



Cisco Nexus 9000 Series NX-OS Programmability Guide, Release 10.4(x)

First Published: 2023-08-18

Last Modified: 2024-03-29

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: <https://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <https://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <https://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2023–2024 Cisco Systems, Inc. All rights reserved.



CONTENTS

| | | |
|------------------|------------------------------------|----------|
| CHAPTER 1 | New and Changed Information | 1 |
| | New and Changed Information | 1 |

| | | |
|------------------|--|----------|
| CHAPTER 2 | Platform Support for Programmability Features | 5 |
| | Platform Support for Programmability Features | 5 |

| | | |
|------------------|---|----------|
| CHAPTER 3 | Overview | 9 |
| | Programmability Overview | 9 |
| | Supported Platforms | 10 |
| | Standard Network Manageability Features | 10 |
| | Advanced Automation Features | 10 |
| | Programmability Support | 10 |
| | NX-API Support | 10 |
| | Python Scripting | 11 |
| | Tel Scripting | 11 |
| | Broadcom Shell | 11 |
| | Bash | 11 |
| | Bash Shell Access and Linux Container Support | 11 |
| | Guest Shell | 11 |
| | Container Tracker Support | 13 |
| | Perl Modules | 13 |
| | FPM Parallel Lookup | 14 |

| | | |
|---------------|-----------------------------|-----------|
| PART I | Shells and Scripting | 17 |
|---------------|-----------------------------|-----------|

| | | |
|------------------|-----------------------------|-----------|
| CHAPTER 4 | Shells and Scripting | 19 |
|------------------|-----------------------------|-----------|

| | |
|---|----|
| About Bash | 19 |
| Guidelines and Limitations | 19 |
| Enabling Consent Token | 20 |
| Accessing Bash | 21 |
| Escalate Privileges to Root | 22 |
| Examples of Bash Commands | 24 |
| Displaying System Statistics | 24 |
| Running Bash from CLI | 24 |
| Managing Feature RPMs | 25 |
| RPM Installation Prerequisites | 25 |
| Installing Feature RPMs from Bash | 25 |
| Upgrading Feature RPMs | 26 |
| Downgrading a Feature RPM | 27 |
| Erasing a Feature RPM | 27 |
| Support for DME Modularity | 28 |
| Installing the DME RPMs | 28 |
| Verifying the Installed RPM | 31 |
| Querying for the RPM in the Local Repo | 32 |
| Downgrading Between Versions of DME RPM | 32 |
| Downgrading to the Base RPM | 34 |
| Managing Patch RPMs | 36 |
| RPM Installation Prerequisites | 36 |
| Adding Patch RPMs from Bash | 37 |
| Activating a Patch RPM | 38 |
| Committing a Patch RPM | 40 |
| Deactivating a Patch RPM | 41 |
| Removing a Patch RPM | 42 |
| Persistently Daemonizing an SDK- or ISO-built Third Party Process | 43 |
| Persistently Starting Your Application from the Native Bash Shell | 44 |
| Synchronize Files from Active Bootflash to Standby Bootflash | 44 |
| Copy Through Kstack | 46 |
| An Example Application in the Native Bash Shell | 46 |

| | |
|---|----|
| About the Guest Shell | 49 |
| Guidelines and Limitations for Guestshell | 50 |
| Accessing the Guest Shell | 56 |
| Resources Used for the Guest Shell | 57 |
| Capabilities in the Guestshell | 57 |
| NX-OS CLI in the Guest Shell | 58 |
| Network Access in Guest Shell | 58 |
| Access to Bootflash in Guest Shell | 60 |
| Python in Guest Shell | 61 |
| Python in Guestshell 2.11 | 61 |
| Python 3 in Guest Shell versions up to 2.10 (CentOS 7) | 62 |
| Python in Guestshell 4.x | 64 |
| Python 2 in Guest Shell 3.0 (Centos 8) | 64 |
| Installing RPMs in the Guest Shell | 65 |
| Security Posture for Virtual ServicesGuest Shell | 66 |
| Kernel Vulnerability Patches | 67 |
| ASLR and X-Space Support | 67 |
| Namespace Isolation | 67 |
| Root-User Restrictions | 68 |
| Resource Management | 69 |
| Guest File System Access Restrictions | 69 |
| Managing the Guest Shell | 70 |
| Disabling the Guest Shell | 73 |
| Destroying the Guest Shell | 74 |
| Enabling the Guest Shell | 74 |
| Replicating the Guest Shell | 76 |
| Exporting Guest Shell rootfs | 76 |
| Importing Guest Shell rootfs | 76 |
| Importing YAML File | 78 |
| show guestshell Command | 81 |
| Verifying Virtual Service and Guest Shell Information | 82 |
| Persistently Starting Your Application From the Guest Shell | 83 |
| Procedure for Persistently Starting Your Application from the Guest Shell | 84 |
| An Example Application in the Guest Shell | 84 |

Troubleshooting Guest Shell Issues 85

CHAPTER 6**Broadcom Shell 87**

About the Broadcom Shell 87

Guidelines and Limitations 87

Accessing the Broadcom Shell (bcm-shell) 87

 Accessing bcm-shell with the CLI API 87

 Accessing the Native bcm-shell on the Fabric Module 89

 Accessing the bcm-shell on the Line Card 90

Examples of Broadcom Shell Commands 91

 Displaying L2 Entries 91

 Displaying Routing Information from FM and LC ASIC Instances 91

 Displaying Spanning Tree Group Entries 92

 Display T2 Counters for Interface xe0 92

 Displaying L3 Information 92

CHAPTER 7**Barefoot Shell 93**

About the Barefoot Shell 93

Guidelines and Limitations 93

Accessing the Barefoot Shell with CLI API 93

CHAPTER 8**Python API 97**

Using Python 97

 Cisco Python Package 97

 Using the CLI Command APIs 99

 Invoking the Python Interpreter from the CLI 100

 Display Formats 100

 Non-Interactive Python 102

 Running Scripts with Embedded Event Manager 103

 Python Integration with Cisco NX-OS Network Interfaces 104

 Cisco NX-OS Security with Python 104

 Examples of Security and User Authority 105

 Example of Running Script with Scheduler 106

CHAPTER 9**Scripting with Tcl 109**

About Tcl 109

Guidelines and Limitations 109

Tclsh Command Help 109

Tclsh Command History 110

Tclsh Tab Completion 110

Tclsh CLI Command 110

Tclsh Command Separation 110

Tcl Variables 111

Telquit 111

Tclsh Security 111

Running the Tclsh Command 112

Navigating Cisco NX-OS Modes from the Tclsh Command 113

Tcl References 114

CHAPTER 10**iPXE 115**

About iPXE 115

Netboot Requirements 115

Guidelines and Limitations for iPXE 116

Boot Mode Configuration 116

Verifying the Boot Order Configuration 118

CHAPTER 11**Kernel Stack 119**

About Kernel Stack 119

Guidelines and Limitations 119

Changing the Port Range 120

About VXLAN with kstack 121

Setting Up VXLAN for kstack 121

Troubleshooting VXLAN with kstack 121

Netdevice Property Changes 122

PART II**Applications 125**

CHAPTER 12

Third-Party Applications 127

- About Third-Party Applications 127
- Guidelines and Limitations 127
- Installing Python2 and Dependent Packages 128
- Installing Third-Party Native RPMs/Packages 128
- Installing Signed RPM 130
 - Checking a Signed RPM 130
 - Installing Signed RPMs by Manually Importing Key 130
 - Installing Signed Third-Party RPMs by Importing Keys Automatically 132
 - Adding Signed RPM into Repo 134
- Persistent Third-Party RPMs 135
- Installing RPM from VSH 136
 - Package Addition 136
 - Package Activation 137
 - Deactivating Packages 138
 - Removing Packages 138
 - Displaying Installed Packages 139
 - Displaying Detail Logs 139
 - Upgrading a Package 139
 - Downgrading a Package 140
- Third-Party Applications 141
 - NX-OS 141
 - DevOps Configuration Management Tools 141
 - V9K 141
 - Automation Tool Educational Content 141
 - collectd 141
 - Ganglia 142
 - Iperf 142
 - LLDP 142
 - Nagios 142
 - OpenSSH 142
 - Quagga 143
 - Splunk 143

tcollector 143
tcpdump 143
TShark 144

CHAPTER 13 **Using Ansible with Cisco NX-OS** 145

Prerequisites 145
About Ansible 145
Cisco Ansible Module 145

CHAPTER 14 **Puppet Agent** 147

About Puppet 147
Prerequisites 148
Puppet Agent NX-OS Environment 148
ciscopuppet Module 148

CHAPTER 15 **Using Chef Client with Cisco NX-OS** 151

About Chef 151
Prerequisites 151
Chef Client NX-OS Environment 152
cisco-cookbook 152

CHAPTER 16 **Nexus Application Development - Yocto** 155

About Yocto 155
Installing Yocto 155

CHAPTER 17 **Nexus Application Development - SDK** 159

About the Cisco SDK 159
Installing the SDK 159
Procedure for Installation and Environment Initialization 160
Using the SDK to Build Applications 161
Using RPM to Package an Application 162
Creating an RPM Build Environment 163
Using General RPM Build Procedure 163

| | |
|--|-----|
| Example to Build RPM for collectd with No Optional Plug-Ins | 164 |
| Example to Build RPM for collectd with Optional Curl Plug-In | 165 |

CHAPTER 18**NX-SDK 167**

| | |
|---------------------------------------|-----|
| About the NX-SDK | 167 |
| Considerations for Go Bindings | 168 |
| About On-Box (Local) Applications | 168 |
| Default Docker Images | 168 |
| Guidelines and Limitations for NX-SDK | 169 |
| About NX-SDK 2.0 | 170 |
| About NX-SDK 2.5 | 170 |
| About Remote Applications | 170 |
| NX-SDK Security | 171 |
| Security Profiles for NX SDK 2.0 | 171 |

CHAPTER 19**Using Docker with Cisco NX-OS 173**

| | |
|--|-----|
| About Docker with Cisco NX-OS | 173 |
| Guidelines and Limitations for Docker | 173 |
| Prerequisites for Setting Up Docker Containers Within Cisco NX-OS | 174 |
| Starting the Docker Daemon | 174 |
| Configure Docker to Start Automatically | 175 |
| Starting Docker Containers: Host Networking Model | 176 |
| Starting Docker Containers: Bridged Networking Model | 177 |
| Mounting the bootflash and volatile Partitions in the Docker Container | 178 |
| Enabling Docker Daemon Persistence on Enhanced ISSU Switchover | 178 |
| Enabling Docker Daemon Persistence on the Cisco Nexus Platform Switches Switchover | 179 |
| Resizing the Docker Storage Backend | 180 |
| Stopping the Docker Daemon | 182 |
| Docker Container Security | 183 |
| Securing Docker Containers With User namespace Isolation | 183 |
| Moving the cgroup Partition | 184 |
| Adding Nodes to a Kubernetes Cluster | 184 |
| Docker Troubleshooting | 187 |
| Docker Fails to Start | 187 |

| | |
|---|-----|
| Docker Fails to Start Due to Insufficient Storage | 187 |
| Failure to Pull Images from Docker Hub (509 Certificate Expiration Error Message) | 188 |
| Failure to Pull Images from Docker Hub (Client Timeout Error Message) | 188 |
| Docker Daemon or Containers Not Running On Switch Reload or Switchover | 189 |
| Resizing of Docker Storage Backend Fails | 189 |
| Docker Container Doesn't Receive Incoming Traffic On a Port | 190 |
| Unable to See Data Port And/Or Management Interfaces in Docker Container | 190 |
| General Troubleshooting Tips | 190 |

PART III
Application Hosting 191

CHAPTER 20
Application Hosting 193

| | |
|--|-----|
| Guidelines and Limitations for Application Hosting | 193 |
| Information About Application Hosting | 194 |
| Need for Application Hosting | 194 |
| Application Hosting Overview | 194 |
| How to Configure Application Hosting | 195 |
| Enabling Application Hosting Feature | 195 |
| Configuring Application Hosting Bridge Connections | 196 |
| Lifecycle of an Application | 198 |
| Upgrading an Application | 199 |
| Configuring Docker Run Time Options | 200 |
| Configuring Application Hosting on the Management Interface | 201 |
| Overriding Application Resource Configuration | 202 |
| Advanced Application Hosting Features | 203 |
| Copying Application Data | 204 |
| Deleting Application Data | 204 |
| Verifying the Application-Hosting Configuration | 204 |
| Configuration Examples for Application Hosting | 208 |
| Example: Enabling AppHosting Feature | 208 |
| Example: Configuring Application Hosting Bridge Connections | 208 |
| Example: Configuring Docker Run Time Options | 208 |
| Example: Configuring Application Hosting on the Management Interface | 208 |
| Example: Overriding App Resource Configuration | 209 |

Additional References 209
 Feature Information for Application Hosting 209

CHAPTER 21

ThousandEyes Enterprise Agent 211
 ThousandEyes Enterprise Agent Overview 211
 Prerequisites for the ThousandEyes Enterprise Agent 211
 Resources Required for the ThousandEyes Enterprise Agent 212
 Installing the ThousandEyes Enterprise Agent 212
 Guidelines and Limitations 212
 Configuring Application Hosting for the ThousandEyes Enterprise Agent 212
 Installing the ThousandEyes Enterprise Agent 216
 Configuration Examples for ThousandEyes Enterprise Agent 217
 Examples: Installing ThousandEyes Enterprise Agent 217
 Sample Configuration for ThousandEyes Enterprise Agent 217

PART IV

NX-API 219

CHAPTER 22

NX-API CLI 221
 About NX-API CLI 221
 Guidelines and Limitations 221
 Transport 222
 Message Format 222
 Security 223
 Using NX-API CLI 223
 Escalate Privileges to Root on NX-API 225
 NX-API Management Commands 226
 Working With Interactive Commands Using NX-API 232
 NX-API Client Authentication 233
 NX-API Client Basic Authentication 233
 NX-API Client Certificate Authentication 233
 Guidelines and Limitations 233
 NX-API Client Certificate Authentication Prerequisites 235
 Configuring NX-API Client Certificate Authentication 235
 Example Python Scripts for Certificate Authentication 236

| | |
|---|-----|
| Example cURL Certificate Request | 237 |
| Validating Certificate Authentication | 237 |
| NX-API Request Elements | 239 |
| NX-API Response Elements | 244 |
| Restricting Access to NX-API | 245 |
| Updating an iptable | 245 |
| Making an Iptable Persistent Across Reloads | 247 |
| Kernel Stack ACL | 248 |
| Table of NX-API Response Codes | 249 |
| JSON and XML Structured Output | 252 |
| About JSON (JavaScript Object Notation) | 252 |
| Examples of XML and JSON Output | 253 |
| Sample NX-API Scripts | 261 |

CHAPTER 23**NX-API REST 263**

| | |
|--|-----|
| About NX-API REST | 263 |
| DME Config Replace Through REST | 264 |
| About DME Full Config Replace Through REST Put | 264 |
| Guidelines and Limitations | 264 |
| Replacing Property-Level Configuration Through REST POST | 265 |
| Replacing Feature-Level Config Through REST PUT | 265 |
| Troubleshooting Config Replace for REST PUT | 266 |

CHAPTER 24**NX-API Developer Sandbox 269**

| | |
|---|-----|
| NX-API Developer Sandbox: NX-OS Releases Prior to 9.2(2) | 269 |
| About the NX-API Developer Sandbox | 269 |
| Guidelines and Limitations | 270 |
| Configuring the Message Format and Command Type | 272 |
| Using the Developer Sandbox | 274 |
| Using the Developer Sandbox to Convert CLI Commands to REST Payloads | 274 |
| Using the Developer Sandbox to Convert from REST Payloads to CLI Commands | 276 |
| NX-API Developer Sandbox: NX-OS Release 9.2(2) and Later | 281 |
| About the NX-API Developer Sandbox | 281 |
| Guidelines and Limitations | 282 |

| | |
|---|-----|
| Configuring the Message Format and Input Type | 284 |
| Using the Developer Sandbox | 287 |
| Using the Developer Sandbox to Convert CLI Commands to REST Payloads | 287 |
| Using the Developer Sandbox to Convert from REST Payloads to CLI Commands | 290 |
| Using the Developer Sandbox to Convert from RESTCONF to json or XML | 295 |

PART V
Model-Driven Programmability 299

CHAPTER 25
Overview 301

| | |
|---|-----|
| About Model-Driven Programmability | 301 |
| About the Programmable Interface Infrastructure | 301 |
| NX-OS Programmability Agents | 302 |
| NX-OS Observability Agents | 303 |
| Model Infrastructure | 303 |
| Native (Device) YANG Model | 303 |
| Common (OpenConfig) YANG Models | 303 |
| Additional YANG References | 304 |

CHAPTER 26
Original and OpenConfig YANG 305

| | |
|--|-----|
| About YANG | 305 |
| Cisco Device YANG | 306 |
| Guidelines and Limitations | 306 |
| Migration from DME to Device YANG | 306 |
| About OpenConfig YANG | 307 |
| Guidelines and Limitations | 307 |
| OpenConfig Paths | 307 |
| Guidelines and Limitations for High Scale Data | 327 |
| Configuring OpenConfig Support | 328 |
| Upgrading or Downgrading YANG Version | 329 |
| RBAC for YANG | 329 |
| Guidelines and Limitations | 329 |
| Configuring YANG RBAC | 330 |
| Troubleshooting YANG | 331 |

CHAPTER 27**NETCONF Agent 333**

- About the NETCONF Agent 333
- Revision History 334
- Guidelines and Limitations for NETCONF 334
- Configuring the NETCONF Agent 336
 - Configuring the NETCONF Agent Over SSH 336
 - Configuring the NETCONF Agent Over TLS 336
 - Generating Key/Certificate Example 338
- Manage a NETCONF Session 339
 - Start Session 339
 - Terminate Session 340
- NETCONF Get / Set 341
- NETCONF Notifications 349
 - About NETCONF Notifications 349
 - Capabilities Exchange 349
 - Event Stream Discovery 350
 - Create Subscriptions 350
 - Receive Notifications 351
 - Terminate Subscriptions 352
- NETCONF Device YANG RPC 352
 - About Model Driven RPC in NETCONF 352
 - Native Device RPC Examples 353
- Netconf Client Examples 355
 - Using the Sandbox to Generate the NETCONF Payload 355
 - Connecting Cisco NX-OS with the ncclient 356
 - Getting Configuration Data 356
 - Getting the Running Configuration and Operational Data 357
 - Creating a New Configuration 357
 - Deleting Configuration 358
- Troubleshooting the NETCONF Agent 359
 - Check Feature Status 359
 - Check Connectivity 359
 - Check TLS 359

| | |
|----------------------------------|-----|
| Accounting Log for NETCONF Agent | 360 |
| About NETCONF Explicit Mode | 362 |
| Guidelines and Limitations | 362 |
| Topic 2.1 | 363 |
| NETCONF Explicit Mode Get/Set | 363 |
| Add Configuration Using CLI | 367 |

CHAPTER 28**RESTCONF Agent 371**

| | |
|--|-----|
| About the RESTCONF Agent | 371 |
| Guidelines and Limitations | 372 |
| Configuring the RESTCONF Agent | 372 |
| Manage a RESTCONF Session | 373 |
| RESTCONF Get / Set | 373 |
| RESTCONF Device YANG RPC | 375 |
| About Operational Commands in RESTCONF | 375 |
| Device YANG RPC Examples | 375 |
| Troubleshooting the RESTCONF Agent | 377 |

CHAPTER 29**gRPC Agent 381**

| | |
|--|-----|
| About gRPC Agent | 381 |
| Revision History | 381 |
| Guidelines and Limitations for gRPC Agent | 381 |
| Configuring the gRPC Agent | 382 |
| Configuring gRPC | 382 |
| Configuring gRPC Client Certificate Authentication | 385 |
| Configuring NGINX proxy for GRPC | 388 |
| Troubleshooting | 389 |
| Check Feature Status | 389 |
| Check Connectivity | 390 |
| Gathering gRPC Agent Logs | 390 |
| Gathering TM-Trace Logs | 390 |
| Gathering MTX-Internal Logs | 390 |

CHAPTER 30**gNMI - Management Interface 393**

| | |
|--|-----|
| About gNMI | 393 |
| Guidelines and Limitations for gNMI | 394 |
| Configuring gNMI | 397 |
| Configuring gNMI Options | 397 |
| gNMI RPCs | 398 |
| About Get | 400 |
| About Set | 401 |
| About Subscribe | 402 |
| About Subscribing Custom Syslog Stream | 405 |
| References | 409 |
| Troubleshooting gNMI | 409 |
| Configuration Examples for gRPC Commands | 412 |
| Gathering Debug Logs | 414 |
| Accounting Log for gNMI | 414 |

CHAPTER 31

| | |
|-------------------------------------|------------|
| gNOI - Operation Interface | 417 |
| About gNOI | 417 |
| Revision History | 418 |
| Guidelines and Limitations for gNOI | 418 |
| Configuring gNOI | 418 |
| System .Proto | 419 |
| OS .Proto | 420 |
| Cert .Proto | 420 |
| File .Proto | 421 |
| Factory Reset .Proto | 421 |
| FactoryReset | 422 |
| Troubleshooting gNOI | 422 |
| Debug gNOI | 422 |
| Show Commands | 422 |
| Example Output | 423 |
| Gathering Debug Logs | 423 |

CHAPTER 32

| | |
|--------------------|------------|
| gRPC Tunnel | 425 |
| About gRPC Tunnel | 425 |

| | |
|---|----------------------|
| Guidelines and Limitations | 425 |
| Topic 2.1 | 426 |
| Configuring gRPC Tunnel | 426 |
| Configuring gRPC Tunnel without authentication | 426 |
| Configuring gRPC Tunnel with Trustpoints | 427 |
| Configuring gRPC Tunnel with VRF | 429 |
| Troubleshooting | 433 |
| <hr/> | |
| CHAPTER 33 | Telemetry 435 |
| About Telemetry | 435 |
| Telemetry Components and Terminology | 435 |
| High Availability of the Telemetry Process | 436 |
| Licensing Requirements for Telemetry | 436 |
| Guidelines and Limitations for Telemetry | 437 |
| Configuring Telemetry Using the CLI | 443 |
| Configuring with CLI | 443 |
| Configuring Cadence for YANG Paths | 447 |
| Configuration Examples for Telemetry Using the CLI | 449 |
| Displaying Telemetry Configuration and Statistics | 453 |
| Configuring Telemetry Using the NX-API | 462 |
| Telemetry Model in the DME | 462 |
| Configuring with NX-API | 463 |
| Configuration Example for Telemetry Using the NX-API | 472 |
| Multicast Flow Path Visibility | 475 |
| Cloud Scale Software Telemetry | 477 |
| About Cloud Scale Software Telemetry | 477 |
| Cloud Scale Software Telemetry Message Formats | 477 |
| Guidelines and Limitations for Cloud Scale Software Telemetry | 477 |
| Telemetry Path Labels | 478 |
| About Telemetry Path Labels | 478 |
| Polling for Data or Receiving Events | 478 |
| Guidelines and Limitations for Path Labels | 479 |
| Configuring the Interface Path to Poll for Data or Events | 479 |
| Configuring the Interface Path for Non-Zero Counters | 481 |

| | |
|---|-----|
| Configuring the Interface Path for Operational Speeds | 483 |
| Configuring the Interface Path with Multiple Queries | 485 |
| Configuring the Environment Path to Poll for Data or Events | 486 |
| Enabling Power Usage Tracking Functionality | 488 |
| Displaying Power Consumption History | 488 |
| Configuring the Resources Path to Poll for Events or Data | 490 |
| Configuring the VXLAN Path to Poll for Events or Data | 492 |
| Verifying the Path Label Configuration | 494 |
| Displaying Path Label Information | 495 |
| Native Data Source Paths | 497 |
| About Native Data Source Paths | 497 |
| Telemetry Data Streamed for Native Data Source Paths | 497 |
| Guidelines and Limitations | 499 |
| Configuring the Native Data Source Path for Routing Information | 499 |
| Configuring the Native Data Source Path for MAC Information | 501 |
| Configuring the Native Data Source Path for all MAC Information | 503 |
| Configuring the Native Data Path for IP Adjacencies | 505 |
| Displaying Native Data Source Path Information | 507 |
| Streaming Syslog | 508 |
| About Streaming Syslog for Telemetry | 508 |
| Configuring the YANG Data Source Path for Syslog Information | 508 |
| Telemetry Data Streamed for Syslog Path | 510 |
| Troubleshooting Telemetry | 513 |
| Displaying Telemetry Log and Trace Information | 513 |
| Additional References | 514 |

CHAPTER 34

| | |
|-------------------------------------|------------|
| Diagnosis and Serviceability | 515 |
| About Diagnosis and Serviceability | 515 |
| Show Commands | 515 |
| Debug Logs | 516 |
| Programmability Agent Logs | 516 |
| YANG Infra Logs | 517 |
| DME Logs | 517 |
| Change the Log Configuration | 517 |

| | |
|--|-----|
| Default Config Example | 521 |
| Change the Log Configuration Using CLI | 523 |
| Diagnosis Suggestions | 524 |

PART VI

XML Management Interface 527

CHAPTER 35

XML Management Interface 529

| | |
|--|-----|
| About the XML Management Interface | 529 |
| Information About the XML Management Interface | 529 |
| NETCONF Layers | 529 |
| SSH xmlagent | 530 |
| Licensing Requirements for the XML Management Interface | 530 |
| Prerequisites to Using the XML Management Interface | 530 |
| Using the XML Management Interface | 531 |
| Configuring the SSH and the XML Server Options Through the CLI | 531 |
| Starting an SSHv2 Session | 532 |
| Sending a Hello Message | 533 |
| Obtaining XML Schema Definition (XSD) Files | 533 |
| Sending an XML Document to the XML Server | 534 |
| Creating NETCONF XML Instances | 534 |
| RPC Request Tag | 535 |
| NETCONF Operations Tags | 536 |
| Device Tags | 537 |
| Extended NETCONF Operations | 539 |
| NETCONF Replies | 543 |
| RPC Response Tag | 543 |
| Interpreting the Tags Encapsulated in the data Tag | 543 |
| Information About Example XML Instances | 544 |
| Example XML Instances | 544 |
| NETCONF Close Session Instance | 545 |
| NETCONF Kill Session Instance | 545 |
| NETCONF Copy Config Instance | 545 |
| NETCONF Edit Config Instance | 546 |
| NETCONF Get Config Instance | 548 |

| | |
|---|-----|
| NETCONF Lock Instance | 548 |
| NETCONF Unlock Instance | 549 |
| NETCONF Commit Instance: Candidate Configuration Capability | 549 |
| NETCONF Confirmed Commit Instance | 550 |
| NETCONF Rollback-On-Error Instance | 550 |
| NETCONF Validate Capability Instance | 551 |
| Additional References | 551 |

| | | |
|-------------------|------------------------------------|------------|
| APPENDIX A | Streaming Telemetry Sources | 553 |
| | About Streaming Telemetry | 553 |
| | Guidelines and Limitations | 553 |
| | Data Available for Telemetry | 553 |

| | | |
|-------------------|-------------------------------|------------|
| APPENDIX B | Websocket Subscription | 555 |
| | WebSocket Subscription | 555 |

| | | |
|-------------------|-----------------------------|------------|
| APPENDIX C | Programmability RFCs | 557 |
| | Programmability RFCs | 557 |



CHAPTER 1

New and Changed Information

- [New and Changed Information](#), on page 1

New and Changed Information

Table 1: New and Changed Features

| Feature | Description | Changed in Release | Where Documented |
|--|--|--------------------|---|
| Sample-based subscription and updates-only support for RIB native path | Added new query conditions to support sample-based subscription and updates-only support for RIB native path | 10.4(3)F | Guidelines and Limitations , on page 499 Configuring the Native Data Source Path for Routing Information , on page 499 |
| Event notifications on all properties - Infra | On-change gnmi subscription support for <code>openconfig/system/processes</code> | 10.4(3)F | About gNMI , on page 393 Guidelines and Limitations for gNMI , on page 394 |
| RESTCONF RFC 8040 | Added support for RESTCONF RFC 8040. | 10.4(3)F | Guidelines and Limitations , on page 372 |
| NETCONF With-Defaults, including explicit capabilities | Added support for NETCONF explicit mode. | 10.4(3)F | About NETCONF Explicit Mode , on page 362 |

| Feature | Description | Changed in Release | Where Documented |
|---------------------------|---|--------------------|--|
| Thurn Reenable DME | Added Support for DME in 92348GC-X. | 10.4(3)F | Guidelines and Limitations for NETCONF, on page 334 Guidelines and Limitations, on page 306 Guidelines and Limitations, on page 372 Guidelines and Limitations for gRPC Agent , on page 381 |
| Telemetry | Added support for telemetry on Cisco Nexus N9K-C9364C-H1 platform switches. | 10.4(3)F | Guidelines and Limitations for Telemetry, on page 437 |
| FPM Parallel lookup | Added support for FPM Parallel lookup on Cisco Nexus N9K-C9364C-H1 platform switches. | 10.4(3)F | FPM Parallel Lookup, on page 14 |
| DME support | Added support for DME on Cisco Nexus N9KC9364C-H1 platform switches. | 10.4(3)F | Support for DME Modularity, on page 28 |
| FPM Parallel lookup | Added support for FPM Parallel lookup on Cisco Nexus 93400LD-H1 platform switches. | 10.4(2)F | FPM Parallel Lookup, on page 14 |
| DME support | Added support for DME on Cisco Nexus 93400LD-H1 platform switches. | 10.4(2)F | Support for DME Modularity, on page 28 |
| Telemetry | Added support for telemetry on Cisco Nexus 93400LD-H1 platform switches. | 10.4(2)F | Guidelines and Limitations for Telemetry, on page 437 |
| Power Consumption History | New functionality is introduced to track power consumption in Nexus 9000 Series switches. | 10.4(1)F | Enabling Power Usage Tracking Functionality, on page 488 Displaying Power Consumption History, on page 488 |

| Feature | Description | Changed in Release | Where Documented |
|--------------------|--|--------------------|--|
| Telemetry | Added support for telemetry on Cisco Nexus 9332D-H2R platform switches. | 10.4(1)F | Guidelines and Limitations for Telemetry, on page 437 |
| DME support | Added support for DME on Cisco Nexus 9332D-H2R platform switches. | 10.4(1)F | Support for DME Modularity, on page 28 |
| Software Telemetry | Added support for software telemetry on the following switches and line cards: <ul style="list-style-type: none"> • Cisco Nexus 9804 platform switches • Cisco Nexus X98900CD-A and X9836DM-A line cards with Cisco Nexus 9808 and 9804 switches | 10.4(1)F | Guidelines and Limitations for Cloud Scale Software Telemetry, on page 477 |
| NETCONF | Added support for NETCONF on the following switches and line cards: <ul style="list-style-type: none"> • Cisco Nexus 9804 platform switches • Cisco Nexus X98900CD-A and X9836DM-A line cards with Cisco Nexus 9808 and 9804 switches | 10.4(1)F | Guidelines and Limitations for NETCONF, on page 334 |

| Feature | Description | Changed in Release | Where Documented |
|-----------|--|--------------------|--|
| DME Infra | Added support for DME infra on the following switches and line cards: <ul style="list-style-type: none"><li data-bbox="672 407 891 462">• Cisco Nexus 9804 platform switches<li data-bbox="672 491 891 672">• Cisco Nexus X98900CD-A and X9836DM-A line cards with Cisco Nexus 9808 and 9804 switches | 10.4(1)F | Support for DME Modularity, on page 28 |



CHAPTER 2

Platform Support for Programmability Features

This chapter defines platform support for features that are not supported across the entire suite of Cisco Nexus platforms.

- [Platform Support for Programmability Features, on page 5](#)

Platform Support for Programmability Features

The following tables list the supported platforms for each feature and the release in which they were first introduced. See the Release Notes for details about the platforms supported in the initial product release.

Barefoot Shell

Return to [Barefoot Shell](#).

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|----------------|--|-------------------------|---------------------|
| Barefoot Shell | Cisco Nexus 3464C switch Cisco Nexus 34180YC switch | Cisco NX-OS 9.2(3) | |

Bash Shell

Return to [About Bash, on page 19](#).

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|-------------------------------|-----------------------------------|-------------------------|----------------------------------|
| Consent token for bash access | Cisco Nexus 9000 Series switches | Cisco NX-OS 10.3(2)F | Cisco Nexus 9808 platform switch |
| DME Modularity | Cisco Nexus 9800 Series switches | Cisco NX-OS 10.3(1)F | |
| DME Modularity | Cisco Nexus 9000 Series switches | Cisco NX-OS 9.3(1) | |

Chef Client

Return to [Using Chef Client with Cisco NX-OS, on page 151](#).

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|------------|---|-------------------------|---------------------|
| Chef Agent | Cisco Nexus 9300 platform switches Cisco Nexus 9500 platform switches and line cards | Cisco NX-OS 7.0(3)I2(5) | N9K-C92348GC |

Docker

Return to [Using Docker with Cisco NX-OS, on page 173](#).

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|---------------------------------------|---|-------------------------|---------------------|
| Docker Container support within NX-OS | Cisco Nexus 9000 Series switches with at least 8 GB of system RAM | Cisco NX-OS 9.2(1) | N9K-C92348GC |

gNMI - gRPC Network Management Interface

Return to [gNMI-gRPC Network Management Interface](#).

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|--------------|-----------------------------------|-------------------------|---------------------|
| gNMI Get/Set | Cisco Nexus 9000 Series switches | Cisco NX-OS 9.3(1) | N9K-C92348GC |

Hardware Telemetry

Return to [Hardware Telemetry](#).

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|----------------------|--|-------------------------|---------------------|
| SSX Counters | Cisco Nexus 9332C and Cisco Nexus 9364C switches | 9.3(3) | |
| Flow Table Analytics | Cisco Nexus 9300-FX and Cisco Nexus 9300-FX2 platform switches Cisco Nexus 9500 platform switches with 9700 -EX/FX line cards | 9.3(1) | |

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|--------------------------------|--|----------------------------|---------------------|
| Flow Table Events (FTE) | Cisco Nexus 9300-FX and Cisco Nexus 9300-FX2 platform switches | Cisco NX-OS 9.2(1) | |
| Flow monitor for VRF filtering | Cisco Nexus 9300-FX/FX2 platform switches and Cisco Nexus 9500 platform switches with 9700-FX line cards | Cisco NX-OS Release 9.3(3) | |

Inband Telemetry

Return to *Inband Network Telemetry*.

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|--|--|-------------------------|---------------------|
| Packet Postcards, support for flow, packet drop, and queue congestion events | Cisco Nexus 3464C switch Cisco Nexus 34180YC switch | Cisco NX-OS 9.2(3) | |
| Packet Postcards, support for the INT source and sink roles | Cisco Nexus 3464C switch | Cisco NX-OS 9.2(3) | |

Model-Driven Telemetry

Return to *Model Driven Telemetry*.

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|-------------------------------|--|-------------------------|---------------------|
| Alias Option for Sensor Path | Cisco Nexus 9200, 9300-EX, 9300-FX/FX2/FXP platform switches Cisco Nexus 9500 platform switches with EX/FX line cards | Cisco NX-OS 9.3(5) | N9K-C92348GC |
| Software Telemetry (dial-out) | Cisco Nexus 9000 platform switches | Cisco NX-OS 7.0(3)I5(1) | N9K-92348GC |
| Software Telemetry | Cisco Nexus 9000 platform switches | Cisco NX-OS 10.3(1)F | |

NETCONF Agent

Return to [NETCONF Agent](#).

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|-----------------|------------------------------------|-------------------------|---------------------|
| NETCONF Support | Cisco Nexus 9800 platform switches | Cisco NX-OS 10.3(1)F | |
| NETCONF Support | Cisco Nexus 9000 platform switches | | N9K-C92348GC |

NX-API REST

Return to [NX-API REST](#), on page 263.

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|--------------------|------------------------------------|-------------------------|---------------------|
| DME Config Replace | Cisco Nexus 9000 platform switches | Cisco NX-OS 9.3(1) | N9K-C92348GC |

Python API

Return to [Python API](#), on page 97.

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|----------|-----------------------------------|-------------------------|---------------------|
| Python 3 | Cisco Nexus 9000 Series switches | Cisco NX-OS 9.3(5) | |

Puppet Agent

Return to [Puppet Agent](#), on page 147.

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|--------------|---|-------------------------|---------------------|
| Puppet Agent | Cisco Nexus 9300 and 9500 platform switches | Cisco NX-OS 7.0(3)I2(5) | N9K-C92348GC |

TCL Scripting

Return to [Scripting with Tcl](#), on page 109.

| Feature | Supported Platforms or Line Cards | First Supported Release | Platform Exceptions |
|-----------|-----------------------------------|-------------------------|---------------------|
| TCL Shell | Cisco Nexus 9000 Series switches | - | |



CHAPTER 3

Overview

- [Programmability Overview, on page 9](#)
- [Supported Platforms, on page 10](#)
- [Standard Network Manageability Features, on page 10](#)
- [Advanced Automation Features, on page 10](#)
- [Programmability Support, on page 10](#)

Programmability Overview

The Cisco NX-OS software running on the Cisco Nexus 9000 Series switches is as follows:

- **Resilient**
Provides critical business-class availability.
- **Modular**
Has extensions that accommodate business needs.
- **Highly Programmatic**
Allows for rapid automation and orchestration through Application Programming Interfaces (APIs).
- **Secure**
Protects and preserves data and operations.
- **Flexible**
Integrates and enables new technologies.
- **Scalable**
Accommodates and grows with the business and its requirements.
- **Easy to use**
Reduces the amount of learning required, simplifies deployment, and provides ease of manageability.

With the Cisco NX-OS operating system, the device functions in the unified fabric mode to provide network connectivity with programmatic automation functions.

Cisco NX-OS contains Open Source Software (OSS) and commercial technologies that provide automation, orchestration, programmability, monitoring, and compliance support.

For more information on Open NX-OS, see <https://developer.cisco.com/site/nx-os/>.

Beginning with Cisco NX-OS Release 10.2.(3)F, the data export of Flow Table (FT) / Flow Table Events (FTE) in V9 format is supported on N9K-X9716D-GX line cards.

Beginning with Cisco NX-OS Release 10.3(1)F, FT/FTE supports MTU drop and TTL match events for multicast flow.

FTE latency will work only for the flows matching the filters.

Supported Platforms

Starting with Cisco NX-OS release 7.0(3)I7(1), use the [Nexus Switch Platform Support Matrix](#) to know from which Cisco NX-OS releases various Cisco Nexus 9000 and 3000 switches support a selected feature.

Standard Network Manageability Features

- SNMP (V1, V2, V3)
- Syslog
- RMON
- NETCONF
- CLI and CLI scripting

Advanced Automation Features

The enhanced Cisco NX-OS on the device supports automation. The platform includes support for Power On Auto Provisioning (POAP).

The enhanced Cisco NX-OS on the device supports automation. The platform includes the features that support automation.

Programmability Support

Cisco NX-OS software on switches support several capabilities to aid programmability.

NX-API Support

Cisco NX-API allows for HTTP-based programmatic access to the switches. This support is delivered by NX-API, an open source webserver. NX-API provides the configuration and management capabilities of the Cisco NX-OS CLI with web-based APIs. The device can be set to publish the output of the API calls in XML or JSON format. This API enables rapid development on the switches.

Python Scripting

Cisco NX-OS supports Python v2.7.5 in both interactive and noninteractive (script) modes.

Beginning in Cisco NX-OS Release 9.3(5), Python 3 is also supported.

The Python scripting capability on the devices provides programmatic access to the switch CLI to perform various tasks, and to Power-On Auto Provisioning (POAP) and Embedded Event Manager (EEM) actions. Responses to Python calls that invoke the Cisco NX-OS CLI return text or JSON output.

The Python interpreter is included in the Cisco NX-OS software.

Tcl Scripting

Cisco Nexus 9000 Series switches support Tcl (Tool Command Language). Tcl is a scripting language that enables greater flexibility with CLI commands on the switch. You can use Tcl to extract certain values in the output of a **show** command, perform switch configurations, run Cisco NX-OS commands in a loop, or define EEM policies in a script.

Broadcom Shell

The Cisco Nexus 9000 Series switch front panel and fabric module line cards contain Broadcom Network Forwarding Engine (NFE). You can access the Broadcom command-line shell (bcm-shell) from these NFEs.

Bash

Cisco Nexus switches support direct Bourne-Again Shell (Bash) access. With Bash, you can access the underlying Linux system on the device and manage the system.

Bash Shell Access and Linux Container Support

Cisco Nexus switches support direct Linux shell access and Linux containers. With Linux shell access, you can access the underlying Linux system on the switch and manage the underlying system. You can also use Linux containers to securely install your own software and to enhance the capabilities of the Cisco Nexus switch. For example, you can install bare-metal provisioning tools like Cobbler on a Cisco Nexus switch to enable automatic provisioning of bare-metal servers from the top-of-rack switch.

Guest Shell

| |
|--|
| In process-hosted deployments, a virtual service execution environment (VSEE) isolates the onePK application from core routing and switching applications that the NOS provides; the environment also isolates its contents from other virtual services on the same host. The Application Developer can allocate specified quantities of CPU time, memory, disk space, and other resources to any particular VSEE. |
|--|

Cisco utilizes the Cisco Secure Development Lifecycle (SDLC) to develop onePK applications that run in VSEEs on closed Cisco systems. The SDLC comprises a continuously evolving set of industry-recognized best practices and tools (including Cisco-proprietary tools) that reduce the vulnerability footprint of Cisco-provided applications, VSEEs, and platforms. These practices include the use of runtime checks that ensure the integrity of Cisco-signed binaries before loading them, and the establishment of a trust domain in which Cisco-provided onePK applications are allowed to run as trusted processes.

onePK applications in a VSEE communicate with the network element using onePK APIs that the onePK SDK provides. A secure communications channel carries messages between the VSEE and the onePK server. The use of the trust domain ensures the integrity of interprocess communications within the Cisco-provided VSEE, and additional security measures protect communications between the VSEE and the Cisco host.

The Network Administrator exercises direct control over the deployment of VSEEs and applications that interact with network elements. The duties of the Network Administrator include responsibility for verifying the integrity and validity of application packages. Cisco provides information (such as digital signatures or MD5 checksums) that enable Network Administrators to use standard software development tools to verify the integrity of application packages before deploying them. In addition, the default settings of the onePK VSEE security infrastructure allow only Cisco-signed application packages to run in process-hosted mode. To allow the deployment of unsigned containers or those containers that signed by third parties, the Network Administrator must take explicit action.

Because onePK allows low-level access to network elements, the Network Administrator must be very selective about user profiles and applications that are given access to onePK.

The Cisco Nexus 9000 Series switches support a virtual service environment that runs inside a secure Linux container (LXC). It isolates the application running in the guest shell of the virtual service environment from other routing and switching applications that the host Cisco NX-OS provides. The environment also isolates its contents from other virtual services on the same host. You can allocate specified quantities of CPU time, memory, disk space, and other resources to any particular virtual service environment.

Cisco utilizes the Cisco Secure Development Lifecycle (SDLC) to develop applications that run in virtual service environments on Cisco Nexus 9000 Series devices. The SDLC comprises a continuously evolving set of industry-recognized best practices and tools (including Cisco-proprietary tools) that reduce the vulnerability footprint of Cisco-provided applications, virtual service environments, and platforms. These practices include the use of runtime checks that ensure the integrity of Cisco-signed binaries before loading them, and the establishment of a trust domain in which Cisco-provided applications are allowed to run as trusted processes.

Applications in a guest shell communicate with the network using APIs. A secure communications channel carries messages between the guest shell and the device. The use of the trust domain ensures the integrity of interprocess communications within the Cisco-provided guest shell, and additional security measures protect communications between the guest shell and the device.

A network administrator controls the deployment of virtual service environments and applications that interact with the network. The duties of a network administrator include the responsibility for verifying the integrity and validity of application packages. Cisco provides information (such as digital signatures or MD5 checksums) that enable network administrators to use standard software development tools to verify the integrity of application packages before deploying them. In addition, the default settings of the guest shell security infrastructure allow only Cisco-signed application packages to run on the host. To allow the deployment of unsigned containers or those containers that are signed by third parties, the network administrator must take explicit action.

The Cisco Nexus 9000 Series switches support a guest shell that provides Bash access into a Linux execution space on the host system that is decoupled from the host Cisco Nexus 9000 NX-OS software. With the guest shell, you can add software packages and update libraries as needed without impacting the host system software.

Applications running in the guest shell have IP connectivity to the host system and the external network through socket APIs.

The guest shell is implemented as a secure Linux container (LXC) and is started automatically when the host system is started. This allows applications in the guest shell to be automatically started when the system is started. The amount of CPU, memory, and bootflash space used for the guest shell can be tuned to balance the resource usage between the guest shell and the host system. The guest shell mounts the system's bootflash to allow access to files on the bootflash using Linux commands.

Container Tracker Support

Cisco NX-OS is configured to communicate with the Kubernetes API Server to understand the capabilities of the containers behind a given switch port.

The following commands communicate with the Kubernetes API Server:

- The **show containers kubernetes** command obtains data from *kube-apiserver* using API calls over HTTP.
- The **kubernetes watch resource** command uses a daemon to subscribe to requested resources and process streaming data from *kube-apiserver*.
- The **action** assigned in the **watch** command is performed on pre-defined triggers. (For example, Add or Delete of a Pod.)

Perl Modules



Note Beginning with Cisco NX-OS Release 9.2(2), support for the Perl modules has been added for the Cisco Nexus 9504 and 9508 switches with -R line cards.

In order to support more applications, the following Perl modules have been added:

- bytes.pm
- feature.pm
- hostname.pl
- lib.pm
- overload.pm
- Carp.pm
- Class/Struct.pm
- Data/Dumper.pm
- DynaLoader.pm

- Exporter/Heavy.pm
- FileHandle.pm
- File/Basename.pm
- File/Glob.pm
- File/Spec.pm
- File/Spec/Unix.pm
- File/stat.pm
- Getopt/Std.pm
- IO.pm
- IO/File.pm
- IO/Handle.pm
- IO/Seekable.pm
- IO/Select.pm
- List/Util.pm
- MIME/Base64.pm
- SelectSaver.pm
- Socket.pm
- Symbol.pm
- Sys/Hostname.pm
- Time/HiRes.pm
- auto/Data/Dumper/Dumper.so
- auto/File/Glob/Glob.so
- auto/IO/IO.so
- auto/List/Util/Util.so
- auto/MIME/Base64/Base64.so
- auto/Socket/Socket.so
- auto/Sys/Hostname/Hostname.so
- auto/Time/HiRes/HiRes.so

FPM Parallel Lookup

Starting with Cisco NX-OS release 10.4(2)F, FPM parallel look is supported on Cisco Nexus 93400LD-H1 platform switches.

Starting with Cisco NX-OS release 10.4(2)F, FPM parallel look is supported on Cisco Nexus N9K-C9364CH1 platform switches.

Nexus 9332D-H2R platform switches support parallel lookup between two successive FPM instances on the die, reducing the overall cycles taken to lookup all 40 FPM instances. To enable parallel lookup between two FPM instances, there should not be any dependency on lookup key or index derivation between the two blocks. For example, parallel lookup cannot be enabled between two FPM instances having IP routes in one and adjacent entries in another, as the adjacency index comes from an IP route lookup. By default, parallel lookup feature is enabled on the switch. The feature cannot be disabled.



PART I

Shells and Scripting

- [Shells and Scripting, on page 19](#)
- [Guest Shell, on page 49](#)
- [Broadcom Shell, on page 87](#)
- [Barefoot Shell, on page 93](#)
- [Python API, on page 97](#)
- [Scripting with Tcl, on page 109](#)
- [iPXE, on page 115](#)
- [Kernel Stack, on page 119](#)



CHAPTER 4

Shells and Scripting

- [About Bash, on page 19](#)
- [Guidelines and Limitations, on page 19](#)
- [Enabling Consent Token, on page 20](#)
- [Accessing Bash, on page 21](#)
- [Escalate Privileges to Root, on page 22](#)
- [Examples of Bash Commands, on page 24](#)
- [Managing Feature RPMs, on page 25](#)
- [Support for DME Modularity, on page 28](#)
- [Managing Patch RPMs, on page 36](#)
- [Persistently Daemonizing an SDK- or ISO-built Third Party Process, on page 43](#)
- [Persistently Starting Your Application from the Native Bash Shell, on page 44](#)
- [Synchronize Files from Active Bootflash to Standby Bootflash, on page 44](#)
- [Copy Through Kstack, on page 46](#)
- [An Example Application in the Native Bash Shell, on page 46](#)

About Bash

In addition to the Cisco NX-OS CLI, Cisco Nexus 30009000 Series switches support access to the Bourne-Again Shell (Bash). Bash interprets commands that you enter or commands that are read from a shell script. Using Bash enables access to the underlying Linux system on the device and to manage the system.

Guidelines and Limitations

The Bash shell has the following guidelines and limitations:

- When you define a link-local address for an interface, Netstack installs a /64 prefix on the net device in the kernel.

When a new link-local address is configured on the kernel, the kernel installs a /64 route in the kernel routing table.

If the peer box's interface is not configured with a link-local address that falls in the same /64 subnet, the **ping** is not successful from the bash prompt. A Cisco NX-OS **ping** works fine.

- The binaries in the `/isan` folder are meant to be run in an environment which is set up differently from the environment of the shell that you enter by the **run bash** command. It is advisable not to use these binaries from the Bash shell as the behavior within this environment isn't predictable.
- When importing Cisco Python modules, don't use Python from the Bash shell. Instead use the more recent Python in NX-OS VSH.
- Some processes and **show** commands can cause a large amount of output. If you are running scripts, and need to terminate long-running output, use Ctrl+C (not Ctrl+Z) to terminate the command output. If you use Ctrl+Z, this key command can generate a SIGCONT (signal continuation) message, which can cause the script to halt. Scripts that are halted through SIGCONT messages require user intervention to resume operation.
- If the **show tech support** command is running and you must kill it, don't use the **clear tech-support lock** command. Use Ctrl+C.

The reason is that **clear tech-support lock** doesn't kill the background VSH session where the actual collection of tech-support information happens. Instead, **clear tech-support lock** command kills only the foreground VSH session where the **show tech support** CLI is called.

To correctly kill the show tech-support session, use Ctrl+C.

If you accidentally used **clear tech-support lock**, perform the following steps to kill the background VSH process:

1. Enter the Bash shell.
 2. Locate the VSH session (**ps -l | more**) for the **show tech support** command.
 3. Kill the PID associated with the VSH for the **show tech support** session, for example, **kill -9 PID**.
- Beginning with Cisco NX-OS Release 10.3(2)F, the consent token for bash access feature provides support for consent token to enable shell access on NX-OS. However, this feature works only in Trust Anchor Module (TAM) based devices. This feature is supported on all Cisco Nexus 9000 Series platform switches except Cisco Nexus 9808 platform switches. The following limitations are applicable:
 - Write-Erase reload is required to disable this configuration.
 - ISSD is supported only on a release that has the consent token feature.
 - When this configuration is enabled, NX-API/Netconf/Restconf do not work and an error indicating the reason appears.
 - Config-replace is allowed only if the command is present in the new configuration or file.
 - Changing boot-variable, copying running-config startup-config, and reloading the device is not recommended when consent token is enabled.

Enabling Consent Token

To enable the consent token that restricts the Bash access, perform the following command.

```
system security consent-token shell-access [ <timeout> ] [force]
```

If you do not provide any value to the timeout parameter, then the default value of 5 minutes is considered. The maximum value for the timeout parameter is 2880 minutes, that is, 2 days.

After the command enables the feature, the device will be in the consent-token secure mode, and access to the shell is granted to the user for the time-period specified in the `<timeout>` parameter of this command.

By default, the command is interactive. Use the **force** keyword to enable the consent-token mode forcefully (noninteractive).



Note The **no** form of this command cannot be used to disable the command (for security reason). Hence, to disable this command, perform write-erase-reload on the device.

To verify the status of the consent token feature, use the **show system security consent-token** command.

Accessing Bash

In Cisco NX-OS, Bash is accessible from user accounts that are associated with the Cisco NX-OS dev-ops role or the Cisco NX-OS network-admin role.

The following example shows the authority of the dev-ops role and the network-admin role:

```
switch# show role name dev-ops

Role: dev-ops
Description: Predefined system role for devops access. This role
cannot be modified.
Vlan policy: permit (default)
Interface policy: permit (default)
Vrf policy: permit (default)
-----
Rule    Perm   Type   Scope                Entity
-----
4       permit command            conf t ; username *
3       permit command            bcm module *
2       permit command            run bash *
1       permit command            python *

switch# show role name network-admin

Role: network-admin
Description: Predefined network admin role has access to all commands
on the switch
-----
Rule    Perm   Type   Scope                Entity
-----
1       permit read-write

switch#
```

Bash is enabled by running the **feature bash-shell** command.

The **run bash** command loads Bash and begins at the home directory for the user.

The following examples show how to enable the Bash shell feature and how to run Bash.

```
switch# configure terminal
switch(config)# feature bash-shell

switch# run?
run          Execute/run program
run-script   Run shell scripts
```

```
switch# run bash?
bash Linux-bash

switch# run bash
bash-4.2$ whoami
admin
bash-4.2$ pwd
/bootflash/home/admin
bash-4.2$
```



Note You can also execute Bash commands with **run bash** *command*.

For instance, you can run **whoami** using **run bash** *command*:

```
run bash whoami
```

You can also run Bash by configuring the user **shelltype**:

```
username foo shelltype bash
```

This command puts you directly into the Bash shell upon login. This does not require **feature bash-shell** to be enabled.

Escalate Privileges to Root

The privileges of an admin user can escalate their privileges for root access.

The following are guidelines for escalating privileges:

- admin privilege user (network-admin / vdc-admin) is equivalent of Linux root privilege user in NX-OS
- Only an authenticated admin user can escalate privileges to root, and password is not required for an authenticated admin privilege user *
- Bash must be enabled before escalating privileges.
- SSH to the switch using `root` username through a non-management interface will default to Linux Bash shell-type access for the root user. Type **vsh** to return to NX-OS shell access.

* From Cisco NX-OS Release 9.2(3) onward, if password prompting is required for some use case even for admin (user with role network-admin) privilege user, enter the **system security hardening sudo prompt-password** command.

NX-OS network administrator users must escalate to root to pass configuration commands to the NX-OS VSH if:

- The NX-OS user has a shell-type Bash and logs into the switch with a shell-type Bash.
- The NX-OS user that logged into the switch in Bash continues to use Bash on the switch.

Run **sudo su 'vsh -c "<configuration commands>"** or **sudo bash -c 'vsh -c "<configuration commands>"**.

The following example demonstrates with network administrator user MyUser with a default shell type Bash using **sudo** to pass configuration commands to the NX-OS:

```
ssh -l MyUser 1.2.3.4
-bash-4.2$ sudo vsh -c "configure terminal ; interface eth1/2 ; shutdown ; sleep 2 ; show
interface eth1/2 brief"
```

```
-----
Ethernet      VLAN      Type Mode      Status Reason      Speed      Port
Interface
-----
Eth1/2        --        eth  routed down  Administratively down  auto(D)  --
```

The following example demonstrates with network administrator user MyUser with default shell type Bash entering the NX-OS and then running Bash on the NX-OS:

```
ssh -l MyUser 1.2.3.4
-bash-4.2$ vsh -h
Cisco NX-OS Software
Copyright (c) 2002-2016, Cisco Systems, Inc. All rights reserved.
Nexus 9000v software ("Nexus 9000v Software") and related documentation,
files or other reference materials ("Documentation") are
the proprietary property and confidential information of Cisco
Systems, Inc. ("Cisco") and are protected, without limitation,
pursuant to United States and International copyright and trademark
laws in the applicable jurisdiction which provide civil and criminal
penalties for copying or distribution without Cisco's authorization.

Any use or disclosure, in whole or in part, of the Nexus 9000v Software
or Documentation to any third party for any purposes is expressly
prohibited except as otherwise authorized by Cisco in writing.
The copyrights to certain works contained herein are owned by other
third parties and are used and distributed under license. Some parts
of this software may be covered under the GNU Public License or the
GNU Lesser General Public License. A copy of each such license is
available at
http://www.gnu.org/licenses/gpl.html and
http://www.gnu.org/licenses/lgpl.html
*****
* Nexus 9000v is strictly limited to use for evaluation, demonstration *
* and NX-OS education. Any use or disclosure, in whole or in part of *
* the Nexus 9000v Software or Documentation to any third party for any *
* purposes is expressly prohibited except as otherwise authorized by *
* Cisco in writing. *
*****
switch# run bash
bash-4.2$ vsh -c "configure terminal ; interface eth1/2 ; shutdown ; sleep 2 ; show interface
eth1/2 brief"
```

```
-----
Ethernet      VLAN      Type Mode      Status Reason      Speed      Port
Interface
-----
Eth1/2        --        eth  routed down  Administratively down  auto(D)  --
```



Note Do not use **sudo su -** or the system hangs.

The following example shows how to escalate privileges to root and how to verify the escalation:

```
switch# run bash
bash-4.2$ sudo su root
bash-4.2# whoami
root
```

```
bash-4.2# exit
exit
```

Examples of Bash Commands

This section contains examples of Bash commands and output.

Displaying System Statistics

The following example displays system statistics:

```
switch# run bash
bash-4.2$ cat /proc/meminfo
<snip>
MemTotal:      16402560 kB
MemFree:       14098136 kB
Buffers:       11492 kB
Cached:        1287880 kB
SwapCached:    0 kB
Active:        1109448 kB
Inactive:      717036 kB
Active(anon):  817856 kB
Inactive(anon): 702880 kB
Active(file):  291592 kB
Inactive(file): 14156 kB
Unevictable:   0 kB
Mlocked:      0 kB
SwapTotal:     0 kB
SwapFree:      0 kB
Dirty:         32 kB
Writeback:     0 kB
AnonPages:    527088 kB
Mapped:       97832 kB
<\snip>
```

Running Bash from CLI

The following example runs `ps` from Bash using `run bash` command:

```
switch# run bash ps -el
F S  UID  PID  PPID  C  PRI  NI ADDR  SZ  WCHAN  TTY          TIME CMD
4 S   0    1    0  0  80   0 -   528 poll_s ?          00:00:03 init
1 S   0    2    0  0  80   0 -    0 kthrea ?          00:00:00 kthreadd
1 S   0    3    2  0  80   0 -    0 run_ks ?          00:00:56 ksoftirqd/0
1 S   0    6    2  0 -40  - -    0 cpu_st ?          00:00:00 migration/0
1 S   0    7    2  0 -40  - -    0 watchd ?          00:00:00 watchdog/0
1 S   0    8    2  0 -40  - -    0 cpu_st ?          00:00:00 migration/1
1 S   0    9    2  0  80   0 -    0 worker ?          00:00:00 kworker/1:0
1 S   0   10    2  0  80   0 -    0 run_ks ?          00:00:00 ksoftirqd/1
```

Managing Feature RPMs

RPM Installation Prerequisites

Use these procedures to verify that the system is ready before installing or adding an RPM.

SUMMARY STEPS

1. switch# `show logging logfile | grep -i "System ready"`
2. switch# `run bash sudo su`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | switch# <code>show logging logfile grep -i "System ready"</code> | Before running Bash, this step verifies that the system is ready before installing or adding an RPM. Proceed if you see output similar to the following: 2018 Mar 27 17:24:22 switch %ASCII-CFG-2-CONF_CONTROL: System ready |
| Step 2 | switch# <code>run bash sudo su</code> Example: switch# <code>run bash sudo su</code> bash-4.2# | Loads Bash. |

Installing Feature RPMs from Bash

Procedure

| | Command or Action | Purpose |
|--------|---|--|
| Step 1 | sudo <code>dnf installed grep platform</code> | Displays a list of the NX-OS feature RPMs installed on the switch. |
| Step 2 | <code>dnf list available</code> | Displays a list of the available RPMs. |
| Step 3 | sudo <code>dnf -y install rpm</code> | Installs an available RPM. |

Example

The following is an example of installing the **bfd** RPM:

```

bash-4.2$ dnf list installed | grep n9000
base-files.n9000                3.0.14-r74.2                installed
bfd.lib32_n9000                1.0.0-r0                    installed
core.lib32_n9000               1.0.0-r0                    installed
eigrp.lib32_n9000              1.0.0-r0                    installed
eth.lib32_n9000                1.0.0-r0                    installed
isis.lib32_n9000               1.0.0-r0                    installed
lACP.lib32_n9000               1.0.0-r0                    installed
linecard.lib32_n9000           1.0.0-r0                    installed
lldp.lib32_n9000               1.0.0-r0                    installed
ntp.lib32_n9000                1.0.0-r0                    installed
nxos-ssh.lib32_n9000           1.0.0-r0                    installed
ospf.lib32_n9000               1.0.0-r0                    installed
perf-cisco.n9000_gdb           3.12-r0                     installed
platform.lib32_n9000           1.0.0-r0                    installed
shadow-securetty.n9000_gdb     4.1.4.3-r1                  installed
snmp.lib32_n9000               1.0.0-r0                    installed
svi.lib32_n9000                1.0.0-r0                    installed
sysvinit-inittab.n9000_gdb     2.88dsf-r14                 installed
tacacs.lib32_n9000             1.0.0-r0                    installed
task-nxos-base.n9000_gdb       1.0-r0                       installed
tor.lib32_n9000                1.0.0-r0                    installed
vtp.lib32_n9000                1.0.0-r0                    installed
bash-4.2$ dnf list available
bgp.lib32_n9000                 1.0.0-r0
bash-4.2$ sudo dnf -y install bfd

```



Note Upon switch reload during boot up, use the **rpm** command instead of **dnf** for persistent RPMs. Otherwise, RPMs initially installed using **dnf bash** or **install cli** shows `reponame` or `filename` instead of `installed`.

Upgrading Feature RPMs

Before you begin

There must be a higher version of the RPM in the dnf repository.

SUMMARY STEPS

1. `sudo dnf -y upgraderpm`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|-------------------------------------|----------------------------|
| Step 1 | <code>sudo dnf -y upgraderpm</code> | Upgrades an installed RPM. |

Example

The following is an example of upgrading the **bfd** RPM:

```
bash-4.2$ sudo dnf -y upgrade bfd
```

Downgrading a Feature RPM

SUMMARY STEPS

1. `sudo dnf -y downgraderpm`

DETAILED STEPS**Procedure**

| | Command or Action | Purpose |
|--------|---------------------------------------|---|
| Step 1 | <code>sudo dnf -y downgraderpm</code> | Downgrades the RPM if any of the dnf repositories has a lower version of the RPM. |

Example

The following example shows how to downgrade the **bfd** RPM:

```
bash-4.2$ sudo dnf -y downgrade bfd
```

Erasing a Feature RPM



Note The SNMP RPM and the NTP RPM are protected and cannot be erased.

You can upgrade or downgrade these RPMs. It requires a system reload for the upgrade or downgrade to take effect.

For the list of protected RPMs, see `/etc/dnf/protected.d/protected_pkgs.conf`.

SUMMARY STEPS

1. `sudo dnf -y eraserpm`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|-----------------------------------|-----------------|
| Step 1 | <code>sudo dnf -y eraserpm</code> | Erases the RPM. |

Example

The following example shows how to erase the **bfd** RPM:

```
bash-4.2$ sudo dnf -y erase bfd
```

Support for DME Modularity

Beginning with NX-OS release 9.3(1), the Cisco NX-OS image supports DME modularity, which interoperates with the switch's RPM manager to enable non-intrusive upgrade or downgrade of DME RPMs. Non-intrusive upgrade or downgrade enables installing RPMs without performing a system restart and prevents disturbing other applications that have their configs in the DME database. DME Modularity enables you to apply model changes to the switch without an ISSU or system reload.



Note After loading the DME RPM, you must restart VSH to enable querying the new MOs.

Beginning with Cisco NX-OS Release 10.3(1)F, DME Infra is supported on the Cisco Nexus 9808 platform switches.

Beginning with Cisco NX-OS Release 10.4(1)F, DME Infra is supported on the Cisco Nexus 9804 platform switches.

Beginning with Cisco NX-OS Release 10.4(1)F, DME is supported on the Cisco Nexus 9332D-H2R platform switches.

Beginning with Cisco NX-OS Release 10.4(1)F, DME Infra is supported on N9KX98900CD-A and N9KX9836DM-A line cards with Cisco Nexus 9808 and 9804 switches.

Beginning with Cisco NX-OS Release 10.4(2)F, DME is supported on the Cisco Nexus 93400LD-H1 platform switches.

Beginning with Cisco NX-OS Release 10.4(2)F, DME is supported on the Cisco Nexus N9K-C9364C-H1 platform switches.

Installing the DME RPMs

By default, the base DME RPM, which is a mandatory upgradeable RPM package, is installed and active when you upgrade to NX-OS release 9.3(1). The DME RPM is installed in the default install directory for RPM files, which is `/rpms`.

If you make code or model changes, you will need to install the DME RPM. To install it, use either the NX-OS RPM manager, which uses the **install** command, or standard RPM tools, such as **dnf**. If you use **dnf**, you will need access to the switch's Bash shell.

Procedure

Step 1 copy *path-to-dme-rpm* bootflash: [//sup-#][/path]

Example:

```
switch-1# copy scp://test@10.1.1.1/dme-2.0.1.0-9.3.1.lib32_n9000.rpm bootflash://
switch-1#
```

Copies the DME RPM to bootflash through SCP.

Step 2 Choose any of the following methods to install or upgrade the DME RPM.

To use the NX-OS **install** command:

- **install add** *path-to-dme-rpm* **activate**

Example:

```
switch-1#install add dme-2.0.1.0-9.3.1.lib32_n9000.rpm activate
Adding the patch (/dme-2.0.1.0-9.3.1.lib32_n9000.rpm)
[#####] 100%
Install operation 90 completed successfully at Fri Jun 7 07:51:58 2019

Activating the patch (/dme-2.0.1.0-9.3.1.lib32_n9000.rpm)
[#####] 100%
Install operation 91 completed successfully at Fri Jun 7 07:52:35 2019
switch-1#
```

- **install add** *path-to-dme-rpm* **activate upgrade**

Example:

```
switch-1#install add dme-2.0.1.0-9.3.1.lib32_n9000.rpm activate upgrade
Adding the patch (/dme-2.0.1.0-9.3.1.lib32_n9000.rpm)
[#####] 100%
Install operation 87 completed successfully at Fri Jun 7 07:18:55 2019

Activating the patch (/dme-2.0.1.0-9.3.1.lib32_n9000.rpm)
[#####] 100%
Install operation 88 completed successfully at Fri Jun 7 07:19:35 2019
switch-1#
```

- **install add** *path-to-dme-rpm* then **install activate** *path-to-dme-rpm*

Example:

```
switch-1#install add bootflash:dme-2.0.1.0-9.3.1.lib32_n9000.rpm
[#####] 100%
Install operation 92 completed successfully at Fri Jun 7 09:31:04 2019
switch-1#install activate dme-2.0.1.0-9.3.1.lib32_n9000.rpm
[#####] 100%
Install operation 93 completed successfully at Fri Jun 7 09:31:55 2019
switch-1#
```

To use **dnf install**:

- **dnf install --add** *path-to-dme-rpm*

```
switch-1# dnf install --add bootflash:///dme-2.0.10.0-9.3.1.lib32_n9000.rpm
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
[#####] 90%Install operation 96 completed successfully at Fri Jun  7 22:58:50
2019.

[#####] 100%
switch-1#
```

- **dnf install --no-persist --nocommit path-to-dme-rpm**

This option requires user intervention, as shown below.

