

Dell EMC Networking N-Series N1100-ON, N1500, N2000, N2100-ON, N3000, N3000-ON, N3100-ON, and N4000 Switches

User's Configuration Guide

Version 6.5.2.x - N2000/N2100-ON/N3000/N3000-ON/
N3100-ON/N4000 Series Switches
Version 6.4.x.x - N1100-ON Series Switches



**Regulatory Models: E04W, E05W, E06W,
E07W, E15W, E16W, E17W, E18W,
PowerConnect 8132, PowerConnect 8132F,
PowerConnect 8164, PowerConnect 8164F**

Notes and Cautions



NOTE: A NOTE indicates important information that helps you make better use of your computer.



CAUTION: A CAUTION indicates potential damage to hardware or loss of data if instructions are not followed.

Information in this publication is subject to change without notice.

Copyright © 2018 Dell EMC Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Dell EMC™ and the Dell EMC logo are trademarks of Dell EMC Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Marketing Models: N1108T-ON, N1108P-ON, N1124T-ON, N1124P-ON, N1148T-ON, N1148P-ON, N1524, N1524P, N1548, N1548P, N2024, N2024P, N2048, N2048P, N2128PX-ON, N3024, N3024F, N3024P, N3048, N3048P, N3024EP-ON, N3024ET-ON, N3024EF-ON, N3048ET-ON, N3048EP-ON, N3132PX-ON, N4032, N4032F, N4064, N4064F

Regulatory Models: E04W, E05W, E06W, E07W, E15W, E16W, E17W, E18W, PowerConnect 8132, PowerConnect 8132F, PowerConnect 8164, PowerConnect 8164F

June 2018 Rev. A06

Contents

| | | |
|---|--|----|
| 1 | Introduction | 57 |
| | About This Document | 57 |
| | Audience | 58 |
| | Document Conventions | 58 |
| | Additional Documentation | 59 |
| 2 | Switch Feature Overview | 61 |
| | System Management Features | 62 |
| | Multiple Management Options | 62 |
| | System Time Management | 62 |
| | Log Messages | 63 |
| | System Reset | 63 |
| | Integrated DHCP Server | 63 |
| | Management of Basic Network Information | 64 |
| | IPv6 Management Features | 64 |
| | Dual Software Images | 64 |
| | File Management | 65 |
| | Switch Database Management Templates | 65 |
| | Automatic Installation of Firmware and Configuration | 65 |
| | sFlow | 66 |
| | SNMP Alarms and Trap Logs | 66 |
| | CDP Interoperability Through ISDP | 67 |
| | Remote Monitoring (RMON) | 67 |
| | N3000 Series Advanced and Advanced-Lite Firmware Images | 67 |

| | |
|---|-----------|
| Stacking Features | 70 |
| Mixed and Single Series Stacking | 70 |
| Single IP Management | 71 |
| Master Failover with Transparent Transition. | 71 |
| Nonstop Forwarding on the Stack | 72 |
| Hot Add/Delete and Firmware Synchronization | 72 |
| Security Features | 73 |
| Configurable Access and Authentication Profiles | 73 |
| Password-Protected Management Access | 73 |
| Strong Password Enforcement | 73 |
| TACACS+ Client. | 73 |
| RADIUS Support | 74 |
| SSH/SSL. | 74 |
| Inbound Telnet Control | 74 |
| Denial of Service | 74 |
| Port Protection | 75 |
| Captive Portal | 76 |
| 802.1X Authentication (IEEE 802.1X) | 76 |
| MAC-Based 802.1X Authentication. | 77 |
| 802.1X Monitor Mode | 77 |
| Port Security | 77 |
| Access Control Lists (ACLs) | 78 |
| Time-Based ACLs. | 78 |
| IP Source Guard (IPSG). | 78 |
| DHCP Snooping | 79 |
| Dynamic ARP Inspection | 79 |
| Protected Ports (Private VLAN Edge). | 79 |
| Green Technology Features | 80 |
| Energy Detect Mode | 80 |
| Energy Efficient Ethernet | 80 |
| Power Utilization Reporting. | 81 |

| | |
|--|-----------|
| Power over Ethernet (PoE) Plus Features | 81 |
| Key PoE Plus Features for the Dell EMC Networking N1108P-ON, N1124P-ON, N1148P-ON, N2024P, N2048P, N2128PX-ON, N3024P, N3048P, N3024EP-ON, N3048EP-ON, and N3132PX-ON Switches | 81 |
| Power Over Ethernet (PoE) Plus Configuration | 82 |
| PoE Plus Support | 82 |
| PoE 60W Support | 83 |
| Powered Device Detection | 83 |
| PoE Power Management Modes | 83 |
| Power Management in Guard Band | 85 |
| PoE Plus Default Settings | 86 |
| | |
| Switching Features | 88 |
| Flow Control Support (IEEE 802.3x) | 88 |
| Head of Line Blocking Prevention | 88 |
| Alternate Store and Forward (ASF) | 88 |
| Jumbo Frames Support | 88 |
| Auto-MDI/MDIX Support | 89 |
| VLAN-Aware MAC-based Switching | 89 |
| Back Pressure Support | 89 |
| Auto-negotiation | 90 |
| Storm Control | 90 |
| Port Mirroring | 90 |
| Static and Dynamic MAC Address Tables | 91 |
| Link Layer Discovery Protocol (LLDP) | 91 |
| Link Layer Discovery Protocol (LLDP) for Media Endpoint Devices | 91 |
| Connectivity Fault Management (IEEE 802.1ag) | 92 |
| Priority-based Flow Control (PFC) | 92 |
| Data Center Bridging Exchange (DBCx) Protocol | 93 |
| Enhanced Transmission Selection | 93 |
| Cisco Protocol Filtering | 93 |

| | |
|---|------------|
| DHCP Layer-2 Relay | 94 |
| Virtual Local Area Network Supported Features | 95 |
| VLAN Support | 95 |
| Port-Based VLANs | 95 |
| IP Subnet-based VLAN | 95 |
| MAC-based VLAN. | 95 |
| IEEE 802.1v Protocol-Based VLANs. | 95 |
| Voice VLAN | 96 |
| GARP and GVRP Support | 96 |
| Guest VLAN | 96 |
| Unauthorized VLAN. | 96 |
| Double VLANs. | 96 |
| Spanning Tree Protocol Features | 98 |
| Spanning Tree Protocol (STP) | 98 |
| Spanning Tree Port Settings | 98 |
| Rapid Spanning Tree | 98 |
| Multiple Spanning Tree. | 98 |
| Bridge Protocol Data Unit (BPDU) Guard. | 99 |
| BPDU Filtering. | 99 |
| RSTP-PV and STP-PV. | 99 |
| Link Aggregation Features. | 100 |
| Link Aggregation | 100 |
| Link Aggregate Control Protocol (LACP) | 101 |
| Multi-Switch LAG (MLAG) | 101 |
| Routing Features | 102 |
| Address Resolution Protocol (ARP) Table Management | 102 |
| VLAN Routing | 102 |
| IP Configuration. | 102 |
| Open Shortest Path First (OSPF) | 103 |
| Border Gateway Protocol (BGP) | 103 |

| | |
|--|------------|
| Virtual Routing and Forwarding (VRF) | 103 |
| BOOTP/DHCP Relay Agent | 104 |
| IP Helper and DHCP Relay | 104 |
| Routing Information Protocol | 104 |
| Router Discovery | 104 |
| Routing Table | 104 |
| Virtual Router Redundancy Protocol (VRRP) | 105 |
| Tunnel and Loopback Interfaces | 105 |
| IPv6 Routing Features | 106 |
| IPv6 Configuration | 106 |
| IPv6 Routes | 106 |
| OSPFv3 | 106 |
| DHCPv6 | 106 |
| Quality of Service (QoS) Features | 107 |
| Differentiated Services (DiffServ) | 107 |
| Class Of Service (CoS) | 107 |
| Auto Voice over IP (VoIP) | 107 |
| Internet Small Computer System Interface (iSCSI) Optimization | 108 |
| Layer-2 Multicast Features | 108 |
| MAC Multicast Support | 108 |
| IGMP Snooping | 108 |
| IGMP Snooping Querier | 109 |
| MLD Snooping | 109 |
| Multicast VLAN Registration | 109 |
| Layer-3 Multicast Features | 110 |
| Distance Vector Multicast Routing Protocol | 110 |
| Internet Group Management Protocol | 110 |
| IGMP Proxy | 110 |
| Protocol Independent Multicast—Dense Mode | 110 |

| | |
|---|------------|
| Protocol Independent Multicast—Sparse Mode | 111 |
| Protocol Independent Multicast—Source Specific Multicast. | 111 |
| Protocol Independent Multicast IPv6 Support | 111 |
| MLD/MLDv2 (RFC2710/RFC3810) | 111 |
| | |
| 3 Hardware Overview | 113 |
| Dell EMC Networking N1100-ON Series Switch | |
| Hardware. | 113 |
| Front Panel | 113 |
| Power Supply | 116 |
| Ventilation System | 116 |
| Thermal Shutdown | 116 |
| LED Definitions | 116 |
| Power Consumption for PoE Switches | 120 |
| Wall Installation. | 121 |
| Dell EMC Networking N1500 Series Switch | |
| Hardware. | 123 |
| Front Panel | 123 |
| Back Panel | 126 |
| LED Definitions | 128 |
| Power Consumption for PoE Switches | 130 |
| Dell EMC Networking N2000 Series Switch | |
| Hardware. | 132 |
| Front Panel | 132 |
| Back Panel | 135 |
| LED Definitions | 137 |
| Power Consumption for PoE Switches | 140 |
| Dell EMC Networking N2100-ON Series Switch | |
| Hardware. | 142 |

| | |
|---|------------|
| Front Panel | 142 |
| Back Panel | 144 |
| LED Definitions | 144 |
| Power Consumption for PoE Switches | 148 |
| Dell EMC Networking N3000/N3000E-ON Series | |
| Switch Hardware | 151 |
| Front Panel | 151 |
| Back Panel | 157 |
| LED Definitions | 159 |
| Power Consumption for PoE Switches | 164 |
| Dell EMC Networking N3100-ON Series Switch | |
| Hardware | 166 |
| Front Panel | 166 |
| Back Panel | 168 |
| LED Definitions | 168 |
| Power Consumption for PoE Switches | 173 |
| PoE Power Budget Limit | 175 |
| Dell EMC Networking N4000 Series Switch | |
| Hardware | 176 |
| Front Panel | 176 |
| Back Panel | 180 |
| LED Definitions | 182 |
| Switch MAC Addresses | 186 |
| | |
| 4 Using Dell EMC OpenManage Switch | |
| Administrator | 189 |
| About Dell EMC OpenManage Switch Administrator | 189 |
| Starting the Application | 190 |
| Understanding the Interface | 191 |

| | | |
|----------|---|------------|
| | Using the Switch Administrator Buttons and Links . . . | 193 |
| | Defining Fields | 194 |
| | Understanding the Device View | 194 |
| | Using the Device View Port Features. | 194 |
| | Using the Device View Switch Locator Feature | 195 |
| 5 | Using the Command-Line Interface | 197 |
| | Accessing the Switch Through the CLI | 197 |
| | Console Connection. | 197 |
| | Telnet Connection. | 198 |
| | Understanding Command Modes | 199 |
| | Entering CLI Commands | 201 |
| | Using the Question Mark to Get Help. | 201 |
| | Using Command Completion | 202 |
| | Entering Abbreviated Commands. | 202 |
| | Negating Commands | 202 |
| | Command Output Paging | 202 |
| | Understanding Error Messages | 203 |
| | Recalling Commands from the History Buffer | 203 |
| 6 | Default Settings | 205 |
| 7 | Setting the IP Address and Other Basic Network Information | 209 |
| | IP Address and Network Information Overview | 209 |
| | What Is the Basic Network Information?. | 209 |

| | |
|--|------------|
| Why Is Basic Network Information Needed? | 210 |
| How Is Basic Network Information Configured?. | 211 |
| What Is Out-of-Band Management and In-Band Management? | 211 |
| Default Network Information | 214 |
| Configuring Basic Network Information (Web) | 215 |
| Out-of-Band Interface | 215 |
| IP Interface Configuration (Default VLAN IP Address) | 216 |
| Route Entry Configuration (Switch Default Gateway) | 218 |
| Domain Name Server. | 220 |
| Default Domain Name | 221 |
| Host Name Mapping | 222 |
| Dynamic Host Name Mapping | 223 |
| Configuring Basic Network Information (CLI) | 224 |
| Enabling the DHCP Client on the OOB Port. | 224 |
| Enabling the DHCP Client on the Default VLAN | 224 |
| Managing DHCP Leases | 225 |
| Configuring Static Network Information on the OOB Port | 226 |
| Configuring Static Network Information on the Default VLAN | 227 |
| Configuring and Viewing Additional Network Information | 228 |
| Basic Network Information Configuration Examples. | 230 |
| Configuring Network Information Using the OOB Port | 230 |
| Configuring Network Information Using the Serial Interface | 232 |

| | | |
|---|---|------------|
| 8 | Managing QSFP Ports | 235 |
| 9 | Stacking | 237 |
| | Stacking Overview | 237 |
| | Dell EMC Networking N1124-ON/N1148-ON, N1500, N2000, N2100-ON, N3000, N3000E-ON, N3100-ON, and N4000 Stacking Compatibility | 243 |
| | How is the Stack Master Selected? | 243 |
| | Adding a Switch to the Stack | 245 |
| | Removing a Switch from the Stack | 246 |
| | How is the Firmware Updated on the Stack? | 246 |
| | What is Stacking Standby? | 247 |
| | What is Nonstop Forwarding? | 247 |
| | Switch Stack MAC Addressing and Stack Design Considerations | 250 |
| | NSF Network Design Considerations | 251 |
| | Why is Stacking Needed? | 251 |
| | Default Stacking Values | 252 |
| | Managing and Monitoring the Stack (Web) | 253 |
| | Unit Configuration | 253 |
| | Stack Summary | 254 |
| | Stack Firmware Synchronization | 255 |
| | Supported Switches | 256 |
| | Stack Port Summary | 257 |
| | Stack Port Counters | 258 |
| | Stack Port Diagnostics | 258 |
| | NSF Summary | 259 |
| | Checkpoint Statistics | 260 |
| | Managing the Stack (CLI) | 261 |
| | Configuring Stack Member, Stack Port, SFS | |

| | |
|---|------------|
| and NSF Settings | 261 |
| Viewing and Clearing Stacking and NSF Information | 263 |
| Connecting to the Management Console from a Stack Member | 264 |
| Stacking and NSF Usage Scenarios. | 264 |
| Basic Failover. | 264 |
| Preconfiguring a Stack Member | 266 |
| NSF in the Data Center | 268 |
| NSF and VoIP | 269 |
| NSF and DHCP Snooping | 270 |
| NSF and the Storage Access Network. | 271 |
| NSF and Routed Access | 273 |
| | |
| 10 Authentication, Authorization, and Accounting | 275 |
| AAA Introduction | 275 |
| Methods | 276 |
| Method Lists | 277 |
| Access Lines | 278 |
| Enabling SSH Access. | 279 |
| Access Lines (AAA) | 279 |
| Access Lines (Non-AAA) | 280 |
| Authentication | 281 |
| Authentication Types. | 281 |
| Authentication Manager | 282 |
| Using RADIUS. | 288 |
| Using TACACS+ Servers to Control Management Access | 293 |
| Dynamic ACL Overview. | 295 |
| Authentication Examples | 302 |

| | |
|---|------------|
| Public Key SSH Authentication Example | 310 |
| Associating a User With an SSH Key. | 318 |
| Authorization | 320 |
| Exec Authorization Capabilities. | 320 |
| Authorization Examples. | 322 |
| RADIUS Change of Authorization. | 324 |
| TACACS Authorization | 328 |
| Accounting | 332 |
| RADIUS Accounting | 332 |
| IEEE 802.1X | 334 |
| What is IEEE 802.1X? | 334 |
| What are the 802.1X Port Authentication Modes? | 335 |
| What is MAC-Based 802.1X Authentication?. | 336 |
| What is the Role of 802.1X in VLAN Assignment? | 339 |
| What is Monitor Mode?. | 341 |
| How Does the Authentication Server Assign DiffServ Policy or ACLs? | 343 |
| What is the Internal Authentication Server? | 344 |
| Default 802.1X Values. | 344 |
| Configuring IEEE 802.1X (Web) | 345 |
| Captive Portal | 370 |
| Captive Portal Overview | 370 |
| Default Captive Portal Behavior and Settings | 378 |
| Configuring Captive Portal (Web). | 380 |
| Configuring Captive Portal (CLI) | 396 |
| Captive Portal Configuration Example | 402 |
| In Case Of Problems in Captive Portal Deployment | 406 |

| | |
|--|------------|
| 11 Monitoring and Logging System Information | 407 |
| System Monitoring Overview | 407 |
| What System Information Is Monitored?. | 407 |
| Why Is System Information Needed? | 408 |
| Where Are Log Messages Sent?. | 408 |
| What Are the Severity Levels? | 409 |
| What Are the System Startup and Operation Logs? | 409 |
| What Is the Log Message Format?. | 410 |
| What Factors Should Be Considered When Configuring Logging?. | 412 |
| Default Log Settings | 413 |
| Monitoring System Information and Configuring Logging (Web) | 414 |
| Device Information | 414 |
| System Health | 416 |
| System Resources | 417 |
| Unit Power Usage History | 418 |
| Integrated Cable Test for Copper Cables. | 419 |
| Optical Transceiver Diagnostics | 420 |
| Log Global Settings. | 421 |
| RAM Log | 422 |
| Log File | 424 |
| SYSLOG Server | 424 |
| Email Alert Global Configuration | 427 |
| Email Alert Mail Server Configuration | 427 |
| Email Alert Subject Configuration | 429 |
| Email Alert To Address Configuration | 430 |
| Email Alert Statistics | 430 |

| | |
|--|------------|
| Monitoring System Information and Configuring Logging (CLI) | 432 |
| Viewing System Information and Enabling the Locator LED | 432 |
| Running Cable Diagnostics | 432 |
| Configuring Local Logging | 434 |
| Configuring Remote Logging | 435 |
| Configuring Mail Server Settings | 436 |
| Configuring Email Alerts for Log Messages | 437 |
| Logging Configuration Examples | 439 |
| Configuring Local and Remote Logging | 439 |
| Configuring Email Alerting | 442 |

12 Managing General System Settings . . . 445

| | |
|---|------------|
| System Settings Overview | 445 |
| Why Does System Information Need to Be Configured? | 447 |
| What Are SDM Templates? | 447 |
| Why is the System Time Needed? | 450 |
| How Does SNTP Work? | 450 |
| What Configuration Is Required for Plug-In Modules? | 451 |
| Default General System Information | 451 |
| Configuring General System Settings (Web) | 452 |
| System Information | 452 |
| CLI Banner | 455 |
| SDM Template Preference | 456 |
| Clock | 457 |
| SNTP Global Settings | 458 |
| SNTP Authentication | 459 |
| SNTP Server | 461 |

| | |
|--|------------|
| Summer Time Configuration | 464 |
| Time Zone Configuration | 465 |
| Card Configuration | 466 |
| Slot Summary. | 467 |
| Supported Cards | 468 |
| Power Over Ethernet Global Configuration. | 469 |
| Power Over Ethernet Unit Configuration | 470 |
| Power Over Ethernet Interface Configuration | 471 |
| Configuring System Settings (CLI). | 473 |
| Configuring System Information | 473 |
| Configuring the Banner. | 474 |
| Managing the SDM Template | 475 |
| Configuring SNTP Authentication and an SNTP Server. | 475 |
| Setting the System Time and Date Manually. | 477 |
| Configuring the Expansion Slots (Dell EMC Networking N3000/N3100-ON/N4000 Series Only). | 478 |
| Viewing Slot Information (Dell EMC Networking N4000 Series Only) | 479 |
| Configuring PoE Settings (Dell EMC Networking N1108P-ON/ N1124P-ON/ N1148P-ON, N1524P/N1548P, N2024P/N2048P/N2128PX-ON, N3024P/N3048P/N3048EP-ON/N3132PX-ON Only). | 479 |
| General System Settings Configuration Examples | 481 |
| Configuring System and Banner Information | 481 |
| Configuring SNTP. | 484 |
| Configuring the Time Manually. | 486 |
| | |
| 13 SNMP. | 487 |
| SNMP Overview | 487 |

| | |
|---|------------|
| What Is SNMP? | 487 |
| What Are SNMP Traps? | 488 |
| Why Is SNMP Needed? | 489 |
| Default SNMP Values | 489 |
| Configuring SNMP (Web) | 491 |
| SNMP Global Parameters | 491 |
| SNMP View Settings | 492 |
| Access Control Group | 494 |
| SNMPv3 User Security Model (USM) | 496 |
| Communities | 499 |
| Notification Filter | 501 |
| Notification Recipients | 502 |
| Trap Flags | 504 |
| OSPFv2 Trap Flags | 505 |
| OSPFv3 Trap Flags | 506 |
| Trap Log | 507 |
| Configuring SNMP (CLI) | 509 |
| Configuring the SNMPv3 Engine ID. | 509 |
| Configuring SNMP Views, Groups, and Users | 510 |
| Configuring Communities | 513 |
| Configuring SNMP Notifications (Traps and Informs) | 515 |
| SNMP Configuration Examples | 518 |
| Configuring SNMPv1 and SNMPv2. | 518 |
| Configuring SNMP Management Station Access | 519 |
| Configuring SNMPv3 | 520 |

14 Images and File Management 525

| | |
|---|------------|
| Image and File Management Overview | 525 |
|---|------------|

| | |
|---|------------|
| What Files Can Be Managed? | 525 |
| Why Is File Management Needed?. | 527 |
| What Methods Are Supported for File Management? | 530 |
| What Factors Should Be Considered When Managing Files? | 531 |
| How Is the Running Configuration Saved?. | 534 |
| Managing Images and Files (Web) | 535 |
| File System | 535 |
| Active Images. | 536 |
| USB Flash Drive. | 537 |
| File Download. | 538 |
| File Upload | 540 |
| Copy Files | 542 |
| Managing Images and Files (CLI) | 543 |
| Downloading and Activating a New Image (TFTP) | 544 |
| Managing Files in Internal Flash | 545 |
| Managing Files on a USB Flash Device | 546 |
| Uploading a Configuration File (SCP). | 547 |
| Managing Configuration Scripts (SFTP) | 548 |
| File and Image Management Configuration | |
| Examples | 549 |
| Upgrading the Firmware | 549 |
| Managing Configuration Scripts | 552 |
| Managing Files by Using the USB Flash Drive | 554 |
| | |
| 15 DHCP and USB Auto-Configuration | 557 |
| Auto Configuration Overview | 557 |
| What Is USB Auto Configuration? | 558 |

| | |
|--|------------|
| What Files Does USB Auto Configuration Use? . . . | 558 |
| How Does USB Auto Configuration Use the | |
| Files on the USB Device? | 559 |
| What Is the Setup File Format? | 561 |
| What Is the DHCP Auto Configuration | |
| Process? | 561 |
| Monitoring and Completing the DHCP Auto | |
| Configuration Process | 567 |
| What Are the Dependencies for DHCP Auto | |
| Configuration?. | 568 |
| Default Auto Configuration Values | 570 |
| Managing Auto Configuration (Web) | 571 |
| Auto-Install Configuration | 571 |
| Managing Auto Configuration (CLI) | 572 |
| Managing Auto Configuration | 572 |
| Auto Configuration Example | 573 |
| Enabling USB Auto Configuration and Auto | |
| Image Download | 573 |
| Enabling DHCP Auto Configuration and Auto | |
| Image Download | 574 |
| Easy Firmware Upgrade via USB | 576 |
| | |
| 16 Monitoring Switch Traffic | 577 |
| Traffic Monitoring Overview. | 577 |
| What is sFlow Technology?. | 577 |
| What is RMON?. | 580 |
| What is Port Mirroring?. | 581 |
| Port Mirroring Behaviors | 583 |
| RSPAN | 585 |

| | |
|---|------------|
| Remote Capture. | 586 |
| Why is Traffic Monitoring Needed? | 587 |
| Default Traffic Monitoring Values. | 587 |
| Monitoring Switch Traffic (Web) | 588 |
| sFlow Agent Summary | 588 |
| sFlow Receiver Configuration | 589 |
| sFlow Sampler Configuration. | 590 |
| sFlow Poll Configuration | 591 |
| Interface Statistics | 592 |
| Etherlike Statistics | 593 |
| GVRP Statistics | 594 |
| EAP Statistics. | 595 |
| Utilization Summary | 596 |
| Counter Summary. | 597 |
| Switchport Statistics | 598 |
| RMON Statistics | 599 |
| RMON History Control Statistics | 599 |
| RMON History Table | 602 |
| RMON Event Control | 603 |
| RMON Event Log | 605 |
| RMON Alarms. | 606 |
| Port Statistics. | 608 |
| LAG Statistics. | 609 |
| Port Mirroring. | 610 |
| Monitoring Switch Traffic (CLI) | 612 |
| Configuring sFlow. | 612 |
| Configuring RMON | 614 |
| Viewing Statistics. | 616 |
| Configuring Port Mirroring | 617 |
| Configuring RSPAN. | 618 |
| Traffic Monitoring Examples | 622 |
| Showing Interface Traffic | 622 |

| | |
|---|------------|
| Configuring sFlow | 623 |
| Configuring RMON | 625 |
| Configuring Remote Capture | 626 |
| Configuring RSPAN | 631 |
| | |
| 17 iSCSI Optimization | 635 |
| iSCSI Optimization Overview | 635 |
| What Does iSCSI Optimization Do? | 635 |
| What Occurs When iSCSI Optimization Is Enabled or Disabled? | 636 |
| How Does the Switch Detect iSCSI Traffic Flows? | 636 |
| How Is Quality of Service Applied to iSCSI Traffic Flows? | 636 |
| How Does iSCSI Optimization Use ACLs? | 637 |
| What Information Does the Switch Track in iSCSI Traffic Flows? | 637 |
| How Does iSCSI Optimization Interact With Dell EqualLogic and Compellent Arrays? | 639 |
| How Does iSCSI Optimization Interact with Other SAN Arrays? | 639 |
| How Does iSCSI Optimization Interact with DCBx? | 640 |
| iSCSI CoS and Priority Flow Control/Enhanced Transmission Selection Interactions | 641 |
| Default iSCSI Optimization Values | 642 |
| Configuring iSCSI Optimization (Web) | 643 |
| iSCSI Global Configuration | 643 |
| Configuring iSCSI Optimization (CLI) | 644 |
| iSCSI Optimization Configuration Examples | 645 |

| | |
|--|------------|
| Configuring iSCSI Optimization Between Servers and a Disk Array | 645 |
| 18 Port Characteristics | 649 |
| Port Overview | 649 |
| What Physical Port Characteristics Can Be Configured? | 649 |
| Auto-Negotiation | 651 |
| Maximum Transmission Unit | 651 |
| What is Link Dependency? | 652 |
| What Interface Types are Supported? | 654 |
| What is Interface Configuration Mode? | 654 |
| What Are the Green Ethernet Features? | 656 |
| Switchport Modes | 657 |
| Default Port Values | 658 |
| Configuring Port Characteristics (Web). | 660 |
| Port Configuration | 660 |
| Link Dependency Configuration | 663 |
| Link Dependency Summary | 665 |
| Port Green Ethernet Configuration | 666 |
| Port Green Ethernet Statistics | 667 |
| Port Green Ethernet LPI History | 669 |
| Configuring Port Characteristics (CLI) | 670 |
| Configuring Port Settings | 670 |
| Configuring Link Dependencies | 672 |
| Configuring Green Features | 673 |
| Port Configuration Examples | 674 |
| Configuring Port Settings | 674 |
| Configuring a Link Dependency Groups | 675 |
| Configuring a Port in Access Mode | 675 |

| | |
|--|------------|
| Configuring a Port in Trunk Mode | 676 |
| Configuring a Port in General Mode | 679 |
| 19 Port and System Security | 681 |
| Port Security | 681 |
| Denial of Service | 688 |
| 20 Access Control Lists | 689 |
| ACL Overview | 689 |
| ACL Counters | 691 |
| What Are MAC ACLs? | 691 |
| What Are IP ACLs? | 692 |
| ACL Actions | 692 |
| What Is the ACL Redirect Function? | 693 |
| What Is the ACL Mirror Function? | 694 |
| What Is ACL Logging | 694 |
| What Are Time-Based ACLs? | 694 |
| ACL Limitations | 695 |
| ACL Configuration Details | 701 |
| How Are ACLs Configured? | 701 |
| Editing Access Lists. | 701 |
| Preventing False ACL Matches. | 701 |
| Using IP and MAC Address Masks. | 703 |
| Policy-Based Routing | 704 |
| Packet Classification | 704 |
| Route-Map Processing | 705 |
| Route-Map Actions | 706 |
| ACLs and Policy Interaction | 708 |
| Limitations. | 709 |

| | |
|---|------------|
| Configuring ACLs (Web) | 712 |
| IP ACL Configuration | 712 |
| IP ACL Rule Configuration | 715 |
| MAC ACL Configuration | 717 |
| MAC ACL Rule Configuration | 719 |
| IPv6 ACL Configuration | 720 |
| IPv6 ACL Rule Configuration | 721 |
| ACL Binding Configuration | 723 |
| Time Range Configuration | 724 |
| Configuring ACLs (CLI) | 726 |
| Configuring an IPv4 ACL | 726 |
| Configuring a MAC ACL | 732 |
| Configuring an IPv6 ACL | 736 |
| Configuring a Time Range | 739 |
| ACL Configuration Examples | 741 |
| Basic Rules | 741 |
| Internal System ACLs | 742 |
| Complete ACL Example | 743 |
| Advanced Examples | 747 |
| Policy-Based Routing Examples | 759 |
| | |
| 21 VLANs | 763 |
| VLAN Overview | 763 |
| VLAN Tagging | 766 |
| GVRP | 767 |
| Double-VLAN Tagging | 768 |
| Voice VLAN | 769 |
| Private VLANs | 775 |
| Additional VLAN Features | 780 |
| | |
| Default VLAN Behavior | 781 |

| | |
|---|------------|
| Configuring VLANs (Web) | 783 |
| VLAN Membership | 783 |
| VLAN Port Settings | 788 |
| VLAN LAG Settings | 789 |
| Bind MAC to VLAN | 791 |
| Bind IP Subnet to VLAN. | 791 |
| GVRP Parameters. | 793 |
| Protocol Group | 795 |
| Adding a Protocol Group | 796 |
| Double VLAN Global Configuration | 798 |
| Double VLAN Interface Configuration | 799 |
| Voice VLAN | 801 |
| | |
| Configuring VLANs (CLI) | 802 |
| Creating a VLAN | 802 |
| Configuring VLAN Settings for a LAG. | 803 |
| Configuring Double VLAN Tagging | 804 |
| Configuring MAC-Based VLANs | 807 |
| Configuring IP-Based VLANs | 809 |
| Configuring a Protocol-Based VLAN | 811 |
| Configuring GVRP. | 814 |
| Configuring Voice VLANs | 816 |
| Configuring a Voice VLAN (Extended Example) | 818 |
| Enterprise Voice VLAN Configuration With QoS | 819 |
| MLAG with RPVST and Voice VLAN | 822 |
| Assigning an 802.1p Priority to VLAN Traffic | 829 |
| Configuring a Private VLAN. | 830 |
| Configuring Inter-Switch Private VLANs | 832 |
| | |
| VLAN Configuration Examples. | 833 |
| Configuring VLANs Using the Dell EMC OpenManage Switch Administrator | 833 |
| Configuring VLANs Using the CLI. | 841 |

22 Spanning Tree Protocol 845

STP Overview 845

| | |
|--|-----|
| What Are Classic STP, Multiple STP, and Rapid STP? | 845 |
| How Does STP Work? | 846 |
| How Does MSTP Operate in the Network? | 847 |
| MSTP with Multiple Forwarding Paths. | 851 |
| MSTP and VLAN IDs | 852 |
| What are the Optional STP Features? | 852 |

RSTP-PV 854

| | |
|--|-----|
| DirectLink Rapid Convergence | 856 |
| IndirectLink Rapid Convergence Feature | 858 |
| Interoperability Between STP-PV and RSTP-PV Modes. | 860 |
| Interoperability With IEEE Spanning Tree Protocols | 860 |
| Configuration Examples | 865 |

Default STP Values 866

Configuring Spanning Tree (Web). 867

| | |
|--|-----|
| STP Global Settings. | 867 |
| STP Port Settings. | 869 |
| STP LAG Settings. | 871 |
| Rapid Spanning Tree | 872 |
| MSTP Settings | 874 |
| MSTP Interface Settings | 876 |
| PVST/RPVST Global Configuration. | 877 |
| PVST/RPVST VLAN Configuration | 878 |
| PVST/RPVST Interface Configuration | 880 |
| PVST/RPVST Statistics | 881 |

Configuring Spanning Tree (CLI) 882

| | |
|--|-----|
| Configuring Global STP Bridge Settings | 882 |
|--|-----|

| | |
|--|------------|
| Configuring Optional STP Features | 883 |
| Configuring STP Interface Settings. | 884 |
| Configuring MSTP Switch Settings. | 885 |
| Configuring MSTP Interface Settings | 886 |
| STP Configuration Examples. | 887 |
| STP Configuration Example. | 887 |
| MSTP Configuration Example. | 889 |
| RSTP-PV Access Switch Configuration Example | 892 |
| 23 Discovering Network Devices | 897 |
| Device Discovery Overview | 897 |
| What Is ISDP?. | 897 |
| What is LLDP?. | 897 |
| What is LLDP-MED? | 898 |
| Why are Device Discovery Protocols Needed?. | 898 |
| Default ISDP and LLDP Values. | 899 |
| Configuring ISDP and LLDP (Web). | 901 |
| ISDP Global Configuration | 901 |
| ISDP Neighbor Table | 902 |
| ISDP Interface Configuration | 903 |
| ISDP Statistics | 904 |
| LLDP Configuration | 905 |
| LLDP Statistics | 907 |
| LLDP Connections. | 908 |
| LLDP-MED Global Configuration | 909 |
| LLDP-MED Interface Configuration. | 910 |
| LLDP-MED Local Device Information. | 911 |
| LLDP-MED Remote Device Information | 911 |

| | |
|---|------------|
| Configuring ISDP and LLDP (CLI) | 912 |
| Configuring Global ISDP Settings | 912 |
| Enabling ISDP on a Port | 913 |
| Viewing and Clearing ISDP Information | 913 |
| Configuring Global LLDP Settings | 914 |
| Configuring Port-based LLDP Settings | 914 |
| Viewing and Clearing LLDP Information | 915 |
| Configuring LLDP-MED Settings | 917 |
| Viewing LLDP-MED Information | 918 |
| Device Discovery Configuration Examples | 918 |
| Configuring ISDP | 918 |
| Configuring LLDP | 919 |
| | |
| 24 Port-Based Traffic Control | 921 |
| | |
| Port-Based Traffic Control Overview | 921 |
| What is Flow Control? | 922 |
| What is Storm Control? | 922 |
| What are Protected Ports? | 923 |
| What is Error Recovery? | 923 |
| What is Link Local Protocol Filtering? | 924 |
| What is Loop Protection? | 925 |
| | |
| Default Port-Based Traffic Control Values | 926 |
| | |
| Configuring Port-Based Traffic Control (Web) | 927 |
| Flow Control (Global Port Parameters) | 927 |
| Storm Control | 928 |
| Protected Port Configuration | 930 |
| LLPF Configuration | 932 |
| | |
| Configuring Port-Based Traffic Control (CLI) | 933 |
| Configuring Flow Control and Storm Control | 933 |
| Configuring Protected Ports | 934 |

| | |
|---|------------|
| Configuring LLPF | 935 |
| Port-Based Traffic Control Configuration Example . . . | 936 |
| 25 Layer-2 Multicast Features | 939 |
| L2 Multicast Overview | 939 |
| Multicast Flooding and Forwarding. | 939 |
| What Are the Multicast Bridging Features? | 940 |
| What Is L2 Multicast Traffic? | 941 |
| What Is IGMP Snooping? | 941 |
| What Is MLD Snooping? | 943 |
| What Is Multicast VLAN Registration? | 945 |
| When Are Layer-3 Multicast Features Required? | 946 |
| What Are GARP and GMRP? | 946 |
| Snooping Switch Restrictions | 948 |
| MAC Address-Based Multicast Group | 948 |
| Topologies Where the Multicast Source Is Not Directly Connected to the Querier | 948 |
| Using Static Multicast MAC Configuration | 948 |
| IGMP Snooping and GMRP | 948 |
| Default L2 Multicast Values | 949 |
| Configuring L2 Multicast Features (Web) | 951 |
| Multicast Global Parameters | 951 |
| Bridge Multicast Group | 952 |
| MFDB Summary. | 955 |
| MRouter Status | 956 |
| General IGMP Snooping | 957 |
| Global Querier Configuration | 960 |
| VLAN Querier | 961 |
| VLAN Querier Status | 963 |

| | |
|--|------------|
| MFDB IGMP Snooping Table | 964 |
| MLD Snooping General. | 965 |
| MLD Snooping Global Querier Configuration | 967 |
| MLD Snooping VLAN Querier | 968 |
| MLD Snooping VLAN Querier Status. | 970 |
| MFDB MLD Snooping Table | 971 |
| MVR Global Configuration | 972 |
| MVR Members | 973 |
| MVR Interface Configuration. | 973 |
| MVR Statistics | 976 |
| GARP Timers | 977 |
| GMRP Parameters | 979 |
| MFDB GMRP Table. | 981 |
| Configuring L2 Multicast Features (CLI). | 982 |
| Configuring Layer-2 Multicasting. | 982 |
| Configuring IGMP Snooping on VLANs | 983 |
| Configuring IGMP Snooping Querier. | 984 |
| Configuring MLD Snooping on VLANs | 985 |
| Configuring MLD Snooping Querier | 986 |
| Configuring MVR | 987 |
| Configuring GARP Timers and GMRP | 989 |
| Case Study on a Real-World Network Topology | 990 |
| Multicast Snooping Case Study | 990 |
| | |
| 26 Connectivity Fault Management | 995 |
| | |
| Dot1ag Overview. | 995 |
| How Does Dot1ag Work Across a Carrier Network? | 996 |
| What Entities Make Up a Maintenance Domain?. | 997 |
| What is the Administrator's Role? | 999 |

| | |
|--|-------------|
| Default Dot1ag Values | 1000 |
| Configuring Dot1ag (Web) | 1001 |
| Dot1ag Global Configuration | 1001 |
| Dot1ag MD Configuration | 1001 |
| Dot1ag MA Configuration | 1002 |
| Dot1ag MEP Configuration | 1003 |
| Dot1ag MIP Configuration | 1004 |
| Dot1ag RMEP Summary | 1005 |
| Dot1ag L2 Ping | 1006 |
| Dot1ag L2 Traceroute | 1006 |
| Dot1ag L2 Traceroute Cache | 1007 |
| Dot1ag Statistics | 1007 |
| Configuring Dot1ag (CLI) | 1009 |
| Configuring Dot1ag Global Settings and Creating Domains | 1009 |
| Configuring MEP Information | 1010 |
| Dot1ag Ping and Traceroute | 1011 |
| Dot1ag Configuration Example | 1012 |
| | |
| 27 Snooping and Inspecting Traffic | 1015 |
| | |
| Traffic Snooping and Inspection Overview | 1015 |
| What Is DHCP Snooping? | 1016 |
| How Is the DHCP Snooping Bindings Database Populated? | 1017 |
| What Is IP Source Guard? | 1020 |
| What is Dynamic ARP Inspection? | 1021 |
| Why Is Traffic Snooping and Inspection Necessary? | 1022 |
| Default Traffic Snooping and Inspection Values | 1022 |

| | |
|---|-------------|
| Configuring Traffic Snooping and Inspection (Web) | 1024 |
| DHCP Snooping Configuration | 1024 |
| DHCP Snooping Interface Configuration | 1025 |
| DHCP Snooping VLAN Configuration | 1027 |
| DHCP Snooping Persistent Configuration | 1028 |
| DHCP Snooping Static Bindings Configuration | 1029 |
| DHCP Snooping Dynamic Bindings Summary | 1030 |
| DHCP Snooping Statistics | 1031 |
| IPSG Interface Configuration | 1032 |
| IPSG Binding Configuration | 1032 |
| IPSG Binding Summary | 1033 |
| DAI Global Configuration | 1034 |
| DAI Interface Configuration | 1035 |
| DAI VLAN Configuration | 1037 |
| DAI ACL Configuration | 1038 |
| DAI ACL Rule Configuration | 1038 |
| DAI Statistics | 1039 |
| Configuring Traffic Snooping and Inspection (CLI) | 1041 |
| Configuring DHCP Snooping | 1041 |
| Configuring IP Source Guard | 1043 |
| Configuring Dynamic ARP Inspection | 1044 |
| Traffic Snooping and Inspection Configuration Examples | 1047 |
| Configuring DHCP Snooping | 1047 |
| Configuring IPSG | 1049 |
| | |
| 28 Link Aggregation | 1051 |
| | |
| Link Aggregation | 1051 |
| Overview | 1051 |
| Default Link Aggregation Values | 1055 |

| | |
|---|-------------|
| Configuring Link Aggregation (Web) | 1056 |
| Configuring Link Aggregation (CLI) | 1062 |
| Link Aggregation Configuration Examples | 1066 |
| Multi-Switch LAG (MLAG) | 1069 |
| Overview | 1069 |
| Deployment Scenarios | 1070 |
| Definitions | 1072 |
| Configuration Consistency | 1073 |
| Operation in the Network | 1076 |
| Layer-2 Configuration Steps | 1080 |
| Switch Firmware Upgrade Procedure | 1083 |
| Static Routing on MLAG Interfaces | 1084 |
| Caveats and Limitations | 1091 |
| Basic Configuration Example | 1097 |
| A Complete MLAG Example | 1105 |
| | |
| 29 Data Center Bridging Features | 1123 |
| | |
| Data Center Bridging Technology Overview | 1123 |
| Default DCB Values | 1124 |
| Priority Flow Control | 1125 |
| PFC Operation and Behavior | 1125 |
| Configuring PFC Using the Web Interface | 1126 |
| Configuring PFC Using the CLI | 1128 |
| PFC Configuration Example | 1130 |
| DCB Capability Exchange | 1132 |
| Interoperability with IEEE DCBx | 1133 |
| DCBx and Port Roles | 1133 |
| Configuration Source Port Selection Process | 1135 |
| Disabling DCBX | 1136 |
| Configuring DCBx | 1137 |

| | |
|--|-------------|
| Enhanced Transmission Selection | 1139 |
| ETS Operation. | 1139 |
| Commands | 1142 |
| ETS Configuration Example. | 1143 |
| ETS Theory of Operation | 1149 |
| | |
| 30 MAC Addressing and Forwarding | 1155 |
| MAC Address Table Overview | 1155 |
| How Is the Address Table Populated? | 1155 |
| What Information Is in the MAC Address Table? | 1156 |
| How Is the MAC Address Table Maintained Across a Stack? | 1156 |
| Default MAC Address Table Values. | 1156 |
| Managing the MAC Address Table (Web) | 1157 |
| Static Address Table | 1157 |
| Global Address Table. | 1159 |
| Managing the MAC Address Table (CLI) | 1160 |
| Managing the MAC Address Table. | 1160 |
| | |
| 31 DHCP Server Settings | 1163 |
| DHCP Overview | 1163 |
| How Does DHCP Work? | 1164 |
| What are DHCP Options? | 1165 |
| What Additional DHCP Features Does the Switch Support? | 1165 |
| Default DHCP Server Values. | 1166 |
| Configuring the DHCP Server (Web). | 1167 |

| | |
|--|-------------|
| DHCP Server Network Properties | 1167 |
| Address Pool | 1169 |
| Address Pool Options. | 1173 |
| DHCP Bindings | 1175 |
| DHCP Server Reset Configuration | 1175 |
| DHCP Server Conflicts Information. | 1176 |
| DHCP Server Statistics | 1177 |
| Configuring the DHCP Server (CLI) | 1178 |
| Configuring Global DHCP Server Settings | 1178 |
| Configuring a Dynamic Address Pool. | 1179 |
| Configuring a Static Address Pool | 1180 |
| Monitoring DHCP Server Information | 1181 |
| DHCP Server Configuration Examples. | 1182 |
| Configuring a Dynamic Address Pool. | 1182 |
| Configuring a Static Address Pool | 1184 |
| | |
| 32 IP Routing. | 1187 |
| IP Routing Overview | 1187 |
| Default IP Routing Values | 1189 |
| IP Path MTU and Path MTU Discovery | 1190 |
| ARP Table | 1191 |
| Configuring IP Routing Features (Web) | 1192 |
| IP Configuration. | 1192 |
| IP Statistics | 1193 |
| ARP Create | 1194 |
| ARP Table Configuration | 1195 |
| Router Discovery Configuration | 1196 |
| Router Discovery Status | 1197 |
| Route Table | 1198 |

| | |
|--|-------------|
| Best Routes Table | 1199 |
| Route Entry Configuration | 1200 |
| Configured Routes | 1202 |
| Route Preferences Configuration | 1203 |
| Configuring IP Routing Features (CLI) | 1204 |
| Configuring Global IP Routing Settings. | 1204 |
| Configuring ARP Settings. | 1205 |
| Configuring Router Discovery (IRDP). | 1206 |
| Configuring Route Table Entries and Route Preferences. | 1207 |
| IP Routing Configuration Example | 1209 |
| Configuring Dell EMC Networking N-Series Switch A | 1210 |
| Configuring Dell EMC Networking N-Series Switch B | 1211 |
| | |
| 33 Routing Interfaces. | 1213 |
| | |
| Routing Interface Overview | 1213 |
| What Are VLAN Routing Interfaces?. | 1213 |
| What Are Loopback Interfaces? | 1214 |
| What Are Tunnel Interfaces?. | 1215 |
| Why Are Routing Interfaces Needed? | 1216 |
| | |
| Default Routing Interface Values | 1218 |
| | |
| Configuring Routing Interfaces (Web). | 1219 |
| IP Interface Configuration | 1219 |
| DHCP Lease Parameters | 1220 |
| VLAN Routing Summary | 1220 |
| Tunnel Configuration | 1221 |
| Tunnels Summary. | 1222 |
| Loopbacks Configuration | 1223 |
| Loopbacks Summary | 1224 |

| | |
|--|-------------|
| Configuring Routing Interfaces (CLI) | 1225 |
| Configuring VLAN Routing Interfaces (IPv4) | 1225 |
| Configuring Loopback Interfaces. | 1227 |
| Configuring Tunnels. | 1228 |
| | |
| 34 Layer-2 and Layer-3 Relay Features | 1229 |
| L2 and L3 Relay Overview | 1229 |
| What Is L2 DHCP Relay? | 1229 |
| What Is L3 DHCP Relay? | 1233 |
| What Is the IP Helper Feature?. | 1234 |
| Default L2/L3 Relay Values. | 1238 |
| Configuring L2 and L3 Relay Features (Web) | 1239 |
| L2 DHCP Relay Global Configuration | 1239 |
| L2 DHCP Relay Interface Configuration | 1240 |
| L2 DHCP Relay Interface Statistics. | 1242 |
| L2 DHCP Relay VLAN Configuration | 1243 |
| DHCP Relay Agent Configuration | 1243 |
| IP Helper (L3 DHCP Relay) Global Configuration | 1245 |
| IP Helper (L3 DHCP Relay) Interface Configuration | 1247 |
| IP Helper Statistics | 1249 |
| Configuring L2 and L3 Relay Features (CLI) | 1250 |
| Configuring L2 DHCP Relay | 1250 |
| Configuring L3 Relay (IP Helper) Settings | 1252 |
| Relay Agent Configuration Example. | 1254 |
| | |
| 35 OSPF and OSPFv3. | 1257 |
| OSPF Overview. | 1258 |

| | |
|---|-------------|
| What Are OSPF Areas and Other OSPF Topology Features? | 1258 |
| What Are OSPF Routers and LSAs? | 1259 |
| How Are Routes Selected? | 1259 |
| How Are OSPF and OSPFv3 Different? | 1259 |
| OSPF Feature Details | 1260 |
| Stub Router | 1260 |
| Static Area Range Cost. | 1262 |
| LSA Pacing | 1263 |
| Flood Blocking | 1264 |
| MTU. | 1265 |
| Default OSPF Values. | 1266 |
| Configuring OSPF Features (Web). | 1268 |
| OSPF Configuration | 1268 |
| OSPF Area Configuration | 1269 |
| OSPF Stub Area Summary | 1272 |
| OSPF Area Range Configuration | 1273 |
| OSPF Interface Statistics | 1274 |
| OSPF Interface Configuration | 1275 |
| OSPF Neighbor Table. | 1276 |
| OSPF Neighbor Configuration | 1277 |
| OSPF Link State Database | 1278 |
| OSPF Virtual Link Configuration | 1278 |
| OSPF Virtual Link Summary. | 1280 |
| OSPF Route Redistribution Configuration | 1281 |
| OSPF Route Redistribution Summary. | 1282 |
| NSF OSPF Configuration | 1283 |
| Configuring OSPFv3 Features (Web) | 1284 |
| OSPFv3 Configuration | 1284 |
| OSPFv3 Area Configuration. | 1284 |
| OSPFv3 Stub Area Summary | 1288 |
| OSPFv3 Area Range Configuration. | 1289 |

| | |
|--|-------------|
| OSPFv3 Interface Configuration | 1290 |
| OSPFv3 Interface Statistics. | 1291 |
| OSPFv3 Neighbors | 1292 |
| OSPFv3 Neighbor Table. | 1293 |
| OSPFv3 Link State Database | 1294 |
| OSPFv3 Virtual Link Configuration | 1295 |
| OSPFv3 Virtual Link Summary | 1297 |
| OSPFv3 Route Redistribution Configuration | 1298 |
| OSPFv3 Route Redistribution Summary | 1299 |
| NSF OSPFv3 Configuration | 1300 |
| Configuring OSPF Features (CLI). | 1301 |
| Configuring Global OSPF Settings | 1301 |
| Configuring OSPF Interface Settings | 1304 |
| Configuring Stub Areas and NSSAs | 1306 |
| Configuring Virtual Links | 1308 |
| Configuring OSPF Area Range Settings | 1310 |
| Configuring NSF Settings for OSPF. | 1312 |
| Configuring OSPFv3 Features (CLI) | 1313 |
| Configuring Global OSPFv3 Settings | 1313 |
| Configuring OSPFv3 Interface Settings. | 1315 |
| Configuring Stub Areas and NSSAs | 1317 |
| Configuring Virtual Links | 1319 |
| Configuring an OSPFv3 Area Range | 1320 |
| Configuring OSPFv3 Route Redistribution Settings | 1321 |
| Configuring NSF Settings for OSPFv3. | 1322 |
| OSPF Configuration Examples | 1323 |
| Configuring an OSPF Border Router and Setting Interface Costs | 1323 |
| Configuring Stub and NSSA Areas for OSPF and OSPFv3 | 1326 |
| Configuring a Virtual Link for OSPF and OSPFv3 | 1329 |

| | | |
|-----------|--|-------------|
| | Interconnecting an IPv4 Backbone and Local IPv6 Network | 1332 |
| | Configuring the Static Area Range Cost | 1335 |
| | Configuring Flood Blocking | 1340 |
| | Configuring OSPF VRFs | 1345 |
| 36 | VRF | 1349 |
| | VRF Resource Sharing | 1350 |
| | VRF ARP Entries | 1350 |
| | VRF Route Entries | 1350 |
| 37 | RIP | 1355 |
| | RIP Overview | 1355 |
| | How Does RIP Determine Route Information? | 1355 |
| | What Is Split Horizon? | 1356 |
| | What RIP Versions Are Supported? | 1356 |
| | Default RIP Values | 1357 |
| | Configuring RIP Features (Web) | 1358 |
| | RIP Configuration | 1358 |
| | RIP Interface Configuration | 1359 |
| | RIP Interface Summary | 1360 |
| | RIP Route Redistribution Configuration | 1361 |
| | RIP Route Redistribution Summary | 1362 |
| | Configuring RIP Features (CLI) | 1363 |
| | Configuring Global RIP Settings | 1363 |
| | Configuring RIP Interface Settings | 1364 |
| | Configuring Route Redistribution Settings | 1365 |
| | RIP Configuration Example | 1367 |

| | | |
|----|---|-------------|
| 38 | VRRP | 1371 |
| | VRRP Overview | 1371 |
| | How Does VRRP Work? | 1371 |
| | What Is the VRRP Router Priority? | 1372 |
| | What Is VRRP Preemption? | 1372 |
| | What Is VRRP Accept Mode? | 1373 |
| | What Are VRRP Route and Interface Tracking? | 1373 |
| | Default VRRP Values | 1375 |
| | Configuring VRRP Features (Web) | 1376 |
| | VRRP Configuration | 1376 |
| | VRRP Virtual Router Status | 1377 |
| | VRRP Virtual Router Statistics | 1378 |
| | VRRP Router Configuration | 1379 |
| | VRRP Route Tracking Configuration | 1380 |
| | VRRP Interface Tracking Configuration | 1382 |
| | Configuring VRRP Features (CLI) | 1384 |
| | Configuring VRRP Settings | 1384 |
| | VRRP Configuration Example | 1386 |
| | VRRP with Load Sharing | 1386 |
| | Troubleshooting VRRP | 1389 |
| | VRRP with Route and Interface Tracking | 1390 |
| | Configuring VRRP in a VRF | 1393 |
| 39 | BGP | 1397 |
| | Overview | 1398 |
| | Autonomous Systems | 1400 |
| | BGP Operations | 1400 |

| | |
|---|-------------|
| Decision Process Overview | 1400 |
| Path Attributes | 1402 |
| BGP Finite State Machine (FSM). | 1404 |
| Detecting Loss of Adjacency. | 1406 |
| Authentication | 1407 |
| Outbound Update Groups. | 1407 |
| Removing Private AS Numbers. | 1408 |
| Templates. | 1408 |
| Resolving Interface Routes. | 1410 |
| Originating BGP Routes. | 1410 |
| Equal Cost Multipath (ECMP). | 1411 |
| BGP Next-Hop Resolution | 1412 |
| Address Aggregation | 1414 |
| Routing Policy. | 1416 |
| Inbound Policy | 1417 |
| Outbound Policy | 1417 |
| Routing Policy Changes | 1418 |
| BGP Timers | 1419 |
| Communities | 1420 |
| Routing Table Overflow. | 1420 |
| Route Reflection | 1421 |
| VRF Support. | 1422 |
| BGP Neighbor Configuration | 1422 |
| Extended Communities | 1422 |
| VPNv4/VRF Route Distribution via MP-BGP . . . | 1425 |
| IPv6 | 1428 |
| BGP Limitations | 1434 |
| BGP Configuration Examples | 1436 |
| Enabling BGP | 1436 |
| BGP Example | 1437 |
| Network Example. | 1438 |
| BGP Redistribution of OSPF Example | 1439 |

| | |
|--|-------------|
| Configuring the Multi-Exit Discriminator in BGP Advertised Routes | 1440 |
| Configuring Communities in BGP | 1441 |
| Configuring a Route Reflector | 1442 |
| Campus Network MP-BGP and OSPF Configuration | 1444 |
| Configuring MP-eBGP and Extended Communities | 1460 |
| | |
| 40 Bidirectional Forwarding Detection | 1467 |
| Overview | 1467 |
| BFD Operational Modes | 1468 |
| Asynchronous Mode | 1468 |
| Demand Mode | 1468 |
| Echo Function | 1469 |
| Limitations | 1469 |
| BFD Example | 1470 |
| | |
| 41 IPv6 Routing | 1473 |
| IPv6 Routing Overview | 1473 |
| How Does IPv6 Compare with IPv4? | 1474 |
| How Are IPv6 Interfaces Configured? | 1474 |
| Default IPv6 Routing Values | 1476 |
| Configuring IPv6 Routing Features (Web) | 1478 |
| Global Configuration | 1478 |
| Interface Configuration | 1479 |
| Interface Summary | 1480 |
| IPv6 Statistics | 1481 |

| | |
|---|-------------|
| IPv6 Neighbor Table | 1482 |
| DHCPv6 Client Parameters | 1483 |
| DHCPv6 Client Statistics | 1484 |
| IPv6 Router Entry Configuration | 1485 |
| IPv6 Route Table | 1486 |
| IPv6 Route Preferences | 1487 |
| Configured IPv6 Routes. | 1488 |
| Configuring IPv6 Routing Features (CLI). | 1489 |
| Configuring Global IP Routing Settings. | 1489 |
| Configuring IPv6 Interface Settings | 1490 |
| Configuring IPv6 Neighbor Discovery | 1491 |
| Configuring IPv6 Route Table Entries and | |
| Route Preferences | 1493 |
| IPv6 Show Commands | 1495 |
| IPv6 Static Reject and Discard Routes | 1496 |
| IPv6 Router Advertisement Guard | 1497 |
| | |
| 42 DHCPv6 Server Settings | 1501 |
| DHCPv6 Overview | 1501 |
| What Is a DHCPv6 Pool? | 1502 |
| What Is a Stateless Server? | 1502 |
| What Is the DHCPv6 Relay Agent Information | |
| Option? | 1502 |
| What Is a Prefix Delegation? | 1502 |
| Default DHCPv6 Server and Relay Values. | 1503 |
| Configuring the DHCPv6 Server and Relay (Web). . . | 1504 |
| DHCPv6 Global Configuration | 1504 |
| DHCPv6 Pool Configuration. | 1505 |
| Prefix Delegation Configuration | 1507 |

| | |
|---|-------------|
| DHCPv6 Pool Summary | 1508 |
| DHCPv6 Interface Configuration | 1509 |
| DHCPv6 Server Bindings Summary. | 1511 |
| DHCPv6 Statistics. | 1512 |
| Configuring the DHCPv6 Server and Relay (CLI). | 1513 |
| Configuring Global DHCP Server and Relay | |
| Agent Settings | 1513 |
| Configuring a DHCPv6 Pool for Stateless | |
| Server Support | 1513 |
| Configuring a DHCPv6 Pool for Specific Hosts | 1514 |
| Configuring DHCPv6 Interface Information. | 1515 |
| Monitoring DHCPv6 Information | 1516 |
| DHCPv6 Configuration Examples | 1517 |
| Configuring a DHCPv6 Stateless Server | 1517 |
| Configuring the DHCPv6 Server for Prefix | |
| Delegation. | 1518 |
| Configuring an Interface as a DHCPv6 Relay | |
| Agent | 1518 |
| | |
| 43 Differentiated Services | 1521 |
| DiffServ Overview | 1521 |
| How Does DiffServ Functionality Vary Based | |
| on the Role of the Switch? | 1522 |
| What Are the Elements of DiffServ | |
| Configuration?. | 1522 |
| Class-Map Processing | 1523 |
| Default DiffServ Values | 1524 |
| Configuring DiffServ (Web) | 1526 |
| DiffServ Configuration | 1526 |
| Class Configuration | 1527 |

