

Part No. 314995-B Rev 00
June 2004

4655 Great America Parkway
Santa Clara, CA 95054

Configuring the Web Switching Module Using Device Manager



* 3 1 4 9 9 5 - B R E V 0 0 *

NORTEL
NETWORKS™

Copyright © 2004 Nortel Networks

All rights reserved. June 2004.

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, the Globemark, Unified Networks, Alteon WebSystems, and Alteon are trademarks of Nortel Networks.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation.

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

Real Server and Real Player are trademarks of Real Networks.

Quicktime Streaming Server and Quicktime Player are trademarks of Apple Inc.

The asterisk after a name denotes a trademarked item.

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 | 37 |
| Before you begin | 37 |
| Text conventions | 37 |
| Related publications | 39 |
| Hard-copy technical manuals | 39 |
| How to get help | 40 |
| | |
| Chapter 1 | |
| Getting started with Device Manager | 41 |
| About Device Manager | 41 |
| Starting Device Manager | 42 |
| Opening a device | 43 |
| Troubleshooting the opening of a device | 45 |
| Understanding the Device Manager window | 46 |
| Menu bar | 47 |
| Accessing the Edit > WSM Card submenu | 48 |
| Accessing the Edit > WSM Card > L4 Switching submenu | 51 |
| Accessing the Graph > WSM Card submenu | 53 |
| Accessing the Actions menu | 54 |
| Accessing the Help menu | 55 |
| Shortcuts | 57 |
| Using the chassis shortcut | 57 |
| Using the WSM shortcut | 58 |
| Using the WSM port shortcut | 60 |
| Toolbar | 61 |
| Device view | 62 |
| Selecting objects | 62 |
| The WSM device view | 64 |

| | |
|---|----|
| Interpreting the status of LEDs and ports | 65 |
| Viewing the port color code legend | 66 |
| Status bar | 66 |
| Using Device Manager dialog boxes | 66 |
| Editing objects | 67 |
| Editing a text field | 68 |
| Formats for entering numeric values | 69 |
| Editing a field from a selection list | 69 |
| Using buttons in dialog boxes | 70 |
| Applying and saving changes | 72 |
| Refreshing the screen | 72 |
| Filtering table content | 73 |
| Adding a table filter | 74 |
| Removing a table filter | 75 |
| Removing all table filters | 75 |
| Replacing a table filter | 76 |
| Browsing a read only dialog box | 76 |
| Opening WSM online Help | 77 |

Chapter 2

Port management..... 81

| | |
|---|----|
| About ports | 81 |
| About the device view | 82 |
| Configuring ports | 83 |
| Editing ports | 83 |
| Naming a port | 83 |
| Setting port parameters | 85 |
| Configuring VLAN tagging | 88 |
| Monitoring link state | 89 |
| Enabling or disabling filtering on a port | 89 |
| Applying filters to a port | 90 |
| Removing a filter from a port | 94 |
| Port trunking | 95 |
| Default multi-link trunk configuration | 95 |
| Dynamic MLTs | 96 |

| | |
|--|------------|
| Viewing trunk group configurations | 96 |
| Configuring a trunk group | 98 |
| Deleting a trunk group | 99 |
| | |
| Chapter 3 | |
| Configuring Layer 2 operations | 101 |
| About VLANs | 101 |
| Management IP address used for VLAN 4093 | 102 |
| VLANs and IP interfaces | 103 |
| About VLAN tagging | 103 |
| About Jumbo frames | 103 |
| Isolating Jumbo frame traffic using VLANs | 104 |
| Routing Jumbo frames to non-Jumbo frame VLANs | 104 |
| Viewing current VLAN membership | 104 |
| Configuring a VLAN | 106 |
| Viewing MultiLink trunk and VLAN membership | 108 |
| About Spanning Tree protocol | 108 |
| VLANs and spanning tree protocol | 109 |
| Understanding multiple spanning trees | 109 |
| Why multiple Spanning Trees? | 110 |
| Example—four-WSM topology with single spanning tree | 111 |
| Example—four-WSM topology with multiple spanning trees | 112 |
| Multiple spanning tree example—determining root cost | 113 |
| Root cost comparison | 114 |
| Configuring spanning tree protocol | 114 |
| Configuring a Spanning Tree Group | 114 |
| About Spanning Tree bridge | 117 |
| Configuring Spanning Tree bridge | 118 |
| Configuring Spanning Tree port | 123 |
| Enabling or disabling spanning tree on a port | 125 |
| | |
| Chapter 4 | |
| IP routing configuration | 129 |
| Configuring IP interfaces | 129 |
| Manually configuring an IP interface | 130 |

| | |
|--|-----|
| Configuring default gateways | 132 |
| Enabling IP forwarding | 136 |
| Viewing IP routes | 137 |
| Configuring domain name system servers | 139 |
| Defining IP Address Ranges for the Local Route Cache | 141 |
| Adding and removing local networks | 141 |
| Configuring routing information protocol | 143 |
| Configuring IP forwarding per port | 146 |
| Configuring static routes | 148 |
| Deleting a static route | 150 |
| Configuring virtual router redundancy protocol | 151 |
| Enabling virtual routing on the WSM | 151 |
| Configuring a virtual router | 153 |
| Configuring VRRP authentication on an IP interface | 159 |
| Configuring VRRP group parameters | 162 |
| Configuring virtual router priority tracking | 167 |

Chapter 5

Network management and diagnostics **169**

| | |
|---|-----|
| Supported software versions | 169 |
| About the SNMP agent | 169 |
| Alteon WebSystems enterprise MIBs supported | 170 |
| SNMP standard MIBs supported | 171 |
| SNMP generic traps supported | 171 |
| SNMP Spanning Tree traps supported | 171 |
| Supported SNMP traps | 172 |
| Configuring the SNMP agent | 173 |
| Generating system messages | 173 |
| Enabling or disabling syslog and SNMP traps | 175 |
| Enabling or disabling authentication traps | 176 |
| Configuring community strings | 178 |
| Configuring boot options | 179 |
| Viewing current software image details | 179 |
| Selecting a software image to run | 181 |
| Selecting a configuration block to run | 182 |

| | |
|--|------------|
| About port mirroring | 182 |
| Configuring port mirroring | 183 |
| Enabling port mirroring | 185 |
| Viewing device information | 187 |
| Chapter 6 | |
| Configuring filters for Layer 4 switching | 189 |
| About Layer 4 filtering | 189 |
| Well-known protocol numbers | 190 |
| Well-known TCP/UDP application port numbers | 191 |
| About default filters | 191 |
| Configuring a default filter | 192 |
| Creating a new filter | 201 |
| Chapter 7 | |
| Server load balancing basics | 203 |
| About server load balancing | 203 |
| Methods of server load balancing | 205 |
| Topology rules for SLB | 206 |
| About load balancing metrics | 207 |
| Minimum Misses | 207 |
| Hash | 208 |
| Least connections | 209 |
| Round robin | 209 |
| Response time | 210 |
| Bandwidth | 210 |
| About connection timeouts for real servers | 211 |
| About real server maximum connections | 212 |
| Assigning backup or overflow servers | 212 |
| Chapter 8 | |
| Virtual server load balancing | 213 |
| How virtual server load balancing works | 213 |
| Before and after SLB example | 214 |
| Configuring virtual server load balancing | 215 |

| | |
|---|-----|
| Configuring each real server | 216 |
| About IP address ranges for SLB | 222 |
| Configuring IP address ranges for SLB | 222 |
| Defining an IP interface on the WSM | 225 |
| Configuring a real server group | 225 |
| Configuring a virtual server | 230 |
| Configuring services for a virtual server | 235 |
| Configuring ports for server load balancing | 240 |
| Enabling or disabling server load balancing | 243 |
| Peers fields | 244 |
| Synchronize tab fields | 244 |

Chapter 9

Extending SLB topologies **245**

| | |
|---|-----|
| Configuring proxy IP addresses | 245 |
| Disabling server processing on a switch port | 246 |
| Configuring a proxy IP address for a port | 247 |
| Mapping Ports | 248 |
| Mapping a virtual server port to a real server port | 248 |
| Mapping multiple real server ports | 250 |
| Multiple port mapping example | 250 |
| Configuration steps for mapping multiple ports | 251 |
| Configuring multiple port mapping | 252 |
| Providing direct access to real servers | 254 |
| About direct server return | 255 |
| Configuring direct server return | 256 |
| Configuring direct access mode | 258 |
| Assigning multiple IP addresses | 260 |
| Using proxy IP addresses | 260 |
| Mapping ports | 260 |
| Monitoring real servers and services | 261 |
| Using delayed binding to prevent DoS attacks | 262 |
| Enabling delayed binding for a server service | 265 |
| Detecting SYN attacks | 266 |
| Configuring SYN attack detection | 267 |

Chapter 10

| | |
|--|------------|
| Load balancing special services | 269 |
| IP server load balancing | 269 |
| Configuring IP server load balancing | 270 |
| FTP server load balancing | 272 |
| Configuring FTP server load balancing | 273 |
| Domain name server load balancing | 277 |
| Configuration tasks for DNS load balancing | 278 |
| Configuring UDP-based DNS server load balancing | 278 |
| Configuring TCP-based DNS server load balancing | 281 |
| Real Time Streaming Protocol server load balancing | 283 |
| How RTSP server load balancing works | 284 |
| Pre-configuration tasks for RTSP SLB | 285 |
| Configuring RTSP Layer 4 load balancing | 286 |
| Configuring RTSP Layer 7 load balancing | 290 |
| Wireless application server load balancing | 292 |
| How WAP SLB works using RADIUS snooping | 292 |
| Configuring WAP SLB using RADIUS snooping | 294 |
| Configuring WAP server load balancing | 294 |
| Configuring the WAP RADIUS snooping filter | 299 |
| Intrusion Detection server load balancing | 302 |
| How Intrusion Detection SLB works | 303 |
| Configuring Intrusion Detection server load balancing | 304 |
| Configuring real servers for IDS SLB | 304 |
| Configuring real server groups for IDS SLB | 306 |
| Configuring ports for IDS SLB | 309 |
| Enabling IDS SLB | 311 |
| WAN link load balancing | 313 |
| How WAN link load balancing works | 313 |
| Configuring WAN link load balancing | 313 |
| Pre-configuration tasks for WAN link load balancing | 314 |
| Configuring real servers for WAN link load balancing | 314 |
| Configuring real server groups for WAN link load balancing | 316 |
| Configuring filters for WAN link load balancing | 318 |

| | |
|---|------------|
| Chapter 11 | |
| Improving WSM performance with VMA | 323 |
| Enabling VMA | 323 |
| Using proxy IP addresses and VMA | 325 |
| Chapter 12 | |
| Firewall server load balancing | 327 |
| Firewall overview | 327 |
| Methods of firewall load balancing | 329 |
| About basic FWLB | 329 |
| About free-metric FWLB | 330 |
| About demilitarized zones | 331 |
| Basic FWLB implementation | 331 |
| Configuring basic FWLB | 335 |
| Configuring the 8600 switch for FWLB | 336 |
| Configuring FWLB on the public-side of the WSM network | 338 |
| Configuring FWLB on the private-side of the WSM network | 343 |
| Configuring the static routes at the firewalls for FWLB | 351 |
| Configuring basic FWLB with free-metric | 352 |
| Configuring basic FWLB with a DMZ | 353 |
| Monitoring the firewall | 353 |
| Monitoring firewall service | 353 |
| Monitoring physical links | 354 |
| Using HTTP Health Checks | 354 |
| Chapter 13 | |
| Virtual private network server load balancing | 357 |
| Defining virtual private network | 357 |
| How VPN load balancing works | 358 |
| Configuring VPN load balancing | 360 |
| VPN load-balancing configuration example | 360 |
| Configuring private-side network devices | 361 |
| Configuring public-side network devices | 365 |
| Monitoring VPN configurations | 372 |

| | |
|---|------------|
| Chapter 14 | |
| Global server load balancing | 373 |
| How GSLB works | 373 |
| Configuring GSLB | 375 |
| Defining remote GSLB peers for the local site | 376 |
| Configuring remote GSLB services for the local site | 378 |
| Enabling GSLB | 382 |
| IP proxy for non-HTTP redirects | 386 |
| How IP proxy works | 387 |
| Configuring IP proxy for non-HTTP redirects | 388 |
| About GSLB network preferences | 390 |
| Configuring GSLB network preferences | 390 |
| | |
| Chapter 15 | |
| Application Redirection | 395 |
| About Application Redirection | 395 |
| Configuring Web Cache Redirection | 396 |
| Configuring delayed binding for Web Cache Redirection | 398 |
| About RTSP WCR | 399 |
| Configuring RTSP for Web Cache Redirection | 399 |
| Configuring IP proxies with Application Redirection | 401 |
| Excluding noncacheable sites | 402 |
| | |
| Chapter 16 | |
| Health Checking | 403 |
| About health checking | 403 |
| About service failures | 404 |
| About server failures | 405 |
| Enabling graceful server failure | 405 |
| Types of real server group health checks | 407 |
| Configuring a health check for a real server group | 408 |
| Modifying the health check interval and retries | 410 |
| ICMP health checking | 412 |
| ARP health checking | 412 |
| Direct server return health checking | 413 |

| | |
|--|-----|
| Configuring DSR health checks | 413 |
| Link health checking | 415 |
| TCP health checking | 415 |
| Script-based health checking | 416 |
| Configuring script-based health checks | 417 |
| Application health checking | 417 |
| HTTP health checking | 417 |
| Configuring HTTP health checks | 419 |
| UDP-based DNS health checking | 419 |
| IMAP server health checking | 420 |
| RADIUS server health checking | 421 |
| Configuring the RADIUS authentication string | 421 |
| HTTPS/SSL server health checking | 423 |
| WAP gateway health checking | 423 |
| Configuring WSP content health checks | 424 |
| Defining WSP content to check | 424 |
| Configuring the UDP service for WSP health check | 426 |
| Configuring WTLS health checks | 428 |
| LDAP health checking | 430 |
| Configuring LDAP health checking | 430 |
| Changing the LDAP version | 430 |

Chapter 17

Maintaining session persistence

433

| | |
|--|-----|
| About session persistence | 433 |
| Effects of using source IP address for session persistence | 434 |
| Using cookies for session persistence | 434 |
| Using SSL session ID for session persistence | 435 |
| About cookie-based persistence | 435 |
| Defining permanent versus temporary cookies | 436 |
| Cookie formats | 437 |
| Cookie properties | 437 |
| Handling client browsers that do not accept cookies | 438 |
| Defining cookie persistence modes | 438 |
| Insert cookie mode | 439 |

| | |
|--|------------|
| Passive cookie mode | 440 |
| Rewrite cookie mode | 441 |
| Configuring cookie-based session persistence | 442 |
| Examples of cookie values | 443 |
| Example 1: Setting the cookie location | 443 |
| Example 2: Parsing the cookie | 444 |
| Example 3: Using passive cookie mode | 446 |
| Example 4: Using rewrite cookie mode | 446 |
| SSL session ID-based persistence | 448 |
| How SSL session ID-based persistence works | 449 |
| Example of SSL session ID-based persistence | 449 |
| Configuring SSL session ID-based persistence | 450 |
| Chapter 18 | |
| Content-intelligent switching | 453 |
| URL-based server load balancing | 453 |
| URL string formats | 455 |
| Configuring URL-based SLB | 456 |
| Defining a string for URL load balancing | 456 |
| Configuring a real server for URL-based load balancing | 458 |
| Enabling URL-based SLB for the server service | 460 |
| Monitoring URL-based SLB | 462 |
| Virtual hosting | 462 |
| Virtual hosting configuration overview | 463 |
| Configuring the host header for virtual hosting | 464 |
| Cookie-based preferential load balancing | 464 |
| Example of cookie-based preferential load balancing | 465 |
| Defining cookie-based criteria | 466 |
| Configuring cookie-based preferential load balancing | 466 |
| Configuring browser-smart load balancing | 467 |
| URL hashing for server load balancing | 468 |
| Configuring URL hashing | 469 |
| Configuring header hash load balancing | 469 |
| DNS load balancing | 470 |
| Configuring DNS query load balancing | 470 |

| | |
|---|------------|
| Layer 7 RTSP load balancing | 471 |
| Configuring RTSP SLB using pattern matching | 471 |
| Content-intelligent Web Cache Redirection | 472 |
| URL-based Web Cache Redirection | 473 |
| Defining URL-based WCR expressions | 474 |
| Network address translation options | 476 |
| Configuring URL-based WCR | 476 |
| Configuring the WSM for SLB | 477 |
| Adding or removing noncacheable expressions | 477 |
| Configuring redirection requests | 479 |
| Configuring filters for Web Cache Redirection | 482 |
| Monitoring WCR statistics | 483 |
| Configuring HTTP header-based Web Cache Redirection | 484 |
| Configuring HTTP header-based redirection requests | 485 |
| Exclusionary string matching for real servers | 485 |
| Configuring exclusionary URL string matching | 486 |
| Regular expression matching | 486 |
| Standard regular expression characters | 487 |
| Rules for using standard regular expressions | 487 |
| Configuring regular expressions | 488 |
| Content precedence lookup | 488 |
| Using the operators Or and And | 489 |
| Configuring content precedence for a service | 491 |
| Assigning multiple strings | 491 |
| Configuring a Layer 7 deny filter | 492 |
| | |
| Chapter 19 | |
| Bandwidth management | 493 |
| About bandwidth management | 493 |
| Statistics and history | 494 |
| Policies | 494 |
| Data pacing | 495 |
| Contracts | 496 |
| Contract classifications | 496 |
| Contract precedence | 497 |

| | |
|--|------------|
| URL-based BWM | 497 |
| HTTP header-based BWM | 498 |
| Cookie-based BWM | 498 |
| Frame Discard | 498 |
| Packet coloring (TOS bits) for burst limit | 499 |
| Using Virtual Matrix Architecture with BWM | 499 |
| Configuring bandwidth management | 499 |
| Configuration prerequisites for BWM | 500 |
| Enabling bandwidth management | 500 |
| Configuring BWM policy | 502 |
| Configuring a BWM contract | 505 |
| Configuring BWM classification | 508 |
| Assigning a BWM contract to a virtual server | 509 |
| Sending BWM statistics to a user | 511 |
| SLB—URL-Based BWM fields | 514 |
| | |
| Chapter 20 | |
| Working with Statistics | 516 |
| About statistics and graphing | 516 |
| Working with data tables | 517 |
| Working with charts | 518 |
| WSM blade statistics | 518 |
| IP statistics | 519 |
| Received ICMP statistics | 521 |
| Sent ICMP statistics | 523 |
| RIP statistics | 525 |
| SNMP statistics | 526 |
| WebOS statistics | 529 |
| Statistics | 531 |
| MP CPU statistics | 533 |
| RADIUS account statistics | 535 |
| Bridge forwarding statistics | 536 |
| Forwarding statistics | 536 |
| Base port statistics | 537 |
| TP port statistics | 539 |

| | |
|--|-----|
| Port statistics | 540 |
| Port interface statistics | 540 |
| Ethernet statistics | 543 |
| Bridge statistics | 546 |
| Load balance statistics | 547 |
| Port CPU statistics | 549 |
| IP routing statistics | 551 |
| Route statistics | 551 |
| ARP statistics | 552 |
| Interface statistics | 553 |
| IP address statistics | 555 |
| TCP statistics | 556 |
| TCP connection statistics | 558 |
| UDP statistics | 560 |
| UDP local statistics | 562 |
| Virtual routing statistics | 563 |
| Virtual routing statistics | 563 |
| Virtual routing state data | 564 |
| Layer 4 statistics | 565 |
| Server load balance statistics | 566 |
| Real server port statistics | 566 |
| Real servers statistics | 567 |
| Groups statistics | 568 |
| Virtual servers statistics | 569 |
| Real servers state data | 571 |
| Maintenance statistics | 572 |
| DNS statistics | 573 |
| TCP Limit statistics | 574 |
| Filter statistics | 575 |
| Remote real servers data | 576 |
| URL statistics | 577 |
| Redirect statistics | 577 |
| URL load balance statistics | 578 |
| RTSP statistics | 580 |
| Wireless Application Protocol statistics | 581 |

| | |
|---|------------|
| Add session statistics | 581 |
| Delete session statistics | 582 |
| Bandwidth management statistics | 584 |
| Traffic contract statistics | 584 |
| Switch port traffic contract statistics | 585 |
| Index | 587 |

Figures

| | | |
|-----------|--|-----|
| Figure 1 | Device Manager window | 43 |
| Figure 2 | Open Device dialog box | 44 |
| Figure 3 | Parts of the Device Manager window | 46 |
| Figure 4 | WSM edit menus | 49 |
| Figure 5 | WSM graph menus | 53 |
| Figure 6 | Actions menu | 54 |
| Figure 7 | Help menu | 56 |
| Figure 8 | Chassis shortcut menu | 58 |
| Figure 9 | WSM shortcut menu | 59 |
| Figure 10 | Port shortcut menu | 60 |
| Figure 11 | Objects in a Passport 8000 Series switch device view | 63 |
| Figure 12 | WSM device view | 65 |
| Figure 13 | Editing a text field | 68 |
| Figure 14 | Editing a field from a selection list | 69 |
| Figure 15 | Adding a table filter | 74 |
| Figure 16 | Read only dialog box example | 77 |
| Figure 17 | WSM online Help | 79 |
| Figure 18 | WSM port layout and default configuration | 82 |
| Figure 19 | WSM device view | 82 |
| Figure 20 | Port—General tab | 84 |
| Figure 21 | Port tab | 86 |
| Figure 22 | Port—Filtering tab | 90 |
| Figure 23 | Port—Filtering dialog box with filters applied | 91 |
| Figure 24 | Select Filters for Port dialog box | 92 |
| Figure 25 | Filter Membership tab | 94 |
| Figure 26 | Trunks tab | 97 |
| Figure 27 | Insert trunks dialog box | 98 |
| Figure 28 | Jumbo frame VLANs | 104 |
| Figure 29 | Virtual LAN—VLAN Membership tab | 105 |

| | | |
|-----------|---|-----|
| Figure 30 | Virtual LAN dialog box | 106 |
| Figure 31 | Virtual LAN, Insert VLAN Membership dialog box | 107 |
| Figure 32 | MLT dialog box | 108 |
| Figure 33 | BDPU message format | 110 |
| Figure 34 | Example of multiple instances of Spanning Tree protocol | 111 |
| Figure 35 | Example of single spanning tree—VLAN 3 blocked | 111 |
| Figure 36 | Example of multiple spanning tree groups | 112 |
| Figure 37 | Bridge—Spanning Tree Group tab | 115 |
| Figure 38 | Select VLANs dialog box | 115 |
| Figure 39 | Bridge—General tab | 118 |
| Figure 40 | Bridge—Spanning Tree tab | 119 |
| Figure 41 | Port—Spanning Tree tab | 123 |
| Figure 42 | Bridge—STG Port tab | 125 |
| Figure 43 | Bridge—STP Port tab | 127 |
| Figure 44 | Interfaces tab | 130 |
| Figure 45 | IP Routing, Insert Interfaces dialog box | 131 |
| Figure 46 | IP Routing—General tab | 133 |
| Figure 47 | Gateways tab | 133 |
| Figure 48 | IP Routing, Insert Gateways dialog box | 134 |
| Figure 49 | Routes tab | 138 |
| Figure 50 | DNS tab | 140 |
| Figure 51 | Local tab | 142 |
| Figure 52 | IP Routing, Insert Local tab | 142 |
| Figure 53 | RIP tab | 144 |
| Figure 54 | IP Routing—Ports tab | 147 |
| Figure 55 | Static Routes tab | 148 |
| Figure 56 | IP Routing, Insert Static Routes dialog box | 149 |
| Figure 57 | Virtual Routing—General tab | 152 |
| Figure 58 | Routers tab | 154 |
| Figure 59 | Virtual Routing, Insert Routers dialog box | 155 |
| Figure 60 | Virtual Routing—Interfaces tab | 160 |
| Figure 61 | Virtual Routing—Insert Interfaces dialog box | 160 |
| Figure 62 | Groups tab | 162 |
| Figure 63 | Virtual Routing, Insert Groups dialog box | 163 |
| Figure 64 | Web Switching Module—Syslog tab | 174 |

| | | |
|-----------|--|-----|
| Figure 65 | Syslog Trap tab | 175 |
| Figure 66 | Trap Hosts tab | 177 |
| Figure 67 | Trap Hosts tab | 178 |
| Figure 68 | Image tab | 180 |
| Figure 69 | Port Mirroring | 183 |
| Figure 70 | Mirror tab | 184 |
| Figure 71 | Insert Mirror dialog box | 184 |
| Figure 72 | General tab | 186 |
| Figure 73 | Layer 4 Switching Filters—Filters tab | 193 |
| Figure 74 | Layer 4 Switching Filters—Insert Filters dialog box (left side) | 194 |
| Figure 75 | Layer 4 Switching Filters—Insert Filters dialog box (right side) | 195 |
| Figure 76 | Traditional versus server load-balanced networks | 204 |
| Figure 77 | Web hosting without server load balancing | 214 |
| Figure 78 | Web hosting with server load balancing | 215 |
| Figure 79 | Server Load Balance—Real Servers tab | 217 |
| Figure 80 | Server Load Balance—Insert Real Servers dialog box | 218 |
| Figure 81 | Server Load Balance—General tab | 223 |
| Figure 82 | Server Load Balance—Real Groups tab | 226 |
| Figure 83 | Insert Real Groups dialog box | 226 |
| Figure 84 | Select Real Servers to include in Group dialog box | 227 |
| Figure 85 | Server Load Balance—Virtual Servers tab | 231 |
| Figure 86 | Insert Virtual Servers dialog box | 232 |
| Figure 87 | Virtual Server Services dialog box | 235 |
| Figure 88 | Virtual Server Services, Insert Services dialog box | 236 |
| Figure 89 | Server Load Balance—Ports tab | 240 |
| Figure 90 | Insert Ports dialog box | 241 |
| Figure 91 | SLB—disabling server processing on a port | 247 |
| Figure 92 | Server Load Balance—Virtual Servers tab | 249 |
| Figure 93 | Virtual server services port mapping | 249 |
| Figure 94 | Basic virtual port to real port mapping configuration | 251 |
| Figure 95 | SLB multiple port mapping | 252 |
| Figure 96 | Server Load Balance—Real Server Port tab | 253 |
| Figure 97 | Server Load Balance, Insert Real Server Port dialog box | 254 |
| Figure 98 | Direct server return | 256 |
| Figure 99 | SLB—enable source MAC substitution for real server | 256 |

| | | |
|------------|--|-----|
| Figure 100 | SLB—Virtual Servers tab | 257 |
| Figure 101 | Virtual Server Services—enable MAC address substitution | 257 |
| Figure 102 | Server Load Balance—enable Direct Access Mode | 259 |
| Figure 103 | Mapped and non-mapped server access | 261 |
| Figure 104 | DoS SYN attacks without delayed binding | 263 |
| Figure 105 | Repelling DoS SYN attacks with delayed binding | 264 |
| Figure 106 | Server Load Balance—Virtual Servers tab | 265 |
| Figure 107 | Virtual Server Services dialog box—delayed binding | 266 |
| Figure 108 | Server Load Balance—SYN Attack tab | 267 |
| Figure 109 | Configuring IP SLB, layer 3 only processing | 270 |
| Figure 110 | Configuring IP service for SLB | 271 |
| Figure 111 | Configuring DAM for FTP server load balancing | 274 |
| Figure 112 | Server Load Balance—Virtual Servers tab | 275 |
| Figure 113 | Configuring FTP server load balancing | 276 |
| Figure 114 | Layer 4 DNS load balancing | 277 |
| Figure 115 | Configuring the UDP DNS service for server load balancing | 280 |
| Figure 116 | Configuring the TCP DNS service for server load balancing | 282 |
| Figure 117 | Configuring RTSP, SLB—General tab | 287 |
| Figure 118 | Configuring L4 RTSP service for server load balancing | 289 |
| Figure 119 | Configuring L7 RTSP service for server load balancing | 291 |
| Figure 120 | Configuring WAP server load balancing, SLB—General tab | 295 |
| Figure 121 | Configuring the WAP service for server load balancing | 297 |
| Figure 122 | WAP—General tab | 298 |
| Figure 123 | Configuring WAP RADIUS snooping filter—Part 1 | 300 |
| Figure 124 | Configuring WAP RADIUS snooping filter—Part 2 | 301 |
| Figure 125 | Configuring real servers for IDS SLB | 305 |
| Figure 126 | Configuring real server groups for IDS SLB | 307 |
| Figure 127 | SLB Real Groups tab | 308 |
| Figure 128 | Adding real servers to IDS SLB real server group | 308 |
| Figure 129 | Configuring ports for IDS SLB | 310 |
| Figure 130 | Enabling Intrusion Detection Server Load Balancing | 312 |
| Figure 131 | Configuring WAN link real server for server load balancing | 315 |
| Figure 132 | Configuring WAN link real server group for SLB | 316 |
| Figure 133 | Adding real servers to WAN link server group | 317 |
| Figure 134 | WAN link filter configuration—left side | 319 |

| | | |
|------------|--|-----|
| Figure 135 | WAN link filter configuration—right side | 320 |
| Figure 136 | Saving WAN link SLB filter configuration | 321 |
| Figure 137 | Enabling VMA | 324 |
| Figure 138 | Firewall configuration without FWLB | 328 |
| Figure 139 | Basic FWLB topology | 329 |
| Figure 140 | Firewall load-balancing topology with DMZ | 331 |
| Figure 141 | Basic FWLB implementation | 332 |
| Figure 142 | Basic FWLB example configuration | 335 |
| Figure 143 | Basic network frame flow and operation | 359 |
| Figure 144 | VPN load balancing configuration example | 360 |
| Figure 145 | DNS resolution with global server load balancing | 373 |
| Figure 146 | Global Server Load Balance—Remote Sites tab | 376 |
| Figure 147 | Insert Remote Sites dialog box | 377 |
| Figure 148 | Configuring real servers for GSLB | 379 |
| Figure 149 | Configuring GSLB real server membership | 380 |
| Figure 150 | Configuring GSLB virtual servers | 381 |
| Figure 151 | Configuring GSLB server services | 381 |
| Figure 152 | Global SLB—General tab | 383 |
| Figure 153 | HTTP and non-HTTP redirects | 386 |
| Figure 154 | POP3 request fulfilled using IP Proxy | 387 |
| Figure 155 | Port configuration—IP proxy address for non-HTTP redirects | 389 |
| Figure 156 | Real server configuration—IP proxy for HTTP redirects | 389 |
| Figure 157 | Global Server Load Balance—Lookup tab | 391 |
| Figure 158 | Global Server Load Balance—Network Preferences tab | 391 |
| Figure 159 | Insert Network Preferences dialog box | 392 |
| Figure 160 | Enabling graceful server failure | 406 |
| Figure 161 | Configuring a health check for a real server group | 409 |
| Figure 162 | Modifying the health check interval and retries | 411 |
| Figure 163 | Enabling VIP health checking for a real server group | 414 |
| Figure 164 | Setting the RADIUS authentication string | 422 |
| Figure 165 | Configuring content for WSP health check | 425 |
| Figure 166 | Configuring the WSP service | 427 |
| Figure 167 | Configuring WTLS port to health check | 429 |
| Figure 168 | Changing the WSM LDAP version | 431 |
| Figure 169 | Understanding cookie-based session persistence | 436 |

| | | |
|------------|--|-----|
| Figure 170 | Session persistence—insert cookie mode | 439 |
| Figure 171 | Session persistence—passive cookie mode | 440 |
| Figure 172 | Session Persistence—rewrite cookie mode | 441 |
| Figure 173 | SSL session ID-based persistence | 450 |
| Figure 174 | URL-based server load balancing | 454 |
| Figure 175 | URL Parsing—Load Balance tab | 457 |
| Figure 176 | URL Parsing, Insert Load Balance dialog box | 457 |
| Figure 177 | Configuring a real server for URL-based SLB | 459 |
| Figure 178 | URL Path Membership dialog box | 459 |
| Figure 179 | Virtual Server Services dialog box | 460 |
| Figure 180 | Configuring a service for URL-based SLB | 461 |
| Figure 181 | Load balancing DNS queries | 470 |
| Figure 182 | URL-based Web Cache Redirection | 475 |
| Figure 183 | URL Parsing—Expressions tab | 478 |
| Figure 184 | URL Parsing, Insert Expression dialog box | 478 |
| Figure 185 | URL Parsing—Redirection tab | 480 |
| Figure 186 | Example: Content precedence lookup | 490 |
| Figure 187 | Example: Content precedence lookup with multiple strings | 491 |
| Figure 188 | Bandwidth management—General tab | 501 |
| Figure 189 | Bandwidth management—Traffic Policies tab | 502 |
| Figure 190 | Insert Traffic Policies dialog box | 503 |
| Figure 191 | Bandwidth management—Traffic Contracts tab | 506 |
| Figure 192 | Insert Traffic Contracts dialog box | 506 |
| Figure 193 | Server Load Balance—Virtual Servers tab | 510 |
| Figure 194 | Insert Virtual Servers dialog box | 510 |
| Figure 195 | Web Switching Module—General tab | 512 |
| Figure 196 | Bandwidth management—General tab | 513 |
| Figure 197 | Server Load Balance—URL-based BWM fields | 514 |
| Figure 198 | Statistics window toolbar | 517 |
| Figure 199 | Chart toolbar | 518 |
| Figure 200 | WSM Blade—IP tab | 519 |
| Figure 201 | WSM Blade—ICMP In tab | 522 |
| Figure 202 | WSM Blade—ICMP Out tab | 524 |
| Figure 203 | WSM Blade—RIP statistics tab | 526 |
| Figure 204 | WSM Blade—SNMP tab | 527 |

| | | |
|------------|--|-----|
| Figure 205 | WSM Blade—WebOS Stats tab | 530 |
| Figure 206 | WSM Blade—Stats tab | 532 |
| Figure 207 | WSM Blade—MP CPU tab | 534 |
| Figure 208 | WSM Blade—RADIUS Account tab | 535 |
| Figure 209 | Bridge forwarding DB—Forwarding tab | 537 |
| Figure 210 | Bridge Forwarding DB—Base Port tab | 538 |
| Figure 211 | Bridge Forwarding DB—TP Port tab | 539 |
| Figure 212 | Port—Interface tab | 541 |
| Figure 213 | Port—Ethernet tab | 544 |
| Figure 214 | Port—Bridge tab | 546 |
| Figure 215 | Port—Load Balance tab | 548 |
| Figure 216 | Port—CPU Statistics tab | 550 |
| Figure 217 | IP Routing—Route tab | 552 |
| Figure 218 | IP Routing—ARP tab | 553 |
| Figure 219 | IP Routing—Interface tab | 554 |
| Figure 220 | IP Routing—IP Address tab | 555 |
| Figure 221 | IP Routing—TCP tab | 557 |
| Figure 222 | IP Routing—TCP Connection tab | 559 |
| Figure 223 | IP Routing—UDP tab | 561 |
| Figure 224 | IP Routing—UDP Local tab | 562 |
| Figure 225 | Virtual Routing—Statistics tab | 564 |
| Figure 226 | Virtual Routing—State tab fields | 565 |
| Figure 227 | Server Load Balancing—Real Server Ports tab | 566 |
| Figure 228 | Server Load Balancing—Real Servers tab | 567 |
| Figure 229 | Server Load Balancing—Groups tab | 568 |
| Figure 230 | Server Load Balancing—Virtual Servers tab | 570 |
| Figure 231 | Server Load Balancing—Real Servers State tab | 571 |
| Figure 232 | Server Load Balancing—Maintenance tab | 572 |
| Figure 233 | Server Load Balancing—DNS tab | 573 |
| Figure 234 | Server Load Balancing—TCP Limit tab | 575 |
| Figure 235 | Graph—L4 Filters dialog box | 576 |
| Figure 236 | Global Server Load Balancing—Remote Real Servers tab | 577 |
| Figure 237 | URL Parsing—Redirect tab | 578 |
| Figure 238 | URL—Load Balance tab | 579 |
| Figure 239 | RTSP Statistics tab | 580 |

| | |
|---|-----|
| Figure 240 Add Sessions tab | 582 |
| Figure 241 Delete Sessions tab | 583 |
| Figure 242 Bandwidth Management—Traffic Contracts tab | 585 |
| Figure 243 Bandwidth Management—Switch Port Traffic Contracts tab | 586 |

Tables

| | | |
|----------|---|-----|
| Table 1 | SNMP community string default values | 45 |
| Table 2 | Device Manager menu bar description | 47 |
| Table 3 | WSM Card Edit submenu selections | 50 |
| Table 4 | WSM Card L4 Switching submenu selections | 52 |
| Table 5 | Actions menu selections | 54 |
| Table 6 | Help menu selections | 56 |
| Table 7 | Chassis shortcut menu options | 58 |
| Table 8 | WSM shortcut menu | 59 |
| Table 9 | Port shortcut menu | 60 |
| Table 10 | Toolbar | 61 |
| Table 11 | Device Manager port color codes | 66 |
| Table 12 | Formats for entering IP and MAC addresses, and time | 69 |
| Table 13 | Device Manager buttons | 70 |
| Table 14 | WSM port descriptions | 81 |
| Table 15 | Port—General tab fields | 85 |
| Table 16 | Port tab fields | 87 |
| Table 17 | Port—Filtering tab fields | 90 |
| Table 18 | Port—Filter Membership tab fields | 93 |
| Table 19 | Initial MLT assignment compared with assignment after reset | 96 |
| Table 20 | Trunks configuration fields | 97 |
| Table 21 | VLAN restrictions for the Web Switching Module | 102 |
| Table 22 | Virtual LAN—VLAN Membership fields | 105 |
| Table 23 | Port, trunk group, VLANs, and Spanning tree relationships | 109 |
| Table 24 | Multiple Spanning Tree groups example—VLAN, WSM, STG | 112 |
| Table 25 | Root cost | 114 |
| Table 26 | Bridge STG tab fields | 116 |
| Table 27 | Select VLANs to Associate with STG dialog box fields | 117 |
| Table 28 | Bridge—General tab fields | 120 |
| Table 29 | Bridge—Spanning Tree tab fields | 121 |

| | | |
|----------|---|-----|
| Table 30 | Port Spanning Tree tab fields | 124 |
| Table 31 | Bridge STG Port tab fields | 126 |
| Table 32 | Bridge STP Port tab fields | 127 |
| Table 33 | Interfaces fields | 132 |
| Table 34 | IP Routing—General tab fields | 135 |
| Table 35 | Gateways fields | 136 |
| Table 36 | Routes tab fields | 138 |
| Table 37 | DNS tab fields | 140 |
| Table 38 | Local Routing Cache Address Ranges | 141 |
| Table 39 | Local IP Configuration fields | 143 |
| Table 40 | RIP tab fields | 145 |
| Table 41 | IP Routing—Ports tab fields | 147 |
| Table 42 | Static route configuration fields | 150 |
| Table 43 | Virtual Routing—General tab fields | 153 |
| Table 44 | Routers fields | 157 |
| Table 45 | Virtual Routing—Interfaces fields | 161 |
| Table 46 | Groups fields | 164 |
| Table 47 | SNMP traps | 172 |
| Table 48 | Syslog tab | 174 |
| Table 49 | Syslog trap fields | 176 |
| Table 50 | Trap Hosts fields | 178 |
| Table 51 | Image tab fields | 180 |
| Table 52 | Port Mirror fields | 185 |
| Table 53 | General tab fields | 186 |
| Table 54 | Filtering tasks | 190 |
| Table 55 | Well-known protocol numbers | 190 |
| Table 56 | Well-known application ports | 191 |
| Table 57 | Layer 4 Switching Filters—Filters fields | 196 |
| Table 58 | Server load balancing methods | 205 |
| Table 59 | General steps for configuring server load balancing | 215 |
| Table 60 | Real Servers fields | 219 |
| Table 61 | SLB—General tab fields | 224 |
| Table 62 | Real server group fields | 228 |
| Table 63 | Select Real Servers to Include in Group dialog box | 230 |
| Table 64 | Virtual server fields | 233 |

| | | |
|----------|--|-----|
| Table 65 | Virtual Server Services tab fields | 237 |
| Table 66 | SLB—Ports fields | 242 |
| Table 67 | Peers fields | 244 |
| Table 68 | Synchronize tab fields | 244 |
| Table 69 | Real Server Port fields | 254 |
| Table 70 | SYN Attack tab fields | 268 |
| Table 71 | RTSP implementation comparison | 285 |
| Table 72 | WAP—General tab fields | 299 |
| Table 73 | Firewall load balancing methods | 329 |
| Table 74 | FWLB example—8648TX VLAN port settings | 336 |
| Table 75 | FWLB example—8648TX MLT settings | 337 |
| Table 76 | FWLB example—8648TX MLT settings | 337 |
| Table 77 | Public-side FWLB network—VLAN settings | 338 |
| Table 78 | Public-side FWLB network—port default VLAN settings | 338 |
| Table 79 | Public-side FWLB network—ports to remove from default VLANs | 339 |
| Table 80 | Public-side FWLB network—IP interface settings | 339 |
| Table 81 | Public-side FWLB network—real server settings | 340 |
| Table 82 | Public-side FWLB network—real server group settings | 341 |
| Table 83 | Public-side FWLB network—firewall filter settings | 341 |
| Table 84 | Public-side FWLB network—redirection filter settings | 342 |
| Table 85 | Public-side FWLB network—port filter settings | 342 |
| Table 86 | Public-side FWLB network—IP static routes settings | 343 |
| Table 87 | Private-side FWLB network—VLAN settings | 343 |
| Table 88 | Private-side FWLB network—port default VLAN settings | 344 |
| Table 89 | Private-side FWLB network—ports to remove from default VLANs | 344 |
| Table 90 | Private-side FWLB network—IP interface settings | 344 |
| Table 91 | Private-side FWLB network—real server settings | 345 |
| Table 92 | Private-side FWLB network—real server group settings | 346 |
| Table 93 | Private-side FWLB network—port SLB settings | 346 |
| Table 94 | Private-side FWLB network—real server settings | 347 |
| Table 95 | Private-side FWLB network—real server group settings | 347 |
| Table 96 | Private-side FWLB network—virtual server settings | 348 |
| Table 97 | Private-side FWLB network—port SLB settings | 348 |
| Table 98 | Private-side FWLB network—firewall filter settings | 348 |
| Table 99 | Private-side FWLB network—firewall redirection filter settings | 349 |

| | | |
|-----------|---|-----|
| Table 100 | Private-side FWLB network—port filter settings | 350 |
| Table 101 | Private-side FWLB network—IP static routes settings | 350 |
| Table 102 | FWLB example—firewall 1 IP static routes settings | 351 |
| Table 103 | FWLB example—firewall 2 IP static routes settings | 352 |
| Table 104 | FWLB redirection filter settings | 354 |
| Table 105 | VPN example—private network IP interface settings | 361 |
| Table 106 | VPN example—private network IP static routes settings | 362 |
| Table 107 | VPN example—private network virtual router settings | 363 |
| Table 108 | VPN example—private network real server settings | 364 |
| Table 109 | VPN example—private network real server group settings | 365 |
| Table 110 | VPN example—public network IP interface settings | 366 |
| Table 111 | VPN example—public network IP static routes settings | 366 |
| Table 112 | VPN example—public network virtual router settings | 367 |
| Table 113 | VPN example—public network real server settings | 368 |
| Table 114 | VPN example—public network real server group settings | 369 |
| Table 115 | VPN example—public network filter settings | 369 |
| Table 116 | VPN example—public network filter settings | 370 |
| Table 117 | VPN example—public network filter settings | 371 |
| Table 118 | VPN example—configuring filter processing on ingress port | 371 |
| Table 119 | Global SLB—Remote Sites fields | 378 |
| Table 120 | Global SLB—General tab fields | 384 |
| Table 121 | HTTP versus non-HTTP redirects | 387 |
| Table 122 | Global SLB—Lookup tab fields | 393 |
| Table 123 | Global SLB—Network Preferences fields | 393 |
| Table 124 | Web Cache Redirection filter settings | 397 |
| Table 125 | Default filter settings for noncached traffic | 398 |
| Table 126 | RTSP Web Cache Redirection filter settings | 399 |
| Table 127 | Default filter settings for noncached traffic | 400 |
| Table 128 | Filter settings to exclude redirection of noncacheable site | 402 |
| Table 129 | Types of real server group health checks | 407 |
| Table 130 | Real server group settings for UDP-based DNShealth checking | 420 |
| Table 131 | Real server group settings for IMAP health checks | 420 |
| Table 132 | Real server group settings for RADIUS health checking | 421 |
| Table 133 | Designating the cookie session ID in an HTTP header | 437 |
| Table 134 | Modes of operation for cookie-based session persistence | 438 |

| | | |
|-----------|---|-----|
| Table 135 | Virtual server settings for HTTP or URI header cookie location | 444 |
| Table 136 | Service 80 settings for HTTP or URI header cookie location | 444 |
| Table 137 | Virtual server settings for parsing the cookie | 445 |
| Table 138 | Service 80 settings for parsing the cookie | 445 |
| Table 139 | Virtual server settings for passive cookie mode | 446 |
| Table 140 | Service 80 settings for passive cookie mode | 446 |
| Table 141 | Virtual server settings for rewrite cookie mode | 447 |
| Table 142 | Service 80 settings for rewrite cookie mode | 448 |
| Table 143 | URL string formats | 455 |
| Table 144 | URL Parsing—Load balance fields | 458 |
| Table 145 | Request disposition in cookie-based preferential load balancing | 465 |
| Table 146 | Examples of user categories for load balancing requests | 466 |
| Table 147 | SLB—RTSP tab fields | 472 |
| Table 148 | Examples of string expressions | 474 |
| Table 149 | Types of WCR NAT | 476 |
| Table 150 | URL Parsing—Expressions fields | 479 |
| Table 151 | URL Parsing—Redirection tab fields | 481 |
| Table 152 | URL-based Web Cache Redirection filter settings | 482 |
| Table 153 | Default filter settings for noncached traffic | 483 |
| Table 154 | HTTP-based Web Cache Redirection filter settings | 484 |
| Table 155 | Standard regular expression characters | 487 |
| Table 156 | Real server content | 492 |
| Table 157 | BWM policy granularity limits | 495 |
| Table 158 | Bandwidth management classifications | 496 |
| Table 159 | Bandwidth Management—General tab fields | 501 |
| Table 160 | Traffic Policies fields | 504 |
| Table 161 | Traffic Contracts fields | 507 |
| Table 162 | Configuring bandwidth management classification | 508 |
| Table 163 | SLB—URL-Based BWM fields | 514 |
| Table 164 | WSM Blade—IP tab fields | 520 |
| Table 165 | WSM Blade—ICMP In tab fields | 522 |
| Table 166 | WSM Blade—ICMP Out tab fields | 524 |
| Table 167 | WSM Blade—RIP tab fields | 526 |
| Table 168 | WSM Blade—SNMP tab fields | 528 |
| Table 169 | WSM Blade—WebOS Stats tab fields | 530 |

| | | |
|-----------|---|-----|
| Table 170 | WSM Blade—Stats tab fields | 532 |
| Table 171 | WSM Blade—MP CPU Stats fields | 534 |
| Table 172 | WSM Blade—Radius account statistics | 536 |
| Table 173 | Bridge Forwarding DB—Forwarding tab | 537 |
| Table 174 | Bridge Forwarding DB—Base Port tab | 538 |
| Table 175 | Bridge Forwarding DB—TP Port tab | 540 |
| Table 176 | Port—Interface tab fields | 541 |
| Table 177 | Port—Ethernet tab fields | 544 |
| Table 178 | Port—Bridge tab fields | 547 |
| Table 179 | Port—Load balance tab fields | 548 |
| Table 180 | CPU Statistics tab fields | 550 |
| Table 181 | IP Routing—Route tab fields | 552 |
| Table 182 | IP Routing—ARP tab fields | 553 |
| Table 183 | IP Routing—Interface tab fields | 554 |
| Table 184 | IP routing—IP Address tab fields | 556 |
| Table 185 | TCP tab fields | 557 |
| Table 186 | IP Routing—TCP Connection tab fields | 560 |
| Table 187 | IP Routing—UDP tab fields | 561 |
| Table 188 | IP Routing—UDP Local tab fields | 563 |
| Table 189 | Virtual Routing—Statistics tab fields | 564 |
| Table 190 | Virtual—Routing State tab fields | 565 |
| Table 191 | SLB—Real Server Ports tab fields | 567 |
| Table 192 | SLB—Real Servers tab fields | 568 |
| Table 193 | SLB—Groups tab fields | 569 |
| Table 194 | SLB—Virtual Servers tab fields | 570 |
| Table 195 | SLB—Real Servers State tab fields | 571 |
| Table 196 | SLB—Maintenance tab fields | 573 |
| Table 197 | SLB—DNS tab fields | 574 |
| Table 198 | SLB—TCP Limit fields | 575 |
| Table 199 | L4 filters—Statistics tab | 576 |
| Table 200 | Global SLB—Remote Real Servers tab fields | 577 |
| Table 201 | URL Parsing—Redirect tab fields | 578 |
| Table 202 | URL—Load Balance tab fields | 579 |
| Table 203 | RTSP—Statistics tab fields | 580 |
| Table 204 | WAP—Add Sessions tab fields | 582 |

| | | |
|-----------|--|-----|
| Table 205 | WAP—Delete Sessions tab fields | 583 |
| Table 206 | BWM—Traffic Contracts tab fields | 585 |
| Table 207 | BWM—Switch Port Traffic Contracts tab fields | 586 |

Preface

Welcome to Device Manager for the Web Switching Module. This guide provides information for using the Nortel Networks* Device Manager graphical user interface (GUI) to configure the Web Switching Module.

Before you begin

This guide is intended for network administrators with the following background:

- Basic knowledge of networks, Ethernet bridging, and IP routing
- Familiarity with networking concepts and terminology
- Basic knowledge of network topologies
- Experience with windowing systems or GUIs

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 <code>ping <ip_address></code> , you enter <code>ping 192.32.10.12</code> |
| bold Courier text | Indicates command names and options and text that you need to enter. Example: Use the info command. Example: Enter show ip {alerts routes} . |

| | |
|------------------------|---|
| braces ({}) | <p>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.</p> <p>Example: If the command syntax is <code>show ip {alerts routes}</code>, you must enter either <code>show ip alerts</code> or <code>show ip routes</code>, but not both.</p> |
| brackets ([]) | <p>Indicate optional elements in syntax descriptions. Do not type the brackets when entering the command.</p> <p>Example: If the command syntax is <code>show ip interfaces [-alerts]</code>, you can enter either <code>show ip interfaces</code> or <code>show ip interfaces -alerts</code>.</p> |
| ellipsis points (...) | <p>Indicate that you repeat the last element of the command as needed.</p> <p>Example: If the command syntax is <code>ethernet/2/1 [<parameter> <value>] ...</code>, you enter <code>ethernet/2/1</code> and as many <code>parameter-value</code> pairs as needed.</p> |
| <i>italic text</i> | <p>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.</p> <p>Example: If the command syntax is <code>show at <valid_route></code>, <code>valid_route</code> is one variable and you substitute one value for it.</p> |
| plain Courier text | <p>Indicates command syntax and system output, for example, prompts and system messages.</p> <p>Example: <code>Set Trap Monitor Filters</code></p> |
| separator (>) | <p>Shows menu paths.</p> <p>Example: <code>Protocols > IP</code> identifies the IP command on the Protocols menu.</p> |
| vertical line () | <p>Separates choices for command keywords and arguments. Enter only one of the choices. Do not type the vertical line when entering the command.</p> <p>Example: If the command syntax is <code>show ip {alerts routes}</code>, you enter either <code>show ip alerts</code> or <code>show ip routes</code>, but not both.</p> |

Related publications

For more information about the Web Switching Module, refer to the following publications:

- *Installing the Web Switching Module for the 8600 series switch*, Part number 314969-A

Provides a description of the WSM's hardware and software features, it's default configuration, installation and set up instructions, 8600 CLI commands that relate to the WSM, mapping of access levels between the 8600 and WSM, etc.

- *Web OS Switch Software 10.0 Application Guide*, part number 212777-A

Provides Web OS networking concepts and design guidelines using the command line interface. This document also presents specific Web OS configuration examples using the command line interface.

- *Web OS Switch Software 10.0 Command Reference*, Part number 212778-A

Provides a reference for the Web OS command line interface.

For more information about using Device Manager, see :

- *Installing and Using Device Manager*, Part number 316341-A
- *Configuring and Managing Security*, Part number 314724-C
- *Configuring Network Management*, Part number 314723-C

For a list of other related publications, see the release notes that accompany your software.

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.

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 Nortel Networks Technical Support. To obtain contact information online, go to the www.nortelnetworks.com/cgi-bin/comments/comments.cgi URL, then click on Technical Support.

From the Technical Support page, you can open a Customer Service Request online or find the telephone number for the nearest Technical Solutions Center. If you are not connected to the Internet, you can call 1-800-4NORTEL (1-800-466-7835) to learn the telephone number for the nearest Technical Solutions Center.

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://www.nortelnetworks.com/help/contact/erc/index.html> URL.

Chapter 1

Getting started with Device Manager

This section describes how to get started using Device Manager to configure and monitor the WSM. It includes the following topics.

- [“About Device Manager” on page 41](#)
- [“Starting Device Manager” on page 42](#)
- [“Opening a device” on page 43](#)
- [“Understanding the Device Manager window” on page 46](#)
- [“Using Device Manager dialog boxes” on page 66](#)
- [“Opening WSM online Help” on page 77](#)

About Device Manager

Device Manager lets you configure and monitor network devices using Simple Network Management Protocol (SNMP).



Note: The 8600 with WSM is supported in Device Manager version 5.5.x and above. Using any earlier version of Device Manager can adversely affect automatic configuration of the WSM.

Device Manager is a graphical user interface (GUI) between you and the supported network devices. It lets you remotely manage a device, just as though you were in the wiring closet.

Device Manager displays a real-time physical view of the front panel of a device. From the front panel view, you can view fault, configuration, and performance information for the device, a module, or a single port.



Note: For information about installing Device Manager software, see *Installing and Using Device Manager*, Part number 316341-A.

Starting Device Manager

To start Device Manager:

→ Do one of the following:

- From the Windows* Start menu, choose Programs > Nortel Networks Device Manager > Device Manager.
- From the UNIX* search path, verify the presence of the Device Manager installation directory; then type **JDM**

The Device Manager window ([Figure 1](#)) opens.



Note: On startup, Device Manager performs a DNS lookup for the machine on which it is running. If the DNS lookup is slow or fails, the initial Device Manager window may take up to 30 seconds to open.

Figure 1 Device Manager window

Opening a device

After starting Device Manager, you can open and view a device. First, you must specify the 8600 community strings, or access level, granted to the device.

For information about how the 8600 access levels map to the WSM access levels, see *Installing the Web Switching Module for the 8600 series switch*, Part number 314969-A.



Note: WSM passwords are synchronized from the 8600 host and cannot be set on the WSM side. Passwords are changed in the 8600 CLI. Passwords cannot be changed in Device Manager.

To open a device:

- 1 Do one of the following:
 - From the menu bar, choose Device > Open.

- From the toolbar, click the Open Device tool.

The Open Device dialog box (Figure 2) opens.

Figure 2 Open Device dialog box



- 2 In the Device Name field, type the DNS name or IP address of the device.
- 3 In the Read Community and Write Community fields, type the [8600 community strings](#).



Note: For read/write/all access, enter the read/write/all community string for both the Read Community and Write Community strings.

- 4 Click Open.

Device Manager determines what version of software the selected device is running, and the device view (Figure 3 on page 46) opens.

Table 1 shows the Device Manager default 8600 community strings.

Table 1 SNMP community string default values

| Access level | Description |
|--------------------|-------------|
| read-only | Public |
| Layer 1 read/write | Private |
| Layer 2 read/write | Private |
| Layer 3 read/write | Private |
| read/write | Private |
| read/write/all | Secret |

Troubleshooting the opening of a device

If a device does not open, Device Manager displays a timeout message.

To troubleshoot opening a device:

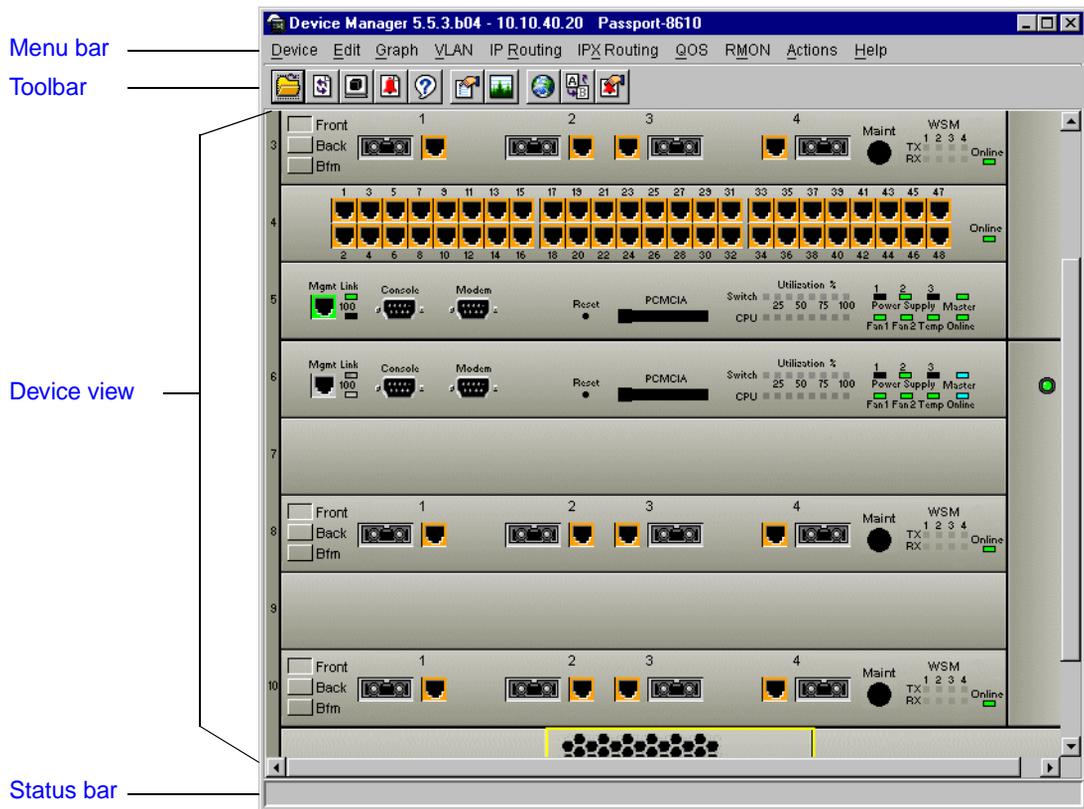
- In slower networks, increase the retransmission retries and timeout interval.
- In the Open Device dialog box, make sure you entered the correct read and write community strings.
- If you cannot reach the switch through an IP address or the management station cannot communicate with the switch, check the following:
 - Is the switch connected to the network?
 - Is the switch turned on?
 - Does the switch have an incorrect IP address?
 - Is the incorrect IP address specified in the Open Device field in Device Manager?
 - Is the network misconfigured?

Understanding the Device Manager window

The Device Manager window (Figure 3) has the following parts.

- “Menu bar” on page 47
- “Toolbar” on page 61
- “Device view” on page 62
- “Status bar” on page 66

Figure 3 Parts of the Device Manager window



Menu bar

Menus ([Figure 3 on page 46](#)) give you access to commands for managing or monitoring devices. You can access the commands for managing the WSM from the Edit, Graph, and Actions menus.

[Table 2](#) describes the menu bar.

Table 2 Device Manager menu bar description

| Menu | Description |
|-----------------------|---|
| Device | The Device menu lets you open a device, refresh the device view, and set polling and SNMP properties. This menu also allows you to open and view the Trap Log and Log. |
| Edit | The Edit menu lets you view parameters for the chassis or for selected objects. The object can be a WSM card, fan, MDA, port, power supply or any other object. This menu also lets you set security parameters, run diagnostic tests, and select all objects in the device. For more information, see “Accessing the Edit > WSM Card submenu” on page 48 . |
| Graph | The Graph menu lets you view Device Manager statistics and produce graphs of the chassis, port, or WSM card statistics. See “Working with Statistics” on page 516 . |
| VLAN | The VLAN menu lets you view information about the 8600 VLANs, spanning tree groups (STGs), MultiLink Trunks, MAC Learning, and SVLAN. |
| IP Routing | The IP Routing menu lets you set up IP routing functions for the switch, including OSPF, RIP, VRRP, BGP, Multicast, IGMP, DVMRP, PIM, PGM, DHCP, UDP forwarding, filters, and policies. |
| IPX Routing | The IPX Routing menu lets you set up IPX routing functions, including RIP, SAP, and policies. |
| QOS | The QOS menu lets you set up and view QoS filters and profiles. |
| RMON | The RMON menu lets you set up RMON alarms and view the alarm log and history log. This menu also allows you to enable or disable RMON history or statistics on all ports. |

Table 2 Device Manager menu bar description (continued)

| Menu | Description |
|-------------------------|---|
| Actions | The Actions menu provides access to the 8600 Web management interface, and lets you update configurations. See “Accessing the Actions menu” on page 54 . |
| Help | The Help menu gives you access to help for the 8600 series switch, Nortel Networks online customer support, a legend of port status color codes, and the running version of Device Manager. For more information, see “Accessing the Help menu” on page 55 . Note: This Help menu does not provide access to WSM online help. To access WSM online Help, see “Opening WSM online Help” on page 77 . |

Accessing the Edit > WSM Card submenu

Commands for managing the WSM ([Figure 4](#)) are accessed from the Edit menu.

To access the WSM Card Edit submenu:

- 1 From the device view, select a Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the menu bar, select Edit > WSM Card.
The [WSM Card Edit submenu](#) ([Figure 4](#)) opens.

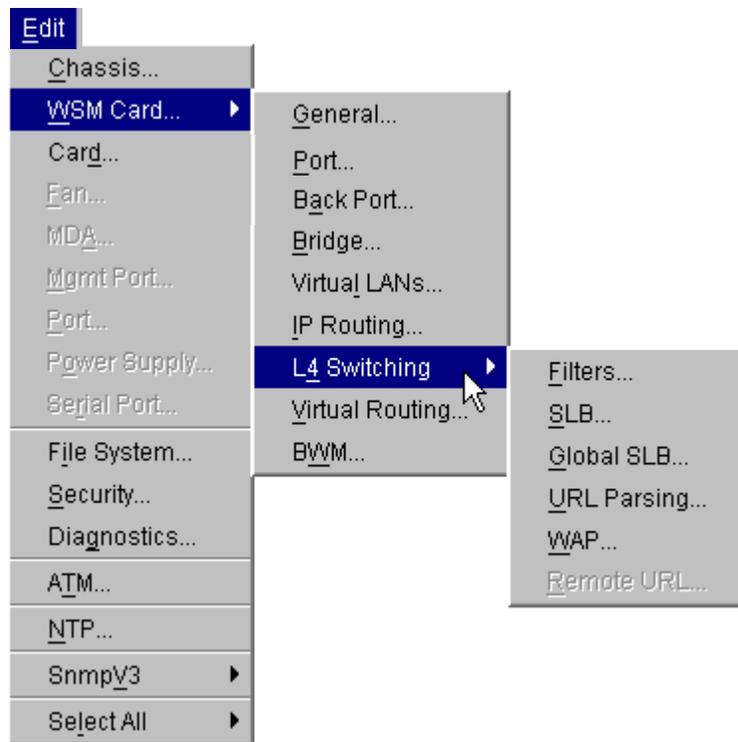
Figure 4 WSM edit menus

Table 3 describes the selections on the Edit submenu.

Table 3 WSM Card Edit submenu selections

| Menu selection | Description |
|-------------------|---|
| General | Opens the Web Switching Module dialog box with the following tabs: <ul style="list-style-type: none"> • “General tab fields” on page 186 • “Syslog tab” on page 174 • “Image tab fields” on page 180 • “Trap Hosts fields” on page 178 • “Syslog trap fields” on page 176 • “Trunks configuration fields” on page 97 • “Port Mirror fields” on page 185 |
| Port Back Port | Opens the Web Switching Module Port dialog box with the following tabs: <ul style="list-style-type: none"> • “Port—General tab fields” on page 85 • “Port tab fields” on page 87 • “Port Spanning Tree tab fields” on page 124 • “Port—Filtering tab fields” on page 90 |
| Bridge | Opens the Web Switching Module Bridge dialog box with the following tabs: <ul style="list-style-type: none"> • “Bridge—General tab fields” on page 120 • “Bridge—Spanning Tree tab fields” on page 121 • “Bridge STG tab fields” on page 116 • “Bridge STG Port tab fields” on page 126 • “Bridge STP Port tab fields” on page 127 |
| Virtual LANs | Opens the Web Switching Module Virtual LAN dialog box with the following tab: <ul style="list-style-type: none"> • “Virtual LAN—VLAN Membership fields” on page 105 |
| IP Routing | Opens the Web Switching Module IP Routing dialog box with the following tabs: <ul style="list-style-type: none"> • “IP Routing—General tab fields” on page 135 • “Interfaces fields” on page 132 • “Routes tab fields” on page 138 • “Gateways fields” on page 136 • “DNS tab fields” on page 140 • “Local IP Configuration fields” on page 143 • “RIP tab fields” on page 145 • “IP Routing—Ports tab fields” on page 147 • “Static route configuration fields” on page 150 |

Table 3 WSM Card Edit submenu selections (continued)

| Menu selection | Description |
|-----------------|---|
| L4 Switching | Opens the Web Switching Module L4 Switching submenu . See “Accessing the Edit > WSM Card > L4 Switching submenu” on page 51 . |
| Virtual Routing | Opens the Web Switching Module Virtual Routing dialog box with the following tabs: <ul style="list-style-type: none"> • “Virtual Routing—General tab fields” on page 153 • “Routers fields” on page 157 • “Virtual Routing—Interfaces fields” on page 161 • “Groups fields” on page 164 |
| BWM | Opens the Web Switching Module Bandwidth Management dialog box with the following tabs: <ul style="list-style-type: none"> • “Bandwidth Management—General tab fields” on page 501 • “Traffic Policies fields” on page 504 • “Traffic Contracts fields” on page 507 |

Accessing the Edit > WSM Card > L4 Switching submenu

Commands for managing WSM L4 switching ([Figure 4 on page 49](#)) are accessed from the Device Manager Edit menu.

To access WSM Card L4 Switching submenus:

- 1 From the device view, select a Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the menu bar, select Edit > WSM Card > L4 Switching.

The [WSM Card Edit L4 Switching submenu](#) ([Figure 4 on page 49](#)) opens.

(Table 4) describes the WSM L4 switching submenu.

Table 4 WSM Card L4 Switching submenu selections

| Menu selection | Description |
|----------------|--|
| Filters | Opens the Web Switching Module Filters dialog box with the following tabs: <ul style="list-style-type: none"> • “Layer 4 Switching Filters—Filters fields” on page 196 • “SLB—URL-Based BWM fields” on page 514 |
| SLB | Opens the Web Switching Module Server Load Balance dialog box with the following tabs: <ul style="list-style-type: none"> • “SLB—General tab fields” on page 224 • “Real server group fields” on page 228 • “Select Real Servers to Include in Group dialog box” on page 230 • “SLB—Ports fields” on page 242 • “Peers fields” on page 244 • “SLB—RTSP tab fields” on page 472 • “Real Servers fields” on page 219 • “Synchronize tab fields” on page 244 • “Virtual server fields” on page 233 • “Virtual Server Services tab fields” on page 237 • “Real Server Port fields” on page 254 • “SYN Attack tab fields” on page 268 • “SLB—URL-Based BWM fields” on page 514 |
| Global SLB | Opens the Web Switching Module Global Server Load Balance dialog box with the following tabs: <ul style="list-style-type: none"> • “Global SLB—General tab fields” on page 384 • “Global SLB—Remote Sites fields” on page 378 • “Global SLB—Lookup tab fields” on page 393 • “Global SLB—Network Preferences fields” on page 393 |
| URL Parsing | Opens the Web Switching Module URL Parsing dialog box with the following tabs: <ul style="list-style-type: none"> • “URL Parsing—Expressions fields” on page 479 • “URL Parsing—Redirection tab fields” on page 481 • “URL Parsing—Load balance fields” on page 458 |
| WAP | Opens the Web Switching Module WAP dialog box with the following tab: <ul style="list-style-type: none"> • “WAP—General tab fields” on page 299 |

Accessing the Graph > WSM Card submenu

You can access commands for graphing WSM statistics (Figure 5) from the Device Manager Edit menu.

To access WSM Card L4 Switching submenus:

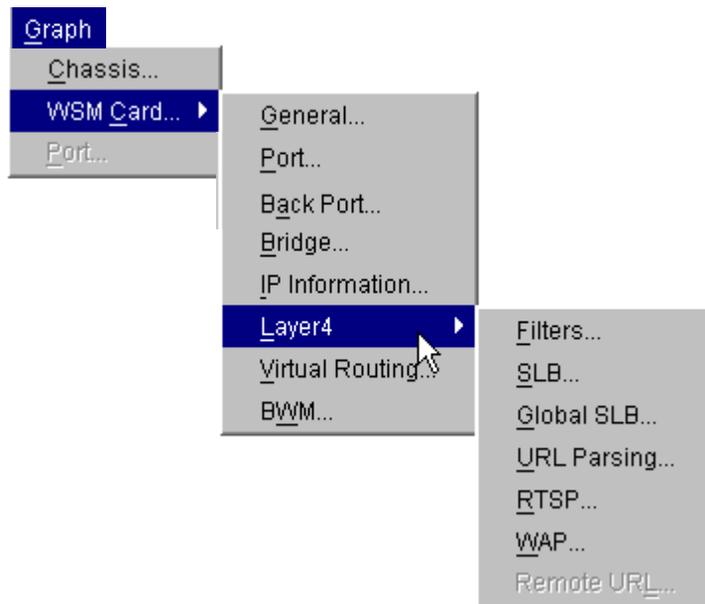
- 1 From the device view, select a Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the menu bar, select Graph > WSM Card.

The WSM Card Graph submenu (Figure 5) opens. For more information, see “Working with Statistics” on page 516

Figure 5 WSM graph menus



Accessing the Actions menu

The Actions menu gives you quick access to the 8600 Web management interface, and lets you update 8600 and WSM configurations.

To access the Actions menu for a WSM:

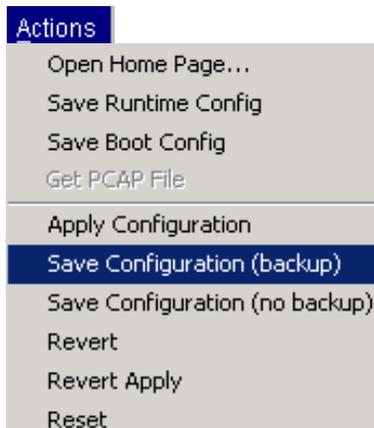
- 1 From the device view, select a Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the menu bar, click Actions.

The [Actions menu](#) (Figure 6) opens.

Figure 6 Actions menu



[Table 5](#) describes the selections in the Actions menu.

Table 5 Actions menu selections

| Menu selection | Description |
|---------------------|---|
| Open Home Page | Opens the 8600 Web management interface. |
| Save Runtime Config | Saves the 8600 runtime configuration. |
| Save Boot Config | Saves the 8600 boot configuration. |
| Get PCAP File | Permits access to the PCAP captured packets file. |

Table 5 Actions menu selections (continued)

| Menu selection | Description |
|--------------------------------|--|
| Apply Configuration | Applies pending WSM configuration changes. When you click <i>Apply Configuration</i> , Device Manager acknowledges that the Apply is complete and reminds you to save the updated configuration. |
| Save Configuration (backup) | Saves changes to the active and backup WSM configuration blocks. When you click <i>Save Configuration (backup)</i> , Device Manager asks you to confirm that you want to overwrite the active configuration. |
| Save Configuration (no backup) | Saves changes to the active WSM configuration block. When you click <i>Save Configuration (no backup)</i> , Device Manager asks you to confirm that you want to save the configuration. |
| Revert | Restores configuration parameters set since last <i>Apply</i> command. When you click <i>Revert</i> , Device Manager asks you to confirm that you want to revert unapplied changes. |
| Revert Apply | Restores configuration parameters set since last <i>Save Configuration</i> command. When you click <i>Revert Apply</i> , Device Manager asks you to confirm that you want to revert applied changes. |
| Reset | Begins a reset of the WSM, restoring default values to settings that have not been saved permanently to the configuration blocks. When you click <i>Reset</i> , Device Manager asks you to confirm that you want to begin a reset. |

Accessing the Help menu

The Help menu gives you access to 8600 online Help.



Note: The WSM and the 8600 series switch have separate online Help systems. For more information, see [“Opening WSM online Help” on page 77](#).

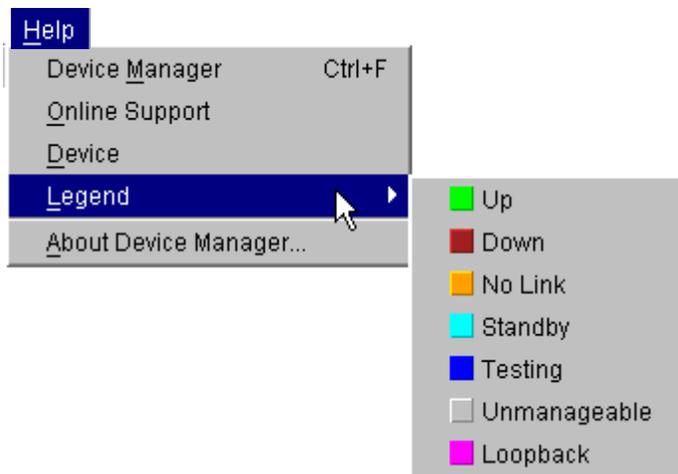
From the Help menu, you can also access the Nortel Networks online customer support Web site, view a legend of port status color codes, and view the running version of Device Manager.

To access the Help menu:

➔ From the menu bar, click Help.

The [Help menu \(Figure 7\)](#) opens.

Figure 7 Help menu



[Table 6](#) describes the selections in the Help menu.



Note: This Help menu does not give you access to WSM online Help. To access WSM online Help, see [“Opening WSM online Help” on page 77](#).

Table 6 Help menu selections

| Menu selection | Description |
|----------------|--|
| Device Manager | Opens a Device Manager Help system in your Web browser that describes elements of the 8600 Device Manager window. Note: This Help menu does not provide access to WSM online help. To access WSM online Help, see “Opening WSM online Help” on page 77 . |
| Online Support | Opens your Web browser to the Nortel Networks online support Web site. |

Table 6 Help menu selections (continued)

| Menu selection | Description |
|----------------------|---|
| Device | Opens online Help for the 8600 series switch in a Web browser. Note: This Help menu does not provide access to WSM online help. To access WSM online Help, see “Opening WSM online Help” on page 77 . |
| Legend | Displays a legend of the device view port colors . For more information see “Interpreting the status of LEDs and ports” on page 65 . |
| About Device Manager | Displays the running version of Device Manager and a list of devices it supports. |

Shortcuts

Objects in the device view such as the chassis, ports, and cards have shortcut menus. Shortcuts provide a faster path for editing objects and applying changes; however, you can access the same options through the menu bar or the toolbar. For information about shortcuts, see:

- [“Using the chassis shortcut,”](#) next
- [“Using the WSM port shortcut” on page 60](#)
- [“Using the WSM shortcut” on page 58](#)

Using the chassis shortcut

To display the chassis shortcut:

- ➔ Select the chassis and right click.

The [chassis shortcut menu](#) ([Figure 8](#)) opens.

Figure 8 Chassis shortcut menu

([Table 7](#)) describes the Chassis shortcut.

Table 7 Chassis shortcut menu options

| Option | Description |
|---------------------|---|
| Edit | Opens the chassis dialog box to edit chassis parameters. |
| Graph | Opens the Graph Chassis dialog box for graphing chassis statistics. |
| Save Runtime Config | Save any 8600 changes made as a run-time configuration. |
| Save Boot Config | Save any 8600 changes made as a boot configuration. |
| Reset Counters | Reset all the statistics counters for the switch. When you click <i>Reset Counters</i> , Device Manager asks you to confirm that you want to reset counters. This can cause monitoring applications to see traffic/error spike. |
| Hard Reset | Performs a hard reset of the 8600. |
| Soft Reset | Performs a soft reset of the 8600. |

Using the WSM shortcut

To display the WSM shortcut menu:

- ➔ Right-click a WSM Card.

The [WSM shortcut menu](#) ([Figure 9](#)) opens.

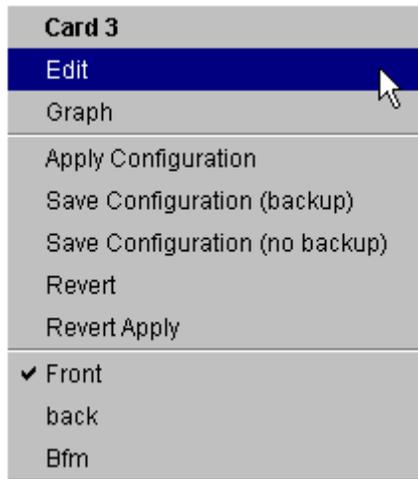
Figure 9 WSM shortcut menu

Table 8 describes the WSM shortcut.

Table 8 WSM shortcut menu

| Option | Description | Menu equivalent |
|--------------------------------|--|--|
| Edit | Opens the Web Switching Module dialog box. This is the menu bar equivalent of choosing | Edit > WSM Card > General |
| Graph | Opens the Graph WSM Blade dialog box. | Graph > WSM Card > General |
| Apply Configuration | Applies pending configuration changes. | Actions > Apply Configuration |
| Save Configuration (backup) | Saves changes to the active and backup configuration blocks. | Actions > Save Configuration (backup) |
| Save Configuration (no backup) | Saves changes to the active configuration block. | Actions > Save Configuration (no backup) |
| Revert | Restores configuration parameters set since last "apply" command. | Actions > Revert |
| Revert Apply | Restores configuration parameters set since last Save Configuration command. | Actions > Revert Apply |

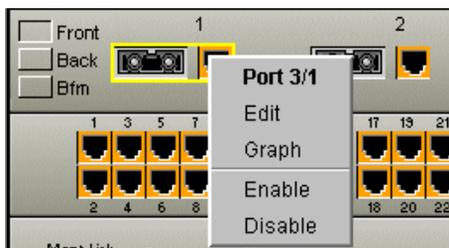
Table 8 WSM shortcut menu (continued)

| Option | Description | Menu equivalent |
|--------|--|---|
| Front | Changes the WSM's port orientation to Front. | Choosing Front on the WSM Card. See "About the device view" on page 82. |
| Back | Changes the WSM's port orientation to Back. | Choosing Back on the WSM Card. See "About the device view" on page 82. |
| BFM | Changes the WSM's port orientation to BFM. | Choosing BFM on the WSM Card. See "About the device view" on page 82. |

Using the WSM port shortcut

To display the WSM port shortcut menu:

- 1 Select a port orientation (Front or Back).
 - 2 Right-click the port you want to edit.
- The [port shortcut menu](#) ([Figure 10](#)) opens.

Figure 10 Port shortcut menu

[Table 9](#) describes the Port shortcut.

Table 9 Port shortcut menu

| Option | Description |
|---------|--------------------------------------|
| Edit | Opens the WSM Port dialog box. |
| Graph | Opens the Graph WSM Port dialog box. |
| Enable | Administratively brings a port up. |
| Disable | Administratively shuts a port down. |

Toolbar

The toolbar provides quick access to commands described in [Table 10](#).

Table 10 Toolbar

| Tool | Name | Description | Menu equivalent |
|---|-----------------------|---|--|
|  | Open Device | Opens the Open Device dialog box. From this dialog box you can ping, Telnet into, or open a device view. | Device > Open Device > Telnet |
|  | Refresh Device Status | Refreshes the device view information. | Device > Refresh Status |
|  | Telnet | Opens a Telnet session with the selected device. | Device > Telnet |
|  | Trap Log | Opens the trap log. | Device > Trap Log |
|  | Help | Opens online Help in a Web browser window. Note: WSM online help can only be accessed from the Edit or Graph menus after first selecting the WSM in the device view. | Help > Device Manager |
|  | Edit Selected | Displays configuration data windows for the selected chassis object. | Edit > Chassis Edit > WSM Card Edit > Card Edit > Fan Edit > MDA Edit > Mgmt Port Edit > Port Edit > Power Supply Edit > Serial Port |
|  | Graph Selected | Opens statistics and graphing windows. | Graph > Chassis Graph > WSM Card Graph > Port |

Table 10 Toolbar (continued)

| Tool | Name | Description | Menu equivalent |
|---|-------------------------|---|-------------------------------|
|  | Open Device's Home Page | Opens the Web management interface home page. | Actions > Open Home Page |
|  | Save Runtime Config | Saves the current run-time configuration. | Actions > Save Runtime Config |
|  | Alarm Manager | Opens the RMON Alarm Manager window. | Rmon > Alarm Manager |

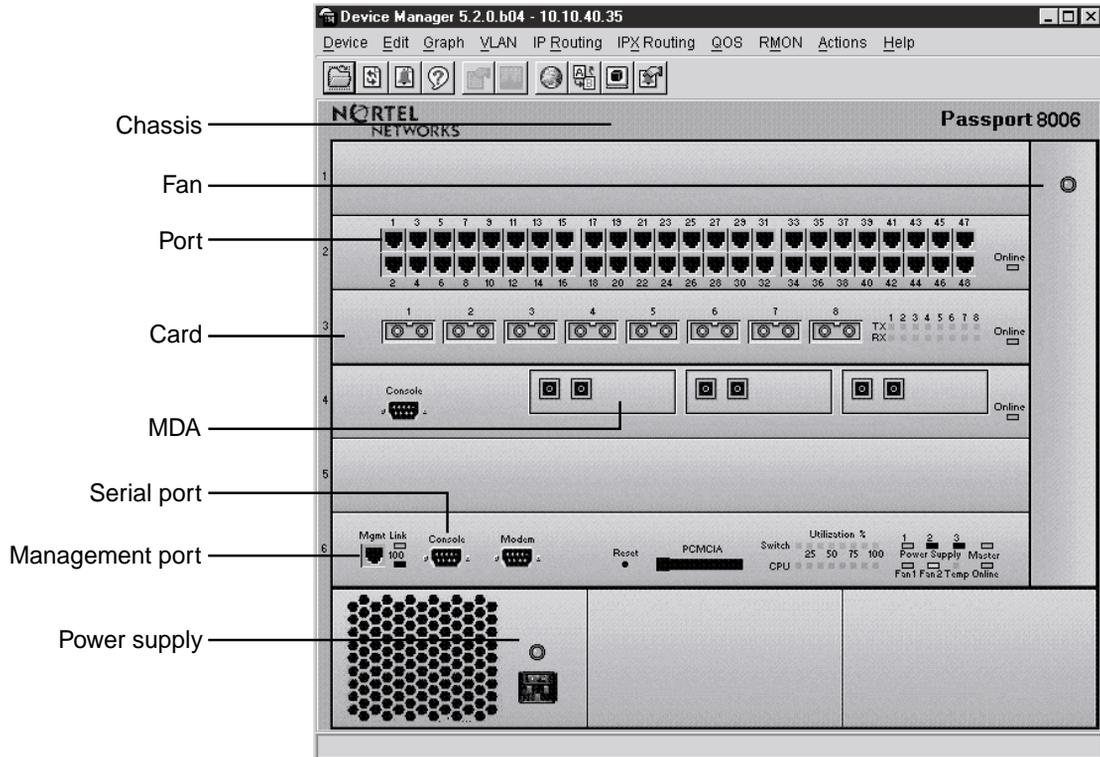
Device view

The device view shows the operating status of the various modules and ports in your hardware configuration. From the device view, you can also configure these devices and objects.

Selecting objects

In the device view ([Figure 11](#)), you can select the following types of objects:

- The entire chassis
- A card (module) or multiple cards
- A port or multiple ports
- A power supply
- A fan
- An MDA
- A management port
- A serial port

Figure 11 Objects in a Passport 8000 Series switch device view

To select a single object, click the edge of the object. The object is outlined in yellow, indicating that it is selected. Subsequent actions in Device Manager refer to the selected object.

To select a block of contiguous ports or modules:

➔ Drag to select the group of objects.

To select multiple ports or modules anywhere in the switch chassis:

➔ CTRL + click the objects anywhere in the device view.

The WSM device view

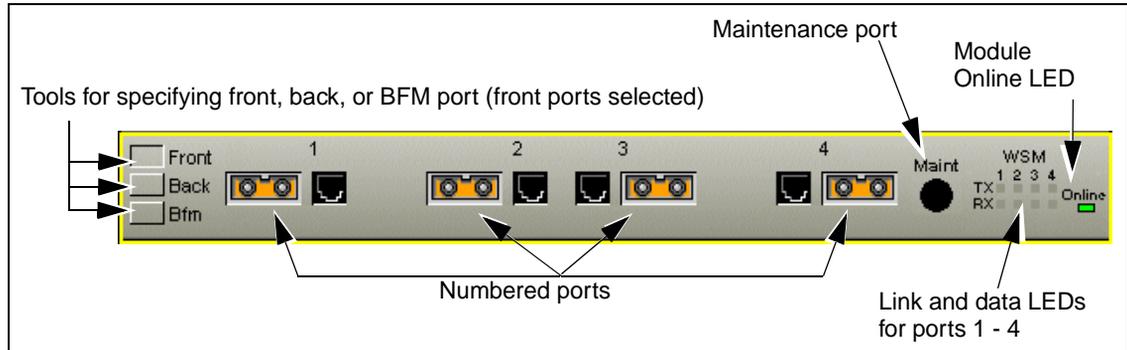
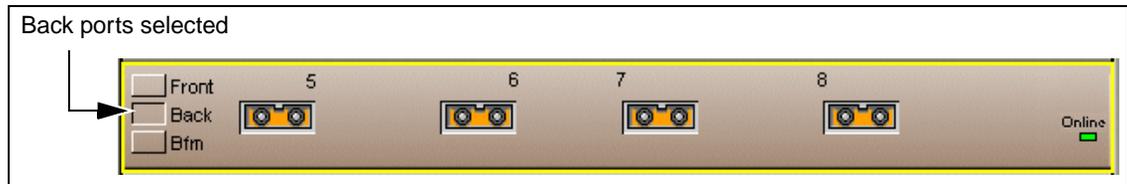
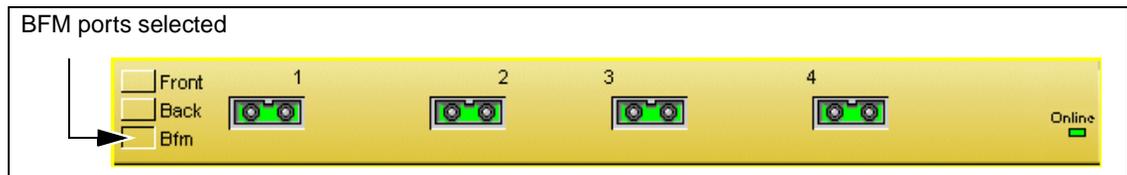
From the device view, you can see the position of the WSM card in the chassis (Figure 12), the numbered ports, the maintenance port, and module LEDs.

On the left side of the device view are tools for selecting the port orientation that you want to view or configure—front, back, or BFM.

To select a port orientation:

- ➔ Click a port orientation tool (Front, Back or BFM) on the left side of the module.

The device view changes to show a view of the front ports, back ports, or backplane fabric module ports.

Figure 12 WSM device view**Front ports view****Back ports view****Backplane fabric module ports view****Interpreting the status of LEDs and ports**

The device view uses conventions that mirror the actual switch in appearance. Module LEDs are in one of three states: on, off, or blinking. For a full description of LEDs and their states, refer to *Installing the Web Switching Module for the 8600 series switch*, Part number 314969-A.

The ports in the device view are color-coded to indicate port status, as shown in [Table 11](#).

Table 11 Device Manager port color codes

| Color | Description |
|------------|--|
| Green | Port is up and operating. |
| Red | Port has been manually disabled. |
| Orange | Port has no link. |
| Light Blue | Port is in standby mode. |
| Dark Blue | Port is being tested. |
| Grey | Port is not reachable by Device Manager. |
| Magenta | Port is in the Loopback mode. |

Viewing the port color code legend

To view a legend of port color codes:

➔ From the menu bar, choose Help > Legend.

A legend ([Figure 7 on page 56](#)) of port color codes opens.

Status bar

At the bottom of the Device Manager window is the status bar. This area displays error and informational messages from the software application. These messages are not related to the device being managed.

Using Device Manager dialog boxes

Many Device Manager dialog boxes contain editable fields for entering values. For example, a port may be set to be enabled or disabled. Other parameters are specified by ranges of user-determined values. For example, the value for an SMTP Host is a host name that you enter in the SMTP Host field.

Editable fields in Device Manager dialog boxes are displayed in white (see [Figure 13 on page 68](#)).

This section includes the following topics:

- [“Editing objects” on page 67](#)
- [“Using buttons in dialog boxes” on page 70](#)
- [“Applying and saving changes” on page 72](#)
- [“Refreshing the screen” on page 72](#)
- [“Filtering table content” on page 73](#)
- [“Browsing a read only dialog box” on page 76](#)

Editing objects

To edit an object:

➔ Do one of the following.

- Select an object, and, from the toolbar, click Edit *<object>*.
- From a shortcut menu, choose Edit.
- Double-click the object.

The edit dialog box opens for that object.

You can set an object’s parameters by editing text fields, choosing a setting from a selection list, or clicking a radio button or checkbox. The following sections provide additional editing information.

- [“Editing a text field,” next](#)
- [“Formats for entering numeric values” on page 69](#)
- [“Editing a field from a selection list” on page 69](#)

Editing a text field

Some text fields support user-specified values, such as the server name on the Real Server tab of the SLB dialog box.

To edit a text field:

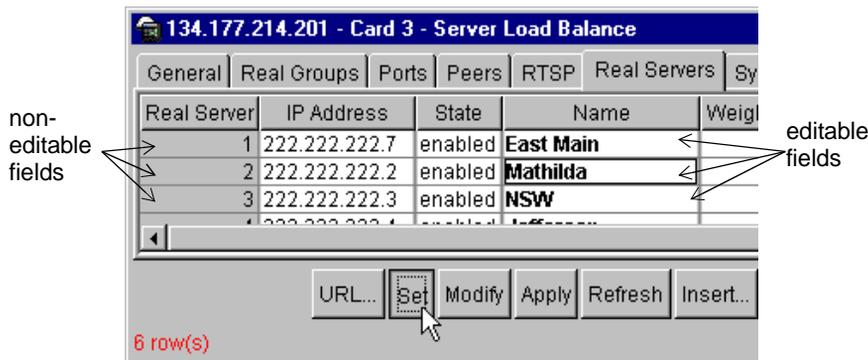
- 1 Double-click the field.

The field becomes editable (Figure 13).



Note: Fields that appear in white can be edited. Fields that are gray cannot be edited.

Figure 13 Editing a text field



- 2 Enter the new value, and click Set > Apply > Close.

The new value is applied to the WSM, and the dialog box closes.

- 3 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Formats for entering numeric values

When entering values for IP addresses, MAC addresses, or time, follow the guidelines in (Table 12).

Table 12 Formats for entering IP and MAC addresses, and time

| Value type | Format |
|-------------|--|
| IP address | Decimal (<xxx> . <xxx> . <xxx> . <xxx>) |
| MAC address | Hexadecimal (xx : xx : xx : xx : xx : xx) |
| Time | Based on the delta from the switch boot-up time. |

Editing a field from a selection list

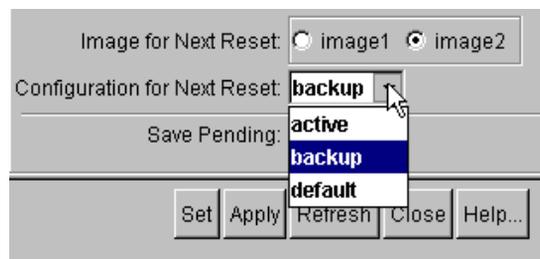
Many fields in Device Manager are editable from selection lists of the available values. For example, you can choose to run the active, backup, or default configuration.

To edit a field from a selection list:

- 1 Click the down arrow on the right-hand side of the field.

A selection list of available options opens (Figure 14) for choosing a value.

Figure 14 Editing a field from a selection list



- 2 Select a new value from the list.
- 3 Click Set > Apply > Close.

The changes are applied to the WSM, and the dialog box closes.

- 4 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Using buttons in dialog boxes

[Table 13](#) describes buttons that appear in Device Manager dialog boxes and tabs. Not all buttons appear in all dialog boxes.

Table 13 Device Manager buttons

| Button | Description |
|-----------------------|---|
| Apply | Applies the changes you have entered to the WSM. The button is grayed out until you change a parameter. Changes are displayed in bold . See “Applying and saving changes” on page 72 . |
| Browse | Opens a read-only dialog box for browsing , usually providing a selection list of available items. See “Browsing a read only dialog box” on page 76 . |
| Close | Closes the tab or dialog box and disregards any changes you have made to fields. |
| Copy | Copies selected items to your computer's memory clipboard. |
| Default Configuration | From the Bridge STG tab, reverts parameters to the default values. |
| Delete | Deletes a selected entry. |
| Export | In a Graph dialog box, saves the current table in ASCII format in a file you specify. The table contains tabs so you can then import this file into a text editor or spreadsheet for further analysis. |
| Export data | Allows you to copy data to external media. |
| Filter... | Opens a dialog box to filter table data and display a defined subset of the total content. See “Filtering table content” on page 73 . |
| Graph | Graphs selected data. |
| Help | Opens the online Help in a browser window, to a topic describing the specific dialog box. |
| Insert | Opens a dialog box to create a new entry for a table; then from the dialog box, inserts the new entry in the table. |
| Modify | Opens an Insert dialog box from which you can edit fields. |

Table 13 Device Manager buttons (continued)

| Button | Description |
|------------------|---|
| Paste | Will paste the contents of your computer's clipboard. |
| Print | From a Graph dialog box, prints the current table. |
| Print Table | Prints the contents of any table that is displayed. |
| Refresh | Refreshes the information in the window. Every time you click on Refresh, new information is polled from the WSM and displayed. |
| Reset changes | Reverts any configuration values you have changed back to their original value. |
| Resize Columns | Resizes table columns to fit the data in them. |
| Send BWM History | From the WSM Bandwidth Management General tab, sends the contents of the history buffer (BWM statistics) to a user when <i>SMTP User Name</i> is also configured. See "Sending BWM statistics to a user" on page 511 . |
| Servers... | From the WSM Real Groups dialog box, opens the <i>Servers to Include in Group</i> dialog box from which you can select the servers to add to a real server group. |
| Services... | From the WSM Virtual Server dialog box, opens the <i>Virtual Server Services</i> dialog box from which you modify or add services for a particular virtual server. |
| Set | Records the change made in a dialog box, but does not apply it to the WSM. The button is grayed out until you change a parameter. In order for the changes to be used by the WSM, you must also click Apply. See "Applying and saving changes" on page 72 . |
| Stop | Stops the current action (polling). |
| URL... | From the WSM Real Server dialog box, opens the <i>URL Path Membership</i> dialog box from which you select the URL paths to apply to a real server. |
| VLANs | From the WSM Bridge STG tab, opens the <i>VLANs to Associate to STG</i> dialog box from which you select VLAN membership for a particular STG. |

Applying and saving changes

When you change a value in a Device Manager field, the change appears in **bold**. This change has not yet been recorded or applied. To record the change in Device Manager, click Set. This records the pending change, but does not apply it to the WSM. To apply the change so that it is used by the WSM, you must also click Apply.



Caution: Changes you make in Device Manager are not used by the WSM until you click Apply. To make the changes permanent, from the Device Manager menu bar, click Actions > Save Configuration.

- To record a change in Device Manager, click Set.
- To apply the change so that it is used by the WSM, click Apply.
- To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Refreshing the screen

Most tabs and dialog boxes contain a Refresh button. If you decide to undo the changes made while working in a dialog box, use Refresh to return a field(s) to its previous setting(s).



Note: You must use Refresh before clicking Set to undo current changes.

In Windows and UNIX environments, changed values are displayed in bold until you click either Set or Refresh.

Filtering table content

Some Device Manager tables display a large amount of information. To reduce the amount of information displayed, you can filter these tables so that they only display a defined subset of the total content. In the Filter dialog box, you define the subset by creating a criteria statement of a variable type, value type, and operator. This statement is used to filter the table content, selecting the information to display.

Operators used in criteria statements include:

- CONTAINS
- IS GREATER THAN
- IS LESS THAN
- EQUALS
- DOES NOT EQUAL

The Filter dialog box supports the following:

- [“Adding a table filter” on page 74](#)
- [“Removing a table filter” on page 75](#)
- [“Removing all table filters” on page 75](#)
- [“Replacing a table filter” on page 76](#)

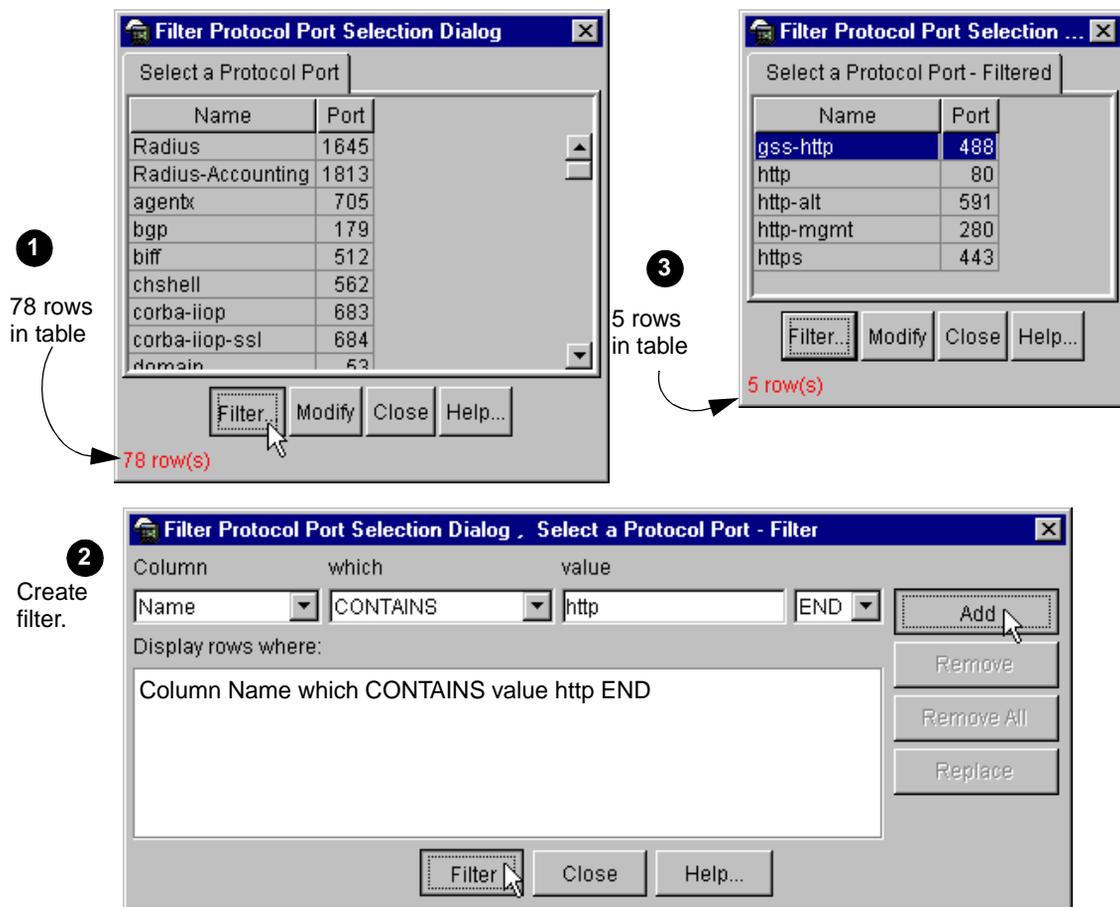
Adding a table filter

To add a table filter:

- 1 From a dialog box (Figure 15), click Filter.

The Filter dialog box opens.

Figure 15 Adding a table filter



- 2 From the Filter dialog box, click the down arrow in the Column field and select a column name from the list.
- 3 Click the down arrow in the Which field and select an [operator](#) from the list.
- 4 Enter a value in the Value field.

- 5 Click the down arrow in the last field and do one of the following:
 - To add another filter statement, choose AND or OR
 - To terminate the filter, choose END.
- 6 Click Add.

The criteria statement is added to the Display field below.
- 7 To add another criteria statement repeat steps 2 - 6.
- 8 To apply the filter to the table, click Filter > Close.

The filter dialog box closes and the filter is applied to the table.
- 9 To apply the changes to the table, click Set > Apply.

The filter is applied to the table. To remove the filter, see [“Removing a table filter,”](#) next.

Removing a table filter

To remove a table filter:

- 1 From a dialog box, click Filter.

The Filter dialog box ([Figure 15 on page 74](#)) opens.
- 2 From the Filter dialog box, select a criteria statement in the Display list.

The statement is highlighted.
- 3 Click Remove > Filter > Close.

The criteria statement is removed from the list, the results are applied to the table, and the Filter dialog box closes.

Removing all table filters

To remove all table filters:

- 1 From a dialog box, click Filter.

The Filter dialog box ([Figure 15 on page 74](#)) opens.
- 2 From the Filter dialog box, click Remove All > Filter > Close.

All filter statements are removed from the list, the results are applied to the table, and the Filter dialog box closes.

Replacing a table filter

To replace a table filter:

- 1 From a dialog box, click Filter.

The Filter dialog box ([Figure 15 on page 74](#)) opens.

- 2 From the Filter dialog box, select a statement in the Display field.

The criteria for the statement appear in the upper fields of the dialog box where you can edit them.

- 3 Modify a field(s).

- 4 Click Replace > Filter > Close.

The modified statement is returned to the Display field below, the results are applied to the table, and the Filter dialog box closes.

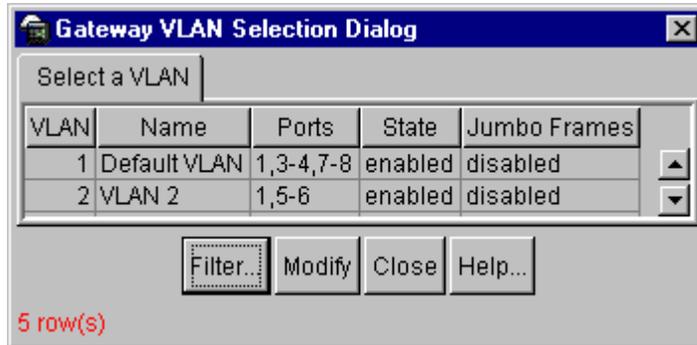
Browsing a read only dialog box

Most WSM dialog boxes have at least one Browse tool that lets you open a read only dialog box containing additional, non-editable information.

To browse and/or select from a read only dialog box:

- 1 From a dialog box, click Browse.

A read only dialog box opens.

Figure 16 Read only dialog box example

- From a selection dialog box, select an item from the list and click Modify > Close.

The Read Only dialog box closes, and your selection is added to the dialog box.



Note: To view a subset of a large selection list, you can filter the table contents by clicking [Filter](#).

Opening WSM online Help

The following Web browsers will display online Help in Device Manager. The Web browser should launch automatically when you click Help from any open dialog box.

- Microsoft Internet Explorer 5.0 or later
- Netscape Navigator 4.7 or later



Note: The WSM and the 8600 series switch have separate Device Manager Help systems. To open online Help for the WSM, you must click Help from a WSM dialog box.

To open WSM online Help:

- Right-click the Web Switching Module.

The WSM shortcut menu opens.

- 2 Click Edit.

The Web Switching Module dialog box opens.

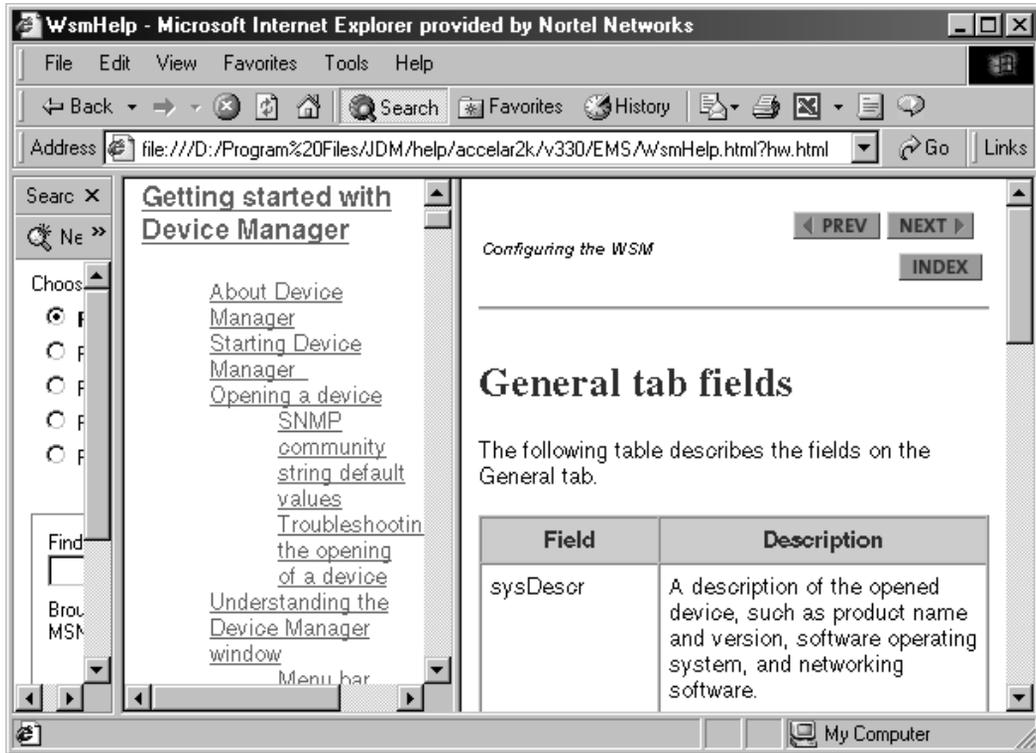
- 3 Click Help.

Your Web browser opens (Figure 17) with the online Help topic for the General tab in the right-hand panel, and the WSM Help table of contents in the left-hand panel.



Note: Be sure you have installed the WSM help files on your system. For information about installing Device Manager software, including online help files, see *Installing and Using Device Manager*, Part number 316341-A.

Figure 17 WSM online Help



Note: If your Web browser does not launch, you can find the WSM Help files in the following directory:

default install directory....JDM/help/accelar2K/v3.xx/EMS

Chapter 2

Port management

This chapter describes the ports in the WSM and how to manage them. It includes the following topics.

- “About ports,” next
- “Configuring ports” on page 83
- “Port trunking” on page 95

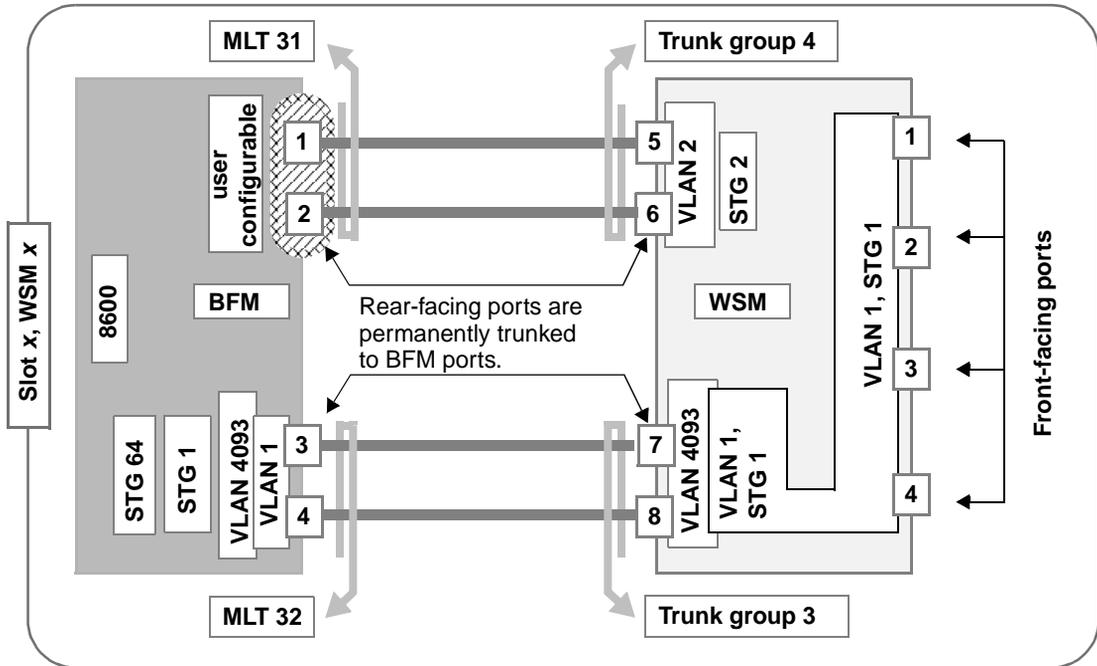
About ports

The WSM has three sets of ports:

Table 14 WSM port descriptions

| Web Switching Module Ports | Description |
|---|--|
| Front ports 1-4 | Dual-media connectors (visible on the Faceplate). |
| Back ports 5-8 | No physical connectors, operate at Gigabit Full-Duplex. Located on the printed circuit board of the WSM and permanently trunked to the BFM ports (Figure 18). Connect to other 8600 I/O modules through the BFM ports. |
| Backplane Fabric Module (BFM) ports 1-4 | Connect to the 8600 backplane. Provide recognition of the WSM by the 8600 switch and are permanently trunked with the Back ports 5-8 (Figure 18). When a WSM is inserted in a slot, the BFM ports are recognized by the 8600 switch. When viewing VLAN port membership on the WSM, the 8600 Device Manager will show the BFM ports. However, when viewing VLAN membership from the WSM Device Manager, it will show the port numbers of the Back ports to which the BFM ports are trunked. |

Figure 18 WSM port layout and default configuration



About the device view

You can configure all port parameters for both front and back ports on the WSM. When configuring a port, be sure to choose the specific port and whether it's a Front or Back port (Figure 19).