Example:

```
switch-1# dnf install --no-persist --nocommit dme-2.0.10.0-9.3.1.lib32_n9000
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages

groups-repo                | 1.1 kB      00:00 ...
localdb                    | 951 B       00:00 ...
localdb/primary            | 6.2 kB      00:00 ...
localdb                    |              2/2
patching                   | 951 B       00:00 ...
thirdparty                 | 951 B       00:00 ...
wrl-repo                   | 951 B       00:00 ...
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package dme.lib32_n9000 0:2.0.10.0-9.3.1 will be updated
---> Package dme.lib32_n9000 0:2.0.10.0-9.3.1 will be an update
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package      Arch          Version           Repository        Size
=====
Updating:
dme          lib32_n9000   2.0.10.0-9.3.1   localdb           45 M

Transaction Summary
=====
Upgrade      1 Package

Total download size: 45 M
Is this ok [y/N]: y
Retrieving key from file:///etc/pki/rpm-gpg/arm-Nexus9k-dev.gpg
Downloading Packages:
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
 /bootflash/.rpmstore/config/etc/pki/rpm-gpg/arm-Nexus9k-dev.gpg
System at HA Standby, running transaction on Standby first
  Updating   : dme-2.0.10.0-9.3.1.lib32_n9000                                1/2
starting pre-install package version mgmt for dme
pre-install for dme complete
ln: failed to create symbolic link /var/run/mgmt/sharedmeta-hash: File exists
ln: failed to create symbolic link /var/run/mgmt/dme-objstores.conf: File exists
ln: failed to create symbolic link /var/run/mgmt/samlog.config: File exists
mgmt/
mgmt/shmetafiles/
mgmt/shmetafiles/sharedmeta-ArgMetaData
mgmt/shmetafiles/sharedmeta-RelsMetaData
mgmt/shmetafiles/sharedmeta-ClassRelMetaData
```

```

mgmt/shmetafiles/sharedmeta-ChunkMetaData
mgmt/shmetafiles/sharedmeta-ConstPropMetaData
mgmt/shmetafiles/sharedmeta-ConstIdMetaData
mgmt/shmetafiles/sharedmeta-ClassMetaData
mgmt/shmetafiles/sharedmeta-PropRefsMetaData
mgmt/shmetafiles/sharedmeta-SvcMetaData
mgmt/shmetafiles/sharedmeta-ActionContextMetaData
mgmt/shmetafiles/sharedmeta-ConstDefTypeMetaData
mgmt/shmetafiles/sharedmeta-ConstArgMetaData
mgmt/shmetafiles/sharedmeta-ClassNamingMetaData
mgmt/shmetafiles/sharedmeta-ConstMetaData
mgmt/shmetafiles/sharedmeta-PropMetaData
mgmt/shmetafiles/sharedmeta-DnMetaData
Cleanup      : dme-2.0.1.0-9.3.1.lib32_n9000                2/2

Updated:
dme.lib32_n9000 0:2.0.10.0-9.3.1

Complete!
switch-1#

```

Verifying the Installed RPM

You can verify that the DME RPM is installed by using either the NX-OS **show install** command or **dnf list**.

Procedure

Choose the method:

- For NX-OS:

show install active

Example:

```

switch-1# show install active
Boot Image:
    NXOS Image: bootflash:///<boot_image.bin>

Active Packages:
    dme-2.0.1.0-9.3.1.lib32_n9000
switch-1#

```

- For **dnf list**, you must log in to the switch's Bash shell (**run bash**) before issuing the **dnf** commands.

dnf list --patch-only installed | grep dme

Example:

```

switch-1# dnf list --patch-only installed | grep dme
dme.lib32_n9000                2.0.1.0-9.3.1                @localdb

```

Querying for the RPM in the Local Repo

You can query the on-switch (local) repo to verify that the RPM is present.

Procedure

Step 1 run bash

Example:

```
switch-1# run bash
bash-4.3$
```

Logs in to the switch's Bash shell.

Step 2 ls /bootflash/.rpmstore/patching/localrepo/dme-2.0.1.0-9.3.1.lib32_n9000.rpm

Example:

```
bash-4.3$ ls /bootflash/.rpmstore/patching/localrepo/dme-2.0.1.0-9.3.1.lib32_n9000.rpm
inactive_feature_rpms.inf
repodata
```

```
bash-4.3$
```

When the base DME RPM is installed, it is in /rpms.

Downgrading Between Versions of DME RPM

You can downgrade from a higher version of DME RPM to a lower version through either the NX-OS **install** command or **dnf**. By downgrading, you retain the DME Modularity functionality.

The DME RPM is protected, so **install deactivate** and **install remove** are not supported.

Procedure

Choose the downgrade method:

For NX-OS:

- **install add path-to-dme-rpm activate downgrade**

Example:

```
switch-1# install add bootflash:dme-2.0.1.0-9.3.1.lib32_n9000.rpm activate downgrade
Adding the patch (/dme-2.0.1.0-9.3.1.lib32_n9000.rpm)
[#####] 100%
Install operation 94 completed successfully at Fri Jun 7 22:48:34 2019

Activating the patch (/dme-2.0.1.0-9.3.1.lib32_n9000.rpm)
[#####] 100%
Install operation 95 completed successfully at Fri Jun 7 22:49:12 2019
switch-1#
```

- **show install active | include dme**

Example:

```
switch-1# show install active | include dme
          dme-2.0.1.0-9.3.1.lib32_n9000
switch-1#
```

In this example, the DME RPM was downgraded to version 2.0.1.0-9.3.1.

For **dnf**, you must run commands in Bash shell as root user (**run bash sudo su**):

- In Bash, run **dnf downgrade dme dme-rpm**.

This option enables you download directly to a lower version of DME RPM in the repository.

This option option requires user intervention to complete as highlighted in the following command output.

Example:

```
bash-4.3# dnf downgrade dme 2.0.1.0-9.3.1
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
Setting up Downgrade Process
groups-repo                | 1.1 kB    00:00 ...
localdb                    |  951 B    00:00 ...
patching                   |  951 B    00:00 ...
thirdparty                 |  951 B    00:00 ...
wrl-repo                   |  951 B    00:00 ...
Resolving Dependencies
--> Running transaction check
---> Package dme.lib32_n9000 0:2.0.1.0-9.3.1 will be a downgrade
---> Package dme.lib32_n9000 0:2.0.10.0-9.3.1 will be erased
--> Finished Dependency Resolution

Dependencies Resolved
=====
Package      Arch           Version           Repository        Size
=====
Downgrading:
dme          lib32_n9000    2.0.10.0-9.3.1   localdb           45 M

Transaction Summary
=====
Downgrade    1 Package

Total download size: 45 M
Is this ok [y/N]: y
Retrieving key from file:///etc/pki/rpm-gpg/arm-Nexus9k-dev.gpg
Downloading Packages:
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
/bootflash/.rpmstore/config/etc/pki/rpm-gpg/arm-Nexus9k-dev.gpg
System at HA Standby, running transaction on Standby first
  Installing : dme-2.0.1.0-9.3.1.lib32_n9000                1/2
starting pre-install package version mgmt for dme
pre-install for dme complete
ln: failed to create symbolic link /var/run/mgmt/sharedmeta-hash: File exists
ln: failed to create symbolic link /var/run/mgmt/dme-objstores.conf: File exists
ln: failed to create symbolic link /var/run/mgmt/samlog.config: File exists
mgmt/
mgmt/shmetafiles/
mgmt/shmetafiles/sharedmeta-ArgMetaData
mgmt/shmetafiles/sharedmeta-RelsMetaData
mgmt/shmetafiles/sharedmeta-ClassRelMetaData
```

```

mgmt/shmetafiles/sharedmeta-ChunkMetaData
mgmt/shmetafiles/sharedmeta-ConstPropMetaData
mgmt/shmetafiles/sharedmeta-ConstIdMetaData
mgmt/shmetafiles/sharedmeta-ClassMetaData
mgmt/shmetafiles/sharedmeta-PropRefsMetaData
mgmt/shmetafiles/sharedmeta-SvcMetaData
mgmt/shmetafiles/sharedmeta-ActionContextMetaData
mgmt/shmetafiles/sharedmeta-ConstDefTypeMetaData
mgmt/shmetafiles/sharedmeta-ConstArgMetaData
mgmt/shmetafiles/sharedmeta-ClassNamingMetaData
mgmt/shmetafiles/sharedmeta-ConstMetaData
mgmt/shmetafiles/sharedmeta-PropMetaData
mgmt/shmetafiles/sharedmeta-DnMetaData
  Cleanup      : dme-2.0.10.0-9.3.1.lib32_n9000                2/2

Removed:
  dme.lib32_n9000 0:2.0.10.0-9.3.1

Installed:
  dme.lib32_n9000 0:2.0.1.0-9.3.1

Complete!

```

Downgrades from one version of DME RPM to a lower version. In this example, version 2.0.10.0-9.3.1 is downgraded to version 2.0.1.0-9.3.1.

- **dnf list --patch-only installed | grep dme**

Example:

```

bash-4.3# dnf list --patch-only installed | grep dme
dme.lib32_n9000                2.0.1.0-9.3.1                @groups-repo
bash-4.3#

```

Displays the installed version of DME RPM.

Downgrading to the Base RPM

You can downgrade from a higher version of the DME RPM to the base DME RPM by either installing the base DME RPM through the NX-OS **install** command or using **dnf downgrade**.

Procedure

Choose the downgrade method:

For NX-OS:

- **install activate *dme-rpm***

Example:

```

switch-1# install activate dme-2.0.0.0-9.2.1.lib32_n9000.rpm
[#####] 100%
Install operation 89 completed successfully at Fri Jun  7 07:21:45 2019
switch-1#

```


- **show install active | dme**

Example:

```
switch-1# show install active | include dme
          dme-2.0.0.0-9.2.1.lib32_n9000
switch-1#
```

For **dnf**, you must run commands in Bash shell as root user (**run bash sudo su**):

- In Bash, run **dnf downgrade dme dme-rpm**.

This option enables downgrading directly to the base DME RPM.

This option requires user intervention to complete as highlighted in the following command output.

Example:

```
bash-4.3# dnf downgrade dme-2.0.0.0-9.3.1.lib32_n9000
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
Setting up Downgrade Process
groups-repo           | 1.1 kB    00:00 ...
localdb               | 951 B    00:00 ...
patching              | 951 B    00:00 ...
thirdparty           | 951 B    00:00 ...
wrl-repo              | 951 B    00:00 ...
Resolving Dependencies
--> Running transaction check
---> Package dme.lib32_n9000 0:2.0.0.0-9.3.1 will be a downgrade
---> Package dme.lib32_n9000 0:2.0.10.0-9.3.1 will be erased
--> Finished Dependency Resolution
```

Dependencies Resolved

```
=====
Package   Arch           Version           Repository        Size
=====
Downgrading:
dme       lib32_n9000    2.0.0.0-9.3.1    groups-repo      44 M
```

Transaction Summary

```
=====
Downgrade      1 Package
```

Total download size: 44 M

Is this ok [y/N]: y

Downloading Packages:

Running Transaction Check

Running Transaction Test

Transaction Test Succeeded

Running Transaction

```
Installing : dme-2.0.0.0-9.3.1.lib32_n9000                                1/2
```

```
starting pre-install package version mgmt for dme
```

```
pre-install for dme complete
```

```
mgmt/
```

```
mgmt/shmetafiles/
```

```
mgmt/shmetafiles/sharedmeta-ChunkMetaData
```

```
mgmt/shmetafiles/sharedmeta-ClassMetaData
```

```
mgmt/shmetafiles/sharedmeta-ArgMetaData
```

```
mgmt/shmetafiles/sharedmeta-ConstMetaData
```

```
mgmt/shmetafiles/sharedmeta-ConstIdMetaData
```

```
mgmt/shmetafiles/sharedmeta-ConstDefTypeMetaData
```

```
mgmt/shmetafiles/sharedmeta-ConstPropMetaData
```

```
mgmt/shmetafiles/sharedmeta-ConstArgMetaData
```

```

mgmt/shmetafiles/sharedmeta-ClassRelMetaData
mgmt/shmetafiles/sharedmeta-DnMetaData
mgmt/shmetafiles/sharedmeta-PropRefsMetaData
mgmt/shmetafiles/sharedmeta-PropMetaData
mgmt/shmetafiles/sharedmeta-RelsMetaData
mgmt/shmetafiles/sharedmeta-ActionContextMetaData
mgmt/shmetafiles/sharedmeta-SvcMetaData
mgmt/shmetafiles/sharedmeta-ClassNamingMetaData
Cleanup      : dme-2.0.10.0-9.3.1.lib32_n9000          2/2

```

```

Removed:
dme.lib32_n9000 0:2.0.10.0-9.3.1

```

```

Installed:
dme.lib32_n9000 0:2.0.0.0-9.3.1

```

```

Complete!
bash-4.3#

```

Installs the base DME RPM.

- **dnf list --patch-only installed | grep dme**

Example:

```

bash-4.3# dnf list --patch-only installed | grep dme
dme.lib32_n9000                2.0.0.0-9.3.1                @groups-repo
bash-4.3#

```

Displays the installed base DME RPM.

Managing Patch RPMs

RPM Installation Prerequisites

Use these procedures to verify that the system is ready before installing or adding an RPM.

SUMMARY STEPS

1. switch# **show logging logfile | grep -i "System ready"**
2. switch# **run bash sudo su**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | switch# show logging logfile grep -i "System ready" | Before running Bash, this step verifies that the system is ready before installing or adding an RPM. Proceed if you see output similar to the following: |

| | Command or Action | Purpose |
|--------|--|--|
| | | 2018 Mar 27 17:24:22 switch %ASCII-CFG-2-CONF_CONTROL: System ready |
| Step 2 | switch# run bash sudo su Example: switch# run bash sudo su bash-4.2# | Loads Bash. |

Adding Patch RPMs from Bash

Procedure

| | Command or Action | Purpose |
|--------|---|---|
| Step 1 | dnf list --patch-only | Displays a list of the patch RPMs present on the switch. |
| Step 2 | sudo dnf install --add <i>URL_of_patch</i> | Adds the patch to the repository, where <i>URL_of_patch</i> is a well-defined format, such as bootflash:/patch , not in standard Linux format, such as /bootflash/patch . |
| Step 3 | dnf list --patch-only available | Displays a list of the patches that are added to the repository but are in an inactive state. |

Example

The following is an example of installing the **nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000** RPM:

```
bash-4.2# dnf list --patch-only
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B    00:00 ...
patching                   | 951 B    00:00 ...
thirdparty                 | 951 B    00:00 ...
bash-4.2#
bash-4.2# sudo dnf install --add
bootflash:/nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000.rpm
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B    00:00 ...
patching                   | 951 B    00:00 ...
thirdparty                 | 951 B    00:00 ...
##### ] 70%Install operation 135 completed successfully at Tue Mar 27 17:45:34
2018.

##### ] 100%
bash-4.2#
```

Once the patch RPM is installed, verify that it was installed properly. The following command lists the patches that are added to the repository and are in the inactive state:

```
bash-4.2# dnf list --patch-only available
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
                : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B    00:00 ...
patching                   | 951 B    00:00 ...
thirdparty                 | 951 B    00:00 ...
nxos.CSCab00001-n9k_ALL.lib32_n9000  1.0.0-7.0.3.I7.3  patching
bash-4.2#
```

You can also add patches to a repository from a tar file, where the RPMs are bundled in the tar file. The following example shows how to add two RPMs that are part of the `nxos.CSCab00002_CSCab00003-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000` tar file to the patch repository:

```
bash-4.2# sudo dnf install --add
bootflash:/nxos.CSCab00002_CSCab00003-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000.tar
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
                : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B    00:00 ...
patching                   | 951 B    00:00 ...
thirdparty                 | 951 B    00:00 ...
[#####] 70%Install operation 146 completed successfully at Tue Mar 27 21:17:39
2018.

[#####] 100%
bash-4.2#
bash-4.2# dnf list --patch-only
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
                : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B    00:00 ...
patching                   | 951 B    00:00 ...
patching/primary          | 942 B    00:00 ...
patching                   |          2/2
thirdparty                 | 951 B    00:00 ...
nxos.CSCab00003-n9k_ALL.lib32_n9000  1.0.0-7.0.3.I7.3  patching
nxos.CSCab00002-n9k_ALL.lib32_n9000  1.0.0-7.0.3.I7.3  patching
bash-4.2#
```

Activating a Patch RPM

Before you begin

Verify that you have added the necessary patch RPM to the repository using the instructions in [#unique_66](#).

Procedure

| | Command or Action | Purpose |
|--------|---|---|
| Step 1 | <code>sudo dnf install <i>patch_RPM</i> --nocommit</code> | Activates the patch RPM, where <i>patch_RPM</i> is a patch that is located in the repository. Do not provide a location for the patch in this step. |

| | Command or Action | Purpose |
|--|-------------------|--|
| | | Note Adding the <code>--nocommit</code> flag to the command means that the patch RPM is activated in this step, but not committed. See Committing a Patch RPM, on page 40 for instructions on committing the patch RPM after you have activated it. |

Example

The following example shows how to activate the `nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000` patch RPM:

```
bash-4.2# sudo dnf install nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000 --nocommit
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
groups-repo                               | 1.1 kB      00:00 ...
localdb                                    | 951 B       00:00 ...
patching                                   | 951 B       00:00 ...
thirdparty                                 | 951 B       00:00 ...
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package nxos.CSCab00001-n9k_ALL.lib32_n9000 0:1.0.0-7.0.3.I7.3 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                                Arch      Version                Repository  Size
=====
Installing:
nxos.CSCab00001-n9k_ALL                lib32_n9000  1.0.0-7.0.3.I7.3      patching   28 k

Transaction Summary
=====
Install      1 Package

Total download size: 28 k
Installed size: 82 k
Is this ok [y/N]: y
Downloading Packages:
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000      1/1
[##### ] 90%error: reading
/var/sysmgr/tmp/patches/CSCab00001-n9k_ALL/isan/bin/sysinfo manifest, non-printable characters
found

Installed:
nxos.CSCab00001-n9k_ALL.lib32_n9000 0:1.0.0-7.0.3.I7.3

Complete!
Install operation 140 completed successfully at Tue Mar 27 18:07:40 2018.

[#####] 100%
bash-4.2#
```

Enter the following command to verify that the patch RPM was activated successfully:

```
bash-4.2# dnf list --patch-only
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B     00:00 ...
patching                   | 951 B     00:00 ...
thirdparty                 | 951 B     00:00 ...
nxos.CSCab00001-n9k_ALL.lib32_n9000  1.0.0-7.0.3.I7.3  installed
bash-4.2#
```

Committing a Patch RPM

Procedure

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | <code>sudo dnf install patch_RPM --commit</code> | Commits the patch RPM. The patch RPM must be committed to keep it active after reloads. |

Example

The following example shows how to commit the `nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000` patch RPM:

```
bash-4.2# sudo dnf install nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000 --commit
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B     00:00 ...
patching                   | 951 B     00:00 ...
thirdparty                 | 951 B     00:00 ...
Install operation 142 completed successfully at Tue Mar 27 18:13:16 2018.

[#####] 100%
bash-4.2#
```

Enter the following command to verify that the patch RPM was committed successfully:

```
bash-4.2# dnf list --patch-only committed
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B     00:00 ...
patching                   | 951 B     00:00 ...
thirdparty                 | 951 B     00:00 ...
nxos.CSCab00001-n9k_ALL.lib32_n9000  1.0.0-7.0.3.I7.3  installed
bash-4.2#
```

Deactivating a Patch RPM

Procedure

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | <code>sudo dnf erase patch_RPM --nocommit</code> | Deactivates the patch RPM. Note Adding the <code>--nocommit</code> flag to the command means that the patch RPM is only deactivated in this step. |
| Step 2 | <code>sudo dnf install patch_RPM --commit</code> | Commits the patch RPM. You will get an error message if you try to remove the patch RPM without first committing it. |

Example

The following example shows how to deactivate the `nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000` patch RPM:

```
bash-4.2# sudo dnf erase nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000 --nocommit
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
                : protect-packages
Setting up Remove Process
Resolving Dependencies
--> Running transaction check
---> Package nxos.CSCab00001-n9k_ALL.lib32_n9000 0:1.0.0-7.0.3.I7.3 will be erased
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch            Version          Repository      Size
=====
Removing:
nxos.CSCab00001-n9k_ALL lib32_n9000      1.0.0-7.0.3.I7.3 @patching      82 k

Transaction Summary
=====
Remove                1 Package

Installed size: 82 k
Is this ok [y/N]: y
Downloading Packages:
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
[#####          ] 30%error: reading
/var/sysmgr/tmp/patches/CSCab00001-n9k_ALL/isan/bin/sysinfo manifest, non-printable characters
found
Erasing      : nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000      1/1
[#####          ] 90%
Removed:
nxos.CSCab00001-n9k_ALL.lib32_n9000 0:1.0.0-7.0.3.I7.3
```

```
Complete!
Install operation 143 completed successfully at Tue Mar 27 21:03:47 2018.

[#####] 100%
bash-4.2#
```

You must commit the patch RPM after deactivating it. If you do not commit the patch RPM after deactivating it, you will get an error message if you try to remove the patch RPM using the instructions in [Removing a Patch RPM, on page 42](#).

```
bash-4.2# sudo dnf install nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000 --commit
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
                : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B    00:00 ...
patching                   | 951 B    00:00 ...
thirdparty                 | 951 B    00:00 ...
Install operation 144 completed successfully at Tue Mar 27 21:09:28 2018.

[#####] 100%
bash-4.2#
```

Enter the following command to verify that the patch RPM has been committed successfully:

```
bash-4.2# dnf list --patch-only
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
                : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B    00:00 ...
patching                   | 951 B    00:00 ...
thirdparty                 | 951 B    00:00 ...
nxos.CSCab00001-n9k_ALL.lib32_n9000  1.0.0-7.0.3.I7.3  patching
bash-4.2#
```

Removing a Patch RPM

Procedure

| | Command or Action | Purpose |
|--------|---|--------------------------------|
| Step 1 | <code>sudo dnf install --remove <i>patch_RPM</i></code> | Removes an inactive patch RPM. |

Example

The following example shows how to remove the `nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000` patch RPM:

```
bash-4.2# sudo dnf install --remove nxos.CSCab00001-n9k_ALL-1.0.0-7.0.3.I7.3.lib32_n9000
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
                : protect-packages
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B    00:00 ...
patching                   | 951 B    00:00 ...
thirdparty                 | 951 B    00:00 ...
[##### ] 50%Install operation 145 completed successfully at Tue Mar 27 21:11:05
2018.
```



```
[#####] 100%
bash-4.2#
```



Note If you see the following error message after attempting to remove the patch RPM:

Install operation 11 "failed because patch was not committed". at Wed Mar 28 22:14:05 2018

Then you did not commit the patch RPM before attempting to remove it. See [Deactivating a Patch RPM, on page 41](#) for instructions on committing the patch RPM before attempting to remove it.

Enter the following command to verify that the inactive patch RPM was removed successfully:

```
bash-4.2# dnf list --patch-only
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
groups-repo                | 1.1 kB      00:00 ...
localdb                    | 951 B       00:00 ...
patching                   | 951 B       00:00 ...
patching/primary          | 197 B       00:00 ...
thirdparty                 | 951 B       00:00 ...
bash-4.2#
```

Persistently Daemonizing an SDK- or ISO-built Third Party Process

Your application should have a startup Bash script that gets installed in `/etc/init.d/application_name`. This startup Bash script should have the following general format (for more information on this format, see <http://linux.die.net/man/8/chkconfig>).

```
#!/bin/bash
#
# <application_name> Short description of your application
#
# chkconfig: 2345 15 85
# description: Short description of your application
#
### BEGIN INIT INFO
# Provides: <application_name>
# Required-Start: $local_fs $remote_fs $network $named
# Required-Stop: $local_fs $remote_fs $network
# Description: Short description of your application
### END INIT INFO
# See how we were called.
case "$1" in
start)
# Put your startup commands here
# Set RETVAL to 0 for success, non-0 for failure
;;
stop)
# Put your stop commands here
# Set RETVAL to 0 for success, non-0 for failure
;;
status)
# Put your status commands here
```

```

# Set RETVAL to 0 for success, non-0 for failure
;;
restart|force-reload|reload)
# Put your restart commands here
# Set RETVAL to 0 for success, non-0 for failure
;;
*)
echo $"Usage: $prog {start|stop|status|restart|force-reload}"
RETVAL=2
esac

exit $RETVAL

```

Persistently Starting Your Application from the Native Bash Shell

Procedure

-
- Step 1** Install your application startup Bash script that you created into `/etc/init.d/application_name`
 - Step 2** Start your application with `/etc/init.d/application_name start`
 - Step 3** Enter `chkconfig --add application_name`
 - Step 4** Enter `chkconfig --level 3 application_name on`
Run level 3 is the standard multi-user run level, and the level at which the switch normally runs.
 - Step 5** Verify that your application is scheduled to run on level 3 by running `chkconfig --list application_name` and confirm that level 3 is set to on
 - Step 6** Verify that your application is listed in `/etc/rc3.d`. You should see something like this, where there is an 'S' followed by a number, followed by your application name (`tcollector` in this example), and a link to your Bash startup script in `../init.d/application_name`
-

```

bash-4.2# ls -l /etc/rc3.d/tcollector
lrwxrwxrwx 1 root root 20 Sep 25 22:56 /etc/rc3.d/S15tcollector -> ../init.d/tcollector
bash-4.2#

```

Synchronize Files from Active Bootflash to Standby Bootflash

Cisco Nexus 9500 platform switches are generally configured with two supervisor modules to provide high availability (one active supervisor module and one standby supervisor module). Each supervisor module has its own bootflash file system for file storage, and the Active and Standby bootflash file systems are generally independent of each other. If there is a need for specific content on the active bootflash, that same content is probably also needed on the standby bootflash in case there is a switchover at some point.

Before the Cisco NX-OS 9.2(2) release, you had to manually manage this content between the Active and Standby supervisor modules. Starting with Cisco NX-OS 9.2(2), certain files and directories on the active

supervisor module, or active bootflash (/bootflash), can be automatically synchronized to the standby supervisor module, or standby bootflash (/bootflash_sup-remote), if the standby supervisor module is up and available. You can select the files and directories to be synchronized by loading Bash on your switch, then adding the files and directories that you would like to have synchronized from the active bootflash to the standby bootflash into the editable file /bootflash/bootflash_sync_list.

For example:

```
switch# run bash
bash-4.2# echo "/bootflash/home/admin" | sudo tee --append /bootflash/bootflash_sync_list
bash-4.2# echo "/bootflash/nxos.7.0.3.I7.3.5.bin" | sudo tee --append
/bootflash/bootflash_sync_list
bash-4.2# cat /bootflash/bootflash_sync_list
/bootflash/home/admin
/bootflash/nxos.7.0.3.I7.3.5.bin
```

```
bash-4.2# echo /bootflash/home/admin >> /bootflash/bootflash_sync_list
```

```
bash-4.2# echo /bootflash/nxos.7.0.3.I7.3.5.bin >>
/bootflash/bootflash_sync_list
```

When changes are made to the files or directories on the active bootflash, these changes are automatically synchronized to standby bootflash, if the standby bootflash is up and available. If the standby bootflash is rebooted, either as a regular boot, switchover or manual standby reload, a catch-up synchronization of changes to the active bootflash is pushed out to the standby bootflash, once the standby supervisor comes online.

Following are the characteristics and restrictions for the editable /bootflash/bootflash_sync_list file:

- The /bootflash/bootflash_sync_list file is automatically created on the first run and is empty at that initial creation state.
- Entries in the /bootflash/bootflash_sync_list file follow these guidelines:
 - One entry per line
 - Entries are given as Linux paths (for example, /bootflash/img.bin)
 - Entries must be within the /bootflash file system
- The /bootflash/bootflash_sync_list file itself is automatically synchronized to the standby bootflash. You can also manually copy the /bootflash/bootflash_sync_list file to or from the supervisor module using the **copy** virtual shell (VSH) command.
- You can edit the /bootflash/bootflash_sync_list file directly on the supervisor module with the following command:

```
run bash vi /bootflash/bootflash_sync_list
```

All output from the synchronization event is redirected to the log file /var/tmp/bootflash_sync.log. You can view or tail this log file using either of the following commands:

```
run bash less /var/tmp/bootflash_sync.log
```

```
run bash tail -f /var/tmp/bootflash_sync.log
```

The synchronization script will not delete files from the standby bootflash directories unless it explicitly receives a delete event for the corresponding file on the active bootflash directories. Sometimes, the standby bootflash might have more used space than the active bootflash, which results in the standby bootflash running out of space when the active bootflash is synchronizing to it. To make the standby bootflash an exact mirror of the active bootflash (to delete any extra files on the standby bootflash), enter the following command:

```
run bash sudo rsync -a --delete /bootflash/ /bootflash_sup-remote/
```

The synchronization script should continue to run in the background without crashing or exiting. However, if it does stop running for some reason, you can manually restart it using the following command:

```
run bash sudo /isan/etc/rc.d/rc.isan-start/S98bootflash_sync.sh start
```

Copy Through Kstack

In Cisco NX-OS release 9.3(1) and later, file copy operations have the option of running through a different network stack by using the **use-kstack** option. Copying files through **use-kstack** enables faster copy times. This option can be beneficial when copying files from remote servers that are multiple hops from the switch. The **use-kstack** option work with copying files from, and to, the switch though standard file copy features, such as **scp** and **sftp**.



Note The **use-kstack** option does not work when the switch is running the FIPS mode feature. If the switch has FIPS mode that is enabled, the copy operation is still successful, but through the default copy method.

To copy through **use-kstack**, append the argument to the end of an NX-OS **copy** command. Some examples:

```
switch-1# copy scp://test@10.1.1.1/image.bin . vrf management use-kstack
switch-1#
switch-1# copy scp://test@10.1.1.1/image.bin bootflash:// vrf management
use-kstack
switch-1#
switch-1# copy scp://test@10.1.1.1/image.bin . use-kstack
switch-1#
switch-1# copy scp://test@10.1.1.1/image.bin bootflash:// vrf default
use-kstack
switch-1#
```

The **use-kstack** option is supported for all NX-OS **copy** commands and file systems. The option is OpenSSL (Secure Copy) certified.

An Example Application in the Native Bash Shell

The following example demonstrates an application in the Native Bash Shell:

```
bash-4.2# cat /etc/init.d/hello.sh
#!/bin/bash

PIDFILE=/tmp/hello.pid
OUTPUTFILE=/tmp/hello
```

```

echo $$ > $PIDFILE
rm -f $OUTPUTFILE
while true
do
    echo $(date) >> $OUTPUTFILE
    echo 'Hello World' >> $OUTPUTFILE
    sleep 10
done
bash-4.2#
bash-4.2#
bash-4.2# cat /etc/init.d/hello
#!/bin/bash
#
# hello Trivial "hello world" example Third Party App
#
# chkconfig: 2345 15 85
# description: Trivial example Third Party App
#
### BEGIN INIT INFO
# Provides: hello
# Required-Start: $local_fs $remote_fs $network $named
# Required-Stop: $local_fs $remote_fs $network
# Description: Trivial example Third Party App
### END INIT INFO

PIDFILE=/tmp/hello.pid

# See how we were called.
case "$1" in
start)
    /etc/init.d/hello.sh &
    RETVAL=$?
    ;;
stop)
    kill -9 `cat $PIDFILE`
    RETVAL=$?
    ;;
status)
    ps -p `cat $PIDFILE`
    RETVAL=$?
    ;;
restart|force-reload|reload)
    kill -9 `cat $PIDFILE`
    /etc/init.d/hello.sh &
    RETVAL=$?
    ;;
*)
echo $"Usage: $prog {start|stop|status|restart|force-reload}"
RETVAL=2
esac

exit $RETVAL
bash-4.2#
bash-4.2# chkconfig --add hello
bash-4.2# chkconfig --level 3 hello on
bash-4.2# chkconfig --list hello
hello          0:off  1:off  2:on   3:on   4:on   5:on   6:off
bash-4.2# ls -al /etc/rc3.d/*hello*
lrwxrwxrwx 1 root root 15 Sep 27 18:00 /etc/rc3.d/S15hello -> ../init.d/hello
bash-4.2#
bash-4.2# reboot

```

After reload

```
bash-4.2# ps -ef | grep hello
root      8790      1  0 18:03 ?          00:00:00 /bin/bash /etc/init.d/hello.sh
root      8973     8775  0 18:04 ttys0    00:00:00 grep hello
bash-4.2#
bash-4.2# ls -al /tmp/hello*
-rw-rw-rw- 1 root root 205 Sep 27 18:04 /tmp/hello
-rw-rw-rw- 1 root root   5 Sep 27 18:03 /tmp/hello.pid
bash-4.2# cat /tmp/hello.pid
8790
bash-4.2# cat /tmp/hello
Sun Sep 27 18:03:49 UTC 2015
Hello World
Sun Sep 27 18:03:59 UTC 2015
Hello World
Sun Sep 27 18:04:09 UTC 2015
Hello World
Sun Sep 27 18:04:19 UTC 2015
Hello World
Sun Sep 27 18:04:29 UTC 2015
Hello World
Sun Sep 27 18:04:39 UTC 2015
Hello World
bash-4.2#
```



CHAPTER 5

Guest Shell

- [About the Guest Shell, on page 49](#)
- [Guidelines and Limitations for Guestshell, on page 50](#)
- [Accessing the Guest Shell, on page 56](#)
- [Resources Used for the Guest Shell, on page 57](#)
- [Capabilities in the Guestshell, on page 57](#)
- [Security Posture for Virtual ServicesGuest Shell, on page 66](#)
- [Guest File System Access Restrictions , on page 69](#)
- [Managing the Guest Shell, on page 70](#)
- [Verifying Virtual Service and Guest Shell Information, on page 82](#)
- [Persistently Starting Your Application From the Guest Shell, on page 83](#)
- [Procedure for Persistently Starting Your Application from the Guest Shell, on page 84](#)
- [An Example Application in the Guest Shell, on page 84](#)
- [Troubleshooting Guest Shell Issues, on page 85](#)

About the Guest Shell

In addition to the NX-OS CLI and Bash access on the underlying Linux environment, switches support access to a decoupled execution space running within a Linux Container (LXC) called the “Guest Shell”.

From within the Guest Shell the network-admin has the following capabilities:

- Access to the network over Linux network interfaces.
- Access to the switch's bootflash.
- Access to the switch's volatile tmpfs.
- Access to the switch's CLI.
- Access to the switch's host file system.
- Access to Cisco NX-API REST.
- The ability to install and run python scripts.
- The ability to install and run 32-bit and 64-bit Linux applications.

Decoupling the execution space from the native host system allows customization of the Linux environment to suit the needs of the applications without impacting the host system or applications running in other Linux Containers.

On NX-OS devices, Linux Containers are installed and managed with the virtual-service commands. The Guest Shell will appear in the virtual-service show command output.



Note By default, the Guest Shell occupies approximately 5 MB of RAM and 200 MB of bootflash when enabled. Beginning with Cisco NX-OS Release 7.0(3)I2(1) the Guest Shell occupies approximately 35 MB of RAM. Use the **guestshell destroy** command to reclaim resources if the Guest Shell is not used.



Note Beginning with Cisco NX-OS 7.0(3)F3(1)NX-OS 7.0(3)I7(1), the Guest Shell is supported on the Cisco Nexus 95083500 switch.

Guidelines and Limitations for Guestshell

Common Guidelines Across All Releases



Important If you have performed custom work inside your installation of the Guestshell, save your changes to the bootflash, off-box storage, or elsewhere outside the Guestshell root file system before performing a `guestshell upgrade`.

The `guestshell upgrade` command essentially performs a `guestshell destroy` and `guestshell enable` in succession.

- Guest Shell is not supported on 3500 models with 4GB of memory (3524, 3548, 3524-X, 3548-X). It is supported on the platforms with higher memory, such as -XL.
- If you are running a third-party DHCPD server in Guestshell, there might be issues with offers reaching the client if used along with SVI. A possible workaround is to use broadcast responses.
- Use the `run guestshell` CLI command to access the Guestshell on the switch: The `run guestshell` command parallels the `run bash` command that is used to access the host shell. This command allows you to access the Guestshell and get a Bash prompt or run a command within the context of the Guestshell. The command uses password-less SSH to an available port on the localhost in the default network namespace.
- When routes are being exchanged between two different VRFs in NXOS (either statically or dynamically), the routing table for the corresponding VRF / Namespace does not get populated with the "shared route" in the guestshell container.
- The `sshd` utility can secure the pre-configured SSH access into the Guestshell by listening on `localhost` to avoid connection attempts from outside the network. The `sshd` has the following features:
 - It is configured for key-based authentication without fallback to passwords.

- Only `root` can read keys use to access the Guestshell after Guestshell restarts.
- Only `root` can read the file that contains the key on the host to prevent a nonprivileged user with host Bash access from being able to use the key to connect to the Guestshell. Network-admin users may start another instance of `sshd` in the Guestshell to allow remote access directly into the Guestshell, but any user that logs into the Guestshell is also given network-admin privilege.



Note Introduced in Guestshell 2.2 (0.2), the key file is readable for whom the user account was created for.

In addition, the Guestshell accounts are not automatically removed, and must be removed by the network administrator when no longer needed.

Guestshell installations before 2.2 (0.2) will not dynamically create individual user accounts.

- Installing the Cisco NX-OS software release on a fresh out-of-the-box switch will automatically enable the Guestshell. Subsequent upgrades to the switch software will not automatically upgrade Guestshell.
- Guestshell releases increment the major number when distributions or distribution versions change.
- Guestshell for NX-OS can access front-panel ports as first-class Linux interfaces.
- Guestshell for NX-OS can access Command shell through `dohost` using local Unix socket to NX-API.
 1. Guestshell for NX-OS: Access to NX-API socket is allowed only for `root/admin` user privilege from 9.3(8) and later.
 2. Guestshell for NX-OS: Access to NX-OS filesystem only as `root/admin` user in 9.3(8) and later.
- Guestshell releases increment the minor number when CVEs have been addressed. The Guestshell updates CVEs only when CentOS makes them publicly available.
- Cisco recommends using **`dnf update`** to pick up third-party security vulnerability fixes directly from the CentOS repository. This provides the flexibility of getting updates as, and when, available without needing to wait for a Cisco NX-OS software update.

Alternatively, using the **`guestshell update`** command would replace the existing Guestshell rootfs. Any customizations and software package installations would then need to be performed again within the context of this new Guestshell rootfs.

CentOS end of life and impact on Guestshell

Guestshell is an **LXC container based on CentOS environment**. As per updates in the open source community, CentOS 8 Project is reaching end of support by December 2021. The CentOS 7 project is to continue through and is targeted to reach end of support by June 2024. Due to this long term support for CentOS 7, the latest Cisco NX-OS software 10.2.x is packaged with Guestshell 2.11 (CentOS 7 based). This replaces Guestshell 3.0 (CentOS 8) which is the default environment in 10.1.x release.

Guestshell 2.11

Beginning with Cisco NX-OS release 10.2(1), CentOS 7 is re-introduced as the default Guestshell environment. See section "*CentOS End of Life*" for a detailed explanation on the reasons.

Guestshell 2.11 comes with python2 and python3.6 support. The functionality between Guestshell 2.11 and Guestshell 3.0 remains the same.



Note The rootfs size of Guestshell 2.11 has increased to approximately 200 MB.

Guestshell 3.0

Guestshell 3.0 is deprecated and is not available from NX-OS 10.2.x. It is recommended to use Guestshell 2.11. However, the 10.2.x software shall remain compatible with Guestshell 3.0 containers and 3.0 guestshell containers running on 10.1.x. software shall continue to function after upgrade to 10.2.x.



Note The rootfs size in Guestshell 3.0 is 220 MB versus the 170 MB in Guestshell 2.0.

Guestshell 4.x

Guestshell 2.x contains Centos 7. End of life for Centos 7 is early 2024. Hence, Guestshell 4.x is a RockyLinux 9 based lxc container that will replace Guestshell 2.x. Guestshell 4.x is available in the following NX-OS releases as downloadable and default options.

| NXOS release | Guestshell default package version | Guestshell downloadable option applicable |
|---------------------|------------------------------------|---|
| 9.3(14) and above | 4.1 or above | Not needed |
| 10.2(6) | 2.15 | 4.0 |
| 10.2(7) | 2.15 | 4.1 |
| 10.2(8) and above | 4.1 or above | Not needed |
| 10.3(4) | 2.15 | 4.0 |
| 10.3(5) and above | 4.1 or above | Not needed |
| 10.4(1) and 10.4(2) | 2.15 | 4.0 |
| 10.4(3) and above | 4.1 or above | Not needed |
| 10.5(1) and above | 4.1 or above | Not needed |



Note The rootfs size in Guestshell 4.x is 400 MB versus the 350 MB in Guestshell 2.x. Guestshell 4.x downloadable OVA is backward compactable to all releases running Guestshell 2.x as default.

Upgrading from Guestshell 1.0 to Guestshell 2.x

Guestshell 2.x is based on a CentOS 7 root file system. If you have an off-box repository of `.conf` files or utilities that pulled the content down into Guestshell 1.0, you must repeat the same deployment steps in Guestshell 2.x. Your deployment script may need to be adjusted to account for the CentOS 7 differences.

Downgrading NX-OS from Jacksonville release Guestshell 3.0

Beginning with Cisco NX-OS release 10.1(1), infrastructure version for Guestshell 3.0 support is increased to 1.11 (check with `show virtual-service` command). Therefore, Guestshell 3.0 OVA cannot be used in previous releases. If used, the **Install all** command will validate version mismatch and throws an error. It is recommended to destroy Guestshell 3.0 before downgrading to previous releases so that Guestshell 3.0 does not come up in previous releases.

Upgrading from Guestshell 2.x to Guestshell 4.x Downloadable OVA

Guestshell 4.x can be downloaded from Cisco's official software download page and can be installed using command `guestshell upgrade` command.

Following table shows the guest shell releases:

Table 2: Guest Shell Releases

| Guest Shell Releases | NX-OS Supported Releases | Python Version(s) Supported |
|----------------------|--------------------------|-----------------------------|
| 2.x | 10.4.1 | python2.7 and python 3.6 |
| 3.0 | 10.1.x | python 3.6 |
| 4.x | 10.4.1 | python 3.9 |

Use below commands to upgrade to Guestshell 4.x:

- Execute command **guestshell enable package <downloaded ova>** when guestshell is not installed.
- Execute command **guestshell upgrade package <downloaded ova>** when guestshell is installed and running.



Note Systems with 4 GB of RAM will not enable Guestshell by default. Use the **guestshell enable** command to install and enable Guestshell.

The **install all** command validates the compatibility between the current Cisco NX-OS image against the target Cisco NX-OS image.

The following is an example output from installing an incompatible image:

```
switch#
Installer will perform compatibility check first. Please wait.
uri is: /
2014 Aug 29 20:08:51 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE:
Successfully activated virtual service 'guestshell+'
Verifying image bootflash:/n9kpregs.bin for boot variable "nxos".
[#####] 100% -- SUCCESS
Verifying image type.
[#####] 100% -- SUCCESS
```

```

Preparing "" version info using image bootflash:/
[#####] 100% -- SUCCESS
Preparing "bios" version info using image bootflash:/
[#####] 100% -- SUCCESS
Preparing "" version info using image bootflash:/
[#####] 100% -- SUCCESS
Preparing "" version info using image bootflash:/
[#####] 100% -- SUCCESS
Preparing "nxos" version info using image bootflash:/
[#####] 100% -- SUCCESS
Preparing "" version info using image bootflash:/
[#####] 100% -- SUCCESS
Preparing "" version info using image bootflash:/
[#####] 100% -- SUCCESS
"Running-config contains configuration that is incompatible with the new image (strict
incompatibility).
Please run 'show incompatibility-all nxos <image>' command to find out which feature
needs to be disabled."
Performing module support checks.
[#####] 100% -- SUCCESS
Notifying services about system upgrade.
[# ] 0% -- FAIL.
Return code 0x42DD0006 ((null)).
"Running-config contains configuration that is incompatible with the new image (strict
incompatibility).
Please run 'show incompatibility-all nxos <image>' command to find out
which feature needs to be disabled."
Service "vman" in vdc 1: Guestshell not supported, do 'guestshell destroy' to remove
it and then retry ISSU
Pre-upgrade check failed. Return code 0x42DD0006 ((null)).
switch#

```



Note As a best practice, remove the Guestshell with the **guestshell destroy** command before reloading an older Cisco NX-OS image that does not support the Guestshell.

Pre-Configured SSHD Service

The Guestshell starts an OpenSSH server upon boot up. The server listens on a randomly generated port on the localhost IP address interface 127.0.0.1 only. This provides the password-less connectivity into the Guestshell from the NX-OS virtual-shell when the guestshell keyword is entered. If this server is killed or its configuration (residing in `/etc/ssh/sshd_config-cisco`) is altered, access to the Guestshell from the NX-OS CLI might not work.

The following steps instantiate an OpenSSH server within the Guestshell as root:

1. Determine which network namespace or VRF you want to establish your SSH connections through.
2. Determine the port that you want OpenSSH to listen on. Use the NX-OS command **show socket connection** to view ports already in use.



Note The Guestshell sshd service for password-less access uses a randomized port starting at 17680 through 49150. To avoid port conflict, choose a port outside this range.

The following steps start the OpenSSH server. The examples start the OpenSSH server for management netns on IP address 10.122.84.34:2222:

1. Create the following files: `/usr/lib/systemd/system/sshd-mgmt.service` and `/etc/ssh/sshd-mgmt_config`. The files should have the following configurations:
2. Copy the Unit and Service contents from the `/usr/lib/systemd/system/ssh.service` file to `sshd-mgmt.service`.
3. Edit the `sshd-mgmt.service` file to match the following:

```
[Unit]
Description=OpenSSH server daemon
After=network.target sshd-keygen.service
Wants=sshd-keygen.service

[Service]
EnvironmentFile=/etc/sysconfig/ssh
ExecStartPre=/usr/sbin/sshd-keygen
ExecStart=/sbin/ip netns exec management /usr/sbin/sshd -f /etc/ssh/sshd-mgmt_config
-D $OPTIONS
ExecReload=/bin/kill -HUP $MAINPID
KillMode=process
Restart=on-failure
RestartSec=42s
[Install]
WantedBy=multi-user.target
```

4. Copy the contents of `/etc/ssh/sshd-config` to `/etc/ssh/sshd-mgmt_config`. Modify the ListenAddress IP and port as necessary.

```
Port 2222
ListenAddress 10.122.84.34
```

5. Start the `systemctl` daemon using the following commands:

```
sudo systemctl daemon-reload
sudo systemctl start sshd-mgmt.service
sudo systemctl status sshd-mgmt.service -l
```

6. (Optional) Check the configuration.

```
ss -tnldp | grep 2222
```

7. SSH into Guestshell:

```
ssh -p 2222 guestshell@10.122.84.34
```

8. Save the configuration across multiple Guestshell or switch reboots.

```
sudo systemctl enable sshd-mgmt.service
```

9. For passwordless SSH/SCP and remote execution, generate the public and private keys for the user ID you want to user for SSH/SCP using the `ssh-keygen -t dsa` command.

The key is then stored in the `id_rsa` and `id_rsa.pub` files in the `/.ssh` directory:

```
[root@node01 ~]# cd ~/.ssh
[root@node02 .ssh]# ls -l
total 8
-rw-----. 1 root root 1675 May 5 15:01 id_rsa
-rw-r--r--. 1 root root 406 May 5 15:01 id_rsa.pub
```

- Copy the public key into the machine you want to SSH into and fix permissions:

```
cat id_rsa.pub >> /root/.ssh/authorized_keys
chmod 700 /root/.ssh
chmod 600 /root/.ssh/*
```

- SSH or SCP into the remote switch without a password:

```
ssh -p <port#> userid@hostname [<remote command>]
scp -P <port#> userid@hostname/filepath /destination
```

Localtime

The Guestshell shares `/etc/localtime` with the host system.



Note If you do not want to share the same localtime with the host, this symlink can be broken and a Guestshell specific `/etc/localtime` can be created.

```
switch(config)# clock timezone PDT -7 0
switch(config)# clock set 10:00:00 27 Jan 2017
Fri Jan 27 10:00:00 PDT 2017
switch(config)# show clock
10:00:07.554 PDT Fri Jan 27 2017
switch(config)# run guestshell
guestshell:~$ date
Fri Jan 27 10:00:12 PDT 2017
```

Accessing the Guest Shell

In Cisco NX-OS, only network-admin users can access the Guest Shell by default. It is automatically enabled in the system and can be accessed using the **run guestshell** command. Consistent with the **run bash** command, these commands can be issued within the Guest Shell with the **run guestshell** *command* form of the NX-OS CLI command.



Note The Guest Shell is automatically enabled on systems with more than 4 GB of RAM.

```
switch# run guestshell ls -al /bootflash/*.ova
-rw-rw-rw- 1 2002 503 83814400 Aug 21 18:04 /bootflash/pup.ova
-rw-rw-rw- 1 2002 503 40724480 Apr 15 2012 /bootflash/red.ova
```



Note When running in the Guest Shell, you have network-admin level privileges.



Note The Guest Shell starting in 2.2(0.2) will dynamically create user accounts with the same as the user logged into switch. However, all other information is NOT shared between the switch and the Guest Shell user accounts.

In addition, the Guest Shell accounts are not automatically removed, and must be removed by the network administrator when no longer needed.

Resources Used for the Guest Shell

By default, the resources for the Guest Shell have a small impact on resources available for normal switch operations. If the network-admin requires additional resources for the Guest Shell, the **guestshell resize** *{cpu | memory | roots}* command changes these limits.

| Resource | Default | Minimum/Maximum |
|----------|---------|-----------------|
| CPU | 1% | 1/620% |
| Memory | 400 MB | 256/3840 MB |
| Storage | 200 MB | 200/2000 MB |

The CPU limit is the percentage of the system compute capacity that tasks running within the Guest Shell are given when there is contention with other compute loads in the system. When there is no contention for CPU resources, the tasks within the Guest Shell are not limited.



Note A Guest Shell reboot is required after changing the resource allocations. This can be accomplished with the **guestshell reboot** command.

Capabilities in the Guestshell

The Guestshell has a number of utilities and capabilities available by default.

The Guestshell is populated with CentOS 7 Linux which provides the ability to yum install software packages built for this distribution. The Guestshell is pre-populated with many of the common tools that would naturally be expected on a networking device including **net-tools**, **iproute**, **tcpdump** and OpenSSH. For Guestshell 2.x, python 2.7.5 is included by default as is the PIP for installing additional python packages. In Guestshell 2.11, by default, python 3.6 is also included.

By default the Guestshell is a 64-bit execution space. If 32-bit support is needed, the `glibc.i686` package can be yum installed.

The Guestshell has access to the Linux network interfaces used to represent the management and data ports of the switch. Typical Linux methods and utilities like **ifconfig** and **ethtool** can be used to collect counters. When an interface is placed into a VRF in the NX-OS CLI, the Linux network interface is placed into a network namespace for that VRF. The name spaces can be seen at `/var/run/netns` and the **ip netns** utility can be used to run in the context of different namespaces. A couple of utilities, **chvrf** and **vrinfo**, are

provided as a convenience for running in a different namespace and getting information about which namespace/vrf a process is running in.

systemd is used to manage services in CentOS 8 environments, including the Guestshell.

NX-OS CLI in the Guest Shell

The Guest Shell provides an application to allow the user to issue NX-OS commands from the Guest Shell environment to the host network element. The **dohost** application accepts any valid NX-OS configuration or exec commands and issues them to the host network element.

When invoking the **dohost** command each NX-OS command may be in single or double quotes:

```
dohost "<NXOS CLI>"
```

The NX-OS CLI can be chained together:

```
[guestshell@guestshell ~]$ dohost "sh lldp time | in Hold" "show cdp global"
Holdtime in seconds: 120
Global CDP information:
CDP enabled globally
Refresh time is 21 seconds
Hold time is 180 seconds
CDPv2 advertisements is enabled
DeviceID TLV in System-Name(Default) Format
[guestshell@guestshell ~]$
```

The NX-OS CLI can also be chained together using the NX-OS style command chaining technique by adding a semicolon between each command. (A space on either side of the semicolon is required.):

```
[guestshell@guestshell ~]$ dohost "conf t ; cdp timer 13 ; show run | inc cdp"
Enter configuration commands, one per line. End with CNTL/Z.
cdp timer 13
[guestshell@guestshell ~]$
```



Note For release 7.0(3)I5(2) using Starting with Guest Shell 2.2 (0.2), commands issued on the host through the **dohost** command are run with privileges based on the effective role of the Guest Shell user.

Prior versions of Guest Shell will run command with network-admin level privileges.

The **dohost** command fails when the number of UDS connections to NX-API are at the maximum allowed.

Network Access in Guest Shell

The NX-OS switch ports are represented in the Guest Shell as Linux network interfaces. Typical Linux methods like view stats in `/proc/net/dev`, through `ifconfig` or `ethtool` are all supported:

The Guest Shell has a number of typical network utilities included by default and they can be used on different VRFs using the **chvrf vrf command** command.

```
[guestshell@guestshell bootflash]$ ifconfig Eth1-47
Eth1-47: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
```



```
inet 13.0.0.47 netmask 255.255.255.0 broadcast 13.0.0.255
ether 54:7f:ee:8e:27:bc txqueuelen 100 (Ethernet)
RX packets 311442 bytes 21703008 (20.6 MiB)
RX errors 0 dropped 185 overruns 0 frame 0
TX packets 12967 bytes 3023575 (2.8 MiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

Within the Guest Shell, the networking state can be monitored, but may not be changed. To change networking state, use the NX-OS CLI or the appropriate Linux utilities in the host bash shell.

The **tcpdump** command is packaged with the Guest Shell to allow packet tracing of punted traffic on the management or switch ports.

The **sudo ip netns exec management ping** utility is a common method for running a command in the context of a specified network namespace. This can be done within the Guest Shell:

```
[guestshell@guestshell bootflash]$ sudo ip netns exec management ping 10.28.38.48
PING 10.28.38.48 (10.28.38.48) 56(84) bytes of data.
64 bytes from 10.28.38.48: icmp_seq=1 ttl=48 time=76.5 ms
```

The **chvrf** utility is provided as a convenience:

```
guestshell@guestshell bootflash]$ chvrf management ping 10.28.38.48
PING 10.28.38.48 (10.28.38.48) 56(84) bytes of data.
64 bytes from 10.28.38.48: icmp_seq=1 ttl=48 time=76.5 ms
```



Note Commands that are run without the **chvrf** command are run in the current VRF/network namespace.

For example, to ping IP address 10.0.0.1 over the management VRF, the command is “**chvrf management ping 10.0.0.1**”. Other utilities such as **scp** or **ssh** would be similar.

Example:

```
switch# guestshell
[guestshell@guestshell ~]$ cd /bootflash
[guestshell@guestshell bootflash]$ chvrf management scp foo@10.28.38.48:/foo/index.html
index.html
foo@10.28.38.48's password:
index.html 100% 1804 1.8KB/s 00:00
[guestshell@guestshell bootflash]$ ls -al index.html
-rw-r--r-- 1 guestshe users 1804 Sep 13 20:28 index.html
[guestshell@guestshell bootflash]$
[guestshell@guestshell bootflash]$ chvrf management curl cisco.com
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>301 Moved Permanently</title>
</head><body>
<h1>Moved Permanently</h1>
<p>The document has moved <a href="http://www.cisco.com/">here</a>.</p>
</body></html>
[guestshell@guestshell bootflash]$
```

To obtain a list of VRFs on the system, use the **show vrf** command natively from NX-OS or through the **dohost** command:

Example:

```
[guestshell@guestshell bootflash]$ dohost 'sh vrf'
VRF-Name    VRF-ID    State    Reason
default     1         Up       --
management 2         Up       --
red         6         Up       --
```

Within the Guest Shell, the network namespaces associated with the VRFs are what is actually used. It can be more convenient to just see which network namespaces are present:

```
[guestshell@guestshell bootflash]$ ls /var/run/netns
default management red
[guestshell@guestshell bootflash]$
```

To resolve domain names from within the Guest Shell, the resolver needs to be configured. Edit the `/etc/resolv.conf` file in the Guest Shell to include a DNS nameserver and domain as appropriate for the network.

Example:

```
nameserver 10.1.1.1
domain cisco.com
```

The nameserver and domain information should match what is configured through the NX-OS configuration.

Example:

```
switch(config)# ip domain-name cisco.com
switch(config)# ip name-server 10.1.1.1
switch(config)# vrf context management
switch(config-vrf)# ip domain-name cisco.com
switch(config-vrf)# ip name-server 10.1.1.1
```

If the switch is in a network that uses an HTTP proxy server, the `http_proxy` and `https_proxy` environment variables must be set up within the Guest Shell also.

Example:

```
export http_proxy=http://proxy.esl.cisco.com:8080
export https_proxy=http://proxy.esl.cisco.com:8080
```

These environment variables should be set in the `.bashrc` file or in an appropriate script to ensure that they are persistent.

Access to Bootflash in Guest Shell

Network administrators can manage files with Linux commands and utilities in addition to using NX-OS CLI commands. By mounting the system bootflash at `/bootflash` in the Guest Shell environment, the network-admin can operate on these files with Linux commands.

Example:

```
find . -name "foo.txt"
rm "/bootflash/junk/foo.txt"
```



Note While the name of the user within the Guest Shell is the same as when on the host, the Guest Shell is in a separate user namespace, and the uid does not match that of the user on the host. The file permissions for group and others will control the type of access the Guest Shell user has on the file.

Python in Guest Shell

Python can be used interactively or python scripts can be run in the Guest Shell.

Example:

```
guestshell:~$ python
Python 2.7.5 (default, Jun 24 2015, 00:41:19)
[GCC 4.8.3 20140911 (Red Hat 4.8.3-9)] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>>
guestshell:~$
```

The pip python package manager is included in the Guest Shell to allow the network-admin to install new python packages.

Example:

```
[guestshell@guestshell ~]$ sudo su
[root@guestshell guestshell]# pip install Markdown
Collecting Markdown
Downloading Markdown-2.6.2-py2.py3-none-any.whl (157kB)
100% |#####| 159kB 1.8MB/s
Installing collected packages: Markdown
Successfully installed Markdown-2.6.2
[root@guestshell guestshell]# pip list | grep Markdown
Markdown (2.6.2)
[root@guestshell guestshell]#
```



Note You must enter the **sudo su** command before entering the **pip install** command.

Python in Guestshell 2.11

Guestshell 2.11 is pre-installed with both Python 2 and Python 3.6. There is no action needed from users to install Python 2 or 3.

```
[admin@guestshell ~]$ python
Python 2.7.5 (default, Nov 16 2020, 22:23:17)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-44)] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

```
[admin@guestshell ~]$ python3
Python 3.6.8 (default, Nov 16 2020, 16:55:22)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-44)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

Python 3 in Guest Shell versions up to 2.10 (CentOS 7)

Guest Shell 2.X provides a CentOS 7.1 environment, which does not have Python 3 installed by default. There are multiple methods of installing Python 3 on CentOS 7.1, such as using third-party repositories or building from source. Another option is using the Red Hat Software Collections, which supports installing multiple versions of Python within the same system.

To install the Red Hat Software Collections (SCL) tool:

1. Install the scl-utils package.
2. Enable the CentOS SCL repository and install one of its provided Python 3 RPMs.

```
[admin@guestshell ~]$ sudo su
[root@guestshell admin]# dnf install -y scl-utils | tail
Running transaction test
Transaction test succeeded
Running transaction
  Installing : scl-utils-20130529-19.el7.x86_64                1/1
  Verifying  : scl-utils-20130529-19.el7.x86_64                1/1

Installed:
  scl-utils.x86_64 0:20130529-19.el7

Complete!

[root@guestshell admin]# dnf install -y centos-release-scl | tail
  Verifying : centos-release-scl-2-3.el7.centos.noarch        1/2
  Verifying : centos-release-scl-rh-2-3.el7.centos.noarch    2/2

Installed:
  centos-release-scl.noarch 0:2-3.el7.centos

Dependency Installed:
  centos-release-scl-rh.noarch 0:2-3.el7.centos

Complete!

[root@guestshell admin]# dnf install -y rh-python36 | tail
warning: /var/cache/dnf/x86_64/7/centos-scl-rh/packages/rh-python36-2.0-1.el7.x86_64.rpm:
Header V4 RSA/SHA1 Signature, key ID f2ee9d55: NOKEY
http://centos.sonn.com/7.7.1908/os/x86_64/Packages/groff-base-1.22.2-8.el7.x86_64.rpm:
[Errno 12] Timeout on
http://centos.sonn.com/7.7.1908/os/x86_64/Packages/groff-base-1.22.2-8.el7.x86_64.rpm: (28,
'Operation too slow. Less than 1000 bytes/sec transferred the last 30 seconds')
Trying other mirror.
Importing GPG key 0xF2EE9D55:
Userid      : "CentOS SoftwareCollections SIG
(https://wiki.centos.org/SpecialInterestGroup/SCLo) <security@centos.org>"
Fingerprint: c4db d535 blfb ba14 f8ba 64a8 4eb8 4e71 f2ee 9d55
Package     : centos-release-scl-rh-2-3.el7.centos.noarch (@extras)
From        : /etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-SIG-SCLo
rh-python36-python-libs.x86_64 0:3.6.9-2.el7
rh-python36-python-pip.noarch 0:9.0.1-2.el7
rh-python36-python-setuptools.noarch 0:36.5.0-1.el7
rh-python36-python-virtualenv.noarch 0:15.1.0-2.el7
rh-python36-runtime.x86_64 0:2.0-1.el7
scl-utils-build.x86_64 0:20130529-19.el7
xml-common.noarch 0:0.6.3-39.el7
zip.x86_64 0:3.0-11.el7

Complete!
```

Using SCL, it is possible to create an interactive bash session with Python 3's environment variables automatically setup.



Note The root user is not needed to use the SCL Python installation.

```
[admin@guestshell ~]$ scl enable rh-python36 bash
[admin@guestshell ~]$ python3
Python 3.6.9 (default, Nov 11 2019, 11:24:16)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-39)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

The Python SCL installation also provides the pip utility.

```
[admin@guestshell ~]$ pip3 install requests --user
Collecting requests
  Downloading
    https://files.pythonhosted.org/packages/51/bd/23c926cd34e8a67db02a00ba99ae0f828ce89d72b2190f27c11d4b7fb/requests-2.22.0-py2.py3-none-any.whl
    (57kB)
    100% |#####| 61kB 211kB/s
Collecting idna<2.9,>=2.5 (from requests)
  Downloading
    https://files.pythonhosted.org/packages/14/2c/cd551d81d8e15200belcf41cd03869a46fe7226e7450af7a6545bfc474c9/idna-2.8-py2.py3-none-any.whl
    (58kB)
    100% |#####| 61kB 279kB/s
Collecting chardet<3.1.0,>=3.0.2 (from requests)
  Downloading
    https://files.pythonhosted.org/packages/bc/a9/01ffeb0b562e42746487b4db1d8c7ca55ec7510b22e4c51f1409844368/chardet-3.0.4-py2.py3-none-any.whl
    (133kB)
    100% |#####| 143kB 441kB/s
Collecting certifi>=2017.4.17 (from requests)
  Downloading
    https://files.pythonhosted.org/packages/b9/63/d50cac98a05b006c55a399c3f1db9a75a24ce7890bc9cf5db9e99/certifi-2019.11.28-py2.py3-none-any.whl
    (156kB)
    100% |#####| 163kB 447kB/s
Collecting urllib3!=1.25.0,!1.25.1,<1.26,>=1.21.1 (from requests)
  Downloading
    https://files.pythonhosted.org/packages/e8/74/6e4f91745020f967d09322b2889d10090957334692b88e4afe91b77f/urllib3-1.25.8-py2.py3-none-any.whl
    (125kB)
    100% |#####| 133kB 656kB/s
Installing collected packages: idna, chardet, certifi, urllib3, requests
Successfully installed certifi-2019.11.28 chardet-3.0.4 idna-2.8 requests-2.22.0
urllib3-1.25.8
You are using pip version 9.0.1, however version 20.0.2 is available.
You should consider upgrading via the 'pip install --upgrade pip' command.
[admin@guestshell ~]$ python3
Python 3.6.9 (default, Nov 11 2019, 11:24:16)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-39)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import requests
>>> requests.get("https://cisco.com")
<Response [200]>
```

The default Python 2 installation can be used alongside the SCL Python installation.

```
[admin@guestshell ~]$ which python3
/opt/rh/rh-python36/root/usr/bin/python3
[admin@guestshell ~]$ which python2
/bin/python2
[admin@guestshell ~]$ python2
Python 2.7.5 (default, Aug 7 2019, 00:51:29)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-39)] on linux2
```

```
Type "help", "copyright", "credits" or "license" for more information.
>>> print 'Hello world!'
Hello world!
```

Software Collections makes it possible to install multiple versions of the same RPM on a system. In this case, it is possible to install Python 3.5 in addition to Python 3.6.

```
[admin@guestshell ~]$ sudo dnf install -y rh-python35 | tail
Dependency Installed:
  rh-python35-python.x86_64 0:3.5.1-13.e17
  rh-python35-python-devel.x86_64 0:3.5.1-13.e17
  rh-python35-python-libs.x86_64 0:3.5.1-13.e17
  rh-python35-python-pip.noarch 0:7.1.0-2.e17
  rh-python35-python-setuptools.noarch 0:18.0.1-2.e17
  rh-python35-python-virtualenv.noarch 0:13.1.2-2.e17
  rh-python35-runtime.x86_64 0:2.0-2.e17
```

Complete!

```
[admin@guestshell ~]$ scl enable rh-python35 python3
Python 3.5.1 (default, May 29 2019, 15:41:33)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-36)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>>
```



Note Creating new interactive bash sessions when multiple Python versions are installed in SCL can cause an issue where the libpython shared object file cannot be loaded. There is a workaround where you can use the **source scl_source enable python-installation** command to properly set up the environment in the current bash session.

The default Guest Shell storage capacity is not sufficient to install Python 3. Use the **guestshell resize rootfs size-in-MB** command to increase the size of the file system. Typically, setting the rootfs size to 550 MB is sufficient.

Python in Guestshell 4.x

Python2 is deprecated, hence will not be available in Guestshell 4.x.

Guestshell 4.x will support python3.9 as default python version.

```
[admin@guestshell ~]$ python
Python 3.9.16 (main, Dec 8 2022, 00:00:00)
[GCC 11.3.1 20221121 (Red Hat 11.3.1-4)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>>
>>>
```

Python 2 in Guest Shell 3.0 (Centos 8)

Guestshell 3.0 provides a Centos 8 environment, which will have python 3 installed by default. Python 2 is reaching end of support and is not part of Guest Shell 3.0.

To provide downward compatibility for users with python 2 based applications, python 2 can be installed as shown in the following example:

```
[admin@guestshell ~]$ sudo su
[root@guestshell admin]# export http_proxy=http://proxy.esl.cisco.com:8080
[root@guestshell admin]# export https_proxy=http://proxy.esl.cisco.com:8080
```

```
[root@guestshell admin]# ip netns exec management yum install -y python2
6/6

Installed:
python2-2.7.17-2.module_el8.3.0+478+7570e00c.x86_64
python2-libs-2.7.17-2.module_el8.3.0+478+7570e00c.x86_64
python2-pip-9.0.3-18.module_el8.3.0+478+7570e00c.noarch
python2-pip-wheel-9.0.3-18.module_el8.3.0+478+7570e00c.noarch
python2-setuptools-39.0.1-12.module_el8.3.0+478+7570e00c.noarch
python2-setuptools-wheel-39.0.1-12.module_el8.3.0+478+7570e00c.noarch

Complete!

[root@guestshell admin]# which python2
/bin/python2
[root@guestshell admin]# python2 --version
Python 2.7.17
[root@guestshell admin]# python2
Python 2.7.17 (default, Aug 31 2020, 21:02:14)
[GCC 8.3.1 20191121 (Red Hat 8.3.1-5)] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>> print "Hello World"
Hello World
>>> exit()
```

Installing RPMs in the Guest Shell

The `/etc/dnf/repos.d/CentOS-Base.repo` file is set up to use the CentOS mirror list by default. Follow instructions in that file if changes are needed.