| | |
|---|-------------|
| Class Criteria | 1528 |
| Policy Configuration | 1530 |
| Policy Class Definition | 1532 |
| Service Configuration | 1535 |
| Service Detailed Statistics | 1536 |
| Flow-Based Mirroring | 1537 |
| Configuring DiffServ (CLI) | 1538 |
| DiffServ Configuration (Global) | 1538 |
| DiffServ Class Configuration for IPv4. | 1539 |
| DiffServ Class Configuration for IPv6. | 1540 |
| DiffServ Protocol Matching | 1542 |
| DiffServ Policy Creation | 1543 |
| Simple DiffServ Policy Attributes Configuration | 1543 |
| DiffServ Service Configuration | 1546 |
| DiffServ Configuration Examples | 1547 |
| Providing Subnets Equal Access to External Network | 1547 |
| Configuring DiffServ Policy Using ACLs | 1549 |
| DiffServ for VoIP | 1551 |
| WRED | 1554 |
| WRED Processing | 1554 |
| WRED Drop Probabilities | 1554 |
| Exponential Weighting Constant | 1555 |
| WRED Color-Aware Processing | 1555 |
| Simple Meter Implementation | 1556 |
| Single Rate Meter Implementation. | 1556 |
| Two-Rate Meter Implementation. | 1557 |

| | |
|--|-------------|
| 44 Class-of-Service | 1559 |
| CoS Overview | 1559 |
| What Are Trusted and Untrusted Port Modes? | 1560 |
| How Is Traffic Shaping Used on Egress Traffic? | 1560 |
| How Are Traffic Queues Defined? | 1561 |
| Which Queue Management Methods Are Supported? | 1561 |
| CoS Queue Usage. | 1563 |
| Default CoS Values. | 1563 |
| Configuring CoS (Web). | 1565 |
| Mapping Table Configuration. | 1565 |
| Interface Configuration | 1567 |
| Interface Queue Configuration | 1568 |
| Interface Queue Drop Precedence Configuration | 1569 |
| Configuring CoS (CLI) | 1571 |
| Mapping Table Configuration. | 1571 |
| CoS Interface Configuration Commands | 1572 |
| Interface Queue Configuration | 1572 |
| Configuring Interface Queue Drop Probability. | 1574 |
| CoS Configuration Example | 1575 |
| Explicit Congestion Notification. | 1578 |
| Enabling ECN in Microsoft Windows. | 1579 |
| Example 1: SLA Configuration | 1580 |
| Example 2: Long-Lived Congestion | 1584 |
| Example 3: Data Center TCP (DCTCP) Configuration | 1584 |

| | | |
|----|--|------|
| 45 | Auto VoIP | 1587 |
| | Auto VoIP Overview | 1587 |
| | How Does Auto VoIP Use ACLs? | 1588 |
| | Default Auto VoIP Values | 1588 |
| | Configuring Auto VoIP (Web) | 1589 |
| | Auto VoIP Global Configuration | 1589 |
| | Auto VoIP Interface Configuration | 1589 |
| | Configuring Auto VoIP (CLI) | 1591 |
| 46 | IPv4 and IPv6 Multicast | 1593 |
| | L3 Multicast Overview | 1593 |
| | What Is IP Multicast Traffic? | 1594 |
| | Multicast Addressing | 1594 |
| | What Multicast Protocols Does the Switch | |
| | Support? | 1595 |
| | What Are the Multicast Protocol Roles? | 1596 |
| | When Is L3 Multicast Required on the | |
| | Switch? | 1596 |
| | What Is the Multicast Routing Table? | 1597 |
| | What Is IGMP? | 1598 |
| | What Is MLD? | 1599 |
| | What Is PIM? | 1599 |
| | What Is DVMRP? | 1610 |
| | Default L3 Multicast Values | 1613 |
| | Configuring General IPv4 Multicast Features | |
| | (Web) | 1615 |
| | Multicast Global Configuration | 1615 |
| | Multicast Interface Configuration | 1616 |

| | |
|---|-------------|
| Multicast Route Table | 1617 |
| Multicast Admin Boundary Configuration | 1618 |
| Multicast Admin Boundary Summary | 1619 |
| Multicast Static MRoute Configuration. | 1619 |
| Multicast Static MRoute Summary. | 1620 |
| Configuring IPv6 Multicast Features (Web). | 1621 |
| IPv6 Multicast Route Table | 1621 |
| Configuring IGMP and IGMP Proxy (Web) | 1622 |
| IGMP Global Configuration | 1622 |
| IGMP Interface Configuration | 1623 |
| IGMP Interface Summary. | 1624 |
| IGMP Cache Information | 1624 |
| IGMP Interface Source List Information | 1625 |
| IGMP Proxy Interface Configuration | 1626 |
| IGMP Proxy Configuration Summary. | 1627 |
| IGMP Proxy Interface Membership Info | 1628 |
| Detailed IGMP Proxy Interface Membership Information | 1629 |
| Configuring MLD and MLD Proxy (Web). | 1630 |
| MLD Global Configuration | 1630 |
| MLD Routing Interface Configuration | 1631 |
| MLD Routing Interface Summary. | 1632 |
| MLD Routing Interface Cache Information. | 1632 |
| MLD Routing Interface Source List Information | 1633 |
| MLD Traffic | 1634 |
| MLD Proxy Configuration | 1635 |
| MLD Proxy Configuration Summary | 1636 |
| MLD Proxy Interface Membership Information | 1637 |
| Detailed MLD Proxy Interface Membership Information | 1638 |

| | |
|--|-------------|
| Configuring PIM for IPv4 and IPv6 (Web) | 1639 |
| PIM Global Configuration | 1639 |
| PIM Global Status | 1641 |
| PIM Interface Configuration | 1642 |
| PIM Interface Summary | 1643 |
| Candidate RP Configuration | 1644 |
| Static RP Configuration | 1646 |
| SSM Range Configuration | 1648 |
| BSR Candidate Configuration | 1650 |
| BSR Candidate Summary | 1651 |
| Configuring DVMRP (Web) | 1652 |
| DVMRP Global Configuration | 1652 |
| DVMRP Interface Configuration | 1653 |
| DVMRP Configuration Summary | 1654 |
| DVMRP Next Hop Summary | 1655 |
| DVMRP Prune Summary | 1656 |
| DVMRP Route Summary | 1656 |
| Configuring L3 Multicast Features (CLI) | 1657 |
| Configuring and Viewing IPv4 Multicast Information | 1657 |
| Configuring and Viewing IPv6 Multicast Route Information | 1659 |
| Configuring and Viewing IGMP | 1660 |
| Configuring and Viewing IGMP Proxy | 1662 |
| Configuring and Viewing MLD | 1663 |
| Configuring and Viewing MLD Proxy | 1664 |
| Configuring and Viewing PIM-DM for IPv4 Multicast Routing | 1665 |
| Configuring and Viewing PIM-DM for IPv6 Multicast Routing | 1666 |
| Configuring and Viewing PIM-SM for IPv4 Multicast Routing | 1667 |
| Configuring and Viewing PIM-SM for IPv6 | |

| | |
|--|-------------|
| Multicast Routing | 1669 |
| Configuring and Viewing DVMRP Information | 1672 |
| L3 Multicast Configuration Examples | 1673 |
| Configuring Multicast VLAN Routing With IGMP and PIM-SM | 1673 |
| Configuring DVMRP | 1677 |
| 47 Audio Video Bridging | 1679 |
| Overview | 1679 |
| MSRP | 1682 |
| MVRP | 1683 |
| MMRP | 1684 |
| IEEE 802.1AS | 1685 |
| Best Master Selection | 1687 |
| Time Synchronization | 1688 |
| Link Delay Measurement | 1689 |
| Caveats and Limitations | 1691 |
| AVB Configuration Example | 1692 |
| 48 OpenFlow | 1695 |
| Dell EMC Networking OpenFlow Hybrid Overview | 1695 |
| Enable Dell EMC Networking OpenFlow Hybrid | 1696 |
| Interaction with OpenFlow Controllers | 1698 |
| Deploy OpenFlow Controller Flows | 1729 |
| Collect Port and Queue Status and Statistics | 1734 |

| | |
|---|-------------|
| Usage Scenarios | 1734 |
| Eligible Interfaces | 1734 |
| OpenFlow Hybrid | 1735 |
| Example Configuration | 1735 |
| Interaction with Other Switch Functions | 1736 |
| OpenSSL | 1736 |
| IP Stack | 1736 |
| VLANs | 1736 |
| LAGs | 1737 |
| Ports | 1737 |
| Network Interface ARP Table | 1737 |
| Routing Interface ARP Table | 1737 |
| QoS | 1737 |
| IP Routing, IP Multicast, and Layer-2 Multicast | 1738 |
| LLDP and Voice VLAN | 1738 |
| Limitations, Restrictions, and Assumptions. | 1739 |
| List of OpenFlow—Dell EMC Networking Component Interferences | 1739 |
| OpenFlow Configuration Example. | 1740 |
| | |
| 49 Dell EMC Networking Python Support. | 1741 |
| | |
| A Appendix | 1749 |
| Feature Limits and Platform Constants | 1749 |
| System Process Definitions | 1762 |
| SupportAssist | 1769 |

Index 1773

Introduction

The switches in the N-Series are stackable layer-2 and layer-3 switches. These switches include the following features:

- 1U form factor, rack-mountable chassis design.
- Support for all data-communication requirements for a multi-layer switch, including layer-2 switching, IPv4 routing, IPv6 routing, IP multicast, quality of service, security, and system management features.
- High availability with automatic failover and checkpointing of dynamic state.

The Dell EMC Networking N-Series includes the following switch models: N1108T-ON, N1108P-ON, N1124T-ON, N1124P-ON, N1148T-ON, N1148P-ON, N1524, N1524P, N1548, N1548P, N2024, N2024P, N2048, N2048P, N2128PX-ON, N3024, N3024F, N3024P, N3048, N3048P, N3024EP-ON, N3024ET-ON, N3024EF-ON, N3048ET-ON, N3048EP-ON, N3132PX-ON, N4032, N4032F, N4064, N4064F



NOTE: Switch administrators are strongly advised to maintain Dell EMC Networking N-Series switches on the latest version of the Dell EMC Networking Operating System. Dell EMC Networking continually improves the features and functions based on feedback from you, the customer. For critical infrastructure, prestaging of a new release into a non-critical portion of the network is recommended to verify network configuration and operation with any new version of Dell EMC Networking N-Series switch firmware.

About This Document

This guide discusses and provides examples on how to configure, monitor, and maintain Dell EMC Networking N-Series switches by using web-based Dell EMC OpenManage Switch Administrator utility or the command-line interface (CLI).

Examples given in this guide may not include complete CLI syntax as the preference is to present CLI syntax relevant to the configuration task. Refer to the Dell EMC Networking N1100-ON, N1500, N2000, N2100-ON, N3000,

N3100-ON, and N4000 Series Switches CLI Reference Guide for definitive syntax for any particular command. The parameter ranges listed in the examples or text may vary from the allowed range on any particular switch due to product limitations. Refer to the Feature Limits and Platform Constants section located in the Appendix of this document for range limits relevant to a particular switch model.

Audience

This guide is for network administrators in charge of managing one or more Dell EMC Networking N-Series switches. To obtain the greatest benefit from this guide, you should have a basic understanding of Ethernet networks and local area network (LAN) concepts.

Document Conventions

Table 1-1 describes the typographical conventions this document uses.

Table 1-1. Document Conventions

| Convention | Description |
|---------------------------|--|
| Bold | Page names, field names, menu options, button names, and CLI commands and keywords. |
| <code>courier font</code> | Command-line text (CLI output) and file names |
| [] | In a command line, square brackets indicate an optional entry. |
| { } | In a command line, inclusive brackets indicate a selection of compulsory parameters separated by the character. One option must be selected. For example: spanning-tree mode {stp rstp mstp} means that for the spanning-tree mode command, stp , rstp , or mstp must be entered. |
| <i>Italic</i> | In a command line, indicates a variable. |
| <Enter> | Any individual key on the keyboard. |
| CTRL + Z | A keyboard combination that involves pressing the Z key while holding the CTRL key. |

Additional Documentation

The following documents for the Dell EMC Networking N-Series switches are available at www.dell.com/support:

- Getting Started Guide—provides information about the switch models in the series, including front and back panel features. It also describes the installation and initial configuration procedures.
- CLI Reference Guide—provides information about the command-line interface (CLI) commands used to configure and manage the switch. The document provides in-depth CLI descriptions, syntax, default values, and usage guidelines.

Switch Feature Overview

This section describes the switch user-configurable software features.



NOTE: Before proceeding, read the release notes for this product. The release notes are part of the firmware download.

The topics covered in this section include:

- System Management Features
- Stacking Features
- Security Features
- Green Technology Features
- Power over Ethernet (PoE) Plus Features
- Switching Features
- Virtual Local Area Network Supported Features
- Spanning Tree Protocol Features
- Link Aggregation Features
- Routing Features
- IPv6 Routing Features
- Quality of Service (QoS) Features
- Layer-2 Multicast Features
- Layer-3 Multicast Features

System Management Features

Multiple Management Options

Any of the following methods can be used to manage the switch:

- Use a web browser to access the Dell EMC OpenManage Switch Administrator interface. The switch contains an embedded Web server that serves HTML pages. Dell EMC Networking N-Series switches support HTTP and HTTPS over IPv4 or IPv6.
- Use a Telnet client, SSH client, or a direct console connection to access the CLI. The CLI syntax and semantics conform as much as possible to common industry practice. Dell EMC Networking N-Series switches support Telnet and SSH access over IPv4 or IPv6.
- Use a network management system (NMS), like the Dell EMC OpenManage Network Manager, to manage and monitor the system through SNMP. The switch supports SNMP v1/v2c/v3 over the UDP/IP transport protocol.

Nearly all switch features support a pre-configuration capability, even when the feature is not enabled or the required hardware is not present. Pre-configured capabilities become active only when enabled (typically via an admin mode control) or when the required hardware is present (or both). For example, a port can be preconfigured with both trunk and access mode information. The trunk mode information is applied only when the port is placed into trunk mode and the access mode information is only applied when the port is placed into access mode. Likewise, OSPF routing can be configured in the switch without being enabled on any port. This capability is present in all of the switch management options.

System Time Management

The switch can be configured to obtain the system time and date through a remote Simple Network Time Protocol (SNTP) server, or the time and date can be set locally on the switch. The time zone and information about time shifts that might occur during summer months can also be configured. When SNTP is used to obtain the time, communications between the switch and the SNTP server can be encrypted.

The Dell EMC Networking SNTP client supports connection to SNTP servers over IPv4 or IPv6.

For information about configuring system time settings, see "Managing General System Settings" on page 445.

Log Messages

The switch maintains in-memory log messages as well as persistent logs. Remote logging can be configured so that the switch sends log messages to a remote syslog server. The switch can also be configured to email log messages to a configured SMTP server. This allows the administrator to receive the log message in a specified e-mail account. Switch auditing messages, CLI command logging, Web logging, and SNMP logging can be enabled or disabled.

Dell EMC Networking N-Series switches support logging to syslog servers over IPv4 or IPv6.

For information about configuring system logging, see "Monitoring and Logging System Information" on page 407.

System Reset

When the switch is reset, it logs the reason in the persistent log, which is displayed in the log on startup. The possible reasons for a switch reset are:

- Switch was reset due to operator intervention.
- Switch was reset due to a software exception.
- Switch was reset due to a watchdog expiration.
- Switch was reset due to a Stack Manager conflict.
- Switch was reset due to software-initiated exit.
- Switch was reset due to power disruption or unexpected restart (error[0x0]).

The last reason code is the default if none of the other conditions are detected.

Integrated DHCP Server



NOTE: This feature is not supported on the Dell EMC Networking N1100-ON/N1500 Series switches.

Dell EMC Networking N-Series switches include an integrated DHCP server that can deliver host-specific configuration information to hosts on the network. The switch DHCP server allows the configuration of IPv4 address pools (scopes), and when a host's DHCP client requests an address, the switch DHCP server automatically assigns the host an address from the pool. For information about configuring the DHCP server settings, see "DHCP Server Settings" on page 1163.

Management of Basic Network Information

The DHCP client on the switch allows the switch to acquire information such as the IPv4 or IPv6 address and default gateway from a network DHCP server. The DHCP client can also be disabled and static network information can be configured instead. Other configurable network information includes a Domain Name Server (DNS), hostname to IP address mapping, and a default domain name.

If the switch detects an IP address conflict on the management interface, it generates a trap and sends a log message.

For information about configuring basic network information, see "Setting the IP Address and Other Basic Network Information" on page 209.

IPv6 Management Features

Dell EMC Networking N-Series switches provide IPv6 support for many standard management features including HTTP, HTTPS/SSL, Telnet, SSH, syslog, SNMP, TFTP, and traceroute on both the in-band and out-of-band management ports.

Dual Software Images

Dell EMC Networking N-Series switches can store up to two software images. The dual image feature enables upgrading the switch without deleting the older software image. One image is designated as the active image and the other image as the backup image.

For information about managing the switch image, see "Images and File Management" on page 525.

File Management

Files, such as configuration files and system images, can be uploaded and downloaded using HTTP (web only), TFTP, Secure FTP (SFTP), or Secure Copy (SCP). Configuration file uploads from the switch to a server are a good way to back up the switch configuration. A configuration file can also be downloaded from a server to the switch to restore the switch to the configuration in the downloaded file.

Files can be copied to and from a USB Flash drive that is plugged into the USB port on the front panel of the switch. Or, the switch can be automatically upgraded by booting it with a newer firmware image on a USB drive plugged into the switch. Dell EMC Networking N-Series switches support file copy protocols to both IPv4 and IPv6 servers.

For information about uploading, downloading, and copying files, see "Images and File Management" on page 525.

Switch Database Management Templates

Switch Database Management (SDM) templates enable reallocating system resources to support a different mix of features based on network requirements. Dell EMC Networking N-Series switches support the following three templates:

- Dual IPv4 and IPv6 (default)
- IPv4 Routing
- IPv4 Data Center

For information about setting the SDM template, see "Managing General System Settings" on page 445.

Automatic Installation of Firmware and Configuration

The Auto Install feature allows the switch to upgrade or downgrade to a newer software image and update the configuration file automatically during device initialization with limited administrative configuration on the device. If a USB device is connected to the switch and contains a firmware image and/or configuration file, the Auto Install feature installs the image or configuration file from USB device. Otherwise, the switch can obtain the necessary information from a DHCP server on the network.



NOTE: Automatic migration of the startup configuration to the next version of firmware from the current and previous versions of firmware is supported; the syntax is automatically updated when it is read into the running-config. Check the release notes to determine if any parts of the configuration cannot be migrated. Save the running-config to maintain the updated syntax. Migration of configuration is not assured on a firmware downgrade. When upgrading or downgrading firmware, check the configuration to ensure that it implements the desired configuration. Meta-configuration data (stack-port and slot configuration) is always reset to the defaults on a downgrade on each stack unit. As an example, Ethernet ports configured as stacking ports default back to Ethernet mode on a downgrade.

Migration of configuration information is never assured when errors are shown while the system is booting. Although the errored lines are displayed, commands that enter a sub-configuration mode followed by an exit command cause the CLI to exit Global Configuration mode, and subsequent configuration commands are ignored. Always hand-edit the startup-config if errors are shown on the screen during bootup.

For information about Auto Install, see "DHCP and USB Auto-Configuration" on page 557.

sFlow

sFlow is the standard for monitoring high-speed switched and routed networks. sFlow technology is built into network equipment and gives complete visibility into network activity, enabling effective management and control of network resources. The Dell EMC Networking N-Series switches support sFlow version 5.

For information about configuring managing sFlow settings, see "Monitoring Switch Traffic" on page 577.

SNMP Alarms and Trap Logs

The system logs events with severity codes and timestamps. The events are sent as SNMP traps to a trap recipient list.

For information about configuring SNMP traps and alarms, see "SNMP" on page 487.

CDP Interoperability Through ISDP

Industry Standard Discovery Protocol (ISDP) allows the Dell EMC Networking N-Series switch to interoperate with Cisco devices running the Cisco Discovery Protocol (CDP). ISDP is a proprietary layer-2 network protocol which inter-operates with Cisco network equipment and is used to share information between neighboring devices (routers, bridges, access servers, and switches).

For information about configuring ISDP settings, see "Discovering Network Devices" on page 897.

Remote Monitoring (RMON)

RMON is a standard Management Information Base (MIB) that defines current and historical MAC-layer statistics and control objects, allowing real-time information to be captured across the entire network.

For information about configuring managing RMON settings, see "Monitoring Switch Traffic" on page 577.

N3000 Series Advanced and Advanced-Lite Firmware Images

There are two N3000/N3100 mixed stacking switch firmware images available. The N3xxx Advanced stacking image supports stacking up to 12 units and supports stacking with the following models: N3132P-ON, N3024EP-ON, N3024ET-ON, N3024EF-ON, N3048ET-ON, N3048EP-ON. The filename for this image is N3000N3100Advv6.5.1.x.itb. The N3xxx Advanced-Lite stacking image supports stacking up to eight units and supports stacking for the following models: N3132P-ON, N3024EP-ON, N3024ET-ON, N3024EF-ON, N3048ET-ON, N3048EP-ON, N3048, N3048P, N3024, N3024P, and N3024F and is named N3000N3100AdvLitev6.5.1.X.itb. The Advanced-Lite firmware supports 1024 VLANs in the range 1-4093 and does not support MMRP/MVRP.

Which image type is installed can be determined by examining the first few lines of the running-config. The following example shows an Advanced version firmware.

```
console#show running-config
!Current Configuration:
!Software Capability "Stack Limit = 12, VLAN Limit = 4093"
!Image File "N3000Advv6.5.1.2"
```

```
!System Description "Dell EMC Networking N3048EP-ON, 6.5.1.2, Linux
3.6.5-d3d24324"
!System Software Version 6.5.1.2
!This firmware supports a stack of up to twelve switches.
!MVRP/MMRP capabilities and up to 4093 VLANs may be configured.
```

When migrating between the two types of images, certain commands in the startup-config may fail to execute because the relevant feature is not available. The switch firmware will identify any failed commands. It is necessary to edit the startup-config if errors are displayed and remove any failed commands. Do not simply save the running-config when commands in the startup-config fail, as the startup-config may contain modal commands that enter into a sub-mode not supported by the firmware. The exit command to exit the sub-mode may, in fact, exit Global Configuration mode, causing all subsequent commands to fail, even though those commands may be valid. When migrating from Adv firmware to AdvLite firmware, Dell recommends clearing the configuration and saving it to avoid any issues caused by exceeding the scaling limits of the AdvLite firmware.

In a mixed stack, all stack members contain the active, and optionally backup, firmware images for both types of units. It is necessary to have both images present on all units to allow any unit to perform the stack master function. A mixed stack may only be updated by installing an .itb firmware image using the copy command from any supported source, e.g. USB or TFTP. A mixed stack may also be upgraded by configuring an .itb firmware image as part of DHCP Auto-configuration that occurs as part of the switch boot process. The USB Auto-configuration upgrade feature should not be used to upgrade a mixed image stack. Do not attempt to update a mixed stack by installing an .stk firmware because a version mismatch may occur later after a stack master failover. In these cases, the recovery procedure is to reinstall the firmware using an .itb image.



NOTE: Switches must be upgraded to firmware version 6.5.1 or later in order to recognize and download an .itb image.

Do not download Adv mixed stack (itb) firmware to a stack containing legacy N3000 Series switches. It is required that a mixed stack containing legacy N3000 Series switches only be loaded with the AdvLite (itb) firmware. If

issues are encountered due to improper download of Adv firmware to a stack containing legacy N3000 Series switches, recover the stack by downloading AdvLite (itb) firmware.

To add a new member to an AdvLite or Adv mixed stack, load the new member with an ITB image, power off the unit, cable the new member into the stack and power it on. Then perform stack firmware synchronization using the boot auto-copy-sw command from the mixed stack master. Do not fail over to the new member prior to performing stack firmware synchronization.

To upgrade an AdvLite mixed stack to Adv mono-culture stack using the .stk firmware, power off the stack, re-cable the stack with the legacy N3000 switches removed from the stack (only N3000E-ON and N3132PX-ON switches can operate with the Adv firmware) and power on the stack starting with the desired stack master unit. Once the stack is fully powered, use the **clear config** command to remove the configuration of the units that are no longer participating in the stack. The **no member/no slot** commands in stack configuration mode also clear the stack member configuration, however, other configuration that references those members is not cleared and may cause issues on a reboot of the stack.

After the member configuration has been removed for all N3000 Series switches and the configuration has been saved, it is possible to download N3000Advv6.5.1.X.stk firmware to the stack master. Once downloaded, the master unit will distribute the images to other members of the stack. Use the **show bootvar [unit]** command to display the stack unit firmware. After code distribution is complete, reboot the stack.

This same technique may be used on the legacy N3000 switches that were configured in an AdvLite stack to allow them to download mono-culture N3000 (stk) firmware.

Stacking Features

For information about creating and maintaining a stack of switches, see "Stacking" on page 237.

Mixed and Single Series Stacking

The Dell EMC Networking N2000, N2100-ON, N3000, N3000E-ON, and N3100-ON Series switches include a stacking feature that allows multiple switches of the same or different series to operate as a single unit.

Dell EMC Networking N1100-ON Series switches stack with other Dell EMC N1100-ON Series switches and Dell EMC Networking N1500 Series switches stack with other Dell EMC N1500 Series switches.

The Dell EMC Networking N1124T-ON/N1124P-ON/N1148P-ON/N1148T-ON switches stack up to four units using 10G Ethernet ports configured for stacking. The Dell EMC Networking N1500 Series switches stack up to four units using 10GB Ethernet links configured as stacking.

Dell EMC Networking N2000 Series switches stack with other Dell EMC Networking N2000 Series switches and with Dell EMC Networking N2100-ON Series switches stack in a stack of up to 12 units. Dell EMC Networking N2000 and N2100-ON Series switches have two fixed mini-SAS stacking connectors at the rear. Any unit may be the stack master. The mixed stacking image name is N2000N2100Stdv6.5.1.X.itb.

Dell EMC Networking N2100-ON and N2000 switch series firmware is also available without mixed stacking capabilities. These images are named as follows:

N2100Stdv6.5.1.X.stk - N2100 only stack

N2000Stdv6.5.1.X.stk - N2000 only stack

Dell EMC Networking N3000 Series switches stack with other Dell EMC Networking N3000/N3000E-ON Series switches and with Dell EMC Networking N3100-ON Series switches stack in a stack of up to eight units. The Dell EMC Networking N3000/N3000E-ON Series switches have two fixed mini-SAS stacking connectors at the rear. The Dell EMC Networking N3100-ON Series switch has a slot in the rear that accepts an optional stacking module. Any unit may be the stack master. The image name is N3000N3100AdvLitev6.5.X.Y.itb.

Dell EMC Networking N3100-ON Series switches may also stack with the Dell EMC Networking N3000E-ON switches in a stack of up to 12 units. The image name is N3000N3100Advv6.5.1.X.itb. Any unit may be the stack master. N3024/N3024P/N3034F/N3048/N3048P units will be recognized if stacked with this image. However, the front panel interfaces will remain detached and inoperable.

Dell EMC Networking N3100-ON and N3000 switch series firmware is also available without mixed stacking capabilities. These images are named as follows:

N3100Advv6.5.1.X.stk - N3100 only stack

N3000Advv6.5.1.X.stk - N3000E-ON only stack

N3000AdvLitev6.5.1.X.stk - N3000 only stack (includes N3000E-ON support)

Dell EMC Networking N4000 Series switches stack with other Dell EMC Networking N4000 Series switches over front-panel ports configured for stacking.

Single IP Management

When multiple switches are connected together through the stack ports, they operate as a single unit with a larger port count. The stack operates and is managed as a single entity. One switch acts as the master, and the entire stack is managed through the management interface (Web, CLI, or SNMP) of the stack master.

Master Failover with Transparent Transition

The stacking feature supports a standby or backup unit that assumes the stack master role if the stack master fails. As soon as a stack master failure is detected, the standby unit initializes the control plane and enables all other stack units with the current configuration. The standby unit maintains a synchronized copy of the running configuration for the stack.

Nonstop Forwarding on the Stack

The Nonstop Forwarding (NSF) feature allows the forwarding plane of stack units to continue to forward packets while the control and management planes restart as a result of a power failure, hardware failure, or software fault on the stack master and allows the standby switch to quickly takeover as the master.

Hot Add/Delete and Firmware Synchronization

Units can be added to and deleted from the stack without cycling the power on the stack. Units to be added to the stack must be powered off prior to cabling into the stack to avoid election of a new master unit and a possible downgrade of the stack. When the newly added unit is powered on, the Stack Firmware Synchronization feature, if enabled, automatically synchronizes the firmware version with the version running on the stack master. The synchronization operation may result in either an upgrade or a downgrade of firmware on the mismatched stack member. Once the firmware is synchronized on a member unit, the running-config on the member is updated to match the master switch. The startup-config on the standby and member switches is not updated to match the master switch due to configuration changes on the master switch. Saving the startup config on the master switch also saves it to the startup config on all the other stack members. The hardware configuration of every switch is updated to match the master switch (unit number, slot configuration, stack member number, etc.).



NOTE: ALWAYS POWER OFF a unit to be added to a stack prior to cabling it into the stack. Newly added units must be powered on one-at-a-time beginning with the unit directly connected to an already powered on stack member.

Security Features

Configurable Access and Authentication Profiles

Rules can be configured to limit access to the switch management interface based on criteria such as access type and source IP address of the management host. The user can also be required to be authenticated locally or by an external server, such as a RADIUS server.

For information about configuring access and authentication profiles, see "Authentication, Authorization, and Accounting" on page 275.

Password-Protected Management Access

Access to the Web, CLI, and SNMP management interfaces is password protected, and there are no default users on the system.

For information about configuring local user accounts, see "Authentication, Authorization, and Accounting" on page 275.

Strong Password Enforcement

The Strong Password feature enforces a baseline password strength for all locally administered users. Password strength is a measure of the effectiveness of a password in resisting guessing and brute-force attacks. The strength of a password is a function of length, complexity and randomness. Using strong passwords lowers overall risk of a security breach.

For information about configuring password settings, see "Authentication, Authorization, and Accounting" on page 275.

TACACS+ Client

The switch has a TACACS+ client. TACACS+ provides centralized security for validation of users accessing the switch. TACACS+ provides a centralized user management system while still retaining consistency with RADIUS and other authentication processes.

For information about configuring TACACS+ client settings, see "Authentication, Authorization, and Accounting" on page 275.

RADIUS Support

The switch has a Remote Authentication Dial In User Service (RADIUS) client and can support up to 32 named authentication and accounting RADIUS servers. The switch also supports configuration of multiple RADIUS Attributes and accepts RADIUS COA termination requests. The switch can also be configured to accept RADIUS-assigned VLANs, ACLs and DiffServ Policies.

For information about configuring RADIUS client settings, see "Authentication, Authorization, and Accounting" on page 275.

SSH/SSL

The switch supports Secure Shell (SSH) for secure, remote connections to the CLI and Secure Sockets Layer (SSL) to increase security when accessing the web-based management interface. The SSH server can be enabled using the **ip ssh server** command or disabled using the **no ip ssh server** command.

For information about configuring SSH and SSL settings, see "Authentication, Authorization, and Accounting" on page 275.

Inbound Telnet Control

By default, the switch allows access over Telnet. The administrator can enable or disable the Telnet server using the **ip telnet server** command. Additionally, the Telnet port number is configurable using the same command.

For information about configuring inbound Telnet settings, see "Authentication, Authorization, and Accounting" on page 275.

Denial of Service

The switch supports configurable Denial of Service (DoS) attack protection for eight different types of attacks.

For information about configuring DoS settings, see "Port and System Security" on page 681.

Port Protection

A port may be put into the error-disabled state for any of the following reasons:

- **BPDU Storm:** By default, if Spanning Tree Protocol (STP) bridge protocol data units (BPDUs) are received at a rate of 15pps or greater for three consecutive seconds on a port, the port will be error-disabled. The threshold is not configurable.
- **Broadcast, Multicast, Unicast Storm:** If broadcast, unknown multicast, or unknown unicast packets are received at a rate greater than the configured limit and the configured action is to disable the port, the port will be error-disabled. Storm control is not enabled by default. See the **storm-control** commands for further information. A trap is issued for ports disabled by Storm Control.
- **DHCP Rate Limit:** If DHCP packets are received on a port at a rate that exceeds 15 pps, the port will be error-disabled. The threshold is configurable up to 300 pps for up to 15s long using the **ip dhcp snooping limit** command. DHCP snooping is disabled by default. The default protection limit is 15 pps. A trap is issued for interfaces disabled by DHCP Snooping.
- **DoS:** Interfaces on which a denial of service attack is detected are error-disabled. Refer to the **dos-control** command for configuration options.
- **ARP Inspection:** By default, if Dynamic ARP Inspection packets are received on a port at a rate that exceeds 15 pps for 1 second, the port will be error-disabled. The threshold is configurable up to 300 pps and the burst is configurable up to 15s long using the **ip arp inspection limit** command. A trap is issued for interfaces disabled by Dynamic ARP Inspection.
- **SFP Mismatch:** Insertion of an unsupported SFP transceiver will error-disable the interface. This behavior can be suppressed using the **service unsupported-transceiver** command.
- **SFP+ transceivers:** SFP+ transceivers are not compatible with SFP slots (N3024F front-panel ports). To avoid damage to SFP+ transceivers mistakenly inserted into SFP ports, the SFP port is error-disabled when an SFP+ transceiver is detected.
- **UDLD:** Interfaces on which unidirectional packet flow is detected are error-disabled.

- ICMP storms: Ports on which ICMP storms are detected are error-disabled. The rate limit and burst sizes are configurable separately for IPv4 and IPv6.
- PML: Interfaces on which the port security violation is configured to shut down the interface are error-disabled when a violation occurs.
- Loop Protect: Loop protection diagnostically disables ports on which a loop is detected. A log message may be issued when a port is disabled by Loop Protection.
- BPDU Guard: An interface that receives a BPDU with BPDU guard enabled is error-disabled. Use the `spanning-tree bpduguard` command to enable BPDU guard.

A port that is error-disabled may be returned to service using the `no shutdown` command. Alternatively, the operator may configure the auto recovery service to return the error disabled ports to service after a configurable period of time. Refer to the `errdisable recovery` command for more information.

Captive Portal

The Captive Portal feature blocks clients from accessing the network until user verification has been established. When a user attempts to connect to the network through the switch, the user is presented with a customized Web page that might contain username and password fields or the acceptable use policy. Users can be required to be authenticated by a local or remote RADIUS database before access is granted.

For information about configuring the Captive Portal features, see "Captive Portal" on page 370.

802.1X Authentication (IEEE 802.1X)

802.1X authentication enables the authentication of network clients through a local internal server or an external server. Only authenticated and approved network clients can transmit and receive frames over the port. Clients are authenticated using the Extensible Authentication Protocol (EAP). EAP-MD5 authentication with no privacy protocol is supported for switch-initiated (server-side) authentication to remote authentication servers. Local (IAS) authentication supports EAP-MD5 only. MAB supports EAP, PAP, and CHAP. Encrypted communication with authentication servers is not

supported; however, the switch will transport encrypted packets, such as PEAP or EAP-TLS packets, between the supplicant and authentication server in support of mutual authentication and privacy.

For information about configuring IEEE 802.1X settings, see "IEEE 802.1X" on page 334.

MAC-Based 802.1X Authentication

MAC-based authentication allows multiple supplicants connected to the same port to each authenticate individually. The switch uses the device's MAC address to restrict access to the port to only the devices that have authenticated. For example, a system attached to the port might be required to authenticate in order to gain access to the network, while a VoIP phone might not need to authenticate in order to send voice traffic through the port.

For information about configuring MAC-based 802.1X authentication, see "IEEE 802.1X" on page 334.

802.1X Monitor Mode

Monitor mode can be enabled in conjunction with 802.1X authentication to allow network access even when the user fails to authenticate. The switch logs the results of the authentication process for diagnostic purposes. The main purpose of this mode is to help troubleshoot the configuration of a 802.1X authentication on the switch without affecting the network access to the users of the switch.

For information about enabling the 802.1X Monitor mode, see "IEEE 802.1X" on page 334.

Port Security

The port security feature limits access on a port to users with specific MAC addresses. These addresses are manually defined or learned on that port. When a frame is seen on a locked port, and the frame source MAC address is not tied to that port, the protection mechanism is invoked.

For information about configuring port security, see "Port and System Security" on page 681.

Access Control Lists (ACLs)

Access Control Lists (ACLs) can help to ensure network availability for legitimate users while blocking attempts to access the network by unauthorized users or to restrict legitimate users from accessing the network. ACLs may be used to provide traffic flow control, restrict contents of routing updates, decide which types of traffic are forwarded or blocked, and above all, provide some level of security for the network. The switch supports the following ACL types:

- IPv4 ACLs
- IPv6 ACLs
- MAC ACLs

For all ACL types, the ACL rule can be configured to filter traffic when a packet enters or exits the Ethernet port, LAG, or VLAN interface. ACLs work only on switched ports. They do not operate on the out-of-band port.

ACLs can be used to implement policy-based routing (PBR) to implement packet routing according to specific organizational policies.

For information about configuring ACLs and PBR, see "Access Control Lists" on page 689.

Time-Based ACLs

With the Time-based ACL feature, the administrator can define when an ACL is in effect and the amount of time it is in effect.

For information about configuring time-based ACLs, see "Access Control Lists" on page 689.

IP Source Guard (IPSG)

IP source guard (IPSG) is a security feature that filters IP packets based on the source ID. The source ID may either be source IP address or a source IP address source MAC address pair as found in the local DHCP snooping database. IPSG depends on DHCP Snooping to associate IP address with MAC addresses.

For information about configuring IPSG, see "Snooping and Inspecting Traffic" on page 1015.

DHCP Snooping

DHCP Snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP server. It filters harmful DHCP messages and builds a bindings database of (MAC address, IP address, VLAN ID, port) tuples that are specified as authorized. DHCP snooping can be enabled globally and on specific VLANs. Ports within the VLAN can be configured to be trusted or untrusted. DHCP servers must be reached through trusted ports.

For information about configuring DHCP Snooping, see "Snooping and Inspecting Traffic" on page 1015.

Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) is a security feature that rejects invalid and malicious ARP packets. The feature prevents a class of man-in-the-middle attacks, where an unfriendly station intercepts traffic for other stations by poisoning the ARP caches of its unsuspecting neighbors. The malicious station sends ARP requests or responses mapping another station's IP address to its own MAC address.

Dynamic ARP Inspection relies on DHCP Snooping.

For information about configuring DAI, see "Snooping and Inspecting Traffic" on page 1015.

Protected Ports (Private VLAN Edge)

Private VLAN Edge (PVE) ports are a layer-2 security feature that provides port-based security between ports that are members of the same VLAN. It is an extension of the common VLAN. Traffic from protected ports is sent only to the uplink ports and cannot be sent to other ports within the VLAN.

For information about configuring IPSC, see "Port-Based Traffic Control" on page 921.

Green Technology Features

For information about configuring Green Technology features, see "Port Characteristics" on page 649.

Energy Detect Mode

When the Energy Detect mode is enabled and the port link is down, the PHY automatically goes down for short period of time and then wakes up periodically to check link pulses. This mode reduces power consumption on the port when no link partner is present. Energy Detect is proprietary and operates independently from EEE.

Energy Efficient Ethernet

Dell EMC Networking switches support IEEE 802.3az Energy Efficient Ethernet (EEE) Lower Power Idle Mode on front panel copper ports, which enables both the send and receive sides of the link to disable some functionality for power savings when the link is lightly loaded. EEE is standardized by the IEEE and operates independently of Energy Detect. EEE requires auto-negotiation to be enabled. Setting a port to a forced speed disables EEE.

EEE and Energy Detect are supported on the Dell EMC Networking N2000, N2100-ON, N3000, and N3100-ON Series 1G copper ports. EEE and energy detect are supported on the Dell EMC Networking N4000 Series 10G copper ports. EEE is supported on Gigabit Ethernet ports 1-8 on the N1108 Series switches, on Gigabit Ethernet ports 5-20 on the N1124 Series switches, and Gigabit Ethernet ports 9-24 and 29-44 on the N1148 Series switches. EEE is supported on Gigabit Ethernet ports 1-17 on the N1524 and Gigabit Ethernet ports 9-41 on the N1548. Energy detect is supported on all Gigabit Ethernet ports on the N1100 and N1500 Series switches.

EEE and Energy Detect are enabled by default on the N-Series copper ports. Energy Detect is enabled by default on the Dell EMC Networking N4000 Series switches 10G copper ports and cannot be disabled. Energy detect is enabled by default on the other Dell EMC Networking N-Series switches. Neither energy-detect nor EEE are supported on out-of-band, 2.5G or 5G NBASE-T ports.

Power Utilization Reporting

The switch displays the current power consumption of the power supply (or power supplies). This information is available from the management interface.

Power over Ethernet (PoE) Plus Features



NOTE: The Dell EMC Networking N1108P-ON/N1124P-ON/N1148P-ON, N1524P/N1548P, N2024P/N2048P/N2128PX-ON and N3024P/N3048P/N3024EP-ON/N3048EP-ON/N3132PX-ON switches support PoE Plus. The N2128PX-ON/N3024P/N3048P/N3024EP-ON/N3048EP-ON/N3132PX-ON switches support PoE 60W on selected ports. The PoE feature does not apply to the other models in the Dell EMC Networking N1100-ON, N1500, N2000, N2100-ON, N3000, N3000E-ON, N3100-ON, and N4000 Series.

For information about configuring PoE Plus features, see "Managing General System Settings" on page 445.

Key PoE Plus Features for the Dell EMC Networking N1108P-ON, N1124P-ON, N1148P-ON, N2024P, N2048P, N2128PX-ON, N3024P, N3048P, N3024EP-ON, N3048EP-ON, and N3132PX-ON Switches

Table 2-1 describes some of the key PoE Plus features.

Table 2-1. PoE Plus Key Features

| Feature | Description |
|-------------------------------|---|
| Global Usage Threshold | Provides the ability to specify a power limit as a percentage of the maximum power available to PoE ports. Setting a limit prevents the PoE switch from reaching an overload condition. |
| Per-Port Power Prioritization | Provides the ability to assign a priority to each PoE port. When the power budget of the PoE switch has been exhausted, the higher-priority ports are given preference over the lower-priority ports. Lower priority ports are automatically stopped from supplying power in order to provide power to higher-priority ports. |
| Per-Port Power Limit | Configurable power limit for each PoE-Plus port. |

Table 2-1. PoE Plus Key Features (Continued)

| Feature | Description |
|------------------------|--|
| Power Management Modes | Supports three power-management modes: <ul style="list-style-type: none">• Static—Reserves a configurable amount of power for a PoE port.• Dynamic—Power is not reserved for the port at any point of time. Power is supplied based upon the detected powered device (PD) signature.• Class-based—Reserves a classed-based amount of power for a PoE port. The final power delivered is determined via LLDP-MED negotiation, which allows for refinement of the power limit. |
| Power Detection Mode | Sets the mode to 802.3at or 802.3at+legacy detection. |

Power Over Ethernet (PoE) Plus Configuration

The Dell EMC Networking N1108P-ON/N1124P-ON/N1148P-ON, N1524P/N1548P, N2024P/N2048P, N2128PX-ON, N3024P/N3048P/N3024EP-ON/N3048EP-ON, and N3132PX-ON switches support PoE Plus configuration for power threshold, power priority, SNMP traps, and PoE legacy device support. Power can be limited on a per-port basis.

PoE Plus Support

The Dell EMC Networking N1108P-ON/N1124P-ON/N1148P-ON, N1524P/N1548P, N2024P/N2048P, N2128PX-ON, N3024P/N3048P/N3024EP-ON/N3048EP-ON, and N3132PX-ON switches implement the PoE Plus specification (IEEE 802.1at), in addition to the IEEE 802.3AF specification. This allows power to be supplied to Class 4 PD devices that require power greater than 15.4 Watts. Each port is capable of delivering up to 34.2W of power. Real-time power supply status is also available on the switch as part of the PoE Plus implementation.

PoE 60W Support

The Dell EMC Networking N3024P/N3048P/N3024EP-ON/N3048EP-ON switches implement 4-pair PoE 60W on the first 12 1G ports. The N3132PX-ON switches implement PoE 60W on the copper 1G and 5G ports. The N2128PX-ON switches implement PoE 60W on the 2.5G ports. The N1108P-ON, N1124P-ON, 1148P-ON, N1524P, N1548P, N2024P, and N2048P switches do not support PoE 60W.

PoE 60W allows power to be supplied to Class 5 powered devices that require power up to 60 watts. PoE 60W power must be configured manually. Class-based and dynamic power allocation is not supported for PoE 60W.

Class D or better cabling is required for feeds in excess of 34.2 watts. Normally, CAT 5E cabling does meet this requirement.

PoE-capable switches that are connected to another PSE supplying power will stop supplying power on the affected ports. PSE capability should be disabled when connecting Dell EMC PoE enabled ports to other PSE equipment.

Powered Device Detection

The switch is capable, based upon configuration, of detecting legacy, AF, or AT devices in two-pair or four-pair modes. AT detection is initiated first, followed by AF detection, and if configured, legacy detection. The switch always supplies full power to the port during power up and prior to performing detection.

PoE Power Management Modes

PoE-capable switches can be configured to manage powered devices (PD) using a dynamic, static, or class-based management. The power management mode is configured using the **power inline management** command.

Static Power Management

In this mode, the power reserved for the port is the configured limit regardless of whether the port is powered or not. The device may draw up to the configured limit. This mode is useful for devices that do not support LLDP-MED.

Available Power = Power Limit of the Sources – Total Configured Power

The total configured power is calculated as the sum of the configured power allocation for each port. Static mode reserves maximum power for the port, for example, 32W for two-pair mode and 60W for four-pair mode, unless a lower limit is configured by the administrator. Power is not reserved until a PD is connected to the port. The powered device may draw up to the configured limit. LLDP-MED packets requesting power are ignored in static mode. Do not configure the powered device to use LLDP-MED to request power in this mode.

Dynamic Power Management

In this mode, power is allocated based upon the detected PD class signature.

Available Power = Power Limit of the Sources – Total Allocated Power

The total allocated power is calculated as the sum of the power consumed by each port. Dynamic mode does not reserve power for the port (the port power limit is 0). Dynamic power management ignores LLDP-MED packets sent by the powered device. Do not configure the powered device to send LLDP-MED packets in this mode. The powered device may draw up to the detected class plus 5%.

Class-Based Power Management

Class-based power management allocates power based on the class selected by the detected powered device signature and LLDP-MED. The detection method must be configured as dot3at+legacy for AF signature devices to be detected.

Available Power = Power Limit of the Sources – Total Class Configured Power

The total class configured power is calculated as the sum of the class-based power allocation for each port. Note that class-based power management mode allocates the class limit for the port. The powered device may draw up to the class maximum based upon the detected powered device signature. The powered device need not draw all of the requested power. The Consumed Power display from the **show power inline** command shows the actual reported power draw and does not take into account the class reserved power. Configure the powered device to send LLDP-MED packets in this mode. It may take up to 60 seconds to fully power up a device in class-based management mode because LLDP-MED packets need to be exchanged in order to configure the desired power.

Power is supplied to the device in class mode per the following table:

| Class | Usage | AF Device (Watts) | AT Device (Watts) |
|-------|----------|-------------------|-------------------|
| 0 | Default | 16.4 | 33 |
| 1 | Optional | 5 | 33 |
| 2 | Optional | 8 | 33 |
| 3 | Optional | 16.4 | 33 |
| 4 | Optional | 16.4 | 33 |

In four-pair mode, twice the power listed in the table above is delivered. For information about the available system power, see the Hardware Overview chapter.

Power Management in Guard Band

The Dell EMC Networking N1100P-ON, N1500P, N2000P, N2100-ON, N3000P, and N3100-ON Series switches support a dynamic guard band, which means that the guard band used varies depending upon the following factors:

- Power management mode
- Class of the device being powered up.

Prior to a device being powered up, the switch calculates the following:

$\text{threshold power} - \text{guard band} - (\text{current power consumption} + \text{computed power draw of the new device})$

If this value is less than zero (which means powering up the new PD device will put the total power draw into the guard band or above the switch power capacity), then the switch does not power up the new device. A device being powered up in class or dynamic mode is always supplied with the maximum power (32 or 64 watts) at startup. Once the device class or power draw is determined, power to the device may be reduced.

The power management mode is configured using the **power inline management** command. The guard band is calculated by the switch as shown below. The user- defined threshold power limit can be found with the **show power inline detailed** command, and is configured with the **power inline usage-threshold** command. Threshold Power is reduced by the guard band when powering up a port.

If the remaining available power (threshold power - guard band - current power consumption) is less than the computed power draw of the new device, the device is not powered up. By default, the guard band is 32 watts.

Regardless of the power management mode, if the device being powered up is a Class 1, 2, or 3 AF device, then the guard band is configured according to the device class.

Dynamic or Static Power Management Mode Guard Band

In these modes, the guard band for the port being powered up is 32 watts.

Class-Based Power Management Mode Guard Band

In this mode, the dynamic guard band for the port being powered up is:

- For Class 0 AF device: 16.4 watts
- For Class 1 AF device: 5 watts
- For Class 2 AF device: 8 watts
- For Class 3 AF device: 16.4 watts
- For Class 4 AF device: 16.4 watts
- If the PD is an AT device, the guard band is 32 watts regardless of the detected class.

PoE Plus Default Settings

The following table shows the default PoE Plus settings for the Dell EMC Networking N1108P-ON/N1124P-ON/ N1148P-ON, N1524P/N1548P, N2024P/N2048P, N2128PX-ON, N3024P/N3048P, and N3132PX-ON switches.

Table 2-2. PoE Plus Key Features (Dell EMC Networking N1108P-ON/N1124P-ON/ N1148P-ON, N1524P/N1548P, N2024P/N2048P, N2128PX-ON, N3024P/N3048P/N3048EP-ON, and N3132PX-ON Only)

| Feature | Description |
|-------------------------------|--|
| Global Usage Threshold | 90% |
| Per-Port Admin Status | Auto |
| Per-Port Power Prioritization | Enabled (globally, per-port priority is Low) |
| Per-Port Power Limit | None |

Table 2-2. PoE Plus Key Features (Dell EMC Networking N1108P-ON/N1124P-ON/ N1148P-ON, N1524P/N1548P, N2024P/N2048P, N2128PX-ON, N3024P/N3048P/N3048EP-ON, and N3132PX-ON Only)

| Feature | Description |
|-----------------------|---------------------|
| Power Management Mode | Dynamic |
| Power Detection Mode | 802.3at plus legacy |
| Power Pairs | alternative-a |

Switching Features

Flow Control Support (IEEE 802.3x)


Flow control enables lower speed switches to communicate with higher speed switches by requesting that the higher speed switch refrain from sending packets for a limited period of time. Transmissions are temporarily halted to prevent buffer overflows.

For information about configuring flow control, see "Port-Based Traffic Control" on page 921.

Head of Line Blocking Prevention

Head of Line (HOL) blocking prevention prevents traffic delays and frame loss caused by traffic competing for the same egress port resources. HOL blocking queues packets, and the packets at the head of the queue are forwarded before packets at the end of the queue.

Alternate Store and Forward (ASF)

 **NOTE:** This feature is available on the Dell EMC Networking N4000 Series switches only.

The Alternate Store and Forward (ASF) feature reduces latency for large packets. When ASF is enabled, the memory management unit (MMU) can forward a packet to the egress port before it has been entirely received on the Cell Buffer Pool (CBP) memory.

AFS, which is also known as cut-through mode, is configurable through the command-line interface. For information about how to configure the AFS feature, see the CLI Reference Guide available at www.dell.com/support.

Jumbo Frames Support

Jumbo frames enable transporting data in fewer frames to ensure less overhead, lower processing time, and fewer interrupts.

For information about configuring the switch MTU, see "Port Characteristics" on page 649.

Auto-MDI/MDIX Support

The switch supports auto-detection between crossed and straight-through cables. Media-Dependent Interface (MDI) is the standard wiring for end stations, and the standard wiring for hubs and switches is known as Media-Dependent Interface with Crossover (MDIX). Auto-negotiation must be enabled for the switch to detect the wiring configuration. NBASE-T ports (2.5G and 5G) do not support auto-detection. Use the correct crossover or straight-through cable on 2.5/5G NBASE-T interfaces.

VLAN-Aware MAC-based Switching

Packets arriving from an unknown source address are sent to the CPU and added to the Hardware Table. Future packets addressed to or from this address are more efficiently forwarded.

Back Pressure Support

On half-duplex links, a receiver may prevent buffer overflows by jamming the link so that it is unavailable for additional traffic. On full-duplex links, a receiver may send a PAUSE frame indicating that the transmitter should cease transmission of frames for a specified period.



NOTE: Dell EMC Networking N2000/N2100-ON/N3000/N3100-ON/N4000 Series switches do not support half-duplex operation.

When flow control is enabled, the Dell EMC Networking N-Series switches will observe received PAUSE frames or jamming signals, but will not issue them when congested.

Auto-negotiation

Auto-negotiation allows the switch to advertise modes of operation. The auto-negotiation function provides the means to exchange information between two switches that share a point-to-point link segment and to automatically configure both switches to take maximum advantage of their transmission capabilities.

Dell EMC Networking N-Series switches enhance auto-negotiation by providing configuration of port advertisement. Port advertisement allows the system administrator to configure the port speeds that are advertised.

For information about configuring auto-negotiation, see "Port Characteristics" on page 649.

Storm Control

When layer-2 frames are processed, broadcast, unknown unicast, and multicast frames are flooded to all ports on the relevant virtual local area network (VLAN). The flooding occupies bandwidth and loads all nodes connected on all ports. Storm control limits the amount of broadcast, unknown unicast, and multicast frames accepted and forwarded by the switch.

For information about configuring Broadcast Storm Control settings, see "Port-Based Traffic Control" on page 921.

Port Mirroring

Port mirroring mirrors network traffic by forwarding copies of incoming and outgoing packets from multiple source ports to a monitoring port. Source ports may be VLANs, Ethernet interfaces, port-channels, or the CPU port. The switch also supports flow-based mirroring, which allows copying certain types of traffic to a single destination port using an ACL. This provides flexibility—instead of mirroring all ingress or egress traffic on a port the switch can mirror a subset of that traffic. The switch can be configured to mirror flows based on certain kinds of layer-2, layer-3, and layer-4 information.

Destination (probe) ports must be connected to a passive monitoring device. Traffic sent from the probe into the switch probe port is dropped. Mirrored traffic sent to the probe device will contain control plane traffic such as spanning-tree, LLDP, DHCP, etc.

Dell EMC Networking N-Series switches support RSPAN destinations where traffic can be tunneled across the operational network. Mirrored traffic is flooded in the RSPAN VLAN from the source(s) to the destination(s) across any intermediate switches. This allows the administrator flexibility in connecting destination (probe) ports to the RSPAN. RSPAN does not support configuration of the CPU port as a source.

For information about configuring port mirroring, see "Monitoring Switch Traffic" on page 577.

Static and Dynamic MAC Address Tables

Static entries can be added to the switch's MAC address table and the aging time can be configured for entries in the dynamic MAC address table. Entries can also be searched in the dynamic table based on several different criteria.

For information about viewing and managing the MAC address table, see "MAC Addressing and Forwarding" on page 1155.

Link Layer Discovery Protocol (LLDP)

The IEEE 802.1AB defined standard, Link Layer Discovery Protocol (LLDP), allows the switch to advertise major capabilities and physical descriptions. This information can be used to help identify system topology and detect bad configurations on the LAN.


For information about configuring LLDP, settings see "Discovering Network Devices" on page 897.

Link Layer Discovery Protocol (LLDP) for Media Endpoint Devices

The Link Layer Discovery Protocol for Media Endpoint Devices (LLDP-MED) provides an extension to the LLDP standard for network configuration and policy, device location, and Power over Ethernet.

For information about configuring LLDP-MED, settings see "Discovering Network Devices" on page 897.


Connectivity Fault Management (IEEE 802.1ag)

 **NOTE:** This feature is available on the Dell EMC Networking N4000 Series switches only.


The Connectivity Fault Management (CFM) feature, also known as Dot1ag, supports Service Level Operations, Administration, and Management (OAM). CFM is the OAM Protocol provision for end-to-end service layer instance in carrier networks. The CFM feature provides mechanisms to help perform connectivity checks, fault detection, fault verification and isolation, and fault notification per service in a network domain.

For information about configuring IEEE 802.1ag settings, see "Connectivity Fault Management" on page 995.

Priority-based Flow Control (PFC)

 **NOTE:** This feature is available on the Dell EMC Networking N4000 Series switches only.

The Priority-based Flow Control feature allows the switch to pause or inhibit transmission of individual priorities within a single Ethernet link. By configuring PFC to pause a congested priority (priorities) independently, protocols that are highly loss sensitive can share the same link with traffic that has different loss tolerances. Priorities are differentiated by the priority field of the 802.1Q VLAN header. The Dell EMC Networking N4000 Series switches support lossless transport of frames on up to two priority classes.

 **NOTE:** An interface that is configured for PFC is automatically disabled for 802.3x flow control.

For information about configuring the PFC feature, see "Data Center Bridging Features" on page 1123.

Data Center Bridging Exchange (DBCx) Protocol



NOTE: This feature is available on the Dell EMC Networking N4000 Series switches only.

The Data Center Bridging Exchange Protocol (DCBx) is used by DCB devices to exchange configuration information with directly connected peers. The protocol is also used to detect misconfiguration of the peer DCB devices and, optionally, for configuration of peer DCB devices. For information about configuring DCBx settings, see "Data Center Bridging Features" on page 1123. DCBx is a link-local protocol and operates only on individual links.

Enhanced Transmission Selection



NOTE: This feature is available on the Dell EMC Networking N4000 Series switches only.

Enhanced Transmission Selection (ETS) allows the switch to allocate bandwidth to traffic classes and share unused bandwidth with lower-priority traffic classes while coexisting with strict-priority traffic classes. ETS is supported on the Dell EMC Networking N4000 Series switches and can be configured manually or automatically using the auto configuration feature. For more information about ETS, see "Enhanced Transmission Selection" on page 1139.