Figure 19 WSM device view

1. When configuring a port, choose either Front or Back.



2. Then choose a port

Configuring ports

This section contains the following topics:

- [“Editing ports” on page 83](#)
- [“Naming a port” on page 83](#)
- [“Setting port parameters” on page 85](#)
- [“Configuring VLAN tagging” on page 88](#)
- [“Enabling or disabling filtering on a port” on page 89](#)
- [“Applying filters to a port” on page 90](#)

Editing ports

To edit a port:

- 1 Select a port orientation (Front or Back).
- 2 Select the port you want to edit.
- 3 Do one of the following:
 - Double-click the port.
 - Right-click the port. From the shortcut menu, choose Edit.
 - From the menu bar, choose Edit > WSM Card > Port/Back Port.

The Port dialog box opens.

Naming a port

You can name a port so that a name appears next to the port number on some Device Manager information and statistics screens.

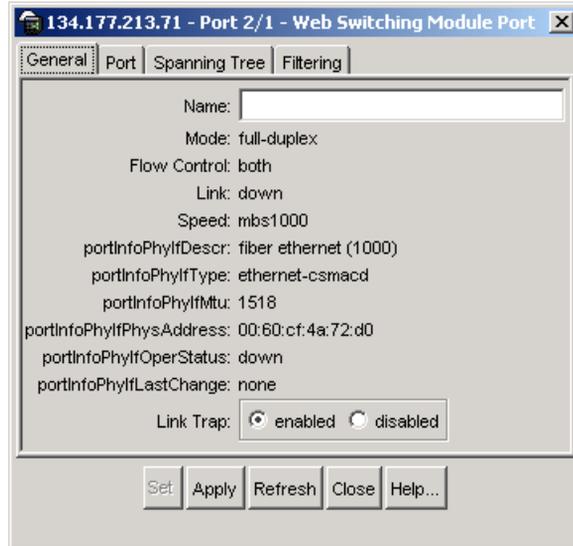
To name a port:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.

The port is highlighted.
- 3 From the menu bar, choose Edit > WSM Card > Port/Back Port.

The Web Switching Module Port dialog box opens to the **General tab** (Figure 20).

Figure 20 Port—General tab



4 To name the Port, enter a name in the Name field.

5 Click Set > Apply > Close.

The port is named and the dialog box closes.

6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 15 describes the fields on the General tab.

Table 15 Port—General tab fields

| Field | Description |
|--------------------|--|
| Name | The name of the port. |
| Mode | The port's mode— full-duplex, half-duplex, full-or-half-duplex, or other. |
| Flow Control | The port's flow control—transmit, receive, both transmit and receive, none, or other. |
| Link | The link's current state—up, down, disabled, or inoperative. |
| Speed | The port's current speed—10Mbps, 100Mbps, 1000Mbps, or other. |
| Description | Description of the Network Interface Card (NIC). |
| Type | The type of interface. |
| MTU | The largest number of bytes of a packet—the Maximum Transmission Unit of the port. |
| MAC Address | The MAC address of the WSM. |
| Operational Status | The port's current state— <ul style="list-style-type: none"> • Up—enabled, port has link and is in a forwarding state. • Down—disabled, port is not in a forwarding state • Testing—allows you to run loopback tests on the BFM port to verify forwarding functionality without impacting the entire WSM. Not recommended, modifying BFM ports can cause communication with the WSM to be lost. |
| Last Change | The date and time the interface entered its current operational state with configuration change or automatically back online—year/month/day-hour:minute:second. If the current state is configured before the network management system, Last Change will be 0 (zero). |
| Link Trap | Enables or disables linkUp/linkDown traps. |

Setting port parameters

You can set port parameters such as speed, flow control, and negotiation mode for the port link through the Port tab of the Port dialog box. You can also enable VLAN tagging, and RMON for the port.

To set port parameters:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.
The port is highlighted.
- 3 From the menu bar, choose Edit > WSM Card > Port/Back Port.
The Web Switching Module Port dialog box opens to the **General** tab.
- 4 Click the Port tab.
The **Port** tab (Figure 21) opens.

Figure 21 Port tab

The screenshot shows a dialog box titled "10.160.17.51 - Port 2/1 - WebSwitch Module Port" with a close button in the top right corner. The dialog has four tabs: "General", "Port", "Spanning Tree", and "Filtering". The "Port" tab is active. The settings are as follows:

- State: enabled disabled
- VLAN: tagged untagged
- Default VLAN: 1..4092
- Gig Auto Negotiate: on off
- Gig Flow Control:
- Auto Negotiate: on off
- Speed:
- Mode:
- Flow Control:
- RMON: on off
- Discard Non-IP: enabled disabled
- Preferred Link: fast-ethernet gigabit-ethernet
- Backup Link:
- BWM Contract: 1..1024

At the bottom of the dialog are five buttons: "Set", "Apply", "Refresh", "Close", and "Help...".

5 Use the fields on the Port tab to make port configurations.

Browse opens a selection list of available choices. See [“Browsing a read only dialog box” on page 76](#).

6 Click Set > Apply > Close.

The configuration is applied to the Port and the dialog box closes.

7 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 16](#) describes the fields on the Port tab.

Table 16 Port tab fields

| Field | Description |
|--------------------|---|
| State | Enables or disables the port. |
| VLAN | Sets the VLAN tag of the port—tagged or untagged. You cannot turn off tagging for trunk groups 3 and 4. |
| Default VLAN | Sets the default VLAN ID for the port—1 to 4,092. |
| Gig Auto Negotiate | Sets autonegotiation for the Gigabit Ethernet connection—on or off. |
| Gig Flow Control | Sets port flow control for the Gigabit Ethernet connection—other, transmit, receive, both, or none. |
| Auto Negotiate | Sets autonegotiation for the Fast Ethernet connection—on or off. |
| Speed | The port speed for Fast Ethernet connection—10Mbps, 100Mbps, or 10/100Mbps. |
| Mode | The port mode for fast Ethernet connection—full-duplex, half-duplex, or full-or-half-duplex. |
| Flow Control | Sets port flow control for Fast Ethernet connection—transmit, receive, both, or none. |
| RMON | Sets the remote monitoring for the port —on or off. The default is off. |
| Discard Non-IP | Enables or disables discarding all non-IP traffic on the port. |

Table 16 Port tab fields (continued)

| Field | Description |
|----------------|--|
| Preferred Link | Sets the preferred link—Fast Ethernet or Gigabit Ethernet. Used by the WSM to select which port to use if both the 100Mbps and 1000Mbps are plugged in on a single port. |
| Backup Link | Sets the backup link—Fast Ethernet, Gigabit Ethernet, or None. Used by the WSM to select backup port if both the 100Mbps and 1000Mbps are plugged in on a single port. Only available on models with dual-media ports. |
| BWM Contract | Sets the number of the bandwidth management contract—1 to 1,024. For more information about Bandwidth Management, see “Bandwidth management” on page 493 . |

Configuring VLAN tagging

You can configure VLANs to tag or untag frames for a port. You cannot turn off VLAN tagging for trunk groups 3 and 4.

To configure VLAN tagging:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.
- 3 From the menu bar, choose Edit > WSM Card > Port/Back Port.
The Port dialog box opens to the [General tab](#).
- 4 Click the Port tab.
The [Port tab](#) ([Figure 21 on page 86](#)) opens.
- 5 In the VLAN field, click Tagged or Untagged.
- 6 Click Set > Apply > Close.
VLANs will tag frames as you have specified for this port.
- 7 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.

- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Monitoring link state

To enable (or disable) monitoring of the link state (up or down) for a port:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.
- 3 From the menu bar, choose Edit > WSM Card > Port/Back Port.

The Port dialog box opens to the [General tab](#) (Figure 20 on page 84).

- 4 In the LinkTrap field click Enabled (or Disabled).
- 5 Click Set > Apply > Close.

Link state (up or down) information will be gathered for the port and displayed on the General tab.

- 6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Enabling or disabling filtering on a port

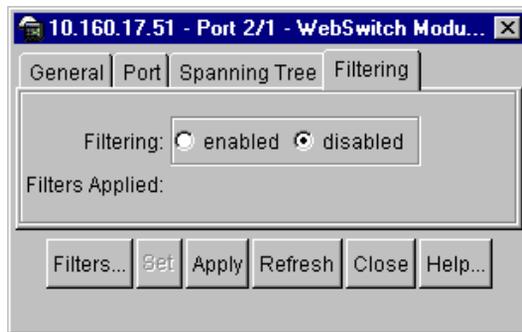
To enable or disable filtering on a port or display the numbers of individual filters applied to a port:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a WSM port.
- 3 From the menu bar, choose Edit > WSM Card > Port/Back Port.

The Port dialog box opens to the [General tab](#).

- 4 Click the Filtering tab.

The [Filtering tab](#) opens.

Figure 22 Port—Filtering tab

5 To enable/disable filtering on the port, click Enable/Disable.

6 Click Set > Apply > Close.

Filtering on the port is enabled/disabled and the dialog box closes.

7 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 17](#) describes the fields on the Filtering tab.

Table 17 Port—Filtering tab fields

| Field | Description |
|-----------------|--|
| Filtering | Enables or disables filtering. |
| Filters Applied | Displays the numbers of the applied filters. |

Applying filters to a port

To apply available filters to a port:

- 1** From the Device view, select a port orientation (Front or Back).
- 2** Select a port.

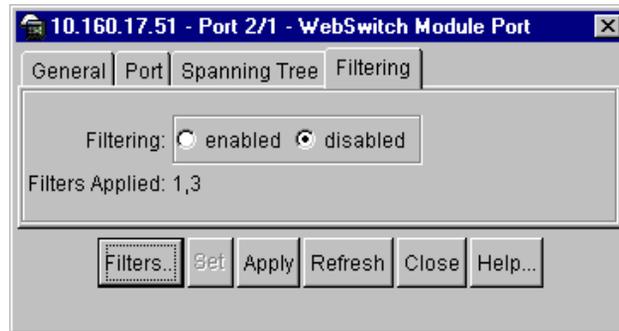
- 3 From the menu bar, choose Edit > WSM Card > Port/Back Port.

The Port dialog box opens to the [General tab](#).

- 4 Click the Filtering tab.

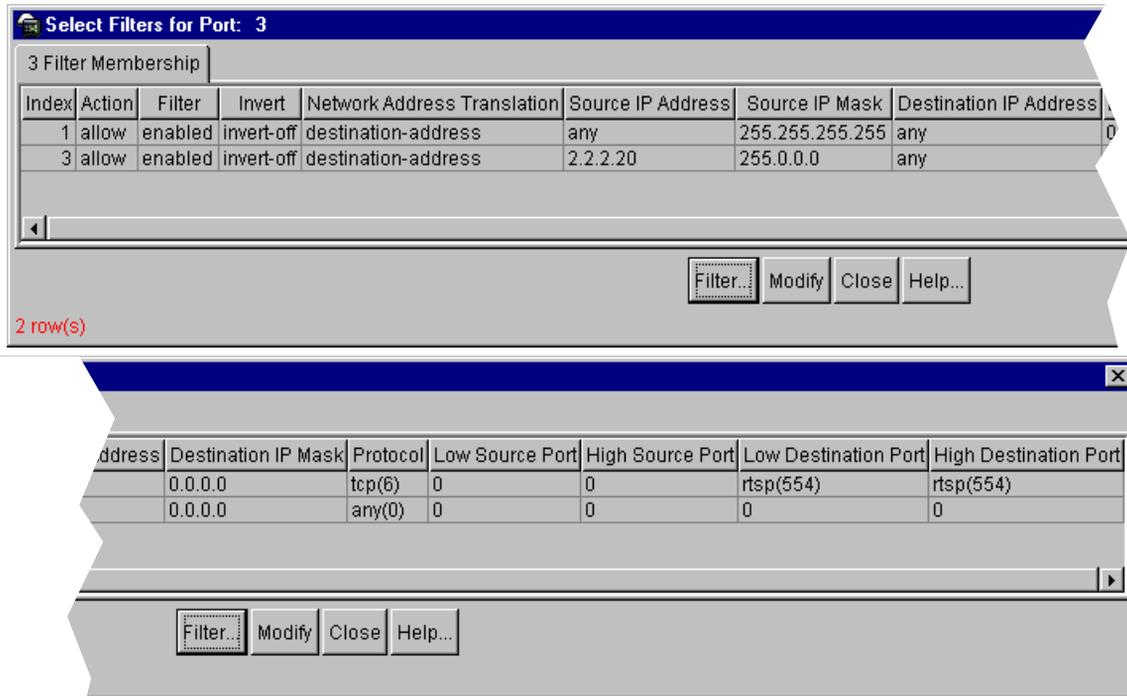
The [Filtering tab](#) ([Figure 23](#)) opens, displaying the filter numbers currently applied to this port in the Filters Applied field.

Figure 23 Port—Filtering dialog box with filters applied



- 5 Click Filters.

The Select Filters for Port dialog box opens displaying available filters on the [Filter Membership tab](#).

Figure 24 Select Filters for Port dialog box

6 Choose ALL filters that you want to apply.



Note: If you have existing filters applied and want to add another, choose both the existing filters and the new filter (CTRL + Click). If you just choose the new filter, when you click Modify, only the new filter is applied, and the previous filters are removed.

7 Click Modify > Close

The filter is applied and the Select Filters for Port dialog box closes.



Note: If you have many filters configured and want to view a subset of this table, see [“Filtering table content” on page 73](#).

8 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 18 describes the Filter Membership fields.

Table 18 Port—Filter Membership tab fields

| Field | Description |
|-----------------------------|--|
| Index | The filter index number. |
| Action | The action for the filtering rule: allow, deny, NAT, or redirect. |
| Filter | The state of the filtering rule: disabled or enabled. |
| Invert | The setting of the invert logic for the filter entry: invert-off or invert-on. |
| Network Address Translation | The selection of destination-address or source-address for Network Address Translation. |
| Source IP Address | The source IP address that is affected by this filter. |
| Source IP Mask | The source IP subnet mask used with the Source IP Address. |
| Destination IP Address | The destination IP address that is affected by this filter. |
| Destination IP Mask | The source IP subnet mask used with the Destination IP Address. |
| Protocol | The selected filter protocol. |
| Low Source Port | The lower source TCP/UDP port number to filter. This applies only when the filter protocol is defined as UDP or TCP. A value of zero (0) indicates no filtering. |
| High Source Port | The higher source TCP/UDP port number to filter. This applies only when the filter protocol is defined as UDP or TCP. A value of zero (0) indicates no filtering. |
| Low Destination Port | The lower destination TCP/UDP port number to filter. This applies only when the filter protocol is defined as UDP or TCP. A value of zero (0) indicates no filtering. |
| High Destination Port | The higher destination TCP/UDP port number to filter. This applies only when the filter protocol is defined as UDP or TCP. A value of zero (0) indicates no filtering. |

Removing a filter from a port

To remove a filter from a port:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.
- 3 From the menu bar, choose Edit > WSM Card > Port/Back Port.

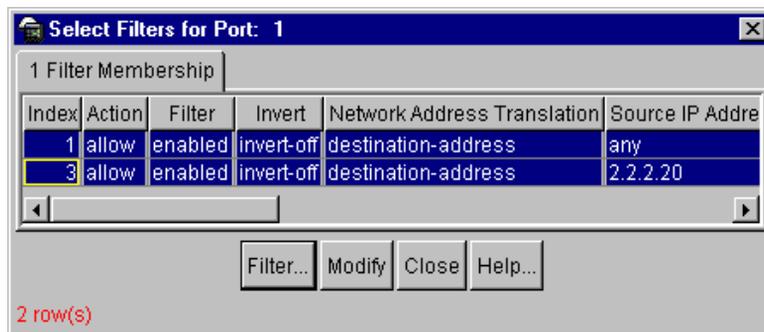
The Port dialog box opens to the **General** tab.

- 4 Click the Filtering tab.

The **Filtering** tab (Figure 22 on page 90) opens, displaying the filter numbers currently applied to this port in the Filters Applied field.

- 5 Click Filters.
- 6 The Select Filters for Port dialog box opens (Figure 25) to the Filter Membership tab with all applied filters highlighted.

Figure 25 Filter Membership tab



- 7 CTRL + Click the filter you want to remove.

The filter is no longer highlighted.

- 8 Click Modify.

The dialog box closes, and the filter number is removed from the Filters Applied field on the Filtering tab.

- 9 From the Filtering tab, click Set > Apply > Close.

The filter is removed from the port and the Port dialog box closes.

- 10** To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Port trunking

Trunk groups provide bandwidth connections between the WSM and the 8600. A trunk is a group of ports that act together, combining their bandwidth to create a single, larger port. Up to four trunk groups can be configured on the WSM, with the following restrictions:

- A physical WSM port can belong to no more than one trunk group.
- Up to four ports can belong to the same trunk group.
- Best performance is achieved when all ports in a trunk group are configured for the same speed.

The Device Manager Trunks tab supports:

- [“Viewing trunk group configurations” on page 96](#)
- [“Configuring a trunk group” on page 98](#)
- [“Deleting a trunk group” on page 99](#)
- [“Filtering table content” on page 73](#)

Default multi-link trunk configuration

The WSM integrates the four Rear-facing ports into the 8600 switch chassis backplane. By default ([Figure 18 on page 82](#)), these ports are configured into two permanent MLT groups with VLAN tagging enabled:

- Back ports 5 and 6 are permanently trunked to BFM ports 1 and 2 respectively.
- Back ports 7 and 8 are permanently trunked to BFM ports 3 and 4 respectively.

The 8600 switch allows a total of 32 MLT groups. Each WSM installed in the 8600 switch chassis will be assigned two available MLT groups from the switch. The WSM allows four trunk groups. Because trunk groups 3 and 4 are pre-configured and cannot be deleted, you can configure trunk groups 1 and 2 only.

Dynamic MLTs

The default MLT groups are assigned dynamically. After inserting a WSM into the 8600 chassis, the highest available MLT group is assigned to BFM ports 3 and 4. The next highest available MLT group is assigned to BFM ports 1 and 2. Based on availability, the MLT group numbers may not always be consecutive.

For example:

- If MLT 32 and 31 are used, the WSM is assigned MLT 30 and 29.
- If MLT 31, 30, and 29 are used, the WSM is assigned MLT 32 and 28.

MLT assignment may change when the WSM is reset, because the WSM will assign the first MLT pair to the WSM that is installed in the lowest slot number.

Table 19 Initial MLT assignment compared with assignment after reset

| Initial installation | After a reset |
|---|--|
| 1. WSM #1 is first inserted in slot 8 and assigned MLT 31 and 32. | 1. WSM #1 in slot 8 is assigned MLT 29 and 30. |
| 2. WSM #2 is then inserted in slot 1 and assigned MLT 29 and 30. | 2. WSM #2 in slot 1 is assigned MLT 31 and 32. |



Note: The physical slots in the chassis are numbered 1 - 10, from top to bottom. Lowest slot number refers to slot 1 at the top of the chassis.

Viewing trunk group configurations

To view the trunk group configurations for the WSM:

- 1 From the Device view, select the Web Switching Module.

The module is highlighted.

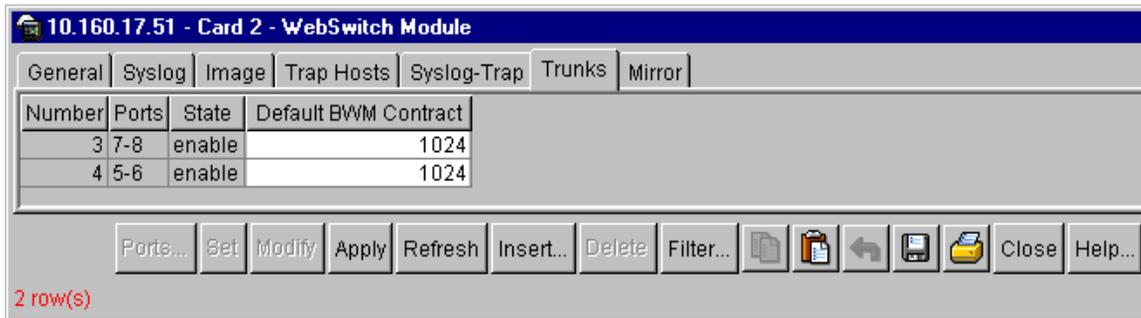
- 2 From the Edit menu, select WSM Card > General.

The Web Switching Module dialog box opens to the General tab.

- 3 Click the Trunks tab.

The [Trunks](#) tab opens listing the trunk group configurations for the Web Switching Module.

Figure 26 Trunks tab



[Table 20](#) describes the Trunks configuration fields.

Table 20 Trunks configuration fields

| Field | Description |
|----------------------|--|
| Number | The number of the trunk group: 1 to 4. This value can be configured by clicking Insert. |
| Ports | The physical ports in the trunk group. The ports can be set by clicking Insert. |
| State | Enables or disables the trunk group. |
| Default BWM Contract | Sets the Bandwidth Management contract number: 1 to 1,024. For more information about Bandwidth Management, see "Bandwidth management" on page 493 . |

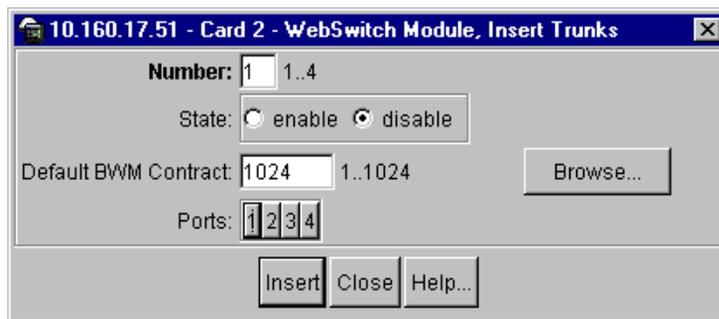
Configuring a trunk group

Because trunk groups 3 and 4 are pre-configured and cannot be deleted, you can configure trunk groups 1 and 2 only. Ports cannot be removed from trunk groups 3 and 4.

To configure a trunk group:

- 1 From the Device view, select the Web Switching Module.
The module is highlighted.
- 2 From the Edit menu, select WSM Card > General.
The Web Switching Module dialog box opens to the General tab.
- 3 Click the Trunks tab.
The **Trunks** tab (Figure 26 on page 97) opens.
- 4 Click Insert.
- 5 The Insert Trunks dialog box opens.

Figure 27 Insert trunks dialog box



- 6 In the number field, type a trunk group number (1-4).
- 7 In the State field, click Enable or Disable to enable or disable the trunk group.
- 8 In the Default BWM Contract field, type the number of a bandwidth management contract, or click Browse to browse a selection list. For more information, see [“Browsing a read only dialog box” on page 76](#).

For more information about bandwidth management, see [“Bandwidth management” on page 493](#).

9 In the Ports field, click the physical ports for the trunk group.

10 Click Insert.

11 To insert more trunks, repeat Steps 6-10.

12 Click Close.

The trunk group is inserted into the Trunks tab with the selected ports, and the Insert Trunks dialog box closes.

13 In the Trunks tab, click Set > Apply > Close.

The trunk group is configured, and the Web Switching Module dialog box closes.

14 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Deleting a trunk group

You cannot delete trunk groups 3 and 4.

To delete a trunk group:

1 From the Device view, select the Web Switching Module.

The module is highlighted.

2 From the Edit menu, select WSM Card > General.

The Web Switching Module dialog box opens to the General tab.

3 Click the Trunks tab.

The [Trunks](#) tab ([Figure 26 on page 97](#)) opens.

4 Select a trunk group and click Delete.

The trunk group is deleted.

5 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Chapter 3

Configuring Layer 2 operations

This chapter describes how to configure Layer 2 operations for the Web Switching Module including Virtual LANs (VLANs) and Spanning Tree Protocol. It includes the following topics:

- [“About VLANs” on page 101](#)
- [“Configuring a VLAN” on page 106](#)
- [“Viewing MultiLink trunk and VLAN membership” on page 108](#)
- [“About Spanning Tree protocol” on page 108](#)
- [“Configuring spanning tree protocol” on page 114](#)

About VLANs

You can configure 246 VLANs on the Web Switching Module using VLAN IDs 1- 4092. During default configuration, VLANs 1, 2, and 4093 (Table 21) are created automatically. For more information, see [“About ports” on page 81](#).

You can change the port memberships or VLAN IDs for VLANs 1 and 2 but not VLAN 4093. If VLANs 1 and 2 are already configured in your network, consider whether Web Switching Module traffic should use VLANs 1 and 2. If so, no change is necessary. If not, assign other VLAN IDs for the Web Switching Module traffic after inserting the module.

Table 21 VLAN restrictions for the Web Switching Module

| VLAN | Restriction |
|-----------|--|
| VLAN 1 | By default, VLAN 1 is assigned to Front-facing ports 1-4, Rear-facing ports 7 and 8, and BFM ports 3 and 4. VLAN 1 cannot be deleted on either the WSM or the 8600 series switch. |
| VLAN 2 | By default, VLAN 2 is assigned to Rear-facing ports 5 and 6 and BFM ports 1 and 2. You cannot delete VLAN 2 on the Web Switching Module; but you can delete it on the 8600 series switch. |
| VLAN 4093 | VLAN 4093 is reserved for management of the Web Switching Module's connection to the Passport backplane. VLAN 4093 is permanently configured on the trunked Rear-facing ports 7 and 8 and BFM ports 3 and 4 and cannot be modified or deleted. If VLAN 4093 is already configured on your network, you must change that VLAN to use another VLAN ID before inserting the Web Switching Module. |



Note: All ports must belong to at least one VLAN. Any port which is removed from a VLAN, and which is not a member of any other VLAN, is automatically added to default VLAN #1. You cannot remove a port from VLAN #1 if the port has no membership in any other VLAN. Also, you cannot add a port to more than one VLAN unless the port has VLAN tagging turned on. See [“Configuring VLAN tagging” on page 88](#).

Management IP address used for VLAN 4093

In order for the 8600 series switch to communicate with the WSM, a Management IP address (172.31.255.240/28) is assigned to VLAN 4093 during default configuration. At the same time, the address 172.31.255.245 is assigned to the 8690SF or 8691SF module. This management IP address is hidden and will not appear in the IP ARP table and IP routing table.

VLANs and IP interfaces

Access to the WSM for remote configuration, trap messages, and other management functions is accomplished through an IP interface (see [“Configuring IP interfaces” on page 129](#)). Access to management functions can also be cut off to any VLAN that doesn’t have an IP interface assigned.

For example, if all IP interfaces are left on VLAN 1 (the default), and all ports are configured for VLANs other than VLAN 1, then WSM management features are effectively cut off. If an IP interface is added to one of the other VLANs, the stations in that VLAN all have access to WSM management features.

About VLAN tagging

The WSM supports 802.1Q VLAN tagging, providing standards-based VLAN support for Ethernet systems. Tagging places the VLAN identifier in the frame header, allowing multiple VLANs per port. When you configure multiple VLANs on a port, you must also enable tagging on that port. Since tagging fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged frames from being transmitted to devices that do not support 802.1Q VLAN tags. To configure VLAN tagging on a port, see [“Configuring VLAN tagging” on page 88](#).

About Jumbo frames

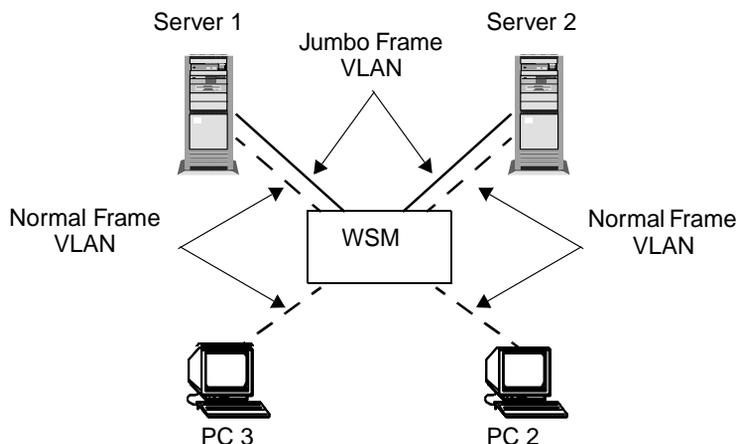
To reduce host frame processing overhead, the WSM supports Jumbo frames of up to 9018 octets. Jumbo frames are received and transmitted between the WSM and any Gigabit network adapter that can handle frame sizes of 9K and higher. By sending one Jumbo frame instead of many smaller frames, the same task requires less processing.

Jumbo frames can be transmitted and received between Gigabit adapter-enabled hosts through the WSM across any VLAN that has Jumbo frames enabled.

Isolating Jumbo frame traffic using VLANs

Do not use Jumbo frames in a VLAN with any device that can't process Jumbo frames (frames that exceed the maximum Ethernet frame size) or on ports that are configured for half-duplex mode. You can assign separate VLANs for non-Jumbo frame support. When attached to a WSM, end-stations with Gigabit adapters that can receive and transmit Jumbo frames, can communicate with both the Jumbo frame VLANs and regular frame VLANs at the same time. In Figure 28, the two servers can handle Jumbo frames but the two clients cannot; therefore Jumbo frames should only be enabled and used on the VLAN represented by the solid lines but not the VLAN with the dashed lines.

Figure 28 Jumbo frame VLANs



Routing Jumbo frames to non-Jumbo frame VLANs

With IP routing between VLANs, the WSM will fragment Jumbo UDP datagrams when routing from a Jumbo frame VLAN to a non-Jumbo frame VLAN. To configure a VLAN using Jumbo frames, see [“Configuring a VLAN” on page 106](#).

Viewing current VLAN membership

The Virtual LAN Membership dialog box displays your current VLAN configuration, including VLAN numbers, port assignments, VLAN state, etc.

To view your current VLAN configuration:

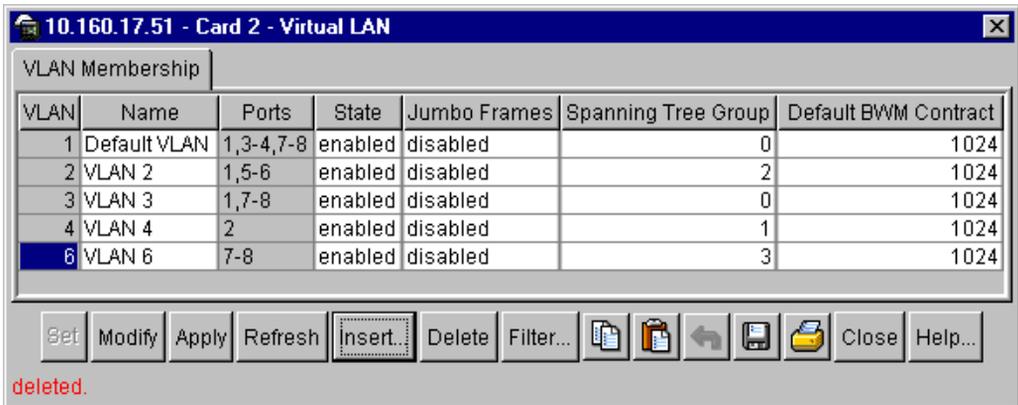
- 1 From the Device View, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > Virtual LANs.

The [Virtual LAN dialog box](#) (Figure 29) opens with the fields described in [Table 22](#).

Figure 29 Virtual LAN—VLAN Membership tab



[Table 22](#) describes the fields on the VLAN Membership tab and the Insert VLAN dialog box.

Table 22 Virtual LAN—VLAN Membership fields

| Field | Description |
|--------------|--|
| VLAN | Sets the VLAN identification number. The number can be set when a VLAN is Inserted or modified. |
| Name | Sets the VLAN name. The default is none, except for the first VLAN name, which defaults to Default VLAN. |
| Ports | Sets the ports in the VLAN. The default is none except for VLAN 1 which defaults to ports 1-9. The port list is set when a VLAN is Inserted. The list can be edited by using the Ports button. |
| State | Enables or disables a VLAN. The default is enabled. |
| Jumbo Frames | Enables or disables Jumbo frame support for the VLAN. The default is disabled. |

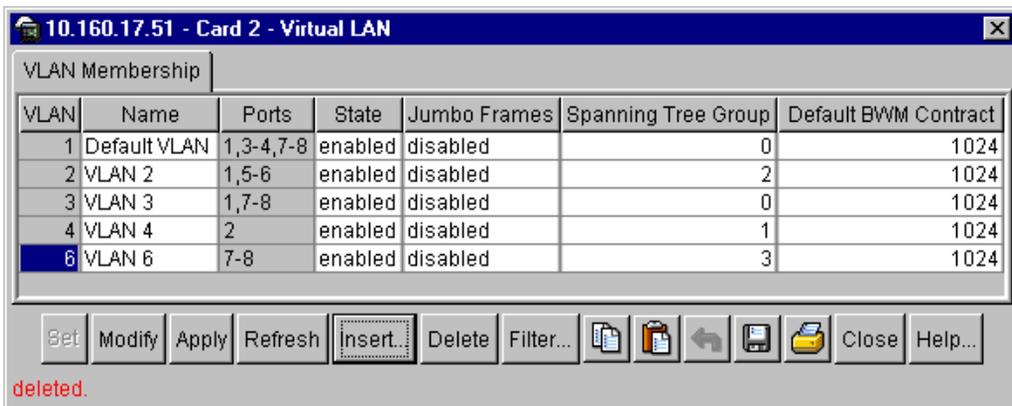
Table 22 Virtual LAN—VLAN Membership fields (continued)

| Field | Description |
|----------------------|---|
| Spanning Tree Group | Sets the number of the Spanning Tree Group (STG) assigned to the VLAN. |
| Default BWM Contract | Sets the number of the default Bandwidth Management (BWM) Contract for the VLAN. The default is 1024. |

Configuring a VLAN

To configure a VLAN:

- 1 From the Device View, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, select WSM Card > Virtual LANs.
The Virtual LAN dialog (Figure 30) box opens.

Figure 30 Virtual LAN dialog box

- 3 Click Insert.

The [Virtual LAN, Insert VLAN Membership dialog box \(Figure 31\)](#) opens.

See [Table 22 on page 105](#) for a description of the VLAN fields.

Figure 31 Virtual LAN, Insert VLAN Membership dialog box

- 4 In the VLAN field, type a VLAN number (1-4092).
- 5 In the Name field, type a VLAN name.
- 6 In the State field, click Enabled to enable the VLAN.
- 7 (Optional) In the Jumbo Frames field, click Enabled to use [Jumbo Frames](#) on the VLAN. For more information, see [“About Jumbo frames” on page 103](#).
- 8 (Optional) In the Default BWM field, type a bandwidth management contract number; or click Browse to choose from a selection list of available BWM contracts. For more information, see [“Browsing a read only dialog box” on page 76](#). For more information about BWM contracts, see [“Contracts” on page 496](#).
- 9 In the Ports field, click the individual port numbers to assign port membership for the VLAN.
- 10 Click Insert.

The VLAN is entered into the VLAN Membership tab, and the Insert VLAN Membership dialog box closes.
- 11 From the VLAN Membership tab, click Set > Apply > Close.

The VLAN is configured and the Virtual LAN dialog box closes.
- 12 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Viewing MultiLink trunk and VLAN membership

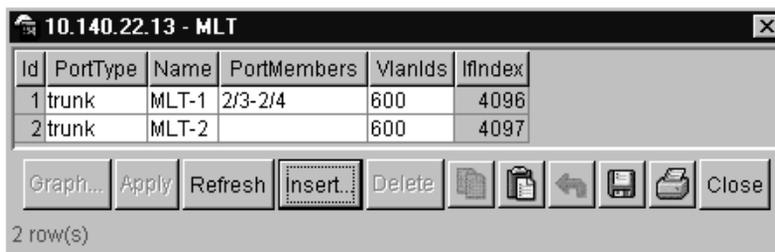
You can view the MultiLink trunk (MLT) groups for each module, the VLAN ids, and the member ports' backplane fabric (BFM) port numbers. For more information about multilink trunking, see [“Port trunking” on page 95](#).

To view MLT VLAN membership:

- 1 From the Device View, choose VLAN > MLT.

The MLT dialog box ([Figure 32](#)) opens displaying the active MLTs.

Figure 32 MLT dialog box



The screenshot shows a window titled "10.140.22.13 - MLT" with a table of MLT configurations. The table has columns for Id, PortType, Name, PortMembers, VlanIds, and IfIndex. Below the table are buttons for Graph..., Apply, Refresh, Insert, Delete, and a set of icons for file operations (copy, paste, undo, redo, save, print) and a Close button. The status bar at the bottom indicates "2 row(s)".

| Id | PortType | Name | PortMembers | VlanIds | IfIndex |
|----|----------|-------|-------------|---------|---------|
| 1 | trunk | MLT-1 | 2/3-2/4 | 600 | 4096 |
| 2 | trunk | MLT-2 | | 600 | 4097 |

About Spanning Tree protocol

To prevent loops in the network topology, the WSM supports IEEE 802.1d Spanning Tree Protocol (STP) with 16 configurable spanning tree groups (STGs). STG 1 is the default STG for all VLANs. You can assign individual VLANs to any of the other 15 STGs.

VLANs and spanning tree protocol

STP detects and eliminates logical loops in a bridged or switched network. STP forces redundant data paths into a standby (blocked) state. When multiple paths exist, STP configures the network so that a WSM uses only the most efficient path. If that path fails, STP automatically sets up another active path on the network to sustain network operations.



Note: Due to STP's sequence of listening, learning, and forwarding or blocking, lengthy delays may occur. For more information on using STP in cross-redundant topologies, see *Web OS Switch Software 10.0 Application Guide*, part number 212777-A..

The following table shows the relationship between port, trunk groups, VLANs, and spanning trees.

Table 23 Port, trunk group, VLANs, and Spanning tree relationships

| WSM element | Belongs to |
|-------------|--------------------------------------|
| Port | Trunk group or 1 or more VLANs |
| Trunk group | 1 or more VLANs |
| VLAN | 1 Spanning tree group |



Note: By default, all newly created VLANs are members of STG 1.

Understanding multiple spanning trees

To create the spanning tree, the root bridge (switch) generates a bridge protocol data unit (BPDU), to forward out of its ports. BPDUs are messages used by bridges implementing STP to learn about the existence of other bridges for the purpose of calculating and maintaining the spanning tree. By exchanging BPDUs,

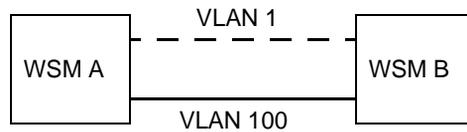
all switches in the Layer 2 network participating in the spanning tree gather information about other switches in the network. The BPDU is used to establish a path, much like a Hello packet in IP routing. [Figure 33](#) shows the information contained in a BPDU about the transmitting bridge and its ports.

Figure 33 BPDU message format

| | Byte | |
|---------------------|------|---------------------|
| Protocol Identifier | 1 | |
| Protocol version | 2 | |
| BPDU Type | 3 | |
| Flags | 4 | |
| Root Identifier | 5 | Current Root Bridge |
| Root Path Cost | 6 | |
| | 13 | |
| Bridge Identifier | 14 | Path Cost to Root |
| | 17 | |
| Port Identifier | 18 | This Bridge ID |
| | 25 | |
| Message Age | 26 | This Port ID |
| | 27 | |
| Max Age | 28 | |
| | 29 | |
| Hello Time | 30 | |
| | 31 | |
| Forward Delay | 32 | |
| | 33 | |
| | 34 | |
| | 35 | |

Why multiple Spanning Trees?

[Figure 34](#) shows a simple example of multiple spanning trees. Two VLANs, VLAN 1 and VLAN 100 exist between WSM A and WSM B. If VLAN 1 and VLAN 100 are in the same Spanning Tree topology or Spanning Tree Group, then a loop exists between the two WSMs. If one of the ports is blocked, then the VLAN loses connectivity between the WSMs. If VLAN 1 and VLAN 100 belong to different Spanning Tree Groups, then the two instances of Spanning Tree separate the topology without forming a loop. Both VLANs can forward packets between WSMs without losing connectivity.

Figure 34 Example of multiple instances of Spanning Tree protocol

Spanning Tree Group 1: VLAN 1

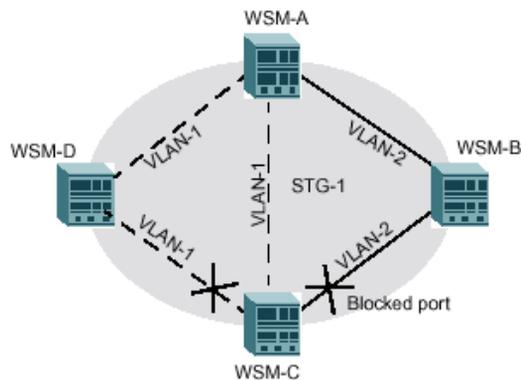
Spanning Tree Group 2: VLAN 100

Example—four-WSM topology with single spanning tree

In the following four-WSM topology example (Figure 2-4), assuming WSM A has the higher priority, you can have at least three loops on the network:

- traffic flowing from WSMs A-B-C-A
- traffic flowing from WSMs A-C-D-A
- traffic flowing from WSMs A-B-C-D-A.

In a single spanning tree environment, depending on the bridge priority, port priority, and port cost, it is possible that the link between WSMs C and D or between WSMs B and C may be blocked. If the block is between WSMs B and C, then the blocked link will inadvertently isolate VLAN 3 altogether.

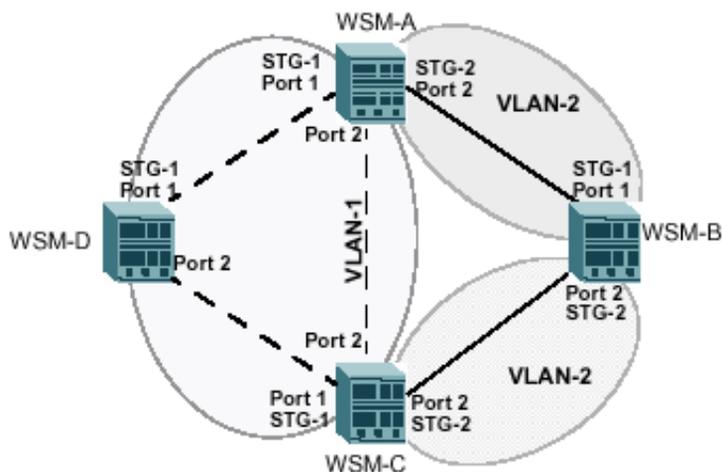
Figure 35 Example of single spanning tree—VLAN 3 blocked

However, if multiple spanning trees are implemented and each VLAN is on a different spanning tree, then elimination of logical loops will not isolate any VLAN altogether.

Example—four-WSM topology with multiple spanning trees

Figure 36 shows multiple spanning trees enabled on the same four-WSM topology found in Figure 35.

Figure 36 Example of multiple spanning tree groups



Multiple Spanning Trees can be enabled on tagged or untagged ports. Three instances of Spanning tree are configured in the example shown in Figure 36. The multiple spanning trees on the four WSMs are configured as shown in (Table 24)

Table 24 Multiple Spanning Tree groups example—VLAN, WSM, STG

| Web Switching Module | VLAN 1 | VLAN 2 | VLAN 3 |
|----------------------|--|--|---------------------------------|
| A | Spanning Tree group 1 Ports 1 and 2 | Spanning Tree group 2 ¹ Port 8 | |
| B | | Spanning Tree group 1 ¹ Port 1 | Spanning Tree group 2 Port 8 |

Table 24 Multiple Spanning Tree groups example—VLAN, WSM, STG

| Web Switching Module | VLAN 1 | VLAN 2 | VLAN 3 |
|----------------------|---------------------------------------|--------|---------------------------------|
| C | Spanning Tree group 1 Port 1 and 2 | | Spanning Tree group 2 Port 8 |
| D | Spanning Tree group 1 Port 1 and 8 | | |

1 On WSM A, VLAN 2 is on STG 2 and on WSM B, VLAN 2 is on STG 1. Spanning tree is switch-centric—each spanning tree decision is based on the configuration of the specific switch.

The WSMs use [BPDU](#) to [determine root cost](#).

Multiple spanning tree example—determining root cost

The following process describes how root cost is determined in [Figure 36](#).

- Once WSMs B and D receive the [BPDU](#) from the root bridge (WSM A), they modify the BPDU by adding their bridge ID numbers. They then forward the BPDU out their egress ports.
- When WSM C receives BPDUs from both B and D, it can immediately determine from the information in the BPDU that they are both part of the same spanning tree. WSM C thus knows that bridging them together will create a loop in the network. At this point, C makes a decision that one port will be used for forwarding, and the other port will be blocked.
- If the ports are tagged, each port sends out a special BPDU containing the tagged information.
- When determining which port to use for forwarding and which port to block, WSM D uses information in the BPDU, including each bridge priority ID. The lowest root cost is then computed to determine the most efficient path for forwarding. For more information on bridge priority, port priority, and port cost, refer to the *Web OS Switch Software 10.0 Command Reference*, Part number 212778-A. .

Root cost comparison

Much like least-cost routing, root cost ([Table 25](#)) assigns lower cost values to high-bandwidth ports, such as Gigabit Ethernet, to encourage their use.

Table 25 Root cost

| Link | Cost |
|---------------------------------|------|
| 10 Mbps | 100 |
| 100 Mbps (Fast Ethernet) | 19 |
| 1000 Mbps (Gigabit Ethernet) | 4 |

The objective is to use the fastest links so that the route with the lowest cost is chosen.

Configuring spanning tree protocol

By default Spanning Tree Groups 2-15 are empty, and Spanning Tree Group 1 contains all configured VLANs until individual VLANs are explicitly assigned to other Spanning Tree Groups. You can assign only one VLAN per STG, except for STG 1.

Configuring a Spanning Tree Group

To assign a VLAN to a Spanning Tree Group:

- 1 From the Device View, select the Web Switching Module.

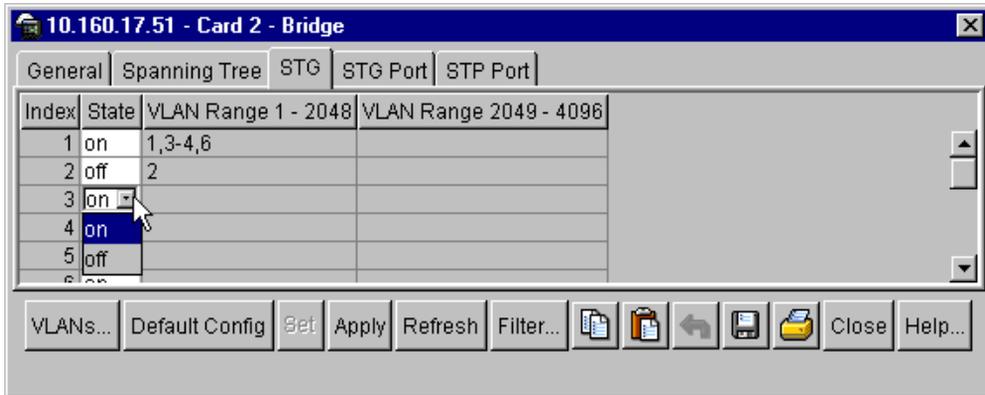
The Web Switching Module is highlighted.

- 2 From the Edit menu, select Edit > WSM Card > Bridge.

The Bridge dialog box opens to the General tab.

- 3 Click the STG tab.

The **STG** tab ([Figure 37](#)) opens with the fields defined in [Table 26](#) on [page 116](#).

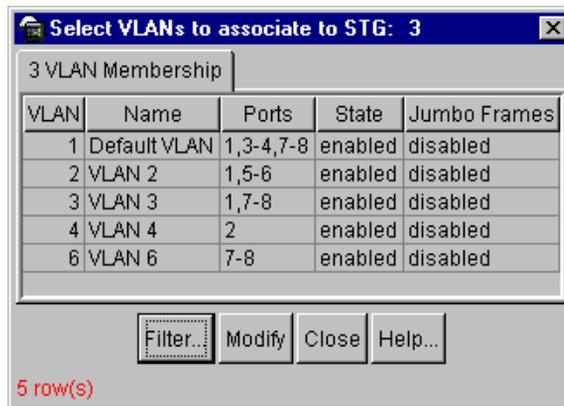
Figure 37 Bridge—Spanning Tree Group tab

- 4 Click the State field and choose On from the drop down list.
- 5 Click VLANs.

The Select VLANs dialog box ([Figure 38](#)) opens.



Note: If you have many VLANs configured and want to view a subset of this table, see [“Filtering table content” on page 73](#).

Figure 38 Select VLANs dialog box

- 6 Choose a VLAN and click Modify.

The Select VLANs dialog box closes, and the VLAN you selected appears in the VLAN Range field for the STG.

7 On the STG tab, click Set > Apply > Close.

The VLAN is assigned to the STG, the STG is enabled, and the dialog box closes.

8 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 26 describes the fields on the bridge STG tab.

Table 26 Bridge STG tab fields

| Field | Description |
|----------------------|--|
| Index | The Spanning Tree Group (STG) index number. You can configure a maximum of 16 STGs on the Web Switching Module. |
| State | The current state of Spanning Tree Protocol for each Spanning Tree Group. |
| VLAN Range 1-2048 | Displays the VLANs in the range 1 to 2048 for each Spanning Tree Group. Only the default Spanning Tree Group (STG 1) may contain more than one VLAN. All other Spanning Tree Groups may contain only one VLAN each. If you select more than one VLAN for Spanning Tree Groups 2-16, an error message will appear when you attempt to apply the configuration. |
| VLAN Range 2049-4096 | Displays the VLANs in the range 2049 to 4096 for each Spanning Tree Group. Only the default Spanning Tree Group (STG 1) may contain more than one VLAN. All other Spanning Tree Groups may contain only one VLAN each. If you select more than one VLAN for Spanning Tree Groups 2-16, an error message will appear when you attempt to apply the configuration. |
| VLANs | Opens the <i>VLANs to Associate to STG</i> dialog box listing all available VLANs. From this dialog box, you select VLAN membership for the STG. |

[Table 27](#) describes the fields on the Select VLANs to Associate with STG dialog box. This dialog box opens when you click VLANs from the Bridge STG tab.

Table 27 Select VLANs to Associate with STG dialog box fields

| Field | Description |
|--------------|---|
| VLAN | The VLAN identification number. The number is set when a VLAN is Inserted or modified. |
| Name | The VLAN name. The name is set when a VLAN is Inserted or modified. |
| Ports | The ports in the VLAN. The port list is set when a VLAN is Inserted. |
| State | The state of the VLAN—enabled or disabled. |
| Jumbo Frames | Enabled = the VLAN is configured for Jumbo frames. Disabled = the VLAN is not configured for Jumbo frames. |

About Spanning Tree bridge

Spanning Tree bridge parameters affect the global STP operation of the Web Switching Module. STP must be enabled for the STG before you can configure Spanning Tree bridge. For more information, see [“Configuring a Spanning Tree Group” on page 114](#).

STP bridge parameters include:

- Bridge aging time
- Bridge priority
- Bridge Hello interval
- Bridge maximum age
- Bridge forwarding delay

Use the following formulas as a guideline when configuring Spanning Tree bridge.

$$2(\text{Hello} + 1) \leq \text{Root Maximum Age}$$

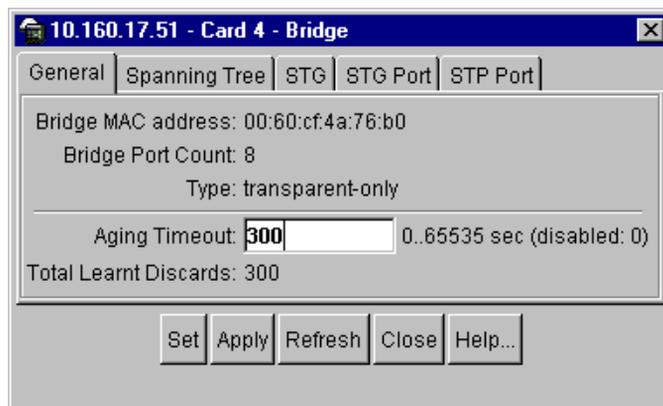
$$2(\text{Fwd} - 1) \geq \text{Root Maximum Age}$$

Configuring Spanning Tree bridge

To configure Spanning Tree bridge:

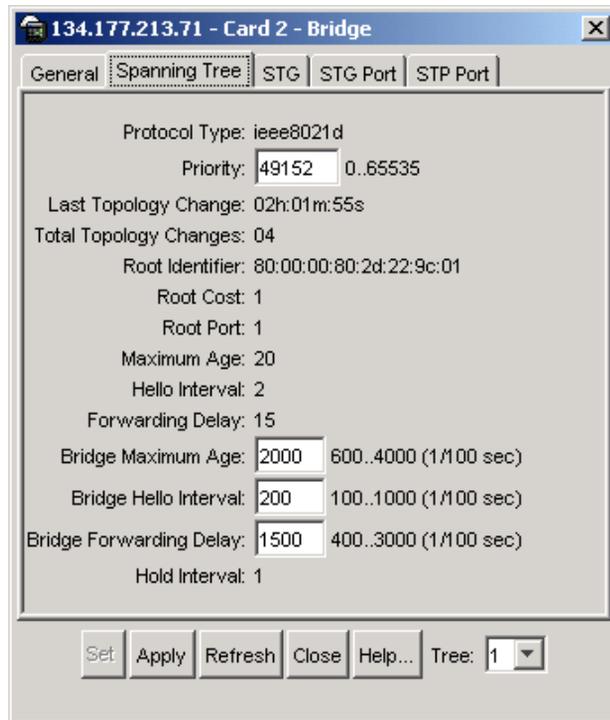
- 1 From the Device View, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, select Edit > WSM Card > Bridge.
The Bridge dialog box opens to the General tab (Figure 39).

Figure 39 Bridge—General tab



- 3 In the Aging Timeout field, type one of the following:
 - the amount of time (1 to 65535 seconds) you want the bridge to wait for a packet from a station before it removes the station from the forwarding database
 - 0 to disable aging
- 4 From the Bridge dialog box, select the Spanning Tree tab.

The [Spanning Tree tab](#) (Figure 40) opens with the fields defined in [Table 29](#) on page 121.

Figure 40 Bridge—Spanning Tree tab

5 In the Tree field, click the down arrow and select an STG.



Note: Spanning Tree Group 1 contains all configured VLANs until you assign individual VLANs to other Spanning Tree Groups. See [“Configuring a Spanning Tree Group” on page 114](#).

- 6 In the Priority field, type the bridge priority—0 to 65,535.
- 7 In the Bridge Maximum Age field, type the maximum amount of time (600 to 4,000 seconds) you want the bridge to wait to receive a BPDU before reconfiguring the STP network.
- 8 In the Bridge Hello Interval field, type the frequency (1 to 1,000 seconds) with which you want the root bridge to transmit a configuration BPDU.
- 9 In the Bridge Forwarding Delay field, type the amount of time (4 to 30 seconds) you want a bridge port to wait before it changes from the listening state to the learning state and from the learning state to the forwarding state.

10 Click Set > Apply > Close.

The bridge parameters are configured for this Spanning Tree Group.

11 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 28 describes the fields on the General tab.

Table 28 Bridge—General tab fields

| Field | Description |
|------------------------|---|
| Bridge MAC Address | The unique MAC address used by the bridge. It is recommended that the bridge address be the numerically smallest MAC address of all ports that belong to the bridge. The only requirement is that the MAC Address be unique. |
| Bridge Port Count | The number of ports controlled by this bridging entity. |
| Type | Indicates the type of bridging this bridge can perform. |
| Aging Timeout | Sets the time period (0 to 65,535 seconds) for aging dynamically-learned forwarding information. The default is 300 seconds. The aging time specifies the amount of time the bridge waits without receiving a packet from a station before removing the station from the forwarding database. The range is 1 to 65535 seconds, and the default is 300 seconds. To disable aging, set this parameter to 0. |
| Total Learned Discards | Displays the total number of Forwarding Database entries that have been discarded because of insufficient storage space in the Forwarding Database. If the counter increases, it indicates that the Forwarding Database is regularly becoming full which can degrade the performance of the subnetwork. If Learned Entry Discards shows a significant value but does not increase, it indicates the problem has occurred but is not persistent. |

Table 29 describes the fields on the Bridge—Spanning Tree tab.

Table 29 Bridge—Spanning Tree tab fields

| Field | Description |
|------------------------|--|
| Protocol Type | The current operating version of the Spanning Tree Protocol. |
| Priority | Sets the Bridge priority—0 to 65,535. The default is 32,768. Controls which bridge on the network is the STP root bridge. To make the WSM the root bridge, configure the bridge priority lower than all other switches and bridges on your network. The lower the value, the higher the bridge priority. |
| Last Topology Change | Time elapsed since a topology change was last detected by the bridge, displayed as days, hours:minutes:seconds. |
| Total Topology Changes | Total number of topology changes, detected by the bridge, since the management entity was last reset or initialized. |
| Root Identifier | The bridge identifier of the root of the spanning tree. This value is used as the Root Identifier parameter in all configuration bridge protocol data units (PDU) that were originated by this node. |
| Root Cost | Cost of the path to the root as seen from this bridge. |
| Root Port | Number of the port that offers the lowest-cost path from this bridge to the root bridge. |
| Maximum Age | Maximum age (in seconds) of Spanning Tree Protocol information learned from the network, on any port, before it is discarded. |
| Hello Interval | Time, in seconds, between the transmission of the configuration bridge PDUs by this node on any port. The time is measured when it is the root of the spanning tree (or trying to become so). The default is 300. |
| Forwarding Delay | Time, in seconds, for a port to change its spanning state when moving towards the Forwarding state. Forwarding Delay determines how long the port stays in each of the Listening and Learning states that precede the Forwarding state. This value is also used for aging all dynamic entries in the Forwarding Database, after a topology change has been initiated. The bridge uses this value, unless the bridge becomes the root. In that case, Forwarding Delay becomes the value that all bridges, including this one, will start using when this bridge becomes the root. |

Table 29 Bridge—Spanning Tree tab fields (continued)

| Field | Description |
|-------------------------|--|
| Bridge Maximum Age | <p>Sets the Maximum bridge age in seconds. When this bridge is acting as the root, all bridges use Maximum Age for Maximum Information Age—600 to 4,000 seconds.</p> <p>The maximum age specifies the maximum time the bridge waits to receive a configuration bridge protocol data unit (BDPU) before it reconfigures the STP network. The range is 6 to 40 seconds, and the default is 20 seconds.</p> <p>The value of Maximum Age is an integer. An agent may return a badValue error if the input value is not a whole number.</p> |
| Bridge Hello Interval | <p>Sets the value, in seconds, that all bridges use for Hello Time when this bridge is acting as the root—100 to 1,000 seconds.</p> <p>The Hello Time specifies how often the root bridge transmits a configuration bridge protocol data unit (BPDU). Any bridge that is not the root bridge uses the root bridge hello value. The range is 1 to 10 seconds, and the default is 2 seconds</p> <p>The value of Bridge Hello Time Period is an integer. An agent may return a badValue error if the input value is not a whole number.</p> |
| Bridge Forwarding Delay | <p>Sets the state change delay value, in seconds as an integer. When this bridge is acting as the root, all bridges use Bridge State Change Delay for State Change Delay—400 to 3,000 seconds.</p> <p>The forward delay parameter specifies the amount of time that a bridge port has to wait before it changes from the listening state to the learning state and from the learning state to the forwarding state. The range is 4 to 30 seconds, and the default is 15 seconds.</p> <p>An agent may return a badValue error if the input value is not a whole number.</p> |
| Hold Interval | <p>The interval (in units of hundredths of a second) during which no more than two Configuration bridge PDUs shall be transmitted by this node.</p> |
| Tree | <p>Sets the Spanning Tree Group for this configuration.</p> <p>Spanning Tree Group 1 contains all configured VLANs until you assign individual VLANs to other Spanning Tree Groups. See “Configuring a Spanning Tree Group” on page 114.</p> |

Configuring Spanning Tree port

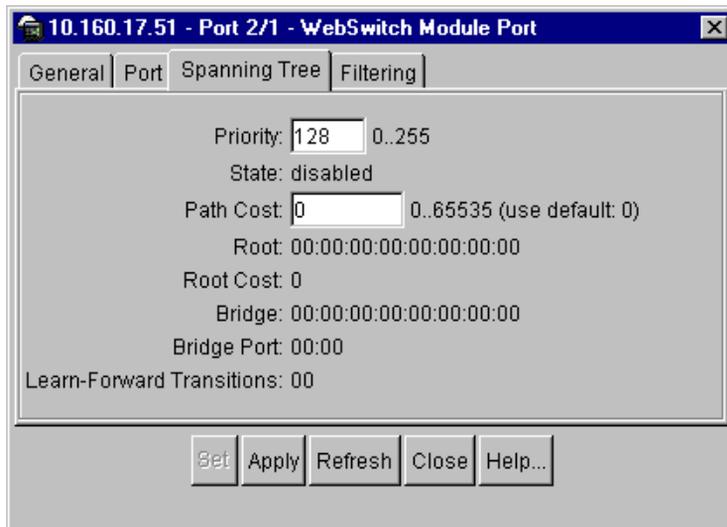
You can enable or disable Spanning Tree protocol on individual ports, and configure their port priority and path cost to help determine the designated port for a segment.

To configure spanning tree port:

- 1 From the Device View, select a port orientation (Front or Back).
- 2 Select a port.
The port is highlighted.
- 3 From the menu bar, choose Edit > WSM Card > Port/Back Port.
The Web Switching Module Port dialog box opens to the General tab.
- 4 Click the Spanning Tree tab.

The [Port—Spanning Tree tab](#) (Figure 41) opens with the fields described in [Table 30](#) on page 124.

Figure 41 Port—Spanning Tree tab



- 5 In the Priority field, type a priority (0 - 255) for the port. The port with the lowest port priority becomes the designated port for the segment.

- 6 In the Path Cost field, type a cost (0 - 65535) for the port path. The faster the port, the lower the cost.



Note: A setting of 0 tells the Web Switching Module to compute a Path Cost after the link speed has been autonegotiated.

- 7 Click Set > Apply > Close.

The port's priority and path cost are configured. To enable spanning tree on the port, see [“Enabling or disabling spanning tree on a port,”](#) next.

[Table 30](#) describes the fields on the Port Spanning Tree tab.

Table 30 Port Spanning Tree tab fields

| Field | Description |
|-----------|---|
| Priority | Sets the value of priority that determines which bridge is the root bridge—0 to 255. The default is 128 The port priority helps determine which bridge port becomes the designated port. In a network topology that has multiple bridge ports connected to a single segment, the port with the lowest port priority becomes the designated port for the segment. |
| State | Displays the port's current state: disabled, blocking, listening, learning, forwarding, or broken. |
| Path Cost | Sets the cost of the port path—1 to 65,535. The faster the port, the lower the cost. The default is 0. The port path cost is used to help determine the designated port for a segment. Generally speaking, the faster the port, the lower the path cost. The range is 1 to 65535. The default is 10 for 100Mbps ports, and 1 for gigabit ports. A setting of 0 indicates that the cost will be set to the appropriate default after the link speed has been auto negotiated. |
| Root | Displays the unique identifier of the bridge that is recorded as the root bridge in the configuration bridge protocol data units (BPDU). |
| Root Cost | Displays the path cost of the designated port of the segment that is connected to the port, and is compared to the Root Path Cost field in the received bridge protocol data units (BPDU). |
| Bridge | Displays the bridge identifier. The port interprets this bridge as the root bridge for the port's segment. |

Table 30 Port Spanning Tree tab fields (continued)

| Field | Description |
|---------------------------|---|
| Bridge Port | Displays the port identifier. This port is the designated bridge for the port's segment. |
| Learn-Forward Transitions | Displays the number of times the port has transitioned from the learning state to the forwarding state. |

Enabling or disabling spanning tree on a port

To enable or disable spanning tree on a port:

- 1 From the Device View, click the Web Switching Module.

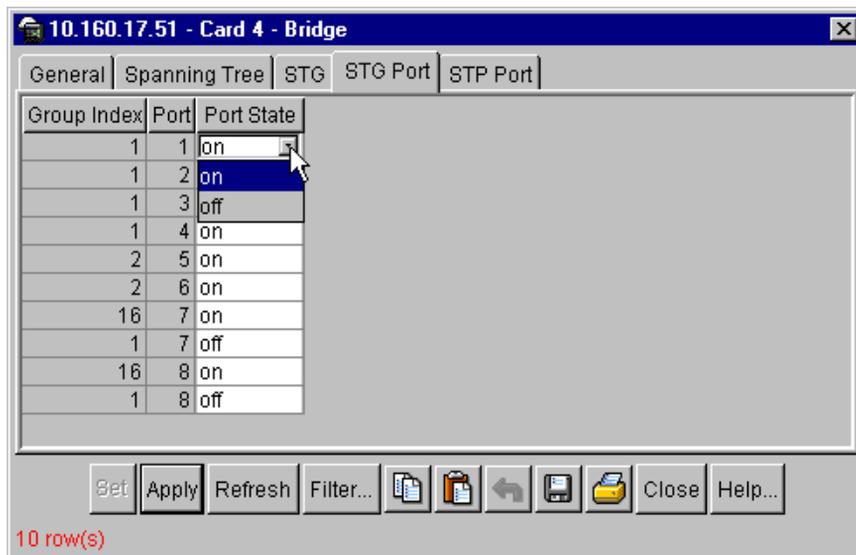
The Web Switching Module is highlighted.

- 2 From the Edit menu, select Edit > WSM Card > Bridge.

The Bridge dialog box opens to the General tab ([Figure 39 on page 118](#)).

- 3 Click the STG Port tab.

The [STG Port tab](#) ([Figure 42](#)) opens with the fields described in [Table 31 on page 126](#).

Figure 42 Bridge—STG Port tab

- 4 To enable (or disable) STP on the port, click its Port State field, and choose On (or Off) from the drop down list.
STP is enabled (or disabled) on the port.
- 5 Click Set > Apply > Close.
STP is enabled (or disabled) for the port and the Bridge dialog box closes.
- 6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

For more information about configuring ports, see [“Configuring ports” on page 83](#).

[Table 31](#) describes the fields on the bridge STG Port tab. .

Table 31 Bridge STG Port tab fields

| Field | Description |
|-------------|--|
| Group Index | The Spanning Tree Group (STG) index number associated with the port. There can be up to 16 Spanning Tree Groups, from 1 to 16. |
| Port | The port number that is associated with the Spanning Tree Group. |
| Port State | Sets the STP port state information—on or off. |

The fields described in this table apply to the Spanning Tree Group selected in the Tree drop-down menu at the bottom right of the window. Sixteen Spanning Tree Groups can be configured using this window.

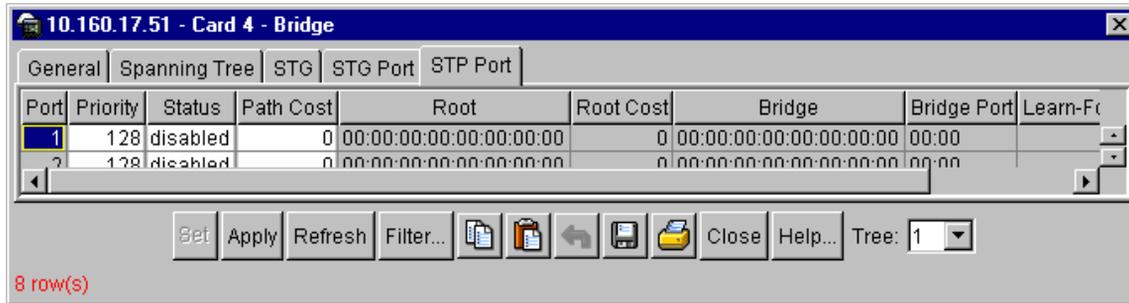
Figure 43 Bridge—STP Port tab

Table 32 describes the fields on the Bridge STP Port tab. .

Table 32 Bridge STP Port tab fields

| Field | Description |
|---------------------------|---|
| Port | The port number. |
| Priority | The value of priority that determines which bridge is the root bridge: 0 to 255. |
| Status | The port's current state: enabled, disabled, blocking, listening, learning, forwarding, or broken. |
| Path Cost | The cost of the port path (1 to 65,535). The faster the port, the lower the cost. The default is 0. |
| Root | The unique identifier of the bridge that is recorded as the root bridge in the configuration bridge protocol data units (BPDU). |
| Root Cost | The cost of the path to the root as seen from this bridge. |
| Bridge | The Bridge Identifier of the bridge that this port considers to be the Designated Bridge for this port's segment. |
| Bridge Port | The Port Identifier of the port on the Designated Bridge for this port's segment. |
| Learn-Forward Transitions | The number of times this port has transitioned from the Learning state to the Forwarding state. |

Chapter 4

IP routing configuration

This chapter describes how to configure IP routing for the WSM. It includes the following topics:

- [“Configuring IP interfaces,”](#) next
- [“Configuring default gateways”](#) on page 132
- [“Enabling IP forwarding”](#) on page 136
- [“Viewing IP routes”](#) on page 137
- [“Configuring domain name system servers”](#) on page 139
- [“Defining IP Address Ranges for the Local Route Cache”](#) on page 141
- [“Configuring routing information protocol”](#) on page 143
- [“Configuring IP forwarding per port”](#) on page 146
- [“Configuring static routes”](#) on page 148
- [“Configuring virtual router redundancy protocol”](#) on page 151

Configuring IP interfaces

IP interfaces define the subnets the WSM belongs to. Up to 255 IP interfaces can be configured on the WSM. The IP address assignments for your IP interfaces provide the WSM with an IP presence on your network. No two IP interfaces can be on the same IP subnet. Interfaces connect to the WSM for remote configuration, and for routing between subnets and VLANs.

If available on your network, a Bootstrap Protocol Relay (BOOTP) server can supply the WSM with IP interface parameters using Dynamic Host Configuration Protocol (DHCP) so that you do not have to enter them manually. BOOTP must be disabled before you can manually configure an IP interface.

For information about IP statistics gathering and analysis, see [“IP routing statistics”](#) on page 551.

Manually configuring an IP interface

To manually configure an IP interface:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the [General](#) tab.

- 3 From the IP Routing dialog box, click the Interfaces tab.

The [Interfaces](#) tab ([Figure 44](#)) opens with the fields described in [Table 33](#) on page 132.



Note: If you have many interfaces configured and wish to view a subset of this table, see [“Filtering table content”](#) on page 73.

Figure 44 Interfaces tab

| Interface Number | IP Address | IP Subnet Mask | IP Broadcast Address | VLAN | State | BOOTP Relay |
|------------------|--------------|----------------|----------------------|------|----------|-------------|
| 2 | 172.21.8.10 | 255.255.0.0 | 172.21.255.255 | 2 | enabled | enabled |
| 3 | 39.1.2.10 | 255.0.0.0 | 39.255.255.255 | 3 | enabled | enabled |
| 4 | 4.4.4.10 | 255.0.0.0 | 4.255.255.255 | 4 | enabled | enabled |
| 6 | 3.3.3.10 | 255.0.0.0 | 3.255.255.255 | 6 | enabled | enabled |
| 10 | 10.10.10.10 | 255.0.0.0 | 10.255.255.255 | 2 | enabled | enabled |
| 172 | 172.21.12.17 | 255.255.0.0 | 172.21.255.255 | 1 | disabled | enabled |

- 4 From the Interfaces tab, click Insert.
- 5 The IP Routing, [Insert Interfaces dialog box](#) ([Figure 45](#)) opens with the fields described in [Table 33](#) on page 132.

Figure 45 IP Routing, Insert Interfaces dialog box

- 6 In the BOOTP Relay field, click Disabled to disable DHCP.
- 7 In the corresponding fields, type the interface number (1- 255), IP address, IP subnet mask, IP broadcast address.
- 8 To assign a VLAN, in the VLAN field, type the VLAN number, or click Browse to open a selection dialog box of available VLANs.

For more information, see [“Browsing a read only dialog box” on page 76](#).

- 9 In the State field, click Enabled.
- 10 Click Insert.

The interface parameters are inserted into the Interfaces tab, and the dialog box closes.
- 11 From the Interfaces tab, click Set > Apply > Close.

The interface is configured and the dialog box closes.
- 12 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 33](#) describes the IP Routing Interfaces fields.

Table 33 Interfaces fields

| Field | Description |
|----------------------|--|
| Number | Displays the interface number of this interface—1 to 255. This number can be set when IP routing is Inserted. |
| IP Address | Configures the IP address of the interface. |
| IP Subnet Mask | Configures the subnet mask of the interface. |
| IP Broadcast Address | Configures the broadcast address of the interface. Note: Make sure that the broadcast address corresponds to the IP subnet mask that you entered. |
| VLAN | Configures the VLAN number for this interface. Each interface can belong to one VLAN, though any VLAN can have multiple IP interfaces in it. |
| State | Enables or disables the interface. |
| BOOTP Relay | Enables or disables the BootP relay. When enabled, the BootP relay allows obtaining configuration from the Dynamic Host Configuration Protocol (DHCP) server |

Configuring default gateways

The WSM can be configured with up to 250 gateways. Gateways 1 - 4 are used for default gateway load balancing. Gateways 5 - 250 are assigned to specific VLANs only (one gateway per VLAN). If a specific VLAN gateway is not available, then default gateways 1 - 4 are used.

To configure a default gateway:

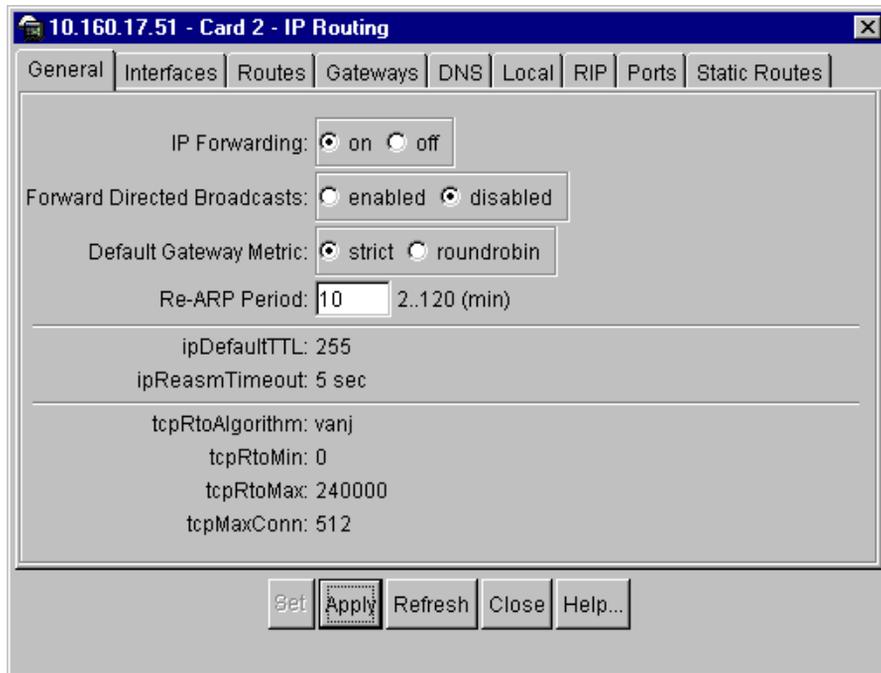
- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > IP Routing.

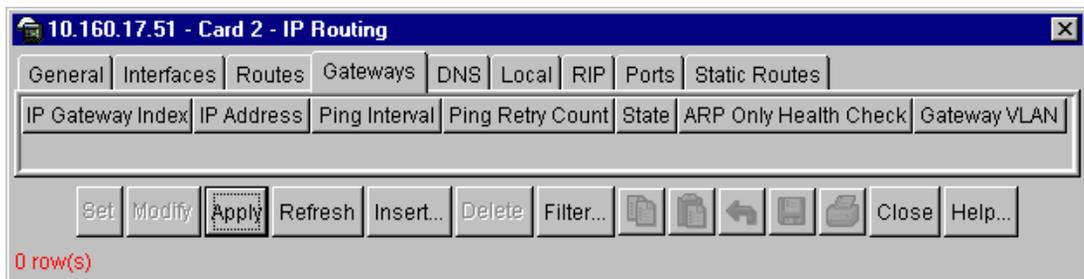
The IP Routing dialog box opens to the [General tab \(Figure 46\)](#).

See [Table 34 on page 135](#) for a description of the General tab fields.

Figure 46 IP Routing—General tab

- 3 In the **Default Gateway Metric field**, click either Strict or Roundrobin.
- 4 Click Set > Apply to apply the selection.
- 5 From the IP Routing dialog box, click the Gateways tab.

The **Gateways tab** (Figure 47) opens with the fields described in Table 35 on page 136.

Figure 47 Gateways tab



Note: If you have many gateways configured and wish to view a subset of this table, see [“Filtering table content”](#) on page 73.

- 6 From the Gateways tab, click Insert.
- 7 The [IP Routing, Insert Gateways dialog box](#) (Figure 48) opens with the fields described in [Table 35](#) on page 136.

Figure 48 IP Routing, Insert Gateways dialog box

- 8 In the corresponding fields, type the IP Gateway Index (1- 250), IP address, Ping Interval (0 - 60), and Ping Retry Count (1 - 120).
- 9 To enable the gateway, click Enabled in the State field.
- 10 To enable address resolution protocol (ARP) health check on the gateway, click Enabled in the ARP Only Health Check field.
- 11 To assign a VLAN to the Gateway, in the VLAN field, type a VLAN number, or click Browse to open a selection list of available VLANs.

For more information, see [“Browsing a read only dialog box”](#) on page 76.

- 12 Click Insert.
The default gateway is entered into the Gateways tab, and the Insert Gateways dialog box closes.
- 13 Click Set > Apply > Close.

The Gateway is configured and the IP Routing dialog box closes.

- 14** To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 34](#) describes the fields on the IP Routing General tab.

Table 34 IP Routing—General tab fields

| Field | Description |
|-----------------------------|--|
| IP Forwarding | Sets IP Forwarding: on or off for the Web Switching Module. |
| Forward Directed Broadcasts | Enables or disables forward directed broadcasts. |
| Default Gateway Metric | <p>Sets the metric—strict or roundrobin.</p> <p>Strict: The gateway number determines its level of preference. Gateway #1 acts as the preferred default IP gateway until it fails or is disabled, at which point the next in line will take over as the default IP gateway.</p> <p>Round robin: This provides basic gateway load balancing. The WSM sends each new gateway request to the next healthy, enabled gateway in line. All gateway requests to the same destination IP address are resolved to the same gateway.</p> |
| Re-ARP Period | Sets the Re-ARP period—2 to 120 minutes. |
| Default Time to Live | The default time-to-live which is the number of router hops the IP datagram can make. |
| Reassembly Timeout | The reassembly time-out. |
| Timeout Algorithm | The time-out algorithm. |
| Minimum Retransmit Timeout | The minimum retransmit time-out. |
| Maximum Retransmit Timeout | The maximum retransmit time-out. |
| Maximum Connection Allowed | The maximum number of allowed connections. |

Table 35 describes the IP Routing Gateways fields.

Table 35 Gateways fields

| Field | Description |
|-----------------------|--|
| IP Gateway Index | The gateway index number—1 to 250. Default gateways 1 to 4 are used for load balancing. Gateways 5 and higher can be assigned to separate VLANs for optimizing resources and segregating traffic. |
| IP Address | Sets the IP address of the default gateway. The default is 0.0.0.0. |
| Ping Interval | The WSM pings the default gateway to verify that it's up. The ping interval sets the time between health checks. The range is from 1 to 120 seconds. The default is 2 seconds. |
| Ping Retry Count | Sets the number of failed health check attempts required before declaring this default gateway inoperative. The range is from 1 to 120 attempts. The default is 8 attempts. |
| State | Enables or disables the default gateway. The default is disabled. |
| ARP Only Health Check | Enables or disables Address Resolution Protocol (ARP) health checks. This command is disabled by default. |
| Gateway VLAN | The VLAN number associated with this default IP gateway. You cannot assign a gateway VLAN number to the default gateways 1-4. You can only assign a Gateway VLAN to any IP Gateway Index of 5 and higher. Select the Browse... button to choose a VLAN. |

Enabling IP forwarding

When you configure the IP interfaces for the subnets attached to your WSM, IP routing between them can be performed entirely within the WSM. This eliminates the need to bounce inter-subnet communication off an external router device. Routing on more complex networks, where subnets may not have a direct presence on the WSM, can be accomplished through configuring static routes or by letting the WSM learn routes dynamically.

To enable IP forwarding:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the [General tab \(Figure 46\)](#) with the fields described in [Table 34 on page 135](#).

- 3 In the IP Forwarding field, click On.
- 4 To enable forwarding of directed broadcasts, click Enabled in that field.
- 5 Click Set > Apply > Close.

IP forwarding is enabled on the Web Switching Module and the IP Routing dialog box closes.

- 6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Viewing IP routes

To view the IP routes for the WSM:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the [General tab](#).

- 3 From the IP Routing dialog box, click the Routes tab.

The [Routes tab \(Figure 49\)](#) opens displaying the configured routes.

See [Table 36 on page 138](#) for a description of the Routes tab fields.

Figure 49 Routes tab

| Dest | Interface | NextHop | Type | Proto | Age | Mask |
|----------------|-----------|----------------|--------|-------|----------|-----------------|
| 10.0.0.0 | 10 | 10.10.10.10 | direct | local | 69156990 | 255.0.0.0 |
| 172.21.0.0 | 2 | 172.21.8.10 | direct | local | 69157074 | 255.255.0.0 |
| 172.31.255.240 | 256 | 172.31.255.246 | direct | local | 69157049 | 255.255.255.240 |



Note: If you have many routes configured and wish to view a subset of this table, see [“Filtering table content”](#) on page 73.

[Table 36](#) describes the fields on the Routes tab.

Table 36 Routes tab fields

| Field | Description |
|------------------------|---|
| Destination IP Address | The destination IP address of this route. The default value is 0.0.0.0. Multiple routes to a single destination can appear in the table if the Destination IP Address has been defined by the Network Management Protocol. |
| Interface | The number of the Interface. |
| Next Hop | The IP address of the next hop of this route. If a route bound to an interface is through a broadcast media, Next Hop Address is the agent's IP address on that interface. |
| Type | The type of route—direct, indirect, invalid, or other. The type invalid disassociates both the destination and the route entry that are identified with this entry. Management stations must be prepared to receive information from agents that correspond to entries that are not currently in use. |
| Protocol | The routing mechanism through which this route was learned. Including values for gateway routing protocols does not indicate that hosts will also support those protocols. |

Table 36 Routes tab fields (continued)

| Field | Description |
|----------------|---|
| Age | Number of seconds since this route was last updated. |
| IP Subnet Mask | The mask that must be logically ANDed with the destination address before it is compared to the destination address of the router. If the value of the destination address is 0.0.0.0 (default value), the mask value is also 0.0.0.0. If the system does not support arbitrary subnet masks, an agent constructs the router mask based on the class of the destination address' network—255.0.0.0 for class A; 255.255.0.0 for class B; 255.255.255.0 for class C. |

Configuring domain name system servers

You can define the primary and secondary domain name system (DNS) servers on your network, and set the default domain name served by the WSM services. DNS parameters must be configured prior to using hostname in ping, traceroute, and tftp commands.

To configure DNS servers for an existing IP route:

- 1 From the Device view, select the Web Switching Module.

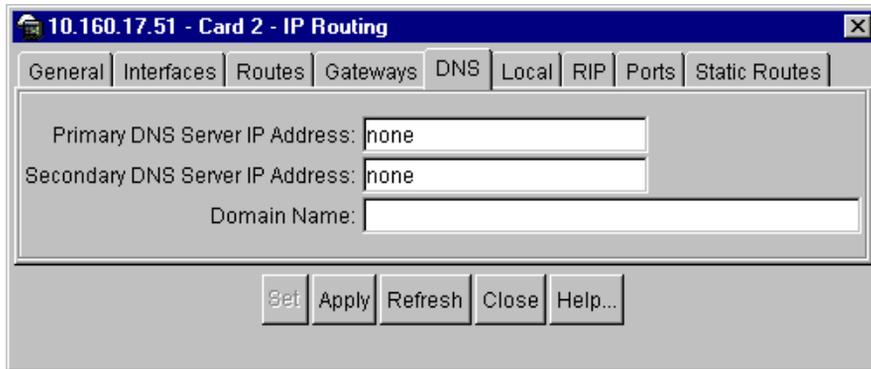
The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the [General tab](#).

- 3 Click the DNS tab.

The [DNS tab](#) ([Figure 50](#)) opens with the fields described in [Table 37](#) on [page 140](#).

Figure 50 DNS tab

- 4 In the corresponding fields, type IP addresses for the primary and secondary DNS servers, and a domain name.
- 5 Click Set > Apply > Close.

The DNS servers are configured, and the dialog box closes.

- 6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 37](#) describes the fields on the DNS tab.

Table 37 DNS tab fields

| Field | Description |
|---------------------------------|--|
| Primary DNS Server IP Address | Sets the DNS primary IP address in the new configuration block, the most recent configurations via SNMP. The default is 0.0.0.0. |
| Secondary DNS Server IP Address | Sets the DNS secondary IP address in the new configuration block. If the primary DNS server fails, the configured secondary DNS server will be used instead. The default is 0.0.0.0. |
| Domain Name | Sets the DNS domain name in the new configuration block, for example, mycompany.com. The default is none. |

Defining IP Address Ranges for the Local Route Cache

The Local Route Cache lets you reduce the size of the ARP table on the WSM by defining a range of addresses that will be cached. The local network address defines the base IP address in the range which will be cached, and the local network mask is the mask which is applied to produce the range. To determine if a route should be added to the memory cache, the destination address is masked (bitwise AND) with the local network mask and checked against the local network address.

By default, the local network address and mask are both set to 0.0.0.0. This produces a range that includes all Internet addresses for route caching: 0.0.0.0 through 255.255.255.255.

Addresses to be cached are subnets that are directly connected to a configured interface. [Table 38](#) shows an example you could configure to limit the route cache to your local hosts.

Table 38 Local Routing Cache Address Ranges

| Local Host Address Range | Address | Mask |
|-----------------------------|------------|-------------|
| 0.0.0.0 - 127.255.255.255 | 0.0.0.0 | 128.0.0.0 |
| 128.0.0.0 - 255.255.255.255 | 128.0.0.0 | 128.0.0.0 |
| 205.32.0.0 - 205.32.255.255 | 205.32.0.0 | 255.255.0.0 |



Note: All addresses that fall outside the defined range are forwarded to the default gateway. The default gateways must be within range.

Adding and removing local networks

You can add and remove local networks by setting the local network address and netmask for the route cache.

To configure the route cache:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

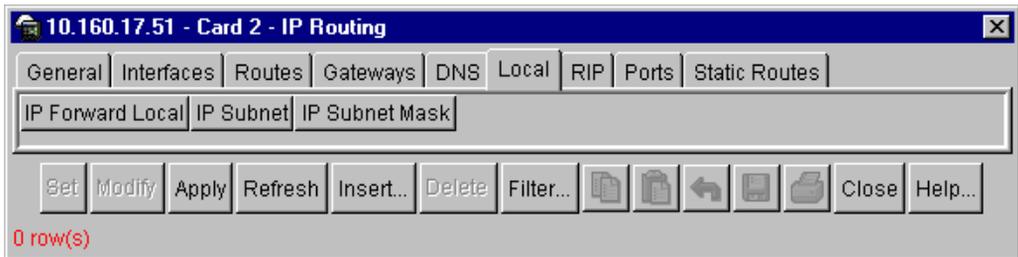
- 2 From the Edit menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the **General** tab.

- 3 Click the Local tab.

The **Local** tab (Figure 51) opens with the fields described in Table 39 on page 143.

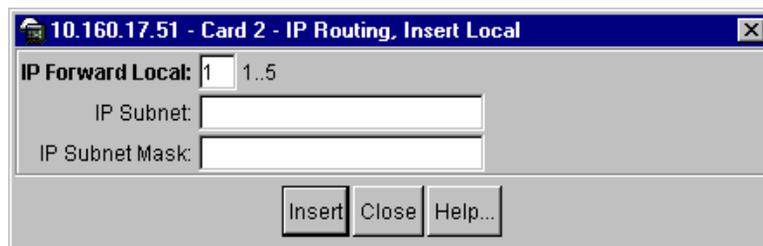
Figure 51 Local tab



Note: If you have many interfaces configured and wish to view a subset of this table, see “[Filtering table content](#)” on page 73.

- 4 From the Local tab, click Insert.
- 5 The **IP Routing, Insert Local** dialog box (Figure 52) opens with the fields listed in Table 39 on page 143.

Figure 52 IP Routing, Insert Local tab



- 6 In the corresponding fields, type the index number (1 - 5), the subnet, and the subnet mask of the local IP route (in dotted decimal format).

7 Click Insert > Close.

The Insert Local dialog box closes and the local route is entered into the Local tab of the IP Routing dialog box.

8 Click Set > Apply > Close.

The route cache is configured and the IP Routing dialog box closes.

9 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 39](#) describes the Local IP Configuration fields.

Table 39 Local IP Configuration fields

| Field | Description |
|------------------|--|
| IP Forward Local | The index number (1 - 5) of the local IP route which is set when a new IP route is Inserted. |
| IP Subnet | The subnet address of the local IP route. |
| IP Subnet Mask | The mask (entered in dotted decimal format) of the local IP route. |

Configuring routing information protocol

You can configure Routing Information Protocol, version 1 (RIP1) parameters for the WSM. This option is turned off by default.



Note: Do not configure RIP1 parameters if your routing equipment uses RIP version 2.

For information about RIP statistics gathering and analysis, see [“RIP statistics” on page 525](#).

To globally configure routing information protocol (RIP) for the WSM:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

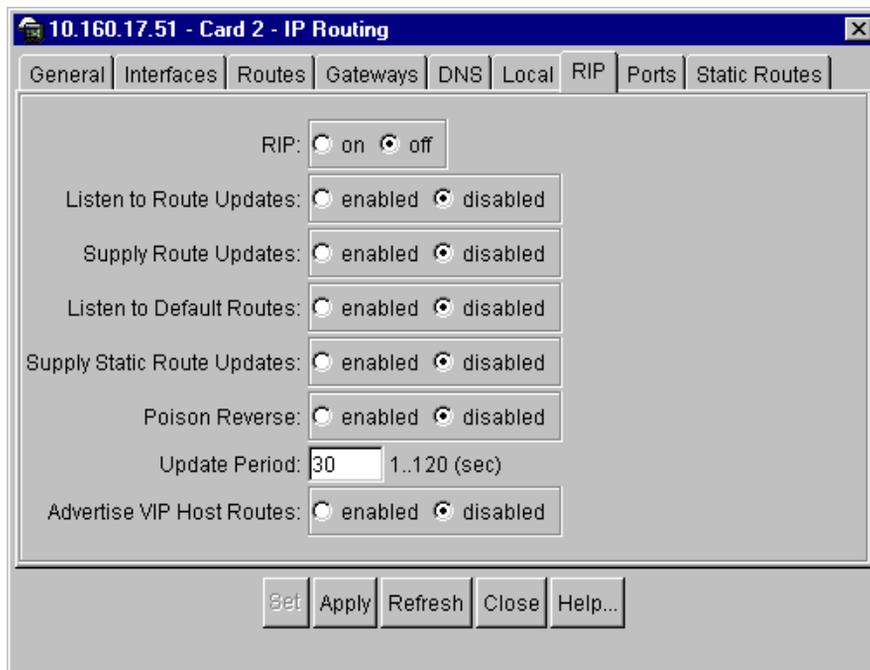
- 2 From the Edit menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the **General** tab.

- 3 Click the RIP tab.

The **RIP** tab (Figure 53) opens with the fields listed in Table 40 on page 145.

Figure 53 RIP tab



- 4 In the RIP field, click On.
- 5 Click Enable or Disable in other fields as required.
- 6 In the Update Period field, type an update period (1 - 120 seconds) for RIP updates.
- 7 To advertise virtual IP addresses as Host Routes, click Enabled in the Advertise VIP Host Routes field.

8 Click Set > Apply > Close.

RIP is configured globally for the WSM, and the IP Routing dialog box closes.

9 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 40 describes the fields on the RIP tab.

Table 40 RIP tab fields

| Field | Description |
|-----------------------------|---|
| RIP | Globally sets the Routing Information Protocol (RIP): on or off. The default is off. |
| Listen to Route Updates | Enables or disables listening to route updates. When enabled, the Web Switching Module learns routes from other routers. The default is disabled. |
| Supply Route Updates | Enables or disables supplying route updates. When enabled, the Web Switching Module supplies routes to other routers. The default is disabled. |
| Listen to Default Routes | Enables or disables listening to default routes. When enabled, the Web Switching Module accepts RIP default routes from other routers and gives them priority over configured default gateways. When disabled (default), the Web Switching Module rejects RIP default routes. |
| Supply Static Route Updates | Enables or disables supplying static route updates. When enabled the Web Switching Module supplies static routes. The default is disabled. |
| Poison Reverse | Enables or disables the RIP poison reverse. When enabled, the Web Switching Module uses split horizon with poisoned reverse. When disabled (default), the Web Switching Module only uses split horizon. |

Table 40 RIP tab fields (continued)

| Field | Description |
|---------------------------|---|
| Update Period | Sets the RIP update period—1 to 120 seconds. The default is 30 seconds. |
| Advertise VIP Host Routes | Enables or disables the advertisement of virtual IP addresses as Host Routes. If a VIP route exists in a routing table, it will always be advertised except when it is included in another network route that is already being advertised. Note: If all real servers behind a VIP go down, the route gets removed from the routing table, and will not be advertised. If all real servers are disabled, the VIP route is not eliminated from the routing table, and the WSM continues to advertise the route. |

Configuring IP forwarding per port

You can turn IP forwarding on or off on a port-by-port basis.

To configure IP forwarding per port:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

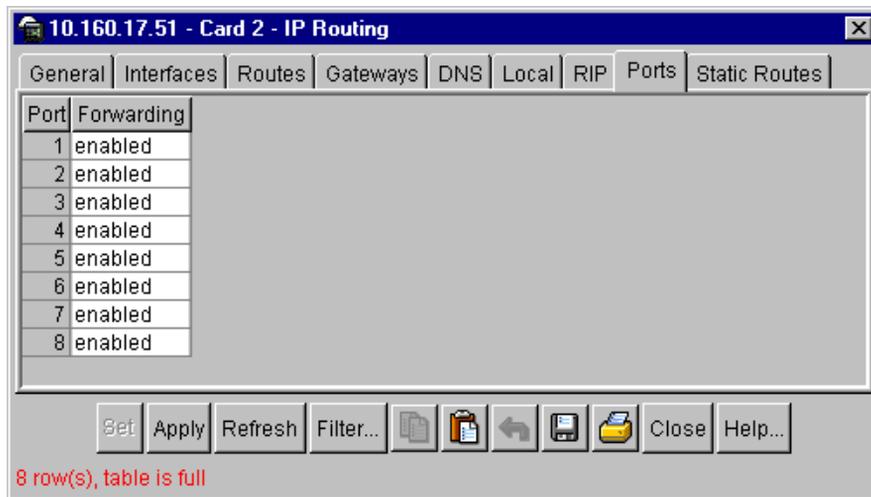
- 2 From the Edit menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the [General tab](#).

- 3 From the IP Routing dialog box, click the Ports tab.

The [Ports tab](#) ([Figure 54](#)) opens, displaying ports configured for IP forwarding.

See [Table 41 on page 147](#) for a description of the Ports tab fields.

Figure 54 IP Routing—Ports tab

- 4 Select a port from the list.
- 5 Click in the Forwarding field and choose Enabled/Disabled from the drop-down list.
- 6 Click Set > Apply > Close.

IP forwarding is configured for the port, and the dialog box closes.
- 7 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 41](#) describes the fields on the IP Routing Ports tab.

Table 41 IP Routing—Ports tab fields

| Field | Description |
|---------------|---|
| Port | Displays the port number of the forwarding state information. |
| IP Forwarding | Enables or disables the forwarding state of the port. |

Configuring static routes

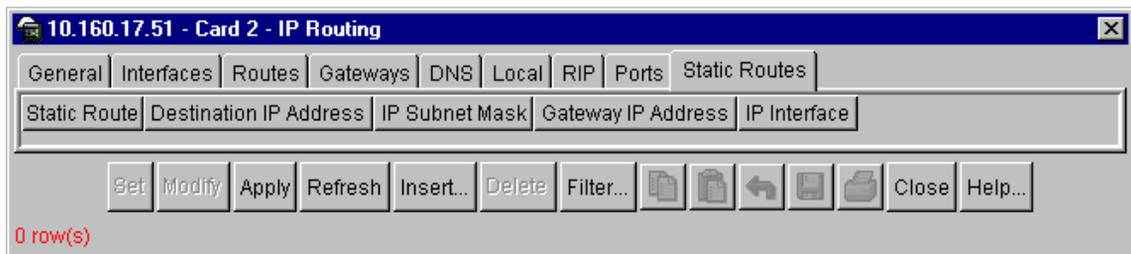
You can configure up to 128 static routes on the WSM.

To configure a static route:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, select WSM Card > IP Routing.
The IP Routing dialog box opens to the [General tab](#).
- 3 From the IP Routing dialog box, click the Static Routes tab.

The [Static Routes tab](#) ([Figure 55](#)) opens with the fields described in [Table 42 on page 150](#).

Figure 55 Static Routes tab



- 4 From the Static Routes tab, click Insert.

The [Insert Static Routes dialog box](#) ([Figure 56](#)) opens with the fields described in [Table 42 on page 150](#).

Figure 56 IP Routing, Insert Static Routes dialog box

The screenshot shows a dialog box titled "10.160.17.51 - Card 2 - IP Routing, Insert Static Routes". It has a standard Windows-style title bar with a close button. The main area contains the following fields and controls:

- Static Route:** A dropdown menu showing "1" and a text box containing "1..128".
- Destination IP Address:** A text box containing "none".
- IP Subnet Mask:** A text box containing "255.255.255.0".
- Gateway IP Address:** A text box containing "none".
- IP Interface:** A dropdown menu showing "1..255" and a "Browse..." button to its right.

At the bottom of the dialog box, there are three buttons: "Insert", "Close", and "Help...".

- 5 In the corresponding fields, type a static route number (1 - 128), destination IP address, IP subnet mask, gateway IP address.
- 6 In the IP Interface field, type an IP interface number (1 - 255), or click [Browse](#) to open a selection list of available interfaces. For more information, see [“Browsing a read only dialog box” on page 76](#).
- 7 Click Insert > Close.

The static route is inserted into the Static Routes tab, and the Insert Static Routes dialog box closes.
- 8 From the Static Routes tab, click Set > Apply > Close.

The static route is configured and the IP Routing dialog box closes.
- 9 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 42](#) describes the fields on the Static Routes tab and the Insert Static Routes dialog box.

Table 42 Static route configuration fields

| Field | Description |
|------------------------|---|
| Static Route | The index number of the static routing table—1 to 128. This number can be set when a new route is inserted. |
| Destination IP Address | Sets the destination IP address for this route. The default is 0.0.0.0. |
| IP Subnet Mask | Sets the subnet IP mask for this route. The default is 255.255.255.0. If the destination IP address is 0.0.0.0, the mask is also 0.0.0.0. |
| Gateway IP Address | Sets the destination IP gateway for this route. |
| IP Interface | Sets the IP interface number of the route that is used as the source IP for routing: 1 to 255. |

Deleting a static route

To delete a static route:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the [General tab](#).

- 3 Click the Static Routes tab.

The [Static Routes tab](#) ([Figure 55](#)) opens with the fields described in [Table 42 on page 150](#).

- 4 Select the static route you want to delete.

- 5 Click Delete > Close.

The static route is deleted and the dialog box closes.

- 6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.

- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Configuring virtual router redundancy protocol

VRRP eliminates single points of failure within a network by enabling redundant router LAN configurations that provide alternate router paths for a host. Each participating VRRP-capable routing device is configured with the same virtual router IP address and ID number.

One of the virtual routers is elected as the master, based on a number of priority criteria, and assumes control of the shared virtual router IP address. If the master fails, one of the backup virtual routers will take control of the virtual router IP address and actively process traffic addressed to it. Because the router associated with a given alternate path supported by VRRP uses the same IP address and MAC address as the routers for other paths, the host's gateway information does not change, no matter what path is used.

VRRP-based redundancy reduces administrative overhead because hosts need not be configured with multiple default gateways.

By default, VRRP is disabled. For VRRP configuration examples, see the High Availability chapter in the *Web OS Switch Software 10.0 Application Guide*, part number 212777-A.

Enabling virtual routing on the WSM

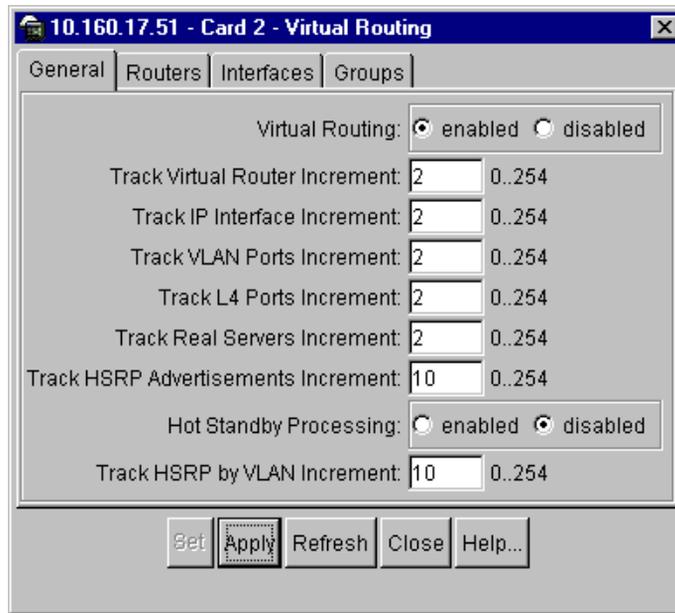
To enable virtual routing on the WSM:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > Virtual Routing.

The Virtual Routing dialog box opens to the [General tab \(Figure 57\)](#) with the fields described in [Table 43 on page 153](#).

Figure 57 Virtual Routing—General tab

3 In the Virtual Routing field, click Enabled.

4 Click Set > Apply > Close.

Virtual routing is enabled, and the Virtual Routing dialog box closes.

5 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 43 describes the fields on the General tab.

Table 43 Virtual Routing—General tab fields

| Field | Description |
|--------------------------------|---|
| Virtual Routing | Globally enables or disables VRRP operation. |
| Track Virtual Router Increment | Sets the increment of VRRP virtual router priority: 0 to 254. This priority is adjusted by tracking the state of other virtual routers. The value 254 provides maximum priority. The default is 2. |
| Track IP Interface Increment | Sets the increment of VRRP virtual router priority: 0 to 254. This priority is adjusted by tracking the number of IP interfaces active on the WSM. The default is 2. |
| Track VLAN Ports Increment | Sets the increment of VRRP virtual router priority: 0 to 254. The priority is adjusted by tracking the port state of those ports that belong to the same virtual LAN as the virtual router. The default is 2. |
| Track L4 Ports Increment | Sets the increment of VRRP virtual router priority: 0 to 254. The priority is adjusted by tracking the Layer 4 port states. This is valid when a virtual server is configured as a VRRP virtual router. The default is 2. |
| Track Real Servers Increment | Sets the increment of VRRP virtual router priority: 0 to 254. The priority is adjusted by tracking the state of real servers under the virtual server that is configured as a VRRP virtual router. The default is 2. |
| Track HSRP Increment | Sets the increment of VRRP virtual router priority: 0 to 254. The priority is adjusted by tracking the HSRP advertisements. The default is 10. |
| Hot Standby Processing | Enables or disables hot standby processing. Note: In a hot-standby configuration, Spanning Tree Protocol (STP) is not needed to eliminate bridge loops. |
| Track HSRP by VLAN Increment | Sets the priority increment value for VLAN port switch tracking: 0 to 254. The default is 2. |

Configuring a virtual router

You can configure up to 256 virtual routers for the WSM. A virtual router is defined by its virtual router ID and an IP address. Each VRRP-capable routing device participating in redundancy for this virtual router is configured to share the same virtual router ID and IP address.

To configure a virtual router:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

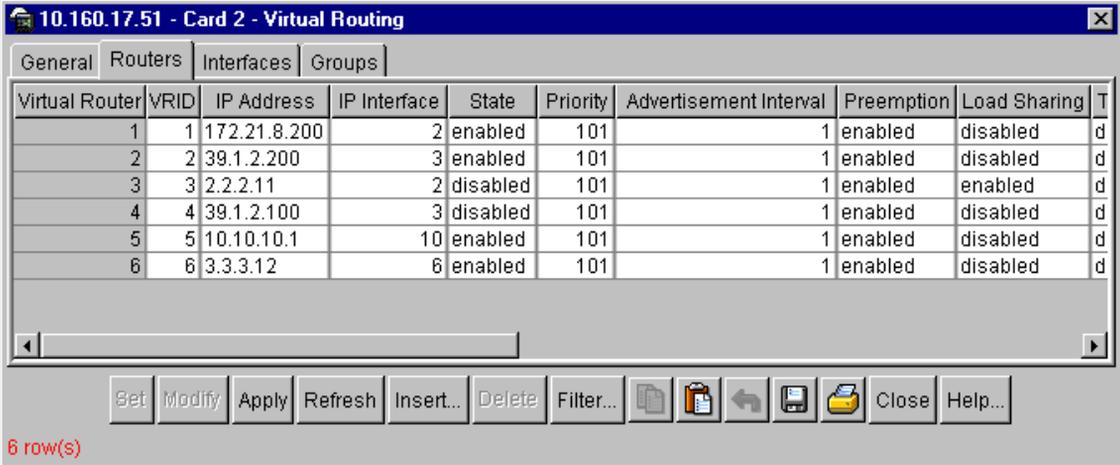
- 2 From the Edit menu, select WSM Card > Virtual Routing.

The Virtual Routing dialog box opens to the [General](#) tab.

- 3 Click the [Routers](#) tab.

The Routers tab ([Figure 58](#)) opens with the fields described in [Table 44](#) on [page 157](#).

Figure 58 Routers tab



| Virtual Router | VRID | IP Address | IP Interface | State | Priority | Advertisement Interval | Preemption | Load Sharing | T |
|----------------|------|--------------|--------------|----------|----------|------------------------|------------|--------------|---|
| 1 | 1 | 172.21.8.200 | 2 | enabled | 101 | 1 | enabled | disabled | d |
| 2 | 2 | 39.1.2.200 | 3 | enabled | 101 | 1 | enabled | disabled | d |
| 3 | 3 | 2.2.2.11 | 2 | disabled | 101 | 1 | enabled | enabled | d |
| 4 | 4 | 39.1.2.100 | 3 | disabled | 101 | 1 | enabled | disabled | d |
| 5 | 5 | 10.10.10.1 | 10 | enabled | 101 | 1 | enabled | disabled | d |
| 6 | 6 | 3.3.3.12 | 6 | enabled | 101 | 1 | enabled | disabled | d |

6 row(s)

- 4 Click Insert.

The Insert Routers dialog box ([Figure 59](#)) opens with the fields described in [Table 44](#) on [page 157](#).

Figure 59 Virtual Routing, Insert Routers dialog box

Virtual Router: 7 1..256

VRID: 1 1..255

IP Address: none

IP Interface: 1 1..255 Browse...

State: enabled disabled

Priority: 100 1..254

Advertisement Interval: 1 1..255

Preemption: enabled disabled

Load Sharing: enabled disabled

Track VRs: enabled disabled

Track IP Interfaces: enabled disabled

Track VLAN Ports: enabled disabled

Track L4 Ports: enabled disabled

Track Real Servers: enabled disabled

Track HSRP Advertisements: enabled disabled

Track HSRP by VLAN: enabled disabled

Insert Close Help...

- 5 In the Virtual Router field, type the index number (1 - 256) of the VRRP virtual router.
- 6 In the VRID field, type the numeric VRRP virtual router identifier (1 - 255).
- 7 In the IP Address field, type the IP address of the VRRP virtual router.
- 8 In the IP Interface field, type the index number (1 - 256) of the IP interface represented by the VRRP virtual router, or click **Browse** to open a selection list of available interfaces. For more information, see [“Browsing a read only dialog box” on page 76](#).