Dnf can be pointed to one or more repositories at any time by modifying the `yumrepo_x86_64.repo` file or by adding a new `.repo` file in the `repos.d` directory.

For applications to be installed inside Guest Shell 3.0, go to the CentOS 8 repo at http://mirror.centos.org/centos/8/BaseOS/x86_64/os/Packages/.

For applications to be installed inside Guest Shell 2.x, go to the CentOS 7 repo at http://mirror.centos.org/centos/7/os/x86_64/Packages/.

For applications to be installed inside Guest Shell 4.x, go to the RockyLinux 9 repo at https://mirrors.rockylinux.org/mirrorlist?arch=x86_64&repo=rocky-BaseOS-9.2. Choose any one of mirror link and view the packages.

Dnf resolves the dependencies and installs all the required packages.

```
[guestshell@guestshell ~]$ sudo chvrf management dnf -y install glibc.i686
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
* base: bay.uchicago.edu
* extras: pubmirrors.dal.corespace.com
* updates: mirrors.cmich.edu
Resolving Dependencies
"--> Running transaction check
"----> Package glibc.i686 0:2.17-78.el7 will be installed
"--> Processing Dependency: libfreebl3.so(NSSRAWHASH_3.12.3) for package:
glibc-2.17-78.el7.i686
"----> Processing Dependency: libfreebl3.so for package: glibc-2.17-78.el7.i686
"--> Running transaction check
"----> Package nss-softokn-freebl.i686 0:3.16.2.3-9.el7 will be installed
"--> Finished Dependency Resolution
```

Dependencies Resolved

Package Arch Version Repository Size

Installing:

glibc i686 2.17-78.el7 base 4.2 M

Installing for dependencies:

nss-softokn-freebl i686 3.16.2.3-9.el7 base 187 k

Transaction Summary

Install 1 Package (+1 Dependent package)

Total download size: 4.4 M

Installed size: 15 M

Downloading packages:

Delta RPMs disabled because /usr/bin/applydeltarpm not installed.

(1/2): nss-softokn-freebl-3.16.2.3-9.el7.i686.rpm | 187 kB 00:00:25

(2/2): glibc-2.17-78.el7.i686.rpm | 4.2 MB 00:00:30

Total 145 kB/s | 4.4 MB 00:00:30

Running transaction check

Running transaction test

Transaction test succeeded

Running transaction

Installing : nss-softokn-freebl-3.16.2.3-9.el7.i686 1/2

Installing : glibc-2.17-78.el7.i686 2/2

error: lua script failed: [string "%triggerin(glibc-common-2.17-78.el7.x86_64)":1: attempt to compare number with nil

Non-fatal "<"unknown">" scriptlet failure in rpm package glibc-2.17-78.el7.i686

Verifying : glibc-2.17-78.el7.i686 1/2

Verifying : nss-softokn-freebl-3.16.2.3-9.el7.i686 2/2

Installed:

glibc.i686 0:2.17-78.el7

Dependency Installed:

nss-softokn-freebl.i686 0:3.16.2.3-9.el7

Complete!



Note When more space is needed in the Guest Shell root file system for installing or running packages, the **guestshell resize roots *size-in-MB*** command is used to increase the size of the file system.



Note Some open source software packages from the repository might not install or run as expected in the Guest Shell as a result of restrictions that have been put into place to protect the integrity of the host system.

Security Posture for Virtual ServicesGuest Shell

Use of the Guest Shell and virtual services in switches are only two of the many ways that the network-admin can manage or extend the functionality of the system. These options are geared toward providing an execution environment that is decoupled from the native host context. This separation allows the introduction of software

into the system that may not be compatible with the native execution environment. It also allows the software to run in an environment that does not interfere with the behavior, performance, or scale of the system.

Use of the Guest Shell in switches is just one of the many ways the network admin can manage or extend the functionality of the system. The Guest Shell is intended to provide an execution environment that is decoupled from the native host context. This separation allows the introduction of software into the system that may not be compatible with the native execution environment. It also allows the software to run in an environment that does not interfere with the behavior, performance, or scale of the system.

Kernel Vulnerability Patches

Cisco responds to pertinent Common Vulnerabilities and Exposures (CVEs) with platform updates that address known vulnerabilities.



Note Cisco tracks the vulnerabilities for Guestshell 4.x (Rocky Linux 9) environment and will include future fixes when they are available from Rocky Linux.

ASLR and X-Space Support

Cisco 30009000 NX-OS supports the use of Address Space Layout Randomization (ASLR) and Executable Space Protection (X-Space) for runtime defense. The software in Cisco-signed packages make use of this capability. If other software is installed on the system, it is recommended that it be built using a host OS and development toolchain that supports these technologies. Doing so reduces the potential attack surface that the software presents to potential intruders.

Namespace Isolation

The Guest Shell environment runs within a Linux container that makes use of various namespaces to decouple the Guest Shell execution space from that of the host. Starting in the NX-OS 9.2(1) release, the Guest Shell is run in a separate user namespace, which helps protect the integrity of the host system, as processes running as root within the Guest Shell are not root of the host. These processes appear to be running as uid 0 within the Guest Shell due to uid mapping, but the kernel knows the real uid of these processes and evaluates the POSIX capabilities within the appropriate user namespace.

When a user enters the Guest Shell from the host, a user of the same name is created within the Guest Shell. While the names match, the uid of the user within the Guest Shell is not the same as the uid on the host. To still allow users within the Guest Shell to access files on shared media (for example, `/bootflash` or `/volatile`), the common NX-OS gids used on the host (for example, `network-admin` or `network-operator`) are mapped into the Guest Shell such that the values are the same and the Guest Shell instance of the user is associated with the appropriate groups based on group membership on the host.

As an example, consider user `bob`. On the host, `bob` has the following uid and gid membership:

```
bash-4.3$ id
uid=2004(bob) gid=503(network-admin) groups=503(network-admin),504(network-operator)
```

When user `bob` is in the Guest Shell, the group membership from the host is set up in the Guest Shell:

```
[bob@guestshell ~]$ id
uid=1002(bob) gid=503(network-admin)
groups=503(network-admin),504(network-operator),10(wheel)
```

Files created by user `bob` in the host Bash shell and the Guest Shell have different owner ids. The example output below shows that the file created from within the Guest Shell has owner id 12002, instead of 1002 as shown in the example output above. This is due to the command being issued from the host Bash shell and the id space for the Guest Shell starting at id 11000. The group id of the file is `network-admin`, which is 503 in both environments.

```
bash-4.3$ ls -ln /bootflash/bob_*
-rw-rw-r-- 1 12002 503 4 Jun 22 15:47 /bootflash/bob_guestshell
-rw-rw-r-- 1 2004 503 4 Jun 22 15:47 /bootflash/bob_host
```

```
bash-4.3$ ls -l /bootflash/bob_*
-rw-rw-r-- 1 12002 network-admin 4 Jun 22 15:47 /bootflash/bob_guestshell
-rw-rw-r-- 1 bob network-admin 4 Jun 22 15:47 /bootflash/bob_host
```

The user is allowed to access the file due to the file permission settings for the `network-admin` group, and the fact that `bob` is a member of `network-admin` in both the host and Guest Shell.

Inside the Guest Shell environment, the example output below shows that the owner id for the file created by `bob` from the host is 65534. This indicates the actual id is in a range that is outside range of ids mapped into the user namespace. Any unmapped id will be shown as this value.

```
[bob@guestshell ~]$ ls -ln /bootflash/bob_*
-rw-rw-r-- 1 1002 503 4 Jun 22 15:47 /bootflash/bob_guestshell
-rw-rw-r-- 1 65534 503 4 Jun 22 15:47 /bootflash/bob_host
```

```
[bob@guestshell ~]$ ls -l /bootflash/bob_*
-rw-rw-r-- 1 bob network-admin 4 Jun 22 15:47 /bootflash/bob_guestshell
-rw-rw-r-- 1 65534 network-admin 4 Jun 22 15:47 /bootflash/bob_host
```

Root-User Restrictions

As a best practice for developing secure code, it is recommend running applications with the least privilege needed to accomplish the assigned task. To help prevent unintended accesses, software added into the Guest Shell should follow this best practice.

All processes within a virtual servicethe Guest Shell are subject to restrictions imposed by reduced Linux capabilities. If your application must perform operations that require root privileges, restrict the use of the root account to the smallest set of operations that absolutely requires root access, and impose other controls such as a hard limit on the amount of time that the application can run in that mode.

The set of Linux capabilities that are dropped for root within virtual servicethe Guest Shell follow:

| | | |
|-------------------|------------------|------------------|
| CAP_SYS_BOOT | CAP_MKNOD | CAP_SYS_PACCT |
| CAP_SYS_MODULE | CAP_MAC_OVERRIDE | CAP_SYS_RESOURCE |
| CAP_SYS_TIME | CAP_SYS_RAWIO | CAP_AUDIT_WRITE |
| CAP_AUDIT_CONTROL | CAP_SYS_NICE | CAP_NET_ADMIN |
| CAP_MAC_ADMIN | CAP_SYS_PTRACE | |

- `cap_audit_control`
- `cap_audit_write`
- `cap_mac_admin`
- `cap_mac_override`
- `cap_mknod`
- `cap_net_broadcast`
- `cap_sys_boot`
- `cap_syslog`
- `cap_sys_module`
- `cap_sys_nice`
- `cap_sys_pacct`
- `cap_sys_ptrace`
- `cap_sys_rawio`
- `cap_sys_resource`
- `cap_sys_time`
- `cap_wake_alarm`

As root within a virtual service, bind mounts may be used as well as tmpfs and ramfs mounts. Other mounts are prevented.

While the `net_admin` capability is not dropped, user namespace and the host ownership of the network namespaces prevents the Guest Shell user from modifying the interface state. As root within the Guest Shell, bind mounts may be used as well as tmpfs and ramfs mounts. Other mounts are prevented.

Resource Management

A Denial-of-Service (DoS) attack attempts to make a machine or network resource unavailable to its intended users. Misbehaving or malicious application code can cause DoS as the result of over-consumption of connection bandwidth, disk space, memory, and other resources. The host provides resource-management features that ensure fair allocation of resources among all virtual services between Guest Shell and services on the host.

Guest File System Access Restrictions

To preserve the integrity of the files within the virtual services, the file systems of the virtual services are not accessible from the NX-OS CLI. If a given virtual-service allows files to be modified, it needs to provide an alternate means by which this can be done (i.e. **yum install**, **scp**, **ftp**, etc).

To preserve the integrity of the files within the Guest Shell, the file systems of the Guest Shell are not accessible from the NX-OS CLI.

The Guest Shell mounts the `bootflash` of the host system at `/bootflash`. The network-admin can access the file using an NX-OS CLI or Linux command from within the Guest Shell.

`bootflash:` and `volatile:` of the host are mounted as `/bootflash` and `/volatile` within the Guest Shell. A network-admin can access files on this media using the NX-OS `exec` commands from the host or using Linux commands from within the Guest Shell.

Managing the Guest Shell

The following are commands to manage the Guest Shell:

Table 3: Guest Shell CLI Commands

| Commands | Description |
|---|--|
| guestshell enable { package [<i>guest shell OVA file</i> <i>rootfs-file-URI</i>]} | <ul style="list-style-type: none"> When <i>guest shell OVA file</i> is specified: Installs and activates the Guest Shell using the OVA that is embedded in the system image. Installs and activates the Guest Shell using the specified software package (OVA file) or the embedded package from the system image (when no package is specified). Initially, Guest Shell packages are only available by being embedded in the system image. When the Guest Shell is already installed, this command enables the installed Guest Shell. Typically this is used after a guestshell disable command. When <i>rootfs-file-URI</i> is specified: Imports a Guest Shell rootfs when the Guest Shell is in a destroyed state. This command brings up the Guest Shell with the specified package. |
| guestshell export rootfs package <i>destination-file-URI</i> | Exports a Guest Shell rootfs file to a local URI (bootflash, USB1, etc.). (7.0(3)I7(1) and later releases) |
| guestshell disable | Shuts down and disables the Guest Shell. |

| Commands | Description |
|--|---|
| <p>guestshell upgrade {package [<i>guest shell OVA file</i> <i>rootfs-file-URI</i>]}</p> | <ul style="list-style-type: none"> • When <i>guest shell OVA file</i> is specified: <p>Deactivates and upgrades the Guest Shell using the specified software package (OVA file) or the embedded package from the system image (if no package is specified). Initially Guest Shell packages are only available by being embedded in the system image.</p> <p>The current rootfs for the Guest Shell is replaced with the rootfs in the software package. The Guest Shell does not make use of secondary filesystems that persist across an upgrade. Without persistent secondary filesystems, a guestshell destroy command followed by a guestshell enable command could also be used to replace the rootfs. When an upgrade is successful, the Guest Shell is activated.</p> <p>You are prompted for a confirmation prior to carrying out the upgrade command.</p> • When <i>rootfs-file-URI</i> is specified: <p>Imports a Guest Shell rootfs file when the Guest Shell is already installed. This command removes the existing Guest Shell and installs the specified package.</p> |
| <p>guestshell reboot</p> | <p>Deactivates the Guest Shell and then reactivates it. You are prompted for a confirmation prior to carrying out the reboot command.</p> <p>Note This is the equivalent of a guestshell disable command followed by a guestshell enable command in exec mode.</p> <p>This is useful when processes inside the Guest Shell have been stopped and need to be restarted. The run guestshell command relies on <code>sshd</code> running in the Guest Shell.</p> <p>If the command does not work, the <code>sshd</code> process may have been inadvertently stopped. Performing a reboot of the Guest Shell from the NX-OS CLI allows it to restart and restore the command.</p> |

| Commands | Description |
|--|---|
| guestshell destroy | <p>Deactivates and uninstalls the Guest Shell. All resources associated with the Guest Shell are returned to the system. The show virtual-service global command indicates when these resources become available.</p> <p>Issuing this command results in a prompt for a confirmation prior to carrying out the destroy command.</p> |
| guestshell run guestshell | Connects to the Guest Shell that is already running with a shell prompt. No username/password is required. |
| guestshell run <i>command</i> run guestshell <i>command</i> | <p>Executes a Linux/UNIX command within the context of the Guest Shell environment.</p> <p>After execution of the command you are returned to the switch prompt.</p> |
| guestshell resize [cpu memory rootfs] | <p>Changes the allotted resources available for the Guest Shell. The changes take effect the next time the Guest Shell is enabled or rebooted.</p> <p>Note Resize values are cleared when the guestshell destroy command is used.</p> |
| guestshell sync | On systems that have active and standby supervisors, this command synchronizes the Guest Shell contents from the active supervisor to the standby supervisor. The network-admin issues this command when the Guest Shell rootfs has been set up to a point that they would want the same rootfs used on the standby supervisor when it becomes the active supervisor. If this command is not used, the Guest Shell is freshly installed when the standby supervisor transitions to an active role using the Guest Shell package available on that supervisor. |
| virtual-service reset force | <p>In the event that the guestshell or virtual-services cannot be managed, even after a system reload, the reset command is used to force the removal of the Guest Shell and all virtual-services. The system needs to be reloaded for the cleanup to happen. No Guest Shell or additional virtual-services can be installed or enabled after issuing this command until after the system has been reloaded.</p> <p>You are prompted for a confirmation prior to initiating the reset.</p> |



Note Administrative privileges are necessary to enable/disable and to gain access to the Guest Shell environment.



Note The Guest Shell is implemented as a Linux container (LXC) on the host system. On NX-OS devices, LXC's are installed and managed with the virtual-service commands. The Guest Shell appears in the virtual-service commands as a virtual service named `guestshell+`.



Note Virtual-service commands that do not pertain to the Guest Shell are being deprecated. These commands have been hidden in the NX-OS 9.2(1) release and will be removed in future releases.

The following exec keywords are being deprecated:

```
# virtual-service ?
connect  Request a virtual service shell
install  Add a virtual service to install database
uninstall Remove a virtual service from the install database
upgrade  Upgrade a virtual service package to a different version
```

```
# show virtual-service ?
detail  Detailed information config)
```

The following config keywords are being deprecated:

```
(config) virtual-service ?
WORD    Virtual service name (Max Size 20)
```

```
(config-virt-serv)# ?
activate Activate configured virtual service
description Virtual service description
```

Disabling the Guest Shell

The `guestshell disable` command shuts down and disables the Guest Shell.

When the Guest Shell is disabled and the system is reloaded, the Guest Shell remains disabled.

Example:

```
switch# show virtual-service list
Virtual Service List:
Name                Status           Package Name
-----
guestshell+        Activated        guestshell.ova
switch# guestshell disable
You will not be able to access your guest shell if it is disabled. Are you sure you want
to disable the guest shell? (y/n) [n] y

2014 Jul 30 19:47:23 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Deactivating virtual
service 'guestshell+'
```

```

2014 Jul 30 18:47:29 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Successfully deactivated
virtual service 'guestshell+'
switch# show virtual-service list
Virtual Service List:
Name                Status                Package Name
guestshell+         Deactivated           guestshell.ova

```



Note The Guest Shell is reactivated with the **guestshell enable** command.

Destroying the Guest Shell

The **guestshell destroy** command uninstalls the Guest Shell and its artifacts. The command does not remove the Guest Shell OVA.

When the Guest Shell is destroyed and the system is reloaded, the Guest Shell remains destroyed.

```

switch# show virtual-service list
Virtual Service List:
Name                Status                Package Name
-----
guestshell+         Deactivated           guestshell.ova

```

```
switch# guestshell destroy
```

```

You are about to destroy the guest shell and all of its contents. Be sure to save your work.
Are you sure you want to continue? (y/n) [n] y
2014 Jul 30 18:49:10 switch %$ VDC-1 %$ %VMAN-2-INSTALL_STATE: Destroying virtual service
'guestshell+'
2014 Jul 30 18:49:10 switch %$ VDC-1 %$ %VMAN-2-INSTALL_STATE: Successfully destroyed
virtual service 'guestshell +'

```

```

switch# show virtual-service list
Virtual Service List:

```



Note The Guest Shell can be re-enabled with the **guestshell enable** command.



Note If you do not want to use the Guest Shell, you can remove it with the **guestshell destroy** command. Once the Guest Shell has been removed, it remains removed for subsequent reloads. This means that when the Guest Shell container has been removed and the switch is reloaded, the Guest Shell container is not automatically started.

Enabling the Guest Shell

The **guestshell enable** command installs the Guest Shell from a Guest Shell software package. By default, the package embedded in the system image is used for the installation. The command is also used to reactivate the Guest Shell if it has been disabled.

When the Guest Shell is enabled and the system is reloaded, the Guest Shell remains enabled.

Example:

```
switch# show virtual-service list
Virtual Service List:
switch# guestshell enable
2014 Jul 30 18:50:27 switch %$ VDC-1 %$ %VMAN-2-INSTALL_STATE: Installing virtual service
'guestshell+'
2014 Jul 30 18:50:42 switch %$ VDC-1 %$ %VMAN-2-INSTALL_STATE: Install success virtual
service 'guestshell+'; Activating
2014 Jul 30 18:50:42 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Activating virtual service
'guestshell+'
2014 Jul 30 18:51:16 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Successfully activated
virtual service 'guestshell+'

switch# show virtual-service list
Virtual Service List:
Name                Status              Package Name
guestshell+         Activated           guestshell.ova
```

Enabling the Guest Shell in Base Boot Mode

Beginning in the NX-OS 9.2(1) release, you can choose to boot your system in *base boot mode*. When you boot your system in base boot mode, the Guest Shell is not started by default. In order to use the Guest Shell in this mode, you must activate the RPMs containing the virtualization infrastructure as well as the Guest Shell image. Once you have done this, the Guest Shell and virtual-service commands will be available.

If the RPM activation commands are run in this order:

1. `install activate guestshell`
2. `install activate virtualization`

The Guest Shell container will be activated automatically as it would have been if the system had been booted in full mode.

If the RPM activation commands are run in the reverse order:

1. `install activate virtualization`
2. `install activate guestshell`

Then the Guest Shell will not be enabled until you run the **guestshell enable** command.

Enabling the Guest Shell on Cisco Nexus 3000 with Compacted Image

The Guest Shell software is not available in a Cisco NX-OS image that has been compacted for the Cisco Nexus 3000 Series switches with 1.6 GB bootflash and 4 GB RAM. You can still use the Guest Shell in this case, but you will need to download the software package from software.cisco.com for the Cisco NX-OS release, then you will need to copy it onto the Cisco Nexus 3000 Series switch and enable it.

For more information on compacted images, refer to the *Cisco Nexus 3000 Series NX-OS Software Upgrade and Downgrade Guide, Release 9.2(1)*.

The Guest Shell software installs onto the bootflash of the switch. To create as much free bootflash space as possible, put the downloaded `guestshell.ova` file onto the `volatile:` storage media. Once the Guest

Shell is successfully activated, the `guestshell.ova` file can be deleted. It will not be needed again unless the Guest Shell is destroyed at some point and needs to be re-installed.

For example:

```
switch# copy scp://admin@1.2.3.4/guestshell.ova volatile: vrf management
guestshell.ova 100% 55MB 10.9MB/s 00:05
Copy complete, now saving to disk (please wait)...
Copy complete.

switch# dir volatile: | inc .ova
57251840 Jun 22 11:56:51 2018 guestshell.ova

switch# guestshell enable package volatile:guestshell.ova
2018 Jun 7 19:13:03 n3x-164 %$ VDC-1 %$ %VMAN-2-INSTALL_STATE: Installing virtual service
'guestshell+'
2018 Jun 7 19:13:56 n3x-164 %$ VDC-1 %$ %VMAN-2-INSTALL_STATE: Install success virtual
service 'guestshell+'; Activating
2018 Jun 7 19:13:56 n3x-164 %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Activating virtual service
'guestshell+'
2018 Jun 7 19:15:34 n3x-164 %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Successfully activated
virtual service 'guestshell+'

switch# del volatile:guestshell.ova
Do you want to delete "/guestshell.ova" ? (yes/no/abort) [y] y

switch# guestshell
[admin@guestshell ~]$_
```

Replicating the Guest Shell

Beginning with Cisco NX-OS release 7.0(3)I7(1), a Guest Shell **rootfs** that is customized on one switch can be deployed onto multiple switches.

The approach is to customize and then export the Guest Shell **rootfs** and store it on a file server. A POAP script can download (import) the Guest Shell **rootfs** to other switches and install the specific Guest Shell across many devices simultaneously.

Exporting Guest Shell **rootfs**

Use the `guestshell export rootfs package destination-file-URI` command to export a Guest Shell **rootfs**.

The `destination-file-URI` parameter is the name of the file that the Guest Shell **rootfs** is copied to. This file allows for local URI options (bootflash, USB1, etc.).

The `guestshell export rootfs package` command:

- Disables the Guest Shell (if already enabled).
- Creates a Guest Shell import YAML file and inserts it into the `/cisco` directory of the **rootfs** ext4 file.
- Copies the **rootfs** ext4 file to the target URI location.
- Re-enables the Guest Shell if it had been previously enabled.

Importing Guest Shell **rootfs**

When importing a Guest Shell **rootfs**, there are two situations to consider:

- Use the **guestshell enable package** *rootfs-file-URI* command to import a Guest Shell **rootfs** when the Guest Shell is in a destroyed state. This command brings up the Guest Shell with the specified package.
- Use the **guestshell upgrade package** *rootfs-file-URI* command to import a Guest Shell **rootfs** when the Guest Shell is already installed. This command removes the existing Guest Shell and installs the specified package.

The *rootfs-file-URI* parameter is the **rootfs** file stored on local storage (bootflash, USB, etc.).

When this command is executed with a file that is on bootflash, the file is moved to a storage pool on bootflash.

As a best practice, you should copy the file to the bootflash and validate the md5sum before using the **guestshell upgrade package** *rootfs-file-URI* command.



Note The **guestshell upgrade package** *rootfs-file-URI* command can be executed from within the Guest Shell.



Note The **rootfs** file is not a Cisco signed package, you must configure to allow unsigned packages before enabling as shown in the example:

```
(config-virt-serv-global)# signing level unsigned
Note: Support for unsigned packages has been user-enabled. Unsigned packages are not endorsed
by Cisco. User assumes all responsibility.
```



Note To restore the embedded version of the **rootfs**:

- Use the **guestshell upgrade** command (without additional parameters) when the Guest Shell has already been installed.
 - Use the **guestshell enable** command (without additional parameters) when the Guest Shell had been destroyed.
-



Note When running this command from within a Guest Shell, or outside a switch using NX-API, you must set **terminal dont-ask** to skip any prompts.

The **guestshell enable package** *rootfs-file-URI* command:

- Performs basic validation of the **rootfs** file.
- Moves the **rootfs** into the storage pool.
- Mounts the **rootfs** to extract the YAML file from the `/cisco` directory.
- Parses the YAML file to obtain VM definition (including resource requirements).
- Activates the Guest Shell.

Example workflow for **guestshell enable** :

```
switch# copy scp://user@10.1.1.1/my_storage/gs_rootfs.ext4 bootflash: vrf management
switch# guestshell resize cpu 8
Note: System CPU share will be resized on Guest shell enable
switch# guestshell enable package bootflash:gs_rootfs.ext4
Validating the provided rootfs
switch# 2017 Jul 31 14:58:01 switch %$ VDC-1 %$ %VMAN-2-INSTALL_STATE: Installing virtual
service 'guestshell+'
2017 Jul 31 14:58:09 switch %$ VDC-1 %$ %VMAN-2-INSTALL_STATE: Install success virtual
service 'guestshell+'; Activating
2017 Jul 31 14:58:09 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Activating virtual service
'guestshell+'
2017 Jul 31 14:58:33 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Successfully activated
virtual service 'guestshell+'
```



Note Workflow for **guestshell upgrade** is preceded by the existing Guest Shell being destroyed.



Note Resize values are cleared when the **guestshell upgrade** command is used.

Importing YAML File

A YAML file that defines some user modifiable characteristics of the Guest Shell is automatically created as a part of the export operation. It is embedded into the Guest Shell **rootfs** in the /cisco directory. It is not a complete descriptor for the Guest Shell container. It only contains some of the parameters that are user modifiable.

Example of a Guest Shell import YAML file:

```
---
import-schema-version: "1.0"
info:
  name: "GuestShell"
  version: "2.2(0.3)"
  description: "Exported GuestShell: 20170216T175137Z"
app:
  apptype: "lxc"
  cpuarch: "x86_64"
  resources:
    cpu: 3
    memory: 307200
    disk:
      - target-dir: "/"
        capacity: 250
...
```

The YAML file is generated when the **guestshell export rootfs package** command is executed. The file captures the values of the currently running Guest Shell.

The info section contains non-operational data that is used to help identify the Guest Shell. Some of the information will be displayed in the output of the **show guestshell detail** command.

The description value is an encoding of the UTC time when the YAML file was created. The time string format is the same as DTSTAMP in RFC5545 (iCal).

The resources section describes the resources required for hosting the Guest Shell. The value "/" for the target-dir in the example identifies the disk as the **rootfs**.



Note If resized values were specified while the Guest Shell was destroyed, those values take precedence over the values in the import YAML file when the **guestshell enable package** command is used.

The cpuarch value indicates the CPU architecture that is expected for the container to run.

You can modify the YAML file (such as the description or increase the resource parameters, if appropriate) after the export operation is complete .

Cisco provides a python script that you can run to validate a modified YAML file with a JSON schema. It is not meant to be a complete test (for example, device-specific resource limits are not checked), but it is able to flag common errors. The python script with examples is located at [Guest Shell Import Export](#). The following JSON file describes the schema for version 1.0 of the Guest Shell import YAML .

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "title": "Guest Shell import schema",
  "description": "Schema for Guest Shell import descriptor file - ver 1.0",
  "copyright": "2017 by Cisco systems, Inc. All rights reserved.",
  "id": "",
  "type": "object",
  "additionalProperties": false,
  "properties": {
    "import-schema-version": {
      "id": "/import-schema-version",
      "type": "string",
      "minLength": 1,
      "maxLength": 20,
      "enum": [
        "1.0"
      ]
    },
    "info": {
      "id": "/info",
      "type": "object",
      "additionalProperties": false,
      "properties": {
        "name": {
          "id": "/info/name",
          "type": "string",
          "minLength": 1,
          "maxLength": 29
        },
        "description": {
          "id": "/info/description",
          "type": "string",
          "minLength": 1,
          "maxLength": 199
        },
        "version": {
          "id": "/info/version",
          "type": "string",
          "minLength": 1,
          "maxLength": 63
        },
        "author-name": {
          "id": "/info/author-name",
```

```

        "type": "string",
        "minLength": 1,
        "maxLength": 199
      },
      "author-link": {
        "id": "/info/author-link",
        "type": "string",
        "minLength": 1,
        "maxLength": 199
      }
    }
  },
  "app": {
    "id": "/app",
    "type": "object",
    "additionalProperties": false,
    "properties": {
      "apptype": {
        "id": "/app/apptype",
        "type": "string",
        "minLength": 1,
        "maxLength": 63,
        "enum": [
          "lxc"
        ]
      },
      "cpuarch": {
        "id": "/app/cpuarch",
        "type": "string",
        "minLength": 1,
        "maxLength": 63,
        "enum": [
          "x86_64"
        ]
      }
    }
  },
  "resources": {
    "id": "/app/resources",
    "type": "object",
    "additionalProperties": false,
    "properties": {
      "cpu": {
        "id": "/app/resources/cpu",
        "type": "integer",
        "multipleOf": 1,
        "maximum": 100,
        "minimum": 1
      },
      "memory": {
        "id": "/app/resources/memory",
        "type": "integer",
        "multipleOf": 1024,
        "minimum": 1024
      },
      "disk": {
        "id": "/app/resources/disk",
        "type": "array",
        "minItems": 1,
        "maxItems": 1,
        "uniqueItems": true,
        "items": {
          "id": "/app/resources/disk/0",
          "type": "object",
          "additionalProperties": false,
          "properties": {

```

```

        "target-dir": {
          "id": "/app/resources/disk/0/target-dir",
          "type": "string",
          "minLength": 1,
          "maxLength": 1,
          "enum": [
            "/"
          ]
        },
        "file": {
          "id": "/app/resources/disk/0/file",
          "type": "string",
          "minLength": 1,
          "maxLength": 63
        },
        "capacity": {
          "id": "/app/resources/disk/0/capacity",
          "type": "integer",
          "multipleOf": 1,
          "minimum": 1
        }
      }
    }
  },
  "required": [
    "memory",
    "disk"
  ]
}
},
"required": [
  "apptype",
  "cpuarch",
  "resources"
]
}
},
"required": [
  "app"
]
}
}

```

show guestshell Command

The output of the **show guestshell detail** command includes information that indicates whether the Guest Shell was imported or was installed from an OVA.

Example of the **show guestshell detail** command after importing **rootfs**.

```

switch# show guestshell detail
Virtual service guestshell+ detail
  State           : Activated
  Package information
    Name          : rootfs_puppet
    Path          : usb2:/rootfs_puppet
  Application
    Name          : GuestShell
    Installed version : 3.0(0.0)
    Description    : Exported GuestShell: 20170613T173648Z
  Signing
    Key type      : Unsigned
    Method        : Unknown

```

```
Licensing
  Name       : None
  Version    : None
```

Verifying Virtual Service and Guest Shell Information

You can verify virtual service and Guest Shell information with the following commands:

| Command | Description |
|--|--|
| <pre>show virtual-service global switch# show virtual-service global Virtual Service Global State and Virtualization Limits: Infrastructure version : 1.11 Total virtual services installed : 1 Total virtual services activated : 1 Machine types supported : LXC Machine types disabled : KVM Maximum VCPUs per virtual service : 1 Resource virtualization limits: Name Quota Committed Available ----- system CPU (%) 20 1 19 memory (MB) 3840 256 3584 bootflash (MB) 8192 200 7992 switch#</pre> | Displays the global state and limits for virtual services. |
| <pre>show virtual-service list switch# show virtual-service list * Virtual Service List: Name Status Package Name ----- guestshell+ Activated guestshell.ova chef Installed chef-0.8.1-n9000-spa-k9.ova</pre> | Displays a summary of the virtual services, the status of the virtual services, and installed software packages. |

| Command | Description |
|--|---|
| <pre> show guestshell detail switch# show guestshell detail Virtual service guestshell+ detail State : Activated Package information Name : guestshell.ova Path : /isan/bin/guestshell.ova Application Name : GuestShell Installed version : 3.0(0.0) Description : Cisco Systems Guest Shell Signing Key type : Cisco key Method : SHA-1 Licensing Name : None Version : None Resource reservation Disk : 400 MB Memory : 256 MB CPU : 1% system CPU Attached devices Type Name Alias ----- Disk _rootfs Disk /cisco/core Serial/shell Serial/aux Serial/Syslog serial2 Serial/Trace serial3 </pre> | <p>Displays details about the guestshell package (such as version, signing resources, and devices).</p> |

Persistently Starting Your Application From the Guest Shell

Your application should have a `systemd / systemctl` service file that gets installed in `/usr/lib/systemd/system/application_name.service`. This service file should have the following general format:

```

[Unit]
Description=Put a short description of your application here

[Service]
ExecStart=Put the command to start your application here
Restart=always
RestartSec=10s

[Install]
WantedBy=multi-user.target

```



Note To run `systemd` as a specific user, add `User=<username>` to the `[Service]` section of your service.

Procedure for Persistently Starting Your Application from the Guest Shell

Procedure

-
- Step 1** Install your application service file that you created above into `/usr/lib/systemd/system/application_name.service`
 - Step 2** Start your application with `systemctl start application_name`
 - Step 3** Verify that your application is running with `systemctl status -l application_name`
 - Step 4** Enable your application to be restarted on reload with `systemctl enable application_name`
 - Step 5** Verify that your application is running with `systemctl status -l application_name`
-

An Example Application in the Guest Shell

The following example demonstrates an application in the Guest Shell:

```
root@guestshell guestshell]# cat /etc/init.d/hello.sh
#!/bin/bash

OUTPUTFILE=/tmp/hello

rm -f $OUTPUTFILE
while true
do
    echo $(date) >> $OUTPUTFILE
    echo 'Hello World' >> $OUTPUTFILE
    sleep 10
done
[root@guestshell guestshell]#
[root@guestshell guestshell]#
[root@guestshell system]# cat /usr/lib/systemd/system/hello.service
[Unit]
Description=Trivial "hello world" example daemon

[Service]
ExecStart=/etc/init.d/hello.sh &
Restart=always
RestartSec=10s

[Install]
WantedBy=multi-user.target
[root@guestshell system]#
[root@guestshell system]# systemctl start hello
[root@guestshell system]# systemctl enable hello
[root@guestshell system]# systemctl status -l hello
hello.service - Trivial "hello world" example daemon
   Loaded: loaded (/usr/lib/systemd/system/hello.service; enabled)
   Active: active (running) since Sun 2015-09-27 18:31:51 UTC; 10s ago
   Main PID: 355 (hello.sh)
```

```
CGroup: /system.slice/hello.service
      ##355 /bin/bash /etc/init.d/hello.sh &
      ##367 sleep 10
```

```
Sep 27 18:31:51 guestshell hello.sh[355]: Executing: /etc/init.d/hello.sh &
[root@guestshell system]#
[root@guestshell guestshell]# exit
exit
[guestshell@guestshell ~]$ exit
logout
switch# reload
This command will reboot the system. (y/n)? [n] y
```

After reload

```
[root@guestshell guestshell]# ps -ef | grep hello
root      20      1  0 18:37 ?        00:00:00 /bin/bash /etc/init.d/hello.sh &
root     123    108  0 18:38 pts/4    00:00:00 grep --color=auto hello
[root@guestshell guestshell]#
[root@guestshell guestshell]# cat /tmp/hello
Sun Sep 27 18:38:03 UTC 2015
Hello World
Sun Sep 27 18:38:13 UTC 2015
Hello World
Sun Sep 27 18:38:23 UTC 2015
Hello World
Sun Sep 27 18:38:33 UTC 2015
Hello World
Sun Sep 27 18:38:43 UTC 2015
Hello World
[root@guestshell guestshell]#
```

Running under `systemd` / `systemctl`, your application is automatically restarted if it dies (or if you kill it). The Process ID is originally 226. After killing the application, it is automatically restarted with a Process ID of 257.

```
[root@guestshell guestshell]# ps -ef | grep hello
root      226      1  0 19:02 ?        00:00:00 /bin/bash /etc/init.d/hello.sh &
root     254    116  0 19:03 pts/4    00:00:00 grep --color=auto hello
[root@guestshell guestshell]#
[root@guestshell guestshell]# kill -9 226
[root@guestshell guestshell]#
[root@guestshell guestshell]# ps -ef | grep hello
root      257      1  0 19:03 ?        00:00:00 /bin/bash /etc/init.d/hello.sh &
root     264    116  0 19:03 pts/4    00:00:00 grep --color=auto hello
[root@guestshell guestshell]#
```

Troubleshooting Guest Shell Issues

Unable to Get Into Guest Shell After Downgrade to 7.0(3)I7

If you downgrade from the NX-OS 9.2(1) release to the NX-OS 7.0(3)7 release image (which does not have user namespace support) while the Guest Shell is in the process of activating or deactivating, you may run into the following condition where the Guest Shell activates, but you are unable to get into the Guest Shell. The reason for this issue is that if a reload is issued while the Guest Shell is in transition, the files within the Guest Shell can't get shifted back into an id range that is usable for NX-OS releases that don't have user namespace support.

```
switch# guestshell
Failed to mkdir .ssh for admin
```

```

admin RSA add failed
ERROR: Failed to connect with Virtual-service 'guestshell+'
switch#
switch# sh virt list

Virtual Service List:
Name           Status      Package Name
-----
guestshell+    Activated   guestshell.ova

switch# run bash ls -al /isan/vdc_1/virtual-instance/guestshell+/rootfs/
drwxr-xr-x  24 11000 11000 1024 Apr 11 10:44 .
drwxrwxrwx   4  root   root    80 Apr 27 20:08 ..
-rw-r--r--   1 11000 11000   0 Mar 21 16:24 .autorelabel
lrwxrwxrwx   1 11000 11000   7 Mar 21 16:24 bin -> usr/bin

```

To recover from this issue without losing the contents of the Guest Shell, reload the system with the previously-running NX-OS 9.2(x) image and let the Guest Shell get to the `Activated` state before reloading the system with the NX-OS 7.0(3)I7 image. Another option is to disable the Guest Shell while running NX-OS 9.2(x) and re-enable it after reloading with 7.0(3)I7.

If you do not have anything to preserve in the Guest Shell and you just want to recover it, you can destroy and recreate it without needing to change images.

Unable to Access Files on bootflash from root in the Guest Shell

You may find that you are unable to access files on bootflash from root in the Guest Shell.

From the host:

```

root@switch# ls -al /bootflash/try.that
-rw-r--r--  1 root  root  0 Apr 27 20:55 /bootflash/try.that
root@switch#

```

From the Guest Shell:

```

[root@guestshellbootflash]# ls -al /bootflash/try.that
-rw-r--r--  1 65534 host-root 0 Apr 27 20:55 /bootflash/try.that
[root@guestshellbootflash]# echo "some text" >> /bootflash/try.that
-bash: /bootflash/try.that: Permission denied
[root@guestshellbootflash]#

```

This may be due to the fact that, because the user namespace is being used to protect the host system, root in the Guest Shell is not actually the root of the system.

To recover from this issue, verify that the file permissions and group-id of the files allow for shared files on bootflash to be accessed as expected. You may need to change the permissions or group-id from the host Bash session.



CHAPTER 6

Broadcom Shell

- [About the Broadcom Shell, on page 87](#)
- [Guidelines and Limitations, on page 87](#)
- [Accessing the Broadcom Shell \(bcm-shell\), on page 87](#)
- [Examples of Broadcom Shell Commands, on page 91](#)

About the Broadcom Shell

The switch's front panel and fabric module line cards contain Broadcom Network Forwarding Engines (NFE). The number of NFEs varies depending upon the specific model of the front panel line card (LC) or the fabric module (FM).

Guidelines and Limitations

You can access and read information from the T2 ASICs without any limitations. However, Cisco does not recommend changing the T2 configuration settings. Use caution when accessing the Broadcom Shell.

Accessing the Broadcom Shell (bcm-shell)

The following sections describe approaches to access the Broadcom Shell (bcm-shell).

Accessing bcm-shell with the CLI API

The bcm-shell commands are passed directly from the Cisco NX-OS CLI to the specific T2 ASIC instance. The T2 ASIC instance can be on the fabric module or on the front panel line card.

The command syntax is as follows:

```
bcm-shell module module_number [instance_number:command]
```

Where

| | |
|----------------------|-------------------------------|
| <i>module_number</i> | Module number in the chassis. |
|----------------------|-------------------------------|

| | |
|------------------------|---|
| <i>instance_number</i> | T2 instance number <ul style="list-style-type: none"> • When not specified, the T2 instance number defaults to 0. • When a wildcard (*) is specified, all T2 instances are processed. |
| <i>command</i> | Broadcom command |



Note Cisco NX-OS command extensions such as ‘pipe include’ or ‘redirect output to file’ can be used to manage command output.



Note Entering commands with the CLI API are recorded in the system accounting log for auditing purposes. Commands that are entered directly from the bcm-shell are not recorded in the accounting log.

The following is an example when the T2 instance number not entered.:

```
switch# bcm-shell module 26 "ports"
Executing ports on bcm shell on module 26
ena/ speed/ link auto STP lrn inter max loop
port link duplex scan neg? state pause discrd ops face frame back
hg0 !ena 42G FD HW No Forward None FA XGMII 16360
hg1 !ena 42G FD HW No Forward None FA XGMII 16360
hg2 !ena 42G FD HW No Forward None FA XGMII 16360
hg3 !ena 42G FD HW No Forward None FA XGMII 16360
hg4 !ena 42G FD HW No Forward None FA XGMII 16360
<snip>
```

The following is an example when accessing the bcm-shell ‘ports’ display on module 26, T2 instance 1, and using the pipe to display only the ports in ‘up’ state:

```
switch# bcm-shell module 26 "1:ports" | inc up
port link duplex scan neg? state pause discrd ops face frame back
hg9 up 42G FD HW No Forward None FA XGMII 16360
hg10 up 42G FD HW No Forward None FA XGMII 16360
hg11 up 42G FD HW No Forward None FA XGMII 16360
```

The following example when the T2 instance number specified as a wildcard (*).:

The first set in the output is the result for T2 instance 0. The second set is the result for instance 1 on the fabric module.

```
switch# bcm-shell module 26 "*:ports" | inc up

port link duplex scan neg? state pause discrd ops face frame back
hg9 up 42G FD HW No Forward None FA XGMII 16360
hg10 up 42G FD HW No Forward None FA XGMII 16360
hg11 up 42G FD HW No Forward None FA XGMII 16360

port link duplex scan neg? state pause discrd ops face frame back
hg9 up 42G FD HW No Forward None FA XGMII 16360
hg10 up 42G FD HW No Forward None FA XGMII 16360
hg11 up 42G FD HW No Forward None FA XGMII 16360
```



Note Alternatively, you can access the bcm-shell on a module (LC or FM) directly from the supervisor. The following is an example of how to access the bcm-shell from the supervisor.

```
switch# bcm-shell module 3
Entering bcm shell on module 3
Available Unit Numbers: 0 1
bcm-shell.0>
```

Accessing the Native bcm-shell on the Fabric Module

An eight-slot line card (LC) chassis can host a maximum of six fabric modules (FMs). These slots are numbered 21 through 26. You must specify the FM that you wish to access the bcm-shell on.

The following example shows how to access the bcm-shell on the FM in slot 24, access context help, and exit the bcm-shell.

- Use the **show module** command to display the FMs.

```
switch# show module
Mod Ports Module-Type Model Status
-----
3 36 36p 40G Ethernet Module N9k-X9636PQ ok
4 36 36p 40G Ethernet Module N9k-X9636PQ ok
21 0 Fabric Module Nexus-C9508-FM ok
22 0 Fabric Module Nexus-C9508-FM ok
23 0 Fabric Module Nexus-C9508-FM ok
24 0 Fabric Module Nexus-C9508-FM ok
25 0 Fabric Module Nexus-C9508-FM ok
26 0 Fabric Module Nexus-C9508-FM ok
27 0 Supervisor Module Nexus-SUP-A active *
29 0 System Controller Nexus-SC-A active
```

- Attach to module 24 to gain access to the command line for the FM in slot 24.

```
switch# attach module 24
Attaching to module 24 ...
To exit type 'exit', to abort type '$.'
```

- Enter the command to gain root access to the fabric module software.

```
module-24# test hardware internal bcm-usd bcm-diag-shell
Available Unit Numbers: 0 1
bcm-shell.0> 1
```

At this point, you are at the Broadcom shell for the fabric module in slot 24, T2 ASIC instance 1. Any commands that you enter are specific to this specific ASIC instance.

- Enter '?' for the context help and to display the available commands.

```
help: "???" or "help" for summary
Commands common to all modes:
?      ??      ASSErt      BackGround  BCM
BCMx   break    BroadSync   CASE        CD
cint   CONFig   CONSOLE    CoPy        CPUDB
CTEcho CTInstall CTSetup     DATE        DBDump
DBParse DeBug    DeBugMod   DELAY       DEvice
```



Note Uppercase characters indicate command abbreviations.



Note The Cisco CLI feature of using TAB to complete a command entry is not available.

- Enter an abbreviated command.

```
bcm-shell.1> conf ?
Usage (CONFIG): Parameters: show <substring> | refresh |
                        save [filename=<filename>] [pattern=<substring>] |
                        add [<var>=<value>] | delete <var>
If no parameters are given, displays the current config vars.
show                display current config vars.
                    Next parameter (optional) maybe a substring to match
refresh             reload config vars from non-volatile storage
save                save config vars to non-volatile storage.
                    it can optionally save current config to any given file
                    providing the optional <pattern> will only save variables

matching the pattern
<var>=<value>       change the value of a config var
add <var>=<value>   create a new config var
delete <var>       deletes a config var
clear              deletes all config vars
=                  prompt for new value for each config var
NOTE: changes are not retained permanently unless saved
```

- Use the exit command to exit the bcm-shell and to detach from the FM.

```
bcm-shell.1> exit
module-24# exit
rlogin: connection closed.
```

Accessing the bcm-shell on the Line Card

When connecting to the T2 ASIC on the line card (LC), you first attach to the module, enter root mode, run the shell access exec, and select the ASIC instance to which you want to attach. The number of available ASICs depends on the model of the line card to which you are attached.

The following example shows how to access the bcm-shell of ASIC instance 1 on the LC in slot 2 and exit the bcm-shell on an LC that contains three T2 instances.

- Attach to module 2 to gain access to the command line for the LC in slot 2.

```
switch# attach module 2
Attaching to module 2 ...
To exit type 'exit', to abort type '$.'
Last login: Wed Aug 7 14:13:15 UTC 2013 from sup27 on tty0
```

- Enter the command to gain root access to the line card software.

```
switch-2# test hardware internal bcm-usd bcm-diag-shell
Available Unit Numbers: 0 1 2
bcm-shell.0> 1
bcm-shell.1>
```

At this point, you are at the Broadcom shell for the line card module in slot 2, T2 ASIC instance 1.

- Use the **exit** command to exit the bcm-shell and detach from the FM.

```
bcm-shell.1> exit
module-2# exit
rlogin: connection closed.
```

Examples of Broadcom Shell Commands

This section contains examples of Broadcom shell commands and output.

Displaying L2 Entries

To compare the L2 entries on an FM and an LC, display the L2 entries on the T2 instance.

- Attach to a T2 instance on an LC in the chassis and enter the show the `l2_mem_entries` command.

```
switch# bcm-shell module 2 "config show l2_mem_entries"
l2_mem_entries=163840
```

- Attach to a T2 instance on an FM in the chassis and enter the show the `l2_mem_entries` command.

```
switch# bcm-shell module 24 "config show l2_mem_entries"
l2_mem_entries=32768
```

In this example, the ASICs have different values that are configured for L2 entries because the operation mode setting of the T2 ASIC is different on the LC and the FM.

Displaying Routing Information from FM and LC ASIC Instances

The following shows how to display the routing information from an FM instance and an LC ASIC instance.



Note

- FM module ASIC instances maintain the longest prefix match (LPM) entries.
- LC module ASIC instances maintain host routes.

- Attach to the bcm-shell of one of the FM ASIC instances and enter the **l3 defip show** command.

```
switch# bcm-shell module 22 "l3 defip show"
Unit 0, Total Number of DEFIP entries: 16385
# VRF Net addr Next Hop Mac INTF MODID PORT PRIO CLASS HIT VLAN
0 1 192.168.1.0/24 00:00:00:00:00:00 10000 4 0 0 0 0 n
0 1 192.168.2.0/24 00:00:00:00:00:00 10000 4 0 0 0 0 n
```

- Attach to the bcm-shell of one of the LC ASIC instances and enter the **l3 defip show** command.

```
switch# bcm-shell module 2 "l3 defip show"
Unit 2, Total Number of DEFIP entries: 8193
# VRF Net addr Next Hop Mac INTF MODID PORT PRIO CLASS HIT VLAN
2048 Override 0.0.0.0/0 00:00:00:00:00:00 10000 4 0 0 0 0 n
```

Displaying Spanning Tree Group Entries

The following shows how to display spanning tree group (STG) entries.

```
switch# bcm-shell module 2 "stg show"
STG 0: contains 7 VLANs (4032-4035,4042,4044,4095)
Forward: xe,hg
STG 1: contains 1 VLAN (4043)
Disable: xe
Forward: hg
STG 4: contains 1 VLAN (1)
Block: xe
Forward: hg
STG 5: contains 1 VLAN (100)
Block: xe
Forward: hg
```

Display T2 Counters for Interface xe0

The following shows how to display T2 counters for interface xe0 that maps to a front panel port.

```
switch# bcm-shell module 2 "show counters xe0"
RDBGCl.xe0 : 38,169 +4
R127.xe0 : 37,994 +4
R511.xe0 : 16,562 +2
RPKT.xe0 : 55,046 +6
RMCA.xe0 : 54,731 +6
<snip>
```



Note The numbers in the output on the far right denoted with '+', indicate the change in the counter value since the last show command was executed.

Displaying L3 Information

The following shows how to display L3 information from the T2.

```
switch# bcm-shell module 2 "l3 l3table show"
Unit 0, free L3 table entries: 147448
Entry VRF IP address Mac Address INTF MOD PORT CLASS HIT
1 1 10.100.100.2 00:00:00:00:00:00 100006 0 0 0 n
2 1 10.101.101.2 00:00:00:00:00:00 100006 0 0 0 y
3 1 10.101.101.1 00:00:00:00:00:00 149151 0 0 0 y (LOCAL ROUTE)
4 1 10.100.100.1 00:00:00:00:00:00 149151 0 0 0 n (LOCAL ROUTE)
5 1 10.101.101.0 00:00:00:00:00:00 100000 0 0 0 n (LOCAL ROUTE)
6 1 10.100.100.0 00:00:00:00:00:00 100000 0 0 0 n (LOCAL ROUTE)
7 1 10.101.101.255 00:00:00:00:00:00 149150 0 0 0 n
8 1 10.100.100.255 00:00:00:00:00:00 149150 0 0 0 n
```



CHAPTER 7

Barefoot Shell

- [About the Barefoot Shell, on page 93](#)
- [Guidelines and Limitations, on page 93](#)
- [Accessing the Barefoot Shell with CLI API, on page 93](#)

About the Barefoot Shell

The Cisco Nexus 3464C and Cisco Nexus 34180YC are top of the rack (ToR) switches that contain the Barefoot Networks Tofino ASICs. Because these switches are TORs, they do not feature any interchangeable fabric modules or line cards. The Barefoot Networks Tofino ASICs exist on the switches' system boards.

On the Cisco Nexus 3464C and Cisco Nexus 34180YC switches, a specific shell enables access directly to the ASICs. This shell is called the Barefoot shell.

Guidelines and Limitations



Note Beginning with Cisco NX-OS Release 9.3(3), this feature is no longer supported.

Following are the guidelines and limitations for the Barefoot shell:

- For notes about platform support, see [Nexus Switch Platform Support Matrix](#).
- The Barefoot shell is for authorized use only. Use it with extreme caution and only when authorized by Cisco.
- You can access and read information from the Barefoot ASICs without any limitations. However, Cisco does not recommend changing the Barefoot configuration settings.

Accessing the Barefoot Shell with CLI API

The Barefoot shell is available directly from EXEC mode. The shell enables you to issue commands directly to the Barefoot Networks Tofino ASICs, which are supported in the Cisco Nexus 3464C and Cisco Nexus 34180YC switches. You can check the model of your switch by issuing the **show module** command.

Example:

```
switch-1# show module
Mod Ports  Module-Type                               Model                               Status
-----  -
1      66      64x100G QSFP28 + 2x10G SFP+ Ethernet  N3K-C3464C                          active *
```

This topic documents how to access the Barefoot shell from NX-OS, use the online help function, and exit the Barefoot shell. For detailed documentation about the Barefoot shell and its commands, consult the Barefoot Networks Tofino documentation.

- To enter the Barefoot shell, issue the **bfshell** command.

Example:

```
switch-1# bfshell
Warning: bfshell access should be used with caution
```

```
*****
*          WARNING: Authorised Access Only          *
*****
```

```
bfshell>
```

The command prompt changes to **bfshell** to indicate that you are in the Barefoot shell.

- To get a list of top-level commands in the hierarchy, type **?** (question mark).

Example:

```
bfshell> ?
cint          C Interpreter
exit          Exit this CLI session
help          Display an overview of the CLI syntax
pd-switch    pd_switch Related Commands
pipemgr      Pipe manager commands
quit         Exit this CLI session
switchapi    switchAPI commands
ucli         UCLI commands
version      Display the SDE version
```

- To get details about the types of help available, type **help** at the prompt.

Example:

```
bfshell> help
```

CONTEXT SENSITIVE HELP

```
[?] - Display context sensitive help. This is either a list of possible
      command completions with summaries, or the full syntax of the
      current command. A subsequent repeat of this key, when a command
      has been resolved, will display a detailed reference.
```

AUTO-COMPLETION

```
The following keys both perform auto-completion for the current command line.
If the command prefix is not unique then the bell will ring and a subsequent
repeat of the key will display possible completions.
```

```
[enter] - Auto-completes, syntax-checks then executes a command. If there is
          a syntax error then offending part of the command line will be
          highlighted and explained.
```

```
[space] - Auto-completes, or if the command is already resolved inserts a space.
```

```

MOVEMENT KEY
[CTRL-A] - Move to the start of the line.
[CTRL-E] - Move to the end of the line.
[up]     - Move to the previous command line held in history.
[down]   - Move to the next command line held in history.
[left]   - Move the insertion point left one character.
[right]  - Move the insertion point right one character.

DELETION KEYS
[CTRL-C] - Delete and abort the current line.
[CTRL-D] - Delete the character to the right on the insertion point.
[CTRL-K] - Delete all the characters to the right of the insertion point.
[CTRL-U] - Delete the whole line.
[backspace] - Delete the character to the left of the insertion point.

ESCAPE SEQUENCES
!! - Substitute the the last command line.
!N - Substitute the Nth command line (absolute as per 'history' command)
!-N - Substitute the command line entered N lines before (relative)
bfshell>

```

- To change to different command subsystems in the hierarchy, enter the top-level command name.

For example, to change to the **pd-switch** commands:

```

bfshell> ?
  cint      C Interpreter
  exit      Exit this CLI session
  help      Display an overview of the CLI syntax
  pd-switch pd_switch Related Commands
  pipemgr   Pipe manager commands
  quit      Exit this CLI session
  switchapi switchAPI commands
  ucli      UCLI commands
  version   Display the SDE version

bfshell> pd-switch
pd-switch:0>

```

Notice the prompt changes to display the subsystem.

New commands are available within each subsystem in the hierarchy.

```

pd-switch:0> ?
  dump_profile  dump action profile entries
  dump_table    dump table entries
  end           End pd-switch sub-commands
  exit          Exit this CLI session
  get           Display value of shell variable
  help          Display an overview of the CLI syntax
  pd            prefix for all pd commands
  quit          Exit this CLI session
  set           Set value of shell variable
  var           Declare new shell variable

```

- To exit the Barefoot shell, issue either the **exit** or **quit** command, which terminates the current Barefoot CLI session and returns you to the NX-OS prompt. You can issue these commands from anywhere in the hierarchy.

Examples:

```

pd-switch:0> exit
switch-1#
bfshell> quit
switch-1#

```




CHAPTER 8

Python API

- [Using Python, on page 97](#)

Using Python

This section describes how to write and execute Python scripts.

Cisco Python Package

Cisco NX-OS provides a Cisco Python package that enables access to many core network-device modules, such as interfaces, VLANs, VRFs, ACLs, and routes. You can display the details of the Cisco Python package by entering the **help()** command. To obtain additional information about the classes and methods in a module, you can run the help command for a specific module. For example, **help(cisco.interface)** displays the properties of the `cisco.interface` module.

For more Python modules, you can install the `python-modules-nxos` RPM (`python-modules-nxos-1.0.0-9.2.1.lib32_x86.rpm`) from https://devhub.cisco.com/artifactory/open-nxos/9.2.1/x86_64/. Refer to the "Manage Feature RPMs" section for instructions on installing an RPM.

The following is an example of how to display information about the Cisco Python package:

```
>>> import cisco
>>> help(cisco)
Help on package cisco:

NAME
    cisco

FILE
    /isan/python/scripts/cisco/__init__.py

PACKAGE CONTENTS
    acl
    bgp
    cisco_secret
    cisco_socket
    feature
    interface
    key
    line_parser
    md5sum
    nxcli
```

```

ospf
routemap
routes
section_parser
ssh
system
tacacs
vrf

CLASSES
__builtin__.object
cisco.cisco_secret.CiscoSecret
cisco.interface.Interface
cisco.key.Key

```

The following is an example of how to display information about the Cisco Python Package for Python 3:

```

switch# python3
Python 3.7.3 (default, Nov 20 2019, 14:38:01)
[GCC 5.2.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import cisco
>>> help(cisco)
Help on package cisco:

NAME
cisco

PACKAGE CONTENTS
acl
bgp
buffer_depth_monitor
check_port_discards
cisco_secret
feature
historys
interface
ipaddress
key
line_parser
mac_address_table
md5sum
nxcli
nxos_cli
ospf
routemap
routes
section_parser
ssh
system
tacacs
transfer
vlan
vrf

CLASSES
builtins.dict(builtins.object)
cisco.history.History
builtins.object
cisco.cisco_secret.CiscoSecret
cisco.interface.Interface
cisco.key.Key

```


Using the CLI Command APIs

The Python programming language uses three APIs that can execute CLI commands. The APIs are available from the Python CLI module.

These APIs are listed in the following table. You must enable the APIs with the **from cli import *** command. The arguments for these APIs are strings of CLI commands. To execute a CLI command through the Python interpreter, you enter the CLI command as an argument string of one of the following APIs:

Table 4: CLI Command APIs

| API | Description |
|--|--|
| cli() Example: <pre>string = cli ("cli-command")</pre> | Returns the raw output of CLI commands, including control or special characters. Note The interactive Python interpreter prints control or special characters 'escaped'. A carriage return is printed as '\n' and gives results that can be difficult to read. The clip() API gives results that are more readable. |
| clid() Example: <pre>json_string = clid ("cli-command")</pre> | Returns JSON output for cli-command, if XML support exists for the command, otherwise an exception is thrown. Note This API can be useful when searching the output of show commands. |
| clip() Example: <pre>clip ("cli-command")</pre> | Prints the output of the CLI command directly to stdout and returns nothing to Python. Note <pre>clip ("cli-command")</pre> is equivalent to <pre>r=cli("cli-command") print r</pre> |

When two or more commands are run individually, the state is not persistent from one command to subsequent commands.

In the following example, the second command fails because the state from the first command does not persist for the second command:

```
>>> cli("conf t")
>>> cli("interface eth4/1")
```

When two or more commands are run together, the state is persistent from one command to subsequent commands.

In the following example, the second command is successful because the state persists for the second and third commands:

```
>>> cli("conf t ; interface eth4/1 ; shut")
```



Note Commands are separated with ";" as shown in the example. The semicolon (;) must be surrounded with single blank characters.

Invoking the Python Interpreter from the CLI

The following example shows how to invoke Python 2 from the CLI:



Note The Python interpreter is designated with the ">>>" or "... " prompt.



Important Python 2.7 is End of Support, Future NX-OS software deprecates Python 2.7 support. We recommend for new scripts to use **python3** instead. Type **python3** to use the new shell.

The following example shows how to invoke Python 3 from the CLI:

```
switch# python3
Python 3.7.3 (default, Nov 20 2019, 14:38:01)
[GCC 5.2.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>>
>>> from cli import *
>>> import json
>>> cli('configure terminal ; interface loopback 1 ; no shut')
''
>>> intflist=json.loads(clid('show interface brief'))
>>> i=0
>>> while i < len(intflist['TABLE_interface']['ROW_interface']):
...     intf=intflist['TABLE_interface']['ROW_interface'][i]
...     i=i+1
...     if intf['state'] == 'up':
...         print(intf['interface'])
...
mgmt0
loopback1
>>>
```

Display Formats

The following examples show various display formats using the Python APIs:

Example 1:

```
>>> from cli import *
>>> cli("conf ; interface loopback 1")
''
>>> clip('where detail')
mode:
username:          admin
vdc:               switch
routing-context vrf: default
```

Example 2:

```
>>> from cli import *
>>> cli("conf ; interface loopback 1")
''
>>> cli('where detail')
' mode:                \n username:                admin\n vdc:
switch\n routing-context vrf: default\n'
>>>
```

Example 3:

```
>>> r = cli('where detail')
>>> print(r)
mode:
username: admin
vdc: switch
routing-context vrf: default

>>>
```

Example 4:

```
>>> from cli import *
>>> import json
>>> out=json.loads(clid('show version'))
>>> for k in out.keys():
... print("%30s - %s" % (k,out[k]))
...
header_str - Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2020, Cisco and/or its affiliates.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their own
licenses, such as open source. This software is provided "as is," and unless
otherwise stated, there is no warranty, express or implied, including but not
limited to warranties of merchantability and fitness for a particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0 or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
http://www.gnu.org/licenses/old-licenses/library.txt.
bios_ver_str - 07.67
kickstart_ver_str - 9.3(5) [build 9.3(4)IIL9(0.879)]
nxos_ver_str - 9.3(5) [build 9.3(4)IIL9(0.879)]
bios_cmpl_time - 01/29/2020
kick_file_name - bootflash:///nxos.9.3.4.IIL9.0.879.bin
nxos_file_name - bootflash:///nxos.9.3.4.IIL9.0.879.bin
kick_cmpl_time - 5/10/2020 21:00:00
nxos_cmpl_time - 5/10/2020 21:00:00
kick_tmstamp - 05/12/2020 07:08:44
nxos_tmstamp - 05/12/2020 07:08:44
chassis_id - Nexus9000 93180YC-EX chassis
cpu_name - Intel(R) Xeon(R) CPU @ 1.80GHz
memory - 24632252
mem_type - kB
proc_board_id - FDO22280FFK
```

```

host_name - switch
bootflash_size - 53298520
kern_uptm_days - 0
kern_uptm_hrs - 0
kern_uptm_mins - 19
kern_uptm_secs - 34
rr_usec - 641967
rr_ctime - Tue May 12 09:52:28 2020
rr_reason - Reset Requested by CLI command reload
rr_sys_ver - 9.4(1)
rr_service - None
plugins - Core Plugin, Ethernet Plugin
manufacturer - Cisco Systems, Inc.
>>>

```

Non-Interactive Python

A Python script can run in non-interactive mode by providing the Python script name as an argument to the Python CLI command. Python scripts must be placed under the bootflash or volatile scheme. A maximum of 32 command-line arguments for the Python script are allowed with the Python CLI command.

The switch also supports the source CLI command for running Python scripts. The `bootflash:scripts` directory is the default script directory for the source CLI command.

This example shows the script first and then executing it. Saving is like bringing any file to the bootflash.

```

switch# show file bootflash:scripts/deltaCounters.py
#!/isan/bin/python3
from cli import *
import sys, time
ifName = sys.argv[1]
delay = float(sys.argv[2])
count = int(sys.argv[3])
cmd = 'show interface ' + ifName + ' counters'
out = json.loads(clid(cmd))
rxuc = int(out['TABLE_rx_counters']['ROW_rx_counters'][0]['eth_inucast'])
rxmc = int(out['TABLE_rx_counters']['ROW_rx_counters'][1]['eth_inmcast'])
rxbc = int(out['TABLE_rx_counters']['ROW_rx_counters'][1]['eth_inbcast'])
txuc = int(out['TABLE_tx_counters']['ROW_tx_counters'][0]['eth_outucast'])
txmc = int(out['TABLE_tx_counters']['ROW_tx_counters'][1]['eth_outmcast'])
txbc = int(out['TABLE_tx_counters']['ROW_tx_counters'][1]['eth_outbcast'])
print ('row rx_ucast rx_mcast rx_bcast tx_ucast tx_mcast tx_bcast')
print ('=====')
print (' %8d %8d %8d %8d %8d %8d' % (rxuc, rxmc, rxbc, txuc, txmc, txbc))
print ('=====')
i = 0
while (i < count):
    time.sleep(delay)
    out = json.loads(clid(cmd))
    rxucNew = int(out['TABLE_rx_counters']['ROW_rx_counters'][0]['eth_inucast'])
    rxmcNew = int(out['TABLE_rx_counters']['ROW_rx_counters'][1]['eth_inmcast'])
    rxbcNew = int(out['TABLE_rx_counters']['ROW_rx_counters'][1]['eth_inbcast'])
    txucNew = int(out['TABLE_tx_counters']['ROW_tx_counters'][0]['eth_outucast'])
    txmcNew = int(out['TABLE_tx_counters']['ROW_tx_counters'][1]['eth_outmcast'])
    txbcNew = int(out['TABLE_tx_counters']['ROW_tx_counters'][1]['eth_outbcast'])
    i += 1
    print ('%-3d %8d %8d %8d %8d %8d' % (i, rxucNew - rxuc, rxmcNew - rxmc, rxbcNew -
    rxbc, txucNew - txuc, txmcNew - txmc, txbcNew - txbc))

switch# python bootflash:scripts/deltaCounters.py mgmt0 1 5
row rx_ucast rx_mcast rx_bcast tx_ucast tx_mcast tx_bcast
=====

```

```

          291      8233      1767      185      57      2
=====
1          1          4          1          1          0          0
2          2          5          1          2          0          0
3          3          9          1          3          0          0
4          4          12         1          4          0          0
5          5          17          1          5          0          0
switch#

```

The following example shows how a source command specifies command-line arguments. In the example, *policy-map* is an argument to the `cgrep python` script. The example also shows that a source command can follow the pipe operator (`|`).

```

switch# show running-config | source sys/cgrep policy-map

policy-map type network-qos nw-pfc
policy-map type network-qos no-drop-2
policy-map type network-qos wred-policy
policy-map type network-qos pause-policy
policy-map type qos foo
policy-map type qos classify
policy-map type qos cos-based
policy-map type qos no-drop-2
policy-map type qos pfc-tor-port

```

Running Scripts with Embedded Event Manager

On Cisco Nexus switches, Embedded Event Manager (EEM) policies support Python scripts.

The following example shows how to run a Python script as an EEM action:

- An EEM applet can include a Python script with an action command.

