 directory.

local area network (LAN)

A data network intended to serve an area of only a small area to optimize data transfer rates.

management information base (MIB)

The set of parameters an SNMP management station can query or set in the SNMP agent of a network device, such as a router.

management IP address

The IP address that is used to manage all system services from a Web browser, such as HTTP, FTP, and SNMP. This address must be accessible from one of the switch's private physical interfaces. To be accessible, the Management IP Address must map to the same network as one of the private interfaces.

management station

The remote workstation that is used to configure and manage the Contivity VPN Switch's integrated Check Point firewall. The FireWall-1 Management Console software runs on Management Station. The GUI portion of the Management software can be installed on the Management Station or on a different system.

medial access control (MAC) address

The hardware address of a device connected to a shared media.

Network Address Translation (NAT)

A mechanism that converts an internal network's private addressing scheme to an acceptable Internet address, thereby enabling the internal systems to communicate on the Internet.

Network Time Protocol (NTP)

Synchronizes the clocks of various devices across networks.

Open Shortest Path First (OSPF)

OSPF is a link-state routing protocol that maintains a database from which a routing table is constructed from the shortest path, using a minimum of routing protocol traffic.

packet

The unit of data sent across a network. Typically, it refers to application data units.

PING

A program used to test reachability of destinations by sending an ICMP echo request and waiting for a reply.

Point-to-Point Protocol (PPP)

A protocol that provides a method for transmitting packets over serial point-to-point links.

Point-to-Point Tunneling Protocol (PPTP)

A tunneling protocol that is used as a security tool.

port

A transport layer demultiplexing value. Each application has a unique port number associated with it.

private default route

The default routes that are used for traffic that comes into the switch via a public interface, through a tunnel, or from the switch's public interface address.

protocol

A formal description of message formats and the rules two computers must follow to exchange those messages. A protocol can describe low-level details of machine-to-machine interfaces or high-level exchanges between allocation programs.

public default route

The default routes that are used for traffic that comes into the switch via a private interface or from the switch's private interface address.

Resource Reservation Protocol (RSVP)

A protocol used to signal QoS requests and confirmations.

route

The path that network traffic takes from its source to its destination.

router

A device that forwards traffic between networks. The forwarding decision is based on network layer information and routing tables.

routing

The process of selecting the correct interface and next hop for a packet being forwarded.

Routing Information Protocol (RIP)

A distance vector, as opposed to link state, routing protocol.

RSA digital signature

A public-key encryptographic system that may be used for encryption and authentication.

server

A provider of resources, such as file servers and name servers.

Simple Network Management Protocol (SNMP)

The Internet standard protocol developed to manage nodes on an IP network.

split horizon

A method that RIP uses to avoid routing problems caused by including routes in updates sent back to the gateway from which they were learned. The simple split horizon scheme omits routes learned from one neighbor in updates sent back to that neighbor. An extension to this method is called split horizon with poisoned reverse. It includes the learned routes, but assigns them a cost of infinity, which causes an update.

static routes

Routes that are manually configured in the switch's routing table.

stub network

A network that only carries packets to and from local hosts. Even if it has paths to more than one other network, it does not carry traffic for other networks.

subnet

A portion of a network, which may be a physically independent network segment, that shares a network address with other portions of the network and is distinguished by a subnet number.

Telnet

A command protocol used to establish login sessions on a remote host.

triggered update

A method used by RIP in which a new routing table is sent almost immediately after a routing change has been made. This is in contrast to the poison reverse method, in which routes are updated after a cost of infinity is reached, a process that can take much time.

User Datagram Protocol (UDP)

An Internet standard transport layer protocol. It is a connectionless protocol that adds a level of reliability to an multiplexing to IP.

Uniform Resource Locator (URL)

A standard for identifying objects on the Internet accessible through the Web.

Virtual Router Redundancy Protocol (VRRP)

A protocol that handles private interface failures. VRRP targets hosts that are configured with static next-hop routing addresses or default gateways. It provides a means of rerouting traffic in the event of a system/interface failure.

Wide Area Network (WAN)

A network, usually constructed with serial lines, that covers a large geographic area.