- 9** In the Priority field, type the priority number (1 - 254) for this VRRP virtual router.
- 10** In the Advertisement Interval field, type the interval (1 - 256) to use between VRRP advertisements.
- 11** In the Preemption field, click Enabled or Disabled to designate whether this VRRP router is to preempt a low priority Master.
- 12** In the Load Sharing field, click Enabled or Disabled to designate whether this WSM will process traffic addressed to this virtual router, even when in backup mode.
- 13** Click Insert > Close.

The VRRP router is inserted into the Routers tab, and the Insert Routers dialog box closes.

- 14** From the Routers tab, click Set > Apply > Close.

The settings are applied to the WSM and the Virtual Routing dialog box closes.

- 15** To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 44 describes the fields on the Virtual Routing Routers tab and the Insert Routers dialog box.

Table 44 Routers fields

| Field | Description |
|----------------|---|
| Virtual Router | The index number of the VRRP virtual router: 1 to 256. |
| VRID | <p>Defines the virtual router ID. This is used in conjunction with the IP Address (below) to define a virtual router on this WSM.</p> <p>To create a pool of VRRP-enabled routing devices which can provide redundancy to each other, each participating VRRP device must be configured with the same virtual router: one that shares the same VRID and IP Address combination.</p> <p>The VRID for standard virtual routers (where the virtual router IP address is not the same as any virtual server) can be any integer between 1 and 255. The default value is 1. All VRID values must be unique within the VLAN to which the virtual router's IP interface belongs.</p> |
| IP Address | Defines the IP address for this virtual router using dotted decimal notation. IP Address is used in conjunction with the VRID (above) to configure the same virtual router on each participating VRRP device. The default address is 0.0.0.0. |
| IP Interface | Selects a WSM IP interface (between 1 and 256). If the IP interface has the same IP address as the IP Address option above, this WSM is considered the "owner" of the defined virtual router. An owner has a special priority of 255 (highest) and will always assume the role of master router, even if it must preempt another virtual router which has assumed master routing authority. This preemption occurs even if the Preemption option below is disabled. The default value is 1. |
| State | Enables or disables the VRRP virtual router. The default is disabled. |
| Priority | Defines the election priority bias for this virtual router. This can be any integer between 1 and 254. The default value is 100. During the master router election process, the routing device with the highest virtual router priority number wins. If there is a tie, the device with the highest IP interface address wins. If this virtual router's IP Address is the same as the one used by the IP interface, the priority for this virtual router will automatically be set to 255 (highest). When priority tracking is used (below), this base priority value can be modified according to a number of performance and operational criteria. |

Table 44 Routers fields (continued)

| Field | Description |
|------------------------|---|
| Advertisement Interval | Defines the time interval between VRRP master advertisements. This can be any integer between 1 and 255 seconds. The default value is 1. |
| Preemption | <p>Enables or disables a higher priority Backup VRRP virtual router to preempt a low-priority Master. The default is enabled.</p> <p>Enables or disables master preemption. When enabled, if this virtual router is in backup mode but has a higher priority than the current master, this virtual router will preempt the lower priority master and assume control. Note that even when Preemption is disabled, this virtual router will always preempt any other master if this WSM is the owner (the IP interface address and virtual router IP Address are the same). By default, this option is enabled.</p> |
| Load Sharing | <p>Enables or disables virtual router sharing.</p> <p>When enabled, this WSM will process any traffic addressed to this virtual router, even when in backup mode. By default, this option is enabled.</p> |
| Track VRs | <p>Enables or disables tracking other virtual routers for priority adjustment. The default is disabled.</p> <p>When enabled, the priority for this virtual router will be increased for each virtual router in master mode on this WSM. This is useful for making sure that traffic for any particular client/server pairing are handled by the same WSM, increasing routing and load balancing efficiency.</p> |
| Track IP Interfaces | <p>Enables or disables tracking the state of IP interfaces active on the WSM. The default is disabled.</p> <p>When enabled, the priority for this virtual router will be increased for each other IP interface active on this WSM. An IP interface is considered active when there is at least one active port on the same VLAN. This helps elect the virtual routers with the most available routes as the master.</p> |
| Track VLAN Ports | <p>Enables or disables tracking the states of VLAN ports for priority adjustment. The default is disabled.</p> <p>When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. A port is considered "active" if it has a link and is forwarding traffic. This helps elect the virtual routers with the most available ports as the master.</p> |

Table 44 Routers fields (continued)

| Field | Description |
|---------------------------|---|
| Track L4 Ports | Enables or disables tracking the state of Layer 4 ports for priority adjustment. This is applied when the virtual server is configured as a VRRP virtual router. The default is disabled. When enabled for virtual server routers, the priority for this virtual router will be increased for each physical WSM port which has active Layer 4 processing on this WSM. This helps elect the main Layer 4 switch as the master. |
| Track Real Servers | Enables or disables tracking the states of real servers for priority adjustment. This is applied when the virtual server is configured as a VRRP virtual router. The default is disabled. When enabled for virtual server routers, the priority for this virtual router will be increased for each healthy real server behind the virtual server IP address of the same IP address as the virtual router on this WSM. This helps elect the switch with the largest server pool as the master, increasing Layer 4 efficiency. |
| Track HSRP Advertisements | Enables or disables tracking HSRP advertisements for priority adjustment. The default is disabled. Hot Standby Router Protocol (HSRP) is used with some types of routers for establishing router failover. In networks where HSRP is used, enable this option to increase the priority of this virtual router for each Layer 4 client-only port that receives HSRP advertisements. Enabling HSRP helps elect the switch closest to the master HSRP router as the master, optimizing routing efficiency. |
| Track HSRP by VLAN | Enables or disables tracking VLANs for priority adjustment. The default is disabled. Hot Standby Router on VLAN works in VLAN-tagged environments. If enabled, increments only that VRRP instance that is on the same VLAN as the tagged HSRP master flagged packet. |

Configuring VRRP authentication on an IP interface

You can configure VRRP authentication parameters for the IP interfaces used with virtual routers.

To configure an interface:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

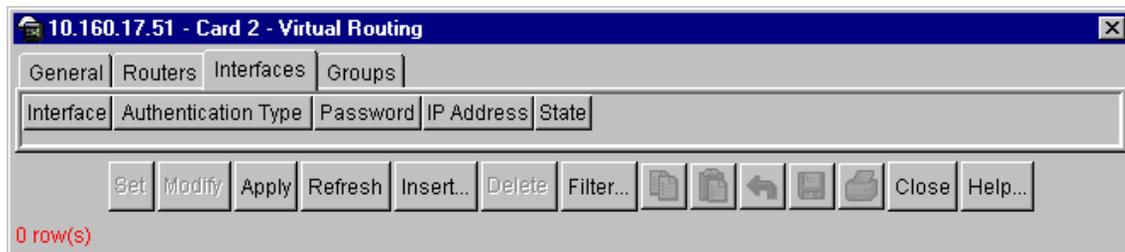
- 2 From the Edit menu, select WSM Card > Virtual Routing.

The Virtual Routing dialog box opens to the [General tab](#).

- 3 Click the Interfaces tab.

The Interfaces tab ([Figure 60](#)) opens with the fields described in [Table 45 on page 161](#).

Figure 60 Virtual Routing—Interfaces tab



- 4 Click Insert.

The [Insert Interfaces dialog box](#) ([Figure 61](#)) opens with the fields described in [Table 45 on page 161](#).

Figure 61 Virtual Routing—Insert Interfaces dialog box



- 5 In the Interface field, type the IP interface number (1 - 256) for which you are configuring authentication parameters.
- 6 In the Authentication field, click either None or Simple text password.
- 7 In the Password field, type a string of up to 8 characters.
- 8 Click Insert > Close.

The Interface is inserted into the Interfaces tab and the Insert Interfaces dialog box closes.

- 9 From the Interfaces tab, click Set > Apply > Close.

The Interface is configured and the Virtual Routing dialog box closes.

- 10 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 45 describes the Virtual Routing interfaces fields.

Table 45 Virtual Routing—Interfaces fields

| Field | Description |
|---------------------|--|
| Interface | The VRRP interface number: 1 to 256. |
| Authentication Type | Sets the type of authentication in use. The default is none. <ul style="list-style-type: none"> • none—No authentication • simple-text-password—use the specified password (Password) for authentication. |
| Password | Sets the password for authentication. The maximum string length is eight characters. The default is none. This password will be added to each VRRP packet transmitted by this interface when password authentication is chosen (see Authentication Type above). |
| IP Address | Displays the IP address. |
| State | Displays the current state. |

Configuring VRRP group parameters

You can associate all virtual routers into a single logical virtual router, which forces all virtual routers on the WSM to either be master or backup as a group. A virtual router is defined by its virtual router ID and an IP address. On each VRRP-capable routing device participating in redundancy for this virtual router, a virtual router will be configured to share the same virtual router ID and IP address.



Note: You must configure VRRP group parameters when using at least two WSMs in a hot-standby failover configuration, where only one WSM is active at any given time.

To configure VRRP group parameters:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

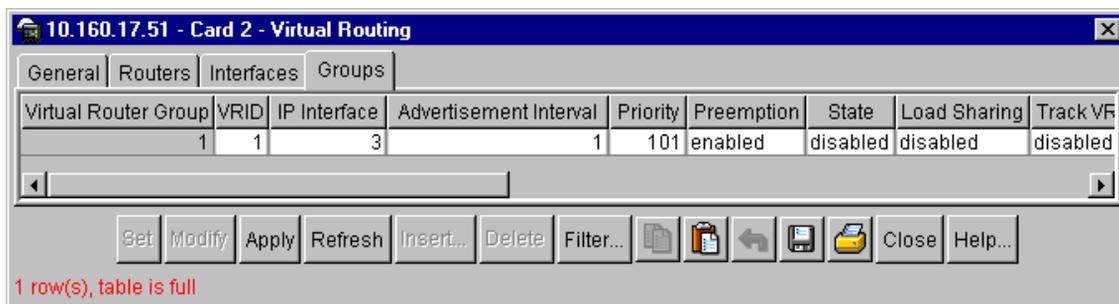
- 2 From the Edit menu, select WSM Card > Virtual Routing.

The Virtual Routing dialog box opens to the [General tab](#).

- 3 Click the Groups tab.

The [Groups tab](#) (Figure 62) opens with the fields described in [Table 46 on page 164](#).

Figure 62 Groups tab



- 4 Click Insert.

The Insert Groups dialog box (Figure 63) opens with the fields described in [Table 46 on page 164](#).

Figure 63 Virtual Routing, Insert Groups dialog box

Virtual Router Group: 1 1..1

VRID: 1 1..255

IP Interface: 1 1..255 Browse...

Advertisement Interval: 1 1..255

Priority: 100 1..254

Preemption: enabled disabled

State: enabled disabled

Load Sharing: enabled disabled

Track VRs: enabled disabled

Track IP Interfaces: enabled disabled

Track VLAN Ports: enabled disabled

Track L4 Ports: enabled disabled

Track Real Servers: enabled disabled

Track HSRP Advertisements: enabled disabled

Track HSRP Advertisements by VLAN: enabled disabled

Insert Close Help...

- 5 In the Virtual Router Group field, type the index number of the VRRP virtual router.
- 6 In the VRID field, type the virtual router group identifier (1 - 255).
- 7 In the IP Interface field, type the WSM IP interface number (1 - 255).
- 8 In the Advertisement Interval field, type an integer (1 - 255) to designate the interval in seconds between VRRP master advertisements.
- 9 In the Priority field, type an integer (1 - 254) to designate the election priority bias for this virtual router group.
- 10 In the Preemption field, click Enabled or Disabled to designate whether this VRRP router is to preempt a low priority Master.

11 In the Load Sharing field, click Enabled or Disabled to designate whether this WSM will process traffic addressed to this virtual router, even when in backup mode.

12 Click Insert > Close.

The VRRP router group is inserted into the Groups tab, and the Insert Groups dialog box closes.

13 From the Groups tab, click Set > Apply > Close.

The changes are applied to the WSM, and the Virtual Routing dialog box.

14 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 46](#) describes the fields on the Virtual Routing Groups tab and the Insert Groups dialog box.

Table 46 Groups fields

| Field | Description |
|------------------------|--|
| Virtual Router Group | The index number of the VRRP virtual router. |
| VRID | The VRRP virtual group identifier: 1 to 255. The default is 1. The VRID for standard virtual routers (where the virtual router IP address is not the same as any virtual server) can be any integer between 1 and 255. All VRID values must be unique within the VLAN to which the virtual router's IP interface (see IP Interface below) belongs. |
| IP Interface | Sets the IP Interface that the VRRP virtual group represents: 1 to 255. The default is 1. |
| Advertisement Interval | Defines the time interval between VRRP master advertisements. This can be any integer between 1 and 255 seconds. The default is 1. |

Table 46 Groups fields (continued)

| Field | Description |
|---------------------|--|
| Priority | <p>Defines the election priority bias for this virtual router group. This can be any integer between 1 and 254. The default value is 100.</p> <p>During the master router election process, the routing device with the highest virtual router priority number wins. If there is a tie, the device with the highest IP interface address wins. If this virtual router's IP address is the same as the one used by the IP interface, the priority for this virtual router will automatically be set to 255 (highest). When priority tracking is used, this base priority value can be modified according to a number of performance and operational criteria.</p> |
| Preemption | <p>Enables or disables whether a higher priority Backup VRRP virtual router can preempt a low priority Master. The default is enabled.</p> <p>When enabled, if the virtual router group is in backup mode but has a higher priority than the current master, this virtual router will preempt the lower priority master and assume control. Note that even when Preemption is disabled, this virtual router will always preempt any other master if this WSM is the owner (the IP interface address and virtual router address are the same).</p> |
| State | <p>Enables or disables the VRRP virtual router. The default is disabled.</p> |
| Load Sharing | <p>Enables or disables load sharing a non-master virtual router. The default is enabled.</p> <p>When enabled, this WSM will process any traffic addressed to this virtual router, even when in backup mode.</p> |
| Track VRs | <p>Enables or disables tracking other virtual routers for priority adjustment. The default is disabled.</p> <p>When enabled, the priority for this virtual router will be increased for each virtual router in master mode on this WSM. This is useful for making sure that traffic for any particular client/server pairing are handled by the same WSM, increasing routing and load balancing efficiency.</p> |
| Track IP Interfaces | <p>Enables or disables tracking of active IP interfaces on the WSM. The default is disabled.</p> <p>When enabled, the priority for this virtual router will be increased for each other IP interface active on this WSM. An IP interface is considered active when there is at least one active port on the same VLAN. This helps elect the virtual routers with the most available routes as the master.</p> |

Table 46 Groups fields (continued)

| Field | Description |
|-----------------------------------|--|
| Track VLAN Ports | <p>Enables or disables tracking the states of VLAN ports for priority adjustment. The default is disabled.</p> <p>When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. A port is considered “active” if it has a link and is forwarding traffic. This helps elect the virtual routers with the most available ports as the master.</p> |
| Track L4 Ports | <p>Enables or disables tracking the state of Layer 4 ports for priority adjustment. This is applied when the virtual server is configured as a VRRP virtual router. The default is disabled.</p> <p>When enabled for virtual server routers, the priority for this virtual router will be increased for each physical WSM port which has active Layer 4 processing on this WSM. This helps elect the main Layer 4 switch as the master.</p> |
| Track Real Servers | <p>Enables or disables tracking the states of real servers for priority adjustment. This is applied when the virtual server is configured as a VRRP virtual router. The default is disabled.</p> <p>When enabled for virtual server routers, the priority for this virtual router will be increased for each healthy real server behind the virtual server IP address of the same IP address as the virtual router on this WSM. This helps elect the switch with the largest server pool as the master, increasing Layer 4 efficiency.</p> |
| Track HSRP Advertisements | <p>Enables or disables tracking HSRP advertisements for priority adjustment. The default is disabled.</p> <p>Hot Standby Router Protocol (HSRP) is used with some types of routers for establishing router failover. In networks where HSRP is used, enable this option to increase the priority of this virtual router for each Layer 4 client-only port that receives HSRP advertisements. Enabling HSRP helps elect the switch closest to the master HSRP router as the master, optimizing routing efficiency.</p> |
| Track HSRP Advertisements by VLAN | <p>Enables or disables tracking VLANs for priority adjustment. The default is disabled.</p> <p>Hot Standby Router on VLAN works in VLAN-tagged environments. If enabled, increments only that VRRP instance that is on the same VLAN as the tagged HSRP master flagged packet.</p> |

Configuring virtual router priority tracking

You can configure a tracking function in the WSM that dynamically modifies the priority of a VRRP router based on its current state. The objective of tracking is to have, whenever possible, the master bidding processes for various virtual routers in a LAN converge on the same switch. Tracking ensures that the selected switch is the one that offers optimal network performance. For tracking to have any effect on virtual router operation, preemption must be enabled.



Note: Tracking only affects hot standby and active-standby configurations. It does not have any effect on active-active sharing configurations.

For more information about virtual router priority tracking, see the *Web OS Switch Software 10.0 Application Guide*, part number 212777-A.

Chapter 5

Network management and diagnostics

This section includes the following topics:

- [“About the SNMP agent,”](#) next
- [“Generating system messages”](#) on page 173
- [“Configuring the SNMP agent”](#) on page 173
- [“Configuring boot options”](#) on page 179
- [“About port mirroring”](#) on page 182
- [“Viewing device information”](#) on page 187

Supported software versions

The WSM is supported by following network management software versions:

- Optivity Switch Manager 2.0
- Optivity Network Management System 9.2

About the SNMP agent

The WSM architecture is a switch within a switch, and it retains its own SNMP agent. The SNMP interface is through the 8600 CPU proxy. A special SNMP community string selects a WSM agent. The SNMP agent on the WSM communicates to the management station through VLAN4093 on the 8600.

The 8600 is the WSM's trap host for receiving SNMP Traps. Traps are sent to the 8600 and then sent to the 8600's configured trap host.



Note: Due to SNMP incompatibility between in the Passport 8600 and the WSM, Device Manager displays an error message if you attempt to configure SNMP V2 on the Passport 8600 (config sys set snmp trap-recv xx.xx.xx.xx v2c public).

Alteon WebSystems enterprise MIBs supported

MIB definitions reside in Device Manager, and allow WSM SNMP functions (Get, Set, and Traps).

MIBs supported on the WSM include:

- alroot.mib
- tigonSwitch.mib
- tigonPhysical.mib
- tigonNetwork.mib
- tigonLayer4.mib
- tigonBwm.mib
- wsm4trap.mib

See the following Alteon WebSystems enterprise MIB documents for detailed WSM SNMP agent MIBs and trap definitions.

- Alroot.mib – Alteon product registrations, which are returned by sysObjectID.
- Altswitch.mib – Alteon enterprise MIB definitions.
- Alttrap.mib – Alteon enterprise trap definitions.

For information about SNMP statistics gathering and analysis, see [“SNMP statistics” on page 526](#).

SNMP standard MIBs supported

The WSM SNMP agent supports the following standard MIBs which reside in Device Manager:

- RFC 1213 - MIB II (System, Interface, Address Translation, IP, ICMP, TCP, UDP, SNMP Groups)
- RFC 1573 - MIB II Extension (IFX table)
- RFC 1643 - EtherLike MIB
- RFC 1493 - Bridge MIB
- RFC 1757 - RMON MIB (Statistics, History, Alarm, Event Groups)
- RFC 1724 RIP2 MIB
- RFC 1398 Ethernet-like MIB
- RFC 2037 Entity MIB.

SNMP generic traps supported

The WSM SNMP agent supports the following generic traps as defined in RFC 1215:

- ColdStart
- WarmStart
- LinkDown
- LinkUp
- AuthenticationFailure

SNMP Spanning Tree traps supported

The WSM SNMP agent supports two Spanning Tree traps as defined in RFC 1493:

- NewRoot
- TopologyChange

Supported SNMP traps

Table 47 describes supported SNMP traps.

Table 47 SNMP traps

| Field | Description |
|-------------------------------------|---|
| altSwDefGwUp | The default gateway defined is alive. |
| altSwDefGwDown | The default gateway defined is down. |
| altSwDefGwInService | The default gateway is up and in service. |
| altSwDefGwNotInService | The default gateway is alive but not in service. |
| altSwSibRealServerUp | The real server is up and operational. |
| altSwSibRealServerDown | The real server is down and out of service. |
| altSwSibRealServerMaxConnReached | The real server has reached maximum connections. |
| altSwSibBkupRealServerAct | The backup real server is activated due to availability of the primary real server. |
| altSwSibBkupRealServerDeact | The backup real server is deactivated due to the primary real server is available. |
| altSwSibBkupRealServerActOverflow | The backup real server is deactivated due to the primary real server is overflowed. |
| altSwSibBkupRealServerDeactOverflow | The backup real server is deactivated due to the primary real server is out from overflow situation. |
| altSwfltFilterFired | The packet received on a switch port matches the filter rule. |
| altSwSibRealServerServiceUp | The service port of the real server is up and operational. |
| altSwSibRealServerServiceDown | The service port of the real server is down and out of service. |
| altSwVrrpNewMaster | The sending agent has transitioned to 'Master' state. |
| altSwVrrpNewBackup | The sending agent has transitioned to 'Backup' state. |
| altSwVrrpAuthFailure | A packet has been received from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. Implementation of this trap is optional. |

Table 47 SNMP traps (continued)

| Field | Description |
|--------------------------|---|
| altSwLoginFailure | Someone failed to enter a valid username/password combination. |
| altSwSlibSynAttack | A SYN attack has been detected. |
| altSwTcpHoldDown | New TCP connection requests from a particular client will be blocked for a pre-determined amount of time since the rate of new TCP connections from that client has reached a pre-determined threshold. |
| altSwTempExceedThreshold | The WSM temperature has exceeded maximum safety limits. |

Configuring the SNMP agent

Device Manager supports the following for the Web Switching Module:

- [“Generating system messages” on page 173](#)
- [“Enabling or disabling syslog and SNMP traps” on page 175](#)
- [“Enabling or disabling authentication traps” on page 176](#)
- [“Configuring community strings” on page 178](#)

Generating system messages

You can configure the syslog facility to generate system messages from the WSM by configuring it to match the [syslog](#) value configured for the 8600 module. The [Syslog Trap tab](#) lists the types of messages generated.

To generate syslog messages from the WSM:

- 1 From the Device View, select the Web Switching Module.
The module is highlighted.
- 2 From the Edit menu, select WSM Card > General.
The Web Switching Module dialog box opens to the General tab.
- 3 Click the Syslog tab.

The **Syslog** tab (Figure 64) opens.

Figure 64 Web Switching Module—Syslog tab



- 4 On the Syslog tab, click the down arrow in the 1st Syslog Facility field and choose a local facility (0 - 7).
- 5 Click Set > Apply > Close.

The change is applied and the dialog box closes.
- 6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 48 describes the Syslog tab.

Table 48 Syslog tab

| Field | Description |
|---------------------|--|
| 1st Syslog Facility | Syslog Facility: messages are dumped from the 1st Syslog Host to the selected bucket: local0 to local7. The default is local0. |
| Console Output | Enables syslog output to the management console. |

Enabling or disabling syslog and SNMP traps

To enable or disable specific syslog and SNMP traps:

- 1 From the Device View, select the Web Switching Module.

The module is highlighted.

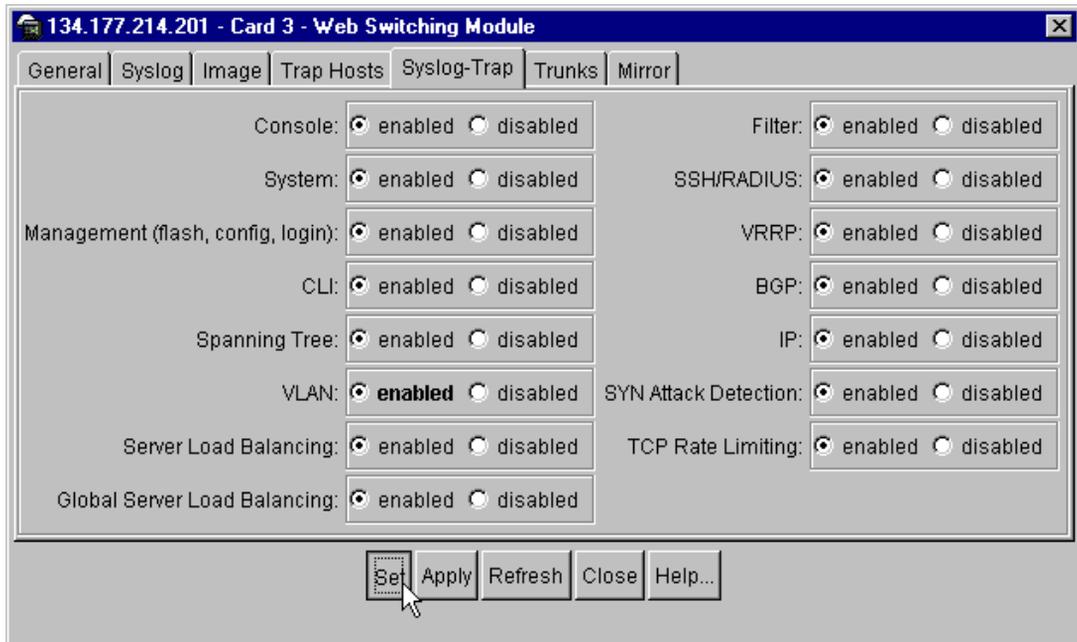
- 2 From the Edit menu, select WSM Card > General.

The Web Switching Module dialog box opens to the General tab.

- 3 Click the Syslog Trap tab.

The [Syslog Trap tab](#) (Figure 65) opens.

Figure 65 Syslog Trap tab



- 4 In the specific trap field, click enabled or disabled to set the Syslog trap.

- 5 Click Set > Apply > Close.

The syslog trap changes are applied and the dialog box closes.

- 6** To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 49 describes the traps that can be enabled or disabled from the Syslog Trap tab.

Table 49 Syslog trap fields

| Field | Enables or disables... |
|------------------------------|---|
| Console | Console syslog and SNMP trap. |
| System | System level syslog and SNMP trap. |
| Management | Management (flash, config, login) syslog and SNMP trap. |
| CLI | CLI generated error syslog and SNMP trap. |
| Spanning Tree | Spanning tree syslog and SNMP trap. |
| VLAN | VLAN syslog and SNMP trap. |
| Server Load Balancing | SLB syslog and SNMP trap. |
| Global Server Load Balancing | GSLB syslog and SNMP trap. |
| Filter | Filter syslog and SNMP trap. |
| SSH/Radius | SSH, RADIUS syslog and SNMP trap. |
| VRRP | VRRP syslog and SNMP trap. |
| BGP | BGP syslog and SNMP trap. |
| IP | IP related syslog and SNMP trap. |
| SYN Attack Detection | SYN Attack Detection syslog and SNMP trap. |
| TCP Rate Limiting | TCP rate limiting syslog and SNMP trap. |

Enabling or disabling authentication traps

To enable or disable authentication traps:

- 1** From the Device View, select the Web Switching Module.

The module is highlighted.

- 2 From the Edit menu, select WSM Card > General.

The Web Switching Module dialog box opens to the General tab.

- 3 In the Web Switching Module dialog box, click the Trap Hosts tab.

The **Trap Hosts tab** (Figure 66) opens.

Figure 66 Trap Hosts tab



- 4 In the Authentication Traps field, click the checkbox.

- 5 Click Set > Apply > Close.

The changes are applied and the dialog box closes.

- 6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 50 describes the Trap Hosts fields.

Table 50 Trap Hosts fields

| Field | Description |
|--------------------------------|---|
| 1st Trap Host Community String | Defines the SNMP community string to use with the trap host. The default is public. |
| Send Authentication Traps | Enables or disables the SNMP agent to generate authentication failure traps. Note: This object overrides any configuration information. All authentication failure traps can be disabled. |

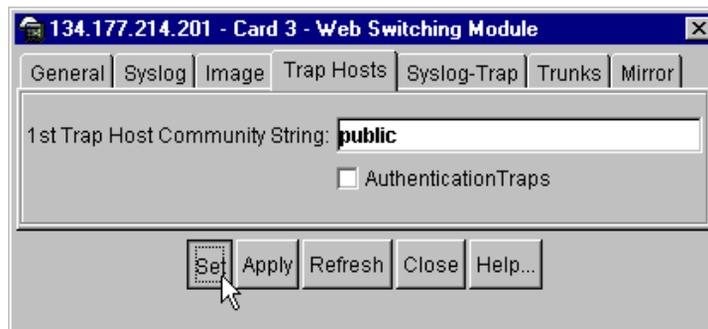
Configuring community strings

You can define the SNMP community string to use with the trap host. The default is public.

To configure the trap host community string:

- 1 From the Device View, select the Web Switching Module.
The module is highlighted.
- 2 From the Edit menu, select WSM Card > General.
The Web Switching Module dialog box opens to the General tab.
- 3 Click the Trap Hosts tab.
The [Trap Hosts tab](#) (Figure 67) opens.

Figure 67 Trap Hosts tab



4 In the 1st Trap Host Community String field, type a community string.

5 Click Set > Apply > Close.

The changes are applied and the dialog box closes.

6 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Configuring boot options

The WSM software image is the executable code running on the WSM. It ships with, and is pre-installed on, the WSM. As new versions of the image are released, you can upgrade the software running on your WSM. The WSM supports 2 unique boot images, identified as Image 1 and Image 2 on the Device Manager Image tab.

The Device Manager Image tab supports the following:

- [“Viewing current software image details” on page 179](#)
- [“Selecting a software image to run” on page 181](#)
- [“Selecting a configuration block to run” on page 182](#)

Viewing current software image details

To view details about the software image currently running for the Web Switching Module:

1 From the Device View, select the WSM.

The module is highlighted.

2 From the Edit menu, select WSM Card > General.

The Web Switching Module dialog box opens to the General tab.

3 In the Web Switching Module dialog box, click the Image tab.

The **Image tab** (Figure 68) opens, displaying image details for the WSM.

Figure 68 Image tab

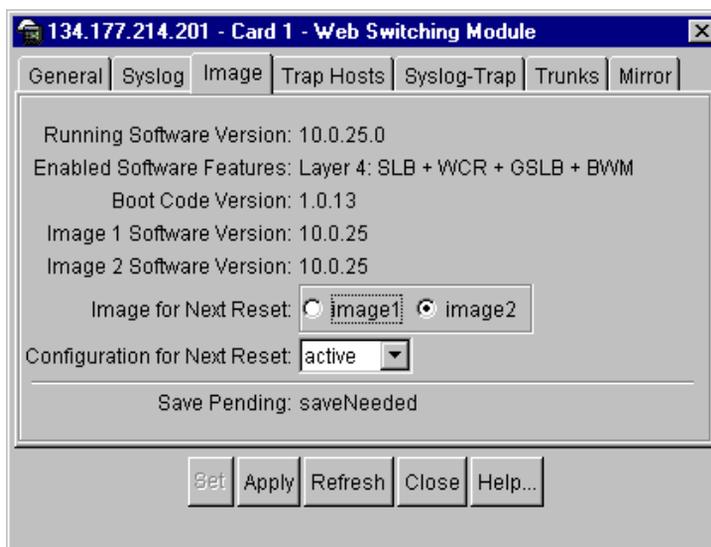


Table 51 describes the Image tab fields.

Table 51 Image tab fields

| Field | Description |
|---------------------------|---|
| Running Software Version | Displays the version of the software image currently running on the system in the following format: <i>major.minor.bugfix</i> (example: 1.1.2) If that information is not available, the field will be blank. |
| Enabled Software Features | Displays the enabled software features: <ul style="list-style-type: none"> • Bandwidth Management (BWM) • Global Server Load Balancing (GSLB) • Server Load Balancing (SLB) • Web Cache Redirection (WCR) |
| Boot Code Version | Displays the version of the boot code in the following format: <i>major.minor.bugfix</i> (example: 1.1.1) If information is not available, the field will be blank. |

Table 51 Image tab fields (continued)

| Field | Description |
|------------------------------|---|
| Image1 Software Version | Displays the version of the software image stored in image 1 in the following format: <i>major.minor.bugfix</i> (example: 1.0.0) If information is not available or if the software image is not valid, the field will be blank. |
| Image2 Software Version | The version of the software image stored in image 2 in the following format: <i>major.minor.bugfix</i> (example: 1.0.1) If information is not available or if the software image is not valid, the field will be blank. |
| Image For Next Reset | Sets the software image to boot for the next reset: image1 or image2. |
| Configuration For Next Reset | Sets the configuration information to load for the next reset: active, backup, or default. |
| Save Pending | Indicates if an action listed in Apply Pending needs to be saved. |

Selecting a software image to run

To select a software image to use when the WSM is next reset:

- 1 From the Device View, select the Web Switching Module.

The module is highlighted.

- 2 From the Edit menu, select WSM Card > General.

The Web Switching Module dialog box opens to the General tab.

- 3 In the Web Switching Module dialog box, click the Image tab.

The **Image tab** (Figure 68) opens with the fields listed in Table 51.

- 4 On the Image tab, click Image 1 or Image 2.

- 5 Click Set > Apply > Close.

The selected image will be loaded at the next reset, and the dialog box closes.

Selecting a configuration block to run

When you make configuration changes to the WSM, you must save them so that they are retained beyond the next time the WSM is reset. When saved, these changes are placed in the active configuration block. The previous configuration is copied into the backup configuration block.

The default factory configuration is the configuration set by the factory when your WSM was manufactured. It is saved in the factory configuration block. You may want to reset the WSM to use this default configuration if, for example, you move it to be reconfigured for a different purpose.

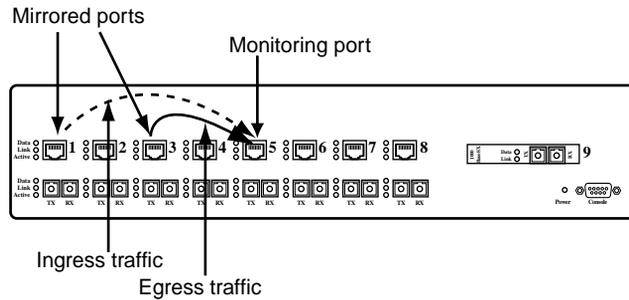
To select a configuration to use when the WSM is next reset:

- 1 From the Device View, select the Web Switching Module.
The module is highlighted.
- 2 From the Edit menu, select WSM Card > General.
The Web Switching Module dialog box opens to the General tab.
- 3 Click the Image tab.
The **Image tab** (Figure 68) opens with the fields listed in Table 51.
- 4 In the Configuration for Next Reset field, click the down arrow and choose a configuration (Active, Backup, or Default).
- 5 Click Set > Apply > Close.
When next reset, the WSM will use this configuration block.

About port mirroring

Port Mirroring lets you gather information about network performance and usage from a monitor port with an attached network analyzer. When Port Mirroring is enabled, the target port's network packets (sent and/or received) are duplicated and sent to the monitor port for analysis.

In Figure 69, a network analyzer is attached to port 5. Port 5 is monitoring both port 1 (*ingress* traffic, or traffic entering the WSM) and port 3 (*egress* traffic, or traffic leaving the WSM).

Figure 69 Port Mirroring

The WSM supports a single port monitoring one or more mirrored ports; but it does not support multiple ports monitoring a single mirrored port.

In the WSM, packets are duplicated and sent to the mirrored ports after client or server port processing (Layer 4 - 7) is completed. Data packets sent from a client to a virtual server are seen at the mirrored port as follows:

- source IP address = client IP address
- destination IP address = real server IP address rather than the virtual server IP address.

The response from the server to the client, however, is seen as follows:

- source IP address = virtual server IP address
- destination IP address = client IP address

Configuring port mirroring

To configure port mirroring:

- 1 From the Device View, select the Web Switching Module.
The module is highlighted.
- 2 From the Edit menu, select WSM Card > General.
The Web Switching Module dialog box opens to the General tab.
- 3 Click the Mirror tab.

The **Mirror** tab (Figure 70) opens with the fields listed in Table 52.

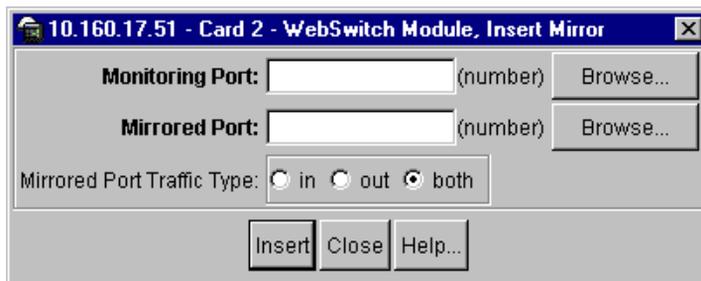
Figure 70 Mirror tab



- 4 Click insert.

The **Insert Mirror dialog box** (Figure 71) opens.

Figure 71 Insert Mirror dialog box



- 5 In the Monitoring port field, type the port number (1-6) of the monitoring port, or click Browse to view a selection list of available ports.
For information about browsing, see [“Browsing a read only dialog box” on page 76](#).
- 6 In the Mirrored Port field, type the port number (1-6) to be mirrored, or click Browse to view a selection list of available ports.



Note: Port Mirroring cannot be used simultaneously with Server Load Balance or Application Redirection on ports connected to a server directly, or through another switch or hub.

- 7 In the Mirrored Port Traffic Type field, specify which packets you want to be received by the monitoring port by clicking In, Out, or Both.
- 8 Click Insert > Close.
- 9 The Insert Mirror dialog box closes.
- 10 On the Mirror tab, click Set > Apply > Close.

The selected ports are configured for port mirroring and the dialog box closes.

[Table 52](#) describes the port mirror fields.

Table 52 Port Mirror fields

| Field | Description |
|----------------------------|---|
| Monitoring Port | Sets the physical port (1 to 6) for monitoring. Receives duplicated packets delivered by the Mirrored Port. Ports 7, 8, and 9 are not available for monitoring. |
| Mirrored Port | Sets the selected physical port (1 to 6) for mirroring. Packets received by or delivered from this port are delivered to the Monitoring Port. Ports 7, 8, and 9 are not available for mirroring. |
| Mirrored Port Traffic Type | Sets which packets are received by the Monitoring Port. <ul style="list-style-type: none"> • In = Packets received by the mirrored port. • Out = Packets transmitted from the mirrored port. • Both = Packets received by or transmitted from the mirrored port. |

Enabling port mirroring

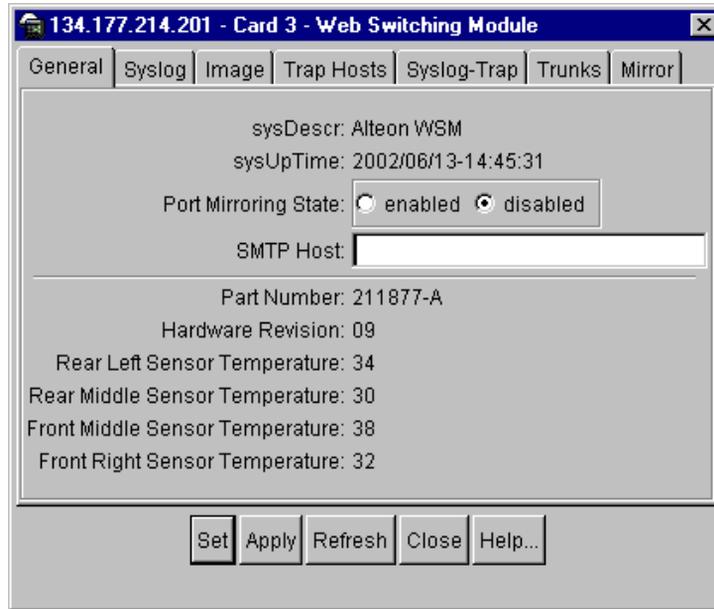
To enable port mirroring on the Web Switching Module:

- 1 From the Device View, select the Web Switching Module.

The module is highlighted.

- 2 From the Edit menu, select WSM Card > General.

The Web Switching Module dialog box opens to the [General tab](#) ([Figure 72](#)).

Figure 72 General tab

- 3 In the Port Mirroring State field, click enabled or disabled.
- 4 Click Set > Apply > Close.

Port mirroring is enabled for the Web Switching Module.

[Table 53](#) describes the fields on the General tab.

Table 53 General tab fields

| Field | Description |
|----------------------|---|
| sysDescr | A description of the opened device, such as product name and version, software operating system, and networking software. |
| sysUpTime | The date and time Device Manager was last initialized. |
| Port Mirroring State | Enables or disables port mirroring for the Web Switching Module. |
| SMTP Host | The domain name or IP address of an SMTP mail gateway. |
| Part Number | The hardware part number of the opened device. |

Table 53 General tab fields (continued)

| Field | Description |
|---------------------------------|---|
| Hardware Revision | The hardware revision. If the revision is not available, the field will be blank. |
| Rear Left Sensor Temperature | Rear left sensor temperature in degrees Celsius |
| Rear Middle Sensor Temperature | Rear middle sensor temperature in degrees Celsius |
| Front Middle Sensor Temperature | Front middle sensor temperature in degrees Celsius |
| Front Right Sensor Temperature | Front right sensor temperature in degrees Celsius |

Viewing device information

To view information about the Web Switching Module, such as part number and hardware revision:

- 1 From the Device View, select the Web Switching Module.

The module is highlighted.

- 2 From the Edit menu, select WSM Card > General.

The Web Switching Module dialog box opens to the [General tab \(Figure 72\)](#).

Chapter 6

Configuring filters for Layer 4 switching

This section contains the following topics:

- [“About Layer 4 filtering” on page 189](#)
- [“About default filters” on page 191](#)
- [“Configuring a default filter” on page 192](#)
- [“Creating a new filter” on page 201](#)

About Layer 4 filtering

You can use address and protocol specifications to configure up to 2048 traffic filters to allow, deny, redirect, or perform Network Address Translation (NAT). You can configure each physical WSM port to use any combination of filters. You can also configure filter logging through syslog. Types of filtering include:

- [IP protocol \(See “Well-known protocol numbers” on page 190\)](#)
- [TCP/UDP application or source ports \(See “Well-known TCP/UDP application port numbers” on page 191\)](#)
- [TCP flags \(See “TCP” on page 199\)](#)
- [ICMP message type \(See “ICMP Type” on page 200\)](#)

Table 54 lists the tasks involved in configuring filtering:

Table 54 Filtering tasks

| Task | For more information, see... |
|---|--|
| Creating a default filter (recommended). | "Configuring a default filter" on page 192 |
| Setting criteria for the filter (address, masks, protocol, etc.). Setting the filter action (Allow, Deny, Redirect, NAT). Enabling the filter. | "Creating a new filter" on page 201 |
| Adding the filter to a port. | "Applying filters to a port" on page 90 |
| Enabling filtering on the port. | "Enabling or disabling filtering on a port" on page 89 |

Well-known protocol numbers

Table 55 lists well-known protocol numbers used in filters.

Table 55 Well-known protocol numbers

| Number | Protocol |
|--------|---|
| 1 | ICMP |
| 2 | IGMP |
| 3 | Gateway-Gateway Protocol (GGP) |
| 6 | TCP |
| 8 | Exterior Gateway Protocol |
| 12 | PARC Universal Packet Protocol (PUP) |
| 17 | UDP |
| 20 | Host Monitoring Protocol (HMP) |
| 27 | Reliable Datagram Protocol (RDP) |
| 46 | Reservation Protocol QOS (RSVP) |
| 47 | General Routing Encapsulation Protocol (PPTP Data over GRE) |
| 50 | Encapsulation Security Payload (ESP) IPsec |
| 51 | Authentication Header (AH) IPsec |
| 66 | MIT Remote Virtual Disk (RVD) |
| 88 | Internet Group management Protocol (IGMP) |

Table 55 Well-known protocol numbers (continued)

| Number | Protocol |
|--------|---|
| 89 | Open Shortest Path First Protocol (OSPF) |
| 112 | Virtual Router Redundancy Protocol (VRRP) |

Well-known TCP/UDP application port numbers

[Table 56](#) lists well-known application port numbers.

Table 56 Well-known application ports

| Number | TCP/UDP Application | Number | TCP/UDP Application | Number | TCP/UDP Application |
|--------|---------------------|--------|---------------------|---------------|---------------------|
| 20 | ftp-data | 79 | finger | 179 | bgp |
| 21 | ftp | 80 | http | 194 | irc |
| 22 | ssh | 109 | pop2 | 220 | imap3 |
| 23 | telnet | 110 | pop3 | 389 | ldap |
| 25 | smtp | 111 | sunrpc | 443 | https |
| 37 | time | 119 | nntp | 520 | rip |
| 42 | name | 123 | ntp | 554 | rtsp |
| 43 | whois | 143 | imap | 1645, 1812 | Radius |
| 53 | domain | 144 | news | 1813 | Radius Acctg |
| 69 | tftp | 161 | snmp | 1985 | hsrp |
| 70 | gopher | 162 | snmptrap | | |

About default filters

Before enabling filtering on any port, you should configure a default filter. Default filters handle traffic not covered by any other filter. You must set all criteria in the default filter to the full range possible (Any).

Default filters are recommended (but not required) when configuring filters for IP traffic control and redirection. Default filters can increase session performance although they use session binding resources. If you observe higher than acceptable binding failures in the server load balance maintenance statistics (from the Device view, choose Graph > L4 > SLB > Maintenance), you can remove some default filters. For more information, see [“Maintenance statistics” on page 572](#).

When configuring a default filter, give it the lowest order of precedence (Filter index 2048), configure the Action to Deny, and set all matching criteria to Any. If no other filter acts on the traffic, Filter 2048 processes it, denying and logging unwanted traffic.

Configuring a default filter

To configure a default filter:

- 1 From the Device View, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > L4 Switching > Filters.

The Filters dialog box opens to the Filters tab ([Figure 73](#)) with the fields defined in [Table 57 on page 196](#).

Figure 73 Layer 4 Switching Filters—Filters tab

| Index | Name | Filter | Action | Invert | Logging | Caching | Client Proxy | VLAN | Source Addr... | Destination Address Filter Type |
|-------|------|---------|--------|------------|----------|---------|--------------|------|----------------|---------------------------------|
| 1 | | enabled | allow | invert-off | disabled | enabled | enabled | any | ip | ip |
| 3 | | enabled | allow | invert-off | disabled | enabled | enabled | any | ip | ip |

| Source IP Addr... | Source IP Mask | Destination IP Addr... | Destination IP Mask | Source MAC Address | Destination MAC Addr... |
|-------------------|-----------------|------------------------|---------------------|--------------------|-------------------------|
| any | 255.255.255.255 | any | 0.0.0.0 | 00:00:00:00:00:00 | 00:00:00:00:00:00 |
| 2.2.2.20 | 255.0.0.0 | any | 0.0.0.0 | 00:00:00:00:00:00 | 00:00:00:00:00:00 |

| Protocol | Low Source Port | High Source Port | Low Destination Port | High Destination Port | Redirection Port |
|----------|-----------------|------------------|----------------------|-----------------------|------------------|
| tcp(6) | 0 | 0 | rtsp(554) | rtsp(554) | 0 |
| any(0) | 0 | 0 | 0 | 0 | 0 |

| Redirection Group | URL Redirec... | Network Address Transl... | NAT Active FTP | NAT Session Timeout | TCP ACK or RST Matching |
|-------------------|----------------|---------------------------|----------------|---------------------|-------------------------|
| 1 | disabled | destination-address | disabled | | 4 disabled |
| 1 | disabled | destination-address | disabled | | 4 disabled |

| TCP ACK or RST Matching | TCP U... | TCP ACK | TCP PSH | TCP RST | TCP SYN | TCP F... | ICMP Type | IP Option | IP TOS |
|-------------------------|----------|---------|---------|---------|---------|----------|-----------|-----------|--------|
| disabled | disable | disable | disable | disable | disable | disable | 255 | disable | 0 |
| disabled | disable | disable | disable | disable | disable | disable | 255 | disable | 0 |

| IP TOS Mask | New IP TOS | TCP Connection Rate Limiting | Maximum Connection for TCP Rate Limiting | Firewall Redirect Hash |
|-------------|------------|------------------------------|--|------------------------|
| 0 | 0 | disabled | | 10 disabled |
| 0 | 0 | disabled | | 10 disabled |

| WAN Link Load Balancing | Intrusion Detection Hash | Hash | BWM Contract | WAP Radius Snooping |
|-------------------------|--------------------------|-----------|--------------|---------------------|
| disabled | sip | Automatic | 2 | disabled |
| disabled | sip | Automatic | 1024 | disabled |

2 row(s)

3 Click Insert.

The [Insert Filters dialog box](#) (Figures 74 and 75) opens.

Figure 74 Layer 4 Switching Filters—Insert Filters dialog box (left side)

10.160.17.51 - Card 2 - Filters, Insert Filters

Index: 1..2048

Name:

Filter: enabled disabled

Action: ▾

Invert: invert-on invert-off

Logging: enabled disabled

Caching: enabled disabled

Client Proxy: enabled disabled

VLAN: 0..4094 (any: 0)

Source Address Filter Type: ip mac

Destination Address Filter Type: ip mac

Source IP Address:

Source IP Mask:

Destination IP Address:

Destination IP Mask:

Source MAC Address:

Destination MAC Address:

Protocol: (number)

Low Source Port: 0..65534 (none: 0)

High Source Port: 0..65534 (none: 0)

Low Destination Port: 0..65534 (none: 0)

High Destination Port: 0..65534 (none: 0)

Redirection Port: 0..65534

Redirection Group: 1..256

Figure 75 Layer 4 Switching Filters—Insert Filters dialog box (right side)

10.160.17.51 - Card 2 - Filters, Insert Filters

URL Redirection: enabled disabled

Network Address Translation: destination-address source-address

NAT Active FTP: enabled disabled

NAT Session Timeout: 4 4..30 (even values only)

TCP ACK or RST Matching: enabled disabled

TCP URG: enable disable

TCP ACK: enable disable

TCP PSH: enable disable

TCP RST: enable disable

TCP SYN: enable disable

TCP FIN: enable disable

ICMP Type: 255 0..255 (any: 255)

IP Option: enable disable

IP TOS: 0 0..255

IP TOS Mask: 0 0..255

New IP TOS: 0 0..255

TCP Connection Rate Limiting: enabled disabled

Maximum Connection for TCP Rate Limiting: 10 0..255

Firewall Redirect Hash: enabled disabled

WAN Link Load Balancing: enabled disabled

Intrusion Detection Hash: sip dip both

Hash: Automatic

BWM Contract: 1024 1..1024 Browse...

WAP Radius Snooping: enabled disabled

Insert Close Help...

- 4 In the Index field, type 2048 to give this default filter the lowest order of precedence.
- 5 In the Source IP Address field, type Any.
- 6 In the Destination IP Address field, type Any.
- 7 In the Protocol field, type Any.
- 8 In the Logging field, click Enabled.
- 9 In the Action field, click the down arrow and choose Deny.
- 10 In the Filter field, click Enabled to enable the default filter.
- 11 Click Insert to insert the filter definition in the Filters tab and close the Insert Filters dialog box.
- 12 In the Filters tab, click Set > Apply > Close to save the changes and close the Filters dialog box.

The Default filter is configured.

[Table 57](#) describes the Filters fields.

Table 57 Layer 4 Switching Filters—Filters fields

| Field | Description |
|--------|---|
| Index | Sets the index number—1 to 2,048. To improve efficiency: <ul style="list-style-type: none">• Place filters used most near the beginning of the list.• Number filters sequentially beginning with 1. When multiple filters are stacked together on a port, the filter's number determines its order of precedence—the filter with the lowest number is checked first. When traffic is encountered at the WSM port, if the filter matches, its configured action takes place and the rest of the filters are ignored. If the filter criteria doesn't match, the next filter is tried. |
| Name | Sets a filter name of up to 31 characters. |
| Filter | Enables or disables the state of this filtering rule. |

Table 57 Layer 4 Switching Filters—Filters fields (continued)

| Field | Description |
|--------------|---|
| Action | <p>Sets the action for the filtering rule.</p> <ul style="list-style-type: none"> • Allow: Allows the frame to pass. • Deny: Discard frames that fit this filter's profile. This can be used for building basic security profiles. • Redirect: Redirect frames that fit this filter's profile, such as for web cache redirection. In addition, Layer 4 processing must be activated. • NAT: Perform generic Network Address Translation (NAT). This can be used to map the source or destination IP address and port information of a private network scheme to/from the advertised network IP address and ports. This is used in conjunction with the <code>nat</code> option below and can also be combined with proxies. |
| Invert | <p>Sets the invert logic for the filter entry—invert-on or invert-off. Used to reverse the filter logic in order to activate the filter whenever the specified conditions are not met.</p> |
| Logging | <p>Enables or disables logging.</p> <p>When enabled, messages are sent to the console port and system host log (<code>syslog</code>). Messages include packet source and destination IP addresses.</p> |
| Caching | <p>Enables or disables caching sessions that match filter. The default is enabled.</p> <p>Exercise caution while applying cache-enabled and cache-disabled filters to the same WSM port. A cache-enabled filter creates a session entry in the WSM, so that the WSM can bypass checking for subsequent frames that match the same criteria. Cache is enabled by default.</p> <p>Note: Cache should be disabled if applying a filter to virtual server IP address while performing UDP load balancing.</p> |
| Client Proxy | <p>Enables or disables client proxy. The default is enabled.</p> <p>Applies only with filter action Redirect or NAT. Enable or disable proxy IP address translation for traffic matching the filter criteria. By default, this is enabled. If disabled, any proxy defined for the WSM port is not performed for traffic meeting the filter criteria. This is useful when certain traffic must retain original IP address information, or when other forms of translation (such as Application Redirection or NAT) are preferred.</p> |
| VLAN | <p>Sets the VLAN associated with this filter.</p> |

Table 57 Layer 4 Switching Filters—Filters fields (continued)

| Field | Description |
|---------------------------------|---|
| Source Address Filter Type | Sets the source filter type to IP or MAC address. |
| Destination Address Filter Type | Sets the destination filter type to IP or MAC address. |
| Source IP Address | Sets the source IP address for the filter. A setting of 0.0.0.0 allows any address to filter through. If defined, traffic with this source IP address will be affected by this filter. Specify an IP address in dotted decimal notation, or “Any”. A range of IP addresses is produced when used with the Source IP Mask below. The default is Any if the source MAC address is Any. |
| Source IP Mask | Sets the source IP subnet mask for filtering. |
| Destination IP Address | Sets the destination IP address to filter. 0.0.0.0 allows any address to filter through. If defined, traffic with this destination IP address will be affected by this filter. Specify an IP address in dotted decimal notation, or “any”. A range of IP addresses is produced when used with the Destination IP Mask below. The default is any if the destination MAC address is any. |
| Destination IP Mask | Sets the IP subnet mask. This IP address mask is used with the Destination IP Address to select traffic which this filter will affect. |
| Source MAC Address | Sets source MAC address to filter. |
| Destination MAC Mask | Sets destination MAC address to filter. |
| Protocol | Sets the protocol to filter. Specify the protocol number (Table 55 on page 190) , name, or “any”. The default is any. If defined, traffic from the specified protocol is affected by this filter. |
| Low Source Port | Sets the lower source TCP/UDP port number to filter. Applies when the filter protocol is defined as UDP or TCP; 0 (zero) indicates no filtering. See Table 56 on page 191 . |
| High Source Port | Sets the higher source TCP/UDP port number to filter. Applies when the filter protocol is defined as UDP or TCP; 0 (zero) indicates no filtering. See Table 56 on page 191 . |
| Low Destination Port | Sets the lower destination TCP/UDP port number to filter. Applies when the filter protocol is defined as UDP or TCP; 0 (zero) indicates no filtering. See Table 56 on page 191 . |

Table 57 Layer 4 Switching Filters—Filters fields (continued)

| Field | Description |
|-----------------------------|--|
| High Destination Port | Sets the higher destination TCP/UDP port number to filter. Applies when the filter protocol is defined as UDP or TCP; 0 (zero) indicates no filtering. See Table 56 on page 191 . |
| Redirection Port | Sets real server port used for redirection. The default is 0. Applies only when Redirect is specified at the filter action. Defines the real server TCP/UDP port number to which redirected traffic will be sent. For valid Layer 4 health checks, this must be configured whenever TCP protocol traffic is redirected. Also, if transparent proxies are used for Network Address Translation (NAT) Redirection Port must be configured for all Application Redirection filters. (See Table 56 on page 191) |
| Redirection Group | Sets real server group to which to redirect. Default is 1. Applies only when Redirect is specified at the filter action. Define a real server group (1 to 16) to which redirected traffic will be sent. |
| URL Redirection | Enables or disables URL redirection. |
| Network Address Translation | Sets the selection of the address for Network Address Translating (NAT): destination-address or source-address. |
| NAT Active FTP | Enables or disables FTP NAT for the active FTP. Disabled by default. Enables or disables active FTP Client Network Address Translation (NAT). When a client in active FTP mode sends a <code>port</code> command to a remote FTP server, the WSM looks into the data part of the frame and replaces the client's private IP address with a proxy IP address. The real server port is replaced with a proxy port. By default, this option is disabled. |
| NAT Session Timeout | Sets time-out for the NAT session—4 to 30, even values only. |
| TCP ACK or RST Matching | Enables or disables filtering on matching TCP ACK (acknowledgement) or RST (reset) flag matching. |
| TCP URG | Enables or disables TCP URG (urgent) flag matching. The default is disabled. |
| TCP ACK | Enables or disables TCP ACK (acknowledgement) flag matching. The default is disabled. |
| TCP PSH | Enables or disables TCP PSH (push) flag matching. The default is disabled. |
| TCP RST | Enables or disables TCP RST (reset) flag matching. The default is disabled. |

Table 57 Layer 4 Switching Filters—Filters fields (continued)

| Field | Description |
|--|---|
| TCP SYN | Enables or disables TCP SYN (synchronize) flag matching. The default is disabled. |
| TCP FIN | Enables or disables TCP FIN (finish) flag matching. The default is disabled. |
| ICMP Type | Sets ICMP message type to filter—0 to 255, or Any. Default is Any. |
| IP Option | Enables or disables the IP option. |
| IP TOS | Sets IP Type of Service (TOS) value to filter—0 to 255. |
| IP TOS Mask | Sets IP TOS mask for filtering—0 to 255. |
| New IP TOS | Overwrites the new IP TOS value when filtering—0 to 255. |
| TCP Connection Rate Limiting | Enables or disables TCP connection rate limiting. |
| Maximum Connection for TCP Rate Limiting | Sets the maximum number of connections (0 to 255) for TCP connection rate limiting. |
| Firewall Redirect Hash | <p>Enables or disables filtering the firewall redirect hash method.</p> <p>To ensure that the stateful inspection behavior of firewalls is maintained, a hashing algorithm is used to ensure that inbound packets and outbound packets for a pair of IPSA/IPDA traverse through the same firewall. If the dport is 80 or 21, enabling this option changes the hash of the filter from a WCR hash to a FWLB hash. By default, this option is disabled.</p> |
| WAN Link Load Balancing | Enables or disables WAN link load balancing. Disabled by default. |
| Intrusion Detection Hash | <ul style="list-style-type: none"> • Sets hash parameter for intrusion detection server load balancing. • sip: source IP address or range • dip: destination IP address or range • both: both source and destination IP address |
| Hash | Sets the hash parameters for this filter: auto, sip (source IP), dip (destination IP) or both. |
| BWM Contract | Sets the Bandwidth Management Contract (0 to 1024). By default, the contract number is set at 1024. For more information, see "Bandwidth management" on page 493 . |
| WAP Radius Snooping | Enables or disables Wireless Application Protocol (WAP) RADIUS snooping. Disabled by default. |

Creating a new filter

To create a new filter:

- 1 From the Device View, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, select WSM Card > L4 Switching > Filters.

The Filters dialog box opens.

- 3 Click Insert.

The [Insert Filters dialog box](#) ([Figure 74 on page 194](#) and [Figure 75 on page 195](#)) opens. See [Table 57](#) for field descriptions.

- 4 In the Index field, type a new filter number (1 to 2048).

- 5 Use the fields in the Insert Filters dialog box to define filter criteria.

For examples of specific types of filters, see the *Web OS Switch Software 10.0 Application Guide* (part number 212777-A).

- 6 When you have defined all criteria for the filter, in the Action field, click the down arrow and choose an action for the new filter (Allow, Deny, Redirect, NAT).

- 7 In the Filter field, click Enabled to enable the new filter rule.

- 8 Click Insert to insert the filter definition in the Filters tab and close the Insert Filters dialog box.
- 9 From the Filters tab, click Set > Apply > Close to save the changes and close the Filters dialog box.

The new filter is configured.

Chapter 7

Server load balancing basics

This section includes the following topics:

- [“About server load balancing” on page 203](#)
- [“Methods of server load balancing” on page 205](#)
- [“Topology rules for SLB” on page 206](#)
- [“About load balancing metrics” on page 207](#)
- [“About connection timeouts for real servers” on page 211](#)

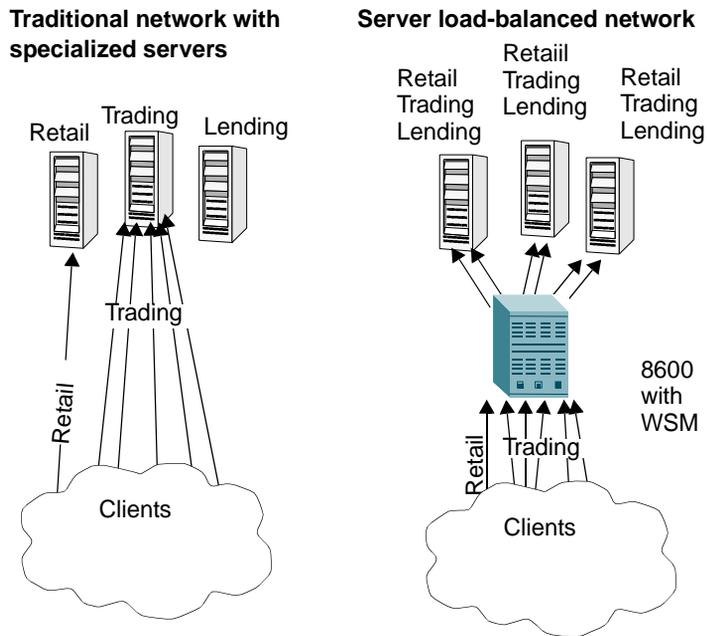
About server load balancing

Server load balancing (SLB) improves network performance by eliminating the potential for overloading servers that provide specialized services. An 8600 with a WSM is aware of the services provided by each server, and uses availability and performance factor metrics to direct and balance session traffic.

In average multiple-server networks without server load balancing, each server usually provides one or two unique services. Servers providing services in high demand are overloaded. This impacts the performance of the entire network as user requests rejected by the server are then resubmitted by user stations. Overutilization of key servers often occurs in networks where other servers are actually available.

SLB maps multiple physical (real) servers providing unique services to one logical (virtual) server connection and applies an algorithm to balance TCP/UDP requests across the server pool.

[Figure 76](#) illustrates the difference between networks using specialized servers and server load-balanced networks.

Figure 76 Traditional versus server load-balanced networks

By front-ending the server pools, the WSM improves network performance, reliability, and scalability as it intercepts and distributes user requests. When a request is received, the WSM uses load-balancing algorithms to bind the session to the real IP address of the best available resource; and maps the session to the virtual connection to protect the server pool.

Basic SLB for specific services is possible when each server in the pool has access to identical content and is front-ended by the 8600 with WSM. Advanced (or Extended) Layer 7 SLB does not require content to be duplicated. For more information, see

Methods of server load balancing

[Table 58](#) lists and describes methods of server load balancing.

Table 58 Server load balancing methods

| Method | Description |
|---|---|
| Virtual server load balancing | <p>The WSM is configured to act as a virtual (logical) server and is given a virtual IP address or range of addresses for each collection of TCP/UDP services it distributes. The WSM supports up to 256 virtual servers, each server distributing up to eight different services (up to a total of 2048 services).</p> <p>Each virtual server is assigned a list of the IP addresses or range of addresses of the real (physical) servers in the pool where its services reside. User stations requesting connections to a service communicate with a virtual server on the WSM. The WSM then binds the session to the IP address of the best available real server and remaps the fields in each frame from virtual addresses to real addresses. IP, FTP, RTSP, IDS, and static session Wireless Application Protocol (WAP) use virtual servers for load balancing.</p> <p>For more information, see “Virtual server load balancing” on page 213.</p> |
| Filter-based server load balancing | <p>Filters allow, deny, or redirect traffic according to IP address, protocol, or Layer 4 port criteria. In filtered-based load balancing, a filter is used to redirect traffic to a real server group. If the group is configured with more than one real server entry, redirected traffic is load balanced among the available real servers in the group. Firewalls and WAN links use redirection filters to load balance traffic.</p> <p>For more information, see:</p> <ul style="list-style-type: none"> • “Firewall server load balancing” on page 327 • “WAN link load balancing” on page 313 • “Configuring WAP SLB using RADIUS snooping” on page 294 |
| Content-intelligent switching | <p>Content-based load balancing uses Layer 7 application data (such as URL, cookies, and Host Headers) to make intelligent load balancing decisions. URL-based load balancing, browser-smart load balancing, and cookie-based preferential load balancing are a few examples of content-based load balancing.</p> <p>For more information, see “Content-intelligent switching” on page 453.</p> |

Topology rules for SLB

Use the following guidelines when configuring SLB topologies.

- In standard SLB, all client requests to a virtual server IP address and all responses from the real servers must pass through the WSM. If there is a path between the client and the real servers that does not pass through the WSM, the WSM can be configured to proxy requests in order to guarantee that responses use the correct path.

For more information, see [“Configuring proxy IP addresses” on page 245](#).

- Make sure that identical content is available to each server in the same pool by either duplicating static applications and data on each real server in the pool, or providing each real server in the pool with access to the same data through use of a shared file system or back-end database server.
- Configure persistence (or use [minimum misses](#) or [hash](#) metrics) for connections that require a series of client requests go to the same real server to retain data between connections. These services include Web search results, multi-page forms that the user fills in, or custom Web-based applications typically created using cgi-bin scripts. For more information, see [“Maintaining session persistence” on page 433](#).

For more information about metrics, see [“About load balancing metrics” on page 207](#).

- Connect clients and servers through the same WSM port. Each WSM port can process client requests, server traffic, or both. You can enable or disable processing on a port independently for client or server traffic.

Layer 4 client processing

Ports configured to process client request traffic provide address translation from virtual server IP to real server IP address.

Layer 4 server processing

Ports configured to process server responses to client requests provide address translation from real server IP address to virtual server IP address. These ports require real servers to be connected to the WSM directly or through a hub, router, or another WSM.



Note: WSM ports configured for Layer 4 client/server processing can simultaneously provide Layer 2 switching and IP routing.

About load balancing metrics

You can configure how traffic is distributed across multiple real servers by choosing a load-balancing metric for the server group. The metric tells the real server group how to select which of its real servers will receive the next client connection. The available metrics are:

- [“Minimum Misses,”](#) next
- [“Hash”](#) on page 208
- [“Least connections”](#) on page 209
- [“Round robin”](#) on page 209
- [“Response time”](#) on page 210
- [“Bandwidth”](#) on page 210

Minimum Misses

The minimum misses metric is optimized for Application Redirection. The minimum misses metric uses the IP address in the client request to select a server. As described below, the minimum misses metric uses either the client’s source or destination IP address.

- For Application Redirection, requests for a specific client destination IP address are sent to the same server. This metric helps maximize successful cache hits. Best statistical load balancing is achieved when the IP address destinations of load-balanced frames are spread across a broad range of IP subnets.

- For SLB, the client source IP address and real server IP address are used. All requests from a specific client are sent to the same server.. This metric is useful for applications where client information must be retained on the server between sessions. With this metric, server load becomes most evenly balanced as the number of active clients with different source or destination addresses increases.

When selecting a server using the minimum misses metric, the WSM calculates a score for each available real server based on the IP address. The server that scores the highest is assigned the connection. This metric attempts to minimize the disruption of persistency when servers are removed from service. This metric should be used only when persistence is a must.



Note: The minimum misses metric cannot be used for firewall load balancing, since the real server IP addresses used in calculating the score for this metric are different on each side of the firewall.

Hash

The hash metric uses IP address information in the client request to select a server. The specific IP address information used depends on the application, as described below.

- For Application Redirection, the client destination IP address is used. All requests for a specific IP destination address will be sent to the same server. This is particularly useful for maximizing successful cache hits.
- For SLB, the client source IP address is used. All requests from a specific client will be sent to the same server. This option is useful for applications where client information must be retained between sessions.
- For FWLB, both the source and destination IP addresses are used to ensure that the two unidirectional flows of a given session are redirected to the same firewall. When selecting a server, a mathematical hash of the relevant IP address information is used as an index into the list of currently available servers. Any given IP address information will always have the same hash result, providing natural persistence, as long as the server list is stable.

However, if a server is added to or leaves the mix, then a different server might be assigned to a subsequent session with the same IP address information even though the original server is still available. Open connections are not cleared.



Note: The hash metric provides more distributed load balancing than minimum misses at any given instant. It should be used if the statistical load balancing achieved using minimum misses is not as optimal as desired. If the load balancing statistics with minimum misses indicate that one server is processing significantly more requests over time than other servers, consider using the hash metric.

Least connections

With the least connections metric, the number of connections currently open on each real server is measured in real time. The server with the fewest current connections is considered to be the best choice for the next client connection request. This option is the most self-regulating, with the fastest servers typically getting the most connections over time.

Round robin

With the roundrobin metric, new connections are issued to each server in turn; that is, the first real server in the group gets the first connection, the second real server gets the next connection, followed by the third real server, and so on. When all the real servers in this group have received at least one connection, the issuing process starts over with the first real server.

Response time

The response metric uses real server response time to assign sessions to servers. The response time between the servers and the WSM is used as the weighting factor. The WSM monitors and records the amount of time it takes for each real server to reply to a health check to adjust the real server weights. The weights are adjusted so they are inversely proportional to a moving average of response time. In such a scenario, a server with half the response time as another server will receive a weight twice as large.



Note: The effects of the response weighting apply directly to the real servers and are not necessarily confined to the real server group. When response time-metered real servers are also used in other real server groups that use the least connections or roundrobin metrics, the response weights are applied on top of these calculations for the affected real servers. Since the response weight changes dynamically, this can produce fluctuations in traffic distribution for the real server groups that use the least connections or roundrobin metrics.

Bandwidth

The bandwidth metric uses real server octet counts to assign sessions to a server. The WSM monitors the number of octets sent between the server and the WSM. Then, the real server weights are adjusted so they are inversely proportional to the number of octets that the real server processes during the last interval.

Servers that process more octets are considered to have less available bandwidth than servers that have processed fewer octets. For example, the server that processes half the amount of octets over the last interval receives twice the weight of the other servers. The higher the bandwidth used, the smaller the weight

assigned to the server. Based on this weighting, the subsequent requests go to the server with the highest amount of free bandwidth. These weights are automatically assigned. The bandwidth metric requires identical servers with identical connections.



Note: The effects of the bandwidth weighting apply directly to the real servers and are not necessarily confined to the real server group. When bandwidth-metered real servers are also used in other real server groups that use the least connections or roundrobin metrics, the bandwidth weights are applied on top of these calculations for the affected real servers. Since the bandwidth weight changes dynamically, this can produce fluctuations in traffic distribution for the real server groups that use the least connections or roundrobin metrics.

About connection timeouts for real servers

Every client-to-server session being load balanced is recorded in the WSM's session table. The typical session completes the following sequence:

- 1 A client makes a request.
- 2 The session is recorded in the table.
- 3 The data is transferred until the client ends the session.
- 4 The session table entry is removed.

In certain circumstances, such as when a client application is abnormally terminated by the client's system, TCP/UDP connections will remain registered in the WSM's binding table. In order to prevent table overflow, these orphaned entries must be aged out. By setting the real server Timeout threshold for the maximum number of minutes an inactive connection can remain open, you can ensure that orphan table entries are removed.

If an open TCP/IP sessions is not closed properly, it remains inactive for 10 minutes, by default, before it is removed from the WSM session table. For example, if the WSM receives the SYN for the session but no FIN, the connection remains open in the session table until it reaches the timeout threshold set for the real server. You can change the default timeout period for a real server. The Timeout threshold must be specified in even-numbered increments.

For more information, see [“Configuring each real server” on page 216](#).

About real server maximum connections

You can configure each real server for the number of open connections it can handle for SLB.

Actual connection values average from approximately 500 HTTP connections for slower servers to 1500 for quicker, multiprocessor servers. The number also depends on the duration of each session and how much CPU capacity is occupied in processing each session. Connections that use a lot of Java or CGI scripts for forms or searches require more server resources and have a lower maximum connections limit. You may wish to use a performance benchmark tool to determine how many connections your real servers can handle.

When a server reaches its maximum connections limit, the WSM no longer sends new connections to the server. When the server drops back below the maximum connections limit, new sessions are again allowed.

The maximum connections setting is described in [Table 60 on page 219](#).

Assigning backup or overflow servers

A real server can backup other real servers and can handle overflow traffic when the maximum connection limit is reached. Each backup real server must be assigned a real server number and real server IP address. It must then be enabled. Finally, the backup server must be assigned to each real server that it will back up. For more information, see [“Configuring each real server” on page 216](#).

A backup/overflow server can also be assigned to a real server group. If all real servers in a real server group fail or overflow, the backup comes online. Real server groups can also use another real server group for backup/overflow. For more information, see [“Configuring a real server group” on page 225](#).

Chapter 8

Virtual server load balancing

This section includes the following topics:

- [“How virtual server load balancing works” on page 213](#)
- [“Before and after SLB example” on page 214](#)
- [“Configuring virtual server load balancing” on page 215](#)

How virtual server load balancing works

With virtual server load balancing (SLB), the Web Switching Module is configured to act as a virtual server and is given a virtual server IP address (or range of addresses) for each collection of services it distributes. The Web Switching Module supports up to 256 virtual servers, each distributing up to eight different services (up to a total of 2048 services).

Each virtual server is assigned a list of the IP addresses (or range of addresses) of the real servers in the pool where its services reside. When user stations request connections to a service, they communicate with a virtual server on the WSM. When the WSM receives the request, it binds the session to the IP address of the best available real server and remaps the fields in each frame from virtual addresses to real addresses.

Some examples of services that use virtual servers for load balancing include:

- Internet protocol (IP)
- File transfer protocol (FTP)
- Real time streaming protocol (RTSP)
- Intrusion detection system (IDS)
- Static session wireless application protocol (WAP)

Before and after SLB example

Figure 77 shows a network where customer Web sites are hosted by a Web hosting company and/or Internet Service Provider (ISP). The Web content is relatively static and is kept on a single NFS server for easy administration. As the customer base increases, the number of simultaneous Web connection requests also increases.

Figure 77 Web hosting without server load balancing

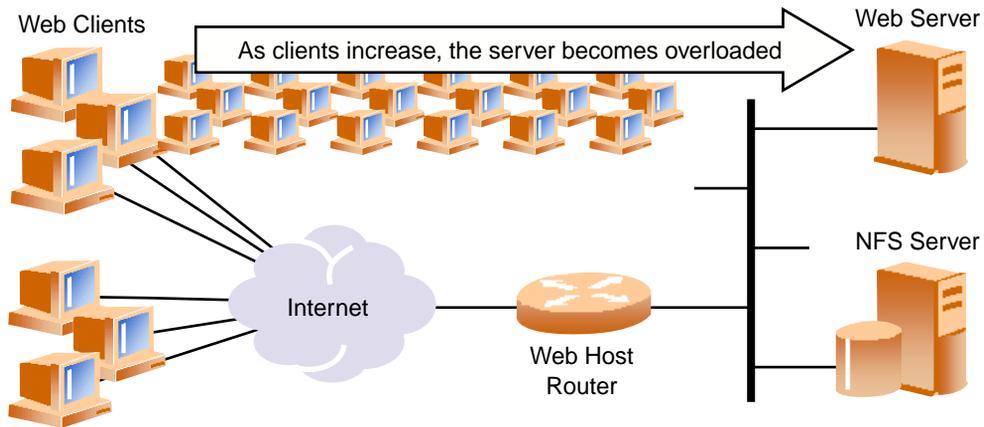
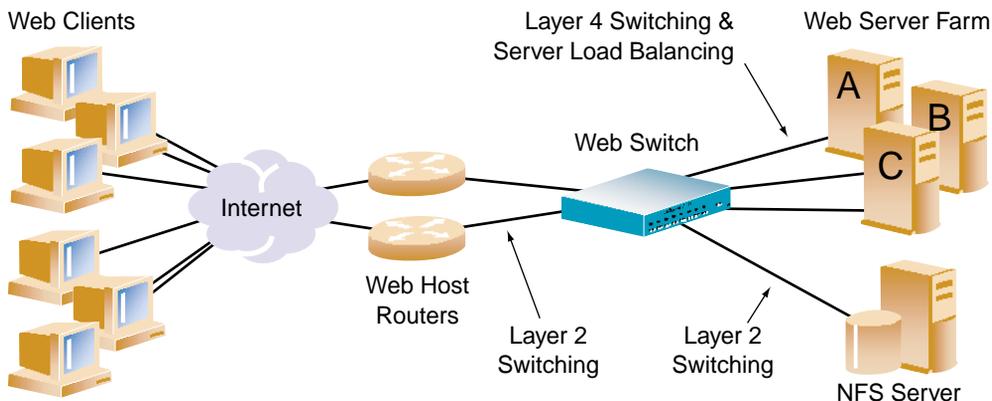


Figure 78 shows the same network after implementing server load balancing. In this example, clients have multiple paths to the Web Switching Module and balanced access to a pool of servers with identical content. If one server fails, the others can take up the additional load. More servers can be added at any time to increase processing power. For ease of maintenance, servers can be added or removed dynamically, with no interruption to shared services.

Figure 78 Web hosting with server load balancing



Configuring virtual server load balancing

The following are general steps for configuring virtual server load balancing (SLB).

Table 59 General steps for configuring server load balancing

| Step | Description |
|---|--|
| "Configuring each real server" on page 216 | Defines the IP route from the real server to the switch that performs SLB. Enables the real server. |
| "Defining an IP interface on the WSM" on page 225 | Defines the IP route from the WSM to all of the real (physical) servers. The WSM uses this path to determine the real servers' TCP/IP reach. |
| "Configuring a real server group" on page 225 | Combines real servers into real server groups. Each real server group should consist of all the real servers which provide a specific service for load balancing. Each group must consist of at least one real server. Each real server can belong to more than one group. |

Table 59 General steps for configuring server load balancing (continued)

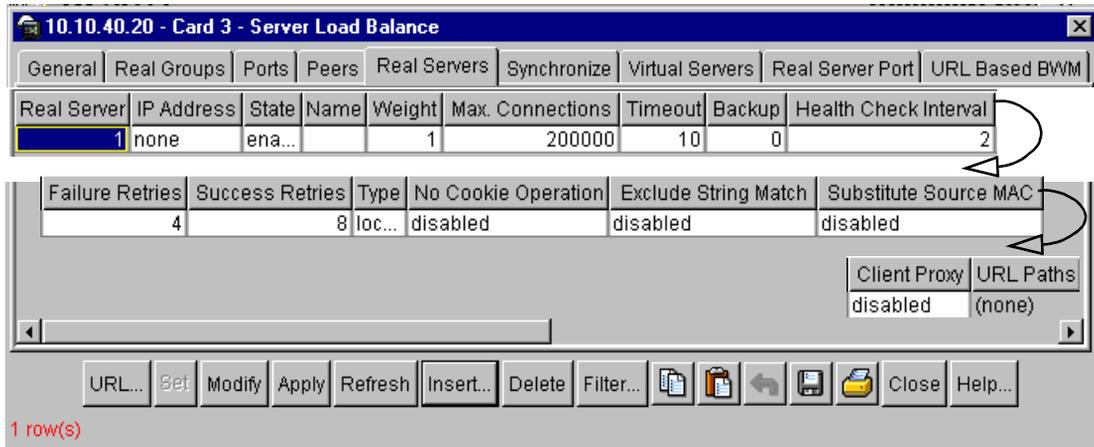
| Step | Description |
|---|--|
| “Configuring a virtual server” on page 230 | Configures the virtual server IP address, adds the TCP/UDP port and real server group to the virtual server configuration, and enables the virtual server. |
| “Configuring services for a virtual server” on page 235 | Configures the services assigned to a virtual server. |
| “Configuring ports for server load balancing” on page 240 | On a per port basis, enables or disables processing independently for each type of Layer 4 traffic (client and server). |
| “Enabling or disabling server load balancing” on page 243 | Globally enables Server Load Balancing |

Configuring each real server

To configure a real server:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the [General tab](#).
- 3 Click the Real Servers tab.

The Real Servers tab ([Figure 79](#)) opens with the fields listed in [Table 60 on page 219](#).

Figure 79 Server Load Balance—Real Servers tab**4** Click Insert

The [Insert Real Servers](#) (Figure 80) dialog box opens with the fields listed in [Table 60](#) on page 219.

Figure 80 Server Load Balance—Insert Real Servers dialog box

Real Server: 6 1..1023

IP Address: none

State: enabled disabled

Name:

Weight: 1 1..48

Max. Connections: 200000 0..200000

Timeout: 10 2..30 (even values only)

Backup: 0 0..1023 (none: 0)

Health Check Interval: 2 0..60 (none: 0)

Failure Retries: 4 1..63

Success Retries: 8 1..63

Type: local-server remote-server

No Cookie Operation: enabled disabled

Exclude String Match: enabled disabled

Substitute Source MAC: enabled disabled

Client Proxy: enabled disabled

Insert Close Help...

- 5 In the Real Server field, type the real server number (1 - 1023).
- 6 In the IP address field, type an IP address for the real server.



Note: The real servers in any given real server group must have an IP route to the switch that performs the SLB functions. This IP routing is most easily accomplished by placing the switches and servers on the same IP subnet, although advanced routing techniques can be used as long as they do not violate the rules outlined in [“Topology rules for SLB” on page 206](#).

- 7 In the State field, click Enabled.
- 8 Use the other fields to configure the real server as needed.

For more information, see [“Modifying the health check interval and retries” on page 410](#)

9 Click Insert.

The server is inserted into the Real Servers tab and the Insert Real Servers dialog box closes.

10 To define other real servers in the server pool, repeat steps 4 - 8.

11 Click Apply > Close.

The real server is defined and the Server Load Balance dialog box closes.



Note: You can define a range of IP addresses for your SLB real and virtual servers. For more information, see [“Configuring IP address ranges for SLB” on page 222](#).

[Table 60](#) describes the fields on the SLB—Real Servers tab and the Insert Real Servers dialog box.

Table 60 Real Servers fields

| Field | Description |
|-------------|--|
| Real Server | The number of the real server: 1 to 1023. |
| IP Address | Sets the IP address of the real server in dotted decimal format. The default is 0.0.0.0. When assigned, the address is PINGed to determine if the server is up, and the administrator will be warned if the server does not respond. |

Table 60 Real Servers fields (continued)

| Field | Description |
|-----------------|---|
| State | <p>Enables or disables the real server. The default is disabled.</p> <p>Enabled—enables this real server for Layer 4 service until it is explicitly disabled. When enabled, the real server can process virtual server requests associated with its real server group.</p> <p>Disabled—disables this real server for Layer 4 service until it is explicitly re-enabled. Any disabled server will no longer process virtual server requests as part of the real server group to which it is assigned. This option <i>does not</i> perform a graceful server shutdown.</p> <p>If you need to reboot a real server, you must make sure that new sessions are not sent to the real server and that old sessions are not discarded. When the session count gets to zero, you may shut down the server by clicking Disabled here. Once maintenance is complete, click Enabled to enable the server again.</p> |
| Name | <p>Defines a 31-character alias for each real server. The default is none. This enables the network administrator to quickly identify the server by a natural language keyword value.</p> |
| Weight | <p>Sets the weighting value (1 to 48) that this real server will be given in the load balancing algorithms. Higher weighting values force the server to receive more connections than the other servers configured in the same real server group. By default, each real server is given a weight setting of 1. A setting of 10 would assign the server roughly 10 times the number of connections as a server with a weight of 1.</p> <p>Weights are not applied when using the hash or minmisses metrics.</p> |
| Max Connections | <p>Sets the maximum number of connections that this server should simultaneously support. By default, the number of maximum connections is set at 200,000. This option sets a threshold as an artificial barrier, such that new connections will not be issued to this server if the maxcon limit is reached. New connections will be issued again to this server once the number of current connections has decreased below the maxcon setting. If all servers in a real server group for a virtual server reach their maxcon limit at the same time, client requests will be sent to the backup/overflow server or backup/overflow server group. If no backup servers/server group are configured, client requests will be dropped by the virtual server.</p> |

Table 60 Real Servers fields (continued)

| Field | Description |
|-----------------------|--|
| TimeOut | Sets the maximum number of minutes (2 - 60, even number only) that an inactive connection remains open. The default is 10. See “About connection timeouts for real servers” on page 211 . |
| Backup | Sets the real server (0 - 1023) used as the backup/overflow server for this real server. The default is 0, which indicates no backup. If the real server becomes inoperative, the WSM activates the backup real server until the original becomes operative again. If the real server reaches its maximum connections limit, the backup comes online to provide additional processing power until the original server becomes de-saturated. The same backup/overflow server may be assigned to more than one real server at the same time |
| Health Check Interval | Sets the interval in seconds (0 - 60) between WSM checks of the real server's health. The default health check interval is 2. A setting of zero disables health check. The type of health check is configured in the real server group settings. See Table 62 on page 228 . For TCP services, the WSM verifies that real servers and their corresponding services are operational by opening a TCP connection to each service, using the defined service ports configured as part of each virtual service. For UDP services, the switch pings servers to determine their status. For more information, see “Modifying the health check interval and retries” on page 410 . |
| Failure Retries | Sets the number of failed attempts (1 - 63) to declare this server down. The default is 4. |
| Success Retries | Sets the number of successful attempts (1 -63) to declare a server up. The default is 8. |
| Type | Sets the server type— local-server or remote-server. The default is local-server. If configured as remote, the server participates in distributed server load balancing. |
| No Cookie Operation | Enables or disables real server handing of client requests that do not contain a URL cookie. The default is disabled. This option is used if you want to designate a specific server to assign cookies only. This server gets the client request, assigns the cookie, and embeds the IP address of the real server that will handle the subsequent requests from the client. |

Table 60 Real Servers fields (continued)

| Field | Description |
|-----------------------|--|
| Exclude String Match | Enables or disables real server handling of requests that do not match the load balance string. The default is disabled. |
| Substitute Source MAC | Enables or disables MAC SA substitution for Layer 4 traffic. The default is disabled. To substitute the MAC SA of client-to-server frames, Substitute Source MAC, must be enabled. |
| Client Proxy | Enables or disables proxy IP address translation. With this option enabled (default), a client request from any application can be proxied using a load-balancing Proxy IP address (PIP). For more information, see "Configuring proxy IP addresses" on page 245 . |
| URL | Opens the URL Path Membership dialog box from which you select the URL paths to apply to a real server. |

About IP address ranges for SLB

You can define a range of IP addresses for your SLB real and virtual servers. The default real and virtual server mask is 255.255.255.255, so that each real and virtual server represents a single IP address. A mask setting of 255.255.255.0 defines a range so that each real and virtual server represents 256 IP addresses. Consider the following example:

- A virtual server is configured with an IP address of 172.16.10.1.
- Real servers 172.16.20.1 and 172.16.30.1 are assigned to service the virtual server.
- The IP address range is set to 255.255.255.0.

When a client request is sent to virtual server IP address 172.16.10.45, the unmasked portion of the address (0.0.0.45) gets mapped directly to the real server IP address selected by the SLB algorithm. Thus, the request would be sent to either 172.16.20.45 or 172.16.30.45.

Configuring IP address ranges for SLB

To define a range of IP addresses for server load balancing:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching> SLB.
The Server Load Balance dialog box opens to the **General** tab (Figure 81) with the fields described in Table 61 on page 224.

Figure 81 Server Load Balance—General tab

The screenshot shows the 'General' tab of the 'Server Load Balance' configuration dialog. The dialog title is '134.177.214.201 - Card 3 - Server Load Balance'. The 'General' tab is selected, and the following settings are visible:

- Server Load Balance: enabled disabled
- Virtual Matrix Architecture: enabled disabled
- Direct Access Mode: enabled disabled
- IP Address Mask: 255.255.255.255
- Management Network: 0.0.0.0
- Management Subnet Mask: 255.255.255.255
- Radius Secret: (empty field)
- Persistent Mask: 255.255.255.255
- Graceful Failover: enabled disabled
- Session Table Fast-Age Period Bit Shift: 1 (range 0..7)
- Session Table Slow-Age Period Bit Shift: 2 (range 0..15)
- Transparent Proxy Cache Protocol: enabled disabled
- Metric (Response & Bandwidth) Update Interval: 10 (range 1..256)
- LDAP Version: version2 version3
- Allow HTTP Health Check on any port?: enabled disabled
- Time Window: 1 (range 1..65535)
- Hold Duration: 1 (range 1..65535)
- Intrusion Detection Real Server Group: 1 (range 1..256)

At the bottom of the dialog, there are five buttons: Set, Apply, Refresh, Close, and Help...

- 3 In the IP address mask field, type the mask that defines a range of IP addresses for your SLB real and virtual servers.
- 4 Click Set > Apply > Close.

The IP address range is defined and the SLB dialog box closes.

Table 61 describes the fields on the SLB—General tab.

Table 61 SLB—General tab fields

| Field | Description |
|---|--|
| Server Load Balance | Enables or disables Server Load Balancing. |
| Virtual Matrix Architecture | Enables or disables Virtual Matrix Architecture. |
| Direct Access Mode | Enables or disables Direct Access Mode. When enabled, Direct Access Mode allows direct access to real servers and any combination of virtual and real servers. |
| IP Address Mask | Sets the virtual and real server IP address mask. The default is 255.255.255.255. |
| Management Network | Sets the IP address, in dotted decimal notation, of the management network. The default is 0.0.0.0. If defined, management traffic with this source IP address will be allowed direct (non-Layer 4) access to the real servers. A range of IP addresses is produced when the Management Subnet Mask field is also defined. |
| Management Subnet Mask | Sets the management subnet mask. The default is 255.255.255.255. This IP address mask is used with the Management Network field (above) to select management traffic which is allowed direct access to real servers. |
| Radius Secret | Sets the RADIUS authentication string. The string is used for generating an encrypted authentication string for the RADIUS health check. The maximum string length is 32 alphanumeric characters. |
| Persistent Mask | Sets the persistent mask. The default is 255.255.255.255. |
| Graceful Failover | Enables or disables graceful server failure. When enabled, the connection to the failure real servers is maintained. See “Enabling graceful server failure” on page 405 . |
| Session Table Fast-Age Period Bit Shift | Sets the session fast-age period bit—0 to 7. |
| Session Table Slow-Age Period Bit Shift | Sets the session slow-age period bit—0 to 15. |
| Transparent Proxy Cache Protocol | Enables or disables TPCP. |

Table 61 SLB—General tab fields (continued)

| Field | Description |
|---|--|
| Metric (Response & Bandwidth) Update Interval | Sets the interval of the response and bandwidth metric updates: 1 to 256. The default is 10. |
| LDAP Version | Sets the LDAP version to 2 or 3. The default is 2. |
| Allow HTTP Health Check on any Port? | Enables or disables HTTP-specific health check on any port. |
| Intrusion Detection real server group | The number (1 to 256) of the real server group configured for Intrusion Detection. For more information, see “Intrusion Detection server load balancing” on page 302 . |
| Time Window | Sets the time window for TCP rate limiting. |
| Hold Duration | Sets the hold down duration for TCP rate limiting. |

Defining an IP interface on the WSM

The WSM must have an IP route to all of the real (physical) servers that receive WSM services. For SLB, the WSM uses this path to determine the real servers’ TCP/IP reach.



Note: The IP interface and the real servers must belong to the same VLAN, if they are in the same subnet.

To configure an IP interface, see [“Configuring IP interfaces” on page 129](#).

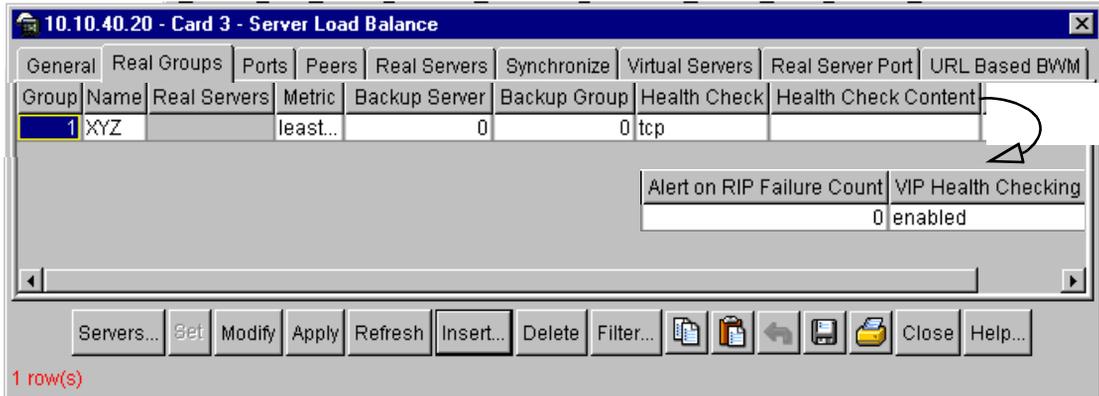
Configuring a real server group

To configure a real (physical) server group:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the [General tab](#).
- 3 Click the Real Groups tab.

The **Real Groups** tab (Figure 82) opens with the fields listed in Table 62 on page 228.

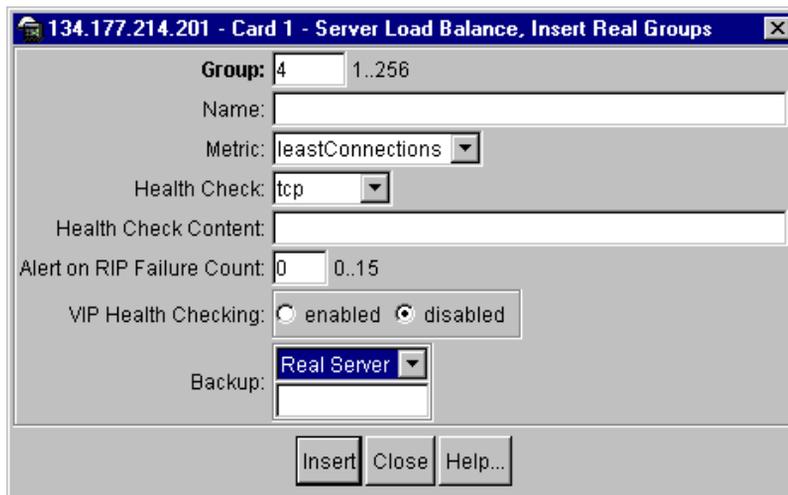
Figure 82 Server Load Balance—Real Groups tab



- 4 Click Insert.

The **Insert Real Groups** dialog box (Figure 83) opens with the fields listed in Table 62 on page 228.

Figure 83 Insert Real Groups dialog box



- 5 In the Group field, type the group number (1 - 256) for the real server group.
- 6 In the Name field, type a name of up to 31 characters for the group.

- 7 In the Metric field, choose a load balancing metric (roundrobin, leastConnections, minmisses, hash, response, or bandwidth) for the group. For more information, see [“About load balancing metrics”](#) on page 207.
- 8 In the Backup field, click the down arrow and choose either Real Server or Real Group; and type the backup server or group number in the text field. For more information, see [“About connection timeouts for real servers”](#) on page 211.
- 9 Use the other fields to further define the real server group.
- 10 Click Insert.
The real server group is inserted into the Real Groups tab and the Insert Real Groups dialog box closes.
- 11 To add the real servers to this group, in the Real Groups tab, select the group.
The row is highlighted.
- 12 Click Servers.
The [Select Real Servers to include in Group dialog box](#) (Figure 84) opens with the fields listed in [Table 63](#) on page 230.

Figure 84 Select Real Servers to include in Group dialog box



- 13 Click the real servers you want to add to the group.

The rows are highlighted.



Caution: If there are already servers in the group, you must use CTRL + Click when selecting new members in this dialog box. If you just choose the new server and click Modify, only the new server is added, and existing server members are removed.

14 Click Modify.

The real servers are added to the group in the Real Groups tab, and the Select Real Servers to include in Group dialog box closes.

15 In the Real Groups tab, click Apply > Close.

The real server group is configured and the SLB dialog box closes.

[Table 62](#) describes the real server group fields.

Table 62 Real server group fields

| Field | Description |
|--------------|--|
| Group | Sets the real server group number—1 to 256. |
| Name | Sets a name of up to 31 characters for the real server group. |
| Real Servers | The numbers of the selected real servers. Servers are added or removed via the Select Real Servers to Include dialog box accessed by clicking Servers. |

Table 62 Real server group fields (continued)

| Field | Description |
|----------------------|---|
| Metric | <p>Sets the load balancing metric used for determining which real server in the group will be the target of the next client request. The default is leastconns.</p> <p>round robin—The WSM sends the request to each server in rotation, regardless of how many connections each server has.</p> <p>leastConnections (leastconns)—This is the default setting. The WSM sends the request to the real server that currently has the least number of connections.</p> <p>minmisses—The WSM sends all requests for a specific IP address to the same real server.</p> <p>hash—The WSM sends all requests from the same client to the same real server.</p> <p>response—The WSM sends the request to the real server with the fastest response time.</p> <p>bandwidth—The WSM sends the request to the real server with the most available free bandwidth.</p> <p>For more information, see “About load balancing metrics” on page 207.</p> |
| Backup Server | <p>Sets a backup real server for this group.</p> <p>If the entire real server group fails, the WSM activates the backup until one of the original real servers becomes operative again.</p> <p>Also applies if all servers in a real server group reach their maximum connections limit. A backup comes online to provide additional processing power until one of the original servers is desaturated.</p> <p>The same backup may be assigned to more than one real server group at the same time.</p> |
| Backup Group | <p>Sets a backup real server group for this group. (See Backup Server, above.)</p> |
| Health Check | <p>Sets the type of health check to be performed—icmp, tcp, http, dns, smtp, pop3, nntp, ftp, imap, radius, sslh, script1 through script 16, link, wsp, wtls, ldap, udpdns, arp. The default is tcp. For more information, see “Types of real server group health checks” on page 407.</p> |
| Health Check Content | <p>Defines the specific content to be examined during health checks. The content depends on the type of health check specified in the Health Check field (above).</p> |

Table 62 Real server group fields (continued)

| Field | Description |
|----------------------------|--|
| Alert on RIP Failure Count | Specifies a minimum threshold for the number of real servers available (1 - 15). If at any time, the number reaches this minimum limit, a SYSLOG ALERT message is sent to the configured SYSLOG servers stating that the real server threshold has been reached for the concerned server load balancing group. The default threshold is 0, which also means the option is disabled |
| Backup | On the Insert Real Groups dialog box, click the down arrow to choose either Real Server or Real Group as a backup for this group. Then type the server or group number in the text field. (See Backup Server, above.) |
| VIP Health Checking | Enables or disables VIP health checking for a service. This feature is enabled by default. However, it works only when the service has DSR (Direct Server Return) feature enabled. When disabled, the switch uses RIP to perform all health checks, whether DSR is enabled or disabled. |
| Servers | Opens the <i>Servers to Include in Group</i> dialog box (below) from which you can select the servers to add to a real server group. |

[Table 63](#) describes the fields on the Select Real Servers to Include in Group dialog box. This dialog box is used to add servers to a server group.

Table 63 Select Real Servers to Include in Group dialog box

| Field | Description |
|-------------|---|
| Real Server | The number assigned to the real server. |
| Name | The name assigned to the real server. |
| IP Address | The real server's IP address. |

Configuring a virtual server

All server load balance client requests are addressed to an IP address on a virtual server. Clients get the virtual server IP address through DNS resolution.

To configure a virtual server:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

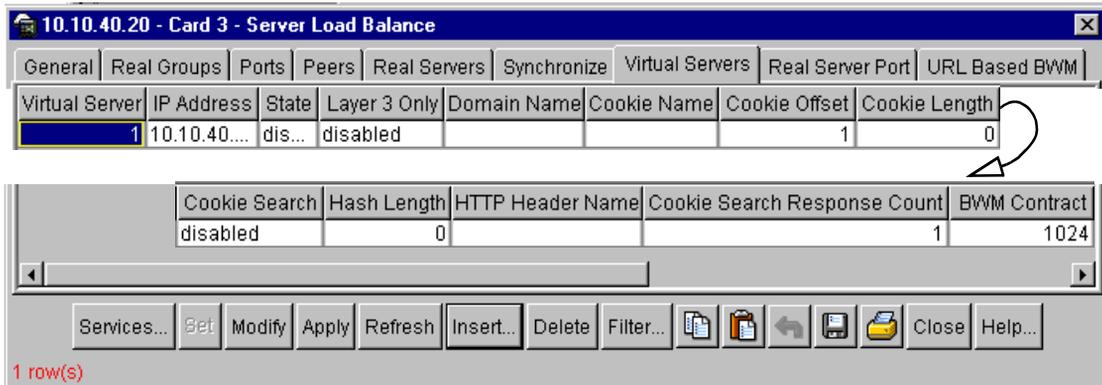
- From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the **General** tab.

- Click the Virtual Server tab.

The **Virtual Server** tab (Figure 85) opens with the fields listed in Table 64 on page 233.

Figure 85 Server Load Balance—Virtual Servers tab



- Click Insert.

The **Insert Virtual Servers** dialog box (Figure 86) opens with the fields listed in Table 64 on page 233.

Figure 86 Insert Virtual Servers dialog box

134.177.214.201 - Card 1 - Server Load Balance, Insert Virtual Servers

Virtual Server: 3 1..256

IP Address: none

State: enabled disabled

Layer 3 Only: layer3Only disabled

Domain Name:

Cookie Name:

Cookie Offset: 1 0..64

Cookie Length: 0 0..64

Cookie Search: enabled disabled

Hash Length: 0 0..255 (none: 0)

HTTP Header Name:

Cookie Search Response Count: 1..16

Expire Cookie: <MM/dd

BWM Contract: 1024 1..1024

- 5 In the Virtual Servers field, type a virtual server number (1 - 256).
- 6 In the IP Address field, type an IP address for the virtual server.
- 7 In the State field, click Enabled.
- 8 Use the other fields to further define the virtual server.
- 9 Click Insert.

The virtual server is entered into the Virtual Servers tab and the Insert Virtual Server dialog box closes.

- 10 Click Apply > Close.

The virtual server is configured and the Server Load Balance dialog box closes.

[Table 64](#) describes the fields on the SLB—Virtual Servers tab and the Insert Virtual Servers dialog box.

Table 64 Virtual server fields

| Field | Description |
|----------------|---|
| Virtual Server | The number of the virtual server: 1 to 256. |
| IP Address | <p>Sets the IP address of the virtual server using dotted-decimal notation. The default is 0.0.0.0.</p> <p>The virtual server created within the switch will respond to ARPs and PINGs from network ports as if it was a normal server. Client requests directed to the virtual server's IP address will be balanced among the real servers available to it through real server group assignments.</p> |
| State | Enables or disables the virtual server. When enabled, the virtual server within the switch services client requests sent to its defined IP address. The default is disabled. |
| Layer 3 Only | <p>Enables or disables Layer 3-only balancing: layer3Only or disable. The default is disable. See “Configuring IP server load balancing” on page 270.</p> <p>Normally, the client IP address is used with the client Layer 4 port number to produce a session identifier. When the Layer 3 Only option is enabled, the WSM uses only the client IP address as the session identifier. It associates all connections from the same client with the same real server while any connection exists between them. This option is necessary for some server applications where state information about the client system is divided across different simultaneous connections, and also in applications where TCP fragments are generated. If the real server to which the client is assigned becomes unavailable, the Layer 4 software will allow the client to connect to a different server.</p> |
| Domain Name | <p>Sets the domain name of the virtual server. The maximum string length is 34 characters. The default is none.</p> <p>Typically includes the name of the company or organization, and the Internet group code (.com, .edu, .gov, .org, and so forth), such as nortelnetworks.com. It does not include the hostname portion (www, www2, ftp, and so forth). This is used with Host Name (see “Virtual Server Services tab fields” on page 237) to create a full host/domain name for individual services.</p> |
| Cookie Name | Sets the cookie name of the virtual server used for cookie load balancing. The maximum string length is 20 characters. The default is none. Use an asterisk in the name for a wildcard, for example, ASPsession*. |

Table 64 Virtual server fields (continued)

| Field | Description |
|------------------------------|--|
| Cookie Offset | Sets the starting byte offset of the cookie value: 0 to 64 bytes. The default is 0. |
| Cookie Length | Sets the number of bytes to extract from the cookie value—0 to 64 bytes. The default is 0. |
| Cookie Search | Enables or disables cookie search in URI. The default is disabled. <ul style="list-style-type: none"> • Enable— looks for cookie name/value pair in the URI. • Disabled—looks for the cookie in the HTTP header. |
| Hash Length | Sets the number of bytes to hash to the server—0 to 255 bytes. A value of 0 (zero) disables URL hashing. The default is 0. |
| HTTP Header Name | Sets the HTTP header name of the virtual server. The default is none. Allowed string length is 0 to 32 characters. |
| Cookie Search Response Count | Sets the number of cookie search responses—1 to 16. The default is 1. |
| Expire Cookie | Sets expiration parameters for the Insert cookie persistence mode as follows: <ul style="list-style-type: none"> • Absolute—a date <MM/dd/yy[@hh:mm]> (e.g. 12/31/01@23:59) Once the expiration date is reached, the cookie is not stored or given out. The date and time is based on RFC 822, RFC 850, RFC 1036, and RFC 1123, for GMT. • Relative—a duration <days[:hours[:minutes]]> (e.g. 45:30:90) Defines the amount of elapsed time from when the cookie was created until it expires. • none—The cookie expires when the user's session ends. When a client sends a request <i>without</i> a cookie, the server responds with the data, and the switch inserts a <i>persistence cookie</i> into the data packet. The switch uses this cookie to bind to the appropriate server. |
| BWM Contract | Sets the virtual server's default Bandwidth Management contract—1 to 1,024. The default is 1024. By default, all services under this virtual server are assigned this BWM contract. For more information about Bandwidth Management, see "Bandwidth management" on page 493. |
| Services | Opens the <i>Virtual Server Services</i> dialog box, listing all services for the virtual server. From the <i>Virtual Server Services dialog box</i> , you modify or add services for the virtual server. |

Configuring services for a virtual server

To configure services for a virtual server:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the [General tab](#).

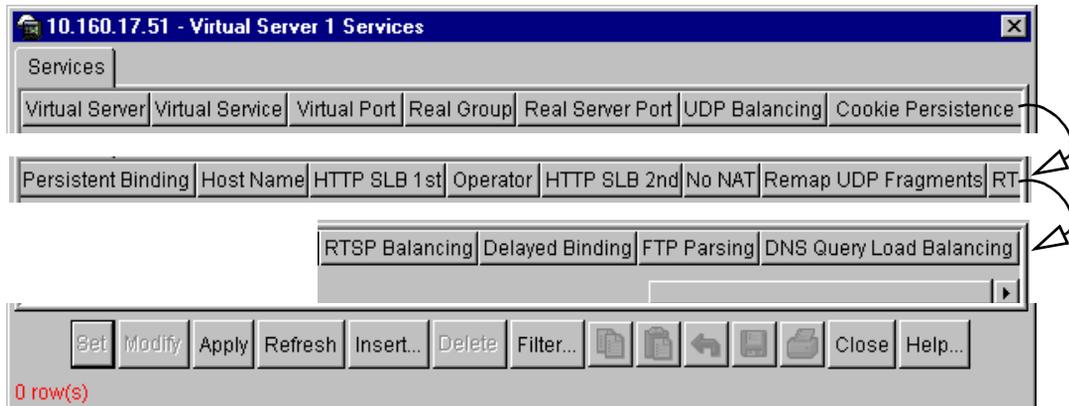
- 3 Click the Virtual Server tab.

The [Virtual Server tab](#) (Figure 85) opens with the fields listed in [Table 64 on page 233](#).

- 4 Click Services.

The Virtual Server Services dialog box (Figure 87) opens.

Figure 87 Virtual Server Services dialog box



- 5 Click Insert.

The [Insert Virtual Server Services dialog box](#) (Figure 88) opens with the fields described in [Table 65 on page 237](#).

Figure 88 Virtual Server Services, Insert Services dialog box

The dialog box is titled "134.177.214.201 - Virtual Server 1 Services, Insert Services". It contains the following fields and options:

- Virtual Server: 1 (range 1..256) with a "Browse..." button.
- Virtual Service: 1.8
- Virtual Port: 2 (range 2..65534)
- Real Group: 1 (range 1..256) with a "Browse..." button.
- Real Server Port: 0 (range 0..65534)
- UDP Balancing: disabled (dropdown)
- Cookie Persistence Mode: disabled (dropdown)
- Persistent Binding: disabled (dropdown)
- Host Name: (text field)
- HTTP SLB 1st: disabled (dropdown)
- Operator: none (dropdown)
- HTTP SLB 2nd: disabled (dropdown)
- No NAT: enabled disabled
- Remap UDP Fragments: enabled disabled
- RTSP Balancing: disabled (dropdown)
- Delayed Binding: enabled disabled
- FTP Parsing: enabled disabled
- DNS Query Load Balancing: enabled disabled
- BWM Contract: 1024 (range 0..1024) with a "Browse..." button.