```

switch# show running-config eem

!Command: show running-config eem
!Running configuration last done at: Thu Jun 25 15:29:38 2020
!Time: Thu Jun 25 15:33:19 2020

version 9.3(5) Bios:version 07.67
event manager applet a1
  event cli match "show clock"
  action 1 cli python bootflash:pydate.py

switch# show file logflash:vdc_1/event_archive_1 | last 33

eem_event_time:06/25/2020,15:34:24 event_type:cli event_id:24 slot:active(1) vdc
:1 severity:minor applets:a1
eem_param_info:command = "exshow clock"
Starting with policy a1
stty: standard input: Inappropriate ioctl for device
Executing the following commands succeeded:
  python bootflash:pydate.py
Completed executing policy a1
Event Id:24 event type:10241 handling completed

```

- You can search for the action that is triggered by the event in the log file by running the `show file logflash:event_archive_1` command.

```

switch# show file logflash:event_archive_1 | last 33

eem_event_time:05/01/2011,19:40:28 event_type:cli event_id:8 slot:active(1)
vdc:1 severity:minor applets:a1
eem_param_info:command = "exshow clock"
Starting with policy a1
Python

2011-05-01 19:40:28.644891
Executing the following commands succeeded:
    python bootflash:pydate.py

PC_VSH_CMD_TLV(7679) with q

```

Python Integration with Cisco NX-OS Network Interfaces

On Cisco Nexus switches, Python is integrated with the underlying Cisco NX-OS network interfaces. You can switch from one virtual routing context to another by setting up a context through the `cisco.vrf.set_global_vrf()` API.

The following example shows how to retrieve an HTML document over the management interface of a device. You can also establish a connection to an external entity over the in-band interface by switching to a desired virtual routing context.

```

switch# python

Warning: Python 2.7 is End of Support, and future NXOS software will deprecate
python 2.7 support. It is recommended for new scripts to use 'python3' instead.
Type "python3" to use the new shell.

Python 2.7.11 (default, Jun  4 2020, 09:48:24)
[GCC 4.6.3] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>> import urllib2
>>> from cisco.vrf import *
>>> set_global_vrf('management')
>>> page=urllib2.urlopen('http://172.23.40.211:8000/welcome.html')
>>> print page.read()
Hello Cisco Nexus 9000
>>>
>>> import cisco
>>> help(cisco.vrf.set_global_vrf)
Help on function set_global_vrf in module cisco.vrf:
set_global_vrf(vrf)
Sets the global vrf. Any new sockets that are created (using socket.socket)
will automatically get set to this vrf (including sockets used by other
python libraries).
Arguments:
vrf: VRF name (string) or the VRF ID (int).
Returns: Nothing
>>>

```

Cisco NX-OS Security with Python

Cisco NX-OS resources are protected by the Cisco NX-OS Sandbox layer of software and by the CLI role-based access control (RBAC).

All users who are associated with a Cisco NX-OS network-admin or dev-ops role are privileged users. Users who are granted access to Python with a custom role are regarded as nonprivileged users. Nonprivileged users

have limited access to Cisco NX-OS resources, such as the file system, guest shell, and Bash commands. Privileged users have greater access to all the resources of Cisco NX-OS.

Examples of Security and User Authority

The following example shows how a privileged user runs commands:

Python 3 example.

```
switch# python3
Python 3.7.3 (default, Nov 20 2019, 14:38:01)
[GCC 5.2.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import os
>>> os.system('whoami')
admin
0
>>> f=open('/tmp/test','w')
>>> f.write('hello from python')
17
>>> f.close()
>>> r=open('/tmp/test','r')
>>> print(r.read())
hello from python
>>> r.close()
>>>
```

The following example shows a nonprivileged user being denied access:

```
switch# python3
Python 3.7.3 (default, Nov 20 2019, 14:38:01)
[GCC 5.2.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import os
>>> os.system('whoami')
system(whoami): rejected!
-1
>>> f=open('/tmp/test','w')
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
PermissionError: [Errno 13] Permission denied: '/tmp/test'
>>>
```

RBAC controls CLI access based on the login user privileges. A login user's identity is given to Python that is invoked from the CLI shell or from Bash. Python passes the login user's identity to any subprocess that is invoked from Python.

The following is an example for a privileged user:

```
>>> from cli import *
>>> cli('show clock')
'Warning: No NTP peer/server configured. Time may be out of sync.\n15:39:39.513 UTC Thu Jun
25 2020\nTime source is NTP\n'
>>> cli('configure terminal ; vrf context myvrf')
''
>>> clip('show running-config l3vm')

!Command: show running-config l3vm
!Running configuration last done at: Thu Jun 25 15:39:49 2020
!Time: Thu Jun 25 15:39:55 2020

version 9.3(5) Bios:version 07.67
```

```
interface mgmt0
  vrf member management
vrf context blue
vrf context management
vrf context myvrf
```

The following is an example for a nonprivileged user:

```
>>> from cli import *
>>> cli('show clock')
'11:18:47.482 AM UTC Sun May 08 2011\n'
>>> cli('configure terminal ; vrf context myvrf2')
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "/isan/python/scripts/cli.py", line 20, in cli
    raise cmd_exec_error(msg)
errors.cmd_exec_error: '% Permission denied for the role\n\nCmd exec error.\n'
```

The following example shows an RBAC configuration:

```
switch# show user-account
user:admin
    this user account has no expiry date
    roles:network-admin
user:pyuser
    this user account has no expiry date
    roles:network-operator python-role
switch# show role name python-role
```

Example of Running Script with Scheduler

The following example shows a Python script that is running the script with the scheduler feature:

```
#!/bin/env python
from cli import *
from nxos import *
import os

switchname = cli("show switchname")
try:
    user = os.environ['USER']
except:
    user = "No user"
    pass

msg = user + " ran " + __file__ + " on : " + switchname
print msg
py_syslog(1, msg)
# Save this script in bootflash:///scripts
```

Python 3 example.

```
#!/bin/env python3
from cli import *
from nxos import *
import os

switchname = cli("show switchname")
try:
    user = os.environ['USER']
except:
```



```

    user = "No user"
    pass

msg = user + " ran " + __file__ + " on : " + switchname
print(msg)
py_syslog(1, msg)

# Save this script in bootflash:///scripts

switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# feature scheduler
switch(config)# scheduler job name testplan
switch(config-job)# python bootflash:///scripts/test.py
switch(config-job)# exit
switch(config)# scheduler schedule name testplan
switch(config-schedule)# job name testplan
switch(config-schedule)# time start now repeat 0:0:4
Schedule starts from Sat Jun 13 04:29:38 2020
switch# 2020 Jun 13 04:29:41 switch %USER-1-SYSTEM_MSG: No user ran /bootflash/scripts/test.py
  on : switch - nxpython
switch# show scheduler schedule
Schedule Name : testplan
-----
User Name : admin
Schedule Type : Run every 0 Days 0 Hrs 4 Mins
Start Time : Sat Jun 13 04:29:38 2020
Last Execution Time : Sat Jun 13 04:29:38 2020
Last Completion Time: Sat Jun 13 04:29:41 2020
Execution count : 1
-----
Job Name Last Execution Status
-----
testplan Success (0)
=====
switch#

switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# feature scheduler
switch(config)# scheduler job name testplan
switch(config-job)# python bootflash:///scripts/testplan.py
switch(config-job)# exit
switch(config)# scheduler schedule name testplan
switch(config-schedule)# job name testplan
switch(config-schedule)# time start now repeat 0:0:4
Schedule starts from Mon Mar 14 16:40:03 2011
switch(config-schedule)# end
switch# term mon
2011 Mar 14 16:38:03 switch %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by admin on
10.19.68.246@pts/2
switch# show scheduler schedule
Schedule Name      : testplan
-----
User Name          : admin
Schedule Type      : Run every 0 Days 0 Hrs 4 Mins
Start Time         : Mon Mar 14 16:40:03 2011
Last Execution Time : Yet to be executed
-----
Job Name           Last Execution Status
-----
testplan           -NA-
=====

```

Example of Running Script with Scheduler

```
switch#  
switch# 2011 Mar 14 16:40:04 switch %USER-1-SYSTEM_MSG: No user ran  
/bootflash/scripts/testplan.py on : switch - nxpython  
2011 Mar 14 16:44:04 switch last message repeated 1 time  
switch#
```



CHAPTER 9

Scripting with Tcl

This chapter contains the following topics:

- [About Tcl, on page 109](#)
- [Running the Tclsh Command, on page 112](#)
- [Navigating Cisco NX-OS Modes from the Tclsh Command, on page 113](#)
- [Tcl References, on page 114](#)

About Tcl

Tcl (pronounced "tickle") is a scripting language that increases flexibility of CLI commands. You can use Tcl to extract certain values in the output of a **show** command, perform switch configurations, run Cisco NX-OS commands in a loop, or define Embedded Event Manager (EEM) policies in a script.

This section describes how to run Tcl scripts or run Tcl interactively on switches.

Guidelines and Limitations

Following are guidelines and limitations for TCL scripting:

Some processes and **show** commands can cause a large amount of output. If you are running scripts, and need to terminate long-running output, use Ctrl+C (not Ctrl+Z) to terminate the command output. If you use Ctrl+Z, a SIGCONT (signal continuation) message can be generated, which can cause the script to halt. Scripts that are halted through SIGCONT messages require user intervention to resume operation.

Tclsh Command Help

Command help is not available for Tcl commands. You can still access the help functions of Cisco NX-OS commands from within an interactive Tcl shell.

This example shows the lack of Tcl command help in an interactive Tcl shell:

```
switch# tclsh
switch-tcl# set x 1
switch-tcl# puts ?
      ^
% Invalid command at '^' marker.
switch-tcl# configure ?
<CR>
  session  Configure the system in a session
```

```
terminal Configure the system from terminal input
switch-tcl#
```



Note In the preceding example, the Cisco NX-OS command help function is still available but the Tcl **puts** command returns an error from the help function.

Tclsh Command History

You can use the arrow keys on your terminal to access commands you previously entered in the interactive Tcl shell.



Note The **tclsh** command history is not saved when you exit the interactive Tcl shell.

Tclsh Tab Completion

You can use tab completion for Cisco NX-OS commands when you are running an interactive Tcl shell. Tab completion is not available for Tcl commands.

Tclsh CLI Command

Although you can directly access Cisco NX-OS commands from within an interactive Tcl shell, you can only execute Cisco NX-OS commands in a Tcl script if they are prepended with the Tcl **cli** command.

In an interactive Tcl shell, the following commands are identical and execute properly:

```
switch-tcl# cli show module 1 | incl Mod
switch-tcl# cli "show module 1 | incl Mod"
switch-tcl# show module 1 | incl Mod
```

In a Tcl script, you must prepend Cisco NX-OS commands with the Tcl **cli** command as shown in the following example:

```
set x 1
cli show module $x | incl Mod
cli "show module $x | incl Mod"
```

If you use the following commands in your script, the script fails and the Tcl shell displays an error:

```
show module $x | incl Mod
"show module $x | incl Mod"
```

Tclsh Command Separation

The semicolon (;) is the command separator in both Cisco NX-OS and Tcl. To execute multiple Cisco NX-OS commands in a Tcl command, you must enclose the Cisco NX-OS commands in quotes ("").

In an interactive Tcl shell, the following commands are identical and execute properly:

```
switch-tcl# cli "configure terminal ; interface loopback 10 ; description loop10"
switch-tcl# cli configure terminal ; cli interface loopback 10 ; cli description loop10
switch-tcl# cli configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

switch(config-tcl)# cli interface loopback 10
switch(config-if-tcl)# cli description loop10
switch(config-if-tcl)#
```

In an interactive Tcl shell, you can also execute Cisco NX-OS commands directly without prepending the Tcl **cli** command:

```
switch-tcl# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

switch(config-tcl)# interface loopback 10
switch(config-if-tcl)# description loop10
switch(config-if-tcl)#
```

Tcl Variables

You can use Tcl variables as arguments to the Cisco NX-OS commands. You can also pass arguments into Tcl scripts. Tcl variables are not persistent.

The following example shows how to use a Tcl variable as an argument to a Cisco NX-OS command:

```
switch# tclsh
switch-tcl# set x loop10
switch-tcl# cli "configure terminal ; interface loopback 10 ; description $x"
switch(config-if-tcl)#
```

Tclquit

The **tclquit** command exits the Tcl shell regardless of which Cisco NX-OS command mode is currently active. You can also press **Ctrl-C** to exit the Tcl shell. The **exit** and **end** commands change Cisco NX-OS command modes. The **exit** command terminates the Tcl shell only from the EXEC command mode.

Tclsh Security

The Tcl shell is executed in a sandbox to prevent unauthorized access to certain parts of the Cisco NX-OS system. The system monitors CPU, memory, and file system resources being used by the Tcl shell to detect events such as infinite loops, excessive memory utilization, and so on.

You configure the initial Tcl environment with the **scripting tcl init** *init-file* command.

You can define the looping limits for the Tcl environment with the **scripting tcl recursion-limit** *iterations* command. The default recursion limit is 1000 iterations.

Running the Tclsh Command

You can run Tcl commands from either a script or on the command line using the **tclsh** command.



Note You cannot create a Tcl script file at the CLI prompt. You can create the script file on a remote device and copy it to the bootflash: directory on the Cisco NX-OS device.

SUMMARY STEPS

1. **tclsh** [**bootflash:filename** [*argument ...*]]

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | tclsh [bootflash:filename [<i>argument ...</i>]] Example: <pre>switch# tclsh ? <CR> bootflash: The file to run</pre> | Starts a Tcl shell. If you run the tclsh command with no arguments, the shell runs interactively, reading Tcl commands from standard input and printing command results and error messages to the standard output. You exit from the interactive Tcl shell by typing tclquit or Ctrl-C . If you run the tclsh command with arguments, the first argument is the name of a script file containing Tcl commands and any additional arguments are made available to the script as variables. |

Example

The following example shows an interactive Tcl shell:

```
switch# tclsh
switch-tcl# set x 1
switch-tcl# cli show module $x | incl Mod
Mod  Ports  Module-Type          Model          Status
1    36      36p 40G Ethernet Module  N9k-X9636PQ   ok
Mod  Sw      Hw
Mod  MAC-Address(es)      Serial-Num

switch-tcl# exit
switch#
```

The following example shows how to run a Tcl script:

```
switch# show file bootflash:showmodule.tcl
set x 1
while {$x < 19} {
```

```
cli show module $x | incl Mod
set x [expr {$x + 1}]
}

switch# tclsh bootflash:showmodule.tcl
Mod  Ports  Module-Type                               Model                Status
1    36      36p 40G Ethernet Module                 N9k-X9636PQ         ok
Mod  Sw      Hw
Mod  MAC-Address(es)                        Serial-Num

switch#
```

Navigating Cisco NX-OS Modes from the Tclsh Command

You can change modes in Cisco NX-OS while you are running an interactive Tcl shell.

SUMMARY STEPS

1. **tclsh**
2. **configure terminal**
3. **tclquit**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | tclsh Example: switch# tclsh switch-tcl# | Starts an interactive Tcl shell. |
| Step 2 | configure terminal Example: switch-tcl# configure terminal switch(config-tcl)# | Runs a Cisco NX-OS command in the Tcl shell, changing modes. Note The Tcl prompt changes to indicate the Cisco NX-OS command mode. |
| Step 3 | tclquit Example: switch-tcl# tclquit switch# | Terminates the Tcl shell, returning to the starting mode. |

Example

The following example shows how to change Cisco NX-OS modes from an interactive Tcl shell:

```

switch# tclsh
switch-tcl# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config-tcl)# interface loopback 10
switch(config-if-tcl)# ?
  description  Enter description of maximum 80 characters
  inherit      Inherit a port-profile
  ip           Configure IP features
  ipv6        Configure IPv6 features
  logging      Configure logging for interface
  no          Negate a command or set its defaults
  rate-limit   Set packet per second rate limit
  shutdown    Enable/disable an interface
  this        Shows info about current object (mode's instance)
  vrf         Configure VRF parameters
  end         Go to exec mode
  exit        Exit from command interpreter
  pop         Pop mode from stack or restore from name
  push        Push current mode to stack or save it under name
  where       Shows the cli context you are in

switch(config-if-tcl)# description loop10
switch(config-if-tcl)# tclquit
Exiting Tcl
switch#

```

Tcl References

The following titles are provided for your reference:

- Mark Harrison (ed), *Tcl/Tk Tools*, O'Reilly Media, ISBN 1-56592-218-2, 1997
- Mark Harrison and Michael McLennan, *Effective Tcl/Tk Programming*, Addison-Wesley, Reading, MA, USA, ISBN 0-201-63474-0, 1998
- John K. Ousterhout, *Tcl and the Tk Toolkit*, Addison-Wesley, Reading, MA, USA, ISBN 0-201-63337-X, 1994.
- Brent B. Welch, *Practical Programming in Tcl and Tk*, Prentice Hall, Upper Saddle River, NJ, USA, ISBN 0-13-038560-3, 2003.
- J Adrian Zimmer, *Tcl/Tk for Programmers*, IEEE Computer Society, distributed by John Wiley and Sons, ISBN 0-8186-8515-8, 1998.



CHAPTER 10

iPXE

- [About iPXE, on page 115](#)
- [Netboot Requirements, on page 115](#)
- [Guidelines and Limitations for iPXE, on page 116](#)
- [Boot Mode Configuration, on page 116](#)
- [Verifying the Boot Order Configuration, on page 118](#)

About iPXE

iPXE is an open source network-boot firmware. iPXE is based on gPXE, which is an open-source PXE client firmware and bootloader derived from Etherboot. Standard PXE clients use TFTP to transfer data whereas gPXE supports more protocols.

Here is a list of additional features that iPXE provides over standard PXE:

- Boots from a web server via HTTP, iSCSI SAN, FCoE, and so on
- Supports both IPv4 and IPv6
- Netboot supports HTTP/TFTP, IPv4, and IPv6
- Supports embedded scripts into the image or served by the HTTP/TFTP, and so on
- Supports stateless address autoconfiguration (SLAAC) and stateful IP autoconfiguration variants for DHCPv6. iPXE supports boot URI and parameters for DHCPv6 options. This depends on IPv6 router advertisement.

In addition, we have disabled some of the existing features from iPXE for security reasons such as:

- Boot support for standard Linux image format such as bzImage+initramfs/initrd, or ISO, and so on
- Unused network boot options such as FCoE, iSCSI SAN, Wireless, and so on
- Loading of unsupported NBP (such as syslinux/pxelinux) because these can boot system images that are not properly code-signed.

Netboot Requirements

The primary requirements are:

- A DHCP server with proper configuration.
- A TFTP/HTTP server.
- Enough space on the device's bootflash because NX-OS downloads the image when the device is PXE booted.
- IPv4/IPv6 support—for better deployment flexibility

Guidelines and Limitations for iPXE

PXE has the following configuration guidelines and limitations:

- While autobooting through iPXE, there is a window of three seconds where you can enter **Ctrl+B** to exit out of the PXE boot. The system prompts you with the following options:

```
Please choose a bootloader shell:
1) . GRUB shell
2) . PXE shell
Enter your choice:
```

- HTTP image download vs. TFTP—TFTP is a UDP-based protocol, and it can be problematic if packet loss starts appearing. TCP is a window-based protocol and handles bandwidth sharing or losses better. As a result, TCP-based protocols support is more suitable given the sizes of the Cisco NX-OS images which are over 250 Mbytes.
- iPXE only allows or boots Cisco signed NBI images. Other standard-image format support is disabled for security reasons.
- On switches that have multiple supervisors, the behavior of supervisor A+ and B+ that are configured to PXE boot is different than the behavior of supervisor A or B.

When supervisor A+ or B+ is configured to boot from PXE boot first and bootflash second, the supervisor continuously attempts to boot from PXE and does not switch over to bootflash (GRUB) after unsuccessful PXE-boot retries. To boot from bootflash, the supervisor requires manual intervention to reload the supervisors.

You can interrupt PXE boot by entering **Ctrl+C**, and then you should get a prompt to stop PXE boot by entering **Ctrl+B**. The supervisors will then boot from bootflash after manually reloading them.

This limitation applies only to supervisor A+ and B+. In a similar configuration, supervisor A and B attempt to PXE boot four times before rebooting automatically and loading from bootflash.

Boot Mode Configuration

VSH CLI

```
switch# configure terminal
switch(conf)# boot order bootflash|pxe [bootflash|pxe]
switch(conf)# end
```



Note The keyword **bootflash** indicates it is Grub based booting.

For example, to do a PXE boot mode only, the configuration command is:

```
switch(conf)# boot order pxe
```

To boot Grub first, followed by PXE:

```
switch(conf)# boot order bootflash pxe
```

To boot PXE first, followed by Grub:

```
switch(conf)# boot order pxe bootflash
```



Note If you set **boot order pxe bootflash** on supervisor A+ or B+, the supervisor continually tries to PXE boot. Supervisor A+ or B+ does not switch over to boot from GRUB without manual intervention.

If you never use the **boot order** command, by default the boot order is Grub.



Note The following sections describe how you can toggle from Grub and iPXE.

Grub CLI

```
bootmode [-g|-p|-p2g|-g2p]
```

| Keyword | Function |
|-------------|--|
| -g | Grub only |
| -p | PXE only |
| -p2g | PXE first, followed by Grub if PXE failed |
| -g2p | Grub first, followed by PXE if Grub failed |

The Grub CLI is useful if you want to toggle the boot mode from the serial console without booting a full Cisco NX-OS image. It also can be used to get a box out of the continuous PXE boot state.

iPXE CLI

```
bootmode [-g|--grub] [-p|--pxe] [-a|--pxe2grub] [-b|--grub2pxe]
```

| Keyword | Function |
|---------------|-----------|
| --grub | Grub only |
| --pxe | PXE only |

| Keyword | Function |
|-------------|--|
| -- pxe2grub | PXE first, followed by Grub if PXE failed |
| -- grub2pxe | Grub first, followed by PXE if Grub failed |

The iPXE CLI is useful if you wish to toggle the boot mode from the serial console without booting a full Cisco NX-OS image. It also can be used to get a box out of continuous PXE boot state.

Verifying the Boot Order Configuration

To display boot order configuration information, enter the following command:

| Command | Purpose |
|------------------------------|--|
| <code>show boot order</code> | Displays the current boot order from the running configuration and the boot order value on the next reload from the startup configuration. |



CHAPTER 11

Kernel Stack

- [About Kernel Stack, on page 119](#)
- [Guidelines and Limitations, on page 119](#)
- [Changing the Port Range, on page 120](#)
- [About VXLAN with kstack, on page 121](#)
- [Netdevice Property Changes, on page 122](#)

About Kernel Stack

Kernel Stack (kstack) uses well known Linux APIs to manage the routes and front panel ports.

Open Containers, like the Guest Shell, are Linux environments that are decoupled from the host software. You can install or modify software within that environment without impacting the host software packages.

Kernel Stack has the following features:

Guidelines and Limitations

- Guest shell, Docker containers, and the host Bash Shell use Kernel Stack (kstack).
- The Guest Shell and the host Bash Shell start in the default network namespace. Docker containers start in the management network namespace by default.
 - Other network namespaces may be accessed by using the **setns** system call
 - The **nsenter** and **ip netns exec** utilities can be used to execute within the context of a different network namespace.
- The interface state may be read from `/proc/net/dev` or retrieved using other typical Linux utilities such as **ip**, **ifconfig**, or **netstat**. The counters are for packets that have initiated or terminated on the switch.
- **ethtool -S** may be used to get extended statistics from the net devices, which includes packets that are switched through the interface.
- Packet capture applications like **tcpdump** may be run to capture packets that are initiated from or terminated on the switch.

- There is no support for networking state changes (interface creation or deletion, IP address configuration, MTU change, and so on) from the Guest Shell.
- IPv4 and IPv6 are supported.
- Raw PF_PACKET is supported.
- Only on stack (Netstack or kstack) at a time can use well-known ports (0-15000), regardless of the network namespace.
- There is no IP connectivity between applications using Nestack and applications running kstack on the same switch. This limitation holds true regardless of whether the kstack applications are being run from the host Bash Shell or within a container.
- Applications within the Guest Shell are not allowed to send packets directly over an Ethernet out-of-band channel (EOBC) interface to communicate with the line cards or standby Sup.
- The management interface (mgmt0) is represented as eth1 in the kernel netdevices.
- Use of the VXLAN overlay interface (NVE x) is not supported for applications utilizing the kernel stack. NX-OS features, including CLI commands, are able to use this interface via netstack.

For more information about the NVE interface, see the [Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide](#).

Changing the Port Range

Netstack and kstack divide the port range between them. The default port ranges are as follows:

- Kstack—15001 to 58000
- Netstack—58001 to 65535



Note Within this range 63536 to 65535 are reserved for NAT.



Note The ports configured with **nxapi use-vrf management** uses kstack and are accessible.

SUMMARY STEPS

1. [no] sockets local-port-range *start-port end-port*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | <code>[no] sockets local-port-range start-port end-port</code> | This command modifies the port range for kstack. This command does not modify the Netstack range. |

Example

The following example sets the kstack port range:

```
switch# sockets local-port-range 15001 25000
```

What to do next

After you have entered the command, be aware of the following issues:

- Reload the switch after entering the command.
- Leave a minimum of 7000 ports unallocated which are used by Netstack.
- Specify the *start-port* as 15001 or the *end-port* as 65535 to avoid holes in the port range.

About VXLAN with kstack

Starting with NX-OS 9.2(1), VXLAN EVPN is supported with kstack to be leveraged by third-party applications. This functionality is supported on the Cisco Nexus 9000 ToR switches.

Setting Up VXLAN for kstack

No additional configuration is required to make the interfaces or network namespaces for VXLAN EVPN accessible to the third-party applications. The VXLAN EVPN routes are programmed automatically in the kernel based on the NX-OS VXLAN EVPN configuration. For more information, see the "Configuring VXLAN BGP EVPN" chapter in the *Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide*.

Troubleshooting VXLAN with kstack

To troubleshoot VXLAN issues, enter the following command to list several critical pieces of information to be collected:

```
switch(config)# show tech-support kstack
```

- Run the **ip route show** command:

```
root@switch(config)# run bash sudo su-
root@switch# ip netns exec evpn-tenant-kk1 ip route show
```

Output similar to the following appears:

```
10.160.1.0/24 dev Vlan1601 proto kernel scope link src 10.160.1.254
10.160.1.1 dev veth1-3 proto static scope link metric 51
10.160.2.0/24 dev Vlan1602 proto kernel scope link src 10.160.2.253
127.250.250.1 dev veth1-3 proto static scope link metric 51
```

Verify that all EVPN routes for the corresponding VRF are present in the kernel.

- Run the **ip neigh show** command:

```
root@switch(config)# run bash sudo su-
root@switch# ip netns exec evpn-tenant-kk1 ip neigh show
```

Output similar to the following appears:

```
10.160.1.1 dev veth1-3 lladdr 0c:75:bd:07:b4:33 PERMANENT
127.250.250.1 dev veth1-3 lladdr 0c:75:bd:07:b4:33 PERMANENT
```

Netdevice Property Changes

Starting with the NX-OS 9.2(2) release, netdevices representing the front channel port interfaces are always in the ADMIN UP state. The final, effective state is determined by the link carrier state.

The following example shows the following interfaces in NX-OS, where eth1/17 is shown as **up** and eth1/1 is shown as **down**:

```
root@kstack-switch# sh int ethernet 1/17 brief
Eth1/17      --      eth  routed up      none                1000(D) -

root@kstack-switch# sh int ethernet 1/1 brief
Eth1/1       --      eth  routed down    Link not connected  auto(D) -
```

The following example shows these same interfaces, but this time as shown in the Bash shell using the **ip link show** command:

```
bash-4.3# ip link show Eth1-17
49: Eth1-17: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP mode
DEFAULT group default qlen 100
    link/ether 00:42:68:58:f8:eb brd ff:ff:ff:ff:ff:ff

bash-4.3# ip link show Eth1-1
33: Eth1-1: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast state DOWN mode
DEFAULT group default qlen 100
    link/ether 00:42:68:58:f8:eb brd ff:ff:ff:ff:ff:ff
```

In this example, Eth1-1 is shown as being **UP**, but is shown as **NO-CARRIER** and **state DOWN**.

The following example shows these same interfaces, but this time as shown in the Bash shell using the **ifconfig** command:

```
bash-4.3# ifconfig Eth1-17
Eth1-17  Link encap:Ethernet  HWaddr 00:42:68:58:f8:eb
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:7388 errors:0 dropped:0 overruns:0 carrier:0
```



```
collisions:0 txqueuelen:100
RX bytes:0 (0.0 B) TX bytes:1869164 (1.7 MiB)

bash-4.3# ifconfig Eth1-1
Eth1-1 Link encap:Ethernet HWaddr 00:42:68:58:f8:eb
inet addr:99.1.1.1 Bcast:99.1.1.255 Mask:255.255.255.0
UP BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:100
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
```

The output from the **ifconfig** command provides different information, where the **RUNNING** keyword is used to represent the final state. By default, all netdevices show the keyword **UP**, which represents the ADMIN state of the netdevice in the kernel.

Following are the changes that are part of the NX-OS 9.2(2) release:

- **IPv4 address on netdevices** — Before the NX-OS 9.2(2) release, the IPv4 address would be plumbed to the netdevice in the kernel even when the corresponding interface in NX-OS was in the **DOWN** state. Starting with the NX-OS 9.2(2) release, the IPv4 address are plumbed to the kernel space only when the interface is in the **UP** state. Once plumbed, the IPv4 address continues to stay with the netdevice in the kernel even if the interface goes **DOWN**. It will be removed only after you have entered the following CLI command to explicitly remove the IP address from the NX-OS interface:

```
Interface Eth1/1
  no ip address IP-address
```

- **IPv6 address on netdevices** — Before the NX-OS 9.2(2) release, the IPv6 address would get flushed from the netdevices in the kernel when the interface was **DOWN**. Starting with the NX-OS 9.2(2) release, the netdevices are always in the admin **UP** state, so the IPv6 addresses will not get flushed from the kernel when the interface goes down.



PART II

Applications

- [Third-Party Applications, on page 127](#)
- [Using Ansible with Cisco NX-OS, on page 145](#)
- [Puppet Agent, on page 147](#)
- [Using Chef Client with Cisco NX-OS, on page 151](#)
- [Nexus Application Development - Yocto, on page 155](#)
- [Nexus Application Development - SDK, on page 159](#)
- [NX-SDK, on page 167](#)
- [Using Docker with Cisco NX-OS, on page 173](#)



CHAPTER 12

Third-Party Applications

- [About Third-Party Applications, on page 127](#)
- [Guidelines and Limitations, on page 127](#)
- [Installing Python2 and Dependent Packages, on page 128](#)
- [Installing Third-Party Native RPMs/Packages, on page 128](#)
- [Installing Signed RPM, on page 130](#)
- [Persistent Third-Party RPMs, on page 135](#)
- [Installing RPM from VSH, on page 136](#)
- [Third-Party Applications, on page 141](#)

About Third-Party Applications

The RPMs for the Third-Party Applications are available in the repository at https://devhub.cisco.com/artifactory/open-nxos/7.0-3-12-1/x86_64/https://devhub.cisco.com/artifactory/open-nxos/9.2.1/. These applications are installed in the native host by using the **dnf** command in the Bash shell or through the NX-OS CLI.

When you enter the **dnf install rpm** command, a Cisco **DNF** plug-in gets executed. This plug-in copies the RPM to a hidden location. On switch reload, the system reinstalls the RPM.

For configurations in `/etc`, a Linux process, **incron**, monitors artifacts that are created in the directory and copies them to a hidden location, which gets copied back to `/etc`.

Guidelines and Limitations

RPMs for the third-party applications have the following guidelines and limitations:

- Starting with Cisco NX-OS Release 9.2(1), the Cisco repository where agents are stored is now located at <https://devhub.cisco.com/artifactory/open-nxos/9.2.1/>. All RPMs hosted in this repository are signed with the release key.
- The NX-OS 10.1(1) release has a new operating system and roots, based on NX-Linux(Cisco's proprietary Linux distribution), so third-party RPMs that were built using WRL5/WRL8 might not be compatible with NX-Linux, so the third-party software might not work. In this case, remove old versions of your apps used with previous releases and replace them with new software that is compatible with NX-Linux, which is available in the repository at <https://devhub.cisco.com/artifactory/open-nxos/10.1.1/>.

- Guidelines and instructions for installing signed RPMs are provided in the *Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide, Release 9.2(x)*, including DNF and VSH CLI options for managing RPMs, signed and nonsigned RPM installations, the clean-up of repositories, and so on.
- The third-party applications are started during switch startup. It is possible that a third-party application could be started before its communication interface is up, or before the routing between the switch and any communication peer or server is established. Therefore, all third-party applications should be written to be robust in case of communication failure, and the application should retry establishing the connection. If an application is not resilient in the presence of a communication failure, a “wrapper” application might be required to establish that any communication peer is reachable before starting the desired application, or restart the desired application if necessary.
- Beginning with Cisco NX-OS Release 10.2(3)F, Python2 and dependent RPMs are removed from NX-OS. However, you can install Python2 and dependent RPMs from devhub site as package group `packagegroup-nxos-64-python-2-deprecated-rpms`.

Installing Python2 and Dependent Packages

The following is the complete workflow of package installation:

```
switch# cat /etc/dnf/repos.d/open-nxos.repo
[open-nxos]
name=open-nxos
baseurl=https://devhub.cisco.com/artifactory/open-nxos/10.2.3/
enabled=1
gpgcheck=0
sslverify=0

dnf info packagegroup-nxos-64-python-2-deprecated-rpms
dnf install packagegroup-nxos-64-python-2-deprecated-rpms
The output of these cmds will be available post KR3F CCO.
```

Installing Third-Party Native RPMs/Packages

The complete workflow of package installation is as follows:

Procedure

Configure the repository on the switch to point to the Cisco repository where agents are stored.

```
bash-4.2# cat /etc/dnf/repos.d/open-nxos.repo
[open-nxos]
name=open-nxos
baseurl=https://devhub.cisco.com/artifactory/open-nxos/7.0-3-I2-1/x86_64/
baseurl=https://devhub.cisco.com/artifactory/open-nxos/9.2.1/
baseurl=https://devhub.cisco.com/artifactory/open-nxos/10.1.1/
enabled=1
gpgcheck=0
sslverify=0
```

Instructions for using the CLIs to import the digital signature are available in the section "Using Install CLIs for Digital Signature Support" in the *Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide, Release 9.2(x)*.

An example of installation of an RPM using *dnf*, with full install log.

Example:

```
bash-4.2# dnf install splunkforwarder
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching, protect-packages
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package splunkforwarder.x86_64 0:6.2.3-264376 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch                Version              Repository            Size
=====
Installing:
splunkforwarder        x86_64              6.2.3-264376        open-nxos             13 M

Transaction Summary
=====
Install                1 Package

Total size: 13 M
Installed size: 34 M
Is this ok [y/N]: y
Downloading Packages:
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : splunkforwarder-6.2.3-264376.x86_64
                                                    1/1

complete

Installed:
  splunkforwarder.x86_64 0:6.2.3-264376

Complete!
bash-4.2#
```

An example of querying the switch for successful installation of the package, and verifying that its processes or services are up and running.

Example:

```
bash-4.2# dnf info splunkforwarder
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching, protect-packages
Fretta                | 951 B      00:00 ...
groups-repo           | 1.1 kB    00:00 ...
localdb               | 951 B     00:00 ...
patching              | 951 B     00:00 ...
thirdparty            | 951 B     00:00 ...
Installed Packages
Name      : splunkforwarder
Arch     : x86_64
Version  : 6.2.3
Release  : 264376
Size     : 34 M
Repo     : installed
From repo : open-nxos
```

```
Summary      : SplunkForwarder
License      : Commercial
Description  : The platform for machine data.
```

Installing Signed RPM

Checking a Signed RPM

Run the following command to check if a given RPM is signed or not.

```
Run, rpm -K rpm_file_name
```

Not a Signed RPM

```
bash-4.2# rpm -K bgp-1.0.0-r0.lib32_n9000.rpm
bgp-1.0.0-r0.lib32_n9000.rpm: (sha1) dsa sha1 md5 OK
```

Signed RPM

```
bash-4.2# rpm -K puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm
puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm: RSA sha1 MD5 NOT_OK
bash-4.2#
```

Signed third-party RPM requires public GPG key to be imported first before the package can be installed otherwise **yum** throws the following error:

```
bash-4.2# yum install puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm -q
Setting up Install Process
warning: rpmts_HdrFromFdno: Header V4 RSA/SHA1 signature: NOKEY, key ID 4bd6ec30
Cannot open: puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm. Skipping.
Error: Nothing to do
```

Installing Signed RPMs by Manually Importing Key

- Copy the GPG keys to `/etc rootfs` so that they are persisted across reboots.

```
bash-4.2# mkdir -p /etc/pki/rpm-gpg
bash-4.2# cp -f RPM-GPG-KEY-puppetlabs /etc/pki/rpm-gpg/
```

- Import the keys using the following command.

```
bash-4.2# rpm --import /etc/pki/rpm-gpg/RPM-GPG-KEY-puppetlabs
```



```

bash-4.2#
bash-4.2# rpm -q gpg-pubkey
gpg-pubkey-4bd6ec30-4c37bb40
bash-4.2# rpm --import /etc/pki/rpm-gpg/RPM-GPG-KEY-puppetlabs
bash-4.2#
bash-4.2# rpm -q gpg-pubkey
gpg-pubkey-4bd6ec30-4c37bb40

```

- Install the signed RPM with `yum` command

```

bash-4.2#
yum install puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm

Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
protect-packages

groups-repo          | 1.1 kB      00:00 ...
.
localdb              | 951 B       00:00 ...

patching             | 951 B       00:00 ...

thirdparty           | 951 B       00:00 ...

Setting up Install Process

Examining puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm:
puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64

Marking puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm to be installed

Resolving Dependencies

--> Running transaction check
---> Package puppet-enterprise.x86_64 0:3.7.1.rc2.6.g6cdc186-1.pe.nxos will be installed

--> Finished Dependency ResolutionDependencies Resolved

=====

Package              Arch      Version                                Repository
Size

=====

Installing:

puppet-enterprise    x86_64    3.7.1.rc2.6.g6cdc186-1.pe.nxos        /puppet-enterprise-
46 M                                                         3.7.1.rc2.6.g6cdc186-1.

```

pe.nxos.x86_64

Transaction Summary

```
=====
```

Install 1 Package

Total size: 46 M

Installed size: 46 M

Is this ok [y/N]: y

Downloading Packages:

Running Transaction Check

Running Transaction Test

Transaction Test Succeeded

Running Transaction

Installing : puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64

1/1

Installed:

puppet-enterprise.x86_64 0:3.7.1.rc2.6.g6cdc186-1.pe.nxos

Complete!

bash-4.2#

Installing Signed Third-Party RPMs by Importing Keys Automatically

Set up the yum repo to point to the keys and RPM.

```
root@switch# cat /etc/yum/repos.d/puppet.repo

[puppet]
name=Puppet RPM
baseurl=file:///bootflash/puppet
enabled=1
gpgcheck=1
gpgkey=http://yum.puppetlabs.com/RPM-GPG-KEY-puppetlabs
metadata_expire=0
cost=500
```

```
bash-4.2# yum install puppet-enterprise
```

```
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
protect-packages
```

```
groups-repo                | 1.1 kB    00:00 ...
localdb                    | 951 B     00:00 ...
patching                   | 951 B     00:00 ...
puppet                    | 951 B     00:00 ...
thirdparty                 | 951 B     00:00 ...
```

```
Setting up Install Process
```

```
Resolving Dependencies
```

```
--> Running transaction check
```

```
---> Package puppet-enterprise.x86_64 0:3.7.1.rc2.6.g6cdc186-1.pe.nxos will be installed
```

```
--> Finished Dependency Resolution
```

```
Dependencies Resolved
```

```
=====
```

```
Package                    Arch      Version                               Repository      Size
```

```
=====
```

```
Installing:
```

```
puppet-enterprise         x86_64    3.7.1.rc2.6.g6cdc186-1.pe.nxos      puppet          14 M
```

```
Transaction Summary
```

```
=====
```

```
Install                1 Package
```

```
Total download size: 14 M
```

```
Installed size: 46 M
```

```
Is this ok [y/N]: y
```

```
Retrieving key from file:///bootflash/RPM-GPG-KEY-puppetlabs
```

```
Importing GPG key 0x4BD6EC30:
```

```
  Userid: "Puppet Labs Release Key (Puppet Labs Release Key) <info@puppetlabs.com>"
```

```
  From   : /bootflash/RPM-GPG-KEY-puppetlabs
```

```
Is this ok [y/N]: y
```

```
Downloading Packages:
```

```

Running Transaction Check

Running Transaction Test

Transaction Test Succeeded

Running Transaction

Warning! Standby is not ready. This can cause RPM database inconsistency.

If you are certain that standby is not booting up right now, you may proceed.

Do you wish to continue?

Is this ok [y/N]: y

Warning: RPMDB altered outside of yum.

Installing : puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64           1/1

/sbin/ldconfig: /usr/lib/libboost_regex.so.1.49.0 is not a symbolic link

Installed:

puppet-enterprise.x86_64 0:3.7.1.rc2.6.g6cdc186-1.pe.nxos

Complete!

```

Adding Signed RPM into Repo

Procedure

-
- Step 1** Copy signed RPM to the repo directory
- Step 2** Import the corresponding key for the create repo to succeed.

```

bash-4.2# ls
puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm  RPM-GPG-KEY-puppetlabs
bash-4.2#
bash-4.2# rpm --import RPM-GPG-KEY-puppetlabs
bash-4.2# createrepo .
1/1 - puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata
bash-4.2#

```

Without importing keys

```

bash-4.2# ls
puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm  RPM-GPG-KEY-puppetlabs
bash-4.2#
bash-4.2# createrepo .
warning: rpmts_HdrFromFdno: Header V4 RSA/SHA1 signature: NOKEY, key ID 4bd6ec30

Error opening package - puppet-enterprise-3.7.1.rc2.6.g6cdc186-1.pe.nxos.x86_64.rpm

Saving Primary metadata

```

```
Saving file lists metadata
Saving other metadata
```

Step 3 Create repo config file under `/etc/yum/repos.d` pointing to this repo.

```
bash-4.2# cat /etc/yum/repos.d/puppet.repo
[puppet]
name=Puppet RPM
baseurl=file:///bootflash/puppet
enabled=1
gpgcheck=1
gpgkey=file:///bootflash/puppet/RPM-GPG-KEY-puppetlabs
#gpgkey=http://yum.puppetlabs.com/RPM-GPG-KEY-puppetlabs
metadata_expire=0
cost=500

bash-4.2# yum list available puppet-enterprise -q
Available Packages
puppet-enterprise.x86_64          3.7.1.rc2.6.g6cdc186-1.pe.nxos
                               puppet
bash-4.2#
```

Persistent Third-Party RPMs

The following is the logic behind persistent third-party RPMs:

- A local **dnf** repository is dedicated to persistent third-party RPMs. The `/etc/yum/repos.d/thirdparty.repo` points to `/bootflash/.rpmstore/thirdparty`.
- Whenever you enter the **dnf install third-party.rpm** command, a copy of the RPM is saved in `//bootflash/.rpmstore/thirdparty`.
- During a reboot, all the RPMs in the third-party repository are reinstalled on the switch.
- Any change in the `/etc` configuration files persists under `/bootflash/.rpmstore/config/etc` and they are replayed during boot on `/etc`.
- Any script that is created in the `/etc` directory persists across reloads. For example, a third-party service script that is created under `/etc/init.d/` brings up the apps during a reload.



Note The rules in iptables are not persistent across reboots when they are modified in a bash-shell.

To make the modified iptables persistent, see [Making an Iptable Persistent Across Reloads, on page 247](#).

Installing RPM from VSH

Package Addition

NX-OS feature RPMs can also be installed by using the VSH CLIs.

SUMMARY STEPS

1. **show install package**
2. **install add ?**
3. **install add rpm-packagename**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|------------------------------------|---|
| Step 1 | show install package | Displays the packages and versions that already exist. |
| Step 2 | install add ? | Determine supported URIs. |
| Step 3 | install add rpm-packagename | The install add command copies the package file to a local storage device or network server. |

Example

The following example shows how to activate the Chef RPM:

```
switch# show install package
switch# install add ?
WORD          Package name
bootflash:   Enter package uri
ftp:         Enter package uri
http:        Enter package uri
modflash:    Enter package uri
scp:         Enter package uri
sftp:        Enter package uri
tftp:        Enter package uri
usb1:        Enter package uri
usb2:        Enter package uri
volatile:    Enter package uri
switch# install add
bootflash:chef-12.0.0alpha.2+20150319234423.git.1608.b6eb10f-1.e15.x86_64.rpm
[#####] 100%
Install operation 314 completed successfully at Thu Aug 6 12:58:22 2015
```

What to do next

When you are ready to activate the package, go to [Package Activation, on page 137](#).



Note Adding and activating an RPM package can be accomplished in a single command:

```
switch#
install add bootflash:chef-12.0.0alpha.2+20150319234423.git.1608.b6eb10f-1.e15.x86_64.rpm
activate
```

Package Activation

Before you begin

The RPM has to have been previously added.

SUMMARY STEPS

1. **show install inactive**
2. **install activate** *rpm-packagename*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | show install inactive | Displays the list of packages that were added and not activated. |
| Step 2 | install activate <i>rpm-packagename</i> | Activates the package. |

Example

The following example shows how to activate a package:

```
switch# show install inactive
Boot image:
  NXOS Image: bootflash:///yumcli6.bin

Inactive Packages:
  sysinfo-1.0.0-7.0.3.x86_64
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching,
               : protect-packages
Available Packages
chef.x86_64          12.0.0alpha.2+20150319234423.git.1608.b6eb10f-1.e15  thirdparty
eigrp.lib32_n9000  1.0.0-r0                                             groups-rep
o
sysinfo.x86_64     1.0.0-7.0.3                                         patching
switch# install activate chef-12.0-1.e15.x86_64.rpm
[#####] 100%
Install operation completed successfully at Thu Aug  6 12:46:53 2015
```

Deactivating Packages

SUMMARY STEPS

1. `install deactivate package-name`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|------------------------------|
| Step 1 | <code>install deactivate package-name</code> | Deactivates the RPM package. |

Example

The following example shows how to deactivate the Chef RPM package:

```
switch# install deactivate chef
```

Removing Packages

Before you begin

Deactivate the package before removing it. Only deactivated RPM packages can be removed.

SUMMARY STEPS

1. `install remove package-name`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|--------------------------|
| Step 1 | <code>install remove package-name</code> | Removes the RPM package. |

Example

The following example shows how to remove the Chef RPM package:

```
switch# install remove chef-12.0-1.e15.x86_64.rpm
```


Displaying Installed Packages

SUMMARY STEPS

1. show install packages

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|-----------------------|--|
| Step 1 | show install packages | Displays a list of the installed packages. |

Example

The following example shows how to display a list of the installed packages:

```
switch# show install packages
```

Displaying Detail Logs

SUMMARY STEPS

1. show tech-support install

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|---------------------------|---------------------------|
| Step 1 | show tech-support install | Displays the detail logs. |

Example

The following example shows how to display the detail logs:

```
switch# show tech-support install
```

Upgrading a Package

SUMMARY STEPS

1. install add *package-name* activate upgrade

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|--------------------|
| Step 1 | install add <i>package-name</i> activate upgrade | Upgrade a package. |

Example

The following example shows how to upgrade a package:

```
switch# install add bootflash:bgp-1.0.1-r0.lib32_n9000.rpm activate ?
downgrade Downgrade package
forced Non-interactive
upgrade Upgrade package
switch# install add bootflash:bgp-1.0.1-r0.lib32_n9000.rpm activate upgrade
[#####] 100%
Install operation completed successfully at Thu Aug 6 12:46:53 2015
```

Downgrading a Package

SUMMARY STEPS

1. install add *package-name* activate downgrade

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|----------------------|
| Step 1 | install add <i>package-name</i> activate downgrade | Downgrade a package. |

Example

The following example shows how to downgrade a package:

```
switch# install add bootflash:bgp-1.0.1-r0.lib32_n9000.rpm activate ?
downgrade Downgrade package
forced Non-interactive
upgrade Upgrade package
switch# install add bootflash:bgp-1.0.1-r0.lib32_n9000.rpm activate downgrade
[#####] 100%
Install operation completed successfully at Thu Aug 6 12:46:53 2015
```

Third-Party Applications

NX-OS

For more information about the Cisco NX-OS repository for other third-party applications, see https://devhub.cisco.com/artifactory/open-nxos/7.0-3-12-1/x86_64/

For more information about NX-API REST API object model specifications, see <https://developer.cisco.com/media/dme/index.html>

DevOps Configuration Management Tools

For DevOps configuration management tools, refer to the following links:

- Ansible 2.0 Release(Nexus Support), [Ansible Release Index](#)
- Ansible NX-OS Sample Modules, [Ansible NX-OS Sample Modules](#)
- Puppet, [Puppet Forge Cisco Puppet](#)
- Cisco Puppet Module(Git), [Cisco Network Puppet Module](#)
- Chef, [Chef Supermarket Cisco Cookbook](#)
- Cisco Chef Cookbook(Git), [Cisco Network Chef Cookbook](#)

V9K

To download a virtual Nexus 9000 switch, for an ESX5.1/5.5, VirtualBox, Fusion, and KVM, go to <https://software.cisco.com/portal/pub/download/portal/select.html?&mdfid=286312239&flowid=81422&softwareid=282088129>.

Automation Tool Educational Content

For a free book on Open NX-OS architecture and automation, see http://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/nexus9000/sw/open_nxos/programmability/guide/Programmability_Open_NX-OS.pdf

collectd

collectd is a daemon that periodically collects system performance statistics and provides multiple means to store the values, such as RRD files. Those statistics can then be used to find current performance bottlenecks (for example, performance analysis) and predict future system load (that is, capacity planning).

For additional information, see <https://collectd.org>.

Ganglia

Ganglia is a scalable distributed monitoring system for high-performance computing systems such as clusters and grids. It is based on a hierarchical design that is targeted at federations of clusters. It leverages widely used technologies such as XML for data representation, XDR for compact, portable data transport, and RRDtool for data storage and visualization. It uses engineered data structures and algorithms to achieve low per-node overheads and high concurrency. The implementation is robust, has been ported to an extensive set of operating systems and processor architectures, and is currently in use on thousands of clusters around the world. It has been used to link clusters across university campuses and around the world and can scale to handle clusters with 2000 nodes.

For additional information, see <http://ganglia.info>.

Iperf

Iperf was developed by NLANR/DAST to measure maximum TCP and UDP bandwidth performance. Iperf allows the tuning of various parameters and UDP characteristics. Iperf reports bandwidth, delay jitter, datagram loss.

For additional information, see <http://sourceforge.net/projects/iperf/> or <http://iperf.sourceforge.net>.

LLDP

The link layer discover protocol (LLDP) is an industry standard protocol that is designed to supplant proprietary link layer protocols such as EDP or CDP. The goal of LLDP is to provide an intervendor compatible mechanism to deliver link layer notifications to adjacent network devices.

For more information, see <https://vincentbernat.github.io/lldpd/index.html>.

Nagios

Nagios is open source software that monitors the following through the Nagios remote plug-in executor (NRPE) and through SSH or SSL tunnels:

- Network services through ICMP, SNMP, SSH, FTP, HTTP, and so on
- Host resources, such as CPU load, disk usage, system logs, and so on
- Alert services for servers, switches, applications
- Services

For more information, see <https://www.nagios.org/>.

OpenSSH

OpenSSH is an open-source version of the SSH connectivity tools that encrypts all traffic (including passwords) to eliminate eavesdropping, connection hijacking, and other attacks. OpenSSH provides secure tunneling capabilities and several authentication methods, and supports all SSH protocol versions.

For more information, see <http://www.openssh.com>.

Quagga

Quagga is a network routing software suite that implements various routing protocols. Quagga daemons are configured through a network accessible CLI called a "vty."



Note Only Quagga BGP has been validated.

For more information, see <http://www.nongnu.org/quagga/>.

Splunk

Splunk is a web-based data collection, analysis, and monitoring tool that has search, visualization, and prepackaged content for use-cases. The raw data is sent to the Splunk server using the Splunk Universal Forwarder. Universal Forwarders provide reliable, secure data collection from remote sources and forward that data into the Splunk Enterprise for indexing and consolidation. They can scale to tens of thousands of remote systems, collecting terabytes of data with a minimal impact on performance.

For additional information, see http://www.splunk.com/en_us/download/universal-forwarder.html.

tcollector

tcollector is a client-side process that gathers data from local collectors and pushes the data to Open Time Series Database (OpenTSDB).

tcollector has the following features:

- Runs data collectors and collates the data.
- Manages connections to the time series database (TSD).
- Eliminates the need to embed TSD code in collectors.
- Deduplicates repeated values.
- Handles wire protocol work.

For additional information, see http://opentsdb.net/docs/build/html/user_guide/utilities/tcollector.html.

tcpdump

tcpdump is a CLI application that prints a description of the contents of packets on a network interface that match a Boolean expression. The description is preceded by a timestamp, printed, by default, as hours, minutes, seconds, and fractions of a second since midnight.

tcpdump can be run with the following flags:

- -w, which causes it to save the packet data to a file for later analysis.
- -r, which causes it to read from a saved packet file rather than to read packets from a network interface.
- -V, which causes it to read a list of saved packet files.

In all cases, tcpdump processes only the packets that match the expression.

For more information, see <http://www.tcpdump.org/manpages/tcpdump.1.html>.

TShark

TShark is a network protocol analyzer on the CLI. Tshark lets you capture packet data from a live network, or read packets from a previously saved capture file. You can print either a decoded form of those packets to the standard output or write the packets to a file. TShark's native capture file format is pcap, the format that is used by **tcpdump** and various other tools also. TShark can be used within the Guest Shell after removing the cap_net_admin file capability.

```
setcap
cap_net_raw=ep /sbin/dumppcap
```



Note This command must be run within the Guest Shell.

For more information, see <https://www.wireshark.org/docs/man-pages/tshark.html>.



CHAPTER 13

Using Ansible with Cisco NX-OS

- [Prerequisites](#), on page 145
- [About Ansible](#), on page 145
- [Cisco Ansible Module](#), on page 145

Prerequisites

Go to https://docs.ansible.com/ansible/latest/getting_started/index.html for installation requirements for supported control environments.

About Ansible

Ansible is an open-source IT automation engine that automates cloud provisioning, configuration management, application deployment, intraservice orchestration, and other IT needs.

Ansible uses small programs that are called Ansible modules to make API calls to your nodes, and apply configurations that are defined in playbooks.

By default, Ansible represents what machines it manages using a simple INI file that puts all your managed machines in groups of your own choosing.

More information can be found from Ansible:

| | |
|--|---|
| Ansible | https://www.ansible.com/ |
| Ansible Automation Solutions. Includes installation instructions, playbook instructions and examples, module lists, and so on. | https://docs.ansible.com/ |

Cisco Ansible Module

There are multiple Cisco NX-OS-supported modules and playbooks for Ansible, as per the following table of links:

| | |
|-------------------------------|--|
| NX-OS developer landing page. | Configuration Management Tools |
|-------------------------------|--|

| | |
|---------------------------------|---|
| Ansible NX-OS playbook examples | Repo for ansible nxos playbooks |
| Ansible NX-OS network modules | nxos network modules |



CHAPTER 14

Puppet Agent

This chapter contains the following topics:

- [About Puppet, on page 147](#)
- [Prerequisites, on page 148](#)
- [Puppet Agent NX-OS Environment, on page 148](#)
- [ciscopuppet Module, on page 148](#)

About Puppet

The Puppet software package, developed by Puppet Labs, is an open source automation toolset for managing servers and other resources. The Puppet software accomplishes server and resource management by enforcing device states, such as configuration settings.

Puppet components include a puppet agent which runs on the managed device (node) and a Puppet Primary (server). The Puppet Primary typically runs on a separate dedicated server and serves multiple devices. The operation of the puppet agent involves periodically connecting to the Puppet Primary, which in turn compiles and sends a configuration manifest to the agent. The agent reconciles this manifest with the current state of the node and updates state that is based on differences.

A puppet manifest is a collection of property definitions for setting the state on the device. The details for checking and setting these property states are abstracted so that a manifest can be used for more than one operating system or platform. Manifests are commonly used for defining configuration settings, but they also can be used to install software packages, copy files, and start services.

More information can be found from Puppet Labs:

| | |
|---------------------------|--|
| Puppet Labs | https://puppetlabs.com |
| Puppet Labs FAQ | https://puppet.com/blog/how-get-started-puppet-enterprise-faq/ https://puppet.com/products/faq |
| Puppet Labs Documentation | https://puppet.com/docs |

Prerequisites

The following are prerequisites for the Puppet Agent:

- You must have the required disk storage available on the device for virtual services installation and deployment of Puppet Agent.
 - A minimum of 450-MB free disk space on the bootflash.
- You must have a Puppet Primary server with Puppet 4.0 or later.
- You must have Puppet Agent 4.0 or later.

Puppet Agent NX-OS Environment

The Puppet Agent software must be installed on a switch in the Guest Shell (the Linux container environment running CentOS). The Guest Shell provides a secure, open execution environment that is decoupled from the host.

Starting with the Cisco NX-OS Release 9.2(1), the Bash-shell (native WindRiver Linux environment underlying Cisco NX-OS) install of Puppet Agent is no longer supported.

The following provides information about agent-software download, installation, and setup:

| | |
|---|---|
| Puppet Agent: Installation & Setup on Cisco Nexus switches (Manual Setup) | https://github.com/cisco/cisco-network-puppet-module/blob/develop/docs/README-agent-install.md |
|---|---|

ciscopuppet Module

The ciscopuppet module is a Cisco developed open-source software module. It interfaces between the abstract resources configuration in a puppet manifest and the specific implementation details of the Cisco NX-OS operating system and platform. This module is installed on the Puppet Primary and is required for puppet agent operation on Cisco Nexus switches.

The ciscopuppet module is available on Puppet Forge.

The following provide additional information about the ciscopuppet module installation procedures:

| | |
|--|---|
| ciscopuppet Module location (Puppet Forge) | Puppet Forge |
| Resource Type Catalog | Cisco Puppet Resource Reference |
| ciscopuppet Module: Source Code Repository | Cisco Network Puppet Module |
| ciscopuppet Module: Setup & Usage | Cisco Puppet Module::README.md |

| | |
|---------------------------------|---|
| Puppet Labs: Installing Modules | https://puppet.com/docs/puppet/7/modules_installing.html https://docs.puppetlabs.com/puppet/latest/reference/modules_installing.html |
| Puppet NX-OS Manifest Examples | Cisco Network Puppet Module Examples |
| NX-OS developer landing page. | Configuration Management Tools |



CHAPTER 15

Using Chef Client with Cisco NX-OS

- [About Chef, on page 151](#)
- [Prerequisites, on page 151](#)
- [Chef Client NX-OS Environment, on page 152](#)
- [cisco-cookbook, on page 152](#)

About Chef

Chef is an open-source software package that is developed by Chef Software, Inc. The software package is a systems and cloud infrastructure automation framework that deploys servers and applications to any physical, virtual, or cloud location, no matter the size of the infrastructure. Each organization consists of one or more workstations, a single server, and every node that the chef-client has configured and is maintaining. Cookbooks and recipes are used to tell the chef-client how each node should be configured. The chef-client, which is installed on every node, does the actual configuration.

A Chef cookbook is the fundamental unit of configuration and policy distribution. A cookbook defines a scenario and contains everything that is required to support that scenario, including libraries, recipes, files, and more. A Chef recipe is a collection of property definitions for setting state on the device. The details for checking and setting these property states are abstracted away so that a recipe may be used for more than one operating system or platform. While recipes are commonly used for defining configuration settings, they also can be used to install software packages, copy files, start services, and more.

The following references provide more information from Chef:

| Topic | Link |
|--------------------------|---|
| Chef home | https://www.chef.io |
| Chef overview | https://docs.chef.io/chef_overview.html |
| Chef documentation (all) | https://docs.chef.io/ |

Prerequisites

The following are prerequisites for Chef:

- You must have the required disk storage available on the device for Chef deployment:

- A minimum of 500 MB of free disk space on bootflash
- You need a Chef server with Chef 12.4.1 or higher.
- You need Chef Client 12.4.1 or higher.

Chef Client NX-OS Environment

The chef-client software must be installed on a switch in the Guest Shell (the Linux container environment running CentOS). This software provides a secure, open execution environment that is decoupled from the host.

Starting with the Cisco NX-OS Release 9.2(1), the Bash-shell (native WindRiver Linux environment underlying NX-OS) install of chef-client is no longer supported.

The following documents provide step-by-step guidance about agent-software download, installation, and setup:

| Topic | Link |
|---|---|
| Chef Client (Native) | Latest information on Client RPM is available here . |
| Chef Client (Guest Shell, CentOS7) | Latest information on Client RPM is available here . |
| Chef Client: Installation and setup on Cisco Nexus platform (manual setup) | cisco-cookbook::README-install-agent.md |
| Chef Client: Installation and setup on a switch (automated installation using the Chef provisioner) | cisco-cookbook::README-chef-provisioning.md |
| Cisco NX-OS developer landing page | https://www.in-github.cisco.com/agents/DevNet-Config-Management |

cisco-cookbook

cisco-cookbook is a Cisco-developed open-source interface between the abstract resources configuration in a Chef recipe and the specific implementation details of the switch. This cookbook is installed on the Chef Server and is required for proper Chef Client operation on switches.

The cisco-cookbook can be found on Chef Supermarket.

The following documents provide more detail for cisco-cookbook and generic cookbook installation procedures:

| Topic | Link |
|--|--|
| cisco-cookbook location | Chef Supermarket Cisco Cookbook |
| Resource Type Catalog | Resource Catalog (by Technology) |
| cisco-cookbook: Source Code Repository | Cisco Network Chef Cookbook |
| cisco-cookbook: Setup and usage | Chef Cookbook Setup and Usage |

| Topic | Link |
|------------------------------|---|
| Chef Supermarket | Chef Supermarket |
| Chef NX-OS Manifest Examples | Cisco Network Chef Cookbook Recipes |



CHAPTER 16

Nexus Application Development - Yocto

- [About Yocto, on page 155](#)
- [Installing Yocto, on page 155](#)

About Yocto

The Cisco NX-OS Release 10.1(1) software is based on Yocto 2.6. More applications can be installed by downloading Yocto 2.6, downloading the new software to be built, building the software, and installing the software on the switch.

Installing Yocto

In the example below, we are building Ruby version 2.2.2 in a Ubuntu 16.04 virtual machine.

Procedure

Step 1 Install all essential packages on the Ubuntu 16.04 virtual machine.

```
sudo apt-get install gawk wget git-core diffstat unzip texinfo gcc-multilib build-essential chrpath  
socat cpio python python3 python3-pip python3-pexpect xz-utils debianutils iputils-ping libssl1.2-dev  
xterm
```

Step 2 Download Yocto 2.6.

```
wget http://downloads.yoctoproject.org/releases/yocto/yocto-2.6/poky-thud-20.0.0.tar.bz2  
tar xjfv poky-thud-20.0.0.tar.bz2  
cd poky-thud-20.0.0
```

Step 3 Source the oe-init-build-env file.

```
source oe-init-build-env
```

Step 4 Use a text editor to edit conf/local.conf to add the following lines:

```
MACHINE = "genericx86-64"  
DEFAULTTUNE = "x86-64"
```

Step 5 Enter the following command:

```
bitbake ruby
```

After the build completes, the RPMs are present in tmp/deploy/rpm/x86_64/*.rpm.

Step 6 Copy the RPMs to the switch.

```
Switch# copy scp://<username>@<IP_address>/ruby-2.2.2-r0.x86_64.rpm bootflash: vrf management
use-kstack
Switch# copy scp://<username>@<IP_address>/libyaml-0-2-0.1.6-r0.x86_64.rpm bootflash: vrf management
use-kstack
```

Step 7 From the Bash shell, enter the following commands.

You will be entering **y** at one point in the install process.

```
bash-4.3# dnf install /bootflash/libyaml-0-2-0.1.6-r0.x86_64.rpm
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching, protect-packages
groups-repo          | 1.1 kB      00:00 ...
localdb              | 951 B       00:00 ...
patching             | 951 B       00:00 ...
thirdparty           | 951 B       00:00 ...
Setting up Install Process
Examining /bootflash/libyaml-0-2-0.1.6-r0.x86_64.rpm: libyaml-0-2-0.1.6-r0.x86_64
Marking /bootflash/libyaml-0-2-0.1.6-r0.x86_64.rpm to be installed
Resolving Dependencies
--> Running transaction check
---> Package libyaml-0-2.x86_64 0:0.1.6-r0 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch             Version           Repository         Size
=====
Installing:
libyaml-0-2            x86_64           0.1.6-r0          /libyaml-0-2-0.1.6-r0.x86_64 119 k

Transaction Summary
=====
Install                1 Package

Total size: 119 k
Installed size: 119 k
Is this ok [y/N]: y
Downloading Packages:
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : libyaml-0-2-0.1.6-r0.x86_64                                1/1
/sbin/ldconfig: /usr/lib/libboost_regex.so.1.49.0 is not a symbolic link

Installed:
  libyaml-0-2-0.1.6-r0.x86_64

Complete!
Install operation 2450 completed successfully at Fri Jul 27 18:54:55 2018.

[#####] 100%
```

Step 8 The following commands provide an example of building Ruby version 2.2.2 in a Ubuntu 16.04 virtual machine. You will be entering **y** at one point in the install process.

```
bash-4.3# dnf install /bootflash/ruby-2.2.2-r0.x86_64.rpm
Loaded plugins: downloadonly, importpubkey, localrpmDB, patchaction, patching, protect-packages
groups-repo          | 1.1 kB    00:00 ...
localdb              | 951 B    00:00 ...
patching             | 951 B    00:00 ...
thirdparty           | 951 B    00:00 ...
thirdparty/primary   | 1.8 kB    00:00 ...
thirdparty            |          2/2
Setting up Install Process
Examining /bootflash/ruby-2.2.2-r0.x86_64.rpm: ruby-2.2.2-r0.x86_64
Marking /bootflash/ruby-2.2.2-r0.x86_64.rpm to be installed
Resolving Dependencies
--> Running transaction check
---> Package ruby.x86_64 0:2.2.2-r0 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package            Arch            Version          Repository        Size
=====
Installing:
ruby                x86_64          2.2.2-r0         /ruby-2.2.2-r0.x86_64 32 M

Transaction Summary
=====
Install            1 Package

Total size: 32 M
Installed size: 32 M
Is this ok [y/N]: y
Downloading Packages:
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : ruby-2.2.2-r0.x86_64                                1/1
/sbin/ldconfig: /usr/lib/libboost_regex.so.1.49.0 is not a symbolic link

Installed:
  ruby.x86_64 0:2.2.2-r0

Complete!
Install operation 2451 completed successfully at Fri Jul 27 18:55:23 2018.