Cisco Protocol Filtering

The Cisco Protocol Filtering feature (also known as Link Local Protocol Filtering) filters Cisco protocols that should not normally be relayed by a bridge. The group addresses of these Cisco protocols do not fall within the IEEE defined range of the 802.1D MAC Bridge Filtered MAC Group Addresses (01-80-C2-00-00-00 to 01-80-C2-00-00-0F).

For information about configuring LLPF, settings see "Port-Based Traffic Control" on page 921.

DHCP Layer-2 Relay

This feature permits layer-3 relay agent functionality in layer-2 switched networks. The switch supports layer-2 DHCP relay configuration on individual ports, link aggregation groups (LAGs) and VLANs.

For information about configuring layer-2 DHCP relay settings see "Layer-2 and Layer-3 Relay Features" on page 1229.

Virtual Local Area Network Supported Features

For information about configuring VLAN features see "VLANs" on page 763.

VLAN Support

VLANs are collections of switching ports that comprise a single broadcast domain. Packets are classified as belonging to a VLAN based on either the VLAN tag or a combination of the ingress port and packet contents. Packets sharing common attributes can be groups in the same VLAN. The Dell EMC Networking N-Series switches are in full compliance with IEEE 802.1Q VLAN tagging.

Port-Based VLANs

Port-based VLANs classify incoming packets to VLANs based on their ingress port. When a port uses 802.1X port authentication, packets can be assigned to a VLAN based on the result of the 802.1X authentication a client uses when it accesses the switch. This feature is useful for assigning traffic to Guest VLANs or Voice VLANs.

IP Subnet-based VLAN

This feature allows incoming untagged packets to be assigned to a VLAN and traffic class based on the source IP address of the packet.

MAC-based VLAN

This feature allows incoming untagged packets to be assigned to a VLAN and traffic class based on the source MAC address of the packet.

IEEE 802.1v Protocol-Based VLANs

VLAN classification rules are defined on data-link layer (layer-2) protocol identification. Protocol-based VLANs are used for isolating layer-2 traffic.

Voice VLAN

The Voice VLAN feature enables switch ports to carry voice traffic with a configured QoS and to optionally authenticate phones on the network. This allows preferential treatment of voice traffic over data traffic transiting the switch. Voice VLAN is the preferred solution for enterprises wishing to deploy VoIP services in their network.

GARP and GVRP Support

The switch supports the Generic Attribute Registration Protocol (GARP). GARP VLAN Registration Protocol (GVRP) relies on the services provided by GARP to provide IEEE 802.1Q-compliant VLAN pruning and dynamic VLAN creation on 802.1Q trunk ports. When GVRP is enabled, the switch registers and propagates VLAN membership on all ports that are part of the active spanning tree protocol topology.

For information about configuring GARP timers see "Layer-2 Multicast Features" on page 939.

Guest VLAN

The Guest VLAN feature allows the administrator to provide service to unauthenticated users, i.e., users that are unable to support 802.1X authentication.

For information about configuring the Guest VLAN see "Guest VLAN" on page 341.

Unauthorized VLAN

The Unauthorized VLAN feature allows the administrator to configure a VLAN for 802.1X-aware hosts that attempt authentication and fail.

Double VLANs



NOTE: DVLAN is not available on the N3000 running the AGREGATION ROUTER image.

The Double VLAN feature (IEEE 802.1QinQ) allows the use of a second tag on network traffic. The additional tag helps differentiate between customers in the Metropolitan Area Networks (MAN) while preserving individual customer's VLAN identification when they enter their own 802.1Q domain.

Spanning Tree Protocol Features

For information about configuring Spanning Tree Protocol features, see "Spanning Tree Protocol" on page 845.

Spanning Tree Protocol (STP)

Spanning Tree Protocol (IEEE 802.1D) is a standard requirement of layer-2 switches that allows bridges to automatically prevent and resolve layer-2 forwarding loops.

Spanning Tree Port Settings

The STP feature supports a variety of per-port settings including path cost, priority settings, Port Fast mode, STP Root Guard, Loop Guard, TCN Guard, and Auto Edge. These settings are also configurable per-LAG.

Rapid Spanning Tree

Rapid Spanning Tree Protocol (RSTP) detects and uses network topologies to enable faster spanning tree convergence after a topology change, without creating forwarding loops. The port settings supported by STP are also supported by RSTP.

Multiple Spanning Tree

Multiple Spanning Tree (MSTP) operation maps VLANs to spanning tree instances. Packets assigned to various VLANs are transmitted along different paths within MSTP Regions (MST Regions). Regions are one or more interconnected MSTP bridges with identical MSTP settings. The MSTP standard lets administrators assign VLAN traffic to unique paths.

The switch supports IEEE 802.1Q-2005, which corrects problems associated with the previous version, provides for faster transition-to-forwarding, and incorporates new features for a port (restricted role and restricted TCN).

Bridge Protocol Data Unit (BPDU) Guard

Spanning Tree BPDU Guard is used to disable the port in case a new device tries to enter the already existing topology of STP. Thus devices, which were originally not a part of STP, are not allowed to influence the STP topology.

BPDU Filtering

When spanning tree is disabled on a port, the BPDU Filtering feature allows BPDU packets received on that port to be dropped. Additionally, the BPDU Filtering feature prevents a port in Port Fast mode from sending and receiving BPDUs. A port in Port Fast mode is automatically placed in the forwarding state when the link is up to increase convergence time.

RSTP-PV and STP-PV

Dell EMC Networking N-Series switches support both Rapid Spanning Tree Per VLAN (RSTP-PV) and Spanning Tree Per VLAN (STP-PV). RSTP-PV is the IEEE 802.1w (RSTP) standard implemented per VLAN. A single instance of rapid spanning tree (RSTP) runs on each configured VLAN. Each RSTP instance on a VLAN has a root switch. STP-PV is the IEEE 802.1s (STP) standard implemented per VLAN.

Link Aggregation Features

For information about configuring link aggregation (port-channel) features, see "Link Aggregation" on page 1051.

Link Aggregation

Up to eight ports can combine to form a single Link Aggregation Group (LAG). This enables fault tolerance protection from physical link disruption, higher bandwidth connections and improved bandwidth granularity. LAGs are formed from similarly configured physical links; i.e., the speed, duplex, auto-negotiation, PFC configuration, DCBX configuration, etc., must be compatible on all member links.

Per IEEE 802.1AX, only links with the identical operational characteristics, such as speed and duplex setting, may be aggregated. Dell EMC Networking N-Series switches aggregate links only if they have the same operational speed and duplex setting, as opposed to the configured speed and duplex setting. This allows operators to aggregate links that use auto-negotiation to set values for speed and duplex or to aggregate ports with SFP+ technology operating at a lower speed, e.g., 1G. Dissimilar ports will not become active in the LAG if their operational settings do not match those of the first member of the LAG.

In practice, some ports in a LAG may auto-negotiate a different operational speed than other ports depending on the far-end settings and any link impairments. Per the above, these ports will not become active members of the LAG. On a reboot or on flapping the LAG links, a lower-speed port may be the first port selected to be aggregated into the LAG. In this case, the higher-speed ports are not aggregated. Use the **lACP port-priority** command to select one or more primary links to lead the formation of the aggregation group.

While it is a requirement of a port-channel that the link members operate at the same duplex and speed settings, administrators should be aware that copper ports have larger latencies than fiber ports. If fiber and copper ports are aggregated together, packets sent over the fiber ports would arrive significantly sooner at the destination than packets sent over the copper ports. This can cause significant issues in the receiving host (e.g., a TCP receiver) as it would be required to buffer a potentially large number of out-

of-order frames. Devices unable to buffer the requisite number of frames will show excessive frame discard. Configuring copper and fiber ports together in an aggregation group is not recommended.


Link Aggregate Control Protocol (LACP)

Link Aggregate Control Protocol (LACP) uses peer exchanges across links to determine, on an ongoing basis, the aggregation capability of various links, and continuously provides the maximum level of aggregation capability achievable between a given pair of systems. LACP automatically determines, configures, binds, and monitors the binding of ports to aggregators within the system.

Multi-Switch LAG (MLAG)

Dell EMC Networking N-Series switches support the MLAG feature to extend the LAG bandwidth advantage across multiple Dell EMC Networking N-Series switches connected to a LAG partner device. The LAG partner device is unaware that it is connected to two peer Dell EMC Networking N-Series switches; instead, the two switches appear as a single switch to the partner. When using MLAG, all links can carry data traffic across a physically diverse topology and, in the case of a link or switch failure, traffic can continue to flow with minimal disruption.

Routing Features

 **NOTE:** The N1100-0N Series switches do not support routing.

Address Resolution Protocol (ARP) Table Management

Static ARP entries can be created, and many settings for the dynamic ARP table can be managed, such as age time for entries, retries, and cache size. The ARP table supports routing by caching MAC addresses corresponding to the IP addresses of attached stations.

For information about managing the ARP table, see "IP Routing" on page 1187.

VLAN Routing

Dell EMC Networking N-Series switches support VLAN routing. A VLAN-routed packet is routed based on a longest prefix match lookup of the destination IP address in the routing table and is forwarded on a different VLAN by rewriting the destination MAC address obtained from the ARP table, decrementing the TTL, recalculating the frame CRC, and transmitting the frame on the VLAN.

For information about configuring VLAN routing interfaces, see "Routing Interfaces" on page 1213.

IP Configuration

The switch IP configuration settings allow the configuration of network information for VLAN routing interfaces, such as the IP address and subnet mask. Global IP configuration settings for the switch allow enabling or disabling the generation of several types of ICMP messages, setting a default gateway, and enabling or disabling inter-VLAN routing of packets.

For information about managing global IP settings, see "IP Routing" on page 1187.

Open Shortest Path First (OSPF)



NOTE: This feature is not available on Dell EMC Networking N1100-ON or N1500 Series switches.

Open Shortest Path First (OSPF) is a dynamic routing protocol commonly used within medium-to-large enterprise networks. OSPF is an interior gateway protocol (IGP) that operates within a single autonomous system.

For information about configuring OSPF, see "OSPF and OSPFv3" on page 1257.

Border Gateway Protocol (BGP)



NOTE: This feature is not available on Dell EMC Networking N1100-ON, N1500, N2000, and N2100-ON Series switches.

BGP is a protocol used for exchanging reachability information between autonomous systems. BGP uses a standardized decision process, which, when used in conjunction with network policies configured by the administrator, support a robust set of capabilities for managing the distribution of routing information.

Dell EMC Networking supports BGP4 configured as an IGP or an EGP. As an IGP, configuration as a source or client route reflector is supported. Both IPv6 and IPv4 peering sessions are supported.

For more information about configuring BGP, see "BGP" on page 1397.

Virtual Routing and Forwarding (VRF)



NOTE: This feature is not available on Dell EMC Networking N1100-ON, N1500, N2000, and N2100-ON switches.

VRF allows multiple independent instances of the forwarding plane to exist simultaneously. This allows segmenting the network without incurring the costs of multiple routers. Each VRF instance operates as an independent VPN. The IP addresses assigned to each VPN may overlap. Static route leaking to and from the global instance is supported. VLANs associated with a VRF may not overlap with other VRF instances.

For more information about configuring VRFs, see "VRF" on page 1349.

BOOTP/DHCP Relay Agent

The switch BootP/DHCP Relay Agent feature relays BootP and DHCP messages between DHCP clients and DHCP servers that are located in different IP subnets.

For information about configuring the BootP/DHCP Relay agent, see "Layer-2 and Layer-3 Relay Features" on page 1229.

IP Helper and DHCP Relay

The IP Helper and DHCP Relay features provide the ability to relay various protocols to servers on a different subnet.

For information about configuring the IP helper and DHCP relay features, see "Layer-2 and Layer-3 Relay Features" on page 1229.

Routing Information Protocol

Routing Information Protocol (RIP), like OSPF, is an IGP used within an autonomous Internet system. RIP is an IGP that is designed to work with moderate-size networks.

For information about configuring RIP, see "RIP" on page 1355.

Router Discovery

For each interface, the Router Discovery Protocol (RDP) can be configured to transmit router advertisements. These advertisements inform hosts on the local network about the presence of the router.

For information about configuring router discovery, see "IP Routing" on page 1187.

Routing Table

The routing table displays information about the routes that have been dynamically learned. Static and default routes and route preferences can be configured. A separate table shows the routes that have been manually configured.

For information about viewing the routing table, see "IP Routing" on page 1187.

Virtual Router Redundancy Protocol (VRRP)

VRRP provides hosts with redundant routers in the network topology without any need for the hosts to reconfigure or know that there are multiple routers. If the primary (master) router fails, a secondary router assumes control and continues to use the virtual router IP (VRIP) address.

VRRP Route Interface Tracking extends the capability of VRRP to allow tracking of specific route/interface IP states within the router that can alter the priority level of a virtual router for a VRRP group.

For information about configuring VRRP settings, see "VRRP" on page 1371.

Tunnel and Loopback Interfaces



NOTE: This feature is not available on Dell EMC Networking N1100-ON or N1500 Series switches.

Dell EMC Networking N-Series switches support the creation, deletion, and management of tunnel and loopback interfaces. Tunnel interfaces facilitate the transition of IPv4 networks to IPv6 networks. A loopback interface is always expected to be up, so a stable IP address can be configured to enable other network devices to contact or identify the switch.

For information about configuring tunnel and loopback interfaces, see "Routing Interfaces" on page 1213.

IPv6 Routing Features



NOTE: This feature is not available on Dell EMC Networking N1100-ON, N1500, N2000, and N2100-ON Series switches.

IPv6 Configuration

The switch supports IPv6, the next generation of the Internet Protocol. IPv6 can be globally enabled on the switch and settings such as the IPv6 hop limit and ICMPv6 rate limit error interval can be configured. The administrator can also control whether IPv6 is enabled on a specific interface. The switch supports the configuration of many per-interface IPv6 settings including the IPv6 prefix and prefix length.

For information about configuring general IPv6 routing settings, see "IPv6 Routing" on page 1473.

IPv6 Routes

Because IPv4 and IPv6 can coexist on a network, the router on such a network needs to forward both traffic types. Given this coexistence, each switch maintains a separate routing table for IPv6 routes. The switch can forward IPv4 and IPv6 traffic over the same set of interfaces.

For information about configuring IPv6 routes, see "IPv6 Routing" on page 1473.

OSPFv3

OSPFv3 provides a routing protocol for IPv6 networking. OSPFv3 is a new routing component based on the OSPF version 2 component. In dual-stack IPv6, both OSPF and OSPFv3 components can be configured and used.

For information about configuring OSPFv3, see "OSPF and OSPFv3" on page 1257.

DHCPv6

DHCPv6 incorporates the notion of the “stateless” server, where DHCPv6 is not used for IP address assignment to a client, rather it only provides other networking information such as DNS, Network Time Protocol (NTP), and/or Session Initiation Protocol (SIP) information.

For information about configuring DHCPv6 settings, see "DHCPv6 Server Settings" on page 1501.

Quality of Service (QoS) Features



NOTE: Some features that can affect QoS, such as ACLs and Voice VLAN, are described in other sections within this chapter.

Differentiated Services (DiffServ)

The QoS Differentiated Services (DiffServ) feature allows traffic to be classified into streams and given certain QoS treatment in accordance with defined per-hop behaviors. Dell EMC Networking N-Series switches support both IPv4 and IPv6 packet classification.

For information about configuring DiffServ, see "Differentiated Services" on page 1521.

Class Of Service (CoS)

The Class Of Service (CoS) queuing feature enables directly configuring certain aspects of switch queuing. This provides the desired QoS behavior for different types of network traffic when the complexities of DiffServ are not required. CoS queue characteristics, such as minimum guaranteed bandwidth and transmission rate shaping, are configurable at the queue (or port) level.

For information about configuring CoS, see "Class-of-Service" on page 1559.

Auto Voice over IP (VoIP)

This feature provides ease of use for the user in setting up VoIP for IP phones on a switch. This is accomplished by enabling a VoIP profile that a user can select on a per port basis.

For information about configuring Auto VoIP, see "Auto VoIP" on page 1587.

This capability is not available on the N3000 Series switches when running the AGGRAGATION ROUTER image.

Internet Small Computer System Interface (iSCSI) Optimization

The iSCSI Optimization feature helps network administrators track iSCSI traffic between iSCSI initiator and target systems. This is accomplished by monitoring, or snooping traffic to detect packets used by iSCSI stations in establishing iSCSI sessions and connections. Data from these exchanges may optionally be used to create classification rules to assign the traffic between the stations to a configured traffic class. This affects how the packets in the flow are queued and scheduled for egress on the destination port.

For information about configuring iSCSI settings, see "iSCSI Optimization" on page 635.

Layer-2 Multicast Features

For information about configuring layer-2 multicast features, see "Layer-2 Multicast Features" on page 939.

MAC Multicast Support

Multicast service is a limited broadcast service that supports one-to-many and many-to-many forwarding behavior. In the layer-2 multicast service, a single frame addressed to a specific multicast address is received and copies of the frame to be transmitted on each relevant port are forwarded.

IGMP Snooping

Internet Group Management Protocol (IGMP) Snooping is a feature that allows a switch to forward multicast traffic intelligently on the switch. Multicast traffic is traffic that is destined to a host group. Host groups are identified by the destination MAC address, i.e. the range 01:00:5e:00:00:00 to 01:00:5e:7f:ff:ff:ff for IPv4 multicast traffic or 33:33:xx:xx:xx:xx for IPv6 multicast traffic. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance.

IGMP Snooping Querier

When Protocol Independent Multicast (PIM) and IGMP are enabled in a network with IP multicast routing, an IP multicast router acts as the IGMP querier. However, if it is desirable to keep the multicast network layer-2 switched only, the IGMP Snooping Querier can perform the query functions of a layer-3 multicast router.

MLD Snooping

In IPv4, layer-2 switches can use IGMP Snooping to limit the flooding of multicast traffic by dynamically configuring layer-2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast address.

In IPv6, MLD snooping performs a similar function. With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports intended to receive the data (instead of being flooded to all of the ports in a VLAN). This list is constructed by snooping IPv6 multicast control packets.

Multicast VLAN Registration

The Multicast VLAN Registration (MVR) protocol, like IGMP Snooping, allows a layer-2 switch to listen to IGMP frames and forward the multicast traffic only to the receivers that request it. Unlike IGMP Snooping, MVR allows the switch to forward multicast frames across different VLANs. MVR uses a dedicated VLAN, which is called the multicast VLAN, to forward multicast traffic over the layer-2 network to the various VLANs that have multicast receivers as members.

Layer-3 Multicast Features

For information about configuring layer-3 (L3) multicast features, see "IPv4 and IPv6 Multicast" on page 1593.



NOTE: This feature is not available on Dell EMC Networking N1100-ON, N1500, N2000, and N2100-ON Series switches.

Distance Vector Multicast Routing Protocol

Distance Vector Multicast Routing Protocol (DVMRP) exchanges probe packets with all DVMRP-enabled routers, establishing two way neighboring relationships and building a neighbor table. It exchanges report packets and creates a unicast topology table, which is used to build the multicast routing table. This multicast route table is then used to route the multicast packets.

Internet Group Management Protocol

The Internet Group Management Protocol (IGMP) is used by IPv4 systems (hosts and routers) to report their IP multicast group memberships to any neighboring multicast routers. Dell EMC Networking N-Series switches perform the “multicast router part” of the IGMP protocol, which means it collects the membership information needed by the active multicast router.

IGMP Proxy

The IGMP Proxy feature allows the switch to act as a proxy for hosts by sending IGMP host messages on behalf of the hosts that the switch discovered through standard IGMP router interfaces.

Protocol Independent Multicast—Dense Mode

Protocol Independent Multicast (PIM) is a standard multicast routing protocol that provides scalable inter-domain multicast routing across the Internet, independent of the mechanisms provided by any particular unicast routing protocol. The Protocol Independent Multicast-Dense Mode (PIM-DM) protocol uses an existing Unicast routing table and a Join/Prune/Graft mechanism to build a tree. PIM-DM creates source-based shortest-path distribution trees, making use of reverse path forwarding (RPF).

Protocol Independent Multicast—Sparse Mode

Protocol Independent Multicast-Sparse Mode (PIM-SM) is used to efficiently route multicast traffic to multicast groups that may span wide area networks, and where bandwidth is a constraint. PIM-SM uses shared trees by default and implements source-based trees for efficiency. This data threshold rate is used to toggle between trees.

Protocol Independent Multicast—Source Specific Multicast

Protocol Independent Multicast—Source Specific Multicast (PIM-SSM) is a subset of PIM-SM and is used for one-to-many multicast routing applications, such as audio or video broadcasts. PIM-SSM does not use shared trees.

Protocol Independent Multicast IPv6 Support

PIM-DM and PIM-SM support IPv6 routes.

MLD/MLDv2 (RFC2710/RFC3810)

MLD is used by IPv6 systems (listeners and routers) to report their IP multicast addresses memberships to any neighboring multicast routers. The implementation of MLD v2 is backward compatible with MLD v1.

MLD protocol enables the IPv6 router to discover the presence of multicast listeners, the nodes that want to receive the multicast data packets, on its directly attached interfaces. The protocol specifically discovers which multicast addresses are of interest to its neighboring nodes and provides this information to the multicast routing protocol that make the decision on the flow of the multicast data packets.

Hardware Overview

This section provides an overview of the switch hardware. It is organized by product type:

- Dell EMC Networking N1100-ON Series Switch Hardware
- Dell EMC Networking N1500 Series Switch Hardware
- Dell EMC Networking N2000 Series Switch Hardware
- Dell EMC Networking N2100-ON Series Switch Hardware
- Dell EMC Networking N3000/N3000E-ON Series Switch Hardware
- Dell EMC Networking N3100-ON Series Switch Hardware
- Dell EMC Networking N4000 Series Switch Hardware
- Switch MAC Addresses

Dell EMC Networking N1100-ON Series Switch Hardware

This section contains information about device characteristics and modular hardware configurations for the Dell EMC Networking N1100-ON switch.

Front Panel

The N1108-ON models are half-width, 1U, rack-mountable switches. To rack-mount the N1108-ON switches, either the Dell EMC Rack Mount Kit or the Dell EMC Tandem Tray Kit is required. The N1124-ON and N1148-ON models are full-width, 1U, rack mountable switches.

The Dell EMC Networking N1108-ON front panel provides eight 10/100/1000BASE-T Ethernet RJ-45 ports, capable of full or half duplex operation, two 100/1000BASE-T RJ45 uplink ports capable of full duplex operation only, and two SFP ports. The SFP ports are capable of 1G full duplex operation only. The N1108P-ON supports two PoE+ or four PoE ports on Gigabit Ethernet ports 1-4. Dell-qualified SFP transceivers are sold separately.

The Dell EMC Networking N1124-ON front panel provides 24 10/100/1000BASE-T Ethernet RJ-45 ports capable of full and half duplex operation, and four SFP+ ports. The N1124P-ON supports six PoE+ or 12 PoE ports on Gigabit Ethernet ports 5-16. Dell EMC-qualified SFP+ transceivers are sold separately.

The Dell EMC Networking N1148-ON front panel provides 48 10/100/1000BASE-T Ethernet RJ-45 ports capable of full and half duplex operation, and four SFP+ ports. The N1148P-ON supports twelve PoE+ or 24 PoE ports on Gigabit Ethernet ports 1-24. Dell EMC-qualified SFP+ transceivers are sold separately.

Figure 3-1. Dell EMC Networking N1108P-ON Switch (Front Panel)





Figure 3-2. Dell EMC Networking N1148P-ON Switch (Front Panel)



Console Port

The console port provides serial communication capabilities, which allows communication using RS-232 protocol. The micro-USB port provides a direct connection to the switch and allows access to the CLI from a console terminal connected to the port through the provided USB cable (with a male USB micro B to male USB type A connector).

The console port is separately configurable and can be run as an asynchronous link from 1200 BAUD to 115,200 BAUD. The Dell EMC CLI supports changing only the speed of the console port. The defaults are 115,200 BAUD, 8 data bits, no parity, 1 stop bit, and no flow control.

USB Port

The Type-A, female USB port supports a USB 2.0-compliant flash memory drive. The Dell EMC Networking N-Series switch can read or write to a flash drive with a single partition formatted as FAT-32. Use a USB flash drive to copy switch configuration files and images between the USB flash drive and the switch. The USB flash drive may be used to move and copy configuration files and images from one switch to other switches in the network. The system does not support the deletion of files on USB flash drives.

The USB port does not support any other type of USB device.

Port and System LEDs

The front panel contains light emitting diodes (LEDs) that indicate the status of port links, power supplies, fans, stacking, and the overall system status. See "LED Definitions" on page 116 for more information.

Stack Master LED

When a switch within a stack is the master unit, the Stack Master LED is solid green. If the Stack Master LED is off, the stack member is not the master unit. If a switch is not part of a stack (in other words, it is a stack of one switch), the Stack Master LED is illuminated.

Information Tag

The front panel includes a slide-out label panel that contains system information, such as the Service Tag, MAC address, and so on.

Power Supply

The internal power supply wattage for the Dell EMC Networking N1100-ON switches is as follows:

- N1108T-ON: 24W
- N1108P-ON: 80W
- N1124T-ON: 40W
- N1124P-ON: 250W
- N1148T-ON: 60W
- N1148P-ON: 500W

For information about power consumption for the N1100-ON PoE switches, see "Power Consumption for PoE Switches" on page 120.

Ventilation System

The N1108T-ON, N1124T-ON, and N1148T-ON switches are fanless. The N1108P-ON has one internal fan, and the N1124P-ON and N1148P-ON each have two internal fans.

Thermal Shutdown

Upon reaching critical temperature, the switch will shut down for 5 minutes and then automatically power on again. This cycle will repeat for as long as the switch is at or above critical temperature. During shutdown, the fans of switches so equipped will remain operational.

LED Definitions

This section describes the LEDs on the front panel of the switch.

Port LEDs

Each port on a Dell EMC Networking N1100-ON Series switch includes two LEDs. One LED is on the left side of the port, and the second LED is on the right side of the port. This section describes the LEDs on the switch ports.

Each 100/1000/10000BASE-T port has two LEDs. Figure 3-16 illustrates the 100/1000/10000BASE-T port LEDs.

Figure 3-3. 100/1000/10000BASE-T Port LEDs




Table 3-19 shows the 100/1000/10000BASE-T port LED definitions.

Table 3-1. 100/1000/10000BASE-T Port LED Definitions

| LED | Color | Definition |
|---------------------------------------|----------------|---|
| Link/SPD LED | Off | There is no link. |
| | Solid amber | The port is operating at 10/100 Mbps. |
| | Solid green | The port is operating at 1000 Mbps. |
| Activity LED (on non-PoE switches) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |
| Activity/PoE LED (on PoE switches) | Off | There is no current transmit/receive activity and PoE power is off. |
| | Blinking green | The port is actively transmitting/receiving and PoE power is off. |
| | Blinking amber | The port is actively transmitting/receiving and PoE power is on. |
| | Solid amber | There is no current transmit/receive activity and PoE power is on. |

Table 3-2. SFP Port LED Definitions (N1108-ON Only)

| LED | Color | Definition |
|-----------------------------|----------------|--|
| Link/SPD LED (Left LED) | Off | There is no link. |
| | Solid green | The port is operating at 1 Gbps. |
| Activity LED (Right LED) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Table 3-3. SFP+ Port LED Definitions (N1124-ON and N1148-ON Only)

| LED | Color | Definition |
|--|----------------|--|
| Link/SPD LED (Left bi-color LED) | Off | There is no link. |
| | Solid green | The port is operating at 10 Gbps. |
| | Solid amber | The port is operating at 1 Gbps. |
| Activity LED (Right single- color LED) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Stacking Port LEDs

Table 3-4. Stacking Port LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is actively transmitting/receiving. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

System LEDs

The system LEDs, located on the front panel, provide information about the power supplies, thermal conditions, and diagnostics.

Table 3-25 shows the System LED definitions for the Dell EMC Networking N1100-ON switches.

Table 3-5. System LED Definitions

| LED | Color | Definition |
|--------------------|----------------|--|
| Status | Solid green | Normal operation. |
| | Blinking green | The switch is booting |
| | Solid amber | A critical system error has occurred. |
| | Blinking amber | A noncritical system error occurred (fan or power supply failure). |
| Power | Off | There is no power or the switch has experienced a power failure. |
| | Solid green | Power to the switch is on. |
| | Solid green | POST is in progress. |
| Master | Off | The switch is not in master mode. |
| | Solid green | The switch is master for the stack. |
| Temp | Solid green | The switch is operating below the threshold temperature. |
| | Solid red | The switch temperature exceeds the threshold limit, or there is a fan failure (if fan-equipped). |
| System Locator LED | Blinking blue | The locator LED has been activated to locate the physical switch. |
| | Off | The beacon LED is idle. |

Power Consumption for PoE Switches

Table 3-6 describes the power consumption for N3132P-ON PoE switches. The PoE power budget is 60W for the N1108P-ON, 185W for the N1124P-ON, and 370W for the N1148P-ON.

Table 3-6. Power Consumption for N3132P-ON PoE Switches


| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Maximum Steady Power (W) |
|-----------|---------------|----------------------------|--|--------------------------|
| N1108P-ON | 100V/60Hz | Main PSU | 0.95A | 88.64W |
| | 110V/60Hz | Main PSU | 0.87A | 88.43W |
| | 120V/60Hz | Main PSU | 0.80A | 88.22W |
| | 220V/50Hz | Main PSU | 0.49A | 89.28W |
| | 240V/50Hz | Main PSU | 0.45A | 89.70W |
| N1124P-ON | 100V/60Hz | Main PSU | 2.66A | 260.66W |
| | 110V/60Hz | Main PSU | 2.38A | 257.95W |
| | 120V/60Hz | Main PSU | 2.16A | 256.27W |
| | 220V/50Hz | Main PSU | 1.18A | 250.52W |
| | 240V/50Hz | Main PSU | 1.10A | 251.25W |
| N1148P-ON | 100V/60Hz | Main PSU | 4.78A | 476.03W |
| | 110V/60Hz | Main PSU | 4.32A | 472.64W |
| | 120V/60Hz | Main PSU | 3.95A | 470.58W |
| | 220V/50Hz | Main PSU | 2.14A | 459.37W |
| | 240V/50Hz | Main PSU | 1.97A | 459.06W |

Wall Installation

To mount the switch on a wall:

- 1 Make sure that the mounting location meets the following requirements:
 - The surface of the wall must be capable of supporting the switch.
 - Allow at least two inches (5.1 cm) space on the sides for proper ventilation and five inches (12.7 cm) at the back for power cable clearance.
 - The location must be ventilated to prevent heat buildup.
- 2 Place the supplied wall-mounting bracket on one side of the switch, verifying that the mounting holes on the switch line up to the mounting holes on the wall-mounting bracket.

Figure 3-4. Bracket Installation for Wall Mounting



- 3 Insert the supplied screws into the wall-mounting bracket holes and tighten with a screwdriver.
- 4 Repeat the process for the wall-mounting bracket on the other side of the switch.

- 5** Place the switch on the wall in the location where the switch is being installed.
- 6** On the wall, mark the locations where the screws to hold the switch must be prepared.
- 7** On the marked locations, drill the holes and place all plugs (not provided) in the holes.
- 8** Secure the switch to the wall with screws (not provided). Make sure that the ventilation holes are not obstructed.

Dell EMC Networking N1500 Series Switch Hardware

This section contains information about device characteristics and modular hardware configurations for the Dell EMC Networking N1500 Series switches.


Front Panel

The Dell EMC Networking N1500 Series front panel includes the following features:

- Switch Ports
- Console Port
- USB Port
- Reset Button
- SFP+ Ports
- Port and System LEDs
- Stack Master LED and Stack Number Display

The following images show the front panels of the switch models in the Dell EMC Networking N1500 Series.


Figure 3-5. Dell EMC Networking N1548 Front-Panel Switch with 48 10/100/1000BASE-T Ports (Front Panel)



In addition to the switch ports, the front panel of each model in the Dell EMC Networking N1500 Series includes the following ports:

- RJ-45 Console port
- USB port for storage

Figure 3-6. Dell EMC Networking N1524P Close-up



The Dell EMC Networking 1524 front panel has status LEDs for over-temperature alarm (left), internal power (middle), and status (right) on the top row. The bottom row of status LEDs displays, from left to right, the Stack Master, redundant power supply (RPS) status, and fan alarm status.

The Dell EMC Networking 1524P front panel, shown in Figure 3-6, has status LEDs for over-temperature alarm, internal power, and status on the top row. The bottom row of status LEDs displays Stack Master, modular power supply (MPS) status, and fan alarm status.

Switch Ports

The Dell EMC Networking N1524/N1524P front panel provides 24 Gigabit Ethernet (10/100/1000BASE-T) RJ-45 ports that support auto-negotiation for speed, flow control, and duplex. The Dell EMC Networking N1500 Series front-panel ports operate in full- or half-duplex mode. The Dell EMC Networking N1524/N1524P models support four SFP+ 10G ports. Dell EMC-qualified SFP+ transceivers are sold separately.

The Dell EMC Networking N1548/N1548P front panel provides 48 Gigabit Ethernet (10BASE-T, 100BASE-TX, 1000BASE-T) RJ-45 ports that support auto-negotiation for speed, flow control, and duplex. The Dell EMC Networking N1500 Series front-panel ports operate in full- or half-duplex mode. The Dell EMC Networking N1548/N1548P supports four SFP+ 10G ports. Dell EMC-qualified SFP+ transceivers are sold separately.

The front-panel switch ports have the following characteristics:

- The switch automatically detects the difference between crossed and straight-through cables on RJ-45 ports and automatically chooses the MDI or MDIX configuration to match the other end.
- SFP+ ports support Dell EMC-qualified transceivers utilizing 10GBASE-SR, 10GBASE-LR, 10GBASE-CR, or 1000BASE-X technologies. The default behavior is to log a message and generate an SNMP trap on insertion or removal of an optic that is not qualified by Dell EMC. The message and trap can be suppressed by using the **service unsupported-transceiver** command.
- RJ-45 front-panel ports support full- or half-duplex mode 10/100/1000 Mbps speeds on standard Category 5 UTP cable. 1000BASE-X and 1000BASE-T operation requires the use of auto-negotiation.
- SFP+ ports support SFP+ transceivers and SFP+ copper twin-ax technology operating at 10G or 1G speeds in full-duplex mode. SFP transceivers are supported in SFP+ ports and operate at 1G full-duplex. SFP transceivers require auto-negotiation to be enabled.

SFP+ ports may be configured to support 16 GB stacking over Ethernet cables. These ports may be configured to support stacking in pairs, e.g., Te1/0/1 and Te1/0/2 may be configured to support stacking, or Te1/0/3 and Te1/0/4 may be configured to support stacking, or all four ports may be configured to support stacking.

- The Dell EMC Networking N1524P/N1548P front-panel ports support PoE (15.4W) and PoE+ (34.2W) as well as legacy capacitive detection for pre-standard powered devices (PDs).

Console Port

The console port provides serial communication capabilities, which allows communication using the RS-232 protocol. The serial port provides a direct connection to the switch and allows access to the CLI from a console terminal connected to the port through the provided serial cable (with RJ45 YOST to female DB-9 connectors).

The console port is separately configurable and can be run as an asynchronous link from 1200 BAUD to 115,200 BAUD. The Dell CLI supports changing the speed only. The defaults are 9600 BAUD, 8 data bits, no parity, 1 stop bit, and no flow control.

USB Port

The Type-A, female USB port supports a USB 2.0-compliant flash memory drive. The Dell EMC Networking N-Series switch can read or write to a flash drive with a single partition formatted as FAT-32. Use a USB flash drive to copy switch configuration files and images between the USB flash drive and the switch. The USB flash drive may be used to move and copy configuration files and images from one switch to other switches in the network. The system does not support the deletion of files on USB flash drives.

The USB port does not support any other type of USB device.

Reset Button

The reset button is accessed through the pinhole and enables performing a hard reset on the switch. To use the reset button, insert an unbent paper clip or similar tool into the pinhole. When the switch completes the boot process after the reset, it resumes operation with the most recently saved configuration. Any changes made to the running configuration that were not saved to the startup configuration prior to the reset are lost.

Port and System LEDs

The front panel contains light emitting diodes (LEDs) that indicate the status of port links, power supplies, fans, stacking, and the overall system status. See "LED Definitions" on page 128 for more information.

Stack Master LED and Stack Number Display

When a switch within a stack is the master unit, the Stack Master LED is solid green. If the Stack Master LED is off, the stack member is not the master unit. The Stack No. panel displays the unit number for the stack member. If a switch is not part of a stack (in other words, it is a stack of one switch), the Stack Master LED is illuminated, and the unit number is displayed.

Back Panel

The following image shows the back panels of the Dell EMC Networking N1500 Series switches.

Figure 3-7. Dell EMC Networking N1500 Series Back Panel



Power Supplies

Dell EMC Networking N1524 and N1548

The Dell EMC Networking N1524 and N1548 Series switches have an internal 100-watt power supply. The additional redundant power supply (Dell EMC Networking RPS720) provides 180 watts of power and gives full redundancy for the switch.

Dell EMC Networking N1524P and N1548P

The Dell EMC Networking N1524P and N1548P switches have an internal 600-watt power supply feeding up to 24 PoE devices at full PoE+ power (500W). An additional modular power supply (MPS1000) provides 1000 watts and gives full power coverage for all 48 PoE devices (1500W).



NOTE: PoE power is dynamically allocated. Not all ports will require the full PoE+ power.



CAUTION: Remove the power cable from the power supplies prior to removing the power supply module itself. Power must not be connected prior to insertion in the chassis.

Ventilation System

Two internal fans cool the Dell EMC Networking N1500 Series switches.

Information Tag

The back panel includes a slide-out label panel that contains system information, such as the Service Tag, MAC address, and so on.

LED Definitions

This section describes the LEDs on the front and back panels of the switch.

Port LEDs

Each port on a Dell EMC Networking N1500 Series switch includes two LEDs. One LED is on the left side of the port, and the second LED is on the right side of the port. This section describes the LEDs on the switch ports.

100/1000/10000BASE-T Port LEDs

Each 100/1000/10000BASE-T port has two LEDs. Figure 3-8 illustrates the 100/1000/10000BASE-T port LEDs.

Figure 3-8. 100/1000/10000BASE-T Port LEDs




Table 3-7 shows the 100/1000/10000BASE-T port LED definitions.

Table 3-7. 100/1000/10000BASE-T Port Definitions

| LED | Color | Definition |
|------------------------------------|-----------------|---|
| Link/SPD LED | Off | There is no link. |
| | Solid yellow | The port is operating at 10/100 Mbps. |
| | Solid green | The port is operating at 1000 Mbps. |
| Activity/PoE LED (on PoE switches) | Off | There is no current transmit/receive activity and PoE power is off. |
| | Blinking green | The port is actively transmitting/receiving and PoE power is off. |
| | Blinking yellow | The port is actively transmitting/receiving and PoE power is on. |
| | Solid yellow | There is no current transmit/receive activity and PoE power is on. |

Stacking Port LEDs

Table 3-8. Stacking Port LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is actively transmitting/receiving. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Console Port LEDs

Table 3-9. Console Port LED Definitions

| LED | Color | Definition |
|--------------|-------------|--------------------|
| Link/SPD LED | Off | There is no link. |
| | Solid green | A link is present. |

System LEDs

The system LEDs, located on the front panel, provide information about the power supplies, thermal conditions, and diagnostics.

Table 3-10 shows the System LED definitions for the Dell EMC Networking N1500 Series switches.

Table 3-10. System LED Definitions

| LED | Color | Definition |
|--------|----------------|--|
| Status | Solid green | Normal operation. |
| | Blinking green | The switch is booting |
| | Solid red | A critical system error has occurred. |
| | Blinking red | A noncritical system error occurred (fan or power supply failure). |
| Power | Off | There is no power or the switch has experienced a power failure. |
| | Solid green | Power to the switch is on. |
| | Blinking green | The switch locator function is enabled. |

Table 3-10. System LED Definitions (Continued)

| LED | Color | Definition |
|---------------------------|-------------|--|
| RPS (on non-PoE switches) | Off | There is no redundant power supply (RPS). |
| | Solid green | Power to the RPS is on. |
| | Solid red | An RPS is detected but it is not receiving power. |
| EPS (on PoE switches) | Off | There is no external power supply (EPS). |
| | Solid green | Power to the EPS is on. |
| | Solid red | An EPS is detected but it is not receiving power. |
| Fan | Solid green | The fan is powered and is operating at the expected RPM. |
| | Solid red | A fan failure has occurred. |
| Stack Master | Off | The switch is not stack master. |
| | Solid green | The switch is master for the stack. |
| Temp | Solid green | The switch is operating below the threshold temperature. |
| | Solid red | The switch temperature exceeds the threshold of 75°C. |
| Stack No. | – | Switch ID within the stack. |

Power Consumption for PoE Switches

Table 3-11 shows power consumption data for the PoE-enabled switches.

Table 3-11. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Max Steady Current Consumption (A) | Max Steady Power (W) |
|----------------------------|---------------|----------------------------|------------------------------------|----------------------|
| Dell EMC Networking N1524P | 100V | Main PSU+EPS PSU | 8.8 | 876.0 |
| | 110V | Main PSU+EPS PSU | 7.9 | 871.0 |
| | 120V | Main PSU+EPS PSU | 7.2 | 865.0 |
| | 220V | Main PSU+EPS PSU | 3.8 | 844.0 |
| | 240V | Main PSU+EPS PSU | 3.5 | 840.0 |

Table 3-11. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Max Steady Current Consumption (A) | Max Steady Power (W) |
|----------------------------|---------------|----------------------------|------------------------------------|----------------------|
| Dell EMC Networking N1548P | 100V | Main PSU+EPS PSU | 17.1 | 1719.0 |
| | 110V | Main PSU+EPS PSU | 15.5 | 1704.0 |
| | 120V | Main PSU+EPS PSU | 14.1 | 1690.0 |
| | 220V | Main PSU+EPS PSU | 7.5 | 1642.4 |
| | 240V | Main PSU+EPS PSU | 6.9 | 1647.0 |

The PoE power budget for each interface is controlled by the switch firmware. The administrator can limit the power supplied on a port or prioritize power to some ports over others. Table 3-12 shows power budget data.

Table 3-12. Dell EMC Networking N1500 Series PoE Power Budget Limit

| Model Name | Internal Only PSU | | MPS Only | | Two PSUs | |
|----------------------------|-------------------------|---|-------------------------|---|--------------------------|---|
| | Max. PSU Output Ability | PoE+ Power Turn-on Limitation | Max. PSU Output Ability | PoE+ Power Turn-on Limitation | Max. PSUs Output Ability | PoE+ Power Turn-on Limitation |
| Dell EMC Networking N1524P | 600W | Power budget is 500W: The total PoE supplied power must not exceed 500W. | 1000W | Power budget is 900W: The total PoE supplied power must not exceed 900W. | 1600W | Power budget is 1350W: All PoE+ ports can supply maximum power. |
| Dell EMC Networking N1548P | 600W | Power budget is 500W: The total PoE supplied power must not exceed 500W. | 1000W | Power budget is 900W: The total PoE supplied power must not exceed 900W. | 1600W | Power budget is 1350W: The total PoE supplied power must not exceed 1350W. |

Dell EMC Networking N2000 Series Switch Hardware

This section contains information about device characteristics and modular hardware configurations for the Dell EMC Networking N2000 Series switches.


Front Panel

The Dell EMC Networking N2000 Series front panel includes the following features:

- Switch Ports
- Console Port
- USB Port
- Reset Button
- SFP+ Ports
- Port and System LEDs
- Stack Master LED and Stack Number Display

The following images show the front panels of the switch models in the Dell EMC Networking N2000 Series.

Figure 3-9. Dell EMC Networking N2048 Switch with 48 10/100/1000BASE-T Ports (Front Panel)



In addition to the switch ports, the front panel of each model in the Dell EMC Networking N2000 Series includes the following ports:

- RJ-45 Console port
- USB port for storage

Figure 3-10. Dell EMC Networking N2024/N2048 Close-up



The Dell EMC Networking N2024/N2048 front panel, shown in Figure 3-10, has status LEDs for over-temperature alarm (left), internal power (middle), and status (right) on the top row. The bottom row of status LEDs displays, from left to right, the Stack Master, redundant power supply (RPS) status, and fan alarm status.

The Dell EMC Networking N2024P/N2048P front panel has status LEDs for over-temperature alarm, internal power and status on the top row. The bottom row of status LEDs displays Stack Master, modular power supply (MPS), status and fan alarm status.

Switch Ports

The Dell EMC Networking N2024/N2024P front panel provides 24 Gigabit Ethernet (10/100/1000BASE-T) RJ-45 ports that support auto-negotiation for speed, flow control, and duplex. The Dell EMC Networking N2024/N2024P models support two SFP+ 10G ports. Dell EMC-qualified SFP+ transceivers are sold separately. Dell EMC Networking N2000 Series switches operate in full-duplex mode only.

The Dell EMC Networking N2048/N2048P front panel provides 48 Gigabit Ethernet (10BASE-T, 100BASE-TX, 1000BASE-T) RJ-45 ports that support auto-negotiation for speed, flow control, and duplex. The Dell EMC Networking N2048/N2048P supports two SFP+ 10G ports. Dell EMC-qualified SFP+ transceivers are sold separately.

The front-panel switch ports have the following characteristics:

- The switch automatically detects the difference between crossed and straight-through cables on RJ-45 ports and automatically chooses the MDI or MDIX configuration to match the other end.
- SFP+ ports support Dell EMC-qualified transceivers. The default behavior is to log a message and generate an SNMP trap on insertion or removal of an optic that is not qualified by Dell. The message and trap can be suppressed by using the **service unsupported-transceiver** command.
- RJ-45 ports support full-duplex mode 10/100/1000 Mbps speeds on standard Category 5 UTP cable. 1000BASE-T operation requires the use of auto-negotiation.
- SFP+ ports support SFP+ transceivers and SFP+ copper twin-ax technology operating at 10G or 1G speeds in full-duplex mode. SFP transceivers are supported in SFP+ ports and operate at 1G full-duplex. SFP transceivers require auto-negotiation to be enabled.
- The Dell EMC Networking N2024P/N2048P front-panel ports support PoE (15.4W) and PoE+ (34.2W) as well as legacy capacitive detection for pre-standard powered devices (PDs).

Console Port

The console port provides serial communication capabilities, which allows communication using RS-232 protocol. The serial port provides a direct connection to the switch and allows access to the CLI from a console terminal connected to the port through the provided serial cable (with RJ45 YOST to female DB-9 connectors).

The console port is separately configurable and can be run as an asynchronous link from 1200 BAUD to 115,200 BAUD. The Dell CLI supports changing only the speed of the console port. The defaults are 9600 BAUD, 8 data bits, no parity, 1 stop bit, and no flow control.

USB Port

The Type-A, female USB port supports a USB 2.0-compliant flash memory drive. The Dell EMC Networking N-Series switch can read or write to a flash drive with a single partition formatted as FAT-32. Use a USB flash drive to copy switch configuration files and images between the USB flash drive and

the switch. The USB flash drive may be used to move and copy configuration files and images from one switch to other switches in the network. The system does not support the deletion of files on USB flash drives.

The USB port does not support any other type of USB device.

Reset Button

The reset button is accessed through the pinhole and enables performing a hard reset on the switch. To use the reset button, insert an unbent paper clip or similar tool into the pinhole. When the switch completes the boot process after the reset, it resumes operation with the most recently saved configuration. Any changes made to the running configuration that were not saved to the startup configuration prior to the reset are lost.

Port and System LEDs

The front panel contains light emitting diodes (LEDs) that indicate the status of port links, power supplies, fans, stacking, and the overall system status. See "LED Definitions" on page 137 for more information.

Stack Master LED and Stack Number Display

When a switch within a stack is the master unit, the Stack Master LED is solid green. If the Stack Master LED is off, the stack member is not the master unit. The Stack No. panel displays the unit number for the stack member. If a switch is not part of a stack (in other words, it is a stack of one switch), the Stack Master LED is illuminated, and the unit number is displayed.

Back Panel

The following images show the back panels of the Dell EMC Networking N2000 Series switches.


Figure 3-11. Dell EMC Networking N2000 Series Back Panel



Figure 3-12. Dell EMC Networking N2024P/N2048P Back Panel

The term mini-SAS refers to the stacking port cable connections shown in Figure 3-13. See "Stacking" on page 237 for information on using the mini-SAS ports to connect switches.

Figure 3-13. Dell EMC Networking N2048 Mini-SAS Stacking Ports and Fans




Power Supplies


Dell EMC Networking N2024 and N2048

The Dell EMC Networking N2024 and N2048 Series switches have an internal 100-watt power supply. The additional redundant power supply (Dell EMC Networking RPS720) provides 180 watts of power and gives full redundancy for the switch.

Dell EMC Networking N2024P and N2048P

The Dell EMC Networking N2024P and N2048P switches have an internal 1000-watt power supply feeding up to 24 PoE devices at full PoE+ power (850W). An additional modular power supply (MPS1000) provides 1000 watts and gives full power coverage for all 48 PoE devices (1800W).

 **NOTE:** PoE power is dynamically allocated. Not all ports will require the full PoE+ power.

 **CAUTION:** Remove the power cable from the power supplies prior to removing the power supply module itself. Power must not be connected prior to insertion in the chassis.

Ventilation System

Two internal fans cool the Dell EMC Networking N2000 Series switches.

Information Tag

The back panel includes a slide-out label panel that contains system information, such as the Service Tag, MAC address, and so on.

LED Definitions

This section describes the LEDs on the front and back panels of the switch.

Port LEDs

Each port on a Dell EMC Networking N2000 Series switch includes two LEDs. One LED is on the left side of the port, and the second LED is on the right side of the port. This section describes the LEDs on the switch ports.

100/1000/10000BASE-T Port LEDs

Each 100/1000/10000BASE-T port has two LEDs. Figure 3-14 illustrates the 100/1000/10000BASE-T port LEDs.

Figure 3-14. 100/1000/10000BASE-T Port LEDs



Table 3-13 shows the 100/1000/10000BASE-T port LED definitions.

Table 3-13. 100/1000/10000BASE-T Port Definitions

| LED | Color | Definition |
|---------------------------------------|-----------------|---|
| Link/SPD LED | Off | There is no link. |
| | Solid yellow | The port is operating at 10/100 Mbps. |
| | Solid green | The port is operating at 1000 Mbps. |
| Activity LED (on non-PoE switches) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |
| Activity/PoE LED (on PoE switches) | Off | There is no current transmit/receive activity and PoE power is off. |
| | Blinking green | The port is actively transmitting/receiving and PoE power is off. |
| | Blinking yellow | The port is actively transmitting/receiving and PoE power is on. |
| | Solid yellow | There is no current transmit/receive activity and PoE power is on. |

Stacking Port LEDs

Table 3-14. Stacking Port LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is actively transmitting/receiving. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Console Port LEDs

Table 3-15. Console Port LED Definitions

| LED | Color | Definition |
|--------------|-------------|--------------------|
| Link/SPD LED | Off | There is no link. |
| | Solid green | A link is present. |

System LEDs

The system LEDs, located on the front panel, provide information about the power supplies, thermal conditions, and diagnostics.

Table 3-16 shows the System LED definitions for the Dell EMC Networking N2000 Series switches.

Table 3-16. System LED Definitions

| LED | Color | Definition |
|---------------------------|----------------|--|
| Status | Solid green | Normal operation. |
| | Blinking green | The switch is booting |
| | Solid red | A critical system error has occurred. |
| | Blinking red | A noncritical system error occurred (fan or power supply failure). |
| Power | Off | There is no power or the switch has experienced a power failure. |
| | Solid green | Power to the switch is on. |
| | Blinking green | The switch locator function is enabled. |
| RPS (on non-PoE switches) | Off | There is no redundant power supply (RPS). |
| | Solid green | Power to the RPS is on. |
| | Solid red | An RPS is detected but it is not receiving power. |
| EPS (on PoE switches) | Off | There is no external power supply (EPS). |
| | Solid green | Power to the EPS is on. |
| | Solid red | An EPS is detected but it is not receiving power. |

Table 3-16. System LED Definitions (Continued)

| LED | Color | Definition |
|--------------|-------------|--|
| Fan | Solid green | The fan is powered and is operating at the expected RPM. |
| | Solid red | A fan failure has occurred. |
| Stack Master | Off | The switch is not stack master. |
| | Solid green | The switch is master for the stack. |
| Temp | Solid green | The switch is operating below the threshold temperature. |
| | Solid red | The switch temperature exceeds the threshold of 75°C. |
| Stack No. | – | Switch ID within the stack. |

Power Consumption for PoE Switches

Table 3-17 shows power consumption data for the PoE-enabled switches.

Table 3-17. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Max Steady Current Consumption (A) | Max Steady Power (W) |
|----------------------------|---------------|----------------------------|------------------------------------|----------------------|
| Dell EMC Networking N2024P | 100V | Main PSU+EPS PSU | 8.9 | 890.0 |
| | 110V | Main PSU+EPS PSU | 8.3 | 913.0 |
| | 120V | Main PSU+EPS PSU | 7.6 | 912.0 |
| | 220V | Main PSU+EPS PSU | 4.0 | 880.0 |
| | 240V | Main PSU+EPS PSU | 3.6 | 873.6 |
| Dell EMC Networking N2048P | 100V | Main PSU+EPS PSU | 17.8 | 1780.0 |
| | 110V | Main PSU+EPS PSU | 15.8 | 1740.2 |
| | 120V | Main PSU+EPS PSU | 14.5 | 1740.0 |
| | 220V | Main PSU+EPS PSU | 7.7 | 1687.4 |
| | 240V | Main PSU+EPS PSU | 7.1 | 1704.0 |

The PoE power budget for each interface is controlled by the switch firmware. The administrator can limit the power supplied on a port or prioritize power to some ports over others. Table 3-18 shows power budget data.

Table 3-18. Dell EMC Networking N2000 Series PoE Power Budget Limit

| Model Name | One PSU | | Two PSUs | |
|----------------------------|-------------------------|---|--------------------------|--|
| | Max. PSU Output Ability | PoE+ Power Turn-on Limitation | Max. PSUs Output Ability | PoE+ Power Turn-on Limitation |
| Dell EMC Networking N2024P | 1000W | Power budget is 850W: The total PoE supplied power must not exceed 850W. | 2000W | Power budget is 1700W: All PoE+ ports can supply maximum power. |
| Dell EMC Networking N2048P | 1000W | Power budget is 850W: The total PoE supplied power must not exceed 850W. | 2000W | Power budget is 1700W: All PoE+ ports can supply maximum power. |


Dell EMC Networking N2100-ON Series Switch Hardware

This section contains information about device characteristics and modular hardware configurations for the Dell EMC Networking N2128PX-ON switch.

Front Panel

All N2128PX-ON PoE models are 1U, rack-mountable switches. The Dell EMC Networking N2128PX-ON front panel provides 24 10/100/1000BASE-T Ethernet RJ-45 ports and four 2.5G NBASE-T Ethernet RJ-45 ports that support auto-negotiation for speed, flow control, and duplex. NBASE-T interfaces require auto-negotiation to be enabled. They will not operate correctly in fixed speed mode. The 2.5G NBASE-T ports support PoE 60W capability. The N2128PX switch front panel ports operate in full duplex mode only. The N2128PX models support two SFP+ 10G ports. The N2128PX-ON provides one RJ-45 console port for local management and a Type-A USB port for storage.

Figure 3-15. Dell EMC Networking N2128PX-ON Switch (Front Panel)



The Dell N2128PX-ON switch is capable of loading an OS via the ONIE boot loader, therefore the port numbering on the front panel labels is consecutive, per the ONIE requirements, regardless of port type.

For the N2128PX-ON switch, ports 1-24 are 1G interfaces, ports 25-28 are 2.5GBASE-T interfaces, ports 29-30 are 10G interfaces, and rear panel ports 31-32 are HiGig stacking interfaces. NBASE-T interfaces require auto-negotiation to be enabled. NBASE-T interfaces require auto-negotiation and will not operate correctly in fixed speed mode.

To remain consistent with prior N-Series devices, CLI and GUI port references will be non-consecutive when the port type changes. Ports labeled 1-28 on the front panel will be referred to in the UI as Gi1/0/X (where X = 1 to 28), ports labeled 29-30 on the front panel will be referred to in the UI as Te1/0/Y (where Y= 1 to 2) and ports labeled 31-32 on the rear panel will be referred to as Tw1/1/W (where W=1 to 2).

Console Port

The console port provides serial communication capabilities, which allows communication using RS-232 protocol. The serial port provides a direct connection to the switch and allows access to the CLI from a console terminal connected to the port through the provided serial cable (with RJ-45 YOST to female DB-9 connectors).

The console port is separately configurable and can be run as an asynchronous link from 1200 BAUD to 115,200 BAUD. The Dell EMC CLI supports changing only the speed of the console port. The defaults are 115,200 BAUD, 8 data bits, no parity, 1 stop bit, and no flow control.

USB Port

The Type-A, female USB port supports a USB 2.0-compliant flash memory drive. The Dell EMC Networking N-Series switch can read or write to a flash drive with a single partition formatted as FAT-32. Use a USB flash drive to copy switch configuration files and images between the USB flash drive and the switch. The USB flash drive may be used to move and copy configuration files and images from one switch to other switches in the network. The system does not support the deletion of files on USB flash drives.

The USB port does not support any other type of USB device.

Reset Button

The reset button is accessed through the pinhole and enables performing a hard reset on the switch. To use the reset button, insert an unbent paper clip or similar tool into the pinhole. When the switch completes the boot process after the reset, it resumes operation with the most recently saved configuration. Any changes made to the running configuration that were not saved to the startup configuration prior to the reset are lost.

Port and System LEDs

The front panel contains light emitting diodes (LEDs) that indicate the status of port links, power supplies, fans, stacking, and the overall system status. See "LED Definitions" on page 159 for more information.

Stack Master LED and Stack Number Display

When a switch within a stack is the master unit, the Stack Master LED is solid green. If the Stack Master LED is off, the stack member is not the master unit. The Stack No. panel displays the unit number for the stack member. If a switch is not part of a stack (in other words, it is a stack of one switch), the Stack Master LED is illuminated, and the unit number is displayed.

Back Panel

The N2128PX-ON has two 21G stacking ports in the rear that accept mini-SAS connectors. It also has a 16-pin connection for a modular power supply (Dell MPS1000) supporting an additional 1000W of power.

Power Supply

The Dell EMC Networking N2128PX-ON switch has an internal 715-watt power supply feeding up to 16 PoE devices at full PoE+ power (500W).

Ventilation System

Two internal fans cool the Dell EMC Networking N2128PX-ON Series switches.

Information Tag

The back panel includes a slide-out label panel that contains system information, such as the Service Tag, MAC address, and so on.

LED Definitions

This section describes the LEDs on the front and back panels of the switch.