Index

A

access control
  subject DN 283
access control filters 51
access hours 35, 49, 104, 105, 109, 139
accessible networks 110, 111, 119, 127
accounting
  data 311
  overview 54
  records 310, 312
ACE 256
actions on rules 239
active sessions 341
ActiveX Scripts 338
administrator 66
  settings 292
administrator privileges 291
advanced configuration 263
advanced routing key 44, 164
Allow All
  enabling 282
anti-spoofing 214, 221
Asymmetric Branch Office tunnel 109
attack detection 215
authentication
  branch office 288
  details 246
  failed 319
  group password 256
  overview 40
authentication methods 35, 49, 109
authentication servers
  configuring 245
automatic backup 295
available rules 143
AXENT tokens 260

B

background images 341
backup
  host server(s) 293
backups 292
bandwidth guarantees 148
bandwidth management 52, 145
banner 134
base group 102
branch office
  access hours 35, 49, 109
  accessible networks 119
  authentication 288
  authentication methods 35, 49, 109
  configuring 109
  encryption settings 112
  indirectly connected switches 111
  sample configurations 110
  sample procedure 113
branch office connections 49
branch office quick start 77
  template 83
branch office quick start paramenters 81
branch office quick start utility 59
branch office tunnels

asymmetric 109
peer to peer 109
browser error messages 339
browsing delays 338

C

call admission
guarantees 53, 148
priority 53, 105, 148
cards
LAN 63
WAN 63
cell menu 233
Certificate Management Protocol (CMP) 278
certificate revocation list (CRL) 283
certificates
Allow All option 282
branch office 288
details 281
owner 281
cestraps.mib 352
CHAP 39
RADIUS 254
Check Point Firewall-1 50, 211, 221
children of base group 102
class attributes RADIUS 258
client 249
client address redistribution 185
configuring 187
sample 186
summarization 186
client atuo connect 133
client failover 132
client policy 135
client screen saver 135
CMP 278
color setting 341
columns
Dst interface 234
Src interface 234
Command Line Interface (CLI) 56
compact flash disk 60, 85
compression 39
configuration
choices 73
firewall 243
initial 222
log 301
perparing for 69
saving current 298
verifying 244
Web interface options 70
configuration checklist 70
configuration file
local 55
storing 54
configurations
restoring 294
configuring
branch offices 109
serial interface 63
Confivity 1010/1050/1100 76
connection types
initiator 117
peer to peer 117
responder 117
connectivity problems
overview 313
solving 316
Contivity 1010/1050/1100 57
branch office quick start 77
compact flash disk 85
default configuration parameters 76
enterprise requirements 78
ISP environment 79
setting up 83
control tunnels 120
branch office 122
creating 124

nailed-up 123
restricted mode 123
sample 122
user 122
conversation 213
critical interface fail-over 195
CRL settings 284

D

data collection records 311
data storage 310
default
login 69
password 69
default group
client authentication 282
default route 179
branch office 111
default rules 230
Defender 256
DES 38, 250
DHCP client 90
DHCP relay 270
DHCP server 269, 271, 327
dial-up
monitor 316
problems 317
Dial-Up Networking 334
Differentiated Services (DiffServ) 52, 146
Diffie-Hellman 249
Diffie-Helman
groups 250
digital certificates
SSL 42
direct connections 33
DNS 39
server 335, 338
DNS proxy 60
DNS server 88
docking station configurations 323
domain controller 327
domain name 135
Domain Name Service (DNS) 60
Dynamic Host Configuration Protocol (DHCP) 58, 271

E

encryption
settings for branch office 112
enhanced routing 44, 154
Entrust 43
Entrust Manager digital certificates 43
equal cost multipath (ECMP) 184
error messages 344
event log 301, 304
sample 305
external
LDAP 252
External DHCP server 342
extinction
interval 329
timeout 329
Extranet Access 30
client monitor 314
connection problems 318

F

failover
client 132
file management 298
filtering 214
filters through a firewall 35
fingerprint
certificate 281
firewall

branch office 110
configuring 243
imbedded 50, 211
installation prerequisites 216
integrated 50, 211
interaction with branch office 111
options 224
routing 157
firewall options 221
flash disk system
compressed files 61
forced logout 131
forwarding priority 53, 105, 147
quality of service 147
full distinguished name
branch office 287
fully qualified domain name
authentication server 252

G

general problems
overview 314
solving 337
getting started 63, 87
group associations 104
group hierarchy
configuration 103
inheritance 103
group password authentication 256
guided configuration 29
requirements 74

H

hardware
health check 304
hardware encryption accelerator 96
HDLC framing 325
header row menu 232
health check 304
health check display 341
historical event logging 301
HMAC 250
HTTP 338
Hyperterminal 317