[#####] 100%
```




CHAPTER 17

Nexus Application Development - SDK

- [About the Cisco SDK, on page 159](#)
- [Installing the SDK, on page 159](#)
- [Procedure for Installation and Environment Initialization, on page 160](#)
- [Using the SDK to Build Applications, on page 161](#)
- [Using RPM to Package an Application, on page 162](#)
- [Creating an RPM Build Environment, on page 163](#)
- [Using General RPM Build Procedure, on page 163](#)
- [Example to Build RPM for collectd with No Optional Plug-Ins, on page 164](#)
- [Example to Build RPM for collectd with Optional Curl Plug-In, on page 165](#)

About the Cisco SDK

The Cisco SDK is a development kit that is based on Yocto 2.0. It contains all the tools to build applications for execution on a Cisco Nexus switch running the Cisco NX-OS Release 9.2(1). The basic components are the C cross-compiler, linker, libraries, and header files that are commonly used in many applications. The list is not exhaustive, and you might need to download and build any dependencies that are needed for any particular application. Some applications are ready to be downloaded and used from the Cisco devhub website and do not require building. The SDK can be used to build RPM packages which may be directly installed on a switch.

Installing the SDK

The following lists the system requirements:

- The SDK can run on most modern 64-bit x86_64 Linux systems. It has been verified on CentOS 7 and Ubuntu 14.04. Install and run the SDK under the Bash shell.
- The SDK includes binaries for both 32-bit and 64-bit architectures, so it must be run on an x86_64 Linux system that also has 32-bit libraries installed.

Procedure

Check if the 32-bit libraries are installed:

Example:

```
bash$ ls /lib/ld-linux.so.2
```

If this file exists, then 32-bit libraries should be installed already. Otherwise, install 32-bit libraries as follows:

- For CentOS 7:

```
bash$ sudo dnf install glibc.i686
```

- For Ubuntu 14.04:

```
bash$ sudo apt-get install gcc-multilib
```

Procedure for Installation and Environment Initialization

The SDK is available for download at: https://devhub.cisco.com/artifactory/open-nxos/10.0.1/nx-linux-x86_64-nxos-rootfs-n9k-sup-toolchain-1.1.0.sh.

This file is a self-extracting archive that installs the SDK into a directory of your choice. You are prompted for a path to an SDK installation directory.

```
bash$ ./wrlinux-8.0.0.25-glibc-x86_64-n9000-nxos-image-rpm-sdk-sdk.sh
Wind River Linux SDK installer version 8.0-n9000
=====
Enter target directory for SDK (default: /opt/windriver/wrlinux/8.0-n9000):
You are about to install the SDK to "/opt/windriver/wrlinux/8.0-n9000". Proceed[Y/n]? Y
Extracting
SDK.....done
Setting it up...done
SDK has been successfully set up and is ready to be used.

. environment-setup-corei7-64-nxos-linux
. environment-setup-corei7-32-nxosmllib32-linux

source environment-setup-corei7-64-nxos-linux
source environment-setup-corei7-32-nxosmllib32-linux
=====
```

Use the **source environment-setup-x86-wrsmlib32-linux** and **source environment-setup-x86_64-wrs-linux** commands to add the SDK-specific paths to your shell environment. Add the SDK-specific paths for each shell you intend to use with the SDK. Adding the SDK-specific paths is the key to setting up the SDK to use the correct versions of the build tools and libraries.

Procedure

Step 1 Browse to the installation directory.

Step 2 Enter the following commands at the Bash prompt:

```
bash$ source environment-setup-x86-wrsml1lib32-linux
bash$ source environment-setup-x86_64-wrs-linux
```

Using the SDK to Build Applications

Many of the common Linux build processes work for this scenario. Use the techniques that are best suited for your situation.

The source code for an application package can be retrieved in various ways. For example, you can get the source code either in tar file form or by downloading from a git repository where the package resides.

The following are examples of some of the most common cases.

(Optional) Verify that the application package builds using standard configure/make/make install.

```
bash$ tar --xvzf example-app.tgz
bash$ mkdir example-lib-install
bash$ cd example-app/
bash$ ./configure --prefix=/path/to/example-app-install
bash$ make
bash$ make install
```

Sometimes it is necessary to pass extra options to the `./configure` script, for example to specify which optional components and dependencies are needed. Passing extra options depends entirely on the application being built.

Example - Build Ganglia and its dependencies

In this example, we build ganglia, along with the third-party libraries that it requires - libexpat, libapr, and libconfuse.

libexpat

```
bash$ wget 'http://downloads.sourceforge.net/project/expat/expat/2.1.0/expat-2.1.0.tar.gz'
bash$ mkdir expat-install
bash$ tar xvzf expat-2.1.0.tar.gz
bash$ cd expat-2.1.0
bash$ ./configure --prefix=/home/sdk-user/expat-install
bash$ make
bash$ make install
bash$ cd ..
```

libapr

```
bash$ wget 'http://www.eu.apache.org/dist/apr/apr-1.5.2.tar.gz'
bash$ mkdir apr-install
bash$ tar xvzf apr-1.5.2.tar.gz
```

```
bash$ cd apr-1.5.2
bash$ ./configure --prefix=/home/sdk-user/apr-install
bash$ make
bash$ make install
bash$ cd ..
```

libconfuse



Note confuse requires the extra `--enable-shared` option to `./configure`, otherwise it builds a statically linked library instead of the required shared library.

```
bash$ wget 'http://savannah.nongnu.org/download/confuse/confuse-2.7.tar.gz'
bash$ mkdir confuse-install
bash$ tar xvzf confuse-2.7.tar.gz
bash$ cd confuse-2.7
bash$ ./configure --prefix=/home/sdk-user/confuse-install --enable-shared
bash$ make
bash$ make install
bash$ cd ..
```

ganglia



Note The locations to all the required libraries are passed to `./configure`.

```
bash$ wget
'http://downloads.sourceforge.net/project/ganglia/ganglia%20monitoring%20core/3.7.2/ganglia-3.7.2.tar.gz'
bash$ mkdir ganglia-install
bash$ tar xvzf ganglia-3.7.2.tar.gz
bash$ cd ganglia-3.7.2
bash$ ./configure --with-libexpat=/home/sdk-user/expat-install
--with-libapr=/home/sdk-user/apr-install/bin/apr-1-config
--with-libconfuse=/home/sdk-user/confuse-install --prefix=/home/sdk-user/ganglia-install
bash$ make
bash$ make install
bash$ cd ..
```

Using RPM to Package an Application

If the application successfully builds using "make", then it can be packaged into an RPM.



Note RPM and spec files

The RPM package format is designed to package up all files (binaries, libraries, configurations, documents, etc) that are needed for a complete install of the given application. The process of creating an RPM file is therefore somewhat non-trivial. To aid in the RPM build process, a `.spec` file is used that controls everything about the build process.



Note Many third-party applications are available on the internet in the form of source code packaged into tarballs. In many cases, these tarballs will include a .spec file to help with RPM build process. Unfortunately, many of these .spec files are not updated as frequently as the source code itself. Even worse, sometimes there is no spec file at all. In these cases the spec file may need editing or even creating from scratch so that RPMs can be built.

Creating an RPM Build Environment

Before using the SDK to build RPMs, an RPM build directory structure must be created, and some RPM macros set.

Procedure

Step 1 Create the directory structure:

```
bash$ mkdir rpmbuild
bash$ cd rpmbuild
bash$ mkdir BUILD RPMS SOURCES SPECS SRPMS
```

Step 2 Set the topdir macro to point to the directory structure created above:

```
bash$ echo "_topdir ${PWD}" > ~/.rpmmacros
```

Note This step assumes that the current user does not already have a .rpmmacros file that is already set up. If it is inconvenient to alter an existing .rpmmacros file, then the following may be added to all rpmbuild command lines:

```
--define "_topdir ${PWD}"
```

Step 3 Refresh the RPM DB:

```
bash$ rm /path/to/sdk/sysroots/x86_64-wrlinuxsdk-linux/var/lib/rpm/__db.*
bash$ rpm --rebuilddb
```

Note The rpm and rpmbuild tools in the SDK have been modified to use /path/to/sdk/sysroots/x86_64-wrlinuxsdk-linux/var/lib/rpm as the RPM database instead of the normal /var/lib/rpm. This modification prevents any conflicts with the RPM database for the host when not using the SDK and removes the need for root access. After SDK installation, the SDK RPM database must be rebuilt through this procedure.

Using General RPM Build Procedure

General RPM Build procedure is as follows:

```
bash$ wget --no-check-certificate --directory-prefix=SOURCES http://<URL of example-app tarball>
bash$ # determine location of spec file in tarball:
```

```

bash$ tar tf SOURCES/example-app.tar.bz2 | grep '.spec$'
bash$ tar xkvf SOURCES/example-app.tar.bz2 example-app/example-app.spec
bash$ mv example-app/example-app.spec SPECS/
bash$ rm -rf example-app
bash$ rpmbuild -v --bb SPECS/example-app.spec

```

The result is a binary RPM in RPMS/ that can be copied to the switch and installed. Installation and configuration of applications can vary. Refer to the application documents for those instructions.

This rpmbuild and installation on the switch is required for every software package that is required to support the application. If a software dependency is required that is not already included in the SDK, the source code must be obtained and the dependencies built. On the build machine, the package can be built manually for verification of dependencies. The following example is the most common procedure:

```

bash$ tar xkzf example-lib.tgz
bash$ mkdir example-lib-install
bash$ cd example-lib/
bash$ ./configure --prefix=/path/to/example-lib-install
bash$ make
bash$ make install

```

These commands place the build files (binaries, headers, libraries, and so on) into the installation directory. From here, you can use standard compiler and linker flags to pick up the location to these new dependencies. Any runtime code, such as libraries, are required to be installed on the switch also, so packaging required runtime code into an RPM is required.



Note There are many support libraries already in RPM form on the Cisco devhub website.

Example to Build RPM for collectd with No Optional Plug-Ins

Download source tarball and extract spec file:

```

bash$ wget --no-check-certificate --directory-prefix=SOURCES
https://collectd.org/files/collectd-5.5.0.tar.bz2
bash$ tar tf SOURCES/collectd-5.5.0.tar.bz2 | grep '.spec$'
collectd-5.5.0/contrib/redhat/collectd.spec
collectd-5.5.0/contrib/aix/collectd.spec
collectd-5.5.0/contrib/sles10.1/collectd.spec
collectd-5.5.0/contrib/fedora/collectd.spec
bash$ tar xkvf SOURCES/collectd-5.5.0.tar.bz2 collectd-5.5.0/contrib/redhat/collectd.spec
bash$ mv collectd-5.5.0/contrib/redhat/collectd.spec SPECS/
bash$ rm -rf collectd-5.5.0

```

There are four spec files in this tarball. The Red Hat spec file is the most comprehensive and is the only one that contains the correct collectd version. We will use it as an example.

This spec file sets the RPM up to use /sbin/chkconfig to install collectd. However on a switch, you will use the /usr/sbin/chkconfig instead. Edit the following edited in the spec file:

```

bash$ sed -r -i.bak 's%(^|\s)/sbin/chkconfig%\1/usr/sbin/chkconfig%' SPECS/collectd.spec

```

collectd has numerous optional plug-ins. This spec file enables many plug-ins by default. Many plug-ins have external dependencies, so options to disable these plug-ins must be passed to the **rpmbuild** command line. Instead of typing out one long command line, we can manage the options in a Bash array as follows:

```
bash$ rpmbuild_opts=()
bash$ for rmdep in \
> amqp apache ascent bind curl curl_xml dbi ipmi java memcached mysql nginx \
> notify_desktop notify_email nut openldap perl pinba ping postgresql python \
> rrdtool sensors snmp varnish virt write_http write_riemann
> do
>   rpmbuild_opts+=("--without")
>   rpmbuild_opts+=(${rmdep})
> done
bash$ rpmbuild_opts+=(--nodeps)
bash$ rpmbuild_opts+=(--define)
bash$ rpmbuild_opts+=("_unpackaged_files_terminate_build 0")
```

It is then passed to **rpmbuild** as follows to start the entire build and RPM package process:

```
bash$ rpmbuild "${rpmbuild_opts[@]}" -bb SPECS/collectd.spec
```

You can then find the resulting RPMs for *collectd* in the RPMS directory.

These RPM files can now be copied to the switch and installed from the switch Bash shell:

```
bash$ rpm --noparentdirs -i /bootflash/collectd-5.5.0-1.ia32e.rpm
```

Example to Build RPM for collectd with Optional Curl Plug-In

The *collectd curl* plug-in has *libcurl* as a dependency.

In order to satisfy this link dependency during the RPM build process, it is necessary to download and build *curl* under the SDK:

```
bash$ wget --no-check-certificate http://curl.haxx.se/download/curl-7.24.0.tar.gz
bash$ tar xkvf curl-7.24.0.tar.gz
bash$ cd curl-7.24.0
bash$ ./configure --without-ssl --prefix /path/to/curl-install
bash$ make
bash$ make install
bash$ cd ..
```



Note The *curl* binaries and libraries are installed to `/path/to/curl-install`. This directory will be created if it does not already exist, so you must have write permissions for the current user. Next, download the source tarball and extract the spec file. This step is exactly the same as in the *collectd* example for no plug-ins.

```
bash$ wget --no-check-certificate --directory-prefix=SOURCES
https://collectd.org/files/collectd-5.5.0.tar.bz2
bash$ tar tf SOURCES/collectd-5.5.0.tar.bz2 | grep '.spec$'
collectd-5.5.0/contrib/redhat/collectd.spec
collectd-5.5.0/contrib/aix/collectd.spec
collectd-5.5.0/contrib/sles10.1/collectd.spec
collectd-5.5.0/contrib/fedora/collectd.spec
bash$ tar xkvf SOURCES/collectd-5.5.0.tar.bz2 collectd-5.5.0/contrib/redhat/collectd.spec
```

```
bash$ mv collectd-5.5.0/contrib/redhat/collectd.spec SPECS/
bash$ rm -rf collectd-5.5.0
```

This spec file sets the RPM up to use `/sbin/chkconfig` to install collectd. However on a switch, you must use `/usr/sbin/chkconfig` instead, so the following can be edited in the spec file:



Note There are four spec files in this tarball. The Red Hat spec file is the most comprehensive, and it is the only one to contain the correct collectd version. We will use that one as an example.

```
bash$ sed -r -i.bak 's%(^|\s)/sbin/chkconfig%\1/usr/sbin/chkconfig%' SPECS/collectd.spec
```

Here a deviation from the previous example is encountered. The collectd rpmbuild process needs to know the location of libcurl. Edit the collectd spec file to add the following.

Find the string `%configure` in `SPECS/collectd.spec`. This line and those following it define the options that rpmbuild will pass to the `./configure` script.

Add the following option:

```
--with-libcurl=/path/to/curl-install/bin/curl-config \
```

Next a Bash array is built again to contain the rpmbuild command options. Note the following differences:

- `curl` is removed from the list of plug-ins not to be built
- The addition of `--with curl=force`

```
bash$ rpmbuild_opts=()
bash$ for rmdep in \
> amqp apache ascent bind curl_xml dbi ipmi java memcached mysql nginx \
> notify_desktop notify_email nut openldap perl pinba ping postgresql python \
> rrdtool sensors snmp varnish virt write_http write_riemann
> do
>   rpmbuild_opts+=("--without")
>   rpmbuild_opts+=(${rmdep})
> done
bash$ rpmbuild_opts+=("--with")
bash$ rpmbuild_opts+=("curl=force")bash$ rpmbuild_opts+=(--nodeps)
bash$ rpmbuild_opts+=(--define)
bash$ rpmbuild_opts+=("_unpackaged_files_terminate_build 0")
```

It is then passed to rpmbuild as follows to start the entire build and RPM package process:

```
bash$ rpmbuild "${rpmbuild_opts[@]}" -bb SPECS/collectd.spec
```

The resulting RPMs in the RPMs directory will now also include collectd-curl. These RPM files can now be copied to the switch and installed from the switch Bash shell:

```
bash$ rpm --noparentdirs -i /bootflash/collectd-5.5.0-1.ia32e.rpm
bash$ rpm --noparentdirs -i /bootflash/collectd-curl-5.5.0-1.ia32e.rpm
```



CHAPTER 18

NX-SDK

- [About the NX-SDK, on page 167](#)
- [About On-Box \(Local\) Applications, on page 168](#)
- [Default Docker Images, on page 168](#)
- [Guidelines and Limitations for NX-SDK, on page 169](#)
- [About NX-SDK 2.0 , on page 170](#)
- [About NX-SDK 2.5, on page 170](#)
- [About Remote Applications, on page 170](#)
- [NX-SDK Security, on page 171](#)
- [Security Profiles for NX SDK 2.0, on page 171](#)

About the NX-SDK

The Cisco NX-OS SDK (NX-SDK) is a C++ abstraction and plugin-library layer that streamlines access to infrastructure for automation and custom application creation, such as generating custom:

- CLIs
- Syslogs
- Event and Error managers
- Inter-application communication
- High availability (HA)
- Route manager

You can use C++, Python, or Go for application development with NX-SDK.

Requirements

The NX-SDK has the following requirements:

- Docker
- A Linux environment (either Ubuntu 14.04, or Centos 6.7). Cisco recommends using the provided NX-SDK Docker containers. For more information, see [Cisco DevNet NX-SDK](#).

Support for Local (On Switch) and Remote (Off Switch) Applications

Applications that are developed with NX-SDK are created or developed off the Cisco Nexus switch in the Docker containers that NX-SDK provides. After the application is created, you have flexibility of where the applications can be deployed:

- Local (on-box) applications run on the switch. For information, see [About On-Box \(Local\) Applications, on page 168](#)
- Remote (off-box) applications run off switch. This option, supported with NX-SDK 2.0 and later, enables you to deploy the application to run anywhere other than on the switch. For information, see [About Remote Applications, on page 170](#).

Related Information

For more information about Cisco NX-SDK, go to:

- [Cisco DevNet NX-SDK](#). Click the `versions.md` link (<https://github.com/CiscoDevNet/NX-SDK/blob/master/versions.md>) to get information about features and details on each supported release.
- [NX-SDK Readmes](#)

As needed, Cisco adds information for NX-SDK to GitHub.

Considerations for Go Bindings

Go bindings are supported at various levels depending on the release of NX-SDK and whether apps are running locally or remotely.

- Go bindings for any version of NX-SDK remote application are pre-EFT quality.
- Go bindings for a local NX-SDK 2.0 application is pre-EFT.
- Go bindings for a local NX-SDK 1.7.5 application or earlier is supported.

For more information, see [GO Bindings for NX-SDK Applications](#).

About On-Box (Local) Applications

With on box (local) applications, you install the NX-SDK, build your application in whichever supported language you choose, package the app as an `.rpm` file which can be installed on the switch, then install and run your applications on the switch. The `.rpm` files can be manually generated or autogenerated.

Application development occurs in the containers that are provided by NX-SDK. You will use a different container and tools for local applications than remote applications. For more information, see [Default Docker Images, on page 168](#).

For information about building, installing, and running local applications, see [Cisco DevNet NX-SDK](#).

Default Docker Images

NX-SDK has the following Docker images and tools by default for local or remote use.

| Usage | Contents |
|---------------------|---|
| On Switch | Cisco ENXOS SDK Wind River Linux (WRL) tool chain for cross compiling Multi-language binding toolkit Beginning with NX-SDK 1.75, a Go compiler |
| Off switch (remote) | NX-SDK multi-language binding Toolkit with pre-built libnxsdk.so A Go compiler RapidJSON gRPC for remote API support |

For more information, see <https://github.com/CiscoDevNet/NX-SDK#readme>https://github.com/CiscoDevNet/NX-SDK/tree/master/readmes/nxsdk_docker_images.md.

Guidelines and Limitations for NX-SDK

NX-SDK has usage guidelines and limitations for running applications locally (on box) or remotely (off box).

For guidelines and limitations, see "Helpful Notes" at [Cisco DevNet NX-SDK](#).

- For remote applications, ports that connect to the SDK server are from 50002 through 50100. Make sure that these ports are open and not used by other services. Port 50051 is blocked for use by an internal application, and cannot be used by remote applications.
- The following limitations apply to developing a custom application through NX-SDK:
 - You must develop applications with any current Linux distribution that have GNU C/C++ toolchains and standard libraries.
 - We recommend developing applications in the provided Docker images. Run the Docker image and associated tools requires a minimum of 8 MB of free memory and 64-bit hosts. See [Default Docker Images, on page 168](#).
- Cisco recommends that you build and test your applications in Bash. When you deploy them in production, build the applications as RPMs and run them in the NX-OS VSH (Vshell).
- Cisco Nexus switches support a maximum of 32 applications in Bash and VSH.
- For NX-SDK 2.0, there is a 1-to-1 limit for remote NX-SDK applications and NX-SDK servers running in the Cisco Nexus switch.
- A maximum of 10 remote NX-SDK applications can run simultaneously in a switch.
- The NX-SDK server accepts remote requests only from a CLI-configured application and not from random applications.
- For any error that occurs, the SDK server throws an exception.
- Some APIs do not run in a remote application. Refer to API documentation for more information.

- Remote client library code is not open source, so NX-SDK remote apps need to run in the NX-SDK remote docker container. To use remote NX-SDK in your own OS, make a request in github with your OS and compiler version. Cisco can generate a remote libnxsdk.so based on your version.

About NX-SDK 2.0

The NX-SDK version 2.0 enables execution-environment flexibility for developers to run their applications wherever needed. With this version of NX-SDK, your applications are still developed off the switch in containers, but you can run the apps either on the switch or off the switch, for example in a cloud.

NX-SDK 2.0 offers the following benefits:

- Easy integration of the switch into the customer environment.
- Scalability to enable the switch to seamlessly operate in data centers, public clouds, and private clouds.
- Decoupling customer apps from switch resources so that changes at the switch-level resources do not require change or rewrite of applications.
- Single library with simple to use APIs for applications to link against, which simplifies switch interactions and allows applications to be written in high-level languages that are easier to write and debug.
- Running Remote services are more secure than on-box applications.

For more information, see https://github.com/CiscoDevNet/NX-SDK/blob/master/readmes/NXSDK_in_NXOS.md Deploying NX-SDK Applications Remotely.

About NX-SDK 2.5

Beginning with Cisco NX-OS Release 9.3(3), support is added for the Streaming Syslog feature.

For more information, see [CiscoDevNet](#).

Table 5: Syslog Events

| Features | Details |
|---------------|--|
| Syslog Events | <ul style="list-style-type: none"> • Ability for custom applications to register for Cisco NX-OS syslog events. • Refer to watchSyslog and postSyslogCb APIs in nx_trace.h for more details. |

About Remote Applications

Remote applications can be on a different switch that is not a Cisco Nexus switch. Remote, or off-box, applications call through the NX-SDK layer to interact with the switch to read information (get) or write information (set).

Both local and remote NX-SDK applications use the same APIs, which offer you the flexibility to deploy NX-SDK applications on- or off-box.

To run remotely, an application must meet specific requirements. For information, see https://github.com/CiscoDevNet/NX-SDK/blob/master/readmes/NXSDK_in_NXOS.md **Deploying NX-SDK Applications Remotely**.

Backward Compatibility for Pre-2.0 NX-SDK Applications

NX-SDK 2.0 has conditional backward compatibility for NX-SDK v1.75 applications depending on how these applications were developed:

- Usually, NX-SDK supports remotely running an app that you created before NX-SDK 2.0 without requiring you to completely rewrite your app. Instead, you can reuse the same app without modifying it to change the API calls. To support older apps in the new NX-SDK 2.0 model, the API call must provide IP and Port parameters. These parameters are not available in NX-SDK 1.75 and earlier, but you can add the IP address and Port information as environment variables that the app can export to the SDK server.
- However, sometimes backward compatibility for pre-NX-SDK 2.0 apps might not be supported. It is possible that some APIs in older apps might not support, or be capable of, running remotely. In this case, the APIs can throw an exception. Depending on how complete and robust the exception-handling is for the original application, the application might operate unpredictably, and in worst cases, possibly crash.

For more information, see https://github.com/CiscoDevNet/NX-SDK/blob/master/readmes/NXSDK_in_NXOS.md **Deploying NX-SDK Applications Remotely**.

NX-SDK Security

Beginning with NX-OS 9.3(1), NX-SDK 2.0 supports the following security features:

- Session security. Remote applications can connect to the NX SDK server on the switch through Transport Layer Service (TLS) to provide encrypted sessions between the applications and the switch's NX SDK server.
- Server certificate security. For new switch deployments with Cisco NX-OS 9.3(1), the NX-SDK server generates a one-day temporary certificate to provide enough time to install a custom certificate.

If your NX-SDK server already has a custom certificate that is installed, for example, if you are upgrading from a previous NX-SDK version to NX-SDK 2.0, your existing certificate is retained and used after upgrade.

- API write-call control. NX-SDK 2.0 introduces security profiles, which enable you to select a pre-defined policy for controlling how much control an application has with the NX-SDK server. For more information about security profiles, see [Security Profiles for NX SDK 2.0, on page 171](#).

Security Profiles for NX SDK 2.0

In previous releases, the APIs for SDK version 1.75 were permitted only to read and get data for events. Beginning in Cisco NX-OS Release 9.3(1), NX-SDK 2.0 supports different types of operations, including write calls.

The ability of an app to read or write to the switch can be controlled through a security profile. A security profile is an optional object that is attached to the applications' service running in the switch. Security profiles control an application's ability to write to the switch, and in turn, control the applications ability to modify, delete, or configure switch functionality. By default, application writes are disallowed, so for each application, you will need to create a security profile that enables write access to the switch.

Cisco's NX-SDK offers the following security profiles.

| Profile | Description | Values |
|----------|---|---|
| Deny | Prevents any API calls from writing to the switch except for adding CLIs. | This is the default profile. |
| Throttle | Allows APIs that modify the switch, but only up to a specified number of calls. This security profile applies throttling to control the number of API calls. The application is allowed to write up to the limit, but when the limit is exceeded, writing stops, and the reply sends an error message. | The throttle is 50 API calls, and the throttle resets after five seconds. |
| Permit | APIs that modify the switch are allowed without restriction | |

For more information about security profiles in NX-SDK, see [Security Profiles for NX-SDK Applications](#).

For additional information about building, installing, and running applications, go to [CiscoDevNet NX-SDK](#).



CHAPTER 19

Using Docker with Cisco NX-OS

- [About Docker with Cisco NX-OS, on page 173](#)
- [Guidelines and Limitations for Docker, on page 173](#)
- [Prerequisites for Setting Up Docker Containers Within Cisco NX-OS, on page 174](#)
- [Starting the Docker Daemon, on page 174](#)
- [Configure Docker to Start Automatically, on page 175](#)
- [Starting Docker Containers: Host Networking Model, on page 176](#)
- [Starting Docker Containers: Bridged Networking Model, on page 177](#)
- [Mounting the bootflash and volatile Partitions in the Docker Container, on page 178](#)
- [Enabling Docker Daemon Persistence on Enhanced ISSU Switchover, on page 178](#)
- [Enabling Docker Daemon Persistence on the Cisco Nexus Platform Switches Switchover, on page 179](#)
- [Resizing the Docker Storage Backend, on page 180](#)
- [Stopping the Docker Daemon, on page 182](#)
- [Docker Container Security, on page 183](#)
- [Adding Nodes to a Kubernetes Cluster, on page 184](#)
- [Docker Troubleshooting, on page 187](#)

About Docker with Cisco NX-OS

Docker provides a way to run applications securely isolated in a container, packaged with all its dependencies and libraries. See <https://docs.docker.com/> for more information on Docker.

Beginning with Cisco NX-OS Release 9.2(1), support is now added for using Docker within Cisco NX-OS on a switch.

The version of Docker that is included on the switch is CE 18.09.0. The Docker daemon is not running by default. You must start it manually or set it up to automatically restart when the switch boots up.

This section describes how to enable and use Docker in the specific context of the switch environment. Refer to the Docker documentation at <https://docs.docker.com/> for details on general Docker usage and functionality.

Guidelines and Limitations for Docker

Following are the guidelines and limitations for using Docker on Cisco NX-OS on a switch:

- If you are running a third-party DHCPD server in Docker, there might be issues with offers reaching the client if used along with SVI. A possible workaround is to use broadcast responses.
- Docker functionality is supported on the Cisco Nexus 9000Cisco Nexus 3000 Series switches with at least 8 GB of system RAM.

Prerequisites for Setting Up Docker Containers Within Cisco NX-OS

Following are the prerequisites for using Docker on Cisco NX-OS on a switch:

- Enable the host Bash shell. To use Docker on Cisco NX-OS on a switch, you must be the root user on the host Bash shell:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# feature bash-shell
```

- If the switch is in a network that uses an HTTP proxy server, the `http_proxy` and `https_proxy` environment variables must be set up in `/etc/sysconfig/docker`. For example:

```
export http_proxy=http://proxy.esl.cisco.com:8080
export https_proxy=http://proxy.esl.cisco.com:8080
```

- Verify that the switch clock is set correctly, or you might see the following error message:

```
x509: certificate has expired or is not yet valid
```

- Verify that the domain name and name servers are configured appropriately for the network and that it is reflected in the `/etc/resolv.conf` file:

```
switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vrf context management
switch(config-vrf)# ip domain-name ?
WORD Enter the default domain (Max Size 64)

switch(config-vrf)# ip name-server ?
A.B.C.D Enter an IPv4 address
A:B::C:D Enter an IPv6 address

root@switch# cat /etc/resolv.conf
domain cisco.com #bleed
nameserver 171.70.168.183 #bleed
root@switch#
```

Starting the Docker Daemon

When you start the Docker daemon for the first time, a fixed-size backend storage space is carved out in a file called `dockerpart` on the bootflash, which is then mounted to `/var/lib/docker`. If necessary, you can adjust the default size of this space by editing `/etc/sysconfig/docker` before you start the Docker daemon for the first time. You can also resize this storage space if necessary as described later on.

To start the Docker daemon:

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Start the Docker daemon.

```
root@switch# service docker start
```

Step 3 Check the status.

```
root@switch# service docker status
dockerd (pid 3597) is running...
root@switch#
```

Note Once you start the Docker daemon, do not delete or tamper with the `dockerpart` file on the bootflash since it is critical to the docker functionality.

```
switch# dir bootflash:dockerpart
20000000000 Mar 14 12:50:14 2018 dockerpart
```

Configure Docker to Start Automatically

You can configure the Docker daemon to always start up automatically when the switch boots up.

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Use the `chkconfig` utility to make the Docker service persistent.

```
root@switch# chkconfig --add docker
root@n9k-2#
```

Step 3 Use the `chkconfig` utility to check the Docker service settings.

```
root@switch# chkconfig --list | grep docker
docker 0:off 1:off 2:on 3:on 4:on 5:on 6:off
root@switch#
```

Step 4 To remove the configuration so that Docker does not start up automatically:

```
root@switch# chkconfig --del docker
root@switch# chkconfig --list | grep docker
```

```
root@switch#
```

Starting Docker Containers: Host Networking Model

If you want Docker containers to have access to all the host network interfaces, including data port and management, start the Docker containers with the `--network host` option. The user in the container can switch between the different network namespaces at `/var/run/netns` (corresponding to different VRFs configured in Cisco NX-OS) using the `ip netns exec <net_namespace> <cmd>`.

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Start the Docker container.

Following is an example of starting an Alpine Docker container on the switch and viewing all the network interfaces. The container is launched into the management network namespace by default.

```
root@switch# docker run --name=alpinerun -v /var/run/netns:/var/run/netns:ro,rslave --rm --network
host --cap-add SYS_ADMIN -it alpine
/ # apk --update add iproute2
fetch http://dl-cdn.alpinelinux.org/alpine/v3.7/main/x86_64/APKINDEX.tar.gz
fetch http://dl-cdn.alpinelinux.org/alpine/v3.7/community/x86_64/APKINDEX.tar.gz
(1/6) Installing libelf (0.8.13-r3)
(2/6) Installing libmnl (1.0.4-r0)
(3/6) Installing jansson (2.10-r0)
(4/6) Installing libnftnl-libs (1.0.8-r1)
(5/6) Installing iptables (1.6.1-r1)
(6/6) Installing iproute2 (4.13.0-r0)
Executing iproute2-4.13.0-r0.post-install
Executing busybox-1.27.2-r7.trigger
OK: 7 MiB in 17 packages
/ #
/ # ip netns list
management
default
/ #
/ # ip address
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127.0.0.1/8 scope host lo
valid_lft forever preferred_lft forever
inet6 ::1/128 scope host
valid_lft forever preferred_lft forever
2: tunl0@NONE: <NOARP> mtu 1480 qdisc noop state DOWN group default
link/ipip 0.0.0.0 brd 0.0.0.0
3: gre0@NONE: <NOARP> mtu 1476 qdisc noop state DOWN group default
link/gre 0.0.0.0 brd 0.0.0.0
...
/ #
/ # ip netns exec default ip address
```

```

1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127.0.0.1/16 scope host lo
valid_lft forever preferred_lft forever
2: dummy0: <BROADCAST,NOARP> mtu 1500 qdisc noop state DOWN group default
link/ether 42:0d:9b:3c:d4:62 brd ff:ff:ff:ff:ff:ff
3: tunl0@NONE: <NOARP> mtu 1480 qdisc noop state DOWN group default
link/ipip 0.0.0.0 brd 0.0.0.0
...

```

Starting Docker Containers: Bridged Networking Model

If you want Docker containers to only have external network connectivity (typically through the management interface) and you don't necessarily care about visibility into a specific data port or other switch interface, you can start the Docker container with the default Docker bridged networking model. This is more secure than the host networking model described in the previous section since it also provides network namespace isolation.

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Start the Docker container.

Following is an example of starting an Alpine Docker container on the switch and installing the `iproute2` package.

```

root@switch# docker run -it --rm alpine
/ # apk --update add iproute2
fetch http://dl-cdn.alpinelinux.org/alpine/v3.7/main/x86_64/APKINDEX.tar.gz
fetch http://dl-cdn.alpinelinux.org/alpine/v3.7/community/x86_64/APKINDEX.tar.gz
(1/6) Installing libelf (0.8.13-r3)
(2/6) Installing libmnl (1.0.4-r0)
(3/6) Installing jansson (2.10-r0)
(4/6) Installing libnftnl-libs (1.0.8-r1)
(5/6) Installing iptables (1.6.1-r1)
(6/6) Installing iproute2 (4.13.0-r0)
Executing iproute2-4.13.0-r0.post-install
Executing busybox-1.27.2-r7.trigger
OK: 7 MiB in 17 packages
/ #
/ # ip netns list
/ #

```

Step 3 Determine if you want to set up user namespace isolation.

For containers using the bridged networking model, you can also set up user namespace isolation to further improve security. See [Securing Docker Containers With User namespace Isolation, on page 183](#) for more information.

You can use standard Docker port options to expose a service from within the container, such as `sshd`. For example:

```
root@switch# docker run -d -p 18877:22 --name sshd_container sshd_ubuntu
```

This maps port 22 from within the container to port 18877 on the switch. The service can now be accessed externally through port 18877, as shown in the following example:

```
root@ubuntu-vm# ssh root@ip_address -p 18877
```

Mounting the bootflash and volatile Partitions in the Docker Container

You can make the `bootflash` and `volatile` partitions visible in the Docker container by passing in the `-v /bootflash:/bootflash` and `-v /volatile:/volatile` options in the run command for the Docker container. This is useful if the application in the container needs access to files shared with the host, such as copying a new NX-OS system image to bootflash.



Note This `-v` command option allows for any directory to be mounted into the container and may result in information leaking or other accesses that may impact the operation of the NX-OS system. Limit this to resources such as `/bootflash` and `/volatile` that are already accessible using NX-OS CLI.

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Pass in the `-v /bootflash:/bootflash` and `-v /volatile:/volatile` options in the run command for the Docker container.

```
root@switch# docker run -v /bootflash:/bootflash -v /volatile:/volatile -it --rm alpine
/# ls /
bin          etc          media        root         srv          usr
bootflash   home        mnt          run          sys          var
dev          lib         proc         sbin        tmp          volatile
/ #
```

Enabling Docker Daemon Persistence on Enhanced ISSU Switchover

You can have both the Docker daemon and any running containers persist on an Enhanced ISSU switchover. This is possible since the bootflash on which the backend Docker storage resides is the same and shared between both Active and Standby supervisors.

The Docker containers are disrupted (restarted) during the switchover, so they will not be running continuously.

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Before starting the switchover, use the `chkconfig` utility to make the Docker service persistent.

```
root@switch# chkconfig --add docker
root@n9k-2#
```

Step 3 Start any containers using the `--restart unless-stopped` option so that they will be restarted automatically after the switchover.

The following example starts an Alpine container and configures it to always restart unless it is explicitly stopped or Docker is restarted:

```
root@switch# docker run -dit --restart unless-stopped alpine
root@n9k-2#
```

The Docker containers are disrupted (restarted) during the switchover, so they will not be running continuously.

Enabling Docker Daemon Persistence on the Cisco Nexus Platform Switches Switchover

You can have both the Docker daemon and any running containers persist on a switchover between two separate physical supervisors with distinct bootflash partitions. However, for the Cisco Nexus switches, the bootflash partitions on both supervisors are physically separate. You will therefore need to copy the `dockerpart` file manually to the standby supervisor before performing the switchover.

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Start any containers using the `--restart unless-stopped` option so that they will be restarted automatically after the switchover.

The following example starts an Alpine container and configures it to always restart unless it is explicitly stopped or Docker is restarted:

```
root@switch# docker run -dit --restart unless-stopped alpine
root@n9k-2#
```

Note that the Docker containers will be disrupted (restarted) during the switchover, so they will not be running continuously.

Step 3 Before starting the switchover, use the `chkconfig` utility to make the Docker service persistent.

```
root@switch# chkconfig --add docker
root@n9k-2#
```

Step 4 Copy the Docker backend storage partition from the active to the standby supervisor bootflash:

```
root@switch# service docker stop
Stopping dockerd: dockerd shutdown

root@switch# cp /bootflash/dockerpart /bootflash_sup-remote/

root@switch# service docker start
```

Resizing the Docker Storage Backend

After starting or using the Docker daemon, you can grow the size of the Docker backend storage space according to your needs.

Procedure

Step 1 Disable the Guest Shell.

If you do not disable the Guest Shell, it may interfere with the resize.

```
switch# guestshell disable
You will not be able to access your guest shell if it is disabled. Are you sure you want to disable
the guest shell? (y/n) [n] y
switch# 2018 Mar 15 17:16:55 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Deactivating virtual
service 'guestshell+'
2018 Mar 15 17:16:57 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Successfully deactivated virtual
service 'guestshell+'
```

Step 2 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 3 Get information on the current amount of storage space available.

```
root@switch# df -kh /var/lib/docker
Filesystem Size Used Avail Use% Mounted on
/dev/loop12 1.9G 7.6M 1.8G 1% /var/lib/docker
root@n9k-2#
```

Step 4 Stop the Docker daemon.

```
root@switch# service docker stop
Stopping dockerd: dockerd shutdown
```

Step 5 Get information on the current size of the Docker backend storage space (`/bootflash/dockerpart`).

```
root@switch# ls -l /bootflash/dockerpart
-rw-r--r-- 1 root root 2000000000 Mar 15 16:53 /bootflash/dockerpart
root@n9k-2#
```

Step 6 Resize the Docker backend storage space.

For example, the following command increases the size by 500 megabytes:

```
root@switch# truncate -s +500MB /bootflash/dockerpart
root@n9k-2#
```

Step 7 Get updated information on the size of the Docker backend storage space to verify that the resizing process was completed successfully.

For example, the following output confirms that the size of the Docker backend storage was successfully increased by 500 megabytes:

```
root@switch# ls -l /bootflash/dockerpart
-rw-r--r-- 1 root root 2500000000 Mar 15 16:54 /bootflash/dockerpart
root@n9k-2#
```

Step 8 Check the size of the filesystem on /bootflash/dockerpart.

```
root@switch# e2fsck -f /bootflash/dockerpart
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/bootflash/dockerpart: 528/122160 files (0.6% non-contiguous), 17794/488281 blocks
```

Step 9 Resize the filesystem on /bootflash/dockerpart.

```
root@switch# /sbin/resize2fs /bootflash/dockerpart
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /bootflash/dockerpart to 610351 (4k) blocks.
The filesystem on /bootflash/dockerpart is now 610351 blocks long.
```

Step 10 Check the size of the filesystem on /bootflash/dockerpart again to confirm that the filesystem was successfully resized.

```
root@switch# e2fsck -f /bootflash/dockerpart
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/bootflash/dockerpart: 528/154736 files (0.6% non-contiguous), 19838/610351 blocks
```

Step 11 Start the Docker daemon again.

```
root@switch# service docker start
Updating certificates in /etc/ssl/certs...
0 added, 0 removed; done.
Running hooks in /etc/ca-certificates/update.d...
done.
Starting dockerd with args '--debug=true':
```

Step 12 Verify the new amount of storage space available.

```
root@switch# df -kh /var/lib/docker
Filesystem Size Used Avail Use% Mounted on
/dev/loop12 2.3G 7.6M 2.3G 1% /var/lib/docker
```

Step 13 Exit out of Bash shell.

```
root@switch# exit
logout
switch#
```

Step 14 Enable the Guest Shell, if necessary.

```
switch# guestshell enable

switch# 2018 Mar 15 17:12:53 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Activating virtual service
'guestshell+'
switch# 2018 Mar 15 17:13:18 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Successfully activated
virtual service 'guestshell+'
```

Stopping the Docker Daemon

If you no longer wish to use Docker, follow the procedures in this topic to stop the Docker daemon.

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Stop the Docker daemon.

```
root@switch# service docker stop
Stopping dockerd: dockerd shutdown
```

Step 3 Verify that the Docker daemon is stopped.

```
root@switch# service docker status
dockerd is stopped
root@switch#
```

Note You can also delete the `dockerpart` file on the bootflash at this point, if necessary:

```
switch# delete bootflash:dockerpart
Do you want to delete "/dockerpart" ? (yes/no/abort) y
switch#
```

Docker Container Security

Following are the Docker container security recommendations:

- Run in a separate user `namespace` if possible.
- Run in a separate network `namespace` if possible.
- Use `cgroups` to limit resources. An existing `cgroup` (`ext_ser`) is created to limit hosted applications to what the platform team has deemed reasonable for extra software running on the switch. Docker allows use of this and limiting per-container resources.
- Do not add unnecessary POSIX capabilities.

Securing Docker Containers With User namespace Isolation

For containers using the bridged networking model, you can also set up user namespace isolation to further improve security. See <https://docs.docker.com/engine/security/usersns-remap/> for more information.

Procedure

Step 1 Determine if a `dockremap` group already exists on your system.

A `dockremap` user must already be set up on your system by default. If the `dockremap` group doesn't already exist, follow these steps to create it.

a) Enter the following command to create the `dockremap` group:

```
root@switch# groupadd dockremap -r
```

b) Create the `dockremap` user, unless it already exists:

```
root@switch# useradd dockremap -r -g dockremap
```

c) Verify that the `dockremap` group and the `dockremap` user were created successfully:

```
root@switch# id dockremap
uid=999(dockremap) gid=498(dockremap) groups=498(dockremap)
root@switch#
```

Step 2 Add the desired re-mapped ID and range to the `/etc/subuid` and `/etc/subgid`.

For example:

```
root@switch# echo "dockremap:123000:65536" >> /etc/subuid
root@switch# echo "dockremap:123000:65536" >> /etc/subgid
```

Step 3 Using a text editor, add the `--usersns-remap=default` option to the `other_args` field in the `/etc/sysconfig/docker` file.

For example:

```
other_args="--debug=true --users-remap=default"
```

Step 4 Restart the Docker daemon, or start it if it is not already running, using `service docker [re]start`.

For example:

```
root@switch# service docker [re]start
```

Refer to the Docker documentation at <https://docs.docker.com/engine/security/users-remap/> for more information on configuring and using containers with user namespace isolation.

Moving the `cgroup` Partition

The `cgroup` partition for third-party services is `ext_ser`, which limits CPU usage to 25% per core. Cisco recommends that you run your Docker container under this `ext_ser` partition.

If the Docker container is run without the `--cgroup-parent=/ext_ser/` option, it can get up to the full 100% host CPU access, which can interfere with the regular operation of Cisco NX-OS.

Procedure

Step 1 Load Bash and become superuser.

```
switch# run bash sudo su -
```

Step 2 Run the Docker container under the `ext_ser` partition.

For example:

```
root@switch# docker run --name=alpinerun -v /var/run/netns:/var/run/netns:ro,rslave --rm --network
host --cgroup-parent=/ext_ser/ --cap-add SYS_ADMIN -it alpine
/ #
```

Adding Nodes to a Kubernetes Cluster

This topic describes how to add nodes to a Kubernetes cluster. In this example:

- The Kubernetes (Ubuntu) primary has an IP address of 10.122.197.246
- The switch software running as Docker containers has an IP address of 10.122.84.24



Note In the following examples, long single lines of text are broken up with the `\` character to improve readability.

Procedure

Step 1 Run the following commands (on? for?) the Kubernetes (Ubuntu) primary.

a) Enter this command:

```
root@switch# docker run -d --net=host gcr.io/google_containers/etcd:2.2.1 /usr/local/bin/etcd
--listen-client-urls=http://0.0.0.0:4001 --advertise-client-urls=http://0.0.0.0:4001
--data-dir=/var/etcd/data
```

b) Enter this command:

```
root@switch# docker run -d --name=api --net=host --pid=host --privileged=true
gcr.io/google_containers/hyperkube:v1.2.2 /hyperkube apiserver --insecure-bind-address=0.0.0.0
--allow-privileged=true --service-cluster-ip-range=10.0.0.1/24 --etcd_servers=http://127.0.0.1:4001
--v=2
```

c) Enter this command:

```
root@switch# docker run -d --name=kubs --volume=/:/rootfs:ro --volume=/sys:/sys:ro
--volume=/dev:/dev --volume=/var/lib/docker/:/var/lib/docker:rw
--volume=/var/lib/kubelet/:/var/lib/kubelet:rw --volume=/var/run:/var/run:rw --net=host --pid=host
--privileged=true gcr.io/google_containers/hyperkube:v1.2.2 /hyperkube kubelet
--allow-privileged=true --hostname-override="127.0.0.1" --address="0.0.0.0"
--api-servers=http://0.0.0.0:8080 --cluster_dns=10.0.0.10 --cluster_domain=cluster.local
--config=/etc/kubernetes/manifests-multi
```

d) Enter this command:

```
root@switch# docker run -d --name=proxy --net=host --privileged
gcr.io/google_containers/hyperkube:v1.2.2 /hyperkube proxy --master=http://0.0.0.0:8080 --v=2
```

e) Enter this command:

```
root@switch# export KUBERNETES_MASTER=http://10.122.197.246:8080
```

f) Enter this command:

```
root@switch# curl -o /usr/bin/kubectl
http://storage.googleapis.com/kubernetes-release/release/v1.2.2/bin/linux/amd64/kubectl
```

g) Enter this command:

```
root@switch# kubectl -s $KUBERNETES_MASTER create -f kube-system.json
```

h) Enter this command:

```
root@switch# kubectl -s $KUBERNETES_MASTER create -f skydns-rc.yaml
```

i) Enter this command:

```
root@switch# kubectl -s $KUBERNETES_MASTER create -f skydns-svc.yaml
```

j) Enter this command:

```
root@switch# kubectl -s $KUBERNETES_MASTER create -f dashboard.yaml
kubectl -s $KUBERNETES_MASTER cluster-info
```

k) Enter this command:

```
root@switch# kubectl -s $KUBERNETES_MASTER cluster-info
```

Step 2 Run the following steps (on? for?) the switch.

a) Enter this command:

```
root@switch# docker run -d --name=kubs --net=host --pid=host --privileged=true --volume=/:/rootfs:ro
--volume=/sys:/sys:ro --volume=/dev:/dev --volume=/var/lib/docker/:/var/lib/docker:rw
--volume=/var/lib/kubelet/:/var/lib/kubelet:rw --volume=/var/run:/var/run:rw
gcr.io/google_containers/hyperkube:v1.2.2 /hyperkube kubelet --allow-privileged=true --containerized
--enable-server --cluster_dns=10.0.0.10 --cluster_domain=cluster.local
--config=/etc/kubernetes/manifests-multi --hostname-override="10.122.84.34" --address=0.0.0.0
--api-servers=http://10.122.197.246:8080
```

b) Enter this command:

```
root@switch# docker run -d --name=proxy --net=host --privileged=true
gcr.io/google_containers/hyperkube:v1.2.2 /hyperkube proxy --master=http://10.122.197.246:8080
--v=2
```

Step 3 Run the following commands (on? for?) the Kubernetes (Ubuntu) primary to deploy an nginx app in a replication controller object.

a) Enter this command:

```
lab@rmbalk-ubuntu1:~$ kubectl get node
NAME                STATUS    AGE
10.122.84.34        Ready     16m
127.0.0.1           Ready     22m
```

b) Enter this command:

```
lab@rmbalk-ubuntu1:~$ kubectl apply -f replication.yaml
replicationcontroller "nginx" created
```

c) Enter this command:

```
lab@rmbalk-ubuntu1:~$ kubectl describe -f replication.yaml
Name:                nginx
Namespace:           default
Image(s):             nginx
Selector:             app=nginx
Labels:               app=nginx
Replicas:             3 current / 3 desired
Pods Status:         3 Running / 0 Waiting / 0 Succeeded / 0 Failed
No volumes.
Events:
  FirstSeen    LastSeen    Count   From              SubobjectPath    Type
  Reason
  -----
  17s          17s         1       {replication-controller }
  SuccessfulCreate Created pod: nginx-zqfpz
```



```

17s          17s          1          {replication-controller }          Normal
SuccessfulCreate Created pod: nginx-ji40
17s          17s          1          {replication-controller }          Normal
SuccessfulCreate Created pod: nginx-loa0g

```

Docker Troubleshooting

These topics describe issues that can arise with Docker containers and provides possible resolutions.

Docker Fails to Start

Problem: Docker fails to start, showing an error message similar to the following:

```

switch# run bash
bash-4.3$ service docker start
Free bootflash: 39099 MB, total bootflash: 51771 MB
Carving docker bootflash storage: 2000 MB
2000+0 records in
2000+0 records out
2000000000 bytes (2.0 GB) copied, 22.3039 s, 89.7 MB/s
losetup: /dev/loop18: failed to set up loop device: Permission denied
mke2fs 1.42.9 (28-Dec-2013)
mkfs.ext4: Device size reported to be zero. Invalid partition specified, or
partition table wasn't reread after running fdisk, due to
a modified partition being busy and in use. You may need to reboot
to re-read your partition table.

```

Failed to create docker volume

Possible Cause: You might be running Bash as an admin user instead of as a root user.

Solution: Determine if you are running Bash as an admin user instead of as a root user:

```

bash-4.3$ whoami
admin

```

Exit out of Bash and run Bash as root user:

```

bash-4.3$ exit
switch# run bash sudo su -

```

Docker Fails to Start Due to Insufficient Storage

Problem: Docker fails to start, showing an error message similar to the following, due to insufficient bootflash storage:

```

root@switch# service docker start
Free bootflash: 790 MB, total bootflash: 3471 MB
Need at least 2000 MB free bootflash space for docker storage

```

Possible Cause: You might not have enough free bootflash storage.

Solution: Free up space or adjust the `variable_dockerstrg` values in `/etc/sysconfig/docker` as needed, then restart the Docker daemon:

```
root@switch# cat /etc/sysconfig/docker
# Replace the below with your own docker storage backend boundary value (in MB)
# if desired.
boundary_dockerstrg=5000

# Replace the below with your own docker storage backend values (in MB) if
# desired. The smaller value applies to platforms with less than
# $boundary_dockerstrg total bootflash space, the larger value for more than
# $boundary_dockerstrg of total bootflash space.
small_dockerstrg=300
large_dockerstrg=2000
```

Failure to Pull Images from Docker Hub (509 Certificate Expiration Error Message)

Problem: The system fails to pull images from the Docker hub with an error message similar to the following:

```
root@switch# docker pull alpine
Using default tag: latest
Error response from daemon: Get https://registry-1.docker.io/v2/: x509: certificate has
expired or is not yet valid
```

Possible Cause: The system clock might not be set correctly.

Solution: Determine if the clock is set correctly or not:

```
root@n9k-2# sh clock
15:57:48.963 EST Thu Apr 25 2002
Time source is Hardware Calendar
```

Reset the clock, if necessary:

```
root@n9k-2# clock set hh:mm:ss { day month | month day } year
```

For example:

```
root@n9k-2# clock set 14:12:00 10 feb 2018
```

Failure to Pull Images from Docker Hub (Client Timeout Error Message)

Problem: The system fails to pull images from the Docker hub with an error message similar to the following:

```
root@switch# docker pull alpine
Using default tag: latest
Error response from daemon: Get https://registry-1.docker.io/v2/: net/http: request canceled
while waiting for connection (Client.Timeout exceeded while awaiting headers)
```

Possible Cause: The proxies or DNS settings might not be set correctly.

Solution: Check the proxy settings and fix them, if necessary, then restart the Docker daemon:

```
root@switch# cat /etc/sysconfig/docker | grep proxy
#export http_proxy=http://proxy.esl.cisco.com:8080
```

```
#export https_proxy=http://proxy.esl.cisco.com:8080
root@switch# service docker [re]start
```

Check the DNS settings and fix them, if necessary, then restart the Docker daemon:

```
root@switch# cat /etc/resolv.conf
domain cisco.com #bleed
nameserver 171.70.168.183 #bleed
root@switch# # conf t
    Enter configuration commands, one per line. End with CNTL/Z.
    switch(config)# vrf context management
    switch(config-vrf)# ip domain-name ?
    WORD Enter the default domain (Max Size 64)

    switch(config-vrf)# ip name-server ?
    A.B.C.D Enter an IPv4 address
    A:B::C:D Enter an IPv6 address
root@switch# service docker [re]start
```

Docker Daemon or Containers Not Running On Switch Reload or Switchover

Problem: The Docker daemon or containers do not run after you have performed a switch reload or switchover.

Possible Cause: The Docker daemon might not be configured to persist on a switch reload or switchover.

Solution: Verify that the Docker daemon is configured to persist on a switch reload or switchover using the `chkconfig` command, then start the necessary Docker containers using the `--restart unless-stopped` option. For example, to start an Alpine container:

```
root@switch# chkconfig --add docker
root@switch#
root@switch# chkconfig --list | grep docker
docker 0:off 1:off 2:on 3:on 4:on 5:on 6:off
root@switch# docker run -dit --restart unless-stopped alpine
```

Resizing of Docker Storage Backend Fails

Problem: An attempt to resize the Docker backend storage failed.

Possible Cause: You might not have Guest Shell disabled.

Solution: Use the following command to determine if Guest Shell is disabled:

```
root@switch# losetup -a | grep dockerpart
root@n9k-2#
```

The command should not display any output if Guest Shell is disabled.

Enter the following command to disable the Guest Shell, if necessary:

```
switch# guestshell disable
```

If you still cannot resize the Docker backend storage, you can delete `/bootflash/dockerpart`, then adjust the `[small_]large_dockerstrg` in `/etc/sysconfig/docker`, then start Docker again to get a fresh Docker partition with the size that you want.

Docker Container Doesn't Receive Incoming Traffic On a Port

Problem: The Docker container doesn't receive incoming traffic on a port.

Possible Cause: The Docker container might be using a netstack port instead of a kstack port.

Solution: Verify that any ephemeral ports that are used by Docker containers are within the kstack range. Otherwise any incoming packets can get sent to netstack for servicing and dropped.

```
switch# show socket local-port-range
Kstack local port range (15001 - 58000)
Netstack local port range (58001 - 63535) and nat port range (63536 - 65535)
switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# sockets local-port-range <start_port> <end_port>
switch# run bash sudo su -
root@switch# cat /proc/sys/net/ipv4/ip_local_port_range
15001 58000
root@switch#
```

Unable to See Data Port And/Or Management Interfaces in Docker Container

Problem: You are unable to see the data port or management interfaces in the Docker container.

Solution:

- Verify that the Docker container is started in the host network namespace with all host namespaces mapped in using the `-v /var/run/netns:/var/run/netns:ro,rslave --network host` options.
- Once in the container, you will be in the management network namespace by default. You can use the `ip netns` utility to move to the default (`init`) network namespace, which has the data port interfaces. The `ip netns` utility might need to be installed in the container using `dnf`, `apk`, or something similar.

General Troubleshooting Tips

Problem: You have other issues with Docker containers that were not resolved using other troubleshooting processes.

Solution:

- Look for `dockerd` debug output in `/var/log/docker` for any clues as to what is wrong.
- Verify that your switch has 8 GB or more of RAM. Docker functionality is not supported on any switch that has less than 8 GB of RAM.



PART **III**

Application Hosting

- [Application Hosting](#), on page 193
- [ThousandEyes Enterprise Agent](#) , on page 211



CHAPTER 20

Application Hosting

A hosted application is a software as a service (SaaS) solution, and it can be run remotely using commands. Application hosting gives administrators a platform for leveraging their own tools and utilities.



Note Application hosting supports only Docker applications.

This module describes the Application Hosting feature and how to enable it.

- [Guidelines and Limitations for Application Hosting, on page 193](#)
- [Information About Application Hosting, on page 194](#)
- [How to Configure Application Hosting, on page 195](#)
- [Copying Application Data, on page 204](#)
- [Deleting Application Data, on page 204](#)
- [Verifying the Application-Hosting Configuration, on page 204](#)
- [Configuration Examples for Application Hosting, on page 208](#)
- [Additional References, on page 209](#)
- [Feature Information for Application Hosting, on page 209](#)

Guidelines and Limitations for Application Hosting

This sections lists the guidelines and limitations for the application hosting feature:

- The application hosting feature does not support IPv6 configuration in Cisco NX-OS Release 10.3(1)F.
- The application hosting feature is supported on Cisco Nexus 9300 series FX, FX2, FX3, GX, and GX2 platforms and Cisco Nexus 9500 series modular switches with FX and GX line cards.
- Only 1 interface per container is supported.
- Application hosting does not support the configuration replace feature.
- Beginning with Cisco NX-OS Release 10.3(3)F, the application hosting feature is supported on Cisco Nexus 9808 platform switches.
- Beginning with Cisco NX-OS Release 10.3(3)F, the application hosting feature is supported on Cisco Nexus 9504 and 9508 platform switches with -R and -R2 line cards. This feature is also supported on Cisco N3K-C36180YC-R, N3K-C3636C-R, and N3K-C36480LD-R2 switches.

- When performing a non-disruptive upgrade on Cisco Nexus 9300-FX2, 9300-FX3, 9300-GX, and 9300-GX2 platform switches from a release which does not support app-hosting to Cisco NX-OS Release 10.3(3)F, post upgrade, enabling feature app-hosting requires a system reload for TCAM entry to take effect.
- Testing connectivity by pinging into the app container from NX-OS will not work because ping from NX-OS uses the netstack process. If bidirectional reachability needs to be verified, use bash ping: **conf t ; feature bash ; run bash ; ping <app-ip-address>**.
- After configuring **system vrf-member-change retain-l3-config** command, when we change vrf membership on interface such that it falls in the same VRF as app-hosting bridge, and HSRP VIP matches subnet of app-hosting bridge IP, then VIP won't be removed from **show running-config** or DME and no faults will be raised. Rather, an HSRP syslog (severity 3) will be thrown notifying the user of misconfiguration.

Information About Application Hosting

This section provides information about Application Hosting.

Need for Application Hosting

The move to virtual environments has brought the need to build applications that are reusable, portable, and scalable. Application hosting gives administrators a platform for leveraging their own tools and utilities. An application, hosted on a network device, can serve a variety of purposes such as automation, configuration management monitoring, and integration with the existing tool chains.



Note In this document, *container* refers to Docker applications.

Application Hosting Overview

The Cisco application-hosting framework is an NX-OS Python process that manages virtualized and container applications that run on devices.

Application hosting provides the following services:

- Launches designated applications in containers.
- Checks available resources (memory, CPU, and storage), and allocates and manages them.
- Provides access to services through REST APIs.
- Provides a CLI endpoint.
- Provides an application-hosting infrastructure referred to as Cisco Application Framework (CAF).
- Helps set up platform-specific networking (packet-path) through a special application bridge interface.

The application-hosting container that is referred to as the virtualization environment is provided to run a guest application on the host operating system. The Cisco NX-OS application hosting feature provides

manageability and networking models for running a guest application. The virtualization infrastructure allows an administrator to define a logical interface that specifies the connectivity between the host and the guest. Cisco NX-OS maps the logical interface into a Virtual Network Interface Card (vNIC) that the guest application uses.

Applications that are to be deployed in the containers are packaged as .tar or .tar.gz files. The configuration that is specific to these applications is also packaged as part of the .tar or .tar.gz files.

How to Configure Application Hosting

The following sections provide information about the various tasks that comprise the configuration of application hosting.

Enabling Application Hosting Feature

Perform this task to enable the Cisco application hosting feature, which enables the user interface command and API interfaces to manage, administer, monitor, and troubleshoot the applications on the host system, and to perform a variety of related activities.

SUMMARY STEPS

1. **configure terminal**
2. **feature app-hosting**
3. **end**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | configure terminal Example: <code>switch# configure terminal</code> | Enters global configuration mode. |
| Step 2 | feature app-hosting Example: <code>switch(config)# feature app-hosting</code> | Enables the Cisco application hosting feature. |
| Step 3 | end Example: <code>switch(config)# end</code> | Exits global configuration mode and returns to privileged EXEC configuration mode. |

Configuring Application Hosting Bridge Connections

Layer 3 connectivity to the application containers requires its own endpoint IPv4 addresses. In NX-OS, a virtual bridge mechanism called app-hosting bridge hosts the application containers inside the Cisco Nexus switch.

The bridge acts as a gateway to the application containers and helps route the traffic to the attached VRF routing context. The bridge forwards the subnet traffic of the application over the switch interface per the VRF context.

The hosting of the application containers with network connections across the switch interface requires a dedicated endpoint IP subnet with a minimum of 2 assignable addresses. One IP address is for the application container guest interface and the other IP address is for the application container gateway.

Internally, the application container guest interface is a Virtual Network Interface Card (vNIC), off the app-hosting virtual bridge.

SUMMARY STEPS

1. **configure terminal**
2. **app-hosting bridge** *bridge-index*
3. **ip address** *ip-address/mask*
4. **vrf member** *name*
5. **exit**
6. **app-hosting appid** *name*
7. **app-vnic gateway bridge** *bridge-index* **guest-interface** *guest-interface-number*
8. **guest-ipaddress** *ip-address/mask*
9. **exit**
10. **app-default-gateway** *ip-address* **guest-interface** *guest-interface*
11. **end**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | configure terminal Example: switch# configure terminal | Enters global configuration mode. |
| Step 2 | app-hosting bridge <i>bridge-index</i> Example: switch(config)# app-hosting bridge 1 | Configures the app-hosting bridge and enters the app-hosting bridge configuration mode. • <1-8> Bridge index |
| Step 3 | ip address <i>ip-address/mask</i> Example: | Configures the app bridge IPv4 address which acts as the gateway to the application container. |

| | Command or Action | Purpose |
|----------------|--|---|
| | <code>switch(config-app-hosting-bridge)# ip address 172.25.44.1/30</code> | Note Beginning with Cisco NX-OS Release 10.3(2)F, the subnet will be rejected if the IP is in use by either an interface or a virtual IP. |
| Step 4 | vrf member <i>name</i> Example: <code>switch(config-app-hosting-bridge)# vrf member overlay-VRF</code> | Sets the VRF context. If not configured, it belongs to VRF <i>default</i> . Note Configuring of management VRF is not supported. Configuring management VRF fails the bridge configuration. |
| Step 5 | exit Example: <code>switch(config-app-hosting-bridge)# exit</code> | Exits app bridge configuration mode and returns to global configuration mode. |
| Step 6 | app-hosting appid <i>name</i> Example: <code>switch(config)# app-hosting appid te_app</code> | Configures an application and enters the application-hosting configuration mode. |
| Step 7 | app-vnic gateway bridge <i>bridge-index</i> guest-interface <i>guest-interface-number</i> Example: <code>switch(config-app-hosting)# app-vnic gateway bridge 1 guest-interface 0</code> | Configures the guest VNIC interface for an application and enters the application-hosting vnic interface mode. Note Beginning with Cisco NX-OS Release 10.3(3)F, configuring more than 1 VNIC is not supported. |
| Step 8 | guest-ipaddress <i>ip-address/mask</i> Example: <code>switch(config-app-hosting-app-vnic)# guest-ipaddress 172.25.44.2/30</code> | Configures one of the available IPv4 addresses from the <i>bridge 1</i> subnet. |
| Step 9 | exit Example: <code>switch(config-app-hosting-app-vnic)# exit</code> | Exits app vnic interface configuration mode and returns to app-hosting configuration mode. |
| Step 10 | app-default-gateway <i>ip-address</i> guest-interface <i>guest-interface</i> Example: <code>switch(config-app-hosting-appid)# app-default-gateway 172.25.44.1 guest-interface 0</code> | Configures the available IPv4 address from the <i>bridge 1</i> subnet. Configures the gateway address that is configured in Step 3 . |
| Step 11 | end Example: <code>switch(config-app-hosting)# end</code> | Exits the application-hosting configuration mode and returns to the privileged EXEC mode. |

Lifecycle of an Application

The following EXEC commands describe the lifecycle of an application.



Note If you make any configuration changes after you install an application, the application in the running state does not reflect these changes. To make changes after starting the application, stop and deactivate the application before making any changes, and then activate and start the application again.

SUMMARY STEPS

1. **enable**
2. **app-hosting install appid** *application-name* **package** *package-path*
3. **app-hosting activate appid** *application-name*
4. **app-hosting start appid** *application-name*
5. **app-hosting stop appid** *application-name*
6. **app-hosting deactivate appid***application-name*
7. **app-hosting uninstall appid** *application-name*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | enable Example: switch# enable | Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | app-hosting install appid <i>application-name</i> package <i>package-path</i> Example: switch# app-hosting install appid te_app package bootflash:my_te_app.tar | Installs an application from the specified location. <ul style="list-style-type: none"> • You can install an application from a local storage location, that is, bootflash. |
| Step 3 | app-hosting activate appid <i>application-name</i> Example: switch# app-hosting activate appid te_app | Activates the application. <ul style="list-style-type: none"> • This command validates all the application resource requests. If all the resources are available, the command activates the application, otherwise the activation fails. |
| Step 4 | app-hosting start appid <i>application-name</i> Example: switch# app-hosting start appid te_app | Starts the application. <ul style="list-style-type: none"> • Activates the application start-up scripts. |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 5 | app-hosting stop appid <i>application-name</i> Example: switch# app-hosting stop appid te_app | (Optional) Stops the application. |
| Step 6 | app-hosting deactivate appid <i>application-name</i> Example: switch# app-hosting deactivate appid te_app | (Optional) Deactivates all the resources that are allocated for the application. |
| Step 7 | app-hosting uninstall appid <i>application-name</i> Example: switch# app-hosting uninstall appid te_app | (Optional) Uninstalls the application. • Uninstalls all the packaging and images stored. Also removes all the changes and updates to the application. |

Upgrading an Application

The following EXEC commands describe how to upgrade an application.

SUMMARY STEPS

1. **enable**
2. **switch# app-hosting upgrade appid** *application-name* **package** *package-path*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | enable Example: switch# enable | Enables privileged EXEC mode. • Enter your password if prompted. |
| Step 2 | switch# app-hosting upgrade appid <i>application-name</i> package <i>package-path</i> Example: switch# app-hosting upgrade appid tea package bootflash:thousandeyes-enterprise-agent-4.1.0.cisco.tar | Upgrades the existing application to a newer version. While doing so, this command stops, upgrades, and reverts the application to the pre-upgrade state. Note <ul style="list-style-type: none"> • If you upgrade an application when it is in a STOPPED state, after a successful upgrade the new app-hosting state changes to ACTIVATED. • You can upgrade an application from a local storage location, that is, bootflash. |

Configuring Docker Run Time Options

You can add a maximum of 30 lines of run time options. The system generates a concatenated string from line 1 though line 30. A string can have more than one Docker run time options.