Buttons at the bottom: Insert, Close, Help...

- 6 In the Virtual Server field, type the number of the virtual server or click Browse to view a selection list of available servers. For more information, see [“Browsing a read only dialog box” on page 76](#).
- 7 In the Virtual Service field, type a number (1 - 8) for the service.
- 8 In the Virtual Port field, type the number of the port (1 - 65534) that will be used for this service.

9 In the Real Group field, type a group number, or click Browse to view a selection list of available groups. For more information, see [“Browsing a read only dialog box” on page 76](#).

10 Click Insert.

The Virtual Server Services, Insert Services dialog box closes and the service appears in the Virtual Server Services tab.

11 From the Virtual Server Services tab, click Apply > Close.

The service is configured and the Virtual Server Services dialog box closes.

12 Click close.

The Server Load Balance dialog box closes.

[Table 65](#) describes the SLB—Virtual Server Services fields.

Table 65 Virtual Server Services tab fields

| Field | Description |
|------------------|--|
| Virtual Server | The number of the virtual server: 1 to 256. |
| Virtual Service | The number of the virtual service: 1 to 8. |
| Virtual Port | The number of the virtual port: 2 to 65,534. |
| Real Group | The number of the real group: 1 to 256. |
| Real Server Port | Defines the real server TCP or UDP port assigned to this service. By default, this is the same as the virtual port (service virtual port). If the real port is different from the virtual port, the switch will map the virtual port to this real port. The port range is 0 to 65,534. |
| UDP Balancing | Enables or disables UDP load balancing for a virtual port (disabled by default). You can configure this option if the service(s) to be load balanced include UDP and TCP. For example, DNS uses UDP and TCP. In those environments, you must activate UDP balancing for the particular virtual servers with which clients will communicate using UDP. The options are enabled, disabled, stateless. The default is disabled. |

Table 65 Virtual Server Services tab fields (continued)

| Field | Description |
|-------------------------|--|
| Cookie Persistence Mode | <p>Sets cookie persistence mode:</p> <ul style="list-style-type: none"> • Rewrite: In active cookie mode (or cookie rewrite mode), the switch, and not the network administrator, generates the cookie value on behalf of the server. The switch intercepts this persistence cookie and rewrites the value to include server-specific information before sending it to the client. NOTE: In Rewrite mode, the Virtual Server parameter Cookie Length can only be 8 or 16. Also, the Virtual Server parameter Cookie Offset is invalid in rewrite mode. • Passive: In this mode, the network administrator configures the Web server to embed a cookie in the server response that the switch looks for in subsequent requests from the same client. • Insert: When a client sends a request without a cookie, the server responds with the data, and the switch inserts a persistence cookie into the data packet. The switch uses this cookie to bind to the appropriate server. NOTE: If set to Insert mode, the Virtual Server parameters Cookie Name, Cookie Offset, Cookie Length and Expire Cookie are invalid. • Disabled. The default is disabled. |
| Persistent Binding | <p>Set persistent binding: enabled, disabled, session id, or cookie. The default is disabled.</p> <p>NOTE: If persistent binding type is not set to Cookie, the Virtual Server parameters Cookie Name, Cookie Offset, Cookie Length and Expire Cookie are invalid.</p> |
| Host Name | <p>Sets the host name for a service. This is used with Domain Name (see “Virtual server fields” on page 233) to create a full host/domain name for individual services. The maximum string length is 9 characters. The default is none.</p> |

Table 65 Virtual Server Services tab fields (continued)

| Field | Description |
|---------------------|--|
| HTTP SLB 1st | <p>Sets the HTTP load balancing precedence as one of the following.</p> <ul style="list-style-type: none"> • disabled— (the default setting) Disables HTTP SLB configuration. • urlslb—Enables URL SLB. • urlhash—Enables hashing based on URL. • cookie—Enables cookie-based SLB for cookie-based preferential load balancing. You must also set the following virtual server settings: Cookie name, starting point of the cookie value, number of bytes to be extracted, enable/disable checking for cookie in URL. • headerhash—hashes on any http header value. • host—Enables virtual hosting. • browser—Enables SLB based on browser type. • others—Requires inputs for a particular header field. <p>Use with the Operator and HTTP SLB 2nd fields to combine or select applications.</p> |
| Operator | <p>Sets the logical operator (and, or, none) used to combine applications selected for the 1st HTTP SLB and the 2nd HTTP SLB. The default is none.</p> |
| HTTP SLB 2nd | <p>Sets 2nd HTTP load balancing consideration as one of the following: disabled, URLSLB, URLhash, cookie, headerhash, host, browser, or others. The default is disabled. See HTTP SLB 1st (above) for definitions. Use with the Operator and HTTP SLB 1st fields to combine or select applications.</p> |
| No NAT | <p>Enables or disables substitution of only the real server's MAC address (disabled by default). This option does not substitute IP addresses. It is used for Direct Server Return (DSR) in an one-armed load balancing setup, so that frames returning from server to the client do not have to pass through the switch. The default is disabled.</p> |
| Remap UDP Fragments | <p>Enables or disables remapping UDP Fragments.</p> |
| RTSP Balancing | <p>Enables (hash or pattern) or disables Real Time Streaming Protocol (RTSP) URL balancing. RTSP Balancing can enable URL hashing using the entire URL except the extension (.xxx) at the end of the URL. The default is disabled.</p> |
| Delayed Binding | <p>Enables or disables Layer 4 Delayed Binding for TCP service and ports. Enabling this command protects the server from Denial of Service (DoS) attacks. This option is disabled by default. The default is disabled.</p> |

Table 65 Virtual Server Services tab fields (continued)

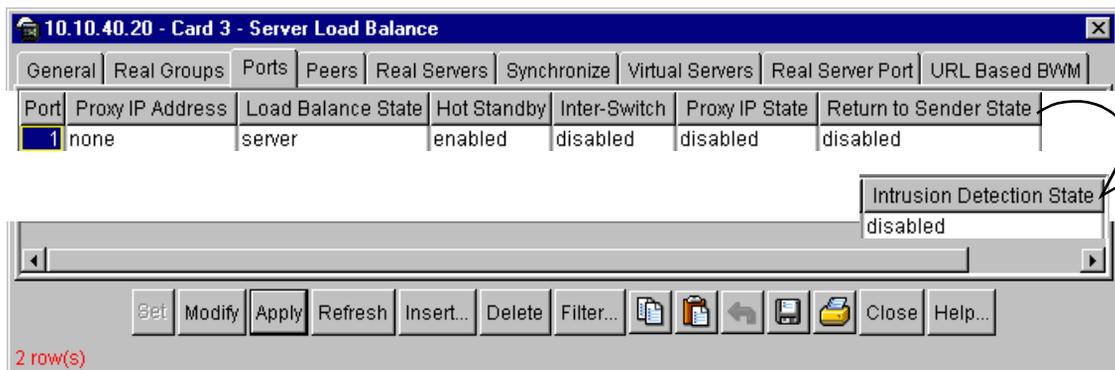
| Field | Description |
|--------------------------|--|
| FTP Parsing | Enables or disables FTP SLB parsing for this virtual server. When this option is enabled, the WSM modifies the appropriate FTP method/command to support FTP servers on a private network for both active and passive FTP modes. To do this, the switch looks deeper into the packet and modifies the port command for active FTP or the “entering the passive mode” command for passive FTP. The default is disabled. |
| DNS Query Load Balancing | Enables or disables DNS-based Layer 7 content load balancing. |
| BWM Contract | Sets the number of the bandwidth management contract. For more information, see “Bandwidth management” on page 493 . |

Configuring ports for server load balancing

To configure ports for server load balancing:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the [General tab](#).
- 3 Click the Ports tab.

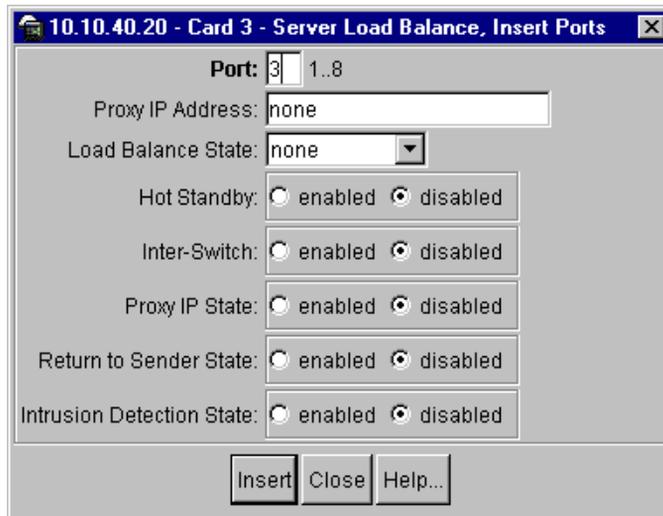
The [Ports tab](#) ([Figure 89](#)) opens with the fields listed in [Table 66 on page 242](#).

Figure 89 Server Load Balance—Ports tab

- 4 Click Insert.

The **Insert Ports dialog box** (Figure 90) opens with the fields listed in Table 66 on page 242.

Figure 90 Insert Ports dialog box



- 5 In the Port field, type the port number.
- 6 In the Load Balance State field, click the down arrow and choose a state from the list (None, Client, Server, or Client-Server).
- 7 Use the other fields on the tab to further define the port.
- 8 Click Insert.
The port is inserted into the Ports tab and the Insert Ports dialog box closes.
- 9 From the Ports tab, click Apply > Close.
The port is configured and the Server Load Balance dialog box closes.

Table 66 describes the Server Load Balance Ports fields.

Table 66 SLB—Ports fields

| Field | Description |
|--------------------|---|
| Port | The port number. |
| Proxy IP Address | Sets the proxy IP address for this port, using dotted decimal notation. When defined, client address information in Layer 4 requests is replaced with this proxy IP address. |
| Load Balance State | The SLB state of the port. <ul style="list-style-type: none"> • none: No SLB port. This is the default setting. • client: Binds servers to clients and provide address translation from the virtual server IP address to the real server IP address, re-mapping virtual server IP addresses and port values to real server IP addresses and ports. Traffic not associated with virtual servers is switched normally. Maximizing the number of these ports on the Layer 4 switch will improve the switch's potential for effective Server Load Balancing. • server: Re-maps real server IP addresses and Layer 4 port values to virtual server IP addresses and Layer 4 ports. Traffic not associated with virtual servers is switched normally. Do not use server processing on ports configured for application redirection or client ports configured with proxy IP addresses. • client-server: SLB client and server port. |
| Hot Standby | Enables or disables hot-standby processing on the switch port. Use in conjunction with VRRP hot-standby failover. Disabled by default. |
| Inter-Switch | Enables or disables inter-switch processing. This option is enabled for ports connected to a peer switch and is disabled by default. Used only when configuring the WSM in VRRP Hotstandby. |
| Proxy IP State | Enables or disables a proxy for traffic that ingresses this port. When defined, client address information in Layer 4 requests is replaced with this proxy IP address. In SLB, this forces response traffic to return through the switch, rather than around it, as is possible in complex routing environments. Proxies are also useful for Application Redirection and Network Address Translation (NAT). When used with Application Redirection filters, each filter's real port must also be defined. Proxy IP State is disabled by default. |

Table 66 SLB—Ports fields (continued)

| Field | Description |
|---------------------------|--|
| Return to Sender State | Enables or disables Return to Sender (RTS) load balancing on this port. This option is used for firewall load balancing or VPN load balancing applications. Enable RTS on all client-side ports to ensure that traffic ingresses and egresses through the same port. This option is disabled by default. |
| Intrusion Detection State | Enables or disables Intrusion Detection System Server Load Balancing. This option is disabled by default. |

Enabling or disabling server load balancing

To enable/disable server load balancing on the WSM:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the **General** tab (Figure 81 on page 223) with the fields described in Table 61 on page 224.

- 3 In the Server Load Balance field, click Enabled or Disabled.

- 4 Click Set > Apply > Close.

SLB is enabled/disabled on the WSM and the Server Load Balance dialog box closes.

Peers fields

[Table 67](#) describes the fields on the SLB—Peers tab and the Insert Peers dialog box.

Table 67 Peers fields

| Field | Description |
|-----------------|---|
| Peer | The index number (1 or 2) of the synchronization for the peer switch. |
| IP Address | The IP address of the peer switch. The default is 0.0.0.0. |
| Synchronization | Enables or disables synchronization with the peer switch. |

Synchronize tab fields

[Table 68](#) describes the fields on the SLB—Synchronize tab.

Table 68 Synchronize tab fields

| Field | Description |
|---------------------------------|---|
| Synch Filter | Enables or disables the configuration filter synchronization. |
| Synch Port | Enables or disables synch port configuration. |
| Synch VRRP Priorities | Enables or disables synch VRRP priorities. |
| Synch Proxy IP Address | Enables or disables synch proxy IP addresses. |
| Sync Stateful Failover | Enables or disables the sync stateful failover. |
| Stateful Failover Update Period | Sets the period (in seconds) of the stateful failover update—1 to 60. |
| Synch Bandwidth Management | Enables or disables synchronization bandwidth management. |

Chapter 9

Extending SLB topologies

In standard server load balancing, all client-to-server requests to a particular virtual server and all related server-to-client responses must pass through the same WSM. To prevent routers from creating alternate paths around the WSM, you can extend its SLB topologies by using the following configurations.

- [“Configuring proxy IP addresses” next](#)
- [“Mapping Ports” on page 248](#)
- [“Providing direct access to real servers” on page 254](#)
- [“Using delayed binding to prevent DoS attacks” on page 262](#)

Configuring proxy IP addresses

When requesting services from the virtual server, a client sends its own IP address for use as a return address. If you configure a proxy IP address for the client port, the WSM replaces the client’s source IP address with the WSM’s own IP address as a proxy before sending the request to the real server. This makes it look to the server like the WSM originated the request.



Note: When a proxy IP address is configured, requests appear to come from the WSM rather than the client. For this reason, the network administrator performing debugging and collecting statistics must be aware that a proxy IP address is configured.

The real server uses the WSM's proxy IP address as the destination address for any response. SLB traffic is forced to return through the proper switch, regardless of alternate paths. Once the switch receives the proxied data, it puts the original client IP address into the destination address and sends the packet to the client. This process is transparent to the client.



Note: If VMA is enabled, all ports must be configured with a unique proxy address when any WSM port is configured with a proxy IP address. If VMA is disabled, only the client port needs a proxy IP address. Port 9 does not require a proxy address when VMA is enabled. See, “[Improving WSM performance with VMA](#)” on page 323.

The proxy IP address can also be used for direct access to the real servers (see “[Providing direct access to real servers](#)” on page 254).

Disabling server processing on a switch port

When implementing proxy IP addresses for client ports, you can reconfigure WSM ports to disable server processing.

To disable server processing on a WSM port:

- 1 From the device view, select the Web Switching Module.

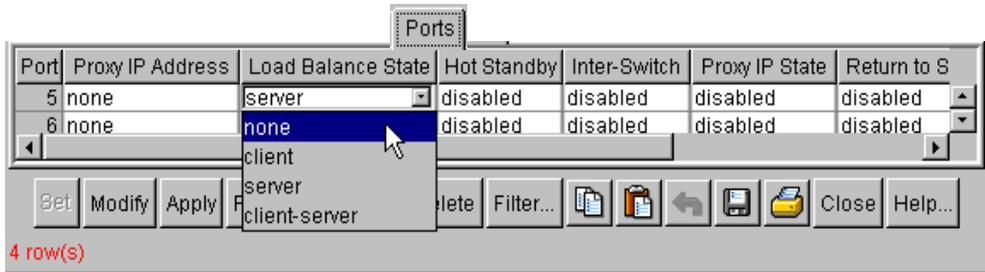
The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the [General tab](#).

- 3 Click the Ports tab.

The [Ports tab](#) ([Figure 91](#)) opens with the fields defined in [Table 66 on page 242](#).

Figure 91 SLB—disabling server processing on a port

- 4 Click the Load Balance State field for the server port, and choose None.
- 5 To disable server processing on other ports, repeat step 4.
- 6 Click Set > Apply > Close.

Server processing is disabled for the port(s) and the Server Load Balance dialog box closes.

Configuring a proxy IP address for a port

You may first want to disable server processing on affected WSM ports when implementing proxy IP addresses. For more information, see [“Disabling server processing on a switch port” on page 246](#).

To configure a proxy IP address for a client port:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the [General](#) tab.
- 3 Click the Ports tab.
The [Ports](#) tab ([Figure 91](#)) opens with the fields defined in [Table 66 on page 242](#).
- 4 Click the Proxy IP Address field for the port, and type a proxy IP address.
- 5 In the Proxy IP State field for the port, click the down arrow and choose Enabled.
- 6 To configure a proxy IP address for other client ports, repeat steps 4 and 5.

- 7 Click Set > Apply > Close.

The proxy IP address is configured and the Server Load Balance dialog box closes.

Mapping Ports

To improve security, you can hide the identity of a port by mapping a virtual server port to a different real server port. This section includes the following port mapping topics.

- [“Mapping a virtual server port to a real server port,”](#) next
- [“Mapping multiple real server ports”](#) on page 250

Mapping a virtual server port to a real server port

Mapping is required if you run real server processes on ports other than the well-known TCP/UDP ports. If your real server processes use well-known TCP/UDP ports, the virtual server ports are mapped directly to real server ports by default and do not require manual mapping. For a list of their port numbers, see [“Well-known TCP/UDP application port numbers”](#) on page 191.



Note: Port mapping is supported with [Direct Access Mode \(DAM\)](#) if you enabled filtering, configured a proxy IP address, or enabled URL parsing on any port. See [“Configuring direct access mode”](#) on page 258.

To map a virtual server port to a real server port:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3 Click the Virtual Servers tab.
The Virtual Servers tab ([Figure 92](#)) opens.

Figure 92 Server Load Balance—Virtual Servers tab

| Virtual Server | IP Address | State | Layer 3 Only | Domain Name | Cookie Name | Cookie Offset | Cookie |
|----------------|-----------------|---------|--------------|-------------|-------------|---------------|--------|
| 1 | 148.148.148.148 | enabled | disabled | | | 1 | |
| 2 | 148.148.148.148 | enabled | disabled | | | 1 | |

Services Set Modify Apply Refresh Insert... Delete Filter... Close Help...

2 row(s)

- 4 Select the virtual server for which you want to map a port.

The Virtual Server number is highlighted.

- 5 Click Services.

The Virtual Server Services dialog box (Figure 87 on page 235) opens with the fields described in Table 65 on page 237.

Figure 93 Virtual server services port mapping

| Virtual Server | Virtual Service | Virtual Port | Real Group | Real Server Port | UDP Balancing | Cookie Pers |
|----------------|-----------------|--------------|------------|------------------|---------------|-------------|
| 1 | 1 | 80 | 1 | 80 | disabled | disabled |

Set Modify Apply Refresh Insert... Delete Filter... Close Help...

1 row(s)

- 6 Double click the Real Server Port field for the service.

The field is editable.

- 7 Enter a port number.

- 8 Click Set > Apply > Close.

The Virtual Server Services dialog box closes and the port is mapped.

- 9 From the Virtual Servers tab, click Close.

The Server Load Balance dialog box closes.

Mapping multiple real server ports

You can map a single WSM virtual server port to as many as 16 real server ports in Layer 4, Layer 7, and in cookie-based and SSL persistence switching environments. When multiple real server ports are mapped to a virtual port, the WSM treats the real server IP address/port mapping combination as a distinct real server.

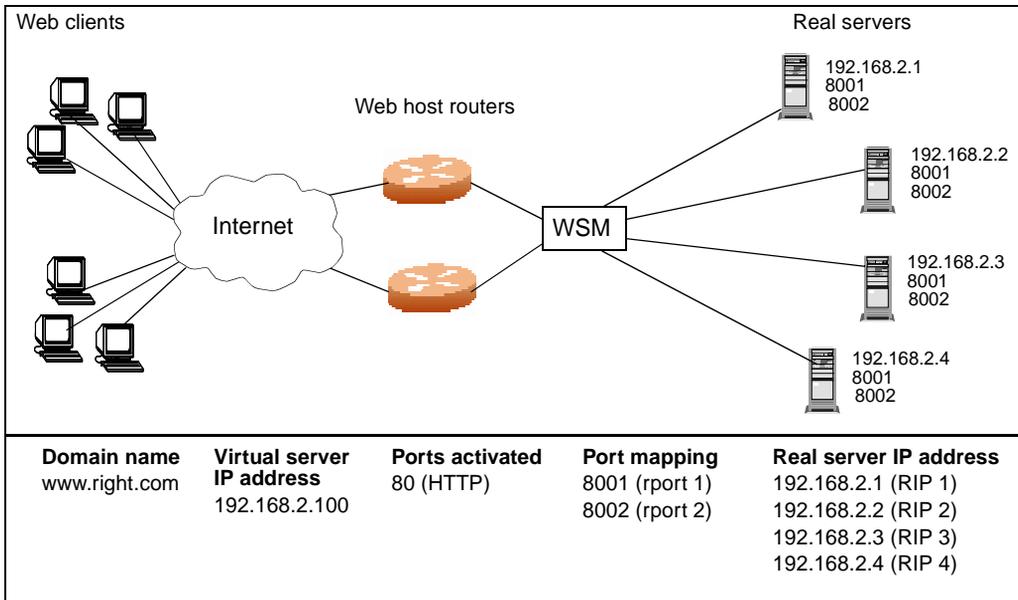


Note: For each real server, you can only configure one service with multiple real ports.

Multiple port mapping example

In [Figure 94](#), four real servers support a single service, HTTP. Clients access HTTP services through a virtual server, with IP address 192.168.2.100, on virtual port 80. Since each real server uses two ports for HTTP services, ports 8001 and 8002, the logical real servers are:

| | | | |
|------------------|------------------|------------------|------------------|
| 192.168.2.1/8001 | 192.168.2.2/8001 | 192.168.2.3/8001 | 192.168.2.4/8001 |
| 192.168.2.1/8002 | 192.168.2.2/8002 | 192.168.2.3/8002 | 192.168.2.4/8002 |

Figure 94 Basic virtual port to real port mapping configuration

The WSM selects a real server using the configured load balancing metric. For more information, see [“About load balancing metrics” on page 207](#).

The real server TCP or UDP port assigned to a service is the same as the virtual port. If you configure a different virtual port via mapping, the WSM maps the virtual port to the real port.



Note: If you map a single virtual port to multiple real ports, set the virtual server’s real server port value to 0. You cannot configure multiple services with multiple real port mapping on this same server.

Configuration steps for mapping multiple ports

The steps for mapping multiple ports to a single virtual port are:

- 1 Configure the real servers whose ports will be mapped to the single virtual port. See [“Configuring each real server” on page 216](#).
- 2 Add these real servers to a group. See [“Configuring a real server group” on page 225](#).

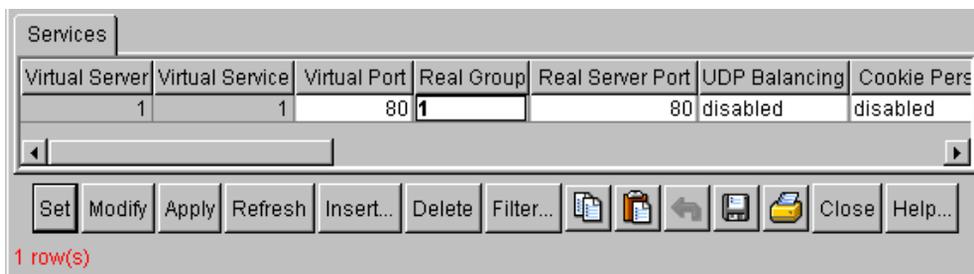
- 3 Configure the virtual server IP address, and enable the virtual server. See “Configuring a virtual server” on page 230.
- 4 Enable multiple real ports for the virtual server by setting its Real Server Port value to 0. See “Configuring multiple port mapping,” next.
- 5 Associate the real server group with the virtual server. See “Configuring multiple port mapping,” next.
- 6 Add the real ports to which the virtual server listens. See “Configuring multiple port mapping,” next.

Configuring multiple port mapping

To map multiple real server ports to a virtual server port:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the **General** tab.
- 3 Click the Virtual Servers tab.
The Virtual Servers tab (Figure 92 on page 249) opens.
- 4 Select the virtual server for which you want to map a port.
The virtual server number is highlighted.
- 5 Click Services.
The Virtual Server Services dialog box (Figure 95) opens with the fields defined in Table 65 on page 237.

Figure 95 SLB multiple port mapping



- 6 Double click the Real Group field.

The field is editable.

- 7 Enter the number of the real server group.

- 8 Double click the Real Server Port field

The field is editable.

- 9 Enter 0 to enable multiple real ports for the virtual server.

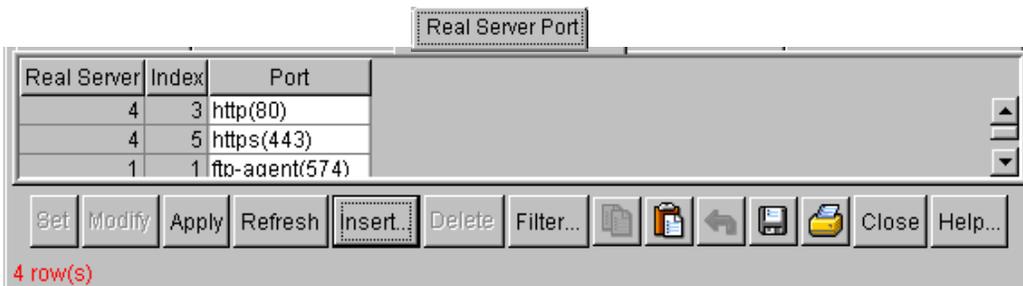
- 10 Click Set > Apply > Close.

Multiple port mapping is enabled for the service and the Virtual Server Services dialog box closes.

- 11 In the Server Load Balance dialog box, click the Real Server Port tab.

The **Real Server Port** tab (Figure 96) opens with the fields defined in Table 69 on page 254.

Figure 96 Server Load Balance—Real Server Port tab



Note: If you have many real servers configured and want to view a subset of this table, see [“Filtering table content” on page 73](#).

- 12 Click Insert.

The Insert Real Server Port dialog box opens.

Figure 97 Server Load Balance, Insert Real Server Port dialog box

- 13** In the Real Server field, type the real server number for the first real server in the group, or click Browse to choose from a selection list.
- 14** In the Index field, type an index number for this server.
- 15** In the Port field, type the port number, or click Browse to choose from a selection list.
- 16** Repeat steps 13 - 15 for all ports and real servers in the group.
- 17** Click Insert.

The real server ports are added to the Real Server Port tab and the Insert Real Server Port dialog box closes.

- 18** In the Server Load Balance dialog box, click Apply > Close.

The real ports to which the virtual server listens are mapped, and the Server Load Balance dialog box closes.

[Table 69](#) describes the real server port fields.

Table 69 Real Server Port fields

| Field | Description |
|-------------|---|
| Real Server | Sets the real server port number—1 to 1023. |
| Index | Sets the index number of the service—1 to 16. |
| Port | Sets the Layer 4 real service port number—2 to 65534. |

Providing direct access to real servers

The following topics describe direct access to real servers and how to configure it:

- [“About direct server return” next](#)
- [“Configuring direct server return” on page 256](#)
- [“Configuring direct access mode” on page 258](#)
- [“Assigning multiple IP addresses” on page 260](#)
- [“Using proxy IP addresses” on page 260](#)
- [“Mapping ports” on page 260](#)
- [“Monitoring real servers and services” on page 261](#)

About direct server return

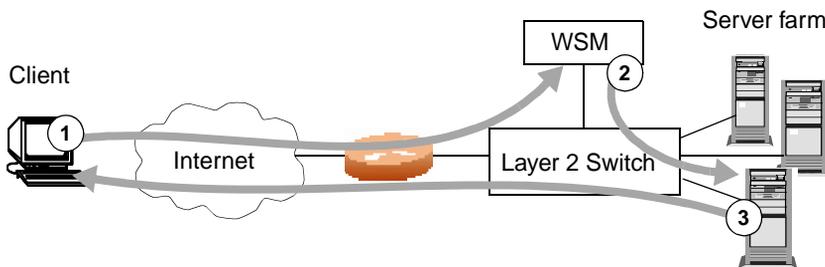
Some clients need direct access to real servers, for example, to monitor a real server from a management workstation. Direct Server Return (DSR) lets the server respond directly to the client. This is useful for sites where large amounts of data flow from servers to clients, such as content providers or portal sites with asymmetric traffic patterns. DSR and content-intelligent Layer 7 switching cannot be performed at the same time because content-intelligent switching requires that all frames go back to the switch for connection splicing.



Note: For DSR, you must set up the server to receive frames that have a destination IP address equal to the virtual server IP address.

Figure 98 illustrates the following DSR operation.

- 1 A client request is forwarded to the WSM.
- 2 Because only MAC addresses are substituted, the WSM forwards the request to the best server, using the configured load-balancing policy.
- 3 The server responds directly to the client, bypassing the WSM, and using the virtual server IP address as the source IP address.

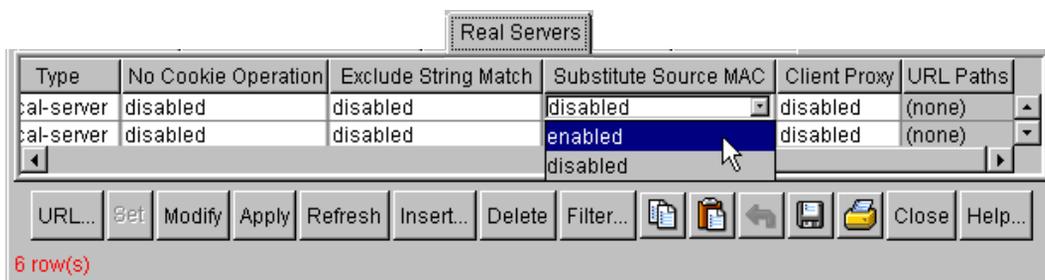
Figure 98 Direct server return

Configuring direct server return

To configure DSR:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3 Click the Real Servers tab.

The Real Servers tab ([Figure 99](#)) opens with the fields listed in [Table 60](#) on [page 219](#).

Figure 99 SLB—enable source MAC substitution for real server

- 4 Click the Substitute Source MAC field for the real server you want to configure, and change the value to Enabled.
- 5 Click Set > Apply.
The real server is configured for source MAC address substitution.

- Click the Virtual Servers tab.

The Virtual Servers tab (Figure 100) opens with the fields listed in Table 64 on page 233.

Figure 100 SLB—Virtual Servers tab

| Virtual Server | IP Address | State | Layer 3 Only | Domain Name | Cookie Name | Coc |
|----------------|-----------------|---------|--------------|-------------|-------------|-----|
| 1 | 148.148.148.148 | enabled | disabled | | | |
| 2 | 48.48.48.48 | enabled | disabled | | | |

Services Set Modify Apply Refresh Insert... Delete Filter... [Document] [Clipboard] [Back] [Refresh]

2 row(s)

- Click the virtual server to configure.

The virtual server is highlighted.

- Click Services.

The Virtual Server Services dialog box (Figure 101) opens with the fields defined in Table 65 on page 237.

Figure 101 Virtual Server Services—enable MAC address substitution

| Binding | Host Name | HTTP SLB 1st | Operator | HTTP SLB 2nd | No NAT | Remap UDP Fragments | RTT |
|---------|-----------|--------------|----------|--------------|----------|---------------------|------|
| | | disabled | none | disabled | disabled | enabled | disc |

Set Modify Apply Refresh Insert... Delete Filter... [Document] [Clipboard] [Back] [Refresh] [Print] Close Help...

1 row(s)

- Click the No NAT field for the service you want to configure, and change the setting to Enabled (disables network address translation, and enables MAC address substitution).

- Click Set > Apply > Close.

Network address translation is turned off for the service and the Services dialog box closes.

- 11 In the Server Load Balance dialog box, click Close.

DSR is configured for the service.

Configuring direct access mode

When you enable Direct Access Mode (DAM) for the WSM, any client can communicate with any real server's load-balanced service; and any number of virtual services can be configured to load balance a real service. Traffic sent directly to the real server IP addresses is excluded from load-balancing decisions. The same clients may also communicate to the virtual server IP address for load-balanced requests.



Note: When DAM is enabled on the WSM, port mapping and default gateway load balancing are supported only when filtering is enabled, a proxy IP address is configured, or URL parsing is enabled on any switch port

To configure DAM:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab ([Figure 102](#)) with the fields described in [Table 61 on page 224](#).

Figure 102 Server Load Balance—enable Direct Access Mode

134.177.214.201 - Card 3 - Server Load Balance

Virtual Servers | Real Server Port | SYN Attack | URL Based BWM

General | Real Groups | Ports | Peers | RTSP | Real Servers | Synchronize

Server Load Balance: enabled disabled

Virtual Matrix Architecture: enabled disabled

Direct Access Mode: **enabled** disabled

IP Address Mask: 255.255.255.255

Management Network: 0.0.0.0

Management Subnet Mask: 255.255.255.255

Radius Secret:

Persistent Mask: 255.255.255.255

Graceful Failover: enabled disabled

Session Table Fast-Age Period Bit Shift: 1 0..7

Session Table Slow-Age Period Bit Shift: 2 0..15

Transparent Proxy Cache Protocol: enabled **disabled**

Metric (Response & Bandwidth) Update Interval: 10 1..256

LDAP Version: version2 version3

Allow HTTP Health Check on any port?: enabled disabled

Time Window: 1 1..65535

Hold Duration: 1 1..65535

Intrusion Detection Real Server Group: 1 1..256

Set Apply Refresh Close Help...

3 In the Direct Access Mode field, click Enabled.

4 Click Set > Apply > Close.

DAM is configured for the Web Switching Module, and the Server Load Balance dialog box closes.

Assigning multiple IP addresses

One way to provide both SLB access and direct access to a real server is to assign multiple IP addresses to the real server. For example, one IP address could be established exclusively for SLB and another could be used for direct access needs.

Using proxy IP addresses

Proxy IP addresses can eliminate SLB topology restrictions in complex networks (see [“Configuring proxy IP addresses” on page 245](#)). Proxy IP addresses can also provide direct access to real servers.

If the WSM port to the client is configured with a proxy IP address, the client can access each real server directly using the real server’s IP address. The WSM port must be connected to the real server and client processing must be disabled (see [“Configuring ports for server load balancing” on page 240](#)). SLB is still accessed using the virtual server IP address.

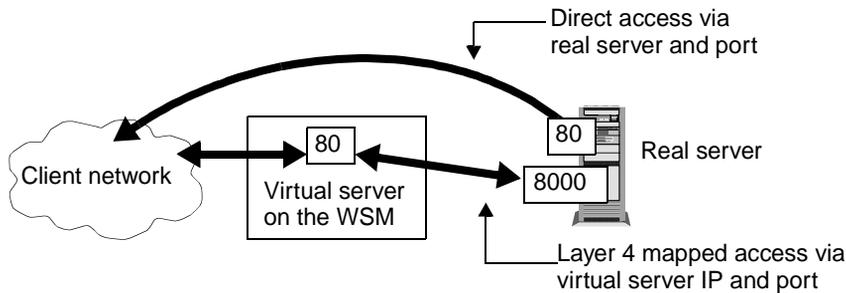
Mapping ports

When SLB is used without proxy IP addresses and without DAM, the WSM must process server-to-client responses.

If clients bypass WSM processing and access the real server IP address and port directly, the server-to-client response could be mishandled by SLB processing as it returns through the WSM. The real server IP address could be remapped back to the virtual server IP address on the WSM.

First, the real server must perform two port processing—one real server port handling the direct traffic, and the other handling SLB traffic. Then, the virtual server port on the WSM must be mapped to the real server port.

In [Figure 103](#), clients can access SLB services through well-known TCP port 80 at the virtual server’s IP address. The WSM, behaving like a virtual server, is mapped to TCP port 8000 on the real server. For direct access that bypasses the virtual server and SLB, clients can specify well-known TCP port 80 as the real server’s IP address.

Figure 103 Mapped and non-mapped server access

Note: Port mapping is supported with DAM when filtering is enabled, a proxy IP address is configured, or URL parsing is enabled on any WSM port.

To configure port mapping, see [“Mapping a virtual server port to a real server port”](#) on page 248.

Monitoring real servers and services

To configure monitoring of real servers and services by the management network:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab ([Figure 102 on page 259](#)) with the fields described in [Table 61 on page 224](#).

- 3 In the Management Network field, enter the source IP address of the management network to allow management traffic direct (non-Layer 4) access to real servers.

- 4 In the Management Subnet Mask field, enter the IP address of the subnet mask that gives management traffic direct access to real servers.



Note: Clients on the management network do not have access to SLB services and cannot access the virtual services being load balanced.

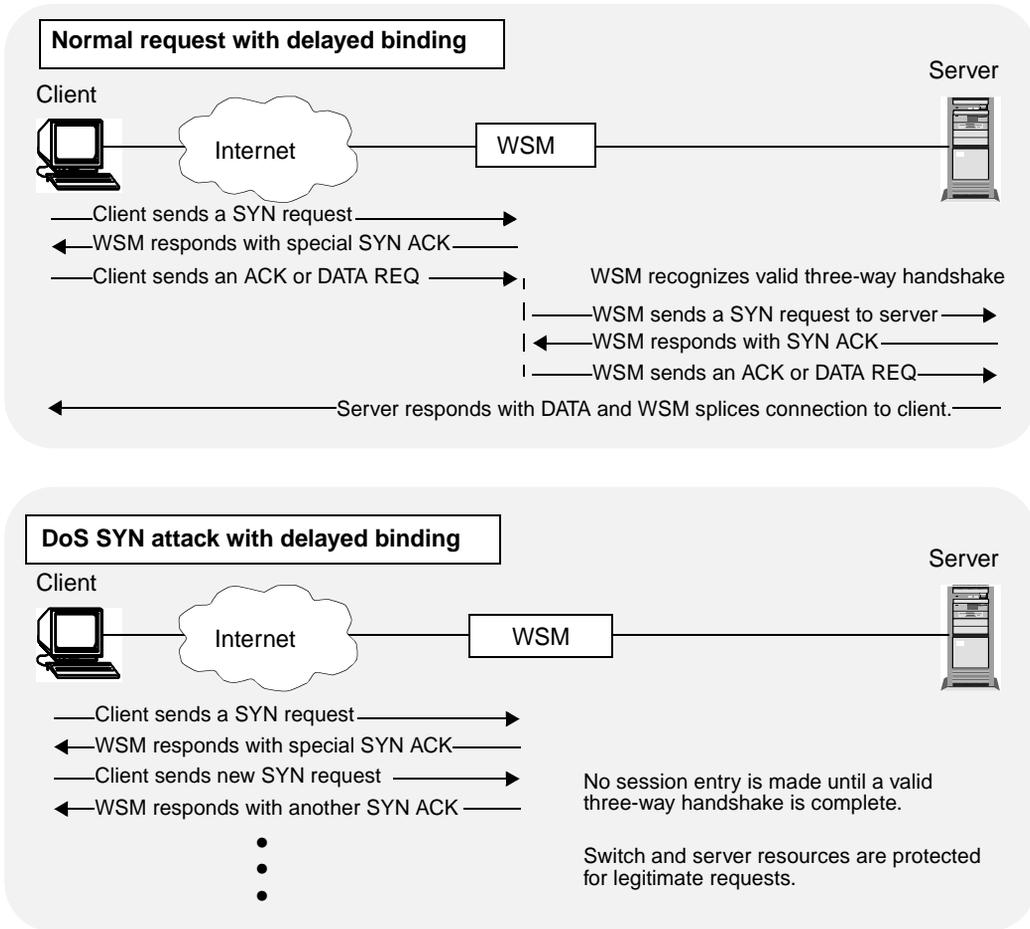
Using delayed binding to prevent DoS attacks

Delayed binding prevents SYN Denial of Service (DoS) attacks which occur when a client saturates a server with repeated SYN requests instead of completing the following three-way handshake as expected.

1. The client sends out a synchronization (SYN) request to the server.
2. The server allocates an area to process the client requests, and acknowledges the client by sending a SYN ACK.
3. The client then completes the three-way handshake by sending an acknowledgement (ACK) back to the server.

Figure 105 illustrates a WSM, configured with delayed binding, intercepting the client SYN request before it reaches the server. The WSM responds to the client with a SYN ACK containing embedded client information. The WSM does not allocate a session until a valid SYN ACK is received from the client or the three-way handshake is complete.

Figure 105 Repelling DoS SYN attacks with delayed binding



Once the WSM receives a valid ACK or DATA REQ from the client, it sends a SYN request to the server on behalf of the client, waits for the server to respond with a SYN ACK, and then forwards the client's DATA REQ to the server. It delays binding the client session to the server until the expected handshakes are

complete. Thus, with delayed binding, two independent TCP connections span a Web session—one from the client to the WSM and the second from the WSM to the selected server. The WSM temporarily terminates each TCP connection until content has been received, preventing the server from being inundated with SYN requests.



Note: Delayed binding is automatically enabled when content intelligent switching features are used. However, if you are not parsing content and you want delayed binding enabled, you must enable it manually.

Enabling delayed binding for a server service

To enable delayed binding for a server service:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3 Click the Virtual Server tab.

The Virtual Servers tab ([Figure 106](#)) opens with the fields defined in [Table 64 on page 233](#).

Figure 106 Server Load Balance—Virtual Servers tab

| Virtual Server | IP Address | State | Layer 3 Only | Domain ... | Cookie ... | Cookie Offset | Cookie Length |
|----------------|-----------------|---------|--------------|------------|------------|---------------|---------------|
| 1 | 148.148.148.148 | enabled | disabled | | | 1 | 0 c |
| 2 | 48.48.48.48 | enabled | disabled | | | 1 | 0 c |

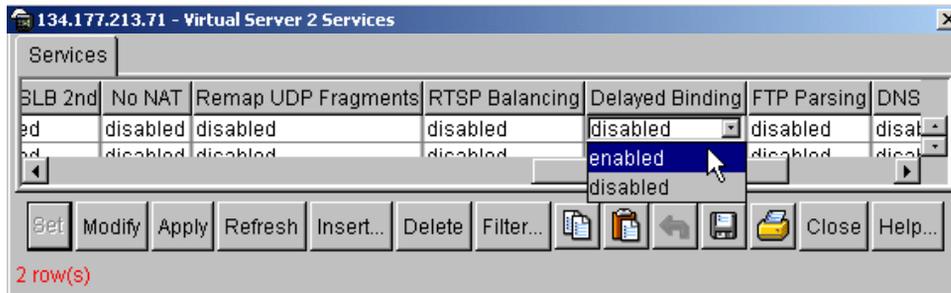
Services... Set Modify Apply Refresh Insert... Delete Filter... [Icons] Close Help...

2 row(s)

- 4 Click the virtual server to configure.
The virtual server is highlighted.
- 5 Click Services.

The Virtual Server Services dialog box (Figure 107) opens with the fields defined in Table 65 on page 237.

Figure 107 Virtual Server Services dialog box—delayed binding



- 6 Click the Delayed Binding field for the virtual service you want to configure, and select Enabled.
- 7 Click Set > Apply > Close.

Delayed binding is enabled for the service.



Note: Enable delayed binding without configuring any HTTP SLB processing or persistent binding types.

For more information, see [“Configuring delayed binding for Web Cache Redirection”](#) on page 398.

Detecting SYN attacks

When delayed binding is enabled for the WSM, SYN attack detection is enabled by default. SYN attack detection tracks half-open connections, triggers a trap when the configured threshold is exceeded, and monitors DoS attacks. For information about viewing the total number of half-open sessions, see [“Statistics”](#) on page 531.

Half-open sessions show an incomplete three-way handshake between the server and the client. The probability of a SYN attack is higher when excessive half-open sessions are generated. To detect SYN attacks, the WSM keeps track of the number of new half-open sessions for a set period of time. If the value exceeds the threshold, then a syslog message and an SNMP trap are generated.

Configuring SYN attack detection

You can specify how frequently you want to check for SYN attacks and modify the default threshold for the number of new half-open sessions per second.

To configure SYN attack detection:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

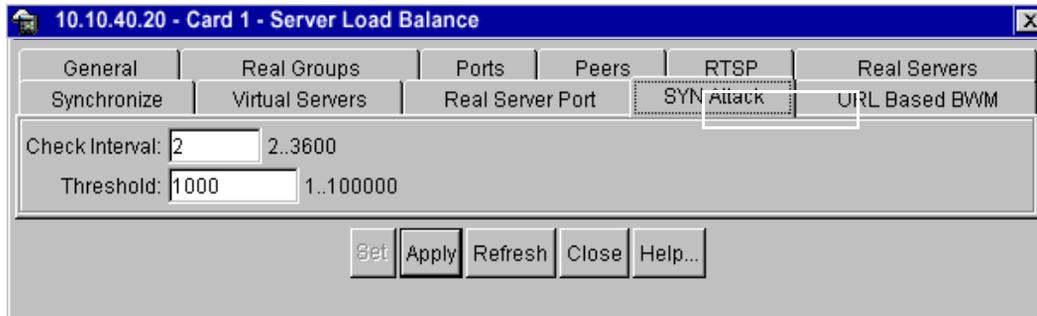
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab.

- 3 Click the SYN Attack tab.

The **SYN Attack** tab (Figure 108) opens with the fields defined in Table 70 on page 268.

Figure 108 Server Load Balance—SYN Attack tab



- 4 In the Check Interval field, type how often (2 seconds to 3,600 seconds) you want the WSM to check for SYN attacks.
- 5 In the Threshold field, type the acceptable threshold per second (1 to 100,000) for new half-open sessions.
- 6 Click Set > Apply > Close.

The SYN attack detection is configured.

[Table 70](#) describes the fields on the SLB—SYN Attack tab.

Table 70 SYN Attack tab fields

| Field | Description |
|----------------|--|
| Check Interval | Sets the interval, in seconds (2 to 3,600), that determines how frequently the WSM checks for SYN attacks. |
| Threshold | Sets the acceptable threshold per second (1 to 100,000) for new half-open sessions. |

Chapter 10

Load balancing special services

This section includes the following topics.

- [“IP server load balancing,”](#) next
- [“FTP server load balancing”](#) on page 272
- [“Domain name server load balancing”](#) on page 277
- [“Real Time Streaming Protocol server load balancing”](#) on page 283
- [“Wireless application server load balancing”](#) on page 292
- [“Intrusion Detection server load balancing”](#) on page 302
- [“WAN link load balancing”](#) on page 313

IP server load balancing

You can configure server load balancing based on a client’s IP address, which is generally used with the client port number to produce a session identifier. You can enable Layer 3 processing so the WSM uses only the client IP address as the session identifier.

When Layer 3 only processing is enabled, the WSM associates all connections from the same client with the same real server while any connection exists between them. This option is necessary for some server applications where state information about the client system is divided across different simultaneous connections, and also in applications where TCP fragments are generated. If the real server to which the client is assigned becomes unavailable, the Layer 4 software will allow the client to connect to a different server.



Note: IP server load balancing must be used if IP traffic is totally encrypted and you do not have access to the content.

Configuring IP server load balancing

To configure IP server load balancing:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

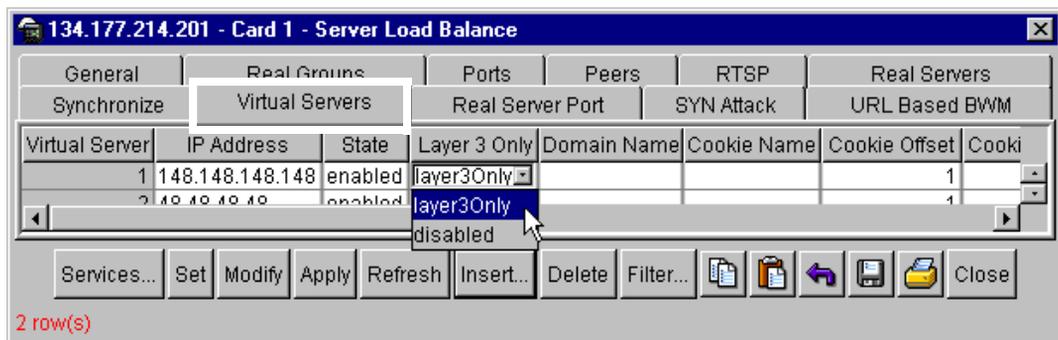
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab.

- 3 Click the Virtual Servers tab.

The Virtual Servers tab (Figure 109) opens with the fields defined in Table 64 on page 233.

Figure 109 Configuring IP SLB, layer 3 only processing



- 4 Click the Layer 3 Only field for the virtual server you want to configure, and change the value to Layer3Only.

- 5 Click Set > Apply.

Layer 3 only processing is enabled for the virtual server so that only the client IP address is used as the session identifier.

- 6 Select the Virtual Server.

The virtual server number is highlighted.

- 7 Click Services.

The Virtual Server Services dialog box opens displaying the services for the selected virtual server.

8 Click Insert.

The Insert Server Services dialog box (Figure 110) opens. For field definitions, see Table 65 on page 237.

Figure 110 Configuring IP service for SLB

The screenshot shows a dialog box titled "134.177.214.201 - Virtual Server 1 Services, Insert Services". It contains the following fields and controls:

- Virtual Server:** 1 1..256 (Browse...)
- Virtual Service:** 1 1..8
- Virtual Port:** 2 2..65534
- Real Group:** 1 1..256 (Browse...)
- Real Server Port:** 1 0..65534
- UDP Balancing:** disabled (dropdown)
- Cookie Persistence Mode:** disabled (dropdown)
- Persistent Binding:** disabled (dropdown)
- Host Name:** (text field)
- HTTP SLB 1st:** disabled (dropdown)
- Operator:** none (dropdown)
- HTTP SLB 2nd:** disabled (dropdown)
- No NAT:** enabled disabled
- Remap UDP Fragments:** enabled disabled
- RTSP Balancing:** disabled (dropdown)
- Delayed Binding:** enabled disabled
- FTP Parsing:** enabled disabled
- DNS Query Load Balancing:** enabled disabled
- BWM Contract:** 1024 0..1024 (Browse...)

At the bottom, there are three buttons: **Insert** (highlighted with a dashed box and a mouse cursor), **Close**, and **Help...**

9 In the Virtual Service field, type a number (1 - 8) for the IP service.

10 In the Real Group field, type the real server group number, or click Browse to choose the real server group from a selection list.

11 In the Real Server Port field, type 1 for the IP Service real server port.

12 Click Insert.

The Virtual Server Services, Insert Services dialog box closes and the service appears in the Virtual Server Services dialog box.

13 Click Apply > Close.

The service is configured and the Virtual Server Services dialog box closes.

14 In the Server Load Balance dialog box, click Close.

IP server load balancing is configured and the Server load Balance dialog box closes.

FTP server load balancing

RFC 959 defines two connections for FTP—one for control information and another for data. Each connection is unique. Unless the client requests a change, the server always uses TCP port 21 (a well-known port) for control information, and TCP port 20 for data. FTP uses TCP for transport. After the initial three-way handshake, a connection is established. When the client requests any data information from the server, it will issue a `PORT` command (such as `ls`, `dir`, `get`, `put`, `mget` and `mput`) via the control port.

There are two modes of FTP operation, active and passive.

Active FTP The FTP server initiates the data connection.

Passive FTP The FTP client initiates the data connection. Because the client also initiates the connection to the control channel, the passive FTP mode does not pose a problem with firewalls and is the most common mode of operation.

FTP network topology has the following restrictions:

- FTP uses both a control channel and a data channel; both channels need to be bound to the same real server.
- The FTP server may initiate FTP data sessions.
- Information exchanged on the control channel is used to determine the IP address and port for data connections between the FTP server and the FTP client.

Configuring FTP server load balancing

To configure FTP server load balancing:

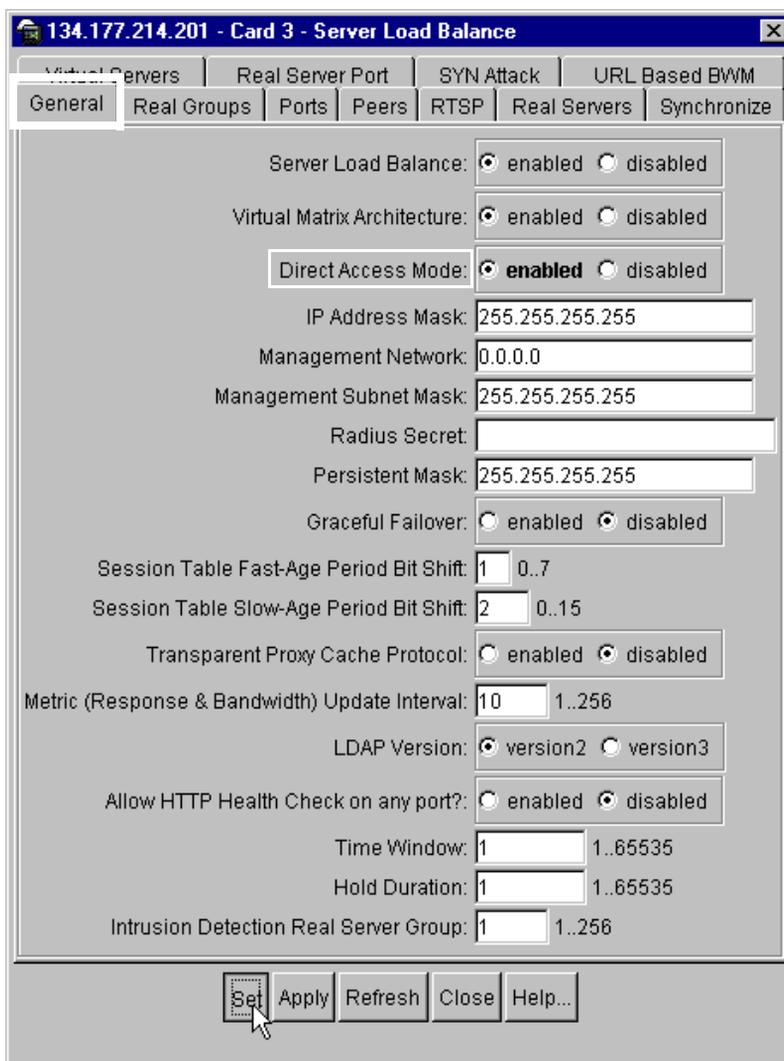
- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab ([Figure 111](#)).

For field definitions, see [Table 61 on page 224](#).

Figure 111 Configuring DAM for FTP server load balancing

3 In the Direct Access Mode field, click Enabled.

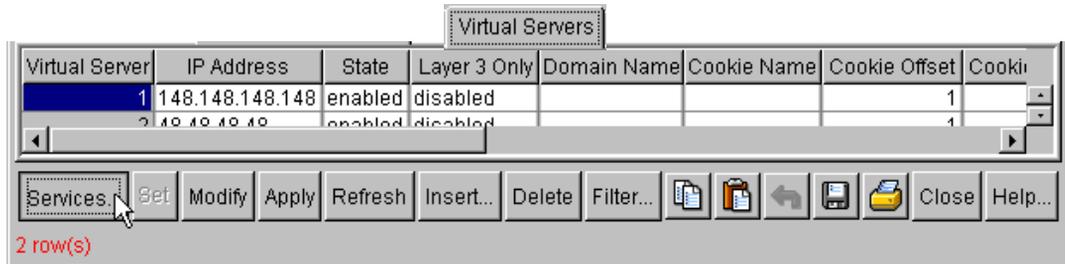


Note: If you do not use DAM, enable a proxy IP address. See [“Configuring a proxy IP address for a port” on page 247.](#)

- 4 Click Set > Apply.
- 5 Direct access mode is configured for the WSM.
- 6 Click the Virtual Servers tab.

The Virtual Servers tab (Figure 112) opens with the fields defined in Table 64 on page 233.

Figure 112 Server Load Balance—Virtual Servers tab



- 7 Select the virtual server to configure.
The virtual server number is highlighted.
- 8 Click Services.
The Virtual Server Services dialog box opens displaying the services for the selected server.
- 9 Click Insert.
- 10 The Insert Virtual Server Services dialog box (Figure 113) opens. For field definitions, see Table 65 on page 237.

Figure 113 Configuring FTP server load balancing

134.177.214.201 - Virtual Server 1 Services, Insert Services

Virtual Server: 1 1..256 Browse...

Virtual Service: 3 1..8

Virtual Port: 21 2..65534

Real Group: 2 1..256 Browse...

Real Server Port: 21 0..65534

UDP Balancing: disabled

Cookie Persistence Mode: disabled

Persistent Binding: disabled

Host Name:

HTTP SLB 1st: disabled

Operator: none

HTTP SLB 2nd: disabled

No NAT: enabled disabled

Remap UDP Fragments: enabled disabled

RTSP Balancing: disabled

Delayed Binding: enabled disabled

FTP Parsing: enabled disabled

DNS Query Load Balancing: enabled disabled

BWM Contract: 1024 0..1024 Browse...

Insert Close Help...

- 11** In the Virtual Service field, type a number (1 - 8) for the FTP service.
- 12** In the Virtual Port field, type 21 for the FTP service.
- 13** In the Real Group field, type the real server group number, or click Browse to choose the real server group from a selection list.
- 14** In the Real Server Port field, type 21 for the FTP service.
- 15** In the FTP Parsing field, click Enabled. FTP parsing must be enabled to translate the Layer 7 server address in the control field.

16 Click Insert.

FTP is added to the list of services in the Virtual Server Services dialog box.

17 In the Services dialog box, click Apply > Close.

The Services dialog box closes.

18 In the Server Load Balance dialog box, click Apply > Close.

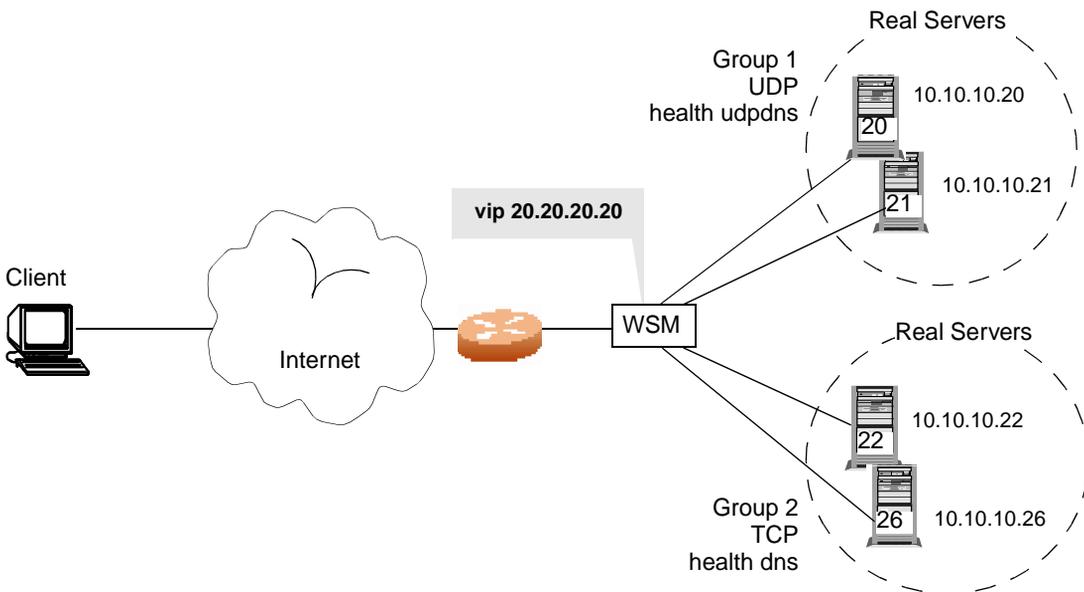
The FTP service is configured and the Server Load Balance dialog box closes.

Domain name server load balancing

You can choose the service for domain name server (DNS) load balancing based on UDP and TCP DNS queries. This enables the WSM to send TCP DNS queries to one group of real servers and UDP DNS queries to another group of real servers. The requests are then load-balanced among the real servers in that group.

[Figure 114](#) shows four real servers load balancing UDP and TCP queries between two groups.

Figure 114 Layer 4 DNS load balancing





Note: You can configure both UDP and TCP DNS queries for the same virtual server IP address.

Configuration tasks for DNS load balancing

To configure DNS load balancing, complete the following tasks:

- 1 Enable server load balancing. See [“Enabling or disabling server load balancing” on page 243](#).
- 2 Configure the real servers and their IP addresses. See [“Configuring each real server” on page 216](#).
- 3 Configure the real servers into groups for TCP and UDP services. See [“Configuring a real server group” on page 225](#).
- 4 Configure a virtual server to handle both UDP and TCP DNS server load balance requests. See [“Configuring a virtual server” on page 230](#).
- 5 Define and enable the server ports and client ports. See [“Configuring ports for server load balancing” on page 240](#).



Note: Some DNS servers initiate upstream requests and must be configured both as server and client.

- 6 Configure the services for DNS server load balancing. See:
 - [“Configuring UDP-based DNS server load balancing” on page 278](#)
 - [“Configuring TCP-based DNS server load balancing” on page 281](#)

Configuring UDP-based DNS server load balancing

To configure UDP-based DNS server load balancing:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab.

- 3 Click the Virtual Servers tab.

The Virtual Servers tab ([Figure 112 on page 275](#)) opens with the fields defined in [Table 64 on page 233](#).

- 4 On the Virtual Servers tab, select the DNS SLB server.

The virtual server number is highlighted.

- 5 Click Services.

The Virtual Server Services dialog box opens displaying the services for the selected virtual server.

- 6 Click Insert.

The Insert Virtual Server Services dialog box ([Figure 115](#)) opens with the fields defined in [Table 65 on page 237](#).

Figure 115 Configuring the UDP DNS service for server load balancing

134.177.214.201 - Virtual Server 1 Services, Insert Services

Virtual Server: 2 1..256 Browse...

Virtual Service: 3 1..8

Virtual Port: 54 2..65534

Real Group: 1 1..256 Browse...

Real Server Port: 54 0..65534

UDP Balancing: enabled

Cookie Persistence Mode: disabled

Persistent Binding: disabled

Host Name:

HTTP SLB 1st: disabled

Operator: none

HTTP SLB 2nd: disabled

No NAT: enabled disabled

Remap UDP Fragments: enabled disabled

RTSP Balancing: disabled

Delayed Binding: enabled disabled

FTP Parsing: enabled disabled

DNS Query Load Balancing: enabled disabled

BWM Contract: 1024 0..1024 Browse...

Insert Close Help...

- 7 In the Virtual Service field, type a number (1 - 8) for the UDP service.
- 8 In the Virtual Port field, type 54 for the UDP DNS service.
- 9 In the Real Group field, type the number of the UDP real server group, or click Browse to choose the group from a selection list. For more information, see [“Browsing a read only dialog box” on page 76](#).
- 10 In the Real Server Port field, type 54 for the UDP DNS service.
- 11 In the UDP Balancing field, click the down arrow and choose Enabled.

- 12** In the Delayed Binding field, click Disabled.
- 13** (Optional) In the DNS Query Load Balancing field, click Enabled.
- 14** Click Insert.

The UDP service is added to the Virtual Server Services for the DNS server, and the Insert Services dialog box closes.
- 15** From the Virtual Server Services dialog box, click Apply > Close.

The UDP service is configured and the Virtual Server Services dialog box closes.
- 16** In the Server Load Balance dialog box, click close.

UDP DNS load balancing is configured, and the Server Load Balance dialog box closes.

Configuring TCP-based DNS server load balancing

To configure TCP-based DNS server load balancing:

- 1** From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.
- 2** From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab.
- 3** In the Direct Access Mode field, click Enabled.
- 4** Click Set > Apply.

DAM is enabled for the WSM.
- 5** Click the Virtual Servers tab.

The Virtual Servers tab ([Figure 112 on page 275](#)) opens with the fields defined in [Table 64 on page 233](#).
- 6** Select the DNS SLB virtual server.

The server number is highlighted.
- 7** Click Services.

The Virtual Server Services dialog box opens displaying the services for the selected virtual server.

8 Click Insert.

The Insert Virtual Server Services dialog box (Figure 116) opens with the fields defined in Table 65 on page 237.

Figure 116 Configuring the TCP DNS service for server load balancing

134.177.214.201 - Virtual Server 1 Services, Insert Services

Virtual Server: 2 1..256 Browse...

Virtual Service: 2 1..8

Virtual Port: 54 2..65534

Real Group: 1 1..256 Browse...

Real Server Port: 54 0..65534

UDP Balancing: enabled

Cookie Persistence Mode: disabled

Persistent Binding: disabled

Host Name:

HTTP SLB 1st: disabled

Operator: none

HTTP SLB 2nd: disabled

No NAT: enabled disabled

Remap UDP Fragments: enabled disabled

RTSP Balancing: disabled

Delayed Binding: enabled disabled

FTP Parsing: enabled disabled

DNS Query Load Balancing: enabled disabled

BWM Contract: 1024 0..1024 Browse...

Insert Close Help...

9 In the Virtual Service field, type a number (1 - 8) for the TCP service.

10 In the Virtual Port field, type 54 for the TCP service.

11 In the Real Group field, type the number of the TCP real server group, or click Browse to choose the group from a selection dialog.

12 In the Real Server Port field, type 54 for the TCP service.

13 In the Delayed Binding field, click Enabled.

14 (Optional) In the DNS Query Load Balancing field, click Enabled.

15 Click Insert.

The TCP service is added to the Virtual Server Services for the DNS server, and the Insert Services dialog box closes.

16 Click Apply > Close.

The TCP service is configured and the Virtual Server Services dialog box closes.

17 Click close.

The Server Load Balance dialog box closes.

Real Time Streaming Protocol server load balancing

Real Time Streaming Protocol (RTSP) is an application-level protocol for controlling delivery of data with real-time properties as documented in RFC 2326. For example, the streams of data in a multimedia presentation—audio, video, and text—must be presented synchronously. A multimedia client like Real Player* or Quick Time Player* downloads these multiple streams of data from multimedia servers and presents them on the player screen. RTSP controls the flow of the multimedia streams. Each presentation uses one RTSP control connection and several other connections to carry the audio/video/text multimedia streams. In this document, the term RTSP server refers to any multimedia server configured for RTSP protocol for multimedia presentation.

The Web Switching Module RTSP implementation supports the following:

- Private addressing on the server side.
- Layer 7 URL-hashing metric, URL pattern matching, and all Layer 4 metrics for load balancing.
- Caching of entire presentations by switching all stream connections and control connections to the same cache server.
- RTSP-compliant applications (excluding Windows Media Player because it is not RTSP-compliant).

How RTSP server load balancing works

The objective of RTSP server load balancing is to intelligently switch an RTSP request, and other media streams associated with a presentation, to a suitable RTSP server using the configured load-balancing metric. The WSM supports one Layer 7 metric (URL hashing and URL pattern matching) and all Layer 4 load-balancing metrics. To configure Layer 7 URL pattern matching, see [“Configuring RTSP SLB using pattern matching” on page 471](#).

RTSP load balancing with the URL hash metric can be used to load balance cache servers that cache multimedia presentations. Since multimedia presentations consume a large amount of Internet bandwidth, and their correct presentation depends upon the real time delivery of the data over the Internet, several caching servers cache the multimedia data. As a result, the data is available quickly from the cache, when required. The Layer 7 metric of URL hashing directs all requests with the same URL to the same cache server, ensuring that no data is duplicated across the cache servers.

Typically, RTSP clients establish a control connection to RTSP servers over TCP port 554 and issue RTSP commands to it. On receiving the first command, the WSM chooses an RTSP server, using the configured Layer 4 metric. The WSM then routes the request by splicing a connection to the server.



Note: RTSP load balancing is not applicable if the streaming media (multimedia) servers use HTTP protocol to tunnel RTSP traffic. To ensure RTSP server load balancing works, make sure the streaming media server is configured for RTSP protocol.

RTSP clients issue sequences of commands to establish connections for each component stream of a presentation. The procedure varies, depending upon the RTSP client and server.

[Table 71](#) describes RTSP implementations for Real Server marketed by Real Networks Corporation*, and Quicktime Streaming Server marketed by Apple Inc.* Their RTSP stream setup sequences differ, and the WSM handles each differently.

Table 71 RTSP implementation comparison

| Implementation | Description |
|----------------------------|---|
| Real Server | Real Server's real media files have the following extensions: .rm, .ram, or .smil. Real Server supports both UDP and TCP transport protocols for the RTSP streams. The actual transport is negotiated during the initialization of the connection. If you specify TCP transport, then all streams of data will flow in the TCP control connection itself. If you choose UDP transport, the client and server negotiate a client UDP port ranging from 6970 to 7000 to set up each stream connection. |
| QuickTime Streaming Server | Apple Inc.'s QuickTime Streaming Server typically runs on Apple platforms. QuickTime files, which can be played over the Internet using RTSP, are specially formatted and are called <i>hinted Quick-Time files</i> . Normal QuickTime files cannot be used for streaming. QuickTime files have the extension .mov. QuickTime uses UDP protocol exclusively for transport and TCP for control connection. Each stream of a QuickTime presentation sends Real Time Protocol (RTP), and Real Time Control Protocol (RTCP) data using two UDP connections. Typically, a QuickTime presentation has two streams and therefore uses four UDP connections and one TCP control connection. QuickTime clients use a UDP port ranging from 6970 to 6999 for setting up stream connections. |

Pre-configuration tasks for RTSP SLB

Complete the following tasks before configuring RTSP load balancing:

- Disable port-based Bandwidth Management. See [“Setting port parameters” on page 85](#).
- Disable Web Cache Redirection. See [“Configuring Web Cache Redirection” on page 396](#).
- Disable proxy IP addressing. See [“Configuring a proxy IP address for a port” on page 247](#).

Configuring RTSP Layer 4 load balancing

To configure RTSP Layer 4 load balancing:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab ([Figure 117](#)).

For field definitions, see [“WAP—General tab fields” on page 299](#).

Figure 117 Configuring RTSP, SLB—General tab

134.177.214.201 - Card 3 - Server Load Balance

Virtual Servers | Real Server Port | SYN Attack | URL Based BWM

General | Real Groups | Ports | Peers | RTSP | Real Servers | Synchronize

Server Load Balance: enabled disabled

Virtual Matrix Architecture: enabled disabled

Direct Access Mode: enabled disabled

IP Address Mask: 255.255.255.255

Management Network: 0.0.0.0

Management Subnet Mask: 255.255.255.255

Radius Secret:

Persistent Mask: 255.255.255.255

Graceful Failover: enabled disabled

Session Table Fast-Age Period Bit Shift: 1 0..7

Session Table Slow-Age Period Bit Shift: 2 0..15

Transparent Proxy Cache Protocol: enabled disabled

Metric (Response & Bandwidth) Update Interval: 10 1..256

LDAP Version: version2 version3

Allow HTTP Health Check on any port?: enabled disabled

Time Window: 1 1..65535

Hold Duration: 1 1..65535

Intrusion Detection Real Server Group: 1 1..256

Set Apply Refresh Close Help...

- 3 In the Virtual Matrix Architecture field, click Enabled.
- 4 In the Direct Access Mode field, click Enabled.
- 5 Click Set > Apply.
- DAM and VMA are enabled for the WSM.
- 6 Click the Virtual Server tab.

The Virtual Server tab opens.

- 7 Select the virtual server to configure.

The virtual server number is highlighted.

- 8 Click Services.

The Virtual Server Services dialog box opens displaying the services for the selected server.

- 9 Click Insert.

The Insert Virtual Server Services dialog box ([Figure 118](#)) opens. For information about this dialog box, see [Table 65 on page 237](#).

Figure 118 Configuring L4 RTSP service for server load balancing

134.177.214.201 - Virtual Server 2 Services, Insert Services

Virtual Server: 2 1..256 Browse...

Virtual Service: 1 1..8

Virtual Port: 554 2..65534

Real Group: 1 1..256 Browse...

Real Server Port: 554 0..65534

UDP Balancing: disabled

Cookie Persistence Mode: disabled

Persistent Binding: disabled

Host Name:

HTTP SLB 1st: disabled

Operator: none

HTTP SLB 2nd: disabled

No NAT: enabled disabled

Remap UDP Fragments: enabled disabled

RTSP Balancing: disabled

Delayed Binding: enabled disabled

FTP Parsing: enabled disabled

DNS Query Load Balancing: enabled disabled

BWM Contract: 1024 0..1024 Browse...

Insert Close Help...

- 10** In the Virtual Service field, type a number (1 - 8) for the RTSP service.
- 11** In the Virtual Port field, type 554 for RTSP.
- 12** In the Real Server Group field, type the number of the real group.
- 13** In the Real Server Port field, type 554 for RTSP.
- 14** Click Insert.

Layer 4 RTSP is added to the list of services and the Insert Services dialog box closes.

- 15** In the Virtual Server Services dialog box, click Apply > Close.

RTSP is configured and the Virtual Server Services dialog box closes.

- 16** From the Virtual Servers tab, click Close.

The Server Load Balance dialog box closes.

Configuring RTSP Layer 7 load balancing

This procedure describes how to configure Layer 7 RTSP load balancing using the hash metric. To configure Layer 7 RTSP load balancing using pattern matching, see [“Configuring RTSP SLB using pattern matching” on page 471](#).

To configure RTSP Layer 7 load balancing using the hash SLB metric:

- 1** From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2** From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab ([Figure 117 on page 287](#)). For field definitions, see [Table 61 on page 224](#).

- 3** In the Virtual Matrix Architecture field, click Enabled.

- 4** In the Direct Access Mode field, click Enabled.

- 5** Click Set > Apply.

DAM and VMA are enabled for the WSM.

- 6** Click the Virtual Servers tab.

The Virtual Servers tab opens.

- 7** Select the virtual server.

The virtual server number is highlighted.

- 8** Click Services.

The Virtual Server Services dialog box opens displaying the services for the selected virtual server.

- 9** Click Insert.

The Insert Virtual Server Services dialog box (Figure 119) opens. For field definitions, see Table 65 on page 237.

Figure 119 Configuring L7 RTSP service for server load balancing

134.177.214.201 - Virtual Server 2 Services, Insert Services

Virtual Server: 2 1..256 Browse...

Virtual Service: 3 1..8

Virtual Port: 0 2..65534

Real Group: 1 1..256 Browse...

Real Server Port: 0 0..65534

UDP Balancing: disabled

Cookie Persistence Mode: disabled

Persistent Binding: disabled

Host Name:

HTTP SLB 1st: disabled

Operator: none

HTTP SLB 2nd: disabled

No NAT: enabled disabled

Remap UDP Fragments: enabled disabled

RTSP Balancing: hash

Delayed Binding: enabled disabled

FTP Parsing: enabled disabled

DNS Query Load Balancing: enabled disabled

BWM Contract: 1024 0..1024 Browse...

Insert Close Help...

- 10 In the Virtual Service field, type a number (1 - 8) for the RTSP service.
- 11 In the Real Server Group field, type the number of the real group.
- 12 In the RTSP Balancing field, click the down arrow and choose Hash.

This enables URL hashing using the entire URL; however, any extension of the type (.xxx) occurring at the end of the URL is omitted from hash computation.

13 Click Insert.

Layer 7 RTSP is added to the list of services and the Insert Services dialog box closes.

14 Click Apply > Close.

Layer 7 RTSP is configured with the Hash metric, and the Virtual Server Services dialog box closes.

15 In the SLB dialog box, click close.

The Server Load Balance dialog box closes.

Wireless application server load balancing

Wireless Application protocol (WAP) carries Internet traffic to mobile devices and allows Web services to be delivered to mobile phones and handsets. Traditional Web servers, known as WAP gateways, perform the translation from HTTP/HTML to WAP/WML (Wireless Markup Language) on the land-based part of the network. WAP supports most wireless networks and is supported by all operating systems. To load balance WAP traffic among available parallel servers, the WSM must provide persistency so that the clients can always go to the same WAP gateway to perform WAP operation. For more information see [“Maintaining session persistence” on page 433](#).

The WSM decides to which real gateway the request should go. WAP SLB is based on RADIUS static session entry or RADIUS snooping.

How WAP SLB works using RADIUS snooping

RADIUS snooping examines RADIUS accounting packets for client information. to update static session entries. The WSM session table uses these to maintain persistency for load balancing. If a static session entry is added using RADIUS snooping, it must also be deleted using RADIUS snooping, instead of aging out. The WSM load balances both the RADIUS and WAP gateway traffic using the same virtual server IP address.

The following is a step-by-step description of how RADIUS snooping works:

- 1** The user is authenticated on dialing.
- 2** The RAS establishes a session with the client and sends a RADIUS Accounting Start message to the RADIUS server which includes the client IP address. Before a session entry is recorded on the WSM, WAP packets for a user can go to any of the real WAP gateways.
- 3** The WSM snoops on the RADIUS accounting packet and adds a session entry only if it finds required information (type of RADIUS Accounting message, client IP address, caller ID, and user's name) in the packet.

If any information is missing, the WSM will not add the session entry.

If a session entry for a client cannot be added because of resource constraints, the subsequent WAP packets for that client will not be load balanced correctly; and the client must drop the connection and reconnect to his wireless service.

- 4** The WSM load balances the WAP traffic to a specific WAP gateway.

The WSM maintains persistency during the session, however, session persistence cannot be maintained if the number of healthy real WAP gateways changes during the session, for example, if a new WAP server comes into service or existing WAP servers are down.

Persistence cannot be maintained if the user moves from one ISP to another, or if the base of the user's session changes (that is, from `CALLING_STATION_ID` to `USER_NAME`, or vice versa). For example, if a user moves out of a roaming area, it is possible that his/her `CALLING_STATION_ID` is not available in the RADIUS Accounting packets. If not, the WSM uses `USER_NAME` to choose a WAP server instead of `CALLING_STATION_ID`, and persistence cannot be maintained.

- 5** When the client terminates the session, the RAS sends an Accounting Stop message to the RADIUS server, and the session entry is deleted from the WSM.

Configuring WAP SLB using RADIUS snooping

Use the following guidelines when configuring WAP RADIUS snooping:

- Use the same virtual server IP address when load balancing both RADIUS accounting traffic and WAP traffic.
- Use the same UDP port for RADIUS accounting services for all RADIUS servers.

The following configurations are required for WAP SLB using RADIUS snooping:

- [“Configuring WAP server load balancing,”](#) next
- [“Configuring the WAP RADIUS snooping filter”](#) on page 299

Configuring WAP server load balancing

To configure WAP server load balancing:

- 1** From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2** From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab ([Figure 120](#)) with the fields defined in [Table 61 on page 224](#).

Figure 120 Configuring WAP server load balancing, SLB—General tab

The screenshot shows a configuration window titled "134.177.214.201 - Card 3 - Server Load Balance". The "General" tab is selected. The settings are as follows:

- Server Load Balance: enabled disabled
- Virtual Matrix Architecture: enabled disabled
- Direct Access Mode: enabled disabled
- IP Address Mask: 255.255.255.255
- Management Network: 0.0.0.0
- Management Subnet Mask: 255.255.255.255
- Radius Secret: (empty)
- Persistent Mask: 255.255.255.255
- Graceful Failover: enabled disabled
- Session Table Fast-Age Period Bit Shift: 1 0..7
- Session Table Slow-Age Period Bit Shift: 2 0..15
- Transparent Proxy Cache Protocol: enabled disabled
- Metric (Response & Bandwidth) Update Interval: 10 1..256
- LDAP Version: version2 version3
- Allow HTTP Health Check on any port?: enabled disabled
- Time Window: 1 1..65535
- Hold Duration: 1 1..65535
- Intrusion Detection Real Server Group: 1 1..256

At the bottom, there are buttons for "Set", "Apply", "Refresh", "Close", and "Help...". The "Set" button is highlighted with a mouse cursor.

- 3 In the Virtual Matrix Architecture field, click Enabled.
- 4 In the Direct Access Mode field, click Disabled.
- 5 In the Transparent Proxy Cache Protocol field, click Enabled.
- 6 Click Set > Apply.

Dam, VMA, and TPCP are configured for WAP.

- 7** Click the Virtual Server tab.

The Virtual Server tab opens.

- 8** Select the virtual server to configure.

The row is highlighted.

- 9** Click Services.

The Virtual Server Services dialog box opens displaying the services for the selected server. For information about this dialog box, see [Table 65 on page 237](#).

- 10** Click Insert.

The Insert Virtual Server Services dialog box ([Figure 121](#)) opens with the fields defined in [Table 65 on page 237](#).

Figure 121 Configuring the WAP service for server load balancing

134.177.214.201 - Virtual Server 1 Services, Insert Services

Virtual Server: 2 1..256 Browse...

Virtual Service: 1 1..8

Virtual Port: 9200 2..65534

Real Group: 1 1..256 Browse...

Real Server Port: 9200 0..65534

UDP Balancing: enabled

Cookie Persistence Mode: disabled

Persistent Binding: disabled

Host Name:

HTTP SLB 1st: disabled

Operator: none

HTTP SLB 2nd: disabled

No NAT: enabled disabled

Remap UDP Fragments: enabled disabled

RTSP Balancing: disabled

Delayed Binding: enabled disabled

FTP Parsing: enabled disabled

DNS Query Load Balancing: enabled disabled

BWM Contract: 1024 0..1024 Browse...

Insert Close Help...

- 11 In the Virtual Service field, type a number (1 - 8) for the WAP service.
- 12 In the Virtual Port field, type 9200 for the WAP service.
- 13 In the Real Server Port field, type 9200 for the WAP service.
- 14 In the Persistent Binding field, click the down arrow and choose Disabled.
- 15 In the UDP Balancing field, click Enabled.
- 16 Click Insert.

The Insert Services dialog box closes and the virtual service appears in the Virtual Server Services dialog box.

- 17** Click Apply > Close.

The WAP service is configured and the Virtual Server Services dialog box closes.

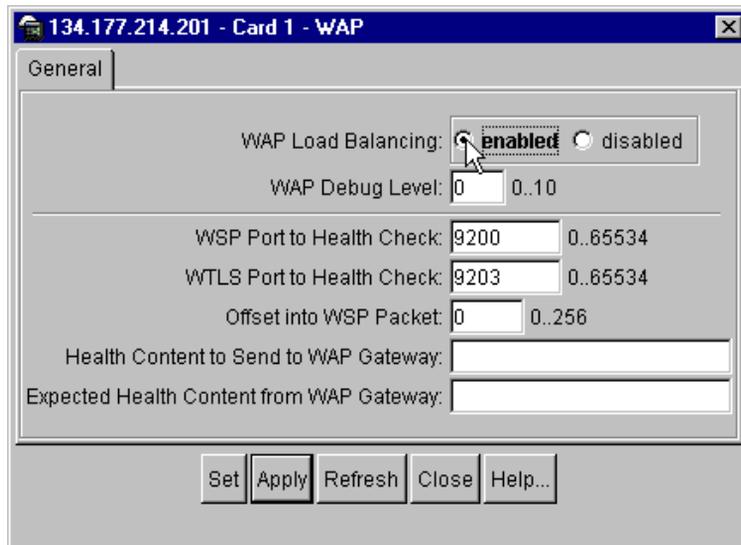
- 18** From the Virtual Server tab, click Close.

The Server Load Balance dialog box closes.

- 19** From the Edit menu, choose WSM Card > L4 Switching > WAP.

The [WAP dialog box \(Figure 122\)](#) opens with the fields defined in [Table 72 on page 299](#).

Figure 122 WAP—General tab



- 20** In the WAP Load Balancing field, click Enabled.

- 21** Click Set > Apply > Close.

WAP server load balancing is configured.

- 22** Continue configuring WAP SLB using RADIUS snooping by [“Configuring the WAP RADIUS snooping filter” on page 299](#).

Table 72 describes the fields on the WAP—General tab.

Table 72 WAP—General tab fields

| Field | Description |
|--|---|
| WAP Load Balancing | Enables or disables wireless application protocol load balancing. Enables/disables WAP TPCP external notification for adding and deleting WAP sessions. |
| WAP Debug Level | Sets the wireless application protocol debug level (0 to 10) for tracing WAP related messages. Level 10 is the most verbose. |
| WSP Port to Health Check | Sets the port number (0 to 65,534) on which WSP health checks will be performed. The default port number = 9200. |
| WTLS Port to Health Check | Sets the port number (0 to 65,534) on which WTLS health checks will be performed. The default port number = 9203. |
| Offset into WSP Packet | Sets the offset (0 - 256) into the received WSP packets. An offset value of 0 (default) sets the WSM to start comparisons at the beginning of the content of the received packet. |
| Health Content to Send to WAP Gateway | Defines content examined during wireless application protocol health checks. The maximum string length is 256 characters. Leave out spaces to allow for larger string. |
| Expected Health Content from WAP Gateway | The selected value for the wireless application protocol gateway response to match. The maximum string length is 256 characters. Leave out spaces to allow for larger string. |

Configuring the WAP RADIUS snooping filter

Before configuring the WAP RADIUS snooping filter, complete the procedure in the section, “[Configuring WAP server load balancing](#)” on page 294.

To configure the filter rule to examine a RADIUS accounting packet:

- 1 From the Device View, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, select WSM Card > L4 Switching > Filters.
The Filters dialog box opens.
- 3 Click Insert.

The Insert Filters dialog box ([Figure 123](#) and [Figure 124](#)) opens. See [Table 57 on page 196](#) for field descriptions.

Figure 123 Configuring WAP RADIUS snooping filter—Part 1

Index: 2 1..2048

Name: WAP RADIUS Snooping

Filter: enabled disabled

Action: redirect

Invert: invert-on invert-off

Logging: enabled disabled

Caching: enabled disabled

Client Proxy: enabled disabled

VLAN: any 0..4094 (any: 0)

Source Address Filter Type: ip mac

Destination Address Filter Type: ip mac

Source IP Address: any

Source IP Mask: 0.0.0.0

Destination IP Address: 10.10.10.100

Destination IP Mask: 255.255.255.255

Source MAC Address: 00:00:00:00:00:00

Destination MAC Address: 00:00:00:00:00:00

Protocol: udp(17) (number)

Low Source Port: http(1813) 0..65534 (none: 0)

High Source Port: http(1813) 0..65534 (none: 0)

Low Destination Port: 0 0..65534 (none: 0)

High Destination Port: 0 0..65534 (none: 0)

Redirection Port: http(1813) 0..65534

Redirection Group: 1 1..256

Figure 124 Configuring WAP RADIUS snooping filter—Part 2

URL Redirection: enabled disabled
 Network Address Translation: destination-address source-address
 NAT Active FTP: enabled disabled
 NAT Session Timeout: 4..30 (even values only)
 TCP ACK or RST Matching: enabled disabled
 TCP URG: enable disable
 TCP ACK: enable disable
 TCP PSH: enable disable
 TCP RST: enable disable
 TCP SYN: enable disable
 TCP FIN: enable disable
 ICMP Type: 0..255 (any: 255)
 IP Option: enable disable
 IP TOS: 0..255
 IP TOS Mask: 0..255
 New IP TOS: 0..255
 TCP Connection Rate Limiting: enabled disabled
 Maximum Connection for TCP Rate Limiting: 0..255
 Firewall Redirect Hash: enabled disabled
 WAN Link Load Balancing: enabled disabled
 Intrusion Detection Hash: sip dip both
 Hash:
 BWM Contract: 1..1024
 WAP Radius Snooping: enabled disabled

- 4 In the Index field, type a new filter number (1 to 2048).
- 5 In the Destination IP Address field, type the IP address to filter.

- 6 In the Destination IP Mask field, type an IP mask to use with the IP address to select the traffic to filter.
- 7 In the Protocol field, click Browse, choose UDP from the selection list, and click Modify.
UDP is inserted into the Protocol field.
- 8 In the Low Destination Port field, type 1813.
- 9 In the High Destination Port field, type 1813.
- 10 In the Action field, click the down arrow and choose Redirect.
The Redirection Port field is active.
- 11 In the Redirection Port field, type 1813.
- 12 In the Redirection Group field, type the number of the real server group to which you want to redirect WAP traffic.
- 13 In the WAP RADIUS Snooping field, click Enabled.
- 14 In the Filter field, click Enabled to enable the new filter rule.
- 15 Click Insert.
The filter definition is inserted in the Filters tab and the Insert Filters dialog box closes.
- 16 From the Filters tab, click Set > Apply > Close.
The new filter is configured and the Filters dialog box closes.

Intrusion Detection server load balancing

An Intrusion Detection System gathers and analyzes computer or network information to identify security breaches from both intrusions (attacks from outside the organization) and misuse (attacks from inside the organization).

Intrusion Detection includes:

- Monitoring and analyzing both user and system activities
- Analyzing system configurations and vulnerabilities
- Assessing system and file integrity
- Recognizing patterns typical of attacks

- Analyzing abnormal activity patterns
- Tracking user policy violations

Intrusion Detection devices inspect every packet for signs of attack before it enters a network. The attacks are recorded and logged in an attempt to guard against future attacks and to record the information about the intruders. IDS Server Load Balancing helps scale Intrusion Detection systems since it is not possible for an individual server to scale information being processed at gigabit speeds.

How Intrusion Detection SLB works

The WSM forwards IP packets to an Intrusion Detection server at the end of the filtering process. If filtering is not enabled, the WSM can forward IP packets to an IDS at the end of client processing. You must enable IDS SLB on the port and allocate a real server group for IDS Server Load Balancing. The IDS SLB-enabled WSM copies all incoming packets to this group of Intrusion Detection servers. For each session entry created on the WSM, an IDS server is selected based on the IDS server load-balancing metric.

The IDS server receives copies of all processed frames forwarded to destination devices. Session entries are maintained so that all the frames of a given session are forwarded to the same IDS server. Each IDS server must be connected directly to a different WSM port or VLAN because no field in the packet header can be substituted. Substituting a field would corrupt the packet that must also be forwarded to its real destination.

The WSM supports the following metrics for IDS load balancing:

- minmisses
- roundrobin (You must disable delayed binding if you select this metric.)
- hash for client processing on the port—hash on the source IP address.
- hash for filter processing on the port—hash on the source IP address, destination IP address, or both.



Note: Leastconns, bandwidth, and response load balancing metrics do not apply to IDS SLB.

Configuring Intrusion Detection server load balancing

Intrusion Detection server load balancing requires the following steps:

- [“Configuring real servers for IDS SLB,”](#) next
- [“Configuring real server groups for IDS SLB”](#) on page 306
- [“Configuring ports for IDS SLB”](#) on page 309
- [“Enabling IDS SLB”](#) on page 311

Configuring real servers for IDS SLB

Since each packet is replicated and forwarded to its destination, the packet cannot be modified without risking corruption. Therefore, all IDS servers must be connected directly to a WSM port. This allows for up to 4 IDS servers per WSM, since the other 4 ports are reserved for the backplane connection.



Note: To configure IDS link health check in stealth mode (mode in which no IP address is used and the IDS just monitors packets), each IDS server must be directly connected to one of the front-facing WSM ports.

To configure IDS server load balancing:



Note: Before configuring IDS SLB, both RTSP SLB and WAP RADIUS Snooping must be disabled.

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3 Click the Real Servers tab.
The Real Servers tab opens.
- 4 Click Insert

The Insert Real Servers (Figure 125) dialog box opens with the fields defined in Table 60 on page 219.

Figure 125 Configuring real servers for IDS SLB

134.177.214.201 - Card 1 - Server Load Balance, Insert Real Servers

Real Server: 6 1..1023

IP Address: 10.10.40.20

State: enabled disabled

Name: IDS

Weight: 1 1..48

Max. Connections: 200000 0..200000

Timeout: 10 2..30 (even values only)

Backup: 0 0..1023 (none: 0)

Health Check Interval: 2 0..60 (none: 0)

Failure Retries: 4 1..63

Success Retries: 8 1..63

Type: local-server remote-server

No Cookie Operation: enabled disabled

Exclude String Match: enabled disabled

Substitute Source MAC: enabled disabled

Client Proxy: enabled disabled

Insert Close Help...

- 5 In the Real Server field, type the real server number (1 - 1023) for your IDS server.



Note: Each IDS server must be connected directly to a different WSM port or VLAN. If the IDS group will be configured for link health check, match the IDS server number to the physical port number (1 to 4) to which it is connected. For ICMP health check, the IDS server number can be between 1 and 31.

- 6 In the IP address field, type an IP address for the IDS real server.



Note: Real servers are configured by providing the IP address of the actual server. If your IDS servers are implemented without an IP address (stealth mode), configure a dummy address for the real server.

- 7 In the State field, click Enabled.

- 8 Click Insert.

The server is inserted into the Real Servers tab and the Insert Real Servers dialog box closes.

- 9 To define other IDS real servers, repeat steps 4 - 8.

- 10 Click Apply > Close.

The real servers are configured for SLB.

- 11 Continue the IDS SLB configuration by [“Configuring real server groups for IDS SLB,”](#) next.

Configuring real server groups for IDS SLB

Before configuring real server groups for IDS SLB, complete the procedure in the section, [“Configuring real servers for IDS SLB”](#) on page 304.

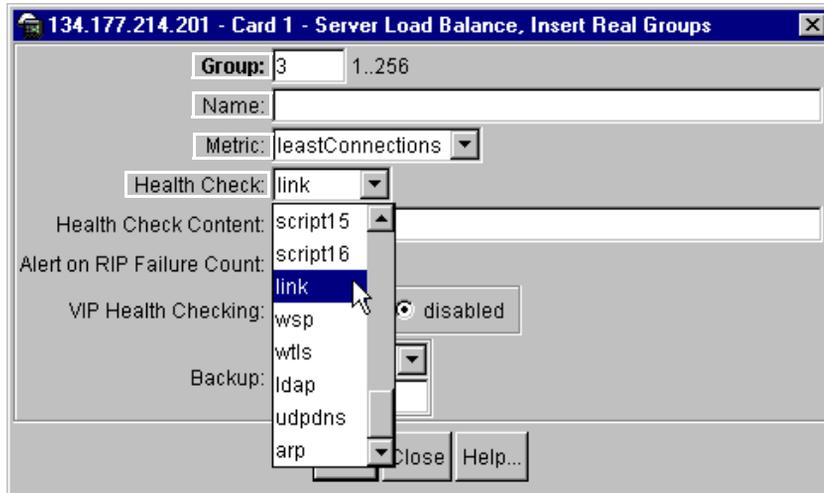
To configure real server groups for IDS SLB:

- 1 From the Server Load Balance dialog box, click the Real Groups tab.

The Real Groups tab opens.

- 2 Click Insert.

The Insert Real Groups dialog box ([Figure 126](#)) opens with the fields defined in [Table 62](#) on page 228.

Figure 126 Configuring real server groups for IDS SLB

- 3 In the Group field, type the group number (2 - 256) for the IDS real server group.



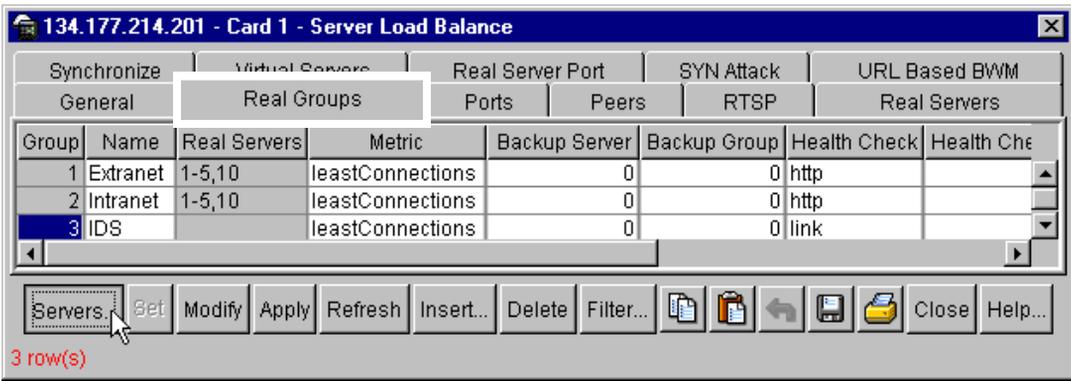
Note: Group 1 cannot be used for IDS SLB and only one IDS load balancing group is allowed per WSM.

- 4 In the Name field, type a name of up to 31 characters for the IDS real server group.
- 5 In the Metric field, choose a load balancing metric (roundrobin, minmisses, or hash) for the IDS real server group.
- 6 In the Health Check Content field, click the down arrow and do one of the following:
 - Choose Link if your IDS servers are configured without an IP address.
 - Choose ICMP if your IDS servers are configured with an IP address.
- 7 Click Insert.

The real server group is inserted into the Real Groups tab and the Insert Real Groups dialog box closes.
- 8 From the Real Groups tab ([Figure 127](#)), click the IDS real server group.

The IDS real server group number is highlighted.

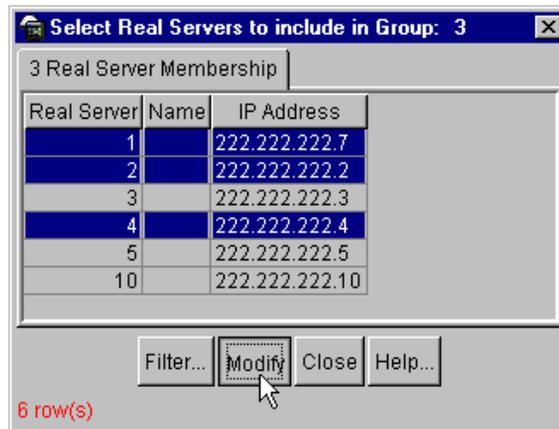
Figure 127 SLB Real Groups tab



9 Click Servers.

10 The Select Real Servers to include in Group dialog box (Figure 128) opens with the fields defined in Table 63 on page 230.

Figure 128 Adding real servers to IDS SLB real server group



11 Click the IDS real server(s) you want to add to the real server group.



Caution: When using this dialog box, you must use CTRL + Click to preserve existing server membership. If you just choose the new server and click Modify, only the new server is added, and existing server members are removed.

12 The row(s) is(are) highlighted.

13 Click Modify.

14 The IDS real servers are added to the Real Groups tab and the Select Real Servers to include in Group dialog box closes.

15 From the Real Groups tab, click Apply.

The real server group is configured for IDS SLB.

16 Continue the IDS configuration by [“Configuring ports for IDS SLB,”](#) next.

Configuring ports for IDS SLB

Before configuring ports for IDS SLB, complete the procedures in the following sections:

- [“Configuring real servers for IDS SLB” on page 304](#)
- [“Configuring real server groups for IDS SLB” on page 306](#)

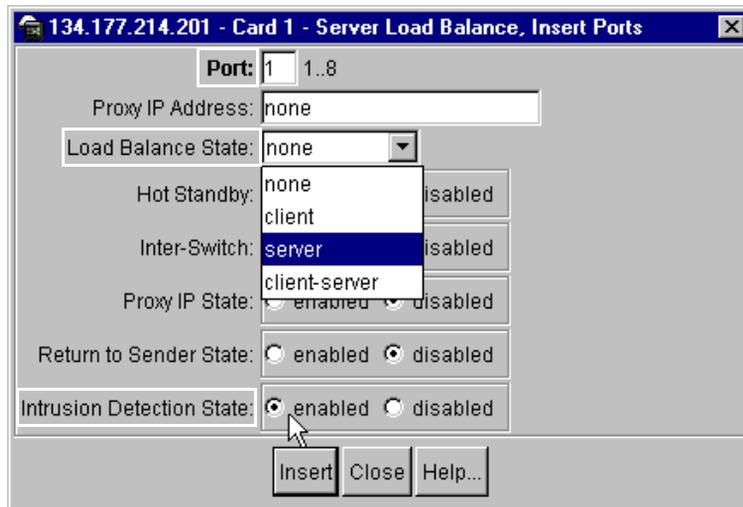
To configure the backplane fabric module ports (1-4) for IDS SLB:

1 From the Server Load Balance dialog box, click the Ports tab.

The Ports tab opens.

2 Click Insert.

The Insert Ports dialog box ([Figure 129](#)) opens with the fields defined in [Table 66 on page 242](#).

Figure 129 Configuring ports for IDS SLB

- 3 In the Port field, type the port number (1-4).



Note: Each IDS server must be connected directly to a different WSM port or VLAN. If the IDS group will be configured for link health check, match the IDS server number to the physical port number (1 - 4) to which it is connected. For ICMP health check, the IDS server number can be between 1 and 31.



Note: An IDS server must be connected to one of the WSM front-facing ports (1-4) to enable link health check in stealth mode (mode in which no IP address is used and the IDS just monitors packets).

- 4 In the Load Balance State field, click the down arrow and choose Server.
- 5 In the Intrusion Detection State field, click Enabled.
- 6 Click Insert.

The port is inserted into the Ports tab and the Insert Ports dialog box closes.

- 7 Repeat steps 2 - 5 for each port.
- 8 Click Apply.

The port is configured for IDS SLB and the Server Load Balance dialog box closes.

- 9 Continue the IDS configuration by [“Enabling IDS SLB,”](#) next.

Enabling IDS SLB

Before configuring ports for IDS SLB, complete the procedures in the following sections:

- [“Configuring real servers for IDS SLB”](#) on page 304
- [“Configuring real server groups for IDS SLB”](#) on page 306
- [“Configuring ports for IDS SLB”](#) on page 309

To enable IDS server load balancing:

- 1 From the Server Load Balance dialog box, click the General tab.

The General tab opens ([Figure 130](#)) with the fields defined in [Table 61](#) on [page 224](#).

Figure 130 Enabling Intrusion Detection Server Load Balancing

The screenshot shows a configuration window titled "134.177.214.201 - Card 3 - Server Load Balance". The window has several tabs: "General", "Real Groups", "Ports", "Peers", "RTSP", "Real Servers", and "Synchronize". The "General" tab is active. The configuration includes the following fields and controls:

- Server Load Balance: enabled disabled
- Virtual Matrix Architecture: enabled disabled
- Direct Access Mode: enabled disabled
- IP Address Mask: 255.255.255.255
- Management Network: 0.0.0.0
- Management Subnet Mask: 255.255.255.255
- Radius Secret: [empty text box]
- Persistent Mask: 255.255.255.255
- Graceful Failover: enabled disabled
- Session Table Fast-Age Period Bit Shift: 1 0..7
- Session Table Slow-Age Period Bit Shift: 2 0..15
- Transparent Proxy Cache Protocol: enabled disabled
- Metric (Response & Bandwidth) Update Interval: 10 1..256
- LDAP Version: version2 version3
- Allow HTTP Health Check on any port?: enabled disabled
- Time Window: 1 1..65535
- Hold Duration: 1 1..65535
- Intrusion Detection Real Server Group: 3 1..256

At the bottom of the window, there are five buttons: "Set", "Apply", "Refresh", "Close", and "Help...". The "Set" button is highlighted with a mouse cursor.

- 2 In the Intrusion Detection Real Server Group field, enter the number of the real server group configured for Intrusion Detection.
- 3 Click Set > Apply > Close.
IDS SLB is enabled for the WSM.

WAN link load balancing

Wide Area Networking (WAN) is a telecommunications network system spread across a broad geographic area. A WAN may be privately owned or rented, but usually includes public (shared user) networks, such as the telephone system. WANs can also be connected through leased lines and satellites. WANs are typically composed of powerful routers and switches that link business enterprises, universities, remote offices, and so on, around the world.

To handle the high volume of data on the Internet, some corporations use more than one Internet Service provider (ISP) to increase Internet reliability. Such enterprises with more than one ISP are referred called multihomed. In addition to reliability, a multihomed network architecture enables enterprises to distribute load among multiple connections.

WAN link load balancing can be used to steer requests initiated within the user's network and his/her responses over the appropriate link at a given time.

How WAN link load balancing works

Redirection filters redirect traffic initiated from within the user's network to a group of devices at the other end of the WAN link (routers, for example). These filters determine the best link at the time of the request. To ensure that the responses traverse the same link, the source IP address of the request is translated to one of the addresses that the selected ISP owns. The design of WAN link load balancing is identical to standard redirection, except that it substitutes the source IP address of each frame with the proxy IP address of the port to which the WAN link is connected.

Configuring WAN link load balancing

WAN link load balancing requires the following tasks:

- [“Pre-configuration tasks for WAN link load balancing,”](#) next
- [“Configuring real servers for WAN link load balancing”](#) on page 314
- [“Configuring real server groups for WAN link load balancing”](#) on page 316
- [“Configuring filters for WAN link load balancing”](#) on page 318

Pre-configuration tasks for WAN link load balancing

Before configuring WAN Link load balancing:

- Disable NAT Web Cache Redirection. WAN Link load balancing and NAT Web Cache Redirection cannot be configured on the same WSM. See [“Application Redirection” on page 395](#).
- Make sure your ports are not configured into trunk groups. See [“Port management” on page 81](#).
- Make sure your WAN link load balancing is not configured with two or more WAN links connected through the same WSM port. This feature uses the proxy IP address of the destination port when translating the source IP address of the requests.

Configuring real servers for WAN link load balancing

Before configuring real servers for WAN link load balancing, complete the [“Pre-configuration tasks for WAN link load balancing” on page 314](#).

To configure real servers for WAN link load balancing:

- 1** From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2** From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3** Click the Real Servers tab.
The Real Servers tab opens.
- 4** Click Insert
The Insert Real Servers ([Figure 131](#)) dialog box opens with the fields defined in [Table 60 on page 219](#).

Figure 131 Configuring WAN link real server for server load balancing

134.177.214.201 - Card 1 - Server Load Balance, Insert Real Servers

Real Server: 6 1..1023

IP Address: 10.10.40.60

State: enabled disabled

Name: WAN Link SLB

Weight: 1 1..48

Max. Connections: 200000 0..200000

Timeout: 10 2..30 (even values only)

Backup: 0 0..1023 (none: 0)

Health Check Interval: 2 0..60 (none: 0)

Failure Retries: 4 1..63

Success Retries: 8 1..63

Type: local-server remote-server

No Cookie Operation: enabled disabled

Exclude String Match: enabled disabled

Substitute Source MAC: enabled disabled

Client Proxy: enabled disabled

Insert Close Help...

- 5 In the Real Server field, type the real server number (1 - 1023) for the WAN link real server.
- 6 In the IP address field, type an IP address for the WAN link real server.
- 7 In the State field, click Enabled.
- 8 In the Client Proxy field, click Disabled.
- 9 Click Insert.

The server is inserted into the Real Servers tab and the Insert Real Servers dialog box closes.

- 10 From the Real Servers tab, click Apply.

- 11 Continue configuring WAN link load balancing by “[Configuring real server groups for WAN link load balancing](#),” next.

Configuring real server groups for WAN link load balancing

Before configuring a real group for WAN link load balancing, complete the procedure for “[Configuring real servers for WAN link load balancing](#)” on page 314.

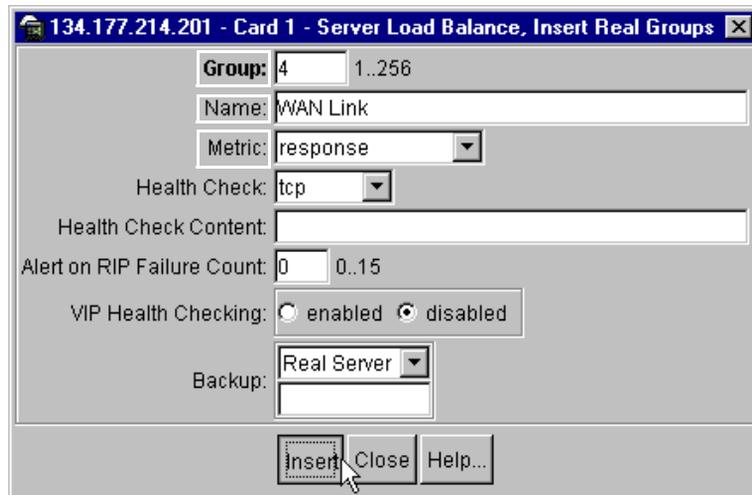
To configure real server groups for WAN link load balancing:

- 1 From the Server Load Balance dialog box, click the Real Groups tab.
The Real Groups tab opens.

- 2 Click Insert.

The Insert Real Groups dialog box ([Figure 132](#)) opens with the fields defined in [Table 62](#) on page 228.

Figure 132 Configuring WAN link real server group for SLB



- 3 In the Group field, type the group number (1 - 256) for the WAN link real server group.
- 4 In the Name field, type a name of up to 31 characters for the WAN link group.
- 5 In the Metric field, click the down arrow, and choose Response.

6 Click Insert.

The real server group is inserted into the Real Groups tab and the Insert Real Groups dialog box closes.

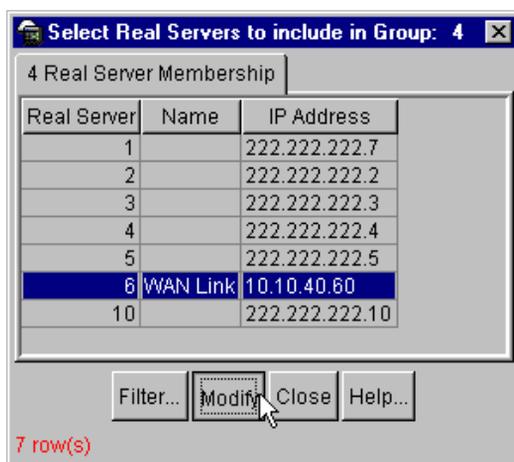
7 To add the WAN link real server to this group, in the Real Groups tab, select the group.