I

ICMP
filter 144
idle timeout 105
IKE 276
IKE encryption 250
implied rules 227, 231
indirectly connected switches 111
inheritance, group 105
inner IP address 269
installation prerequisites 216
interface
classifiers 143
interface filter 91
interface groups 195
interfaces 213
interface-specific rules 228
internal
LDAP 252
internal address pool 342
Internet Explorer 340
Internetwork Packet Exchange 443
interval
session update 268, 269
IP accept policies 176
IP address
assigning 47
branch office 118
currently assigned 91
reconfiguring values 67
IP address pool 269

IP announce policies 176
IP packets 144
ipconfig command 316
IPSec 36, 37, 38, 249
  password 319
  tunnel filters 35
  username 319
IPsec
  third-party clients 130
IPSec certificate credentials 287
IPSec client 128
IPSecShm 323
IPX 48, 443
IPX client 444
ISAKMP 38, 43
ISAKMP/IPSec filters 35
ISP 95
  tunnels 102

J

Java 337, 338, 340
Java 2
  Netscape 218
Java 2 Runtime Environment
  Internet Explorer 217
  Netscape 6 220
  Netscape on Solaris 220
JavaScript 338
jetpack.exe 330

L

L2F 36, 37, 40
L2TP 36, 37
  configuring over IPSec 135
L2TP/IPSec 39
L2TP/IPSec client 140
L2TP/IPSec gateway 141
L2TP/L2F
  filters 35
LAN
  card 91
LAN interfaces 88
last name search 109
LCP options 326
LDAP 251
  attribute search 109
  authentication 40
  database authentication 250
  database files 54
  directory 251
  full vendors 251
  Netscape Directory Server 266
  overview 41
  server port number 285
LDAP authentication 245
LDAP certificates 276
LDAP server authentication 252
licenses 73
Lightweight Directory Access Protocol 251
log column 240
logging 54, 56
  displays 315
login 69
  ignored 340
  not allowed 319
logs
  events 304
  levels 240
  system 309
loopback test 324

M

MAC Pause 92
MAC Pause Ticks 92
Manage CRL Servers 285
management

Extranet Access Switch 74
tasks 29
management IP address 67, 87
manager WAN statistics display 325
master browser 330
maximum number of sessions 319
MBONE 144
MD5 38
menus
cell 233
header row 232
row 232
MIB support 347
Microsoft
auto disconnect feature 321
client troubleshooting tools 316
Internet Explorer 337
Knowledge Base 336
networking tips 327
modem hardware errors 326
MS-CHAP 39
MS-DOS naming convention 342
Multicast 183
multicast relay 182
multiple Help windows 340

N

NAT
branch office 204
configure new policy 208
double 203
dynamic 200
dynamic address translation 198
dynamic pooled 202
dynamic port 200
editing existing set 208
interface NAT 207
rules 207
sets 206
static address translation 198
translation type 207
NAT Traversal 209
navigational menu 74
NetBEUI 322, 328
NetBIOS 322, 327, 328, 333
Netscape
Communicator 337
Netscape Directory Server
LDAP 262, 266
Netscape directory server 261
netstats command 316
NetWare client 444
network access server (NAS) 39
Network Address Translation (NAT) 112, 197
network address translation (NAT) 50
network configurations 32
Network Neighborhood 328
network objects 236
Network Time Protocol (NTP) 53, 98
newoak.mib 354
Nortel Networks MIB 299
Novell intraNetWare client 444

O

OSPF
configuration 166
OSPF overview 45, 163
outer IP address 269
override rules 227

P

PAP 39
RADIUS 254
password 69
password storage 134

performance problems
overview 314
solving 326
Permit All Interface 183
ping command 318, 321
Point to Point Protocol over Ethernet (PPPoE) 59, 90, 91
poison reverse 161
policies
actions 225
adding 241
components 225
copying 242
creating 226, 241
deleting 242
editing 226
renaming 242
selecting 241
pooled translation type 207
port speed 92
ports
RADIUS accounting 269
post-implied rules 227, 231
power failure 341
PPP 57
layer 325
PPTP 37, 39
control connection filters 35
tunnels filters 35
white papers 334
PPTP and L2TP authentication 289
pre-shared key 288
primary administrator 66, 292
primary WINS server 323
private LAN 46, 88
proxy ARP 150
public data network 89
public data network (PDN) 46
Public Key Infrastructure (PKI) 276
public key infrastructure (PKI) 43
public key sizes 279

Q

quality of service 52, 143
forwarding priority 53, 147
RSVP 53, 150
Quick Start 74
configuration 29