Note To change the run time option, the application must be in a deactivated state.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **app-hosting appid** *application-name*
4. **app-resource docker**
5. **run-opts** *options*
6. **end**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | enable Example: switch# enable | Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted. |
| Step 2 | configure terminal Example: switch# configure terminal | Enters global configuration mode. |
| Step 3 | app-hosting appid <i>application-name</i> Example: switch(config)# app-hosting appid te_app | Configures an application and enters application-hosting configuration mode. |
| Step 4 | app-resource docker Example: switch(config-app-hosting)# app-resource docker | Enters application-hosting docker-configuration mode to specify application resource updates. |
| Step 5 | run-opts <i>options</i> Example: switch(config-app-hosting-docker)# run-opts 1 "-v \$(APP_DATA) :/data" | Specifies the Docker run time options. |
| Step 6 | end Example: | Exits application-hosting docker-configuration mode and returns to privileged EXEC mode. |

| | Command or Action | Purpose |
|--|---|---------|
| | <code>switch(config-app-hosting-docker)# end</code> | |

Configuring Application Hosting on the Management Interface

NX-OS allows application containers to share the network connections over the Cisco NX-OS management interface. You can internally set up a virtual NAT bridge and assign a private IP address to the guest vNIC interface. The guest interface private IP address gets automatically assigned by the Apphosting framework.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **app-hosting appid** *name*
4. **app-vnic management guest-interface** *network-interface*
5. **end**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | enable Example: <code>switch# enable</code> | Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted. |
| Step 2 | configure terminal Example: <code>switch# configure terminal</code> | Enters global configuration mode. |
| Step 3 | app-hosting appid <i>name</i> Example: <code>switch(config)# app-hosting appid te_app</code> | Configures an application and enters application-hosting configuration mode. |
| Step 4 | app-vnic management guest-interface <i>network-interface</i> Example: <code>switch(config-app-hosting)# app-vnic management guest-interface 0</code> | Connects the guest interface to the management port, and enters application-hosting management-gateway configuration mode. <ul style="list-style-type: none">• The management keyword specifies the Cisco NX-OS <i>interface mgmt0</i> that connects to the container through private IPNAT mode.• The guest-interface <i>network-interface</i> keyword-argument pair specifies the container's internal Ethernet interface number that connects to the Cisco NX-OS management interface <i>mgmt0</i>. The example here uses <i>guest-interface 0</i> for the container's Ethernet 0 interface. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 5 | end Example: <pre>switch(config-app-hosting-mgmt-gateway)# end</pre> | Exits application-hosting management-gateway configuration mode and returns to privileged EXEC mode. |

Overriding Application Resource Configuration

For resource changes to take effect, you must first stop and deactivate an app using the **app-hosting stop** and **app-hosting deactivate** commands, and then restart the app using the **app-hosting activate** and **app-hosting start** commands.

You can use these commands to reset both resources and the app-hosting appid configuration.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **app-hosting appid** *name*
4. **app-resource profile** *name*
5. **cpu** *unit*
6. **memory** *memory*
7. **end**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | enable Example: <pre>switch# enable</pre> | Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | configure terminal Example: <pre>switch# configure terminal</pre> | Enters global configuration mode. |
| Step 3 | app-hosting appid <i>name</i> Example: <pre>switch(config)# app-hosting appid te_app</pre> | Enables application hosting and enters application-hosting configuration mode. |
| Step 4 | app-resource profile <i>name</i> Example: <pre>switch(config-app-hosting)# app-resource profile custom</pre> | Configures the custom application resource profile, and enters the custom application resource profile configuration mode. <ul style="list-style-type: none"> • Only the custom profile name is supported. |

| | Command or Action | Purpose |
|--------|--|--|
| Step 5 | cpu unit Example: <pre>switch(config-app-resource-profile-custom) # cpu 7400</pre> | Changes the default CPU allocation for the application. <ul style="list-style-type: none"> Resource values are application specific, and any adjustment to these values must ensure that the application can run reliably with the changes. |
| Step 6 | memory memory Example: <pre>switch(config-app-resource-profile-custom) # memory 2048</pre> | Changes the default memory allocation. |
| Step 7 | end Example: <pre>switch(config-app-resource-profile-custom) # end</pre> | Exits the custom application resource profile configuration mode and returns to the privileged EXEC mode. |

Advanced Application Hosting Features

By default, App-hosting feature allows only Cisco supported and signed application packages. To install non-Cisco signed application docker images, the sign verification functionality must be turned off. This is a global configuration and affects all the applications that are being installed. The **app-hosting signed-verification [disable | enable]** command disables the sign verification and helps install the non-Cisco Docker applications.

When the application hosting feature is configured, it reserves 2 GB of file space from the bootflash as application storage space. If more space is required for a particular application, then you can increase the partition size. Alternatively, you can decrease the space based on the space requirements of the applications using the **app-hosting bootflash backend storage limit size** global configuration command. The application restarts.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **app-hosting signed-verification [disable | enable]**
4. **app-hosting bootflash backend storage limit size**
5. **end**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|---|--|
| Step 1 | enable Example: <pre>switch# enable</pre> | Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted. |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 2 | configure terminal Example: switch# configure terminal | Enters global configuration mode. |
| Step 3 | app-hosting signed-verification [disable enable] Example: switch(config)# app-hosting signed-verification disable | Disables the package verification to allow non-Cisco applications. <ul style="list-style-type: none"> • Sign verification is enabled by default. |
| Step 4 | app-hosting bootflash backend storage limit size Example: switch(config)# app-hosting bootflash backend storage limit 600 | Configures the applications storage size required considering all the applications going to be installed. <ul style="list-style-type: none"> • 2048 MB would be used by default. • Size in MB; size must be less than the available free space in bootflash. |
| Step 5 | end Example: switch(config-app-resource-profile-custom)# end | Exits the custom application resource profile configuration mode and returns to the privileged EXEC mode. |

Copying Application Data

To copy application data into an application's persistent data mount, use the **app-hosting data appid <appid> copy** command in the privileged EXEC mode.

```
app-hosting data appid tea copy bootflash:src dest
```

where,

`src` is the source file from bootflash, and `dest` is the destination file path.

Deleting Application Data

To delete application data from the application's persistent data mount, use the **app-hosting data appid <appid> delete** command in the privileged EXEC mode.

```
app-hosting data appid tea delete file
```

where,

`file` is the file to be deleted from the application's persistent data mount.

Verifying the Application-Hosting Configuration

Use these **show** commands to verify the configuration. You can use these commands in any order.

SUMMARY STEPS

1. **show app-hosting infra**
2. **show app-hosting list**
3. **show app-hosting bridge**
4. **show app-hosting detail**
5. **show app-hosting resource**
6. **show app-hosting app-hosting utilization appid <app-name>**
7. **show-tech app-hosting**

DETAILED STEPS**Procedure****Step 1 show app-hosting infra**

Displays a summary of the app-hosting infra.

Note Move CAF to running state before performing further operation.

Example:

```
switch(config)# show app-hosting infra
App signature verification: disabled
Docker partition size: 0 MB
Inband packet rate limit: 0 PPS
Services
-----
CAF 1.16.0.0 : Running
HA : Running
App Manager : Running
Libvirtd 4.7.0 : Running
Dockerd 18.09.0-ce : Running
Linux kernel 5.10.126 : Running
```

Step 2 show app-hosting list

Displays the list of apps that are running.

Example:

```
switch(config)# show app-hosting list
App id           State
-----
nginx_1         started
```

Step 3 show app-hosting bridge

Displays the list of app-hosting bridges.

Example:

```
switch(config)# show app-hosting bridge
Bridge ID  VRF      IP Address      IPv6 Address
-----
1          blue     172.10.23.45/24  ::/0
```

Step 4 show app-hosting detail

Displays detailed information about app-hosting.

Example:

```
switch(config)# show app-hosting detail
App id : nginx_1
Owner : appmgr
State : started
Application
Type : docker
Name : nginx
Version : latest
Description :
Author :
Path : /bootflash/nginx.tar.gz
URL Path :
Activated profile name : default

Resource reservation
Memory : 64 MB
Disk : 10 MB
CPU : 200 units

Platform resource profiles
Profile Name CPU(unit) Memory(MB) Disk(MB)
-----
Attached devices
Name          Type          Alias
-----
iox_trace     serial/trace  serial3
iox_syslog    serial/syslog serial2
iox_console_aux serial/aux    serial1
iox_console_shell serial/shell  serial0

Network interfaces
-----
eth0:
MAC address : 5254.9999.0000
IPv4 address : 192.168.10.130
IPv6 address : fe80::5054:99ff:fe99:0/64
Network name : iox-nat_docker0
Tx Packets : 9
Tx Bytes : 726
Tx Errors : 0
Rx Packets : 0
Rx Bytes : 0
Rx Errors : 0

Docker
-----
Run-time information
Command :
Entry-point : /docker-entrypoint.sh nginx -g 'daemon off;'
Run options in use : --publish=40080:80
Package run options :
Application health information
Status : 0
Last probe error :
Last probe output :
```

Step 5 **show app-hosting resource**

Displays information about the resources for app-hosting.

Example:

```
switch(config)# show app-hosting resource
CPU:
Total: 7400 units
Available: 7200 units
VCPU:
Application Hosting
Additional References
Application Hosting
46
Count: 1
Memory:
Total: 3840 (MB)
Available: 3776 (MB)
Storage space:
Total: 110745 (MB)
Available: 93273 (MB)
vice
```

Step 6 **show app-hosting app-hosting utilization appid <app-name>**

Displays the utilization for an application.

Example:

```
switch(config)# show app-hosting utilization appid nginx_1
Application: nginx_1
CPU Utilization:
CPU Allocation: 200 units
CPU Used: 0 %
Memory Utilization:
Memory Allocation: 64 MB
Memory Used: 7000 KB
Disk Utilization:
Disk Allocation: 10 MB
Disk Used: 0 MB
```

Step 7 **show-tech app-hosting**

Displays all the app-hosting logs and the dependent component logs that are relevant.

This show-tech command collects the details for the following show commands:

Example:

```
show system internal app-hosting
show system internal app-hosting event-history debug
show system internal app-hosting event-history error
show system internal app-hosting event-history msgs
show app-hosting list
show app-hosting detail
show app-hosting utilization
show app-hosting infra
show app-hosting resource
show app-hosting bridge
show routing appmgr vrf all
show routing ipv6 appmgr vrf all
```

Configuration Examples for Application Hosting

The following are the various examples pertaining to the configuration of the Application Hosting feature.

Example: Enabling AppHosting Feature

This example shows how to enable Cisco Apphosting feature.

```
switch# configure terminal
switch(config)# feature app-hosting
switch(config)# end
```

Example: Configuring Application Hosting Bridge Connections

This example shows how to configure application hosting bridge connections.

```
switch(config)# configure terminal
switch(config)# app-hosting bridge 1
switch(config-app-hosting-bridge)# ip address 172.25.44.1/30
switch(config-app-hosting-bridge)# vrf member overlay-VRF
switch(config-app-hosting-bridge)# exit
switch(config)# app-hosting appid te_app
switch(config-app-hosting)# app-vnic bridge 1 guest-interface 0
switch(config-app-hosting-app-vnic)# guest-ipaddress 172.25.44.2/30
switch(config-app-hosting-app-vnic)# exit
switch(config-app-hosting-appid)# app-default-gateway 172.25.44.1 guest-interface 0
switch(config-app-hosting)# end
```

Example: Configuring Docker Run Time Options

The following example shows how to configure docker run time options.

```
switch> enable
switch# configure terminal
switch(config)# app-hosting appid te_app
switch(config-app-hosting)# app-resource docker
switch(config-app-hosting-docker)# run-opts 1 "-v $(APP_DATA) :/data"
switch(config-app-hosting-docker)# end
```

Example: Configuring Application Hosting on the Management Interface

This example shows how to configure application hosting on the management interface.

```
switch> enable
switch# configure terminal
switch(config)# app-hosting appid te_app
switch(config-app-hosting)# app-vnic management guest-interface 0
switch(config-app-hosting)# end
```

Example: Overriding App Resource Configuration

This example shows how to override an app resource configuration.

```
switch> enable
switch# configure terminal
switch(config)# app-hosting appid te_app
switch(config-app-hosting)# app-resource profile custom
switch(config-app-resource-profile-custom)# cpu 7400
switch(config-app-resource-profile-custom)# memory 2048
switch(config-app-resource-profile-custom)# end
```

Additional References

Related Documents

| Related Topic | Document Title |
|------------------------|---|
| Configuring Apphosting | Cisco Nexus 3000 and 9000 Series NX-API REST SDK User Guide and API Reference |
| ThousandEyes URL | https://app.thousandeyes.com |

Technical Assistance

| Description | Link |
|---|---|
| <p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p> | http://www.cisco.com/support |

Feature Information for Application Hosting

The following table provides release information about the feature or features described in this module. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn, you do not need to have an account with Cisco.com.

Table 6: Feature Information for Application Hosting

| Feature Name | Release | Feature Information |
|---|------------------------------|--|
| Support for Cisco Application Hosting Framework (CAF) on Cisco Nexus 9808 platform switches | Cisco NX-OS Release 10.3(3)F | The application hosting feature is now supported on Cisco Nexus 9808 platform switches. |
| Support for Cisco Application Hosting Framework (CAF) on Cisco Nexus 9504 and 9508 platform switches with -R and -R2 line cards | Cisco NX-OS Release 10.3(3)F | The application hosting feature is now supported on Cisco Nexus 9504 and 9508 platform switches with -R and -R2 line cards. This feature is also supported on Cisco N3K-C36180YC-R, N3K-C3636C-R, and N3K-C36480LD-R2 switches |
| Cisco Application Hosting Framework (CAF) | Cisco NX-OS Release 10.3(1)F | <p>A hosted application is a software as a service (SaaS) solution, and you can execute and operate this solution entirely from the cloud. This module describes the Cisco application hosting feature and how to enable it.</p> <p>The application hosting feature is supported on Cisco Nexus 9300 series FX, FX2, GX, and GX2 platforms and Cisco Nexus 9500 series modular switches with FX and GX line cards.</p> |
| Application Hosting: ThousandEyes Integration | Cisco NX-OS Release _____ | <p>ThousandEyes is a cloud-ready, enterprise network-monitoring tool that provides an end-to-end view across networks and services.</p> <ul style="list-style-type: none"> • In Cisco NX-OS Release _____, this feature is implemented on Cisco Nexus 9000 Series Switches. |



CHAPTER 21

ThousandEyes Enterprise Agent

- [ThousandEyes Enterprise Agent Overview, on page 211](#)
- [Prerequisites for the ThousandEyes Enterprise Agent, on page 211](#)
- [Resources Required for the ThousandEyes Enterprise Agent, on page 212](#)
- [Installing the ThousandEyes Enterprise Agent, on page 212](#)
- [Configuration Examples for ThousandEyes Enterprise Agent, on page 217](#)

ThousandEyes Enterprise Agent Overview

ThousandEyes Enterprise Agent is an enterprise network-monitoring tool that provides you an end-to-end view across networks and services that impact your business. It monitors the network traffic paths across internal, external, carrier, and internet networks in real time, to provide network performance data. Enterprise Agents are commonly installed in branch sites and data centers to provide a detailed understanding of WAN and internet connectivity.

The ThousandEyes Enterprise Agent provides the following:

- Benchmarking the performance of networks and applications.
- Detailed hop-by-hop metrics.
- End-to-end path visualization from branch or campus to data center or cloud.
- Outage detection and resolution.
- User-experience analysis.
- Visualization of the traffic-flow pattern.

Prerequisites for the ThousandEyes Enterprise Agent

- The ThousandEyes Enterprise Agent image available at the ThousandEyes site must be signed by the same certificate authority (CA) that is used by <http://www.cisco.com> for HTTPS downloads; without a username or a password.
- Installation of the Enterprise Agent requires internet connectivity, or a proxy server. For more information, see the *ThousandEyes documentation* at: <https://docs.thousandeyes.com/product-documentation/enterprise-agents>.

- The Enterprise Agent application can only be used after the user's license privileges are validated.
- Only Docker-based applications are supported.

Resources Required for the ThousandEyes Enterprise Agent

This table describes the required resources for installing the ThousandEyes Enterprise Agent:

Table 7: Resources Required for the ThousandEyes Enterprise Agent

| App Media | Maximum Resource | Minimum Supported Release |
|-----------|--|------------------------------|
| Bootflash | <ul style="list-style-type: none"> • CPU: 2 vCPU • Memory: 2G RAM • Storage: 1G for persistent logging by applications. | Cisco NX-OS Release 10.4(2)F |

Installing the ThousandEyes Enterprise Agent

Guidelines and Limitations

ThousandEyes Enterprise Agent has the following guidelines and limitations:

- NX-OS supports ThousandEyes Enterprise Agent version 4.4.2 and above.
- It is recommended to use interface eth0 for the container interface in 10.4(2)F and above. If the interface other than the eth0 is used, the app comes up with proper network configuration right after downgrade, but if the app is stopped, deactivated, or started, the default gateway will not be installed and connectivity to the app is not established. If eth1 is used in 10.4(2)F prior to downgrade, after the downgrade, the runtime configuration incorrectly reflects the VNIC as management instead of gateway (front panel port) even though the configuration within the container is correct. To resolve this issue, reconfigure the ThousandEyes Enterprise Agent. See [Installing the ThousandEyes Enterprise Agent, on page 212](#).

Configuring Application Hosting for the ThousandEyes Enterprise Agent

Beginning with Cisco NX-OS 10.4(2)F, you can configure up to 4 IPv4 and IPv6 interfaces within the ThousandEye Enterprise agent.

Follow the steps to install the ThousandEyes Enterprise Agent.

SUMMARY STEPS

1. **configure terminal**
2. **app-hosting bridge** *bridge-index*
3. **ip address** *ip-address/mask*

4. **vrf member** *name*
5. **exit**
6. **app-hosting appid** *name*
7. **app-vnic gateway bridge** *bridge-index* **guest-interface** *guest-interface-number*
8. **guest-ipaddress** *ip-address/mask*
9. **exit**
10. **app-default-gateway** *ip-address* **guest-interfacenumber**
11. **nameserver#** *ip-address*
12. **app-resource docker**
13. **run-opts** *options*
14. **prepend-pkg-opts**
15. **end**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | configure terminal Example: switch# configure terminal | Enters global configuration mode. |
| Step 2 | app-hosting bridge <i>bridge-index</i> Example: switch(config)# app-hosting bridge 1 | Configures the App-hosting bridge and enters App Hosting Bridge configuration mode. <1-8> Bridge index |
| Step 3 | ip address <i>ip-address/mask</i> Example: switch(config-app-hosting-bridge)# ip address 172.25.44.1/30 | Configures the App Bridge IPv4 address which acts as the gateway to the Application Container. |
| Step 4 | vrf member <i>name</i> Example: switch(config-app-hosting-bridge)# vrf member overlay-VRF | Sets the VRF context. If not configured, it would be part of VRF default. |
| Step 5 | exit Example: switch(config-app-hosting-bridge)# exit | Exits App Bridge configuration mode and returns to global configuration mode. |
| Step 6 | app-hosting appid <i>name</i> Example: switch(config)# app-hosting appid te_app | Configures an application and enters application-hosting configuration mode. |

| | Command or Action | Purpose |
|----------------|--|---|
| Step 7 | app-vnic gateway bridge <i>bridge-index</i> guest-interface <i>guest-interface-number</i> Example: <pre>switch(config-app-hosting)# app-vnic bridge 1 guest-interface 0</pre> | Configures the guest VNIC interface for an application and enters application-hosting vnic interface mode. |
| Step 8 | guest-ipaddress <i>ip-address/mask</i> Example: <pre>switch(config-config-app-hosting-app-vnic)# guest-ipaddress 172.25.44.2/30</pre> | Configures one of the available IPv4 address from the bridge 1 subnet. |
| Step 9 | exit Example: <pre>switch(config-config-app-hosting-vlan-access-ip)# exit</pre> | Exits App vnic interface configuration mode and returns to app-hosting configuration mode. |
| Step 10 | app-default-gateway <i>ip-address</i> <i>guest-interfacenumber</i> Example: <pre>switch(config-app-hosting-)# app-default-gateway 172.25.44.1</pre> | Configures the available IPv4 address from the bridge1 subnet. |
| Step 11 | nameserver# <i>ip-address</i> Example: <pre>Device(config-app-hosting)# name-server0 10.2.2.2</pre> | Configures the DNS server. |
| Step 12 | app-resource docker Example: <pre>switch(config-app-hosting)# app-resource docker</pre> | Enters application-hosting docker-configuration mode to specify application resource updates. |
| Step 13 | run-opts <i>options</i> Example: <pre>switch(config-app-hosting-docker)# run-opts 1 "-e TEAGENT_ACCOUNT_TOKEN=[account-token]" run-opts 10 "-e TEAGENT_DEF_IPV4_GW_ETH1=172.25.44.65" run-opts 11 "-e TEAGENT_DEF_IPV6_GW_ETH1=2001:420:287:2003:7:110:1205:1" run-opts 12 "-e TEAGENT_DEF_IPV4_GW_ETH2=172.25.44.73" run-opts 13 "-e TEAGENT_DEF_IPV6_GW_ETH2=2001:420:287:2003:7:110:1206:1" run-opts 14 "-e TEAGENT_DEF_IPV4_GW_ETH3=172.25.44.81" run-opts 15 "-e TEAGENT_DEF_IPV6_GW_ETH3=2001:420:287:2003:7:110:1207:1"</pre> | <p>Specifies the Docker run time options.</p> <p>You can configure additional default gateways through the run-opts option. The TEAGENT_DEF_IPV[4/6]_GW_ETH[0-3] command is used to indicate IPv4 and IPv6 and the index is between 0-3. You can only configure 1 gateway using the app-default-gateway command.</p> |
| Step 14 | prepend-pkg-opts Example: | <p>Merges the package options with the Docker runtime options.</p> <p>Any duplicate variable is overwritten.</p> |

| | Command or Action | Purpose |
|----------------|---|--|
| | switch(config-app-hosting-docker)# prepend-pkg-opts | |
| Step 15 | end Example: switch(config-app-hosting-docker)# end | Exits application-hosting docker-configuration mode and returns to privileged EXEC mode. |

Example

Following is the example for ThousandEye multi interface configuration.

```

app-hosting bridge 1
ip address 172.25.44.113/29
ipv6 address 2001:420:287:2003:7:110:1210:1/125
app-hosting bridge 2
vrf member apphosting
ip address 172.25.44.97/30
ipv6 address 2001:420:287:2003:7:110:1208:1/126
app-hosting bridge 3
vrf member te
ip address 172.25.44.105/29
ipv6 address 2001:420:287:2003:7:110:1209:1/125
app-hosting bridge 4
vrf member vxlan_blue
ip address 172.25.44.33/30
ipv6 address 2001:420:287:2003:7:110:1203:1/114
app-hosting appid tea
app-vnic gateway bridge 1 guest-interface 0
guest-ipaddress 172.25.44.114/29
guest-ipv6address 2001:420:287:2003:7:110:1210:2/125
app-vnic gateway bridge 2 guest-interface 1
guest-ipaddress 172.25.44.98/30
guest-ipv6address 2001:420:287:2003:7:110:1208:2/126
app-vnic gateway bridge 3 guest-interface 2
guest-ipaddress 172.25.44.106/29
guest-ipv6address 2001:420:287:2003:7:110:1209:2/125
app-vnic gateway bridge 4 guest-interface 3
guest-ipaddress 172.25.44.34/30
guest-ipv6address 2001:420:287:2003:7:110:1203:2/114
app-default-gateway 172.25.44.97 guest-interface 1
app-default-ipv6-gateway 2001:420:287:2003:7:110:1208:1 guest-interface 1
app-resource docker
prepend-pkg-opts
run-opts 1 "-e TEAGENT_ACCOUNT_TOKEN=[account-token]"
run-opts 2 "--hostname=southlake2-1-Multi-UsrVRF"
run-opts 5 "-e TEAGENT_PROXY_TYPE=STATIC"
run-opts 6 "-e TEAGENT_PROXY_LOCATION=proxy.domainname.com:80"
run-opts 7 "-e TEAGENT_PROXY_BYPASS_LIST=T*.domainname.com"
run-opts 8 "--dns 8:8:8:8"
run-opts 9 "--dns 8::8"
run-opts 10 "-e TEAGENT_DEF_IPV4_GW_ETH0=172.25.44.113"
run-opts 11 "-e TEAGENT_DEF_IPV6_GW_ETH0=2001:420:287:2003:7:110:1210:1"
run-opts 14 "-e TEAGENT_DEF_IPV4_GW_ETH2=172.25.44.105"
run-opts 15 "-e TEAGENT_DEF_IPV6_GW_ETH2=2001:420:287:2003:7:110:1209:1"
run-opts 16 "-e TEAGENT_DEF_IPV4_GW_ETH3=172.25.44.33"
run-opts 17 "-e TEAGENT_DEF_IPV6_GW_ETH3=2001:420:287:2003:7:110:1203:1"

```

Installing the ThousandEyes Enterprise Agent

Before you begin

You can install the ThousandEyes Enterprise Agent from a file in the bootflash of the switch.

SUMMARY STEPS

1. **app-hosting install appid** *application-name* **package** *package-path*
2. **app-hosting activate appid** *application-name*
3. **app-hosting start appid** *application-name*
4. **end**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | app-hosting install appid <i>application-name</i> package <i>package-path</i> Example: switch# app-hosting install appid lkeys package bootflash:[file path] | Installs an application from the specified location. |
| Step 2 | app-hosting activate appid <i>application-name</i> Example: switch# app-hosting activate appid lkeys | Activates the application hosting configuration mode |
| Step 3 | app-hosting start appid <i>application-name</i> Example: switch# app-hosting start appid lkeys | (Optional) Starts the application. |
| Step 4 | end Example: switch# end | Exits application hosting configuration mode and returns to privileged EXEC mode. |

The following is sample output from the **show app-hosting list** command:

```
switch# show app-hosting list

App id                               State
-----
lkeys                                 RUNNING
```

Configuration Examples for ThousandEyes Enterprise Agent

Examples: Installing ThousandEyes Enterprise Agent

The following example shows how to enable apphosting feature:

```
switch# configure terminal
switch(config)# feature app-hosting
switch(config)# end
```

The following example shows how to configure AppHosting:

```
switch# configure terminal
switch(config)# app-hosting bridge 2

switch(config)# ip address 172.25.44.1/28
switch(config)# app-hosting appid lkeys
switch(config-app-hosting)# app-vnic gateway bridge 1 guest-interface 0
switch(config-config-app-hosting)# guest-ipaddress 172.25.44.3/28
switch(config-config-app-hosting)# exit
switch(config-app-hosting)# app-default-gateway 172.25.44.1 guest-interface 0
switch(config-app-hosting)# name-server0 10.2.2.2
switch(config-app-hosting)# app-resource docker
switch(config-app-hosting-docker)# run-opts 1 "-e TEAGENT_ACCOUNT_TOKEN=[account-token]"
switch(config-app-hosting-docker)# prepend-pkg-opts
switch(config-app-hosting-docker)# end
```

The following example shows how to configure the SVI:

```
switch(config)# interface Vlan606
switch(config-if)# no shutdown
switch(config-if)# vrf member red
switch(config-if)# no ip redirects
switch(config-if)# ip address 172.30.2.193/26
switch(config-if)# ip proxy-arp
switch(config-if)# ip local-proxy-arp
switch(config-if)# interface Ethernet1/15
switch(config-if)# switchport
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 606
switch(config-if)# no shutdown
```

Sample Configuration for ThousandEyes Enterprise Agent

The following is sample output from the `show app-hosting detail` command:

```
switch# show app-hosting detail
App id : lkeys
Owner : appmgr
State : DEPLOYED
Application
Type : docker
Name : ThousandEyes Enterprise Agent
Version : 4.0.2
Description :
Author : ThousandEyes <support@thousandeyes.com>
Path : /bootflash/thousandeyes-enterprise-agent-4.0.2.cisco.tar.gz
URL Path :
```

```

Activated profile name : custom
Resource reservation
Memory : 2048 MB
Disk : 51 MB
CPU : 7400 units
Platform resource profiles
Profile Name CPU(unit) Memory(MB) Disk(MB)
-----
Attached devices
Name Type Alias
-----
iox_trace serial/trace serial3
iox_syslog serial/syslog serial2
iox_console_aux serial/aux serial1
iox_console_shell serial/shell serial0
Network interfaces
-----
Docker
-----
Run-time information
Command :
Entry-point :
Run options in use :
Package run options :
Application health information
Status : 0
Last probe error :
Last probe output :

```

The following sample output from the **show running-configuration** command without proxy server:

```

switch# show running-config app-hosting

feature app-hosting
app-hosting signed-verification disable
app-hosting bridge 2
ip address 172.25.44.33/28
app-hosting appid lkeys
app-vnic gateway bridge 2 guest-interface 0
guest-ipaddress 172.25.44.35/28
app-default-gateway 172.25.44.33 guest-interface 0
name-server0 171.70.168.183
name-server1 173.36.131.10
app-resource docker
prepend-pkg-opts
run-opts 1 "-e TEAGENT_ACCOUNT_TOKEN = [account-token]"
run-opts 2 "--hostname=southlake2-1"
run-opts 3 "--cap-add=NET_ADMIN"
run-opts 4 "--mount type=tmpfs,destination=/var/log/agent,tmpfs-size=140m"
run-opts 5 "-e TEAGENT_PROXY_TYPE=STATIC"
run-opts 6 "-e TEAGENT_PROXY_LOCATION=proxy.domainname.com:80"
run-opts 7 "-e TEAGENT_PROXY_BYPASS_LIST=T*.domainname.com"

```




PART **IV**

NX-API

- [NX-API CLI, on page 221](#)
- [NX-API REST, on page 263](#)
- [NX-API Developer Sandbox, on page 269](#)



CHAPTER 22

NX-API CLI

- [About NX-API CLI, on page 221](#)
- [Using NX-API CLI, on page 223](#)
- [Kernel Stack ACL, on page 248](#)
- [Table of NX-API Response Codes, on page 249](#)
- [JSON and XML Structured Output, on page 252](#)
- [Sample NX-API Scripts, on page 261](#)

About NX-API CLI

NX-API CLI is an enhancement to the Cisco NX-OS CLI system, which supports XML output. NX-API CLI also supports JSON output format for specific commands.

On Cisco Nexus switches, command-line interfaces (CLIs) are run only on the switch. NX-API CLI improves the accessibility of these CLIs by making them available outside of the switch by using HTTP/HTTPS. You can use this extension to the existing Cisco NX-OS CLI system on the switches. NX-API CLI supports **show** commands, configurations, and Linux Bash.

NX-API CLI supports JSON-RPC.

Guidelines and Limitations

- NX-API CLI spawns VSH to execute Cisco NX-OS CLIs on a switch. The VSH timeout limit is 5 minutes. If the Cisco NX-OS CLIs take longer than 5 minutes to execute, the commands fail with the message: "Back-end processing error.". This is governed by the NX-API command timeout, which governs how long a command requested via NX-API can run. It is fixed at 300s and cannot be changed.
- Beginning with Cisco NX-OS Release 10.2(1)F, can use **system server session cmd-timeout** to increase the timeout.
- NX-API spawns the worker processes and load balances the request between the worker processes.
 - The number of nginx backend worker processes is 4.
 - The number of nginx backend worker processes in N3k and the low memory-based platform is 2.
- Each worker process maintains a pool of 5 persistent VSH sessions. Each VSH session is uniquely identified with a combination of username and remote IP from the incoming request. Whenever a new

request comes, the worker process checks if a matching username and remote IP entry is already present, if yes then use the corresponding VSH session else a new VSH session is created based on the availability in the pool and a new entry is added into the pool. If a worker process is already running with the max allowed VSH sessions, then the new request will be rejected, and an appropriate error message will be returned in the response.

- The number of VSH sessions per worker process is a hardcoded value and cannot be configured. The total number of sessions that can exist at any point in time is 20.

Chunk-mode

- Chunk mode supports only 2 concurrent sessions. If the chunk option is selected, then it can be given only in 2 parallel sessions at a time.
- The maximum size of response supported for chunk mode is 200MB until the release 10.3(1)F release.
- After the 10.3(1)F release, the chunk mode supports the response size, until the space is available in the volatile (which is approximately 2.0GB). The size of chunk mode response supports depends on the space in the volatile. Once volatile is 90% full, chunk mode returns failure when first the show output is collected to file. The chunk size supported for each response is 10MB.

Transport

NX-API uses HTTP/HTTPS as its transport. CLIs are encoded into the HTTP/HTTPS POST body.

Starting with Cisco NX-OS Release 9.2(1), the NX-API feature is enabled by default on HTTPS port 443. HTTP port 80 is disabled.

NX-API is also supported through UNIX Domain Sockets for applications running natively on the host or within Guest Shell.

The NX-API backend uses the Nginx HTTP server. The Nginx process, and all its children processes, are under the Linux cgroup protection where the CPU and memory usage is capped. The NX-API processes are part of the cgroup `ext_ser_nginx`, which is limited to 2,147,483,648 bytes of memory. If the Nginx memory usage exceeds the cgroup limitations, the Nginx process is restarted and the NX-API configuration (the VRF, port, and certificate configurations) is restored.

Message Format

NX-API is an enhancement to the Cisco NX-OS CLI system, which supports XML output. NX-API also supports JSON output format for specific commands.



Note

- NX-API XML output presents information in a user-friendly format.
- NX-API XML does not map directly to the Cisco NX-OS NETCONF implementation.
- NX-API XML output can be converted into JSON.

Security

- NX-API supports HTTPS. All communication to the device is encrypted when you use HTTPS.
- NX-API does not support insecure HTTP by default.
- NX-API does not support weak TLSv1 protocol by default.

NX-API is integrated into the authentication system on the device. Users must have appropriate accounts to access the device through NX-API. NX-API uses HTTP basic authentication. All requests must contain the username and password in the HTTP header.



Note You should consider using HTTPS to secure your user's login credentials.

You can enable NX-API by using the **feature** manager CLI command. NX-API is disabled by default.

NX-API provides a session-based cookie, **nxapi_auth** when users first successfully authenticate. With the session cookie, the username and password are included in all subsequent NX-API requests that are sent to the device. The username and password are used with the session cookie to bypass performing the full authentication process again. If the session cookie is not included with subsequent requests, another session cookie is required and is provided by the authentication process. Avoiding unnecessary use of the authentication process helps to reduce the workload on the device.



Note A **nxapi_auth** cookie expires in 600 seconds (10 minutes). This value is a fixed and cannot be adjusted.



Note NX-API performs authentication through a programmable authentication module (PAM) on the switch. Use cookies to reduce the number of PAM authentications, which reduces the load on the PAM.

Using NX-API CLI

The commands, command type, and output type for the Cisco Nexus 9000 Series switches are entered using NX-API by encoding the CLIs into the body of a HTTP/HTTPS POST. The response to the request is returned in XML or JSON output format.



Note For more details about NX-API response codes, see [Table of NX-API Response Codes, on page 249](#).

NX-API CLI is enabled by default for local access. The remote HTTP access is disabled by default.

The following example shows how to configure and launch the NX-API CLI:

- Enable the management interface.

```
switch# conf t
Enter configuration commands, one per line.
```

```

End with CNTL/Z.
switch(config)# interface mgmt 0
switch(config-if)# ip address 10.126.67.53/25
switch(config-if)# vrf context management
switch(config-vrf)# ip route 0.0.0.0/0 10.126.67.1
switch(config-vrf)# end
switch#

```

- Enable the NX-API **nxapi** feature.

```

switch# conf t
switch(config)# feature nxapi

```

The following example shows a request and its response in XML format:

Request:

```

<?xml version="1.0" encoding="ISO-8859-1"?>
<ins_api>
  <version>0.1</version>
  <type>cli_show</type>
  <chunk>0</chunk>
  <sid>session1</sid>
  <input>show switchname</input>
  <output_format>xml</output_format>
</ins_api>

```

Response:

```

<?xml version="1.0"?>
<ins_api>
  <type>cli_show</type>
  <version>0.1</version>
  <sid>eoc</sid>
  <outputs>
    <output>
      <body>
        <hostname>switch</hostname>
      </body>
      <input>show switchname</input>
      <msg>Success</msg>
      <code>200</code>
    </output>
  </outputs>
</ins_api>

```

The following example shows a request and its response in JSON format:

Request:

```

{
  "ins_api": {
    "version": "0.1",
    "type": "cli_show",
    "chunk": "0",
    "sid": "session1",
    "input": "show switchname",
    "output_format": "json"
  }
}

```

Response:

```

{
  "ins_api": {
    "type": "cli_show",
    "version": "0.1",
    "sid": "eoc",
    "outputs": {
      "output": {
        "body": {
          "hostname": "switch"
        },
        "input": "show switchname",
        "msg": "Success",
        "code": "200"
      }
    }
  }
}

```



Note There is a known issue where an attempt to delete a user might fail, resulting in an error message similar to the following appearing every 12 hours or so:

```
user delete failed for username:userdel: user username is currently logged in - securityd
```

This issue might occur in a scenario where you try to delete a user who is still logged into a switch through NX-API. Enter the following command in this case to try to log the user out first:

```
switch(config)# clear user username
```

Then try to delete the user again. If the issue persists after attempting this workaround, contact Cisco TAC for further assistance.

Escalate Privileges to Root on NX-API

For NX-API, the privileges of an admin user can escalate their privileges for root access.

The following are guidelines for escalating privileges:

- Only an admin user can escalate privileges to root.
- Escalation to root is password protected.

The following examples show how an admin escalates privileges to root and how to verify the escalation. Note that after becoming root, the **whoami** command shows you as admin; however, the admin account has all the root privileges.

First example:

```

<?xml version="1.0"?>
<ins_api>
  <version>1.0</version>
  <type>bash</type>
  <chunk>0</chunk>
  <sid>sid</sid>
  <input>sudo su root ; whoami</input>
  <output_format>xml</output_format>
</ins_api>

```

```
<?xml version="1.0" encoding="UTF-8"?>
<ins_api>
  <type>bash</type>
  <version>1.0</version>
  <sid>eoc</sid>
  <outputs>
    <output>
      <body>admin </body>
      <code>200</code>
      <msg>Success</msg>
    </output>
  </outputs>
</ins_api>
```

Second example:

```
<?xml version="1.0"?>
<ins_api>
  <version>1.0</version>
  <type>bash</type>
  <chunk>0</chunk>
  <sid>sid</sid>
  <input>sudo cat path_to_file </input>
  <output_format>xml</output_format>
</ins_api>
```

```
<?xml version="1.0" encoding="UTF-8"?>
<ins_api>
  <type>bash</type>
  <version>1.0</version>
  <sid>eoc</sid>
  <outputs>
    <output>
      <body>[Contents of file]</body>
      <code>200</code>
      <msg>Success</msg>
    </output>
  </outputs>
</ins_api>
```

NX-API Management Commands

You can enable and manage NX-API with the CLI commands listed in the following table.

Table 8: NX-API Management Commands

| NX-API Management Command | Description |
|--|----------------------|
| feature nxapi | Enables NX-API. |
| no feature nxapi | Disables NX-API. |
| nxapi {http https} port <i>port</i> | Specifies a port. |
| no nxapi {http https} | Disables HTTP/HTTPS. |

| NX-API Management Command | Description |
|---|--|
| show nxapi | <p>Displays port and certificate information.</p> <p>Note The "show nxapi" command doesn't display certificate/config information for network-operator role.</p> |
| nxapi certificate {httpsrt certfile httpskey keyfile} filename | <p>Specifies the upload of the following:</p> <ul style="list-style-type: none"> • HTTPS certificate when httpsrt is specified. • HTTPS key when httpskey is specified. <p>Example of HTTPS certificate:</p> <pre>nxapi certificate httpsrt certfile bootflash:cert.crt</pre> <p>Example of HTTPS key:</p> <pre>nxapi certificate httpskey keyfile bootflash:privkey.key</pre> |
| nxapi certificatehttpskey keyfile filename password passphrase | <p>Installs NX-API certificates with encrypted private keys:</p> <p>Note The passphrase for decrypting the encrypted private key is pass123!.</p> <p>Example:</p> <pre>nxapi certificate httpskey keyfile bootflash:encr-cc.pem password pass123!</pre> |
| nxapi certificate enable | Enables a certificate. |

| NX-API Management Command | Description |
|--|--|
| nxapi certificate trustpoint <trustpoint label> | <p>Beginning with Cisco NX-OS release 10.2(3)F, the user can now import the certificate or use the CA certificate for the NX-API using the trustpoint infra.</p> <p>Note Refer to the <i>Cisco Nexus 9000 Security Configuration Guide</i> to configure the crypto ca import trustpoint to first import certificate.</p> <p>Note Currently only pkcs12 certificate import is supported in this form. The NX-API certificate enable/NX-API certificate trustpoint and NX-API certificate sudi are mutually exclusive and each configuration will overwrite the certificate/key.</p> <p>Note The maximum size of cert/key supported with NX-API certificate enable is 8k. If the size is >8k, use NX-API certificate trustpoint to import the certificate.</p> <p>Note If you have configured a custom certificate in NX-API using trustpoint infra, upon entering the reload ascii command the configuration is lost. It will revert to the default day-1 NX-API certificate. After entering the reload ascii command, the switch will reload. Once the switch is up again, you need to reconfigure the NX-API certificate trustpoint configuration.</p> <p>Note Beginning with Cisco NX-OS Release 10.3(1)F, support for ascii trustpoint reload is added.</p> <p>Note Config-replace will fail if the current running-config do not contain the trustpoint and certificate imported, but the target config contains the creation of trustpoint "crypto ca trustpoint <trustpoint name>" and "nxapi certificate trustpoint <trustpoint-name>" CLI. If trustpoint is not present, then first you need to create trustpoint and import certificate before attempting "nxapi certificate trustpoint <trustpoint-label>".</p> |

| NX-API Management Command | Description |
|----------------------------------|--|
| nxapi certificate sudi | <p>This CLI provides a secure way of authenticating to the device by using Secure Unique Device Identifier (SUDI).</p> <p>The SUDI based authentication in nginx will be used by the CISCO SUDI compliant controllers.</p> <p>SUDI is an IEEE 802.1AR-compliant secure device identity in an X.509v3 certificate which maintains the product identifier and serial number of Cisco devices. The identity is implemented at manufacturing and is chained to a publicly identifiable root certificate authority.</p> <p>Note When NX-API comes up with the SUDI certificate, it is not accessible by any third-party applications like browser, curl, and so on.</p> <p>Note "nxapi certificate sudi" will overwrite the custom certificate/key if configured, and there is no way to get the custom certificate/key back.</p> <p>Note "nxapi certificate sudi" and "nxapi certificate trustpoint" and "nxapi certificate enable" are mutually exclusive , and configuring one will delete the other configuration.</p> <p>Note NX-API do not support SUDI certificate-based client certificate authentication. If client certificate authentication is needed, then Identity certificate need to be used.</p> <p>Note As NX-API certificate CLI is not present in show run output, CR/Rollback case currently does not go back to the custom certificate once it is overwritten with "nxapi certificate sudi" options.</p> |
| no nxapi certificate sudi | This will disable the SUDI and NX-API will come with a default self-signed certificate. |
| nxapi ssl-ciphers weak | Beginning with Cisco NX-OS Release 9.2(1), weak ciphers are disabled by default. Running this command changes the default behavior and enables the weak ciphers for NGINX. The no form of the command changes it to the default (by default, the weak ciphers are disabled). |

| NX-API Management Command | Description |
|--|--|
| nxapi ssl-protocols {TLSv1.0 TLSv1.1 TLSv1.2 TLSv1.3} | <p>Beginning with Cisco NX-OS Release 10.2(4)M, TLSv1.3 is supported on Cisco Nexus 9000 series platform switches. Running this command enables the TLS versions specified in the string. Beginning with Cisco NX-OS Release 9.3(2), only TLSv1.2 is enabled by default.</p> <p>The no form of the command changes the TLS version to the default version.</p> <ul style="list-style-type: none"> If you want to enable particular TLS version, specify only that respective TLS version. <p>For example, if you need TLSv1.3, use the following command:</p> <pre>switch(config)# nxapi ssl protocols TLSv1.3</pre> <ul style="list-style-type: none"> If you want to enable multiple TLS versions for backward compatibility at a later stage, specify all the required TLS versions that are supported. <p>For example:</p> <ul style="list-style-type: none"> If you need TLSv1.1 through TLSv1.3, use the following command to enable all required TLS versions: <pre>switch(config)# nxapi ssl protocols TLSv1.2 TLSv1.3</pre> <ul style="list-style-type: none"> When you need backward compatibility, use the following command to enable that version: <pre>switch(config)# nxapi ssl protocols TLSv1.2</pre> <p>Note</p> <ul style="list-style-type: none"> It is recommended to use TLSv1.2 and TLSv1.3 for backward compatibility. <pre>switch(config)# nxapi ssl protocols TLSv1.2 TLSv1.3</pre> <p>For example, if you are :</p> <ul style="list-style-type: none"> Before configuring TLSv1.3, validate the server and client certificates for TLSv1.3 support. NX-API server side SUDI certificate is not supported with TLSv1.3. |
| nxapi use-vrf <i>vrf</i> | <p>Specifies the default VRF, management VRF, or named VRF.</p> <p>Note In Cisco NX-OS Release 7.0(3)I2(1) NGINX listens on only one VRF.</p> |

| NX-API Management Command | Description |
|--|--|
| system server session cmd-timeout <timeout> | Beginning with Cisco NX-OS release, 10.2(3)F, in NGINX server, the default timeout to run any command is 5 minutes. The users can increase the timeout to the desired value from 60 seconds (1 minute) to 3600 seconds (1 hour) according to their need and time taken for executing the commands. |
| ip netns exec management iptables | <p>Implements any access restrictions and can be run in management VRF.</p> <p>Note You must enable feature bash-shell and then run the command from Bash Shell. For more information on Bash Shell, see the chapter on Bash.</p> <p><i>Iptables is a command-line firewall utility that uses policy chains to allow or block traffic and almost always comes pre-installed on any Linux distribution.</i></p> <p>Note For more information about making iptables persistent across reloads when they are modified in a bash-shell, see Making an Iptable Persistent Across Reloads, on page 247.</p> |
| nxapi idle-timeout <timeout> | Starting with Release 9.3(5), you can configure the amount of time before an idle NX-API session is invalidated. The time can be 1 - 1440 minutes. The default time is 10 minutes. Return to the default value by using the no form of the command: no nxapi idle-timeout <timeout> |

The following is an example for NX-API output for SUDI:

```
switch(config)# nxapi certificate sudi
switch# show nxapi
nxapi enabled
NXAPI timeout 10
NXAPI cmd timeout 300
HTTP Listen on port 80
HTTPS Listen on port 443
Certificate Information:
  Issuer:  issuer=CN = High Assurance SUDI CA, O = Cisco
  Expires: Aug  9 20:58:26 2099 GMT
switch#
switch#
switch# show run | sec nxapi
feature nxapi
nxapi http port 80
nxapi certificate sudi
switch#
```

The following is an example for trustpoint configuration:

```
switch(config)# crypto ca trustpoint ngx
switch(config-trustpoint)# crypto ca import ngx pkcs12 bootflash:server.pfx cisco123
switch(config)# nxapi certificate trustpoint ngx
switch(config)# show nxapi
nxapi enabled
NXAPI timeout 10
NXAPI cmd timeout 300
HTTP Listen on port 80
```

```
Trustpoint label ngx
HTTPS Listen on port 443
Certificate Information:
Issuer: issuer=C = IN, ST = KA, L = bang, O = cisco, OU = nxpi, CN = %username%@cisco.com,
    emailAddress = %username%@cisco.com
Expires: Jan 13 06:13:50 2023 GMT
switch(config)#
switch(config)# show run | sec nxapi
feature nxapi
nxapi http port 80
nxapi certificate trustpoint ngx
```

Following is an example of a successful upload of an HTTPS certificate:

```
switch(config)# nxapi certificate https crt certfile certificate.crt
Upload done. Please enable. Note cert and key must match.
switch(config)# nxapi certificate enable
switch(config)#
```



Note You must configure the certificate and key before enabling the certificate.

Following is an example of a successful upload of an HTTPS key:

```
switch(config)# nxapi certificate httpskey keyfile bootflash:privkey.key
Upload done. Please enable. Note cert and key must match.
switch(config)# nxapi certificate enable
switch(config)#
```

The following is an example of how to install an encrypted NXAPI server certificate:

```
switch(config)# nxapi certificate https crt certfile bootflash:certificate.crt
switch(config)# nxapi certificate httpskey keyfile bootflash:privkey.key password pass123!

switch(config)#nxapi certificate enable
switch(config)#
```

In some situations, you might get an error message saying that the key file is encrypted:

```
switch(config)# nxapi certificate https crt certfile bootflash:certificate.crt
switch(config)# nxapi certificate httpskey keyfile bootflash:privkey.key
ERROR: Unable to load private key!
Check keyfile or provide pwd if key is encrypted, using 'nxapi certificate httpskey keyfile
<keyfile> password <passphrase>'.
```

In this case, the passphrase of the encrypted key file must be specified using **nxapi certificatehttpskey keyfile filename password passphrase**.

If this was the reason for the issue, you should now be able to successfully install the certificate:

```
switch(config)# nxapi certificate httpskey keyfile bootflash:privkey.key password pass123!
switch(config)# nxapi certificate enable
switch(config)#
```

Working With Interactive Commands Using NX-API

To disable confirmation prompts on interactive commands and avoid timing out with an error code 500, prepend interactive commands with **terminal dont-ask**. Use **;** to separate multiple interactive commands, where each **;** is surrounded with single blank characters.

Following are several examples of interactive commands where **terminal dont-ask** is used to avoid timing out with an error code 500:

```
terminal dont-ask ; reload module 21
terminal dont-ask ; system mode maintenance
```

NX-API Client Authentication

NX-API Client Basic Authentication

NX-API clients can authenticate with the NGINX server on the switch through basic authentication over SSL/TLS. This authentication method is supported by configuring a username and password that is saved to a database on the switch. When the NX-API client initiates a connection request, it sends the Hello message which contains the username and password. Assuming the username and password exist in the database, the switch responds by sending the Hello response, which contains a cookie. After this initial handshake is complete, the communication session is open, and the client can begin sending API calls to the switch. For additional information, see [Security, on page 223](#).

For additional information about basic authentication, including how to configure the username and password on the switch, refer to the [Cisco Nexus 9000 Series NX-OS Security Configuration Guide](#).

NX-API Client Certificate Authentication

Beginning with NX-OS 9.3(3), NX-API supports client-initiated certificate-based authentication. Certificate-based authentication offers stronger security by mutually authenticating both the client, using a trusted party—the Certificate Authority (CA)—and the server during the TLS handshake. Certificate-based authentication allows for human authentication, as well as machine authentication, for accessing the NX-OS switch.

Client certificate authentication is supported by using an X509 SSL certificate that is assigned through a valid CA (certificate authority) and stored on the NX-API client. A certificate is assigned to each NX-API username.

When the NX-API client initiates a connection request with a Hello message, the server Hello response contains the list of valid CAs. The client's response contains additional information elements, including the certificate for the specific username that the NX-API client is using.

You can configure the NX-API client to use either basic authentication, certificate authentication, or give priority to certificate but fallback to basic authentication if the certificate authentication method is not available.

Guidelines and Limitations

Certificate authentication has the following guidelines and limitations:

- The NX-API client must be configured with a user name and password.
- The NX-API client and the switch communicate over HTTP by default on its well-known port. For flexibility HTTP is also supported on its well-known port. However, you can configure additional ports.
- Python scripting of client certificate authentication is supported. If the client certificate is encrypted with a passphrase, python successfully prompts for the passphrase. However, the passphrase cannot be passed into the script due to a current limitation with the Python requests library.
- The NX-API client and switch must use the same trustpoint.
- The maximum number of trustpoints supported is 26 for each switch.

- The list of trusted CAs must be the same for all NX-API clients and the switch. Separate lists of trusted CAs are not supported.
- Certificate authentication is not supported for the NX-API sandbox.
- The following conditions determine if the NX-API sandbox loads on the switch:
 - The NX-API sandbox loads only when **nxapi client certificate authentication optional** or **no nxapi client certificate authentication** are configured.
 - The NX-API sandbox does not load for **strict** and **two-step** authentication modes unless a valid client certificate is presented to the browser when a connection is being established.
- The switch has an embedded NGINX server. If multiple trustpoints are configured, but a certificate revocation list (CRL) is installed for only one of the trustpoints, NX-API client certificate authentication fails because of an NGINX limitation. To workaround this limitation, configure CRLs for all trustpoints.
- Certificates can expire or become out of date, which can affect the validity of the CRL set by the CA (trustpoint). To ensure the switch uses valid CRLs, always install CRLs for all of the configured trustpoints. If no certificates were revoked by the trustpoints, an empty CRL should be generated, installed, and updated periodically, for example, once a week.

After you update the CRLs through the crypto CLIs, issue **nxapi client cert authentication** to reapply the newly updated CRLs.
- If you use ASCII reload when NX-API client certificate authentication is enabled, you must issue **nxapi client certificate authentication** after the reload is complete.
- The certificate path must terminate with a trusted CA certificate.
- Server certificates that are presented for TLS must have the Server Authentication purpose (id-kp 1 with OID 1.3.6.1.5.5.7.3.1) in the `extendedKeyUsage` field.
- Client certificates that are presented for TLS must have the Server Authentication purpose (id-kp 1 with OID 1.3.6.1.5.5.7.3.2) in the `extendedKeyUsage` field.
- The feature supports CRLs (certificate revocation lists). Online Certificate Status Protocol (OSCP) is not supported.
- Follow the additional Guidelines and Limitations in the NX-OS Security Guide.
 - Use both certificate and basic authentication. By doing so, the correct user and password is still required if the certificate somehow gets compromised.
 - Keep private keys private, as the servers public key is accessible to anyone attempting a connection.
 - CRLs should be downloaded from the central CA and kept current. Out-of-date CRLs can lead to a security risk.
 - Keep trustpoints updated. When a trust point or configuration change is made to the certificate authentication feature, explicitly disable then reenble the feature to reload the updated information.
- There is a maximum file size limit of 8K for the client certificate identity file associated to NX-API with **nxapi certificate httpscert certfile bootflash:<> " CLI."** This is a day-1 limitation.
- In the NX-API Management Commands Table 1 for the row associated with the command `nxapi certificate {httpscert certfile | httpskey keyfile} filename`, the maximum certfile size supported is less than 8K.

NX-API Client Certificate Authentication Prerequisites

Before configuring certificate authentication, make sure the following are present on the switch:

1. Configure the client with a username and password. For information see [Configuring User Accounts and RBAC](#).
2. Configure the CA(s) (trustpoint) and CRL(s) (if any).

If no certificates were revoked by a trustpoint, create a blank CRL for each trustpoint.

For information, see the [Cisco Nexus 9000 Series NX-OS Security Configuration Guide](#).

Configuring NX-API Client Certificate Authentication

You can configure the NX-API certificate authentication through the **nxapi client certificate authentication** command. The command supports restriction options that control how authentication occurs.

You can disable this feature by using **no nxapi client certificate authentication**.

To configure certificate authentication for NX-API clients, follow this procedure:

SUMMARY STEPS

1. Make sure the prerequisites for the feature are complete.
2. **config terminal**
3. **nxapi client certificate authentication** [{optional | strict | two-step}]

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | Make sure the prerequisites for the feature are complete. | See NX-API Client Certificate Authentication Prerequisites, on page 235 . |
| Step 2 | config terminal Example: <pre>switch-1# config terminal</pre> Enter configuration commands, one per line. End with CNTL/Z. <pre>switch-1(config)#</pre> | Enters configuration mode. |
| Step 3 | nxapi client certificate authentication [{optional strict two-step}] Example: <pre>switch-1# nxapi client certificate authentication strict</pre> <pre>switch-1(config)#</pre> | Enables certificate authentication in any of the following modes: <ul style="list-style-type: none"> • optional requests a client certificate: <ul style="list-style-type: none"> • If the client provides a certificate, mutual verification occurs between the client and the server. |

| | Command or Action | Purpose |
|--|-------------------|--|
| | | <ul style="list-style-type: none"> • If the client provides an invalid certificate, authentication fails and fall back to basic authentication does not occur. • If the client does not provide a certificate, authentication falls back to basic authentication (username and password). • strict enables client certificate verification and requires a valid client certificate to be presented for authentication. • two-step enables two-step verification in which both the basic authentication and certificate authentication methods are required. <p>Note If no trustpoints are configured on the switch, this feature cannot be enabled, and the switch displays an onscreen error message.</p> <p>No trustpoints configured! Please configure trustpoint using 'crypto ca trustpoint <trustpoint-label>' and associated commands, and then enable this feature.</p> |

Example Python Scripts for Certificate Authentication

The following example shows a Python script with a client certificate for authentication.

```
import requests
import json

"""
Modify these please
"""
switchuser='USERID'
switchpassword='PASSWORD'
mgmtip='NXOS MANAGEMENT IP/DOMAIN NAME'

client_cert_file='PATH_TO_CLIENT_CERTIFICATE'
client_key_file='PATH_TO_CLIENT_KEY_FILE'
ca_cert='PATH_TO_CA_CERT_THAT_SIGNED_NXAPI_SERVER_CERT'

url='https://' + mgmtip + '/ins'
myheaders={'content-type':'application/json-rpc'}
payload=[
    {
        "jsonrpc": "2.0",
        "method": "cli",
        "params": {
            "cmd": "show clock",
            "version": 1
        },
        "id": 1
    }
]
```

```
response = requests.post(url,data=json.dumps(payload),
headers=myheaders,auth=(switchuser,switchpassword),cert=(client_cert_file_path,client_key_file),verify=ca_cert).json()
```

If needed, you can change the script:

- Depending on the client certificate authentication mode, you can omit the switch password by setting the switch password to a null value (`switchpassword=`):
 - For **optional** and **strict** modes, the `switchpassword=` can be left blank. In this situation, NX-API authenticates the client based on username and client certificate alone.
 - For **two-step** mode, a password is required, so you must specify a value for `switchpassword=`.
- You can bypass verifying that the NX-API server's certificate is valid by setting `verify=False` in the POST command.

Example cURL Certificate Request

The following example shows a correctly structured cURL certificate request for NX-API client authentication.

```
/usr/bin/curl --user admin: --tlsv1.2 --cacert ./ca.pem --cert ./user.crt:pass123! --key
./user.key -v -X POST -H "Accept: application/json" -H "Content-type: application/json"
--data '{"ins_api":{"version": "1.0", "type": "cli_show", "chunk": "0", "sid": "1", "input":
"show clock","output_format": "json"}}' https://<device-management-ip>:443/ins
```

Syntax Elements

The following table shows the parameters that are used in this request.

| Parameter | Description |
|-----------------|--|
| --user | Takes the username that the user wants to log in as, which should be same as the common name in user.crt). To provide a password for user, specify it after a colon, for example: --user username:password |
| --cacert | Takes the path to the CA that signed the NX-API server certificate. If the server certificate does not need to be verified, specify cURL with the -k (insecure) option, for example: /usr/bin/curl -k |
| --cert | Takes the path to the client certificate. If the client certificate is encrypted, specify the password after a colon, for example: --cert user.crt:pass123! |
| --key | Takes the path to the client certificate's private key. |

Validating Certificate Authentication

When correctly configured, certificate authentication occurs and the NX-API clients can access the switch.

If the NX-API client cannot access the switch, you can use the following guidelines to assist with troubleshooting:

SUMMARY STEPS

1. Check user or cookie errors.
2. Check for client or certificate errors.
3. If errors occur, flap the feature to reload any changes to the trustpoint, CA, CRL, or NX-OS certificate feature, by issuing **no nxapi client certificate authentication** , then **nxapi client certificate authentication** .

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | Check user or cookie errors. | <p>If any of the following errors occur:</p> <ul style="list-style-type: none"> • No username provided in auth header and no valid cookie provided • Incorrect user provided in auth header • Invalid cookie provided • Mismatch between username in auth header and username in client certificate's CN field <p>You will see specific errors depending on the NX-API method used:</p> <ul style="list-style-type: none"> • For JSON/XML, a 401 Authentication failure - user not found. error occurs. For example: <pre> {{{ "code": "400", "msg": "Authentication failure - user not found." }}}</pre> • For JSON RPC 2.0, a -32004 Invalid username or password error occurs. For example: <pre> {{ "code": -32004, "message": "Invalid username or password" }}</pre> |
| Step 2 | Check for client or certificate errors. | <p>Look for HTTPs 400 errors which can indicate the following:</p> <ul style="list-style-type: none"> • If an invalid or revoked client certificate was provided. • If the CRL configured on the switch has expired. <p>For example:</p> |

| | Command or Action | Purpose |
|---------------|--|---|
| | | <pre><html> <head><title>400 The SSL certificate error</title></head> <body bgcolor="white"> <center><h1>400 Bad Request</h1></center> <center>The SSL certificate error</center> <hr<center>nginx/1.7.10</center> </body> </html></pre> |
| Step 3 | If errors occur, flap the feature to reload any changes to the trustpoint, CA, CRL, or NX-OS certificate feature, by issuing no nxapi client certificate authentication , then nxapi client certificate authentication . | Disables, then reenables certificate authentication. |

NX-API Request Elements

NX-API request elements are sent to the device in XML format or JSON format. The HTTP header of the request must identify the content type of the request.

You use the NX-API elements that are listed in the following table to specify a CLI command:



Note Users need to have permission to execute "configure terminal" command. When JSON-RPC is the input request format, the "configure terminal" command will always be executed before any commands in the payload are executed.

Table 9: NX-API Request Elements for XML or JSON Format

| NX-API Request Element | Description |
|------------------------|-------------------------------|
| version | Specifies the NX-API version. |

| NX-API Request Element | Description |
|------------------------|---|
| <i>type</i> | <p>Specifies the type of command to be executed.</p> <p>The following types of commands are supported:</p> <ul style="list-style-type: none"> • cli_show CLI show commands that expect structured output. If the command does not support XML output, an error message is returned. • cli_show_array CLI show commands that expect structured output. Only for show commands. Similar to cli_show, but with cli_show_array, data is returned as a list of one element, or an array, within square brackets []. • cli_show_ascii CLI show commands that expect ASCII output. This aligns with existing scripts that parse ASCII output. Users are able to use existing scripts with minimal changes. • cli_conf CLI configuration commands. • bash Bash commands. Most non-interactive Bash commands are supported by NX-API. <p>Note</p> <ul style="list-style-type: none"> • Each command is only executable with the current user's authority. • The pipe operation is supported in the output when the message type is ASCII. If the output is in XML format, the pipe operation is not supported. • A maximum of 10 consecutive show commands are supported. If the number of show commands exceeds 10, the 11th and subsequent commands are ignored. • No interactive commands are supported. |

| NX-API Request Element | Description | | | | |
|------------------------|---|---|----------------------|---|---------------|
| <i>chunk</i> | <p>Some show commands can return a large amount of output. For the NX-API client to start processing the output before the entire command completes, NX-API supports output chunking for show commands.</p> <p>Enable or disable chunk with the following settings:</p> <table border="1" data-bbox="982 485 1516 596"> <tr> <td data-bbox="982 485 1068 539">0</td> <td data-bbox="1068 485 1516 539">Do not chunk output.</td> </tr> <tr> <td data-bbox="982 539 1068 596">1</td> <td data-bbox="1068 539 1516 596">Chunk output.</td> </tr> </table> <p>Note</p> <ul style="list-style-type: none"> • Only show commands support chunking. When a series of show commands are entered, only the first command is chunked and returned. • The output message format options are XML or JSON. • For the XML output message format, special characters, such as < or >, are converted to form a valid XML message (< is converted into &lt; > is converted into &gt;). <p>You can use XML SAX to parse the chunked output.</p> <ul style="list-style-type: none"> • When the output message format is JSON, the chunks are concatenated to create a valid JSON object. <p>Note When chunking is enabled, the maximum message size supported is currently 200MB of chunked output.</p> | 0 | Do not chunk output. | 1 | Chunk output. |
| 0 | Do not chunk output. | | | | |
| 1 | Chunk output. | | | | |
| <i>rollback</i> | <p>Valid only for configuration CLIs, not for show commands. Specifies the configuration rollback options. Specify one of the following options.</p> <ul style="list-style-type: none"> • Stop-on-error—Stops at the first CLI that fails. • Continue-on-error—Ignores and continues with other CLIs. • Rollback-on-error—Performs a rollback to the previous state the system configuration was in. <p>Note The rollback element is available in the <code>cli_conf</code> mode when the input request format is XML or JSON.</p> | | | | |

| NX-API Request Element | Description | | | | | | |
|------------------------|---|-----------------------|--|-----------------------|--|-------------------|--|
| <i>sid</i> | <p>The session ID element is valid only when the response message is chunked. To retrieve the next chunk of the message, you must specify a <i>sid</i> to match the <i>sid</i> of the previous response message.</p> <p>NX-OS release 9.3(1) introduces the <i>sid</i> option <code>clear</code>. When a new chunk request is initiated with the <i>sid</i> set to <code>clear</code>, all current chunk requests are discarded or abandoned.</p> <p>When you receive response code 429: Max number of concurrent chunk request is 2, use <i>sid</i> <code>clear</code> to abandon the current chunk requests. After using <i>sid</i> <code>clear</code>, subsequent response codes operate as usual per the rest of the request.</p> | | | | | | |
| <i>input</i> | <p>Input can be one command or multiple commands. However, commands that belong to different message types should not be mixed. For example, show commands are <code>cli_show</code> message type and are not supported in <code>cli_conf</code> mode.</p> <p>Note Except for bash, multiple commands are separated with " ; ". (The ; must be surrounded with single blank characters.)</p> <p>Prepend commands with <code>terminal dont-ask</code> to avoid timing out with an error code 500. For example:</p> <pre>terminal dont-ask ; cli_conf ; interface Eth4/1 ; no shut ; switchport</pre> <p>For bash, multiple commands are separated with ";". (The ; is not surrounded with single blank characters.)</p> <p>The following are examples of multiple commands:</p> <p>Note</p> <table border="1" data-bbox="883 1205 1481 1459"> <tbody> <tr> <td data-bbox="883 1205 992 1283"><code>cli_show</code></td> <td data-bbox="992 1205 1481 1283"><code>show version ; show interface brief ; show vlan</code></td> </tr> <tr> <td data-bbox="883 1283 992 1381"><code>cli_conf</code></td> <td data-bbox="992 1283 1481 1381"><code>interface Eth4/1 ; no shut ; switchport</code></td> </tr> <tr> <td data-bbox="883 1381 992 1459"><code>bash</code></td> <td data-bbox="992 1381 1481 1459"><code>cd /bootflash;mkdir new_dir</code></td> </tr> </tbody> </table> | <code>cli_show</code> | <code>show version ; show interface brief ; show vlan</code> | <code>cli_conf</code> | <code>interface Eth4/1 ; no shut ; switchport</code> | <code>bash</code> | <code>cd /bootflash;mkdir new_dir</code> |
| <code>cli_show</code> | <code>show version ; show interface brief ; show vlan</code> | | | | | | |
| <code>cli_conf</code> | <code>interface Eth4/1 ; no shut ; switchport</code> | | | | | | |
| <code>bash</code> | <code>cd /bootflash;mkdir new_dir</code> | | | | | | |

| NX-API Request Element | Description | | | | |
|------------------------|--|-----|---------------------------------|------|----------------------------------|
| <i>output_format</i> | <p>The available output message formats are the following:</p> <p>Note</p> <table border="1" data-bbox="950 338 1516 453"> <tr> <td data-bbox="950 338 1144 392">xml</td> <td data-bbox="1144 338 1516 392">Specifies output in XML format.</td> </tr> <tr> <td data-bbox="950 392 1144 453">json</td> <td data-bbox="1144 392 1516 453">Specifies output in JSON format.</td> </tr> </table> <p>Note The Cisco NX-OS CLI supports XML output, which means that the JSON output is converted from XML. The conversion is processed on the switch.</p> <p>To manage the computational overhead, the JSON output is determined by the amount of output. If the output exceeds 1 MB, the output is returned in XML format. When the output is chunked, only XML output is supported.</p> <p>The content-type header in the HTTP/HTTPS headers indicate the type of response format (XML or JSON).</p> | xml | Specifies output in XML format. | json | Specifies output in JSON format. |
| xml | Specifies output in XML format. | | | | |
| json | Specifies output in JSON format. | | | | |

When JSON-RPC is the input request format, use the NX-API elements that are listed in the following table to specify a CLI command:

Table 10: NX-API Request Elements for JSON-RPC Format

| NX-API Request Element | Description |
|------------------------|--|
| <i>jsonrpc</i> | <p>A string specifying the version of the JSON-RPC protocol.</p> <p>Version must be 2.0.</p> |
| <i>method</i> | <p>A string containing the name of the method to be invoked.</p> <p>NX-API supports either:</p> <ul style="list-style-type: none"> • cli—show or configuration commands • cli_ascii—show or configuration commands; output without formatting • cli_array—only for show commands; similar to cli, but with cli_array, data is returned as a list of one element, or an array, within square brackets, []. |
| <i>params</i> | <p>A structured value that holds the parameter values used during the invocation of a method.</p> <p>It must contain the following:</p> <ul style="list-style-type: none"> • cmd—CLI command • version—NX-API request version identifier |

| NX-API Request Element | Description |
|------------------------|---|
| <i>rollback</i> | Valid only for configuration CLIs, not for show commands. Configuration rollback options. You can specify one of the following options. <ul style="list-style-type: none"> • Stop-on-error—Stops at the first CLI that fails. • Continue-on-error—Ignores the failed CLI and continues with other CLIs. • Rollback-on-error—Performs a rollback to the previous state the system configuration was in. |
| <i>validate</i> | Configuration validation settings. This element allows you to validate the commands before you apply them on the switch. This enables you to verify the consistency of a configuration (for example, the availability of necessary hardware resources) before applying it. Choose the validation type from the Validation Type drop-down list. <ul style="list-style-type: none"> • Validate-Only—Validates the configurations, but does not apply the configurations. • Validate-and-Set—Validates the configurations, and applies the configurations on the switch if the validation is successful. |
| <i>lock</i> | An exclusive lock on the configuration can be specified, whereby no other management or programming agent will be able to modify the configuration if this lock is held. |
| <i>id</i> | An optional identifier established by the client that must contain a string, number, or null value, if it is specified. The value should not be null and numbers contain no fractional parts. If a user does not specify the id parameter, the server assumes that the request is simply a notification, resulting in a no response, for example, <i>id</i> : 1 |

NX-API Response Elements

The NX-API elements that respond to a CLI command are listed in the following table:

Table 11: NX-API Response Elements

| NX-API Response Element | Description |
|-------------------------|--|
| version | NX-API version. |
| type | Type of command to be executed. |
| sid | Session ID of the response. This element is valid only when the response message is chunked. |

| NX-API Response Element | Description |
|-------------------------|---|
| outputs | Tag that encloses all command outputs. When multiple commands are in <code>cli_show</code> or <code>cli_show_ascii</code> , each command output is enclosed by a single output tag. When the message type is <code>cli_conf</code> or <code>bash</code> , there is a single output tag for all the commands because <code>cli_conf</code> and <code>bash</code> commands require context. |
| output | Tag that encloses the output of a single command output. For <code>cli_conf</code> and <code>bash</code> message types, this element contains the outputs of all the commands. |
| input | Tag that encloses a single command that was specified in the request. This element helps associate a request input element with the appropriate response output element. |
| body | Body of the command response. |
| code | Error code returned from the command execution. NX-API uses standard HTTP error codes as described by the Hypertext Transfer Protocol (HTTP) Status Code Registry (http://www.iana.org/assignments/http-status-codes/http-status-codes.xhtml). |
| msg | Error message associated with the returned error code. |

Restricting Access to NX-API

There are two methods for restricting HTTP and HTTPS access to a device: ACLs and iptables. The method that you use depends on whether you have configured a VRF for NX-API communication using the `nxapi use-vrf <vrf-name>` CLI command.