The row is highlighted.

8 Click Servers.

The Select Real Servers to include in Group dialog box (Figure 133) opens with the fields defined in Table 63 on page 230.

Figure 133 Adding real servers to WAN link server group

**9** Click the WAN link real server to add to the group.

Caution: When using this dialog box, you must use CTRL + Click to preserve existing server membership. If you just choose the new server and click Modify, only the new server is added, and existing server members are removed.

The row is highlighted.

10 Click Modify > Close.

The real server is added to the group in the Real Groups tab, and the Select Real Servers to include in Group dialog box closes.

- 11 From the Real Groups tab, click Apply > Close.

The real server group is configured and the SLB dialog box closes.

- 12 Continue the WAN link load balance configuration by [“Configuring filters for WAN link load balancing,”](#) next.

Configuring filters for WAN link load balancing

To configure filters for WAN link load balancing:

- 1 From the Edit menu, select WSM Card > L4 Switching > Filters.

The Filters dialog box opens.

- 2 Click Insert.

The Insert Filters dialog box ([Figure 134 on page 319](#) and [Figure 135 on page 320](#)) opens with the fields defined in [Table 57 on page 196](#).

Figure 134 WAN link filter configuration—left side

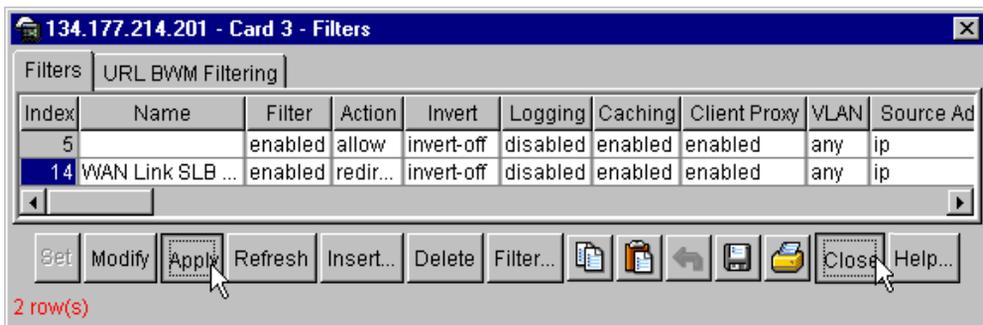
| | | |
|----------------------------------|---|---|
| Index: | 14 | 1..2048 |
| Name: | WAN Link SLB filter | |
| Filter: | <input checked="" type="radio"/> enabled <input type="radio"/> disabled | |
| Action: | redirect ▼ | |
| Invert: | <input type="radio"/> invert-on <input checked="" type="radio"/> invert-off | |
| Logging: | <input type="radio"/> enabled <input checked="" type="radio"/> disabled | |
| Caching: | <input checked="" type="radio"/> enabled <input type="radio"/> disabled | |
| Client Proxy: | <input checked="" type="radio"/> enabled <input type="radio"/> disabled | |
| VLAN: | 0 | 0..4094 (any: 0) <input type="button" value="Browse..."/> |
| Source Address Filter Type: | <input checked="" type="radio"/> ip <input type="radio"/> mac | |
| Destination Address Filter Type: | <input checked="" type="radio"/> ip <input type="radio"/> mac | |
| Source IP Address: | any | |
| Source IP Mask: | 0.0.0.0 | |
| Destination IP Address: | any | |
| Destination IP Mask: | 0.0.0.0 | |
| Source MAC Address: | 00:00:00:00:00:00 | |
| Destination MAC Address: | 00:00:00:00:00:00 | |
| Protocol: | 0 | (number) <input type="button" value="Browse..."/> |
| Low Source Port: | 0 | 0..65534 (none: 0) <input type="button" value="Browse..."/> |
| High Source Port: | 0 | 0..65534 (none: 0) <input type="button" value="Browse..."/> |
| Low Destination Port: | 0 | 0..65534 (none: 0) <input type="button" value="Browse..."/> |
| High Destination Port: | 0 | 0..65534 (none: 0) <input type="button" value="Browse..."/> |
| Redirection Port: | 0 | 0..65534 <input type="button" value="Browse..."/> |
| Redirection Group: | 1 | 1..256 |

Figure 135 WAN link filter configuration—right side

| | |
|---|---|
| URL Redirection: | <input type="radio"/> enabled <input checked="" type="radio"/> disabled |
| Network Address Translation: | <input checked="" type="radio"/> destination-address <input type="radio"/> source-address |
| NAT Active FTP: | <input type="radio"/> enabled <input checked="" type="radio"/> disabled |
| NAT Session Timeout: | <input type="text" value="4"/> 4..30 (even values only) |
| TCP ACK or RST Matching: | <input type="radio"/> enabled <input checked="" type="radio"/> disabled |
| TCP URG: | <input type="radio"/> enable <input checked="" type="radio"/> disable |
| TCP ACK: | <input type="radio"/> enable <input checked="" type="radio"/> disable |
| TCP PSH: | <input type="radio"/> enable <input checked="" type="radio"/> disable |
| TCP RST: | <input type="radio"/> enable <input checked="" type="radio"/> disable |
| TCP SYN: | <input type="radio"/> enable <input checked="" type="radio"/> disable |
| TCP FIN: | <input type="radio"/> enable <input checked="" type="radio"/> disable |
| ICMP Type: | <input type="text" value="255"/> 0..255 (any: 255) |
| IP Option: | <input type="radio"/> enable <input checked="" type="radio"/> disable |
| IP TOS: | <input type="text" value="0"/> 0..255 |
| IP TOS Mask: | <input type="text" value="0"/> 0..255 |
| New IP TOS: | <input type="text" value="0"/> 0..255 |
| TCP Connection Rate Limiting: | <input type="radio"/> enabled <input checked="" type="radio"/> disabled |
| Maximum Connection for TCP Rate Limiting: | <input type="text" value="10"/> 0..255 |
| Firewall Redirect Hash: | <input type="radio"/> enabled <input checked="" type="radio"/> disabled |
| WAN Link Load Balancing: | <input checked="" type="radio"/> enabled <input type="radio"/> disabled |
| Intrusion Detection Hash: | <input checked="" type="radio"/> sip <input type="radio"/> dip <input type="radio"/> both |
| Hash: | <input type="text" value="Automatic"/> |
| BWM Contract: | <input type="text" value="1024"/> 1..1024 |
| WAP Radius Snooping: | <input type="radio"/> enabled <input checked="" type="radio"/> disabled |

- 3 In the Index field, type a new filter number (1 to 2048).
- 4 In the Name field, type a name for the filter.
- 5 In the WAN Link Load Balancing field, click Enabled.
- 6 In the Client Proxy field, click Enabled (enables proxy IP address translation for WAN link traffic).
- 7 In the Action field, click the down arrow and choose Redirect (redirects WAN link frames).
- 8 In the Filter field, click Enabled to enable the new filter rule.
- 9 Click Insert to insert the filter definition in the Filters tab and close the Insert Filters dialog box.
- 10 From the Filters tab (Figure 136), click Apply > Close to save the changes and close the Filters dialog box.

Figure 136 Saving WAN link SLB filter configuration



The WAN link server load balancing filter is configured.

•

Chapter 11

Improving WSM performance with VMA

Virtual Matrix Architecture (VMA) improves WSM performance and increases session capacity by distributing the workload over multiple processors. VMA also removes topology constraints introduced by Direct Access Mode (DAM). With VMA, each WSM can have up to 512K concurrent sessions.

This section includes the following topics:

- [“Enabling VMA,”](#) next
- [“Using proxy IP addresses and VMA”](#) on page 325

Enabling VMA

VMA is enabled on the WSM by default. For better performance and higher session capacities, it is recommended that VMA remain enabled, especially when using Bandwidth Management and Content Intelligent Switching for processing multiple frames (up to 4500 bytes).



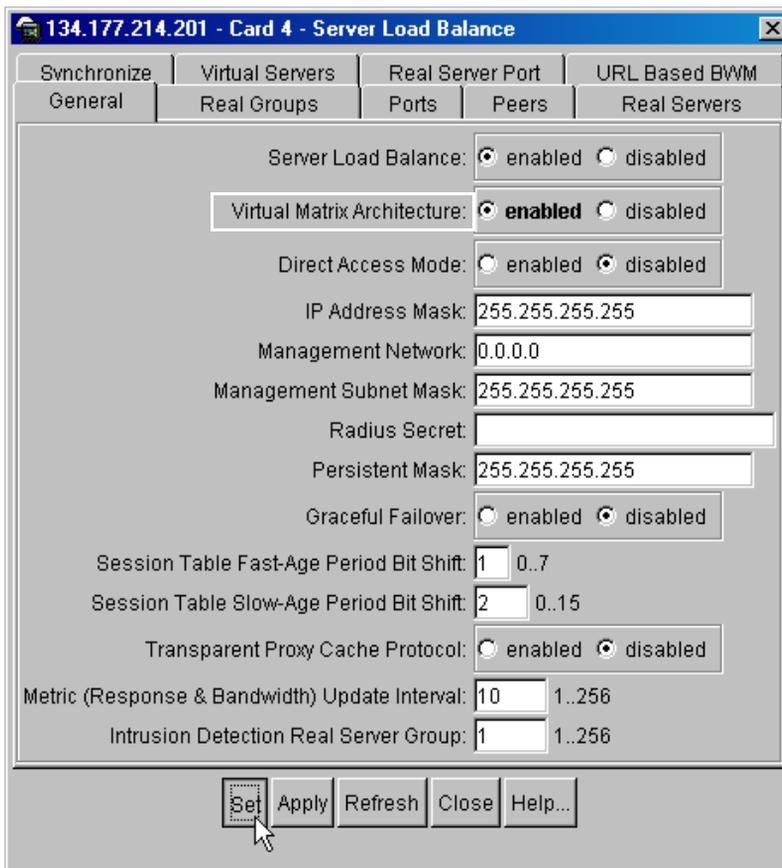
Note: VMA must be enabled if you are setting up Firewall Load Balancing (FWLB) with private-side WSMs performing server load balancing (SLB) or URL-based SLB and Direct Access Mode (DAM) is enabled.

To enable VMA on the WSM:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab (Figure 137) with the fields described in Table 61 on page 224.

Figure 137 Enabling VMA



3 In the Virtual Matrix Architecture field, click Enabled.

4 Click Set > Apply > Close.

VMA is enabled on the WSM and the Server Load Balance dialog box closes.

5 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.

- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Using proxy IP addresses and VMA

Although VMA is enabled by default on the WSM, if you have previously configured a proxy IP address for a port, VMA may be disabled initially following upgrade to the current WebOS.

Prior to enabling VMA, if any port is configured with a proxy IP address, then all ports must be. The proxy IP address identifies the processor to process responses. With VMA, a global session table replaces the concept of a per-port session table.

With proxy enabled on a WSM port with a proxy IP address, frames ingressing the port can be processed using a proxy IP address by any WSM port. The client source address is substituted with the proxy IP address of the port processing the request.

If proxy is not enabled for a WSM port with a proxy IP address, frames ingressing the port are processed by other ports configured with a proxy IP address; but the client source address will not be replaced with a proxy IP address before it is forwarded to a server.

For more information see [“Configuring proxy IP addresses” on page 245](#).

Chapter 12

Firewall server load balancing

Firewall load balancing (FWLB) lets you configure multiple active firewalls in parallel on the WSM. Parallel operation lets users maximize firewall productivity, scale firewall performance without forklift upgrades, and eliminate the firewall as a single point-of-failure.

This section includes the following topics:

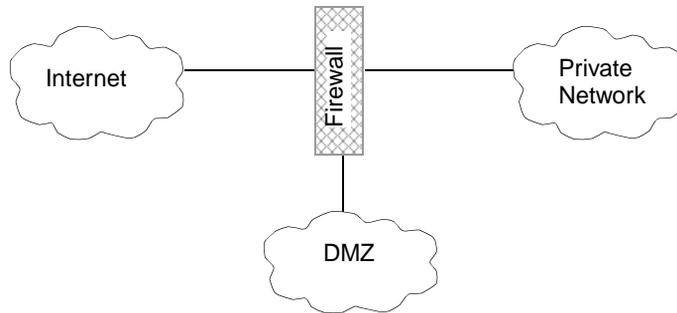
- [“Firewall overview,”](#) next
- [“Methods of firewall load balancing”](#) on page 329
- [“About basic FWLB”](#) on page 329
- [“About free-metric FWLB”](#) on page 330
- [“About demilitarized zones”](#) on page 331
- [“Basic FWLB implementation”](#) on page 331
- [“Configuring basic FWLB”](#) on page 335
- [“Monitoring the firewall”](#) on page 353

For information about four-subnet firewall load balancing, see the *Web OS Switch Software 10.0 Application Guide*, part number 212777-A.

Firewall overview

Firewall devices protect network resources from unauthorized access. Prior to FWLB, firewalls could become bottlenecks or single points-of-failure for your network.

[Figure 138](#) depicts a firewall configuration without FWLB.

Figure 138 Firewall configuration without FWLB

One network interface card on the firewall is connected to the public-side of the network, often to an Internet router. Another network interface card on the firewall is connected to the side of the network with the resources that must be protected.

In [Figure 138](#), all traffic passing between the public, private, and DMZ networks must traverse the firewall, which examines each individual packet. The firewall is configured with rules controlling traffic to allow and traffic to deny. In heavy traffic the firewall can turn into a serious bottleneck. The firewall is also a single point-of-failure device. If it goes out of service, external clients can no longer reach your services and internal clients can no longer reach the Internet.

Sometimes, a Demilitarized Zone (DMZ) is attached to the firewall or between the Internet and the firewall. Typically, a DMZ contains its own servers that provide external clients with access to services, making it unnecessary for public traffic to use private resources. WSM with FWLB provides a variety of options that enhance firewall performance and resolve typical firewall problems.

Methods of firewall load balancing

Table 73 describes the methods of firewall load balancing.

Table 73 Firewall load balancing methods

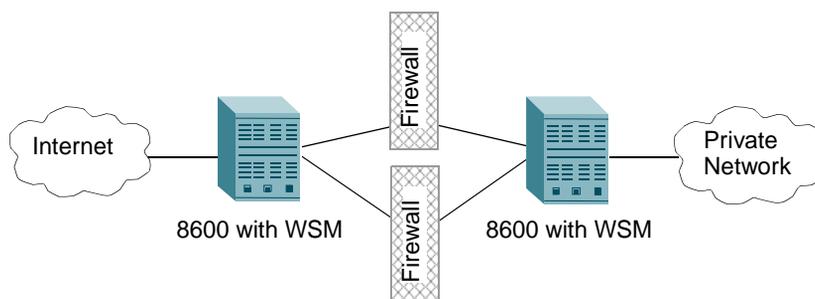
| Method | Description |
|------------------|---|
| Basic FWLB | For simple networks, this method uses a combination of static routes and redirection filters. A WSM filter on the public-side splits incoming traffic into streams headed for different firewalls. To ensure persistence of session traffic through the same firewall, distribution is based on a mathematical hash of the IP source and destination addresses. For more information about basic FWLB, see “Basic FWLB” on page 306. |
| Four-Subnet FWLB | For larger networks, the four-subnet method is more often used in networks that require high-availability solutions. This method adds Virtual Router Redundancy Protocol (VRRP) to the configuration. As in basic FWLB, four-subnet FWLB uses the hash metric to distribute firewall traffic and maintain persistence. For more information, see “Four-Subnet FWLB” in the <i>Web OS Switch Software 10.0 Application Guide</i> , part number 212777-A. |

About basic FWLB

Basic FWLB lets multiple active firewalls operate in parallel by combining static routes and redirection filters.

Figure 139 shows a basic FWLB topology.

Figure 139 Basic FWLB topology



The firewalls being load balanced are in the middle of the network, requiring a minimum of two WSMs—one on the public side of the network and one on the private side.

A redirection filter on the public-side WSM splits incoming client traffic into multiple streams. Each stream is routed through a different firewall. The valid client traffic in each stream is forwarded to a virtual server on the private-side WSM. The private-side WSM is configured with a server load balancing (SLB) metric to select a real server on the internal network for each incoming request. The same process is used for outbound server responses; a redirection filter on the private-side WSM splits the traffic, and static routes forward each stream through a different firewall and then back to the client.

The distribution of firewall load-balanced traffic within each stream is normally based on a mathematical hash of the IP source and destination addresses. This ensures that each client request and its related responses will use the same firewall (a feature known as persistence) and that the streams will be roughly equal in traffic load. Although basic firewall load-balancing techniques can support more firewalls as well as multiple WSMs on the private and public sides for redundancy, the configuration complexity increases dramatically. The four-subnet FWLB solution is usually preferred in larger scale, high-availability topologies. For information about four-subnet firewall load balancing, see the *Web OS Switch Software 10.0 Application Guide*, part number 212777-A.

About free-metric FWLB

Free-metric FWLB means that you can assign any SLB metric to the real server group. With free-metric, your FWLB configuration is not limited to the hash metric. Free-metric FWLB requires that you enable Return to Sender (RTS) on ports connected to the firewall on the public side of the network. This ensures that traffic enters and leaves through the same port. Free-metric FWLB can be used with basic FWLB or four-subnet FWLB networks.

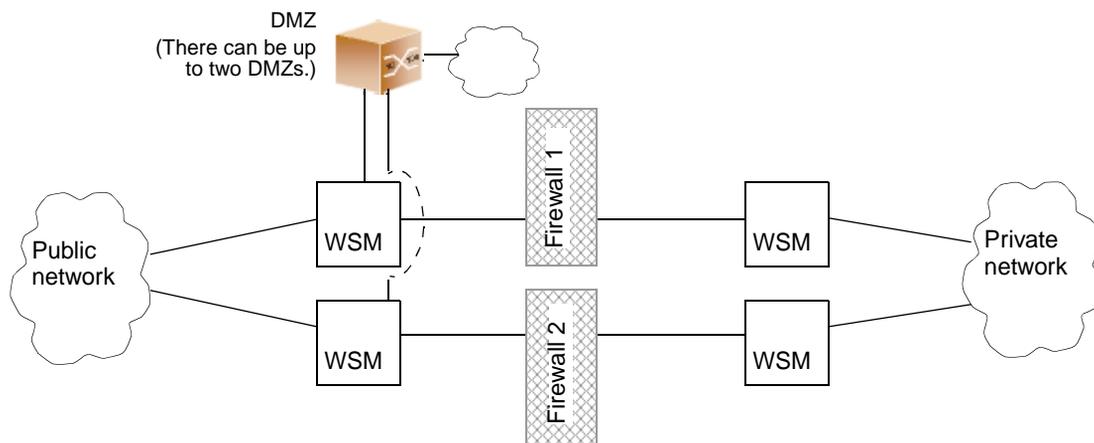
For more information, see [“About load balancing metrics” on page 207](#).

For information about four-subnet firewall load balancing, see the *Web OS Switch Software 10.0 Application Guide*, part number 212777-A.

About demilitarized zones

A demilitarized zone (DMZ) lets the WSM do the traffic filtering instead of the firewall. A FWLB DMZ is created by adding another real server group and a redirection filter toward the DMZ subnets. The DMZ servers can be connected to the WSM on the public side of the firewall. [Figure 140](#) shows a typical firewall load balancing configuration with a DMZ.

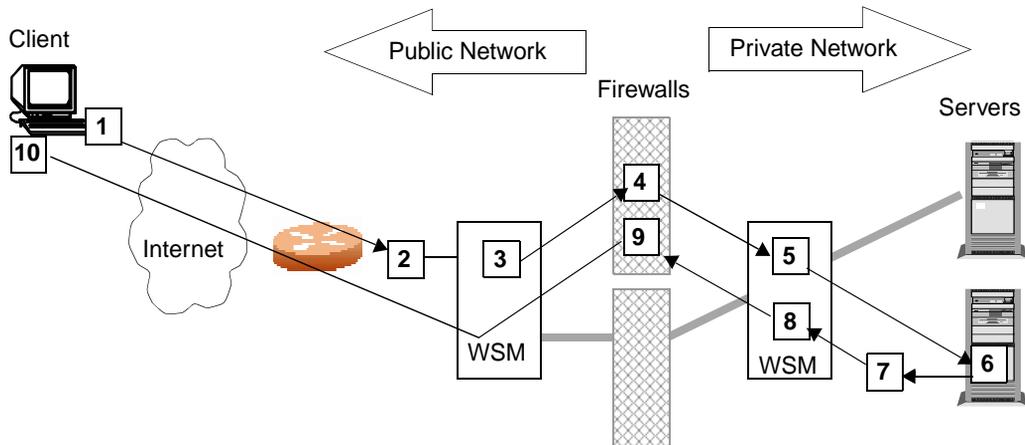
Figure 140 Firewall load-balancing topology with DMZ



The DMZ servers can be attached to the WSM directly or through an intermediate hub or switch. The WSM is then configured with filters to permit or deny access to the DMZ servers. In this manner, two levels of security are implemented: one that restricts access to the DMZ through the use of WSM filters, and another that restricts access to the private-side network through the use of stateful inspection performed by the firewalls.

Basic FWLB implementation

In [Figure 141](#) traffic is load-balanced among the available firewalls.

Figure 141 Basic FWLB implementation

The following steps describe the basic FWLB process in [Figure 141](#).

1 The client requests data.

The external clients intend to connect to services at the publicly-advertised IP address assigned to a virtual server on the private-side WSM.

2 A redirection filter balances incoming requests among different IP addresses.

When the client request arrives at the public-side WSM, a filter redirects it to a real server group that consists of a number of different IP addresses. This redirection filter splits the traffic into balanced streams: one for each IP address in the real server group. For FWLB, each IP address in the real server group represents an IP Interface (IF) on a different private-side WSM subnet.

3 Requests are routed to the firewalls.

On the public-side WSM, one static route is needed for each traffic stream. For instance, the first static route will lead to an IP interface on the private-side WSM using the first firewall as the next hop. A second static route will lead to a second private-side IP interface using the second firewall as the next hop, and so on. By combining the redirection filter and static routes, traffic is load balanced among all active firewalls.



Note: More than one stream can be routed through a particular firewall. You can weight the load to favor one firewall by increasing the number of static routes that traverse it.

All traffic between specific IP source/destination address pairs flows through the same firewall, ensuring that sessions established by the firewalls persist for their duration.

- 4 The firewalls decide if they should allow the packets and, if so, forwards them to a virtual server on the private-side WSM.

Client requests are forwarded or discarded according to rules configured for each firewall.



Note: Rules must be consistent across all firewalls.

-
- 5 The private-side WSM performs normal SLB functions.

Packets forwarded from the firewalls are sent to the original destination address, that is, the virtual server on the private-side WSM. There, they are load balanced to the real servers using standard SLB configuration.

- 6 The real server responds to the client request.
- 7 Redirection filters on the private-side WSM balance responses among different IP addresses.

Redirection filters are needed on all ports on the private-side WSM that attach to real servers or internal clients on the private-side of the network. Filters on these ports redirect the Internet-bound traffic to a real server group that consists of a number of different IP addresses. Each IP address represents an IP interface on a different subnet on the public-side WSM.

- 8 Outbound traffic is routed to the firewalls.

Static routes are configured on the private-side WSM. One static route is needed for each stream that was configured on the public-side WSM. For instance, the first static route would be configured to lead to the first public-side IP interface using the first firewall as the next hop. The second

static route would lead to the second public-side IP interface using the second firewall as the next hop, and so on. Since WSMs intelligently maintain state information, all traffic between specific IP source/destination addresses flows through the same firewall, maintaining session persistence.



Note: If Network Address Translation (NAT) software is used on the firewalls, FWLB session persistence requires RTS to be enabled on the WSM. For more information, see [“Configuring basic FWLB with free-metric” on page 352](#).

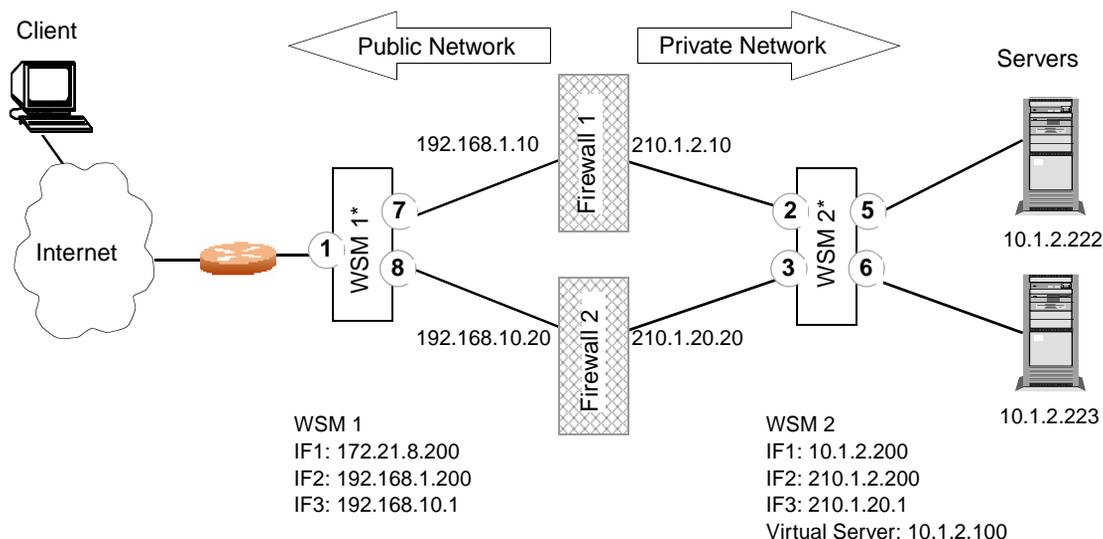
- 9 The firewall decides if it should allow the packet and, if so, forwards it to the public-side WSM.

Each firewall forwards or discards the server responses according to the rules that are configured for it. Forwarded packets are sent to the public-side WSM and out to the Internet.
- 10 The client receives the server response.

Configuring basic FWLB

Figure 142 shows a simple network topology with FWLB. This example uses two WSMs and two 8648TX modules in a single 8600. The public-side WSM is installed in slot 4, and the private-side WSM is installed in slot 10. The upstream router, Web server farm, and firewalls are physically connected to the 8648TX modules. The WSMs, connected to the Passport 8648TX backplane, perform the firewall load balancing.

Figure 142 Basic FWLB example configuration



*Both WSMs are installed in the same 8600

The configuration in Figure 142, requires the following:

- “Configuring the 8600 switch for FWLB,” next
- “Configuring FWLB on the public-side of the WSM network” on page 338
- “Configuring FWLB on the private-side of the WSM network” on page 343
- “Configuring the static routes at the firewalls for FWLB” on page 351
- “Configuring basic FWLB with free-metric” on page 352 (optional)
- “Configuring basic FWLB with a DMZ” on page 353

Configuring the 8600 switch for FWLB

To configure the 8600 in [Figure 142 on page 335](#):

- 1 Configure new 8648TX VLANs using the settings in [Table 74](#). Assume all new VLANs use the same spanning-tree group. For more information, see *Creating a port-based VLAN* in the publication, *Configuring Layer 2 Operations*, part number 314725-A.

Table 74 FWLB example—8648TX VLAN port settings

| VLAN field in Device Manager | Setting for 8648TX |
|---------------------------------|---|
| Id | 11 |
| Name | FWLB WSM |
| StgId | (Assume all new VLANs use same STG) |
| Type | byPort |
| Port Members | 1/33, 1/48 (connect to interface 1 of firewall 1 and 2) |
| Id | 12 |
| Name | FWLB WSM |
| StgId | (Assume all new VLANs use same STG) |
| Type | byPort |
| Port Members | 1/1, 1/16 (connect to all SLB clients) |
| Id | 13 |
| Name | FWLB WSM |
| StgId | (Assume all new VLANs use same STG) |
| Type | byPort |
| Port Members | 7/33, 7/48 (connect to interface 2 of firewalls 1 and 2) |
| Id | 14 |
| Name | FWLB WSM |
| StgId | (Assume all new VLANs use same STG) |
| Type | byPort |
| Port Members | 7/1, 7/16 (connect to all SLB servers) |

- 2 Add the newly created VLANs to MLT groups using the settings in [Table 75](#). For more information, see *Configuring a multilink trunk* in the publication, *Configuring Layer 2 Operations*, part number 314725-A.

Table 75 FWLB example—8648TX MLT settings

| MLT field in Device Manager | Setting for 8648TX |
|-----------------------------|--------------------|
| Id | 32 |
| VlanId | 11 |
| Id | 31 |
| VlanId | 12 |
| Id | 30 |
| VlanId | 13 |
| Id | 29 |
| VlanId | 14 |

- 3 Remove VLAN 1 and VLAN 2 from the default MLT configuration.

Table 76 FWLB example—8648TX MLT settings

| MLT field in Device Manager | Setting for 8648TX |
|-----------------------------|--------------------|
| Id | 32 |
| VlanId | 1 (remove) |
| Id | 30 |
| VlanId | 1 (remove) |

Configuring FWLB on the public-side of the WSM network

To configure the public-side network in [Figure 142 on page 335](#):

- 1 Configure VLANs on public-side WSM-1 using the settings in [Table 77](#). For more information, see [“Configuring a VLAN” on page 106](#).

Table 77 Public-side FWLB network—VLAN settings

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| VLAN | 11 |
| Name | FWLB |
| State | Enabled |
| Ports | 7, 8 |
| VLAN | 12 |
| Name | FWLB |
| State | Enabled |
| Ports | 5, 6 |

- 2 On each public-side WSM port, set the default VLAN number which will be used to forward frames which are not VLAN tagged. Use the settings in [Table 78](#). For more information, see [“Setting port parameters” on page 85](#).

Table 78 Public-side FWLB network—port default VLAN settings

| Port | Field in Device Manager | Settings for WSM-1 |
|------|-------------------------|--------------------|
| 5 | Default VLAN | 12 |
| 6 | Default VLAN | 12 |
| 7 | Default VLAN | 11 |
| 8 | Default VLAN | 11 |

- 3 Remove the rear-facing ports from default VLANs 1 and 2.

Table 79 Public-side FWLB network—ports to remove from default VLANs

| VLAN | Ports to remove |
|------|-----------------|
| 1 | 7 and 8 |
| 2 | 5 and 6 |

- 4 Define IP interfaces for the public-side network using the settings in [Table 80](#). For more information see, [See “Manually configuring an IP interface” on page 130](#).

There must be an IP interface for general WSM management (IF1), in addition to a public-side IP interface for each firewall path being load balanced (IF2 and IF3).

Table 80 Public-side FWLB network—IP interface settings

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| Interface Number | 1 |
| IP Address | 172.21.8.200 |
| IP Subnet Mask | 255.255.255.0 |
| IP Broadcast Address | 172.21.8.255 |
| VLAN | 12 |
| State | Enabled |
| Interface Number | 2 |
| IP Address | 192.168.1.200 |
| IP Subnet Mask | 255.255.255.0 |
| IP Broadcast Address | 192.168.1.255 |
| State | Enabled |
| VLAN | 11 |
| Interface Number | 3 |
| IP Address | 192.168.10.1 |
| IP Subnet Mask | 255.255.255.0 |
| IP Broadcast Address | 192.168.10.255 |

Table 80 Public-side FWLB network—IP interface settings (continued)

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| State | Enabled |
| VLAN | 11 |

- 5 Enable SLB on WSM-1 for the public-side network. For more information, see [“Enabling or disabling server load balancing”](#) on page 243.
- 6 On the public-side WSM, create two real servers using the IP address of each private-side FWLB IP interface. Use the settings in [Table 81](#) . For more information, see [“Configuring each real server”](#) on page 216.

Later in this procedure, you’ll configure one private-side IP interface on a different subnet for each firewall path being load balanced.



Note: Each of the four IFs used for FWLB (two on each WSM) in this example must be configured for a different IP subnet.

Table 81 Public-side FWLB network—real server settings

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| Real Server | 200 |
| IP Address | 210.1.2.200 |
| State | Enabled |
| Name | FWLB Server |
| Real Server | 1 |
| IP Address | 210.1.20.1 |
| State | Enabled |
| Name | FWLB Server |

- 7 Place the IP interface real servers into a real server group using the settings in [Table 82](#). For more information, see [“Configuring a real server group”](#) on page 225.

Using the hash metric, all traffic between specific IP source/destination address pairs flows through the same firewall. This ensures that sessions established by the firewalls are maintained for their duration.



Note: Other load balancing metrics such as least connections, roundrobin, minimum misses, response, and bandwidth can be used by enabling Return to Sender (RTS). For more information, see [“Configuring basic FWLB with free-metric” on page 352](#).

Table 82 Public-side FWLB network—real server group settings

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| Group | 1 |
| Name | FWLB Group |
| Metric | hash |
| Health Check | icmp |
| Real Servers | 1 and 200 |

- 8 Create a filter to allow local subnet traffic on the public side of the firewalls to reach the firewall interfaces. Use the settings in [Table 83](#). For more information, see [“Creating a new filter” on page 201](#).

Table 83 Public-side FWLB network—firewall filter settings

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| Index | 10 |
| Name | FWLB |
| Filter | Enabled |
| Action | Allow |
| Source IP Address | Any |
| Destination IP Address | 172.21.8.0 |
| Destination IP Mask | 255.255.255.0 |

- 9 Create the FWLB redirection filter, using the settings in [Table 84](#). For more information, see [“Creating a new filter” on page 201](#).

This filter redirects all inbound traffic, and load balances it among the defined real servers in the group. In this network, the real servers in group 1 represent IP interfaces on the private-side WSM.

Table 84 Public-side FWLB network—redirection filter settings

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| Index | 100 |
| Name | FWLB Redirect |
| Filter | Enabled |
| Action | Redirect |
| Source IP Address | any |
| Destination IP Address | any |
| Protocol | any |
| Redirection Group | 1 |

- 10 Add filters to the ingress ports using the settings [Table 85](#). For more information, see [“Enabling or disabling filtering on a port” on page 89](#) and [“Applying filters to a port” on page 90](#).

Table 85 Public-side FWLB network—port filter settings

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| Port | 5 |
| Filtering | Enabled |
| Filters Applied | 10 and 100 |
| Port | 6 |
| Filtering | Enabled |
| Filters Applied | 10 and 100 |

- 11** Define static routes from the public to the private-side IP interfaces, using the firewalls as gateways. Use the settings in the following table. For more information, see [“Configuring static routes”](#) on page 148.

Table 86 Public-side FWLB network—IP static routes settings

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| Static Route | [1 - 128] |
| Destination IP Address | 210.1.2.200 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 192.168.1.10 |
| IP Interface | 2 |
| Static Route | [1 - 128] |
| Destination IP Address | 210.1.20.1 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 192.168.10.20 |
| IP Interface | 3 |

Configuring FWLB on the private-side of the WSM network

To configure the private-side network in [Figure 142](#):

- 1 Configure VLANs on WSM-2 (private-side network) using the settings in [Table 87](#). For more information, see [“Configuring a VLAN”](#) on page 106.

Table 87 Private-side FWLB network—VLAN settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| VLAN | 13 |
| Name | FWLB |
| State | Enabled |
| Ports | 7, 8 |
| VLAN | 14 |
| Name | FWLB |

Table 87 Private-side FWLB network—VLAN settings (continued)

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| State | Enabled |
| Ports | 5, 6 |

- On each WSM port for the private-side network, set the default VLAN number which will be used to forward frames which are not VLAN tagged. Use the settings in [Table 78](#). For more information, see [“Setting port parameters” on page 85](#).

Table 88 Private-side FWLB network—port default VLAN settings

| Port | Field in Device Manager | Settings for WSM-2 |
|------|-------------------------|--------------------|
| 5 | Default VLAN | 14 |
| 6 | Default VLAN | 14 |
| 7 | Default VLAN | 13 |
| 8 | Default VLAN | 13 |

- Remove the rear-facing ports from default VLANs 1 and 2 of the private-side network. For more information, see [“Configuring a VLAN” on page 106](#)

Table 89 Private-side FWLB network—ports to remove from default VLANs

| VLAN | Ports to remove |
|------|-----------------|
| 1 | 7 and 8 |

- Define the private-side IP interfaces using the settings in [Table 90](#). For more information, see [“Manually configuring an IP interface” on page 130](#).

Create one private-side IP interface on a different subnet for each firewall being load-balanced.

Table 90 Private-side FWLB network—IP interface settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Interface Number | 1 |
| IP Address | 10.1.2.200 |

Table 90 Private-side FWLB network—IP interface settings (continued)

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| IP Subnet Mask | 255.255.255.0 |
| IP Broadcast Address | 10.1.2.255 |
| VLAN | 14 |
| State | Enabled |
| Interface Number | 2 |
| IP Address | 210.1.2.200 |
| IP Subnet Mask | 255.255.255.0 |
| IP Broadcast Address | 210.1.2.255 |
| State | Enabled |
| VLAN | 13 |
| Interface Number | 3 |
| IP Address | 210.1.20.1 |
| IP Subnet Mask | 255.255.255.0 |
| IP Broadcast Address | 210.1.20.255 |
| State | Enabled |
| VLAN | 13 |

- 5 Create two real servers on the private-side WSM, using the IP address of each public-side IP interface. Use the settings in [Table 91](#) . For more information, see [“Configuring each real server” on page 216](#).

You should already have configured a public-side IP interface on a different subnet for each firewall path being load-balanced.

Table 91 Private-side FWLB network—real server settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Real Server | 1 |
| IP Address | 192.168.10.1 |
| State | Enabled |
| Name | FWLB Server |
| Real Server | 200 |

Table 91 Private-side FWLB network—real server settings (continued)

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| IP Address | 192.168.1.200 |
| State | Enabled |
| Name | FWLB Server |

- 6 Place the real servers (public-side IP interfaces) into a real server group using the settings in [Table 92](#). For more information, see [“Configuring a real server group”](#) on page 225.



Note: The private-side WSM must use the same metric defined on the public side.

Table 92 Private-side FWLB network—real server group settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Group | 1 |
| Name | FWLB Group |
| Metric | hash |
| Health Check | icmp |
| Real Servers | 1 and 200 |

- 7 Configure client processing on ports 7 and 8, which are connected to the private-side of the firewalls. For more information, see [“Configuring ports for server load balancing”](#) on page 240.

Table 93 Private-side FWLB network—port SLB settings

| Port | Field in Device Manager | Setting for WSM-2 |
|------|-------------------------|-------------------|
| 7 | Load Balanced State | client |
| 8 | Load Balanced State | client |

- 8 Configure the real servers to which traffic will be load-balanced. These are the actual servers on the network. Use the settings in [Table 94](#). For more information, see [“Configuring each real server” on page 216](#).

Table 94 Private-side FWLB network—real server settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Real Server | 222 |
| IP Address | 10.1.2.222 |
| State | Enabled |
| Name | FWLB Server |
| Real Server | 223 |
| IP Address | 10.1.2.223 |
| State | Enabled |
| Name | FWLB Server |

- 9 Place the real servers into a real server group using the settings in the following table. For more information, see [“Configuring a real server group” on page 225](#).

Table 95 Private-side FWLB network—real server group settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Group | 200 |
| Name | FWLB Group |
| Metric | hash |
| Health Check | icmp |
| Real Servers | 222 and 223 |

- 10** Configure the virtual server that will load balance the real servers using the settings in Table 96. For more information, see [“Configuring a virtual server” on page 230](#) and [“Configuring services for a virtual server” on page 235](#).

Table 96 Private-side FWLB network—virtual server settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Virtual Server | 100 |
| IP Address | 10.1.2.100 |
| State | Enabled |
| Virtual Service | [1 - 8] |
| Real Group | 200 |

- 11** Configure ports 5 and 6, which are connected to the real servers for server processing using the settings in the following table. For more information, see [“Configuring ports for server load balancing” on page 240](#).

Table 97 Private-side FWLB network—port SLB settings

| Port | Field in Device Manager | Setting for WSM-2 |
|------|-------------------------|-------------------|
| 5 | Load Balanced State | server |
| 6 | Load Balanced State | server |

- 12** Enable server load balancing on the WSM. For more information, see [“Enabling or disabling server load balancing” on page 243](#).
- 13** Create a filter to prevent server-to-server traffic from being redirected. Use the settings in Table 98. For more information, see [“Creating a new filter” on page 201](#).

Table 98 Private-side FWLB network—firewall filter settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Index | 50 |
| Name | FWLB |

Table 98 Private-side FWLB network—firewall filter settings (continued)

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Filter | Enabled |
| Action | Allow |
| Source IP Address | Any |
| Destination IP Address | 10.1.2.0 |
| Destination IP Mask | 255.255.255.0 |

- 14** Create the redirection filter for the private-side network using the settings in the following table. For more information, see [“Creating a new filter” on page 201](#).

This filter will redirect outbound traffic, load-balancing it among the defined real servers in the group. In this case, the real servers represent IP interfaces on the public-side WSM.

Table 99 Private-side FWLB network—firewall redirection filter settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Index | 100 |
| Name | FWLB Redirect |
| Filter | Enabled |
| Action | Redirect |
| Source IP Address | Any |
| Destination IP Address | Any |
| Proto | Any |
| Redirection Group | 1 |

- 15** Add the filters to the ingress ports for the outbound packets using the settings in [Table 85](#). For more information, see [“Enabling or disabling filtering on a port” on page 89](#) and [“Applying filters to a port” on page 90](#).

Redirection filters are needed on all the ingress ports on the private-side WSM. Ingress ports attach to real servers or internal clients on the private-side of the network. In this case, two real servers are attached to the private-side WSM on rear-facing ports 5 and 6.

Table 100 Private-side FWLB network—port filter settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Port | 5 |
| Filtering | Enabled |
| Filters Applied | 50 and 100 |
| Port | 6 |
| Filtering | Enabled |
| Filters Applied | 50 and 100 |

- 16** Define static routes to the public-side IP interfaces, using the firewalls as gateways. Use the settings in [Table 101](#). For more information, see [“Configuring static routes” on page 148](#).

One static route is required for each firewall path being load balanced. In this case, two paths are required: Interface 2, which leads to public-side IF 2 (192.168.1.200) through the first firewall (210.1.2.10) as its gateway, and Interface 3, which leads to public-side IF 3 (192.168.10.1) through the second firewall (210.1.20.20) as its gateway.

Table 101 Private-side FWLB network—IP static routes settings

| Field in Device Manager | Setting for WSM-2 |
|-------------------------|-------------------|
| Static Route | [1 - 128] |
| Destination IP Address | 192.168.1.200 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 210.1.2.10 |
| IP Interface | 2 |
| Static Route | [1 - 128] |
| Destination IP Address | 192.168.10.1 |
| IP Subnet Mask | 255.255.255.255 |

Table 101 Private-side FWLB network—IP static routes settings (continued)

| Field in Device Manager | Setting for WSM-2 |
|----------------------------|----------------------|
| Gateway IP Address | 210.1.20.20 |
| IP Interface | 3 |



Note: Configuring static routes for FWLB does not require that IP forwarding be turned on.

Configuring the static routes at the firewalls for FWLB

To configure static routes at the firewalls in [Figure 142 on page 335](#):

- 1 At firewall 1, define the network and mask for the gateway using the settings in the following table. For more information, see [“Configuring static routes” on page 148](#).

Table 102 FWLB example—firewall 1 IP static routes settings

| Field in Device Manager | Setting for firewall 1 |
|----------------------------|---------------------------|
| Static Route | [1 - 128] |
| Destination IP Address | 172.21.8.0 |
| IP Subnet Mask | 255.255.255.0 |
| Gateway IP Address | 192.168.1.200 |
| Static Route | [1 - 128] |
| Destination IP Address | 10.1.2.0 |
| IP Subnet Mask | 255.255.255.0 |
| Gateway IP Address | 210.1.2.200 |

- At firewall 2, define the network and mask for the gateway using the settings in the following table. For more information, see [“Configuring static routes” on page 148](#).

Table 103 FWLB example—firewall 2 IP static routes settings

| Field in Device Manager | Setting for firewall 2 |
|-------------------------|------------------------|
| Static Route | [1 - 128] |
| Destination IP Address | 172.21.8.0 |
| IP Subnet Mask | 255.255.255.0 |
| Gateway IP Address | 192.168.10.1 |
| Static Route | [1 - 128] |
| Destination IP Address | 10.1.2.0 |
| IP Subnet Mask | 255.255.255.0 |
| Gateway IP Address | 210.1.20.1 |

Configuring basic FWLB with free-metric

Free-metric FWLB lets you use other load balancing metrics besides hash. For more information, see [“About free-metric FWLB” on page 330](#).

To configure free-metric for the [basic FWLB](#) figure on [page 335](#):

- On the private-side WSM, enable return to sender (RTS) on the ports attached to the firewalls (ports 7 and 8). For more information, see [“Configuring ports for server load balancing” on page 240](#).
- On the public-side WSM, remove the redirection filter from the ports attached to the real servers (ports 5 and 6), but make sure filter processing is enabled. For more information, see [“Creating a new filter” on page 201](#) and [“Enabling or disabling filtering on a port” on page 89](#).
- On the public-side WSM, set the FWLB metric (hash, least connections, round robin, minimum misses, response, or bandwidth). See [“About load balancing metrics” on page 207](#).

Configuring basic FWLB with a DMZ

To add the filters required for a DMZ:

- 1 On the public-side WSM, create a filter to allow HTTP traffic to reach the DMZ Web servers. In this example, the DMZ Web servers use IP addresses 205.178.29.0/24. For more information, see [“Creating a new filter” on page 201](#).
- 2 Create another filter to deny all other traffic to the DMZ Web servers. For more information, see [“Creating a new filter” on page 201](#).
- 3 Add the filters to the traffic ingress ports. See [“Enabling or disabling filtering on a port” on page 89](#) and [“Applying filters to a port” on page 90](#).

Monitoring the firewall

Basic FWLB health checking is automatic. No special configuration is necessary unless you wish to tune the health checking parameters. For more information, see [“Health Checking” on page 403](#).

Monitoring firewall service

To maintain high availability, WSMs monitor firewall health status and send packets only to healthy firewalls. There are two methods of firewall service monitoring: ICMP and HTTP. Each WSM monitors the health of the firewalls at scheduled intervals.

If a WSM IP interface fails to respond to a user-specified health check interval, it (and, by implication, the associated firewall), is placed in a Server Failed state. At this time, the partner WSM stops routing traffic to that IP interface and, instead, distributes it across the remaining healthy WSM IP interfaces and firewalls.

When a WSM IP interface is in the Server Failed state, its partner WSM continues to forward health checks at user-defined intervals. After a specified number of successful checks, the IP interface (and its associated firewall) is brought back into service. You can configure the WSM to allow, for example, one-second intervals between health checks or pings, two failed health checks to remove the firewall, and four successful health checks to restore the firewall to the real server group. For more information, see [“Health Checking” on page 403](#).

Monitoring physical links

The WSM also monitors physical link status of the ports connected to firewalls. If the physical link to a firewall goes down, that firewall is placed immediately in the *Server Failed* state. When a WSM detects that a failed physical link to a firewall has been restored, it brings the firewall back into service. For more information, see [“Health Checking” on page 403](#).

Using HTTP Health Checks

For those firewalls that do not permit ICMP pings to pass through, WSMs can be configured to perform HTTP health checks, as described below.

- 1 Set the health check type to HTTP instead of ICMP. See [“Configuring a health check for a real server group” on page 408](#).
- 2 Configure a “dummy” redirection filter as the last filter (after the *redirect all* filter) to force the HTTP health checks to activate. For more information, see [“Creating a new filter” on page 201](#). Use the settings in Table 104.

Table 104 FWLB redirection filter settings

| Field in Device Manager | WSM setting |
|----------------------------|---------------|
| Index | 224 |
| Name | FWLB Redirect |
| Filter | Enabled |
| Action | Redirect |
| Proto | tcp |

Table 104 FWLB redirection filter settings (continued)

| Field in Device Manager | WSM setting |
|----------------------------|-------------|
| Redirection Port | 80 |
| Redirection Group | 1 |



Note: Make sure that the number of each real filter is lower than the number of the “dummy” redirect filter.

Chapter 13

Virtual private network server load balancing

This section describes using VPN load balancing, and includes the following topics:

- [“Defining virtual private network,”](#) next
- [“How VPN load balancing works”](#) on page 358
- [“Configuring VPN load balancing”](#) on page 360
- [“VPN load-balancing configuration example”](#) on page 360

Defining virtual private network

A virtual private network (VPN) is a connection that has the appearance and advantages of a dedicated link, but occurs over a shared network. Using tunneling, data packets are transmitted across a routed network, such as the Internet, in a private tunnel that simulates a point-to-point connection. This enables network traffic from many sources to travel via separate tunnels across the infrastructure. It also enables traffic from many sources to be differentiated, for directing to specific destinations at specific levels of service.

VPNs provide:

- the security features of a firewall
- network address translation
- data encryption
- authentication and authorization

Since most of the data sent over a VPN is encrypted, network devices cannot use information inside the packet to make intelligent routing decisions.

How VPN load balancing works

The WSM can simultaneously load balance up to 255 VPN devices while ensuring that the traffic returns back to the same VPN server from which it started.

Traffic coming from the Internet is usually addressed to the VPNs, with the real destination encrypted inside the datagram. Traffic from the VPNs to the intranet contains the real destination in the clear. Using the hash algorithm on the source and destination address may not be possible in many VPN/firewall configurations because the address may be encrypted inside the datagram. Also, the source/destination IP address of the packet may change as the packet traverses from the public-side WSM to private-side WSM and back.

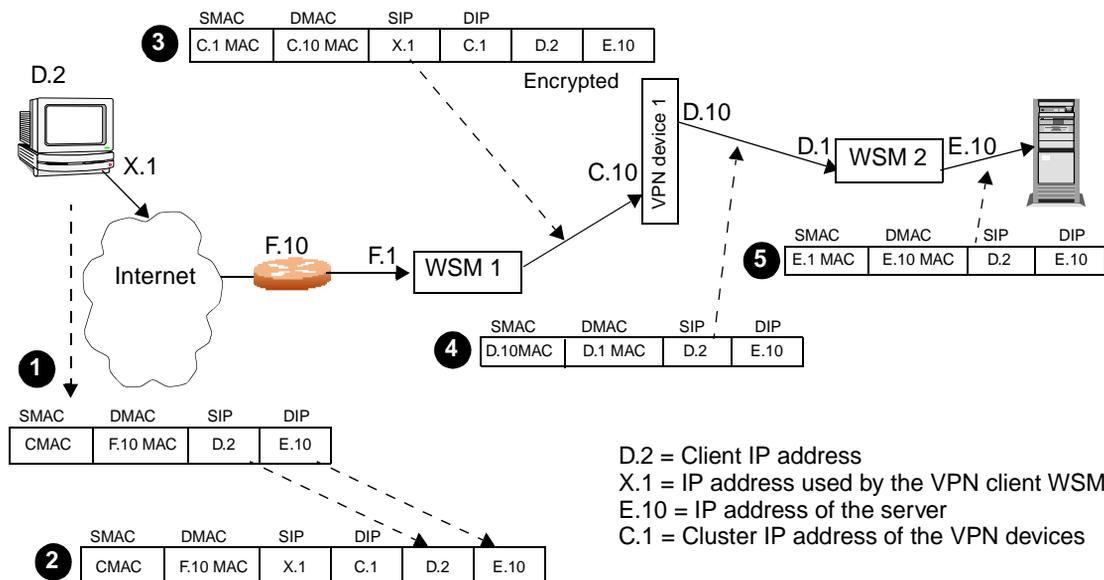
The WSM records state information in a session table for frames entering the WSM to and from the VPNs. This session table ensures that the same VPN server handles all the traffic between an inside host and an outside client for a particular session.



Note: VPN load balancing is supported for connecting from remote sites to the network behind the VPN cluster IP address. Connections initiated from clients internal to the VPN gateways is not supported.

Figure 143 illustrates the basic frame flow of a request arriving from the Internet. An external client is accessing an internal server. No network address translation (NAT) is performed by the VPN devices.

Figure 143 Basic network frame flow and operation



Configuring VPN load balancing

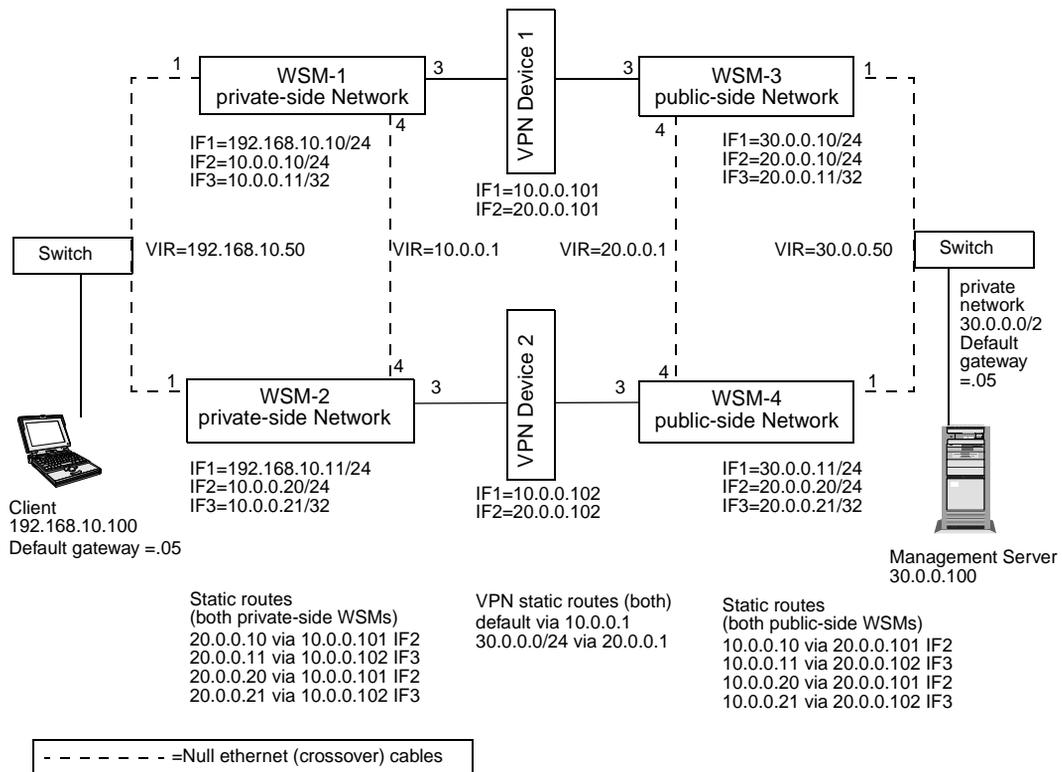
To configure VPN server load balancing:

- Configure the WSM for firewall load balancing. For more information, see [“Firewall server load balancing” on page 327](#).
- Enable the Return to Sender (RTS) feature on the ports attached to the VPN devices. See [“Configuring ports for server load balancing” on page 240](#).

VPN load-balancing configuration example

Figure 144 uses four WSMs, four subnets, and two VPN devices for VPN load balancing.

Figure 144 VPN load balancing configuration example



To implement the VPN in [Figure 144](#), build the topology and configure the private and public sides of the network as described in the following sections.

- “[Configuring private-side network devices](#),” next
- “[Configuring public-side network devices](#)” on page 365

Configuring private-side network devices

To configure the private-side network devices (WSM-1 and WSM-2) in the [Figure 144](#):

- 1 Define and enable VLAN 2 for ports 3 and 4. See “[Configuring a VLAN](#)” on page 106.
- 2 Turn off Spanning Tree Protocol (STP) on ports 3 and 4. See “[Enabling or disabling spanning tree on a port](#)” on page 125.
- 3 Define the private-side network IP interfaces using the settings in [Table 105](#). See “[Manually configuring an IP interface](#)” on page 130.

Create one private-side network IP interface on a different subnet for each VPN device being load-balanced.

Table 105 VPN example—private network IP interface settings

| Field in Device Manager | Setting for WSM-1 | Setting for WSM-2 |
|-------------------------|-------------------|-------------------|
| Interface Number | 1 | 1 |
| IP Address | 30.0.0.10 | 30.0.0.11 |
| IP Subnet Mask | 255.255.255.0 | 255.255.255.0 |
| VLAN | 1 | 1 |
| State | Enabled | Enabled |
| BOOTP Relay | Disabled | Disabled |
| Interface Number | 2 | 2 |
| IP Address | 20.0.0.10 | 20.0.0.20 |
| IP Subnet Mask | 255.255.255.0 | 255.255.255.0 |
| VLAN | 2 | 2 |
| State | Enabled | Enabled |
| BOOTP Relay | Disabled | Disabled |

Table 105 VPN example—private network IP interface settings (continued)

| Field in Device Manager | Setting for WSM-1 | Setting for WSM-2 |
|-------------------------|-------------------|-------------------|
| Interface Number | 3 | 3 |
| IP Address | 20.0.0.11 | 20.0.0.21 |
| IP Subnet Mask | 255.255.255.255 | 255.255.255.255 |
| VLAN | 2 | 2 |
| State | Enabled | Enabled |
| BOOTP Relay | Disabled | Disabled |

- 4 Configure routes for each of the IP interfaces using the VPN devices as gateways. Configure one static route for each VPN device being load-balanced using the settings in [Table 106e](#). See “[Configuring static routes](#)” on page 148.

Table 106 VPN example—private network IP static routes settings

| Field in Device Manager | Setting for WSM-1 and WSM-2 |
|-------------------------|-----------------------------|
| Static Route | [1 - 128] |
| Destination IP Address | 10.0.0.10 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 20.0.0.101 |
| IP Interface | 2 |
| Static Route | [1 - 128] |
| Destination IP Address | 10.0.0.11 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 20.0.0.102 |
| IP Interface | 3 |
| Static Route | [1 - 128] |
| Destination IP Address | 10.0.0.20 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 20.0.0.101 |
| IP Interface | 2 |
| Static Route | [1 - 128] |

Table 106 VPN example—private network IP static routes settings (continued)

| Field in Device Manager | Setting for WSM-1 and WSM-2 |
|-------------------------|-----------------------------|
| Destination IP Address | 10.0.0.21 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 20.0.0.102 |
| IP Interface | 3 |

- 5 Enable virtual routing for the WSM. See [“Enabling virtual routing on the WSM” on page 151](#).
- 6 Configure VRRP for virtual routers 1 and 2 by making the settings in [Table 107](#). See [“Configuring a virtual router” on page 153](#).

Table 107 VPN example—private network virtual router settings

| Field in Device Manager | Setting for WSM-1 and WSM-2 |
|------------------------------|-----------------------------|
| Virtual Router Number | 1 |
| Virtual Router ID | 1 |
| IP Address | 30.0.0.50 |
| IP Interface | 1 |
| State | Enabled |
| Priority | 101 |
| Load Sharing | Disabled |
| Track VRs | Enabled |
| Track VLAN Ports | Enabled |
| Virtual Router Number | 2 |
| Virtual Router ID | 2 |
| IP Address | 20.0.0.1 |
| IP Interface | 2 |
| State | Enabled |
| Priority | 101 |
| Load Sharing | Disabled |

Table 107 VPN example—private network virtual router settings (continued)

| Field in Device Manager | Setting for WSM-1 and WSM-2 |
|----------------------------|--------------------------------|
| Track VRs | Enabled |
| Track VLAN Ports | Enabled |

- 7 Enable Server Load Balancing (SLB) on the private-side network WSMs. See [“Enabling or disabling server load balancing” on page 243](#).
- 8 Configure real servers for health checking VPN devices using the settings in [Table 108](#). See [“Configuring each real server” on page 216](#).

Table 108 VPN example—private network real server settings

| Field in Device Manager | Setting for WSM-1 and WSM-2 |
|----------------------------|--------------------------------|
| Real Server | 1 |
| IP Address | 10.0.0.10 |
| State | Enabled |
| Name | VPN Server 1 |
| Real Server | 2 |
| IP Address | 10.0.0.11 |
| State | Enabled |
| Name | VPN Server 2 |
| Real Server | 3 |
| IP Address | 10.0.0.20 |
| State | Enabled |
| Name | VPN Server 3 |
| Real Server | 4 |
| IP Address | 10.0.0.21 |
| State | Enabled |
| Name | VPN Server 4 |

- 9 Configure real server group 1 with the hash metric; and add real servers 1, 2, 3, and 4 to the group. Use the settings in [Table 109](#). See “[Configuring a real server group](#)” on page 225.

Table 109 VPN example—private network real server group settings

| Field in Device Manager | Setting for WSM-1 and WSM-2 |
|-------------------------|-----------------------------|
| Group | 1 |
| Name | VPN Group |
| Metric | hash |
| Real Servers | 1, 2, 3, 4 |

- 10 Enable RTS on the necessary ports (ports 3 and 4). See “[Configuring ports for server load balancing](#)” on page 240.
- 11 Enable filter processing on the server ports (port 1) so that the responses from the real server will be looked up in the VPN session table. See “[Enabling or disabling filtering on a port](#)” on page 89.
- 12 Apply and save the configuration, and reboot the WSM.

Configuring public-side network devices

To configure the public-side network devices (WSM-3 and WSM-4) in [Figure 144](#):

- 1 Define and enable VLAN 2 for ports 3 and 4. See “[Configuring a VLAN](#)” on page 106.
- 2 Turn off Spanning Tree Protocol (STP) for ports 3 and 4. See “[Enabling or disabling spanning tree on a port](#)” on page 125.
- 3 Define the public-side network IP interfaces using the settings in [Table 105](#). See “[Manually configuring an IP interface](#)” on page 130.

Create one public-side network IP interface on a different subnet for each VPN device being load-balanced.

Table 110 VPN example—public network IP interface settings

| Field in Device Manager | Setting for WSM-3 | Setting for WSM-4 |
|-------------------------|-------------------|-------------------|
| Interface Number | 1 | 1 |
| IP Address | 192.168.10.10 | 192.168.10.11 |
| IP Subnet Mask | 255.255.255.0 | 255.255.255.0 |
| VLAN | 1 | 1 |
| State | Enabled | Enabled |
| BOOTP Relay | Disabled | Disabled |
| Interface Number | 2 | 2 |
| IP Address | 10.0.0.10 | 10.0.0.20 |
| IP Subnet Mask | 255.255.255.0 | 255.255.255.0 |
| VLAN | 2 | 2 |
| State | Enabled | Enabled |
| BOOTP Relay | Disabled | Disabled |
| Interface Number | 3 | 3 |
| IP Address | 10.0.0.11 | 10.0.0.21 |
| IP Subnet Mask | 255.255.255.255 | 255.255.255.255 |
| VLAN | 2 | 2 |
| State | Enabled | Enabled |
| BOOTP Relay | Disabled | Disabled |

- 4 Configure routes for each of the IP interfaces using the VPN devices as gateways. Configure one static route for each VPN device being load-balanced. See [“Configuring static routes” on page 148](#).

Table 111 VPN example—public network IP static routes settings

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| Static Route | [1 - 128] |
| Destination IP Address | 20.0.0.10 |

Table 111 VPN example—public network IP static routes settings (continued)

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 10.0.0.101 |
| IP Interface | 2 |
| Static Route | [1 - 128] |
| Destination IP Address | 20.0.0.11 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 10.0.0.102 |
| IP Interface | 3 |
| Static Route | [1 - 128] |
| Destination IP Address | 20.0.0.20 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 10.0.0.101 |
| IP Interface | 2 |
| Static Route | [1 - 128] |
| Destination IP Address | 20.0.0.21 |
| IP Subnet Mask | 255.255.255.255 |
| Gateway IP Address | 10.0.0.102 |
| IP Interface | 3 |

- 5 Enable virtual routing on WSM-3 and WSM-4. See [“Enabling virtual routing on the WSM” on page 151](#).
- 6 Configure VRRP for virtual routers 1 and 2 using the settings in [Table 112](#). See [“Configuring a virtual router” on page 153](#).

Table 112 VPN example—public network virtual router settings

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|------------------------------|-----------------------------|
| Virtual Router Number | 1 |
| Virtual Router ID | 1 |
| IP Address | 192.168.10.50 |

Table 112 VPN example—public network virtual router settings (continued)

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|------------------------------|-----------------------------|
| IP Interface | 1 |
| State | Enabled |
| Priority | 101 |
| Load Sharing | Disabled |
| Track VRs | Enabled |
| Track VLAN Ports | Enabled |
| Virtual Router Number | 2 |
| Virtual Router ID | 2 |
| IP Address | 10.0.0.1 |
| IP Interface | 2 |
| State | Enabled |
| Priority | 101 |
| Load Sharing | Disabled |
| Track VRs | Enabled |
| Track VLAN Ports | Enabled |

- 7 Enable Server Load Balancing (SLB) on WSM-3 and WSM-4. See [“Enabling or disabling server load balancing”](#) on page 243.
- 8 Configure real servers for health checking VPN devices using the settings in [Table 113](#). See [“Configuring each real server”](#) on page 216.

Table 113 VPN example—public network real server settings

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| Real Server | 1 |
| IP Address | 20.0.0.10 |
| State | Enabled |
| Name | VPN Server 1 |
| Real Server | 2 |
| IP Address | 20.0.0.11 |

Table 113 VPN example—public network real server settings (continued)

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| State | Enabled |
| Name | VPN Server 2 |
| Real Server | 3 |
| IP Address | 20.0.0.20 |
| State | Enabled |
| Name | VPN Server 3 |
| Real Server | 4 |
| IP Address | 20.0.0.21 |
| State | Enabled |
| Name | VPN Server 4 |

- 9 Configure real server group 1 using the hash metric. Add real servers 1, 2, 3, and 4 to the group. See [“Configuring a real server group” on page 225](#).

Table 114 VPN example—public network real server group settings

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| Group | 1 |
| Name | VPN Group |
| Metric | hash |
| Real Servers | 1, 2, 3, 4 |

- 10 Configure filters that allow local subnet traffic on the public side of the VPN device to reach the VPN device interfaces. Use the settings in [Table 115](#). See [“Creating a new filter” on page 201](#).

Table 115 VPN example—public network filter settings

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| Index | 100 |
| Name | VPN |

Table 115 VPN example—public network filter settings (continued)

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| Filter | Enabled |
| Action | Allow |
| Source IP Address | Any |
| Destination IP Address | 192.168.10.0 |
| Destination IP Mask | 255.255.255.0 |
| Index | 110 |
| Name | VPN |
| Filter | Enabled |
| Action | Allow |
| Source IP Address | Any |
| Destination IP Address | 224.0.0.0 |
| Destination IP Mask | 255.0.0.0 |

- 11** Create a filter to allow the management firewall (Policy Server) to reach the VPN firewall using the settings in [Table 116](#).

Table 116 VPN example—public network filter settings

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| Index | 120 |
| Name | VPN |
| Filter | Enabled |
| Action | Allow |
| Source IP Address | 192.168.10.120 |
| Source IP Mask | 255.255.255.255 |
| Destination IP Address | 10.0.0.0 |
| Destination IP Mask | 255.255.255.0 |

- 12** Create the redirection filter and enable firewall load balancing using the settings in [Table 117](#). This filter will redirect inbound traffic among the defined real servers in the group.

Table 117 VPN example—public network filter settings

| Field in Device Manager | Setting for WSM-3 and WSM-4 |
|-------------------------|-----------------------------|
| Index | 224 |
| Name | VPN Redirect FWLB |
| Filter | Enabled |
| Action | Redirect |
| Source IP Address | any |
| Destination IP Address | any |
| Protocol | any |
| Redirection Port | 1 |
| Firewall Redirect Hash | Enabled |

- 13** Enable filter processing on the ingress port (port 1). See [“Enabling or disabling filtering on a port” on page 89](#).
- 14** Add filters to the ingress port (port 1). For more information, see [“Applying filters to a port” on page 90](#).

Table 118 VPN example—configuring filter processing on ingress port

| Field in Device Manager | Setting for WSM-1 |
|-------------------------|-------------------|
| Port | 1 |
| Filtering | Enabled |
| Filters Applied | 100, 110, 224 |

- 15** Apply and save the configuration, and reboot the WSM.

Monitoring VPN configurations

You can monitor the following statistics for VPN configurations:

- [“Route statistics” on page 551](#)
- [“Filter statistics” on page 575](#)
- [“Real server port statistics” on page 566](#)
- [“Real servers statistics” on page 567](#)
- [“Groups statistics” on page 568](#)
- [“Virtual servers statistics” on page 569](#)
- [“Real servers state data” on page 571](#)
- [“Maintenance statistics” on page 572](#)
- [“Redirect statistics” on page 577](#)
- [“Virtual routing statistics” on page 563](#)

Chapter 14

Global server load balancing

Global server load balancing (GSLB) lets you balance server traffic across multiple physical sites. Remote sites exchange state information using the Distributed Site State Protocol (DSSP). Load balancing decisions are then based on the health, response time, and geographic location of each site at a given point in time. GSLB supports all IP protocols.

This section contains the following topics:

- [“How GSLB works,”](#) next
- [“Configuring GSLB”](#) on page 375
- [“IP proxy for non-HTTP redirects”](#) on page 386
- [“About GSLB network preferences”](#) on page 390

How GSLB works

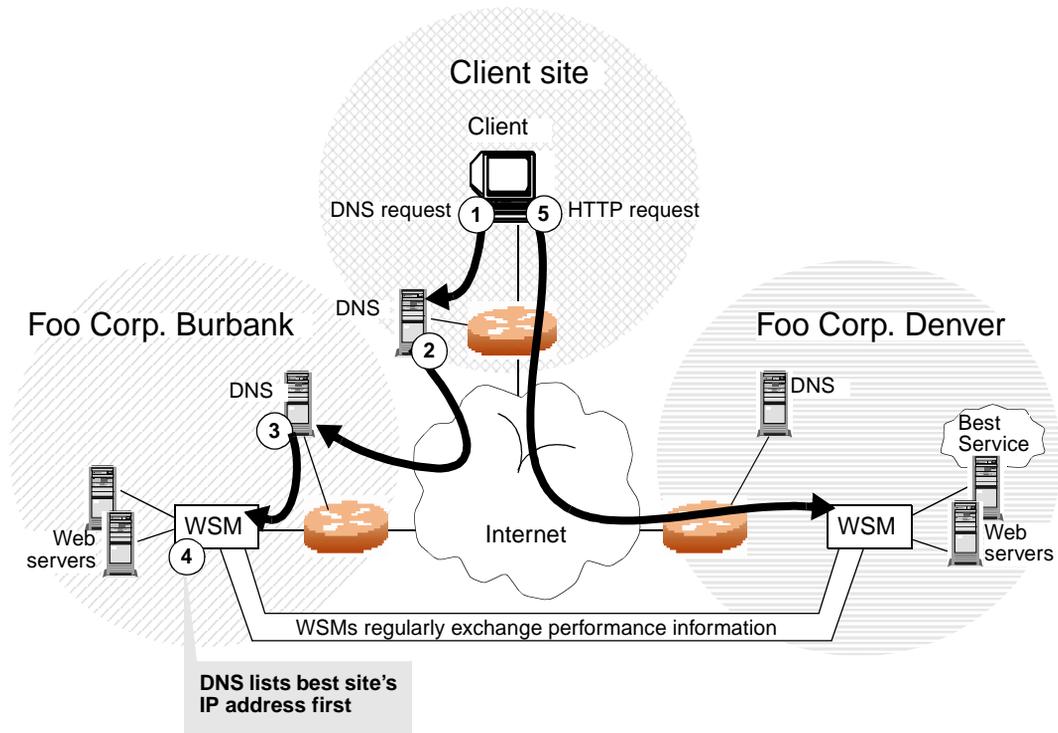
GSLB is based on the following:

- Domain Name System (DNS)
- Site proximity by source IP address

[Figure 145](#) shows a client using a browser to view the Web site for the Foo Corporation at *www.foocorp.com*.

Figure 145 DNS resolution with global server load balancing

The Foo Corporation has two Web sites with WSMs configured for GSLB, and identical content and services—one in Burbank, and one in Denver. The WSMs are also configured as the Authoritative Name Servers for *www.foocorp.com*.



The following steps describe DNS resolution in [Figure 145](#):

- 1 The client Web browser requests the *www.foocorp.com* IP address from the local DNS.
- 2 Client's DNS asks its upstream DNS, which in turn asks the next, and so on, until the address is resolved. The request reaches either:
 - An upstream DNS server with the requested IP address information on hand, or
 - A Foo Corporation DNS server
- 3 The Foo Corporation's Burbank DNS has been configured to use the local WSM as the authoritative name server for *www.foocorp.com*.
- 4 The Burbank WSM responds to the DNS request with the current best service IP address.

Since each WSM regularly checks and communicates health and performance information with its peers, each can determine which site(s) are best able to serve the client's needs. Each can respond with a list of IP addresses for the Foo Corporation's distributed sites, which are prioritized by performance, geography, and other criteria. In this case, the Burbank WSM knows that Denver currently provides better service, and lists Denver's virtual server IP address first when responding to the DNS request.

- 5 The client connects to Denver for the best service.

The client's Web browser uses IP address information obtained from the DNS request to open a connection to the best available site. The IP addresses represent virtual servers which are locally load-balanced according to regular SLB configuration. If the site serving the client HTTP content suddenly experiences a failure (no healthy real servers) or becomes overloaded with traffic (all real servers reach their maximum connection limit), the WSM issues an HTTP Redirect and transparently causes the client to connect to another peer site.

Configuring GSLB

To configure GSLB:

- 1 Log into the 8600 with rwa access, and connect to the WSM. For more information, see the *Installing the Web Switching Module for the 8600 series switch*, Part number 314969-A
- 2 Activate the GSLB software key. For more information, see the *Web OS Switch Software 10.0 Command Reference*, Part number 212778-A.
- 3 Configure the WSM IP interface. See [“Configuring IP interfaces” on page 129](#).
- 4 Configure the default gateways. See [“Configuring default gateways” on page 132](#).
- 5 Configure the WSM at each site for local SLB. See [“Configuring virtual server load balancing” on page 215](#).
- 6 Configure each WSM so that it recognizes its remote peers. See [“Defining remote GSLB peers for the local site,”](#) next.

- 7 Configure the WSM at each site to act as the DNS server for each service that is hosted on its virtual servers. Also, configure the local DNS server to recognize the WSM as the authoritative DNS server for the hosted services. See [“Configuring remote GSLB services for the local site” on page 378](#).
- 8 Enable GSLB for each site. See [“Enabling GSLB” on page 382](#).

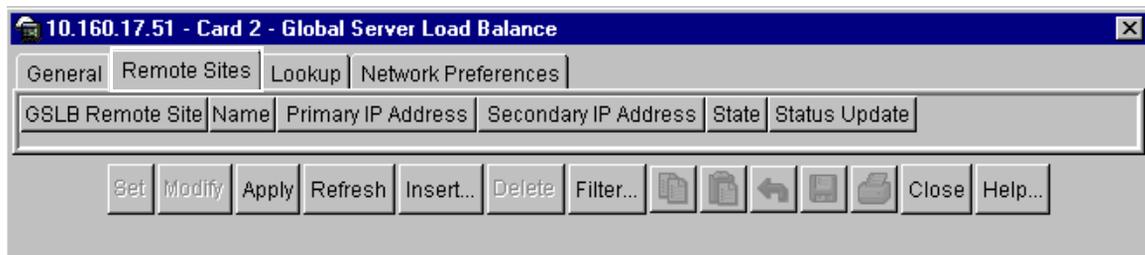
Defining remote GSLB peers for the local site

For each site in your GSLB configuration, you must add and enable the IP addresses of all remote sites. You can enable up to 64 remote sites with a total of 2048 service/site combinations.

To configure a remote GSLB site:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > Global SLB.
The Global Server Load Balance dialog box opens to the [General tab](#).
- 3 Click the Remote Sites tab.
The [Remote Sites tab](#) (Figure 146) opens with the fields defined in [Table 119 on page 378](#).

Figure 146 Global Server Load Balance—Remote Sites tab



- 4 Click Insert.
The [Insert Remote Sites tab](#) (Figure 147) opens with the fields defined in [Table 119 on page 378](#).

Figure 147 Insert Remote Sites dialog box

- 5 In the GSLB Remote Site field, type a site number (1 - 64).
- 6 In the Name field, type the remote site's name (up to 31 characters).
- 7 In the Primary Address field, type the IP address of the remote interface.
- 8 (optional) In the Secondary IP Address field, type the IP address of the remote site's redundant WSM, if equipped.
- 9 In the State field, click Enabled.
- 10 In the Status Update field, click Enabled.



Note: If status updates are enabled, the remote site sends periodic updates to all peer sites. If your local firewall does not permit this traffic, disable updates.

- 11 Click Insert.
- 12 The Insert Remote Sites dialog box closes and the remote site is added to the Remote Sites tab.
- 13 To add another remote site, repeat steps 4 - 12.
- 14 From the Remote Sites tab, click Apply > Close.

The remote site is configured and the GSLB dialog box closes.

Table 119 describes the GSLB remote sites fields.

Table 119 Global SLB—Remote Sites fields

| Field | Description |
|----------------------|---|
| Number | The index number (1 - 64) of the GSLB remote sites table. |
| Name | Sets the Name of the GSLB remote site. The maximum string length is 31 characters. |
| Primary IP Address | Sets the IP interface IP address of the GSLB remote site primary WSM. Use dotted decimal notation. The default is 0.0.0.0. |
| Secondary IP Address | For remote sites configured with a redundant WSM, sets the IP address of the IP interface for the remote secondary WSM. If the remote site primary WSM fails, the local WSM will address the remote site secondary WSM instead. The default is 0.0.0.0. |
| State | Enables this remote site for use with Global Server Load Balancing. |
| Status Update | Enables or disables remote site updates. If enabled (default), this site will send regular Distributed Site State Protocol (DSSP) updates to its remote peers using HTTP port 80. If disabled, the WSM will not send state updates. If your local firewall does not permit this traffic, disable the updates. |

Configuring remote GSLB services for the local site

Each GSLB local site must be configured to recognize the services offered at the remote sites. To do this, configure one real server on the local WSM for each virtual server located at each remote site.

To configure the remote GSLB services:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3 Click the Real Servers tab.
The Real Servers tab opens.

4 Click Insert

The **Insert Real Servers** (Figure 148) dialog box opens with the fields listed in Table 60 on page 219.

Figure 148 Configuring real servers for GSLB

5 In the Real Server field, type a real server number (1 - 1023).

6 In the IP address field, type the IP address of the remote virtual server.



Note: Do not confuse the IP address of the remote virtual server with that of the IP interface address on the remote WSM.

7 In the State field, click Enabled.

- 8 In the Health Check Interval field, type an interval of 30 - 60 seconds to avoid generating excess traffic across the Internet.
- 9 In the Type field, click Remote-Server.
- 10 Click Insert.

The server is inserted into the Real Servers tab and the Insert Real Servers dialog box closes.

- 11 From the SLB dialog box, click the Real Groups tab.

The Real Groups tab opens.

- 12 Select the group associated with the intended virtual server service.

The group is highlighted.

- 13 Click Servers.

The [Select Real Servers to include in Group dialog box](#) (Figure 149) opens with the fields defined in [Table 63 on page 230](#).

Figure 149 Configuring GSLB real server membership



- 14 CTRL + Click the real server you want to add to the group.

The row is highlighted.



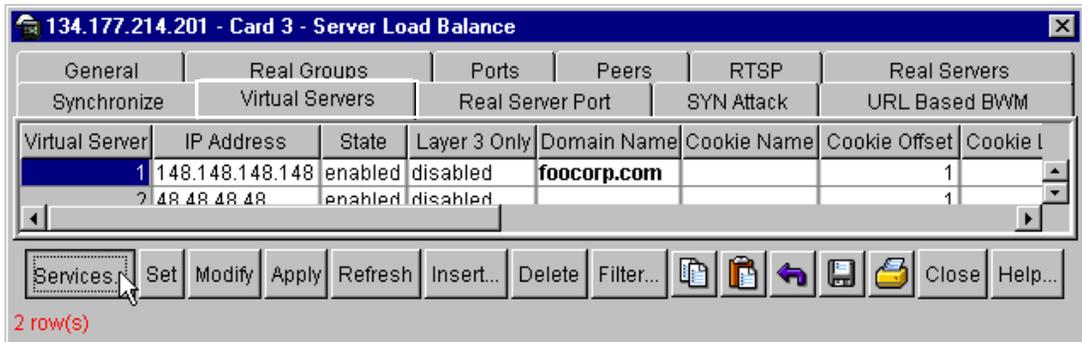
Caution: You must use CTRL + Click to preserve existing server membership. If you just choose the new server and click Modify, only the new server is added, and existing server members are removed.

15 Click Modify.

The real server is added to the group in the Real Groups tab, and the Select Real Servers to include in Group dialog box closes.

16 From the SLB dialog box, click the Virtual Servers tab.

The Virtual Servers tab opens.

Figure 150 Configuring GSLB virtual servers**17** Double-click the Domain Name field for the intended virtual server and enter a domain name (such as foocorp.com).**18** Click Services.

The Virtual Server Services dialog box opens.

Figure 151 Configuring GSLB server services**19** Double-click the Host Name field and enter the host name for the service (such as www).**20** Click Set > Apply > Close

The host name is added and the Virtual Server Services dialog box closes.

21 Repeat steps 17 - 20 to define the domain name and host name for each service hosted on each virtual server.

22 From the SLB dialog box, click Set > Apply > Close.

Remote GSLB services are configured and the SLB dialog box closes.

Enabling GSLB

GSLB must be enabled at each site.

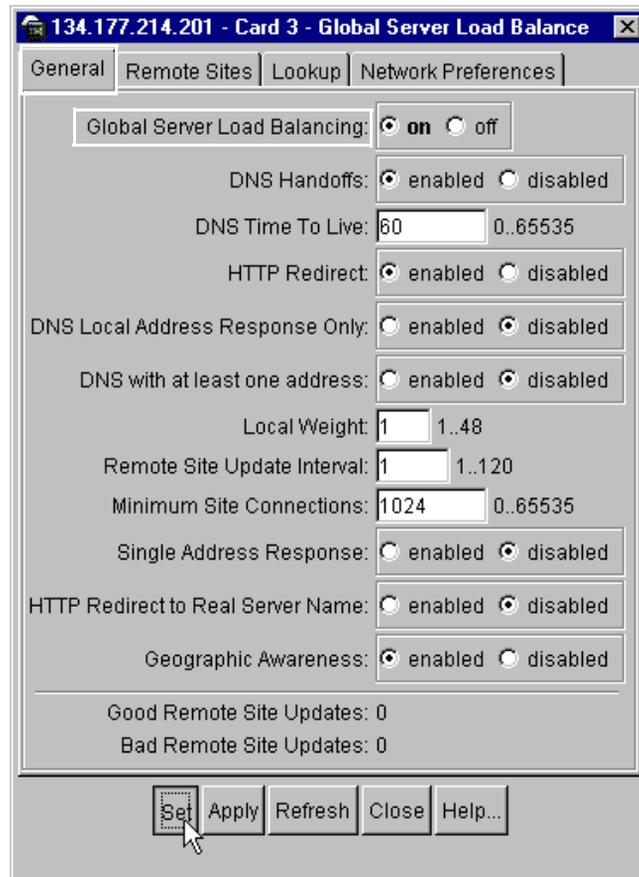
To enable GSLB:

1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

2 From the Edit menu, choose WSM Card > L4 Switching > Global SLB.

The Global Server Load Balance dialog box opens to the [General tab \(Figure 152\)](#) with the fields defined in [Table 120 on page 384](#).

Figure 152 Global SLB—General tab

3 In the Global Server Load Balancing field, click On.

4 Click Set > Apply > Close.

GSLB is enabled for the local site and the GSLB dialog box closes.

Table 120 describes the fields on the Global SLB—General tab.

Table 120 Global SLB—General tab fields

| Field | Description |
|---------------------------------|--|
| Global Server Load Balancing | <p>Sets the state (on or off) of global SLB. The default is off.</p> <p>Activates Global Server Load Balancing (GSLB) for this WSM. This option can be performed only after the optional GSLB software is activated (For more information, see “Activating Optional Software” in <i>Web OS Switch Software 10.0 Command Reference</i>, Part number 212778-A).</p> <p>If you turn GSLB off for this WSM, any active remote sites will still perform GSLB services with each other, but will not hand off requests to this WSM.</p> |
| DNS Handoffs | <p>Enables or disables DNS handoffs to peer sites by this WSM. This should be enabled for proper GSLB operation. If disabled, whenever the WSM receives a DNS request for a configured service, it will respond only with its own virtual server IP address, regardless of performance or load considerations. This option is enabled by default.</p> |
| DNS Time to Live | <p>Sets the duration in seconds (0 to 65535, with default at 60) that the DNS response from the WSM (indicating site of best service) will remain in the cache of DNS servers. A lower value may increase the ability of the GSLB system to adjust to sudden changes in traffic load, but will generate more DNS traffic. Higher numbers may reduce the amount of DNS traffic, but may slow GSLB’s response to sudden traffic changes.</p> |
| HTTP Redirect | <p>Enables or disables HTTP redirects to peer sites by this WSM. When enabled (default), this WSM will redirect client requests to peer sites if its own real servers fail or have reached their maximum connection limits. If disabled, the WSM will not perform HTTP Redirects, but will instead drop requests for new connections and cause the client’s browser to eventually issue a new DNS request.</p> |
| DNS Local Address Response Only | <p>Enables or disables WSM responses to DNS queries with local virtual server IP addresses. This option is disabled by default. When enabled, the WSM will always respond to DNS queries by providing a local virtual server IP address, as long as the virtual server IP address has healthy real servers with an aggregate number of available connections equal to the total from each server’s configured maximum connections, minus the server’s current number of connections. When the real servers for the local virtual server IP addresses are unavailable or saturated, the WSM will respond to DNS requests using normal GSLB rules.</p> |

Table 120 Global SLB—General tab fields (continued)

| Field | Description |
|-----------------------------------|---|
| DNS with at Least One Address | Enables or disables DNS responses (with) at least one address. At least one IP address is included in each DNS response. Even if all remote sites cannot handle another request, the local VIP is returned in DNS response to eliminate long DNS timeouts caused by an empty response. This option is disabled by default. |
| Local Weight | Sets the local weight (1 - 48). The default is 1. The higher the weight value, the more connections that will be directed to the local site. The response time of this site is divided by this weight before the best site is assigned to a client. Remote site response times are divided by the real server weight before selection occurs. |
| Remote Site Update Interval | Sets the time in minutes (1 - 120) between Distributed Site State Protocol (DSSP) updates between this WSM and its peers. The default is 1 minute. |
| Minimum Site Connections | Sets the minimum number (0 - 65,535) of available site connections. The default is set at 1024. If the site's available sessions fall below this value, traffic won't be redirected to the site. A site is not eligible for more requests (such as DNS or HTTP redirects) once the number of available connections at a site drops below this threshold. |
| Single Address Response | Enables or disables DNS responses with only one address. At most one IP address is included in each DNS response. This option is disabled by default. |
| HTTP Redirect to Real Server Name | Enables or disables an HTTP redirect to a real server name. When a site redirects a client to another site using an HTTP redirect, the client is redirected to the new site's IP address. This option is disabled by default. If enabled, the client will be redirected to the domain name specified by the remote real server name plus virtual server domain name: <i><remote real server name>.<virtual server domain name></i> |
| Geographic Awareness | Enables or disables geographic awareness, such as the IANA table. This option is enabled by default. If this option is disabled, all clients and sites will be assumed to exist in the same geographic region, allowing all sites to be eligible for each client. |
| Good Remote Site Updates | Displays the number of good remote site updates that were received. |
| Bad Remote Site Updates | Displays the number of bad remote site updates that were received. |

IP proxy for non-HTTP redirects

The HTTP redirect function uses the best response and least load for the requested content to redirect requests to an alternate site. However, if the client requests a non-HTTP application, such as FTP, POP3, or SMTP, and resources are unavailable at the first site, then you can configure a proxy IP address on the client port to redirect requests.



Note: This feature should be a last resort in topologies where remote servers are usually virtual server IP addresses in other WSMs.

Figure 153 illustrates the packet-flow of HTTP and non-HTTP redirects in a GSLB environment.

Figure 153 HTTP and non-HTTP redirects

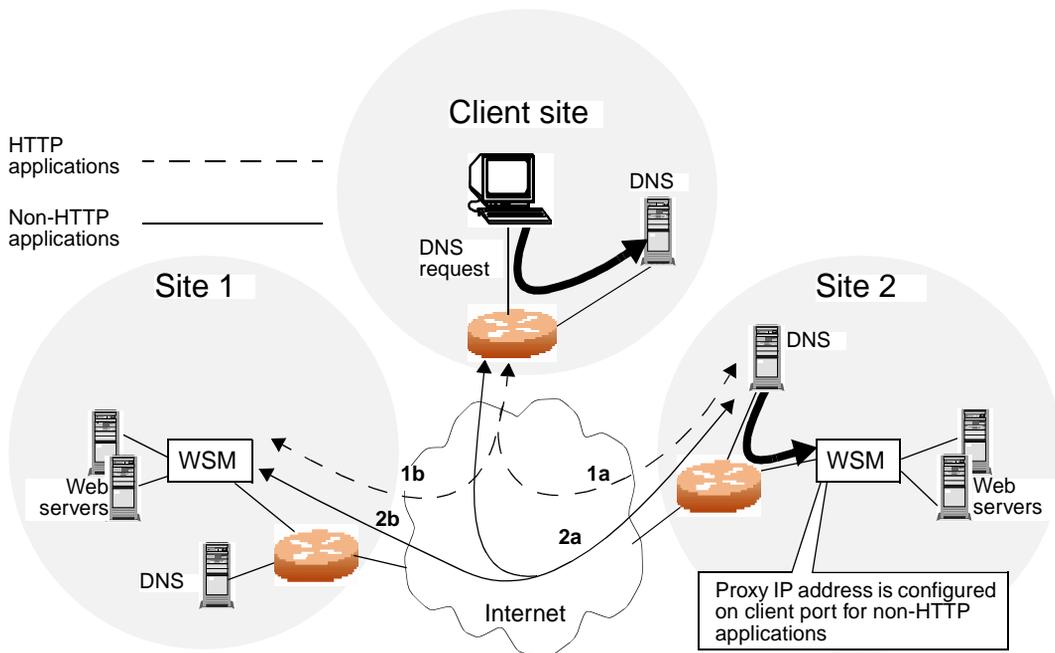


Table 121 describes the redirects shown in Figure 153.

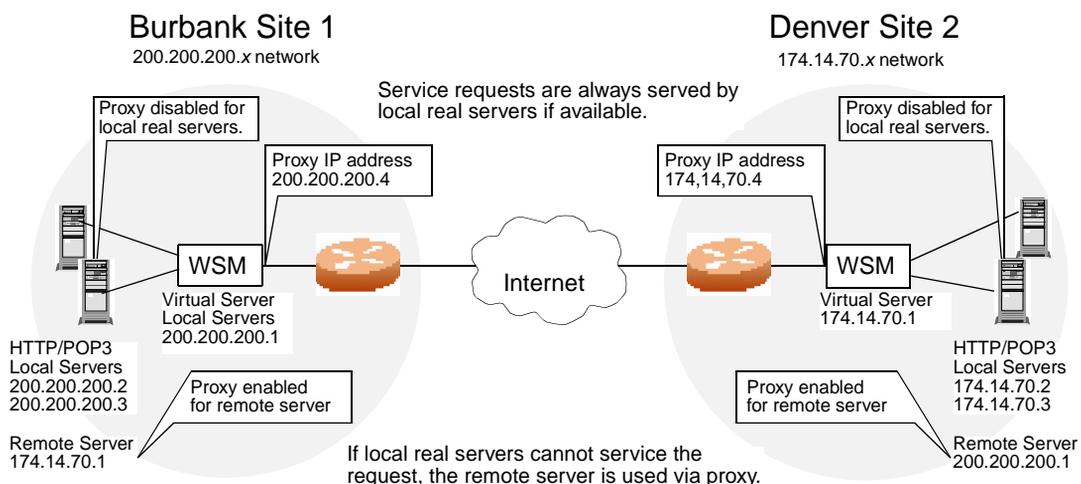
Table 121 HTTP versus non-HTTP redirects

| Application | Site 2 WSM | Site 1 WSM |
|-------------|--|--|
| HTTP | 1a. Client DNS request reaches Site 2, where resources are unavailable. Site 2 sends response to client with Site 1's virtual server IP address. | 1b. Client re-sends request to site 1, where resources are available. Site 1 completes TCP three-way handshake with client. |
| Non-HTTP | 2a. Client DNS request reaches Site 2, where resources are unavailable. Site 2 sends a request to Site 1 with Site 2's proxy IP address as the source IP address and the virtual server IP address of Site 1 as the destination IP address. | 2b. Site 1 processes the client proxy IP request. Resources are available at Site 1. Site 1 returns request to proxy IP port on Site 2. Site 2 completes the three-way handshake with the client. |

How IP proxy works

Figure 154 shows two GSLB sites—one in Burbank and one in Denver—load-balancing HTTP and POP3. Any request that cannot be serviced locally is sent to the peer site. HTTP requests are sent to the peer site using HTTP Redirect. Other application requests are sent to the peer site using IP proxy.

Figure 154 POP3 request fulfilled using IP Proxy



The procedure outlined below explains the three-way handshake between the two sites and the client for a non-HTTP application (POP3). When operator error at site 1 terminates POP3 processes, POP3 requests are fulfilled as follows:

- 1 The site 1 virtual server received a POP3 TCP SYN request from a user. The WSM determines there are no local resources to handle the request.
- 2 The site 1 WSM rewrites the request to contain a client proxy IP address as the source IP address, and the Site 2 virtual server IP address as the destination IP address.
- 3 The site 2 virtual server receives the POP3 TCP SYN request. The request looks like a normal SYN frame, so the WSM performs normal local load-balancing.
- 4 The site 2 WSM and real servers exchange information. The TCP SYN ACK from Site 2's local real server is sent back to the IP address specified by the proxy IP address.
- 5 The site 2 WSM sends the TCP SYN ACK frame to Site 1, with Site 2's virtual server IP address as the source IP address and Site 1's proxy IP address as the destination IP address.
- 6 The site 1 WSM receives the frame and translates it, using Site 1's virtual server IP address as the source IP address and the client's IP address as the destination IP address.

This cycle continues for the remaining frames to transmit all the client's mail, until a FIN frame is received.

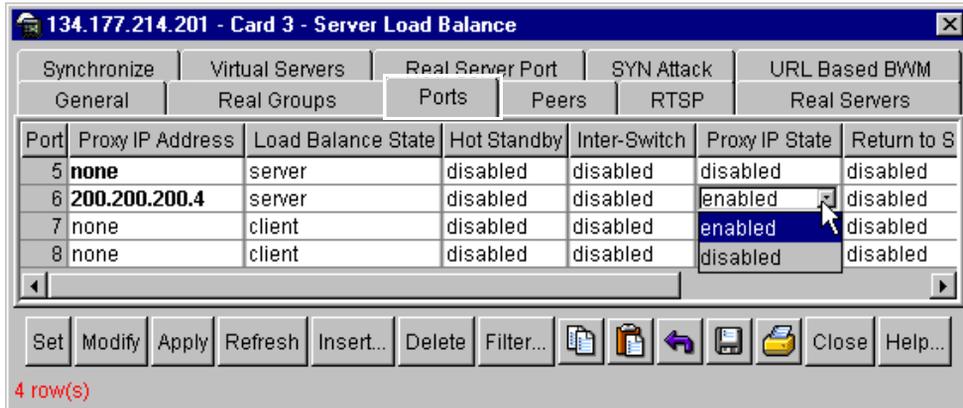
Configuring IP proxy for non-HTTP redirects

To configure IP proxy for non-HTTP redirects:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3 Click the Ports tab.

The **Ports** tab (Figure 155) opens with the fields defined in Table 66 on page 242.

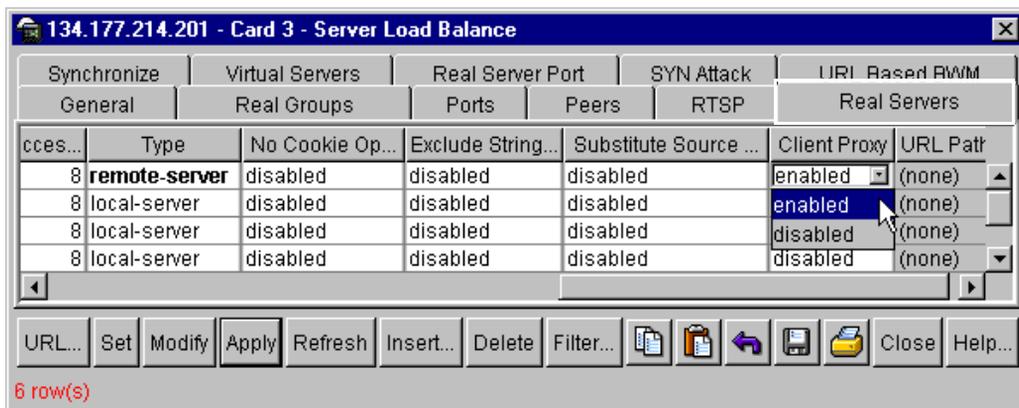
Figure 155 Port configuration—IP proxy address for non-HTTP redirects



- 4 For the port you want to proxy, double-click the Proxy IP Address field, and enter the proxy IP address.
- 5 For the port you want to proxy, double-click the Proxy IP State field and choose enabled from the drop-down list.
- 6 Click the Real Servers tab.

The Real Servers tab (Figure 156) opens with the fields defined in Table 60 on page 219.

Figure 156 Real server configuration—IP proxy for HTTP redirects



- 7 For local servers, double-click the Client Proxy field, and choose Disabled from the drop-down list.
- 8 For remote servers, double click the Client Proxy field and choose Enabled from the drop-down list.
- 9 Click Set > Apply > Close.

The ports and servers are configured and the Server Load Balance dialog box closes.

About GSLB network preferences

The Internet Assigned Numbers Authority (IANA), the central coordinator for the assignment of unique parameter values for Internet protocols, does not provide sufficient geographic separation of client proximity information. Because of this, large ISPs cannot use their own geographic data to determine GSLB site selection based on client location. However, the WSM lets you overwrite the IANA table and configure client-to-site mapping to select GSLB sites based on client location. The WSM proximity database is limited to 128 entries and is supported for HTTP protocol.

Configuring GSLB network preferences

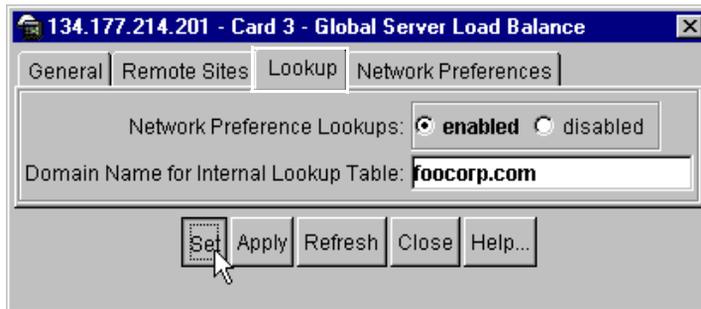
You can configure network preferences for a single domain only. For each client subnet you can configure up to two preferred sites. The WSM forwards the client request based on the minimum available sessions and response time between the two preferred sites.

To configure network preferences:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > Global SLB.
The Global Server Load Balance dialog box opens to the [General tab](#).
- 3 Click the Lookup tab.

The **Lookup** tab (Figure 157) opens with the fields defined in Table 122 on page 393.

Figure 157 Global Server Load Balance—Lookup tab



- 4 In the Network Preference Lookups field, click Enabled.
- 5 In the Domain Name for Internal Lookup Table, type a domain name.

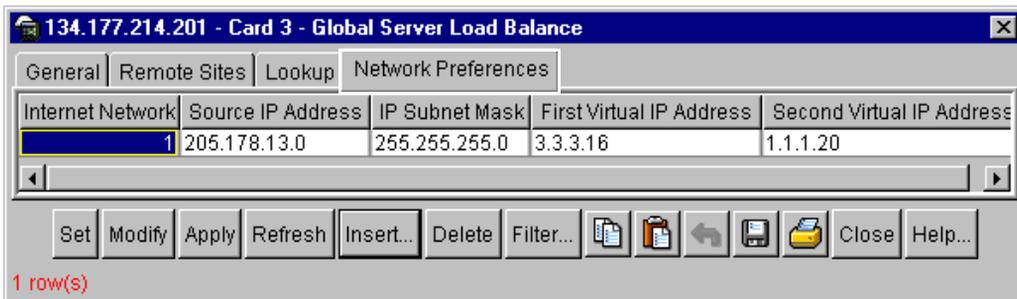


Note: You can configure network preferences for a single domain only.

- 6 Click Set > Apply.
- 7 Click the Network Preferences tab.

The Network Preferences tab opens.

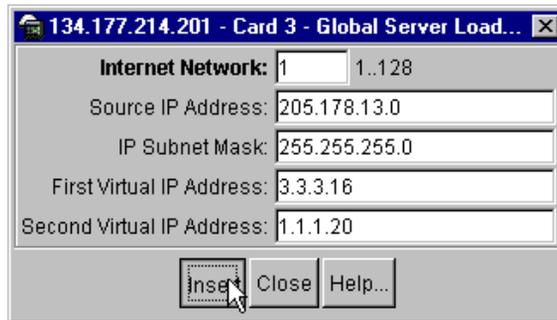
Figure 158 Global Server Load Balance—Network Preferences tab



- 8 Click Insert.

The Insert Network Preferences dialog box opens.

Figure 159 Insert Network Preferences dialog box



- 9 In the Internet Network field, type a number (1 - 128) for this network.
- 10 In the Source IP Address field, type the client source IP address for this client-to-site mapping.
- 11 In the IP Subnet Mask field, type a subnet mask.
- 12 In the First Virtual IP Address field, type the IP address of the first choice virtual server for this client-to-site mapping.
- 13 In the Second Virtual Address field, type the IP address of the second choice virtual server for this client-to-site mapping.
- 14 Click Insert.

The client-to-site mapping appears in the Network Preferences tab and the Insert Network Preferences dialog box closes.

- 15 From the Network Preferences tab, click Apply > Close.

The lookup table is updated and the GSLB dialog box closes.

[Table 122](#) describes the Global SLB—Lookup tab fields.

Table 122 Global SLB—Lookup tab fields

| Field | Description |
|---------------------------------------|---|
| Network Preference Lookups | Enables or disables network preference lookups. If enabled, the WSM responds to DNS requests based on the configured domain name and network preferences . For information about network preferences, see Table 123 on page 393 . |
| Domain Name for Internal Lookup Table | Sets a domain name of up to 34 characters for the internal lookup table. The default is none. |

[Table 123](#) describes the fields on the Global SLB—Network Preferences tab and the Insert Network Preferences dialog box.

Table 123 Global SLB—Network Preferences fields

| Field | Description |
|---------------------------|---|
| Number | The index number of the Internet network preference table: 1 to 128. |
| Source IP Address | Sets the source IP address. The default is 0.0.0.0. Specify an IP address in dotted decimal notation, or “any”. A range of IP addresses is produced when used with the mask option, below. |
| IP Subnet Mask | Sets the IP Subnet mask of the network table. The default is 255.255.255.0. This IP address mask is used with the source IP SIP address to find a correct virtual server IP address to respond to a DNS request. |
| First Virtual IP Address | Sets the IP address of the primary virtual server. The default is 0.0.0.0. The address can either be a local or remote virtual server. The WSM returns the VIP address with the least response time that is over the minimum number of available connections. |
| Second Virtual IP Address | Sets the IP address of the secondary virtual server. The default is 0.0.0.0. |

Chapter 15

Application Redirection

Application Redirection improves speed of access to common Web or application content by using filters to redirect traffic to cache and application servers.

This section includes the following topics:

- [“About Application Redirection” on page 395](#)
- [“Configuring Web Cache Redirection” on page 396](#)
- [“Configuring delayed binding for Web Cache Redirection” on page 398](#)
- [“Configuring RTSP for Web Cache Redirection” on page 399](#)
- [“Configuring IP proxies with Application Redirection” on page 401](#)
- [“Excluding noncacheable sites” on page 402](#)

About Application Redirection

Much of the information downloaded from the Internet is not unique, and clients often access a Web page many times for additional information or to explore other links. When duplicate information is requested, the Web site’s pictures, buttons, frames, text, and so on are reloaded from page to page. Redundant requests can consume a large amount of your available bandwidth to the Internet.

Application redirection helps reduce the traffic congestion during peak loads. When WSM application redirection filters are properly configured, outbound client requests for Internet data are intercepted and redirected to a group of application or Web cache servers on your network. The servers duplicate and store

inbound Internet data that has been requested by your clients. If the servers recognize a client's outbound request as one that can be filled with cached information, the servers supply the information rather than send the request across the Internet.



Note: WAN Link load balancing and NAT Web Cache Redirection cannot be configured on the same WSM.

Configuring Web Cache Redirection

To configure Web Cache Redirection:

- 1 Install transparent Web cache software on all three Web cache servers.
- 2 Define an IP interface on the WSM. See [“Defining an IP interface on the WSM” on page 225](#).

Since, by default, the WSM only remaps destination MAC addresses, it must have an IP interface on the same subnet as the three Web cache servers.

- 3 Define each Web cache real server. See [“Configuring each real server” on page 216](#). Set the following parameters:
 - Real Server = *<assign a real server number, 1-1024>*
 - IP Address = *<actual IP Address of the Web cache server>*
 - State = Enabled



Note: The Web cache real servers must be in the same VLAN and must have an IP route to the WSM that performs the WCR. In addition, the path from the WSM to the real servers must not contain a router. The router stops HTTP requests from reaching the Web cache servers and, instead, directs them back out to the Internet.

- 4 Define a real server group and add the Web cache servers to it. See [“Configuring a real server group” on page 225](#). Set the following parameters:
 - Metric = minmisses
 - Real Servers = *<Web cache real server numbers>*

- 5 Configure the ports supporting application redirection. See [“Configuring ports for server load balancing” on page 240](#). Set the following parameter:
- Load Balance State = None (this is the default setting)



Note: Do not enable server processing on a port with Application Redirection enabled. Server processing is used only with SLB.

- 6 Create a filter that will intercept and redirect all client HTTP requests. See [“Creating a new filter” on page 201](#). Set the parameters in the [Table 124](#).

The filter must intercept all TCP traffic for the HTTP destination port and redirect it to a real server TCP or UDP port. This port is called the redirection port, and is used when performing Layer 4 health checks of TCP services.

Table 124 Web Cache Redirection filter settings

| Filter field | Setting |
|------------------------|--|
| Index | <Filter index number> |
| Name | Web Cache Redirection |
| Filter | Enabled |
| Action | Redirect |
| Source IP Address | Any |
| Destination IP Address | Any |
| Protocol | TCP |
| Source Port | Any |
| Destination Port | <To an HTTP destination port> |
| Redirection Port | <The redirection port number> |
| Redirection Group | <The Web Cache Redirection Real Server Group number> |



Note: If the transparent proxy occurs on the WSM, make sure to use the service port required by the specific software package. For more about IP proxy addresses, see [“Configuring IP proxies with Application Redirection” on page 401](#).

- 7 Create a default filter to allow noncached traffic to proceed normally. See [“Configuring a default filter” on page 192](#). Set the parameters in [Table 125](#).

Table 125 Default filter settings for noncached traffic

| Filter field | Setting |
|------------------------|-----------------------|
| Index | <Filter index number> |
| Name | Default |
| Filter | Enabled |
| Action | Allow |
| Source IP Address | Any |
| Destination IP Address | Any |
| Protocol | Any |



Note: When the filter’s Protocol field is not set to TCP or UDP, then its Source Port and Destination Port are ignored.

- 8 Configure the filters on the client ports. See [“Enabling or disabling filtering on a port” on page 89](#) and [“Applying filters to a port” on page 90](#).
- 9 Enable SLB. See [“Enabling or disabling server load balancing” on page 243](#). SLB must be enabled for application redirection to work properly.

Configuring delayed binding for Web Cache Redirection

To configure delayed binding for Web Cache Redirection only:

- ➔ Enable URL Redirection on the WCR filter. See [“Creating a new filter” on page 201](#).

For more information about delayed binding, see [“Using delayed binding to prevent DoS attacks” on page 262](#).

About RTSP WCR

RTSP WCR is similar to HTTP WCR. Several local caching servers are needed to ensure a high quality of data for multimedia presentations.

RTSP WCR redirects cached data transparently and balances the load among the cache servers. If there is no cache server, the request is directed to the origin server. Internet Service Providers use RTSP WCR to cache the multimedia data of a customer site locally. Since the requests for this data are directed to the local cache, they are served faster.

You can also configure certain URL content to be non-cacheable. The requests for non-cacheable URLs bypasses the cache server and is sent across the Internet to the origin server. The client packets are relayed to the server by using Layer 4 server load balancing.

Configuring RTSP for Web Cache Redirection

For information about configuring RTSP load balancing, see [“Real Time Streaming Protocol server load balancing” on page 283](#).

To configure RTSP for web cache redirection:

- 1 Configure an RTSP redirection filter to cache data and balance the load among the cache servers. See [“Creating a new filter” on page 201](#). Set the parameters in [Table 126](#).