Port LEDs

Each port on a Dell EMC Networking N2100-ON Series switch includes two LEDs. One LED is on the left side of the port, and the second LED is on the right side of the port. This section describes the LEDs on the switch ports.

Each 100/1000/10000BASE-T port has two LEDs. Figure 3-16 illustrates the 100/1000/10000BASE-T port LEDs.

Figure 3-16. 100/1000/10000BASE-T Port LEDs




Table 3-19 shows the 100/1000/10000BASE-T port LED definitions.

Table 3-19. 100/1000/10000BASE-T Port LED Definitions

| LED | Color | Definition |
|------------------------------------|-----------------|---|
| Link/SPD LED | Off | There is no link. |
| | Solid yellow | The port is operating at 10/100 Mbps. |
| | Solid green | The port is operating at 1000 Mbps. |
| Activity/PoE LED (on PoE switches) | Off | There is no current transmit/receive activity and PoE power is off. |
| | Blinking green | The port is actively transmitting/receiving and PoE power is off. |
| | Blinking yellow | The port is actively transmitting/receiving and PoE power is on. |
| | Solid yellow | There is no current transmit/receive activity and PoE power is on. |

Table 3-20. 2500BASE-T Port LED Definitions

| LED | Color | Definition |
|--|-----------------|--|
| Link/SPD LED (Left bi-color LED) | Off | There is no link. |
| | Solid green | The port is operating at 2.5 Gbps. |
| | Solid amber | The port is operating at 100 Mbps or 1 Gbps. |
| Activity/PoE LED (Right bi-color LED) | Off | There is no current transmit/receive activity, and PoE power is off. |
| | Blinking green | The port is actively transmitting/receiving, and PoE power is off. |
| | Blinking yellow | The port is actively transmitting/receiving, and PoE power is on. |
| | Solid yellow | There is no current transmit/receive activity, and PoE power is on. |

Table 3-21. SFP+ Port LED Definitions

| LED | Color | Definition |
|--|----------------|--|
| Link/SPD LED (Left bi-color LED) | Off | There is no link. |
| | Solid green | The port is operating at 10 Gbps. |
| | Solid amber | The port is operating at 1 Gbps. |
| Activity LED (Right single-color LED) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Table 3-22. QSFP Port LED Definitions

| LED | Color | Definition |
|--|----------------|--|
| Link/SPD LED (Left single-color LED) | Off | There is no link. |
| | Solid green | The port is operating at 40 Gbps. |
| Activity LED (Right single-color LED) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Stacking Port LEDs

Table 3-23. Stacking Port LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is actively transmitting/receiving. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Console Port LEDs

Table 3-24. Console Port LED Definitions

| LED | Color | Definition |
|--------------|-------------|--------------------|
| Link/SPD LED | Off | There is no link. |
| | Solid green | A link is present. |

System LEDs

The system LEDs, located on the front panel, provide information about the power supplies, thermal conditions, and diagnostics.

Table 3-25 shows the System LED definitions for the Dell EMC Networking N2128PX-ON switches.

Table 3-25. System LED Definitions

| LED | Color | Definition |
|--------|----------------|--|
| Status | Solid green | Normal operation. |
| | Blinking green | The switch is booting |
| | Solid red | A critical system error has occurred. |
| | Blinking red | A noncritical system error occurred (fan or power supply failure). |
| Power | Off | There is no power or the switch has experienced a power failure. |
| | Solid green | Power to the switch is on. |
| | Blinking green | The switch locator function is enabled. |

Table 3-25. System LED Definitions (Continued)

| LED | Color | Definition |
|-----------------------|-------------|--|
| EPS (on PoE switches) | Off | There is no external power supply (EPS). |
| | Solid green | Power to the EPS is on. |
| | Solid red | An EPS is detected but it is not receiving power. |
| Fan | Solid green | The fan is powered and is operating at the expected RPM. |
| | Solid red | A fan failure has occurred. |
| Stack Master | Off | The switch is not stack master. |
| | Solid green | The switch is master for the stack. |
| Temp | Solid green | The switch is operating below the threshold temperature. |
| | Solid red | The switch temperature exceeds the threshold of 75°C. |
| Stack No. | – | Switch ID within the stack. |

Power Consumption for PoE Switches

Table 3-26 shows power consumption data for the PoE-enabled N2128PX-ON switch when the power budget is 800W for the main power supply.

Table 3-26. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Max Steady Power (W) |
|--------------------------------|---------------|----------------------------|--|----------------------|
| Dell EMC Networking N2128PX-ON | 100V/60Hz | Main PSU | 9.73A | 965.5W |
| | 110V/60Hz | Main PSU | 8.75A | 960.4W |
| | 120V/60Hz | Main PSU | 8.03A | 958.3 |
| | 220V/50Hz | Main PSU | 4.33A | 931W |
| | 240V/50Hz | Main PSU | 3.97A | 928.7W |

Table 3-27 shows power consumption data for the PoE-enabled N2128PX-ON switch when the power budget is 800W for the MPS.

Table 3-27. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Max Steady Power (W) |
|--------------------------------|----------------------|-----------------------------------|---|-----------------------------|
| Dell EMC Networking N2128PX-ON | 100V/60Hz | MPS | 9.92A | 986.5W |
| | 110V/60Hz | MPS | 8.93A | 975.7W |
| | 120V/60Hz | MPS | 8.01A | 955.4W |
| | 220V/50Hz | MPS | 4.44A | 945.4W |
| | 240V/50Hz | MPS | 4.08A | 951.4W |

Table 3-28 shows power consumption data for the PoE-enabled N2128PX-ON switch when the power budget is 1600W for the main power supply and the MPS.

Table 3-28. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Max Steady Power (W) |
|--------------------------------|----------------------|-----------------------------------|---|-----------------------------|
| Dell EMC Networking N2128PX-ON | 100V/60Hz | Main PSU + MPS | 11.83A | 1175W |
| | 110V/60Hz | Main PSU + MPS | 10.71A | 1169W |
| | 120V/60Hz | Main PSU + MPS | 9.84A | 1168.9W |
| | 220V/50Hz | Main PSU + MPS | 5.4A | 1138.4W |
| | 240V/50Hz | Main PSU + MPS | 5.93A | 1141W |

The PoE power budget for each interface is controlled by the switch firmware. The administrator can limit the power supplied on a port or prioritize power to some ports over others. Table 3-29 shows power budget data.

Table 3-29. Dell EMC Networking N2100-ON Series PoE Power Budget Limit

| Model Name | One PSU | | Two PSUs | |
|--------------------------------|-------------------------|---|--------------------------|--|
| | Max. PSU Output Ability | PoE+ Power Turn-on Limitation | Max. PSUs Output Ability | PoE+ Power Turn-on Limitation |
| Dell EMC Networking N2128PX-ON | 1000W | Power budget is 800W: The total PoE supplied power must not exceed 800W. | 2000W | Power budget is 1600W: All PoE+ ports can supply maximum power. |

Dell EMC Networking N3000/N3000E-ON Series Switch Hardware

This section contains information about device characteristics and modular hardware configurations for the Dell EMC Networking N3000/N3000E-ON Series switches.


Front Panel

The Dell EMC Networking N3000/N3000E-ON Series front panel includes the following features:

- Switch Ports
- Console Port
- Out-of-Band Management Port
- USB Port
- SFP+ Ports
- Reset Button
- Port and System LEDs
- Stack Master LED and Stack Number Display


The following images show the front panels of the switch models in the Dell EMC Networking N3000 Series.

Figure 3-17. Dell EMC Networking N3024F/N3024EF-ON with 24 10/100/1000BASE-T Ports (Front Panel)




The Dell EMC Networking N3000/N3000E-ON Series switch includes two combo ports. The combo ports are SFP on the Dell EMC Networking N3000 Series and 1000BASE-T on the Dell EMC Networking N3024F/N3024EF-ON switch. The combo ports are set to prefer fiber. If using the 1000BASE-T port, remove any transceivers from the SFP port.

Figure 3-18. Dell EMC Networking N3048/N3048ET-ON with 48 10/100/1000BASE-T Ports (Front Panel)



The additional ports are on the right side of the front panel, as shown in Figure 3-18 and Figure 3-19.

Figure 3-19. Additional Dell EMC Networking N3000/N3000E-ON Series Ports



The Dell EMC Networking N3048/N3048P/N3048ET-ON/N3048EP-ON front panel provides 48 Gigabit Ethernet (10BASE-T, 100BASE-TX, 1000BASE-T) RJ-45 ports that support auto-negotiation for speed, flow control, and duplex. The Dell EMC Networking N3048/N3048P/N3048ET-ON/N3048P-ON support two SFP+ 10G ports. Dell EMC-qualified SFP+ transceivers are sold separately.

The front-panel switch ports have the following characteristics:

- The switch automatically detects the difference between crossed and straight-through cables on RJ-45 ports and automatically chooses the MDI or MDIX configuration to match the other end.
- SFP+ ports support Dell EMC-qualified transceivers. The default behavior is to log a message and generate an SNMP trap on insertion or removal of an optic that is not qualified by Dell. The message and trap can be suppressed by using the **service unsupported-transceiver** command.
- RJ-45 ports support full-duplex mode 10/100/1000 Mbps speeds on standard Category 5 UTP cable.
- SFP ports support 1000BASE-X and 100BASE-X (100BASE-FX/100BASE-LX/100BASE-SX) transceivers. SFP ports with 1000BASE-X transceivers require auto-negotiation to be enabled but also allow configuration of forced speeds with auto-negotiation disabled. However, the SFP ports do not support auto-negotiation for 100BASE-X transceivers. When the switch detects a 100BASE-X transceiver, it resets the port to use 4B/5B NRZI line coding from the default 8B/10B NRZ coding. This causes the link to flap momentarily. The flap may be observed at the link partner upon insertion of an SFP transceiver and on switch reboot. Administrators should ensure that the link partner is set to accept a single link flap on 100BASE-X interfaces.
- SFP+ ports support SFP+ transceivers and SFP+ copper twin-ax technology operating at 10G/1G speeds in full-duplex mode. SFP transceivers are supported in SFP+ ports and operate at 1G full-duplex. SFP transceivers in an SFP+ port require auto-negotiation to be enabled per the IEEE 802.3 standard but may be configured to use a forced speed with auto-negotiation disabled.
- The Dell EMC Networking N3024P/N3048P/N3048ET-ON/N3048EP-ON front-panel ports support PoE (15.4W) and PoE+ (34.2W) as well as legacy capacitive detection for pre-standard PDs.

- Additionally, ports 1–12 support PoE 60W power when configured in high-power mode.

Combo Ports

Combo ports automatically select the active media and always choose fiber (SFP) media if both copper and fiber are active. Copper combo ports do not support 10 Mbps forced mode. Auto-negotiation is not supported for 100BASE-X (FX/SX/LX) transceivers. If using the 1000BASE-T port, remove any transceivers from the SFP port.

Console Port

The console port provides serial communication capabilities, which allows communication using RS-232 protocol. The serial port provides a direct connection to the switch and allows access to the CLI from a console terminal connected to the port through the provided serial cable (with RJ45 YOST to female DB-9 connectors).

The console port is separately configurable and can be run as an asynchronous link from 1200 BAUD to 115,200 BAUD.

The Dell CLI only supports changing the speed.

The defaults are 9600 BAUD, 8 data bits, no parity, 1 stop bit, and no flow control. The N3048ET-ON/N3048EP-ON defaults to 115,200 BAUD.

Out-of-Band Management Port

The Out-of-Band (OOB) management port is a 10/100/1000BASE-T Ethernet port connected directly to the switch CPU and dedicated to switch management. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to or from the operational network. In addition, ACLs (including management ACLS), do not operate on the out-of-band port. Connect the out-of-band port only to a physically secure network.

USB Port

The Type-A, female USB port supports a USB 2.0-compliant flash memory drive. The Dell EMC Networking N-Series switch can read or write to a flash drive with a single partition formatted as FAT-32. Use a USB flash drive to copy switch configuration files and images between the USB flash drive and

the switch. It is also possible to use the USB flash drive to move and copy configuration files and images from one switch to other switches in the network. The system does not support the deletion of files on attached USB flash drives.

The USB port does not support any other type of USB device.

Reset Button

The reset button is accessed through the pinhole and enables performing a hard reset on the switch. To use the reset button, insert an unbent paper clip or similar tool into the pinhole. When the switch completes the boot process after the reset, it resumes operation with the most recently saved configuration. Any changes made to the running configuration that were not saved to the startup configuration prior to the reset are lost.

Port and System LEDs

The front panel contains light emitting diodes (LEDs) that indicate the status of port links, power supplies, fans, stacking, and the overall system status.

Stack Master LED and Stack Number Display

When a switch within a stack is the master unit, the Stack Master LED is solid green. If the Stack Master LED is off, the stack member is not the master unit. The Stack No. panel displays the unit number for the stack member. If a switch is not part of a stack (in other words, it is a stack of one switch), the Stack Master LED is illuminated and the unit number is displayed.

Back Panel

The following images show the back panels of the Dell EMC Networking N3000 Series switches.

Figure 3-20. Dell EMC Networking N3000/N3000E-ON Series Back Panel




Figure 3-21. Dell EMC Networking N3024P/N3048P/N3024EP-ON/N3048EP-ON Back Panel



Figure 3-22. Dell EMC Networking N3048/N3048ET-ON Mini-SAS Stacking Ports Close-up



The term mini-SAS refers to the stacking port cable connections shown in Figure 3-22. See "Stacking" on page 237 for information on using the mini-SAS ports to connect switches.

Expansion Slots for Plug-in Modules

One expansion slot is located on the back of the Dell EMC Networking N3000 Series models and can support the following modules:

- 10GBASE-T module
- SFP+ module

Each plug-in module has two ports. The plug-in modules include hot-swap support, so a switch reboot is not needed after a new module is installed. Issue a **no slot** command after removing the original module and prior to inserting a new type of module. If the module is not recognized, issue the **no slot** command, then remove and re-insert the module.

Power Supplies

Dell EMC Networking N3024, N3024F and N3048

Dell EMC Networking N3024, N3024F and N3048 switches support two 200-watt Field Replaceable Unit (FRU) power supplies which give full power redundancy for the switch. The Dell EMC Networking N3024, N3024F, and N3048 switches offer the V-lock feature for users desiring the need to eliminate accidental power disconnection. The V-lock receptacle on the Power Supply Unit (PSU) allows for the use of a power cord that has the V-lock feature to create an integral secure locking connection.

Dell EMC Networking N3024P, N3048P, and N3048EP-ON

Dell EMC Networking N3024P, N3048P, and N3048EP-ON switches support one or two 1100-watt FRU power supplies. The Dell EMC Networking N3024P switch is supplied with a single 715-watt power supply (the default configuration) and supports an additional 1100-watt supply. For the Dell EMC Networking N3048P, and N3048EP-ON switches, a single 1100-watt power supply is supplied and another 1100 watt power supply can be added.

A single 1100-watt power supply can feed up to 24 PoE devices at full PoE+ power (950W). Dual-equipped switches will feed up to 48 PoE devices at full PoE+ power (1800W), as well as provide power supply redundancy.



NOTE: PoE power is dynamically allocated by default. Not all ports will require the full PoE+ power.

⚠ CAUTION: Remove the power cable from the power supplies prior to removing the power supply module itself. Power must not be connected prior to insertion in the chassis.

Ventilation System

Two fans cool the Dell EMC Networking N3000 Series switches. The Dell EMC Networking N3000 Series switches additionally have a fan in each internal power supply. The Dell EMC Networking N3000 Series fan is field-replaceable.

Information Tag

The back panel includes a slide-out label panel that contains system information, such as the Service Tag, MAC address, and other information.

LED Definitions

This section describes the LEDs on the front and back panels of the switch.

Port LEDs

Each port on a Dell EMC Networking N3000 Series switch includes two LEDs. One LED is on the left side of the port, and the second LED is on the right side of the port. This section describes the LEDs on the switch ports.

100/1000/10000BASE-T Port LEDs

Each 100/1000/10000BASE-T port has two LEDs. Figure 3-23 illustrates the 100/1000/10000BASE-T port LEDs.

Figure 3-23. 100/1000/10000BASE-T Port LEDs



Table 3-30 shows the 100/1000/10000BASE-T port LED definitions.

Table 3-30. 100/1000/10000BASE-T Port Definitions

| LED | Color | Definition |
|---------------------------------------|-----------------|---|
| Link/SPD LED | Off | There is no link. |
| | Solid yellow | The port is operating at 10/100 Mbps. |
| | Solid green | The port is operating at 1000 Mbps. |
| Activity LED (on non-PoE switches) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |
| Activity/PoE LED (on PoE switches) | Off | There is no current transmit/receive activity and PoE power is off. |
| | Blinking green | The port is actively transmitting/receiving and PoE power is off. |
| | Blinking yellow | The port is actively transmitting/receiving and PoE power is on. |
| | Solid yellow | There is no current transmit/receive activity and PoE power is on. |

Module Bay LEDs

The following tables describe the purpose of each of the module bay LEDs when SFP+ and 10GBASE-T modules are used.

Table 3-31. SFP+ Module LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link/SPD LED | Off | There is no link. |
| | Solid green | The port is operating at 10 Gbps. |
| | Solid amber | The port is operating at 1000 Mbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Table 3-32. 10GBASE-T Module LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link/SPD LED | Off | There is no link. |
| | Solid green | The port is operating at 10 Gbps. |
| | Solid amber | The port is operating at 100/1000 Mbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Stacking Port LEDs**Table 3-33. Stacking Port LED Definitions**

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is actively transmitting/receiving. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Out-of-Band Port LEDs**Table 3-34. OOB Port LED Definitions**

| LED | Color | Definition |
|--------------|----------------|---|
| Link/SPD LED | Off | There is no link. |
| | Solid green | The port is actively transmitting/receiving at 1000 Mbps. |
| | Solid amber | The port is actively transmitting/receiving at 10/100 Mbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Console Port LEDs

Table 3-35. Console Port LED Definitions

| LED | Color | Definition |
|--------------|-------------|--------------------|
| Link/SPD LED | Off | There is no link. |
| | Solid green | A link is present. |

System LEDs

The system LEDs, located on the front panel, provide information about the power supplies, thermal conditions, and diagnostics.

Table 3-36 shows the System LED definitions for the Dell EMC Networking N3000 Series switches.

Table 3-36. System LED Definitions

| LED | Color | Definition |
|--------|----------------|--|
| Status | Solid green | Normal operation. |
| | Blinking green | The switch is booting |
| | Solid red | A critical system error has occurred. |
| | Blinking red | A noncritical system error occurred (fan or power supply failure). |

Table 3-36. System LED Definitions

| LED | Color | Definition |
|---------------------|----------------|--|
| Power 1, Power 2 | Off | There is no power or the switch has experienced a power failure. |
| | Solid green | Power to the switch is on. |
| | Blinking green | The switch locator function is enabled. |
| Fan | Solid green | The fan is powered and is operating at the expected RPM. |
| | Solid red | A fan failure has occurred. |
| Stack Master | Off | The switch is in stand-alone mode. |
| | Solid green | The switch is master for the stack. |
| Temp | Solid green | The switch is operating below the threshold temperature. |
| | Solid red | The switch temperature exceeds the threshold of 75°C. |
| Stack No. | – | Switch ID within the stack. |

Power Consumption for PoE Switches

Table 3-37 shows power consumption data for the PoE-enabled switches.

Table 3-37. Dell EMC Networking N3000 Series Power Consumption

| Model | Input Voltage | Power Supply Configuration | Max Steady Current Consumption (A) | Max Steady Power (W) |
|---------------------------------------|---------------|----------------------------|------------------------------------|----------------------|
| Dell EMC Networking N3024P | 100V | PSU1+PSU2 | 13.1 | 1310.0 |
| | 110V | PSU1+PSU2 | 11.7 | 1287.0 |
| | 120V | PSU1+PSU2 | 10.6 | 1272.0 |
| | 220V | PSU1+PSU2 | 5.6 | 1232.0 |
| | 240V | PSU1+PSU2 | 5.2 | 1240.8 |
| Dell EMC Networking N3048P/N3048EP-ON | 100V | PSU1+PSU2 | 21.8 | 2180.0 |
| | 110V | PSU1+PSU2 | 19.5 | 2145.0 |
| | 120V | PSU1+PSU2 | 17.8 | 2136.0 |
| | 220V | PSU1+PSU2 | 9.31 | 2048.2 |
| | 240V | PSU1+PSU2 | 8.6 | 2064.0 |

The PoE power budget for each interface is controlled by the switch firmware. The administrator can limit the power supplied on a port or prioritize power to some ports over others. Table 3-38 shows the power budget data.

Table 3-38. Dell EMC Networking N3000 Series PoE Power Budget Limit

| Model Name | One PSU | | Two PSUs | |
|--|--------------------------------|---|---------------------------------|--|
| | Max. PSU Output Ability | PoE+ Power Turn-on Limitation | Max. PSUs Output Ability | PoE+ Power Turn-on Limitation |
| Dell EMC Networking N3024P | 715W | Power budget is 550W: The total PoE supplied power must not exceed 550W. | 715W | Power budget is 1100W: All PoE+ ports can supply maximum power. |
| Dell EMC Networking N3048P/N3048 EP-ON | 1100W | Power budget is 950W: The total PoE supplied power must not exceed 950W. | 2200W | Power budget is 1900W: All PoE+ ports can supply maximum power. |


Dell EMC Networking N3100-ON Series Switch Hardware

Front Panel

All N3132PX-ON models are 1U, rack-mountable switches. The N3132PX-ON front panel provides twenty-four 10/100/1000BASE-T Ethernet RJ-45 ports and eight 5G NBASE-T Ethernet RJ-45 ports that support auto-negotiation for speed, flow control, and duplex. NBASE-T interfaces require auto-negotiation to be enabled. They will not operate correctly in fixed speed mode. The N3132PX-ON switch front panel ports operate in full duplex mode only. N3132PX-ON front panel copper ports support PoE 60W capability. The N3132PX-ON models support four SFP+ 10G ports. The SFP+ 10G ports support SFP+ transceivers or SFP transceivers, but not both types simultaneously. Use either all SFP+ transceivers or all SFP transceivers. One RJ-45 console port provides serial communication capabilities, which allows communication using RS-232 protocol. The front panel has a Type-A USB port for storage.

The following image shows the front panel of the Dell EMC Networking N3132PX-ON switch.

Figure 3-24. Dell EMC Networking N3132PX-ON with 24 10/100/1000BASE-T Ports (Front Panel)



Console Port

The console port provides serial communication capabilities, which allows communication using RS-232 protocol. The serial port provides a direct connection to the switch and allows access to the CLI from a console terminal connected to the port through the provided serial cable (with RJ45 YOST to female DB-9 connectors).

The console port is separately configurable and can be run as an asynchronous link from 1200 BAUD to 115,200 BAUD. The Dell EMC CLI only supports changing the speed. The defaults are 115,200 BAUD, 8 data bits, no parity, 1 stop bit, and no flow control.

USB Port

The Type-A female USB port supports a USB 2.0-compliant flash memory drive. The Dell EMC Networking N-Series switch can read or write to a flash drive with a single partition formatted as FAT-32. Use a USB flash drive to copy switch configuration files and images between the USB flash drive and the switch. The USB flash drive may be used to move and copy configuration files and images from one switch to other switches in the network. The system does not support the deletion of files on USB flash drives.

The USB port does not support any other type of USB device.

Reset Button

The reset button is accessed through the pinhole and enables performing a hard reset on the switch. To use the reset button, insert an unbent paper clip or similar tool into the pinhole. When the switch completes the boot process after the reset, it resumes operation with the most recently saved configuration. Any changes made to the running configuration that were not saved to the startup configuration prior to the reset are lost.

Port and System LEDs

The front panel contains light emitting diodes (LEDs) that indicate the status of port links, power supplies, fans, stacking, and the overall system status. See "LED Definitions" on page 168 for more information.

Stack Master LED and Stack Number Display

When a switch within a stack is the master unit, the Stack Master LED is solid green. If the Stack Master LED is off, the stack member is not the master unit. The Stack No. panel displays the unit number for the stack member. If a switch is not part of a stack (in other words, it is a stack of one switch), the Stack Master LED is illuminated, and the unit number is displayed.

Back Panel

The N3132PX-ON rear panel has an expansion slot which accepts a 2 x 40G QSFP+ module or a 2 x 21G stacking module. The QSFP+ module supports SR4, LR4, and copper CR4 technologies in 40G mode only.

Power Supply

The N3132PX-ON rear panel has one 715W field-replaceable power supply. A redundant power supply may be added in the available slot. Optional 715W and 1100W power supplies are available.



CAUTION: Remove the power cable from the power supplies prior to removing the power supply module itself. Power must not be connected prior to insertion in the chassis.

Ventilation System

Two internal fans in a single field replaceable unit (FRU) cool the Dell EMC Networking N3132PX-ON Series switches.

Information Tag

The back panel includes a slide-out label panel that contains system information, such as the Service Tag, MAC address, and so on.

LED Definitions

Each port on a N3132PX-ON Series switch includes two LEDs. One LED is on the left side of the port, and the second LED is on the right side of the port. This section describes the LEDs on the switch ports.

Port LEDs

Each 100/1000/10000BASE-T port has two LEDs. Figure 3-25 illustrates the 100/1000/10000BASE-T port LEDs.

Figure 3-25. 100/1000/10000BASE-T Port LEDs




Table 3-39, Table 3-40, and Table 3-41 show the port LED definitions.

Table 3-39. 100/1000/10000BASE-T Port LED Definitions

| LED | Color | Definition |
|---------------------------------------|-----------------|---|
| Link/SPD LED | Off | There is no link. |
| | Solid yellow | The port is operating at 10/100 Mbps. |
| | Solid green | The port is operating at 1000 Mbps. |
| Activity LED (on non-PoE switches) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |
| Activity/PoE LED (on PoE switches) | Off | There is no current transmit/receive activity and PoE power is off. |
| | Blinking green | The port is actively transmitting/receiving and PoE power is off. |
| | Blinking yellow | The port is actively transmitting/receiving and PoE power is on. |
| | Solid yellow | There is no current transmit/receive activity and PoE power is on. |

Table 3-40. 5000BASE-T Port LED Definitions

| LED | Color | Definition |
|--|-----------------|---|
| Link/SPD LED (Left bi-color LED) | Off | There is no link. |
| | Solid green | The port is operating at 2.5/5 Gbps. |
| | Solid amber | The port is operating at 100 Mbps or 1 Gbps. |
| Activity/PoE LED (Right bi-color LED) | Off | There is no current transmit/receive activity and PoE power is off. |
| | Blinking green | The port is actively transmitting/receiving and PoE power is off. |
| | Blinking yellow | The port is actively transmitting/receiving and PoE power is on. |
| | Solid yellow | There is no current transmit/receive activity and PoE power is on. |

Table 3-41. SFP+ Port LED Definitions

| LED | Color | Definition |
|--|----------------|--|
| Link/SPD LED (Left bi-color LED) | Off | There is no link. |
| | Solid green | The port is operating at 10 Gbps. |
| | Solid amber | The port is operating at 1 Gbps. |
| Activity LED (Right single-color LED) | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Module Bay LEDs

The following tables describe the purpose of each of the module bay LEDs when a QSFP or a Stacking module is installed.

Table 3-42. QSFP Module LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link/SPD LED | Off | There is no link. |
| | Solid green | The port is operating at 40 Gbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Table 3-43. Stacking Module Port LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port detects a link. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Out-of-Band Port and Console Port LEDs

Table 3-44 shows the OOB port LED definitions, and Table 3-45 shows the console port LED definitions.

Table 3-44. OOB Port LED Definitions

| LED | Color | Definition |
|--------------|----------------|---|
| Link/SPD LED | Off | There is no link. |
| | Solid green | The port is actively transmitting/receiving at 1000 Mbps. |
| | Solid amber | The port is actively transmitting/receiving at 10/100 Mbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Table 3-45. Console Port LED Definitions

| LED | Color | Definition |
|--------------|-------------|--------------------|
| Link/SPD LED | Off | There is no link. |
| | Solid green | A link is present. |

System LEDs

The system LEDs, located on the front panel, provide information about the power supplies, thermal conditions, and diagnostics.

Table 3-46 shows the System LED definitions for the Dell EMC Networking N3132PX-ON Series switches.

Table 3-46. System LED Definitions

| LED | Color | Definition |
|---------------------|----------------|--|
| Status | Solid green | Normal operation. |
| | Blinking green | The switch is booting |
| | Solid red | A critical system error has occurred. |
| | Blinking red | A noncritical system error occurred (fan or power supply failure). |
| Power 1, Power 2 | Off | There is no power or the switch has experienced a power failure. |
| | Solid green | Power to the switch is on. |
| | Blinking green | The switch locator function is enabled. |
| Fan | Solid green | The fan is powered and is operating at the expected RPM. |
| | Solid red | A fan failure has occurred. |
| Stack Master | Off | The switch is in stand-alone mode. |
| | Solid green | The switch is master for the stack. |
| Temp | Solid green | The switch is operating below the threshold temperature. |
| | Solid red | The switch temperature exceeds the threshold of 75°C. |
| Stack No. | – | Switch ID within the stack. |

Power Consumption for PoE Switches

Table 3-47 shows power consumption data for the PoE-enabled N3132PX-ON switch when the power budget is 500W for one 715W power supply.

Table 3-47. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Max Steady Power (W) |
|--------------------------------|---------------|----------------------------|--|----------------------|
| Dell EMC Networking N3132PX-ON | 100V/60Hz | One 715W | 6.47A | 647.3W |
| | 110V/60Hz | One 715W | 5.79A | 636.1W |
| | 120V/60Hz | One 715W | 5.12A | 611.9W |
| | 220V/50Hz | One 715W | 2.85A | 621.7W |
| | 240V/50Hz | One 715W | 2.62A | 618.7W |

Table 3-48 shows power consumption data for the PoE-enabled N3132PX-ON switch when the power budget is 1200W for two 715W power supplies.

Table 3-48. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Max Steady Power (W) |
|--------------------------------|---------------|----------------------------|--|----------------------|
| Dell EMC Networking N3132PX-ON | 100V/60Hz | Two 715W | 14.37A | 1429.8W |
| | 110V/60Hz | Two 715W | 12.95A | 1417.6W |
| | 120V/60Hz | Two 715W | 11.78A | 1409.1W |
| | 220V/50Hz | Two 715W | 6.35A | 1374.8W |
| | 240V/50Hz | Two 715W | 5.84A | 1372.5W |

Table 3-49 shows power consumption data for the PoE-enabled N3132PX-ON switch when the power budget is 750W for one 1100W power supply.

Table 3-49. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Max Steady Power (W) |
|--------------------------------|----------------------|-----------------------------------|---|-----------------------------|
| Dell EMC Networking N3132PX-ON | 100V/60Hz | One 1100W | 9.41A | 937.1W |
| | 110V/60Hz | One 1100W | 8.48A | 929.7W |
| | 120V/60Hz | One 1100W | 7.69A | 918.3W |
| | 220V/50Hz | One 1100W | 4.16A | 904.3W |
| | 240V/50Hz | One 1100W | 3.81A | 902.3W |

Table 3-50 shows power consumption data for the PoE-enabled N3132PX-ON switch when the power budget is 1700W for two 1100W power supplies.

Table 3-50. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Max Steady Power (W) |
|--------------------------------|----------------------|-----------------------------------|---|-----------------------------|
| Dell EMC Networking N3132PX-ON | 100V/60Hz | Two 1100W | 19.16A | 1911.2W |
| | 110V/60Hz | Two 1100W | 17.24A | 1892W |
| | 120V/60Hz | Two 1100W | 15.68A | 1873W |
| | 220V/50Hz | Two 1100W | 8.37A | 1819W |
| | 240V/50Hz | Two 1100W | 7.7A | 1819.2W |

Table 3-51 shows power consumption data for the PoE-enabled N3132PX-ON switch when the power budget is 1440W for one 1100W power supply + one 715W power supply.

Table 3-51. Power Consumption

| Model | Input Voltage | Power Supply Configuration | Maximum Steady Current Consumption (A) | Max Steady Power (W) |
|--------------------------------|---------------|----------------------------|--|----------------------|
| Dell EMC Networking N3132PX-ON | 100V/60Hz | 1100W + 715W | 17.51A | 1748W |
| | 110V/60Hz | 1100W + 715W | 15.7A | 1722.3W |
| | 120V/60Hz | 1100W + 715W | 14.36A | 1704.2W |
| | 220V/50Hz | 1100W + 715W | 7.63A | 1663.1W |
| | 240V/50Hz | 1100W + 715W | 6.99A | 1656.3W |

PoE Power Budget Limit

The PoE power budget for each interface is controlled by the switch firmware. The administrator can limit the power supplied on a port or prioritize power to some ports over others. Table 3-38 shows the power budget data.

Table 3-52. Dell EMC Networking N3132PX-ON PoE Power Budget Limit

| Model Name | One PSU | | Two PSUs | |
|--------------------------------|-------------------------|---|--------------------------|--|
| | Max. PSU Output Ability | PoE+ Power Turn-on Limitation | Max. PSUs Output Ability | PoE+ Power Turn-on Limitation |
| Dell EMC Networking N3132PX-ON | 715W | Power budget is 500W: The total PoE supplied power must not exceed 500W. | 715W | Power budget is 1200W: All PoE+ ports can supply maximum power. |

Dell EMC Networking N4000 Series Switch Hardware



NOTE: Both the Dell EMC Networking PC8100 and N4000 Series switches can run firmware versions 6.0.0.8 and beyond. The Dell EMC Networking N4000 Series switches cannot run firmware prior to version 6.0.0.8.

This section contains information about device characteristics and modular hardware configurations for the Dell EMC Networking N4000 Series switches.

Front Panel

The Dell EMC Networking N4000 Series front panel includes the following features:

- Switch ports
- Module bay that supports the following modules:
 - 2 x 40 Gig QSFP (each QSFP may be configured as 4 x 10 Gig ports)
 - 4 x SFP+ module
 - 4 x 10GBASE-T module

See "Hot-Pluggable Interface Modules" on page 178 for more information.

- USB port
- Reset button
- Port and system LEDs
- Stack LED

The Dell EMC Networking N4032 front panel provides 24 x 10GbE copper ports that support up to 100M of CAT-6A UTP cabling. The Dell EMC Networking N4032F provides 24 SFP+ ports supporting SFP+ and SFP transceivers.

Figure 3-26. Dell EMC Networking N4032 Front Panel





Figure 3-27. Dell EMC Networking N4032F Front Panel



Dell EMC Networking N4032 and N4032F switches can be stacked with other Dell EMC Networking N4000 Series switches using 10G or 40G SFP+ or QSFP modules in the module bay.

The Dell EMC Networking N4064 front panel provides 48 x 10GbE copper ports and two fixed QSFP ports, each supporting 4 x 10G or 1 x 40G connections. The Dell EMC Networking N4064F front panel provides 48 SFP+ ports supporting SFP+ and SFP transceivers plus two fixed QSFP ports, each supporting 4 x 10G or 1 x 40G connections.

All Dell EMC Networking N4000 Series ports operate in full-duplex mode only. The copper ports require that auto-negotiation be enabled. Auto-negotiation should be disabled for Dell EMC Networking N4000 Series SFP+ ports and enabled for SFP+ ports with SFP transceivers installed.

Figure 3-28. Dell EMC Networking N4064 Front Panel





Figure 3-29. Dell EMC Networking N4064F Front Panel



The Dell EMC Networking N4064 and N4064F switches can be stacked with other Dell EMC Networking N4000 Series switches using the 10G or 40G SFP+ or QSFP modules in the module bay or fixed QSFP ports.

Hot-Pluggable Interface Modules

The Dell EMC Networking N4032, N4032F, N4064, and N4064F switches support the following hot-pluggable interface modules:

- N4000-QSFP — 2 x 40G QSFP port module - defaults to 2 x 40G
- N4000-SFP+ — 4 x SFP+ port module - defaults to 4 x 10G mode
- N4000-10GBT — 4 x 10GBASE-T ports module - defaults to 4 x 10G mode

- Blank module — defaults to 10G mode

A reboot is not necessary when a hot-pluggable Ethernet module is replaced with an Ethernet module of different type. Issue the **no slot** command after removing the module and prior to installing the new module. Plug-in modules with any port configured as a stacking port are not hot-swappable.

A **no slot** command must be executed prior to inserting the new Ethernet module. Changing the role of a port from stacking to Ethernet or vice-versa requires a switch reboot.

If a **no slot** command is not issued prior to inserting a module, a message such as the following will appear:

```
Card Mismatch: Unit:1 Slot:1 Inserted-Card: Dell 2 Port QSFP  
Expansion Card Config-Card: Dell 4 Port 10GBase-T Expansion Card
```

The following sections provides details on each module.

Quad-Port SFP (QSFP) Uplink Module

The QSFP module supports features four ports that support 10G SFP+ transceivers. The QSFP module supports the following features:

- Four 10G ports with quad-breakout/QBO cable or one 40G port
- Front-panel port status LEDs

The QSFP interfaces can be used for stacking. Stacking is supported at distances of up to 100M.

Quad-Port SFP+ Uplink Module

The N4000-SFP+ module features four SFP+ ports, each providing the following features:

- SFP+ optical interfaces
- SFP+ copper twinax interface
- Front-panel port status LEDs

The SFP+ connections can be used for stacking. Stacking is supported at distances of up to 100M.

10GBASE-T Copper Uplink Module

The 10GBASE-T copper module features four copper ports that can support 10GbE/1GbE/100MbE switching and provides following features:

- Complies with IEEE802.3z, IEEE 802.3, IEEE802.3u, IEEE802.3ab, IEEE802.3az, IEEE802.3an
- Four 10GBASE-T/1GBASE-T/100MBASE-T copper ports.
- front-panel port status LEDs

USB Port

The Type-A, female USB port supports a USB 2.0-compliant flash memory drive. The Dell EMC Networking N4000 Series switch can read or write to a flash drive with a single partition formatted as FAT-32. Use a USB flash drive to copy switch configuration files and images between the USB flash drive and the switch. The USB flash drive may be used to move and copy configuration files and images from one switch to other switches in the network. Deletion of files on the USB drive is not supported.

The USB port does not support any other type of USB device.

Port and System LEDs

The front panel contains light emitting diodes (LEDs) to indicate port status. For information about the status that the LEDs indicate, see "LED Definitions" on page 182.

SFP+ and QSFP+ Ports

SFP+ and QSFP+ ports support Dell EMC-qualified transceivers. The default behavior is to log a message and generate an SNMP trap on insertion or removal of an optic that is not qualified by Dell EMC. This message and trap can be suppressed by using the **service unsupported-transceiver** command.


Back Panel

The Dell EMC Networking N4000 Series back panel has the following features:

- Console port
- Out-of-band management port
- Power Supplies
- Ventilation System

The following image shows the back panel of the Dell EMC Networking N4000 Series switches.

Figure 3-30. Dell EMC Networking N4000 Series Back Panel



Console Port

The console port is for management through a serial interface. This port provides a direct connection to the switch and provides access to the CLI from a console terminal connected to the port through the provided serial cable (RJ-45 to female DB-9 connectors).

The console port supports asynchronous data of eight data bits, one stop bit, no parity bit, and no flow control. The default BAUD is 9600 bps.

Out-of-Band Management Port

The Out-of-Band (OOB) management port is a 10/100/1000BASE-T Ethernet port connected directly to the switch CPU and dedicated to switch management. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to or from the operational network. In addition, ACLs (including management ACLS), do not operate on the out-of-band port. Only connect the out-of-band port to a physically secure network.

Power Supplies

Each Dell EMC Networking N4000 Series switch has two power supplies for redundant or loadsharing operation. Each power supply can support 300W.

CAUTION: Remove the power cable from the modules prior to removing the module itself. Power must not be connected prior to insertion in the chassis.

Ventilation System

The Dell EMC Networking N4000 Series switches have two fans. Each switch also has four thermal sensors and a fan speed controller, which can be used to control FAN speeds. Verify operation by observing the LEDs.

LED Definitions

This section describes the LEDs on the front and back panels of the switch.

Port LEDs

Each port on a Dell EMC Networking N4000 Series switch includes two LEDs. One LED is on the left side of the port, and the second LED is on the right side of the port. This section describes the LEDs on the switch ports.

100/1000/10000BASE-T Port LEDs

Each 100/1000/10000BASE-T port has two LEDs. Figure 3-31 illustrates the 100/1000/10000BASE-T port LEDs.

Figure 3-31. 100/1000/10000BASE-T Port LEDs




Table 3-53 shows the 100/1000/10000BASE-T port LED definitions.

Table 3-53. 100/1000/10000BASE-T Port LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is operating at 10 Gbps. |
| | Solid amber | The port is operating at 100/1000 Mbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Module Bay LEDs

The following tables describe the purpose of each of the module bay LEDs when SFP+, 10GBASE-T, and QSFP modules are used.

Table 3-54. SFP+ Module LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is operating at 10 Gbps. |
| | Solid amber | The port is operating at 100/1000 Mbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Table 3-55. 10GBASE-T Module LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is operating at 10 Gbps. |
| | Solid amber | The port is operating at 100/1000 Mbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Table 3-56. QSFP Module LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is operating at 40 Gbps. |
| | Solid amber | The port is operating at other speeds. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

Out-of-Band Ethernet Management Port LEDs

Table 3-57 shows the LED definitions for the OOB Ethernet management port.

Table 3-57. OOB Ethernet Management Port LED Definitions

| LED | Color | Definition |
|--------------|----------------|--|
| Link LED | Off | There is no link. |
| | Solid green | The port is operating at 1000 Mbps. |
| | Solid amber | The port is operating at 10/100 Mbps. |
| Activity LED | Off | There is no current transmit/receive activity. |
| | Blinking green | The port is actively transmitting/receiving. |

System LEDs

The system LEDs, located on the front panel, provide information about the power supplies, thermal conditions, and diagnostics.

Table 3-58 shows the System LED definitions for the Dell EMC Networking N4000 Series switches.

Table 3-58. System LED Definitions—Dell EMC Networking N4000 Series Switches

| LED | Color | Definition |
|---------|----------------|--|
| System | Blinking blue | The switch is booting |
| | Solid red | A critical system error has occurred. |
| | Blinking red | A noncritical system error occurred (fan or power supply failure). |
| Temp | Off | The switch is operating at normal temperature. |
| | Solid amber | The thermal sensor's system temperature threshold of 75°C has been exceeded. |
| Diag | Off | The switch is operating normally |
| | Blinking green | A diagnostic test is running. |
| Fan | Solid green | The fan is powered and is operating at the expected RPM. |
| | Solid red | A fan failure has occurred. |
| Stack | Solid blue | The switch is in stacking master mode. |
| | Solid amber | The switch is in stacking slave mode. |
| | Off | The switch is in stand-alone mode. |
| Locator | Blinking green | The locator function is enabled. |
| | Solid green | The locator function is disabled. |

Switch MAC Addresses

The switch allocates MAC addresses from the Vital Product Data information stored locally in flash. MAC addresses are used as follows:

Table 3-59. MAC Address Use

| | |
|----------|---|
| Base | Switch address, layer 2 |
| Base + 1 | Out-of-band port (not available on Dell EMC Networking N1100-ON/N1500/N2000/N2100-ON Series switches) |
| Base + 3 | Layer 3 |

Shown below are three commands that display the MAC addresses used by the switch:

```
console#show system
```

```
System Description: Dell Ethernet Switch
System Up Time: 0 days, 00h:05m:11s
System Contact:
System Name:
System Location:
Burned In MAC Address: 001E.C9F0.004D
System Object ID: 1.3.6.1.4.1.674.10895.3042
System Model ID: N4032
Machine Type: N4032
Temperature Sensors:
```

| Unit | Description | Temperature (Celsius) | Status |
|------|------------------|--------------------------|--------|
| ---- | ----- | ----- | ----- |
| 1 | MAC | 32 | Good |
| 1 | CPU | 31 | Good |
| 1 | PHY (left side) | 26 | Good |
| 1 | PHY (right side) | 29 | Good |

```
Fans:
```

| Unit | Description | Status |
|------|-------------|----------|
| ---- | ----- | ----- |
| 1 | Fan 1 | OK |
| 1 | Fan 2 | OK |
| 1 | Fan 3 | OK |
| 1 | Fan 4 | OK |
| 1 | Fan 5 | OK |
| 1 | Fan 6 | No Power |

Power Supplies:

| Unit | Description | Status | Average Power (Watts) | Current Power (Watts) | Since Date/Time |
|------|-------------|----------|-----------------------|-----------------------|---------------------|
| 1 | System | OK | 42.0 | 43.4 | |
| 1 | Main | OK | N/A | N/A | 04/06/2001 16:36:16 |
| 1 | Secondary | No Power | N/A | N/A | 01/01/1970 00:00:00 |

USB Port Power Status:

Device Not Present

console#show ip interface out-of-band

IP Address..... 10.27.21.29
Subnet Mask..... 255.255.252.0
Default Gateway..... 10.27.20.1
Configured IPv4 Protocol..... DHCP
Burned In MAC Address..... 001E.C9F0.004E

console#show ip interface vlan 1

Routing Interface Status..... Down
Primary IP Address..... 1.1.1.2/255.255.255.0
Method..... Manual
Routing Mode..... Enable
Administrative Mode..... Enable
Forward Net Directed Broadcasts..... Disable
Proxy ARP..... Enable
Local Proxy ARP..... Disable
Active State..... Inactive
MAC Address..... 001E.C9F0.0050
Encapsulation Type..... Ethernet
IP MTU..... 1500
Bandwidth..... 10000 kbps
Destination Unreachables..... Enabled
ICMP Redirects..... Enabled

Using Dell EMC OpenManage Switch Administrator

Dell EMC Networking N-Series Switches

This section describes how to use the Dell EMC OpenManage Switch Administrator application. The topics covered in this section include:

- About Dell EMC OpenManage Switch Administrator
- Starting the Application
- Understanding the Interface
- Using the Switch Administrator Buttons and Links
- Defining Fields

About Dell EMC OpenManage Switch Administrator

Dell EMC OpenManage Switch Administrator is a web-based tool for managing and monitoring Dell EMC Networking N-Series switches. Table 4-1 lists the web browsers that are compatible with Dell EMC OpenManage Switch Administrator. The browsers have been tested on a PC running the Microsoft Windows operating system.

Table 4-1. Compatible Browsers

| Browser | Version |
|-------------------|---------|
| Internet Explorer | v9 |
| Mozilla Firefox | v14 |
| Safari | v5.0 |
| Chrome | v21 |




NOTE: Additional operating systems and browsers might be compatible but have not been explicitly tested with Dell EMC OpenManage Switch Administrator.


Starting the Application

To access the Dell EMC OpenManage Switch Administrator and log on to the switch:

- 1 Open a web browser.
- 2 Enter the IP address of the switch in the address bar and press <Enter>. For information about assigning an IP address to a switch, see "Setting the IP Address and Other Basic Network Information" on page 209.
- 3 When the **Login** window displays, enter a username and password. Passwords and usernames are both case sensitive and alpha-numeric.

Figure 4-1. Login Screen



 **NOTE:** The switch is not configured with a default user name or password. The administrator must connect to the CLI by using the console port to configure the initial user name and password. For information about connecting to the console, see "Console Connection" on page 197. For information about creating a user and password, see "Authentication, Authorization, and Accounting" on page 275.


- 4 Click **Submit**.
- 5 The **Dell EMC OpenManage Switch Administrator** home page displays.
The home page is the **Device Information** page, which contains a graphical representation of the front panel of the switch. For more information about the home page, see "Device Information" on page 414.

Understanding the Interface

The Dell EMC OpenManage Switch Administrator interface contains the following components:

- **Navigation panel** — Located on the left side of the page, the navigation pane provides an expandable view of features and their components.
- **Configuration and status options** — The main panel contains the fields used to configure and monitor the switch.
- **Page tabs** — Some pages contain tabs that allow the administrator to access additional pages related to the feature.
- **Command buttons** — Command buttons are located at the bottom of the page. Use the command buttons to submit changes, perform queries, or clear lists.
- **Save, Print, Refresh, and Help buttons** — These buttons appear on the top-right side of the main panel and are on every page.
- **Support, About, and Logout links** — These links appear at the top of every page.

Figure 4-2. Switch Administrator Components



Using the Switch Administrator Buttons and Links

Table 4-2 describes the buttons and links available from the Dell EMC OpenManage Switch Administrator interface.

Table 4-2. Button and Link Descriptions

| Button or Link | Description |
|----------------------------|---|
| Support | Opens the Dell Support page at www.dell.com/support . |
| About | Contains the version and build number and Dell copyright information. |
| Log Out | Logs out of the application and returns to the login screen. |
| Save | Saves the running configuration to the startup configuration. When a user clicks Apply , changes are saved to the running configuration. When the system boots, it loads the startup configuration. Any changes to the running configuration that were not saved to the startup configuration are lost across a power cycle. |
| Print | Opens the printer dialog box that enables printing the current page. Only the main panel prints. |
| Refresh | Refreshes the screen with the current information. |
| Help | Online help that contains information to assist in configuring and managing the switch. The online help pages are context sensitive. For example, if the IP Addressing page is open, the help topic for that page displays if the user clicks Help . |
| Apply | Updates the running configuration on the switch with the changes. Configuration changes take effect immediately. |
| Clear | Resets statistic counters and log files to the default configuration. |
| Query | Queries tables. |
| Left arrow and Right arrow | Moves information between lists. |



NOTE: A few pages contain a button that occurs only on that page. Page-specific buttons are described in the sections that pertain to those pages.

Defining Fields

User-defined fields can contain 1–159 characters, unless otherwise noted on the Dell EMC OpenManage Switch Administrator web page.


All characters may be used except for the following:

- \
- /
- :
- *
- ?
- <
- >
- |

Understanding the Device View

The Device View shows various information about switch. This graphic appears on the Dell EMC OpenManage Switch Administrator **Home** page, which is the page that displays after a successful login. The graphic provides information about switch ports and system health.

Figure 4-3. Dell EMC Networking N3048 Device View



Using the Device View Port Features

The switching-port coloring indicates if a port is currently active. Green indicates that the port has a link, red indicates that an error has occurred on the port, and blue indicates that the link is down. Ethernet ports configured for stacking show as gray. Each port image is a hyperlink to the **Port Configuration** page for the specific port.

Using the Device View Switch Locator Feature

The Device View graphic includes a **Locate** button and a drop-down menu of timer settings. When the user clicks **Locate**, the switch locator LED blinks for the number of seconds selected from the timer menu. The blinking LED can help the administrator or a technician near the switch identify the physical location of the switch within a room or rack full of switches. After the user clicks the **Locate** button, it turns green on the screen and remains green while the LED is blinking. For information about the locator LED on a specific switch, including color and physical placement, see the hardware description for the switch model in "Hardware Overview" on page 113.



NOTE: The `locate` command in the CLI can be used to enable the locator LED.

Using the Command-Line Interface

Dell EMC Networking N-Series Switches

This section describes how to use the Command-Line Interface (CLI) on Dell EMC Networking N-Series switches.

The topics covered in this section include:

- Accessing the Switch Through the CLI
- Understanding Command Modes
- Entering CLI Commands

Accessing the Switch Through the CLI

The CLI provides a text-based way to manage and monitor the Dell EMC Networking N-Series switches. The CLI can be accessed using a direct connection to the console port or by using a Telnet or SSH client.

To access the switch by using Telnet or Secure Shell (SSH), the switch must have an IP address, and the management station used to access the device must be able to ping the switch IP address.

For information about assigning an IP address to a switch, see "Setting the IP Address and Other Basic Network Information" on page 209.

Console Connection

Use the following procedures to connect to the CLI by connecting to the console port. For more information about creating a serial connection, see the Getting Started Guide available at www.dell.com/support.

- 1 Connect the DB-9 connector of the supplied serial cable to a management station, and connect the RJ-45 connector to the switch console port. The N1100-ON switches utilize a USB cable to access the console.

On Dell EMC Networking N1500, N2000, N2100-ON, N3000 and N3100-ON Series switches, the console port is located on the right side of the front panel and is labeled with the |O|O| symbol. On the Dell EMC Networking N4000 Series switches, it is located on the back panel above

the OOB Ethernet port. On the N1100-ON Series switches, the USB console port is located in the bottom right corner of the front panel.



NOTE: For a stack of switches, be sure to connect to the console port on the Master switch. The Master LED is illuminated on the stack Master. Alternatively, use the connect command to access the console session.

- 2 Start the terminal emulator, such as Microsoft HyperTerminal, and select the appropriate serial port (for example, COM 1) to connect to the console.
- 3 Configure the management station serial port with the following settings:
 - Data rate — 9600 BAUD (115,200 for the N1100-ON, N2128PX-ON, N3048EP-ON, and N3132PX-ON switches).
 - Data format — 8 data bits
 - Parity — None
 - Stop bits — 1
 - Flow control — None
- 4 Power on the switch (or stack).

After the boot process completes, the `console>` prompt displays, and CLI commands can be entered.



NOTE: By default, no authentication is required for console access. However, if an authentication method has been configured for console port access, the User: login prompt displays.

Telnet Connection

Telnet is a terminal emulation TCP/IP protocol. ASCII terminals can be virtually connected to the local device through a TCP/IP protocol network.

Telnet connections are enabled by default, and the Telnet port number is 23. All CLI commands can be used over a Telnet session. Use the **terminal monitor** command to receive asynchronous notification of system events in the telnet session.



NOTE: SSH, which is more secure than Telnet, is disabled by default.

To connect to the switch using Telnet, the switch must have an IP address, and the switch and management station must have network connectivity. Any Telnet client on the management station can be used to connect to the switch.

A Telnet session can also be initiated from the Dell EMC OpenManage Switch Administrator. For more information, see "Initiating a Telnet Session from the Web Interface" on page 453.

Understanding Command Modes

The CLI groups commands into modes according to the command function. Each of the command modes supports specific software commands. The commands in one mode are not available until the user switches to that particular mode, with the exception of the User Exec mode commands. The User Exec mode commands can be executed in the Privileged Exec mode.

To display the commands available in the current mode, enter a question mark (?) at the command prompt. In each mode, a specific command is used to navigate from one command mode to another.

The main command modes include the following:

- User Exec (0) — Commands in this mode permit connecting to remote devices, changing terminal settings on a temporary basis, performing basic tests, and listing limited system information.
- Privileged Exec (1) — Commands in this mode enable viewing all switch settings and entering Global Configuration mode.
- Global Configuration (15) — Commands in this mode manage the device configuration on a global level and apply to system features, rather than to a specific protocol or interface.
- Interface Configuration (15) — Commands in this mode configure the settings for a specific interface or range of interfaces.
- VLAN Configuration (15) — Commands in this mode create and remove VLANs and configure IGMP/MLD Snooping parameters for VLANs.

The CLI includes many additional command modes. For more information about the CLI command modes, including details about all modes, see the CLI Reference Guide.

Table 5-1 describes how to navigate between CLI Command Mode and lists the prompt that displays in each mode.

Table 5-1. Command Mode Overview

| Command Mode | Access Method | Command Prompt | Exit or Access Previous Mode |
|-------------------------|---|-----------------------|---|
| User Exec | The user is automatically in User Exec mode unless the user is defined as a privileged user. | console> | logout |
| Privileged Exec | From User Exec mode, enter the enable command | console# | Use the exit command, or press Ctrl-Z to return to User Exec mode. |
| Global Configuration | From Privileged Exec mode, use the configure command. | console(config)# | Use the exit command, or press Ctrl-Z to return to Privileged Exec mode. |
| Interface Configuration | From Global Configuration mode, use the interface command and specify the interface type and ID. | console(config-if)# | To exit to Global Configuration mode, use the exit command, or press Ctrl-Z to return to Privileged Exec mode. |

Entering CLI Commands

The switch CLI provides several techniques to help users enter commands.

Using the Question Mark to Get Help

Enter a question mark (?) at the command prompt to display the commands available in the current mode.

```
console(config-vlan)#?
```

```
exit           To exit from the mode.
help           Display help for various special keys.
ip             Configure IP parameters.
ipv6           Configure IPv6 parameters.
protocol       Configure the Protocols associated with
               particular Group Ids.
vlan           Create a new VLAN or delete an existing
               VLAN.
```

Enter a question mark (?) after entering each word to display available command keywords or parameters.

```
console(config)#vlan ?
```

```
<vlan-list>   <1-4093> - separate non-consecutive IDs with ',' and
               no spaces; Use '-' for range.
makestatic     Change the VLAN type from 'Dynamic' to 'Static'.
protocol       Configure the protocol based VLAN settings.
```

If there are no additional command keywords or parameters, or if additional parameters are optional, the following message appears in the output:

```
<cr>          Press enter to execute the command.
```

Typing a question mark (?) after one or more characters of a word shows the available command or parameters that begin with the characters, as shown in the following example:

```
console#show po?
```

```
policy-map           port           ports
```

Using Command Completion

The CLI can complete partially entered commands when the <Tab> or <Space> key are pressed.

```
console#show run<Tab>  
console#show running-config
```

If the characters entered are not enough for the switch to identify a single matching command, continue entering characters until the switch can uniquely identify the command. Use the question mark (?) to display the available commands matching the characters already entered.

Entering Abbreviated Commands

To execute a command, enter enough characters so that the switch can uniquely identify a command. For example, to enter Global Configuration mode from Privileged Exec mode, enter **conf** instead of **configure**.

```
console#conf  
  
console(config)#
```

Negating Commands

For many commands, the prefix keyword **no** is entered to cancel the effect of a command or reset the configuration to the default value. Many configuration commands have this capability.

Command Output Paging

Lines are printed on the screen up to the configured terminal length limit (default 24). Use the space bar to show the next page of output or the carriage return to show the next line of output. Setting the terminal length to zero disables paging. Command output displays until no more output is available.

Understanding Error Messages

If a command is entered and the system is unable to execute it, an error message appears. Table 5-2 describes the most common CLI error messages.

Table 5-2. CLI Error Messages

| Message Text | Description |
|---|--|
| % Invalid input detected at '^' marker. | Indicates that an incorrect or unavailable command was entered. The carat (^) shows where the invalid text is detected. This message also appears if any of the parameters or values are not recognized. |
| Command not found / Incomplete command. Use ? to list commands. | Indicates that the required keywords or values were not entered. |
| Ambiguous command | Indicates that not enter enough letters were entered to uniquely identify the command. |

If you attempt to execute a command and receive an error message, use the question mark (?) to help determine the possible keywords or parameters that are available.

Recalling Commands from the History Buffer

Every time a command is entered in the CLI, it is recorded in an internally managed Command History buffer. By default, the history buffer is enabled and stores the last 10 commands entered. These commands can be recalled, reviewed, modified, and reissued. This buffer is not preserved after switch resets.