Use ACLs to restrict HTTP or HTTPS access to a device only if you have not configured NXAPI to use a specific VRF. For information about configuring ACLs, see the *Cisco Nexus Series NX-OS Security Configuration Guide* for your switch family.

If you have configured a VRF for NX-API communication, however, ACLs will not restrict HTTP or HTTPS access. Instead, create a rule for an iptable. For more information about creating a rule, see [Updating an iptable, on page 245](#).

Updating an iptable

An iptable enables you to restrict HTTP or HTTPS access to a device when a VRF has been configured for NX-API communication. This section demonstrates how to add, verify, and remove rules for blocking HTTP and HTTPS access to an existing iptable.

Procedure

Step 1 To create a rule that blocks HTTP access:

```
bash-4.3# ip netns exec management iptables -A INPUT -p tcp --dport 80 -j DROP
```

Note The **management** mentioned in this step is the VRF name. It can be **management** | **default** | **custom vrf name**.

Step 2 To create a rule that blocks HTTPS access:

```
bash-4.3# ip netns exec management iptables -A INPUT -p tcp --dport 443 -j DROP
```

Step 3 To verify the applied rules:

```
bash-4.3# ip netns exec management iptables -L
```

```
Chain INPUT (policy ACCEPT)
target     prot opt source                destination            tcp dpt:http
DROP      tcp  --  anywhere              anywhere               tcp dpt:https
DROP      tcp  --  anywhere              anywhere               tcp dpt:https

Chain FORWARD (policy ACCEPT)
target     prot opt source                destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
```

Step 4 To create and verify a rule that blocks all traffic with a 10.155.0.0/24 subnet to port 80:

```
bash-4.3# ip netns exec management iptables -A INPUT -s 10.155.0.0/24 -p tcp --dport 80 -j DROP
bash-4.3# ip netns exec management iptables -L
```

```
Chain INPUT (policy ACCEPT)
target     prot opt source                destination            tcp dpt:http
DROP      tcp  --  10.155.0.0/24        anywhere               tcp dpt:http

Chain FORWARD (policy ACCEPT)
target     prot opt source                destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
```

Step 5 To remove and verify previously applied rules:

This example removes the first rule from INPUT.

```
bash-4.3# ip netns exec management iptables -D INPUT 1
bash-4.3# ip netns exec management iptables -L
```

```
Chain INPUT (policy ACCEPT)
target     prot opt source                destination

Chain FORWARD (policy ACCEPT)
target     prot opt source                destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
```

What to do next

The rules in iptables are not persistent across reloads when they are modified in a bash-shell. To make the rules persistent, see [Making an Iptable Persistent Across Reloads, on page 247](#).

Making an Iptable Persistent Across Reloads

The rules in iptables are not persistent across reloads when they are modified in a bash-shell. This section explains how to make a modified iptable persistent across a reload.

Before you begin

You have modified an iptable.

Procedure

Step 1 Create a file called `iptables_init.log` in the `/etc` directory with full permissions:

```
bash-4.3# touch /etc/iptables_init.log; chmod 777 /etc/iptables_init.log
```

Step 2 Create the `/etc/sys/iptables` file where your iptables changes will be saved:

```
bash-4.3# ip netns exec management iptables-save > /etc/sysconfig/iptables
```

Step 3 Create a startup script called `iptables_init` in the `/etc/init.d` directory with the following set of commands:

```
#!/bin/sh

### BEGIN INIT INFO

# Provides:          iptables_init
# Required-Start:
# Required-Stop:
# Default-Start:    2 3 4 5
# Default-Stop:
# Short-Description: init for iptables
# Description:      sets config for iptables
#                   during boot time

### END INIT INFO

PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin
start_script() {
    ip netns exec management iptables-restore < /etc/sysconfig/iptables
    ip netns exec management iptables
    echo "iptables init script executed" > /etc/iptables_init.log
}
case "$1" in
    start)
```

```

    start_script
    ;;
stop)
    ;;
restart)
    sleep 1
    $0 start
    ;;
*)
    echo "Usage: $0 {start|stop|status|restart}"
    exit 1
esac
exit 0

```

Step 4 Set the appropriate permissions to the startup script:

```
bash-4.3# chmod 777 /etc/init.d/iptables_init
```

Step 5 Set the iptables_init startup script to on with the chkconfig utility:

```
bash-4.3# chkconfig iptables_init on
```

The iptables_init startup script will now execute each time that you perform a reload, making the iptable rules persistent.

Kernel Stack ACL

The Kernel Stack ACL is a common CLI infrastructure to configure ACLs for management of inband and outband components.

The Kernel Stack ACL uses NX-OS ACL CLI to secure management applications on management and front panel ports. Configuring a single ACL must be able to secure all management applications on NX-OS.

Kernel Stack ACL is the component that fixes the manual intervention of the user and automatically programs iptable entries when the ACL is applied to mgmt0 interface.

The following is an example for configuring Kernel Stack ACL:

```

switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ip access-list kacl1
switch(config-acl)# statistics per-entry
switch(config-acl)# 10 deny tcp any any eq 443
switch(config-acl)# 20 permit ip any any
switch(config-acl)# end
switch#

switch(config-if)# interface mgmt0
switch(config-if)# ip access-group acl1 in
switch(config-if)# ipv6 traffic-filter acl6 in
switch(config-if)#

switch# sh ip access-lists kacl1
IP access list kacl1
statistics per-entry
10 deny tcp any any eq 443 [match=136]
20 permit ip any any [match=44952]
switch(config)#

```

The following is the Kernel Stack filtering for iptables entries based on the configuration:

```

bash-4.4# ip netns exec management iptables -L -n -v --line-numbers
Chain INPUT (policy ACCEPT 0 packets, 0 bytes)
num pkts bytes target prot opt in out source destination
1 9 576 DROP tcp -- * * 0.0.0.0/0 0.0.0.0/0 tcp dpt:443
2 0 0 ACCEPT all -- * * 0.0.0.0/0 0.0.0.0/0
3 0 0 DROP all -- * * 0.0.0.0/0 0.0.0.0/0

Chain FORWARD (policy ACCEPT 0 packets, 0 bytes)
num pkts bytes target prot opt in out source destination

Chain OUTPUT (policy ACCEPT 0 packets, 0 bytes)
num pkts bytes target prot opt in out source destination
bash-4.4#

```

The following are the limitations for the Kernel Stack ACL support:

- This feature is supported only on mgmt0 interface and not on other inband interfaces.
- Five tuples (protocol, source-ip, destination-ip, source-port, and destination-port) of the ACL entry are programmed in the iptables. Rest of the options provided in the ACL entry are not programmed in the iptables and throws a warning syslog in such instances.

For example, "WARNING: Some ACL options are not supported in kstack. Only partial rule will be installed".

- If the device user has host bash access, then the user can manually update the iptables. This update could potentially corrupt the iptable rules for which they are programmed.
- The verified maximum number of ACEs is 100 for IPv4 traffic and an additional 100 for IPv6 traffic. Throughput may be impacted if more than this scale is applied.

Table of NX-API Response Codes

The following are the possible NX-API errors, error codes, and messages of an NX-API response.

The following are the possible NX-API errors, error codes, and messages of an NX-API response.

When the request format is in XML or JSON format, the following are the possible NX-API errors, error codes, and messages of an NX-API response.



Note The standard HTTP error codes are at the Hypertext Transfer Protocol (HTTP) Status Code Registry (<http://www.iana.org/assignments/http-status-codes/http-status-codes.xhtml>).

Table 12: NX-API Response Codes

| NX-API Response | Code | Message |
|-------------------------|------|---|
| SUCCESS | 200 | Success. |
| CUST_OUTPUT_PIPED | 204 | Output is piped elsewhere due to request. |
| BASH_CMD_ERR | 400 | Bash command error. |
| CHUNK_ALLOW_ONE_CMD_ERR | 400 | Chunking honors only one command. |

| | | |
|-----------------------------------|-----|---|
| CLI_CLIENT_ERR | 400 | CLI execution error. |
| CLI_CMD_ERR | 400 | Input CLI command error. |
| EOC_NOT_ALLOWED_ERR | 400 | The eoc value is not allowed as session Id in the request. |
| IN_MSG_ERR | 400 | Incoming message is invalid. |
| INVALID_REMOTE_IP_ERR | 400 | Unable to retrieve remote ip of request. |
| MSG_VER_MISMATCH | 400 | Message version mismatch. |
| NO_INPUT_CMD_ERR | 400 | No input command. |
| SID_NOT_ALLOWED_ERR | 400 | Invalid character that is entered as a session ID. |
| PERM_DENY_ERR | 401 | Permission denied. |
| CONF_NOT_ALLOW_SHOW_ERR | 405 | Configuration mode does not allow show . |
| SHOW_NOT_ALLOW_CONF_ERR | 405 | Show mode does not allow configuration. |
| EXCEED_MAX_SHOW_ERR | 413 | Maximum number of consecutive show commands exceeded. The maximum is 10. |
| MSG_SIZE_LARGE_ERR | 413 | Response size too large. |
| RESP_SIZE_LARGE_ERR | 413 | Response size stopped processing because it exceeded the maximum message size. The maximum is 200 MB. |
| EXCEED_MAX_INFLIGHT_CHUNK_REQ_ERR | 429 | Maximum number of concurrent chunk requests is exceeded. The maximum is 2. |
| MAX_SESSIONS_ERR | 429 | Max sessions reached. If you are a new user/client, please try again later. |
| OBJ_NOT_EXIST | 432 | Requested object does not exist. |
| BACKEND_ERR | 500 | Backend processing error. |
| CREATE_CHECKPOINT_ERR | 500 | Error creating a checkpoint. |
| DELETE_CHECKPOINT_ERR | 500 | Error deleting a checkpoint. |
| FILE_OPER_ERR | 500 | System internal file operation error. |
| LIBXML_NS_ERR | 500 | System internal LIBXML NS error. This is a request format error. |
| LIBXML_PARSE_ERR | 500 | System internal LIBXML parse error. This is a request format error. |
| LIBXML_PATH_CTX_ERR | 500 | System internal LIBXML path context error. This is a request format error. |

| | | |
|-----------------------------------|-----|---|
| MEM_ALLOC_ERR | 500 | System internal memory allocation error. |
| ROLLBACK_ERR | 500 | Error executing a rollback. |
| SERVER_BUSY_ERR | 500 | Request is rejected because the server is busy. |
| USER_NOT_FOUND_ERR | 500 | User not found from input or cache. |
| VOLATILE_FULL | 500 | Volatile memory is full. Free up memory space and retry. |
| XML_TO_JSON_CONVERT_ERR | 500 | XML to JSON conversion error. |
| BASH_CMD_NOT_SUPPORTED_ERR | 501 | Bash command not supported. |
| CHUNK_ALLOW_XML_ONLY_ERR | 501 | Chunking allows only XML output. |
| CHUNK_ONLY_ALLOWED_IN_SHOW_ERR | 501 | Response chunking allowed only in <code>show</code> commands. |
| CHUNK_TIMEOUT | 501 | Timeout while generating chunk response. |
| CLI_CMD_NOT_SUPPORTED_ERR | 501 | CLI command not supported. |
| JSON_NOT_SUPPORTED_ERR | 501 | JSON not supported due to a potential large amount of output. |
| MALFORMED_XML | 501 | Malformed XML output. |
| MSG_TYPE_UNSUPPORTED_ERR | 501 | Message type not supported. |
| OUTPUT_REDIRECT_NOT_SUPPORTED_ERR | 501 | Output redirection is not supported. |
| PIPE_OUTPUT_NOT_SUPPORTED_ERR | 501 | Pipe operation is not supported. |
| PIPE_XML_NOT_ALLOWED_IN_INPUT | 501 | Pipe XML for this command is not allowed in input. |
| PIPE_NOT_ALLOWED_IN_INPUT | 501 | Pipe is not allowed for this input type. |
| RESP_BIG_JSON_NOT_ALLOWED_ERR | 501 | Response has large amount of output. JSON not supported. |
| RESP_BIG_USE_CHUNK_ERR | 501 | Response is greater than the allowed maximum. The maximum is 10 MB. Use XML or JSON output with chunking enabled. |
| STRUCT_NOT_SUPPORTED_ERR | 501 | Structured output unsupported. |
| ERR_UNDEFINED | 600 | Unknown error. |

JSON and XML Structured Output

The NX-OS supports redirecting the standard output of various **show** commands in the following structured output formats:

- XML
- JSON. The limit for JSON output is 60 MB.
- JSON Pretty, which makes the standard block of JSON-formatted output easier to read. The limit for JSON output is 60 MB.
- Introduced in NX-OS release 9.3(1), JSON Native and JSON Pretty Native displays JSON output faster and more efficiently by bypassing an extra layer of command interpretation. JSON Native and JSON Pretty Native preserve the data type in the output. They display integers as integers instead of converting them to a string for output.

Converting the standard NX-OS output to any of these formats occurs on the NX-OS CLI by "piping" the output to a JSON or XML interpreter. For example, you can issue the **show ip access** command with the logical pipe (|) and specify the output format. If you do, the NX-OS command output is properly structured and encoded in that format. This feature enables programmatic parsing of the data and supports streaming data from the switch through software streaming telemetry. Most commands in Cisco NX-OS support JSON, JSON Pretty, JSON Native, JSON Native Pretty, and XML output. Some, for example, consistency checker commands, do not support all formats. Consistency checker commands support XML, but not any variant of JSON.



Note To avoid validation error, use file redirection to redirect the JSON output to a file, and use the file output.

Example:

```
Switch#show version | json > json_output ; run bash cat /bootflash/json_output
```

Selected examples of this feature follow.

About JSON (JavaScript Object Notation)

JSON is a light-weight text-based open standard that is designed for human-readable data and is an alternative to XML. JSON was originally designed from JavaScript, but it is language-independent data format. JSON and JSON Pretty format, as well as JSON Native and JSON Pretty Native, are supported for command output.

The two primary Data Structures that are supported in some way by nearly all modern programming languages are as follows:

- Ordered List :: Array
- Unordered List (Name/Value pair) :: Objects

JSON or XML output for a **show** command can be accessed through the NX-API sandbox also.

CLI Execution

```
switch-1-vxlan-1# show cdp neighbors | json
{"TABLE_cdp_neighbor_brief_info": {"ROW_cdp_neighbor_brief_info": [{"ifindex": "83886080", "device_id": "SW-SWITCH-1", "intf_id": "mgmt0", "ttl": "148", "capability": ["switch", "IGMP_cnd_filtering"], "platform_id": "cisco AA-C0000 S-29-L", "port_id": "GigabitEthernet1/0/24"}, {"ifindex": "436207616", "device
```

```
_id": "SWITCH-1-VXLAN-1(FOC1234A01B)", "intf_id": "Ethernet1/1", "ttl": "166", "capability": ["router", "switch", "IGMP_cnd_filtering", "Supports-STP-Dispute"], "platform_id": "N3K-C3132Q-40G", "port_id": "Ethernet1/1"}}}
BLR-VXLAN-NPT-CR-179#
```

Examples of XML and JSON Output

This section documents selected examples of NX-OS commands that are displayed as XML and JSON output.

This example shows how to display the unicast and multicast routing entries in hardware tables in JSON format:

```
switch(config)# show hardware profile status | json
{"total_lpm": ["8191", "1024"], "total_host": "8192", "max_host4_limit": "4096",
 "max_host6_limit": "2048", "max_mcast_limit": "2048", "used_lpm_total": "9", "used_v4_lpm": "6", "used_v6_lpm": "3", "used_v6_lpm_128": "1", "used_host_lpm_total": "0", "used_host_v4_lpm": "0", "used_host_v6_lpm": "0", "used_mcast": "0", "used_mcast_oif1": "2", "used_host_in_host_total": "13", "used_host4_in_host": "12", "used_host6_in_host": "1", "max_ecmp_table_limit": "64", "used_ecmp_table": "0", "mfib_fd_status": "Disabled", "mfib_fd_maxroute": "0", "mfib_fd_count": "0"}
switch(config)#
```

This example shows how to display the unicast and multicast routing entries in hardware tables in XML format:

```
switch(config)# show hardware profile status | xml
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns="http://www.cisco.com/nxos:1.0:fib">
  <nf:data>
    <show>
      <hardware>
        <profile>
          <status>
            <_XML_OPT_Cmd_dynamic_tcam_status>
              <_XML_OPT_Cmd_dynamic_tcam_status__readonly__>
                <_readonly__>
                  <total_lpm>8191</total_lpm>
                  <total_host>8192</total_host>
                  <total_lpm>1024</total_lpm>
                  <max_host4_limit>4096</max_host4_limit>
                  <max_host6_limit>2048</max_host6_limit>
                  <max_mcast_limit>2048</max_mcast_limit>
                  <used_lpm_total>9</used_lpm_total>
                  <used_v4_lpm>6</used_v4_lpm>
                  <used_v6_lpm>3</used_v6_lpm>
                  <used_v6_lpm_128>1</used_v6_lpm_128>
                  <used_host_lpm_total>0</used_host_lpm_total>
                  <used_host_v4_lpm>0</used_host_v4_lpm>
                  <used_host_v6_lpm>0</used_host_v6_lpm>
                  <used_mcast>0</used_mcast>
                  <used_mcast_oif1>2</used_mcast_oif1>
                  <used_host_in_host_total>13</used_host_in_host_total>
                  <used_host4_in_host>12</used_host4_in_host>
                  <used_host6_in_host>1</used_host6_in_host>
                  <max_ecmp_table_limit>64</max_ecmp_table_limit>
                  <used_ecmp_table>0</used_ecmp_table>
                  <mfib_fd_status>Disabled</mfib_fd_status>
                  <mfib_fd_maxroute>0</mfib_fd_maxroute>
                  <mfib_fd_count>0</mfib_fd_count>
```

```

        </__readonly__>
    </__XML__OPT_Cmd_dynamic_tcam_status__readonly__>
</__XML__OPT_Cmd_dynamic_tcam_status>
</status>
</profile>
</hardware>
</show>
</nf:data>
</nf:rpc-reply>
]]>]]>
switch(config)#

```

This example shows how to display LLDP timers that are configured on the switch in JSON format:

```

switch(config)# show lldp timers | json
{"ttl": "120", "reinit": "2", "tx_interval": "30", "tx_delay": "2", "hold_mplier": "4", "notification_interval": "5"}
switch(config)#

```

This example shows how to display LLDP timers that are configured on the switch in XML format:

```

switch(config)# show lldp timers | xml
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns="http://www.cisco.com/nxos:1.0:lldp">
  <nf:data>
    <show>
      <lldp>
        <timers>
          <__XML__OPT_Cmd_lldp_show_timers__readonly__>
            <__readonly__>
              <ttl>120</ttl>
              <reinit>2</reinit>
              <tx_interval>30</tx_interval>
              <tx_delay>2</tx_delay>
              <hold_mplier>4</hold_mplier>
              <notification_interval>5</notification_interval>
            </__readonly__>
          </__XML__OPT_Cmd_lldp_show_timers__readonly__>
        </timers>
      </lldp>
    </show>
  </nf:data>
</nf:rpc-reply>
]]>]]>
switch(config)#

```

This example shows how to display ACL statistics in XML format.

```

switch-1(config-acl)# show ip access-lists acl-test1 | xml
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns="http://www.cisco.com/nxos:1.0:aclmgr" xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0">
  <nf:data>
    <show>
      <__XML__OPT_Cmd_show_acl_ip_ipv6_mac>
        <ip_ipv6_mac>ip</ip_ipv6_mac>
        <access-lists>
          <__XML__OPT_Cmd_show_acl_name>
            <name>acl-test1</name>
          <__XML__OPT_Cmd_show_acl_capture>
            <__XML__OPT_Cmd_show_acl_expanded>
              <__XML__OPT_Cmd_show_acl__readonly__>

```

```

    <__readonly__>
    <TABLE_ip_ipv6_mac>
    <ROW_ip_ipv6_mac>
    <op_ip_ipv6_mac>ip</op_ip_ipv6_mac>
    <show_summary>0</show_summary>
    <acl_name>acl-test1</acl_name>
    <statistics>enable</statistics>
    <frag_opt_permit_deny>permit-all</frag_opt_permit_deny>
    <TABLE_seqno>
    <ROW_seqno>
    <seqno>10</seqno>
    <permitdeny>permit</permitdeny>
    <ip>ip</ip>
    <src_ip_prefix>192.0.2.1/24</src_ip_prefix>
    <dest_any>any</dest_any>
    </ROW_seqno>
    </TABLE_seqno>
    </ROW_ip_ipv6_mac>
    </TABLE_ip_ipv6_mac>
    </__readonly__>
  </__XML_OPT_Cmd_show_acl__readonly__>
</__XML_OPT_Cmd_show_acl_expanded>
</__XML_OPT_Cmd_show_acl_capture>
</__XML_OPT_Cmd_show_acl_name>
</access-lists>
</__XML_OPT_Cmd_show_acl_ip_ipv6_mac>
</show>
</nf:data>
</nf:rpc-reply>
]]>]]>
switch-1(config-acl)#

```

This example shows how to display ACL statistics in JSON format.

```

switch-1(config-acl)# show ip access-lists acl-test1 | json
{"TABLE_ip_ipv6_mac": {"ROW_ip_ipv6_mac": {"op_ip_ipv6_mac": "ip", "show_summary": "0", "acl_name": "acl-test1", "statistics": "enable", "frag_opt_permit_deny": "permit-all", "TABLE_seqno": {"ROW_seqno": {"seqno": "10", "permitdeny": "permit", "ip": "ip", "src_ip_prefix": "192.0.2.1/24", "dest_any": "any"}}}}}
switch-1(config-acl)#

```

The following example shows how to display the switch's redundancy status in JSON format.

```

switch-1# show system redundancy status | json
{"rdn_mode_admin": "HA", "rdn_mode_oper": "None", "this_sup": "(sup-1)", "this_sup_rdn_state": "Active, SC not present", "this_sup_sup_state": "Active", "this_sup_internal_state": "Active with no standby", "other_sup": "(sup-1)", "other_sup_rdn_state": "Not present"}
nxosv2#
switch-1#

```

This example shows how to display the switch's redundancy information in JSON Pretty Native format.

```

switch-1# show system redundancy status | json-pretty native
{
  "rdn_mode_admin": "HA",
  "rdn_mode_oper": "None",
  "this_sup": "(sup-1)",
  "this_sup_rdn_state": "Active, SC not present",
  "this_sup_sup_state": "Active",
  "this_sup_internal_state": "Active with no standby",
  "other_sup": "(sup-1)",
  "other_sup_rdn_state": "Not present"
}
switch-1#

```

The following example shows how to display the switch's OSPF routing parameters in JSON Native format.

```
switch-1# show ip ospf | json native
{"TABLE_ctx":{"ROW_ctx":[{"ptag":"Blah","instance_number":4,"cname":"default","rid":"0.0.0.0","stateful_ha":"true","gr_ha":"true","gr_planned_only":"true","gr_grace_period":"PT60S","gr_state":"inactive","gr_last_status":"None","support_tos0_only":"true","support_opaque_lsa":"true","is_abr":"false","is_asbr":"false","admin_dist":110,"ref_bw":40000,"spf_start_time":"PT0S","spf_hold_time":"PT1S","spf_max_time":"PT5S","lsa_start_time":"PT0S","lsa_hold_time":"PT5S","lsa_max_time":"PT5S","min_lsa_arr_time":"PT1S","lsa_aging_pace":10,"spf_max_paths":8,"max_metric_adver":"false","asext_lsa_cnt":0,"asext_lsa_crc":"0","asopaque_lsa_cnt":0,"asopaque_lsa_crc":"0","area_total":0,"area_normal":0,"area_stub":0,"area_nssa":0,"act_area_total":0,"act_area_normal":0,"act_area_stub":0,"act_area_nssa":0,"no_discard_rt_ext":"false","no_discard_rt_int":"false"},{"ptag":"100","instance_number":3,"cname":"default","rid":"0.0.0.0","stateful_ha":"true","gr_ha":"true","gr_planned_only":"true","gr_grace_period":"PT60S","gr_state":"inactive","gr_last_status":"None","support_tos0_only":"true","support_opaque_lsa":"true","is_abr":"false","is_asbr":"false","admin_dist":110,"ref_bw":40000,"spf_start_time":"PT0S","spf_hold_time":"PT1S","spf_max_time":"PT5S","lsa_start_time":"PT0S","lsa_hold_time":"PT5S","lsa_max_time":"PT5S","min_lsa_arr_time":"PT1S","lsa_aging_pace":10,"spf_max_paths":8,"max_metric_adver":"false","asext_lsa_cnt":0,"asext_lsa_crc":"0","asopaque_lsa_cnt":0,"asopaque_lsa_crc":"0","area_total":0,"area_normal":0,"area_stub":0,"area_nssa":0,"act_area_total":0,"act_area_normal":0,"act_area_stub":0,"act_area_nssa":0,"no_discard_rt_ext":"false","no_discard_rt_int":"false"},{"ptag":"111","instance_number":1,"cname":"default","rid":"0.0.0.0","stateful_ha":"true","gr_ha":"true","gr_planned_only":"true","gr_grace_period":"PT60S","gr_state":"inactive","gr_last_status":"None","support_tos0_only":"true","support_opaque_lsa":"true","is_abr":"false","is_asbr":"false","admin_dist":110,"ref_bw":40000,"spf_start_time":"PT0S","spf_hold_time":"PT1S","spf_max_time":"PT5S","lsa_start_time":"PT0S","lsa_hold_time":"PT5S","lsa_max_time":"PT5S","min_lsa_arr_time":"PT1S","lsa_aging_pace":10,"spf_max_paths":8,"max_metric_adver":"false","asext_lsa_cnt":0,"asext_lsa_crc":"0","asopaque_lsa_cnt":0,"asopaque_lsa_crc":"0","area_total":0,"area_normal":0,"area_stub":0,"area_nssa":0,"act_area_total":0,"act_area_normal":0,"act_area_stub":0,"act_area_nssa":0,"no_discard_rt_ext":"false","no_discard_rt_int":"false"},{"ptag":"112","instance_number":2,"cname":"default","rid":"0.0.0.0","stateful_ha":"true","gr_ha":"true","gr_planned_only":"true","gr_grace_period":"PT60S","gr_state":"inactive","gr_last_status":"None","support_tos0_only":"true","support_opaque_lsa":"true","is_abr":"false","is_asbr":"false","admin_dist":110,"ref_bw":40000,"spf_start_time":"PT0S","spf_hold_time":"PT1S","spf_max_time":"PT5S","lsa_start_time":"PT0S","lsa_hold_time":"PT5S","lsa_max_time":"PT5S","min_lsa_arr_time":"PT1S","lsa_aging_pace":10,"spf_max_paths":8,"max_metric_adver":"false","asext_lsa_cnt":0,"asext_lsa_crc":"0","asopaque_lsa_cnt":0,"asopaque_lsa_crc":"0","area_total":0,"area_normal":0,"area_stub":0,"area_nssa":0,"act_area_total":0,"act_area_normal":0,"act_area_stub":0,"act_area_nssa":0,"no_discard_rt_ext":"false","no_discard_rt_int":"false"}]}}
```

The following example shows how to display OSPF routing parameters in JSON Pretty Native format.

```
switch-1# show ip ospf | json-pretty native
{
  "TABLE_ctx": {
    "ROW_ctx": [{
      "ptag": "Blah",
      "instance_number": 4,
      "cname": "default",
      "rid": "0.0.0.0",
      "stateful_ha": "true",
      "gr_ha": "true",
      "gr_planned_only": "true",
      "gr_grace_period": "PT60S",
      "gr_state": "inactive",
      "gr_last_status": "None",
      "support_tos0_only": "true",
```

```

    "support_opaque_lsa": "true",
    "is_abr": "false",
    "is_asbr": "false",
    "admin_dist": 110,
    "ref_bw": 40000,
    "spf_start_time": "PT0S",
    "spf_hold_time": "PT1S",
    "spf_max_time": "PT5S",
    "lsa_start_time": "PT0S",
    "lsa_hold_time": "PT5S",
    "lsa_max_time": "PT5S",
    "min_lsa_arr_time": "PT1S",
    "lsa_aging_pace": 10,
    "spf_max_paths": 8,
    "max_metric_adver": "false",
    "asext_lsa_cnt": 0,
    "asext_lsa_crc": "0",
    "asopaque_lsa_cnt": 0,
    "asopaque_lsa_crc": "0",
    "area_total": 0,
    "area_normal": 0,
    "area_stub": 0,
    "area_nssa": 0,
    "act_area_total": 0,
    "act_area_normal": 0,
    "act_area_stub": 0,
    "act_area_nssa": 0,
    "no_discard_rt_ext": "false",
    "no_discard_rt_int": "false"
  }, {
    "ptag": "100",
    "instance_number": 3,
    "cname": "default",
    "rid": "0.0.0.0",
    "stateful_ha": "true",
    "gr_ha": "true",
    "gr_planned_only": "true",
    "gr_grace_period": "PT60S",
    "gr_state": "inactive",

    ... content deleted for brevity ...

    "max_metric_adver": "false",
    "asext_lsa_cnt": 0,
    "asext_lsa_crc": "0",
    "asopaque_lsa_cnt": 0,
    "asopaque_lsa_crc": "0",
    "area_total": 0,
    "area_normal": 0,
    "area_stub": 0,
    "area_nssa": 0,
    "act_area_total": 0,
    "act_area_normal": 0,
    "act_area_stub": 0,
    "act_area_nssa": 0,
    "no_discard_rt_ext": "false",
    "no_discard_rt_int": "false"
  }
}
switch-1#

```

The following example shows how to display the IP route summary in XML format.

```

switch-1# show ip route summary | xml
<?xml version="1.0" encoding="ISO-8859-1"?> <nf:rpc-reply
xmlns="http://www.cisco.com/nxos:1.0:urib" xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0">

<nf:data>
<show>
<ip>
<route>
  <__XML__OPT_Cmd_urib_show_ip_route_command_ip>
  <__XML__OPT_Cmd_urib_show_ip_route_command_unicast>
  <__XML__OPT_Cmd_urib_show_ip_route_command_topology>
  <__XML__OPT_Cmd_urib_show_ip_route_command_l3vm-info>
  <__XML__OPT_Cmd_urib_show_ip_route_command_rpf>
  <__XML__OPT_Cmd_urib_show_ip_route_command_ip-addr>
  <__XML__OPT_Cmd_urib_show_ip_route_command_protocol>
  <__XML__OPT_Cmd_urib_show_ip_route_command_summary>
  <__XML__OPT_Cmd_urib_show_ip_route_command_vrf>
  <__XML__OPT_Cmd_urib_show_ip_route_command__readonly__>
  <__readonly__>
  <TABLE_vrf>
  <ROW_vrf>
  <vrf-name-out>default</vrf-name-out>
  <TABLE_addrf>
  <ROW_addrf>
  <addrf>ipv4</addrf>
  <TABLE_summary>
  <ROW_summary>
  <routes>938</routes>
  <paths>1453</paths>
  <TABLE_unicast>
  <ROW_unicast>
  <clientnameuni>am</clientnameuni>
  <best-paths>2</best-paths>
  </ROW_unicast>
  <ROW_unicast>
  <clientnameuni>local</clientnameuni>
  <best-paths>105</best-paths>
  </ROW_unicast>
  <ROW_unicast>
  <clientnameuni>direct</clientnameuni>
  <best-paths>105</best-paths>
  </ROW_unicast>
  <ROW_unicast>
  <clientnameuni>broadcast</clientnameuni>
  <best-paths>203</best-paths>
  </ROW_unicast>
  <ROW_unicast>
  <clientnameuni>ospf-10</clientnameuni>
  <best-paths>1038</best-paths>
  </ROW_unicast>
  </TABLE_unicast>
  <TABLE_route_count>
  <ROW_route_count>
  <mask_len>8</mask_len>
  <count>1</count>
  </ROW_route_count>
  <ROW_route_count>
  <mask_len>24</mask_len>
  <count>600</count>
  </ROW_route_count>
  <ROW_route_count>
  <mask_len>31</mask_len>
  <count>13</count>
  </ROW_route_count>

```



```

        <ROW_route_count>
        <mask_len>32</mask_len>
        <count>324</count>
    </ROW_route_count>
</TABLE_route_count>
</ROW_summary>
</TABLE_summary>
</ROW_addrf>
</TABLE_addrf>
</ROW_vrf>
</TABLE_vrf>
</__readonly__>
</__XML_OPT_Cmd_urib_show_ip_route_command__readonly__>
</__XML_OPT_Cmd_urib_show_ip_route_command_vrf>
</__XML_OPT_Cmd_urib_show_ip_route_command_summary>
</__XML_OPT_Cmd_urib_show_ip_route_command_protocol>
</__XML_OPT_Cmd_urib_show_ip_route_command_ip-addr>
</__XML_OPT_Cmd_urib_show_ip_route_command_rpf>
</__XML_OPT_Cmd_urib_show_ip_route_command_l3vm-info>
</__XML_OPT_Cmd_urib_show_ip_route_command_topology>
</__XML_OPT_Cmd_urib_show_ip_route_command_unicast>
</__XML_OPT_Cmd_urib_show_ip_route_command_ip>
</route>
</ip>
</show>
</nf:data>
</nf:rpc-reply>
]]>]]>
switch-1#

```

The following example shows how to display the IP route summary in JSON format.

```

switch-1# show ip route summary | json
{"TABLE_vrf": {"ROW_vrf": {"vrf-name-out": "default", "TABLE_addrf": {"ROW_addrf": {"addrf": "ipv4", "TABLE_summary": {"ROW_summary": {"routes": "938", "paths": "1453", "TABLE_unicast": {"ROW_unicast": [{"clientnameuni": "am", "best-paths": "2"}, {"clientnameuni": "local", "best-paths": "105"}, {"clientnameuni": "direct", "best-paths": "105"}, {"clientnameuni": "broadcast", "best-paths": "203"}, {"clientnameuni": "ospf-10", "best-paths": "1038"}]}}, "TABLE_route_count": {"ROW_route_count": [{"mask_len": "8", "count": "1"}, {"mask_len": "24", "count": "600"}, {"mask_len": "31", "count": "13"}, {"mask_len": "32", "count": "324"}]}}}}}}}
switch-1#

```

The following example shows how to display the IP route summary in JSON Pretty format.

```

switch-1# show ip route summary | json-pretty
{
  "TABLE_vrf": {
    "ROW_vrf": {
      "vrf-name-out": "default",
      "TABLE_addrf": {
        "ROW_addrf": {
          "addrf": "ipv4",
          "TABLE_summary": {
            "ROW_summary": {
              "routes": "938",
              "paths": "1453",
              "TABLE_unicast": {
                "ROW_unicast": [
                  {
                    "clientnameuni": "am",
                    "best-paths": "2"
                  },
                  {
                    "clientnameuni": "local",

```



```
      "addrf": "ipv4",
      "TABLE_summary": {
        "ROW_summary": [{
          "routes": 3,
          "paths": 3,
          "TABLE_unicast": {
            "ROW_unicast": [{
              "clientnameuni": "broadcast",
              "best-paths": 3
            }]
          },
          "TABLE_route_count": {
            "ROW_route_count": [{
              "mask_len": 8,
              "count": 1
            }, {
              "mask_len": 32,
              "count": 2
            }]
          }
        }]
      }
    }
  }
}
switch-1(config)#
```

Sample NX-API Scripts

You can access sample scripts that demonstrate how to use a script with NX-API. To access a sample script, click the following link then choose the directory that corresponds to the required software release: [Cisco Nexus 9000 NX-OS NX-API](#)



CHAPTER 23

NX-API REST

This chapter contains the following topics:

- [About NX-API REST, on page 263](#)
- [DME Config Replace Through REST, on page 264](#)

About NX-API REST

NX-API REST

In Release 7.0(3)I2(1), the NX-API REST SDK has been added.

On Cisco Nexus switches, configuration is performed using command-line interfaces (CLIs) that run only on the switch. NX-API REST improves the accessibility of the Cisco Nexus configuration by providing HTTP/HTTPS APIs that:

- Make specific CLIs available outside of the switch.
- Enable configurations that would require issuing many CLI commands by combining configuration actions in relatively few HTTP/HTTPS operations.

NX-API REST supports **show** commands, basic and advanced switch configurations, and Linux Bash.

NX-API REST uses HTTP/HTTPS as its transport. CLIs are encoded into the HTTP/HTTPS POST body. The NX-API REST backend uses the Nginx HTTP server. The Nginx process, and all of its children processes, are under Linux cgroup protection where the CPU and memory usage is capped. The NX-API processes are part of the cgroup `ext_ser_nginx`, which is limited to 2,147,483,648 bytes of memory. If the Nginx memory usage exceeds the cgroup limitations, the Nginx process is restarted and the NX-API configuration (the VRF, port, and certificate configurations) is restored.

For more information about the Cisco Nexus 3000 and 9000 Series NX-API REST SDK, see <https://developer.cisco.com/docs/nx-os-n3k-n9k-api-ref/>.

DME Config Replace Through REST

About DME Full Config Replace Through REST Put

Beginning with Cisco NX-OS Release 9.3(1), Cisco NX-OS supports model-based full config replace through REST PUT operations. This method of replacing configurations uses the Cisco DME model.

The DME Full Config replace feature enables you to use the REST programmatic interface to replace the switch running configuration. The feature provides the following extra benefits: DME full config replace occurs through a PUT operation. All parts of the config tree (system-level, subtree, and leaf) support DME full config replace.

- Supports non-disruptive replacement of the switch configuration
- Supports automation
- Offers the ability to selectively modify features without affecting other features or their configs.
- Simplifies config changes and eliminates human error by enabling you to specify the final config outcome. The switch calculates the differences and pushes them to the affected parts of config tree.



Note Although not accomplished through a programmatic interface, you can also achieve a full config replace by using the **config replace config-file-name** Cisco NX-OS CLI command.

Guidelines and Limitations

The following are the guidelines and limitations for the DME full config replace feature:

- For information about supported platforms, see the [Nexus Switch Platform Matrix](#).
- For information about supported platforms for Cisco NX-OS prior to release 9.3(x), see [Platform Support for Programmability Features, on page 5](#). Starting with Cisco NX-OS release 9.3(x), for information about supported platforms, see the [Nexus Switch Platform Matrix](#).
- DME is supported on N9K-92348GC-X.
- It is important for you to know the tree and know where you are applying the config replace. If you are using the Sandbox for the config replace operation, the Sandbox defaults to the subtree, so you might need to change the URI to target the correct node in the config tree.
- If you use the NX-OS Sandbox to Convert (for Replace), you must use the POST operation because of the presence of the `status: 'replaced'` attribute in the request. If you are using any other conversion option, you can use the PUT operation.
- If you use the REST PUT option for this feature on a subtree node, config replace operation is applied to the entire subtree. The target subtree node is correctly changed with the config replace data in the PUT, but be aware that leaf nodes of the subtree node are also affected by being set to default values.

If you do not want the leaf nodes to be affected, do not use a PUT operation. Instead, you can use a POST operation with the `status: 'replaced'` attribute.

If you are applying the config replace to a leaf node, the PUT operation operates predictably.

Replacing Property-Level Configuration Through REST POST

Cisco's DME model supports property-level config replace for CLI-based features through a REST POST operation. You can replace the config for the property of a feature through the NX-OS Sandbox by generating a request payload and sending it to the switch through a REST POST operation. For information about the NX-OS Sandbox, see [NX-API Developer Sandbox](#).

Procedure

- Step 1** Connect to the switch through NX-OS Sandbox through HTTPS and provide your login credentials.
- Step 2** In the work area, enter the CLI for the feature that you want to change.
- Step 3** In the field below the work area, set the URI to the MO in the tree for the feature that you want to configure. This MO level is where you will send the Put request.
- Step 4** For Method, select `NX-API (DME)`.
- Step 5** For Input Type, select `CLI`.
- Step 6** From the Convert drop-down list, select `Convert (for replace)` to generate the payload in the Request pane.
- Step 7** Click the request using a **POST** operation to the switch..

Note Property-level config replace can fail if the config is a default config because the replace operation tries to delete all the children MOs and reset all properties to default.

Replacing Feature-Level Config Through REST PUT

Cisco DME supports replacing feature-level configurations through REST PUT operations. You can replace the configuration for specific features by sending a PUT at the feature level of the model.

Use the following procedure:

Procedure

- Step 1** From the client, issue a REST PUT operation at the model object (MO) of the feature:
 - a) The Put must specify the URL from the top System level to the MO of the feature.
For example, for a BGP `/api/mo/sys/bgp.json`

The payload must be a valid config, and the config must be retrievable from the switch at any time by issuing a GET on the DN of the feature. For example, for BGP,
`/api/mo/sys/bgp.json?rsp-subtree=full&rsp-prop-include=set-config-only`.
 - b) The payload for the feature should start with the MO that you want to replace (for example, `bgp`).
For example:

```

{
  "bgpInst": {
    "attributes": {
      "asn": "100",
      "rn": "inst"
    },
    "children": [
      ... content removed for brevity ...
      {
        "bgpDom": {
          "attributes": {
            "name": "vrf1",
            "rn": "dom-vrf1"
          },
          "children": [
            {
              "bgpPeer": {
                "attributes": {
                  "addr": "10.1.1.1",
                  "inheritContPeerCtrl": "",
                  "rn": "peer-[10.1.1.1]"
                }
              }
            }
          ]
        },
        {
          "bgpDom": {
            "attributes": {
              "name": "default",
              "rn": "dom-default",
              "rtrId": "1.1.1.1"
            }
          }
        }
      ]
    ]
  }
}

```

Step 2 Send a GET on the DN you used for the config replace by using
 /api/mo/sys/bgp.json?rsp-subtree=full&rsp-prop-include=set-config-only.

Step 3 (Optional) Compare the payload that you sent with the GET on the DN you replaced. The payload of the GET should be the same as the payload you sent.

Troubleshooting Config Replace for REST PUT

The following are steps to help troubleshoot if config replace through a REST Put operation is not successful.

Procedure

Step 1 Check if the request is valid.

The URL, operation, and payload should be valid. For example, if the URL is `api/mo/sys/foo.json` then the payload should start with `foo`

Step 2 Make sure the payload is valid and contains only the config properties which are:

- Successfully set
- Taken from a valid device config

To get only the config properties, use a GET that filters for `rsp-subtree=full&rsp-prop-include=set-config-only`

Step 3 To validate the payload, send it to the switch using a DME POST operation.

Step 4 Check the error to verify that it has the name of the MO and property.

Step 5 Validate the payload also has the name of the MO and property.



CHAPTER 24

NX-API Developer Sandbox

- NX-API Developer Sandbox: NX-OS Releases Prior to 9.2(2), on page 269
- NX-API Developer Sandbox: NX-OS Release 9.2(2) and Later, on page 281


NX-API Developer Sandbox: NX-OS Releases Prior to 9.2(2)

About the NX-API Developer Sandbox

The NX-API Developer Sandbox is a web form hosted on the switch. It translates NX-OS CLI commands into equivalent XML or JSON payloads, and converts NX-API REST payloads into their CLI equivalents.

The web form is a single screen with three panes — Command (top pane), Request, and Response — as shown in the figure.

Figure 1: NX-API Developer Sandbox with Example Request and Output Response



Controls in the Command pane allow you to choose a message format for a supported API, such as NX-API REST, and a command type, such as XML or JSON. The available command type options vary depending on the selected message format.

When you type or paste one or more CLI commands into the Command pane, the web form converts the commands into an API payload, checking for configuration errors, and displays the resulting payload in the Request pane. If you then choose to post the payload directly from the Sandbox to the switch, using the POST button in the Command pane, the Response pane displays the API response.

Conversely, when you type an NX-API REST designated name (DN) and payload into the Command pane and select the **nx-api rest** Message format and the **model** Command type, Developer Sandbox checks the payload for configuration errors, then the Response pane displays the equivalent CLIs.

Guidelines and Limitations

Following are the guidelines and limitations for the Developer Sandbox:

- Clicking **Send** in the Sandbox commits the command to the switch, which can result in a configuration or state change.
- Some feature configuration commands are not available until their associated feature has been enabled. For example, configuring a BGP router requires first enabling BGP with the **feature bgp** command. Similarly, configuring an OSPF router requires first enabling OSPF with the **feature ospf** command. This also applies to **evpn esi multihoming**, which enables its dependent commands such as **evpn multihoming core-tracking**. For more information about enabling features to access feature dependent commands, see the [Cisco Nexus 9000 Configuration Guides](#) and [Cisco Nexus 3000 Configuration Guides](#).
- Using Sandbox to convert with DN is supported only for finding the DN of a CLI config. Any other workflow, for example, using DME to convert DN for CLI configuration commands is not supported.
- CLI to model or xml conversion will not happen for OSPFv2 interface commands until you explicitly enable OSPF on interface by configuring router instance and area using **[no] ip router ospf <tag> area {<area-id-ip> | <area-id-int>} [secondaries none]** command.
- The Command pane (the top pane) supports a maximum of 10,000 individual lines of input.
- When you use XML or JSON as the Message Type for CLI input, you can use semicolon to separate multiple commands on the same line. However, when you use JSON RPC as the Message Type for CLI input, you cannot enter multiple commands on the same line and separate them with a semicolon (;).

For example, assume that you want to send **show hostname** and **show clock** commands through JSON RPC as the following.

In the Sandbox, you enter the CLIs as follows.

```
show hostname ; show clock
```

In the JSON RPC request, the input is formatted as follows.

```
[
  {
    "jsonrpc": "2.0",
    "method": "cli",
    "params": {
      "cmd": "show hostname ; show clock",
      "version": 1
    },
    "id": 1
  }
]
```

When you send the request, the response returns the following error.

```
{
  "jsonrpc": "2.0",
  "error": {
    "code": -32602,
    "message": "Invalid params",
    "data": {
      "msg": "Request contains invalid special characters"
    }
  },
  "id": 1
}
```

This situation occurs because the Sandbox parses each command in a JSON RPC request as individual items and assigns an ID to each. When using JSON RPC requests, you cannot use internal punctuation to separate multiple commands on the same line. Instead, enter each command on a separate line and the request completes successfully.

Continuing with the same example, enter the commands as follows in the NX-API CLI.

```
show hostname
show clock
```

In the request, the input is formatted as follows.

```
[
  {
    "jsonrpc": "2.0",
    "method": "cli",
    "params": {
      "cmd": "show hostname",
      "version": 1
    },
    "id": 1
  },
  {
    "jsonrpc": "2.0",
    "method": "cli",
    "params": {
      "cmd": "show clock",
      "version": 1
    },
    "id": 2
  }
]
```

The response completes successfully.

```
[
  {
    "jsonrpc": "2.0",
    "result": {
      "body": {
        "hostname": "switch-1"
      }
    },
    "id": 1
  },
  {
    "jsonrpc": "2.0",
    "result": {
      "body": {
        "simple_time": "12:31:02.686 UTC Wed Jul 10 2019\n",
        "time_source": "NTP"
      }
    },
    "id": 2
  }
]
```

```

    "id": 2
  }
]

```

Configuring the Message Format and Command Type

The **Message Format** and **Command Type** are configured in the upper right corner of the Command pane (the top pane). For **Message Format**, choose the format of the API protocol that you want to use. The Developer Sandbox supports the following API protocols:

Table 13: NX-OS API Protocols

| Protocol | Description |
|-------------|--|
| json-rpc | A standard lightweight remote procedure call (RPC) protocol that can be used to deliver NX-OS CLI commands in a JSON payload. The JSON-RPC 2.0 specification is outlined by jsonrpc.org . |
| xml | Cisco NX-API proprietary protocol for delivering NX-OS CLI or bash commands in an XML payload. |
| json | Cisco NX-API proprietary protocol for delivering NX-OS CLI or bash commands in a JSON payload. |
| nx-api rest | Cisco NX-API proprietary protocol for manipulating and reading managed objects (MOs) and their properties in the internal NX-OS data management engine (DME) model. For more information about the Cisco Nexus 3000 and 9000 Series NX-API REST SDK, see https://developer.cisco.com/site/cisco-nexus-nx-api-references/ . |
| nx yang | The YANG ("Yet Another Next Generation") data modeling language for configuration and state data. |

When the **Message Format** has been chosen, a set of **Command Type** options are presented just below the **Message Format** control. The **Command Type** setting can constrain the input CLI and can determine the **Request** and **Response** format. The options vary depending on the **Message Format** selection. For each **Message Format**, the following table describes the **Command Type** options:

Table 14: Command Types

| Message format | Command type |
|----------------|---|
| json-rpc | <ul style="list-style-type: none"> cli — show or configuration commands cli-ascii — show or configuration commands, output without formatting |

| Message format | Command type |
|----------------|---|
| xml | <ul style="list-style-type: none"> • cli_show — show commands. If the command does not support XML output, an error message will be returned. • cli_show_ascii — show commands, output without formatting • cli_conf — configuration commands. Interactive configuration commands are not supported. • bash — bash commands. Most non-interactive bash commands are supported. <p>Note The bash shell must be enabled in the switch.</p> |
| json | <ul style="list-style-type: none"> • cli_show — show commands. If the command does not support XML output, an error message will be returned. • cli_show_ascii — show commands, output without formatting • cli_conf — configuration commands. Interactive configuration commands are not supported. • bash — bash commands. Most non-interactive bash commands are supported. <p>Note The bash shell must be enabled in the switch.</p> |
| nx-api rest | <ul style="list-style-type: none"> • cli — configuration commands • model — DN and corresponding payload. |
| nx yang | <ul style="list-style-type: none"> • json — JSON structure is used for payload • xml — XML structure is used for payload |

Output Chunking

In order to handle large show command output, some NX-API message formats support output chunking for show commands. In this case, an **Enable chunk mode** checkbox appears below the **Command Type** control along with a session ID (**SID**) type-in box.

When chunking is enabled, the response is sent in multiple "chunks," with the first chunk sent in the immediate command response. In order to retrieve the next chunk of the response message, you must send an NX-API request with **SID** set to the session ID of the previous response message.

Using the Developer Sandbox

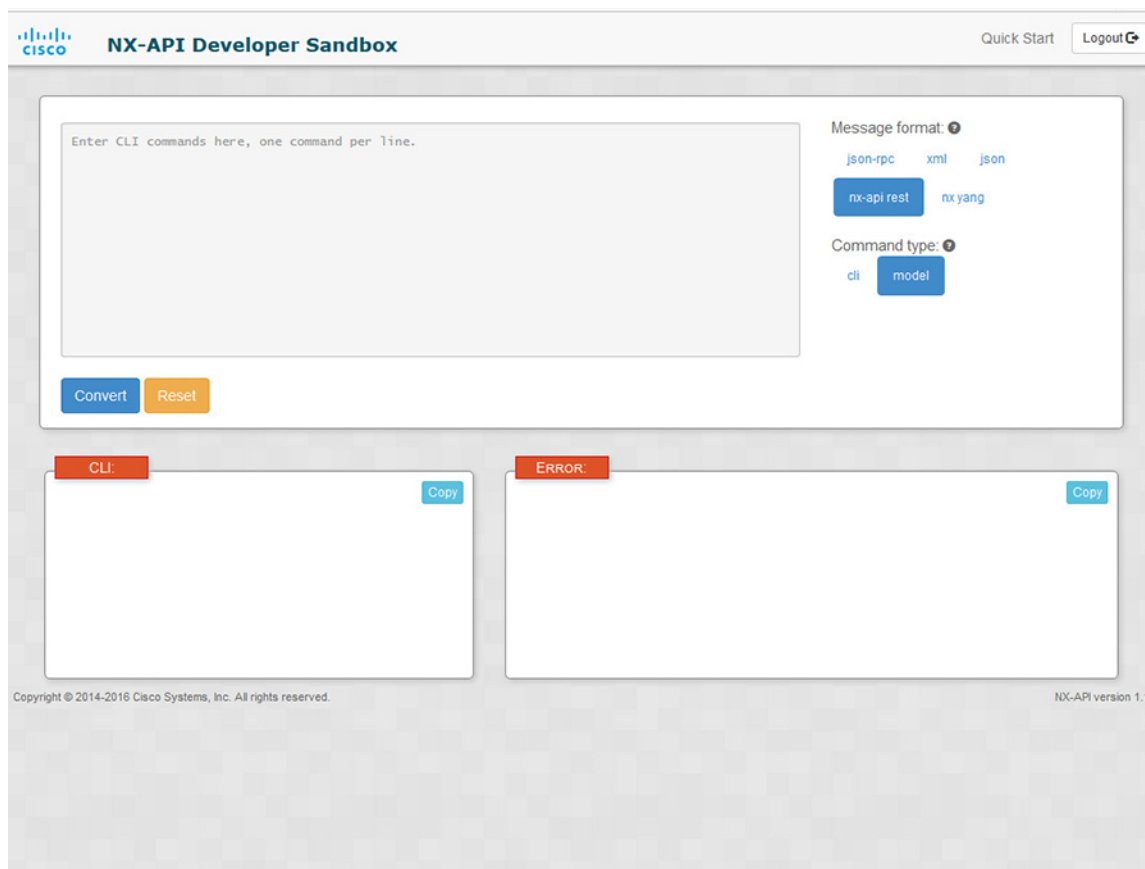
Using the Developer Sandbox to Convert CLI Commands to REST Payloads



Tip Online help is available by clicking **Quick Start** in the upper right corner of the Sandbox window. Additional details, such as response codes and security methods, can be found in the chapter "NX-API CLI". Only configuration commands are supported.

Procedure

- Step 1** Configure the **Message Format** and **Command Type** for the API protocol you want to use. For detailed instructions, see [Configuring the Message Format and Command Type, on page 272](#).
- Step 2** Type or paste NX-OS CLI configuration commands, one command per line, into the text entry box in the top pane. You can erase the contents of the text entry box (and the **Request** and **Response** panes) by clicking **Reset** at the bottom of the top pane.



Step 3 Click the **Convert** at the bottom of the top pane.

If the CLI commands contain no configuration errors, the payload appears in the **Request** pane. If errors are present, a descriptive error message appears in the **Response** pane.

The screenshot displays the NX-API Developer Sandbox interface. At the top, there is a header with the Cisco logo, the text "NX-API Developer Sandbox", and links for "Quick Start" and "Logout". The main area is divided into several sections:

- Request Pane:** Contains a text area with a JSON payload:

```
api/mo/sys.json
{
  "topSystem": {
    "attributes": {
      "name": "REST2CLI"
    }
  }
}
```

 To the right of this area are two dropdown menus: "Message format:" with options "json-rpc", "xml", "json", "nx-api rest" (selected), and "nx yang"; and "Command type:" with options "cli" and "model" (selected). Below the text area are "Convert" and "Reset" buttons.
- CLI Pane:** Labeled "CLI:" in a red header, it shows the resulting CLI command: "hostname REST2CLI". A "Copy" button is located to the right.
- ERROR Pane:** Labeled "ERROR:" in a red header, it is currently empty. A "Copy" button is located to the right.

At the bottom of the interface, there is a status bar with the text "Waiting for bam.nr-data.net..." and a copyright notice: "Copyright © 2014-2016 Cisco Systems, Inc. All rights reserved." and "NX-API version 1.1".

Step 4 When a valid payload is present in the **Request** pane, you can click **POST** to send the payload as an API call to the switch.

The response from the switch appears in the **Response** pane.

Warning Clicking **POST** commits the command to the switch, which can result in a configuration or state change.

Step 5 You can copy the contents of the **Request** or **Response** pane to the clipboard by clicking **Copy** in the pane.

Step 6 You can obtain a Python implementation of the request on the clipboard by clicking **Python** in the **Request** pane.

Using the Developer Sandbox to Convert from REST Payloads to CLI Commands



Tip Online help is available by clicking **Quick Start** in the upper right corner of the Sandbox window. Additional details, such as response codes and security methods, can be found in the chapter "NX-API CLI".

SUMMARY STEPS

1. Select **nx-api rest** as the **Message Format** and **model** as the **Command Type**.
2. Enter a DN and payload into the text entry box in the top pane. Then click on the **Convert** button below the top pane.

DETAILED STEPS

Procedure

Step 1 Select **nx-api rest** as the **Message Format** and **model** as the **Command Type**.

Example:

The screenshot shows the NX-API Developer Sandbox interface. At the top left is the Cisco logo and the title "NX-API Developer Sandbox". At the top right are "Quick Start" and "Logout" links. The main area contains a text input box with the placeholder text "Enter CLI commands here, one command per line." Below this box are "Convert" and "Reset" buttons. To the right of the input box are two dropdown menus: "Message format" with options "json-rpc", "xml", "json", "nx-api rest", and "nx yang"; and "Command type" with options "cli" and "model". Below the main area are two empty output boxes: "CLI" and "ERROR", each with a "Copy" button. The footer contains the text "Copyright © 2014-2016 Cisco Systems, Inc. All rights reserved." and "NX-API version 1.1".

Step 2 Enter a DN and payload into the text entry box in the top pane. Then click on the **Convert** button below the top pane.

Example:

For this example, the DN is `api/mo/sys.json` and the NX-API REST payload is:

```
{
  "topSystem": {
    "attributes": {
      "name": "REST2CLI"
    }
  }
}
```

The screenshot displays the NX-API Developer Sandbox interface. At the top left is the Cisco logo and the text "NX-API Developer Sandbox". At the top right are "Quick Start" and "Logout" links. The main area contains a text editor with a REST payload:

```
api/mo/sys.json
{
  "topSystem": {
    "attributes": {
      "name": "REST2CLI"
    }
  }
}
```

Below the text editor are "Convert" and "Reset" buttons. To the right, there are controls for "Message format" (with options: json-rpc, xml, json, nx-api rest, nx.yang) and "Command type" (with options: cli, model). Below these are two output panes: "CLI:" and "ERROR:". The "CLI:" pane is currently empty, and the "ERROR:" pane contains the text "Waiting for bam.nr-data.net...". At the bottom left, there is a copyright notice: "Copyright © 2014-2016 Cisco Systems, Inc. All rights reserved." and at the bottom right, "NX-API version 1.1".

When you click on the **Convert** button, the CLI equivalent appears in the **CLI** pane as shown in the following image.

NX-API Developer Sandbox
Quick Start [Logout](#)

```
api/mo/sys.json
{
  "topSystem": {
    "attributes": {
      "name": "REST2CLI"
    }
  }
}
```

Message format: ?

[json-rpc](#)
[xml](#)
[json](#)

[nx-api rest](#)
[nx yang](#)

Command type: ?

[cli](#)
[model](#)

Convert
Reset

CLI:

```
hostname REST2CLI
```

[Copy](#)

ERROR:

[Empty error box]

[Copy](#)

Copyright © 2014-2016 Cisco Systems, Inc. All rights reserved. NX-API version 1.1

Waiting for bam.nr-data.net...

Cisco Nexus 9000 Series NX-OS Programmability Guide, Release 10.4(x)

279

Note The Developer Sandbox cannot convert all payloads into equivalent CLIs, even if the Sandbox converted the CLIs to NX-API REST payloads. The following is a list of possible sources of error that can prevent a payload from completely converting to CLI commands:

Table 15: Sources of REST2CLI Errors

| Payload Issue | Result |
|---|---|
| <p>The payload contains an attribute that does not exist in the MO.</p> <p>Example:</p> <pre>api/mo/sys.json { "topSystem": { "children": [{ "interfaceEntity": { "children": [{ "l1PhysIf": { "attributes": { "id": "eth1/1", "fakeattribute": "totallyFake" } } }] } }] } }</pre> | <p>The Error pane will return an error related to the attribute.</p> <p>Example:</p> <p>CLI</p> <p>Error unknown attribute 'fakeattribute' in element 'l1PhysIf'</p> |
| <p>The payload includes MOs that aren't yet supported for conversion:</p> <p>Example:</p> <pre>api/mo/sys.json { "topSystem": { "children": [{ "dhcpEntity": { "children": [{ "dhcpInst": { "attributes": { "SnoopingEnabled": "yes" } } }] } }] } }</pre> | <p>The Error Pane will return an error related to the unsupported MO.</p> <p>Example:</p> <p>CLI</p> <p>Error The entire subtree of "sys/dhcp" is not converted.</p> |

NX-API Developer Sandbox: NX-OS Release 9.2(2) and Later

About the NX-API Developer Sandbox

The Cisco NX-API Developer Sandbox is a web form hosted on the switch. It translates NX-OS CLI commands into equivalent XML or JSON payloads and converts NX-API REST payloads into their CLI equivalents.

The web form is a single screen with three panes — Command (top pane), Request (middle pane), and Response (bottom pane) — as shown in the figure below. The designated name (DN) field is located between the Command and Request panes (seen in the figure below located between the **POST** and **Send** options).

The Request pane also has a series of tabs. Each tab represents a different language: **Python**, **Python3**, **Java**, **JavaScript**, and **Go-Lang**. Each tab enables you to view the request in the respective language. For example, after converting CLI commands into an XML or JSON payload, click the **Python** tab to view the request in Python, which you can use to create scripts.

Figure 2: NX-API Developer Sandbox with Example Request and Output Response

The screenshot displays the NX-API Developer Sandbox interface. At the top, there is a header with the Cisco logo and the text "NX-API Sandbox". To the right of the header are links for "Quick Start", "Command Reference", and "Logout".

The main interface is divided into three panes:

- Command pane (top):** Contains the text "show version". To the right of this pane are three dropdown menus: "Method:" (set to "NX-API-CLI"), "Message format:" (set to "json-enc"), and "Input type:" (set to "cli_ascii").
- Request pane (middle):** Features a "POST" label, a text input field containing "/ins", and three buttons: "Send", "Reset", and "Output Schema". Below this is a tabbed interface with tabs for "Request", "Python", "Python3", "Java", "JavaScript", and "Go-Lang". The "Request" tab is active, showing a JSON payload:

```
{
  "jsonrpc": "2.0",
  "method": "cli_ascii",
  "params": {
    "cmd": "show version",
    "version": 1
  },
  "id": 1
}
```

 A green "Copy" button is located to the right of the JSON text.
- Response pane (bottom):** Labeled "Response:", it shows the JSON response:

```
{
  "jsonrpc": "2.0",
  "result": {
    "msg": "Cisco Nexus Operating System (NX-OS) Software\nTAC support: http://www.cisco.com/tac\nDocuments: http://www.cisco.com/en/US/products/ps9372/tsd_products_support_series_home."
  },
  "id": 1
}
```

 A green "Copy" button is also present to the right of the response text.

Controls in the Command pane enable you to choose a supported API, such as NX-API REST, an input type, such as model (payload) or CLI, and a message format, such as XML or JSON. The available options vary depending on the chosen method.