Table 126 RTSP Web Cache Redirection filter settings

| Filter field | Setting |
|------------------------|-----------------------|
| Index | <Filter index number> |
| Name | RTSP WCR |
| Filter | Enabled |
| Action | Redirect |
| Source IP Address | Any |
| Destination IP Address | Any |
| Protocol | TCP/UDP |

Table 126 RTSP Web Cache Redirection filter settings (continued)

| Filter field | Setting |
|-------------------|-------------------------------------|
| Source Port | Any |
| Destination Port | RTSP |
| Redirection Port | <The redirection port number> |
| Redirection Group | <The RTSP Real Server Group number> |
| Client Proxy | Disable |

- 2 Configure a default filter to allow all noncached traffic to proceed normally. See [“Configuring a default filter” on page 192](#). Set the parameters in [Table 127](#).

Table 127 Default filter settings for noncached traffic

| Filter field | Setting |
|------------------------|-----------------------|
| Index | <Filter index number> |
| Name | Default |
| Filter | Enabled |
| Action | Allow |
| Source IP Address | Any |
| Destination IP Address | Any |
| Protocol | Any |



Note: When the filter’s Protocol field is not set to TCP or UDP, then its Source Port and Destination Port are ignored.

- 3 Configure the filters on the client ports. See [“Enabling or disabling filtering on a port” on page 89](#) and [“Applying filters to a port” on page 90](#).

Configuring IP proxies with Application Redirection

When used with Application Redirection, transparent proxies for NAT redirect client requests to servers located on any subnet, anywhere. They are transparent to the user and substitute all source and destination addresses, including remapping destination ports.



Note: Application redirection is automatically enabled when a Redirect filter is applied to a port.

This configuration assumes that Application Redirection is configured for your network.

To configure a proxy IP address for NAT:

- 1 Add proxy IP addresses to the SLB ports as follows. See [“Configuring a proxy IP address for a port” on page 247](#).
 - If VMA is not enabled, add proxy IP addresses to the redirection ports. Each port using a redirection filter requires a unique proxy IP address.
 - If VMA is enabled, add proxy IP addresses for all WSM ports (except port 9). For information about VMA, see [“Improving WSM performance with VMA” on page 323](#).

Set the following SLB port parameters.

- Port = *<port number>*
 - Proxy IP Address = *<proxy IP address>*
 - Proxy IP State = Enabled
- 2 Configure the application redirection filter to use the real server TCP or UDP port to which redirected traffic will be sent. See [“Creating a new filter” on page 201](#). Set the following parameter.
 - Redirection Port = *<the proxy redirection port number>*
 - 3 Enable proxies on the real servers. See [“Configuring each real server” on page 216](#). Set the following parameter.
 - Client Proxy = Enabled

- 4 Verify that traffic has been redirected. See [“Filter statistics” on page 575](#).

Excluding noncacheable sites

Some Web sites provide browser-based games or applications that keep real-time session information, or authenticate by client IP address— content which is not well-suited for redirection to cache servers.

To exclude noncacheable sites from redirection filters:

- 1 Create a filter to allow specific traffic to pass normally through the WSM. See [“Creating a new filter” on page 201](#). Set the parameters in [Table 128](#).

Table 128 Filter settings to exclude redirection of noncacheable site

| Filter field | Setting |
|------------------------|--|
| Index | <Filter index number> (Must be a lower filter number than the application redirection filter) |
| Name | Exclude <site name> |
| Filter | Enabled |
| Action | Allow |
| Source IP Address | Any |
| Destination IP Address | <IP address of Web site to redirect> |
| Destination Mask | 255.255.255.0 (entire subnet range) |
| Protocol | TCP |
| Source Port | Any |
| Destination Port | HTTP |
| Redirection Port | <The redirection port number> |

- 2 Add the filter to the necessary ports. See [“Enabling or disabling filtering on a port” on page 89](#) and [“Applying filters to a port” on page 90](#).

Chapter 16

Health Checking

This chapter describes the WSM health checks, how they affect the server and application port states, and how to configure health checks. This section includes the following topics:

- [“About health checking,”](#) next
- [“Configuring a health check for a real server group”](#) on page 408
- [“Modifying the health check interval and retries”](#) on page 410
- [“ICMP health checking”](#) on page 412
- [“ARP health checking”](#) on page 412
- [“Direct server return health checking”](#) on page 413
- [“Link health checking”](#) on page 415
- [“TCP health checking”](#) on page 415
- [“Script-based health checking”](#) on page 416
- [“Application health checking”](#) on page 417

About health checking

A WSM configured for Server Load Balancing (SLB) monitors the servers in the real server group and the load-balanced application(s) running on them. If a WSM detects that a server or application has failed, it will not direct any new connection requests to that server. When a service fails, a WSM can remove the individual service from the load-balancing algorithm without affecting the server’s other services.

The WSM checks the status of each service on each real server every two seconds (the default setting). Real servers that are busy processing connections may not respond to health checks. If a service does not respond to four consecutive health checks (the default), the WSM declares the service unavailable. When a server is

declared unavailable, the WSM continues its health checks. Following 8 successful health checks (the default), the server is returned to service. You can modify both the health check interval and the number of retries. For more information, see [“Modifying the health check interval and retries” on page 410](#).



Note: Health checks are performed sequentially when used in conjunction with a virtual server configured with multiple services and groups. As a result, the actual health-check interval could vary significantly from its configured value.

About service failures

If a certain number of connection requests for a particular service fail, the WSM places the service into the Service Failed state. While in this state, no new connection requests are sent to the server for this service. However, if you enable [graceful real server failure](#), state information about existing sessions is maintained and traffic associated with existing sessions continues to be sent to the server. See [“Enabling graceful server failure” on page 405](#). Connection requests to, and traffic associated with, other load-balanced services continue to be processed by the server.

For example, a real server is configured to support HTTP and FTP within two real server groups. If the WSM detects an HTTP service failure on the real server, it removes that real server group from the load-balancing algorithm for HTTP but keeps the real server in the mix for FTP. Removing only the failed service from load balancing allows users access to all healthy servers supporting a given service. When a service on a server is in the Service Failed state, the WSM sends Layer 4 connection requests for the failed service to the server. When the WSM has successfully established a connection to the failed service, the service is restored to the load-balancing algorithm.

About server failures

If all load-balanced services supported on a server fail to respond to connection requests within the specified number of [attempts](#), then the server is placed in the Server Failed state. While in this state, no new connection requests are sent to the server. However, if [graceful real server failure](#) is enabled, state information about existing sessions is maintained and traffic associated with existing sessions continues to be sent to the server.



Note: All load-balanced services on a server must fail before the WSM places the server in the server failed state. The server is brought back into service as soon as the first service is proven to be healthy. Additional services are brought online as they are subsequently proven to be healthy.

For more information, see [“Enabling graceful server failure,”](#) next, and [“Modifying the health check interval and retries”](#) on page 410.

Enabling graceful server failure

You can configure the WSM so that when a server fails, traffic and state information for that server’s existing connections are still maintained. By default, this option is disabled.

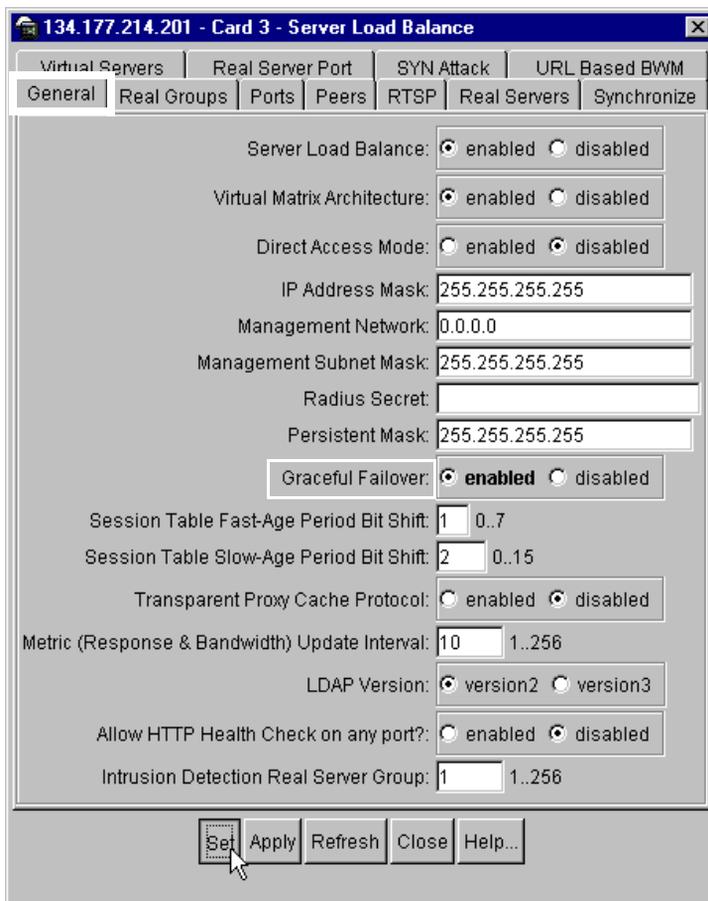
To enable graceful server failure for the WSM:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab ([Figure 160](#)) with the fields described in [Table 61 on page 224](#).

Figure 160 Enabling graceful server failure

- 3 In the Graceful Failover field, click Enabled.
- 4 Click Set > Apply.
An Information dialog box opens reminding you to save the configuration.
- 5 Click OK.
The Information dialog box closes.
- 6 From the General tab, click Close.
Graceful server failure is configured and the SLB dialog box closes.

- 7 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Types of real server group health checks

Table 129 lists and defines the types of health checks that can be configured for a real server group.

Table 129 Types of real server group health checks

| Health check | Description |
|----------------------|--|
| link | For IDSLB group only, checks status of port for each server. |
| arp | For Layer 2 health checking, sends an ARP request. |
| icmp | For Layer 3 health checking, pings the server. |
| tcp | For TCP service, open and close a TCP/IP connection to the server. |
| http | For HTTP service, use HTTP 1.1 GETS when a HOST: header is required to check that the URL content is specified in content command. Otherwise, an HTTP/1.0 GET occurs. If content is not specified in the Health Check Content field, the health check will revert to TCP on the port that is being load balanced. |
| dns | For Domain Name Service, check that the domain name specified in content can be resolved by the server. |
| pop3 | For user mail service, check that the user:password account specified in content exists on the server. |
| smtp | For mail-server services, check that the user specified in content is accessible on the server. |
| nntp | For news group services, check that the news group name specified in content is accessible on the server. |
| ftp | For FTP services, check that the filename specified in content is accessible on the server through anonymous login. |
| imap | For user mail service, check that the user:password value specified in content exists on the server. |

Table 129 Types of real server group health checks (continued)

| Health check | Description |
|--------------|--|
| radius | For RADIUS remote access server authentication, check that the user:password value specified in content exists on the WSM and the server. To perform application health checking to a RADIUS server, you must also configure the Secret field on the SLB General tab. |
| sslh | Query the health of the SSL servers by sending an SSL client Hello packet and then verifying the contents of the server's Hello response. During the handshake, the user and server exchange security certificates, negotiate an encryption and compression method, and establish a session ID for each session. |
| script | Check for application and content availability using one of 16 script-based health checks in send/expect format. <n> denotes the health script number (1-16). The script is configured in the CLI, and can be enabled in JDM for a real server group. |
| wsp | Health checks on connectionless WSP content for WAP gateways. The content must also be configured. For more information, see "Configuring WAP SLB using RADIUS snooping" on page 294 . |
| udpdns | Allows the user to perform health checking using UDP DNS queries. |
| wtls | Provides Wireless Transport Layer Security (WTLS) Hello-based health check for encrypted and connection-oriented WTLS traffic on port 9203. |
| ldap | Sets the health check type to LDAP. |

Configuring a health check for a real server group

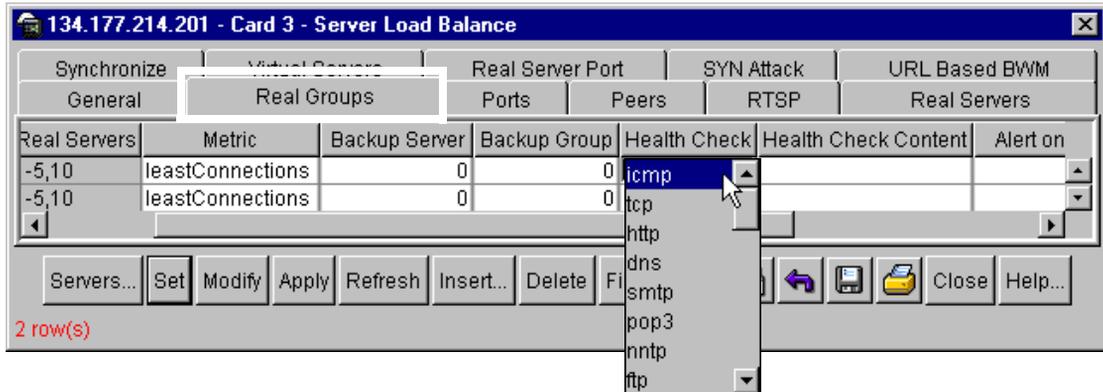
You can configure the [type](#) of health check performed on a real server group. The default health check is [TCP](#).

To configure the type of health check to perform for a real server group:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3 Click the Real Groups tab.

The Real Groups tab (Figure 161) opens with the fields defined in Table 62 on page 228.

Figure 161 Configuring a health check for a real server group



- 4 Double-click the Health Check field for the real server group and choose a health check from the drop-down selection list.
- 5 (Optional) Double-click the Health Check Content field for the real server group, and type a string of up to 34 characters.
- 6 Click Set > Apply.

An Information dialog box opens reminding you to save the configuration.
- 7 Click OK.

The Information dialog box closes.
- 8 From the Real Groups tab, click Close.

The health check is configured for the real server group and the SLB dialog box closes.
- 9 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

For information about scheduling the frequency of health checks, see [“Modifying the health check interval and retries” on page 410](#).

Modifying the health check interval and retries

For TCP services, the WSM verifies that real servers and their services are operational by opening a TCP connection to each service, using the defined service ports. For UDP services, the WSM pings servers to determine their status.

You can modify the following settings for real server health checks:

- the interval between health checks (default = 2)
- the threshold for the number of failed attempts before declaring a server down (default = 4)
- the number of successful attempts required before declaring a server up (default = 8)

The [type](#) of health check is configured in the real server group settings (see [“Configuring a health check for a real server group” on page 408](#)).

To modify the health check interval and retries for a real server:

- 1** From the device view, select the Web Switching Module.

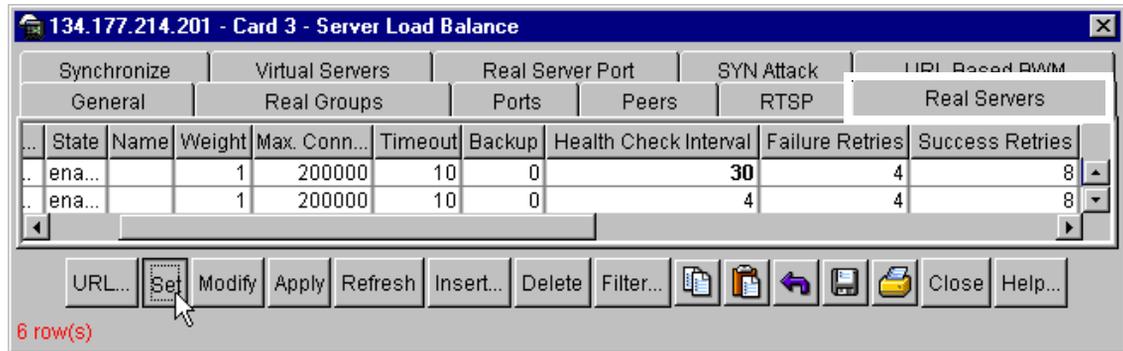
The Web Switching Module is highlighted.

- 2** From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab.

- 3** Click the Real Servers tab.

The Real Servers tab ([Figure 162](#)) opens with the fields defined in [Table 60 on page 219](#).

Figure 162 Modifying the health check interval and retries

- 4 Double-click the Health Check Interval field for the real server, and enter a value (0 - 60 seconds) representing the interval between health checks for this real server.
- 5 Double-click the Failure Retries field for the real server and enter a value (1 - 63) to represent the number of failed attempts required to declare this server down.
- 6 Double-click the Success Retries field for the real server and enter a value (1 - 63) representing the number of successful attempts required to declare this server up.
- 7 Click Set > Apply.
An Information dialog box opens reminding you to save the configuration.
- 8 Click OK.
The Information dialog box closes.
- 9 From the Real Servers tab, click Close.
The health check interval and retries are configured for the real server and the SLB dialog box closes.
- 10 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

ICMP health checking

You can verify that a real server is alive by configuring the ICMP health check for the real server group. The Layer 3 echo - echo reply health check is used for UDP services or when ICMP health checks are configured.

To configure the ICMP health check, choose ICMP from the health check selection list for the real server group. See [“Configuring a health check for a real server group” on page 408](#).

To configure the number of retries required to declare a server up or down, see [“Modifying the health check interval and retries” on page 410](#).

ARP health checking

Address Resolution Protocol (ARP) is the TCP/IP protocol that resides within the Internet layer. ARP resolves a physical address from an IP address. ARP queries machines on the local network for their physical addresses. ARP also maintains IP to physical address pairs in its cache memory. In any IP communication, the ARP cache is consulted for the presence of the IP address of the computer or the router. If present, ARP resolves the physical address to send a packet. In the WSM, this feature allows the user to health check the Intrusion Detection Server (IDS) by sending an ARP query. The ARP health check consists of:

- 1 Accessing the ARP table.
- 2 Looking for the session entry in the ARP table. If the entry exists in the table, that means the real server is up, otherwise the health check has failed.
- 3 If present, checking the timestamp to find out if the last used time is greater than the ARP health check interval. If it is, then delete the query, as this means that the health check has failed.
- 4 Sending another ARP request and repeating the above process until the timestamp shows the last used time smaller than the ARP health check interval.

To configure ARP health checks, choose ARP from the health check selection list for the real server group. See [“Configuring a health check for a real server group” on page 408](#).

Direct server return health checking

Direct Server Return (DSR) health checking verifies the existence of a server service where the server replies directly back to the client without responding through the virtual server IP address.

The server is configured with a real server IP address and virtual server IP address. The virtual server IP address is configured to be the same address as the user's virtual server IP address. When DSR health checking is configured, the health check is sent originating from one of the WSM's configured IP interfaces, and is destined to the virtual server IP address with the MAC address that was acquired from the real server IP address's Address Resolution Protocol (ARP) entry. See [“About direct server return” on page 255](#).

The WSM verifies that the server correctly responds to requests made to the virtual server IP address as required in DSR configurations. To perform this function, the real server IP address is replaced with the virtual server IP address in the health check packets that are forwarded to the real servers for health checking. With this feature enabled, the health check will fail if the real server is not properly configured with the virtual server IP address.

Configuring DSR health checks

Although virtual server IP address (VIP) health checking is enabled by default, it only works when the real server is configured with DSR.

To configure DSR health checks:

- 1 From the device view, select the Web Switching Module.

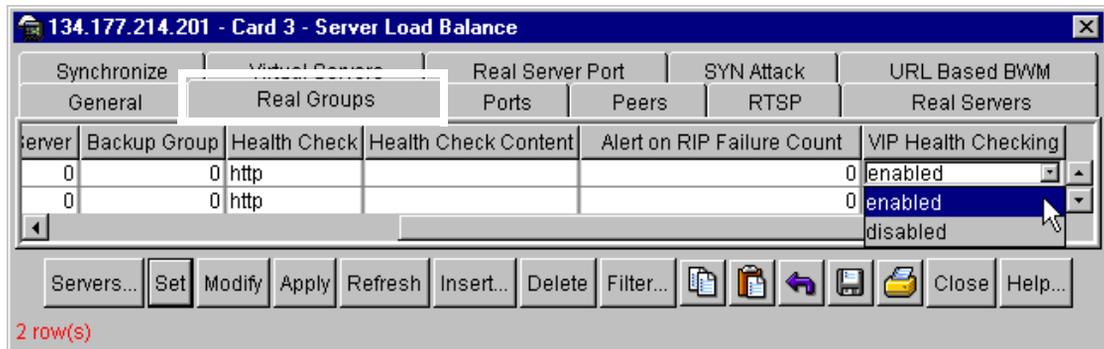
The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching> SLB.

The Server Load Balance dialog box opens to the General tab.

- 3 Click the Real Groups tab.

The Real Groups tab ([Figure 163](#)) opens with the fields defined in [Table 62 on page 228](#).

Figure 163 Enabling VIP health checking for a real server group

- 4 Double-click the VIP Health Checking field for the real server group and choose Enabled from the selection list.
- 5 Click Set > Apply.
An Information dialog box opens reminding you to save the configuration.
- 6 Click OK.
The Information dialog box closes.
- 7 From the Real Groups tab, click Close.
DSR health checking is enabled for the real server group and the SLB dialog box closes.
- 8 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Link health checking

You can check the health of an Intrusion Detection Server (IDS) by configuring link health checks. An IDS often has two physical interfaces—one to detect intrusions, and one to generate logging. The first interface detects intrusions but since it does not have TCP/IP stack, it can only support Layer 1 (physical) health checking. As long as the physical link between the WSM and the IDS is up, it indicates the IDS is alive.

To enable link health checking, choose Link from the real server group health check selection list. See [“Configuring a health check for a real server group” on page 408](#). The real server number determines which port the server is connected to. For example, real server 1 is assumed to be connected to port 1. The valid IDS real server numbers are from 1 to 9 when health check is in use. The server is considered to be up when the link is present and the server is considered to be down when the link is absent.

TCP health checking

TCP is the WSM default health check. TCP health checks verify user-specific TCP applications that cannot be scripted. The WSM monitors the health of servers and applications by sending scheduled Layer 4 connection requests (TCP SYN packets) for each load-balanced TCP service to each server in the server group.

You can configure the rate at which the connection requests are sent. See [“Modifying the health check interval and retries” on page 410](#). Connection requests identify both failed servers and failed services on a healthy server. When a connection request succeeds, the WSM quickly closes the connection by sending a TCP FIN (finished) packet.

Script-based health checking

Script-based health checks run a sequence of tests to dynamically verify application and content availability. You can configure the WSM to send a series of ASCII health check requests to real servers or real server groups and monitor the responses.



Note: Only TCP services can be health checked using scripts, since UDP protocols are usually not ASCII based.

Note: You can enable script-based health checks from JDM, but the actual scripts must be configured from the CLI.

Script-based health checking provides:

- Multiple command format
- Checking for any return string
- Testing for availability of different applications
- Testing for availability of multiple domains or Web sites
- 1024 bytes per script per WSM
- 16 scripts per WSM
- Approximately 10 to 15 health check statements (HTTP Get and Expect strings) per script

The script-based health check is structured to perform the following:

- 1 Open a connection to a specific TCP port.
- 2 Send an ASCII request to the server.
- 3 Expect an ASCII string.

The WSM searches each response packet for this string. If it is not found before the real server health check interval expires, the server does not pass the expect step and fails the health check.

- 4 Close the connection.

A script can contain any number of these commands, up to the allowable number of characters that a script supports.

Configuring script-based health checks

You can enable script-based health checks from JDM, but the actual scripts must be configured from the CLI.

To enable a specific script-based health check, choose the script (script1-script16) from the list of available health checks for the real server group. See [“Configuring a health check for a real server group” on page 408](#).

Application health checking

A Layer 7 health check is an application-aware health check designed for the specific application. For a list of application ports known to the WSM, see [“Well-known TCP/UDP application port numbers” on page 191](#). This section includes the following topics:

- [“HTTP health checking” on page 417](#)
- [“UDP-based DNS health checking” on page 419](#)
- [“IMAP server health checking” on page 420](#)
- [“RADIUS server health checking” on page 421](#)
- [“HTTPS/SSL server health checking” on page 423](#)
- [“WAP gateway health checking” on page 423](#)
- [“LDAP health checking” on page 430](#)

HTTP health checking

The WSM sends HTTP GET or HEAD requests to HTTP servers specifying a page (identified by a URL) on the server. The WSM examines the page, compares the content of the page to user-defined content, and marks the port as either Active or Failed.

HTTP-based health checks can include the hostname for HOST: headers. If the HOST: header is required, an HTTP/1.1 GET will occur. Otherwise, an HTTP/1.0 GET will occur. An HTTP health check is successful if you get return code 200.



Note: If content is not specified for the HTTP health check, the port being load-balanced will revert to a TCP health check.

The following examples of HTTP health checks show how the check is performed depending on the criteria you define.

Example 1:

Virtual server service host name = everest
Virtual server domain name = alteonwebsystems.com
Real server group content = index.html
Health check is performed using: GET /index.html HTTP/1.1
Host: everest.alteonwebsystems.com

Example 2:

Virtual server service host name = (none)
Virtual server domain name = raleighduram.cityguru.com
Real server group content = /page/gen/?_template=alteon
Health check is performed using: GET /page/gen/?_template=alteon HTTP/1.1
Host: raleighduram.cityguru.com

Example 3:

Virtual server service host name = (none)
Virtual server domain name = jansus
Real server group content = index.html
Health check is performed using: GET /index.html HTTP/1.1
Host: jansus

Example 4:

Virtual server service host name = (none)

Virtual server domain name = (none)

Real server group content = index.html

Health check is performed using: GET /index.html HTTP/1.0
(since no HTTP Host: header is required)

Example 5:

Virtual server service host name = (none)

Virtual server domain name = (none)

Real server group content = //everest/index.html

Health check is performed using: GET /index.html HTTP/1.1
Host: everest

Configuring HTTP health checks

HTTP health checking includes the following configurations:

- 1 [“Configuring a health check for a real server group” on page 408.](#)
- 2 (Optional) [“Modifying the health check interval and retries” on page 410.](#)
- 3 Configuring a domain name for the virtual server. See [“Configuring a virtual server” on page 230.](#)
- 4 Specifying a host name for the virtual service. See [“Configuring services for a virtual server” on page 235.](#)

UDP-based DNS health checking

A UDP-based DNS health check is performed by sending a UDP-based query (for example, www.nortelnetworks.com), and watching for the server’s reply. You can modify the domain name to be queried in the Health Check Content field. If you enable UDP-based DNS health checks, you can send TCP-based queries to one real server group and UDP-based queries to another real server group.

To configure UDP-based DNS health checks, see [“Configuring a health check for a real server group” on page 408](#). Configure the fields in [Table 130](#).

Table 130 Real server group settings for UDP-based DNS health checking

| Real server group field | Setting |
|-------------------------|---------------|
| Health Check | UDPDNS |
| Health Check Content | <domain name> |



Note: If no host name is configured, the health check is performed by sending a UDP-based query from a dummy host and watching for the server’s reply. The reply, even though negative (for example, *Server not found* if from a dummy host), satisfies the health check.

IMAP server health checking

Internet Message Access Protocol (IMAP) is a mail server protocol used between a client system and a mail server that lets you retrieve and manipulate mail messages. IMAP is not used for mail transfers between mail servers. IMAP servers listen to TCP port 143. You configure a username and password for the real server group. The IMAP health check matches this username and password content in its user database.

To configure IMAP server health checking, see [“Configuring a health check for a real server group” on page 408](#). Use the settings in [Table 132](#).

Table 131 Real server group settings for IMAP health checks

| Real server group field | Setting |
|-------------------------|-------------------------|
| Health Check | IMAP |
| Health Check Content | <username and password> |

RADIUS server health checking

For RADIUS authentication, the WSM sends an authentication request with a user name, password, and key to the RADIUS server. The account information need not be valid for the server to pass the health check. In fact, by using invalid information for the health check, you can safeguard account information. You can configure a check for either the well-known RADIUS port number 1812 or 1645, but you cannot configure a health check for both on the same server.

RADIUS server health checks require the following configurations:

- “[Configuring a health check for a real server group](#)” on page 408. Use the settings in [Table 132](#).

Table 132 Real server group settings for RADIUS health checking

| Real server group field | Setting |
|-------------------------|-------------------------|
| Health Check | RADIUS |
| Health Check Content | <username and password> |

- “[Configuring the RADIUS authentication string](#),” next.

Configuring the RADIUS authentication string

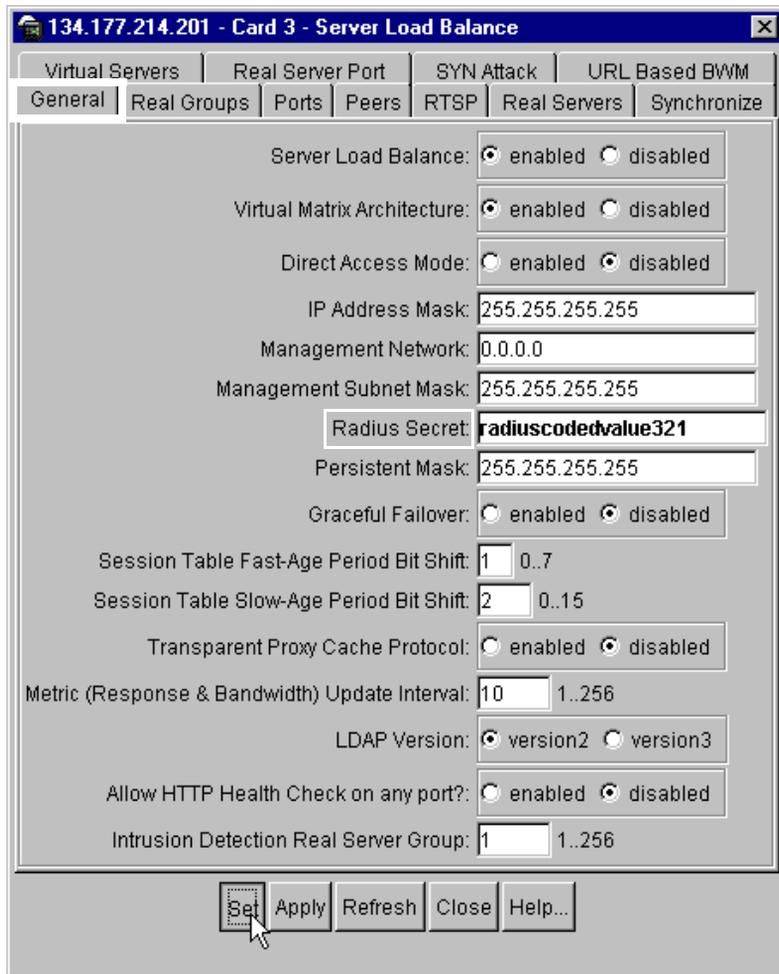
To configure the RADIUS authentication string:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching> SLB.

The Server Load Balance dialog box opens to the General tab ([Figure 160](#)) with the fields described in [Table 61 on page 224](#).

Figure 164 Setting the RADIUS authentication string

- 3 In the RADIUS Secret field, type an authentication string of up to 32 alphanumeric characters.
- 4 Click Set > Apply.
An Information dialog box opens reminding you to save the configuration.
- 5 Click OK.
The Information dialog box closes.
- 6 From the General tab, click Close.

The RADIUS authentication is configured and the SLB dialog box closes.

- 7 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

HTTPS/SSL server health checking

You can configure the WSM to send an SSL client Hello packet and verify the contents of the server's Hello response to check the health of the SSL server. This SSLH health check uses the configured real server port.

The SSL enhanced health check behavior is summarized below:

- The WSM sends an SSL Hello packet to the SSL server.
- If it is up and running, the SSL server responds with a Server Hello message.
- The WSM verifies fields in the response and marks the service as being Up if the fields are OK.

During the handshake, the user and server exchange security certificates, negotiate an encryption and compression method, and establish a session ID for each session. To configure SSLH, see [“Configuring a health check for a real server group” on page 408](#).

WAP gateway health checking

Wireless application protocol (WAP) carries Internet traffic to mobile devices for delivering Web services to mobile phones and handsets. Servers on the land-based part of the network, known as WAP gateways, translate the HTTP/HTML to WAP/WML (wireless markup language).

WAP devices can communicate in two ways:

- [Wireless Session Protocol \(WSP\) content health checks sending unencrypted WML traffic \(similar to HTTPS\)](#). See [“Configuring WSP content health checks, next.”](#)

- [Wireless Transport Layer Security \(WTLS\) health checks](#), sending encrypted WML traffic (similar to HTTP). See [“Configuring WTLS health checks” on page 428](#).

Configuring WSP content health checks

The WSP content health check requires the following configurations:

- [“Defining WSP content to check,”](#) next.
- [“Configuring the UDP service for WSP health check” on page 426](#).
- [“Configuring a health check for a real server group” on page 408](#).

Defining WSP content to check

The content of the WSP/UDP packet sent to the gateway can be configured as a hexadecimal string, encapsulated in a UDP packet, and shipped to the server. This byte string should include all applicable WSP headers.

The content that the WSM expects to receive from the gateway is also specified in a hexadecimal byte string. The WSM matches each byte of this string with the received content.

If there is a mismatch of even a single byte on the received content, the gateway fails the health check. You can also configure an offset for the received WSP packet: a byte index to the location in the WSP response content where the byte match can be performed.

To define the WSP content to check:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching > WAP.

The WAP dialog box opens to the General tab ([Figure 165](#)) with the fields described in [Table 72 on page 299](#).

Figure 165 Configuring content for WSP health check

134.177.214.201 - Card 3 - WAP

General

WAP Load Balancing: enabled disabled

WAP Debug Level: 0 0..10

WSP Port to Health Check: 9200 0..65534

WTLS Port to Health Check: 9203 0..65534

Offset into WSP Packet: 0 0..256

Health Content to Send to WAP Gateway: 01 42 15 68 74 74 70 3a 2f 77 77 77 2e 6e 6f 6b 61 6d 00

Expected Health Content from WAP Gateway: 01 04 60 0e 03 94

Set Apply Refresh Close Help...

- 3 In the WSP Port to Health Check field, type the WSP port number (9200 is the default) on which WSP health checks are performed.
- 4 In the Offset into WSP Packet field, type an offset from 0 - 256. If set to 0 (the default) comparisons start at the beginning of the received packet's content.
- 5 In the Health Content to Send to WAP Gateway field, type a hexadecimal string of up to 255 characters to be examined during the health check.
- 6 In the Expected Health Content from WAP Gateway field, type a hexadecimal string of up to 255 characters for the gateway response to match.
- 7 Click Set > Apply.
An Information dialog box opens reminding you to save the configuration.
- 8 Click OK.
The Information dialog box closes.
- 9 From the General tab, click Close.
The content for the WSP health check is configured.
- 10 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

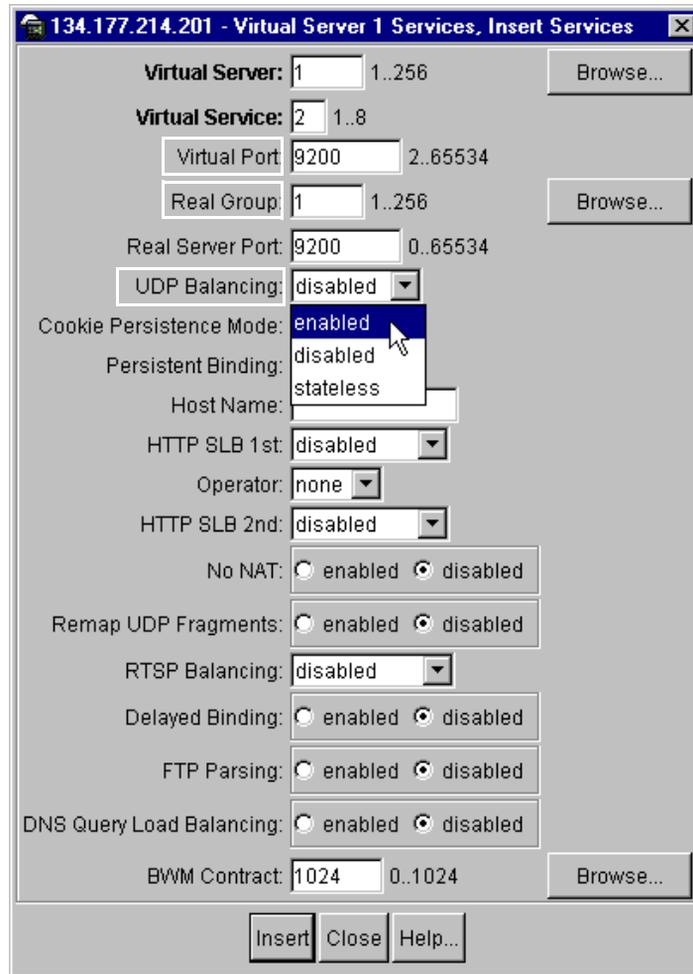
- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Configuring the UDP service for WSP health check

Connectionless WSP runs on UDP/IP protocol, port 9200 and WSMs can load balance these gateways.

To configure the UDP service for WSP health checking:

- 1** From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 2** Click the Virtual Servers tab.
The Virtual Servers tab opens
- 3** Click in the Virtual Server field for the server you want to configure.
The Virtual Server number is highlighted.
- 4** Click Services.
The Virtual Server Services dialog box opens, listing the services for this virtual server.
- 5** Click Insert.
- 6** The Insert Virtual Server Services dialog box ([Figure 166](#)) opens with the fields defined in [Table 65 on page 237](#).

Figure 166 Configuring the WSP service

- 7 In the Virtual Port field, type 9200 for the WSP service.
- 8 In the Real Group field, type the real server group number.
- 9 In the UDP Balancing field, click the down arrow and choose Enabled from the selection list.
- 10 Click Insert.

The Insert Virtual Server Services dialog box closes and the WSP service is added to the list of services for the virtual server.

11 In the Virtual Server Services dialog box, click Apply > Close.

The Virtual Server Services dialog box closes.

12 From the Virtual Servers tab, click Set > Apply.

An Information dialog box opens reminding you to save the configuration.

13 Click OK.

The Information dialog box closes.

14 From the Virtual Servers tab, click Close.

The SLB dialog box closes and the UDP service is configured for the WSP health check.

15 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

Configuring WTLS health checks

The WSM performs a WTLS Hello-based health check for connection-oriented WTLS traffic on port 9203. It sends a new WTLS Client Hello to the WAP gateway, and checks to see if it receives a valid WTLS Server Hello back from the WAP Gateway. You can configure a WTLS health check on the real server group, and you can change the port on which your gateway listens to WTLS traffic.

To configure WTLS health checks for a real server group, see [“Configuring a health check for a real server group” on page 408](#).

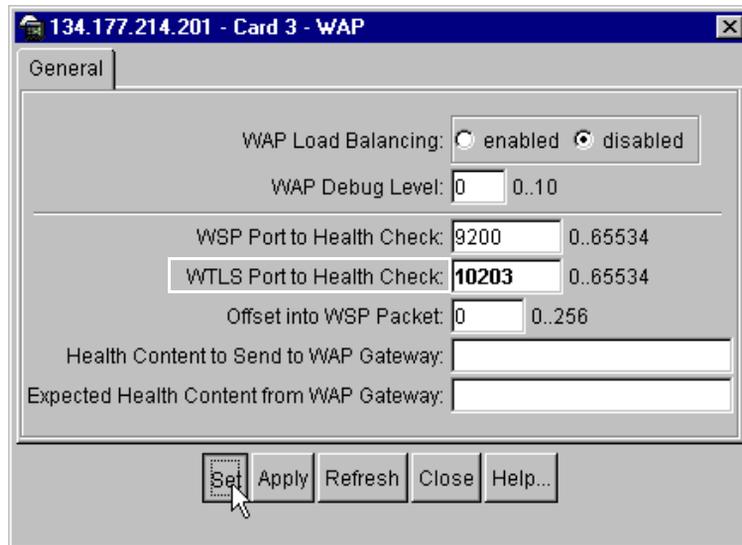
To change the port on which your gateway is listening to WTLS traffic:

1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

2 From the Edit menu, choose WSM Card > L4 Switching > WAP.

The WAP dialog box opens to the General tab ([Figure 165](#)) with the fields described in [Table 72 on page 299](#).

Figure 167 Configuring WTLS port to health check

- 3 In the WTLS Port to Health Check field, type a port number (9203 is the default) on which WTLS health checks are performed.
- 4 Click Set > Apply.
An Information dialog box opens reminding you to save the configuration.
- 5 Click OK.
The Information dialog box closes.
- 6 From the General tab, click Close.
The WTLS port is configured.
- 7 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

LDAP health checking

Lightweight Directory Access Protocol (LDAP) health checks enable the WSM to determine if the LDAP server is alive. LDAP versions 2 and 3 are described in RFC 1777 and RFC 2251. The LDAP health check process consists of three LDAP messages over one TCP connection:

- **Bind request**—The WSM first creates a TCP connection to the LDAP server on port 339, which is the default port. After the connection is established, the WSM initiates an LDAP protocol session by sending an anonymous bind request to the server.
- **Bind response**—On receiving the bind request, the server sends a bind response to the WSM. If the result code indicates that the server is alive, the WSM marks the server as up. Otherwise, the WSM marks the server as down even if the WSM did this because the server did not respond within the timeout window.
- **Unbind request**—If the server is alive, the WSM sends a request to unbind the server. This request does not require a response. It is necessary to send an unbind request as the LDAP server may crash if too many protocol sessions are active.

If the server is up, the WSM closes the TCP connection after sending an unbind request. If the server is down, the connection is torn down after a bind response, if one arrives. The connection will also be torn down if it crosses the timeout limit, regardless of the server's condition.

Configuring LDAP health checking

LDAP health checking requires that you configure LDAP as the health check type for the real server group. You may also need to change the LDAP version specified for the WSM. The default is LDAP version 2. For more information, see:

- [“Configuring a health check for a real server group” on page 408](#)
- [“Changing the LDAP version, next”](#)

Changing the LDAP version

The default LDAP version is 2.

To change the LDAP version:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Edit menu, choose WSM Card > L4 Switching> SLB.

The Server Load Balance dialog box opens to the General tab (Figure 168) with the fields described in Table 61 on page 224.

Figure 168 Changing the WSM LDAP version

134.177.214.201 - Card 3 - Server Load Balance

| Virtual Servers | Real Server Port | SYN Attack | URL Based BWM |
|-----------------|------------------|-------------|---------------|
| General | Real Groups | Ports | Peers |
| RTSP | Real Servers | Synchronize | |

Server Load Balance: enabled disabled

Virtual Matrix Architecture: enabled disabled

Direct Access Mode: enabled disabled

IP Address Mask: 255.255.255.255

Management Network: 0.0.0.0

Management Subnet Mask: 255.255.255.255

Radius Secret:

Persistent Mask: 255.255.255.255

Graceful Failover: enabled disabled

Session Table Fast-Age Period Bit Shift: 1 0..7

Session Table Slow-Age Period Bit Shift: 2 0..15

Transparent Proxy Cache Protocol: enabled disabled

Metric (Response & Bandwidth) Update Interval: 10 1..256

LDAP Version: version2 version3

Allow HTTP Health Check on any port?: enabled disabled

Intrusion Detection Real Server Group: 1 1..256

Set Apply Refresh Close Help...

- 3** In the LDAP Version field, click the version you are running.
- 4** Click Set > Apply.
An Information dialog box opens reminding you to save the configuration.
- 5** Click OK.
The Information dialog box closes.
- 6** From the General tab, click Close.
The LDAP version is configured and the SLB dialog box closes.
- 7** To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Chapter 17

Maintaining session persistence

This section contains the following topics:

- [“About session persistence,”](#) next
- [“About cookie-based persistence”](#) on page 435
- [“Configuring cookie-based session persistence”](#) on page 442
- [“Examples of cookie values”](#) on page 443
- [“SSL session ID-based persistence”](#) on page 448
- [“Configuring SSL session ID-based persistence”](#) on page 450

About session persistence

In an SLB environment, client traffic comes across the Internet to a WSM virtual server IP address. The WSM balances this traffic among the available real servers.

In any authenticated Web-based application, the WSM must provide a persistent connection between a client and the Web server to which it is connected. Because HTTP does not carry state information for these applications, it is important that the browser be mapped to the same real server for each HTTP request until the transaction is complete. This ensures that server load balancing does not send mid-session client traffic to a different real server, forcing the user to restart the entire transaction.

Persistence-based SLB enables you to configure the network to redirect requests from a client to the same real server that initially handled the request. Persistence is an important consideration for administrators of e-commerce Web sites, where a server may have data associated with a specific user that is not dynamically shared with other servers at the site. In the WSM, persistence can be based on source IP address, cookies, and the Secure Sockets Layer (SSL) session ID.

Effects of using source IP address for session persistence

Until recently, the only way to achieve TCP/IP session persistence was to use the source IP address as the key identifier. Two concerns with basing session persistence on a packet's IP source address are:

- Many clients sharing the same source IP address (proxied clients)
Proxied clients appear to the WSM as a single source IP address and do not take advantage of SLB on the WSM. When many individual clients behind a firewall use the same proxied source IP address, requests are directed to the same server, without the benefit of load balancing the traffic across multiple servers. Persistence is supported without the capability of effectively distributing traffic load. Also, persistence is broken if you have multiple proxy servers behind the WSM performing SLB. The WSM changes the client's address to different proxy addresses as attempts are made to load balance client requests.
- Single client sharing a pool of source IP addresses
When individual clients share a pool of source IP addresses, persistence for any given request cannot be assured. Although each source IP address is directed to a specific server, the source IP address itself is randomly selected, thereby making it impossible to predict which server will receive the request. SLB is supported, but without persistence for any given client.

Using cookies for session persistence

Cookies are strings passed via HTTP from servers to browsers, and are inserted by either the WSM or the server. After a client receives a cookie, a server can poll that cookie with a GET command, which allows the querying server to positively identify the client as the one that received the cookie earlier. The cookie-based persistence feature solves the proxy server problem and gives better load distribution at the server site. In the WSM, cookies are used to route client traffic back to the same physical server to maintain session persistence.

Using SSL session ID for session persistence

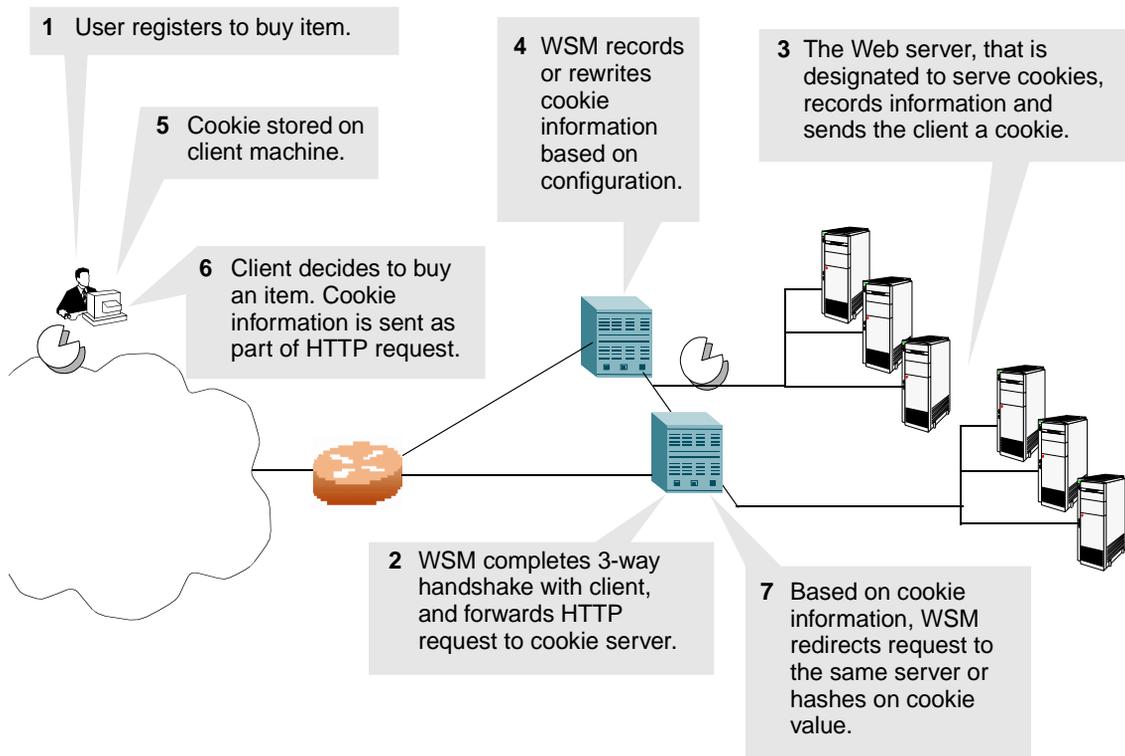
The SSL session ID is effective only when the server is running SSL transactions. Because of the heavy processing load required to maintain SSL connections most network configurations use SSL only when it is necessary. Persistence based on SSL Session ID ensures completion of complex transactions in proxy server environments. However, this type of persistence does not scale on servers because of their computational requirements.

About cookie-based persistence

Cookies maintain state data between clients and servers. When a server receives a client request, it issues a cookie, or token, to the client. The client returns the cookie to the server on all subsequent requests. Using cookies, the server does not require authentication, the client IP address, or any other mechanism to determine it is the same user that sent the original request. The cookie may be:

- A customer ID.
- A token of trust, allowing the user to skip authentication while his or her cookie is valid.
- A key that associates the user with additional state data that is kept on the server, such as a shopping cart and its contents.
- Encoded so that it actually contains more data than just a single key or an identification number.
- The user's preferences for a site that allows their pages to be customized.

[Figure 169](#) illustrates cookie-based persistence.

Figure 169 Understanding cookie-based session persistence

Defining permanent versus temporary cookies

Cookies can either be permanent or temporary.

- A permanent cookie is stored on the client's browser, as part of the response from a Web site's server. It is sent by the browser when the client makes subsequent requests to the same site, even after the browser has been shut down.
- A temporary cookie is only valid for the current browser session. Similar to a SSL Session- based ID, the temporary cookie expires when you shut down the browser. Based on RFC 2109, any cookie without an expiration date is a temporary cookie.

Cookie formats

It is a good practice to define a cookie in the HTTP header. A cookie can also be placed in the URL for hashing. [Table 133](#) shows how to designate a cookie session ID in an HTTP header.

Table 133 Designating the cookie session ID in an HTTP header

| Designation | Description |
|---|----------------------------|
| Session ID | SessionID-1234 |
| Active Server Page (ASP) session ID | ASP_SESSIONID=POIUHKJHLKHD |
| Application-specific cookie that records the name of the client | name=john_smith |

You can also designate the cookie session ID within the URL, as follows:

`http://www.mysite.com/reservations/SessionID=1234`

Cookie properties

Cookies are configured on the WSM by defining the following properties:

- A name of up to 20 bytes
- The offset of the cookie value within the cookie string

For enhanced security, the real cookie value can be embedded somewhere within a longer string. The offset directs the WSM to the starting point of the real cookie value within the longer cookie string.

- The length of the cookie value

This defines the number of bytes to extract for the cookie value within a longer cookie string.

- Where to find the cookie value—in the HTTP header (the default) or the URI
- Values of up to 64 bytes for hashing

Hashing on cookie values is used only with the passive cookie mode (see [“Passive cookie mode” on page 440](#)), using a temporary cookie. The WSM mathematically calculates the cookie value using a hash algorithm to determine which real server should receive the request.

- An asterisk (*) in cookie names for wildcard
For example, Cookie name = ASPsession*

Handling client browsers that do not accept cookies

If a client browser is not configured to accept cookies, you must use the hash load-balancing metric to maintain session persistence. With cookie-based persistence enabled, session persistence for browsers that do not accept cookies will be based on the source IP address. However, individual client requests coming from a proxy firewall will appear to be coming from the same source IP address. Therefore, the requests will be directed to a single server, resulting in traffic being concentrated on a single real server instead of load-balanced across all available real servers.

Defining cookie persistence modes

[Table 134](#) describes the modes of operation for cookie-based session persistence.

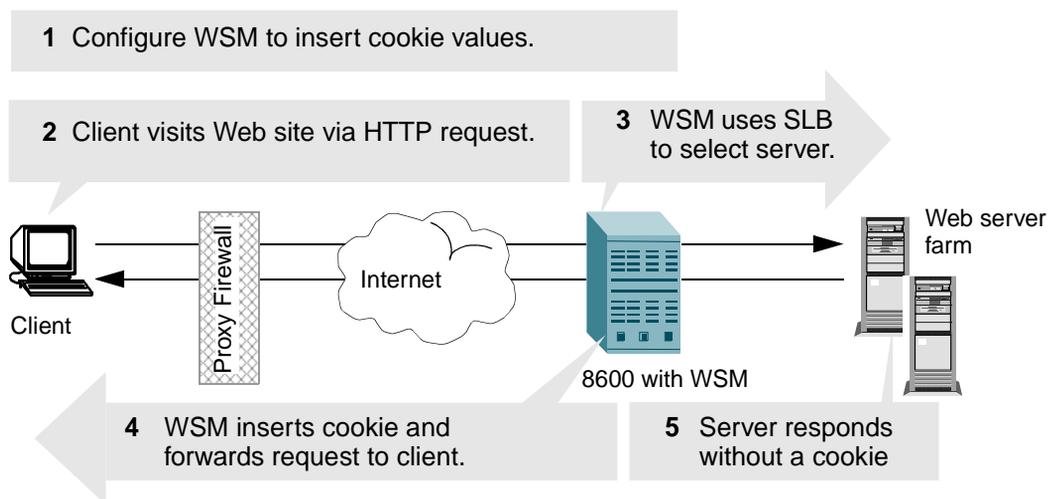
Table 134 Modes of operation for cookie-based session persistence

| Mode | Configuration | Location | Uses WSM session entry |
|---|----------------|--------------------|------------------------|
| "Insert cookie mode" on page 439 | WSM only | HTTP header | No |
| "Passive cookie mode" on page 440 | server and WSM | HTTP header or URL | Yes |
| "Rewrite cookie mode" on page 441 | server and WSM | HTTP header | No |

Insert cookie mode

Figure 170 illustrates the insert cookie mode. The client sends a request to visit the Web site. The WSM performs load-balancing and selects a real server. The real server responds without a cookie. The WSM inserts a cookie and forwards the new request with the cookie to the client. Because the real servers don't configure cookies, they don't require cookie server software.

Figure 170 Session persistence—insert cookie mode



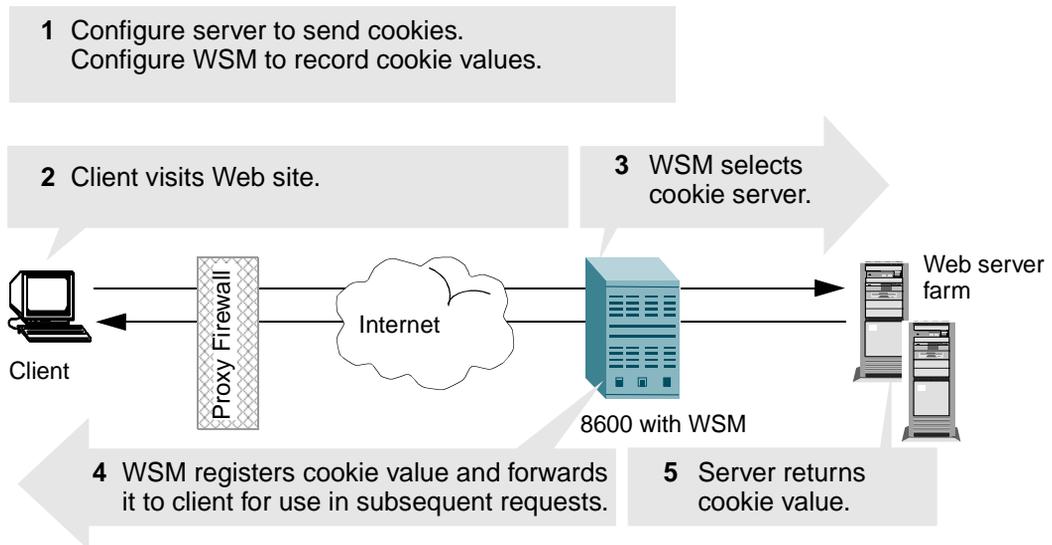
Passive cookie mode

Figure 171 illustrates the passive cookie mode. When the client first makes a request, the WSM selects the server based on the load-balancing metric. The real server embeds a cookie in its response to the client. The WSM records the cookie value and matches it in subsequent requests from the same client. Subsequent requests from this client with the same cookie value will be sent to the same real server.



Note: The passive cookie mode is recommended for temporary cookies. However, you can use this mode for permanent cookies if the server is embedding an IP address.

Figure 171 Session persistence—passive cookie mode



Rewrite cookie mode

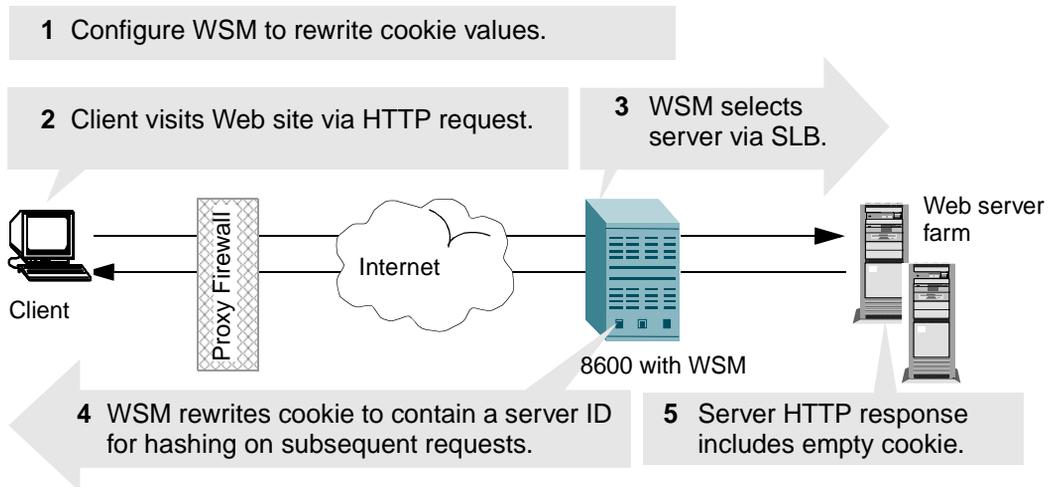
Figure 172 illustrates the rewrite cookie mode. The WSM generates the cookie value on behalf of the server. The server is configured to return a special persistence cookie which the WSM is configured to recognize. The WSM intercepts this persistence cookie and rewrites the value to include server-specific information before sending it on to the client. Subsequent requests from the same client with the same cookie value are sent to the same real server.

Rewrite cookie mode requires at least eight bytes in the cookie header. An additional eight bytes must be reserved if you are using cookie-based persistence with Global Server Load Balancing (GSLB).



Note: Rewrite cookie mode only works for cookies defined in the HTTP header, not cookies defined in the URL.

Figure 172 Session Persistence—rewrite cookie mode



Note: The WSM rewrites the cookie to an encoded value that represents the responding server. The value can be hashed into a real server ID or it can be the real server IP address.

Configuring cookie-based session persistence

To configure the WSM for cookie-based persistence:

- 1 Define an IP interface on the WSM. See [“Defining an IP interface on the WSM” on page 225](#).
- 2 Configure each real server with its IP address, name, weight, and so forth. See [“Configuring each real server” on page 216](#).
- 3 Assign servers to real server groups. See [“Configuring a real server group” on page 225](#).
 - If embedding an IP address in the cookie, select the roundrobin or leastconns load-balancing metric.
 - If not embedding an IP address in the cookie, select hash as the load-balancing metric in conjunction with a cookie assignment server.



Note: While you may experience traffic concentration using the hash metric with a cookie assignment server, using a hash metric without a cookie assignment server will cause traffic concentration on your real servers.

- 4 Define the virtual servers. See [“Configuring a virtual server” on page 230](#).
 - Set the following parameters on the virtual server:
 - Cookie Name
 - Cookie Offset
 - Cookie Length
 - Cookie Search
 - Expire Cookie (for Insert mode only)
 - Cookie Search Response Count (for passive mode only)



Note: In complex server configurations, the persistence cookie might not be sent in the first server response. If this happens, you can configure the WSM to look through multiple HTTP responses from the server. The WSM looks for the persistence cookie in the specified number of server responses (each of them can be multi-frame).

- 5 Define the virtual server services. See [“Configuring services for a virtual server” on page 235](#).
 - Set the following parameters on the virtual server service:
 - Persistent binding field=Cookie
 - Cookie Persistence Mode field=Rewrite, Passive, or Insert
- 6 Enable client or server processing on the ports. See [“Configuring ports for server load balancing” on page 240](#).
- 7 Do one of the following:
 - Enable direct access mode (DAM) on the WSM. See [“Configuring direct access mode” on page 258](#).
 - Configure client ports with proxy IP addresses. See [“Configuring proxy IP addresses” on page 245](#).
- 8 Enable server load balancing for the WSM. See [“Enabling or disabling server load balancing” on page 243](#).

Cookie-based session persistence is configured for the WSM.

Examples of cookie values

This section includes the following examples:

- [“Example 1: Setting the cookie location” on page 443](#)
- [“Example 2: Parsing the cookie” on page 444](#)
- [“Example 3: Using passive cookie mode” on page 446](#)
- [“Example 4: Using rewrite cookie mode” on page 446](#)

Example 1: Setting the cookie location

This example illustrates how to configure the WSM to search for a cookie in an HTTP header or a URI.

The following client request has two different cookies labeled UID—one in the HTTP header and the other in the URI.

```
GET /product/switch/UID=12345678;ck=1234...  
Host: www.alteonwebsystems.com  
Cookie: UID=87654321
```

To configure the WSM to search for the UID or HTTP cookie, use the following virtual server ([Table 135](#)) and service ([Table 136](#)) settings.

Table 135 Virtual server settings for HTTP or URI header cookie location

| Virtual Server setting | Description |
|-------------------------------------|---|
| Cookie Search = Enabled or Disabled | <ul style="list-style-type: none">Enabled— looks for cookie name/value pair in the URI.Disabled—looks for the cookie in the HTTP header. |
| Cookie Name = UID | Embeds this cookie name in the server response. The WSM searches for this cookie name in subsequent requests from the same client. |
| Cookie Offset = 1 | Tells the WSM to start searching at the first byte of the cookie. |
| Cookie Length = 8 | Tells the WSM to extract 8 bytes. |

Table 136 Service 80 settings for HTTP or URI header cookie location

| Virtual Server Service 80 setting | Description |
|-----------------------------------|---|
| Persistent Binding = Cookie | Tells the WSM to look for the cookie embedded in the server response in subsequent requests from the same client. |
| Cookie Persistence Mode = Passive | Enables passive cookie-based persistence. |

Example 2: Parsing the cookie

This example includes three WSM configurations using hashing or wild cards to determine which part of the following cookie to use for selecting a real server.

```
Cookie: sid=0123456789abcdef; name1=value1
```

To direct the WSM to use the entire sid cookie as a hashing key for selecting the real server, use the following virtual server ([Table 137](#)) and service ([Table 138](#)) settings.

Table 137 Virtual server settings for parsing the cookie

| Virtual Server setting | Description |
|--------------------------|--|
| Cookie Search = Disabled | Disabled—looks for the cookie in the HTTP header. |
| Cookie Name = sid | Embeds this cookie name in the server response. The WSM searches for this cookie name in subsequent requests from the same client. |
| Cookie Offset = 1 | Tells the WSM to start searching at the 1st byte of the cookie. |
| Cookie Length = 16 | Tells the WSM to extract the full 16 bytes. |

Table 138 Service 80 settings for parsing the cookie

| Virtual Server Service 80 setting | Description |
|-----------------------------------|---|
| Persistent Binding = Cookie | Tells the WSM to look for the cookie embedded in the server response in subsequent requests from the same client. |
| Cookie Persistence Mode = Passive | Enables passive cookie-based persistence. |

- To direct the WSM to start with the 8th byte in the value and use only four bytes of the sid cookie (789a) as a hashing key for selecting the real server, change the following virtual server ([Table 137](#)) fields.

Cookie Offset = 8
Cookie Length = 4

- To direct the WSM to use a wildcard in selecting the cookie, change the following virtual server ([Table 137](#)) field.

Cookie Name = ASPSESSIONID*

The WSM will look for a cookie name that starts with ASPSESSIONID. ASPSESSIONID123, ASPSESSIONID456, and ASPSESSIONID789 will all be seen by the switch as the same cookie name. If more than one cookie matches, only the first one will be used.

Example 3: Using passive cookie mode

This example shows how to configure passive cookie mode.

If you are using passive cookie mode, the WSM examines the server's `Set-Cookie:` value and directs all subsequent connections to the server that assigned the cookie. For example, if Server 1 sets the cookie as `Set-Cookie: sid=12345678`, then all traffic from a particular client with cookie `sid=12345678` will be directed to Server 1.

To direct the WSM to examine the server's `Set-Cookie:` value for selecting the real server, use the following virtual server ([Table 139](#)) and service ([Table 140](#)) settings.

Table 139 Virtual server settings for passive cookie mode

| Virtual Server setting | Description |
|--------------------------|--|
| Cookie Search = Disabled | Tells the WSM to look for the cookie in the HTTP header. |
| Cookie Name = sid | Embeds this cookie name in the server response. The WSM searches for this cookie name in subsequent requests from the same client. |
| Cookie Offset = 1 | Tells the WSM to start searching at the first byte of the cookie. |
| Cookie Length = 8 | Tells the WSM to extract 8 bytes. |

Table 140 Service 80 settings for passive cookie mode

| Virtual Server Service 80 setting | Description |
|-----------------------------------|---|
| Persistent Binding = Cookie | Tells the WSM to look for the cookie embedded in the server response in subsequent requests from the same client. |
| Cookie Persistence Mode = Passive | Enables passive cookie-based persistence for the service. |

Example 4: Using rewrite cookie mode

This example shows how to configure the rewrite cookie mode with an 8-byte or 16-byte cookie length.

- 8-byte cookie length:

If you configure an 8-byte cookie length, the WSM rewrites the placeholder cookie value with the encrypted real server IP address.

If the server is configured to include a placeholder cookie (Set-Cookie: sid=alteonpersistence), the WSM rewrites the first eight bytes of the cookie to include the server's encrypted IP address (Set-Cookie: sid=cdb20f04rsistence). Subsequent traffic from a specific client with this cookie will be directed to the same real server.

- 16-byte cookie length:

If you configure a 16-byte cookie length, the WSM will rewrite the cookie value with the encrypted real server IP address and virtual server IP address.

If the server is configured to include a placeholder cookie (Set-Cookie: sid=alteonwebcookies) the WSM will rewrite the first 16 bytes of the cookie to include the encrypted real server IP address and virtual server IP address (Set-Cookie: sid=cdb20f04cdb20f0a). Subsequent traffic from a specific client to the particular virtual server IP address with this cookie will be directed to the same real server.

To direct the WSM to rewrite the server cookie, use the following virtual server (Table 141) and service (Table 142) settings:

Table 141 Virtual server settings for rewrite cookie mode

| Virtual Server setting | Description |
|--------------------------|--|
| Cookie Search = Disabled | Tells the WSM to look for the cookie in the HTTP header. |
| Cookie Name = sid | Embeds this cookie name in the server response. The WSM searches for this cookie name in subsequent requests from the same client. |
| Cookie Offset = 1 | Tells the WSM to start searching at the first byte of the cookie. This setting is not effective in rewrite mode. |
| Cookie Length = 8 or 16 | Tells the WSM to extract 8 bytes to rewrite the encrypted real server IP address, or 16 bytes to rewrite both the encrypted real server and virtual server IP addresses. |

Table 142 Service 80 settings for rewrite cookie mode

| Virtual Server Service 80 setting | Description |
|-----------------------------------|---|
| Persistent Binding = Cookie | Tells the WSM to look for the cookie embedded in the server response in subsequent requests from the same client. |
| Cookie Persistence Mode = Rewrite | Tells the WSM to generate the cookie value on behalf of the server. The WSM intercepts this persistence cookie and rewrites it with the real server's encrypted IP address before sending it to the client. NOTE: In Rewrite mode, the cookie length can only be 8 or 16 and the cookie offset parameter is not effective. |

SSL session ID-based persistence

Secure sockets layer (SSL) is a set of protocols built on top of TCP/IP that allows an application server and client to communicate over an encrypted HTTP session, providing authentication, non-repudiation, and security. The client and server complete the SSL protocol handshake using clear (unencrypted) text. The content data is then encrypted (using an algorithm exchanged during the handshake) prior to being transmitted.

Using the SSL session ID, the WSM forwards the client request to the same real server to which it was bound during the last session. Because SSL protocol allows many TCP connections to use the same session ID from the same client to a server, key exchange needs to be done only when the session ID expires. This reduces server overhead and provides a mechanism, even when the client IP address changes, to send all sessions to the same real server.



Note: You can configure the destination port number to monitor for SSL traffic.

How SSL session ID-based persistence works

- All SSL sessions that present the same session ID (32 random bytes chosen by the SSL server) will be directed to the same real server.



Note: The SSL session ID can only be read by the WSM after the TCP three-way handshake. In order to make a forwarding decision, the WSM must terminate the TCP connection to examine the request.

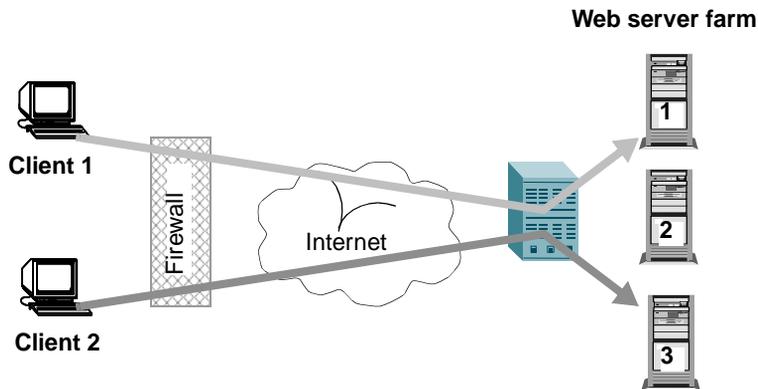
- New sessions are sent to the real server based on the metric selected (hash, roundrobin, leastconns, minmisses, response, and bandwidth).
- If no session ID is presented by the client, the WSM picks a real server based on the metric for the real server group and waits until a connection is established with the real server and a session ID is received.
- The session ID is stored in a session hash table. Subsequent connections with the same session ID are sent to the same real server. This binding is preserved even if the server changes the session ID mid-stream. A change of session ID in the SSL protocol will cause a full three-way handshake to occur.
- Session IDs are kept on the WSM until an idle time equal to the configured server time-out (a default of 10 minutes) for the selected real server has expired.

Example of SSL session ID-based persistence

Figure 173 illustrates the following:

- 1 Client 1 and Server 1 complete SSL Hello handshake through the WSM.
- 2 Server 1 assigns SSL session ID to Client 1.
- 3 The WSM records the SSL session ID.
- 4 The WSM selects a real server based on the existing SLB settings.

As a result, subsequent connections from Client 1 with the same SSL session ID are directed to Server 1.

Figure 173 SSL session ID-based persistence

- 5 Client 2 appears to the WSM to have the same source IP address as Client 1 because they share the same proxy firewall.

However, the WSM does not automatically direct Client 2 traffic to Server 1 based on the source IP address. Instead the server assigns a SSL session ID for the new traffic. Based on the configured SLB metric, the connection from Client 2 is spliced to Server 3. As a result, subsequent connections from Client 2 with the same SSL session ID are directed to Server 3.

Configuring SSL session ID-based persistence

To configure SSL session ID-based persistence for a real server:

- 1 Define each real server and assign an IP address to each real server in the server pool. See [“Configuring each real server” on page 216](#).
- 2 Define a real server group and set up health checks for the group. See [“Configuring a real server group” on page 225](#), and [“Configuring a health check for a real server group” on page 408](#).
- 3 Define virtual servers. See [“Configuring a virtual server” on page 230](#).
- 4 Define the server service on the virtual port for HTTPS (for example, port 443) and assign a real server group to service it. See [“Configuring services for a virtual server” on page 235](#).
 - In the Persistent Binding field, choose sessid.

- 5 Enable SLB on the WSM. See [“Enabling or disabling server load balancing” on page 243](#).
- 6 Enable client processing on the port connected to the client. See [“Configuring ports for server load balancing” on page 240](#).
- 7 If a proxy IP address is not configured on the client port, enable DAM for real servers. See [“Configuring direct access mode” on page 258](#).

SSL session ID-based persistence is configured for the real server.

Chapter 18

Content-intelligent switching

This section provides advanced load balancing solutions using Layer 7 content switching, including the following topics:

- “URL-based server load balancing” on page 453
- “Virtual hosting” on page 462
- “Cookie-based preferential load balancing” on page 464
- “Configuring browser-smart load balancing” on page 467
- “URL hashing for server load balancing” on page 468
- “Configuring header hash load balancing” on page 469
- “DNS load balancing” on page 470
- “Layer 7 RTSP load balancing” on page 471
- “Content-intelligent Web Cache Redirection” on page 472
- “Exclusionary string matching for real servers” on page 485
- “Regular expression matching” on page 486
- “Content precedence lookup” on page 488
- “Configuring a Layer 7 deny filter” on page 492

URL-based server load balancing

URL-based SLB directs the WSM to make load-balancing decisions on the entire path and filename of each URL.



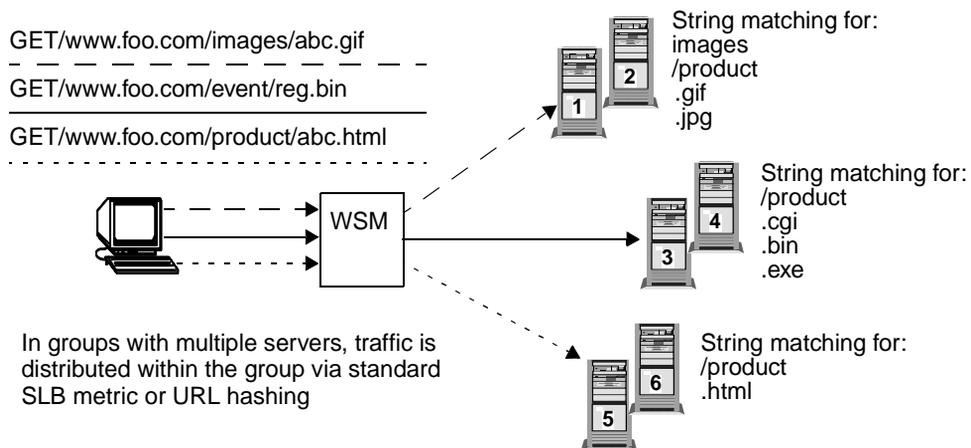
Note: The WSM supports both HTTP 1.0 and HTTP 1.1 requests.

For URL matching, you can configure up to 128 strings of 40 bytes each. Each URL Web request is then examined against the URL strings defined for each real server. URL requests are load-balanced among multiple servers matching the URL, according to the load balancing metric configured for the real server group (Least Connections is the default).

Figure 174 illustrates the delivery of the following requests:

- Requests with “.cgi” in the URL go to real servers 3 and 4.
- Requests with the string “images” in the URL go to real servers 1 and 2.
- Requests for URLs starting with “/product:” go to real servers 2, 3, and 5.
- Requests containing URLs with anything else go to real servers 1, 2, 3, and 4. These servers have been defined with the “any” string.

Figure 174 URL-based server load balancing



URL string formats

Table 143 describes how URL string formats are defined for a server to influence load balancing decisions.

Table 143 URL string formats

| String | Description |
|-------------------------|---|
| any (or no string) | If you don't configure a string for the server, or configure the string <i>any</i> , the server handles any request. |
| / (forward slash) | The server only processes requests to the root directory. For example, the server processes files in the root directory including: / /index.htm /default.asp /index.shtm |
| / <i><string></i> | The server processes requests starting with <i>/<string></i> . For example, if the server is configured with the string <i>/images</i> , the server processes requests that include: /images/product/b.gif /images/company/a.gif /images/testing/c.jpg The server will not handle requests such as: /company/images/b.gif /product/images/c.gif /testing/images/a.gif |
| <i><string></i> | The server processes requests that contain <i><string></i> . For example, if the server is configured with the string <i>images</i> , the server processes requests that include: /images/product/b.gif /images/company/a.gif /images/testing/c.jpg /company/images/b.gif /product/images/c.gif /testing/images/a.gif |

Configuring URL-based SLB

URL-based load balancing configuration requires the following tasks:

- Configuring the WSM for server load balancing. See [“Configuring virtual server load balancing” on page 215](#).
- Enabling DAM for the WSM or configuring a proxy IP address on the client port. See [“Configuring direct access mode” on page 258](#) or [“Configuring a proxy IP address for a port” on page 247](#).



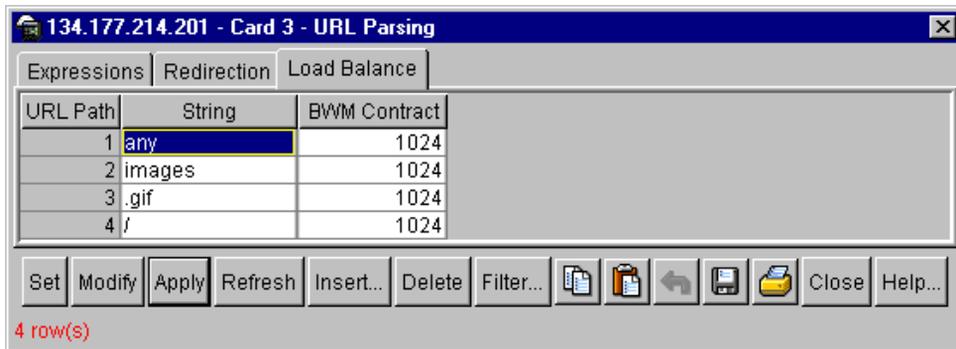
Note: When URL-based SLB is used in an active/active redundant setup, use a proxy IP address instead of Direct Access Mode (DAM) to enable URL parsing. See [“Configuring proxy IP addresses” on page 245](#).

- [“Defining a string for URL load balancing,” next](#).
- [“Configuring a real server for URL-based load balancing” on page 458](#).
- [“Enabling URL-based SLB for the server service” on page 460](#).

Defining a string for URL load balancing

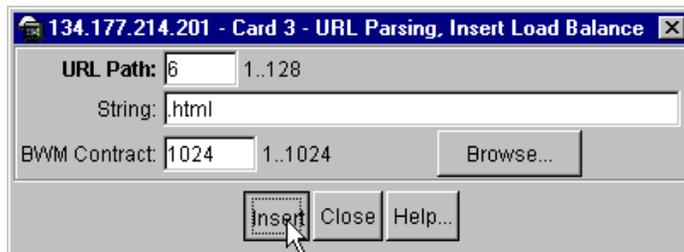
To define a [string](#) for URL load balancing:

- 1** From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2** From the Edit menu, choose WSM Card > L4 Switching > URL Parsing.
The URL Parsing dialog box opens to the Expressions tab.
- 3** Click the Load Balance tab.
The [Load Balance tab \(Figure 175\)](#) opens.

Figure 175 URL Parsing—Load Balance tab

4 Click Insert

The Insert Load Balance dialog box (Figure 176) opens with the fields defined in Table 144 on page 458.

Figure 176 URL Parsing, Insert Load Balance dialog box

- 5** In the URL Path field, type a number (1 - 128) for this URL path. This number is used for easy identification of the string for monitoring statistics.
- 6** In the **String** field, type a string of up to 40 characters for this URL path.
- 7** (Optional) In the BWM Contract field, specify a bandwidth management contract number for this URL path, or click Browse to choose from a selection list of available contracts.
- 8** Click Insert > Close.

The URL path is inserted into the Load Balance tab.

- 9** Click Apply > Close.

The URL path is configured and the URL Parsing dialog box closes.

Table 144 describes the fields on the URL Parsing—Load Balance tab and the Insert Load Balance dialog box.

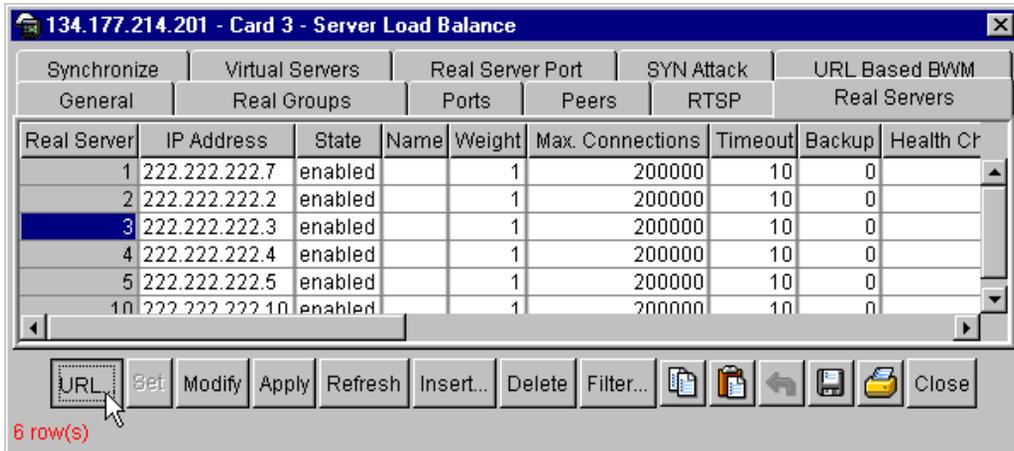
Table 144 URL Parsing—Load balance fields

| Field | Description |
|--------------|---|
| URL Path | The number of the URL Path: 1 to 128. This number is used by the real server to make a URL-based switching decision. |
| String | Sets the string for the URL path. The maximum string length is 40 characters. Any—The server can handle all URL or Web-cache requests. /—The server will only handle requests to the root directory. /<string>—The server handle the specified string. |
| BWM Contract | Sets the bandwidth management contract number: 1 to 1,024. |

Configuring a real server for URL-based load balancing

To configure a real server for URL-based load balancing:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.
The Server Load Balance dialog box opens to the General tab.
- 3 Click the Real Servers tab.
The Real Servers tab opens.

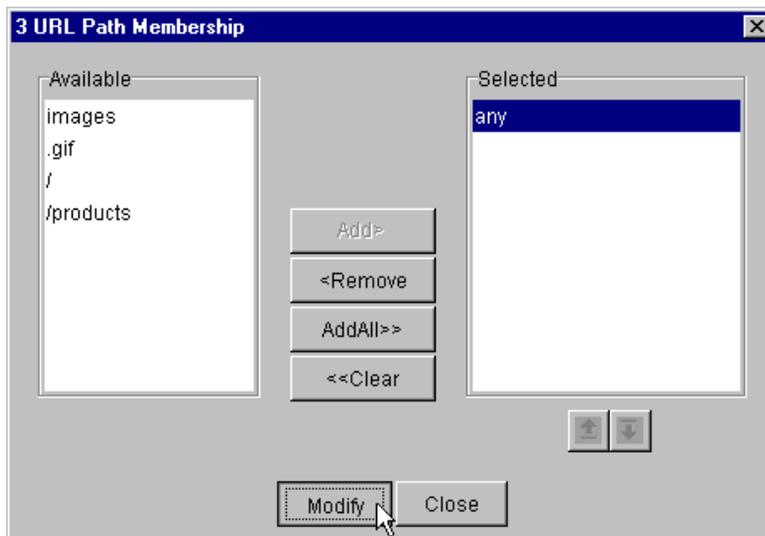
Figure 177 Configuring a real server for URL-based SLB

- Click the Real Server field for the server you want to configure.

The Real Server number is highlighted.

- Click URL.

The URL Path Membership dialog box (Figure 178) opens, listing both the available and selected URL paths for this real server.

Figure 178 URL Path Membership dialog box

- 6 From the Available list, select the URL path(s) that you want this real server to use in load balancing decisions.

The path(s) are highlighted.

- 7 Click Add.

The path(s) move from the Available list into the Selected list.

- 8 Click Modify.

The URL Path Membership dialog box closes.

- 9 In the Real Servers tab, click Apply > Close.

The URL paths are configured for the real server, and the SLB dialog box closes.

Enabling URL-based SLB for the server service

To enable URL-based SLB for the service:

- 1 From the device view, select the Web Switching Module.

The Web Switching Module is highlighted.

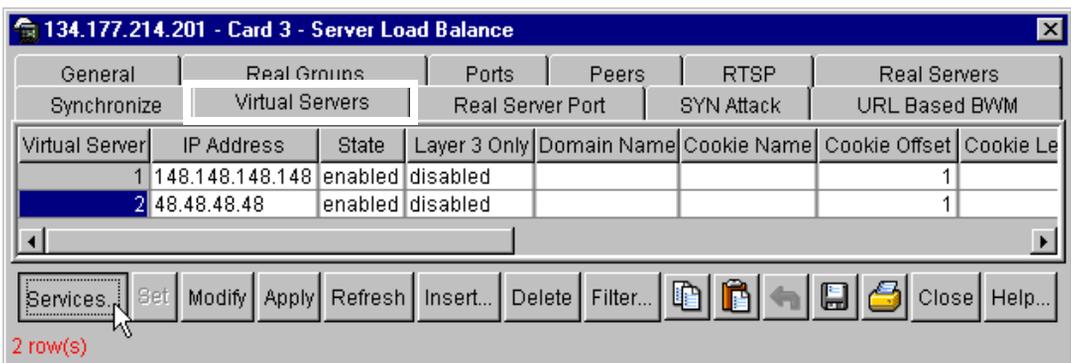
- 2 From the Edit menu, choose WSM Card > L4 Switching > SLB.

The Server Load Balance dialog box opens to the General tab.

- 3 Click the Virtual Servers tab.

The Virtual Servers tab opens.

Figure 179 Virtual Server Services dialog box



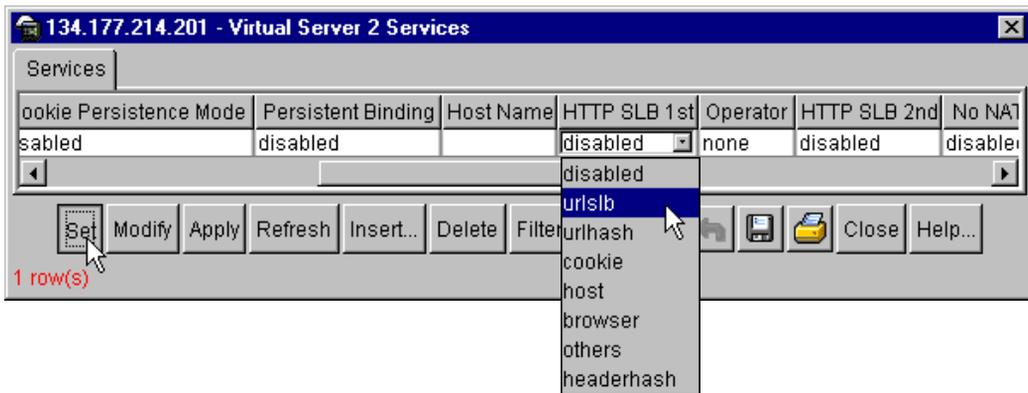
- 4 Click the Virtual Server field for the server you want to configure.

The Virtual Server number is highlighted.

- 5 Click Services.

The Virtual Server Services dialog box (Figure 179) opens listing the services configured for this virtual server. See Table 65 on page 237 for field definitions.

Figure 180 Configuring a service for URL-based SLB



- 6 For the service you want to configure, double-click in either the HTTP SLB 1st or HTTP SLB 2nd field and choose URLSLB from the selection list.

For content precedence, configure the following:

- HTTP 1st = *<method>*
- Operator = And or Or
- HTTP 2nd = *<method>*

- 7 Click Set > Apply > Close.

URL-based SLB is configured for the service, and the Virtual Server Services dialog box closes.

- 8 In the SLB dialog box, click Apply > Close.

The SLB dialog box closes.

Monitoring URL-based SLB

To view URL-based SLB statistics, including the number of hits that are load balanced because of URL path matches, see [“URL load balance statistics” on page 578](#).

Virtual hosting

Web OS lets individuals and companies have a presence on the Internet in the form of a dedicated Web site address. For example, you can have a *www.site-a.com* and *www.site-b.com* instead of *www.hostsite.com/site-a* and *www.hostsite.com/site-b*. Service providers, on the other hand, do not want to deplete the pool of unique IP addresses by dedicating an individual IP address for each home page they host. By supporting an extension in HTTP 1.1 to include the host header, Web OS enables service providers to create a single virtual server IP address to host multiple Web sites per customer, each with their own host name.



Note: For SLB, one HTTP header is supported per virtual server.

With virtual hosting:

- An HTTP/1.0 request sent to an origin server (*not* a proxy server) is a partial URL instead of a full URL.

An example of the request that the origin server would see as follows:

```
GET /products/180/ HTTP/1.0
```

```
User-agent: Mozilla/3.0
```

```
Accept: text/html, image/gif, image/jpeg
```

The GET request does not include the host name. From the TCP/IP headers, the origin server knows the requests host name, port number, and protocol.

- With the extension to HTTP/1.1 to include the HTTP HOST: header, the above request to retrieve the URL */www.nortelnetworks.com/products/180* would look like this:

```
GET /products/180/ HTTP/1.1
Host: www.nortelnetworks.com
User-agent: Mozilla/3.0
Accept: text/html, image/gif, image/jpeg
```

The Host: header carries the hostname used to generate the IP address of the site.

- Based on the Host: header, the WSM will forward the request to servers representing different customers' Web sites.
- You must define a domain name as part of the 128 supported URL strings.
- The WSM performs string matching; that is, the string *nortelnetworks.com* or *www.nortelnetworks.com* will match *www.nortelnetworks.com*.

Virtual hosting configuration overview

The process for configuring virtual hosting based on HTTP Host: headers is described below.

- 1 The network administrator defines a domain name as part of the 128 supported URL strings.

Both domain names *www.company-a.com* and *www.company-b.com* resolve to the same IP address. In this example, the IP address is for a virtual server on the WSM.
- 2 *www.company-a.com* and *www.company-b.com* are defined as URL strings.
- 3 Server group 1 is configured with servers 1 through 8.

Servers 1 through 4 belong to *www.company-a.com* and servers 5 through 8 belong to *www.company-b.com*.
- 4 The network administrator assigns string *www.company-a.com* to servers 1 through 4 and string *www.company-b.com* to servers 5 through 8.
- 5 The WSM inspects the HTTP host header in requests received from the client.
 - If the host header is *www.company-a.com*, the WSM directs requests to one of the servers 1 through 4.
 - If the host header is *www.company-b.com*, the WSM directs requests to one of the servers 5 through 8.

Configuring the host header for virtual hosting

Before configuring virtual hosting, configure the WSM for server load balancing. See [“Configuring virtual server load balancing” on page 215](#).

- 1 Enable virtual hosting on virtual service 80. See [“Configuring services for a virtual server” on page 235](#). Set the following parameter.
 - HTTP SLB 1st = Host
- 2 Define the [host names](#). See [“Defining a string for URL load balancing” on page 456](#).
- 3 Configure the real server(s) to handle the host string(s). See [“Configuring a real server for URL-based load balancing” on page 458](#).



Note: If you do not add a defined string (or add the defined string *any*), the server will handle any request.

Cookie-based preferential load balancing

Preferential services for customers can be based on cookies, ensuring that certain users are offered better access to resources than other users when site resources are scarce.



Note: Cookie-based persistent load balancing is described in [“About cookie-based persistence” on page 435](#).

Cookie-based preferential services:

- Redirect higher priority users to a larger server or server group.
- Identify a user group and redirect them to a particular server.
- Serve content based on user identity.
- Prioritize access to scarce resources on a Web site.
- Provide better services to repeat customers, based on access count.

Example of cookie-based preferential load balancing

For example, a Web server authenticates a user via a password and then sets cookies to classify customers as *Gold*, *Silver*, or *Bronze*. Servers are configured as follows:

- Real Server 1 handles gold requests.
- Real Server 2 handles silver request.
- Real Server 3 handles bronze request.
- Real Server 4 handles any request that does not have a cookie or matching cookie.

With servers defined to handle the requests listed above, here is what happens:

Table 145 Request disposition in cookie-based preferential load balancing

| Request # | Cookie | Request is forwarded to... |
|-----------|----------|---|
| 1 | None | Real Server 4 to get cookie assigned |
| 2 | Gold | Real Server 1 |
| 3 | Silver | Real Server 2 |
| 4 | Bronze | Real Server 3 |
| 5 | Titanium | Real Server 4, since it does not have an exact cookie match (matches with “any” configured at Real Server 4). |

Defining cookie-based criteria

Clients can be categorized so that some benefit from preferential service based on pre-defined criteria. One or more of the criteria in [Table 146](#) can be defined to load balance requests to different server groups.

Table 146 Examples of user categories for load balancing requests

| User category | Description |
|------------------|--|
| Individual User | An individual user is distinguished by IP address, login authentication, or permanent HTTP cookie. |
| User Communities | Some set of users, such as <i>Premium Users</i> for service providers who pay higher than normal membership fees are identified by source address range, login authentication, or permanent HTTP cookie. |
| Applications | Users are identified by the specific application they use. For example, priority is given to HTTPS traffic that is performing credit card transactions versus HTTP browsing traffic. |
| Content | Users are identified by the specific content they access. |

Configuring cookie-based preferential load balancing

Before configuring preferential load balancing, configure the WSM for server load balancing. See [“Configuring virtual server load balancing”](#) on page 215.

To configure cookie-based preferential load balancing:

- 1 Turn on URL parsing for the virtual server. See [“Configuring a virtual server”](#) on page 230. Set the following parameters on the virtual server:
 - Cookie Name
 - Cookie Offset
 - Cookie Length
 - Cookie Search
- 2 Enable cookie-based preferential load balancing on virtual service 80. See [“Configuring services for a virtual server”](#) on page 235. Set the following parameter.
 - HTTP SLB 1st = Cookie

- 3 Define the cookie values. See [“Defining a string for URL load balancing” on page 456](#).
- 4 Configure the real server(s) to handle the cookie string(s). See [“Configuring a real server for URL-based load balancing” on page 458](#).



Note: If you do not add a defined string (or add the defined string *any*), the server will handle any request.

- 5 Enable DAM for the WSM or configure a proxy IP address on the client port. See [“Configuring direct access mode” on page 258](#) or [“Configuring a proxy IP address for a port” on page 247](#).



Note: If VMA is enabled, configure a proxy IP address on ports 1-8. If VMA is disabled, you need only one proxy IP address.

- 6 Enable proxy load balancing on the port used for cookie-based preferential load balancing. See [“Configuring ports for server load balancing” on page 240](#). Set the following parameter:
 - Proxy IP Sate = Enabled



Note: If Virtual Matrix Architecture (VMA) is enabled on the WSM, you can disable proxy IP on the remaining ports.

Configuring browser-smart load balancing

HTTP requests can be directed to different servers based on browser type by inspecting the User-Agent header. For example:

```
GET /products/180/ HTTP/1.0
User-agent: Mozilla/3.0
Accept: text/html, image/gif, image/jpeg
```

Before configuring browser-smart load balancing, configure the WSM for server load balancing. See [“Configuring virtual server load balancing” on page 215](#).

To configure the WSM to perform load balancing based on browser type:

- 1 Enable browser-based load balancing on virtual service 80. See [“Configuring services for a virtual server” on page 235](#). Set the following parameter.
 - HTTP SLB 1st = Browser
- 2 Specify the browser names to be load balanced. See [“Defining a string for URL load balancing” on page 456](#).
- 3 Configure the real server(s) to handle the browser string(s). See [“Configuring a real server for URL-based load balancing” on page 458](#).



Note: If you do not add a defined string (or add the defined string *any*), the server will handle any request.

URL hashing for server load balancing

By enabling URL hashing, requests going to the same page of an origin server are redirected to the same real server or cache server.

WSM SLB hashes, by default, on the IP source address and/or IP destination address (depending on the application area) to determine content location. The WSM may not send Web requests for the same origin server to the same proxy cache server. For example, different clients requesting *http://www.nortelnetworks.com/products* may be sent to different caches.

You can direct the same URL request to the same cache or proxy server by using a virtual server IP address to load balance proxy requests. By configuring the hash or minmisses SLB metric, the WSM uses the number of bytes in the URI to calculate the hash key. If the host field exists and the WSM is configured to look into the Host: header, it calculate the hash key based on the Host: header field.

Configuring URL hashing

Before configuring URL hashing, configure the WSM for server load balancing. See [“Configuring virtual server load balancing” on page 215](#).

- 1 Enable URL hashing for virtual service 80. See [“Configuring services for a virtual server” on page 235](#). Set the following parameter.
 - HTTP SLB 1st = URLHASH
- 2 Configure the virtual server for URL hashing. See [“Configuring a virtual server” on page 230](#). Set the following parameter.
 - Hash Length = 25
- 3 Set the SLB metric for the real server group. [“Configuring a real server group” on page 225](#). Set the following parameter.
 - Metric = MinMisses or Hash

Configuring header hash load balancing

Before configuring HTTP header hashing, configure the WSM for server load balancing. See [“Configuring virtual server load balancing” on page 215](#).

To configure the WSM to hash on a selected HTTP header:

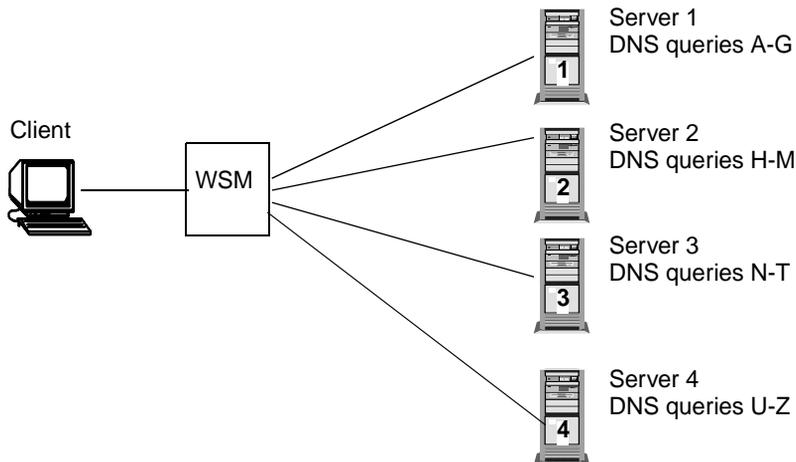
- 1 Enable header hashing for virtual service 80. See [“Configuring services for a virtual server” on page 235](#). Set the following parameter.
 - HTTP SLB 1st = HeaderHash
- 2 Configure the virtual server for HTTP header hashing. See [“Configuring a virtual server” on page 230](#). Set the following parameters.
 - Hash Length = 25
 - HTTP Header Name = User-Agent
- 3 Set the SLB metric for the real server group. [“Configuring a real server group” on page 225](#). Set the following parameter.
 - Metric = MinMisses or Hash

DNS load balancing

The Internet name registry has become so large that a single server cannot keep track of all the entries. This is resolved by splitting the registry and saving it on different servers. If you have large DNS server farms, the WSM lets you load balance traffic based on DNS names. To load balance DNS names, the host name is extracted from the query, processed by the regular expressions engine, and the request is sent to the real server.

Figure 181 shows a DNS server farm load balancing DNS queries based on DNS names. Requests having DNS names that begin with A through G go to Server 1; DNS names that begin with H through M go to Server 2; DNS names that begin with N through T go to Server 3; DNS names that begin with U through Z go to Server 4.

Figure 181 Load balancing DNS queries



Configuring DNS query load balancing

Before configuring load balancing based on DNS queries, configure the WSM for server load balancing. See [“Configuring virtual server load balancing” on page 215](#).

To configure load balancing based on DNS queries:

- 1 Enable DNS query load balancing for the DNS service (service 53). See [“Configuring services for a virtual server” on page 235](#). Set the following parameter.
 - DNS Query Load Balancing = Enabled
 - Delayed Binding = Enabled
- 2 Specify the host name [strings](#) to be load balanced. See [“Defining a string for URL load balancing” on page 456](#). For example, set the following parameters.
 - string = www.[abcdefg]*.com
 - string = www.[hijklm]*.com
 - string = www.[nopqrst]*.com
 - string = www.[uvwxyz]*.com
- 3 Configure the real server(s) to handle the host name string(s). See [“Configuring a real server for URL-based load balancing” on page 458](#).



Note: If you do not add a defined string (or add the defined string *any*), the server will handle any request.

Layer 7 RTSP load balancing

Hashing binds a client’s request to a real server for Layer 7 load balancing. As a result, all real servers carry the same content. In addition to hashing, you can segregate the requests based on the string pattern, match the strings in the requests, and direct the requests to the assigned servers. For more information on RTSP, see [“Real Time Streaming Protocol server load balancing” on page 283](#).

To configure hashing, see [“Configuring RTSP Layer 7 load balancing” on page 290](#).

Configuring RTSP SLB using pattern matching

Before configuring RTSP SLB using pattern matching, configure the WSM for server load balancing. See [“Configuring virtual server load balancing” on page 215](#).

To configure RTSP load balancing using pattern matching:

- 1 Define the URL [strings](#) to be load balanced. See [“Defining a string for URL load balancing” on page 456](#). Set the following parameter.

— string = *<string>*

- 2 Configure the real server(s) to handle the URL string(s). See [“Configuring a real server for URL-based load balancing” on page 458](#).



Note: If you do not add a defined string (or add the defined string *Any*), the server will handle any request.

- 3 Enable pattern matching for the RTSP service. See [“Configuring services for a virtual server” on page 235](#). Set the following parameter.

— RTSP Balancing = Pattern Match

[Table 147](#) describes the fields on the SLB—RTSP tab.

Table 147 SLB—RTSP tab fields

| Field | Description |
|------------|--|
| Index | The index number (1...32) assigned to each RTSP URL expression. |
| Expression | The non-cacheable RTSP URL expression. A maximum of 20 characters can be used. |

Content-intelligent Web Cache Redirection

The WSM lets you redirect Web cache requests for browser-smart load balancing, based on different HTTP header information, such as Host: header or User-Agent.

For more information on Layer 4 Web cache redirection, see [“Application Redirection” on page 395](#).

The No Cache/Cache Control for Web Cache Redirection (WCR) off loads non-cacheable content from cache servers. A Cache-Control header in an HTTP 1.1 request identifies the client's caching request. If the Cache-Control: no cache directive is enabled, HTTP 1.1 GET requests are forwarded directly to the origin servers.



Note: Origin server is the server originally specified in the request.

The HTTP 1.0 Pragma: no-cache header is equivalent to the HTTP 1.1 Cache-Control header. By enabling the Pragma: no-cache header, requests are forwarded to the origin server.



Note: For WCR, at any given time, one HTTP header is supported globally for the entire WSM.

This section includes the following topics:

- [“URL-based Web Cache Redirection” on page 473](#)
- [“Configuring HTTP header-based Web Cache Redirection” on page 484](#)

URL-based Web Cache Redirection

URL parsing for Web Cache Redirection is similar to URL-based server load balancing. For information, see [“URL-based server load balancing” on page 453](#). However, in WCR, the target of all IP/HTTP requests is a WSM virtual server.

By using URL parsing to separate static and dynamic content requests, requests with specific URLs or URL strings are sent to designated cache servers. URL-based WCR off loads overhead processing from the cache servers by only sending defined requests to the cache server farm.



Note: The WSM supports both HTTP 1.0 and HTTP 1.1 requests.

Each request is examined and handled as described below:

- If the request is a non-GET request, such as HEAD, POST, PUT, or HTTP with cookies, it is not sent to the cache.
- If the request is an ASP or CGI request or a dynamically generated page, it is not sent to the cache.
- If the request contains a Cookie, it can optionally bypass the cache.

Defining URL-based WCR expressions

You can configure up to 32 URL expressions, each 8 bytes long, for noncacheable content types. You can use up to 128 strings of 40 bytes each for URL string matching on each WSM. As each URL Web request is examined, noncacheable requests are forwarded to the origin server while requests with matching strings are redirected to a cache server.

[Table 148](#) shows examples of string expressions.

Table 148 Examples of string expressions

| Expression | Description |
|------------|---|
| /product | Any URL that starts with <i>/product</i> , including any information in the <i>/product</i> directory |
| product | Any URL that has the string <i>product</i> . |

The WSM is pre-configured with a list of 13 noncacheable expressions, which you can add, delete, or modify. Expressions are either known dynamic content file extensions or dynamic URL parameters, as described below.

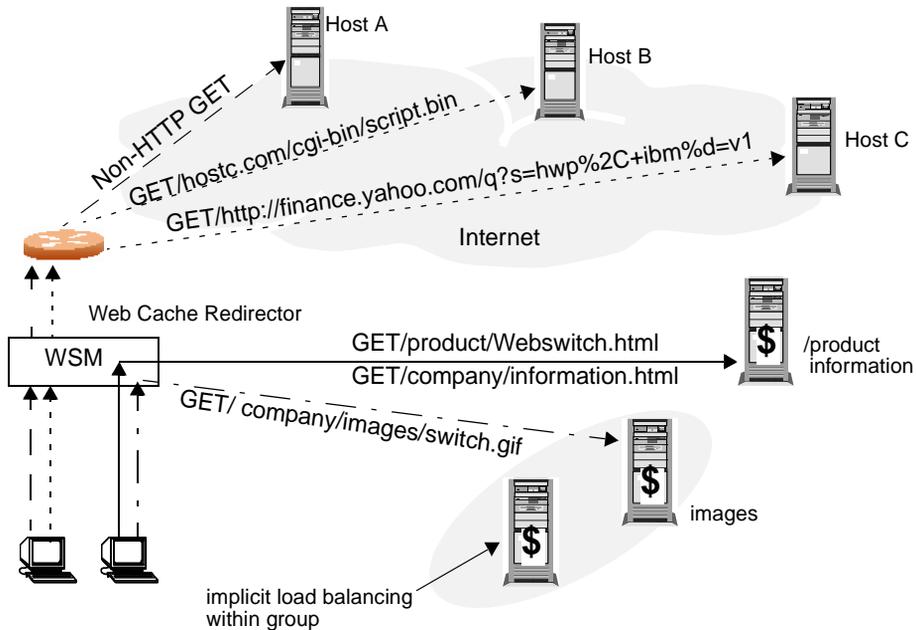
- Dynamic content files
 - Common gateway interface files (.cgi)
 - Cold fusion files (.cfm)
 - ASP files (.asp)
 - BIN directory
 - CGI-BIN directory
 - SHTML (scripted html)
 - Microsoft HTML extension files (.htx)

— Executable files (.exe)

- Dynamic URL parameters: + ! % = &

Figure 182 shows how, with WCR, requests matching the URL are load balanced among the multiple servers, depending on the metric selected for the real server group (leastconns is the default metric).

Figure 182 URL-based Web Cache Redirection



Network address translation options

Table 149 describes the types of URL-based Network Address Translation supported by WCR. For more information about NAT options, see [“Layer 4 Switching Filters—Filters fields” on page 196](#).

Table 149 Types of WCR NAT

| Type | Description |
|----------|---|
| No NAT | Traffic is redirected to the Web cache with the destination MAC address replaced by the MAC address of the cache. The destination IP address remains unchanged, and no modifications are made to the IP address or the MAC address of the source or origin server. This works well for transparent cache servers, which process traffic destined to their MAC address but with the IP address of some other device. For No NAT, choose Disabled in the Client Proxy field of the Filter dialog box. |
| Half NAT | The destination IP address is replaced by the IP address of the Web cache, and the destination MAC address is replaced by the MAC address of the Web cache. Both the IP address and the MAC address of the source remain unchanged. This is the type most commonly used. For Half NAT, choose Enabled in the Client Proxy field of the Filter dialog box. |
| Full NAT | The source IP address and the source MAC address are replaced by the IP address and MAC address of the Web cache. This method works well for proxy cache servers. For Full NAT, choose Enabled in the Client Proxy field of the Filter dialog box, and configure filtering on the SLB port. |

Configuring URL-based WCR

Use the following tasks to configure URL-based WCR.

- [“Configuring the WSM for SLB,” next.](#)
- [“Adding or removing noncacheable expressions” on page 477](#)
- [“Configuring redirection requests” on page 479](#)
- [“Defining a string for URL load balancing” on page 456.](#)
- Adding the defined string(s) to the real servers. See [“Configuring a real server for URL-based load balancing” on page 458.](#)
- Defining a real server group and adding the real servers to it. See [“Configuring a real server group” on page 225.](#)

- Configuring filters to support basic WCR. See [“Configuring filters for Web Cache Redirection” on page 482](#).
- Enabling SLB. See [“Enabling or disabling server load balancing” on page 243](#).