Table 5-3. History Buffer Navigation

| Keyword | Source or Destination |
|---------------------------------|--|
| Up-arrow key or <Ctrl>+<P> | Recalls commands in the history buffer, beginning with the most recent command. Repeats the key sequence to recall successively older commands. |
| Down-arrow key or <Ctrl>+<N> | Returns to more recent commands in the history buffer after recalling commands with the up-arrow key. Repeating the key sequence recalls more recent commands in succession. |

Default Settings

This section describes the default settings for many of the software features on the Dell EMC Networking N-Series switches.

Table 6-1. Default Settings

| Feature | Default |
|----------------------------|---|
| IP address | DHCP on OOB interface, if equipped. DHCP on VLAN1 if no OOB interface |
| Subnet mask | None |
| Default gateway | None |
| DHCP client | Enabled on out-of-band (OOB) interface or VLAN 1 if no OOB interface. |
| VLAN 1 Members | All switch ports |
| SDM template | Dual IPv4 and IPv6 routing |
| Users | None |
| Minimum password length | 8 characters |
| IPv6 management mode | Enabled |
| SNTP client | Disabled |
| Global logging | Enabled |
| Switch auditing | Enabled |
| CLI command logging | Disabled |
| Web logging | Disabled |
| SNMP logging | Disabled |
| Console logging | Enabled (Severity level: warnings and above) |
| Monitor logging: | Disabled |
| Buffer (In-memory) logging | Enabled (Severity level: informational and above) |
| Persistent (flash) logging | Enabled (Severity level: emergencies and above) |

Table 6-1. Default Settings (Continued)

| Feature | Default |
|-------------------------------------|---------------------------------|
| DNS | Enabled (No servers configured) |
| SNMP | Enabled (SNMPv1) |
| SNMP Traps | Enabled |
| Auto Configuration | Enabled |
| Auto Save | Disabled |
| Stacking | Enabled |
| Nonstop Forwarding on the Stack | Enabled |
| sFlow | Disabled |
| ISDP | Enabled (Versions 1 and 2) |
| RMON | Enabled |
| TACACS+ | Not configured |
| RADIUS | Not configured |
| SSH/SSL | Disabled |
| Telnet | Enabled |
| Denial of Service Protection | Disabled |
| Captive Portal | Disabled |
| IEEE 802.1X Authentication | Disabled |
| Access Control Lists (ACL) | None configured |
| IP Source Guard (IPSG) | Disabled |
| DHCP Snooping | Disabled |
| Dynamic ARP Inspection | Disabled |
| Protected Ports (Private VLAN Edge) | None |
| Energy Detect Mode | Enabled |
| EEE Lower Power Mode | Enabled |
| PoE Plus (POE switches) | Auto |
| Flow Control Support (IEEE 802.3x) | Enabled |
| Maximum Frame Size | 1518 bytes |

Table 6-1. Default Settings (Continued)

| Feature | Default |
|---------------------------------------|---|
| Auto-MDI/MDIX Support | Enabled |
| Auto-negotiation | Enabled |
| Advertised Port Speed | Maximum Capacity |
| Broadcast Storm Control | Disabled |
| Port Mirroring | Disabled |
| LLDP | Enabled |
| LLDP-MED | Disabled |
| Loop Protection | Disabled |
| MAC Table Address Aging | 300 seconds (Dynamic Addresses) |
| Cisco Protocol Filtering (LLPF) | No protocols are blocked |
| DHCP Layer-2 Relay | Disabled |
| Default VLAN ID | 1 |
| Default VLAN Name | Default |
| GVRP | Disabled |
| GARP Timers | Leave: 60 centiseconds Leave All: 1000 centiseconds Join: 20 centiseconds |
| Interface Auto-Recovery (err-disable) | Disabled for all causes |
| Voice VLAN | Disabled |
| Guest VLAN | Disabled |
| RADIUS-assigned VLANs | Disabled |
| Double VLANs | Disabled |
| Spanning Tree Protocol (STP) | Enabled |
| STP Operation Mode | IEEE 802.1w Rapid Spanning Tree |
| Optional STP Features | Disabled |
| STP Bridge Priority | 32768 |
| Multiple Spanning Tree | Disabled |

Table 6-1. Default Settings (Continued)

| Feature | Default |
|--------------------------------|--|
| Link Aggregation | No LAGs configured |
| LACP System Priority | 1 |
| Routing Mode | Disabled |
| OSPF Admin Mode | Disabled |
| OSPF Router ID | 0.0.0.0 |
| IP Helper and UDP Relay | Disabled |
| RIP | Disabled |
| VRRP | Disabled |
| Tunnel and Loopback Interfaces | None |
| IPv6 Routing | Disabled |
| DHCPv6 | Disabled |
| OSPFv3 | Disabled |
| DiffServ | Enabled |
| Auto VoIP | Disabled |
| Auto VoIP Traffic Class | 6 |
| PFC | Disabled; no classifications configured. |
| DCBx version | Auto detect |
| iSCSI | Enabled |
| MLD Snooping | Enabled |
| IGMP Snooping | Enabled |
| IGMP Snooping Querier | Disabled |
| GMRP | Disabled |
| IPv4 Multicast | Disabled |
| IPv6 Multicast | Disabled |
| OpenFlow | Disabled |

Setting the IP Address and Other Basic Network Information

Dell EMC Networking N-Series Switches

This chapter describes how to configure basic network information for the switch, such as the IP address, subnet mask, and default gateway. The topics in this chapter include:

- IP Address and Network Information Overview
- Default Network Information
- Configuring Basic Network Information (Web)
- Configuring Basic Network Information (CLI)
- Basic Network Information Configuration Examples

IP Address and Network Information Overview

What Is the Basic Network Information?

The basic network information includes settings that define the Dell EMC Networking N-Series switches in relation to the network. Table 7-1 provides an overview of the settings this chapter describes.

Table 7-1. Basic Network Information

| Feature | Description |
|-------------|--|
| IP Address | On an IPv4 network, the a 32-bit number that uniquely identifies a host on the network. The address is expressed in dotted-decimal format, for example 192.168.10.1. |
| Subnet Mask | Determines which bits in the IP address identify the network, and which bits identify the host. Subnet masks are also expressed in dotted-decimal format, for example 255.255.255.0. |

Table 7-1. Basic Network Information (Continued)

| Feature | Description |
|---------------------------------|--|
| Default Gateway | Typically a router interface that is directly connected to the switch and is in the same subnet. The switch sends IP packets to the default gateway when it does not recognize the destination IP address in a packet. |
| DHCP Client | Requests network information from a DHCP server on the network. |
| Domain Name System (DNS) Server | Translates hostnames into IP addresses. The server maintains a domain name databases and their corresponding IP addresses. |
| Default Domain Name | Identifies your network, such as dell.com. If a hostname is entered without the domain name information, the default domain name is automatically appended to the hostname. |
| Host Name Mapping | Allows statically mapping an IP address to a hostname. |

Additionally, this chapter describes how to view host name-to-IP address mappings that have been dynamically learned by the system.

Why Is Basic Network Information Needed?

Dell EMC Networking N-Series switches are layer-2/3 managed switches. To manage the switch remotely by using a web browser or Telnet client, the switch must have an IP address, subnet mask, and default gateway. A username and password is required to be able to log into the switch from a remote host. For information about configuring users, see "Authentication, Authorization, and Accounting" on page 275. If managing the switch by using the console connection only, configuring an IP address and user is not required. In this case, disabling the Telnet server using the **no ip telnet** command is recommended.



NOTE: The configuration example in this chapter includes commands to create an administrative user with read/write access.

Configuring the DNS information, default domain name, and host name mapping help the switch identify and locate other devices on the network and on the Internet. For example, to upgrade the switch software by using a TFTP

server on the network, the TFTP server must be identified. If configuring the switch to use a DNS server to resolve hostnames into IP addresses, it is possible to enter the hostname of the TFTP server instead of the IP address. It is often easier to remember a hostname than an IP address, and if the IP address is dynamically assigned, it might change from time-to-time.

How Is Basic Network Information Configured?

A console-port connection is required to perform the initial switch configuration. When booting the switch for the first time, if there is no startup configuration file, the Dell Easy Setup Wizard starts. The Dell Easy Setup Wizard is a CLI-based tool to help the administrator perform the initial switch configuration. If no response to the Dell Easy Setup Wizard prompt is received within 60 seconds, the `console>` prompt appears, and the switch enters User Configuration mode.

For more information about performing the initial switch configuration by using the wizard, see the Getting Started Guide at www.dell.com/support.

If the wizard is not used to supply the initial configuration information, the administrator can manually enable the DHCP client on the switch to obtain network information from a DHCP server via the in-band ports or the out-of-band port. Alternatively, the network configuration can be statically configured.

After configuring the switch with an IP address and creating a user account, continue to use the console connection to configure basic network information, or log on to the switch by using a Telnet client or a web browser. It is possible at this point to change the IP address information and configure additional network information from the remote system.

What Is Out-of-Band Management and In-Band Management?

The Dell EMC Networking N3000, N3100-ON, and N4000 Series switches have an external port intended solely for management of the switch. This port is the out-of-band (OOB) management port. Traffic received on the OOB port is never switched or routed to any in-band port and is not rate limited. Likewise, traffic received on any in-band port is never forwarded or routed over the OOB port. The only applications available on the OOB port are protocols required to manage the switch, for example Telnet, SSH, DHCP client, and TFTP. If using the out-of-band management port, it is strongly

recommended that the port be connected only to a physically isolated secure management network. The OOB port is a layer-3 interface that uses an internal non-user-configurable VLAN.

The out-of-band port is a logical management interface. The IP stack's routing table contains both IPv4/IPv6 routes associated with these management interfaces and IPv4/IPv6 routes associated with routing interfaces. If routes to the same destination (such as a default route) are learned or configured on both the OOB interface and a routing interface, the routing interface route is preferred. If a directly connected subnet is configured on an out-of-band interface, it cannot also be configured on an in-band interface. If a default gateway is configured on routing interfaces (front-panel ports), then IP addresses not in the OOB port subnet will not be reachable via the OOB port. It is never recommended that the switch default gateway be configured on the out-of-band port subnet.

Dell recommends that, if used, the OOB port be used for remote management on a physically independent management network and be assigned an IP address from the non-routable private IP address space. The following list highlights some advantages of using OOB management instead of in-band management:

- Traffic on the OOB port is passed directly to the switch CPU, bypassing the switching silicon. The OOB port is implemented as an independent NIC, which allows direct access to the switch CPU from the management network.
- If the production network is experiencing problems, administrators can still access the switch management interface and troubleshoot issues.
- Because the OOB port is intended to be physically isolated from the production network or deployed behind a firewall, configuration options are limited to just those protocols needed to manage the switch. Limiting the configuration options makes it difficult to accidentally cut off management access to the switch.

Alternatively, network administrators may choose to manage their network via the production network. This is in-band management. Because in-band management traffic is mixed in with production network traffic, it is subject to all of the filtering rules usually applied on a switched/routed port, such as ACLs and VLAN tagging, and may be rate limited to protect against DoS attacks.

The administrator can assign an IPv4 address or an IPv6 address to the OOB management port and to any VLAN. By default, all ports (other than the OOB port) are members of VLAN 1. If an IP address is assigned to VLAN 1, it is possible to connect to the switch management interface by using any of the front-panel switch ports. Assignment of an IP address to a VLAN associated to a front panel interface is required to manage the Dell EMC Networking, N1100-ON, N1500, N2000, and N2100-ON Series switches. The use of VLAN 1 for switch administration presents some security risks. Alternatively, a management VLAN can be assigned as the native VLAN for a limited set of front-panel ports and an IP address can be assigned to that VLAN. The use of ACLs to restrict access to switch management is strongly recommended.


DHCP can be enabled on the OOB interface and VLAN interfaces simultaneously, or they can be configured with static information. To configure static address information on the default VLAN (or the management VLAN), set the IP address and subnet mask on the VLAN interface and configure a global default gateway for the switch to use front panel interfaces (not the OOB interface). If a default gateway is configured on routing interfaces (front-panel ports), then IP addresses not in the OOB port subnet will not be reachable via the OOB port. The switch sends the Vendor Class Identifier (Option 60) in the DHCP discover messages to assist DHCP server administrators in distinguishing Dell EMC switches from other devices in the network. This is a text string of the form "DellEMC;<switch model>;<firmware version>;<serial number>" where the switch model number is the specific switch model.

Adjusting the Management Interface MTU

When logging into the Dell EMC Networking N-Series switch using TCP, the switch negotiates the TCP Maximum Segment Size (MSS) using the minimum of the requested MSS or the MTU setting of the port. TCP packets are transmitted from the switch with the DF (Don't Fragment) bit set in order to receive notification of fragmentation from any transit routers. Upon receiving an ICMP Destination Unreachable, Fragmentation needed but DF set notification, the switch will reduce the MSS. However, many firewalls block ICMP Destination Unreachable messages, which causes the destination to request the packet again until the connection times out.


To resolve this issue, reduce the TCP MSS setting to a more appropriate value on the local host or alternatively, set the system MTU to a smaller value.

Default Network Information

 **NOTE:** Dell EMC Networking, N1100-ON, N1500, N2000, and N2100-ON Series switches do not have an out-of-band interface.

By default, no network information is configured. The DHCP client is enabled on the OOB interface by default on Dell EMC Networking N3000, N3100-ON, and N4000 Series switches. The DHCP client is enabled on VLAN 1 by default on the Dell EMC Networking, N1100-ON, N1500, N2000, and N2100-ON Series switches. DNS is enabled, but no DNS servers are configured. VLAN 1 does not have an IP address, subnet mask, or default gateway configured on Dell EMC Networking N3000, N3100-ON, and N4000 Series switches.

Configuring Basic Network Information (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring basic network information on the Dell EMC Networking N-Series switches. For details about the fields on a page, click  at the top of the Dell EMC OpenManage Switch Administrator web page.

Out-of-Band Interface



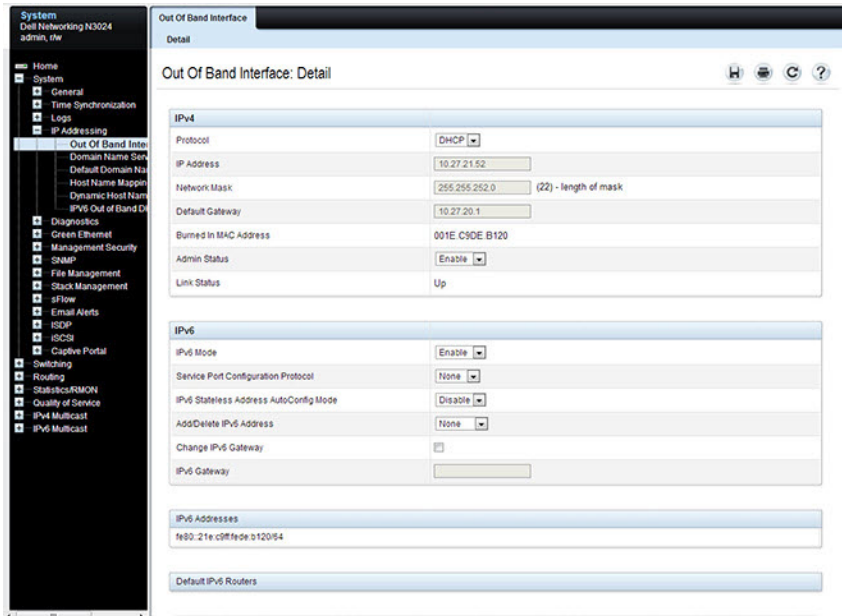
NOTE: Dell EMC Networking, N1100-ON, N1500, N2000, and N2100-ON Series switches do not have an out-of-band interface.

Use the **Out of Band Interface** page to assign the out-of-band interface IP address and subnet mask or to enable/disable the DHCP client for address information assignment. DHCP is enabled by default on the OOB interface. The OOB interface must be configured on a subnet separate from the front-panel port routing interfaces. The system default gateway must not share an address range/subnet with the OOB interface.

The out-of-band interface may also be assigned an IPv6 address, either statically or via DHCP. In addition, the out-of-band port may be assigned an IPv6 address via the IPv6 auto-configuration process.

To display the **Out of Band Interface** page, click **System** → **IP Addressing** → **Out of Band Interface** in the navigation panel.

Figure 7-1. Out of Band Interface




To enable the DHCP client and allow a DHCP server on your network to automatically assign the network information to the OOB interface, select DHCP from the **Protocol** menu. If the network information is statically assigned, ensure that the **Protocol** menu is set to None.

IP Interface Configuration (Default VLAN IP Address)

Use the **IP Interface Configuration** page to assign the default VLAN IP address and subnet mask, the default gateway IP address, and to assign the boot protocol.

To display the **IP Interface Configuration** page, click **Routing** → **IP** → **IP Interface Configuration** in the navigation panel.

Figure 7-2. IP Interface Configuration (Default VLAN)



Assigning Network Information to the Default VLAN

To assign an IP Address and subnet mask to the default VLAN:

- 1 From the **Interface** menu, select VLAN 1.
- 2 From the **Routing Mode** field, select **Enable**.
- 3 From the **IP Address Configuration Method** field specify whether to assign a static IP address (Manual) or use DHCP for automatic address assignment.
- 4 If **Manual** is selected for the configuration method, then the **IP Address** and **Subnet Mask** can be entered in the appropriate fields.
- 5 Click **Apply**.



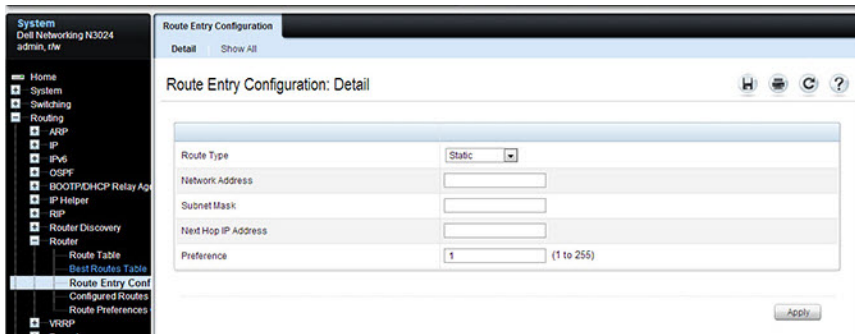
NOTE: No additional fields on the page must be configured. For information about VLAN routing interfaces, see "Routing Interfaces" on page 1213.

Route Entry Configuration (Switch Default Gateway)

Use the **Route Entry Configuration** page to configure the default gateway for the switch. The default VLAN uses the switch default gateway as its default gateway. The switch default gateway must not be on the same subnet as the OOB management port, as the OOB management port cannot route packets received on the front-panel ports.

To display the **Route Entry Configuration** page, click **Routing** → **Router** → **Route Entry Configuration** in the navigation panel.

Figure 7-3. Route Configuration (Default VLAN)

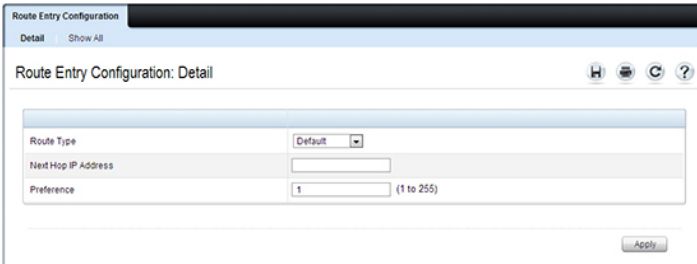


Configuring a Default Gateway for the Switch:

To configure the switch default gateway:

- 1 Open the **Route Entry Configuration** page.
- 2 From the **Route Type** field, select **Default**.

Figure 7-4. Default Route Configuration (Default VLAN)



The screenshot shows the 'Route Entry Configuration' window with the 'Detail' tab selected. The window title is 'Route Entry Configuration: Detail'. The configuration fields are as follows:

| Field | Value |
|---------------------|--------------|
| Route Type | Default |
| Next Hop IP Address | |
| Preference | 1 (1 to 255) |

An 'Apply' button is located at the bottom right of the configuration area.

- 3 In the **Next Hop IP Address** field, enter the IP address of the default gateway.
- 4 Click **Apply**.


For more information about configuring routes, see "IP Routing" on page 1187.

Domain Name Server

Use the **Domain Name Server** page to configure the IP address of the DNS server. The switch uses the DNS server to translate hostnames into IP addresses.


To display the **Domain Name Server** page, click **System** → **IP Addressing** → **Domain Name Server** in the navigation panel.

Figure 7-5. DNS Server



To configure DNS server information, click the **Add** link and enter the IP address of the DNS server in the available field.

Figure 7-6. Add DNS Server




Default Domain Name

Use the **Default Domain Name** page to configure the domain name the switch adds to a local (unqualified) hostname.

To display the **Default Domain Name** page, click **System** → **IP Addressing** → **Default Domain Name** in the navigation panel.

Figure 7-7. Default Domain Name




Host Name Mapping

Use the **Host Name Mapping** page to assign an IP address to a static host name. The **Host Name Mapping** page provides one IP address per host.


To display the **Host Name Mapping** page, click **System** → **IP Addressing** → **Host Name Mapping**.

Figure 7-8. Host Name Mapping



To map a host name to an IP address, click the **Add** link, type the name of the host and its IP address in the appropriate fields, and then click **Apply**.

Figure 7-9. Add Static Host Name Mapping




Use the **Show All** link to view all configured host name-to-IP address mappings.

Dynamic Host Name Mapping

Use the **Dynamic Host Name Mapping** page to view dynamic host entries the switch has learned. The switch learns hosts dynamically by using the configured DNS server to resolve a hostname. For example, if you ping `www.dell.com` from the CLI, the switch uses the DNS server to lookup the IP address of `dell.com` and adds the entry to the Dynamic Host Name Mapping table.

To display the **Dynamic Host Name Mapping** page, click **System** → **IP Addressing** → **Dynamic Host Name Mapping** in the navigation panel.


Figure 7-10. View Dynamic Host Name Mapping



Configuring Basic Network Information (CLI)

This section provides information about the commands used for configuring basic network information on the Dell EMC Networking N-Series switches. For more information about these commands, see the Dell EMC Networking N1100-ON, N1500, N2000, N2100-ON, N3000, N3100-ON, and N4000 Series Switches CLI Reference Guide at www.dell.com/support.

Enabling the DHCP Client on the OOB Port

 **NOTE:** Dell EMC Networking N1100-ON, N1500, N2000, and N2100-ON Series switches do not have an out-of-band interface.

Use the following commands to enable the DHCP client on the OOB port.

| Command | Purpose |
|--|--|
| <code>configure</code> | Enter Global Configuration mode. |
| <code>interface out-of-band</code> | Enter Interface Configuration mode for the OOB port. |
| <code>ip address dhcp</code> | Enable the DHCP client. |
| <code>CTRL + Z</code> | Exit to Privileged Exec mode. |
| <code>show ip interface out-of-band</code> | Display network information for the OOB port. |

Enabling the DHCP Client on the Default VLAN

Use the following commands to enable the DHCP client on the default VLAN, which is VLAN 1. As a best practice, it is recommended that a separate VLAN other than one used for client traffic be used for in-band management of the switch. In general, using VLAN 1, or any other VLAN carrying client traffic, for in-band management introduces a security vulnerability.

| Command | Purpose |
|-------------------------------|--|
| <code>configure</code> | Enter Global Configuration mode. |
| <code>interface vlan 1</code> | Enter Interface Configuration mode for VLAN 1. |
| <code>ip address dhcp</code> | Enable the DHCP client. |


| Command | Purpose |
|---------------------------------------|---|
| <code>ipv6 address dhcp</code> | Enable the DHCPv6 client. |
| <code>CTRL + Z</code> | Exit to Privileged Exec mode. |
| <code>show ip interface vlan 1</code> | Display network information for VLAN 1. |

Managing DHCP Leases

Use the following commands to manage and troubleshoot DHCP leases on the switch.


| Command | Purpose |
|--|---|
| <code>show dhcp lease interface [interface]</code> | Display IPv4 addresses leased from a DHCP server. |
| <code>show ipv6 dhcp interface vlan [interface]</code> | Display information about the IPv6 DHCP information for all interfaces or for the specified interface. |
| <code>debug dhcp packet</code> | Display debug information about DHCPv4 client activities and to trace DHCPv4 packets to and from the local DHCPv4 client. |
| <code>debug ipv6 dhcp</code> | Display debug information about DHCPv6 client activities and to trace DHCPv6 packets to and from the local DHCPv6 client. |
| <code>ipv6 address { [prefix/prefixlen] autoconfig dhcp }</code> | Set the IPv6 address of the management interface or enables auto-configuration or DHCP. |
| <code>ip default-gateway ipv4- address</code> | Configure a global default gateway. Only one IPv4 gateway may be configured per switch. |
| <code>ipv6 gateway ipv6- address</code> | Set the global IPv6 default gateway address. Only one IPv6 gateway may be configured per switch. |
| <code>ipv6 enable</code> | Enable IPv6 functionality on the interface. |
| <code>show ipv6 interface out- of-band</code> | Show settings for the interface. |

Configuring Static Network Information on the OOB Port

 **NOTE:** Dell EMC Networking N1100-ON, N1500, N2000, and N2100-ON Series switches do not have an out-of-band interface.

Use the following commands to configure a static IP address, subnet mask, and default gateway on the OOB port. If no default gateway is configured, then the zero subnet (0.0.0.0) is used. In this configuration, the OOB port can reach hosts in the local subnet only, because the OOB port will not be able to issue ARP requests to the default gateway. Configuring a default gateway address on the OOB port allows the OOB port to issue ARPs and address traffic to hosts on other subnets; however, if routing is enabled, routing will use the gateway on the OOB port for front-panel ARP requests. The OOB port subnet may not overlap with any in-band VLAN subnet.

| Command | Purpose |
|---|--|
| <code>configure</code> | Enter Global Configuration mode. |
| <code>interface out-of-band</code> | Enter Interface Configuration mode for the OOB port. |
| <code>ip address ip_address subnet_mask [gateway_ip]</code> | Configure a static IP address and subnet mask. Optionally, a default gateway can also be configured. |
| <code>ipv6 address prefix/prefix-length</code> | Configure an IPv6 prefix for the OOB port |
| <code>ipv6 address enable</code> | Enable IPv6 addressing on the OOB port |
| <code>ipv6 address autoconfig</code> | Enable IPv6 auto-configuration for the OOB port |
| <code>ipv6 address dhcp</code> | Enable DHCP address assignment for the OOB port. |
| <code>CTRL + Z</code> | Exit to Privileged Exec mode. |
| <code>show ip interface out-of-band</code> | Verify the network information for the OOB port. |
| <code>show ipv6 interface out-of-band</code> | Verify the IPv6 network information for the OOB port. |

 **NOTE:** The out-of-band port also supports IPv6 address assignment, including IPv6 auto-configuration and an IPv6 DHCP client.

Configuring Static Network Information on the Default VLAN

Use the following commands to configure a static IP address, subnet mask, and default gateway on the default VLAN. Alternatively, a DHCP server may be used to obtain a network address. The switch also supports IPv6 address auto-configuration.

IP subnets on in-band ports (configured on switch VLANs) may not overlap with the OOB port subnet. If configuring management access on the front-panel ports, it is recommended that:

- A VLAN other than the default VLAN be used to avoid attack vectors enabled by incorrect cabling.
- Both ACLs and Management ACLs be utilized on front-panel ports to reduce the possibility of DoS attacks or intruders gaining access to the switch management console. Management ACLs provide software filtering with deep inspection of packets, whereas ACLs provide hardware filtering with a more limited set of capabilities.

| Command | Purpose |
|--|---|
| <code>configure</code> | Enter Global Configuration mode. |
| <code>vlan 10</code> | Create a management VLAN and enter VLAN Configuration mode. |
| <code>exit</code> | Exit VLAN Configuration mode |
| <code>interface vlan 10</code> | Enter Interface Configuration mode for VLAN 10. VLAN 10 is the management VLAN. |
| <code>ip address ip_address subnet_mask</code> | Enter the IP address and subnet mask. |
| <code>ipv6 address prefix/prefix-length [eui64]</code> | Enter the IPv6 address and prefix. |
| <code>ipv6 enable</code> | Enable IPv6 on the interface. |
| <code>exit</code> | Exit to Global Configuration mode |
| <code>ip default-gateway ip_address</code> | Configure the IPv4 default gateway. Only one IPv4 gateway may be configured per switch. |
| <code>ipv6 gateway ip_address</code> | Configure the default gateway for IPv6. Only one IPv6 gateway may be configured per switch. |
| <code>exit</code> | Exit to Privileged Exec mode. |

| Command | Purpose |
|-----------------------------|--|
| show ip interface vlan 10 | Verify the network information for VLAN 10. |
| show ipv6 interface vlan 10 | Verify IPv6 network information for VLAN 10. |
| interface Gi1/0/24 | Enter physical Interface Configuration mode for the specified interface. |
| switchport access vlan 10 | Allow access to the management VLAN over this port. |
| exit | Exit Interface Configuration mode. |

Configuring and Viewing Additional Network Information

Use the following commands to configure a DNS server, the default domain name, and a static host name-to-address entry. Use the **show** commands to verify configured information and to view dynamic host name mappings. Remember to assign VLANs to interfaces.

| Command | Purpose |
|--------------------------------|---|
| configure | Enter Global Configuration mode. |
| ip domain-lookup | Enable IP DNS-based host name-to-address translation. |
| ip name-server ip_address | Enter the IP address of an available name server to use to resolve host names and IP addresses. Up to eight DNS servers may be configured. |
| ip domain-name name | Define a default domain name to complete unqualified host names. |
| ip host name ip_address | Use to configure static host name-to-address mapping in the host cache. |
| ip address-conflict-detect run | Trigger the switch to run active address conflict detection by sending gratuitous ARP packets for IPv4 addresses on the switch. |
| CTRL + Z | Exit to Privileged Exec mode. |
| show ip interface vlan 1 | Verify the network information for VLAN 1. |
| show ipv6 interface vlan 1 | Verify the network information for VLAN 1. |
| show hosts | Verify the configured network information and view the dynamic host mappings. |

| Command | Purpose |
|---|--|
| <code>show ip address-conflict</code> | View the status information corresponding to the last detected address conflict. |
| <code>clear ip address-conflict-detect</code> | Clear the address conflict detection status in the switch. |

Basic Network Information Configuration Examples

Configuring Network Information Using the OOB Port

In this example, an administrator at a Dell office in California decides not to use the Dell Easy Setup Wizard to perform the initial switch configuration. The administrator configures Dell EMC Networking N3000, N3100-ON, and N4000 Series switches to obtain information from a DHCP server on the management network and creates the administrative user with read/write access. The administrator also configures the following information:

- Primary DNS server: 10.27.138.20
- Secondary DNS server: 10.27.138.21
- Default domain name: sunny.dell.com

The administrator also maps the administrative laptop host name to its IP address. The administrator uses the OOB port to manage the switch.

To configure the switch:



NOTE: Dell EMC Networking N1100-ON, N1500, N2000, and N2100-ON Series switches do not have an out-of-band interface.

- 1 Connect the OOB port to the management network. DHCP is enabled by default on the switch OOB interface by default on Dell EMC Networking N3000, N3100-ON, and N4000 Series switches. DHCP is enabled on VLAN 1 on the Dell EMC Networking N1100-ON/N1500/N2000/N2100-ON Series switches, as they do not support an OOB interface. If the DHCP client on the switch has been disabled, use the following commands to enable the DHCP client on the OOB port.

```
console#configure
console(config)#interface out-of-band
console(config-if)#ip address dhcp
console(config-if)#exit
```

- 2 Configure the administrative user.

```
console(config)#username admin password secret123 privilege 15
```

- 3 Configure the DNS servers, default domain name, and static host mapping.

```
console(config)#ip name-server 10.27.138.20 10.27.138.21
console(config)#ip domain-name sunny.dell.com
console(config)#ip host admin-laptop 10.27.65.103
console(config)#exit
```

- 4 View the network information that the DHCP server on the network dynamically assigned to the switch.

```
console#show ip interface out-of-band

IP Address..... 10.27.22.153
Subnet Mask..... 255.255.255.0
Default Gateway..... 10.27.22.1
Protocol Current..... DHCP
Burned In MAC Address..... 001E.C9AA.AA08
```

- 5 View additional network information.

```
console#show hosts

Host name:
Default domain: sunny.dell.com dell.com
Name/address lookup is enabled
Name servers (Preference order): 10.27.138.20, 10.27.138.21
Configured host name-to-address mapping:
Host                               Addresses
-----
admin-laptop                       10.27.65.103

cache: TTL (Hours)
```

```
Host          Total    Elapsed Type          Addresses
-----
No hostname is mapped to an IP address
```

- 6 Verify that the static hostname is correctly mapped.

```
console#ping admin-laptop
Pinging admin-laptop with 0 bytes of data:

Reply From 10.27.65.103: icmp_seq = 0. time <10 msec.
Reply From 10.27.65.103: icmp_seq = 1. time <10 msec.
```

Configuring Network Information Using the Serial Interface

In this example, the administrator configures a Dell EMC Networking N1100-ON/N1500/N2000/N2100-ON Series switch via the serial interface while using the same DHCP server and address configuration as given in the previous example.

- 1 Connect a front-panel port (e.g., `gil/0/24`) to the management network. Use the following commands to create a management VLAN, disable DHCP on VLAN 1, and disable L3 addressing on VLAN 1, and enable the DHCP client on the management VLAN.

```
console#configure
console(config)#vlan 4093
console(config-vlan4093)#interface vlan 1
console(config-if-vlan1)#no ip address
console(config-if-vlan1)#exit
console(config)#no interface vlan 1
console(config-)#interface vlan 4093
console(config-if-vlan4093)#ip address dhcp
```

- 2 Assign the management VLAN to an interface connected to the management network.

```
console(config-if-vlan4093)#interface gil/0/24
console(config-if-Gil/0/24)#switchport access vlan 4093
console(config-if-Gil/0/24)#exit
```

- 3 Configure the administrative user.

```
console(config)#username admin password secret123 privilege 15
```

- 4 Configure the DNS servers, default domain name, and static host mapping.

```
console(config)#ip name-server 10.27.138.20 10.27.138.21
console(config)#ip domain-name sunny.dell.com
console(config)#ip host admin-laptop 10.27.65.103
console(config)#exit
```

- 5 View the network information that the DHCP server on the network dynamically assigned to the switch.

```
console#show ip interface vlan 4093
```

```
Routing interface status..... Up
Primary IP Address..... 10.27.22.150/255.255.252.0
Method..... DHCP
Routing Mode..... Enable
Administrative Mode..... Enable
```



```
Forward Net Directed Broadcasts..... Disable
Proxy ARP..... Enable
Local Proxy ARP..... Disable
Active State..... Active
MAC Address..... 001E.C9DE.B77A
Encapsulation Type..... Ethernet
IP MTU..... 1500
Bandwidth..... 10000 kbps
Destination Unreachables..... Enabled
ICMP Redirects..... Enabled
```

Refer to the Access Control Lists section for information on restricting access to the switch management interface.

Managing QSFP Ports

Dell EMC Networking N3100-ON and N4000 Series Switches

QSFP ports available on Dell EMC Networking N4000 Series switches can operate in 1 x 40G mode or in 4 x 10G mode. Appropriate cables must be used that match the selected mode. When changing from one mode to another, a switch reboot is required. The QSFP ports also support stacking over the interfaces in either 1 x 40G or 4 x 10G mode. Changing from Ethernet mode to stacking mode and vice-versa requires a reboot as well.

Two QSFP ports are available on Dell EMC Networking N3100-ON Series switches in a plug-in module. These port can operate in 1 x 40G mode only and do not support stacking. A separate stacking module is available.

The ports on a QSFP plugin module are named Fo1/1/1-2 in 40-Gigabit mode and Te1/1/1-8 in 10-Gigabit mode. On the N4064, the fixed QSFP ports are named Fo1/0/1-2 in 40-Gigabit mode and Te1/0/49-56 in 10-Gigabit mode. All of the possible populated or configured interfaces will show in the **show interfaces status** command regardless of the port mode, i.e. 40-Gigabit or 10-Gigabit. Unpopulated or unconfigured interfaces for plug in modules do not show in the **show interfaces status** command.

The default setting for a 40-Gigabit Ethernet interface is nonstacking, 40-Gigabit Ethernet (1 x 40G).

The commands to change 1 x 40G and 4 x 10G modes are always entered on the 40-Gigabit interfaces.

The commands to change the Ethernet/stack mode are entered on the appropriate interface (tengigabitethernet or fortygigabitethernet). It is possible to configure some of the 10G ports in a 40G interface as stacking and not others.

To reconfigure a QSFP port, select the 40-Gigabit port to change in Interface Config mode and enter the **hardware profile portmode** command with the selected mode. For example, to change a 1 x 40G port to 4 x 10G mode, enter the following commands on the forty-Gigabit interface:

```
console(config)#interface fo1/1/1  
console(config-if-Fo1/1/2)#hardware profile portmode 4x10g
```

This command will not take effect until the switch is rebooted.

```
console(config-if-Fo1/1/2)#do reload
```

Are you sure you want to reload the stack? (y/n)

To change a 4 x 10G port to 1 x 40G mode, enter the following commands on the 40-Gigabit interface:

```
console(config)#interface Fo2/1/1
console(config-if-Fo2/1/1)#hardware profile portmode 1x40g
```

This command will not take effect until the switch is rebooted.

```
console(config-if-Fo1/1/2)#do reload
```

Are you sure you want to reload the stack? (y/n)

Attempting to change the port mode on the tengigabit interface will give the error “An invalid interface has been used for this function.”

Stacking

Dell EMC Networking N-Series Switches

This chapter describes how to configure and manage a stack of switches.

The topics covered in this chapter include:

- Stacking Overview
- Default Stacking Values
- Managing and Monitoring the Stack (Web)
- Managing the Stack (CLI)
- Stacking and NSF Usage Scenarios

Stacking Overview

The Dell EMC Networking N2000, N2100-ON, N3000, N3000E-ON, N3100-ON, and N4000 Series switches include a stacking feature that allows up to 12 switches to operate as a single unit. Dell EMC Networking N3000 series switches can stack up to eight units as of firmware release 6.5.1. The Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON, and N1500 Series switches stack up to four units using 10GB Ethernet links configured as stacking. Dell EMC Networking N2000, N2100-ON, N3000 and N3000E-ON Series switches have two fixed mini-SAS connectors at the rear for stacking. Dell EMC Networking N3100-ON Series switches have optional an stacking module.

Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON Series switches only stack with other Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON Series switches. Dell EMC Networking N1108P-ON/N1108T-ON Series switches do not stack.

Dell EMC Networking N2000 Series switches stack with other Dell EMC Networking N2000 Series switches, and with Dell EMC Networking N2100-ON Series switches using 21G stacking links.

Dell EMC Networking N3000 Series switches stack with other Dell EMC Networking N3000 Series switches (including the N3000E-ON Series) and Dell EMC Networking N3100-ON Series switches, using the optional

stacking module. Beginning with the 6.5.1 release, any stack containing any N3000 Series switch (other than the N3000E-ON) is limited to a maximum of eight units.

Dell EMC Networking N4000 Series switches stack with other Dell EMC Networking N4000 Series switches over front-panel ports configured for stacking.

Dell EMC Networking N1500 Series switches stack with other N1500 Series switches using the 10G SFP+ front-panel ports.

Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON switches and N1500 Series switches support high-performance stacking over the 10G front-panel ports, allowing increased capacity to be added as needed, without affecting network performance and providing a single point of management. Up to four Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON switches and N1500 Series switches can be stacked using any 10G port as long as the link bandwidth for parallel stacking links is the same. Note that configuring a 10G port for stacking also configures the adjacent partner 10G port for stacking.

A stack of four Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON switches and N1500 Series switches has an aggregate throughput capacity of 192 Gbps. Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON and N1500 Series stacking links operate at 10 Gbps or 5.2% of total aggregate throughput capacity of a full stack; therefore, it is recommended that operators provision large stacking topologies such that it is unlikely that a significant portion of the stack capacity will transit stacking links. One technique for achieving this is to distribute uplinks evenly across the stack vs. connecting all uplinks to a single stack unit or to adjacent stacking units.

A stack of twelve 48-port Dell EMC Networking N2000, N2100-ON, N3048EP-ON, or N3100-ON Series switches has an aggregate throughput capacity of 576 Gbps. Dell EMC Networking N2000/N2100-ON/N3000E-ON/N3100-ON Series stacking links operate at 21 Gbps or 3.6% of total aggregate throughput capacity of a twelve high stack; N2100-ON/N3100-ON stack links can operate at 21 Gbps/40 Gbps; therefore, it is recommended that operators provision large stacking topologies such that it is unlikely that a

significant portion of the stack capacity will transit stacking links. One technique for achieving this is to distribute uplinks evenly across the stack vs. connecting all uplinks to a single stack unit or to adjacent stacking units.

NOTE: Beginning with the 6.5.1 release, any stack containing any N3000 Series switch (other than the N3000E-ON) is limited to a maximum of eight units.

Dell EMC Networking N2100-ON Series switches have two fixed stacking ports in the rear that accept mini-SAS cables. Dell EMC Networking N3100-ON Series switches support an optional 2x21G or 2x40G stacking module in the rear slot.

Dell EMC Networking N4000 Series switches support high performance stacking over front-panel ports, allowing increased capacity to be added as needed, without affecting network performance and providing a single point of management. Up to twelve Dell EMC Networking N4000 Series switches can be stacked using any port as long as the link bandwidth for parallel stacking links is the same. In other words, all the port types on the Dell EMC Networking N4000 Series switches can be used for stacking. Additional stacking connections can be made between adjacent switch units to increase the stacking bandwidth provided that all redundant stacking links have the same port speed. It is strongly recommended that the stacking bandwidth be kept equal across all stacking connections; that is, avoid mixing single and double stacking connections within a stack. Up to eight redundant stacking links operating at the same speed can be configured on a Dell EMC Networking N4000 Series stack unit (four in each direction).

A stack of twelve Dell EMC Networking N4000 Series switches has an aggregate front panel capacity of 5.760 terabits (not including the 40G ports). Provisioning for 5% inter-stack capacity requires 280 Gigabit of bandwidth dedicated to stacking or all four 40G ports plus another twelve 10G ports. Therefore, it is recommended that operators provision large stacking topologies such that it is unlikely that a significant portion of the stack capacity will transit stacking links. One technique for achieving this is to distribute downlinks and transit links evenly across the stack vs. connecting all downlinks/transit links to a single stack unit or to adjacent stacking units.

If Priority Flow Control (PFC) is enabled on any port in a Dell EMC Networking N4000 Series stack, stacking is supported at distances up to 100 meters on the stacking ports. If PFC is not enabled, stacking is supported up

to the maximum distance supported by the transceiver on the stack links. Note that PFC cannot be enabled on stacking ports — the system handles the buffering and flow control automatically.

A single switch in the stack manages all the units in the stack (the stack master), and the stack is managed by using a single IP address. The IP address of the stack does not change, even if the stack master changes.

A stack is created by daisy-chaining stacking links on adjacent units. If available, up to eight links per stack unit can be used for stacking (four in each direction). A stack of units is manageable as a single entity when the units are connected together. If a unit cannot detect a stacking partner on any port enabled for stacking, the unit automatically operates as a standalone unit. If a stacking partner is detected, the switch always operates in stacking mode. One unit in the stack is designated as the stack master. The master manages all the units in the stack. The stack master runs the user interface and switch software, and propagates changes to the member units. To manage a stack using the serial interface, the administrator must connect to the stack master via the **connect** command or by physically connecting the cable to the stack master.

A second switch is designated as the standby unit, which becomes the master if the stack master is unavailable. The unit to be selected as the standby can be manually configured, or the system can select the standby automatically.

When units are in a stack, the following activities occur:


- All units are checked for software version consistency.
- The switch Control Plane is active only on the master. The Control Plane is a software layer that manages system and hardware configuration and runs the network control protocols to set system configuration and state.
- The switch Data Plane is active on all units in the stack, including the master. The Data Plane is the set of hardware components that forward data packets without intervention from a control CPU.
- The running configuration is propagated to all units and the application state is synchronized between the master and standby during normal stacking operation. The startup configuration and backup configuration on the stack members are not overwritten with the master switch configuration.

Dell strongly recommends connecting the stack in a ring topology so that each switch is connected to two other switches. Connecting switches in a ring topology allows the stack to utilize the redundant communication path to each switch. If a switch in a ring topology fails, the stack can automatically establish a new communications path to the other switches. Switches not stacked in a ring topology may split into multiple independent stacks upon the failure of a single switch or stacking link.

Additional stacking connections can be made between adjacent switch units to increase the stacking bandwidth, provided that all redundant stacking links have the same bandwidth. It is strongly recommended that the stacking bandwidth be kept equal across of all stacking connections; that is, avoid mixing single and double stacking connections within a stack. Up to eight redundant stacking links can be configured on a stacking unit (four in each direction).

Figure 9-1 shows a stack with three switches as stack members connected in a ring topology.

Figure 9-1. Connecting a Stack of Switches



The stack in Figure 9-1 has the following physical connections between the switches:

- The lower stacking port on Unit 1 is connected to the upper stacking port on Unit 2.
- The lower stacking port on Unit 2 is connected to the upper stacking port on Unit 3.
- The lower stacking port on Unit 3 is connected to the upper stacking port on Unit 1.

Dell EMC Networking N1124-ON/N1148-ON, N1500, N2000, N2100-ON, N3000, N3000E-ON, N3100-ON, and N4000 Stacking Compatibility

Dell EMC Networking N1100-ON, N1500, and N4000 Series switches do not stack with different Dell EMC Networking Series switches or other Dell EMC Networking switches. Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON Series switches only stack with other Dell EMC Networking N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON Series switches. Dell EMC Networking N1108T-ON/N1108P-ON switches do not stack. Dell EMC Networking N1500 Series switches only stack with other Dell EMC Networking N1500 Series switches.

Dell EMC Networking N2000 Series switches stack with Dell EMC Networking N2100 Series switches. Dell EMC Networking N3100-ON Series switches stack with Dell EMC Networking N3000E-ON switches up to twelve units. The maximum stack size may vary depending on loaded firmware (Adv/AdvLite). Any stack containing an N3000 Series switch must use the AdvLite firmware and is limited to a maximum of eight units. The N3000E-ON Series switches ship with the Adv firmware. Be sure to load AdvLite firmware prior to attempt to connect an N3000E-ON Series switch into a mixed stack with N3000 Series switches.

Dell EMC Networking N3000E-ON Series switches stack with Dell EMC Networking N3000 Series switches up to eight units.

Dell EMC Networking N3000E-ON/N3100-ON Series switches stack with Dell EMC Networking N3000 Series switches up to eight units.

Dell EMC Networking N4000 Series switches only stack with other Dell EMC Networking N4000 Series switches.

How is the Stack Master Selected?

A stack master is elected or re-elected based on the following considerations, in order:

- 1** The switch is currently the stack master.
- 2** The switch has the higher MAC address.
- 3** A unit is selected as standby by the administrator, and a fail over action is manually initiated or occurs due to stack master failure.

In most cases, a switch that is added to an existing stack will become a stack member, and not the stack master. When a switch is added to the stack, one of the following scenarios takes place regarding the management status of the new switch:

- If the switch has the stack master function enabled but another stack master is already active, then the switch changes its configured stack master value to disabled.
- If the stack master function is unassigned and there is another stack master in the system then the switch changes its configured stack master value to disabled.
- If the stack master function is enabled or unassigned and there is no other stack master in the system, then the switch becomes stack master.
- If the stack master function is disabled, the unit remains a non-stack master.

If the entire stack is powered OFF and ON again, the unit that was the stack master before the reboot will remain the stack master after the stack resumes operation.

The unit number for the switch can be manually configured. To avoid unit-number conflicts, one of the following scenarios takes place when a new member is added to the stack:

- If the switch has a unit number that is already in use, then the unit that is added to the stack changes its configured unit number to the lowest unassigned unit number.
- If the added switch does not have an assigned unit number, then the switch sets its configured unit number to the lowest unassigned unit number.
- If the unit number is configured and there are no other devices using the unit number, then the switch starts using the configured unit number.
- If the switch detects that the maximum number of units already exist in the stack making it unable to assign a unit number, then the switch sets its unit number to unassigned and does not participate in the stack.

Adding a Switch to the Stack

When adding a new member to a stack, make sure that only the stack cables, and no network cables, are connected before powering up the new unit. Stack port configuration is stored on the member units. If stacking over Ethernet ports (Dell EMC Networking N1100-ON, N1500 and N4000 Series only), configure the ports on the unit to be added to the stack as stacking ports and power the unit off prior to connecting the stacking cables. Make sure the links are not already connected to any ports of that unit. This is important because if STP is enabled and any links are UP, the STP reconvergence will take place as soon as the link is detected.

After the stack cables on the new member are connected to the stack, the units can be powered up, beginning with the unit directly attached to the currently powered-up unit. Always power up new stack units closest to an existing powered unit first. Do not connect a new member to the stack after it is powered up. Never connect two functional, powered-up stacks together. Hot insertion of units into a stack is not supported.

If a new switch is added to a stack of switches that are powered and running and already have an elected stack master, the newly added switch becomes a stack member rather than the stack master. Use the **boot auto-copy-sw** command on the stack master to enable automatic firmware upgrade of newly added switches. If a firmware mismatch is detected, the newly added switch does not fully join the stack and holds until it is upgraded to the same firmware version as the master switch. After firmware synchronization finishes, the running configuration of the newly added unit is overwritten with the stack master configuration. Stack port configuration is always stored on the local unit and may be updated with preconfiguration information from the stack master when the unit joins the stack.

Information about a stack member and its ports can be preconfigured before the unit is added to the stack. The preconfiguration takes place on the stack master. If there is saved configuration information on the stack master for the newly added unit, the stack master applies the configuration to the new unit; otherwise, the stack master applies the default configuration to the new unit.

Removing a Switch from the Stack

Prior to removing a member from a stack, check that other members of the stack will not become isolated from the stack due to the removal. Check the stack-port error counters to ensure that a stack configured in a ring topology can establish a communication path around the member to be removed.

The main point to remember when removing a unit from the stack is to disconnect all the links on the stack member to be removed. Also, be sure to take the following actions:

- Remove all the STP participating ports and wait to stabilize the STP.
- Remove all the member ports of any Port-Channels (LAGs) so there will not be any control traffic destined to those ports connected to this member.
- Statically re-route any traffic going through this unit.

When a unit in the stack fails, the stack master removes the failed unit from the stack. The failed unit reboots with its original running-config. If the stack is configured in a ring topology, then the stack automatically routes around the failed unit. If the stack is not configured in a ring topology, then the stack may split, and the isolated members will reboot and re-elect a new stack master. No changes or configuration are applied to the other stack members; however, the dynamic protocols will try to reconverge as the topology could change because of the failed unit.

If you remove a unit and plan to renumber the stack, issue a **no member unit** command in Stack Configuration mode to delete the removed switch from the configured stack member information.

How is the Firmware Updated on the Stack?

When adding a new switch to a stack, the Stack Firmware Synchronization feature, if enabled, automatically synchronizes the firmware version with the version running on the stack master per the configuration on the master switch. The synchronization operation may result in either upgrade or downgrade of firmware on the mismatched stack member. Use the **boot auto-copy-sw** command to enable stack firmware synchronization (SFS).

Upgrading the firmware on a stack of switches is the same as upgrading the firmware on a single switch. After downloading a new image by using the File Download page or **copy** command, the downloaded image is distributed to all the connected units of the stack. For more information about downloading

and installing images, see "Images and File Management" on page 525. When copying firmware onto the switch in a stacked configuration, use the `show sfs` and `show version` commands to check the status of stack firmware synchronization prior to a reboot.

What is Stacking Standby?

The standby unit may be preconfigured or automatically selected. If the current stack master fails, the standby unit becomes the stack master. If no switch is preconfigured as the standby unit, the software automatically selects a standby unit from among the existing stack units.

When the failed master resumes normal operation, it joins the stack as a member (not as the master) if the new stack master has already been elected.

The stack master copies its running configuration to the standby unit whenever it changes (subject to some restrictions to reduce overhead). This enables the standby unit to take over the stack operation with minimal interruption if the stack master becomes unavailable.

Operational state synchronization also occurs:

- when the running configuration is saved to the startup configuration on the stack master.
- when the standby unit changes.

What is Nonstop Forwarding?

Networking devices, such as the Dell EMC Networking N-Series switches, are often described in terms of three semi-independent functions called the forwarding plane, the control plane, and the management plane. The forwarding plane forwards data packets and is implemented in hardware. The control plane is the set of protocols that determine how the forwarding plane should forward packets, deciding which data packets are allowed to be forwarded and where they should go. Application software on the stack master acts as the control plane. The management plane is application software running on the stack master that provides interfaces allowing a network administrator to configure the device.

The Nonstop Forwarding (NSF) feature allows the forwarding plane of stack units to continue to forward packets while the control and management planes restart as a result of a power failure, hardware failure, or software fault

on the stack master. This type of operation is called nonstop forwarding (NSF). When the stack master fails, only the switch ASICs and processor on the stack master need to be restarted.

To prevent adjacent networking devices from rerouting traffic around the restarting device, the NSF feature uses the following three techniques:

- 1 A protocol can distribute a part of its control plane across stack units so that the protocol can give the appearance that it is still functional during the restart.
- 2 A protocol may enlist the cooperation of its neighbors through a technique known as graceful restart.
- 3 A protocol may simply restart after the failover if neighbors react slowly enough that they will not normally detect the outage.

The NSF feature enables the stack master unit to synchronize the running-config within 60 seconds after a configuration change has been made.

However, if a lot of configuration changes happen concurrently, NSF uses a back-off mechanism to reduce the load on the switch. In this case, the stack master will attempt resynchronization no more often than once every 120 seconds.

The **show nsf** command output includes information about when the next running-config synchronization will occur.

Initiating a Failover

The NSF feature allows the administrator to initiate a failover using the **initiate failover** command. This method is preferred over the **reload** unit command as it ensures synchronization of the stack master and standby unit.

Initiating a failover reloads the stack master, triggering the standby unit to take over. Before the failover, the stack master pushes application data and other important information to the standby unit. Although the handoff is controlled and causes minimal network disruption, some ephemeral application state is lost, such as pending timers and other pending internal events. Use the **show nsf** command to view the stack checkpoint status prior to reloading a stack member. Do not reload while a checkpoint operation is in progress.

Always check the stack health before failing over to the standby unit. Use the **show switch stack-ports counters** command to verify that the stack ports are up and no errors are present. Resolve any error conditions prior to failing over

a stack master. Use the **show switch stack-ports stack-path** command to verify the reachability of all stack units. If any units are not reachable, the stack may split during a failover.

Checkpointing

Switch applications (features) that build up a list of data such as neighbors or clients can significantly improve their restart behavior by remembering this data across a warm restart. This data can either be stored persistently, as in the case of configuration data, or the stack master can checkpoint this data directly to the standby unit active processes, as in the case of operational data.

Use the **show nsf** command to view the stack checkpoint status prior to reloading a stack member. Do not reload while a checkpoint operation is in progress.

The NSF checkpoint service allows the stack master to communicate startup configuration data to the standby unit in the stack. When the stack selects a standby unit, the checkpoint service notifies applications to start a complete checkpoint. After the initial checkpoint is done, applications checkpoint changes to their data every 120 seconds.



NOTE: The switch cannot guarantee that a standby unit has exactly the same data that the stack master has when it fails. For example, the stack master might fail before the checkpoint service gets data to the standby if an event occurs shortly before a failover.

Table 9-1 lists the applications on the switch that checkpoint data and describes the type of data that is checkpointed.

Table 9-1. Applications that Checkpoint Data


| Application | Checkpointed Data |
|----------------|-------------------------------|
| ARP | Dynamic ARP entries |
| Auto VOIP | Calls in progress |
| Captive Portal | Authenticated clients |
| DHCP server | Address bindings (persistent) |
| DHCP snooping | DHCP bindings database |
| DOT1Q | Internal VLAN assignments |

Table 9-1. Applications that Checkpoint Data

| Application | Checkpointed Data |
|---------------------|---|
| DOT1S | Spanning tree port roles, port states, root bridge, etc. |
| 802.1X | Authenticated clients |
| DOT3ad | Port states |
| IGMP/MLD Snooping | Multicast groups, list of router ports, last query data for each VLAN |
| IPv6 NDP | Neighbor cache entries |
| iSCSI | Connections |
| LLDP | List of interfaces with MED devices attached |
| OSPFv2 | Neighbors and designated routers |
| OSPFv3 | Neighbors and designated routers |
| Route Table Manager | IPv4 and IPv6 dynamic routes |
| SIM | The system's MAC addresses. System up time. IP address, network mask, default gateway on each management interface, DHCPv6 acquired IPv6 address. |
| Voice VLAN | VoIP phones identified by CDP or DHCP (not LLDP) |

Switch Stack MAC Addressing and Stack Design Considerations

The switch stack uses the MAC addresses assigned to the stack master.

 **NOTE:** Each switch is assigned four consecutive MAC addresses. A stack of switches uses the MAC addresses assigned to the stack master.

If the backup unit assumes control due to a stack master failure or warm restart, the backup unit continues to use the original stack master's MAC addresses. This reduces the amount of disruption to the network because ARP and other layer-2 entries in neighbor tables remain valid after the failover to the backup unit.

Stack units should always be connected with a ring topology (or other redundant topology), so that the loss of a single stack link does not divide the stack into multiple stacks. If a stack is partitioned such that some units lose all connectivity to other units, then both parts of the stack start using the same MAC addresses. This can cause severe problems in the network.

If removing the stack master from a stack for use in a different place in the network, make sure to power down the whole stack before redeploying the stack master so that the stack members do not continue to use the MAC address of the redeployed master switch.

NSF Network Design Considerations

A network can be designed to take maximum advantage of NSF. For example, by distributing a LAG's member ports across multiple units, the stack can quickly switch traffic from a port on a failed unit to a port on a surviving unit. When a unit fails, the forwarding plane of surviving units removes LAG members on the failed unit so that it only forwards traffic onto LAG members that remain up. If a LAG is left with no active members, the LAG goes down. To prevent a LAG from going down, configure LAGs with members on multiple units within the stack, when possible. If a stack unit fails, the system can continue to forward on the remaining members of the stack.

If the switch stack performs VLAN routing, another way to take advantage of NSF is to configure multiple “best paths” to the same destination on different stack members. If a unit fails, the forwarding plane removes Equal Cost Multipath (ECMP) next hops on the failed unit from all unicast forwarding table entries. If the cleanup leaves a route without any next hops, the route is deleted. The forwarding plane only selects ECMP next hops on surviving units. For this reason, try to distribute links providing ECMP paths across multiple stack units.


Why is Stacking Needed?