R

RADIUS 245
accounting 245, 311
authentication 40, 255
class attributes 258
configuring server 259
overview 41
RADIUS accounting 268
RADIUS server authentication 254
RADIUS service 273
RC4-128 275
RC4-40 275
recovery diskette 296
register 73, 74
relative distinguished name 104
remarks 240
remote access 31
branch office 141
configuring 135
configuring the CES 135
configuring Windows 2000 140
remote identity 288
renewal interval 329
reports 303
reservation control filters 35
restart failure 342
RIP 45, 159
RFCs 160

using 160, 164
route redistribution 179
route selection 175
routes
default 179
dynamic 44, 170
static 44, 170
routing
dynamic 155
enhanced 157
integrated firewall 155
loops 161
overview 43, 153
route lookup 174
routing table types 174
rules of redistribution 179
table 44, 170
traffic patterns 180
VPN 44, 156
routing policy
configuring 178
routing policy service 44, 176
routing table 44, 170
lookup 174
saving as file 172
searching 172
status 173
row menu 232
RSA digital signature
certificates 43, 276
RSVP quality of service 53, 150
rule column 234
rules
default 230
implied and post-implied 227, 231
in policies 213
interface-specific 228
navigating 226, 231
override 227

S

Safe mode 68
search for users 109
Secure Socket Layer 275
SecurID 249
security association 249
Security Dynamics 249
serial interface
configuration 63
prerequisites 63
serial PPP
dial-up networking 368
establishing a connection 367
option settings 373
overview 367
setting up the switch 369
troubleshooting 371
server certificate 277, 279
server certificates 280
PKCS #7 and #10 277
servers
external RADIUS 245
internal LDAP 245
LDAP authentication 245
RADIUS 245
service objects 238
Service Pack 2 334
services 73
services routing 44
sessions 302
SHA 38
SHA-1 249
SMNP version 1
traps 298
SNMP 298
MIB II 30
overview 30
traps 30

split horizon 161
split proxy DNS 60
split tunneling
  enable 128
  third-party clients 131
split-horizon DNS 321
SSL
  port number 285
SSL digital certificates 42
stateful inspection 212
  application 212
  TCP 212
static
  translation type 207
static address
  NAT 199
static routes 158
static translation type 207
static tunnel failover 133
statistics 90, 304
statistics display 315
status 240, 300
subject DN 136, 140, 282
subnet mask 91
subnetworks 35, 49, 109
Switch concepts 29
Symmetric Branch Office tunnel 109
synchronize RADIUS servers 252
system
  log 301
  shutdown 299
  status 303
system identity 87
system log 309
system messages 375
  branch office 380
  certificate 344, 375
  database 382
  hardware 406
  ISAKMP 377
  RADIUS accounting 393
  RADIUS authentication 396
  routing 400
  security 383
  SSL 381
system requirements 216

T

T1/V.35 interface 324
TCP
  filter 144
Telco expenses reduced 32
template 83
terminal emulator 63
third-party clients
  authentication methods 131
  split tunneling 131
third-party IPSec clients 130
tokens
  card 256
  security 249
tools
  ARP 296
  ping 295
  traceroute 296
tracert command 316
traffic conditioning 147
traffic patterns 180
translation types 207
traps
  hardware 355
  information for all 362, 363
  intrusion-related 362
  login-related 361
  server-related 359
  software-related 361
  system-related 362

triggered updates 161
triple DES 38
troubleshooting
  client address redistribution 343
  client connection problems 318
  Extranet Access Manager 337
  Internet service provider login problems 336
  modem and dial-up problems 317, 336
  overview 313
  PPTP connectivity 337
  routing 343
  toolbox 314
  WAN link problems 324
trusted CA certificate 278, 279
trusted CA certificates 279, 281
tunnel types 106
tunnels, configuring 101

U

UDP
  filter 144
Uniform Resource Locator (URL) 68
upgrades 297
upgrading software 339
user
  ID search 109
user control tunnel
  serial interface 127
user groups
  adding 107
  searching 108
user profile
  adding 107
user tunnels 106
Utunnel 185

V

valid issuer certificate authority
  branch office 288
verify interval 329
virtual links 165
VPN connector
  digital certificates 43
VPN routing 44
VRRP
  configuring 193
  failover 190
  high availability 189
  master interface 192
VRRP overview 45, 189

W

WAN interfaces
  configuration 95
  currently installed 96
  display 325
Web browser
  problems 337, 341
web browser interface 68
Web Connector
  digital certificates 43
Web interface options 70
Welcome display 72
winipcfg command 316
WINS 39
  secondary servers 329
  server 329, 333
  settings 329
Winsock DNS Update 335

X

X.500 directory search base 285
X.509 certificates 43, 276, 277