Configuring the WSM for SLB

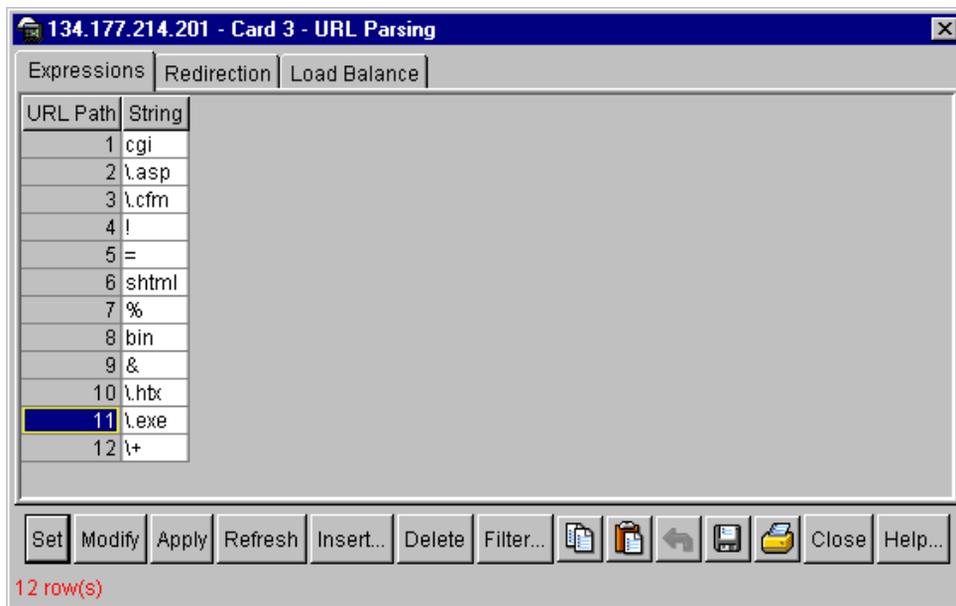
Before configuring URL-based WCR, configure the WSM for SLB using the following tasks:

- Assign an IP address to each of the real servers in the server pool. See [“Configuring each real server” on page 216](#).
- Define an IP interface on the switch. See [“Defining an IP interface on the WSM” on page 225](#).
- Define each real server. See [“Configuring each real server” on page 216](#).
- Configure the switch to support basic WCR. See [“Configuring Web Cache Redirection” on page 396](#).

Adding or removing noncacheable expressions

To add or remove expressions that should not be cacheable:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > URL Parsing.
The URL Parsing dialog box opens to the [Expressions tab \(Figure 183\)](#).

Figure 183 URL Parsing—Expressions tab

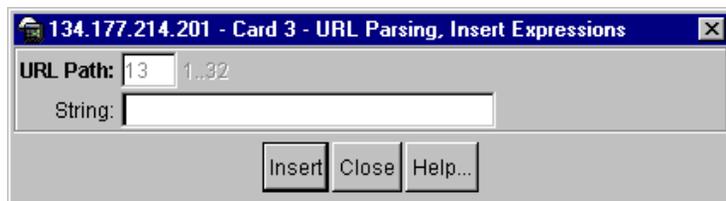
3 To delete or change an expression, do one of the following:

- To delete an expression, select it, and click Delete > Apply > Close.
- To change an expression, double-click the String field, enter the change, and click Set > Apply > Close.

The URL Parsing dialog box closes.

4 To add an expression, Click Insert

The Insert Expressions dialog box ([Figure 184](#)) opens with the fields defined in [Table 150 on page 479](#).

Figure 184 URL Parsing, Insert Expression dialog box

- 5 In the URL Path field, type a number (1 - 32) for this URL path. This number is used for easy identification of the string for monitoring statistics.
- 6 In the **String** field, type a string of up to 40 characters for this URL path.
The URL path is inserted into the Expressions tab and the Insert Expressions dialog box closes.
- 7 Click Apply > Close.
The expression is configured and the URL Parsing dialog box closes.

Table 150 describes the fields on the URL Parsing—Expressions tab and the Insert Expressions dialog box.

Table 150 URL Parsing—Expressions fields

| Field | Description |
|----------|--|
| URL Path | The index number of the URL path: 1 to 32. |
| String | <p>Strings for URL parsing. The maximum string length is 20 characters. The default strings are:</p> <p>Dynamic content files:</p> <ul style="list-style-type: none"> • Common gateway interface files (.cgi) • Cold fusion files (.cfm) • ASP files (.asp) • BIN directory • CGI-BIN directory • SHTML (scripted html) • Microsoft HTML extension files (.htx) • Executable files (.exe) <p>Dynamic URL parameters: + ! % = &</p> |

Configuring redirection requests

To configure redirection requests:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > URL Parsing.
The URL Parsing dialog box opens to the Expressions tab.

3 Click the Redirection tab.

The **Redirection tab** (Figure 185) opens with the fields defined in [Table 151](#) on page 481.

Figure 185 URL Parsing—Redirection tab

- 4** In the Redirect non-GETs to Origin Server field, click Enabled to redirect all non-GET requests to the origin server.
- 5** In the Redirect Cookies to Origin Server field, click Enabled to redirect all requests that contain cookie: in the HTTP header to the origin server.
- 6** In the Redirect no-cache to Origin Server field, click Enabled to redirect all requests that contain Cache-control: no cache in the HTTP 1.1 header or Pragma:no cache in the HTTP 1.0 header to the origin server.
- 7** Click Set > Apply > Close.

The noncacheable requests are defined and the URL Parsing dialog box closes.

Table 151 describes the fields on the URL Parsing—Redirection tab.

Table 151 URL Parsing—Redirection tab fields

| Field | Description |
|---|--|
| Redirect HTTP non-GETs to Origin Server | <p>Enables or disables redirecting packets to the origin server when HTTP non-GETs are detected.</p> <ul style="list-style-type: none"> • Enabled: WSM redirects all non-GET requests to the origin server. • Disabled: WSM compares the URI against the expression table to determine whether all non-GET requests should be redirected to a cache server or origin server. <p>Default = Enabled</p> |
| Redirect Cookies to Origin server | <p>Enables or disables redirecting packets to the origin server when user cookies are detected.</p> <ul style="list-style-type: none"> • Enabled: WSM redirects all requests that contain <i>Cookie</i>: in the HTTP header to the origin server. • Disabled: WSM compares the URI against the expression table to determine whether it should redirect all requests that contain <i>Cookie</i>: in the HTTP header to a cache server or origin server. <p>Default = Disabled</p> |
| Redirect no-cache to Origin Server | <p>Enables or disables redirecting packets to the origin server when no-cache headers are detected.</p> <ul style="list-style-type: none"> • Enabled: The WSM redirects all requests that contain <i>Cache-Control: no-cache</i> in HTTP/1.1 header, or <i>Pragma: no-cache</i> in HTTP/1.0 header to the origin server. • Disabled: the WSM compares the URI against the expression table to determine whether it should redirect requests that contain <i>Cache-Control: no-cache</i> in HTTP/1.1 header, or <i>Pragma: no-cache</i> in HTTP/1.0 header to a cache server or origin server. <p>Default = Enabled</p> |
| URI Hash Length | <p>Enables or disables URL hashing based on the URI. 0 = Disabled (default)</p> <ul style="list-style-type: none"> • Enabled: The length (0 - 255 bytes) of URI that will be used to hash into the cache server. • Disabled: WSM will only use the host header field to calculate the hash key. <p>Default= Disabled</p> |
| Redirect Based on URI Header | <p>Enables or disables URI header-based redirection.</p> |

Table 151 URL Parsing—Redirection tab fields (continued)

| Field | Description |
|------------------|--|
| HTTP Header Name | Sets the HTTP header name. The maximum string length is 32 characters. |
| Message | User-defined message that is delivered to the client when the switch cannot bind the client's request to server. The maximum string length is 64 characters. |

Configuring filters for Web Cache Redirection

Two filters are required for Web Cache Redirection:

- A filter to intercept all TCP traffic for the HTTP destination port and redirect it to the proper port in the real server group.
- A default filter to allow the noncached traffic to proceed normally.

To configure a Web Cache Redirection filter:

- 1 Create a filter that will intercept all TCP traffic for the HTTP destination port and redirect it to the port in the real server group. See [“Creating a new filter” on page 201](#). Set the parameters in the [Table 152](#).

Table 152 URL-based Web Cache Redirection filter settings

| Filter field | Setting |
|------------------------|--|
| Index | <Filter index number> |
| Name | Web Cache Redirection |
| Filter | Enabled |
| Action | Redirect |
| Client Proxy | Enabled or Disabled See “Network address translation options” on page 476 . |
| Source IP Address | Any |
| Destination IP Address | Any |
| Protocol | TCP |
| Source Port | Any |
| Destination Port | <The HTTP destination port> |

Table 152 URL-based Web Cache Redirection filter settings (continued)

| Filter field | Setting |
|-------------------|--|
| Redirection Port | <The HTTP redirection port number> |
| Redirection Group | <The Web Cache Redirection Real Server Group number> |
| URL Redirection | Enabled |

- If no client proxy is used, enable Direct Access Mode. See [“Configuring direct access mode” on page 258](#).
- Create a default filter to allow noncached traffic to proceed normally. See [“Configuring a default filter” on page 192](#). Set the parameters in [Table 153](#).

Table 153 Default filter settings for noncached traffic

| Filter field | Setting |
|------------------------|-----------------------|
| Index | <Filter index number> |
| Name | Default |
| Filter | Enabled |
| Action | Allow |
| Source IP Address | Any |
| Destination IP Address | Any |
| Protocol | Any |



Note: When the filter’s Protocol field is not set to TCP or UDP, then its Source Port and Destination Port are ignored.

- Configure the filters on the client ports. See [“Enabling or disabling filtering on a port” on page 89](#) and [“Applying filters to a port” on page 90](#).

Monitoring WCR statistics

You can view WCR statistics, including the number of hits to the cache server or origin server. See [“Redirect statistics” on page 577](#).

Configuring HTTP header-based Web Cache Redirection

Before configuring RTSP SLB using pattern matching, configure the WSM for server load balancing. See [“Configuring virtual server load balancing”](#) on page 215.

To configure HTTP header-based WCR.

- 1 Turn on URL parsing for the filter. See [“Creating a new filter”](#) on page 201. Set the following parameters.

Table 154 HTTP-based Web Cache Redirection filter settings

| Filter field | Setting |
|------------------------|----------------------------------|
| Index | <Filter index number> |
| Name | HTTP-based Web Cache Redirection |
| Filter | Enabled |
| Action | Redirect |
| Client Proxy | |
| Source IP Address | |
| Destination IP Address | |
| Protocol | |
| Source Port | |
| Destination Port | |
| Redirection Port | |
| Redirection Group | |
| URL Redirection | Enabled |

- 2 Enable header load balancing for the Host: header. See [“Configuring HTTP header-based redirection requests”](#) on page 485.
- 3 Define the host names. See [“Defining a string for URL load balancing”](#) on page 456.
- 4 Add the defined string(s) to the real servers. See [“Configuring a real server for URL-based load balancing”](#) on page 458.

- 5 Define a real server group and add the real servers to it. See [“Configuring a real server group” on page 225](#).
- 6 Configure the WSM to use the host header to determine whether requests are cacheable.

Configuring HTTP header-based redirection requests

To configure HTTP header-based redirection requests:

- 1 From the device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Edit menu, choose WSM Card > L4 Switching > URL Parsing.
The URL Parsing dialog box opens to the Expressions tab.
- 3 Click the Redirection tab.
The [Redirection tab \(Figure 185 on page 480\)](#) opens with the fields defined in [Table 151 on page 481](#).
- 4 In the HTTP Header Name field, enter the Host: name(s) to enable HTTP header-based WCR.
- 5 Click Set > Apply > Close.
The redirection requests are defined and the URL Parsing dialog box closes.

Exclusionary string matching for real servers

Exclusionary string matching lets you define a server to accept any requests regardless of the URL, except those with a specific string.

Once exclusionary string matching is enabled, clients cannot access the URL strings that are added to that real server. This means you cannot configure a dedicated server to receive a certain string and, at the same time, have it exclude other URL strings. The exclusionary feature is enabled per server, not per string.

For example, the following strings are assigned to a real server:

- string 1 = cgi

- string 2 = NOT cgi/form_A
- string 3 = NOT cgi/form_B

As a result, all cgi scripts are matched except form_A and form_B.

Configuring exclusionary URL string matching

To configure exclusionary URL string matching:

- 1 Configure the WSM for SLB. See [“Configuring virtual server load balancing” on page 215](#).
- 2 Define the strings to be excluded. See [“Defining a string for URL load balancing” on page 456](#).
- 3 Add the strings to the real server and enable exclusionary string matching on the real server. See [“Configuring a real server for URL-based load balancing” on page 458](#). Configure the following parameters:
 - Exclude String Match = Enable
 - URL Paths = *<strings>*



Caution: If you configured a string *any* and enabled the exclusion option, the server will not handle any requests. This has the same effect as disabling the server.

Regular expression matching

Regular expressions describe patterns for string matching that enable you to match the exact string, such as URLs, host names, or IP addresses. Both Layer 7 HTTP SLB and Web Cache Redirection can use regular expressions.

Standard regular expression characters

Table 155 lists standard regular expression special characters that are supported in the WSM.

Table 155 Standard regular expression characters

| Construction | Description |
|--------------|--|
| * | Matches any string of zero or more characters |
| . | Matches any single character |
| + | Matches one or more occurrences of the pattern it follows |
| ? | Matches zero or one occurrences of its followed pattern |
| \$ | Matches the end of a line |
| \ | Escape the following special character |
| [abc] | Matches any of the single characters inside the bracket |
| [^abc] | Matches any single character except those inside the bracket |

Rules for using standard regular expressions

Use the following rules when creating patterns for string matching:

- Use only one layer of parenthesis.
- Use only a single \$ (match at end of line), which must appear at the end of the string. For example, abc\$*def is not supported.
- The size of the user input string must be 40 characters or less.
- The size of the regular expression structure after compilation cannot exceed 43 bytes for load balancing strings and 23 bytes for Web Cache Redirection. The size of regular expression after compilation varies, based on regular expression characters used in the user input string.
- Use / at the beginning of the regular expression. Otherwise a regular expression will have * prefixed to it automatically. For example, html/**.htm appears as *html/**.htm.
- Incorrectly or ambiguously formatted regular expressions are rejected instantly. For example:
 - Where a + or ? follows a special character like the *
 - A single + or ? sign

— Unbalanced brackets and parenthesis

Configuring regular expressions

The regular expression feature is applicable to both URL SLB path strings and URL SLB redirected expression strings.

To configure regular expressions, use the following tasks:

- 1 [“Defining a string for URL load balancing” on page 456.](#)
- 2 [“Configuring HTTP header-based redirection requests” on page 485.](#)

Content precedence lookup

WSM Layer 7 content precedence lookup lets you give precedence to one Layer 7 request over another and decide which should be analyzed first. For the server service, you can specify which of the following types of Layer 7 content the WSM combines and examines, the order in which they are examined, and a logical operator (And/Or) for their evaluation.

- URL SLB
- HTTP Host
- Cookie
- Browsers (User agent)
- URL hash
- Header hash

Using this combination, WSM refines HTTP-based server load balancing multiple times on a single client HTTP request for binding to a server. When you combine two content types with an operator (And/Or), URL hash and Header hash are combined with host, cookie, or browser content types.

For example, the WSM can load balance the following types of content using precedence lookup:

- Virtual host and/or URL-based load balancing

- Cookie persistence and URL-based load balancing
- Cookie load balancing and/or URL-based load balancing
- Cookie persistence and HTTP SLB together in the same service
- Multiple HTTP SLB process type on the same service



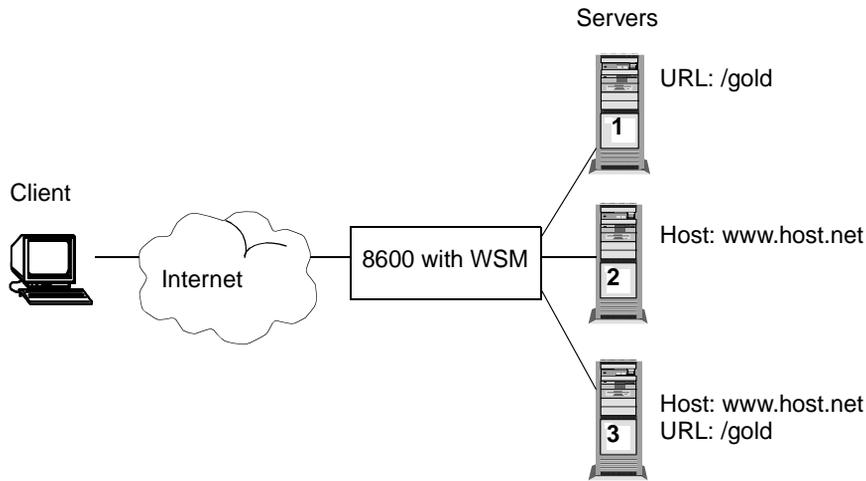
Note: Cookie persistence can also be combined with the Layer 7 content types. For more information on cookie persistence, see [“Maintaining session persistence” on page 433](#).

Content precedence lookup can be used as follows:

- If the client request is sent without a cookie and if no HTTP SLB is configured, then the WSM binds the request to the real server using normal SLB.
- If the client request is sent without a cookie, but HTTP SLB is configured on the WSM, then the request is bound to real server based on HTTP SLB.
- If the client request is sent with a cookie, and a real server associated to the cookie is found in the local session table, then the request will stay bound to that real server.

Using the operators *Or* and *And*

[Figure 186](#) shows a network with real servers 1 and 3 configured for URL SLB and real servers 2 and 3 configured for HTTP Host SLB.

Figure 186 Example: Content precedence lookup

In [Figure 186](#), with content precedence lookup, the request from the client is processed as follows, depending on the configured operator (And/Or):

- HTTP Host *or* URL SLB

The HTTP Host header takes precedence because it is specified first. If there is no Host Header information and the *Or* operator is used, the URL string is examined next.

- If a request from a client contains no Host Header but has a URL string (such as “/gold”), the request is load balanced among Servers 1 or 3.
- If a request from a client contains a Host Header, then the request is load balanced among Servers 2 and 3. The URL string is ignored because the HTTP Host was specified and matched first.

- HTTP Host *and* URL SLB

The HTTP Host header takes precedence because it is specified first. Because the *And* operator is used, both a Host Header and URL string are required. If neither is available, the request is dropped.

- If a request from a client contains a URL string (/gold) but not a Host Header, it is not served by any real server.
- If a request from a client contains a URL string (/gold) and Host Header, it is served only by real server 3.

Configuring content precedence for a service

To configure a server service for content precedence, see “[Configuring services for a virtual server](#)” on page 235. Set the following parameters:

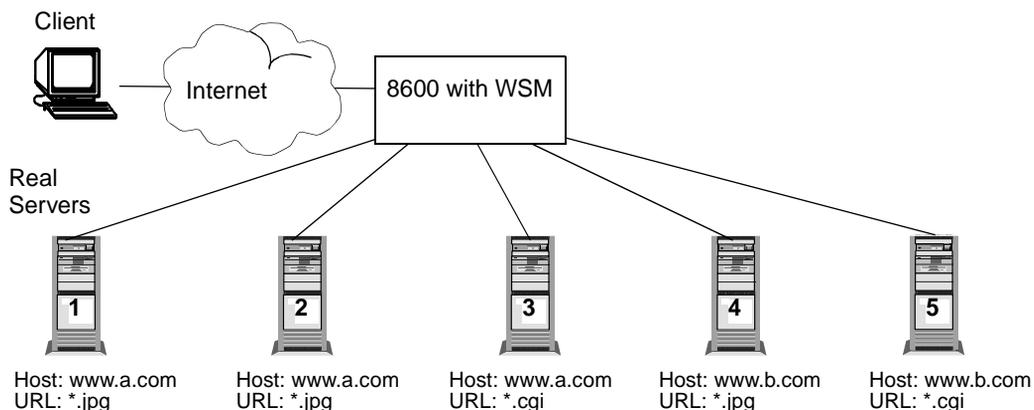
- HTTP SLB 1st = *<urlslb, HTTP Host, Cookie, Browsers (User agent), URL hash, or Header hash>*
- Operator = *<And or Or>*
- HTTP SLB 2nd = *<urlslb, HTTP Host, Cookie, Browsers (User agent), URL hash, or Header hash>*

Assigning multiple strings

In [Figure 187](#), customer A uses *www.a.com* as their domain name, and customer B uses *www.b.com*. Company C has a limited number of public IP addresses and wants to assign them conservatively. They implement virtual hosting, advertising a single virtual server IP address that includes Web sites for both company A and company B. Additionally, the hosting company assigns only one service (HTTP port 80) to support the virtual server.

The virtual hosting company wants to place different types of content on different servers. To efficiently use their servers, they separate them into two groups—their fastest servers process dynamic content (such as *.cgi* files) and their slower servers process static content (such as *.jpg* files).

Figure 187 Example: Content precedence lookup with multiple strings



In [Figure 187](#), the hosting company groups all the real servers into one real server group even though different servers provide services for different customers and different types of content ([Table 156](#)).

Table 156 Real server content

| Server | Customer | Content |
|--------|----------|--------------------|
| 1 | A | Static .jpg files |
| 2 | A | Static .jpg files |
| 3 | A | Dynamic .cgi files |
| 4 | B | Static .jpg files |
| 5 | B | Dynamic .cgi files |

When the virtual hosting company receives a client request with *www.a.com* in the Host Header and *.jpg* in the URL, it is load balanced between server 1 and server 2. For this configuration, you must assign multiple strings (a Host Header string and a URL string) for each real server.

Configuring a Layer 7 deny filter

You can secure the WSM from virus attacks by configuring a list of potentially offending string patterns (HTTP URL request). The WSM examines the HTTP content of the incoming client request for the matching string pattern. If the matching virus pattern is found, then the packet is dropped and a reset frame is sent to the offending client. SYSLOG messages and SNMP traps are generated warning operators of a possible attack.

To configure a Layer 7 deny filter, see the *Web OS Switch Software 10.0 Application Guide*, part number 212777-A.

Chapter 19

Bandwidth management

This section includes the following topics:

- [“About bandwidth management” on page 493](#)
- [“Configuring bandwidth management” on page 499](#)

About bandwidth management

Bandwidth Management (BWM) enables Web site managers to allocate a certain portion of available bandwidth for specific users or applications. It allows companies to guarantee a higher priority for critical business traffic, such as e-commerce. Traffic [classification](#) can be based on user or application information. BWM [policies](#) can be configured to set lower and upper bounds on bandwidth allocation.

Bandwidth management occurs on the egress port of the WSM—that is, the port from which the frame is leaving. However, in the case of multiple routes or trunk groups, the egress port can actually be one of several ports (from the point-of-view of where the queues are located).

Rate management is controlled by using queues for each [contract](#) and by scheduling when frames are sent from each queue. Each frame is put into a managed buffer and placed on a contract queue. The time that the next frame is supposed to be transmitted for the contract queue is calculated according to the configured rate of the contract, the current egress rate of the ports, and the buffer size set for the contract queue. The scheduler then organizes all the frames to be sent according to their time-based ordering and meters them out to the port.

Statistics and history

BWM statistics are maintained so that WSM owners can bill for bandwidth usage. You can configure statistics for frequency and count. Statistics are kept in the individual Switch Processors (SP) and then collected every second by the Management Processor (MP). The MP then combines the statistics, as statistics for some classifications may be spread across multiple SPs. The MP maintains some global statistics, such as total octets and history.

History is maintained only for the contracts for which the history option is enabled on the [Traffic Contracts](#) tab. When the history buffer of 128K is ready to overflow, it can be e-mailed to a user for long-term storage. Device Manager can be configured to [send the contents of the history buffer](#) to the user (see [“Sending BWM statistics to a user” on page 511](#)).

For more information about bandwidth management statistics, see [“Bandwidth management statistics” on page 584](#).

Policies

Up to 64 bandwidth policies can be defined for the Web Switching Module. A BWM policy establishes bandwidth limitations for a set of frames that specifies guaranteed bandwidth rates. A bandwidth policy is often based on a rate structure in which a Web host or co-location provider charges a customer for bandwidth utilization. Bandwidth is configured at one of three rates:

- Committed Information Rate (CIR)/Reserved Limit
- Soft Limit (or Best Effort)
- Hard Limit (see [Table 160 on page 504](#)).

A queue depth associated with the policy establishes the size of the queue that holds the BWM data. It can be adjusted to accommodate delay-sensitive traffic (such as audio) versus drop-sensitive traffic (such as FTP).

Bandwidth limits are usually entered in Mbps. For better [granularity](#), rates can be entered in Kbps by appending “K” to the entered number. For example, 1 Mbps can be entered as either “1” or as “1024k.”

Table 157 lists the granularity of policy limits.

Table 157 BWM policy granularity limits

| Bandwidth range | Interval | Bandwidth range | Interval |
|-----------------------|----------|-----------------------|----------|
| 250 Kbps to 5000 Kbps | 250 Kbps | 50 Mbps to 150 Mbps | 10 Mbps |
| 1 Mbps to 20 Mbps | 1 Mbps | 150 Mbps to 500 Mbps | 25 Mbps |
| 20 Mbps to 50 Mbps | 5 Mbps | 500 Mbps to 1000 Mbps | 50 Mbps |

Data pacing

Data pacing keeps individual traffic flows under control. It is based on the concept of a virtual clock and theoretical departure times (TDT). The actual calculation of the TDT is based initially on the soft limit rate. The soft limit, also called Best Effort, can be thought of as a target limit for the ISP's customer. If bandwidth is available and the classification queue is not being filled at a rate greater than the soft limit, the TDT will be met for both incoming frames and outgoing frames and no borrowing or bandwidth limitation will be necessary.

If the data is arriving more quickly than it can be transmitted at the soft limit and sufficient bandwidth is still available, the rate is adjusted upwards based on the depth of the queue, until either the rate is fast enough to reduce the queue depth or the hard limit is reached.

If the data cannot be transmitted at the soft limit, then the rate is adjusted downward until the data can be transmitted or the committed information rate (CIR) is hit. If the CIR is overcommitted among all the contracts configured for the WSM, graceful degradation will reduce each CIR until the total bandwidth allocated fits within the total bandwidth available.

Each BWM contract is assigned a bandwidth policy index and (optionally) a name. Contracts can be enabled and disabled. For frames qualifying for multiple classifications, precedence of contracts is also specified per contract. If no precedence is specified, the default order is used.

- When both filter Types-Of-Service (TOS) and BWM TOS are applied, BWM TOS has precedence.

- BWM configurations will optionally be synchronized during VRRP synchronization except the port contracts and VLAN contracts, which will not be synchronized.

Contracts

You can create up to 1024 bandwidth management contracts to limit individual traffic flows. Each contract consists of:

- A classification policy grouping certain frames together
- A bandwidth policy specifying usage limitations applied to these frames

Contract classifications

[Table 158](#) shows how you can classify BWM contracts.

Table 158 Bandwidth management classifications

| Classification | Description |
|----------------|---|
| Server output | <ul style="list-style-type: none">• Physical Port - All frames from a specified physical port.• VLAN - All frames from a specified VLAN. Even if a VLAN translation occurs, the bandwidth policy is based on the ingress VLAN.• IP Source Address - All frames with a specified IP source address or range of addresses defined with a subnet mask.• IP Destination Address - All frames with a specified IP destination address or range of addresses defined with a subnet mask.• WSM services on the virtual servers• Layer 4 services, including a single virtual server, a group of virtual servers, or a service for a particular virtual server.• Layer 7 services, including single URL path, a group of URL paths, or a single cookie, |

Table 158 Bandwidth management classifications (continued)

| Classification | Description |
|----------------|---|
| Application | Applications can be identified and grouped, such as: <ul style="list-style-type: none"> • TCP Port Number - All frames with a specific TCP source or destination port number • UDP - All UDP frames • UDP Port Number - All frames with a specific UDP source or destination port number |
| Combinations | Combinations group items together into a contract. For example, if you wanted to have three different virtual servers associated with a contract, you would specify the same contract index on each of the three virtual server IP addresses. You can also combine filters in this manner |

Contract precedence

You can specify per-contract precedence or allow default ordering. If a contract is not assigned a precedence value, then the default ordering is:

- 1 Layer 7 application (URL, HTTP header, cookie, and so forth)
- 2 Layer 4 service on the virtual server
- 3 Filter
- 4 VLAN
- 5 Source Port/Default Assignment

URL-based BWM

URL-based BWM lets the network administrator or Web site manager control bandwidth based on URL, HTTP header, or cookie. All three types of BWM are accomplished by following the configuration guidelines on content switching. You must assign a contract to each defined string, where the string is contained in a URL, an HTTP header, or a cookie. BWM based on URLs lets you:

- Allocate bandwidth based on the type of request.

The WSM allocates bandwidth based on certain strings in the incoming URL request. For example, if a Web site has 10Mbs of bandwidth, the site manager can allocate 1 Mbs of bandwidth for static HTML content, 3Mbs of bandwidth for graphic content and 4Mbs of bandwidth for dynamic transactions, such as URLs with `cgi-bin` requests and `.asp` requests.

- Prioritize transactions or applications.

By allocating bandwidth, the WSM can guarantee that certain applications and transactions get better response time.

- Allocate bandwidth for requests that can be cached.

Users can allocate a certain percentage of bandwidth for Web cache requests by using URL parsing and bandwidth management.

HTTP header-based BWM

HTTP header-based BWM allows Web site managers to allocate bandwidth based on header value. Thus, they can allocate bandwidth based on browser type, cookie value, and so forth.

Cookie-based BWM

Cookie-based BWM enables Web site managers to prevent network abuse by bandwidth-hogging users. Bandwidth can be allocated by type of user or other user-specific information available in the cookie. Cookie-based bandwidth management empowers service providers to create tiered services. For example, Web site managers can classify users as first class, business class, and coach and allocate a larger share of the bandwidth for preferred classes.



Note: Cookie-based BWM does not apply to cookie-based persistency or cookie passive/active mode applications.

Frame Discard

When packets in a contract queue have not yet been sent and the buffer size set for the queue is full, any new frames attempting to be placed in the queue will be discarded.

Packet coloring (TOS bits) for burst limit

Whenever the soft limit for the traffic policy type of service is exceeded, optional packet coloring can be done to allow downstream routers to use diff-serv mechanisms (that is, writing the Type-Of-Service (TOS) byte of the IP header) to delay or discard these out-of-profile frames. Frames that are not out-of-profile are marked with a different, higher priority value. This feature can be enabled or disabled on a per-contract basis, by enabling/disabling overwriting IP TOS on the [Traffic Contracts tab](#). The actual values used by the WSM for overwriting TOS values (depending on whether traffic is over or under the soft TOS limit) are set on the [Traffic Policies tab](#). The values allowed are 0-255. Typically, the values specified should match the appropriate diff-serv specification but could be different, depending on the customer environment.

Using Virtual Matrix Architecture with BWM

Virtual Matrix Architecture (VMA) is recommended when Bandwidth Management is enabled. If the traffic classification is filter-based or SLB traffic, then the classification occurs on the designated port.

When VMA is enabled, traffic classification that is not based on filters or Server Load Balancing (SLB) is done on the ingress port—that is, the port on which the frame is received (not the client port or the server port).

When VMA is not enabled, bandwidth classification is done on the ingress side of the WSM (at the ingress port or designated port) and can be based on the following: source port, VLAN, filters, Virtual Internet Protocol (VIP) address, service on the Virtual server, URL, and so on.

For more information about Virtual Matrix Architecture, see [Chapter 11, “Improving WSM performance with VMA,”](#) on page 323.

Configuring bandwidth management

This section contains general information about configuring bandwidth management for the Web Switching Module. It contains the following topics:

- [“Configuration prerequisites for BWM”](#) on page 500

- [“Enabling bandwidth management” on page 500](#)
- [“Configuring BWM policy” on page 502](#)
- [“Configuring a BWM contract” on page 505](#)
- [“Configuring BWM classification” on page 508](#)

Configuration prerequisites for BWM

Bandwidth Management requires a software key. There are two operational keys for BWM—a standard key and a demo key. The demo key automatically expires after a demo time period. These keys may only be enabled if Layer 4 services have been enabled.

Before you begin the configuration tasks for bandwidth management, you must complete the following tasks to configure the WSM for SLB:

- [“Configuring each real server” on page 216](#)
- [“Defining an IP interface on the WSM” on page 225](#)
- [“Configuring a real server group” on page 225](#)
- [“Configuring a virtual server” on page 230](#)
- [“Configuring ports for server load balancing” on page 240](#)
- [“Enabling or disabling server load balancing” on page 243](#)

For more information about server load balancing, see [“Server load balancing basics” on page 203](#).

Enabling bandwidth management

To enable bandwidth management on the Web Switching Module:

- 1** From the Device Manager main window, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2** Select Edit > WSM Card > BWM.
The Bandwidth Management dialog box opens to the [General tab](#).

Figure 188 Bandwidth management—General tab

- 3 In the Bandwidth Management field, click On.
- 4 Click Set > Apply.
Bandwidth Management is enabled and the dialog box closes.
- 5 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

Table 159 describes the Bandwidth Management dialog box General tab.

Table 159 Bandwidth Management—General tab fields

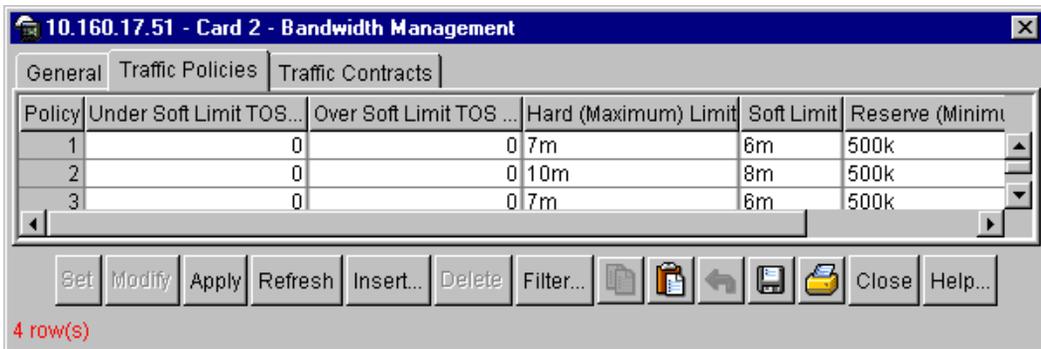
| Field | Description |
|----------------------|---|
| Bandwidth Management | Sets the current state of BWM: on or off. The default is off. |
| Enforce BWM Policies | Enables or disables the BWM enforce policy. |
| SMTP User Name | The email address where the mail gateway will send BWM History. The maximum string length is 127 characters. The default is none. |
| Send BWM History | Sends the contents of the history buffer (BWM statistics) to a user when <i>SMTP User Name</i> (above) is configured. |

Configuring BWM policy

To configure bandwidth management (BWM) policy:

- 1 From the Device Manager main window, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 Select Edit > WSM Card > BWM.
The Bandwidth Management dialog box opens to the **General** tab.
- 3 Click the Traffic Policies tab.
The **Traffic Policies** tab (Figure 189) opens.

Figure 189 Bandwidth management—Traffic Policies tab



- 4 To Insert a new policy, click Insert.
The Insert Policy dialog box opens.

Figure 190 Insert Traffic Policies dialog box

- 5 In the Policy field, type the bandwidth policy number.
Each policy must have a unique number from 1 to 64.
- 6 In the Limit fields, type the Hard, Soft, and Reserve limits for the policy.
Typically, charges are applied for burst rates between the soft and hard limit. Each limit must be set between 256K-1000M.
- 7 (Optional) In the Soft Limit TOS Value fields, type the TOS byte values, between 0-255, for the policy under limit and over limit.

There are two parameters for specifying the TOS bits: under limit and over limit. These TOS values are used to overwrite the TOS values of IP packets if the traffic for a contract is under or over the soft limit, respectively. These values only have significance to a contract if TOS overwrite is enabled.

You should set the TOS values considering their greater impact on the downstream routers.
- 8 In the buffer field, type the buffer limit for the policy.
Set a value between 8192-128000 bytes. The buffer depth for a BWM contract should be set to a multiple of the packet size. Keep in mind that the total buffer limit for the Bandwidth Management policy is 128K.
- 9 Click Insert.

The Traffic Policies dialog box closes and the BWM policy is inserted into the Traffic Policies tab.
- 10 In the Traffic Policies tab, click Set > Apply.

The BWM policy is configured and the Bandwidth Management dialog box closes.

11 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 160](#) describes the fields on the Bandwidth Management Traffic Policies tab and Insert Traffic Policies dialog box.

Table 160 Traffic Policies fields

| Field | Description |
|----------------------------|---|
| Policy | The unique index number of the BWM traffic policy: 1 to 64. |
| Under Soft Limit TOS Value | The profile state of IP Type of Service (TOS) in the traffic policy: 0 to 255. The default is 0. |
| Over Soft Limit TOS Value | The profile state of IP Type of Service (TOS) in the traffic policy: 0 to 255. The default is 0. |
| Hard (Maximum) Limit | The hard speed limit for the traffic policy: 250k to 5,000k or 1M to 1,000M. The default is 2,000k. This is a “never exceed” rate. A bandwidth class is never allowed to transmit above this rate. Typically, traffic bursts between the soft limit and the hard limit are charged a premium. The maximum hard limit for a bandwidth policy is 1 Gbps, even when multiple Gigabit ports are trunked. |
| Soft Limit | The soft speed limit for the traffic policy: 250k to 5,000k or 1M to 1,000M. The default is 1,000k. This is the desired bandwidth rate, that is, the rate the customer has agreed to pay on a regular basis. When output bandwidth is available, a bandwidth class will be allowed to send data at this rate. No exceptional condition will be reported when the data rate does not exceed this limit. |

Table 160 Traffic Policies fields (continued)

| Field | Description |
|-------------------------|--|
| Reserve (Minimum) Limit | The reservation speed limit for the traffic policy: 250k to 5,000k or 1M to 1,000M. The default is 500k. Also called Committed Information Rate (CIR), this is a rate that a bandwidth classification is always guaranteed. In configuring BWM contracts, ensure that the sum of all committed information rates never exceeds the link speeds associated with ports on which the traffic is transmitted. In a case where the total CIRs exceed the out-bound port bandwidth, the WSM will perform a graceful degradation of all traffic on the associated ports. |
| Buffer Limit | The buffer limit for the traffic policy: 8,192 to 128,000 bytes. The default is 16,320. |

Configuring a BWM contract

To configure a BWM contract:

- 1** From the Device Manager main window, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2** Select Edit > WSM Card > BWM.
The Bandwidth Management dialog box opens to the [General tab](#).
- 3** Click the Traffic Contracts tab.
The [Traffic Contracts tab](#) opens.

Figure 191 Bandwidth management—Traffic Contracts tab

| Contract | Name | State | Policy | Precedence | TOS Overwrite | Keep Statistics on TFTP Server |
|----------|-----------|---------|--------|------------|---------------|--------------------------------|
| 1 | port7_8 | enabled | 1 | 1 | disabled | enabled |
| 2 | filter | enabled | 1 | 1 | disabled | enabled |
| 3 | service80 | enabled | 2 | 1 | disabled | enabled |
| 4 | vlan | enabled | 4 | 5 | disabled | enabled |
| 5 | trunk | enabled | 1 | 10 | disabled | enabled |
| 1024 | Default | enabled | 64 | 0 | disabled | enabled |

6 row(s)

4 To Insert a new contract, click Insert.

The [Insert Traffic Contracts dialog box](#) (Figure 192) opens.

Figure 192 Insert Traffic Contracts dialog box

10.160.17.51 - Card 2 - Bandwidth Management, Insert Traffic Contracts

Contract: 1..1024 (reserved: last index)

Name:

State: enabled disabled

Policy: 1..64

Precedence: 1..255

TOS Overwrite: enabled disabled

Keep Statistics on TFTP Server: enabled disabled

5 In the Contract field, type the Traffic Contract number.

Each contract must have a unique number from 1 to 1024.

6 (Optional) In the Name field, type a name for the contract of up to 15 characters in length.

7 (Optional) In the Precedence field, type a precedence value for the BWM contract.

- 8 (Optional) To enable TOS overwriting for the BWM contract, click enable in the TOS Overwrite field.
- 9 In the Policy field, select a policy for this contract. To view the available policies, click Browse.
Each bandwidth management contract must be assigned a bandwidth policy.
- 10 (Optional) In the Keep Statistics on TFTP Server field, click Enabled to save the statistics. This field is enabled by default.
- 11 In the State field, click Enabled to enable the BWM contract.
- 12 Click Insert.
The BWM traffic contract is inserted into the Traffic contracts tab and the Insert Traffic Contracts dialog box closes.
- 13 Click Set > Apply > Close.
The BWM traffic contract is configured and the Bandwidth Management dialog box closes.
- 14 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:
 - Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
 - Save configuration (no backup)—Saves changes to the active WSM configuration block.

[Table 161](#) describes the fields on the Traffic Contracts tab and the Insert Traffic Contracts dialog box.

Table 161 Traffic Contracts fields

| Field | Description |
|----------|---|
| Contract | The contract number of the Bandwidth Management (BWM) traffic contract: 1 to 1,024. |
| Name | The traffic contract name. The maximum string length is 15 characters. The default is none. |
| State | Enables or disables the state of the traffic contract. |
| Policy | Sets the policy number of the traffic contract: 1 to 64. The default is 64. |

Table 161 Traffic Contracts fields (continued)

| Field | Description |
|--------------------------------|---|
| Precedence | <p>Sets the precedence value of the traffic contract: 1 to 255. The default is 1.</p> <p>Each contract can have a precedence value from 1-255. The higher the number, the higher the precedence. If a frame is applicable to different classifications, then the contract with higher precedence will be assigned to the frame. If the precedence is the same for the applicable contracts, then contracts will be assigned to the frame in the following order:</p> <p>(1) URL/Cookie, (2) Service on the Virtual server, (3) Filter, (4) VLAN, (5) Source Port/Default.</p> |
| TOS Overwrite | Enables or disables Type of Service (TOS) overwriting for the traffic contract. The default is disabled. |
| Keep Statistics on TFTP Server | <p>Enables or disables saving contract statistics on the Trivial File Transfer Protocol (TFTP) server of the traffic contract. The default is disabled.</p> <p>The total number of octets sent, octet discards, and times over the soft limit are kept for each contract. The history buffer maintains the average queue size for the time interval and the average rate for the interval.</p> |

Configuring BWM classification

Each BWM contract is assigned a classification policy. The classification can be based on a filter or service(s) on the Virtual server. Filters are used to create classification policies based on the IP source address, IP destination address, TCP port number, UDP, and UDP port number.

Any item configured with a filter can be used for bandwidth management. To associate a particular classification with a contract, enter the contract number into the Device Manager tab's contract field as shown in [Table 162](#).

Table 162 Configuring bandwidth management classification

| BWM classification item | Where to configure in Device Manager |
|--|--|
| IP destination address IP source address TCP port number | Edit > WSM Card > L4 Switching > Filters > Filters tab See " Creating a new filter " on page 201. |
| UDP | Edit -> WSM Card -> L4 Switching -> Filters -> Protocol = 17 See " Creating a new filter " on page 201. |

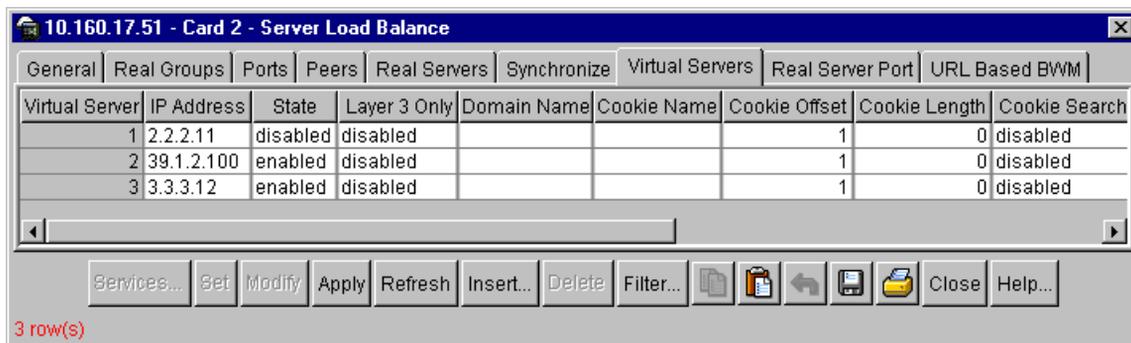
Table 162 Configuring bandwidth management classification (continued)

| BWM classification item | Where to configure in Device Manager |
|-------------------------|---|
| UDP Port | Edit -> WSM Card -> L4 Switching -> Filters -> Low Source Port = number (ex 53 for DNS) High Source Port = number (ex 53 for DNS) Low Destination Port = number (ex 53 for DNS) High Destination Port = number (ex 53 for DNS) See "Creating a new filter" on page 201 . |
| Virtual server | Edit > WSM Card > L4 Switching > SLB > Virtual Servers tab See "Configuring a virtual server" on page 230 . |
| Port | Edit > WSM Card > Port > Port tab Note: A BWM contract cannot be applied to a single port in a trunk group. It must be applied to the trunk group itself. See "Setting port parameters" on page 85 . |
| Trunk | Edit > WSM Card > General > Trunks tab See "Configuring a trunk group" on page 98 . |
| VLAN | Edit > WSM Card > Virtual LANs > VLAN Membership tab See "Configuring a VLAN" on page 106 . |
| URL path | Edit > WSM Card > L4 Switching > Filters > URL BWM Filtering tab Edit > WSM Card > L4 Switching > SLB > URL based BWM tab Edit > WSM Card > L4 Switching > URL parsing > Load Balance tab |

Assigning a BWM contract to a virtual server

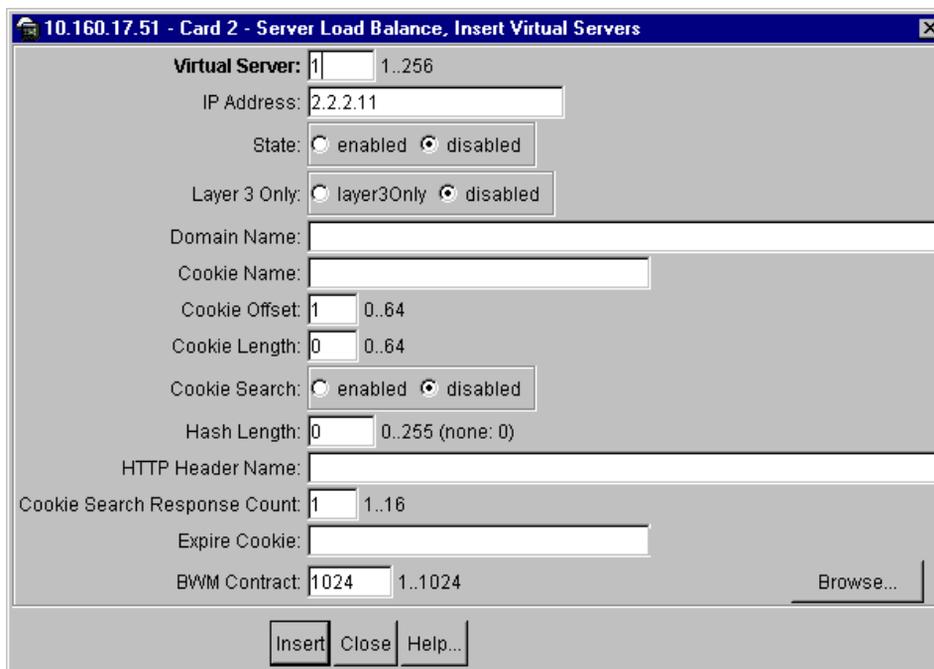
To assign a BWM contract to a virtual server:

- 1 From the Device view, select the Web Switching Module.
 The Web Switching Module is highlighted.
- 2 Select Edit > WSM Card > L4 Switching > SLB.
 The Server Load Balance dialog box opens to the [General tab](#).
- 3 Click the Virtual Servers tab.
 The Virtual Servers tab opens.

Figure 193 Server Load Balance—Virtual Servers tab

- 4 From the Virtual Servers tab, select the virtual server to which you want to assign a BWM contract.
- 5 Click Modify.

The Insert Virtual Servers dialog box opens.

Figure 194 Insert Virtual Servers dialog box

6 Click **Browse**.

The Virtual Server BWM Selection dialog box opens.

7 Choose a contract from the selection list and click **Modify**.

The BWM contract is assigned and the dialog box closes.

8 In the **Insert Virtual Servers** dialog box, click **Insert**.

The BWM contract is inserted into the Virtual Servers tab and the **Insert Virtual Servers** dialog box closes.

9 From the Virtual Servers tab, click **Set > Apply > Close**.

The BWM contract is configured for the virtual server and the **Server Load Balance** dialog box closes.

10 To save the change, right-click the WSM card and choose one of the following from the shortcut menu:

- **Save configuration (backup)**—Saves changes to the active and backup WSM configuration blocks.
- **Save configuration (no backup)**—Saves changes to the active WSM configuration block.

Sending BWM statistics to a user

To send bandwidth management statistics to a user:

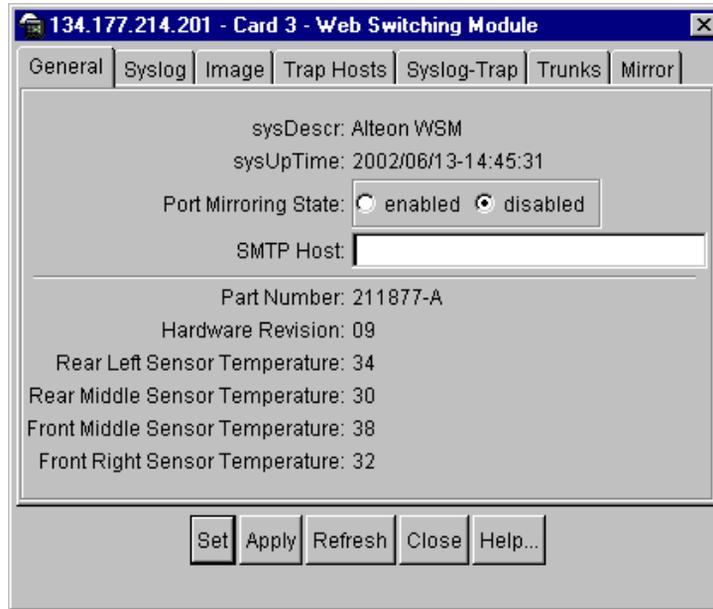
1 From the **Device** view, select the **Web Switching Module**.

The **Web Switching Module** is highlighted.

2 From the **Edit** menu, choose **WSM Card > General**.

The **Web Switching Module** dialog box opens to the [General tab](#).

See [“General tab fields” on page 186](#) for a description of the fields on this tab.

Figure 195 Web Switching Module—General tab

3 In the SMTP Host field, type the domain name or IP address of an SMTP mail gateway (the mail relay agent that sends the email).

4 Click Set > Apply > Close.

The SMTP host is configured and the Web Switching Module dialog box closes.

5 From the Edit menu, choose WSM Card > BWM.

The Bandwidth Management dialog box opens to the [General tab](#) (Figure 196).

Figure 196 Bandwidth management—General tab

6 In the SMTP User Name field, type the user's email address.

7 Click Set > Apply.

The email address where the statistics will be sent is configured.

8 Click Send BWM History.

The contents of the bandwidth management history buffer are sent to the user.

9 Click Close.

The Bandwidth Management dialog box closes.

10 To save your changes, right-click the WSM card and choose one of the following from the shortcut menu:

- Save configuration (backup)—Saves changes to the active and backup WSM configuration blocks.
- Save configuration (no backup)—Saves changes to the active WSM configuration block.

The configuration for sending bandwidth management statistics to a user is saved.

SLB—URL-Based BWM fields

Figure 197 Server Load Balance—URL-based BWM fields

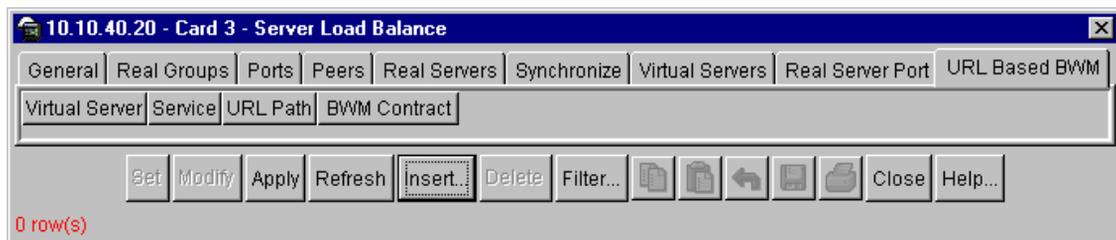


Table 163 describes the SLB—URL-Based BWM fields.

Table 163 SLB—URL-Based BWM fields

| Field | Description |
|----------------|--|
| Virtual Server | Number of the virtual server: 1 to 256. Click the Browse... button to select an available virtual server. |
| Service | Number of the virtual service: 1 to 8. Click the Browse... button to select an available virtual service. |
| URL Path | Number of the URL path: 1 to 8. Click the Browse... button to select an available URL path. |
| BWM Contract | Number of the bandwidth management contract: 1 to 1,024. Click the Browse... button to select an available BWM Contract. |

Chapter 20

Working with Statistics

This section describes the Device Manager statistics gathering and analysis functions, and includes the following topics:

- [“About statistics and graphing” on page 516](#)
- [“WSM blade statistics” on page 518](#)
- [“Bridge forwarding statistics” on page 536](#)
- [“Port statistics” on page 540](#)
- [“IP routing statistics” on page 551](#)
- [“Virtual routing statistics” on page 563](#)
- [“Layer 4 statistics” on page 565](#)
- [“Bandwidth management statistics” on page 584](#)

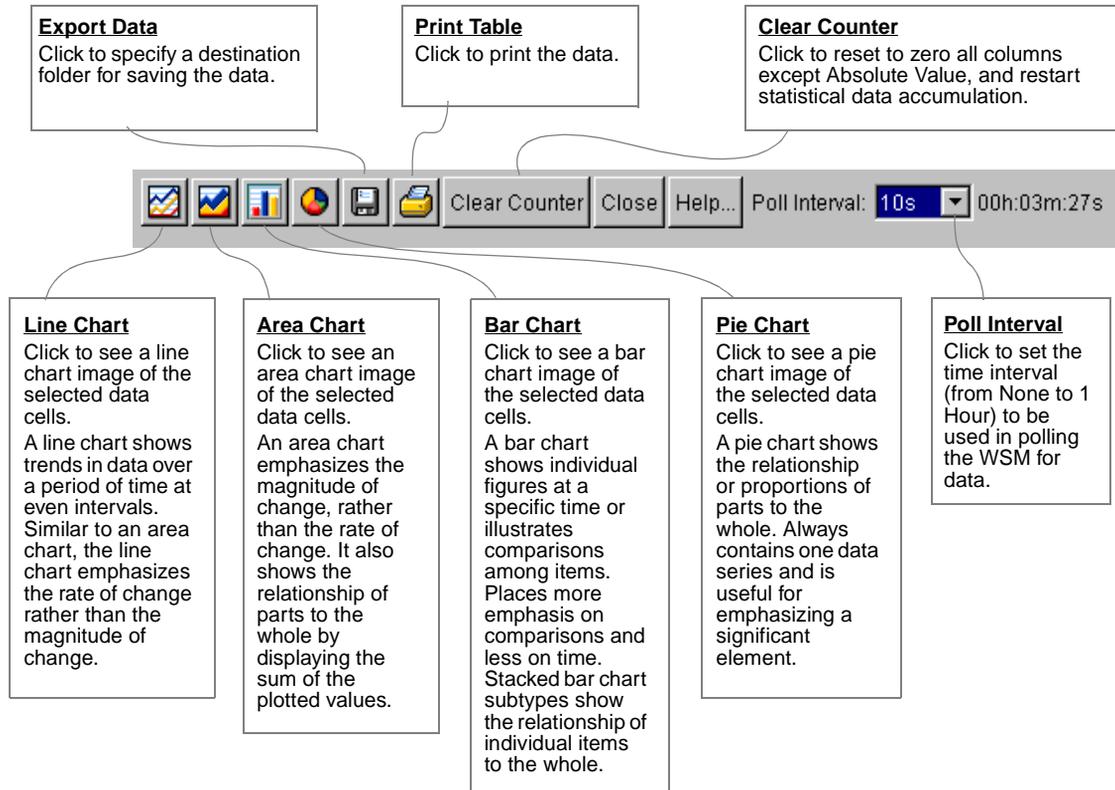
About statistics and graphing

Device Manager accumulates statistics for the WSM and ports. Some statistical windows let you view, export, and print statistical charts. You can also view and print visual representations of data relationships. This section describes the Device Manager data and graphing functions and defines the statistics collected.

Working with data tables

Figure 198 shows the tools in Device Manager data tables.

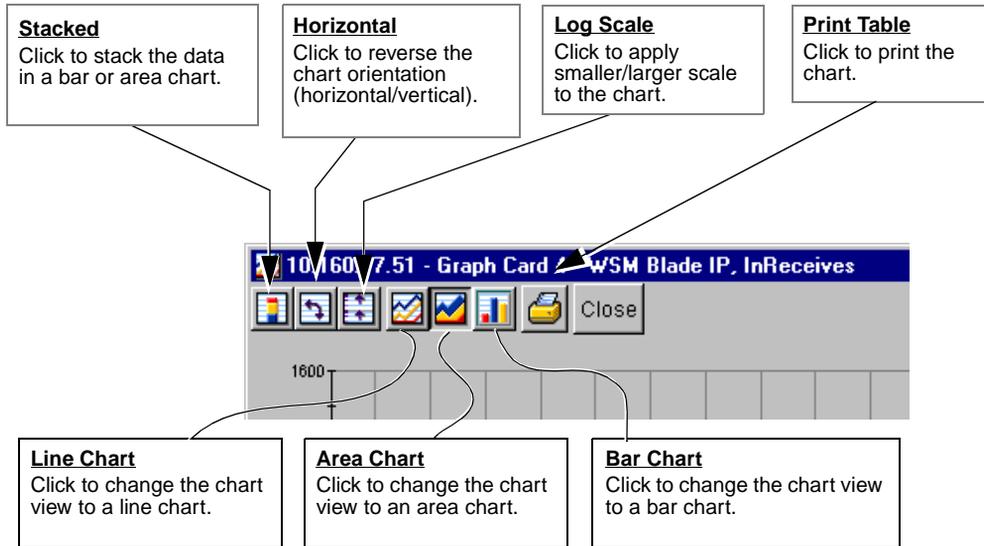
Figure 198 Statistics window toolbar



Working with charts

Figure 199 shows the tools in Device Manager chart windows.

Figure 199 Chart toolbar



WSM blade statistics

WSM blade statistics include:

- “IP statistics”, next
- “Received ICMP statistics” on page 521
- “Sent ICMP statistics” on page 523
- “RIP statistics” on page 525
- “SNMP statistics” on page 526
- “WebOS statistics” on page 529
- “Statistics” on page 531
- “MP CPU statistics” on page 533
- “RADIUS account statistics” on page 535

IP statistics

To graph IP statistics:

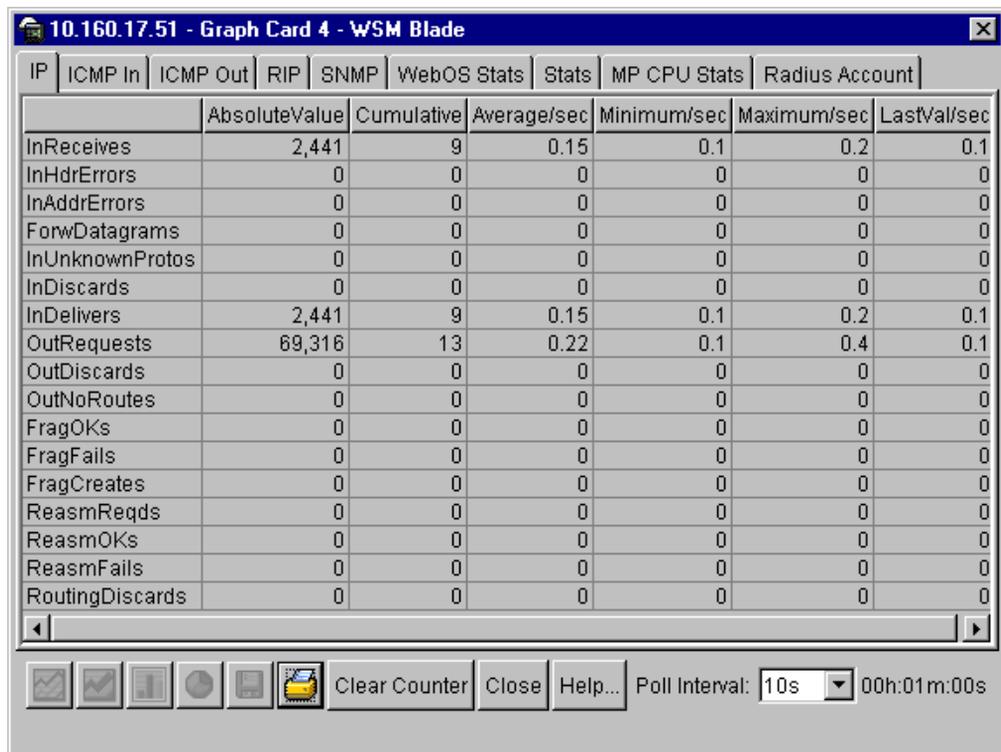
- 1 From the device view, select the Web Switching Module.

The module is highlighted.

- 2 From the Graph menu, select WSM Card > General.

The WSM Blade dialog box opens to the **IP tab** (Figure 200) opens displaying the fields described in Table 164.

Figure 200 WSM Blade—IP tab



- 3 Click a cell(s) to graph.
- 4 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing”](#) on page 516.

Table 164 describes the fields on the WSM Blade—IP tab.

Table 164 WSM Blade—IP tab fields

| Field | Description |
|----------------------------------|---|
| In Receives | The number of input datagrams received from interfaces, including those received in error. |
| In Header Errors | The number of input datagrams that were discarded because of errors in the IP headers. Errors: bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, and so on. |
| In Address Errors | The number of input datagrams that were discarded because the IP address in the IP header's destination field was not a valid address at this WSM. Invalid addresses: 0.0.0.0, addresses of unsupported Classes such as Class E, and so forth. For entities that are not IP Gateways that do not forward datagrams, the count includes datagrams that were discarded because the destination address was not a local address. |
| Forward Datagrams | The number of input datagrams for which this entity was not their final IP destination. An attempt was made to find a route to forward them to that final destination. In entities that do not act as IP Gateways, the count only includes packets that were Source-Routed via this entity, and that the Source-Route option processing was successful. |
| Packets In with Unknown Protocol | The number of locally-addressed datagrams that were received successfully, but were discarded because of an unknown or an unsupported protocol. |
| In Discards | The number of input IP datagrams that were discarded, although no errors were identified. This can occur because of insufficient buffer space. This counter does not include any datagrams that were discarded while waiting for re-assembly. |
| In Delivers | The total number of input datagrams successfully delivered to IP user-protocols, including ICMP. |
| Out Requests | The total number of IP datagrams that local IP user-protocols, including ICMP, supplied to IP in requests for transmission. This counter does not include any datagrams that were counted in Packets Routed. |
| Out Discards | The number of output IP datagrams that were discarded, although no problems were noted. This can occur because of insufficient buffer space. This counter includes datagrams that were counted in Packets Routed if the packets met this discard criterion. |

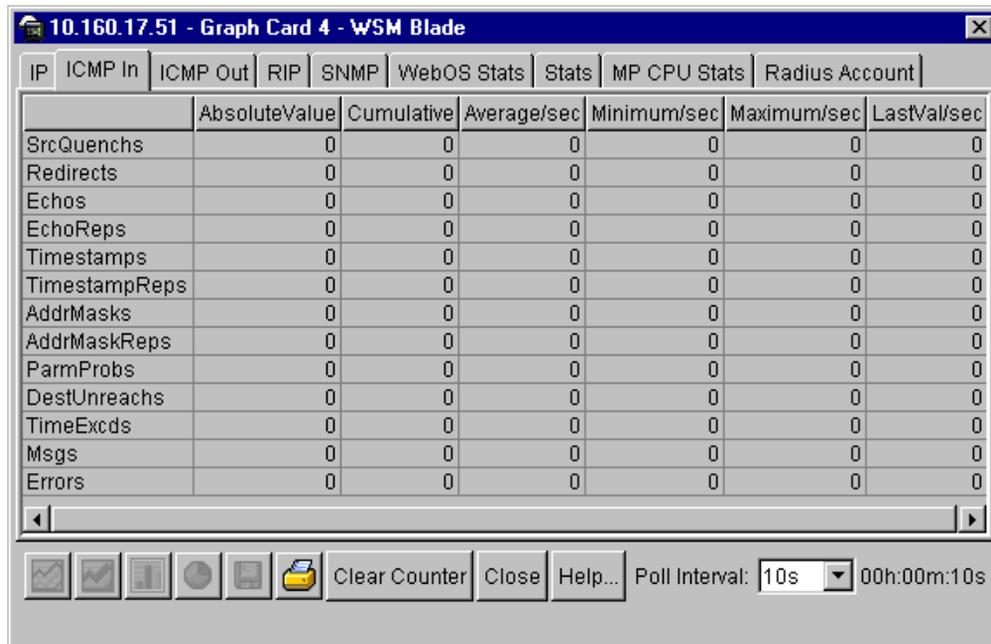
Table 164 WSM Blade—IP tab fields (continued)

| Field | Description |
|------------------------------|--|
| Out Discards No Routes | The number of IP datagrams discarded because no route was available for transmitting them to their destinations. This counter includes any packets counted in Packets Routed that meet this no-route criterion. Also included are any datagrams that a host cannot route because all default gateways are down. |
| Fragments OK | The number of IP datagrams that have been successfully fragmented. |
| Fragments Failed | The number of IP datagrams that have been discarded because they could not be fragmented, such as when the Don't Fragment flag has been set. |
| Fragments Created | The number of IP datagrams that have been fragmented. |
| Fragment Reassembly Required | The number of received IP fragments that needed to be reassembled. |
| Fragment Reassembly OK | The number of IP datagrams successfully reassembled. |
| Fragment Reassembly Failed | The number of failures detected by the IP re-assembly algorithm. Possible failures include timed out, errors, and so on. Note: This is not necessarily a count of discarded IP fragments. Some algorithms, notably the algorithm in RFC 815, can lose track of the number of fragments by combining them as they are received. |
| Routing Discards | The number of dropped packets. |

Received ICMP statistics

To view and graph received ICMP statistics:

- 1 From the device view, select the Web Switching Module.
The module is highlighted.
- 2 From the Graph menu, select WSM Card > General.
The WSM Blade dialog box opens to the IP tab.
- 3 Click the ICMP In tab.
- 4 The **ICMP In tab** (Figure 201) opens displaying the fields described in Table 165.

Figure 201 WSM Blade—ICMP In tab

- 5 Click a cell(s) to graph.
- 6 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 165](#) describes the fields on the WSM Blade—ICMP In tab.

Table 165 WSM Blade—ICMP In tab fields

| Field | Description |
|--------------------------------|---|
| Source Quench Packets In | The number of received Internet Control Message Protocol (ICMP) Source Quench messages. |
| Redirect Packets In | The number of received ICMP Redirect messages. |
| Echo (Ping) Request Packets In | The number of received ICMP Echo (request) messages. |
| Echo (Ping) Reply Packets In | The number of received ICMP Echo Reply messages. |
| Timestamp Request Packets In | The number of received ICMP Timestamp (request) messages. |

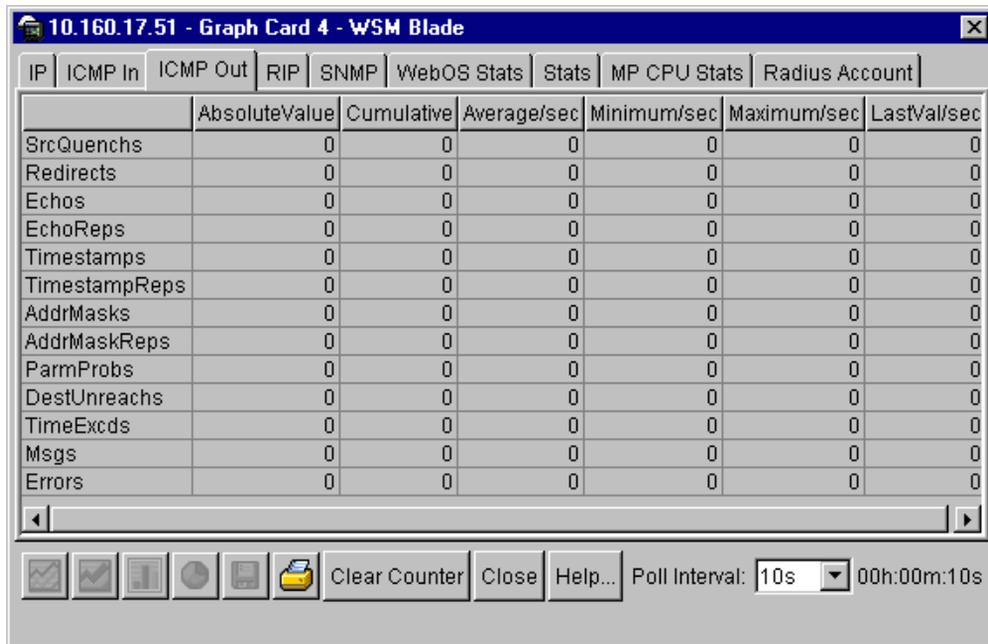
Table 165 WSM Blade—ICMP In tab fields (continued)

| Field | Description |
|------------------------------------|---|
| Timestamp Reply Packets In | The number of received ICMP Timestamp Reply messages. |
| Address Mask Request Packets In | The number of received ICMP Address Mask Request messages. |
| Address Mask Reply Packets In | The number of received ICMP Address Mask Reply messages. |
| Parameter Problem Packets In | The number of received ICMP Parameter Problem messages. |
| Destination Unreachable Packets In | The number of received ICMP Destination Unreachable messages. |
| Time Exceeded Packets In | The number of received ICMP Time Exceeded messages. |
| Messages In | The number of received ICMP messages. |
| Error Packets In | The number of received ICMP Time error messages. |

Sent ICMP statistics

To view and graph sent ICMP statistics:

- 1 From the device view, select the Web Switching Module.
The module is highlighted.
- 2 From the Graph menu, select WSM Card > General.
The WSM Blade dialog box opens to the IP tab.
- 3 [Click the ICMP Out tab \(Figure 202\)](#) opens displaying the fields described in [Table 166](#).

Figure 202 WSM Blade—ICMP Out tab

- 4 Click a cell(s) to graph.
- 5 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 166](#) describes the fields on the WSM Blade—ICMP Out tab.

Table 166 WSM Blade—ICMP Out tab fields

| Field | Description |
|---------------------------------|---|
| Source Quench Packets Out | The number of Internet Control Message Protocol (ICMP) Source Quench messages sent. |
| Redirect Packets Out | The number of ICMP Redirect messages sent. For a host, this object will always be 0 (zero) since hosts do not send redirects. |
| Echo (Ping) Request Packets Out | The number of transmitted ICMP Echo request messages. |
| Echo Reply Packets Out | The number of transmitted ICMP Echo Reply messages. |

Table 166 WSM Blade—ICMP Out tab fields (continued)

| Field | Description |
|-------------------------------------|--|
| Timestamp Request Packets Out | The number of transmitted ICMP Timestamp request messages. |
| Timestamp Reply Packets Out | The number of transmitted ICMP Timestamp Reply messages. |
| Address Mask Request Packets Out | The number of transmitted ICMP Address Mask Request messages. |
| Address Mask Reply Packets Out | The number of transmitted ICMP Address Mask Reply messages. |
| Parameter Problem Packets Out | The number of transmitted ICMP Parameter Problem messages. |
| Destination Unreachable Packets Out | The number of transmitted ICMP Destination Unreachable messages. |
| Time Exceeded Packets Out | The number of transmitted ICMP Time Exceeded messages. |
| Packets Out | The total number of delivered ICMP packets. |
| Error Packets Out | The number of ICMP packets delivered with error messages. |

RIP statistics

To view and graph RIP statistics:

- 1 From the Device view, select the Web Switching Module.

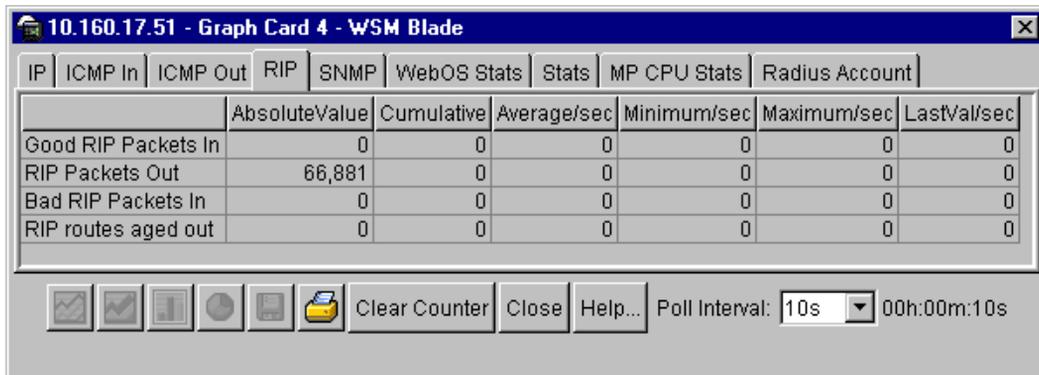
The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > General.

The IP Routing dialog box opens to the IP tab.

- 3 Click the RIP tab.

The [RIP tab \(Figure 203\)](#) opens with the fields listed in [Table 167 on page 526](#).

Figure 203 WSM Blade—RIP statistics tab

- 4 Click a cell(s) to graph.
- 5 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 167](#) describes the fields on the WSM Blade—RIP tab.

Table 167 WSM Blade—RIP tab fields

| Field | Description |
|---------------------|---|
| Good RIP Packets In | The number of good routing information protocol (RIP) packets that were received. |
| RIP Packets Out | The number of transmitted RIP packets that were transmitted. |
| Bad RIP Packets In | The number of received error RIP packets. |
| RIP routes aged out | The number of aged out RIP routes. |

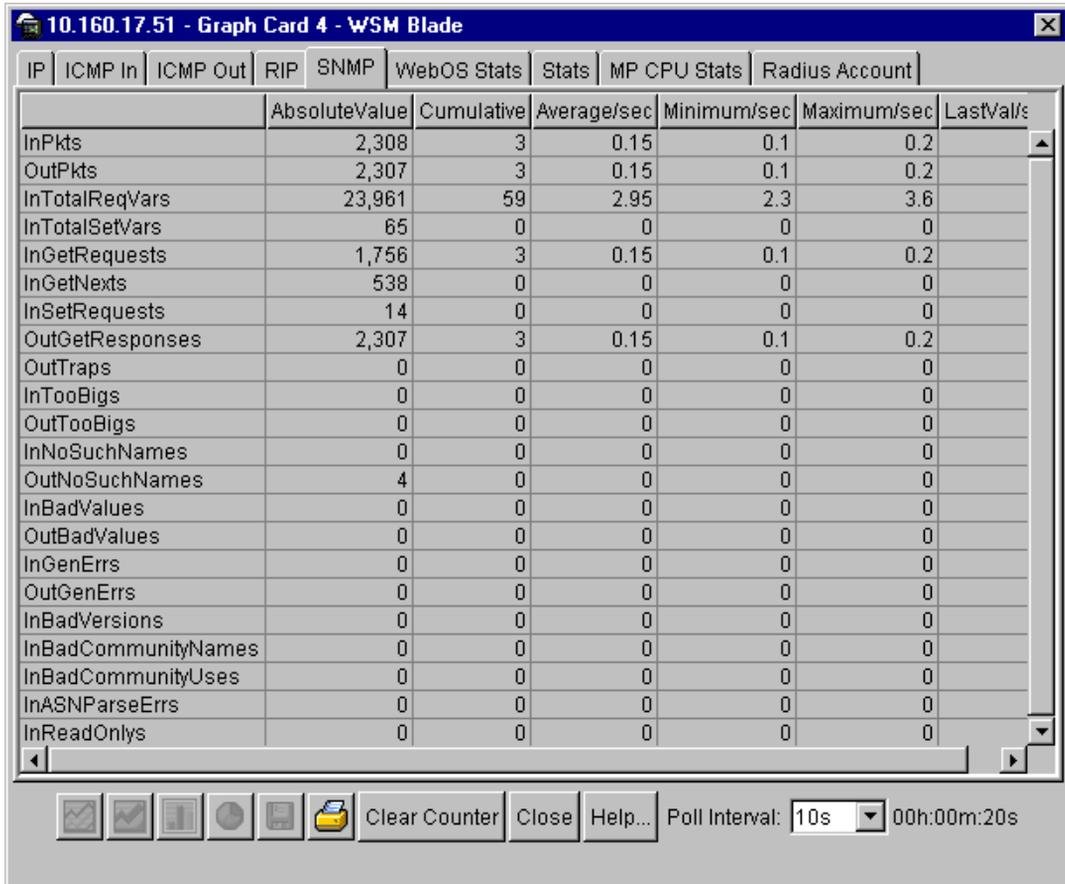
SNMP statistics

To view and graph SNMP statistics:

- 1 From the device view, select the Web Switching Module.
The module is highlighted.
- 2 From the Graph menu, select WSM Card > General.
The WSM Blade dialog box opens to the IP tab.

- 3 Click the SNMP tab.
- 4 The **SNMP tab** (Figure 204) opens displaying the fields described in Table 168.

Figure 204 WSM Blade—SNMP tab



- 5 Click a cell(s) to graph.
- 6 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

Table 168 describes the WSM Blade—SNMP tab fields.

Table 168 WSM Blade—SNMP tab fields

| Field | Description |
|-------------------------|--|
| Packets In | The number of messages delivered to the SNMP switch from the transport. |
| Packets Out | The number of NMP Messages passed from the SNMP switch to the transport service. |
| MIB Variables Retrieved | The number of MIB objects successfully retrieved by the SNMP switch after valid SNMP get-request and get-next protocol data units (PDU). |
| MIB Variables Modified | The number of MIB objects that were successfully altered by the SNMP stack after valid SNMP set-request PDUs. |
| Get Requests In | The number of SNMP get-request PDUs that were accepted and processed by the SNMP stack. |
| Get NEXT Requests In | The number of SNMP get-next PDUs that were accepted and processed by the SNMP stack. |
| Set Requests In | The number of SNMP set-request PDUs that were accepted and processed by the SNMP stack. |
| Get Responses Out | The total number of SNMP get-response PDUs generated by the SNMP protocol entity. |
| Get Responses In | The total number of SNMP get-response PDUs that were accepted and processed by the SNMP stack. |
| Traps Out | The number of SNMP Trap PDUs that were generated by the SNMP entity. |
| Too Big Errors In | The number of SNMP PDUs that were delivered to the SNMP entity when Too Big Errors In occurred. |
| Too Big Errors Out | The number of SNMP PDUs that were generated by the SNMP entity when Too Big Errors Out occurred. |
| No Such Names In | The number of SNMP PDUs that were delivered to the SNMP entity when the No Such Names In error occurred. |
| No Such Names Out | The number of SNMP PDUs that were generated by the SNMP entity when No Such Names Out error occurred. |
| Bad Value Errors In | The number of SNMP PDUs that were delivered to the SNMP entity when Bad Value Errors occurred. |
| Bad Value Errors Out | The number of SNMP PDUs that were generated by the SNMP entity when Bad Value Errors Out occurred. |
| Generic Errors In | The number of SNMP PDUs that were delivered to the SNMP entity when Generic Errors In occurred. |

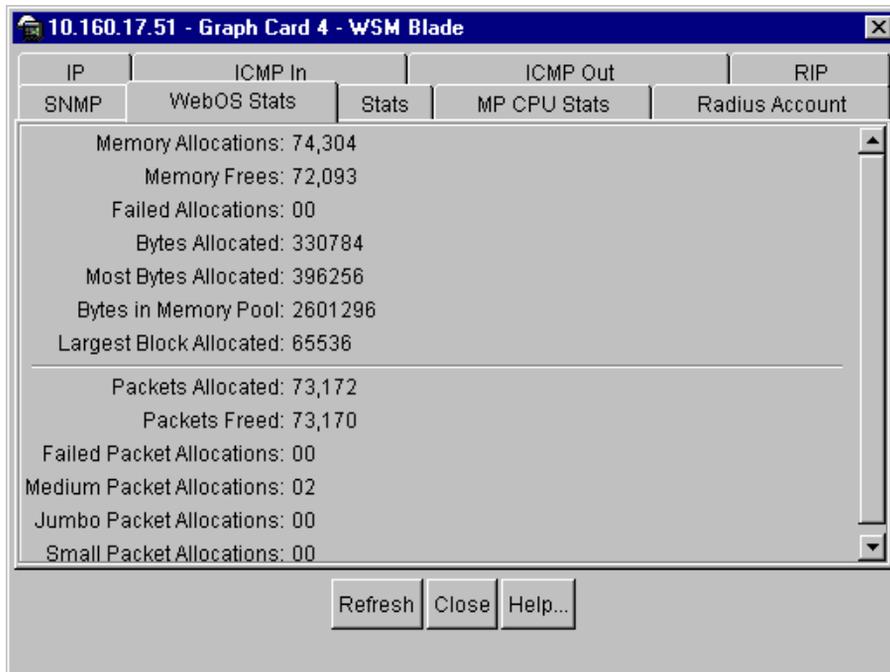
Table 168 WSM Blade—SNMP tab fields (continued)

| Field | Description |
|--|--|
| Generic Errors Out | The number of SNMP PDUs that were generated by the SNMP when Generic Errors Out occurred. |
| Packets Using Unsupported SNMP Version | The number of SNMP messages that were received for an unsupported SNMP version. |
| Packets with Unknown Community String | The number of SNMP messages received that used an unknown SNMP community name. |
| Packets with Wrong Community String | The number of SNMP messages that represented an SNMP operation that was not allowed by the SNMP community named in the message. |
| ASN1 Decode Errors | The number of ASN.1 or BER errors that occurred while the SNMP was decoding received SNMP messages. |
| Read Only Errors In | The number of valid SNMP PDUs that were delivered to the SNMP entity when the Read Only Errors In occurred. This is a protocol error that generates an SNMP PDU that contains readOnly in the error status field. This method detects incorrect implementations of the SNMP. |

WebOS statistics

To view WebOS statistics:

- 1 From the device view, select the Web Switching Module.
The module is highlighted.
- 2 From the Graph menu, select WSM Card > General.
The WSM Blade dialog box opens to the IP tab.
- 3 Click the WebOS Stats tab.
- 4 The [WebOS Stats tab](#) (Figure 205) opens displaying the fields described in [Table 169](#).

Figure 205 WSM Blade—WebOS Stats tab

[Table 169](#) describes the WSM Blade—WebOS Stats tab fields.

Table 169 WSM Blade—WebOS Stats tab fields

| Field | Description |
|-------------------------|--|
| Memory Allocations | The outstanding bytes of memory that have been allocated. |
| Memory Frees | The total number of bytes of memory frees. |
| Failed Allocations | The total number of failed memory allocations. |
| Bytes Allocated | The total number bytes of memory allocations. |
| Most Bytes Allocated | The number of allocated bytes with the high water mark. |
| Bytes in Memory Pool | The number of bytes in the memory pool. |
| Largest Block Allocated | The number of bytes in the largest allocated memory block. |
| Packets Allocated | The total number of allocated packets. |
| Packets Freed | The total number of freed allocated packets. |

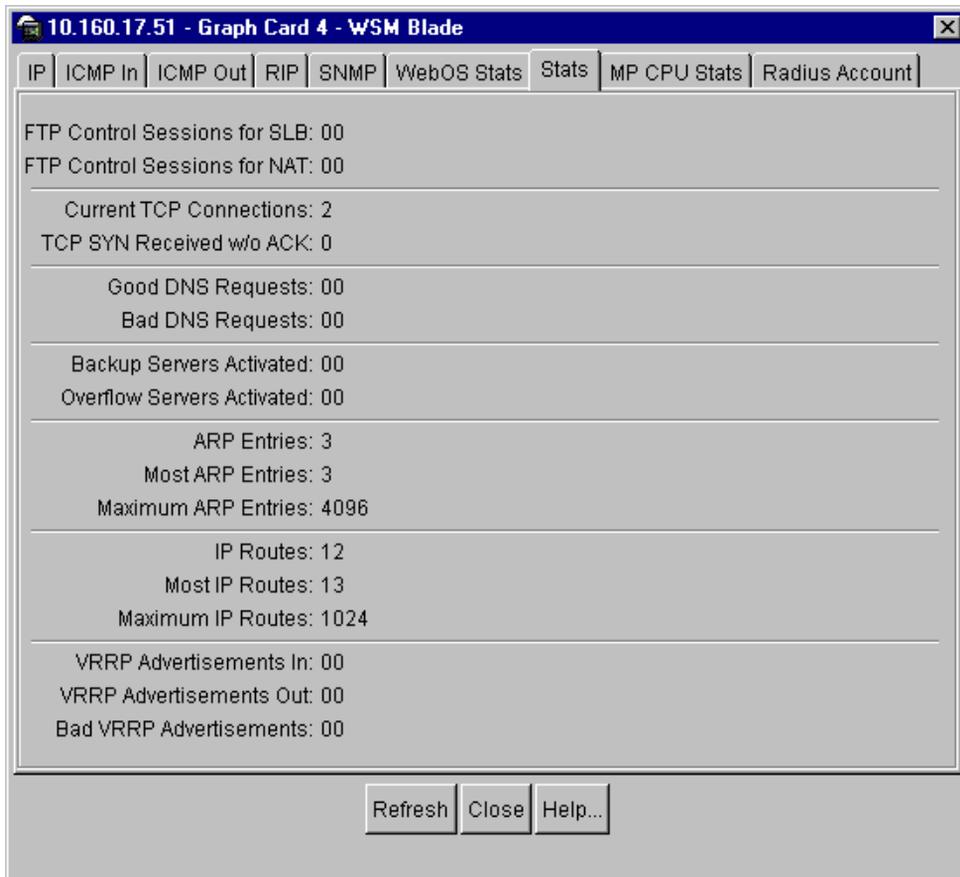
Table 169 WSM Blade—WebOS Stats tab fields (continued)

| Field | Description |
|---------------------------|--|
| Failed Packet Allocations | The total number of failed allocated packets. |
| Medium Packet Allocations | The number of allocated medium size packets. A medium packet contains 1,536 bytes. |
| Jumbo Packet Allocations | The number of allocated jumbo size packets. A jumbo packet contains 9,014 bytes. |
| Small Packet Allocations | The number of allocated small size packets. A small packet contains 128 bytes. |

Statistics

To view data on the Stats tab:

- 1 From the device view, select the Web Switching Module.
The module is highlighted.
- 2 From the Graph menu, select WSM Card > General.
The WSM Blade dialog box opens to the IP tab.
- 3 Click the Stats tab.
- 4 The [Stats tab \(Figure 206\)](#) opens displaying the fields described in [Table 170](#).

Figure 206 WSM Blade—Stats tab

[Table 170](#) describes the fields on the Stats tab.

Table 170 WSM Blade—Stats tab fields

| Field | Description |
|------------------------------|---|
| FTP Control Sessions for SLB | The total number of received FTP control sessions received for SLB after FTP parsing was turned on. |
| FTP Control Sessions for NAT | The total number of received FTP control sessions received for Network Address Translation (NAT) after FTP parsing was turned on. |
| Current TCP Connections | The number of established outstanding TCP current connections. |

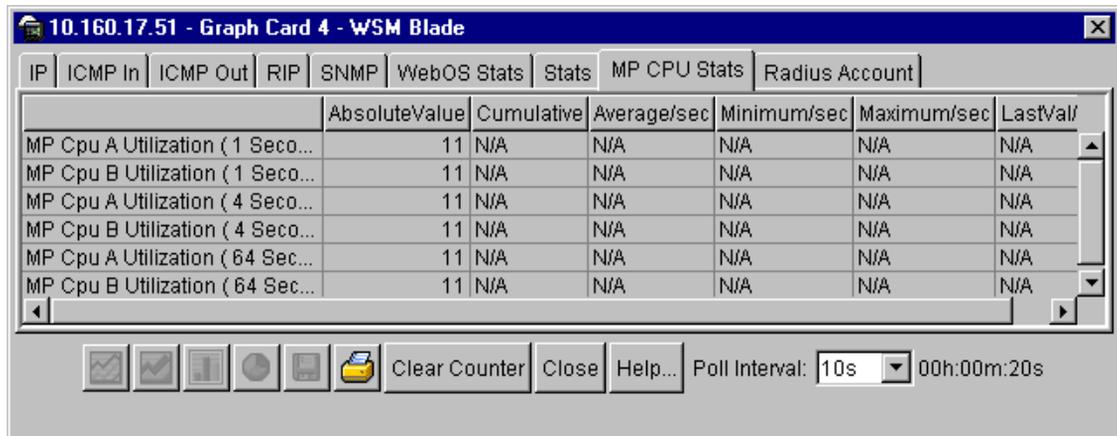
Table 170 WSM Blade—Stats tab fields (continued)

| Field | Description |
|----------------------------|--|
| TCP SYN Received w/o ACK | The number of half-open TCP connections. Valid only when URL parsing has been turned on. It is incremented when the WSM responds to a TCP SYN packet, and is decremented upon receiving a TCP SYN ACK packet from the requester. |
| Good DNS Requests | The number of good DNS requests received. |
| Bad DNS Requests | The number of bad DNS requests received. |
| Backup Servers Activated | The total number of backup servers that were activated when the primary servers failed. |
| Overflow Servers Activated | The total number of times the backup servers became active when the primary servers overflowed. |
| ARP Entries | The current number of ARP entries. |
| Most ARP Entries | The highest number of ARP entries. |
| Maximum ARP Entries | The maximum number of ARP entries. |
| IP Routes | The current number of IP routes. |
| Most IP Routes | The highest number of IP routes. |
| Maximum IP Routes | The maximum number of IP routes. |
| VRRP Advertisements In | The number of received good VRRP advertisements. |
| VRRP Advertisements Out | The number of transmitted good VRRP advertisements. |
| Bad VRRP Advertisements | The number of received bad VRRP advertisements. |

MP CPU statistics

To view and graph MP CPU statistics:

- 1 From the device view, select the Web Switching Module.
The module is highlighted.
- 2 From the Graph menu, select WSM Card > General.
The WSM Blade dialog box opens to the IP tab.
- 3 Click the MP CPU Stats tab.
- 4 The **MP CPU Stats tab** (Figure 207) opens displaying the fields described in Table 171.

Figure 207 WSM Blade—MP CPU tab

- 5 Click a cell(s) to graph.
- 6 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing”](#) on page 516.

[Table 171](#) describes the WSM Blade—MP CPU Stats fields.

Table 171 WSM Blade—MP CPU Stats fields

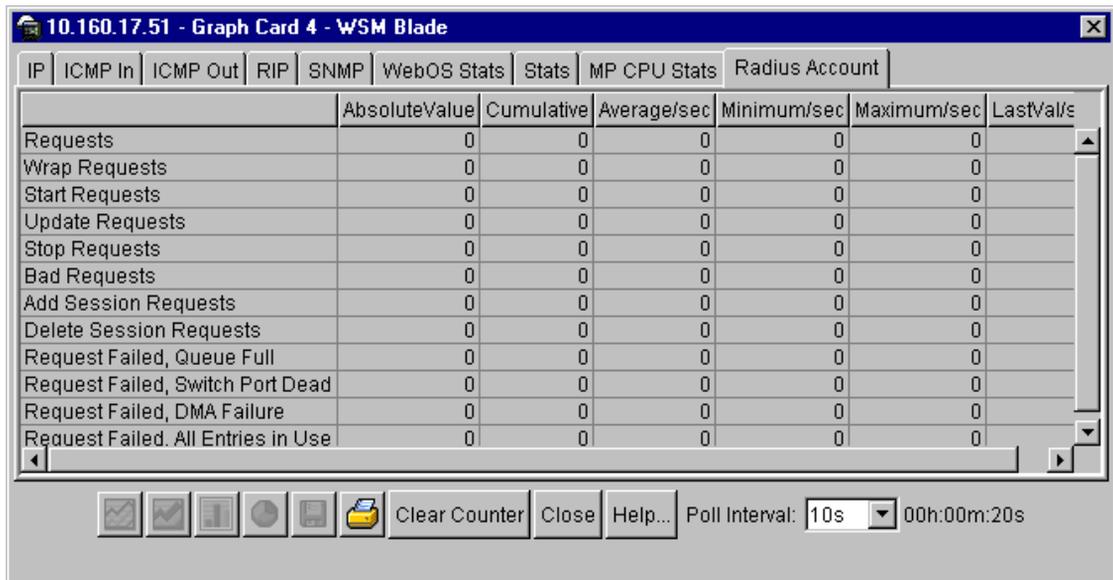
| Field | Description |
|--------------------------------------|---|
| MP CPU A Utilization (1 Second Avg) | The percentage of time CPU A was utilized more than one second. |
| MP CPU B Utilization (1 Second Avg) | The percentage of time CPU B was utilized more than one second. |
| MP CPU A Utilization (4 Second Avg) | The percentage of time CPU A was utilized more than four consecutive seconds. |
| MP CPU B Utilization (4 Second Avg) | The percentage of time CPU B was utilized more than four consecutive seconds. |
| MP CPU A Utilization (64 Second Avg) | The percentage of time CPU A was utilized more than 64 consecutive seconds. |
| MP CPU B Utilization (64 Second Avg) | The percentage of time CPU B was utilized more than 64 consecutive seconds. |

RADIUS account statistics

To view and graph RADIUS account statistics:

- 1 From the device view, select the Web Switching Module.
The module is highlighted.
- 2 From the Graph menu, select WSM Card > General.
The WSM Blade dialog box opens to the IP tab.
- 3 Click the RADIUS Account tab.
- 4 The **RADIUS Account tab** (Figure 208) opens displaying the fields described in Table 172.

Figure 208 WSM Blade—RADIUS Account tab



- 5 Click a cell(s) to graph.
- 6 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing”](#) on page 516.

[Table 172](#) describes the fields on the WSM Blade—Radius Account tab.

Table 172 WSM Blade—Radius account statistics

| Field | Description |
|------------------------------------|---|
| Requests | The number of RADIUS accounting requests. |
| Wrap Requests | The number of RADIUS accounting requests in an internally wrapped buffer. |
| Start Requests | The number of RADIUS accounting start requests. |
| Update Requests | The number of RADIUS accounting update requests. |
| Stop Requests | The number of RADIUS accounting stop requests. |
| Bad Requests | The number of RADIUS accounting requests in bad format. |
| Add Session Requests | The number of RADIUS accounting add session requests. |
| Delete Session Requests | The number of RADIUS accounting delete session requests. |
| Request Failed, Queue Full | The number of requests that failed because the DMA queue was full. |
| Request Failed, Switch Port Dead | The number of requests that failed because the WSM port was dead. |
| Request Failed, DMA Failure | The number of requests that failed because of DMA write failures. |
| Request Failed, All Entries In Use | The number of requests that failed because all entries were in use. |

Bridge forwarding statistics

Bridge forwarding statistics include:

- [“Forwarding statistics” on page 536](#)
- [“Base port statistics” on page 537](#)
- [“TP port statistics” on page 539](#)

Forwarding statistics

To view bridge forwarding statistics:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- From the Graph menu, select WSM Card > Bridge.

The Bridge Forwarding dialog box opens to the **Forwarding tab** (Figure 209) with the fields defined in Table 173.

Figure 209 Bridge forwarding DB—Forwarding tab

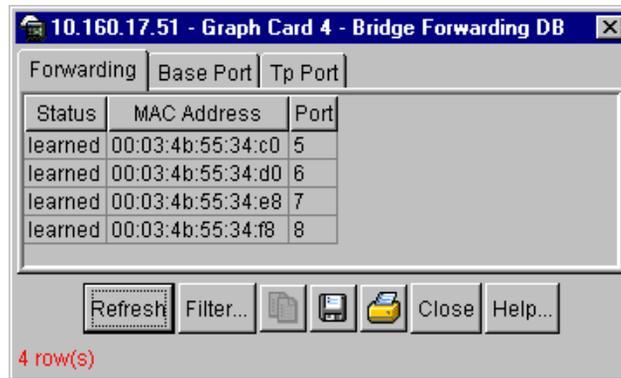


Table 173 describes the fields on the Bridge Forwarding DB—Forwarding tab.

Table 173 Bridge Forwarding DB—Forwarding tab

| Field | Description |
|-------------|--|
| Status | The status of the bridge: learned, self, or other. |
| MAC Address | The MAC address of the bridge. |
| TP Port | The physical port on which the MAC address is located. |

Base port statistics

To view base port statistics:

- From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- From the Graph menu, select WSM Card > Bridge.

The Bridge Forwarding DB dialog box opens to the Forwarding tab.

- Click the Base Port tab.

The **Base Port** tab (Figure 210) opens with the fields defined in Table 174.

Figure 210 Bridge Forwarding DB—Base Port tab

| Port | Port Interface | Port Circuit | Exceeded Delay Discarded Frames | Exceeded Size Discarded Frames |
|------|----------------|--------------|---------------------------------|--------------------------------|
| 1 | 1 | .0.0 | 00 | 00 |
| 2 | 2 | .0.0 | 00 | 00 |
| 3 | 3 | .0.0 | 00 | 00 |
| 4 | 4 | .0.0 | 00 | 00 |
| 5 | 5 | .0.0 | 00 | 00 |
| 6 | 6 | .0.0 | 00 | 00 |
| 7 | 7 | .0.0 | 00 | 00 |
| 8 | 8 | .0.0 | 00 | 00 |

- Click the Tree field and choose a spanning tree group from the drop-down selection list.

The data for the selected spanning tree group display in the table.

Table 174 describes the fields on the Bridge Forwarding DB—Base Port tab.

Table 174 Bridge Forwarding DB—Base Port tab

| Field | Description |
|---------------------------------|---|
| Port | The physical port. |
| Port Interface | The interface number to the physical port. |
| Port Circuit | The circuit number of the physical port. |
| Exceeded Delay Discarded Frames | The number of frames that were discarded because of excessive transit delay through the bridge. Discards Due to Transit Delay is incremented by transparent and source route bridges. |
| Exceeded Size Discarded Frames | The number of excessively large frames that were discarded. Discards Due to Excessive Size is incremented by transparent and source route bridges. |

TP port statistics

To view TP port statistics:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Bridge.
The Bridge Forwarding DB dialog box opens to the Forwarding tab.
- 3 Click the TP Port tab.
The **TP Port** tab (Figure 210) opens with the fields defined in Table 174.

Figure 211 Bridge Forwarding DB—TP Port tab

| Port | Maximum size of INFO | Frames In | Frames Out | Discarded Frames In |
|------|----------------------|-----------|------------|---------------------|
| 1 | 9008 | 00 | 00 | 00 |
| 2 | 9008 | 00 | 00 | 00 |
| 3 | 9008 | 00 | 00 | 00 |
| 4 | 9008 | 00 | 00 | 00 |
| 5 | 9008 | 767,385 | 00 | 00 |
| 6 | 9008 | 767,386 | 00 | 00 |
| 7 | 9008 | 1,315,535 | 01 | 00 |
| 8 | 9008 | 1,319,936 | 71,338 | 00 |

- 4 Click the Tree field and choose a spanning tree group from the drop-down selection list.

The data for the selected spanning tree group display in the table.

[Table 175](#) describes the fields on the Bridge Forwarding DB—TP Port tab.

Table 175 Bridge Forwarding DB—TP Port tab

| Field | Description |
|----------------------|---|
| Port | Port number of the WSM. |
| Maximum size of INFO | Maximum size of the INFO (non-MAC) field that this port will receive or transmit. |
| Frames In | Number of frames that have been received by this port from its segment. Note that a frame received on the interface corresponding to this port is only counted by this object if and only if it is for a protocol being processed by the local bridging function, including bridge management frames. |
| Frames Out | Number of frames that have been transmitted by this port to its segment. Note that a frame transmitted on the interface corresponding to this port is only counted by this object if and only if it is for a protocol being processed by the local bridging function, including bridge management frames. |
| Discarded Frames In | Count of valid frames received which were discarded (filtered) by the Forwarding Process. |

Port statistics

Port statistics include:

- [“Port interface statistics” on page 540](#)
- [“Ethernet statistics” on page 543](#)
- [“Bridge statistics” on page 546](#)
- [“Load balance statistics” on page 547](#)
- [“Port CPU statistics” on page 549](#)

Port interface statistics

To view port interface statistics:

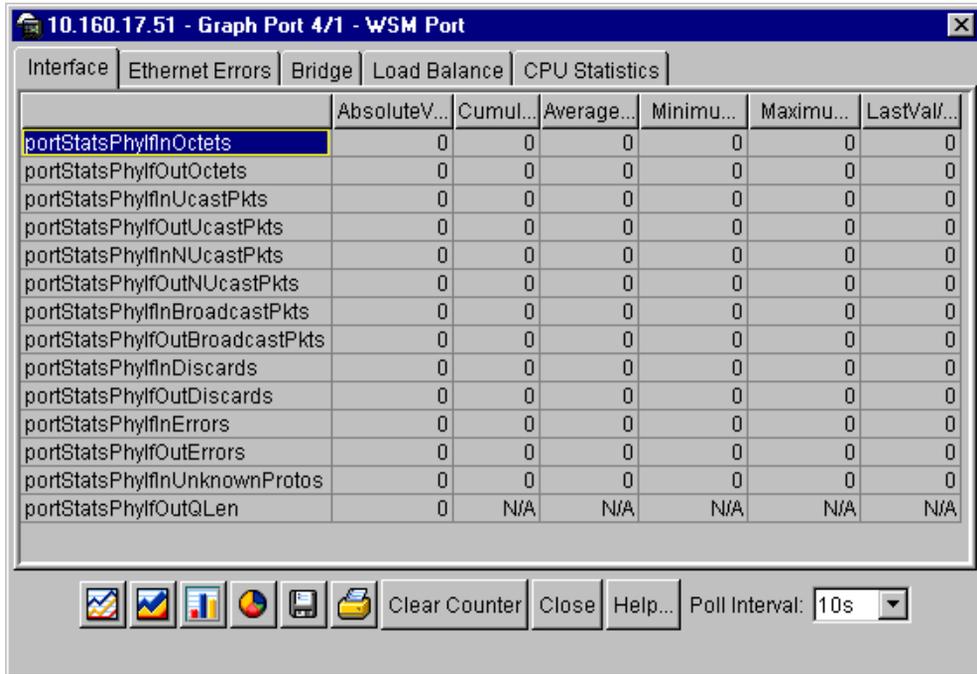
- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.

The port is highlighted.

- From the menu bar, choose Graph > WSM Card > Port/Back Port.

The Web Switching Module Port dialog box opens to the [Interface](#) tab ([Figure 212](#)) with the fields described in [Table 176](#).

Figure 212 Port—Interface tab



- Click a cell(s) to graph.
- Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing”](#) on page 516.

[Table 176](#) describes the fields on the Port—Interface tab.

Table 176 Port—Interface tab fields

| Field | Description |
|-----------|---|
| Bytes In | The number of bytes received on the interface, including framing characters. |
| Bytes Out | The number of bytes transmitted out of the interface, including framing characters. |

Table 176 Port—Interface tab fields (continued)

| Field | Description |
|-----------------------|--|
| Unicast Packets In | The number of packets, delivered by this sublayer to a higher layer that were not addressed to a multicast or a broadcast address at this sublayer. |
| Unicast Packets Out | The number of packets that higher-level protocols requested to transmit that were not addressed to a multicast or broadcast address at this sublayer. The count includes the packets that were discarded or not delivered. |
| Multicasts In | The number of packets, delivered by this sublayer to a higher layer that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this includes Group and Functional addresses. |
| Multicasts Out | The number of packets that higher-level protocols requested to transmit that were addressed to a multicast address at this sublayer. The count includes the packets that were discarded or not sent. For a MAC layer protocol, this includes Group and Functional addresses. |
| Broadcast In | The number of packets, delivered by this sublayer to a higher layer, that were addressed to a broadcast address at this sublayer. |
| Broadcast Out | The number of packets that higher-level protocols requested to transmit, that were addressed to a broadcast address at this sublayer. The count includes the packets that were discarded or not delivered. Discarded Packets The number of inbound packets that were discarded, although no errors had been detected to prevent their delivery to a higher-layer protocol. This can occur to free up buffer space. |
| Inbound Discards | The number of discarded inbound packets. |
| Outbound Discards | The number of discarded outbound packets. No errors were detected that would prevent their transmission. This can occur to free up buffer space. |
| Error Packets | For packet-oriented interfaces, the number of inbound packets with errors that prevented their delivery to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units with errors that prevented their delivery to a higher-layer protocol. |
| Not Sent Due to Error | For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors. |

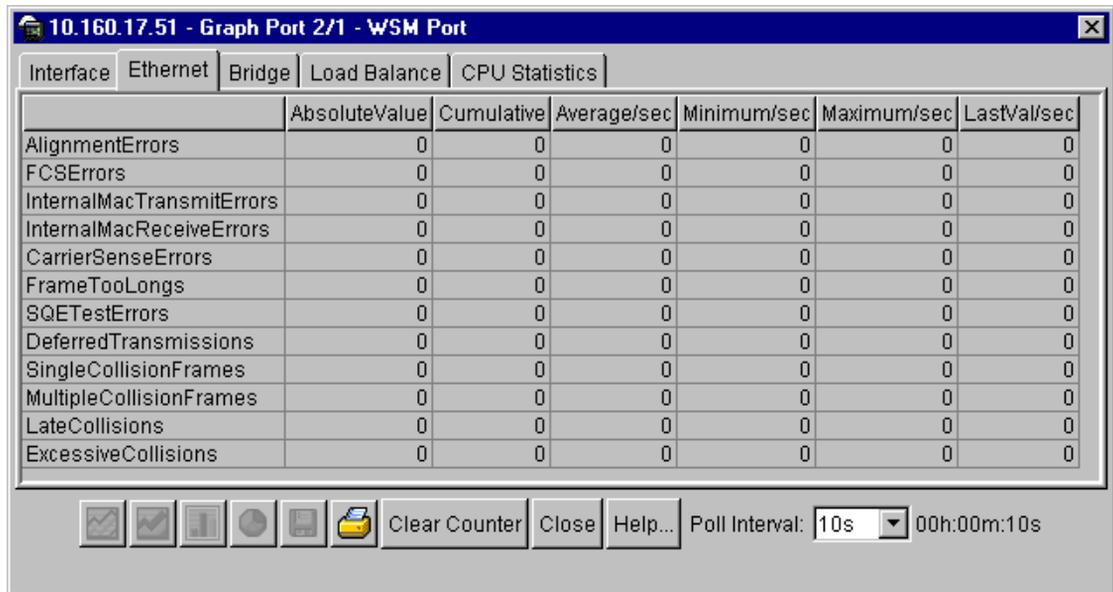
Table 176 Port—Interface tab fields (continued)

| Field | Description |
|------------------------------|--|
| Unknown Protocol Packets | For packet-oriented interfaces, the number of packets received via the interface that were discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received via the interface that were discarded because of an unknown or unsupported protocol. If the interface does not support protocol multiplexing, Unknown Protocol Packets will be 0 (zero). |
| Outbound Packet Queue Length | The number of packets in the output queue. |

Ethernet statistics

To graph port ethernet statistics:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.
The port is highlighted.
- 3 From the menu bar, choose Graph > WSM Card > Port/Back Port.
The Web Switching Module Port dialog box opens to the Interface tab.
- 4 Click the Ethernet tab.
- 5 The [Ethernet tab \(Figure 213\)](#) opens with the fields described in [Table 177](#).

Figure 213 Port—Ethernet tab

- 6 Click a cell(s) to graph.
- 7 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing”](#) on page 516.

[Table 177](#) describes the fields on the Port—Ethernet tab.