Stacking increases port count without requiring additional configuration. If you have multiple Dell EMC Networking N-Series switches, stacking them helps make management of the switches easier because you configure the stack as a single unit and do not need to configure individual switches.

Default Stacking Values

Stacking is always enabled on Dell EMC Networking N-Series switches.


On the Dell EMC Networking N1100-ON/N1500/N4000 Series switches, by default, the 10G SFP+ ports are in Ethernet mode and must be configured to be used as stacking ports. Ports that are configured in stacking mode show as “detached” in the output of the **show interfaces status** command.


 **NOTE:** N1124T-ON/N1148T-ON/N1124P-ON/N1148P-ON/N1500 10G SFP+ ports may only be configured as stacking in adjacent pairs, e.g. Te1/0/1 and Te1/0/2. If configuring all four ports as stacking, the pairs of stacking links must be connected to the same unit, i.e. both Te1/0/1-2 must connect to a single adjacent stack unit.

Configuring an Ethernet port as a stacking port changes the default configuration of the port. To determine the stacking configuration of a port, use the **show switch stack-ports** command. On the Dell EMC Networking N2000/N2100-ON/N3000 Series switches, there are two fixed stacking ports in the rear of the switch. The N3100-ON supports a pluggable stacking module in the rear. Stacking on Ethernet ports is not supported. The fixed stacking ports show as TwentyGigabitStacking and are abbreviated Tw.

NSF is enabled by default. NSF can be disabled to redirect the CPU resources consumed by data checkpointing; however, this is ill-advised, as checkpointing consumes almost no switch resources. Checkpointing only occurs when a backup unit is elected, so there is no need to disable the NSF feature on a standalone switch. When a new unit is added to the stack, the new unit is given the configuration of the stack, including the NSF setting. OSPF implements a separate graceful restart control that enables NSF for OSPF. OSPF graceful restart is not enabled by default.

Managing and Monitoring the Stack (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring stacking on Dell EMC Networking N1100-ON, N1500, N2000, N2100-ON, N3000, N3100-ON, and N4000 Series switches. For details about the fields on a page, click  at the top of the Dell EMC OpenManage Switch Administrator web page.

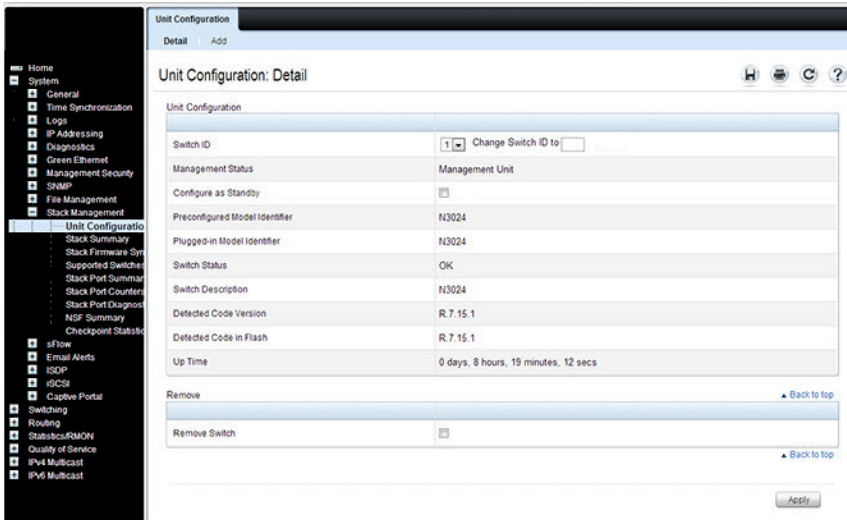
 **NOTE:** Changes made on the Stacking configuration pages take effect only after the device is reset.

Unit Configuration

Use the **Unit Configuration** page to change the unit number and unit type (Management, Member, or Standby).

To display the **Unit Configuration** page, click **System** → **Stack Management** → **Unit Configuration** in the navigation panel. For the N30xx series switches, stack size is limited to 8.

Figure 9-2. Stack Unit Configuration




Changing the ID or Switch Type for a Stack Member

To change the switch ID or type:

- 1 Open the **Unit Configuration** page.
- 2 Click **Add** to display the **Add Unit** page.

For the N30xx series switches, stack size is limited to 8.

Figure 9-3. Add Remote Log Server Settings




- 3 Specify the switch ID, and select the model number of the switch.
- 4 Click **Apply**.

Stack Summary

Use the **Stack Summary** page to view a summary of switches participating in the stack.

To display the **Stack Summary** page, click **System** → **Stack Management** → **Stack Summary** in the navigation panel.

Figure 9-4. Stack Summary

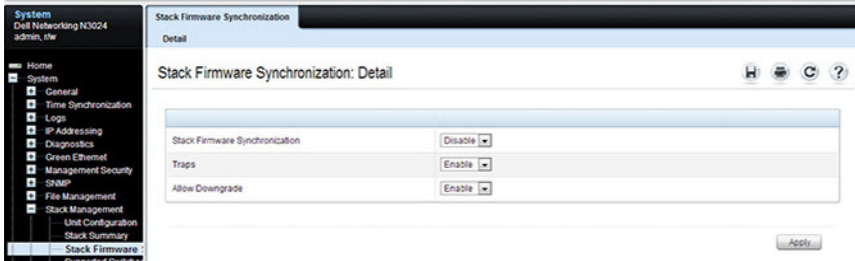


Stack Firmware Synchronization

Use the **Stack Firmware Synchronization** page to control whether the firmware image on a new stack member can be automatically upgraded or downgraded to match the firmware image of the stack master.

To display the **Stack Firmware Synchronization** page, click **System** → **Stack Management** → **Stack Firmware Synchronization** in the navigation panel.

Figure 9-5. Stack Firmware Synchronization

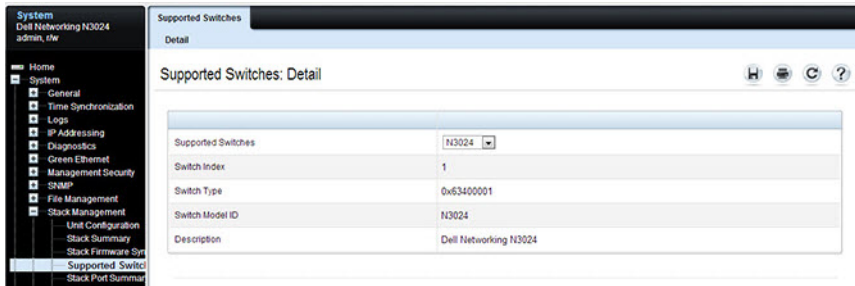


Supported Switches

Use the **Supported Switches** page to view information regarding each type of supported switch for stacking, and information regarding the supported switches.

To display the **Supported Switches** page, click **System** → **Stack Management** → **Supported Switches** in the navigation panel.

Figure 9-6. Supported Switches



Stack Port Summary

Use the **Stack Port Summary** page to configure the stack-port mode and to view information about the stackable ports. This screen displays the unit, the stackable interface, the configured mode of the interface, the running mode as well as the link status and link speed of the stackable port.



NOTE: By default the ports are configured to operate as Ethernet ports. To configure a port as a stack port on the N1124-ON/N1148-ON, N1500, or N4000 Series switches, the Configured Stack Mode setting must be changed from Ethernet to Stack.

To display the **Stack Port Summary** page, click **System** → **Stack Management** → **Stack Port Summary** in the navigation panel.

Figure 9-7. Stack Port Summary

| Interface | Configured Stack-mode | Running Stack-mode | Link Status | Link Speed (Gbps) | Edit |
|-----------|-----------------------|--------------------|-------------|-------------------|--------------------------|
| Tw10/01 | Stack | Stack | Link Down | 21 | <input type="checkbox"/> |
| Tw10/02 | Stack | Stack | Link Down | 21 | <input type="checkbox"/> |

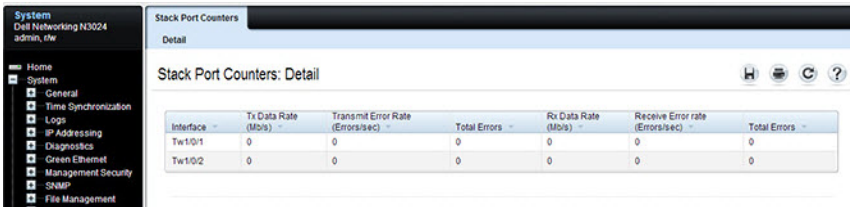
Apply

Stack Port Counters

Use the **Stack Port Counters** page to view the transmitted and received statistics, including data rate and error rate.

To display the **Stack Port Counters** page, click **System** → **Stack Management** → **Stack Port Counters** in the navigation panel.

Figure 9-8. Stack Port Counters



| Interface | Tx Data Rate (Mbps) | Transm Error Rate (Errors/sec) | Total Errors | Rx Data Rate (Mbps) | Receive Error rate (Errors/sec) | Total Errors |
|-----------|---------------------|--------------------------------|--------------|---------------------|---------------------------------|--------------|
| Tw10/1 | 0 | 0 | 0 | 0 | 0 | 0 |
| Tw10/2 | 0 | 0 | 0 | 0 | 0 | 0 |

Stack Port Diagnostics

The **Stack Port Diagnostics** page is intended for Field Application Engineers (FAEs) only.

NSF Summary

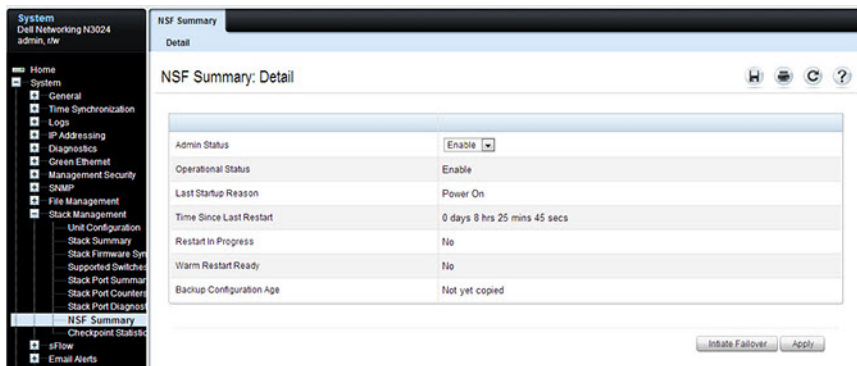
Use the **NSF Summary** page to change the administrative status of the NSF feature and to view NSF information.



NOTE: The OSPF feature uses NSF to enable the hardware to continue forwarding IPv4 packets using OSPF routes while a backup unit takes over stack master responsibility. To configure NSF on a stack that uses OSPF or OSPFv3, see "NSF OSPF Configuration" on page 1283 and "NSF OSPFv3 Configuration" on page 1300.

To display the **NSF Summary** page, click **System** → **Stack Management** → **NSF Summary** in the navigation panel.

Figure 9-9. NSF Summary



To cause the master unit to failover to the standby unit, click **Initiate Failover**. The failover results in a warm restart of the stack master. Initiating a failover reloads the stack master, triggering the backup unit to take over.

Checkpoint Statistics

Use the Checkpoint Statistics page to view information about checkpoint messages generated by the stack master.

To display the Checkpoint Statistics page, click **System** → **Stack Management** → **Checkpoint Statistics** in the navigation panel.

Figure 9-10. Checkpoint Statistics

The screenshot shows the 'Checkpoint Statistics: Detail' page. The left navigation pane is expanded to 'Stack Management' > 'Checkpoint Statistics'. The main content area displays a table with the following data:

| | |
|--------------------------------|------------------------------|
| Messages Checkpointed | 0 |
| Bytes Checkpointed | 0 |
| Time Since Counters Cleared | 0 days 8 hrs 26 mins 29 secs |
| Checkpoint Message Rate | 0.000 msg/sec |
| Last 10-second Message Rate | 0.0 msg/sec |
| Highest 10-second Message Rate | 0.0 msg/sec |

A 'Clear' button is located at the bottom right of the table area.

Managing the Stack (CLI)

This section provides information about the commands for managing the stack and viewing information about the switch stack. For more information about these commands, see the Dell EMC Networking N1100-ON, N1500, N2000, N2100-ON, N3000, N3100-ON, and N4000 Series Switches CLI Reference Guide at www.dell.com/support.

Configuring Stack Member, Stack Port, SFS and NSF Settings

Use the following commands to configure stacking and SFS settings.

| Command | Purpose |
|--|---|
| <code>configure</code> | Enter Global Configuration mode. |
| <code>switch current_ID</code> <code>renumber new_ID</code> | Change the switch ID number. Changing the ID number causes all switches in the stack to be reset to perform stack master renumbering. The running configuration is cleared when the units reset. |
| <code>stack</code> | Enter Global Stack Configuration mode. |
| <code>initiate failover</code> | Move the management switch functionality from the master switch to the standby switch. |
| <code>standby unit</code> | Specify the stack member that will come up as the master if a stack failover occurs. |
| <code>set description unit</code> <code><text></code> | Configure a description for the specified stack member. |

| Command | Purpose |
|---|--|
| <code>member unit SID</code> | <p>Add a switch to the stack and specify the model of the new stack member.</p> <ul style="list-style-type: none"> • <code>unit</code> - The switch unit ID • <code>SID</code> - The index into the database of the supported switch types, indicating the type of the switch being preconfigured. <p>Note: Member configuration displayed in the running config may be learned from the physical stack. Member configuration is not automatically saved in the startup configuration. Save the configuration to retain the current member settings.</p> <p>To view the SID associated with the supported switch types, use the show supported switcheype command in Privileged Exec mode.</p> |
| <code>stack-port {tengigabitethernet twentygigabitethernet fortygigabitethernet} unit/slot/port {ethernet stack shutdown} [speed {40g 21g}]</code> | <p>Set the mode of the port to either Ethernet or stacking (Dell EMC Networking N1124-ON/N1148-ON, N1500 and N4000 Series only). The speed option is only available on the N2100/N3100 Series switches.</p> |
| <code>nsf</code> | <p>Enable nonstop forwarding on the stack. (Enabled by default.)</p> |
| <code>exit</code> | <p>Exit to Global Config mode.</p> |
| <code>boot auto-copy-sw</code> | <p>Enable the Stack Firmware Synchronization feature.</p> |
| <code>boot auto-copy-sw allow-downgrade</code> | <p>Allow the firmware version on the newly added stack member to be downgraded if the firmware version on manager is older. Config migration is not assured for firmware downgrade.</p> |
| <code>exit</code> | <p>Exit to Privileged Exec mode.</p> |
| <code>show auto-copy-sw</code> | <p>View the Stack Firmware Synchronization settings for the stack.</p> |
| <code>reload unit</code> | <p>If necessary, reload the specified stack member.</p> |



NOTE: The OSPF feature uses NSF to enable the hardware to continue forwarding IPv4 packets using OSPF routes while a backup unit takes over stack master responsibility. Additional NSF commands are available in OSPF and OSPFv3 command modes. For more information, see "NSF OSPF Configuration" on page 1283 and "NSF OSPFv3 Configuration" on page 1300

Viewing and Clearing Stacking and NSF Information

Use the following commands to view stacking information and to clear NSF statistics.

| Command | Purpose |
|--|--|
| <code>show switch [stack-member-number]</code> | View information about all stack members or the specified member. |
| <code>show switch stack-standby</code> | View the ID of the switch that will assume the role of the stack master if it goes down. |
| <code>show switch stack-ports</code> | View information about the stacking ports. |
| <code>show switch stack-ports counters</code> | View the statistics about the data the stacking ports have transmitted and received. |
| <code>show switch stack-ports stack-path { <unit> all }</code> | View the path that packets take from one stack member to another. |
| <code>show supported switchtype [<switchindex>]</code> | View the Dell EMC Networking models that are supported in the stack and the switch index (SID) associated with each model. The information may vary, depending on the loaded firmware (Adv/AdvLite). |
| <code>show nsf</code> | View summary information about the NSF state of the master and standby switches. |
| <code>show checkpoint statistics</code> | View information about checkpoint messages generated by the stack master. |
| <code>clear checkpoint statistics</code> | Reset the checkpoint statistics counters to zero. |

Connecting to the Management Console from a Stack Member

From the CLI Unavailable prompt, use the following command to connect the console session to the local unit.

| Command | Purpose |
|-----------------------------|--|
| <code>connect [unit]</code> | Connect the console on the remote unit to the local unit |

Stacking and NSF Usage Scenarios

Only a few settings are available to control the stacking configuration, such as the designation of the standby unit or enabling/disabling NSF. The examples in this section describe how the stacking and NSF feature act in various environments.

This section contains the following examples:

- Basic Failover
- Preconfiguring a Stack Member
- NSF in the Data Center
- NSF and VoIP
- NSF and DHCP Snooping
- NSF and the Storage Access Network
- NSF and Routed Access

Basic Failover

In this example, the stack has four members that are connected in a ring topology, as Figure 9-11 shows.

Figure 9-11. Basic Stack Failover



When all four units are up and running, the `show switch` CLI command gives the following output:

```
console#show switch
```

| SW | Management Status | Standby Status | Preconfig Model ID | Plugged-in Model ID | Switch Status | Code Version |
|----|-------------------|----------------|--------------------|---------------------|---------------|--------------|
| 1 | Stack Member | Opr Stby | N3048 | N3048 | OK | 6.0.0.0 |
| 2 | Stack Member | | N3048 | N3048 | OK | 6.0.0.0 |
| 3 | Mgmt Switch | | N3048 | N3048 | OK | 6.0.0.0 |
| 4 | Stack Member | | N3048 | N3048 | OK | 6.0.0.0 |

At this point, if Unit 2 is powered off or rebooted due to an unexpected failure, `show switch` gives the following output:

```
console#show switch
```

| SW | Management Status | Standby Status | Preconfig Model ID | Plugged-in Model ID | Switch Status | Code Version |
|----|-------------------|----------------|--------------------|---------------------|---------------|--------------|
| 1 | Stack Member | Opr Stby | N3048 | N3048 | OK | 6.0.0.0 |
| 2 | Unassigned | | N3048 | | Not Present | 0.0.0.0 |
| 3 | Mgmt Switch | | N3048 | N3048 | OK | 6.0.0.0 |
| 4 | Stack Member | | N3048 | N3048 | OK | 6.0.0.0 |

When the failed unit resumes normal operation, the previous configuration that exists for that unit is reapplied by the stack master.

To permanently remove the unit from the stack, enter into Stack Config Mode and use the member command, as the following example shows.

```
console#configure
console(config)#stack
console(config-stack)#no member 2
console(config-stack)#exit
console(config)#exit
console#show switch
```

| SW | Management Status | Standby Status | Preconfig Model ID | Plugged-inSwitch Model ID | Status | Code Version |
|----|-------------------|----------------|--------------------|---------------------------|--------|--------------|
| 1 | Mgmt Sw | | N2128PX | N2128PX | OK | 6.3.6.4 |
| 3 | Stack Mbr | | N2128PX | N2128PX | OK | 6.3.6.4 |
| 4 | Stack Mbr | | N2128PX | N2128PX | OK | 6.3.6.4 |

Preconfiguring a Stack Member

To preconfigure a stack member before connecting the physical unit to the stack, use the **show supported switchtype** command to obtain the switch model ID (SID) of the unit to be added.

The example in this section demonstrates pre-configuring a stand-alone Dell EMC Networking N-Series switch.

To configure the switch:

- 1 View the list of SIDs to determine which SID identifies the switch to preconfigure. The following is the output on the switches. The supported switch types vary by switch series and loaded image.

```
console#show supported switchtype
```

| SID | Switch Model ID |
|-----|-----------------|
| 1 | N3024 |
| 2 | N3024F |
| 3 | N3024P |
| 4 | N3048 |
| 5 | N3048P |
| 6 | N3048EP-ON |
| 7 | N3132PX-ON |

The following is the output on Dell EMC Networking N1500 Series switches:

```
console#show supported switchtype
```

```
SID Switch Model ID
-----
1   N1524
2   N1524P
3   N1548
4   N1548P
```

- 2 Preconfigure the switch (SID = 2) as member number 2 in the stack.

```
console#configure
console(config)#stack
console(config-stack)#member 2 2
console(config-stack)#exit
console(config)#exit
```

- 3 Confirm the stack configuration. Some of the fields have been omitted from the following output due to space limitations.

```
console#show switch
```


| Management SW | Standby Status | Preconfig Model ID | Plugged-in Model ID | Switch Status | Code Version |
|---------------|----------------|--------------------|---------------------|---------------|--------------|
| --- | --- | ----- | ----- | ----- | ----- |
| 1 | Mgmt Sw | N3048 | N3048 | OK | 6.0.0.0 |

Preconfigured switches may be removed from the configuration using the **no member** command. Only switches that are not active members of the stack may be removed.

NSF in the Data Center

Figure 9-12 illustrates a data center scenario, where the stack of two Dell EMC Networking N-Series switches acts as an access switch. The access switch is connected to two aggregation switches, AS1 and AS2. The stack has a link from two different units to each aggregation switch, with each pair of links grouped together in a LAG. The two LAGs and link between AS1 and AS2 are members of the same VLAN. Spanning tree is enabled on the VLAN. Assume spanning tree selects AS1 as the root bridge. Assume the LAG to AS1 is the root port on the stack and the LAG to AS2 is discarding. Unit 1 is the stack master. If unit 1 fails, the stack removes the Unit 1 link to AS1 from its LAG. The stack forwards outgoing packets through the Unit 2 link to AS1 during the failover. During the failover, the stack continues to send BPDUs and LAG PDUs on its links on Unit 2. The LAGs stay up (with one remaining link in each), and spanning tree on the aggregation switches does not see a topology change.


Figure 9-12. Data Center Stack Topology



NSF and VoIP

Figure 9-13 shows how NSF maintains existing voice calls during a stack master failure. Assume the top unit is the stack master. When the stack master fails, the call from phone A is immediately disconnected. The call from phone B continues. On the uplink, the forwarding plane removes the failed LAG member and continues using the remaining LAG member. If phone B has learned VLAN or priority parameters through LLDP-MED, it continues to use those parameters. The stack resumes sending LLDPDUs with MED TLVs once the control plane restarts. Phone B may miss an LLDPDU from the stack, but should not miss enough PDUs to revert its VLAN or priority, assuming the administrator has not reduced the LLDPDU interval or hold count. If phone B is receiving quality of service from policies installed in the hardware, those policies are retained across the stack master restart.


Figure 9-13. NSF and VoIP



NSF and DHCP Snooping

Figure 9-14 illustrates a layer-2 access switch running DHCP snooping. DHCP snooping only accepts DHCP server messages on ports configured as trusted ports. DHCP snooping listens to DHCP messages to build a bindings database that lists the IP address the DHCP server has assigned to each host. IP Source Guard (IPSG) uses the bindings database to filter data traffic in hardware based on source IP address and source MAC address. Dynamic ARP Inspection (DAI) uses the bindings database to verify that ARP messages contain a valid sender IP address and sender MAC address. DHCP snooping checkpoints its bindings database.

Figure 9-14. NSF and DHCP Snooping



If the stack master fails, all hosts connected to that unit lose network access until that unit reboots. The hardware on surviving units continues to enforce source filters IPSG installed prior to the failover. Valid hosts continue to communicate normally. During the failover, the hardware continues to drop data packets from unauthorized hosts so that security is not compromised.

If a host is in the middle of an exchange with the DHCP server when the failover occurs, the exchange is interrupted while the control plane restarts. When DHCP snooping is enabled, the hardware traps all DHCP packets to the CPU. The control plane drops these packets during the restart. The DHCP client and server retransmit their DHCP messages until the control plane has resumed operation and messages get through. Thus, DHCP snooping does not miss any new bindings during a failover.

As DHCP snooping applies its checkpointed DHCP bindings, IPSP confirms the existence of the bindings with the hardware by reinstalling its source IP address filters.

If Dynamic ARP Inspection is enabled on the access switch, the hardware traps ARP packets to the CPU on untrusted ports. During a restart, the control plane drops ARP packets. Thus, new traffic sessions may be briefly delayed until after the control plane restarts.

If IPSP is enabled and a DHCP binding is not checkpointed to the backup unit before the failover, that host will not be able to send data packets until it renews its IP address lease with the DHCP server.

NSF and the Storage Access Network


Figure 9-15 illustrates a stack of three Dell EMC Networking N-Series switches connecting two servers (iSCSI initiators) to a disk array (iSCSI targets). There are two iSCSI connections as follows:

Session A: 10.1.1.10 to 10.1.1.3

Session B: 10.1.1.11 to 10.1.1.1

An iSCSI application running on the stack master (the top unit in the diagram) has installed priority filters to ensure that iSCSI traffic that is part of these two sessions receives priority treatment when forwarded in hardware.

Figure 9-15. NSF and a Storage Area Network




When the stack master fails, session A drops. The initiator at 10.1.1.10 detects a link down on its primary NIC and attempts to reestablish the session on its backup NIC to a different IP address on the disk array. The hardware forwards the packets to establish this new session, but assuming the session is established before the control plane is restarted on the backup unit, the new session receives no priority treatment in the hardware.

Session B remains established and fully functional throughout the restart and continues to receive priority treatment in the hardware.

NSF and Routed Access

Figure 9-16 shows a stack of three units serving as an access router for a set of hosts. Two LAGs connect the stack to two aggregation routers. Each LAG is a member of a VLAN routing interface. The stack has OSPF and PIM adjacencies with each of the aggregation routers. The top unit in the stack is the stack master.

Figure 9-16. NSF and Routed Access



If the stack master fails, its link to the aggregation router is removed from the LAG. When the control plane restarts, both routing interfaces come back up by virtue of the LAGs coming up. OSPF sends grace LSAs to inform its OSPF neighbors (the aggregation routers) that it is going through a graceful restart.



NOTE: The graceful restart feature for OSPF is disabled by default. For information about the web pages and commands to configure NSF for OSPF or OSPFv3, see "OSPF and OSPFv3" on page 1257.

The grace LSAs reach the neighbors before they drop their adjacencies with the access router. PIM starts sending hello messages to its neighbors on the aggregation routers using a new generation ID to prompt the neighbors to quickly resend multicast routing information. PIM neighbors recognize the new generation ID and immediately relay the group state back to the restarting router. IGMP sends queries to relearn the hosts' interest in multicast groups. IGMP tells PIM the group membership, and PIM sends

JOIN messages upstream. The control plane updates the driver with checkpointed unicast routes. The forwarding plane reconciles layer-3 hardware tables.

The OSPF graceful restart finishes, and the control plane deletes any stale unicast routes not relearned at this point. The forwarding plane reconciles layer-3 multicast hardware tables. Throughout the process, the hosts continue to receive their multicast streams, possibly with a short interruption as the top aggregation router learns that one of its LAG members is down. The hosts see no more than a 50 ms interruption in unicast connectivity.

Authentication, Authorization, and Accounting

Dell EMC Networking N-Series Switches

This chapter describes how to control access to the switch management interface using authentication and authorization. These services can also be used to restrict or allow network access when used in conjunction with IEEE 802.1x. It also describes how to record this access using accounting. Together the three services are referred to by the acronym AAA.

The topics covered in this chapter include:

- AAA Introduction
- Authentication
- Authorization
- Accounting
- IEEE 802.1X
- Captive Portal

AAA Introduction

AAA is a framework for configuring management security in a consistent way. Three services make up AAA:

- Authentication—Validates the user identity. Authentication takes place before the user is allowed access to switch services.
- Authorization—Determines which services the user is allowed to access. Examples of services are access to the switch management console and access to network services.
- Accounting—Collects and sends security information about switch management console users and switch management commands

Each service is configured using method lists. Method lists define how each service is to be performed by specifying the methods available to perform the service. The first method in a list is tried first. If the first method returns an

error, the next method in the list is tried. This continues until all methods in the list have been attempted. If no method can perform the service, then the service fails. A method may return an error due to lack of network access, misconfiguration of a server, and other reasons. If there is no error, the method returns success if the user is allowed access to the service and failure if the user is not.

AAA gives the user flexibility in configuration by allowing different method lists to be assigned to different access lines. In this way, it is possible to configure different security requirements for the serial console than for Telnet, for example.

Methods

A method performs authentication or authorization for the configured service. Not every method is available for every service. Some methods require a username and password and other methods only require a password.

Table 10-1 summarizes the various methods:

Table 10-1. AAA Methods

| Method | Username? | Password? | Can Return an Error? |
|---------------|------------------|------------------|-----------------------------|
| enable | no | yes | yes |
| ias | yes | yes | no |
| line | no | yes | yes |
| local | yes | yes | yes |
| none | no | no | no |
| radius | yes | yes | yes |
| tacacs | yes | yes | yes |

Methods that never return an error cannot be followed by any other methods in a method list.

- The **enable** method uses the enable password. If there is no enable password defined, then the enable method will return an error.
- The **ias** method is a special method that is only used for 802.1X. It uses an internal database (separate from the local user database) that acts like an 802.1X authentication server. This method never returns an error. It will always pass or deny a user.
- The **line** method uses the password for the access line on which the user is accessing the switch. If there is no line password defined for the access line, then the line method will return an error.
- The **local** method uses the local user database. If the user password does not match, then access is denied. This method returns an error if the user name is not present in the local user database.
- The **none** method does not perform any service, but instead always returns a result as if the service had succeeded. This method never returns an error. If none is configured as a method, the user will always be authenticated and allowed to access the switch.
- The **radius** and **tacacs** methods communicate with servers running the RADIUS and TACACS+ protocols, respectively. These methods can return an error if the switch is unable to contact the server.

Method Lists

The method lists shown in Table 10-2 are defined by default. They cannot be deleted, but they can be modified. Using the “no” command on these lists will return them to their default configuration.

Table 10-2. Default Method Lists

| AAA Service (type) | List Name | List Methods |
|-------------------------|------------------|--------------|
| Authentication (login) | defaultList | none |
| Authentication (login) | networkList | local |
| Authentication (enable) | enableList | enable none |
| Authentication (enable) | enableNetList | enable |
| Authorization (exec) | dfltExecAuthList | none |

Table 10-2. Default Method Lists (Continued)

| AAA Service (type) | List Name | List Methods |
|---------------------------|------------------|---------------------|
| Authorization (commands) | dfltCmdAuthList | none |
| Accounting (exec) | dfltExecList | tacacs (start-stop) |
| Accounting (commands) | dfltCmdList | tacacs (stop-only) |

Access Lines

There are five access lines: console, Telnet, SSH, HTTP, and HTTPS. HTTP and HTTPS are not configured using AAA method lists. Instead, the authentication list for HTTP and HTTPS is configured directly (authorization and accounting are not supported). The default method lists for both the HTTP and HTTPS access lines consist of only the local method. Each of the other access lines may be assigned method lists independently for the AAA services.

The SSH line has built-in authentication beyond that configured by the administrator.

In the SSH protocol itself, there are multiple methods for authentication. These are not the authentication methods configured in AAA, but are internal to SSH itself. When an SSH connection is attempted, the challenge-response method is specified in the connection request.

The methods available for authentication using SSH are: host-based authentication, public key authentication, challenge-response authentication, and password authentication. Authentication methods are tried in the order specified above, although SSH-2 has a configuration option to change the default order.

Host-based SSH authentication is not supported by Dell EMC Networking N-Series switches. Use the Management ACL capability to perform the equivalent function.

Public key SSH authentication operates as follows:

The administrator first generates a pair of encryption keys, the “public” key and the “private” key. Messages encrypted with the private key can be decrypted only by the public key, and vice-versa. The administrator keeps the private key on his/her local machine, and loads the public key on to the switch. When the administrator attempts to log into the switch, the protocol sends a brief message, encrypted with the public key. If the switch can decrypt

the message (and can send back some proof that it has done so) then the response proves that switch must possess the public key, and user is authenticated without giving a username/password.

The public key method is implemented in the Dell EMC Networking N-Series switch as opposed to an external server. If the user does not present a certificate, it is not considered an error and authentication will continue with challenge-response authentication.

Challenge-response SSH authentication works as follows:

The switch sends an arbitrary “challenge” text and prompts for a response. SSH-2 allows multiple challenges and responses; SSH-1 is restricted to one challenge/response only. Examples of challenge-response authentication include BSD Authentication.

Finally, if all other authentication methods fail, SSH prompts the user for a password.

Enabling SSH Access

The following example enables the switch to be accessed using SSH. If RSA or DSA keys exist, the switch will prompt to overwrite the keys as shown below. The RSA and DSA keys are used to negotiate the symmetric encryption algorithm used for the SSH session.

```
console(config)#crypto key generate rsa
Do you want to overwrite the existing RSA keys? (y/n):y
RSA key generation started, this may take a few minutes...
RSA key generation complete.
console(config)#crypto key generate dsa
Do you want to overwrite the existing DSA keys? (y/n):y
DSA key generation started, this may take a few minutes...
DSA key generation complete.
console(config)#ip ssh server
```

Access Lines (AAA)

Table 10-3 shows the method lists assigned to the various access lines by default.

Table 10-3. Default AAA Methods

| AAA Service (type) | Console | Telnet | SSH |
|---------------------------|------------------|------------------|------------------|
| Authentication (login) | defaultList | networkList | networkList |
| Authentication (enable) | enableList | enableList | enableList |
| Authorization (exec) | dfltExecAuthList | dfltExecAuthList | dfltExecAuthList |
| Authorization (commands) | dfltCmdAuthList | dfltCmdAuthList | dfltCmdAuthList |
| Accounting (exec) | none | none | none |
| Accounting (commands) | none | none | none |

Access Lines (Non-AAA)

Table 10-4 shows the default configuration of the access lines that do not use method lists.

Table 10-4. Default Configuration for Non-AAA Access Lines

| Access Line | Authentication | Authorization |
|--------------------|-----------------------|----------------------|
| HTTP | local | n/a |
| HTTPS | local | n/a |
| 802.1X | none | none |

Authentication

Authentication is the process of validating a user's identity. During the authentication process, only identity validation is done. There is no determination made of which switch services the user is allowed to access. This is true even when RADIUS is used for authentication; RADIUS cannot perform separate transactions for authentication and authorization. However, the RADIUS server can provide attributes during the authentication process that are used in the authorization process.

Authentication Types

There are three types of authentication:

- **Login**—Login authentication grants access to the switch if the user credentials are validated. Access is granted only at privilege level one.
- **Enable**—Enable authentication grants access to a higher privilege level if the user credentials are validated for the higher privilege level. When RADIUS is used for enable authentication, the username for this request is always \$enab15\$. The username used to log into the switch is not used for RADIUS enable authentication.
- **802.1X**—802.1X authentication is used to grant an 802.1X supplicant access to the network. For more information about 802.1X, see "Port and System Security" on page 681.

Table 10-5 shows the valid methods for each type of authentication:

Table 10-5. Valid Methods for Authentication Types

| Method | Login | Enable | 802.1x |
|--------|-------|--------|--------|
| enable | yes | yes | no |
| ias | no | no | yes |
| line | yes | yes | no |
| local | yes | no | no |
| none | yes | yes | yes |
| radius | yes | yes | yes |
| tacacs | yes | yes | no |

Authentication Manager

Overview

The Authentication Manager supports the hierarchical configuration of host authentication methods on an interface. Use of the Authentication Manager is optional, but it is recommended when using multiple types of authentication on an interface, e.g., Captive Portal in conjunction with MAB or IEEE 802.1X. Dell switches support the following host authentication methods:

- IEEE 802.1x
- MAC Authentication Bypass (MAB)
- Captive portal

Using the Authentication Manager, the administrator can configure an authentication method list on a per-port basis. Authentication can be enabled or disabled. If authentication is disabled, then no authentication method is applied and the port is provided with open access. The default behavior is that authentication is disabled for all ports.

The configured authentication methods are attempted in the configured order. If an authentication method times out (an error), then the next configured method is attempted. If an authentication method fails, i.e., an incorrect password was entered, then the next method is not attempted and authentication begins again from the first method. If all the methods return an error, then the Authentication Manager starts a timer for reauthentication. The value of the timer is equal to the re-authentication restart timer. Failure in this context means that host authentication was attempted and the host was unable to successfully authenticate. At the expiry of the timer, the Authentication Manager starts the authentication process again from the first method in the list.

The Authentication Manager supports configuring a priority for each authentication method on a port. The authentication priority allows a higher priority method (not currently running) to interrupt an authentication in progress with a lower-priority method. If a client is already authenticated, an interrupt from a higher-priority method can cause a client previously authenticated using a lower priority method to reauthenticate.

By default, Dell switches are configured with a method list that contains the methods (in order) 802.1x, MAB as the default methods for all the ports. Dell switches restrict the configuration such that no method is allowed to follow the Captive Portal method, if configured.

The authentication manager controls only the order in which the authentication methods are executed. The switch administrator is responsible for implementing the required configuration for the respective methods to authenticate successfully.

Authentication Restart

Authentication restarts from the first configured method on any of the following events:

- Link flap
- Authentication fails for all configured methods
- Authentication priority (802.1X packet received when a lower priority method is active)

802.1X Interaction

By default, 802.1X drops all traffic (other than LLDP/CDP) prior to successful 802.1X (or MAB) authentication. If Captive Portal is configured as a method, authentication allows certain traffic types, such as DHCP or DNS, access to the network during the Captive Portal method invocation.

Authentication Priority

The default authentication priority of a method is equivalent to its position in the order of the authentication list. If authentication method priorities are not configured, then the relative priorities (first is highest) are in the same order as that of the per-port based authentication list.

Authentication priority allows a higher-priority method (not currently running) to interrupt an authentication in progress with a lower-priority method. Alternatively, if the client is already authenticated, an interrupt from a higher-priority method can cause a client, which was previously authenticated using a lower-priority method, to reauthenticate.

For example, if a client is already authenticated using a method other than 802.1X (MAB or Captive Portal) and 802.1X has higher priority than the authenticated method, and if an 802.1X frame is received, then the existing

authenticated client is removed and the authentication process begins again from the first method in the order. If 802.1X has a lower priority than the authenticated method, then the client is not removed and the 802.1X frames are ignored.

If administrator changes the priority of the methods, then all the users who are authenticated using a lower-priority method are forced to reauthenticate. If an authentication session is in progress and the administrator changes the order of the authentication methods, then the configuration will take effect for the next session onwards.

Configuration Example—802.1X and MAB

In this scenario, the authentication manager selects the first authentication method, 802.1X. If authentication using 802.1X is successful, then the client is allowed network access. If authentication using 802.1X errors out, then authentication manager selects the next authentication method: MAB. If authentication using MAB returns an error, then the port is unauthorized. The authentication manager will start a timer to re-authenticate the client. At the expiry of the timer, the authentication manager restarts authentication by selecting the 802.1X method.

- 1 Enter global configuration mode and define the RADIUS server.

```
console#configure
console(config)#aaa new-model
console(config)#radius server auth 10.10.10.10
console(config-auth-radius)#name BigRadius
console(config-auth-radius)#primary
console(config-auth-radius)#usage 802.1x
console(config-auth-radius)#exit
```

- 2 Define the global RADIUS server key.

```
console(config)#radius server key thatsyoursecret-keepit-keepit
```

- 3 Enable authentication and globally enable 802.1x client authentication via RADIUS:

```
console(config)#authentication enable
console(config)#aaa authentication dot1x default radius
console(config)#dot1x system-auth-control
```

- 4 On the interface, enable MAC based authentication mode, enable MAB, and set the order of authentication to 802.1X followed by MAC authentication. Configure the switch to send CHAP attributes to the RADIUS server. Set the format of the User-Name sent to the RADIUS server to XXXX.XXXX.XXXX. Also enable periodic re-authentication.

```
console(config)#mab request format attribute 1 groupsize 4
separator . uppercase
console(config)#vlan 2
console(config-vlan2)#interface gil/0/4
console(config-if-Gil/0/4)#switchport mode general
console(config-if-Gil/0/4)#switchport general pvid 2
console(config-if-Gil/0/4)#dot1x port-control mac-based
console(config-if-Gil/0/4)#mab
console(config-if-Gil/0/4)#default mab chap
console(config-if-Gil/0/4)#authentication order dot1x mab
console(config-if-Gil/0/4)#dot1x reauthentication
console(config-if-Gil/0/4)#exit
```

Configuration Example—MAB Client

This example shows how to configure a MAB client on interface Gil/0/2 using the IAS database for authentication.

- 1 Enter global configuration mode and create VLAN 3.

```
console#configure
console(config)#configure
console(config)#vlan 3
console(config-vlan3)#exit
```

- 2 Enable the authentication manager and globally enable 802.1x.

```
console(config)#authentication enable
console(config)#dot1x system-auth-control
```

- 3 Set IEEE 802.1x to use the local IAS user database.

```
console(config)#aaa authentication dot1x default ias
```

- 4 Configure the IAS database with the client MAC address as the user name and password. The password MUST be entered in upper case or the authentication will fail with an MD5 Validation Failure, as the MD5 password hashes would not match.

```
console(config)#aaa ias-user username F8B1562BA1D9
console(config-ias-user)#password F8B1562BA1D9
console(config-ias-user)#exit
```

- 5 Configure interface gi1/0/2 to use VLAN 3 in general mode. General mode is required for MAC-based authentication.

```
console(config)#interface Gi1/0/2
console(config-if-Gi1/0/2)#switchport mode general
console(config-if-Gi1/0/2)#switchport general pvid 3
```

- 6 On the interface, configure the port to use MAC based authentication and enable MAB. The authentication manager is configured to only use MAB and the priority is set to MAB.

```
console(config-if-Gi1/0/2)#dot1x port-control mac-based
console(config-if-Gi1/0/2)#mab
console(config-if-Gi1/0/2)#authentication order mab
console(config-if-Gi1/0/2)#authentication priority mab
console(config-if-Gi1/0/2)#exit
```

If it is possible that an 802.1x aware client may be connected, it is advisable to configure a re-authentication timer on the port using the `dot1x timeout re-authperiod` command.

The following command shows the 802.1x configuration on the interface:

```
console(config-if-Gi1/0/1)#show dot1x interface gi1/0/2
```

```
Administrative Mode..... Enabled
Dynamic VLAN Creation Mode..... Disabled
VLAN Assignment Mode..... Disabled
Monitor Mode..... Disabled
```

| Port | Admin Mode | Oper Mode | Reauth Control | Reauth Period |
|---------|------------|------------|----------------|---------------|
| ----- | ----- | ----- | ----- | ----- |
| Gi1/0/2 | mac-based | Authorized | FALSE | 3600 |

```
Quiet Period..... 60
Transmit Period..... 30
Maximum Requests..... 2
Max Users..... 64
Guest-vlan Timeout..... 90
Server Timeout (secs)..... 30
MAB mode (configured)..... Enabled
MAB mode (operational)..... Enabled
```

| Logical Supplicant Filter | AuthPAE | Backend | VLAN Username |
|---------------------------|-------------|---------|---------------|
| Port | MAC-Address | State | State |
| Id | | | Id |

```
-----  
-----  
64      F8B1.562B.A1D9 Authenticated      Idle      3  
F8B1562BA1D9
```

```
console(config-if-Gi1/0/1)#show dot1x clients all
```

```
Clients Authenticated using Monitor Mode..... 0  
Clients Authenticated using Dot1x..... 1  
Interface..... Gi1/0/2  
User Name..... F8B1562BA1D9  
Supp MAC Address..... F8B1.562B.A1D9  
Session Time..... 1240  
Filter Id.....  
DACL Name.....  
RADIUS Framed IPv4/IPv6 address.....  
VLAN Assigned..... 3
```

Using RADIUS

The RADIUS client on the switch supports multiple RADIUS servers. When multiple authentication servers are configured, they can help provide redundancy. One server can be designated as the primary and the other(s) will function as backup server(s). The switch attempts to use the primary server first. If the primary server does not respond, the switch attempts to use the backup servers. A priority value can be configured to determine the order in which the backup servers are contacted.

How Does RADIUS Control Management Access?

Many networks use a RADIUS server to maintain a centralized user database that contains per-user authentication information. RADIUS servers provide a centralized authentication method for:

- Network Access (IEEE 802.1X)
- User Manager (Management access)
- Captive Portal


Like TACACS+, RADIUS access control utilizes a database of user information on a remote server. Making use of a single database of accessible information—as in an Authentication Server—can greatly simplify the authentication and management of users in a large network. One such type of Authentication Server supports the Remote Authentication Dial In User Service (RADIUS) protocol as defined by RFC 2865.

For authenticating users, the RADIUS standard has become the protocol of choice by administrators of large networks. To accomplish the authentication in a secure manner, the RADIUS client and RADIUS server must both be configured with the same shared password or “secret”. This “secret” is used to generate one-way encrypted authenticators that are present in all RADIUS packets. The “secret” is never transmitted over the network.

RADIUS conforms to a secure communications client/server model using UDP as a transport protocol. It is extremely flexible, supporting a variety of methods to authenticate and statistically track users. RADIUS is also extensible, allowing for new methods of authentication to be added without disrupting existing functionality.

As a user attempts to connect to the switch management interface, the switch first detects the contact and prompts the user for a name and password. The switch encrypts the supplied information, and a RADIUS client transports the request to a preconfigured RADIUS server.

Figure 10-1. RADIUS Topology



The server can authenticate the user itself or make use of a back-end device to ascertain authenticity. In either case a response may or may not be forthcoming to the client. If the server accepts the user, it returns a positive result with attributes containing configuration information. If the server rejects the user, it returns a negative result. If the server rejects the client or the shared secrets differ, the server returns no result. If the server requires additional verification from the user, it returns a challenge, and the request process begins again.

If using a RADIUS server to authenticate users, the RADIUS administrator must configure user attributes in the user database on the RADIUS server. The user attributes include the user name, password, and privilege level.



NOTE: To set the user privilege level at login, it is required that the Service-Type attribute be used for RADIUS instead of the Cisco AV pair priv-lvl attribute. The Cisco AV priv-lvl is supported only for TACACS authorization.

Which RADIUS Attributes Does the Switch Support?

Table 10-6 lists the RADIUS attributes that the switch supports and indicates whether the 802.1X feature, User Manager feature, or Captive Portal feature supports the attribute. The RADIUS administrator must configure these attributes on the RADIUS server(s) when utilizing the switch RADIUS service, and may also need to enable processing of the specific attribute on the switch.

Table 10-6. Supported RADIUS Attributes

| Type | RADIUS Attribute Name | 802.1X | User Manager | Captive Portal |
|------|-----------------------|------------|--------------|----------------|
| 1 | User-Name | Yes | Yes | No |
| 2 | User-Password | Yes | Yes | No |
| 3 | CHAP-Password | Yes | No | No |
| 4 | NAS-IP-Address | Yes | Yes | No |
| 5 | NAS-Port | Yes | No | No |
| 6 | Service-Type | Yes | Yes | No |
| 8 | Framed-IP-Address | Auth. only | Yes | No |
| 11 | Filter-Id | Yes | No | No |
| 12 | Framed-MTU | Yes | No | No |
| 15 | Login-Service | No | Yes | No |
| 18 | Reply-Message | Auth. only | Yes | No |
| 24 | State | Yes | Yes | No |
| 25 | Class | Yes | Yes | No |
| 26 | Vendor-Specific | Yes | Yes | Yes |
| 27 | Session-Timeout | Yes | No | Yes |
| 28 | Idle-Timeout | No | No | Yes |
| 29 | Termination-Action | Yes | No | No |
| 30 | Called-Station-Id | Yes | No | No |
| 31 | Calling-Station-Id | Yes | No | Yes |
| 32 | NAS-Identifier | No | Yes | No |
| 40 | Acct-Status-Type | Acct. only | Yes | No |

Table 10-6. Supported RADIUS Attributes (Continued)

| Type | RADIUS Attribute Name | 802.1X | User Manager | Captive Portal |
|------|------------------------|------------|--------------|----------------|
| 41 | Acct-Delay-Time | Acct. only | No | No |
| 42 | Acct-Input-Octets | Yes | No | No |
| 43 | Acct-Output-Octets | Yes | No | No |
| 44 | Acct-Session-Id | Acct. only | Yes | No |
| 46 | Acct-Session-Time | Yes | Yes | No |
| 49 | Acct-Terminate-Cause | Yes | No | No |
| 52 | Acct-Input-Gigawords | Yes | No | No |
| 53 | Acct-Output-Gigawords | Yes | No | No |
| 61 | NAS-Port-Type | Yes | No | Yes |
| 64 | Tunnel-Type | Yes | No | No |
| 65 | Tunnel-Medium-Type | Yes | No | No |
| 79 | EAP-Message | Yes | No | No |
| 80 | Message-Authenticator | Auth. only | Yes | No |
| 81 | Tunnel-Privategroup-Id | Yes | No | No |
| 168 | Framed-IPv6-Address | Acct. only | No | No |

How Are RADIUS Attributes Processed on the Switch?

The following attributes are processed in the RADIUS Access-Accept message received from a RADIUS server:

- **REPLY-MESSAGE**
Trigger to respond to the Access-Accept message with an EAP notification.
- **STATE**
RADIUS server state. Transmitted in Access-Request messages.
- **SERVICE-TYPE**
The Service-Type attribute may be validated in the Access-Accept packet received from the RADIUS server. Only the Login-User(1), Administrative-User(6), and Call-Check(10) values are considered valid

for Service-Type in the Access-Accept message returned from the RADIUS server.

- **SESSION-TIMEOUT**
Session time-out value for the session (in seconds). Used by both 802.1x and Captive Portal.
- **TERMINATION-ACTION**
Indication as to the action taken when the service is completed.
- **EAP-MESSAGE**
Contains an EAP message to be sent to the user. This is typically used for MAB clients.
- **VENDOR-SPECIFIC**
The following Cisco AV Pairs are supported:
 - shell:priv-lvl
 - shell:roles
 - ip:inacl={standard-access-control-list-name | extended-access-control-list-name}
 - ipv6:inacl={standard-access-control-list-name | extended-access-control-list-name}
 - ip:inacl[#number]={extended-access-control-list}
 - ip:outacl[#number]={extended-access-control-list}
 - ipv6:inacl[#number]={extended-access-control-list}
 - ipv6:outacl[#number]={extended-access-control-list}
 - ip:traffic-class={existing ACL name}
- **FILTER-ID**
Name of an existing ACL or DiffServ policy for this user.
- **FRAMED-IP-ADDRESS**
The IP address assigned to the host accessing the network. Cached and transmitted in accounting packets.
- **FRAMED-IPv6-ADDRESS**

The IPv6 address assigned to the host accessing the network. Cached and transmitted in accounting packets.

- **TUNNEL-TYPE**

Used to indicate that a VLAN is to be assigned to the user when set to tunnel type VLAN (13).

- **TUNNEL-MEDIUM-TYPE**

Used to indicate the tunnel medium type. Must be set to medium type 802 (6) to enable VLAN assignment.

- **TUNNEL-PRIVATE-GROUP-ID**

Used to indicate the VLAN to be assigned to the user. May be a string which matches a preconfigured VLAN name or a VLAN ID. If a VLAN ID is given, the string must contain only decimal digits.


Using TACACS+ Servers to Control Management Access

TACACS+ (Terminal Access Controller Access Control System) provides access control for networked devices via one or more centralized servers. TACACS+ simplifies authentication by making use of a single database that can be shared by many clients on a large network. TACACS+ uses TCP to ensure reliable delivery and a shared key configured on the client and daemon server to encrypt all messages.

If TACACS+ is configured as the authentication method for user login and a user attempts to access the user interface on the switch, the switch prompts for the user login credentials and requests services from the TACACS+ client. The client then uses the configured list of servers for authentication, and provides results back to the switch.

Figure 10-2 shows an example of access management using TACACS+.

Figure 10-2. Basic TACACS+ Topology



The TACACS+ server list can be configured with one or more hosts defined via their network IP addresses. Each can be assigned a priority to determine the order in which the TACACS+ client will contact the servers. TACACS+ contacts the server when a connection attempt fails or times out for a higher priority server.

Each server host can be configured with a specific connection type, port, timeout, and shared key, or the server hosts can be globally configured with the key and timeout.

The TACACS+ server can do the authentication itself, or redirect the request to another back-end device. All sensitive information is encrypted and the shared secret is never passed over the network; it is used only to encrypt the data.

Which TACACS+ Attributes Does the Switch Support?

Table 10-7 lists the TACACS+ attributes that the switch supports and indicates whether the authorization or accounting service supports sending or receiving the attribute. The authentication service does not use attributes. The following attributes can be configured on the TACACS+ server(s) when utilizing the switch TACACS+ service.

Table 10-7. Supported TACACS+ Attributes

| Attribute Name | Exec Authorization | Command Authorization | Accounting |
|-----------------------|---------------------------|------------------------------|-------------------|
| cmd | both (optional) | sent | sent |
| cmd-arg | | sent | |
| elapsed-time | | | sent |
| priv-lvl | received | | |
| protocol | | | sent |
| roles | both (optional) | | |
| service=shell | both | sent | sent |
| start-time | | | sent |
| stop-time | | | sent |

Dynamic ACL Overview

NOTE: This feature is only supported in 802.1X auto mode configuration.

Dynamic ACLs allow operators to administer bespoke network access policies from a central location (the RADIUS server). Access policies are enforced via the use of ACLs installed for the duration of the user session. Unique policies can be assigned based upon the user credentials/location/time of day and other information presented to the RADIUS server during the authentication process. The benefit to the end user is that the policy can follow the user around the network, regardless of where the network is accessed. The benefit to the network administrator is that policy can be configured once for the user and does not need to be configured on multiple devices.

IEEE 802.1X auto mode ports may be configured to accept 802.1X authentication for both the data VLAN and voice VLAN. In this case, both authentications may contain DACL references or definitions. The DACLs are applied and removed for each authentication session independently of the other sessions, however, the DACLs are applied at the port level and are capable of filtering any matching ingress traffic, regardless of which authentication session actually instantiated the DACL.

Note that 802.1X auto mode ports are restricted to a single data device and a single voice device by default. This restriction is enforced by implicitly filtering incoming traffic based upon the MAC address of the authenticating client.

DACLs contained in an 802.1X re-authentication Access-Accept replace the DACLs instantiated in the existing session. DACLs are never applied to hosts authenticated into the Guest or Unauthenticated VLAN. DACLs are compatible with RADIUS VLAN assignment.

Filter-ID Support

The switch supports the association of preconfigured access-lists to an 802.1X authenticated port as presented in the IETF Filter-ID (11) RADIUS attribute (RFC 2865) in an Access-Accept message if configured to accept same. The port must be configured in 802.1X auto mode. If DACL capability is not enabled, or the port is not configured for 802.1X auto mode, Filter-ID attributes are ignored (as if they are not present in the message) and authentication proceeds in the normal manner. Other RADIUS attributes (forexample, Tunnel-Medium-Type, Tunnel-Type, Tunnel-Private-Group-ID, and so on) are processed in the normal manner. The named ACL must exist on the switch and can be of any ACL type (MAC, IPv4, or IPv6).

When the identified ACL is applied, all statically-configured ACLs on the port are removed and the new ACL is configured prior to 802.1X authorizing the port. When the 802.1X session terminates, the dynamic ACL is removed and the pre-existing ACLs are restored to the port.

If no Access list exists matching the Filter-ID, the Access-Accept is treated as an Access-Reject and the port is not authorized. A log message indicating same is issued (Interface X/X/X not authorized. Filter-ID XXXX selected by server x.y.z.x is not present on switch) . No Acct-Start packet is sent and an EAP-Failure is sent to the 802.1X client. Note that the name in a Filter-ID may be a number of an ACL in the form of <ACL#.in>, such as 100.in. If both a Filter-ID and a Cisco AV-Pair (26) are present in the Access-Accept, the Access-Accept is treated as an Access-Reject and the port is not authorized. A log message indicating same is issued (Interface X/X/X not authorized. RADIUS Access-Accept/COA-Request contains both Filter-ID(11)and AV-Pair(26)attributes) . No Acct-Start packet is sent and an EAP-Failure is sent to the 802.1X client.

Dynamic ACLs using the Filter-ID syntax are always enabled.

Filter-ID syntax:

Named ACL - printable character string of the form <ACLNAME>, <Direction>, for example, Filter-id="test_static.in"

Filter-ID example:

Named_ACL - printable character string of the form Filter-id="test_static.in"

Preconfigured or Dynamic ACLs

The switch also supports the application of preconfigured ACLs or the configuration and application of dynamically-created Access Lists to an 802.1X authenticated port as presented in a series of Cisco VSA (009/001) av-pair (26) attributes in a RADIUS Access-Accept. If dynamic ACL capability is not enabled, VSA 26 attributes are ignored as if they are not present in the message and authentication proceeds in the normal manner. Other RADIUS attributes (for example, Tunnel-Medium-Type, Tunnel-Type, Tunnel-Private-Group-ID, and so on) are processed in the normal manner.

Dynamic ACLs using the VSA AV-Pair syntax may be enabled by configuring the **radius server vsa send authentication** command.

The switch will configure the rules in IPv4 or IPv6 Extended Access Lists named IP-DACL-IN-<port id>#d where <port-id> is the user presentable short form port name, such as Tel0/1. The corresponding IPv6 naming convention is IPV6-DACL-IN-<port-id>#d. DACLs for Voice VLAN are named IP-V-DACL-IN-<port id>#d. Note that the # sign is not an acceptable character for an ACL name which prevents the DACL from being edited or removed via the UI. The original ACL, if any, is restored to the port after the 802.1X session terminates. Only ingress ACLs are supported.

If there is an error applying the ACL to the port, a WARN log message indicating same is issued (Interface X/X/X not authorized. Application of downloaded ACL XXX did not complete due to resource exhaustion) and the Access-Accept is treated as an Access-Reject. The port is not authorized. Any previously configured ACLs are added back to the port. If Accounting is enabled, the Acct-Start packet is not sent and an EAP-Failure packet is sent to the 802.1X client.

The VSA av-pair is coded as follows: Attribute 26, Vendor ID 9, Vendor type 9.

Predefined or Dynamic ACL Selection

Send the following Cisco VSA (009/001) av-pair (26) attribute syntax from the RADIUS server in the Access-Accept message to select an ACL that is already configured on the switch. The ACL must be preconfigured on the switch. The extended-access-control-list-name is the name or number of an existing ACL. The standard-access-control-list-name is the number of an existing ACL. The ACL need not be statically preconfigured on the port prior to RADIUS configuring the ACL when authorizing the port. All statically-configured ACLs on a port are removed prior to configuring the dynamic ACL and authorizing the port. The ACL applied is considered state, not configuration and is not shown in the running-config.

Syntax

```
ip:inacl={standard-access-control-list-name | extended-access-control-list-name }
```

```
ipv6:inacl={standard-access-control-list-name | extended-access-control-list-name }
```

- The ip before the colon indicates an existing IPv4 ACL name or number follows the equals sign.
- The ipv6 before the colon indicates an IPv6 ACL name or number follows the equals sign.
- The token standard-access-control-list-name means a Dell EMC Standard ACL identified by the decimal number after the equals sign.
- The token extended-access-control-list-name means a Dell EMC IP/IPv6 Extended ACL identified by the decimal number or the name of an preconfigured ACL. The range numbers are not restricted to ranges as in other vendor implementations.
- The tokens ip:inacl and ipv6:inacl are in lower case and are followed by an equals sign with no intervening white space.

Predefined ACL Examples

```
ip:inacl=Named_ACL
```

```
ipv6:inacl=Named_IPv6_ACL
```

Dynamic ACL Creation

Send the following Cisco VSA (009/001) av-pair (26) attribute syntax from the RADIUS server in the Access-Accept message to create an ACL that does not exist on the switch. The ACL need not be statically preconfigured on the port prior to RADIUS creating the ACL and authorizing the port. All statically configured ACLs on a port are removed prior to configuring the dynamic ACL. The ACL applied is considered state, not configuration and is not shown in the running-config.

Syntax

```
ip:inacl[#number]={extended-access-control-list}  
ipv6:inacl[#number]={ extended-access-control-list}
```

- where ip indicates an IPv4 ACL definition follows the equals sign and ipv6 indicates an IPv6 ACL definition follows the equals sign.
- #number is the ACL sequence number in decimal format. Range 1-2147483647.
- The tokens ip:inacl and ipv6:inacl are in lower case and are followed by an equals sign with no intervening white space.
- The token extended-access-control-list means a Dell EMC IPv4/IPv6 Extended ACL CLI rule definition beginning with the {permit|deny} tokens followed by the protocol { eigrp | gre | icmp | igmp | ip | ipinip | ospf | pim | tcp | udp | 0-55} et. seq., as described in the CLI Reference Guide for the permit/deny commands.

Dynamic ACL Example (Extended syntax, for example, ip access-list extended ...):

```
ip:inacl#100=permit ip any 209.165.0.0 0.0.255.255  
ip:inacl#110=permit ip any 209.166.0.0 0.0.255.255  
ip:inacl=permit ip any 209.167.0.0 0.0.255.255
```

Restrictions and Caveats

Only ingress ACLs are supported. Dynamic ACLs are supported only for ports in General or Access mode when configured in 802.1X auto mode.

The processing of dynamic ACLs VSAs is controlled by the [no] radius server vsa send authentication syntax. The default is disabled. No other VSAs (such as, voice VLAN) are affected by this configuration.

Either traffic-class av-pairs or multiple ip:inacl/ipv6:inacl av-pairs may be present in the RADIUS message, but not both. If both are present, or there are syntax errors in the received ACLs (other than duplicate rules), the ACL rules are not applied, the RADIUS Access-Accept is treated as an Access-Reject, and a WARN log message or Interface X/X/X not authorized. Application of downloaded ACL did not complete due to invalid syntax XXXXX is issued indicating that a received RADIUS rule is misconfigured with invalid syntax or configured with both ip:traffic-class and in acl rules, and identifying the RADIUS server and the affected interface. If Accounting is enabled, the Acct-Start packet is not sent. An EAP-Failure is sent to the 802.1X client.

The VSAs may appear in any order in the RADIUS message. A mixture of in/out and IPv4/IPv6 rules may be present in the RADIUS message to be parsed into the four two Access-Groups. Rules are separated by newlines (either CR or CR/LF). Upper and lower case shall be accepted. The strings ip:traffic-class, ip:inacl, ... are always in lower case. The optional digits following the # symbol indicate the ACL number in the access list.

The rules are applied in the order they appear in the RADIUS packet (the ACL numbers indicate the relative internal priority). Duplicate entries (identical number) in the Access-Accept message follow the same behavior as exists in the UI today (overwrite the previous entry). Conflicting rules are handled in the same manner as if configured via the CLI.

RADIUS-supplied dynamic ACLS are applied at the access-group level after removing all statically configured access groups/traffic filters on the port and before any policies specified in Filter-ID. The following order is observed for application of the access-groups: IPv6-DACL-IN, IP-DACL-IN, IPv6-V-DACL-IN, IP-V-DACL-IN. Empty rules sets are not applied to the port. The words statically configured access-groups do not include denial of service or storm control configurations as they use different internal hardware.

The dynamic ACLS exist only for the duration of the 802.1X session. They are removed when the 802.1X session is terminated (including for COA bounce-host-port or COA termination requests) or when the port goes down (unplugged or shut down). Any static ACLs previously removed from the port are restored when the last 802.1X session ends. Note that the port is unauthorized when the session ends, so the static rules are not actually written into hardware. They are available for application if the RADIUS server does not send an ACL or the port otherwise becomes authorized. The

administrator can override the port configuration and add a manually configured ACL. If the administrator adds an ACL, only the DACL is removed when the session ends.

The switch does not alter the dynamic ACL IP address filter; IP source addresses in the DACL are not altered to use the supplicant IP address.

The dynamic ACL is supported for 802.1X auto mode for a port configured in access or general mode. Dynamic ACLs are ignored/rejected on ports configured for multi-session or MAC-based mode. The Access-Accept is treated as an Access-Reject and the port is not authorized. A log message indicating same is issued (`Interface x/x/x not authorized. Dynamic ACL XXXX not supported in 802.1X MAC-based mode`). No Acct-Start packet is sent and an EAP-Failure is sent to the 802.1X client.

Only one dynamic IPv4 ACL and one dynamic IPv6 ACL may be associated with an 802.1X session (for a total of two access-groups per 802.1X session). Only two named ACLs (one IPv4 and one IPv6) are supported (for a total of two access groups per 802.1X session) per received Access-Accept.

Dynamic ACLs are supported for ports configured in 802.1X Monitor Mode. Syntax errors are logged in the Monitor Mode log. Monitor mode behavior is not altered, for example, if sufficient information to allow access the host to the port is present, the host is allowed access to the port.

Dynamic ACLs are subject to the same hardware scale limitations as static ACLs. If the ACL cannot be applied (resource limitation), then the Access-Accept is treated as an Access-Reject and the port is not authorized. A log message indicating same is issued (`Interface x/x/x not authorized. ACL received from RADIUS server exceeds available resources`). No Acct-Start packet is sent and an EAP-Failure is sent to the 802.1X client.

Dynamic ACLs may not exceed the size of a single RADIUS Access-Accept packet. There is no support for multiple packet ACLs. (Max dynamic ACL is 4000 ASCII characters). There is no support for Downloadable ACLs where the NAS sends a request to the RADIUS server to retrieve an ACL.

Authentication Examples

It is important to understand that during authentication, all that happens is that the user is validated. If any attributes are returned from the server, they are not processed during authentication. In the examples below, it is assumed that the default configuration of authorization—that is, no authorization—is used.

Local Authentication Example

Use the following configuration to require local authentication when logging in over a Telnet connection:

- 1 Create a login authentication list called “loc” that contains the method local:

```
console#config  
console(config)#aaa authentication login "loc" local
```

- 2 Enter the configuration mode for the Telnet line:

```
console(config)#line telnet
```

- 3 Assign the loc login authentication list to be used for users accessing the switch via Telnet:

```
console(config-telnet)#login authentication loc  
console(config-telnet)#exit
```

- 4 Allow Telnet and SSH users access to Privileged Exec mode. It is required that an enable password be configured to allow local access users to elevate to privileged exec level:

```
console(config)#enable password PaSSW0rd
```

- 5 Create a user with the name “guest” and password “password”. A simple password can be configured here, since strength-checking has not yet been enabled:

```
console(config)#username guest password password
```

- 6 Set the minimum number of numeric characters required when password strength checking is enabled. This parameter is enabled only if the **passwords strength minimum character-classes** parameter is set to something greater than its default value of 0:

```
console(config)#passwords strength minimum numeric-characters  
2
```

- 7 Set the minimum number of character classes that must be present in the password. The possible character classes are: upper-case, lower-case, numeric and special:

```
console(config)#passwords strength minimum character-classes 4
```

- 8 Enable password strength checking:

```
console(config)#passwords strength-check
```

- 9 Create a user with the name “admin” and password “paSS1&word2”. This user is enabled for privilege level 15. Note that, because password strength checking was enabled, the password was required to have at least two numeric characters, one uppercase character, one lowercase character, and one special character:

```
console(config)#username admin password paSS1&word2 privilege 15
```

- 10 Configure the switch to lock out a local user after three failed login attempts:

```
console(config)#passwords lock-out 3
```

This configuration allows either user to log into the switch. Both users will have privilege level 1. If no enable password was configured, neither user would be able to successfully execute the **enable** command, which grants access to Privileged Exec mode, because there is no enable password set by default (the default method list for Telnet enable authentication is only the “enable” method).



NOTE: It is recommended that the password strength checking and password lockout features be enabled when configuring local users.

RADIUS Authentication Example

Use the following configuration to require RADIUS authentication to login over a Telnet connection:

- 1 Create a login authentication list called “rad” that contains the method radius. If this method returns an error, the user will fail to login:

```
console#config  
console(config)#aaa authentication login "rad" radius
```

- 2 Create an enable authentication list called “raden” that contains the method radius. If this method fails, then the user will be unable to execute the enable command:

```
console(config)#aaa authentication enable "raden" radius
```

- 3 The following command is the first step in defining a RADIUS authentication server at IP address 1.2.3.4. The `automate-tester username` parameter is a dummy User ID that is NOT configured on the RADIUS server, and is used to verify server liveness. The result of this command is to place the user in radius server configuration mode to allow further configuration of the server:

```
console(config)#radius server auth 1.2.3.4  
console(config-auth-radius)#name Radius-Server  
console(config-auth-radius)#automate-tester username  
DummyLogin idle-time 30
```

- 4 Define the shared secret. This must be the same as the shared secret defined on the RADIUS server:

```
console(config-auth-radius)#key "secret"  
console(config-auth-radius)#exit
```

- 5 Enter the configuration mode for the Telnet line:

```
console(config)#line telnet
```

- 6 Assign the rad login authentication method list to be used for users accessing the switch via Telnet:

```
console(config-telnet)#login authentication rad
```

- 7 Assign the raden enable authentication method list to be used for users executing the enable command when accessing the switch via Telnet:

```
console(config-telnet)#enable authentication raden  
console(config)#exit
```


ACL Using Authentication Manager to Configure MAB with RADIUS Server

The following is a relatively complex example of using an ACL to control access to Gi1/0/1, using the Authentication Manager to configure MAB in conjunction with a RADIUS server.

- 1 Create VLAN 60 which will be used for management access via Gi1/0/1:

```
console#config
console(config)#vlan 60
console(config-vlan60)#exit
```

- 2 Enable the authentication manager:

```
console(config)#authentication enable
```

- 3 Create an access list limiting IP communication exclusively to host 172.25.129.299. All other IP addresses are excluded. This address is in the Bogons address space:

```
console(config)#ip access-list RADIUSCAP
console(config-ip-acl)#permit ip any 172.25.129.229 0.0.0.0
console(config-ip-acl)#permit ip 172.25.129.229 0.0.0.0 any
console(config-ip-acl)#deny ip any any
console(config-ip-acl)#permit every
console(config-ip-acl)#exit
```

- 4 Set a default gateway for the switch:

```
console(config)#ip default-gateway 172.25.128.254
```

- 5 Set a default route with administrative distance 253:

```
console(config)#ip route 0.0.0.0 0.0.0.0 172.25.128.254 253
```

- 6 Assign an IP address to the management VLAN:

```
console(config)#interface vlan 60
console(config-vlan60)#ip address 172.25.128.214 255.255.0.0
console(config-vlan60)#exit
```

- 7 Enable 802.1x client authentication via RADIUS and allow VLAN assignment to 802.1x clients:

```
console(config)#dot1x system-auth-control
console(config)#aaa authentication dot1x default radius
console(config)#aaa authorization network default radius
```

- 8 Allow 802.1x client VLANs to be dynamically created via RADIUS:

```
console(config)#dot1x dynamic-vlan enable
```

- 9 Configure the primary RADIUS server:

```
console(config)#radius server auth 172.25.129.229
```

```
console(config-auth-radius)#name Default-Radius-Server
console(config-auth-radius)#primary
console(config-auth-radius)#usage 802.1x
console(config-auth-radius)#key "dellSecret"
console(config)#exit
```

- 10** Configure the management interface and bypass 802.1x authentication for the connected management host:

```
console(config)#interface Gi1/0/1
console(config-if-Gi1/0/1)#switchport access vlan 60
console(config-if-Gi1/0/1)#dot1x port-control force-authorized
console(config-if-Gi1/0/1)#ip access-group RADIUSCAP in 1
console(config)#exit
```

- 11** Configure a dedicated printer port. This port is enabled for MAB only. The VLAN is assigned by the RADIUS server:

```
console(config)#interface Gi1/0/21
console(config-if-Gi1/0/21)#switchport mode general
console(config-if-Gi1/0/21)#dot1x port-control mac-based
console(config-if-Gi1/0/21)#mab
console(config-if-Gi1/0/21)#authentication order mab
console(config-if-Gi1/0/21)#authentication priority mab
console(config-if-Gi1/0/21)#exit
```

- 12** Configure a port for 802.1x access using MAB. This port will periodically re-authenticate connected clients using the configured timer values. The selected timer values are intended to reduce the time required to authenticate:

```
console(config)#interface Gi1/0/22
console(config-if-Gi1/0/22)#switchport mode general
console(config-if-Gi1/0/22)#dot1x port-control mac-based
console(config-if-Gi1/0/22)#dot1x reauthentication
console(config-if-Gi1/0/22)#dot1x timeout quiet-period 10
console(config-if-Gi1/0/22)#dot1x timeout re-authperiod 300
console(config-if-Gi1/0/22)#dot1x timeout tx-period 7
console(config-if-Gi1/0/22)#dot1x timeout guest-vlan-period 5
console(config-if-Gi1/0/22)#dot1x timeout server-timeout 6
console(config-if-Gi1/0/22)#mab
console(config-if-Gi1/0/22)#exit
```

Combined RADIUS, CoA, MAB and 802.1x Example

The following example configures RADIUS in conjunction with IEEE 802.1X to provide network access to switch clients.

- 1 Enable 802.1x:

```
console#config  
console(config)#dot1x system-auth-control
```

- 2 Configure 802.1x clients to use RADIUS services:

```
console(config)#aaa authentication dot1x default radius
```

- 3 Enable CoA for RADIUS:

```
console(config)#aaa server radius dynamic-author
```

- 4 Configure the remote RADIUS server for COA requests at 10.130.191.89 with “shared secret” as the key:

```
console(config-radius-da)#client 10.130.191.89 server-key  
“shared secret”
```

- 5 Specify that any CoA request with a matching key identifies a client:

```
console(config-radius-da)#auth-type any  
console(config-radius-da)#exit
```

- 6 Configure a group of RADIUS clients (switches) to appear as a single large RADIUS client (by using the same NAS-IP-Address):

```
console(config)#radius server attribute 4 10.130.65.4
```

- 7 Specify that the RADIUS server for host authentication/network access is located at 10.130.191.89:

```
console(config)#radius server auth 10.130.191.89  
console(config-auth-radius)#name Default-RADIUS-Server
```

- 8 Configure the RADIUS shared secret as “shared secret”:

```
console(config-auth-radius)#key “shared secret”  
console(config-auth-radius)#exit
```

- 9 Configure Gi1/0/7 to use MAC based authentication. This allows multiple hosts sharing the same network port to be individually allowed or denied access to network resources. CoA requests to terminate a host session can be issued by the RADIUS server. This means that if the RADIUS server terminates the host session and subsequently refuses to authorize the host (based upon the MAC address), the host is denied access to the network:

```
console(config)#interface Gi1/0/7  
console(config-if-Gi1/0/7)#dot1x port-control mac-based
```

```
console(config-if-Gi1/0/7)#exit
```

- 10 Configure Gi1/0/6 to allow connected hosts access to network resources, regardless of RADIUS configuration. RADIUS CoA disconnect requests are ignored for clients on this port:

```
console(config)#interface Gi1/0/6
console(config-if-Gi1/0/6)#dot1x port-control force-authorized
console(config-if-Gi1/0/6)#exit
```

- 11 Configure Gi1/0/5 to use standard 802.1x authentication:

```
console(config)#interface Gi1/0/5
console(config-if-Gi1/0/5)#dot1x port-control auto
console(config-if-Gi1/0/5)#exit
```

TACACS+ Authentication Example

Use the following configuration to require TACACS+ authentication when logging in over a Telnet connection:

- 1 Create a login authentication list called “tacplus” that contains the method tacacs. If this method returns an error, the user will fail to login:

```
console#config
console(config)#aaa authentication login "tacplus" tacacs
```

- 2 Create an enable authentication list called “tacp” that contains the method tacacs. If this method fails, then the user will fail to execute the enable command:

```
console(config)#aaa authentication enable "tacp" tacacs
```

- 3 The following command is the first step in defining a TACACS+ server at IP address 1.2.3.4. The result of this command is to place the user in tacacs-server mode to allow further configuration of the server:

```
console(config)#tacacs-server host 1.2.3.4
```

- 4 Define the shared secret. This must be the same as the shared secret defined on the TACACS+ server:

```
console(config-tacacs)#key "secret"
console(config-tacacs)#exit
```

- 5 Enter the configuration mode for the Telnet line.

```
console(config)#line telnet
```

- 6 Assign the tacplus login authentication method list to be used for users accessing the switch via Telnet:

```
console(config-telnet)#login authentication tacplus
```

- 7 Assign the tacp enable authentication method list to be used for users executing the enable command when accessing the switch via Telnet:

```
console(config-telnet)#enable authentication tacp
console(config-telnet)#exit
```



NOTE: A user logging in with this configuration would be placed in User Exec mode with privilege level 1. To access Privileged Exec mode with privilege level 15, use the enable command.



NOTE: Dell EMC Networking TACACS supports setting the maximum user privilege level in the authorization response. Configure the TACACS server to send `priv-lvl=X`, where X is either 1 (Non-privileged mode), or 15 (Privileged mode).

Public Key SSH Authentication Example

The following is an example of a public key configuration for SSH login. Using a tool such as `putty` and a private/public key infrastructure, one can enable secure login to the Dell EMC Networking N-Series switch without a password. Instead, a public key is used with a private key kept locally on the administrator's computer. The public key can be placed on multiple devices, allowing the administrator secure access without needing to remember multiple passwords. It is strongly recommended that the private key be protected with a password.

This configuration requires entering a public key, which can be generated by a tool such as `PuTTYgen`. Be sure to generate the correct type of key. In this case, we use an RSA key with the SSH-2 version of the protocol.

Switch Configuration

- 1 Create a switch administrator:

```
console#config
console(config)#username "admin" password
f4d77eb781360c5711ecf3700a7af623 privilege 15 encrypted
```

- 2 Set the login and enable methods for line to NOAUTH.

```
console(config)#aaa authentication login "NOAUTH" line
console(config)#aaa authentication enable "NOAUTH" line
```

- 3 Generate an internal RSA key. This step is not required if an internal RSA key has been generated before on this switch:

```
console(config)#crypto key generate rsa
```

- 4 Set SSH to use a public key for the specified administrator login. The user login is specified by the `username` command, not the `ias-user` command:

```
console(config)#crypto key pubkey-chain ssh user-key "admin"
rsa
```

- 5 Enter the public key obtained from a key authority or from a tool such as PuTTYGen. This command is entered as a single line, not as multiple lines as it appears in the following text.

```
console(config-pubkey-key)#key-string row
AAAAB3NzaC1yc2EAAAABJQAAAIBor6DPjYDpSy8Qcji68xrs/4Lf8c9Jq4xXKI
Z5Pvv20AkRFE0ifVI9EH4jyZagR3wzH5Xl9dyja6bTuqMgN15C1xJC1159FU88
JaY7ywGdRppmoaJrNRPM7RZtQPdVIunzm3eMr9PywwQ0umsHWGNexUrDYHFWR
IAmJp689AAxw==
console(config)#exit
```

- 6 Set the line method to SSH:

```
console(config)#line ssh
```

- 7 Configure the authentication method to the networkList. The networkList contains a single method — local — which is equivalent to password authentication. Since the authentication is provided by the public key, a second layer of authentication is not required:

```
console(config-ssh)#login authentication networkList
console(config-ssh)#exit
```

- 8 The following three lines enable the SSH server, configure it to use public key authentication, and specify use of the SSH-2 protocol.

```
console(config)#ip ssh server
console(config)#ip ssh pubkey-auth
console(config)#ip ssh protocol 2
```

The following command shows the configured authentication methods:

```
console (config)#show authentication methods
```

```
Login Authentication Method Lists
```

```
-----
defaultList      : none
networkList      : local
NOAUTH           : line
```

```
Enable Authentication Method Lists
```

```
-----
enableList       : enable  none
enableNetList    : enable
NOAUTH           : line
```

| Line | Login Method List | Enable Method List |
|---------|-------------------|--------------------|
| Console | defaultList | enableList |
| Telnet | networkList | enableList |


```
SSH      defaultList      enableList

HTTPS    :local
HTTP     :local
DOT1X    :
```


PUTTY Configuration

Main Screen


On the following screen, the IP address of the switch is configured and SSH is selected as the secure login protocol.




On the next screen, PUTTY is configured to use SSH-2 only. This is an optional step that accelerates the login process.




The following screen is the key to the configuration. It is set to display the authentication banner, disable authentication with Pageant, disable keyboard-interactive authentication (unless desired), disable attempted changes of user name, and select the private key file used to authenticate with the switch.



The following screen configures the user name to be sent to the switch. A user name is always required. Alternatively, leave Auto-login name blank and the system will prompt for a user name.



After configuring Putty, be sure to save the configuration. The following screen shows the result of the login process. The user name is entered automatically and the switch confirms that public key authentication occurs.



Authenticating with a Public Key from Linux

The following example configures the switch to allow administrative access without a password for Linux users with correctly configured SSH clients. Dell EMC Networking SSH is configured to require a password on administrator accounts. This example shows how to generate a public/private key pair on Linux, configure Linux SSH, and configure the switch to authenticate SSH connections.

- 1 Log in to your Linux account and generate the RSA key pair. DSA keys are considered weak.

```
ssh-keygen -t rsa
```

- 2 In the `~/.ssh` subdirectory in your Linux account, create an SSH configuration file "ssh_config" with the following contents:

```
User admin
PubkeyAuthentication yes
IdentityFile /home/jmclendo/.ssh/id_rsa
```

Substitute the login ID of the switch administrator for the User admin parameter above, and set the correct path to your account for the IdentityFile parameter.

- 3 On the switch, generate the ephemeral encryption keys to enable the SSH server to run, create the admin user, and configure the SSH server and the authentication key as shown below, making the appropriate substitutions for the login ID:

```
console(config)#crypto key generate rsa
Do you want to overwrite the existing RSA keys? (y/n):y
RSA key generation started, this may take a few minutes...
RSA key generation complete.
console(config)#crypto key generate dsa
Do you want to overwrite the existing DSA keys? (y/n):y
DSA key generation started, this may take a few minutes...
DSA key generation complete.
console(config)#username "admin" password
5f4dcc3b5aa765d61d8327deb882cf99
privilege 15 encrypted
console(config)#ip ssh server
console(config)#ip ssh pubkey-auth
console(config)#ip ssh protocol 2
console(config)#line ssh
console(config-ssh)#login authentication networkList
console(config-ssh)#exit
console(config)#crypto key pubkey-chain ssh user-key admin rsa
console(config-pubkey-key)#Key-String "ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAQEAvcChaxFl4sMoWMZAAwtX/pcVbljY6moer3C
T231M47dgZDPFJ
1qf7/fuDwmES72FmIJAqq8cTufT55BrI0r3vk05QJu0nnhcNjW6c98mNL9wxfx
7TWybySs3zJJpS
NhcZ9JM+OJ104n4oS4izIzY7NSSNa+LQgg5j0mw9jdITY8SicImenLCjluILrp
i6YA9WtC9RHGpi
xLzIRFQ/Kmf5SWcXiSRft4gUJP7Xp69SF3VAAuoUFQove5RMr6paLXUizfwzDk
HA8F4WHaDyHCtx
ESLXnZuQQjCiowl18Q2Nq5YXnu/ZEUJTyof1Uc8S13aP2rr+6Ndzbn6khBmSSg
QnVw==
jsmith@x1-rtp-02"
console(config-pubkey-key)#exit
```

The Key-String above is the contents of the ~/.ssh/id_rsa.pub file enclosed in quotes. This file was generated by the ssh-keygen command as shown above.


Also, ensure that the private key `~/.ssh/id_rsa` is not readable by others by executing the `chmod 0600 ~/.ssh/id_rsa` command in Linux. Authentication will fail if the file is readable by others.

The command string to log into the switch (substituting the correct IP address) from a Linux account is:

```
ssh -2 -i ~/.ssh/id_rsa -F ~/.ssh/ssh_config 10.27.21.70
```

Authenticating Without a Public Key

When authenticating without the public key, the switch prompts for the user name and password. This is an SSH function, not a switch function. If the user knows the administrator login and password, then they are able to authenticate in this manner.



Associating a User With an SSH Key

The following example shows how to associate a user with an externally generated SSH key. SSH, RSA, or DSA keys can be generated by using the `ssh-keygen` command on a Unix system or with other publicly available utilities.

- 1 Create the local user:

```
console#config
console(config)#username mylogin password XXXXXXXX privilege
15
```

- 2 Enter the externally generated key:

```
console(config)#crypto key pubkey-chain ssh
```

- 3 Associate the key with the newly added user login:

```
console(config-pubkey-chain)#user-key mylogin dsa
```

- 4 Add the externally generated key. All of the key information is entered between double quotes.

```
console(config-pubkey-key)#key-string "ssh-dss
AAAAAB3NzaC1kc3MAAACBAJRwUAD3AuRACp1MObeh1AgyZb18wf9Btdip+t+1C
bAqiQNEh4lBiew184DSKk0T6SnSSXuCN+bJnQPXJeiQt+OFnmjiYhnHcvI04Q5
KnQhloZcEFgSsmQ7zJnReWtLvUQI0QvBISTanzedmQVGHvDrQ5X2R729ToSH0i
bBrnYtAAAAFQDNord7S9EJvUkKKxVBpWE6/skCmQAAAIBMjMO+BPP5KXzNwfZz
qAhxBSobvif/z6pzi9xWLLy99A03zmRYCpcGIoLWiRHsR7NVpxFqwqbqvez8KS
0CDJ5aoKKLrpBlpg5ETkYEew/utZ14lQQRBrzPwGBfxvTXKCWiI2j5KFa/WKLS
nmWJX0/98qpxW/LMXoXsA9iK4pnMKwAAAIB4Jrt6jmoLybpzqOPOI0DsJ7jQwW
acinD0j1lz8k+qzCpanhd2Wh+DEdJ/xO2sFRfnYlME3hmXoB+7NByVUtheVjuQ
2CWhcGFIKm9tbuPC6DtXh1xxT0NJ7rspvLgb0s6y/0tk+94ZP5RCoAtLZ7wirs
hy3/KJ4RE0y2SFZjIVjQ=="
```

```
console(config-pubkey-key)#exit
console(config-pubkey-chain)#exit
console(config)#exit
```

- 5 Use the following command to show the user and SSH association:

```
console#show crypto key pubkey-chain ssh username mylogin
Username : mylogin
ssh-dss
AAAAAB3NzaC1kc3MAAACBAJRwUAD3AuRACp1MObeh1AgyZb18wf9Btdip+t+1C
bAqiQNEh4lBiew184DSKk0T6SnSSXuCN+bJnQPXJeiQt+OFnmjiYhnHcvI04Q5
KnQhloZcEFgSsmQ7zJnReWtLvUQI0QvBISTanzedmQVGHvDrQ5X2R729ToSH0i
bBrnYtAAAAFQDNord7S9EJvUkKKxVBpWE6/skCmQAAAIBMjMO+BPP5KXzNwfZz
qAhxBSobvif/z6pzi9xWLLy99A03zmRYCpcGIoLWiRHsR7NVpxFqwqbqvez8KS
0CDJ5aoKKLrpBlpg5ETkYEew/utZ14lQQRBrzPwGBfxvTXKCWiI2j5KFa/WKLS
nmWJX0/98qpxW/LMXoXsA9iK4pnMKwAAAIB4Jrt6jmoLybpzqOPOI0DsJ7jQwW
acinD0j1lz8k+qzCpanhd2Wh+DEdJ/xO2sFRfnYlME3hmXoB+7NByVUtheVjuQ
2CWhcGFIKm9tbuPC6DtXh1xxT0NJ7rspvLgb0s6y/0tk+94ZP5RCoAtLZ7wirs
hy3/KJ4RE0y2SFZjIVjQ==
Fingerprint : d9:d1:21:ad:26:41:ba:43:b1:dc:5c:6c:b9:57:07:6c
```

Authorization

Authorization is used to determine which services the user is allowed to access. For example, the authorization process may assign a user's privilege level, which determines the set of commands the user can execute. There are three kinds of authorization: commands, exec, and network.

- **Commands:** Command authorization determines which CLI commands the user is authorized to execute.
- **Exec:** Exec authorization determines what the user is authorized to do on the switch; that is, the user's privilege level and an administrative profile.
- **Network:** Network authorization enables a RADIUS server to assign a particular 802.1X supplicant to a VLAN. For more information about 802.1X, see "Port and System Security" on page 681.

Table 10-8 shows the valid methods for each type of authorization:

Table 10-8. Authorization Methods

| Method | Commands | Exec | Network |
|--------|----------|------|---------|
| local | no | yes | no |
| none | yes | yes | no |
| radius | no | yes | yes |
| tacacs | yes | yes | no |

Exec Authorization Capabilities

Dell EMC Networking N-Series switches support two types of service configuration with exec authorization: privilege level and administrative profiles.

Privilege Level

By setting the privilege level during exec authorization, a user can be placed directly into Privileged Exec mode when they log into the command line interface.

Administrative Profiles

The Administrative Profiles feature allows the network administrator to define a list of rules that control the CLI commands available to a user. These rules are collected in a “profile.” The rules in a profile can define the set of commands, or a command mode, to which a user is permitted or denied access.

Within a profile, rule numbers determine the order in which the rules are applied. When a user enters a CLI command, rules within the first profile assigned to the user are applied in descending order until there is a rule that matches the input. If no rule permitting the command is found, then the other profiles assigned to the user (if any) are searched for rules permitting the command. Rules may use regular expressions for command matching. All profiles have an implicit “deny all” rule, such that any command that does not match any rule in the profile is considered to have been denied by that profile.

A user can be assigned to more than one profile. If there are conflicting rules in profiles, the “permit” rule always takes precedence over the “deny” rule. That is, if any profile assigned to a user permits a command, then the user is permitted access to that command. A user may be assigned up to 16 profiles.

A number of profiles are provided by default. These profiles cannot be altered by the switch administrator. See "Administrative Profiles" on page 321 for the list of default profiles.

If the successful authorization method does not provide an administrative profile for a user, then the user is permitted access based upon the user's privilege level. This means that, if a user successfully passes enable authentication or if exec authorization assigns a privilege level, the user is permitted access to all commands. This is also true if none of the administrative profiles provided are configured on the switch. If some, but not all, of the profiles provided in the authentication are configured on the switch, then the user is assigned the profiles that exist, and a message is logged that indicates which profiles could not be assigned.

The administrative profiles shown in Table 10-9 are system-defined and may not be deleted or altered. To see the rules in a profile, use the **show admin-profiles name** profile name command.

Table 10-9. Default Administrative Profiles

| Name | Description |
|------------------|---|
| network-admin | Allows access to all commands. |
| network-security | Allows access to network security features such as 802.1X, Voice VLAN, Dynamic ARP Inspection and IP Source Guard. |
| router-admin | Allows access to Layer 3 features such as IPv4 Routing, IPv6 Routing, OSPF, RIP, etc. |
| multicast-admin | Allows access to multicast features at all layers, this includes L2, IPv4 and IPv6 multicast, IGMP, IGMP Snooping, etc. |
| dhcp-admin | Allows access to DHCP related features such as DHCP Server and DHCP Snooping. |
| CP-admin | Allows access to the Captive Portal feature. |
| network-operator | Allows access to all User Exec mode commands and show commands. |

Authorization Examples

Authorization allows the administrator to control which services a user is allowed to access. Some of the things that can be controlled with authorization include the user's initial privilege level and which commands the user is allowed to execute. When authorization fails, the user is denied access to the switch, even though the user has passed authentication.

The following examples assume that the configuration used in the previous examples has already been applied.

Local Authorization Example—Direct Login to Privileged Exec Mode

Apply the following configuration to use the local user database for authorization, such that a user can enter Privileged Exec mode directly:

```
aaa authorization exec "locex" local
line telnet
authorization exec locex
exit
```

With the users that were previously configured, the guest user will still log into user Exec mode, since the guest user only has privilege level 1 (the default). The admin user will be able to login directly to Privileged Exec mode since his privilege level was configured as 15.

RADIUS Authorization Example—Direct Login to Privileged Exec Mode

Apply the following configuration to use RADIUS for authorization, such that a user can enter Privileged Exec mode directly:

```
aaa authorization exec "rad" radius
line telnet
authorization exec rad
exit
```

Configure the RADIUS server so that the RADIUS attribute Service Type (6) is sent with value Administrative. Any value other than Administrative is interpreted as privilege level 1.

The following describes each line in the above configuration:

- The `aaa authorization exec "rad" radius` command creates an exec authorization method list called "rad" that contains the method radius.
- The `authorization exec rad` command assigns the rad exec authorization method list to be used for users accessing the switch via Telnet.



NOTES:

- If the privilege level is zero (that is, blocked), then authorization will fail and the user will be denied access to the switch.
- If the privilege level is higher than one, the user will be placed directly in Privileged Exec mode. Note that all commands in Privileged Exec mode require privilege level 15, so assigning a user a lower privilege level will be of no value.
- A privilege level greater than 15 is invalid and treated as if privilege level zero had been supplied.

RADIUS Authorization Example—Administrative Profiles

The switch should use the same configuration as in the previous authorization example.

The RADIUS server should be configured such that it will send the Cisco AV Pair attribute with the “roles” value. For example:

```
shell:roles=router-admin
```

The above example attribute gives the user access to the commands permitted by the router-admin profile.

RADIUS Change of Authorization

Dell EMC Networking N-Series switches support the Change of Authorization Disconnect-Request per RFC 3576. The Dell EMC Networking N-Series switch listens for the Disconnect-Request on UDP port 3799. The Disconnect-Request identifies the user session to be terminated using any or all of the following attributes:

- User-Name (IETF attribute #1)
- NAS-Port (IETF attribute #5)
- Framed-IP-Address (IETF attribute #8)
- Acct-Session-Id (IETF attribute #44)
- Calling-Station-Id (IETF attribute #31, which contains the host MAC address)

For CLI-based sessions (Console, Telnet and SSH), the supported Session Identification Attributes are User-Name and Framed-IP-Address.

The Calling-Station-ID must be a string of upper or lower case hexadecimal digits in one of the following formats:

- Raw notation, for example, AbCD01234567 - length 12
- Dotted quad notation, for example, BADC.1010.1234 - length 14
- Colon separated hex digits, for example, AB:cd:01:23:45:67 - length 17
- Dash separated hex digits: 01-23-45-67-89-Ab - length 17

The RADIUS Disconnect message may also contain the Acct-Terminate-Cause attribute (IETF #49).

The following messages from RFC 3576 are supported:

- 40 – Disconnect-Request
- 41 – Disconnect-ACK
- 42 – Disconnect-NAK

A CoA Disconnect-Request terminates the session without disabling the switch port. Instead, a CoA Disconnect-Request termination causes reinitialization of the authenticator state machine for the specified host. MAC-based authentication can be enabled for 802.1X sessions in conjunction with CoA. In this case, if the RADIUS server successfully terminates an 802.1X host session and subsequently does not re-authorize the host MAC address to access network resources, the host is effectively denied network access.

If the session cannot be located, the device returns a Disconnect-NAK message with the “Session Context Not Found” error-cause attribute. If the session is located, the device terminates the 802.1X session. After the session has been completely removed, the device returns a Disconnect-ACK message. The attributes returned within a CoA ACK can vary based on the CoA Request.

The administrator can configure whether all or any of the session attributes are used to identify a client session. If all is configured, all session identification attributes included in the CoA Disconnect-Request must match a session or the device returns a Disconnect-NAK or CoA-NAK with the “Invalid Attribute Value” error-code attribute. All attributes in the Disconnect-Request are treated as mandatory attributes, except Acct-Terminate-Cause. Unsupported attributes generate a Disconnect-NAK with error-cause Unsupported Service.

Dell EMC Networking N-Series switches support the following attributes in responses:

- User-Name (IETF attribute #1)
- NAS-Port (IETF attribute #5)
- Framed-IP-Address (IETF attribute #8)
- Calling-Station-ID (IETF attribute #31)
- Acct-Session-ID (IETF attribute #44)
- Message-Authenticator (IETF attribute #80)
- Error-Cause (IETF attribute #101)

A CoA NAK message is not sent for all CoA requests with a key mismatch. The message is sent only for the first three requests for a client. After that, all the packets from that client are dropped. When there is a key mismatch, the response authenticator sent with the CoA NAK message is calculated from a dummy key value.

The Dell EMC Networking N-Series switch will start listening to the 802.1X client again based on the re-authentication timer.

RADIUS COA Example

The following example configures the Dell EMC Networking N-Series switch to listen for and respond to RADIUS COA messages. This example does not configure any ports to use 802.1X or enable 802.1X. See "IEEE 802.1X" on page 334 for information on configuring 802.1X on interfaces.

- 1 Configure the switch to use the new model CLI command set. Dell EMC Networking N-Series switches do not support old model commands:

```
console#config  
console(config)#aaa new-model
```

- 2 Configure the switch to listen to RADIUS CoA requests.

```
console(config)#aaa server radius dynamic-author
```

- 3 Configure a local RADIUS client connection to RADIUS server 10.11.12.13 using the shared secret "secret sauce". The default port number is used.

```
console(config-radius-da)#client 10.11.12.13 server-key  
"secret sauce"
```

- 4 Disconnect-request client identification must match on all keys present in the request.

```
console(config-radius-da)#auth-type all  
console(config-radius-da)#exit
```

RADIUS COA Example with Telnet and SSH

The following example configures telnet and SSH clients in conjunction with RADIUS CoA.

- 1 Configure a login list named "login-list" that uses RADIUS as the only method:

```
console#config  
console(config)#aaa authentication login "login-list" radius
```

2 Enable RADIUS COA:

```
console(config)#aaa server radius dynamic-author
```

3 Enable the switch RADIUS client connecting to the RADIUS server at 10.130.191.89:

```
console(config-radius-da)#client 10.130.191.89 server-key  
"shared secret"
```

4 Allow matching of the client session on any of the key values present in the RADIUS disconnect:

```
console(config-radius-da)#auth-type any  
console(config-radius-da)#exit
```

5 Configure the RADIUS server attribute 4 (NAS-IP-Address). This attribute is sent in the RADIUS message to the RADIUS server but does not change the source IP address sent in the RADIUS messages. It allows a group of NASs to simulate a large RADIUS NAS:

```
console(config)#radius server attribute 4 10.130.65.4
```

6 Configure the remote RADIUS server address with name Default-RADIUS-Server and key "shared secret":

```
console(config)#radius server auth 10.130.191.89  
console(config-auth-radius)#name Default-RADIUS-Server  
console(config-auth-radius)#key "shared secret"  
console(config-auth-radius)#exit
```

7 Configure telnet sessions to the switch to use RADIUS authentication (the only login-list method):

```
console(config)#line telnet  
console(config-telnet)#login authentication login-list  
console(config-telnet)#exit
```

8 Configure SSH sessions to the switch to use RADIUS authentication:

```
console(config)#line ssh  
console(config-ssh)#login authentication login-list  
console(config-ssh)#exit
```

9 Enable the SSH server (the telnet server is enabled by default):

```
console(config)#ip ssh server
```

TACACS Authorization

TACACS+ Authorization Example—Direct Login to Privileged Exec Mode

Apply the following configuration to use TACACS+ for authorization, such that a user can enter Privileged Exec mode directly:

- 1 Create an exec authorization method list called “tacex” which contains the method tacacs.

```
console#config
console(config)#aaa authorization exec "tacex" tacacs
```

- 2 Assign the tacex exec authorization method list to be used for users accessing the switch via Telnet.

```
console(config)#line telnet
console(config-telnet)#authorization exec tacex
console(config-telnet)#exit
```

- 3 Configure the TACACS+ server so that the shell service is enabled and the priv-lvl attribute is sent when user authorization is performed. For example:

```
shell:priv-lvl=15
```

NOTES:

- If the privilege level is zero (that is, blocked), then authorization will fail and the user will be denied access to the switch.
- If the privilege level is higher than one, the user will be placed directly in Privileged Exec mode. Note that all commands in Privileged Exec mode require privilege level 15, so assigning a user a lower privilege level will be of no value.
- A privilege level greater than 15 is invalid and treated as if privilege level zero had been supplied.
- The shell service must be enabled on the TACACS+ server. If this service is not enabled, authorization will fail and the user will be denied access to the switch.

TACACS+ Authorization Example—Administrative Profiles

The switch should use the same configuration as for the previous authorization example.

The TACACS+ server should be configured such that it will send the “roles” attribute. For example:

```
shell:roles=router-admin
```


The above example attribute will give the user access to the commands permitted by the router-admin profile.



NOTE: If the priv-lvl attribute is also supplied, the user can also be placed directly into Privileged Exec mode.

TACACS+ Authorization Example—Custom Administrative Profile

This example creates a custom profile that allows the user to control user access to the switch by configuring a administrative profile that only allows access to AAA related commands. Use the following commands to create the administrative profile:

- 1 Create an administrative profile called “aaa” and place the user in admin-profile-config mode.

```
console#config
console(config)#admin-profile aaa
```

- 2 Enter rule number **permit command** regex commands to allows any command that matches the regular expression.

The command rules use regular expressions as implemented by Henry Spencer's regex library (the POSIX 1003.2 compliant version). In the regular expressions used in this example, the caret (^) matches the null string at the beginning of a line, the period (.) matches any single character, and the asterisk (*) repeats the previous match zero or more times.

```
console(config)#rule 99 permit command ^show aaa .*"
console(admin-profile)#rule 98 permit command ^show
authentication .*"
console(admin-profile)#rule 97 permit command ^show
authorization .*"
console(admin-profile)#rule 96 permit command ^show
accounting .*"
console(admin-profile)#rule 95 permit command ^show tacacs
.*"
console(admin-profile)#rule 94 permit command ^aaa .*"
console(admin-profile)#rule 93 permit command ^line .*"
console(admin-profile)#rule 92 permit command ^login .*"
console(admin-profile)#rule 91 permit command ^authorization
.*"
console(admin-profile)#rule 90 permit command ^accounting .*"
console(admin-profile)#rule 89 permit command ^configure .*"
```

```
console(admin-profile)#rule 88 permit command "^password .*"
console(admin-profile)#rule 87 permit command "^username .*"
console(admin-profile)#rule 86 permit command "^show user.*"
console(admin-profile)#rule 85 permit command "^radius server
.*"
console(admin-profile)#rule 84 permit command "^tacacs-server
.*"
```

- 3 Enter rule number **permit mode** mode-name commands to allows all commands in the named mode.

```
console(admin-profile)#rule 83 permit mode radius-auth-config
console(admin-profile)#rule 82 permit mode radius-acct-config
console(admin-profile)#rule 81 permit mode tacacs-config
console(admin-profile)#exit
```

- 4 Assign this profile to a user by configuring the TACACS+ server so that it sends the following “roles” attribute for the user:

```
shell:roles=aaa
```

If it is desired to also permit the user access to network-operator commands (basically, all the command in User Exec mode), then the “roles” attribute would be configured as follows:

```
shell:roles=aaa,network-operator
```

TACACS+ Authorization Example—Per-command Authorization

An alternative method for command authorization is to use the TACACS+ feature of per-command authorization. With this feature, every time the user enters a command, a request is sent to the TACACS+ server to ask if the user is permitted to execute that command. Exec authorization does not need to be configured to use per-command authorization.

Apply the following configuration to use TACACS+ to authorize commands:

- 1 Creates a command authorization method list called “taccmd” that includes the method tacacs.

```
console#config
console(config)#aaa authorization commands "taccmd" tacacs
```

- Assigns the taccmd command authorization method list to be used for users accessing the switch via Telnet.

```
console(config)#line telnet
console(config-telnet)#authorization commands taccmd
console(config-telnet)#exit
```

The TACACS+ server must be configured with the commands that the user is allowed to execute. If the server is configured for command authorization as “None”, then no commands will be authorized. If both administrative profiles and per-command authorization are configured for a user, any command must be permitted by both the administrative profiles and by per-command authorization.

TACACS Authorization—Privilege Level

Dell EMC Networking TACACS supports setting the maximum user privilege level in the TACACS authorization response. Configure the TACACS server to send `priv-lvl=X`, where X is either 1 (Non-privileged mode), or 15 (Privileged Exec mode).

Accounting

Accounting is used to record security events, such as a user logging in or executing a command. Accounting records may be sent upon completion of an event (stop-only) or at both the beginning and end of an event (start-stop). There are three types of accounting: commands, Dot1x, and exec.

- **Commands**—Sends accounting records for command execution.
- **Dot1x**—Sends accounting records for network access.
- **Exec**—Sends accounting records for management access (logins).

For more information about the data sent in accounting records, see "Which RADIUS Attributes Does the Switch Support?" on page 290 and "Using TACACS+ Servers to Control Management Access" on page 293.

Table 10-10 shows the valid methods for each type of accounting:

Table 10-10. Accounting Methods

| Method | Commands | Dot1x | Exec |
|--------|----------|-------|------|
| radius | no | yes | yes |
| tacacs | yes | no | yes |

RADIUS Accounting

Dell EMC Networking N-Series switches support RADIUS accounting. The supported accounting types are start-only or start-stop.

The following attributes may be sent in the Accounting Stop record that is sent to the RADIUS server when the switch is configured for 802.1X accounting:

- User-Name (1)
- NAS-IP-Address (4)
- Framed-IP-Address (8)
- Called-Station-Id (30)
- Calling-Station-Id (31)
- NAS-Port-Type (61)
- Acct-Terminate-Cause(49)

- Class (25)
- Acct-Session Time(46)
- Acct-Input-Octets (42)
- Acct-Output-Octets (43)
- Acct-Input-Gigawords(52)
- Acct-Output-Gigawords (53)
- Framed-IPv6-Address (168)
- Acct-Delay-Time (41)
- Acct-Session-Id (44)
- NAS-Port-Id (87)

Certain of the attributes above are sent only if received from the RADIUS server during the Access Request process, for example, Class.

The following attributes are sent in the Accounting Start record sent to the RADIUS server when the switch is configured for 802.1x accounting:

- User-Name (1)
- NAS-IP-Address (4)
- Framed-IP-Address (8)
- Class (25)
- Called-Station-Id (30)
- Calling-Station-Id (31)
- NAS-Port-Type (61)
- Tunnel-Private-Group-Id (81) - VLAN ID
- Framed-IPv6-Address (168)
- Acct-Session-Id (44)
- NAS-Port-Id (87)

The Framed-IP-Address or Framed-IPv6-Address are only sent if available.

IEEE 802.1X

What is IEEE 802.1X?


The IEEE 802.1X standard provides a means of preventing unauthorized access by supplicants (clients) to the services the switch offers, such as access to the LAN.

The 802.1X network has three components:

- **Supplicant** — The client connected to the authenticated port that requests access to the network.
- **Authenticator** — The network device that prevents network access prior to authentication.
- **Authentication Server** — The network server (such as a RADIUS server) that performs the authentication on behalf of the authenticator, and indicates whether the user is authorized to access system services. Dell EMC Networking supports interoperability with a variety of external authentication servers. Refer to "Authentication, Authorization, and Accounting" on page 275 for more information.

Figure 10-3 shows the 802.1X network components.

Figure 10-3. IEEE 802.1X Network



As shown in Figure 10-3, the Dell EMC Networking switch is the authenticator and ensures that the supplicant (a PC) that is attached to an 802.1X-controlled port is authenticated by an authentication server (a RADIUS server). The result of the authentication process determines whether the supplicant is authorized to access network services on that controlled port. Dell EMC Networking N-Series switches support 802.1X authentication using remote RADIUS or using a local authentication service (IAS).

Supported security methods for supplicant communication with remote authentication servers include MD5, PEAP, EAP-TTL, EAP-TTLS, and EAP-TLS. Only EAP-MD5 is supported when using the local authentication server (IAS) for communication with the supplicant.

For a list of RADIUS attributes that the switch supports, see "Using RADIUS" on page 288.

What are the 802.1X Port Authentication Modes?

The 802.1X port authentication mode determines whether to allow or prevent network traffic on the port. A port can be configured to be in one of the following 802.1X authentication modes:

- Auto (default)
- MAC-based
- Force-authorized
- Force-unauthorized

These modes control the behavior of the port. The port state is either Authorized or Unauthorized. 802.1X auto mode may be configured on ports in general or access mode. 802.1X is not supported on trunk mode ports.

If the port is in the force-authorized mode, the port state is Authorized and the port sends and receives normal traffic without client port-based authentication. When a port is in a forced-unauthorized mode, the port state is Unauthorized and the port ignores supplicant authentication attempts and does not provide authentication services. By default, when 802.1X is globally enabled on the switch, all ports are in auto authentication mode, which means the port will be unauthorized until a successful authentication exchange has taken place. Auto mode (port based mode) is suitable for authentication of a single supplicant attached to a port. If multiple devices

are attached to a port configured in auto mode, they will all be allowed access to network resources as soon as any 802.1X-aware device on the port authenticates.

The port security feature can be utilized if it is desired to limit access on auto mode configured ports. To limit access to a phone and laptop configuration using Voice VLAN, the port security limit should be set to 3 as many IP phones also utilize the data VLAN during power up. For more information on port security, see "Port and System Security" on page 681.

In addition to force-authorized, force-unauthorized, and auto modes, the 802.1X mode of a port can be MAC based, as the following section describes.



NOTE: Only MAC-Based and Auto modes use 802.1X and RADIUS to authenticate. Force-authorized and Force-unauthorized modes are manual overrides.

What is MAC-Based 802.1X Authentication?

MAC-based authentication allows multiple supplicants connected to the same port to authenticate individually. For example, a 5-port hub might be connected to a single port on the switch. Each host connected to the hub must authenticate separately in order to gain access to the network. Hosts that do not authenticate (or are not configured with MAB or a guest or unauthenticated VLAN) are denied access to the network, or are placed into a restricted VLAN such as the guest or unauthenticated VLAN, if configured. MAC-based authentication is only supported for ports configured in general mode.

The hosts are distinguished by their MAC addresses. Internally, the switch adds an ACL to the port to allow packets from the host MAC address to pass into the switch. For this reason, enabling port security on an interface configured for MAC-based authentication is neither necessary nor desirable.

When multiple hosts (for example, a PC, a printer, and a phone in the same office) are connected to the switch on the same port, each of the connected hosts authenticates separately with the RADIUS server.

If a port uses MAC-based 802.1X authentication, the option to use MAC Authentication Bypass (MAB) is available. MAB is a supplemental authentication mechanism that allows 802.1X unaware clients – such as printers, fax machines, and some IP phones — to authenticate to the network using the client MAC address as an identifier.

The known and allowable MAC address and corresponding access rights of the client must be pre-populated in the authentication server. Both MAB and MAC-based authentications are supported on a port simultaneously.

When a port configured for MAB receives traffic from an unauthenticated client, the switch (Network Authentication Server or NAS):

- Sends a EAP Request packet to the unauthenticated client
- Waits a pre-determined period of time for a response
- Retries – resends the EAP Request packet up to three times
- Considers the client to be 802.1X unaware client (if it does not receive an EAP response packet from that client)

The NAS sends a request to the authentication server with the MAC address of the client in a hexadecimal format as the username and the MD5 hash of the MAC address as the password. The authentication server checks its database for the authorized MAC addresses and returns an Access-Accept or an Access-Reject response, depending on whether the MAC address is found in the database. If an Access-Accept is received by the NAS, an internal ACL is applied to the port using the MAC address of the authenticated device allowing it to access the network. Any other devices wishing to access the network must authenticate individually. MAB also allows 802.1X-unaware clients to be placed in a RADIUS-assigned VLAN or to apply a specific Filter ID to the client traffic.

The following information is sent to the RADIUS authenticator for MAB clients using EAP-MD5 authentication:

1 - User-Name—MAC address of MAB device (AA:BB:CC:DD:EE:FF)

Attribute 2 is not sent if Auth type is EAP-MD5.

4 - NAS-IP-Address—IP address of the switch

5 - NAS-Port—switch internal port number (ifIndex)

6 - Service Type 10 (Call-Check)

12 - Framed-MTU - port/switch MTU - header length (e.g. 1500)

30 - Called Station ID —MAC address of device (xx:xx:xx:xx:xx:xx format)

31 - Calling Station ID—Switch MAC address

61 - NAS-Port-Type (Ethernet 15)

80 - Message Authenticator

87- NAS-Port-Id (such as Gigabitethernet 1/0/15)

79-EAP-Message

The format of the Calling-Station-ID for MAB clients may be altered using the **attribute 31** command. The format of the User-Name attribute for MAB clients may be altered using the **attribute 1** command.

By default, MAB clients are authenticated to the authentication server using EAP-MD5. MAB clients may optionally be configured to use CHAP or PAP to authenticate the MAB device. For CHAP or PAP, the following attributes are sent to the RADIUS server:

1 - User-Name—MAC address of MAB device

2 - User Password (PAP only)

3 - CHAP-Password - = Encrypted MAC address (CHAP) only or unencrypted (PAP) User Name

4 - NAS-IP-Address—IP address of the switch

5 - NAS-Port—switch internal port number (ifIndex)

6 - Service Type is set to 10 for MAB (Call-Check)

12 - Framed-MTU - port/switch MTU - header length (e.g. 1500)

30 - Called Station ID—MAC address of device (in xx:xx:xx:xx:xx:xx format)

31 - Calling Station ID—Switch MAC address

60 - CHAP Challenge (CHAP only)

61 - NAS-Port-Type (Ethernet 15)

80 - Message Authenticator

87 - NAS-Port-ID



NOTE: MAB initiates only after the dot1x guest VLAN period times out. If the client responds to any of the EAPOL identity requests, MAB does not initiate for that client.

What is the Role of 802.1X in VLAN Assignment?

Dell EMC Networking N-Series switches allow a port to be placed into a particular VLAN based on the result of the authentication. The authentication server can provide information to the switch about which VLAN to assign the supplicant or the administrator can configure the level of access provided when authentication fails or is never attempted.

When a host connects to a switch that uses an authentication server to authenticate, the host authentication will have one of three outcomes:

- The host is authenticated.
- The host attempts to authenticate but fails because it lacks certain security credentials.
- The host does not try to authenticate at all (802.1X unaware).

Three separate VLANs can be created on the switch to handle a host depending on whether the host authenticates, fails the authentication, or does not attempt authentication. The RADIUS server informs the switch of the selected VLAN as part of the authentication.

Authenticated VLANs

Hosts that authenticate normally are assigned a VLAN that includes access to network resources. This VLAN may be assigned by the RADIUS server. Hosts that fail authentication may be denied access to the network or placed into an unauthenticated VLAN, if configured. Hosts that do not attempt authentication may be placed into a guest VLAN, if configured. The network administrator can configure the type of access provided to the authenticated, guest, and unauthenticated VLANs.

Much of the configuration to assign authenticated hosts to a particular VLAN takes place on the 802.1X authenticator server (for example, a RADIUS server). If an external RADIUS server is used to manage VLANs, configure the server to use Tunnel attributes in Access-Accept messages in order to inform the switch about the selected VLAN. These attributes are defined in RFC 2868 and their use for dynamic VLAN is specified in RFC 3580.

The VLAN attributes defined in RFC3580 and required for VLAN assignment via RADIUS are as follows:

- Tunnel-Type (64) = VLAN (13)
- Tunnel-Medium-Type (65) = 802 (6)

- Tunnel-Private-Group-ID (81) = VLANID

The tag value for the Tunnel-Private-Group-ID is parsed as the length of the VLAN ID. The VLAN ID may consist of a VLAN name (not to exceed 32 characters) or a numeric value in ASCII (no alphabetic characters are allowed) in the range 1–4093.

Dynamic VLAN Creation

If RADIUS-assigned VLANs are enabled through the Authorization Network RADIUS configuration option, the RADIUS server is expected to include the VLAN ID in the 802.1X tunnel attributes of its response message to the switch. If dynamic VLAN creation is enabled on the switch and the RADIUS-assigned VLAN does not exist, then the assigned VLAN is dynamically created and the port PVID or native VLAN is set to the RADIUS-assigned VLAN ID. Trunk mode ports are also made members of the created VLAN.

If the VLAN is already created on the switch, the port PVID or native VLAN is set to the VLAN ID. This implies that the client can connect from any port and be assigned to the appropriate VLAN based on the RADIUS server configuration. This gives flexibility for clients to move around the network without much additional configuration required on the switches in the network. Dynamic VLAN assignment requires that the port be configured in general or access mode.

Unauthenticated VLAN

The network administrator may choose to configure an unauthenticated VLAN. Hosts that attempt authentication and fail are placed in the unauthenticated VLAN, if configured. Once in the unauthenticated VLAN, authentication is not reattempted until:

- the re-authentication timer expires
- the supplicant disconnects from the port
- the port is shut down and re-enabled

The number of re-authentication failures required to place a supplicant in the unauthenticated VLAN is not configurable.

The network administrator can configure the unauthenticated VLAN to provide the desired level of network access, i.e., a black hole or a guest VLAN type of access.

Guest VLAN

The Guest VLAN feature provides a mechanism to allow users access to a guest VLAN. For example, the administrator might provide a guest VLAN to visitors and contractors to permit network access that allows visitors to connect to external network resources, such as the Internet, with no ability to access information on the internal LAN.

As an example, on a port configured in auto authentication mode (**dot1x port-control auto**) and connected to a client that does not support 802.1X, the client does not respond to the 802.1X requests from the switch. The port remains in the unauthorized state and the client is not granted access to the network. If a guest VLAN is configured for that port, the port is placed in the configured guest VLAN and moved to the authorized state, allowing access to the client over the guest VLAN.



NOTE: MAB and the guest VLAN feature are mutually exclusive on a port. If MAB is enabled on a port concurrently with guest VLAN, the port will not move to the authorized state.

When the guest VLAN capability is disabled, users authorized by the guest VLAN are removed from the VLAN and denied network access.

What is Monitor Mode?

The monitor mode is a special mode that can be enabled in conjunction with 802.1X authentication. Monitor mode provides a way for network administrators to identify possible issues with the 802.1X configuration on the switch without affecting the network access to the users of the switch. It allows network access even in case where there is a failure to authenticate but logs the results of the authentication process for diagnostic purposes.