Table 177 Port—Ethernet tab fields

| Field | Description |
|------------------|--|
| Alignment Errors | The number of frames received on this interface that were not of integral length and did not pass the FCS check. The count is incremented when the Alignment Error status is returned by the MAC service to the LLC, or other MAC user. Frames with multiple errors are counted exclusively. |
| FCS Errors | The number of frames received on this interface that failed the FCS health check because of length. The count is incremented when the Frame Check Error status is returned by the MAC service to the LLC, or other MAC user. Frames with multiple errors are counted exclusively. |

Table 177 Port—Ethernet tab fields (continued)

| Field | Description |
|--|--|
| Internal MAC Transmission Errors | The number of frames transmitted on this interface that failed because of an internal MAC sublayer transmit error. This error is only counted if it was not counted for Late Collisions, Excessive Collisions, or Carrier Sense Errors. May represent a number of transmission errors that were not otherwise recorded. |
| Internal MAC Receive Errors | The number of frames on this interface that could not be accepted because of an internal MAC sublayer error. This frame error is only counted if it was not counted for Received Frames > Maximum Length, Alignment Errors, or FCS Errors. May represent a number of transmission errors that were not otherwise recorded. |
| Carrier Sense Errors | The number of times the carrier sense condition was lost or was never asserted while attempting to transmit a frame on this interface. Incremented once per transmission attempt. |
| Received Frames Exceeding Maximum Length | The number of frames received that exceeded the maximum frame size on this interface. The count is incremented when the Received Frames > Maximum Length status is returned by the MAC service to the LLC, or other MAC user. Received frames that have multiple errors are counted exclusively. |
| SQE Test Errors | The number of times the SQE TEST ERROR message was generated by the PLS sublayer for this interface. |
| Deferred Transmissions | The number of frames for which the first transmission attempt was delayed because the medium was busy on this interface. Does not include frames involved in collisions. |
| Single Collision Frames | The number of successfully transmitted frames where transmission was inhibited by a single collision on this interface. A frame counted here can also be counted in the number of Unicast Packets Out, Multicasts Out, or Broadcast Out, but is not counted in the number of Multiple Collision Frames. |
| Multiple Collision Frames | A count of successfully transmitted frames where transmission was inhibited by more than one collision on this interface. A frame counted here can also be counted in the number of Unicast Packets Out, Multicasts Out, or Broadcast Out, but is not counted in the number of Single Collision Frames. |

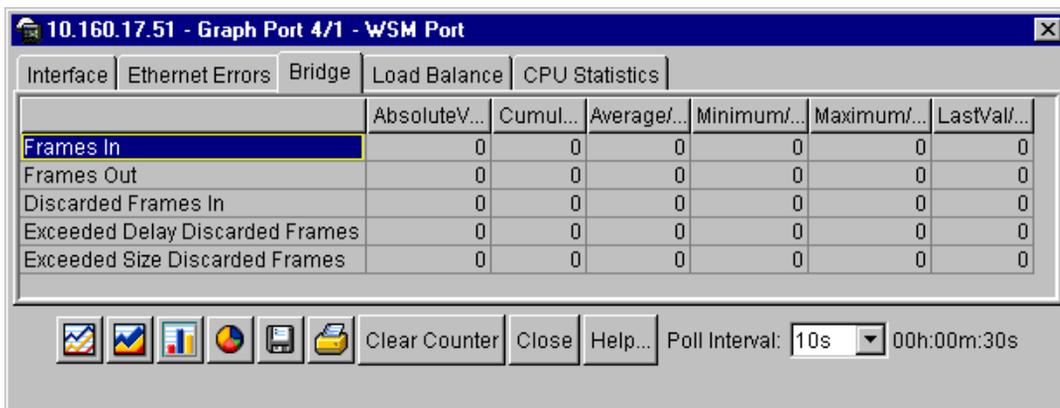
Table 177 Port—Ethernet tab fields (continued)

| Field | Description |
|----------------------|---|
| Late Collisions | The number of times a late collision (later than 512 bit-times into the transmission of a packet) is detected on this interface. A late collision counted here can be considered as a generic collision for other statistics. Bit-times vary per system. Example: On a 10Mbps system, 512 bit-times represents 51.2 microseconds. |
| Excessive Collisions | The number of frames in which transmission failed because of excessive collisions on this interface. |

Bridge statistics

To graph port bridge statistics:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.
The port is highlighted.
- 3 From the menu bar, choose Graph > WSM Card > Port/Back Port.
The Web Switching Module Port dialog box opens to the Interface tab.
- 4 Click the Bridge tab.
- 5 The **Bridge tab** (Figure 214) opens with the fields described in Table 178.

Figure 214 Port—Bridge tab

- 6 Click a cell(s) to graph.
- 7 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 178](#) describes the fields on the Port—Bridge tab.

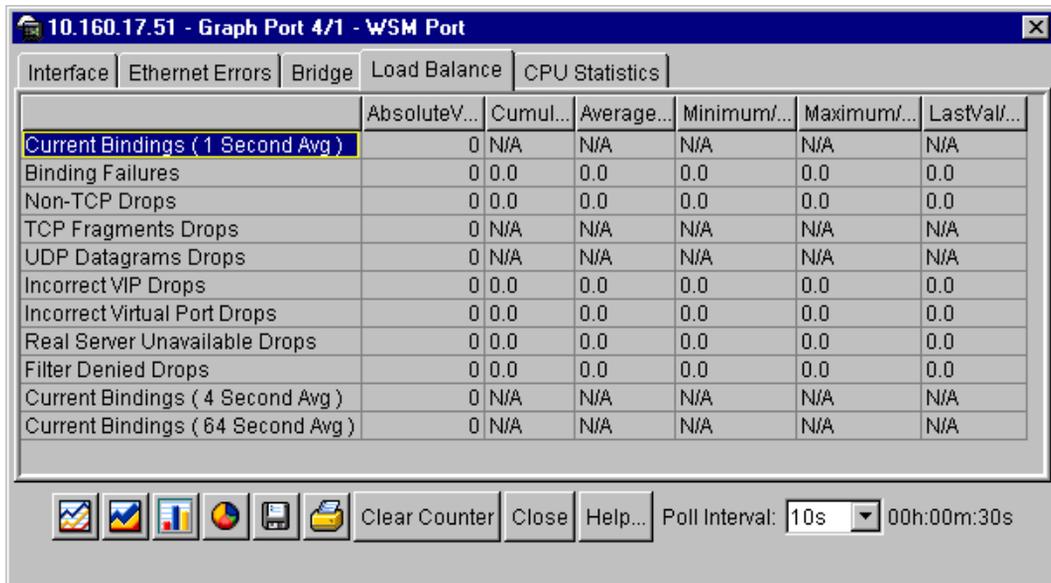
Table 178 Port—Bridge tab fields

| Field | Description |
|--------------------------------|---|
| Packets In | The number of frames that have been received by this port from its segment. Only counts frames that are for a protocol being processed by the local bridging function, including bridge management frames. |
| Packets Out | The number of frames that have been transmitted by this port to its segment. Only counts frames that are for a protocol being processed by the local bridging function, including bridge management frames. |
| Discarded/Filtered Packets | The number of valid, received frames that were discarded by the Forwarding Process. |
| Discards Due to Transit Delay | The number of frames that were discarded because of excessive transit delay through the bridge. Incremented by transparent and source route bridges. |
| Discards Due to Excessive Size | The number of excessively large frames that were discarded. Incremented by transparent and source route bridges. |

Load balance statistics

To graph port load balance statistics:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.
The port is highlighted.
- 3 From the menu bar, choose Graph > WSM Card > Port/Back Port.
The Web Switching Module Port dialog box opens to the Interface tab.
- 4 Click the Load Balance tab.
- 5 The [Load Balance tab](#) ([Figure 215](#)) opens with the fields described in [Table 179](#).

Figure 215 Port—Load Balance tab

- 6 Click a cell(s) to graph.
- 7 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 179](#) describes the fields on the Port—Load Balance tab.

Table 179 Port—Load balance tab fields

| Field | Description |
|---------------------------------|--|
| Current Bindings (1 Second Avg) | The current number of binding sessions that average one second. |
| Binding Failures | The number of times the port has run out of binding table entries. |
| Non-TCP Drops | The number of non-TCP/IP frames that were dropped from the port. |
| TCP Fragment Drops | The number of TCP fragments that were dropped from the port. |
| UDP Datagrams Drops | The number of UDP datagrams that were dropped from the port. |

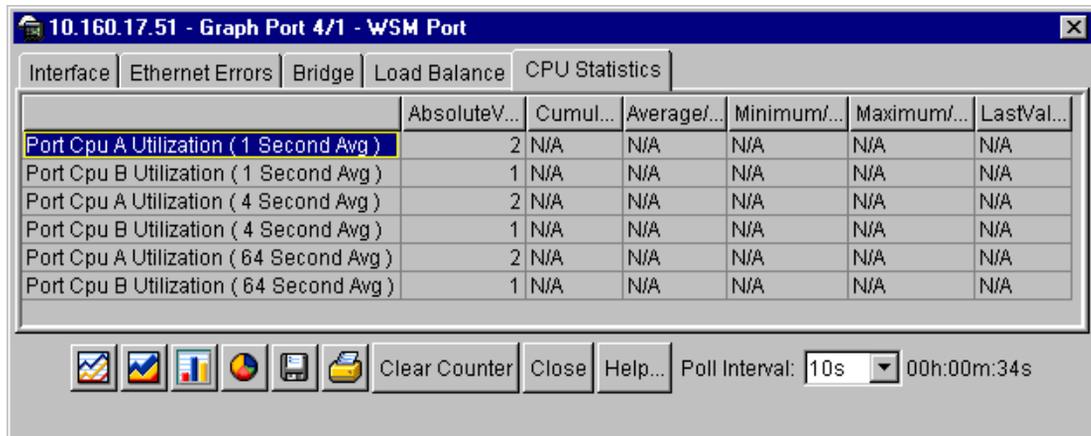
Table 179 Port—Load balance tab fields (continued)

| Field | Description |
|----------------------------------|---|
| Incorrect VIP Drops | The number of frames with incorrect virtual IP (VIP) addresses that were dropped from the port. |
| Incorrect Virtual Port Drops | The number of frames with incorrect Virtual Port that were dropped from the port. |
| Real Server Unavailable Drops | The number of frames that were dropped from the port because no real server was available. |
| Filtered Denied Frames | The number of frames on the port that were denied by the filter. |
| Current Bindings (4 Second Avg) | The current number of bindings that averaged four seconds. |
| Current Bindings (64 Second Avg) | The number of bindings that averaged 64 seconds. |

Port CPU statistics

To graph port CPU statistics:

- 1 From the Device view, select a port orientation (Front or Back).
- 2 Select a port.
The port is highlighted.
- 3 From the menu bar, choose Graph > WSM Card > Port/Back Port.
The Web Switching Module Port dialog box opens to the Interface tab.
- 4 Click the CPU Statistics tab.
- 5 The [CPU Statistics tab](#) (Figure 216) opens with the fields described in [Table 180](#).

Figure 216 Port—CPU Statistics tab

- 6 Click a cell(s) to graph.
- 7 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing”](#) on page 516.

[Table 180](#) describes the fields on the Port—CPU Statistics tab.

Table 180 CPU Statistics tab fields

| Field | Description |
|--|---|
| Port CPU A Utilization (1 Second Avg) | The percentage of total operation time that SP CPU A was utilized more than one second. |
| Port CPU B Utilization (1 Second Avg) | The percentage of total operation time that SP CPU B was utilized more than one second. |
| Port CPU A Utilization (4 Second Avg) | The percentage of total operation time that SP CPU A was utilized more than four seconds. |
| Port CPU B Utilization (4 Second Avg) | The percentage of total operation time that SP CPU B was utilized more than four seconds. |
| Port CPU A Utilization (64 Second Avg) | The percentage of total operation time that SP CPU A was utilized more than 64 seconds. |
| Port CPU B Utilization (64 Second Avg) | The percentage of total operation time that SP CPU B was utilized more than 64 seconds. |

IP routing statistics

IP routing statistics include:

- [“Route statistics” on page 551](#)
- [“ARP statistics” on page 552](#)
- [“Interface statistics” on page 553](#)
- [“IP address statistics” on page 555](#)
- [“TCP statistics” on page 556](#)
- [“TCP connection statistics” on page 558](#)
- [“UDP statistics” on page 560](#)
- [“UDP local statistics” on page 562](#)

Route statistics

To view route statistics for IP routing:

- 1** From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2** From the Graph menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the [Route tab \(Figure 217\)](#) with the fields defined in [Table 181](#).

Figure 217 IP Routing—Route tab

| Route | ARP | Interface | IP Address | TCP | TCP Connection | UDP | UDP Local |
|-------|------------------------|---------------------|-----------------|-----------|----------------|-----------|-----------|
| Route | Destination IP Address | Destination IP Mask | Next-Hop Router | Tag Type | Route Type | Interface | |
| 1 | 0.0.0.0 | 255.255.255.255 | 0.0.0.0 | broadcast | broadcast | 1 | |
| 2 | 255.255.255.255 | 255.255.255.255 | 255.255.255.255 | broadcast | broadcast | 0 | |
| 3 | 10.10.10.0 | 255.255.255.0 | 10.10.10.80 | fixed | direct | 1 | |
| 4 | 10.10.10.80 | 255.255.255.255 | 10.10.10.80 | addr | local | 1 | |
| 5 | 10.10.10.255 | 255.255.255.255 | 0.0.0.0 | broadcast | broadcast | 1 | |
| 6 | 127.0.0.0 | 255.0.0.0 | 0.0.0.0 | martian | martian | 0 | |
| 7 | 224.0.0.0 | 224.0.0.0 | 0.0.0.0 | martian | martian | 0 | |
| 8 | 224.0.0.5 | 255.255.255.255 | 0.0.0.0 | addr | multicast | 0 | |
| 9 | 224.0.0.6 | 255.255.255.255 | 0.0.0.0 | addr | multicast | 0 | |

9 row(s)

[Table 181](#) describes the fields on the IP Routing—Route tab.

Table 181 IP Routing—Route tab fields

| Field | Description |
|------------------------|--|
| Route | The index number of the routing table. |
| Destination IP Address | The destination IP address of this route. |
| Destination IP Mask | The IP mask of this route. |
| Next-Hop Router | The gateway of this route. |
| Tag Type | The tag type: ICMP, static, SNMP, addr, RIP, broadcast, martian, or multicast. |
| Route Type | The type of route: indirect, direct, local, broadcast, martian, multicast, or other. |
| Interface | The IP interface of this route that is used as the source IP for routing. |

ARP statistics

To view ARP statistics for IP routing:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > IP Routing.
The IP Routing dialog box opens to the Route tab.
- 3 Click the ARP tab.
- 4 The **ARP tab** (Figure 218) opens with the fields defined in Table 182.

Figure 218 IP Routing—ARP tab

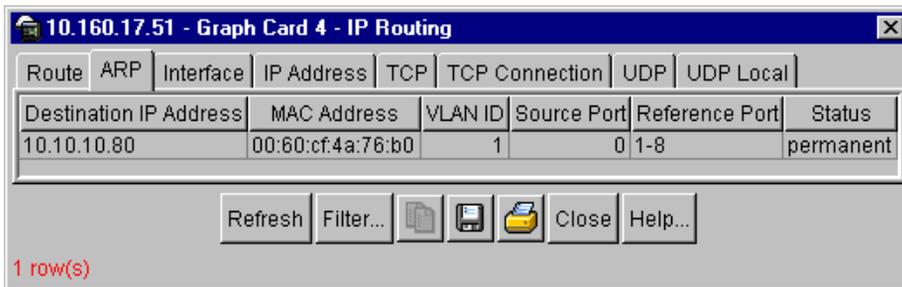


Table 182 describes the fields on the IP Routing—ARP tab.

Table 182 IP Routing—ARP tab fields

| Field | Description |
|------------------------|---|
| Destination IP Address | The destination IP address of the ARP. |
| MAC Address | The MAC address for the Address Resolution Protocol (ARP) entry. |
| VLAN ID | The VLAN identifier for the ARP. |
| Source Port | The port number. |
| Reference Port | The reference ports that are associated with this ARP. |
| Status | The flag status of this ARP: clear, unresolved, permanent, or indirect. |

Interface statistics

To view interface statistics for IP routing:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the Route tab.

- 3 Click the Interface tab.
- 4 The **Interface tab** (Figure 219) opens with the fields defined in Table 183.

Figure 219 IP Routing—Interface tab

| Index | Descr | Type | Mtu | Speed | PhysAddress | AdminSt... | OperSt... | LastChan... | Specific |
|-------|-------|-------------|-----|-------|-------------------|------------|-----------|-------------|----------|
| 1 | net0 | propVirtual | 0 | 0Kbps | 00:60:cf:4a:72:90 | up | up | 06h:26m:... | .0.0 |
| 2 | net1 | propVirtual | 0 | 0Kbps | 00:60:cf:4a:72:90 | up | up | 06h:26m:... | .0.0 |

10 row(s)

Table 183 describes the fields on the IP Routing—Interface tab.

Table 183 IP Routing—Interface tab fields

| Field | Description |
|----------------------|---|
| Index | The index number of the Interface table. |
| Description | A description of the network interface card (NIC). |
| Type | The type of network interface. |
| MTU (Largest Packet) | The largest number of bytes of a packet -- the Maximum Transmission Unit (MTU) of the port. |
| Speed | The current speed of the port: 10Mbps, 100Mbps, 1,000Mbps, or other. |
| MAC Address | The MAC address of the WSM. |
| Admin State | The current status. |
| Operational Status | The current status: up, down, or testing. |

Table 183 IP Routing—Interface tab fields (continued)

| Field | Description |
|-------------------|---|
| Last Change | Date and time since the interface entered its current operational state because of configuration change or automatically back online: year/monthday-hour:minute:second If the current state was configured before the network management system, Last Change will be 0. |
| MIB Specification | A reference to MIB definitions that are specific to the media that realizes the interface. Example: if the interface is realized by Ethernet, then MIB Specific refers to a document that defines Ethernet objects. |

IP address statistics

To view IP address statistics for IP routing:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > IP Routing.
The IP Routing dialog box opens to the Route tab.
- 3 Click the IP Address tab.

The **IP Address** tab (Figure 220) opens with the fields defined in Table 184.

Figure 220 IP Routing—IP Address tab

[Table 184](#) describes the fields on the IP Address tab.

Table 184 IP routing—IP Address tab fields

| Field | Description |
|-------------------------|---|
| IP Address | The IP address. |
| Interface | The index number of the interface. |
| IP Subnet Mask | The subnet mask of the IP address. The subnet mask is an IP address with all the network bits set to 1 and all the hosts bits set to 0. |
| Broadcast LSB | The broadcast address of the interface. |
| Maximum Reassembly Size | The size of the largest IP datagram that can be re-assembled from fragmented IP datagrams. Range: 0 to 65,535. |

TCP statistics

To view TCP statistics for IP routing:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > IP Routing.
The IP Routing dialog box opens to the Route tab.
- 3 Click the TCP tab.
The [TCP tab \(Figure 221\)](#) opens with the fields defined in [Table 185](#).

Figure 221 IP Routing—TCP tab

| Route | AbsoluteValue | Cumulative | Average/sec | Minimum/sec | Maximum/sec | LastVal/sec |
|--------------|---------------|------------|-------------|-------------|-------------|-------------|
| ActiveOpens | 1 | 0 | 0 | 0 | 0 | 0 |
| PassiveOpens | 0 | 0 | 0 | 0 | 0 | 0 |
| AttemptFails | 0 | 0 | 0 | 0 | 0 | 0 |
| EstabResets | 0 | 0 | 0 | 0 | 0 | 0 |
| CurrEstab | 1 | N/A | N/A | N/A | N/A | N/A |
| InSegs | 158 | 0 | 0 | 0 | 0 | 0 |
| OutSegs | 158 | 0 | 0 | 0 | 0 | 0 |
| RetransSegs | 0 | 0 | 0 | 0 | 0 | 0 |
| InErrs | 0 | 0 | 0 | 0 | 0 | 0 |
| OutRsts | 0 | 0 | 0 | 0 | 0 | 0 |

- 4 Click a cell(s) to graph.
- 5 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).



Note: When you try to graph a TCP connection state, the error message tcpConnState:noSuchName may display.

Table 185 describes the fields on the TCP tab.

Table 185 TCP tab fields

| Field | Description |
|-----------------|---|
| Active Opens | The number of TCP connections that were a direct transition to the SYN-SENT state from the CLOSED state. |
| Passive Opens | The number of TCP connections that were a direct transition to the SYN-RCVD state from the LISTEN state. |
| Failed Attempts | The number of TCP connections that were a direct transition to the CLOSED state from the SYN-SENT state or the SYN-RCVD state, and a direct transition to the LISTEN state from the SYN-RCVD state. |

Table 185 TCP tab fields (continued)

| Field | Description |
|-------------------------------|---|
| Resets In | The number of TCP connections that made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state. |
| Established | The number TCP connections that were a direct transition to the CLOSED state from the ESTABLISHED state or the CLOSE-WAIT state. |
| Segments In | The total number of segments received, including those received in error. This count includes segments received on currently established connections. |
| Segments Out | The total number of transmitted segments, including those on current connections, but excluding those that contain only retransmitted bytes. |
| Retransmitted Segments | The total number of retransmitted segments (TCP segments transmitted that contain one or more previously transmitted bytes). |
| Segments Received with Errors | The total number of received segments, including errors. This count includes segments received on currently established connections. |
| Resets Out | The number of transmitted TCP segments that contain the RST flag. |

TCP connection statistics

To view TCP connection statistics for IP routing:

- 1 From the Device view, select the Web Switching Module.

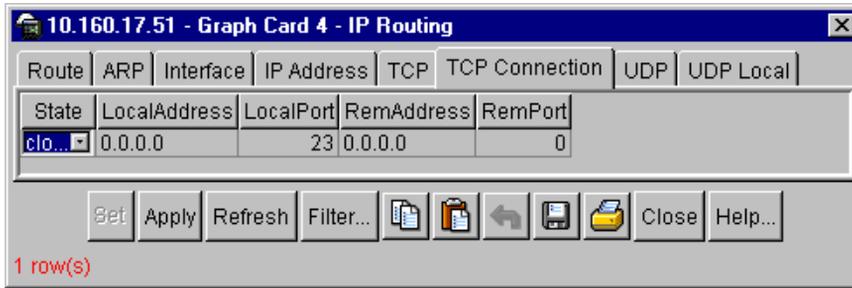
The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the Route tab.

- 3 Click the TCP Connection tab.

The [TCP Connection tab](#) (Figure 222) opens with the fields defined in [Table 186](#).

Figure 222 IP Routing—TCP Connection tab

The screenshot shows a window titled "10.160.17.51 - Graph Card 4 - IP Routing". It features several tabs: "Route", "ARP", "Interface", "IP Address", "TCP", "TCP Connection", "UDP", and "UDP Local". The "TCP Connection" tab is active, displaying a table with the following columns: "State", "LocalAddress", "LocalPort", "RemAddress", and "RemPort". A single row is visible with the values "clo...", "0.0.0.0", "23", "0.0.0.0", and "0". Below the table is a toolbar with buttons for "Set", "Apply", "Refresh", "Filter...", and several icons (document, folder, back, forward, printer). The text "1 row(s)" is displayed in red at the bottom left of the window.

| State | LocalAddress | LocalPort | RemAddress | RemPort |
|--------|--------------|-----------|------------|---------|
| clo... | 0.0.0.0 | 23 | 0.0.0.0 | 0 |

Table 186 describes the fields on the TCP Connection tab.

Table 186 IP Routing—TCP Connection tab fields

| Field | Description |
|-------------------|--|
| State | <p>The state of this TCP connection.</p> <ul style="list-style-type: none"> • LISTEN - waiting for a connection request from any remote TCP and port. • SYN-SENT - waiting for a matching connection request after having sent a connection request. • SYN-RECEIVED - waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTABLISHED - an open connection, data received can be delivered to the user. The normal state for the data transfer phase of the connection. • FIN-WAIT-1 - waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent. • FIN-WAIT-2 - waiting for a connection termination request from the remote TCP. • CLOSE-WAIT - waiting for a connection termination request from the local user. • CLOSING - waiting for a connection termination request acknowledgment from the remote TCP. • LAST-ACK - waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request). • TIME-WAIT - waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request. • CLOSED - no connection state at all. • Delete TCB - The only value which may be set by a management station, terminates the connection. |
| Local IP Address | The local IP address. |
| Local TCP Port | The number of the local port. |
| Remote IP Address | The remote IP address. |
| Remote TCP Port | The number of the remote port. |

UDP statistics

To view TCP connection statistics for IP routing:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

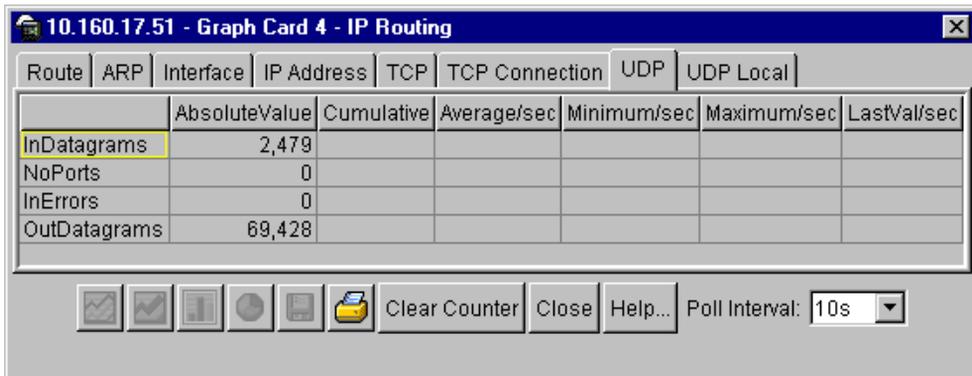
- 2 From the Graph menu, select WSM Card > IP Routing.

The IP Routing dialog box opens to the Route tab.

- 3 Click the UDP tab.

The [UDP tab](#) (Figure 223) opens with the fields defined in [Table 187](#).

Figure 223 IP Routing—UDP tab



- 4 Click a cell(s) to graph.
- 5 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing”](#) on page 516.

[Table 187](#) describes the fields on the UDP tab.

Table 187 IP Routing—UDP tab fields

| Field | Description |
|------------------------|---|
| Datagrams In | The total number of UDP datagrams delivered to UDP users. |
| No Application at Port | The total number of received UDP datagrams when no application was at the destination port. |

Table 187 IP Routing—UDP tab fields (continued)

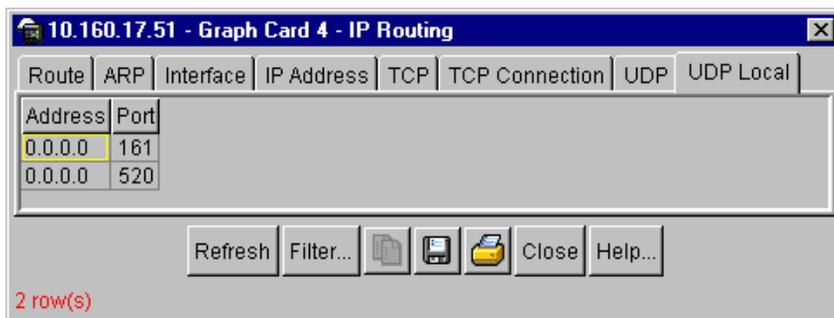
| Field | Description |
|-------------------|--|
| Dropped Datagrams | The number of received UDP datagrams that could not be delivered for reasons other than the absence of an application at the destination port. |
| Datagrams Out | The total number of delivered UDP datagrams. |

UDP local statistics

To view UDP local statistics for IP routing:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > IP Routing.
The IP Routing dialog box opens to the Route tab.
- 3 Click the UDP Local tab.

The [UDP Local tab](#) (Figure 224) opens with the fields defined in [Table 188](#).

Figure 224 IP Routing—UDP Local tab

[Table 188](#) describes the fields on the UDP Local tab.

Table 188 IP Routing—UDP Local tab fields

| Field | Description |
|---------|---|
| Address | The local IP address for the UDP listener. When the UDP listener accepts datagrams for any IP interface associated with the node, the address is 0.0.0.0. |
| Port | The local port number for the UDP listener. |

Virtual routing statistics

Virtual routing statistics include:

- [“Virtual routing state data” on page 564](#)
- [“Virtual routing statistics” on page 563](#)

Virtual routing statistics

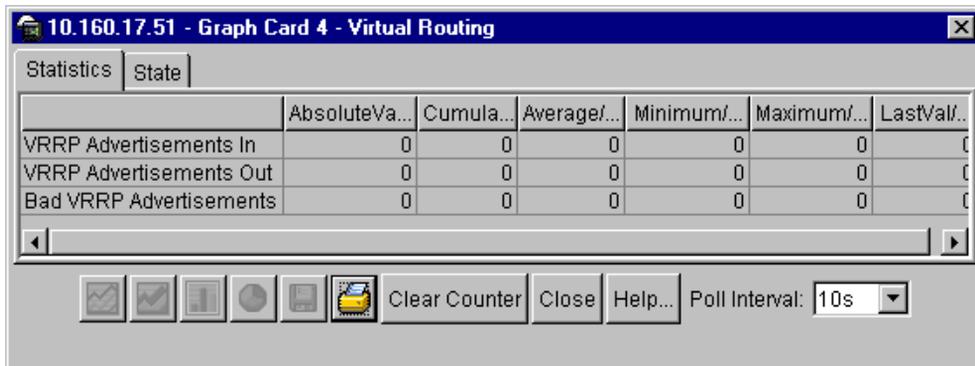
To view statistics for virtual routing:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > Virtual Routing.

The Virtual Routing dialog box opens to the [Statistics tab \(Figure 225\)](#) with the fields defined in [Table 189](#).

Figure 225 Virtual Routing—Statistics tab

- 3 Click a cell(s) to graph.
- 4 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 189](#) describes the fields on the Virtual Routing—Statistics tab.

Table 189 Virtual Routing—Statistics tab fields

| Field | Description |
|-------------------------|---|
| VRRP Advertisements In | The number of good VRRP advertisements that were received. |
| VRRP Advertisements Out | The number of good VRRP advertisements that were transmitted. |
| Bad VRRP Advertisements | The number of bad VRRP advertisements that were received. |

Virtual routing state data

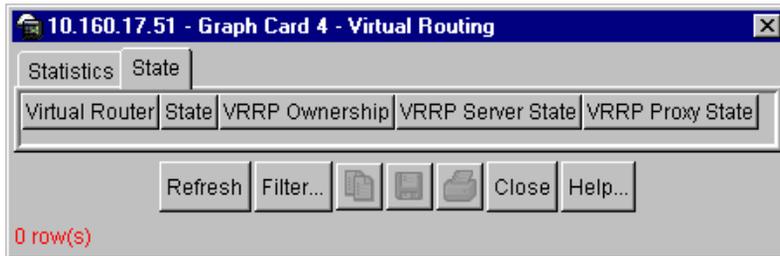
To view statistics for virtual routing:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Virtual Routing.
The Virtual Routing dialog box opens to the Statistics tab.

3 Click the State tab.

The **State** tab (Figure 226) opens with the fields defined in Table 190.

Figure 226 Virtual Routing—State tab fields



describes the fields on the Virtual Routing—State tab.

Table 190 Virtual—Routing State tab fields

| Field | Description |
|-------------------|--|
| Virtual Router | The index number of the VRRP virtual router. |
| State | The current state of the virtual router. <ul style="list-style-type: none"> initialize -- Waiting for a startup event. backup -- Monitoring the state/availability of the master router. master -- Forwarding IP addresses associated with this virtual router. |
| VRRP Ownership | The VRRP virtual router ownership status. |
| VRRP Server State | Identifies virtual routers that support Layer 4 services. |
| VRRP Proxy State | Identifies virtual proxy routers. |

Layer 4 statistics

Layer 4 statistics include:

- “Server load balance statistics” on page 566
- “Filter statistics” on page 575
- “Remote real servers data” on page 576
- “URL statistics” on page 577

Server load balance statistics

Layer 4 server load balance statistics include:

- “Real server port statistics” on page 566
- “Real servers statistics” on page 567
- “Groups statistics” on page 568
- “Virtual servers statistics” on page 569
- “Real servers state data” on page 571
- “Maintenance statistics” on page 572
- “DNS statistics” on page 573
- “TCP Limit statistics” on page 574

Real server port statistics

To view statistics for real server ports:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > Layer4 > SLB.

The Server Load Balancing dialog box (Figure 227) opens to the **Real Server Ports** tab with the fields defined in Table 191.

Figure 227 Server Load Balancing—Real Server Ports tab

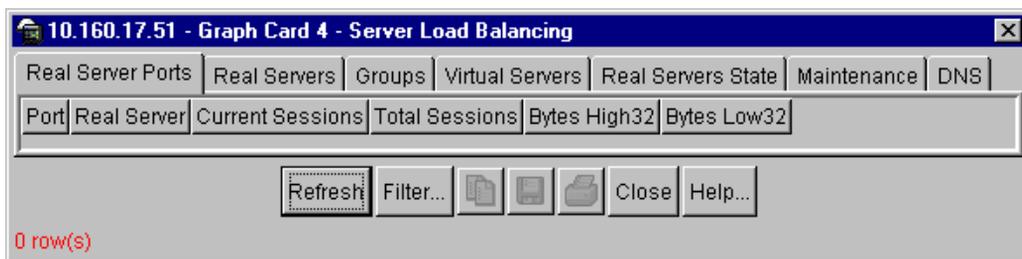


Table 191 describes the fields on the SLB—Real Server Ports tab.

Table 191 SLB—Real Server Ports tab fields

| Field | Description |
|------------------|--|
| Port | The index number that identifies the WSM port. |
| Real Server | The index number that identifies the real server. |
| Current Sessions | The current number of sessions handled by the real server. |
| Total Sessions | The total number of sessions handled by the real server. |
| Bytes High32 | The upper 32-bit value of all bytes received and transmitted through a specific port of the real server. |
| Bytes Low32 | The lower 32-bit value of all bytes received and transmitted through a specific port of the real server. |

Real servers statistics

To view server load balancing real servers statistics:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Layer4 > SLB.
The Server Load Balancing dialog box (Figure 235) opens to the Real Server Ports tab.
- 3 Click the Real Servers tab.
The Real Servers tab (Figure 228) opens with the fields defined in Table 192.

Figure 228 Server Load Balancing—Real Servers tab



Table 192 describes the fields on the SLB—Real Servers tab.

Table 192 SLB—Real Servers tab fields

| Field | Description |
|--------------------|---|
| Real Server | The index number that identifies the real server. |
| Current Sessions | The number of sessions currently handled by the real server. |
| Total Sessions | The total number of sessions handled by the real server. |
| Highest Sessions | The highest number of sessions handled by the real server. |
| Bytes High32 | The upper 32-bit value of the bytes received and transmitted through the real server. |
| Bytes Low32 | The lower 32-bit value of the bytes received and transmitted through the real server. |
| Times Claimed Down | The total number of times the real server was down. |

Groups statistics

To view statistics for server load balancing groups:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Layer4 > SLB.
The Server Load Balancing dialog box opens to the Real Server Ports tab.
- 3 Click the Groups tab.
The **Groups tab** (Figure 229) opens with the fields defined in Table 193.

Figure 229 Server Load Balancing—Groups tab

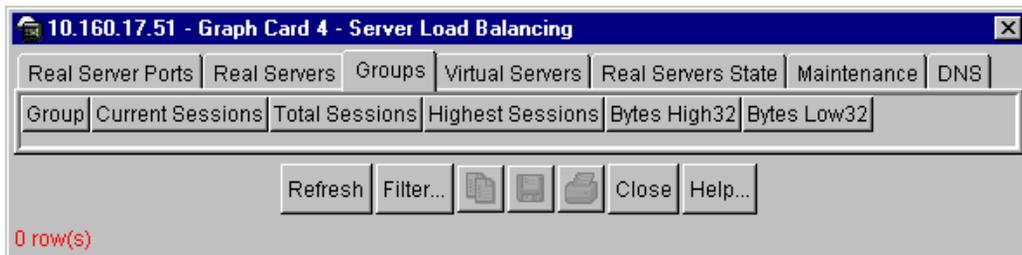


Table 193 describes the fields on the SLB—Groups tab.

Table 193 SLB—Groups tab fields

| Field | Description |
|------------------|--|
| Group | The index number of the real server group. |
| Current Sessions | The number of sessions currently handled by the real server group. |
| Total Sessions | The total number of sessions handled by the real server group. |
| Highest Sessions | The highest number of sessions that were handled by the real server group. |
| Bytes High32 | The upper 32-bit value of the bytes received and transmitted through the real server group. To obtain the real value in SNMP v1, both lower and upper bytes must be retrieved. |
| Bytes Low32 | The lower 32-bit value of the bytes received and transmitted through the real server group. To obtain the real value in SNMP v1, both lower and upper bytes must be retrieved. |

Virtual servers statistics

To view statistics for server load balancing virtual servers:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Layer4 > SLB.
The Server Load Balancing dialog box opens to the Real Server Ports tab.
- 3 Click the Virtual Servers tab.
The [Virtual Servers tab](#) (Figure 230) opens with the fields defined in Table 194.

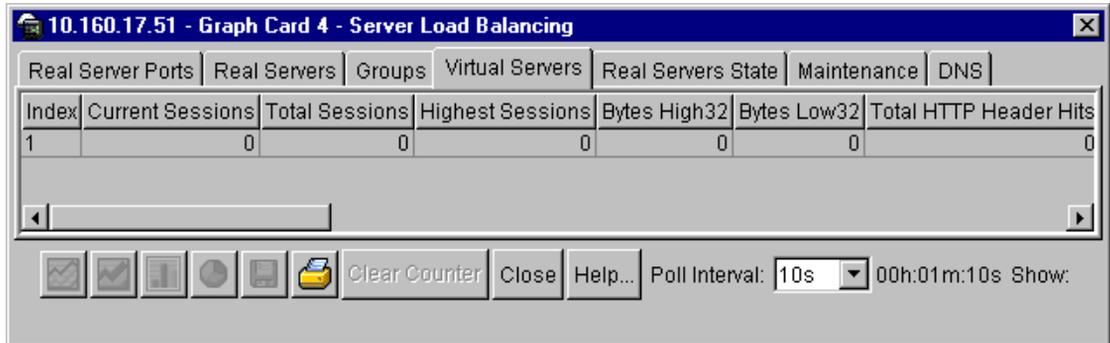
Figure 230 Server Load Balancing—Virtual Servers tab

Table 194 describes the fields on the SLB—Virtual Servers tab.

Table 194 SLB—Virtual Servers tab fields

| Field | Description |
|----------------------------|---|
| Index | The index number of the virtual server. |
| Current Sessions | The number of sessions that are currently handled by the virtual server. |
| Total Sessions | The total number of sessions that were handled by the virtual server. |
| Highest Sessions | The highest number of sessions that were handled by the real server. |
| Bytes High32 | The upper 32-bit value of that were received and transmitted through the real server group. |
| Bytes Low32 | The lower 32-bit value of bytes that were received and transmitted through the real server group. |
| Total HTTP Header Hits | The total number of HTTP header hits. |
| Current HTTP Header Hits | The current number of HTTP header hits. |
| Total HTTP Header Misses | The total number of HTTP header misses. |
| Current HTTP Header Misses | The current number of HTTP header misses. |
| Total HTTP Header Sessions | The total number of HTTP header sessions. |
| Total Cookie Rewrites | The total number of cookie rewrites. |
| Current Cookie Rewrites | The current number of cookie rewrites. |

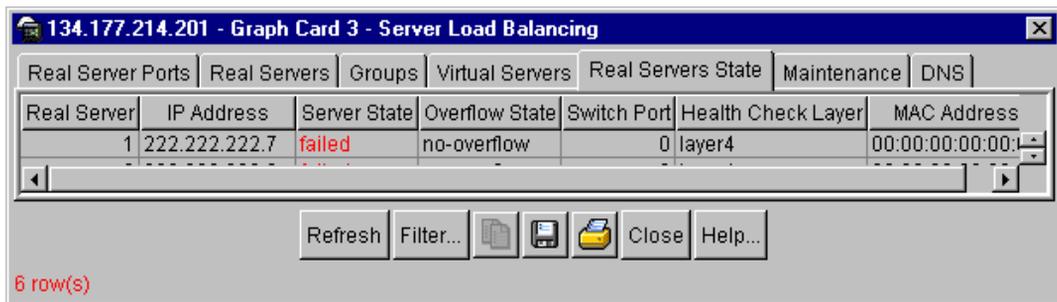
Table 194 SLB—Virtual Servers tab fields (continued)

| Field | Description |
|------------------------|---------------------------------------|
| Total Cookie Inserts | The total number of cookie inserts. |
| Current Cookie Inserts | The current number of cookie inserts. |

Real servers state data

To view real servers state data:

- From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- From the Graph menu, select WSM Card > Layer4 > SLB.
The Server Load Balancing dialog box opens to the Real Server Ports tab.
- Click the Real Servers State tab.
The [Real Servers State tab](#) (Figure 231) opens with the fields defined in [Table 195](#).

Figure 231 Server Load Balancing—Real Servers State tab

[Table 195](#) describes the fields on the SLB—Real Servers State tab.

Table 195 SLB—Real Servers State tab fields

| Field | Description |
|-------------|--------------------------------------|
| Real Server | The index number of the real server. |
| IP Address | The IP address of the real server. |

Table 195 SLB—Real Servers State tab fields (continued)

| Field | Description |
|--------------------|---|
| Server State | The state of the real server: running, failed, disabled, or other. |
| Overflow State | The overflow state of the real server: overflow or no-overflow. |
| Switch Port | The WSM port through which the real server is connected. |
| Health Check Layer | The OSI layer where the health of the real server is checked: layer3, layer4, or other. |
| MAC Address | The MAC address of the real server. |

Maintenance statistics

To view maintenance statistics for server load balance:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Layer4 > SLB.
The Server Load Balancing dialog box opens to the Real Server Ports tab.
- 3 Click the Maintenance tab.

The [Maintenance tab](#) (Figure 232) opens with the fields defined in [Table 196](#).

Figure 232 Server Load Balancing—Maintenance tab

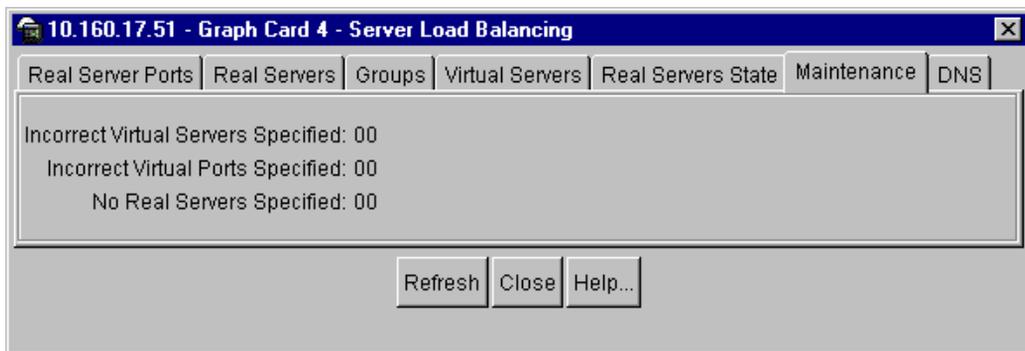


Table 196 describes the fields on the SLB—Maintenance tab.

Table 196 SLB—Maintenance tab fields

| Field | Description |
|-------------------------------------|--|
| Incorrect Virtual Servers Specified | The number of incorrectly specified virtual servers for server load balancing. |
| Incorrect Virtual Ports Specified | The number of incorrectly specified virtual ports for server load balancing. |
| No Real Servers Specified | No real servers are specified for server load balancing. |

DNS statistics

To view server load balancing DNS statistics:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

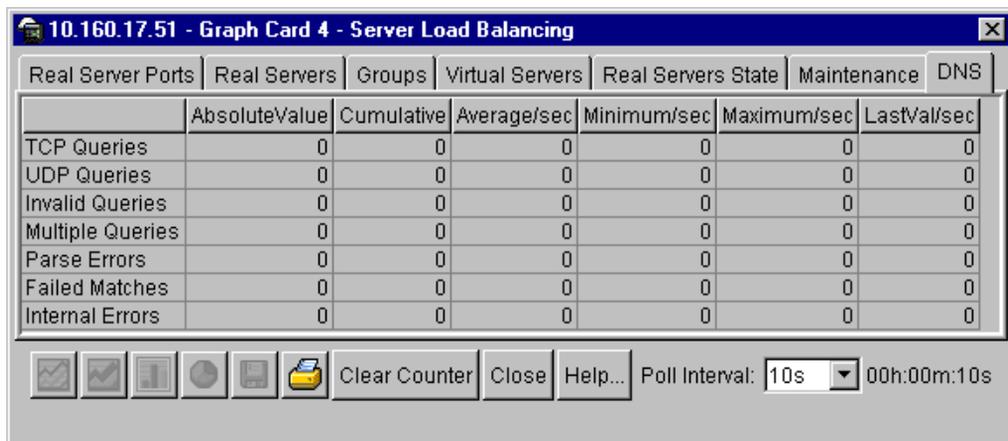
- 2 From the Graph menu, select WSM Card > Layer4 > SLB.

The Server Load Balancing dialog box opens to the Real Server Ports tab.

- 3 Click the Real Servers State tab.

The **DNS tab** (Figure 233) opens with the fields defined in Table 197.

Figure 233 Server Load Balancing—DNS tab



- 4 Click a cell(s) to graph.
- 5 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing”](#) on page 516.

[Table 197](#) describes the fields on the SLB—DNS tab.

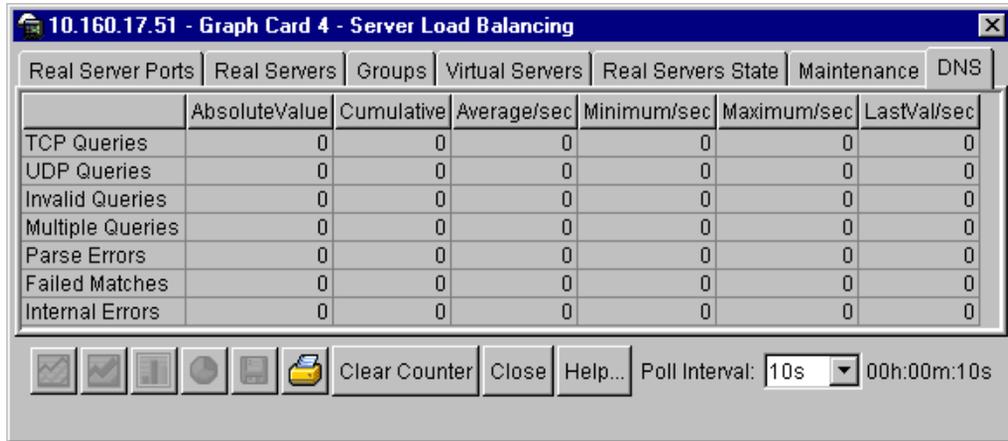
Table 197 SLB—DNS tab fields

| Field | Description |
|------------------|--|
| TCP Queries | Number of TCP DNS queries. |
| UDP Queries | Number of UDP DNS queries. |
| Invalid Queries | Number of invalid DNS queries. |
| Multiple Queries | Number of multiple DNS queries. |
| Parse Errors | Number of domain name parse errors. |
| Failed Matches | Number of failed real server name matches. |
| Internal Errors | Number of DNS parsing internal errors. |

TCP Limit statistics

To view server load balancing TCP Limit statistics:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Layer4 > SLB.
The Server Load Balancing dialog box opens to the Real Server Ports tab.
- 3 Click the TCP Limit tab.
The [TCP Limit tab](#) ([Figure 234](#)) opens with the fields defined in [Table 198](#).

Figure 234 Server Load Balancing—TCP Limit tab

- 4 Click a cell(s) to graph.
- 5 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 197](#) describes the fields on the SLB—TCP Limit tab.

Table 198 SLB—TCP Limit fields

| Field | Description |
|--------------------------|---|
| Client Hold Down Trigger | The number of client hold downs that were triggered. |
| Client State Entries | The current number of TCP rate limiting per-client state entries. |

Filter statistics

From the Statistics tab, you can monitor and graph statistics about the Layer 4 filters you have defined.

To configure Layer 4 filter statistics:

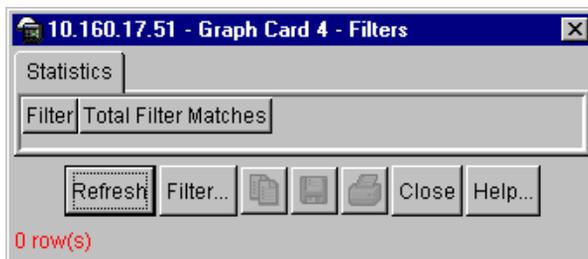
- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- From the Graph menu, select WSM Card > Layer4 > Filters.

The Filters dialog box (Figure 235) opens to the [Statistics tab](#) with the fields defined in [Table 199](#).

Figure 235 Graph—L4 Filters dialog box



[Table 199](#) describes the fields on the L4 Filters—Statistics tab.

Table 199 L4 filters—Statistics tab

| Field | Description |
|----------------------|--|
| Filter | The index number of the filters. |
| Total Filter Matches | The total number of received packets that matched the filter rule. |

Remote real servers data

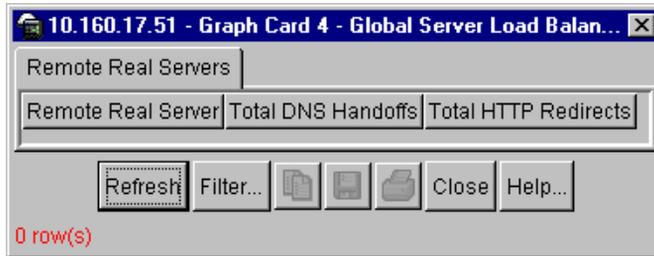
To view remote real servers data for global server load balancing:

- From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- From the Graph menu, select WSM Card > Layer4 > Global SLB.

The Global Server Load Balancing dialog box opens to the [Remote Real Servers tab](#) (Figure 236) with the fields defined in [Table 200](#).

Figure 236 Global Server Load Balancing—Remote Real Servers tab

[Table 200](#) describes the fields on the GSLB—Remote Real Servers tab.

Table 200 Global SLB—Remote Real Servers tab fields

| Field | Description |
|----------------------|--|
| Remote Real Server | The index number of the remote real server. |
| Total DNS Handoffs | The total number of DNS Handoffs by the remote real server. |
| Total HTTP Redirects | The total number of HTTP redirections by the remote real server. |

URL statistics

URL statistics include:

- [“URL load balance statistics” on page 578](#)
- [“Redirect statistics” on page 577](#)

Redirect statistics

To view statistics for URL path redirects:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Layer4 > URL Parsing.
The URL dialog box opens to the [Redirect](#) tab ([Figure 238](#)) with the fields defined in [Table 202](#).

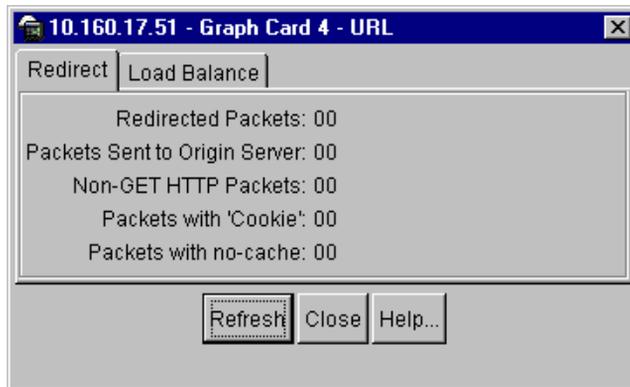
Figure 237 URL Parsing—Redirect tab

Table 201 describes the fields on the URL Parsing—Redirect tab.

Table 201 URL Parsing—Redirect tab fields

| Field | Description |
|-------------------------------|--|
| Redirected Packets | The number of received packets with matching specified URL expressions that were redirected to a specified group of real servers. |
| Packets Sent to Origin Server | The number of received packets that were delivered to the original server. This occurs when there is a mis-match with the specified URL expressions, or detecting HTTP non-GETs, user cookies, and no-cache as user configuration. |
| Non-GET HTTP Packets | The number of received packets that contained non-GETs methods, such as POST, HEAD, PUT, and so forth. |
| Packets with 'Cookie' | The number of received packets that contained a Cookie: header |
| Packets with No-Cache | The number of received packets that contained a no-cache header value. |

URL load balance statistics

To view statistics for URL load balance:

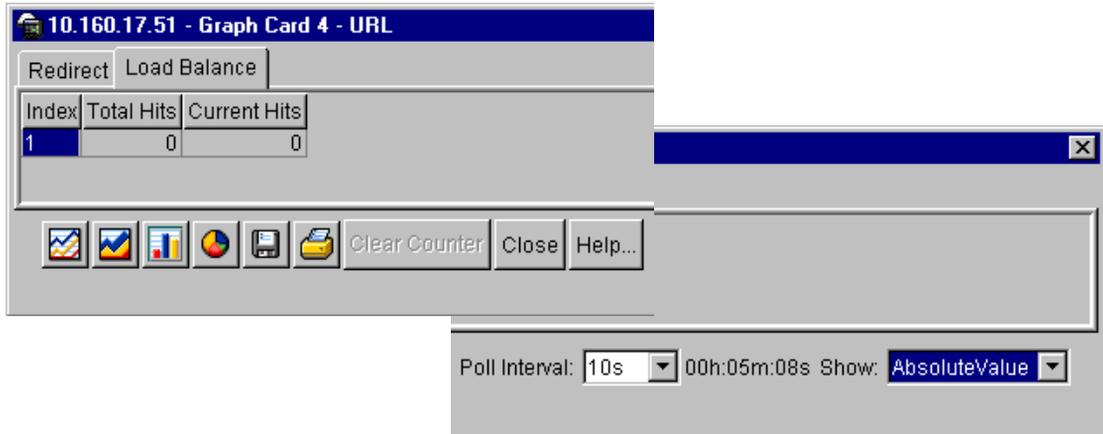
- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- From the Graph menu, select WSM Card > Layer4 > URL Parsing.

The URL dialog box opens to the **Load Balance** tab (Figure 238) with the fields defined in Table 202.

Figure 238 URL—Load Balance tab



- Click the Show field and choose a value to display (Absolute Value, Cumulative, Average/Second, Minimum/second, Maximum/second, or Last Value/second).
- Click a cell(s) to graph.
- Click a graph tool to view a graphic representation of the data. For more information, see “About statistics and graphing” on page 516.

Table 202 describes the fields on the URL Parsing—Load Balance tab.

Table 202 URL—Load Balance tab fields

| Field | Description |
|--------------|--|
| Index | The URL path table index. |
| Total Hits | The total number of instances that are load balanced because the URL path matches. |
| Current Hits | The current number of instances that are load balanced because of the particular URL path matches. |

RTSP statistics

To view statistics for Real Time Streaming Protocol (RTSP):

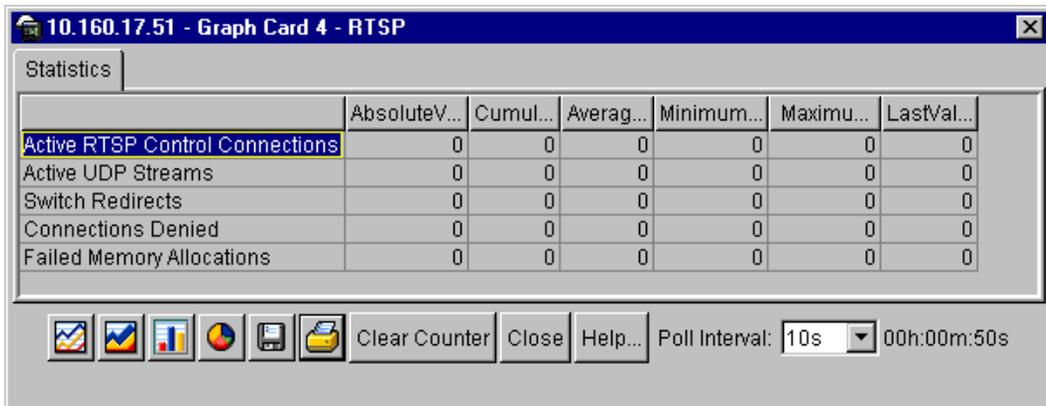
- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > Layer4 > RTSP.

The RTSP dialog box opens to the **Statistics** tab (Figure 239) with the fields defined in Table 203.

Figure 239 RTSP Statistics tab



- 3 Click a cell(s) to graph.
- 4 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 203](#) describes the fields on the RTSP—Statistics tab.

Table 203 RTSP—Statistics tab fields

| Field | Description |
|---------------------------------|--|
| Active RTSP Control Connections | The number of active RTSP control connections. |
| Active UDP Streams | The number of active User Datagram Protocol (UDP) streams. |
| Switch Redirects | The number of WSM redirects. |

Table 203 RTSP—Statistics tab fields (continued)

| Field | Description |
|---------------------------|--|
| Connections Denied | The number of connections denied because of the RTSP connection limit. |
| Failed Memory Allocations | The cases of heap memory failure allocations. |

Wireless Application Protocol statistics

Wireless Application Protocol statistics include:

- [“Add session statistics” on page 581](#)
- [“Delete session statistics” on page 582](#)

Add session statistics

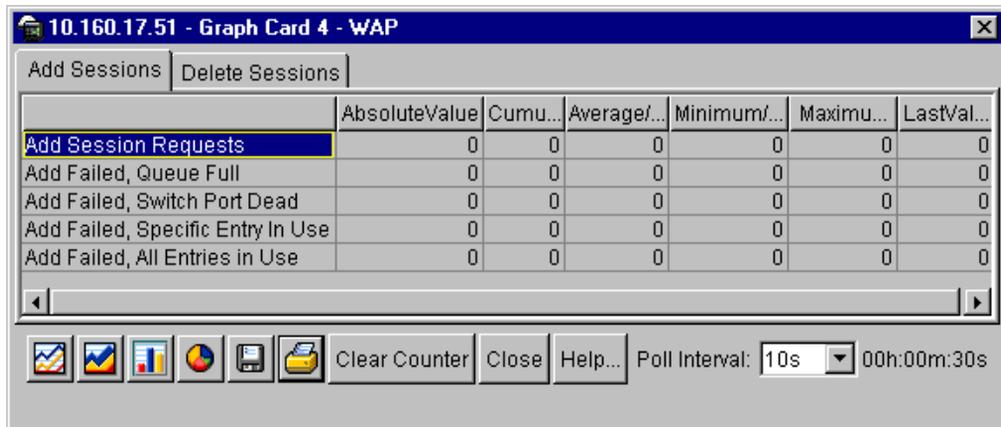
To view WAP add session statistics:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > Layer4 > WAP.

The WAP dialog box opens to the [Add Sessions tab \(Figure 240\)](#) with the fields defined in [Table 204](#).

Figure 240 Add Sessions tab

- 3 Click a cell(s) to graph.
- 4 Click a graph tool to view a graphic representation of the data. For more information, see [“About statistics and graphing” on page 516](#).

[Table 204](#) describes the fields on the WAP—Add Sessions tab.

Table 204 WAP—Add Sessions tab fields

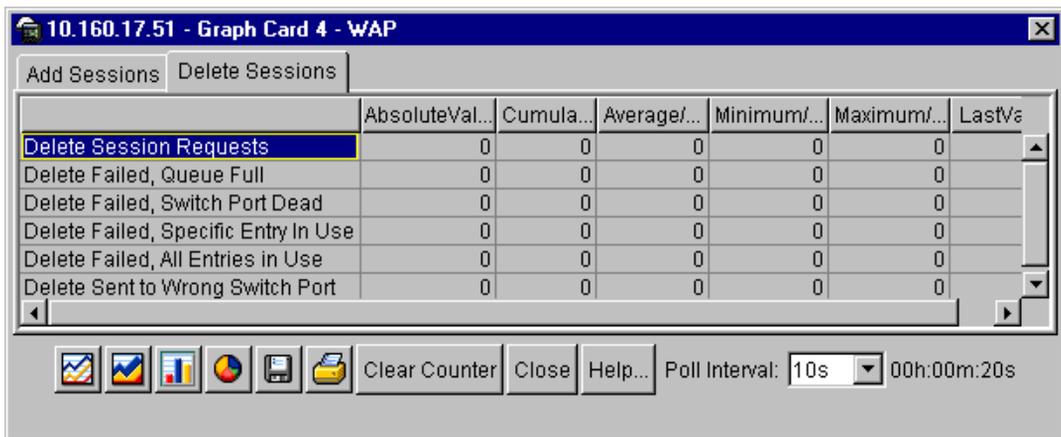
| Field | Description |
|-----------------------------------|--|
| Add Session Requests | The number of add session requests for Wireless Application Protocol (WAP). |
| Add Failed, Queue Full | The number of Add Session Requests that failed because the queue was full. |
| Add Failed, Switch Port Dead | The number of Add Session Requests that failed because the port was dead. |
| Add Failed, Specific Entry in Use | The number of Add Session Requests that failed because the specified entry was in use. |
| Add Failed, All Entries in Use | The number of Add Session Requests that failed because all entries were in use. |

Delete session statistics

To view WAP delete session statistics:

- 1 From the Device view, select the Web Switching Module.
The Web Switching Module is highlighted.
- 2 From the Graph menu, select WSM Card > Layer4 > WAP.
The WAP dialog box opens to the Add Sessions tab.
- 3 Click the Delete Sessions tab.
- 4 The Delete Sessions tab (Figure 241) opens with the fields defined in Table 205.

Figure 241 Delete Sessions tab



- 5 Click a cell(s) to graph.
- 6 Click a graph tool to view a graphic representation of the data. For more information, see “About statistics and graphing” on page 516.

Table 205 describes the WAP—Delete Sessions tab fields.

Table 205 WAP—Delete Sessions tab fields

| Field | Description |
|---------------------------|--|
| Delete Session Requests | The number of delete session requests for Wireless Application Protocol (WAP). |
| Delete Failed, Queue Full | The number of Delete Session Requests that failed because the queue was full. |

Table 205 WAP—Delete Sessions tab fields (continued)

| Field | Description |
|---------------------------------------|--|
| Delete Failed, Switch Port Dead | The number of Delete Session Requests that failed because the port was dead. |
| Deleted Failed, Specific Entry in Use | The number of Delete Session Requests that failed because the specified entry was in use. |
| Delete Failed, All Entries in Use | The number of Delete Session Requests that failed because all entries were in use. |
| Delete Sent to Wrong Switch Port | The number of Delete Session Requests that failed because they were delivered to the wrong port. |

Bandwidth management statistics

Bandwidth management statistics include:

- [“Traffic contract statistics” on page 584](#)
- [“Switch port traffic contract statistics” on page 585](#)

Traffic contract statistics

To view traffic contract statistics for bandwidth management:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > BWM.

The IP Routing dialog box opens to the [Traffic Contracts tab \(Figure 242\)](#) with the fields defined in [Table 206](#).

Figure 242 Bandwidth Management—Traffic Contracts tab

[Table 206](#) describes the fields on the BWM—Traffic Contracts tab.

Table 206 BWM—Traffic Contracts tab fields

| Field | Description |
|------------------|--|
| Contract | Number of the bandwidth management contract. |
| Name | Name of the traffic contract queue, such as URL, service, URLcont, host, or default. |
| Bytes Out | Total bytes sent out from the traffic contract queue. |
| Bytes Dropped | Total bytes dropped from the traffic contract queue. |
| Buffers Used | Total buffers used by the traffic contract queue. |
| Buffers Assigned | Total buffers assigned to the traffic contract queue. |

Switch port traffic contract statistics

To view WSM port traffic contract statistics for bandwidth management:

- 1 From the Device view, select the Web Switching Module.

The Web Switching Module is highlighted.

- 2 From the Graph menu, select WSM Card > BWM.

The IP Routing dialog box opens to the [Switch Port Traffic Contracts tab](#) ([Figure 243](#)) with the fields defined in [Table 207](#).

Figure 243 Bandwidth Management—Switch Port Traffic Contracts tab

| Switch Port | Contract | Name | Bytes Out | Bytes Dropped | Buffers Used | Buffers Assigned |
|-------------|----------|---------|-----------|---------------|--------------|------------------|
| 1 | 1024 | Default | 00 | 00 | 00 | 16,320 |
| 2 | 1024 | Default | 00 | 00 | 00 | 16,320 |
| 3 | 1024 | Default | 00 | 00 | 00 | 16,320 |
| 4 | 1024 | Default | 00 | 00 | 00 | 16,320 |
| 5 | 1024 | Default | 00 | 00 | 00 | 16,320 |
| 6 | 1024 | Default | 00 | 00 | 00 | 16,320 |
| 7 | 1024 | Default | 00 | 00 | 00 | 16,320 |
| 8 | 1024 | Default | 00 | 00 | 00 | 16,320 |

8 row(s)

Table 207 describes the fields on the BWM—Switch Port Traffic Contracts tab.

Table 207 BWM—Switch Port Traffic Contracts tab fields

| Field | Description |
|------------------|--|
| Switch Port | Number of the monitored WSM port. |
| Contract | Contract number. |
| Name | Name of the traffic contract queue, such as URL, service, URLcont, host, or default. |
| Bytes Out | Total number of bytes sent out from the traffic contract queue. |
| Bytes Dropped | Total number of dropped bytes from the traffic contract queue. |
| Buffers Used | Total number of buffers used by the traffic contract queue. |
| Buffers Assigned | Total number of buffers assigned to the traffic contract queue. |

Index

Numbers

- 1st and 2nd HTTP SLB precedence 239
- 1st syslog facility 174
- 1st trap host community string 178
- 4093, VLAN management IP address 102
- 802.1Q VLAN tagging 103

A

- action
 - L4 filter 197
 - port filter rule 93
- Actions menu 54
- active configuration 181
- active FTP 272
- advertise VIP host routes, RIP 146
- advertisement
 - interval
 - virtual router 158
 - VRRP group 164
- age
 - of route since last update 139
- aging timeout
 - bridge 118, 120
- Alarm Manager, toolbar 62
- Alert on RIP Failure Count, real server group 230
- allow filter
 - about 197
 - port setting 93
- Alteon WebSystems
 - MIB support 170
- altroot.mib 170

- application health check, about 417
- application port numbers 191
- Application Redirection
 - about 395
 - configuring IP proxy with 401
- Apply button 70
- applying changes 72
- area chart tool 518
- ARP
 - health check 412
- ARP only health check, gateway 136
- ARP statistics 552
- ARP table
 - reducing size 141
- authentication
 - failure trap, send 178
 - setting for VRRP interface 161
 - string, RADIUS secret 224
 - traps
 - enabling or disabling 176
- autonegotiation, port 87

B

- back ports 64, 81
- backup
 - configuration 181
 - link, port 88
 - or overflow servers, about 212
 - real server 221
 - real server group 229
 - server, for real server and group 229
 - specify either server or group 230

- Backup Group assignment, real server group 229
- bad remote site updates 385
- bandwidth
 - SLB metric
 - about 210
 - define for server group 229
 - update interval 225
- bandwidth management
 - about 493
 - burst limit 499
 - configuring 499
 - configuring contracts 505, 508
 - contract precedence 497
 - contracts 496
 - cookie-based 498
 - data pacing 495
 - frame discard 498
 - HTTP header-based 498
 - packer coloring (TOS bits) 499
 - policies 494
 - policy
 - configuring 502
 - statistics and history 494
 - synchronize 244
 - tabs in Device Manager
 - General 501
 - Policies 504
 - Traffic Contracts 507
 - URL-based 497
 - with Virtual Matrix Architecture 499
- bar chart tool 518
- Base port statistics 537
- basic FWLB 329
- BDPU
 - determining root cost 113
 - message format 109
 - root cost comparison 114
- best effort, BWM 495
- BFM ports 64, 81
- BGP syslog and SNMP trap 176
- binding, delayed 239, 265
- binding, persistent 238
- boot
 - code version 180
 - options, about 179
- BOOTP relay
 - about 129
 - enable or disable 132
- bridge
 - assign VLANs to STG 117
 - designated for STP port 127
 - designated SPT port 127
 - identifier, spanning tree port 124
 - MAC address 120
 - number of ports 120
 - port identifier, spanning tree port 125
 - spanning tree
 - about 117
 - configuring 118
 - statistics 546
 - STP Port tab fields 126
 - tabs in Device Manager
 - General 120
 - Spanning Tree 121
 - STG 116
 - STG Port 126
- Bridge Forwarding DB
 - tabs in Device Manager
 - Base Port 538
 - Forwarding 537
 - TP Port 540
- Browse button 70
- browser-based SLB 239
- browsers, online help 77
- browser-smart SLB, configuring 467
- burst limit, BWM 499
- BWM
 - configuring contracts 505, 508
 - policy
 - configuring 502
 - tabs in Device Manager
 - General 501

- Policies 504
- Traffic Contracts 507

BWM contract

- L4 switching filter 200
- port 88
- server service 240
- trunk group 97
- URL parsing 458
- virtual server 234
- virtual server, assigning 509
- VLAN 106

C

caching

- enable or disable for filter matches 197

changes, applying and saving 72

chart toolbar 518

charts, working with 518

chassis

- selecting 63
- shortcut menu 58
- slot numbering 96

check interval, SYN attack detection 268

CLI

- generated error syslog and SNMP trap 176

client

- on management network 262
- proxy
 - about 260
 - configuring IP 245
 - for redirect or NAT L4 filters 197
 - real server 222
 - state for SLB port 242

client-server session, about 211

client-server state for SLB port 242

Close button 70

color-codes, ports 66

community string

- 1st trap host 178

community strings

- configuring 178
- configuring for Device Manager log in 44
- default 45
- for WSM with 8600 169

configuration block 182

configuration to load for next reset 181

connection

- rate limiting, TCP 200
- timeouts, real servers, about 211

connections

- maximum allowed for WSM 135
- maximum for TCP rate limiting 200

console

- logging for L4 filters 197
- output, syslog 174
- syslog and SNMP trap 176

content load balancing

- enable or disable 240

content precedence lookup 488

contract

- BWM, assign to virtual server 509

contract precedence, bandwidth management 497

contract, BWM

- about 496
- L4 switching filter 200
- trunk group 97
- url-based 514
- VLAN 106

contracts, BWM, configuring 505, 508

conventions, text 37

cookie

- based BWM 498
- browsers that don't accept 438
- configuring persistence 442
- examples of values 443
- formats 437
- insert mode 439
- name 437
- offset 437
- passive mode 440
- permanent versus temporary 436

- [persistence mode, for virtual service](#) 238
- [persistence, about](#) 435
- [preferential SLB](#)
 - [configuring](#) 466
 - [define criteria](#) 466
 - [example](#) 465
- [preferential SLB, about](#) 464
- [rewrite mode](#) 441
- [using for session persistence](#) 434
- [virtual server](#) 233
 - [expire](#)
 - [length](#)
 - [name](#)
 - [offset](#)
 - [search](#)
 - [search response count](#)
- [cookie-based SLB](#) 239
- [Copy button](#) 70
- [CPU proxy, 8600](#) 169
- [current software image, viewing](#) 179
- [customer support](#) 40

D

DAM

- [enabling or disabling](#) 224, 259
- [data pacing](#) 495
- [data pacing, bandwidth management](#) 495
- [default BWM contract](#)
 - [trunk group](#) 97
 - [VLAN](#) 106
- [Default BWM Contract field](#)
 - [VLAN membership](#) 107
- [Default Configuration button](#) 70
- [default domain name, setting](#) 139
- [default filters](#)
 - [about](#) 191
 - [configuring](#) 192
- [default gateway](#)
 - [forwarding to](#) 141
- [default gateway metric, configure](#) 135

- [default gateways](#)
 - [configuring](#) 132
- [default STG](#) 108
- [default VLAN](#)
 - [name](#) 105
 - [port assignment](#) 87
- [delayed binding](#)
 - [configuring for service](#) 239, 265
 - [configuring for Web Cache Redirection](#) 398
 - [preventing DoS attacks](#) 262
- [Delete button](#) 70
- [demilitarized zone](#) 331
- [deny filter](#)
 - [about](#) 197
 - [port setting](#) 93
- [deny filter, for Layer 7](#) 492
- [description](#)
 - [port, NIC](#) 85
- [description of device opened in Device Manager](#) 186
- [destination address](#)
 - [NAT setting for port filter](#) 93
- [destination address filter type, MAC or IP](#) 198
- [destination IP address](#)
 - [for port filter](#) 93
- [destination IP address, L4 switching filter](#) 198
- [destination IP mask](#)
 - [for port filter](#) 93
- [destination IP mask, L4 switching filter](#) 198
- [destination MAC mask, L4 switching filter](#) 198
- [destination port range](#)
 - [L4 switching filter](#) 198
- [device](#)
 - [opening](#) 43
 - [viewing](#) 62
- [Device Manager](#)
 - [about](#) 41
 - [dialog boxes](#) 66
 - [help files](#) 77

- menu bar 47
 - opening a device 43
 - starting 42
 - toolbar 61
 - window, parts of 46
 - Device menu 47
 - DHCP 129
 - dialog boxes
 - about 66
 - buttons in 70
 - Direct Access Mode
 - enabling or disabling 224, 259
 - direct IP route 138
 - direct server return
 - about 255
 - configuring 256
 - health check, about 413
 - health check, configuring 413
 - disable server processing on a port 246
 - discarding non-IP traffic, port 87
 - distributed site state updates, GSLB 385
 - DMZ 331
 - DNS
 - configuring servers 139
 - domain name
 - DNS, IP routing 140
 - handoff to peer sites, GSLB 384
 - health check 407
 - IP configuration fields 140
 - local address response only, GSLB 384
 - query load balancing 240
 - response with at least one address, GSLB 385
 - server load balancing
 - about 277
 - configuring TCP-based 281
 - configuring UDP-based 278
 - tasks 278
 - single address response, GSLB 385
 - SLB 470
 - time to live response, GSLB 384
 - domain name
 - for internal lookup table, GSLB 393
 - SMTP mail gateway 186
 - virtual server 233
 - domain name, configuration block 140
 - DoS
 - monitor 266
 - using delayed binding 262
 - DSR
 - about 255
 - configuring 256
 - health check, about 413
 - DSR health check, configuring 413
 - DSSP updates, GSLB 385
 - duplicating content, SLB 206
- ## E
- Edit Selected, toolbar 61
 - Edit WSM Card menu 50
 - election priority, virtual router 157
 - enabled software version 180
 - enforce policy, bandwidth management 501
 - Ethernet statistics, port 543
 - examples of cookie values 443
 - exclude string match
 - real server 222
 - excluding non-cacheable sites 402
 - exclusionary string matching 485
 - expire cookie, parameters
 - Export button 70
 - Export data button 70
 - expression
 - for URL parsing 479
 - non-cacheable for URL-based WCR 474
 - non-cacheable RTSP URL 472
- ## F
- failure
 - server, about 405

- service, about 404
- failure retries, real server 221
- fast-age period, session table 224
- FE port autonegotiation 87
- filter
 - button 70
 - creating 201
 - default
 - about 191
 - configuring 192
 - fields 196
 - for VLAN 197
 - for Web Cache Redirection 482
 - L4, about 189
 - Layer 7 deny 492
 - membership fields, port 93
 - monitoring and graphing 575
 - port
 - apply 90
 - enable or disable 89
 - removing 94
 - protocol numbers used in 190
 - redirect
 - configuring for WAN link SLB 318
 - show for port 90
 - state 196
 - statistics 576
 - synchronize configuration 244
 - syslog and SNMP trap 176
 - turn on for port 90
- Filter Membership tab, port 93
- filter settings
 - port 90
- filtering tables 73
 - add filter 74
 - removing 75
 - removing all 75
 - replacing 76
- firewall redirect hash, L4 switching filter 200
- firewall SLB
 - about 327
 - basic 329
 - basic, configuring 335
 - demilitarized zone 331
 - free-metric 330
 - methods of 329
 - monitoring 353
 - operation 331
- first virtual IP address, GSLB 393
- flag matching, L4 switching filter 199
- flow control
 - port FE 87
- flow control, port 85
- format for URL string 455
- formats, numeric 69
- forward directed broadcasts, enable or disable for WSM 135
- forwarding delay, spanning tree bridge 121
- Forwarding statistics 536
- fragment, UDP remapping 239
- Frame discard 498
- frame flow, VPN 359
- free-metric firewall SLB 330
- front ports 64, 81
- FTP health check 407
- FTP parsing 240
- FTP server load balancing
 - configuring 273
- FTP server load balancing, about 272
- Full NAT, URL-based WCR 476
- FWLB
 - about 327
 - basic 329
 - basic, configuring 335
 - demilitarized zone 331
 - free-metric 330
 - methods of 329
 - monitoring 353
 - operation 331

G

- gateway
 - destination for static route 150
 - VLAN assigned 136
- gateway, metric 135
- gateways fields, IP routing 136
- General tab
 - bridge 120
 - port 85
- geographic awareness, GSLB 385
- GET operation 169
- Gigabit Ethernet autonegotiation, port setting 87
- Gigabit Ethernet flow control, port 87
- global SLB
 - about 373
 - configuring 375
 - syslog and SNMP trap 176
 - tabs in Device Manager
 - General 384
 - Lookup 393
 - Network Preferences 393
 - Remote Real Servers 577
 - Remote Sites 378
- good remote site updates, GSLB 385
- graceful failover
 - enabling or disabling 224
- Graph button 70
- Graph Selected, toolbar 61
- Graph WSM Card menu 53
- graphing, about 516
- group number, real server group 228
- group parameters, VRRP 162
- GSLB
 - about 373
 - configuring 375
 - network preferences, about 390
 - syslog and SNMP trap 176
 - tabs in Device Manager
 - General 384

- Lookup 393
- Network Preferences 393
- Remote Real Servers 577
- Remote Sites 378

H

- Half NAT, URL-based WCR 476
- half-open sessions
 - threshold 268
- handoff to peer sites, DNS, for GSLB 384
- hardware revision number of opened device 187
- hash
 - filter setting 200
- hash length
 - virtual server 234
- hash SLB metric
 - about 208
 - define for server group 229
- hashing based on url 239
- hashing, URL SLB 468
- header hash HTTP SLB 239
 - configuring 469
- health check
 - about 403
 - configuring for real server group 408
 - configuring interval and retries for 410
 - content to send to WAP gateway 299
 - default gateway 136
 - define content for group 229
 - expected content from WAP gateway 299
 - interval, real server 221
 - type, real server group 229
 - WSP port 299
 - WTLS port 299
- Hello interval
 - spanning tree bridge 121
- Help
 - button in dialog box 70
 - menu 56
 - opening 77

- toolbar button 61
- high destination port
 - for TCP/UDP port filter 93
- high source port
 - for TCP/UDP port filter 93
- history, bandwidth management 494
- history, sending BWM 501
- hold duration for TCP rate limiting 225
- hold interval, spanning tree bridge 122
- horizontal graph tool 518
- host
 - virtual, about 462
- host name
 - virtual service 238
- hosting, virtual SLB 239
- hot standby processing
 - SLB port 242
- HTTP header name 482
 - virtual server 234
- HTTP header, bandwidth management 498
- HTTP health check 417
 - configuring 419
- HTTP health check, allow on any port 225
- HTTP redirect to real server, GSLB 385
- HTTP redirect, GSLB peer sites 384
- HTTP SLB precedence, for server services 239
- HTTP SLB, configuring header hash 469

I

- ICMP
 - message type to filter 200
- ICMP health check 412
- ICMP received statistics 521
- image 1 and 2 software versions 181
- image, select software to run 181
- IMAP health check 407
 - configuring 420

- index
 - number of internal network preference table 393
 - number of URL path for parsing 479
 - RTSP URL expression 472
- index number
 - IP gateway 136
 - L4 filter 196
 - local IP route 143
 - port filter 93
 - service for real server port 254
 - spanning tree group 116, 126
 - static route 150
 - virtual router 157
 - virtual router group 164
- indirect IP route 138
- Insert button 70
- insert mode, cookie persistence 238
- interface
 - number for an IP route 138
 - number, VRRP 161
- interface configuration fields 132
- interface statistics, IP routing 553
- internal network preference table, number 393
- inter-switch processing
 - for SLB port 242
- interval
 - health check 410
 - remote site updates, GSLB 385
 - update, SLB metric 225
- Intrusion Detection
 - hash setting, L4 switching filters 200
 - metrics supported 303
 - real server group 225
 - SLB
 - about 302
 - configurations required for 304
 - configuring real servers for 304
 - enabling 311
 - SLB enable or disable for port 243
- invalid IP route 138

- invert logic
 - L4 switching filter setting 197
 - port filter setting 93
- IP
 - filter option 200
 - syslog and SNMP trap 176
 - TOS filter settings 200
- IP address
 - assigning multiple 260
 - configure proxy for port 247
 - configuring a proxy 245
 - destination filter type 198
 - for management network 224, 261
 - for management subnet mask 224, 262
 - for route destination 138
 - for VLAN 4093 management 102
 - format to use in Device Manager 69
 - gateway 136
 - GSLB remote site 378
 - interface 132
 - mask for virtual and real server 224
 - next hop for a route 138
 - of gateway for static route 150
 - of static route destination 150
 - persistent mask 224
 - primary DNS server 140
 - primary virtual server, GSLB 393
 - range for local route cache 141
 - range for SLB 222
 - real server 218, 219
 - real server to include in group 230
 - secondary DNS server 140
 - secondary virtual server, GSLB 393
 - server load balancing
 - configuring 270
 - server load balancing, about 269
 - SLB peer 244
 - SMTP mail gateway 186
 - source filter type 198
 - source, network preferences 393
 - statistics 555
 - virtual router 157
 - virtual server 233
 - VRRP interface 161
- IP broadcast address 132
- IP forward local 143
- IP forwarding
 - enabling 136
- IP forwarding state
 - disable for port 147
- IP forwarding, enable or disable for the WSM 135
- IP gateway index number
 - IP gateway index 136
- IP interface
 - about 129
 - configure VRRP authentication 159
 - configuring 130
 - defining for the WSM 225
 - number for source static route 150
 - VRRP group 164
- IP port statistics 539
- IP proxy
 - configuring with Application Redirection 401
- IP proxy for non-HTTP redirects
 - about 386
 - configuring 388
 - example of 387
- IP routes, viewing 137
- IP routing
 - configuring RIP 143
 - graphing tabs in Device Manager
 - ARP 553
 - Interface 554
 - IP Address 556
 - Route 552
 - TCP 557
 - TCP Connection 560
 - UDP 561
 - UDP Local 563
 - IP forwarding per port 146
 - menu 47
 - tabs in Device Manager
 - DNS 140

- Gateways 136
 - General 135
 - Interfaces 132
 - Local 143
 - Ports 147
 - RIP 145
 - Routes 138
 - Static Routes 150
- IP statistics 519
- IP subnet mask
- for route 139
 - for static IP route 150
 - interface 132
 - local IP route 143
 - of the network table 393
- IPX routing menu 47
- ## J
- Jumbo frames
- about 103
 - about isolating 104
 - field
 - VLAN membership 107
 - routing to non-Jumbo frame VLANs 104
 - select VLANs for STG 117
 - VLAN, enable or disable 105
- ## L
- L4 Switching
- menu 52
- L4 switching
- tabs in Device Manager
 - Filters 196
- last change, port operational state 85
- last topology change, spanning tree bridge 121
- Layer 3 only,virtual server 233
- Layer 4 filters
- statistics 576
- Layer 7 deny filter 492
- Layer 7 RTSP SLB 471
- LDAP health check 408, 430
- LDAP version 225
- learned discards, bridge 120
- learn-forward transitions
- bridge STP port 127
- learn-forward transitions, spanning tree port 125
- least connections SLB metric
- about 209
 - define for a server group 229
- LEDs 64, 65
- Legend of port colors 57
- viewing 66
- line chart tool 518
- link
- health check, about 415
 - root cost comparison 114
- link state
- monitoring 89
 - port 85
- link trap, port 85
- link, port backup 88
- list, edit selection 69
- listen to default routes 145
- listen to route updates 145
- load balance state
- configure 246
 - field for SLB port 242
- Load Balance tab, URL parsing 458
- load sharing, virtual routing 158
- load sharing, master virtual router 165
- local host address range 141
- local IP configuration fields 143
- local networks, add or remove from route
- cache 141
- local route cache, set IP address range 141
- log scale tool 518
- logging
- send filter info to console port 197

-
- lookup
 - content precedence 488
 - Lookup tab, GSLB 393
 - loops, preventing 110
 - low destination port
 - for TCP/UDP port filter 93
 - low source port
 - for TCP/UDP port filter 93
 - lowest cost port 121
 - M**
 - MAC address
 - bridge 120
 - destination filter type 198
 - format in Device Manager 69
 - source filter type 198
 - substitution 222, 239
 - WSM, port 85
 - management
 - syslog and SNMP trap 176
 - management console
 - syslog output 174
 - management IP address, default for VLAN 4093 102
 - management network
 - IP address for 224, 261
 - monitoring real servers and services 261
 - management subnet mask 224, 262
 - mapping
 - about multiple real server ports 250
 - configurations required for multiple ports 251
 - multiple port example 250
 - multiple ports to virtual server port 252
 - ports
 - about 260
 - virtual server port to real server port 248
 - mask
 - for virtual and real server 224
 - max connections
 - real server 220
 - maximum age
 - spanning tree bridge 121
 - maximum connections, real server 212
 - menu bar 47
 - message, non-binding, to send to client 482
 - methods, SLB 205
 - metric, default gateway, configure 135
 - metric, SLB
 - about 207
 - bandwidth
 - about 210
 - update interval 225
 - define for server group 229
 - hash, about 208
 - least connections, about 209
 - minimum misses, about 207
 - response time
 - about 210
 - update interval 225
 - round robin, about 209
 - metrics, SLB
 - for Intrusion Detection 303
 - MIBs
 - AlteonWebSystems 170
 - standard 171
 - tigon 170
 - minimum GSLB site connections 385
 - minimum misses SLB metric
 - about 207
 - define for server group 229
 - mirrored port 185
 - MLT
 - configuring 98
 - default configuration 95
 - deleting 99
 - dynamic assignment 96
 - restrictions 95
 - trunk
 - configuration fields 97
 - viewing configuration 96
 - VLAN membership, viewing 108
-

- mode
 - port 85
 - port FE 87
 - Modify button 70
 - monitoring port for port mirroring 185
 - MP CPU statistics 533
 - MTU, port 85
 - multiple spanning trees
 - about 109
 - example 110, 112
 - multiple strings, assigning 491
- ## N
- Name
 - VLAN membership 107
 - name
 - cookie 437
 - DNS domain 140
 - GSLB remote site 378
 - host for virtual service 238
 - L4 filter 196
 - port 85
 - real server 220
 - real server group 228
 - real servers to include in group 230
 - select VLANs for STG 117
 - SMTP host 186
 - SMTP user, bandwidth management 501
 - VLAN 105
 - name HTTP header 482
 - naming ports 83
 - NAT
 - filter
 - about 197
 - port setting 93
 - URL-based for WCR 476
 - NAT active FTP, enable or disable 199
 - NAT Session Timeout setting 199
 - Network Address Translation
 - disable for server service 239
 - port filter 93
 - source or destination address for L4 switching filter 199
 - URL-based for WCR 476
 - network preference lookups 393
 - network preferences fields, GSLB 393
 - network preferences, GSLB 390
 - NewRoot trap 171
 - next hop, IP address 138
 - NNTP health check 407
 - no cookie operation
 - real server 221
 - No NAT server services 239
 - No NAT, URL-based WCR 476
 - non-cacheable expressions 474
 - non-cacheable sites, excluding from redirection 402
 - number
 - for L4 filter 196
 - GSLB remote sites 378
 - hardware revision of device 187
 - index of static route 150
 - interface 132
 - internal network preference table 393
 - IP interface for a route 138
 - local IP route 143
 - part, of opened device 186
 - peer index assignment 244
 - port for virtual service 237
 - port, spanning tree bridge 127
 - real group for virtual service 237
 - real server 218, 219
 - real server assignment for port 254
 - service for real server port index 254
 - service port 254
 - spanning tree group index 116, 126
 - TCP/UDP application port 191
 - trunk group 97
 - virtual port 237
 - virtual router 157
 - virtual router group 164

virtual server, for a virtual service 237
 virtual server, URL-based BWM 514
 virtual service 237
 VLAN ID 105
 VLAN, for an interface 132
 numbers, how to format 69

O

objects
 editing 67
 selecting 63
 offset into WSP packet 299
 offset, cookie 437
 online Help 77
 ONMS version support 169
 Open Device dialog box 44
 Open Device Home Page, toolbar 62
 Open Device, toolbar 61
 opening a device 43
 troubleshooting 45
 operational status
 port 85
 operator
 for URL SLB 239
 orientation, port 64, 82
 OSM version support 169
 overflow servers, assigning 212

P

packet coloring type of service, BWM 499
 packets, received by monitoring port 185
 parsing, FTP 240
 part number of opened device 186
 passive FTP 272
 passive mode
 cookie persistence 238
 password

VRRP interface authentication 161
 Paste button 71
 path cost
 bridge STP port 127
 spanning tree port 124
 path, URL-based BWM 514
 pattern matching, RTSP SLB 471
 peer, index number 244
 persistence 443
 about 433
 configuring cookie-based 442
 cookie
 formats 437
 name 437
 offset 437
 permanent versus temporary 436
 cookie-based 435
 insert cookie mode 439
 mode 238
 passive cookie mode 440
 rewrite cookie mode 441
 SSL session-based
 about 448
 configuring 450
 example 449
 using cookie for 434
 using SSL ID for 435
 persistent binding, virtual service 238
 persistent mask 224
 ping interval, default gateway 136
 ping retry count, gateway 136
 poison reverse 145
 policies, bandwidth management 494
 policy
 bandwidth management 504
 pop3 health check 407
 port
 about multiple VLANs on 103
 allow HTTP health check on any 225
 back 81

- client processing, enable 242
 - client versus server processing 206
 - color, legend 57
 - viewing 66
 - configuration fields 87
 - configure IP proxy for 247
 - configure server processing 246
 - configure spanning tree protocol 123
 - configure VLAN tagging 88
 - configuring 85
 - configuring for SLB 240
 - CPU statistics 549
 - default trunk configuration 95
 - dynamic MLT assignment 96
 - editing 83
 - enable or disable UDP balancing 237
 - field
 - assign for SLB 242
 - filter
 - apply 90
 - enable or disable 89
 - removing 94
 - filter membership fields 93
 - front, back, BFM 81
 - general tab fields 85
 - interface statistics 540
 - IP forwarding configuration 147
 - map virtual server to real server 248
 - mapping
 - about 260
 - mapping multiple example 250
 - mapping multiple to virtual server port 252
 - mapping multiple, configurations required 251
 - monitor link state 89
 - naming 83
 - number
 - bridge STP port 127
 - for service 254
 - for virtual service 237
 - spanning tree group 126
 - virtual 237
 - orientation 64, 82
 - range for TCP/UDP L4 switching filter 198
 - redirection, L4 switching filter 199
 - server processing, enable 242
 - shortcut menu 60
 - spanning tree fields 124
 - spanning tree protocol
 - enable 125
 - specify for WSP health check 299
 - specify for WTLS health check 299
 - state
 - spanning tree group 126
 - statistics tabs in Device Manager
 - Bridge 547
 - CPU Statistics 550
 - Ethernet 544
 - Interface 541
 - Load Balance 548
 - status 66
 - synchronize configuration 244
 - tabs in Device Manager
 - Filter Membership 93
 - TCP/UDP application numbers 191
 - trunk group assigned 97
 - trunking 95
 - VLAN membership 105
 - VLAN membership, about 101
 - VLANs and spanning trees 109
 - with multiple VLANs 103
- port mirroring
 - about 182
 - configuration settings 185
 - configuring 183
 - enabling 185
 - state 186
 - ports
 - about multiple mapping for real server 250
 - select VLANs for STG 117
 - Ports field
 - VLAN membership 107
 - precedence, for HTTP SLB 239
 - preemption
 - master router election, VRRP group 165
 - preemption, virtual router 158
 - preferential SLB

- cookie-based
 - configuring 466
 - define criteria 466
 - cookie-based, about 464
 - example 465
- preferred link, port setting 88
- primary DNS server 139
- primary DNS server IP address 140
- primary IP address, GSLB remote site 378
- primary virtual server, GSLB 393
- Print button 71
- Print Table button 71
- print table tool 518
- priority
 - master router election, VRRP group 165
 - root bridge, STP port 127
 - spanning tree bridge 119, 121
 - virtual router 157
- priority field
 - spanning tree port 124
- priority tracking, VRRP 167
- product support 40
- protocol
 - for L4 switching filter 198
 - for port filter 93
- protocol numbers 190
- protocol type, spanning tree bridge 121
- protocol, routing 138
- proxy
 - about 260
 - client, for redirect or NAT L4 filters 197
 - configure for port 247
 - configuring IP address 245
 - IP address
 - enable translation 222
 - for non-HTTP redirects 386
 - setting for port 242
 - synchronize 244
 - with VMA 246, 325
 - IP state for SLB port 242

- Transparent Proxy Cache Protocol 224
- publications
 - hard copy 39
 - related 39

Q

- QOS menu 47

R

RADIUS

- account statistics 535
- health check 408
 - configuring 421
- secret, authentication string 224
 - configuring 421
- snooping
 - configuring filter for 299
 - configuring with WAP 294
 - WAP, enable for L4 filter 200
 - with WAP, about 292
- range, IP address for SLB 222
- rate limiting, TCP connection 200
- Read Community, SNMP 44
- Read-Write-All access 44
- real server
 - about mapping multiple ports 250
 - backup, about 212
 - configuration fields 219
 - configuring 216
 - configuring for Intrusion Detection SLB 304
 - configuring WAN link SLB for 314
 - connection timeouts 211
 - enable or disable 218, 220
 - exclusionary string matching 485
 - maximum connections 212
 - modifying health check interval and retries 410
 - monitoring by management network 261
 - number 218, 219
 - number for port 254
 - numbers assigned to group 228
 - port mapping 248

- to include in group 230
- real server group
 - configuring 225
 - configuring a health check for 408
 - configuring for WAN link SLB 316
 - define metric for 229
 - fields 228
 - group number 228
 - Intrusion Detection 225
 - number
 - for virtual service 237
 - redirection filter 199
 - select servers to include 230
- Real Time Streaming Protocol
 - Layer 7 SLB 471
 - tab in Device Manager 472
 - URL balancing 239
 - WCR
 - about 399
 - configuring 399
- re-ARP period, setting for WSM 135
- reassembly timeout for WSM 135
- receive flow control, port FE 87
- received ICMP statistics 521
- redirect
 - application
 - configuring IP proxy 401
 - application, about 395
 - based on URI header 481
 - excluding non-cacheable sites from 402
 - graphing data 577
 - monitoring 577
 - non-HTTP, IP proxy for 386
 - requests for URL WCR 479
 - statistics 577
 - to origin server
 - cookies 481
 - HTTP GETs 481
 - no-cache 481
 - to peer site, GSLB 384
 - to real server name, GSLB 385
 - URL, Web cache 473
- redirect filter
 - about 197
 - for WAN link SLB
 - about 313
 - port setting 93
 - real server group setting 199
 - real server port setting 199
 - WAN link SLB
 - configuring 318
- Redirection tab 481
- Refresh button 71
- Refresh Device, toolbar 61
- regular expression matching 486
- remap UDP fragments 239
- remote monitoring, port 87
- remote real server statistics 576
- remote site update interval, GSLB 385
- remote site updates, GSLB 385
- Remote Sites tab, GSLB 378
- Reset changes button 71
- Resize Columns button 71
- response time SLB metric
 - about
 - about 210
 - define for server group 229
 - update interval 225
- retransmit timeouts for WSM 135
- retries, health check 410
- return to sender,SLB port 243
- rewrite mode, cookie persistence 238
- RFC 959 272
- RIP statistics 525
- RIP, configuring 143
- RIP, enable or disable 145
- RMON
 - menu 47
 - port 87
- root

- bridge port priority 124
- bridge STP port 127
- cost
 - bridge STP port 127
 - spanning tree bridge 121
- cost comparison, link speed 114
- cost, BDPU 113
- cost, spanning tree port 124
- identifier, spanning tree bridge 121
- port, spanning tree bridge 121
- spanning tree port 124
- round robin
 - default gateway metric 135
 - SLB metric
 - about 209
- round robin SLB metric
 - define for server group 229
- route cache
 - adding or removing local networks 141
- route configuration fields 138
- route statistics 551
- route type 138
- Routing Information Protocol
 - configuration fields 145
 - enable or disable 145
- RTS, SLB port 243
- RTSP
 - implementation comparison 285
 - Layer 7 SLB 471
 - server load balancing
 - about 283
 - configuring for layer 4 286
 - configuring for layer 7 290
 - enable or disable 239
 - pre-configuration 285
 - tab in Device manager 472
- RTSP WCR
 - about 399
 - configuring 399
- running software version 180
- run-time configuration, saving 62

S

- save pending actions 181
- SaveRun-Time Config, toolbar 62
- saving changes 72
- script health check 408, 416
 - configuring 417
- second virtual server, GSLB 393
- secondary DNS server 139
- secondary DNS server IP address 140
- secondary IP address, GSLB remote site 378
- select browser type for SLB 239
- selection list, editing 69
- Send BWM History button 71
- send history, bandwidth management 501
- sent ICMP statistics 523
- server failure
 - about 405
- server failure, graceful 224
- Server Load Balance
 - Layer 7 RTSP 471
 - persistence-based 433
 - port statistics 547
 - tabs in Device Manager
 - URL-based BWM 514
- Server Load Balancing
 - about 203, 213
 - configuring ports 240
 - DNS 277, 470
 - enabling or disabling 224, 243
 - example 214
 - FTP 272
 - general steps for configuring 215
 - Global, about 373
 - graphing tabs in Device Manager
 - DNS 574, 575
 - Groups 569
 - Maintenance 573
 - Real Server Ports 567
 - Real Servers 568

- Real Servers State 571
- Virtual Servers 570
- Intrusion Detection 302
- IP address ranges
 - about 222
 - configuring 222
- IP address-based 269
- Layer 3 only 233
- methods 205
- metric
 - about 207
 - bandwidth, about 210
 - define for group 229
 - hash, about 208
 - least connections, about 209
 - minimum misses, about 207
 - response time, about 210
 - round robin, about 209
- real server
 - configuring 216
- real server group
 - configuring 225
- RTSP 283
- syslog and SNMP trap 176
- tabs in Device Manager
 - General 224
 - Peers 244
 - Ports 242
 - Real Server Group 228
 - Real Server Port 254
 - Real Servers 219
 - RTSP 472
 - SYN Attack 268
 - Synchronize 244
 - Virtual Server 233
 - Virtual Server Services 237
- topology rules 206
- URL-based 453
- virtual server
 - configuring 230
 - configuring services for 235
- WAN Link 313
- server state for SLB port 242
- Servers button 71
- Servers button, real server group 230
- service
 - configuring SLB
 - IP 270
 - UDP-based DNS 278
 - delayed binding for 265
 - DNS
 - about SLB for 277
 - configurations for SLB 278
 - configuring TCP-based SLB 281
 - configuring UDP-based SLB 278
 - FTP
 - about SLB for 272
 - configuring SLB for 273
 - monitoring by management network 261
 - port number for 254
 - RTSP
 - about 283
 - configuring layer 4 SLB 286
 - configuring layer 7 SLB 290
 - implementation comparison 285
 - preconfiguration tasks 285
 - URL-based BWM 514
 - WAP
 - about 292
 - enable or disable 299
 - WAP RADIUS snooping
 - about 292
 - configuring 294
- service failure
 - about 404
- Services button 71, 234
- services, configuring for a virtual server 235
- session persistence
 - about 433
 - configuring cookie-based 442
 - cookie
 - formats 437
 - name 437
 - offset 437
 - permanent versus temporary 436

- examples of cookie values 443
- insert cookie mode 439
- passive cookie mode 440
- rewrite cookie mode 441
- SSL-based 435
 - about 448
 - configuring 450
 - example 449
 - using cookie for 434
- session table aging period 224
- session, client-server 211
- Set button 71
- SET operation 169
- shortcuts
 - about 57
 - chassis 58
 - port 60
 - WSM 59
- simple text password authentication, VRRP
 - interface 161
- single address response, GSLB 385
- single spanning tree example 111
- site connections, minimum, GSLB 385
- SLB peers
 - configuration fields 244
- slot numbering, chassis 96
- slow-age period, session table 224
- SMTP health check 407
- SMTP mail host 186
- SMTP user name, BWM 501
- SNMP
 - 1st trap host 178
 - agent 169
 - generate authentication failure traps 178
 - generic supported traps 171
 - Read Community 44
 - spanning tree traps 171
 - syslog 173
 - enabling or disabling traps 175
 - traps 172
 - enabling or disabling 175
 - Write Community 44
- SNMP statistics 526
- soft limit, BWM 495
- software
 - image details 180
 - select image to run 181
 - version, about 179
- source address filter type, MAC or IP 198
- source and destination port range
 - TCP/UDP port filter 93
- source IP address
 - for port filter 93
 - L4 Switching, Filters 198
 - using for session persistence 434
- source IP mask
 - for port filter 93
- source IP mask, L4 switching filter 198
- source MAC address, L4 switching filter 198
- source port range
 - L4 switching filter 198
- source-address
 - NAT setting for port filter 93
- spanning tree
 - syslog and SNMP trap 176
- Spanning Tree Group
 - VLAN assignment 106
- Spanning Tree Protocol
 - about 108
 - and VLANs 109
 - assigning STG 114
 - configuring 114
 - configuring spanning tree port 123
 - default STG 114
 - enable on port 125
 - multiple spanning trees
 - about 109
 - example 110, 112
 - root cost
 - comparison 114

- example 113
- single spanning tree example 111
- spanning tree bridge 117
 - configuring 118
- Tabs in Device Manager
 - Bridge STG 116
 - Bridge STP Port 126
- tabs in Device Manager
 - Bridge STG Port 126
 - Port Spanning Tree 124
- spanning tree traps 171
- speed
 - port 85
 - port FE 87
- SSH/RADIUS
 - syslog and SNMP trap 176
- SSH/Radius field, syslog trap 176
- SSL session ID-based persistence 435
 - about 448
 - configuring 450
 - example 449
- SSLH health check 408
- stacked filters 196
- stacked graph tool 518
- state
 - bandwidth management 501
 - filter 196
 - filter rule, enable or disable for port 93
 - gateway, enabled or disabled 136
 - GSLB, for WSM 384
 - GSLB, remote site 378
 - interface, enable or disable 132
 - port mirroring 186
 - port, on or off 87
 - port, operational 85
 - real server 218, 220
 - spanning tree group 116
 - spanning tree port 124, 126
 - trunk group 97
 - virtual server 233
 - VLAN 105
 - VLAN membership 107
 - VLANs for STG 117
 - VRRP group 165
 - VRRP interface 161
- stateful failover
 - synchronize 244
 - update period 244
- static route
 - configuring 148
 - deleting 150
 - tab 150
- statistics, about 516
- statistics, bandwidth management 494
- statistics, BWM, sending to user 511
- Stats 531
- status
 - bridge STP port 127
- status bar, Device Manager 66
- status update, GSLB remote sites 378
- stealth mode, Intrusion Detection 304
- STG
 - about multiple 109
 - configuration fields 116
 - configuring 114
 - default 108, 114
 - example of single 111
 - select VLAN for 117
 - VLAN membership 116
- Stop button 71
- strict default gateway metric 135
- string
 - for URL parsing 479
 - for URL path 458
 - formats for URLs 455
 - health content expected from WAP gateway 299
 - health content to send to WAP gateway 299
- substitute source MAC address 222, 239
- success retries
 - real server 221
- supply route updates 145

- supply static route updates 145
- support, Nortel Networks 40
- SYN attack
 - about detecting 266
 - configuring detection 267
 - detection interval 268
 - detection, syslog and SNMP trap 176
- SYN requests 262
- synchronization, enable with SLB peer switch 244
- synchronizing fields 244
- sysDescr 186
- syslog
 - 1st facility 174
 - console output 174
 - generating system messages 173
 - traps, enabling or disabling 175
- system
 - syslog and SNMP trap 176
- sysUpTime 186
- T**
- table filters 73
 - adding 74
 - removing 75
 - removing all 75
 - replacing 76
- TCP connection statistics 558
- TCP health check 415
- TCP rate limiting
 - hold duration for 225
 - time window for 225
- TCP Rate Limiting field, syslog trap 176
- TCP statistics 556
- TCP-based SLB, configuring 281
- technical publications 39
- technical support 40
- Telnet Display, toolbar 61
- temperature sensor readings 187
- text conventions 37
- text field, editing 68
- three-way handshake 262, 266
- threshold
 - half-open sessions 268
- tigon MIBs 170
- Time format 69
- time to live DNS response, GSLB 384
- time to live, default for WSM 135
- time window, for TCP rate limiting 225
- timeout
 - real server 221
- timeout algorithm for WSM 135
- toolbar 61
- toolbar for graphs 518
- topology changes, spanning tree bridge 121
- topology rules, SLB 206
- TopologyChange trap 171
- TPCP 224
- tracking
 - enabling for virtual router 158
 - enabling for VRRP group 165
 - setting increments for virtual routing 153
- traffic type to mirror 185
- transmission distance, port maximum 85
- transmit flow control, port FE 87
- Transparent Proxy Cache Protocol 224
- Trap Log, toolbar 61
- traps, supported generic 171
- tree
 - select spanning tree to configure 119, 122
- trunk
 - BFM ports to back ports 81
 - group
 - configuring 98
 - default configuration 95
 - deleting 99
 - dynamic MLT assignment 96

- restrictions 95
 - viewing configuration 96
 - tab 97
 - VLANs and spanning trees 109
 - type
 - authentication, VRRP interface 161
 - bridging 120
 - filter destination address, MAC or IP 198
 - filter source address, MAC or IP 198
 - ICMP message to filter 200
 - of IP route 138
 - of packets to mirror 185
 - port interface 85
 - real server 221
 - URL-based NAT for WCR 476
- ## U
- UDP balancing 237
 - UDP fragment remapping 239
 - UDP local statistics 562
 - UDP statistics 560
 - UDP-based SLB, configuring 278
 - UDPDNS health check 408
 - configuring 419
 - update interval, metric 225
 - update period, RIP 146
 - updates, GSLB remote site 385
 - upgrade, WSM software 179
 - URI hash length 481
 - URL
 - bandwidth management 497
 - button 71
 - for technical documentation 39
 - graphics tabs in Device Manager
 - Redirect 578
 - hashing for SLB 468
 - hashing, SLB precedence 239
 - NAT for WCR 476
 - path
 - path membership, real server 222
 - path, index number 479
 - path, URL parsing 458
 - redirection, enable or disable 199
 - RTSP balancing 239
 - SLB 453
 - configuring 456
 - monitoring 462
 - SLB precedence 239
 - string formats 455
 - WCR 473
 - URL parsing
 - tabs in Device Manager
 - Expressions 479
 - Load Balance 458
 - Redirection 481
 - URL-based WCR
 - configuring 476
- ## V
- version, LDAP 225
 - VIP Health Checking, enable 230
 - virtual hosting
 - about 462
 - configure host header 464
 - configuring 463
 - virtual hosting SLB 239
 - Virtual LAN, membership fields 105
 - Virtual Matrix Architecture
 - and proxy IP addresses 246
 - bandwidth management 499
 - enabling or disabling 224
 - using with proxy IP address 325
 - Virtual Matrix Architecture, enabling 323
 - virtual private network
 - about 357
 - frame flow illustration 359
 - monitoring 372
 - SLB example 360
 - SLB, configuring 360
 - virtual router

- tracking 153
 - Virtual Routing
 - graphing tabs in device manager
 - State 565
 - Statistics 564
 - tabs in Device Manager
 - General 153
 - Groups 164
 - Interfaces 161
 - virtual server
 - about 213
 - add or modify services 234
 - assign number 233
 - configuring 230
 - configuring services for 235
 - fields for configuring 233
 - GSLB 393
 - port mapping 248
 - services configuration 237
 - URL-based BWM 514
 - VLAN
 - 4093, management IP address 102
 - 802.1Q tagging, about 103
 - button 71, 116
 - configure STG for 114
 - configure tagging 88
 - configuring 106
 - current members, viewing 104
 - for L4 switching filter 197
 - gateway assigned to 136
 - ID number 105
 - ID number for STG 117
 - IP interface requirement 103
 - membership 107
 - membership in STG 116
 - menu 47
 - MLT membership
 - viewing 108
 - number for an interface 132
 - port tag setting 87
 - ranges for spanning tree group 116
 - restrictions 101
 - select STG for 117
 - spanning trees and trunks 109
 - STG membership, default 109
 - syslog and SNMP trap 176
 - tagging
 - BDPU 113
 - using Jumbo frames 104
 - VMA
 - and proxy IP addresses 246
 - using with proxy IP address 325
 - VMA, enabling 323
 - VPN
 - about 357
 - frame flow, illustration 359
 - monitoring 372
 - SLB example 360
 - SLB, configuring 360
 - VRID 157, 164
 - VRRP
 - about 151
 - authentication, configuring 159
 - configuring 153
 - enabling 151
 - group parameters, configuring 162
 - priority tracking, configuring 167
 - synchronize priorities 244
 - syslog and SNMP trap 176
 - tabs in Device Manager
 - Routers 157
- ## W
- WAN link SLB
 - about 313
 - configurations required 313
 - configuring filters for 318
 - configuring real server groups for 316
 - configuring real servers for 314
 - enable or disable 200
 - preconfiguration 314
 - WAP
 - about 292
 - configuring SLB 294
 - debug level, setting 299

- enable or disable SLB 299
- gateway
 - health content expected from 299
 - health content to send 299
- gateway health check 423
- health check
 - WSP content 424
- RADIUS snooping
 - configuring filter for 299
 - enable or disable 200
- Web Cache Redirection
 - about RTSP redirection 399
 - configuring 396
 - configuring delayed binding for 398
 - content-intelligent 472
 - monitoring statistics 483
 - RTSP, configuring 399
 - URL-based 473
 - configuring 476
- Web interface, opening 62
- WebOS statistics 529
- weight
 - bandwidth, server 210
 - for directing GSLB connections to local site 385
 - real server 220
 - response time, server 210
- Wireless Application Protocol
 - about 292
 - debug level, setting 299
 - gateway
 - health content expected from 299
 - health content to send 299
 - General tab 299
 - health check 423
 - WSP content 424
 - RADIUS snooping
 - about 292
 - configuring 294
 - configuring filter for 299
 - enable or disable for L4 switching filter 200
 - SLB
 - configuring 294
 - enable or disable 299
 - Write Community, SNMP 44
 - WSM device view 64
 - WSM shortcut menu 59
 - WSM statistics
 - tabs in Device Manager
 - ICMP 522
 - ICMP Out 524
 - IP 520
 - MP CPU Stats 534
 - RADIUS Account 536
 - RIP 526
 - SNMP 528
 - Stats 532
 - WebOS Stats 530
 - wsp health check 408
 - WSP port to health check 299
 - WTLS health check 408
 - WTLS port to health check 299