The monitor mode can be configured globally on a switch. If the switch fails to authenticate a user for any reason (for example, RADIUS access reject from RADIUS server, RADIUS timeout, or the client itself is dot1x-unaware), the client is authenticated and is undisturbed by the failure condition(s). The reasons for failure are logged for tracking purposes.

Table 10-11 provides a summary of the 802.1X Monitor Mode behavior.

Table 10-11. IEEE 802.1X Monitor Mode Behavior

| Case | Sub-case | Regular 802.1X | 802.1X Monitor Mode |
|------------------------|----------------------------|--|--|
| RADIUS/IAS Success | Success | Port State: Permit VLAN: Assigned Filter: Assigned | Port State: Permit VLAN: Assigned Filter: Assigned |
| | Incorrect NAS Port | Port State: Deny | Port State: Permit VLAN: Assigned |
| | Invalid VLAN Assignment | Port State: Deny | Port State: Permit VLAN: Default PVID of the port |
| | Invalid Filter-ID | Port State: Deny | Port State: Permit VLAN: Assigned |
| | Invalid DACL | Port State: Deny | Port State: Permit DACL: Not Assigned VLAN: Assigned |
| | Bad RADIUS packet | Port State: Deny | Port State: Permit VLAN: Default PVID of the port |
| RADIUS/IAS Failure | Default behavior | Port State: Deny | Port State: Permit VLAN: Default PVID of the port |
| | Unauth VLAN enabled | Port State: Permit VLAN: Unauth | Port State: Permit VLAN: Unauth |
| RADIUS Timeout | Default behavior | Port State: Deny | Port State: Permit VLAN: Default PVID of the port |
| | Unauth VLAN enabled | Port State: Deny | Port State: Permit VLAN: Unauth |
| Critical Voice VLAN | Default behavior | Port State: Deny | Port State: Permit VLAN: Critical Voice VLAN |
| EAPOL Timeout | Default behavior | Port State: Deny | Port State: Permit |

Table 10-11. IEEE 802.1X Monitor Mode Behavior (Continued)

| Case | Sub-case | Regular 802.1X | 802.1X Monitor Mode |
|--|-----------------------------------|--|--|
| 3 × EAPOL Timeout (Guest VLAN timer expiry or MAB timer expiry) | Guest VLAN enabled | Port State: Permit VLAN: Guest | Port State: Permit VLAN: Guest |
| | MAB Success Case | Port State: Permit VLAN: Assigned Filter: Assigned | Port State: Permit VLAN: Assigned Filter: Assigned |
| | MAB Fail Case | Port State: Deny | Port State: Permit VLAN: Default PVID of the port |
| Supplicant Timeout | | Port State: Deny | Port State: Deny |
| Port/Client Authenticated on Guest VLAN | Delete Guest VLANID through Dot1Q | Port State: Deny | Port State: Permit VLAN: Default PVID of the port |

How Does the Authentication Server Assign DiffServ Policy or ACLs?

The Dell EMC Networking N-Series switches allow the external 802.1X Authenticator or RADIUS server to assign ACL or DiffServ policies to users that authenticate to the switch. When a host (supplicant) attempts to connect to the network through a port, the switch contacts the 802.1X authenticator or RADIUS server, which then provides information to the switch about which ACL or DiffServ policy to assign the host (supplicant). The application of the policy is applied to the host after the authentication process has completed. The ACL or DiffServ policy is always applied for the “in” direction of the interface and applies to the interface as a whole.

For additional guidelines about using an authentication server to assign DiffServ policies, see "Configuring Authentication Server Dynamic ACL or DiffServ Policy Assignments" on page 367.

What is the Internal Authentication Server?

The Internal Authentication Server (IAS) is a dedicated local database for authentication of users for network access through 802.1X. In this database, the switch maintains a list of username and password combinations to use for 802.1X authentication. Entries can be created in the database manually, or the IAS information can be uploaded to the switch.

If the authentication method for 802.1X is IAS, the switch uses the locally stored list of username and passwords to provide port-based authentication to users instead of using an external authentication server. Authentication using the IAS supports the EAP-MD5 method only.



NOTE: The IAS database does not support VLAN assignments or DiffServ policy/ACL assignments.

Default 802.1X Values

Table 10-12 lists the default values for the 802.1X features.


Table 10-12. Default Port-Based Security Values

| Feature | Description |
|---|--------------|
| Global 802.1X status | Disabled |
| 802.1X authentication method | None |
| Per-port 802.1X status | Disabled |
| Port authentication mode | Auto mode |
| Port authentication state | Unauthorized |
| Periodic reauthentication | Disabled |
| Seconds between reauthentication attempts | 3600 |
| Authentication server timeout | 30 seconds |
| Resending EAP identity Request | 30 seconds |
| Quiet period | 60 seconds |
| Supplicant timeout | 30 seconds |
| Max EAP request | 2 times |

Table 10-12. Default Port-Based Security Values

| Feature | Description |
|--|--|
| Maximum number of supplicants per port for MAC-based authentication mode | 64 (32 for N1100-ON and N1500 Series switches) |
| Guest VLAN | Disabled |
| Unauthenticated VLAN | Disabled |
| Dynamic VLAN creation | Disabled |
| RADIUS-assigned VLANs | Disabled |
| IAS users | none configured |
| Port security | Unlocked |
| Port security traps | Enabled |
| Maximum learned MAC addresses | 100 (when locked) |
| Monitor mode | Disabled |

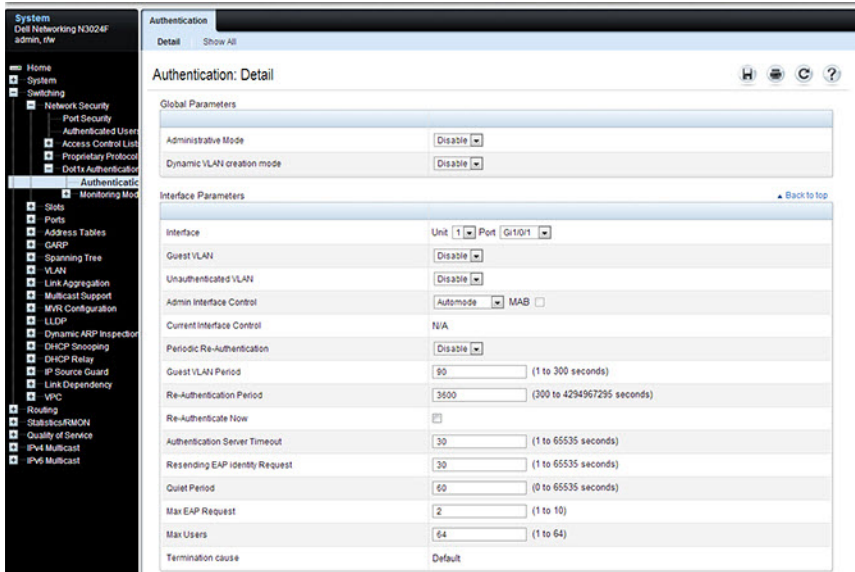
Configuring IEEE 802.1X (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the IEEE 802.1X features and Port Security on Dell EMC Networking N1100-ON, N1500, N2000, N2100-ON, N3000, N3100-ON, and N4000 Series switches. For details about the fields on a page, click  at the top of the Dell EMC OpenManage Switch Administrator web page.

Dot1x Authentication

Use the **Dot1x Authentication** page to configure the 802.1X administrative mode on the switch and to configure general 802.1X parameters for a port. To display the **Dot1x Authentication** page, click **Switching** → **Network Security** → **Dot1x Authentication** → **Authentication** in the navigation panel.


Figure 10-4. Dot1x Authentication



To configure 802.1X authentication on multiple ports:

- 1 Open the **Dot1x Authentication** page.
- 2 Click **Show All** to display the **Dot1x Authentication Table** page.
- 3 In the Ports list, select the check box in the **Edit** column for the port to configure.
- 4 Select the desired settings to change for all ports that are selected for editing.

Figure 10-5. Configure Dot1x Settings



5 Click **Apply**.

To reauthenticate a port:

1 Open the **Dot1x Authentication** page.

2 Click **Show All**.

The **Dot1x Authentication Table** displays.

3 Check **Edit** to select the **Unit/Port** to re-authenticate.

4 Check **Re-authenticate Now**.

5 Click **Apply**.

The authentication process is restarted on the specified port.

To reauthenticate multiple ports:

1 Open the **Dot1x Authentication** page.

2 Click **Show All**.

The **Dot1x Authentication Table** displays.

3 Check **Edit** to select the **Units/Ports** to re-authenticate.

4 To re-authenticate on a periodic basis, set **Periodic Re-Authentication** to **Enable**, and specify a **Re-Authentication Period** for all desired ports.

5 To re-authenticate immediately, check **Re-authenticate Now** for all ports to be re-authenticated.

6 Click **Apply**.

The authentication process is restarted on the specified ports (either immediately or periodically).

To change the administrative port control:

- 1 Open the **Dot1x Authentication** page.
- 2 Click **Show All**.

The **Dot1x Authentication Table** displays.

- 3 Scroll to the right side of the table and select the **Edit** check box for each port to configure. Change **Admin Interface Control** to **Authorized**, **Unauthorized**, **MAC-based**, or **Automode** as needed for chosen ports. Only **MAC-based** and **Automode** actually use 802.1X to authenticate. **Authorized** and **Unauthorized** are manual overrides.
- 4 Click **Apply**.


Admin Port Control is updated for the specified ports, and the device is updated.

Authenticated Users

The **Authenticated Users** page is used to display lists of ports that have authenticated users.


To display the **Authenticated Users** page, click **Switching** → **Network Security** → **Authenticated Users** in the navigation panel.

Figure 10-6. Network Security Authenticated Users




Port Access Control Configuration

Use the **Port Access Control Configuration** page to globally enable or disable RADIUS-assigned VLANs and to enable Monitor Mode to help troubleshoot 802.1X configuration issues.

 **NOTE:** The VLAN Assignment Mode field is the same as the Admin Mode field on the System → Management Security → Authorization Network RADIUS page.

To display the **Port Access Control Configuration** page, click **Switching** → **Network Security** → **Dot1x Authentication** → **Monitor Mode** → **Port Access Control Configuration** in the navigation panel.

Figure 10-7. Port Access Control Configuration




Port Access Control History Log

Use the **Port Access Control History Log** page to view log messages about 802.1X client authentication attempts. The information on this page can help you troubleshoot 802.1X configuration issues.

To display the **Port Access Control History Log Summary** page, click **Port Access Control Configuration** page, click **Switching** → **Network Security** → **Dot1x Authentication** → **Monitor Mode** → **Port Access Control History Log Summary** in the navigation panel.

Figure 10-8. Port Access Control History Log

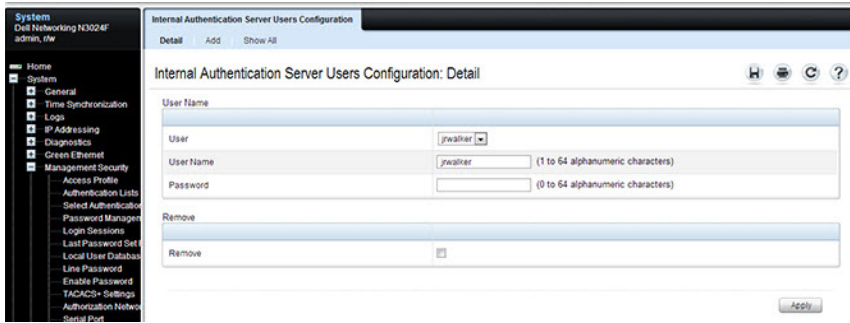


Internal Authentication Server Users Configuration

Use the **Internal Authentication Server Users Configuration** page to add users to the local IAS database and to view the database entries.

To display the **Internal Authentication Server Users Configuration** page, click **System** → **Management Security** → **Internal Authentication Server Users Configuration** in the navigation panel.

Figure 10-9. Internal Authentication Server Users Configuration




NOTE: If no users exist in the IAS database, the IAS Users Configuration Page does not display the fields shown in the image.

To add IAS users:

- 1 Open the **Internal Authentication Server Users Configuration** page.
- 2 Click **Add** to display the **Internal Authentication Server Users Add** page.
- 3 Specify a username and password in the appropriate fields.

Figure 10-10. Adding an IAS User



- 4 Click **Apply**.

To view the Internal Authentication Server Users Table page, click **Show All**. To delete an IAS user:

- 1 Open the **Internal Authentication Server Users Configuration** page.
- 2 From the **User** menu, select the user to remove, select the user to remove.
- 3 Select the **Remove** check box.

Figure 10-11. Removing an IAS User



- 4 Click **Apply**.

Configuring IEEE 802.1X (CLI)

This section provides information about commands you use to configure 802.1X and Port Security settings. For additional information about the commands in this section, see the Dell EMC Networking N1100-ON, N1500, N2000, N2100-ON, N3000, N3100-ON, and N4000 Series Switches CLI Reference Guide at www.dell.com/support.

Configuring Basic 802.1X Authentication Settings

Use the following commands to enable and configure 802.1X authentication on the switch.

| Command | Purpose |
|-----------|----------------------------------|
| configure | Enter Global Configuration mode. |

| Command | Purpose |
|---|--|
| <code>aaa authentication dot1x default method1</code> | Specify the authentication method to use to authenticate 802.1X clients that connect to the switch. method1—The method keyword can be radius , none , or ias . |
| <code>dot1x system-auth-control</code> | Globally enable 802.1X authentication on the switch. |
| <code>interface interface</code> | Enter interface configuration mode for the specified interface. The interface variable includes the interface type and number, for example tengigabitethernet 1/0/3 . A range of interfaces can be specified using the interface range command. For example, interface range tengigabitethernet 1/0/8-12 configures interfaces 8, 9, 10, 11, and 12. |

| Command | Purpose |
|--|--|
| <pre>dot1x port-control {force-authorized force-unauthorized auto mac-based}</pre> | <p>Specify the 802.1X mode for the port.</p> <p>NOTE: For standard 802.1X implementations in which one client is connected to one port, use the <code>dot1x port-control auto</code> command to enable 802.1X authentication on the port.</p> <ul style="list-style-type: none"> • auto — Enables 802.1X authentication on the interface and causes the port to transition to the authorized or unauthorized state based on the 802.1X authentication exchange between the switch and the client. Once the port is authenticated by any host, additional hosts on the port will have access to network resources using the port PVID. • force-authorized — Disables 802.1X authentication on the interface and causes the port to transition to the authorized state without any authentication exchange required. The port sends and receives normal traffic without 802.1X-based authentication of the client. • force-unauthorized — Denies all access through this interface by forcing the port to transition to the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface. • mac-based — Enables 802.1X authentication on the interface and allows multiple hosts to individually authenticate on a single port. The hosts are distinguished by their MAC addresses. |
| <code>mab</code> | If the 802.1X mode on the interface is mac-based , this command can be used to enable MAB on the interface. |
| <code>CTRL + Z</code> | Exit to Privileged Exec mode. |
| <code>show dot1x</code> | View the current 802.1X status. |
| <code>show dot1x clients {all interface}</code> | View information about 802.1X clients that have successfully authenticated and are connected to the switch. The interface variable includes the interface type and number. |
| <code>show dot1x users [username username]</code> | View the 802.1X authenticated users for the switch. |



NOTE: To enable 802.1X Monitor Mode to help troubleshoot authentication issues, use the `dot1x system-auth-control monitor` command in Global Configuration mode. To view 802.1X authentication events and information, use the `show dot1x authentication-history {interface | all} [failed-auth-only] [detail]` command. To clear the history, use the `clear dot1x authentication-history` command in Privileged Exec mode.

Configuring Additional 802.1X Interface Settings

Use the following commands to configure 802.1X interface settings such as the reauthentication period and switch-to-client retransmission time.

| Command | Purpose |
|---|--|
| <code>configure</code> | Enter Global Configuration mode. |
| <code>interface interface</code> | Enter interface configuration mode for the specified interface. The interface variable includes the interface type and number, for example <code>tengigabitethernet 1/0/3</code> . A range of interfaces can be specified using the interface range command. For example, interface range tengigabitethernet 1/0/8-12 configures interfaces 8, 9, 10, 11, and 12. |
| <code>dot1x reauthentication</code> | Enable periodic re-authentication of the client. |
| <code>dot1x timeout re-authperiod</code> seconds | Set the number of seconds between re-authentication attempts. |
| <code>dot1x timeout server-timeout</code> seconds | Set the time that the switch waits for a response from the authentication server. |
| <code>dot1x timeout tx-period</code> seconds | Set the number of seconds that the switch waits for a response to an Extensible Authentication Protocol (EAP)-request/identity frame from the client before resending the request. |
| <code>dot1x timeout quiet-period</code> seconds | Set the number of seconds that the switch remains in the quiet state following a failed authentication exchange (for example, the client provided an invalid password). |
| <code>dot1x timeout supp-timeout</code> seconds | Set the time that the switch waits for a response before retransmitting an Extensible Authentication Protocol (EAP)-request frame to the client. |


| Command | Purpose |
|--|---|
| <code>dot1x max-req count</code> | Set the maximum number of times that the switch sends an Extensible Authentication Protocol (EAP)-request frame (assuming that no response is received) other than Request-Identity to the client before restarting the authentication process. |
| <code>dot1x max-reauth-req count</code> | Set the maximum number of times that the switch sends an Extensible Authentication Protocol (EAP)-Request Identify frame to client with no response before restarting the authentication process. |
| <code>dot1x max-users users</code> | Set the maximum number of clients supported on the port when MAC-based 802.1X authentication is enabled on the port. |
| CTRL + Z | Exit to Privileged Exec mode. |
| <code>dot1x re-authenticate [interface]</code> | Manually initiate the re-authentication of all 802.1X-enabled ports or on the specified 802.1X-enabled port. The interface variable includes the interface type and number. |
| <code>dot1x initialize [interface]</code> | Start the initialization sequence on all ports or on the specified port. NOTE: This command is valid only if the port-control mode for the specified port is auto or MAC-based. |
| <code>show dot1x [interface interface]</code> | View 802.1X settings for the switch or for the specified interface. |
| <code>show dot1x interface interface statistics</code> | View 802.1X statistics for the specified interface. |

Configuring 802.1X Settings for RADIUS-Assigned VLANs

Use the following commands to configure 802.1X settings that affect the RADIUS-assigned VLAN.

| Command | Purpose |
|---|--|
| <code>configure</code> | Enter Global Configuration mode. |
| <code>aaa authorization network default radius</code> | Allow the RADIUS server to assign VLAN IDs to clients. |

| Command | Purpose |
|--|--|
| <code>dot1x dynamic-vlan enable</code> | If the RADIUS assigned VLAN does not exist on the switch, allow the switch to dynamically create the assigned VLAN. |
| <code>interface interface</code> | Enter interface configuration mode for the specified interface. The interface variable includes the interface type and number, for example <code>tengigabitethernet 1/0/3</code> . A range of interfaces can be specified using the <code>interface range</code> command. For example, <code>interface range tengigabitethernet 1/0/8-12</code> configures interfaces 8, 9, 10, 11, and 12. |
| <code>dot1x guest-vlan vlan-id</code> | Specify the guest VLAN. |
| <code>dot1x unauth-vlan vlan-id</code> | Specify the unauthenticated VLAN. |
| <code>CTRL + Z</code> | Exit to Privileged Exec mode. |
| <code>show dot1x advanced interface</code> | View the current 802.1X configuration. |

 **NOTE:** When dynamically creating VLANs, the uplink port should be in trunk mode so that it will automatically participate in all dynamically-created VLANs. Otherwise, the supplicant may be placed in a VLAN that does not extend beyond the switch because no other ports are participating.

Configuring Internal Authentication Server Users

Use the following commands to add users to the IAS database and to use the database for 802.1X authentication.

| Command | Purpose |
|--|--|
| <code>configure</code> | Enter Global Configuration mode. |
| <code>aaa ias-user username user</code> | Add a user to the IAS user database. This command also changes the mode to the IAS User Config mode. |
| <code>password password [encrypted]</code> | Configure the password associated with the user. |
| <code>CTRL + Z</code> | Exit to Privileged Exec mode. |
| <code>show aaa ias-users</code> | View all configured IAS users. |
| <code>clear aaa ias-users</code> | Delete all IAS users from the database. |

IEEE 802.1X Configuration Examples

This section contains the following examples:

- Configuring 802.1X Authentication
- Controlling Authentication-Based VLAN Assignment
- Allowing Dynamic Creation of RADIUS-Assigned VLANs
- Configuring Authentication Server Dynamic ACL or DiffServ Policy Assignments

Configuring 802.1X Authentication

The network in this example requires clients to use 802.1X authentication to access the network through the switch ports. The administrator must configure the following settings on systems other than the switch before configuring the switch:

- 1 Add the users to the client database on the Authentication Server, such as a RADIUS server with Cisco[®] Secure Access Control Server (ACS) software.
- 2 Configure the settings on the client, such as a PC running Microsoft[®] Windows, to require 802.1X authentication.

The switch uses an authentication server with an IP address of 10.10.10.10 to authenticate clients. Port 7 is connected to a printer in the unsecured area. The printer is an 802.1X unaware client, so Port 7 is configured to use MAC-based authentication with MAB.




NOTE: The printer requires an entry in the client database that uses the printer MAC address as the username.

An IP phone is directly connected to Port 8, and a PC is connected to the IP phone. Both devices are authenticated through MAC-based authentication, which allows multiple hosts to authenticate on a single port. The hosts are distinguished by their MAC addresses, and hosts authenticate separately with the RADIUS server.

Port 9 is connected to a server in a part of the network that has secure physical access (i.e. the doors to the wiring closet and data center are locked), so this port is set to the Authorized state, meaning that the device connected to this port does not need to authenticate using 802.1X. Port 24 is the uplink to a router and is also in the Authorized state.

Figure 10-12. 802.1X Example



The following example shows how to configure the example shown in Figure 10-12.

- 1 Configure the RADIUS server IP address and a global shared secret (secret).

```
console#configure
console(config)#radius server auth 10.10.10.10
console(config-auth-radius)#name Default-RADIUS-Server
console(config-auth-radius)#exit
console(config)#radius server key secret
console(config)#exit
```

- 2 Enable 802.1X port-based access control on the switch.

```
console(config)#dot1x system-auth-control
```

- 3 Configure ports 9 and 24 to be in the Authorized state, which allows the devices to connect to these ports to access the switch services without authentication.

```
console(config)#interface range Gi1/0/9,Gi1/0/24
```

```
console(config-if)#dot1x port-control force-authorized
console(config-if)#exit
```

- 4 Configure Port 7 to require MAC-based authentication with MAB. By default, EAP-MD5 authentication is used.

```
console(config)#interface gil/0/7
console(config-if-Gil/0/7)#dot1x port-control mac-based
console(config-if-Gil/0/7)#mab
```

- 5 Configure the port in general mode. General mode is required for MAC-based authentication.

```
console(config-if-Gil/0/7)#switchport mode general
console(config-if-Gil/0/7)#exit
```

- 6 Enable MAC-based authentication on port 8 and limit the number of devices that can authenticate on that port to 2.

```
console(config)#interface gil/0/8
console(config-if-Gil/0/8)#dot1x port-control mac-based
console(config-if-Gil/0/8)#dot1x max-users 2
```

- 7 Configure the port in general mode. General mode is required for MAC-based authentication.

```
console(config-if-Gil/0/8)#switchport mode general
console(config-if-Gil/0/8)#exit
console(config)#exit
```

- 8 View the client connection status.

When the clients on Ports 1, 3, and 7 (supplicants), attempt to communicate via the switch, the switch challenges the supplicants for 802.1X credentials. The switch encrypts the provided information and transmits it to the RADIUS server. If the RADIUS server grants access, the system sets the 802.1X port state of the interface to authorized and the supplicants are able to access network resources.

```
console#show dot1x clients all
```

```
Interface..... Gil/0/1
User Name..... barneyr
Supp MAC Address..... 0012.1753.031A
Session Time..... 756
Filter Id.....
DACL Name.....
RADIUS Framed IPv4/IPv6 address.....
VLAN Assigned..... 1 (Default)
```



```

Interface..... Gil/0/3
User Name..... fredf
Supp MAC Address..... 0004.5A55.EFAD
Session Time..... 826
Filter Id.....
DACL Name.....
RADIUS Framed IPv4/IPv6 address.....
VLAN Assigned..... 1 (Default)

Interface..... Gil/0/7
User Name..... 0006.6B33.06BA
Supp MAC Address..... 0006.6B33.06BA
Session Time..... 826
Filter Id.....
DACL Name.....
RADIUS Framed IPv4/IPv6 address.....
VLAN Assigned..... 1 (Default)

```

9 View a summary of the port status.

```

console#show dot1x
Administrative Mode..... Enabled
Dynamic VLAN Creation Mode..... Disabled
VLAN Assignment Mode..... Disabled
Monitor Mode..... Disabled
EAPOL Flood Mode..... Disabled

```

| Port | Admin Mode | Oper Mode | Reauth Control | Reauth Period |
|----------|------------------|------------|----------------|---------------|
| Gil/0/1 | auto | Authorized | FALSE | 3600 |
| Gil/0/2 | auto | N/A | FALSE | 3600 |
| Gil/0/3 | auto | Authorized | FALSE | 3600 |
| Gil/0/4 | auto | N/A | FALSE | 3600 |
| Gil/0/5 | auto | N/A | FALSE | 3600 |
| Gil/0/6 | auto | N/A | FALSE | 3600 |
| Gil/0/7 | mac-based | Authorized | FALSE | 3600 |
| Gil/0/8 | mac-based | N/A | FALSE | 3600 |
| Gil/0/9 | force-authorized | Authorized | FALSE | 3600 |
| Gil/0/10 | force-authorized | Authorized | FALSE | 3600 |
| Gil/0/11 | auto | N/A | FALSE | 3600 |

10 View 802.IX information about Port 8.

```
console#show dot1x interface Gi1/0/8
```

```
Administrative Mode..... Enabled
Dynamic VLAN Creation Mode..... Enabled
VLAN Assignment Mode..... Disabled
Monitor Mode..... Disabled

Port      Admin          Oper          Reauth      Reauth
  Mode          Mode          Control      Period
-----
Gi1/0/8   mac-based     Authorized    FALSE       3600

Quiet Period..... 60
Transmit Period..... 30
Maximum Request-Identities..... 2
Maximum Requests..... 2
Max Users..... 2
VLAN Assigned..... 1 (Default)
Supplicant Timeout..... 30
Guest-vlan Timeout..... 90
Server Timeout (secs)..... 30
MAB mode (configured)..... Disabled
MAB mode (operational)..... Disabled
MAB Protocol..... EAP
Authentication Server Dead action for Voice.... None
Authentication Server Alive action..... None
Authenticator PAE State..... Initialize
Backend Authentication State..... Idle
```

Controlling Authentication-Based VLAN Assignment

The network in this example uses three VLANs to control access to network resources. When a client connects to the network, it is assigned to a particular VLAN based on one of the following events:

- It attempts to contact the 802.IX server and is authenticated.
- It attempts to contact the 802.IX server and fails to authenticate.
- It does not attempt to contact the 802.IX server.

The following table describes the three VLANs:

| VLAN ID | VLAN Name | VLAN Purpose |
|---------|------------|------------------------------|
| 100 | Authorized | Data from authorized clients |

| VLAN ID | VLAN Name | VLAN Purpose |
|---------|--------------|--|
| 200 | Unauthorized | Data traffic from clients that fail the authentication with the RADIUS server |
| 300 | Guest | Data traffic from clients that do not attempt to authenticate with the RADIUS server |



NOTE: Dynamic VLAN creation applies only to authorized ports. The VLANs for unauthorized and guest users must be configured on the switch and cannot be dynamically created based on RADIUS-based VLAN assignment.



NOTE: RADIUS VLAN assignment is supported for all port modes other than trunk mode.

The commands in this example show how to configure the switch to control VLAN assignment for the example network. This example also contains commands to configure the uplink, or trunk, port (a port connected to a router or the internal network), and to configure the downlink, or access, ports (ports connected to one or more hosts). Ports 1–23 are downstream ports. Port 24 is an uplink port. An external RADIUS server handles the VLAN assignment.



NOTE: The configuration to control the VLAN assignment for authorized users is done on the external RADIUS server.

To configure the switch:

- 1 Create the VLANs and configure the VLAN names.

```
console(config)#vlan 100  
console(config-vlan100)#name Authorized  
console(config-vlan100)#exit
```

```
console(config)#vlan 200  
console(config-vlan200)#name Unauthorized  
console(config-vlan200)#exit
```

```
console(config)#vlan 300  
console(config-vlan300)#name Guest  
console(config-vlan300)#exit
```

- 2 Configure information about the external RADIUS server the switch uses to authenticate clients. The RADIUS server IP address is 10.10.10.10, and the global shared secret is qwerty123.

```
console(config)#radius server key qwerty123  
console(config)#radius server auth 10.10.10.10  
console(config-auth-radius)#name MyRadius  
console(config-auth-radius)#exit
```

- 3 Enable 802.1X on the switch.

```
console(config)#dot1x system-auth-control
```

- 4 Create a default authentication login list and use the RADIUS server for port-based authentication for connected clients.

```
console(config)#aaa authentication dot1x default radius
```

- 5 Allow the switch to accept VLAN assignments by the RADIUS server.

```
console(config)#aaa authorization network default radius
```

- 6 Enter interface configuration mode for the downlink ports.

```
console(config)#interface range Gi1/0/1-23
```

- 7 Set the downlink ports to the access mode because each downlink port connects to a single host that belongs to a single VLAN. Set the port control mode to auto (default) to allow VLAN assignment from the RADIUS server.

```
console(config-if)#switchport mode access  
console(config-if)#dot1x port-control auto
```

- 8 Enable periodic reauthentication of the client on the ports and set the number of seconds to wait between reauthentication attempts to 300 seconds. Reauthentication is enabled to increase security. If the client information is removed from the RADIUS server after it has been authenticated, the client will be denied access when it attempts to reauthenticate.

```
console(config-if)#dot1x reauthentication
console(config-if)#dot1x timeout re-authperiod 300
```

- 9 Set the unauthenticated VLAN on the ports to VLAN 200 so that any client that connects to one of the ports and fails the 802.1X authentication is placed in VLAN 200.

```
console(config-if)#dot1x unauth-vlan 200
```

- 10 Set the guest VLAN on the ports to VLAN 300. This command automatically enables the Guest VLAN Mode on the downlink ports. Any client that connects to the port and does not attempt to authenticate is placed on the guest VLAN.

```
console(config-if)#dot1x guest-vlan 300
console(config-if)#exit
```

- 11 Enter Interface Configuration mode for port 24, the uplink (trunk) port.

```
console(config)#interface Gi1/0/24
```

- 12 Disable 802.1X authentication on the interface. This causes the port to transition to the authorized state without any authentication exchange required. This port does not connect to any end-users, so there is no need for 802.1X-based authentication.

```
console(config-if-Gi1/0/24)#dot1x port-control force-authorized
```

- 13 Set the uplink port to trunk mode so that it accepts tagged traffic and transmits it to the connected device (another switch or router).

```
console(config-if-Gi1/0/24)#switchport mode trunk
```

Allowing Dynamic Creation of RADIUS-Assigned VLANs

The network in this example uses a RADIUS server to provide VLAN assignments to host that connect to the switch. In this example, the VLANs are not configured on the switch. Instead, the switch is configured to allow the dynamic creation of VLANs when a RADIUS-assigned VLAN does not already exist on the switch.

In this example, Ports 1–23 are configured as downlink, or access, ports, and Port 24 is the trunk port. As a trunk port, Port 24 is automatically added as a member to all VLANs that are statically or dynamically configured on the switch. However, the network administrator in this example has determined that traffic in VLANs 1000–2000 should not be forwarded on the trunk port, even if the RADIUS server assigns a connected host to a VLAN in this range, and the switch dynamically creates the VLAN.



NOTE: The configuration to control the VLAN assignment for hosts is done on the external RADIUS server.

To configure the switch:

- 1 Configure information about the external RADIUS server the switch uses to authenticate clients. The RADIUS server IP address is 10.10.10.10, and the global shared secret is qwerty123.

```
console(config)#radius server key qwerty123
console(config)#radius server 10.10.10.10
console(config-auth-radius)#name MyRadius
console(config-auth-radius)#exit
```

- 2 Enable 802.1X on the switch.

```
console(config)#dot1x system-auth-control
```

- 3 Create a default authentication login list and use the RADIUS server for port-based authentication for connected clients.

```
console(config)#aaa authentication dot1x default radius
```

- 4 Allow the switch to accept VLAN assignments by the RADIUS server.

```
console(config)#aaa authorization network default radius
```

- 5 Allow the switch to dynamically create VLANs when a RADIUS-assigned VLAN does not exist on the switch.

```
console(config)#dot1x dynamic-vlan enable
```

- 6 Enter interface configuration mode for the downlink ports.

```
console(config)#interface range Gi1/0/1-23
```

- 7 Set the downlink ports to the access mode because each downlink port connects to a single host that belongs to a single VLAN. Set the port-control mode to auto (the default) to allow assignment of the dynamically created VLANs to the host connected port.

```
console(config-if)#switchport mode access
console(config-if)#dot1x port-control auto
console(config-if)#exit
```

- 8 Enter Interface Configuration mode for port 24, the uplink (trunk) port.

```
console(config)#interface Gi1/0/24
```

- 9 Disable 802.1X authentication on the interface. This causes the port to transition to the authorized state without any authentication exchange required. This port does not connect to any end-users, so there is no need for 802.1X-based authentication.

```
console(config-if-Gi1/0/24)#dot1x port-control force-authorized
```

- 10 Set the uplink port to trunk mode so that it accepts tagged traffic and transmits it to the connected device (another switch or router). The trunk port will automatically become a member of any dynamically created VLANs unless configured to exclude them.

```
console(config-if-Gi1/0/24)#switchport mode trunk
```

- 11 Forbid the trunk from forwarding traffic that has VLAN tags for any VLAN from 1000–2000, inclusive.

```
console(config-if-Gi1/0/24)#switchport trunk allowed vlan
remove 1000-2000
console(config-if-Gi1/0/24)#exit
```

Configuring Authentication Server Dynamic ACL or DiffServ Policy Assignments

To enable Dynamic ACL or DiffServ policy assignment by an external server, the following conditions must be true:

- The RADIUS or 802.1X server must specify the name of the ACL or policy to assign.

For example, if the DiffServ policy to assign is named `internet_access`, include the following attribute in the RADIUS server configuration:

```
Filter-id (11) = "internet_access"
```

If it is desired that an existing ACL be configured, include the following attribute in the RADIUS server configuration:

```
Filter-ID(11) = "Existing_ACL"
```

- The ACL or DiffServ policy specified in the attribute must already be configured on the switch, and the ACL/policy names must be identical to the one sent by the RADIUS server.

For information about configuring a DiffServ policy, see "DiffServ Configuration Examples" on page 1547. For information about configuring a Dynamic ACL, see "Dynamic ACL Overview" on page 295. The example "Providing Subnets Equal Access to External Network" on page 1547, describes how to configure a policy named `internet_access`.

If you use an authentication server to assign ACLs or DiffServ policies to an authenticated user, note the following guidelines:

- If the policy or ACL specified within the server Filter-ID attribute does not exist on the switch, authentication will fail.
- If a DiffServ policy and ACL have the same name, the named ACL is applied.
- Do not delete policies or ACLs used as the Filter-ID by the RADIUS server while 802.1X is enabled.
- Do not use the DiffServ **service-policy** command to apply the filter to an interface if you configure the RADIUS server or 802.1X authenticator to assign the DiffServ filter.

In the following example, Company XYZ uses IEEE 802.1X to authenticate all users. Contractors and temporary employees at Company XYZ are not permitted to have access to SSH ports, and data rates for Web traffic is limited. When a contractor is authenticated by the RADIUS server, the server assigns a DiffServ policy to control the traffic restrictions.

The network administrator configures two DiffServ classes: `cl-ssh` and `cl-http`. The class `cl-ssh` matches all incoming SSH packets. The class `cl-http` matches all incoming HTTP packets. Then, the administrator configures a traffic policy called `con-pol` and adds the `cl-ssh` and `cl-http`. The policy is configured so that SSH packets are to be dropped, and HTTP data rates are limited to 1 MB with a burst size of 64 Kbps. HTTP traffic that exceeds the limit is dropped. The host ports, ports 1–23, are configured to use MAC-based dot1x authentication to allow the DiffServ policy to be applied. Finally, the administrator configures the RADIUS server with the attribute Filter-id (11) = "con-pol" (steps not shown).

To configure the switch:

- 1 Configure the DiffServ traffic class that matches SSH traffic.

```
console#configure
console(config)#class-map match-all cl-ssh
console(config-classmap)#match dst14port 22
console(config-classmap)#exit
```

- 2 Configure the DiffServ traffic class that matches HTTP traffic.

```
console(config)#class-map match-all cl-http
console(config-classmap)#match dst14port 80
console(config-classmap)#exit
```

- 3 Configure the DiffServ policy.

```
console(config)#policy-map con-pol in
console(config-policy-map)#class cl-ssh
console(config-policy-classmap)#drop
console(config-policy-classmap)#exit
console(config-policy-map)#class cl-http
console(config-policy-classmap)#police-simple 1000000 64
console(config-policy-classmap)#police-action transmit violate-action drop
console(config-policy-classmap)#exit
console(config-policy-map)#exit
```

- 4 Enable DiffServ on the switch. (Optional as DiffServ is enabled by default.)

```
console(config)#diffserv
```

- 5 Configure information about the external RADIUS server the switch uses to authenticate clients. The RADIUS server IP address is 10.10.10.10, and the global shared secret is qwerty123.

```
console(config)#radius server key qwerty123
console(config)#radius server 10.10.10.10
console(config-auth-radius)#name MyRadius
console(config-auth-radius)#exit
```

- 6 Enable 802.1X on the switch.

```
console(config)#dot1x system-auth-control
```

- 7 Create a default authentication login list and use the RADIUS server for port-based authentication for connected clients.

```
console(config)#aaa authentication dot1x default radius
```

- 8 Enter Interface Configuration mode for ports 1–23 and enable MAC-based authentication.

```
console(config)#interface range Gi1/0/1-23
console(config-if)#dot1x port-control mac-based
```

- 9 Set the ports to access mode (default VLAN 1). Enable the policy on the ports.

```
console(config-if)#switchport mode access
console(config-if)#service-policy in con-pol
console(config-if)#exit
console(config)#exit
```

Captive Portal

This section describes how to configure the Captive Portal feature.

The topics covered in this section include:

- Captive Portal Overview
- Default Captive Portal Behavior and Settings
- Configuring Captive Portal (Web)
- Configuring Captive Portal (CLI)
- IEEE 802.1X Configuration Examples

Captive Portal Overview

A Captive Portal (CP) helps manage or restrict network access. CPs are often used in locations that provide wired Internet access to customers, such as business centers and hotels. For example, a hotel might provide an Ethernet port in each room so that guests can connect to the Internet during their stay. The hotel might charge for Internet use, or the hotel might allow guests to connect only after they indicate that they have read and agree to the acceptable use policy.


What Does Captive Portal Do?

The CP feature allows you to require a user to enter login information on a custom web page before gaining access to the network. When the user connects to the port and opens a browser, the user is presented with a welcome screen. To gain network access, the user must enter a username (for guest access) or a username and password (for authenticated access) and

accept the terms of use. The network administrator can also configure the CP feature to redirect the user to another web page after successful authentication, for example a company home page.

CP is supported in IPv4 networks only.

Figure 10-13. Connecting to the Captive Portal



Default Captive Portal Welcome Screen (Displays in Captive Portal User's Browser)

The CP feature blocks hosts connected to the switch from most network access until user verification has been established. Access to 802.1X, DHCP, ARP, NetBIOS, and DNS services is allowed. The network administrator can configure CP verification to allow access for both guest and authenticated users. Authenticated users must be validated against a database of authorized CP users before access is granted. The database can be stored locally on the switch or on a RADIUS server.

Is Captive Portal Dependent on Any Other Feature?

If security procedures require RADIUS authentication, the administrator must configure the RADIUS server information on the switch (see "Using RADIUS" on page 288). The RADIUS administrator must also configure the RADIUS attributes for CP users on the RADIUS server. For information about the RADIUS attributes to configure, see Table 10-15.

For a list of RADIUS attributes that the switch supports, see "Which RADIUS Attributes Does the Switch Support?" on page 290.

To support redirection of user entered URLs from a web browser, a DNS server must be configured in the network. If routing is enabled on the switch, IP helper should be enabled to allow hosts to obtain an IP address via DHCP. A DHCP server must be available if it is expected that hosts will obtain IP addresses dynamically. In addition, if routing is enabled, DHCP relay must be configured.

The only type of interface where CP can be enabled is a physical port. CP is not supported on multi-access VLANs or on LAGs.

A physical port's VLAN membership does not affect CP. A physical port enabled for CP can be a member of any VLAN or multiple VLANs, which can be switching or routing VLANs.

A port enabled for CP may be directly connected to a single client (e.g., an access switch), or the port may serve many clients (e.g., a port on an aggregation switch).

Port security and CP cannot both be enabled on the same interface.

If a physical port configured with CP is made a member of a LAG, CP is disabled on the port.

Dell EMC Networking does not support configuring spanning tree on a CP port. BPDUs received on a port enabled for CP will not receive their normal prioritization.

CP can coexist on an interface with DHCP snooping and Dynamic ARP Inspection (DAI).

The administrator can configure the switch to send SNMP trap messages to any enabled SNMP Trap Receivers for several CP events, such as when a CP user has an authentication failure or when a CP user successfully connects to

the network. If traps are enabled, the switch also writes a message to the trap log when the event occurs. To enable the CP traps, see "Configuring SNMP Notifications (Traps and Informs)" on page 515.

What Factors Should Be Considered When Designing and Configuring a Captive Portal?

Before enabling the CP feature, decide what type (or types) of authentication will be supported. Since Dell EMC Networking N-Series switches support up to 10 different CP instances, it is possible to configure one CP that requires a username and password and another that only requires the username. For each CP, the administrator can customize the welcome screen, including the colors and logo.

If network policy requires authentication, consider the number of users that must exist in the user database. The local user database supports up to 128 users. If there is a need to support more than 128 authenticated users, use a remote RADIUS server for authentication.

The administrator can specify whether the CP uses HTTP or HTTPS as the protocol during the user verification process. HTTP does not use encryption during verification, and HTTPS uses the Secure Sockets Layer (SSL), which requires a certificate to provide encryption. The certificate is presented to the user at connection time.

If the authenticating user requires DNS or DHCP services, these will need to be configured in the network and the switch will need to relay DHCP packets.

The initial Web page that a user sees when he or she connects to the CP can be customized. The logo, color schemes, welcome messages, and all text on the page can be customized, including the field and button labels. The welcome page the user sees after a successful verification or authentication can also be customized.

Figure 10-14. Customized Captive Portal Welcome Screen



How Does Captive Portal Work?

When a port is enabled for CP, all the traffic coming onto the port from the unverified clients is dropped except for the ARP, DHCP, NetBIOS, and DNS packets. These packets are forwarded by the switch so that the unverified clients can get an IP address and are able to resolve host or domain names. If an unverified client opens a web browser and tries to connect to the network, CP redirects all the HTTP/HTTPS traffic from the unverified client to the authenticating server on the switch. If the network administrator has configured an additional web server port, packets with this destination TCP port number are also forwarded to the authenticating server. A CP web page is sent back to the unverified client. If the verification mode for the CP associated with the port is Guest, the client can be verified without providing authentication information. If the verification mode is Local or RADIUS, the client must provide credentials that are compared against the information in the Local or RADIUS client database. After the user successfully provides the required information, the CP feature grants access to the network.

What Captive Portal Pages Can Be Customized?

The following three CP pages can be customized:

- Authentication Page —This page displays when a client attempts to connect to the network. The images, text, and colors that display on this page can be customized.

- **Logout Page** — If the user logout mode is enabled, this page displays in a pop-up window after the user successfully authenticates. This window contains the logout button.
- **Logout Success Page** — If the user logout mode is enabled, this page displays after a user clicks the logout button and successfully deauthenticates.

Understanding User Logout Mode

The User Logout Mode feature allows a user who successfully authenticates to the network through the CP to explicitly deauthenticate from the network. When User Logout Mode is disabled or the user does not specifically request logout, the connection status will remain authenticated until the CP deauthenticates the user based on the configured session timeout value. In order for the user logout feature to function properly, the client browser must have JavaScript enabled and must allow popup windows.

Localizing Captive Portal Pages

The CP localization feature allows you to create up to three language-specific web pages for each CP as long as all pages use the same verification type; either guest or authorized user web pages. This allows you to create pages in a variety of languages to accommodate a diverse group of users.

To customize the pages that the user sees, click the language tab. By default, the English tab is available. The settings for the **Authentication Page** display.

Captive Portal IP Address Selection

CP automatically associates with one of the IP addresses assigned to the switch. The automatic IP address selection algorithm is outlined below:

- 1 On switching-only devices or when routing is disabled, CP uses the out-of-band interface IP address, if available.
- 2 If routing is enabled, CP uses a loopback interface if one is defined, and a routing interface as the second choice.
- 3 If routing is enabled and no active routing interface is available, the CP goes down.
- 4 If the CP IP address changes due to administrator action or due to an interface going down, then the CP is automatically disabled and re-enabled. All active sessions are dropped.

Captive Portal and DNS

CP allows unauthenticated users access to DNS services on TCP and UDP destination port 53. CP inspects all DNS traffic to ensure that it conforms with the DNS protocol (RFC 1035/1996). CP checks the format of DNS messages and discards packets that do not conform to the minimum standards. Specifically, CP performs the following checks on a DNS packet:

- The packet must have a full-size header and at least one question field
- The packet must have a valid DNS response code
- The first question field must not exceed 63 octets in length, nor must the length field be greater than 63
- The first question class field must be valid.

Captive Portal Troubleshooting

The following table explains the status values for CP authentication sessions and the resulting actions taken, if any. CP global status, interface status, and session status are available in the user interfaces.

Table 10-13. Captive Portal Status Values

| Status Value | Description | Browser Action |
|--------------|---|--|
| Default | Initial request from the client. | Used to detect initial request. |
| Serve | Default serve. | Used when serving the initial connection page. |
| Validate | Actual validation request. | Indicates that the user has submitted credentials and requests authentication. |
| WIP | Indicates that validation is in progress. | The validation page begins to poll the server until the status flag changes. The actual poll request is the same http(s) request used to “validate” as described above. While waiting between polls, the browser displays an “authorization in process” message. |


Table 10-13. Captive Portal Status Values (Continued)

| Status Value | Description | Browser Action |
|---------------------|---|---|
| RADIUS_WIP | Indicates that RADIUS validation is in progress. | The browser action is the same as for the WIP status. |
| Success | Indicates that authentication is a success. | Displays either the customized welcome page or an external URL. |
| Denied | Indicates that the user has failed to enter credentials that match the expected configuration. | The default serve page is resubmitted and includes the appropriate failure message. |
| Resource | Indicates that the system has rejected authentication due to system resource limitations or session timeout. | The default serve page is resubmitted and includes the appropriate failure message. |
| No Accept | Indicates that the user did not accept the acceptance use policy. | The default serve page is resubmitted and includes the appropriate failure message. |
| Timeout | Indicates that the authentication transaction took too long. This could be due to user input time, or a timeout due to the overall transaction. | The default serve page is resubmitted and includes the appropriate failure message. |

Default Captive Portal Behavior and Settings

CP is disabled by default. If you enable CP, no interfaces are associated with the default CP. After you associate an interface with the CP and globally enable the CP feature, a user who connects to the switch through that interface is presented with the CP Welcome screen shown in Figure 10-15.

Figure 10-15. Default Captive Portal Welcome Screen



The user types a name in the Username field, selects the Acceptance Use Policy check box, and clicks **Connect** to gain network access. By default, the user does not need to be defined in a database or enter a password to access the network because the default verification mode is Guest. Note that duplicate Username entries can exist in this mode because the client IP and MAC addresses are obtained for identification.

Table 10-14 shows the default values for the CP feature.


Table 10-14. Default Captive Portal Values

| Feature | Value |
|--|--|
| Global Captive Portal Operational Status | Disabled |
| Additional HTTP or HTTPS Ports | Disabled CP can be configured to use an additional HTTP and/or HTTPS port (in support of Proxy networks). |
| Authentication Timeout | 300 seconds |

Table 10-14. Default Captive Portal Values

| Feature | Value |
|--------------------------------|--|
| Configured Captive Portals | 1 |
| Captive Portal Name | Default |
| Protocol Mode | HTTP |
| Verification Mode | Guest |
| URL Redirect Mode | Off |
| User Group | 1-Default |
| Session Timeout | 86400 seconds |
| Local Users | None configured |
| Interface associations | None |
| Interface status | Not blocked If the CP is blocked, users cannot gain access to the network through the CP. Use this function to temporarily protect the network during unexpected events, such as denial of service attacks. |
| Supported Captive Portal users | 1024 |
| Supported local users | 128 |
| Supported Captive Portals | 10 |

Configuring Captive Portal (Web)


This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring CP settings on Dell EMC Networking N1100-ON, N1500, N2000, N2100-ON, N3000, N3100-ON, and N4000 Series switches. For details about the fields on a page, click  at the top of the Dell EMC OpenManage Switch Administrator web page.

Captive Portal Global Configuration

Use the **Captive Portal Global Configuration** page to control the administrative state of the CP feature and configure global settings that affect all CPs configured on the switch.

To display the **Captive Portal Global Configuration** page, click **System** → **Captive Portal** → **Global Configuration**.

Figure 10-16. Captive Portal Global Configuration




Captive Portal Configuration

Use the **Captive Portal Configuration** page to view summary information about CPs on the system, add a CP, and configure existing CPs.

The switch supports 10 CP configurations. CP configuration 1 is created by default and cannot be deleted. Each CP configuration can have unique guest or group access modes and a customized acceptance use policy that displays when the client connects.


To display the **Captive Portal Configuration** page, click **System** → **Captive Portal** → **Configuration**.

Figure 10-17. Captive Portal Configuration




From the Captive Portal Configuration page, click **Add** to create a new CP instance.

Figure 10-18. Add Captive Portal Configuration



From the Captive Portal Configuration page, click **Summary** to view summary information about the CP instances configured on the switch.

Figure 10-19. Captive Portal Summary



Customizing a Captive Portal

The procedures in this section customize the pages that the user sees when he or she attempts to connect to (and log off of) a network through the CP. These procedures configure the English version of the Default Captive Portal.

To configure the switch:

- 1 From the **Captive Portal Configuration** page click the **(English)** tab. The settings for the **Authentication Page** display, and the links to the CP customization appear.
- 2 Click **Download Image** to download one or more custom images to the switch. A downloaded custom image can be used for the branding logo (default: Dell logo) on the Authentication Page and Logout Success page, for the account image (default: blue banner with keys) on the Authentication Page, and for the background image (default: blank) on the Logout Success Page.



 **NOTE:** The image to download must be accessible from your local system. The image should be 5 KB max, 200x200 pixels, GIF or JPG format.

Figure 10-20. Captive Portal Download Image Page



- 3** Make sure **Download** is selected in the **Available Images** menu, and click **Browse**.
- 4** Browse to the directory where the image to be downloaded is located and select the image.
- 5** Click **Apply** to download the selected file to the switch.
- 6** To customize the **Authentication Page**, which is the page that a user sees upon attempting to connect to the network, click the **Authentication Page** link.

Figure 10-21. Captive Portal Authentication Page

The screenshot shows a web-based configuration interface for a Captive Portal Authentication Page. The interface is titled "Configuration: Language Authentication Page" and is organized into three main sections: "Greeting and Resources", "Textual Content", and "Messages".

- Greeting and Resources:** This section includes fields for "Captive Portal ID" (Default), "Branding Image" (del_logo.gif), "Fonts" (arial, sans-serif), "Browser Title" (Captive Portal), "Page Title" (Welcome to the Network), "Separator Color" (#003366), "Foreground Color" (#999999), and "Background Color" (#FFFFFF).
- Textual Content:** This section includes fields for "Account Image" (login_key.jpg), "Account Title" (Enter your Username), "User Label" (Username), "Password Label" (Password), "Button Label" (Connect), "Acceptance Use Policy" (Acceptance Use Policy), and "Acceptance Message" (Check here to indicate that you have read and accepted the).
- Messages:** This section includes fields for "Instructional Text" (To start using this service, enter your credentials and click the Connect button), "Denied Message" (Error: Invalid Credentials, please try again!), "Resource Message" (Error: Limited Resources, please reconnect and try again later!), "Timeout Message" (Error: Timed Out, please reconnect and try again!), "Busy Message" (Connecting, please be patient), "No Accept Message" (Error: You must acknowledge the Acceptance Use Policy before connecting!), "Welcome Title" (Congratulations!), and "Welcome Text" (You are now authorized and connected to the network).

At the bottom right of the interface, there are three buttons: "Clear", "Preview", and "Apply".

- 7 Select the branding image to use and customize other page components such as the font for all text the page displays, the page title, and the acceptance use policy.
- 8 Click **Apply** to save the settings to the running configuration or click **Preview** to view what the user will see. To return to the default views, click **Clear**.

- Click the **Logout Page** link to configure the page that contains the logout window.


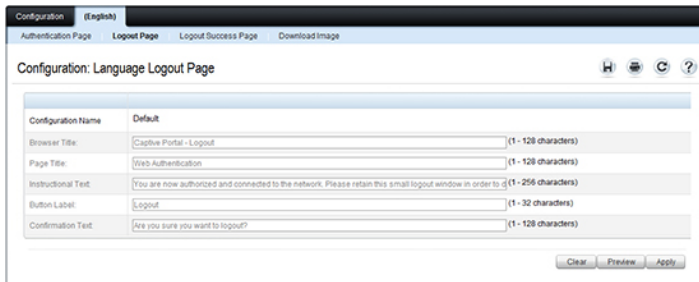
 **NOTE:** The Logout Page settings can be configured only if the User Logout Mode is selected on the Configuration page. The User Logout Mode allows an authenticated client to deauthenticate from the network.

Figure 10-22. Captive Portal Logout Page



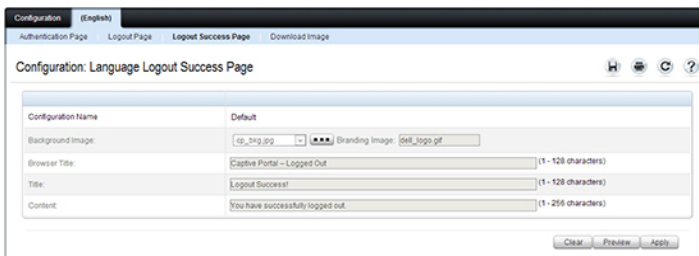
The screenshot shows the configuration page for the 'Language Logout Page'. The page title is 'Configuration: Language Logout Page'. The configuration table is as follows:

| Configuration Name | Default | |
|--------------------|---|----------------------|
| Browser Title | Captive Portal - Logout | (1 - 128 characters) |
| Page Title | Web Authentication | (1 - 128 characters) |
| Instructional Text | You are now authorized and connected to the network. Please retain this small logout window in order to | (1 - 256 characters) |
| Button Label | Logout | (1 - 32 characters) |
| Confirmation Text | Are you sure you want to logout? | (1 - 128 characters) |

Buttons: Clear, Preview, Apply

- Customize the look and feel of the Logout Page, such as the page title and logout instructions.
- Click **Apply** to save the settings to the running configuration or click **Preview** to view what the user will see. To return to the default views, click **Clear**.
- Click the **Logout Success Page** link to configure the page that contains the logout window. A user is required to logout only if the User Logout Mode is selected on the **Configuration** page.

Figure 10-23. Captive Portal Logout Success Page



The screenshot shows the configuration page for the 'Language Logout Success Page'. The page title is 'Configuration: Language Logout Success Page'. The configuration table is as follows:

| Configuration Name | Default | |
|--------------------|-----------------------------------|------------------------------|
| Background Image | cp_bg.jpg | Branding Image: SWL_logo.gif |
| Browser Title | Captive Portal - Logged Out | (1 - 128 characters) |
| Title | Logout Successful | (1 - 128 characters) |
| Content | You have successfully logged out. | (1 - 256 characters) |

Buttons: Clear, Preview, Apply

- 13 Customize the look and feel of the Logout Page, such as the background image and successful logout message.
- 14 Click **Apply** to save the settings to the running configuration or click **Preview** to view what the user will see. To return to the default views, click **Clear**.

Local User

A portal can be configured to accommodate guest users and authorized users. Guest users do not have assigned user names and passwords. Authorized users provide a valid user name and password that must first be validated against a local database or RADIUS server. Authorized users can gain network access once the switch confirms the user's credentials.

By default, each CP instance contains the default group. The default group can be renamed, or a different group can be created and assigned to each CP instance. A CP instance can be associated to one user group only. A user, however, can be assigned to multiple groups.

The **Local User** page allows you to add authorized users to the local database, which can contain up to 128 user entries. Users can be added to and deleted from the local database using the **Local User** page.


To display the **Local User** page, click **System** → **Captive Portal** → **Local User**.

Figure 10-24 shows the **Local User** page after a user has been added. If no users have been added to the switch, many of the fields do not display on the screen.




NOTE: Multiple user groups can be selected by holding the CTRL key down while clicking the desired groups.

Figure 10-24. Local User Configuration




From the **Local User** page, click **Add** to add a new user to the local database.

Figure 10-25. Add Local User



From the **Local User** page, click **Show All** to view summary information about the local users configured in the local database.

Figure 10-26. Captive Portal Local User Summary



To delete a configured user from the database, select the Remove check box associated with the user and click **Apply**.

Configuring Users in a Remote RADIUS Server

A remote RADIUS server client authorization can be used. All users must be added to the RADIUS server. The local database does not share any information with the remote RADIUS database.

Table 10-15 indicates the RADIUS attributes you use to configure authorized CP clients. The table indicates both RADIUS attributes and vendor-specific attributes (VSA). VSAs are denoted in the Attribute column and are comma delimited (vendor ID, attribute ID).

Table 10-15. Captive Portal User RADIUS Attributes

| Attribute | Number | Description | Range | Usage | Default |
|----------------------------|-----------|--|-------------------|----------|---|
| User-Name | 1 | User name to be authorized | 1-32 characters | Required | None |
| User-Password | 2 | User password | 8-64 characters | Required | None |
| Session-Timeout | 27 | Logout once session timeout is reached (seconds). If the attribute is 0 or not present then use the value configured for the CP. | Integer (seconds) | Optional | 0 |
| Dell-Captive-Portal-Groups | 6231, 127 | A comma-delimited list of group names that correspond to the configured CP instance configurations. | String | Optional | None. The default group is used if not defined here |