When you choose the NX-API-REST (DME) method, type or paste one or more CLI commands into the Command pane, and click **Convert**, the web form converts the commands into a REST API payload, checking for configuration errors, and displays the resulting payload in the Request pane. If you then choose to post the payload directly from the sandbox to the switch (by choosing the **POST** option and clicking **SEND**), the Response pane displays the API response. For more information, see [Using the Developer Sandbox to Convert CLI Commands to REST Payloads, on page 287](#)

Conversely, the Cisco NX-API Developer Sandbox checks the payload for configuration errors then displays the equivalent CLIs in the Response pane. For more information, see [Using the Developer Sandbox to Convert from REST Payloads to CLI Commands, on page 290](#)

Guidelines and Limitations

Following are the guidelines and limitations for the Developer Sandbox:

- Clicking **Send** in the Sandbox commits the command to the switch, which can result in a configuration or state change.
- Some feature configuration commands are not available until their associated feature has been enabled. For example, configuring a BGP router requires first enabling BGP with the **feature bgp** command. Similarly, configuring an OSPF router requires first enabling OSPF with the **feature ospf** command. This also applies to **evpn esi multihoming**, which enables its dependent commands such as **evpn multihoming core-tracking**. For more information about enabling features to access feature dependent commands, see the [Cisco Nexus 9000 Configuration Guides](#) [Cisco Nexus 3000 Configuration Guides](#).
- Using Sandbox to convert with DN is supported only for finding the DN of a CLI config. Any other workflow, for example, using DME to convert DN for CLI configuration commands is not supported.
- CLI to model or xml conversion will not happen for OSPFv2 interface commands until you explicitly enable OSPF on interface by configuring router instance and area using **[no] ip router ospf <tag> area <area-id-ip> | <area-id-int> [secondaries none]** command.
- The Command pane (the top pane) supports a maximum of 10,000 individual lines of input.
- When you use XML or JSON as the Message Type for CLI input, you can use semicolon to separate multiple commands on the same line. However, when you use JSON RPC as the Message Type for CLI input, you cannot enter multiple commands on the same line and separate them with a semicolon (;).

For example, assume that you want to send **show hostname** and **show clock** commands through JSON RPC as the following.

In the Sandbox, you enter the CLIs as follows.

```
show hostname ; show clock
```

In the JSON RPC request, the input is formatted as follows.

```
[
  {
    "jsonrpc": "2.0",
    "method": "cli",
    "params": {
      "cmd": "show hostname ; show clock",
      "version": 1
    },
    "id": 1
  }
]
```


When you send the request, the response returns the following error.

```
{
  "jsonrpc": "2.0",
  "error": {
    "code": -32602,
    "message": "Invalid params",
    "data": {
      "msg": "Request contains invalid special characters"
    }
  },
  "id": 1
}
```

This situation occurs because the Sandbox parses each command in a JSON RPC request as individual items and assigns an ID to each. When using JSON RPC requests, you cannot use internal punctuation to separate multiple commands on the same line. Instead, enter each command on a separate line and the request completes successfully.

Continuing with the same example, enter the commands as follows in the NX-API CLI.

```
show hostname
show clock
```

In the request, the input is formatted as follows.

```
[
  {
    "jsonrpc": "2.0",
    "method": "cli",
    "params": {
      "cmd": "show hostname",
      "version": 1
    },
    "id": 1
  },
  {
    "jsonrpc": "2.0",
    "method": "cli",
    "params": {
      "cmd": "show clock",
      "version": 1
    },
    "id": 2
  }
]
```

The response completes successfully.

```
[
  {
    "jsonrpc": "2.0",
    "result": {
      "body": {
        "hostname": "switch-1"
      }
    },
    "id": 1
  },
  {
    "jsonrpc": "2.0",
    "result": {
      "body": {
        "simple_time": "12:31:02.686 UTC Wed Jul 10 2019\n",
        "time_source": "NTP"
      }
    }
  }
]
```

```

    }
  },
  "id": 2
}
]

```

Configuring the Message Format and Input Type

The **Method**, **Message format**, and **Input type** are configured in the upper right corner of the Command pane (the top pane). For **Method**, choose the format of the API protocol that you want to use. The Cisco NX-API Developer Sandbox supports the following API protocols:

Table 16: NX-OS API Protocols

| Protocol | Description |
|------------------|---|
| NXAPI-CLI | Cisco NX-API proprietary protocol for delivering NX-OS CLI or bash commands in an XML or a JSON payload. |
| NXAPI-REST (DME) | <p>Cisco NX-API proprietary protocol for manipulating and reading managed objects (MOs) and their properties in the internal NX-OS data management engine (DME) model. The NXAPI-REST (DME) protocol displays a drop-down list that enables you to choose from the following methods:</p> <ul style="list-style-type: none"> • POST • GET • PUT • DELETE <p>For more information about the Cisco Nexus 3000 and 9000 Series NX-API REST SDK, see https://developer.cisco.com/site/cisco-nexus-nx-api-references/.</p> |
| RESTCONF (Yang) | <p>The YANG ("Yet Another Next Generation") data modeling language for configuration and state data.</p> <p>The RESTCONF (Yang) protocol displays a drop-down list that enables you to choose from the following methods:</p> <ul style="list-style-type: none"> • POST • GET • PUT • PATCH • DELETE |

When you choose the **Method**, a set of **Message format** or **Input type** options are displayed in a drop-down list. The **Message format** can constrain the input CLI and determine the **Request** and **Response** format. The options vary depending on the **Method** you choose.

The following table describes the **Input/Command type** options for each **Message format**:

Table 17: Command Types

| Method | Message format | Input/Command type |
|-----------|----------------|--|
| NXAPI-CLI | json-rpc | <ul style="list-style-type: none"> • cli — show or configuration commands • cli-ascii — show or configuration commands, output without formatting • cli-array — show commands. Similar to cli, but with cli_array, data is returned as a list of one element, or an array, within square brackets, []. |
| NXAPI-CLI | xml | <ul style="list-style-type: none"> • cli_show — show commands. If the command does not support XML output, an error message will be returned. • cli_show_ascii — show commands, output without formatting • cli_conf — configuration commands. Interactive configuration commands are not supported. • bash — bash commands. Most non-interactive bash commands are supported. <p>Note The bash shell must be enabled in the switch.</p> |
| NXAPI-CLI | json | <ul style="list-style-type: none"> • cli_show — show commands. If the command does not support XML output, an error message will be returned. <p>Note Beginning with Cisco NX-OS Release 9.3(3), the cli_show_array command is recommended over the cli_show command.</p> <ul style="list-style-type: none"> • cli_show_array — show commands. Similar to cli_show, but with cli_show_array, data is returned as a list of one element, or an array, within square brackets []. • cli_show_ascii — show commands, output without formatting • cli_conf — configuration commands. Interactive configuration commands are not supported. • bash — bash commands. Most non-interactive bash commands are supported. <p>Note The bash shell must be enabled in the switch.</p> |

| Method | Message format | Input/Command type |
|------------------|--|---|
| NXAPI-REST (DME) | | <ul style="list-style-type: none"> cli — CLI to model conversion model — Model to CLI conversion. |
| RESTCONF (Yang) | <ul style="list-style-type: none"> json — JSON structure is used for payload xml — XML structure is used for payload | |

Output Chunking

JSON and XML NX-API message formats enable you to receive large show command responses in 10-MB chunks. When received, the chunks are concatenated to create a valid JSON object or XML structure. To view a sample script that demonstrates output chunking, click the following link and choose the directory that corresponds to Release 9.3x: [Cisco NX-OS NXAPI](#).



Note For chunk JSON mode, the browser or python script part does not provide the valid JSON output (there will be no closing tags). To use chunk mode and get valid JSON, use the script provided in the directory.

You receive the first chunk in the immediate command response, which also includes a **sid** field that contains a session Id. To retrieve the next chunk, you enter the session Id from the previous chunk in the **SID** text box. You repeat the process until reaching the last response, which is indicated by the **eoc** (end of content) value in the **sid** field.

Chunk mode is available when using the **NXAPI-CLI** method with the **JSON** or **XML** format type and the **cli_show**, **cli_show_array**, or **cli_show_ascii** command type. For more information about configuring the chunk mode, see the *Chunk Mode Fields* table.



Note NX-API supports a maximum of 2 chunking sessions.

Table 18: Chunk Mode Fields

| Field Name | Description |
|--------------------------|--|
| Enable Chunk Mode | Click to place a check mark in the Enable Chunk Mode check box to enable chunking. When you enable chunk mode, responses that exceed 10 MB are sent in multiple chunks of up to 10 MB in size. |
| SID | Enter the session Id of the previous response in the SID text box to retrieve the next chunk of the response message. Note Only alphanumeric characters and ‘_’ are allowed. Invalid characters receive an error. |

Using the Developer Sandbox

Using the Developer Sandbox to Convert CLI Commands to REST Payloads

**Tip**

- Online help is available by clicking the help icons (?) next to the field names located in the upper-right corner of the Cisco NX-API Developer Sandbox window.
- For additional details, such as response codes and security methods, see the *NX-API CLI* chapter.
- Only configuration commands are supported.

The Cisco NX-API Developer Sandbox enables you to convert CLI commands to REST payloads.

Procedure

-
- Step 1** Click the **Method** drop-down list and choose **NXAPI-REST (DME)**.
The **Input** type drop-down list appears.
- Step 2** Click the **Input** type drop-down list and choose **cli**.
- Step 3** Type or paste NX-OS CLI configuration commands, one command per line, into the text entry box in the top pane.
You can erase the contents of the text entry box (and the **Request** and **Response** panes) by clicking **Reset** at the bottom of the top pane.

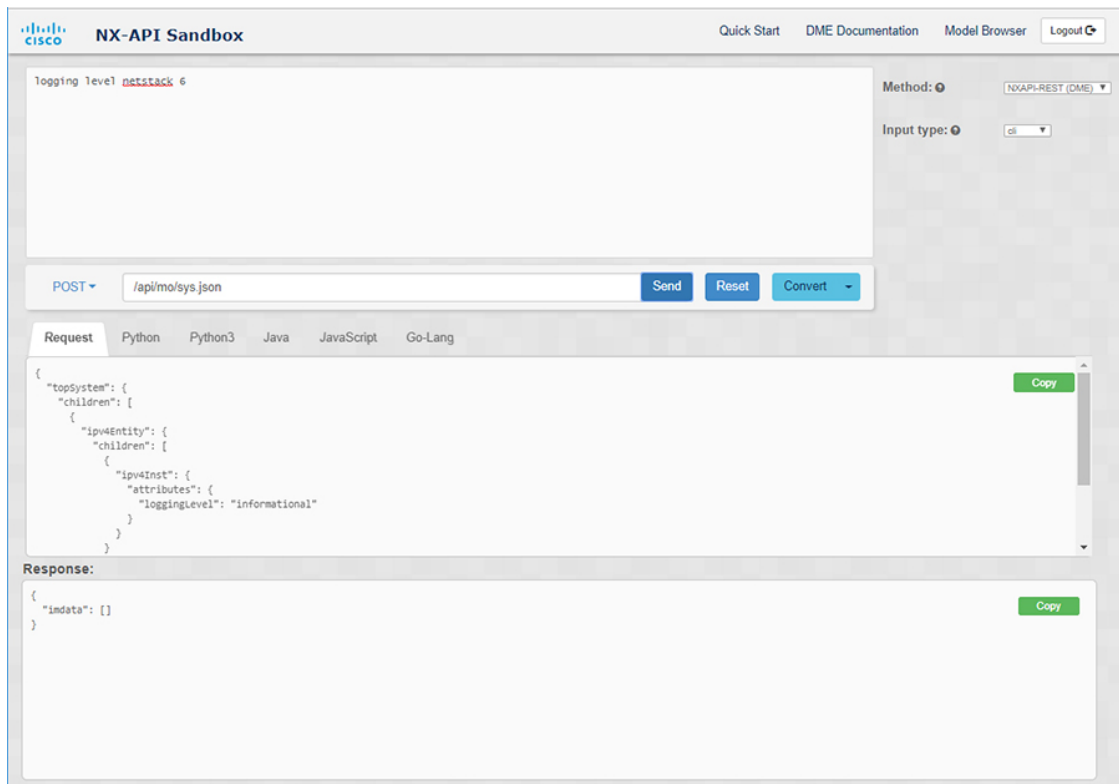
Step 4 Click **Convert**.

If the CLI commands contain no configuration errors, the payload appears in the **Request** pane. If errors are present, a descriptive error message appears in the **Response** pane.

Step 5 (Optional) To send a valid payload as an API call to the switch, click **Send**.

The response from the switch appears in the **Response** pane.

Warning Clicking **Send** commits the command to the switch, which can result in a configuration or state change.



Step 6 (Optional) To obtain the DN for an MO in the payload:

- a. From the **Request** pane, choose **POST**.
- b. Click the **Convert** drop-down list and choose **Convert (with DN)**.

The payload appears with with a **dn** field that contains the DN that corresponds to each MO in the payload.

Step 7 (Optional) To overwrite the current configuration with a new configuration:

- a. Click the **Convert** drop-down list and choose **Convert (for Replace)**. The **Request** pane displays a payload with a **status** field set to **replace**.
- b. From the **Request** pane, choose **POST**.
- c. Click **Send**.

The current configuration is replaced with the posted configuration. For example, if you start with the following configuration:

```
interface eth1/2
  description test
  mtu 1501
```

Then use **Convert (for Replace)** to POST the following configuration:

```
interface eth1/2
  description testForcr
```

The `mtu` configuration is removed and only the new description (`testForCr`) is present under the interface. This change is confirmed when entering **show running-config**.

Step 8 (Optional) To copy the contents of a pane, such as the **Request** or **Response** pane, click **Copy**. The contents of the respective pane is copied to the clipboard.

Step 9 (Optional) To convert the request into one of the formats listed below, click on the appropriate tab in the **Request** pane:

- **Python**
- **Python3**
- **Java**
- **JavaScript**
- **Go-Lang**

Using the Developer Sandbox to Convert from REST Payloads to CLI Commands

The Cisco NX-API Developer Sandbox enables you to convert REST payloads to corresponding CLI commands. This option is only available for the NXAPI-REST (DME) method.



Tip

- Online help is available by clicking help icons (?) next to the Cisco NX-API Developer Sandbox field names. Click a help icon to get information about the respective field.

For additional details, such as response codes and security methods, see the chapter *NX-API CLI*.

- The top-right corner of the Cisco NX-API Developer Sandbox contains links for additional information. The links that appear depend on the **Method** you choose. The links that appear for the NXAPI-REST (DME) method:

- **NX-API References**—Enables you to access additional NX-API documentation.
- **DME Documentation**—Enables you to access the NX-API DME Model Reference page.
- **Model Browser**—Enables you to access Visore, the Model Browser. Note that you might have to manually enter the IP address for your switch to access the Visore page:

`https://management-ip-address/visore.html`.

Procedure

Step 1 Click the **Method** drop-down list and choose **NXAPI-REST (DME)**.

Example:


The screenshot shows the NX-API Sandbox interface. At the top, there is a header with the Cisco logo and the text "NX-API Sandbox". To the right of the header are links for "Quick Start", "DME Documentation", "Model Browser", and a "Logout" button. Below the header is a large text area for entering a DME payload, with the placeholder text "Enter DME payload here.". To the right of this area are two dropdown menus: "Method" set to "NX-API-REST (DME)" and "Input type" set to "model". Below these is a text input field containing the path "/api/mo/sys.json", followed by "Send", "Reset", and "Convert" buttons. Underneath is a tabbed interface with "Request" selected, and other tabs for "Python", "Python3", "Java", "JavaScript", and "Go-Lang". The "Request" tab shows a large empty area with a "Copy" button in the top right corner. Below the request area is a "Response:" section, also with a "Copy" button in the top right corner.

- Step 2** Click the **Input Type** drop-down list and choose **model**.
- Step 3** Enter the designated name (DN) that corresponds to the payload in the field above the Request pane.
- Step 4** Enter the payload in the Command pane.
- Step 5** Click **Convert**.

Example:

For this example, the DN is `/api/mo/sys.json` and the NX-API REST payload is:

```
{
  "topSystem": {
    "attributes": {
      "name": "REST2CLI"
    }
  }
}
```



When you click on the **Convert** button, the CLI equivalent appears in the **CLI** pane as shown in the following image.

The screenshot displays the Cisco NX-API Developer Sandbox interface. At the top, there is a navigation bar with links for "Quick Start", "DME Documentation", "Model Browser", and "Logout". The main area is divided into several sections:

- JSON Payload:** A text area containing a REST payload:


```
{
  "sysSystem": {
    "attributes": {
      "name": "REST2CLI"
    }
  }
}
```
- Method:** A dropdown menu set to "NX-API-REST (DME)".
- Input type:** A dropdown menu set to "model".
- URL:** A text input field containing "/api/mo/sys.json".
- Buttons:** "Send" (blue), "Reset" (orange), and "Convert" (blue).
- Request Tab:** A tabbed interface with "Request" selected. Below it, the converted CLI command is shown: "hostname REST2CLI". A green "Copy" button is located to the right.
- Response Tab:** A section labeled "Response:" with an empty text area and a green "Copy" button.

Note The Cisco NX-API Developer Sandbox cannot convert all payloads into equivalent CLIs, even if the sandbox converted the CLIs to NX-API REST payloads. The following is a list of possible sources of error that can prevent a payload from completely converting to CLI commands:

Table 19: Sources of REST2CLI Errors

| Payload Issue | Result |
|---|---|
| <p>The payload contains an attribute that does not exist in the MO.</p> <p>Example:</p> <pre>api/mo/sys.json { "topSystem": { "children": [{ "interfaceEntity": { "children": [{ "l1PhysIf": { "attributes": { "id": "eth1/1", "fakeattribute": "totallyFake" } } }] } }] } }</pre> | <p>The Error pane will return an error related to the attribute.</p> <p>Example:</p> <p>CLI</p> <p>Error unknown attribute 'fakeattribute' in element 'l1PhysIf'</p> |
| <p>The payload includes MOs that aren't yet supported for conversion:</p> <p>Example:</p> <pre>api/mo/sys.json { "topSystem": { "children": [{ "dhcpEntity": { "children": [{ "dhcpInst": { "attributes": { "SnoopingEnabled": "yes" } } }] } }] } }</pre> | <p>The Error Pane will return an error related to the unsupported MO.</p> <p>Example:</p> <p>CLI</p> <p>Error The entire subtree of "sys/dhcp" is not converted.</p> |

Using the Developer Sandbox to Convert from RESTCONF to json or XML



Tip

- Online help is available by clicking the help icon (?) in the upper-right corner of the Cisco NX-API Developer Sandbox window.
- Click on the **Yang Documentation** link in the upper right corner of the Sandbox window to go to the Model Driven Programmability with Yang page.
- Click on the **Yang Models** link in the upper right corner of the Sandbox window to access the YangModels GitHub site.

Procedure

Step 1 Click the **Method** drop-down list and choose **RESTCONF (Yang)**.

Example:

The screenshot displays the Cisco NX-API Sandbox interface. At the top, there is a navigation bar with links for 'Quick Start', 'Yang Documentation', 'Yang Models', and 'Logout'. The main area contains a text input field with the command 'logging level netstack 4'. Below this, there is a 'POST' dropdown menu and a text entry box containing 'restconf/data/Cisco-NX-OS-device:System/'. To the right of the input field, there are 'Send', 'Reset', and 'Convert' buttons. On the right side of the interface, there are two dropdown menus: 'Method' set to 'RESTCONF (Yang)' and 'Message format' set to 'json'. Below the input field, there are tabs for 'Request', 'Python', 'Python3', 'Java', 'JavaScript', and 'Go-Lang'. The 'Request' tab is active, showing a large empty text area with a 'Copy' button. Below the 'Request' tab, there is a 'Response' section with another large empty text area and a 'Copy' button.

Step 2 Click **Message format** and choose either **json** or **xml**.

Step 3 Enter a command in the text entry box in the top pane.

Step 4 Choose a message format.

Step 5 Click **Convert**.

Example:

For this example, the command is **logging level netstack 6** and the message format is json:


The screenshot shows the NX-API Sandbox interface. At the top, there is a header with the Cisco logo and the text "NX-API Sandbox". To the right of the header are links for "Quick Start", "Yang Documentation", "Yang Models", and a "Logout" button. Below the header, there is a text input field containing the command "logging level netstack 6". To the right of this field, there are two dropdown menus: "Method:" set to "RESTCONF (Yang)" and "Message format:" set to "json". Below the input field, there is a "POST" dropdown menu, a text input field containing the URL "restconf/data/Cisco-NX-OS-device:System/", and three buttons: "Send" (orange), "Reset" (blue), and "Convert" (blue). Below the "Convert" button, there are tabs for "Request", "Python", "Python3", "Java", "JavaScript", and "Go-Lang". The "Request" tab is selected, and it displays a JSON object:


```
{
  "ipv4-items": {
    "inst-items": {
      "loggingLevel": "informational"
    }
  }
}
```

 To the right of the JSON object is a green "Copy" button. Below the "Request" tab, there is a "Response:" label and an empty text area with a green "Copy" button to its right.

Example:

For this example, the command is **logging level netstack 6** and the message format is xml:



Note When converting a negated CLI to a Yang payload using the XML or JSON message format, the sandbox throws a warning and disables the **Send** option. The warning message that appears depends on the message format:

- For the XML message format — "This is a Netconf payload as it is being generated for DELETE operation(s), hence SEND option is disabled for Restconf!"
- For the JSON message format—"This is a gRPC payload as it is being generated for DELETE operation(s), hence SEND option is disabled for Restconf!"

Step 6 You can also convert the request into the following formats by clicking on the appropriate tab in the **Request** pane:

- Python
- Python3
- Java
- JavaScript
- Go-Lang

Note The Java-generated script does not work if you choose the PATCH option from the drop-down menu in the area above the Request tab. This is a known limitation with Java and is expected behavior.



PART **V**

Model-Driven Programmability

- [Overview, on page 301](#)
- [Original and OpenConfig YANG, on page 305](#)
- [NETCONF Agent, on page 333](#)
- [RESTCONF Agent, on page 371](#)
- [gRPC Agent, on page 381](#)
- [gNMI - Management Interface, on page 393](#)
- [gNOI - Operation Interface, on page 417](#)
- [gRPC Tunnel, on page 425](#)
- [Telemetry, on page 435](#)
- [Diagnosis and Serviceability, on page 515](#)



CHAPTER 25

Overview

- [About Model-Driven Programmability, on page 301](#)
- [About the Programmable Interface Infrastructure, on page 301](#)

About Model-Driven Programmability

The model-driven programmability (MDP) of Cisco NX-OS device allows you to automate the configuration and control of the device. It also allows you to monitor the state changes of the device.

Data modeling provides a programmatic and standards-based method of writing configurations to the network device, replacing the process of manual configuration. Data models are written in a standard, industry-defined language. Although configuration using a CLI may be more human-friendly, automating the configuration using data models results in better scalability.

The Cisco NX-OS device supports the YANG data modeling language. YANG is a data modeling language used to describe configuration and operational data, remote procedure calls, and notifications for network devices.

The following are standards-based programmable interfaces supported by NX-OS:


- NETCONF
- RESTCONF
- gRPC

The following are observability interfaces supported by NX-OS:

- NETCONF
- gRPC
- Telemetry

About the Programmable Interface Infrastructure

This section provides a brief overview of the NX-OS Programmable Interface infrastructure.



When a request is received through NETCONF, RESTCONF, or gRPC, the request is converted into an abstract message object. That message object is distributed to the underlying model infrastructure based on the model namespace in the request. Using the namespace, the appropriate model is selected, and the request is passed to it for processing. The model infrastructure executes the request (read or write) on the underlying NX-OS device datastore. The results are returned to the agent of origin for response transmission back to the requesting client.

NX-OS Programmability Agents

Agents provide an interface between the NX-OS device and the external clients. They specify the specific transport, the protocol, and the encoding of the communications with the networking device. NX-OS Programmable Interfaces support three agents: NETCONF, RESTCONF, and gRPC, each providing different interfaces for configuration management via YANG models.



Note Supported YANG models for each Cisco NX-OS release are provided at <https://devhub.cisco.com/ui/native/open-nxos-agents>.

Table 20: NX-OS Programmable Interface Agents

| Agent | Transport | Protocol | Encoding |
|----------|-----------|---|-----------------|
| NETCONF | SSH | NA | XML |
| RESTCONF | HTTP | draft-ietf-netconf-restconf-10 ¹ | XML or JSON |
| gRPC | HTTP | gRPC Protocol ^[2] | Google Protobuf |

The protocol specifications are described in the following documents:

- [RESTCONF Agent, on page 371](#)
- [gRPC Agent, on page 381](#)

NX-OS Observability Agents

Note that there exist read-only streaming interfaces which serve to monitor the switch events. NX-OS supports standard-based NETCONF, gRPC, and the proprietary software telemetry.

Table 21: NX-OS Observability Interface Agents

| Agent | Transport | Protocol |
|-----------|-----------------|---------------|
| NETCONF | SSH | NA |
| gRPC | HTTP | gRPC Protocol |
| Telemetry | HTTP, UDP, GRPC | NA |

Model Infrastructure

The Model Infrastructure takes requests that are received from the Agent, determines the namespace that is associated with the YANG model in the request, and selects the model component matching the namespace to process the request. When the selected model component completes request processing, the processing results are sent to the requesting Agent for transmission back to the client. The Model Infrastructure is also responsible for handling protocol initiation requests involving authentication, handshaking, and so on, as specified by the Agent protocol.

Native (Device) YANG Model

The YANG model, where the Device Configuration is described, is called a device YANG model. The Device Model is manifested in the Model Infrastructure as another model component with the device namespace.

Common (OpenConfig) YANG Models

A Common Model is another kind of model component defined by the public community. The YANG paths in the common model could translate to the equivalent Device Model elements. These equivalent Device Model elements are used to read and write Device Model data in the Device YANG context.



Note Cisco NX-OS maintains a centralized DME database. The Device and OpenConfig YANG are the wrapping representation of the DME database. This guarantees consistency across all programmability interfaces. Users of different interfaces would see the same config/state of the switch without discrepancy

Additional YANG References

For additional information about YANG, see <https://handwiki.org/wiki/YANG>.



CHAPTER 26

Original and OpenConfig YANG

- [About YANG, on page 305](#)
- [Cisco Device YANG, on page 306](#)
- [Guidelines and Limitations, on page 306](#)
- [Migration from DME to Device YANG, on page 306](#)
- [About OpenConfig YANG, on page 307](#)
- [Upgrading or Downgrading YANG Version, on page 329](#)
- [RBAC for YANG , on page 329](#)
- [Guidelines and Limitations, on page 329](#)
- [Configuring YANG RBAC, on page 330](#)
- [Troubleshooting YANG, on page 331](#)

About YANG

Yet Another Next Generation (YANG) is a data modeling language for the definition of data that are transferred over network management protocols. YANG defines a primary construct to define a tree data structure and various commonly used types to represent devices configuration and operational status.

YANG is a textual language, it can be used to define a model. Cisco NX-OS supports the following YANG models:

- **Cisco NX-OS-device**

This is a Cisco NX-OS proprietary YANG model, which reflects the DME REST definitions. This model is referred as 'Device' or 'Original' model in this document. If you prefer to have full access over NX-OS device, it is suggested to use the 'Original' model. For more information, see the [NX-API REST](#) guide.

- **Open Config**

The primary purpose of this model is to abstract the common functionalities across networking management. The output is a collection of YANG models, which is referred as the Open Config model. For more information, see [About OpenConfig YANG](#).

See [Yang NX Repository](#) guide to view models supported for NX-OS. Open Config YANG models are grouped based on Cisco NX-OS release.

This chapter describes the usage and limitations of Original and Open Config models.

Cisco Device YANG

Cisco NX-OS retains a single DME database to represent a centralized view of the configuration and state. You can access DME through a secured NX-API REST interface. If you prefer to operate based on standard YANG semantics, you can access the device through below Original YANG namespace:

Cisco-NX-OS-device - <http://cisco.com/ns/yang/cisco-nx-os-device>.



Note The device YANG URL is used to define a unique namespace identifier. The above URL is an example and not the accessible HTTP URL.

Guidelines and Limitations

The original YANG represents the DME database and inherits some behavior which is not in accordance with standard YANG operation. The following are the guidelines and limitations for YANG model:

- If a property is a numeric value, DME supports to define string aliases for specific values. You can use an alias or the numeric value to configure the property and DME converts the alias to a numeric value. If you prefer to adhere to the YANG standards, you can avoid using an alias.
- If a property is bitmask, DME supports special keywords (+,-) to gradually change the number in the bitmask range. This syntax is shown in Device YANG and it is not complaint to YANG. If you prefer to adhere to YANG standards, you can avoid using special keywords.
- In Device YANG, ephemeral paths are identified with comment **Ephemeral data**. These paths are NX-OS specific nonpersistent high volume data and are managed differently from the rest of the models. You can retrieve only when `<GET> query's <FILTER>` parameters marked specific to the elements marked with the comment. For more information, see *Ephemeral data support*.
- Beginning with Cisco NX-OS Release 10.4(3)F, YANG is supported on 92348GC-X.

Migration from DME to Device YANG

If you are familiar with Cisco NX-OS DME interfaces, you can migrate existing DME to Device YANG. As the switches maintain one to one correspondence between DME and Device YANG models.

The following table shows representation of VRFs in DME and the Device YANG.

| DME | Device YANG |
|---|--|
| <pre> sys +- name +- inst +- name = default +- inst +- name = management </pre> | <pre> System +- name +- inst-items +- Inst-list +- name = default +- Inst-list +- name = management </pre> |

Make sure that you follow the below mentioned guidelines for migrating DME model to Device YANG model:

- The root element 'sys' MO is represented as 'System' root element in the YANG model.
- The DME MO is represented with the suffix '- items' in the YANG model.
- If DME MO has children of the same type, the Device YANG model adds an intermediate parent node for each child and use the suffix '-list' to represent the lists. This behavior is the same as OpenConfig.
- The DME property name remains the same in the Device YANG model.

About OpenConfig YANG

OpenConfig YANG model supports advanced networking standards, such as declarative configuration and model-driven management and operations. OpenConfig is a business data model for configuring and monitoring the network. This data model helps with moving from a pull model to a push model, with subscriptions and event update streaming.

Cisco NX-OS supports a wide range of functional areas, such as BGP, OSPF, Interface L2 and L3, VRFs, VLANs, and TACACs.

Guidelines and Limitations

The following are the guidelines and limitations of the OpenConfig YANG model:

Still guidelines and limitations required

- In networking protocol, you can add L2 MAC in a different way. When you use MAC leaf switch property values as input to perform an `NETCONF GET` operation, it is recommended to input the letters in MAC addresses in the uppercase in the forma of (AA:AA:AA:AA:AA:AA)". For an example, source-mac: 0A:0B:0C:0D:0E:0F.

OpenConfig Paths

Table 22: OpenConfig IP

| Path | Description |
|---|--|
| <pre>.../oc-ip:ip .../oc-ip:prefix_length</pre> | <p>For IPv4 and IPv6 addresses, you must provide the same operation to remove and delete the IP address field (oc-ip:ip and oc-ip:prefix_length)</p> <p>For example:</p> <pre>oc-ip:ip: remove oc-ip:prefix_length: remove</pre> |

Table 23: OpenConfig Network Instance

| Path | Description |
|------|-------------|
| bgp | |

| Path | Description |
|------|--|
| | <p>With OpenConfig YANG network-instance (OCNI), and if you want to delete only BGP configuration of default VRF instead of BGP routing instance, you cannot delete BGP information at protocol or BGP level. You can enter an autonomous system number in the payload and delete only the configuration of the default VRF and not the BGP routing instance.</p> <p>The following is an example for payload to delete the configuration in the default VRF OF BGP.</p> <pre><rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101"> <edit-config> <target> <running/> </target> <config> <network-instances xmlns="http://openconfig.net/yang/network-instance"> <network-instance> <name>default</name> <protocols> <protocol> <identifier>BGP</identifier> <name>bgp</name> <bgp xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" nc:operation="delete"> <global> <config> <as>100</as> </config> </global> </bgp> </protocol> </protocols> </network-instance> </network-instances> </config> </edit-config> </rpc></pre> <p>Expected Behavior: The BG routing instance is deleted, which is similar to configuring <code>no route bgp 100</code>.</p> <p>Actual Behaviour: Only the BGP configuration in default VRF is deleted, and there is no similar CLI configuration.</p> <p>The following mentioned below is the running configuration before configuring delete operation:</p> <pre>router bgp 100 router-id 1.2.3.4 address-family ipv4 unicast vrf abc address-family ipv4 unicast maximum-paths 2</pre> <p>The following mentioned below is the running</p> |

| Path | Description |
|---------------------|--|
| | <p>configuration after configuring delete operation:</p> <pre>router bgp 100 vrf abc address-family ipv4 unicast maximum-paths 2</pre> |
| .../set-med | <p>Make sure that you don't configure BGP actions with <code>set-med</code> and OSPF commands with <code>metric</code> in the same route-map through OpenCofig NETCONF. As the OSPF command metrics succeed over BGP <code>set-med</code> properties.</p> <p>You must use two different route maps to set metrics in OSPF actions. Use <code>set-med</code> in BGP actions using separate route maps.</p> <p>Cisco recommends you not to change the metric of BGP actions to OSPF actions or OSPF actions to BGP actions of a route map in a single payload.</p> |
| .../is-reachability | <p>When you search for IS-IS <code>is-reachability</code> with <code>system-id</code> as key, the original DME returns all the entries. But, OpenConfig returns one entry with the highest metric value.</p> |

| Path | Description |
|--------------------------|-------------|
| .../bgp/global/config/as | |

| Path | Description |
|------|---|
| | <p>You must enter an autonomous system (AS) number to have a valid BGP instance. As there is no default value for an AS number, deleting NETCONF/OPENCONFIG<asn> without removing BGP instance results in the following highlighted error message:</p> <pre> 764 <nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="urn:uuid:1ea09de2-605e-46aa-984b-9dfdad03354d"> <nc:edit-config> <nc:target> <nc:running/> </nc:target> <nc:config> <network-instances xmlns="http://openconfig.net/yang/network-instance"> <network-instance> <name>default</name> <protocols> <protocol> <identifier>BGP</identifier> <name>bgp</name> <bgp> <global> <config nc:operation="delete"> <as>100</as> </config> </global> <neighbors> <neighbor> <neighbor-address>1.1.1.1</neighbor-address> <enable-bfd xmlns="http://openconfig.net/yang/bfd"> <config> <enabled>>true</enabled> </config> </enable-bfd> </neighbor> </neighbors> </bgp> </protocol> </protocols> </network-instance> </network-instances> </nc:config> </nc:edit-config> </nc:rpc> ## Received: <rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="urn:uuid:1ea09de2-605e-46aa-984b-9dfdad03354d"> <rpc-error> <error-type>protocol</error-type> <error-tag>operation-failed</error-tag> <error-severity>error</error-severity> <error-message xml:lang="en">invalid property value , for property asn, class bgpInst</error-message> </pre> |

| Path | Description |
|------|---|
| | <pre><error-path>/config/network-instances</error-path> </rpc-error> <rpc-error> <error-type>protocol</error-type> <error-tag>operation-failed</error-tag> <error-severity>error</error-severity> <error-message xml:lang="en">invalid property value , for property asn, class bgpInst Commit Failed</error-message> <error-path>/config/network-instances</error-path> </rpc-error> </rpc-reply></pre> |

| Path | Description |
|-------------------|-------------|
| .../protocol/ospf | |

| Path | Description |
|------|--|
| | <ul style="list-style-type: none"> • If you configure and remove an area configuration in OSPF, the deleted area shows in DME. These areas are shown in <code>GETCONFIG/GET</code> output in OpenConfig YANG. • You can configure one are in OpenConfig YANG in the OSPF policy <code>match ospf-area</code> configuration. Though you can configure multiple areas, such as <code>match ospf-area 100 101</code>. But in OpenConfig YANG, you can configure only one area. For an example, <code>match ospf-area 100</code>. • You cant configure both area virtual-link and area interface configurations payload in the same area list. Ensure that you split the area container payload as Virtual ink area and interface area in the same payload. • You cannot configure the MD5 authentication string in OSPF OpenConfig YANG. In the OSPF model, <code>Authentication-type</code> is defined for the Authentication: <pre>leaf authentication-type { type string; description "The type of authentication that should be used on this interface"; }</pre> • OSPF OpenConfig YANG doesn't support an option for password authentication. • The OSPF area authentication configuration is not supported. For example, <code>area 0.0.0.200 authentication message-digest</code> cannot be configured from OpenConfig YANG. • You cannot delete the OSPF or BGP instance configuration, which is in default VRF. For example, router ospf 1/router bgp 1, you cannot delete protocols container with the default network instance. • When you add an interface through OpenConfig YANG, OSPFv2 can send an error message. If there is an issue, you cannot add an interface, and the RPC reply has list merge failed errors as the following: <pre><rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"</pre> |

| Path | Description |
|------|---|
| | <pre> message-id="urn:uuid:39507023-8569-4cf8-869c-e19eaf76a260"> <rpc-error> <error-type>protocol</error-type> <error-tag>operation-failed</error-tag> <error-severity>error</error-severity> <error-message xml:lang="en">List Merge Failed: operation-failed</error-message> </rpc-error> </rpc-reply> </pre> |

Table 24: OpenConfig Routing Policy

| Path | Description |
|--|--|
| <pre> /bgp-defined-sets/community-sets/community-set/ /bgp-defined-sets/as-path-sets/as-path-set/ </pre> | <p>Action type is always permit for community-set and as-path-set.</p> <p>In OpenConfig YANG, there is no action type concept as there is in the CLI for community-set and as-path-set. Therefore, the action type is always permit for community-set and as-path-set.</p> |
| <pre> /bgp-defined-sets/community-sets/community-set/ /bgp-defined-sets/as-path-sets/as-path-set/ </pre> | <p>Action type is always permit for community-set and as-path-set.</p> <p>In OpenConfig YANG, there is no action type concept as there is in the CLI for community-set and as-path-set. Therefore, the action type is always permit for community-set and as-path-set.</p> |

| Path | Description |
|---|---|
| /bgp-defined-sets/community-sets/community-set/ | <p>In the CLI, <code>community-list</code> has two different types, such as standard and expanded. In the OpenConfig YANG model, <code>community-set-name</code> has no differentiation.</p> <p>When you create the <code>community-set-name</code> through OpenConfig YANG, the following changes are configured:</p> <ul style="list-style-type: none"> • The <code>_std</code> suffix is appended after <code>community-set-name</code> if <code>community-member</code> is in the standard form (AS:NN). • The <code>_exp</code> suffix is appended after <code>community-set-name</code> if <code>community-member</code> is in the expanded form (regex): <pre data-bbox="1019 768 1528 968"> <community-set> <community-set-name>oc_commset1d</community-set-name> <config> <community-set-name>oc_commset1d</community-set-name> <community-member>0:1</community-member> <community-member>_1_</community-member> </config> </community-set> </pre> <p>The preceding OpenConfig YANG configuration is mapped to the following CLI:</p> <pre data-bbox="967 1083 1528 1184"> ip community-list expanded oc_commset1d_exp seq 5 permit "_1_" ip community-list standard oc_commset1d_std seq 5 permit 0:1 </pre> |

| Path | Description |
|---|--|
| /bgp-defined-sets/community-sets/community-set/ | <p>In the CLI, <code>community-list</code> has two different types, such as standard and expanded. In the OpenConfig YANG model, <code>community-set-name</code> has no differentiation.</p> <p>When you create the <code>community-set-name</code> through OpenConfig YANG, the following changes are configured:</p> <ul style="list-style-type: none"> • The <code>_std</code> suffix is appended after <code>community-set-name</code> if <code>community-member</code> is in the standard form (AS:NN). • The <code>_exp</code> suffix is appended after <code>community-set-name</code> if <code>community-member</code> is in the expanded form (regex): <pre data-bbox="980 766 1490 968"> <community-set> <community-set-name>oc_commsetld</community-set-name> <config> <community-set-name>oc_commsetld</community-set-name> <community-member>0:1</community-member> <community-member>_1_</community-member> </config> </community-set> </pre> <p>The preceding OpenConfig YANG configuration is mapped to the following CLI:</p> <pre data-bbox="927 1079 1490 1178"> ip community-list expanded oc_commsetld_exp seq 5 permit "_1_" ip community-list standard oc_commsetld_std seq 5 permit 0:1 </pre> |

| Path | Description |
|---|-------------|
| /bgp-conditions/match-community-set/config/community-set/ | |

| Path | Description |
|------|---|
| | <p>OpenConfig YANG can only map to one <code>community-set</code>, while the CLI can match to multiple instances of the <code>community-set</code>:</p> <ul style="list-style-type: none"> • In the CLI: <pre>ip community-list standard 1-1 seq 1 permit 1:1 ip community-list standard 1-2 seq 1 permit 1:2 ip community-list standard 1-3 seq 1 permit 1:3 route-map To_LC permit 10 match community 1-1 1-2 1-3</pre> • The corresponding OpenConfig YANG payload follows: <pre><config> <routing-policy xmlns="http://openconfig.net/yang/routing-policy"> <defined-sets> <bgp-defined-sets xmlns="http://openconfig.net/yang/bgp-policy"> <community-sets> <community-set> <community-set-name>cs</community-set-name> <config> <community-set-name>cs</community-set-name> <community-member>1:1</community-member> <community-member>1:2</community-member> <community-member>1:3</community-member> </config> </community-set> </community-sets> </bgp-defined-sets> </defined-sets> <policy-definitions> <policy-definition> <name>To_LC</name> <statements> <statement> <name>10</name> <conditions> <bgp-conditions xmlns="http://openconfig.net/yang/bgp-policy"> <match-community-set> <config> <community-set>cs</community-set> </config> </match-community-set> </bgp-conditions> </conditions> </statement> </statements> </policy-definition> </policy-definitions> </routing-policy> </config></pre> <p>As a workaround, you can create one community with multiple statements through OpenConfig</p> |

| Path | Description |
|--|--|
| | <p>YANG:</p> <pre>ip community-list standard cs_std seq 5 permit 1:1 ip community-list standard cs_std seq 10 permit 1:2 ip community-list standard cs_std seq 15 permit 1:3 route-map To_LC permit 10 match community cs_std</pre> |
| /bgp-conditions/state/next-hop-in | <p>In OpenConfig YANG, the <code>next-hop-in</code> type is an IP address, but in the CLI, it is an IP prefix. While creating the <code>next-hop-in</code> through OpenConfig YANG, the IP address is converted to a "/32" mask prefix in the CLI configuration. The following is an example of <code>next-hop-in</code> in the OpenConfig YANG payload:</p> <pre><policy-definition> <name>sc0</name> <statements> <statement> <name>5</name> <conditions> <bgp-conditions xmlns="http://openconfig.net/yang/bgp-policy"> <config> <next-hop-in>2.3.4.5</next-hop-in> </config> </bgp-conditions> </conditions> </statement> </statements> </policy-definition></pre> <p>The following is an example of the same information in the CLI:</p> <pre>ip prefix-list IPV4_PFX_LIST_OPENCONFIG_sc0_5 seq 5 permit 2.3.4.5/32 route-map sc0 permit 5 match ip next-hop prefix-list IPV4_PFX_LIST_OPENCONFIG_sc0_5</pre> |
| /bgp-actions/set-community/config/method | enum "REFERENCE" is not supported |
| /bgp-actions/config/set-next-hop | enum "SELF" is not supported |
| /bgp-conditions/match-community-set/config/community-set | Get mapped only to <code>match community</code> <code><community-set>_std</code> , only a standard community is supported. Match to an expanded community set is not supported. |

| Path | Description |
|-------------------|---|
| .../match-tag-set | <p>There is a limitation in replacing <code>match-tag-set</code> because defined sets for <code>tag-sets</code> are not currently implemented.</p> <p>Currently, replacing <code>match-tag-set</code> appends the values. To replace <code>match-tag-set</code>, delete it, then create it again.</p> |

Table 25: OpenConfig Interfaces

| Path | Description |
|--|--------------------|
| <i>interfaces/interface/ethernet/switched-vlan/config/interface-mode</i> | |

| Path | Description |
|------|--|
| | <p>You cannot configure successfully configuring both trunk-mode interface and trunk VLANs simultaneously in the same OpenConfig payload. If you split the payload and configure the trunk-mode interface and then the trunk VLANs, the configuration is successful.</p> <p>On Cisco NX-OS interfaces, the default interface mode is access. To implement any trunk-related configurations, you must first change the interface mode to trunk, then configure the trunk VLAN ranges. Configure these in separate payloads.</p> <p>The following examples show the separate payloads for configuring trunk mode and VLAN ranges.</p> <p>The following example shows the payload configuring the interface to trunk mode.</p> <pre data-bbox="922 821 1484 1682"> <rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101"> <edit-config> <target> <running/> </target> <config> <interfaces xmlns="http://openconfig.net/yang/interfaces"> <interface> <name>eth1/47</name> <subinterfaces> <subinterface> <index>0</index> <config> <index>0</index> </config> </subinterface> </subinterfaces> <ethernet xmlns="http://openconfig.net/yang/interfaces/ethernet"> <switched-vlan xmlns="http://openconfig.net/yang/vlan"> <config> <interface-mode>TRUNK</interface-mode> </config> </switched-vlan> </ethernet> </interface> </interfaces> </config> </edit-config> </rpc> </pre> <p>The following shows payload configuring the VLAN ranges:</p> <pre data-bbox="922 1780 1484 1841"> <rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101"> </pre> |

| Path | Description |
|---|--|
| | <pre> <edit-config> <target> <running/> </target> <config> <interfaces xmlns="http://openconfig.net/yang/interfaces"> <interface> <name>eth1/47</name> <subinterfaces> <subinterface> <index>0</index> <config> <index>0</index> </config> </subinterface> </subinterfaces> <ethernet xmlns="http://openconfig.net/yang/interfaces/ethernet"> <switched-vlan xmlns="http://openconfig.net/yang/vlan"> <config> <native-vlan>999</native-vlan> <trunk-vlans xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" nc:operation="delete">1..4094</trunk-vlans> <trunk-vlans>401</trunk-vlans> <trunk-vlans>999</trunk-vlans> </config> </switched-vlan> </ethernet> </interface> </interfaces> </config> </edit-config> </rpc> </pre> |
| <i>interfaces/interface/ethernet/switched-vlan/config/trunk-vlans</i> | <p>Due to restriction of OpenConfig YANG, ensure that configuring VLANs must not overlap between VLANs in the payload and the VLANs configured on an interface. If there is an overlap, the configuration through OpenConfig is not successful..</p> <p>Make sure that the VLANs configured on an interface are different from the VLANs in the OpenConfig payload and ensure about appropriate about the VLAN range.</p> |
| | <p>The following for switch port, shut/no shut, MTU, and mac-address:</p> <p>You must reload ASCII, while configuring switch port, shut/no shut, MTU, and mac-address. If you reload a binary result, the configuration is lost.</p> |

Table 26: OpenConfig LACP

| Path | Description |
|---|--|
| <i>lACP/interfaces/interface/config/lACP-mode</i> | <p>OC-LACP enables configuring the port-channel mode on the port-channel interface. But, for the NXOS-CLI, the port-channel mode is configured on the member interface using channel-group mode active or passive.</p> <p>Though OC-LACP explicitly configures the port-channel mode on a port-channel interface, issuing the NX-OS show running-config command on a port-channel interface does not show the port-channel mode configuration for either empty or nonempty port-channels.</p> <p>When you add a member to the port-channel, show running interface ethernet <> shows the port-channel mode configuration as a channel-group mode active or passive.</p> <p>Note All port-channels that you create through OpenConfig must continue to be managed by OpenConfig.</p> |
| <i>lACP/interfaces/interface/config/interval</i> | <p>You can change port channel interval only when members are in <code>shut</code> state.</p> <p>The OC-LACP interval is per port-channel. The NX-OS LACP interval is per port-channel member. Due to this difference, the following behavior is expected:</p> <ul style="list-style-type: none"> • If you configure the port-channel interval through OpenConfig, all members in the port-channel get the same configuration applied to them. <p>If you configure the port-channel interval through OpenConfig and then you add a member. Ensure that you configure the interval through OpenConfig for applying configuration to the new member.</p> |
| <i>lACP/interfaces/interface/config/system-id-mac</i> | does not support <code>system-id-mac</code> per port-channel |
| LACP | <p>The following has member-state data only when a port is in <code>admin up</code> state:</p> <ul style="list-style-type: none"> • LACP • Interface • Interfaces • Member • State |

The following state Containers are implemented for OpenConfig ACL at the interface-ref level.

Table 27: OpenConfig ACL

| Path | Description |
|--|-------------|
| <i>acl/interfaces/interface/interface-ref/stateforacl/interfaces/state</i> | N/A |
| <i>acl/interfaces/interface/interface-ref/state/interface</i> | read-only |
| <i>acl/interfaces/interface/interface-ref/state/subinterface</i> | read-only |

Table 28: OpenConfig QoS

| Path | Description |
|------|---|
| QoS | <ul style="list-style-type: none"> • Queuing stats for HIG (ii) ports is not supported. • You do not see the tx-packets, or bytes, and drop-packets per unicast, multicast, or broadcast queue. The stats that display in the OC response are the sum of the ucast, mcast, and bcast queues per qos-group. • Does not support stats for a QoS policy that is applied at the VLAN level. • The ingress queue drop count that can be retrieved through OC can be displayed at the slice/port/queue level depending on the platform. |

The following system config containers are implemented for domain-name, login-banner, and motd-banner.

Table 29: OpenConfig System

| Path | Description |
|-----------------------------------|--|
| <i>system/config/domain-name</i> | System/dns-items/ prof-items/Prof-list/dom-items/name |
| <i>system/config/login-banner</i> | System/userext-items/postloginbanner-items/message |
| <i>system/config/motd-banner</i> | System/userext-items/preloginbanner-items/message |

Guidelines and Limitations for High Scale Data

Cisco NX-OS supports a new set of operational state OpenConfig path which has high scale data. The following are the guidelines and limitations:

- For optimal performance, you must retrieve data by providing the exact path. Parent-level path queries will not provide the same performance.
- The high scale paths are supported only for gNMI and not for RESTCONF or NETCONF.

- The high-scale paths do not support suppress-redundant.
- The high scale paths do not support the gNMI ON_CHANGE subscription.

| Path | Description |
|--|---|
| <i>/network-instances/network-instance/fdb/l2rib/mac-table</i> | Parent level queries for l2rib are supported at l2rib level. For example, you can query till <code>network-instances/network-instance/fdb/l2rib</code> but not at fdb level <code>network-instances/network-instance/fdb</code> . |
| <i>/interfaces/interface/routed-vlan/ipv4/neighbors/neighbor/state</i> | |
| <i>/interfaces/interface/routed-vlan/ipv6/neighbors/neighbor/state</i> | <p>For parent-level queries, the infrastructure retrieves all the keys for the list items and a request is sent to populate the rest of the data for each of these list items. This means that the infrastructure must have the same view of the tree as the back end.</p> <p>For example, if the infrastructure has a track of static entries, while the back end has static and dynamic entries. For the list, walk the infrastructure sends requests for each static entry which results in incomplete data. The paths with this limitation in the current release are:</p> <p><code>"/interfaces/interface/routed-vlan/ipv6/neighbors/neighbor/state"</code> and <code>"/interfaces/interface/routed-vlan/ipv4/neighbors/neighbor/state"</code>.</p> <p>The data contains both dynamic and static ARP and ND entries. If the exact path is provided, it contains the static entries if the parent path given.</p> |
| <i>/network-instances/network-instance/protocols/protocol/bgp/rib/sets/dynamic-entries</i> | |
| <i>network-instances/network-instance/protocols/protocol/bgp/rib/attr-sets</i> | |
| <i>/network-instances/network-instance/protocols/protocol/bgp/rib/communities</i> | |
| <i>/network-instances/network-instance/protocols/protocol/bgp/rib/ext-communities</i> | |
| <i>/network-instances/network-instance/protocols/protocol/bgp/rib/neighbors/neighbor/paths</i> | |
| <i>/network-instances/network-instance/protocols/protocol/bgp/rib/neighbors/neighbor/paths</i> | |

Configuring OpenConfig Support

To enable or disable OpenConfig support on the programmability agents such as (NETCONF, RESTCONF and gRPC).

SUMMARY STEPS

1. `configureterminal`
2. `featureopenconfig`

DETAILED STEPS**Procedure**

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | configureterminal Example: <code>switch# configure terminal</code> | Enters global configuration mode. |
| Step 2 | featureopenconfig Example: | Enters global configuration mode. <code>switch(config)# feature netconf</code> |

Upgrading or Downgrading YANG Version

YANG models are a collection of logical supported configuration or states. This model doesn't support earlier or future compatibility releases.

Cisco NX-OS doesn't support multiple YANG versions for a single release. Each NX-OS release supports a specific and authorized YANG collection. See [Yang NX Repository](#), for respective NX-OS releases.

To upgrade or downgrade NX-OS releases, there exist few modifications in the supported YANG models. You must review the YANG models and perform appropriate actions.

| Changes | Recommendation |
|----------------|---|
| Paths added | You can evaluate the way to handle the additional YANG paths |
| Paths modified | This is due to the property type changes, from integer to float. You must up |
| Paths removed | In OpenConfig, the community decides to remove paths which are not app if you use the corresponding functionality, access through other interfaces, |

RBAC for YANG

Cisco NX-OS supports write permission to YANG paths for non-admin user roles.

Guidelines and Limitations

The following are the guidelines and limitations for RBAC of YANG:

- A user with role network-admin has write access. This user can provide other edit permission to nonadmin groups.

- User role network-admin only has the read access.
- RBAC doesn't support providing root system access to nonadmin user roles.
- RBAC doesn't support providing access to paths *System/yangrbacdb-items* or *System/rbacdb-items* to nonadmin user roles. These user roles are restricted for modifying the access for their roles.
- RBAC Honors is supported in running configuration, not for candidate. Candidate configuration is supported in NETCONF and not in RESTCONF or gNMI. Cisco recommends configuring RBAC on the running configuration in a single transaction. If you want to change RBAC through candidate configuration, it is recommended to change and commit RBAC itself. Do not change RBAC configuration and regular configuration simultaneously, in such cases RBAC changes in the candidate configuration will not change.
- Make sure that you provide access to the last element in the path and its subtree. For an example, for path "x/y/z", to provide access to the user to modify the last path element and the child tree.
- If you want to provide write access to specific parent or child nodes in the same path, there is no syntax to merge rules in a single rule. You must configure it separately.

For example, the following mentioned below rule provides access to priv-9, to change description for all interfaces. But, it can change two properties of a particular SVI interface. These rules cannot be merged.

- interfaces/interface/description (priv-9)
 - interfaces/interface[id=vlan100]/type (priv-9)
 - interfaces/interface[id=vlan100]/enabled (priv-9)
- RBAC supports a maximum of 512 rules.
 - RBAC user role supports forward reference.
 - The user role can specify a role which doesn't exist.
 - The user role can specify a yang path which doesn't exist.
 - Make sure that you check the following syntax:
 - Don't add specific wildcard characters.
 - Make sure that you add the namespace string at the first element only.
 - Don't add leading '/' in the path.
 - Path with string only 'system' is not allowed.
 - RBAC has no impact on the existing limitations for NETCONF, none operation.

Configuring YANG RBAC

To configure yang path to a specific user role.

SUMMARY STEPS

1. `configure terminal`
2. `rbac yang-rule yang-pathrole`
3. `allow-writes`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | configure terminal Example: <pre>switch# configure terminal</pre> | Enters global configuration mode. |
| Step 2 | rbac yang-rule yang-pathrole Example: <pre>switch(config)# rbac yang-rule System/intf-items/aggr-items/AggrIf-list[id=po101] network-operator</pre> Example: <pre>switch(config)# rbac yang-rule openconfig-interface:interfaces/interface[name=vlan100]/desc priv-9</pre> | Configures yang path to the user role. The path can be either a Device or OpenConfig YANG path. |
| Step 3 | allow-writes Example: <pre>switch(config)# allow-writes</pre> | Enables the write permission. |

Troubleshooting YANG

You can use the following commands to verify YANG configurations. Specifically to verify the enablement of “feature OpenConfig”.

For issues related to read or write operations for specific agents, see troubleshooting guide for respective agents.

| Command | Description |
|---|--|
| <code>show running-config openconfig</code> | To check whether OpenConfig is enabled. |
| <code>show openconfig nxsdk event-history { event errors }</code> | To check the debug of the OpenConfig feature |
| <code>show telemetry yang direct-path cisco-nxos-device</code> | To display the paths which are supported. |



CHAPTER 27

NETCONF Agent

- [About the NETCONF Agent, on page 333](#)
- [Revision History, on page 334](#)
- [Guidelines and Limitations for NETCONF, on page 334](#)
- [Configuring the NETCONF Agent, on page 336](#)
- [Manage a NETCONF Session, on page 339](#)
- [NETCONF Get / Set, on page 341](#)
- [NETCONF Notifications, on page 349](#)
- [NETCONF Device YANG RPC, on page 352](#)
- [Netconf Client Examples, on page 355](#)
- [Troubleshooting the NETCONF Agent, on page 359](#)
- [About NETCONF Explicit Mode, on page 362](#)
- [Guidelines and Limitations, on page 362](#)
- [NETCONF Explicit Mode Get/Set, on page 363](#)
- [Add Configuration Using CLI, on page 367](#)

About the NETCONF Agent

The Network Configuration Protocol (NETCONF) is a network management protocol defined by [RFC 6241](#). Cisco NX-OS provides a NETCONF agent which is a client-facing interface that provides secure transport over SSH or TLS for the client requests and server responses in the form of a YANG model, encoded in XML.

NETCONF defines configuration datastores and a set of Create, Read, Update, and Delete (CRUD) operations that allow manipulation and query on these datastores. Three datastores are supported on NX-OS: running, startup, and candidate. Here's a brief description of the operations that are supported:

Table 30: Supported Operations

| Operation | Description |
|---------------|--|
| get | Retrieve running configuration and operational state |
| get-config | Retrieve configuration from specified datastore |
| edit-config | Load specified configuration to the specified target datastore |
| close-session | Request graceful termination of a session |
| kill-session | Force the termination of a session |

| Operation | Description |
|---------------------|---|
| copy-config | Create or replace datastore with the contents of another datastore |
| lock | Lock the datastore |
| unlock | Unlock the datastore |
| validate | Validate the contents of the specified configuration |
| commit | Commit the candidate configuration as the new current running configuration |
| cancel-commit | Cancel an ongoing confirmed commit |
| discard-changes | Revert the candidate configuration to the current running configuration |
| create-subscription | Subscribe to event notification stream |

Revision History

| Release | Description |
|---------|--|
| 9.3(1) | Beginning with NX-OS 9.3(1), NETCONF get and get-config requests from the NETCONF client to the switch must contain an explicit namespace and filter |
| 9.3(3) | RFC 6241 Compliant |
| 9.3(5) | Support for OpenConfig models in NETCONF notifications |
| 10.2(1) | Support for RPC operations: checkpoint, rollback, install, import ca certificate, module reload, individual module reload, and copy file |
| 10.3(1) | Support Cisco Nexus 9800 platform switches |
| 10.3(2) | Support RBAC for edit-config |
| 10.3(3) | Support Netconf transport over TLS |

Guidelines and Limitations for NETCONF

The NETCONF Agent has the following guideline and limitation:

Netconf YANG Operations

- NETCONF is [RFC 6241](#) compliant with the following exceptions:
 - Sibling content match nodes are logically combined in an "OR" expression instead of an "AND" expression. (Section 6.2.5)
 - Once a candidate datastore has been edited, the running configuration for the same property must not be edited.

- Cisco NX-OS NETCONF does not fully support enhanced Role-Based Access Control (RBAC) as specified in [RFC 6536](#). Instead, it allows to grant per path EDIT permission to non “network-admin” roles.
- NETCONF `get` and `get-config` requests from the NETCONF client to the switch must contain an explicit namespace and filter. This requirement affects requests to the OpenConfig YANG and NETCONF Device models. If you see an error response that is like the following, the requests are not carrying a namespace:

Request without namespace and filter is an unsupported operation

The following example shows a proper request and response with the correct behavior in NX-OS release 9.3(1) and later.

```
<get>
<filter>
<System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device"> </System>
</filter>
</get>
```

- The `<edit-config>` "replace" operation sometimes might not work due to run-time default values and behaviors that are implemented by the affected system component. Therefore, it's better to base the configuration to replace on the configuration obtained through the `<get-config>` query instead of the NX-API Developer Sandbox.
- In a single Get request, the number of objects that are supported is 250,000. If you see the following error, it means that the data requested is more than 250,000. To avoid this error, send requests with filters querying for a narrower scope of data.

```
too many objects(459134 > 250000) to query the entire device model
```

Native Device RPC

- The Cisco NX-OS supports Device YANG RPC operations: checkpoint, rollback, install, import ca certificate, module reload, individual module reloads, and copy file are supported.

Subscription / Notification

- The Cisco NX-OS supports a maximum of five subscriptions, one subscription per client session.
- Per [RFC 5277](#), autonomous notifications support NETCONF, SYSLOG, and SNMP streams for event sources. In this release, Cisco NX-OS supports NETCONF streams only.
- Cisco NX-OS does not support the Replay option for subscriptions. Because Start Time and Stop Time options are part of Replay, they are not supported.
- For a stream subscription and filtering, support is only for subtree filtering. XPath filtering is not supported.
- When the Cisco NX-OS NETCONF Agent is operating under a heavy load, it is possible that some event notifications can get dropped.
- Beginning with Cisco NX-OS Release 10.4(3)F, NETCONF is supported on 92348GC-X.

Configuring the NETCONF Agent

Configuring the NETCONF Agent Over SSH

This procedure describes how to enable and configure the NETCONF Agent over SSH.

Before you begin

Before communicating with the switch using NETCONF, the NETCONF Agent must be enabled. The NETCONF Agent is enabled or disabled by entering the **[no] feature netconf** command.

SUMMARY STEPS

1. **configure terminal**
2. **feature netconf**
3. (Optional) **netconf idle-timeout *it-num***

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | configure terminal Example: switch# configure terminal | Enters global configuration mode. |
| Step 2 | feature netconf Example: switch(config)# feature netconf | Enables NETCONF services. |
| Step 3 | (Optional) netconf idle-timeout <i>it-num</i> Example: switch(config)# netconf idle-timeout 5 | Specifies the timeout in minutes after which idle client sessions are disconnected. The range of <i>it-num</i> is 0-1440 minutes. The default timeout is 5 minutes. A value of 0 disables timeout. |

Configuring the NETCONF Agent Over TLS

This procedure describes how to enable and configure the NETCONF Agent over TLS. This enables an additional TLS transport channel to run side-by-side together with the SSH transport.

Before you begin

Prepare and sign the required certificate files for both the server and client-side authentication.



Note This is not specific to Netconf, so you can re-use the existing trustpoint files.

SUMMARY STEPS

1. **configure terminal**
2. (Optional) **crypto ca trustpoint** *server-trustpoint*
3. (Optional) **crypto ca trustpoint** *server-trustpoint* **pkcs12 bootflash:** *server-ca-file pkcs-password*
4. (Optional) **crypto ca trustpoint** *client-root-trustpoint*
5. (Optional) **rsa keypair** *client-key*
6. (Optional) **crypto ca authenticate** *client-root-trustpoint*
7. **netconf tls certificate** *server-trustpoint*
8. **netconf tls client root certificate** *client-root-trustpoint*
9. **netconf tls port** *port*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | configure terminal Example: <code>switch# configure terminal</code> | Enters global configuration mode. |
| Step 2 | (Optional) crypto ca trustpoint <i>server-trustpoint</i> Example: <code>switch(config)# crypto ca trustpoint tls_server_trustpoint</code> | Creates a trustpoint for server authentication Note The Steps 2-6 is optional if you already have an usable server and client trustpoints. |
| Step 3 | (Optional) crypto ca trustpoint <i>server-trustpoint</i> pkcs12 bootflash: <i>server-ca-file pkcs-password</i> Example: <code>switch(config)# crypto ca import tls_server_trustpoint pkcs12 bootflash:server.pfx test</code> | Imports the server pkcs12 file to the trustpoint. |
| Step 4 | (Optional) crypto ca trustpoint <i>client-root-trustpoint</i> Example: <code>switch(config)# crypto ca trustpoint tls_client_trustpoint</code> | Creates a trustpoint for client authentication. |
| Step 5 | (Optional) rsa keypair <i>client-key</i> Example: <code>switch(config)# rsa keypair client-key</code> | Generates a rsa key pair for the client trustpoint. |

| | Command or Action | Purpose |
|---------------|---|---|
| Step 6 | (Optional) crypto ca authenticate <i>client-root-trustpoint</i> Example: switch(config)# crypto ca authenticate tls_client_trustpoint | Imports the client certificate. This step requires manual copy paste. Please follow the instruction<need info on which instruction?.> |
| Step 7 | netconf tls certificate <i>server-trustpoint</i> Example: switch(config)# netconf tls certificate tls_server_trustpoint | Enters the trustpoint to host the server identity certificate. |
| Step 8 | netconf tls client root certificate <i>client-root-trustpoint</i> Example: switch(config)# netconf tls client root certificate tls_client_trustpoint | Enters the trustpoint to host the client CA root certificate. |
| Step 9 | netconf tls port <i>port</i> Example: switch(config)# netconf tls port 7000 | Specify the TLS port. The default port is 6513. |

Generating Key/Certificate Example

The following is an example for generating a self-signed key/certificate in the switch bash shell.

Note that this is only for experimental usage. For more information on generating identity certificates, see the **Installing Identity Certificates** section of the *Cisco Nexus 9000 Series NX-OS Security Configuration Guide*.



Note This task is an example of how a certificate can be generated on a NX-OS switch. You can also generate a certificate in any Linux environment. In a production environment, the user should consider using a CA signed certificate.

1. Generate the self-signed key and pem files.

```
switch# run bash sudo su
bash-4.3# openssl req -x509 -newkey rsa:2048 -keyout self_sign2048.key -out
self_sign2048.pem -days 365 -nodes
```

2. After generating the key and pem files, you must bundle the key and pem files for use in the trustpoint CA Association.

```
switch# run bash sudo su bash-4.3# cd /bootflash/
bash-4.3# openssl pkcs12 -export -out self_sign2048.pfx -inkey self_sign2048.key -in
self_sign2048.pem -certfile self_sign2048.pem -password pass:Ciscolab123!
bash-4.3# exit
```

3. Set up the trustpoint CA Association by inputting in the pkcs12 bundle into the trustpoint.

```
switch(config)# crypto ca trustpoint mytrustpoint
switch(config-trustpoint)# crypto ca import mytrustpoint pkcs12 self_sign2048.pfx
Ciscolab123!
```

4. Verify the setup.

```
Verify the setup.
switch(config)# show crypto ca certificates Trustpoint: mytrustpoint certificate:
subject= /C=US/O=Cisco Systems, Inc./OU=CSG/L=San Jose/ST=CA/street=3700 Cisco
Way/postalCode=95134/CN=ems.cisco.com/serialNumber=FGE18420K0R
issuer= /C=US/O=Cisco Systems, Inc./OU=CSG/L=San
Jose/ST=CA/street=3700 Cisco
Way/postalCode=95134/CN=ems.cisco.com/serialNumber=FGE18420K0R serial=0413
notBefore=Nov 5 16:48:58 2015 GMT notAfter=Nov 5 16:48:58 2035 GMT
SHA1 Fingerprint=2E:99:2C:CE:2F:C3:B4:EC:C7:E2:52:3A:19:A2:10:D0:54:CA:79:3E
purposes: sslserver sslclient
CA certificate 0: subject= /C=US/O=Cisco Systems, Inc./OU=CSG/L=San
Jose/ST=CA/street=3700 Cisco
Way/postalCode=95134/CN=ems.cisco.com/serialNumber=FGE18420K0R
issuer= /C=US/O=Cisco Systems, Inc./OU=CSG/L=San
Jose/ST=CA/street=3700 Cisco
Way/postalCode=95134/CN=ems.cisco.com/serialNumber=FGE18420K0R serial=0413
notBefore=Nov 5 16:48:58 2015 GMT notAfter=Nov 5 16:48:58 2035 GMT
SHA1 Fingerprint=2E:99:2C:CE:2F:C3:B4:EC:C7:E2:52:3A:19:A2:10:D0:54:CA:79:3E
purposes: sslserver sslclient
```

Manage a NETCONF Session

NETCONF is a connection-oriented protocol requiring a persistent connection between client and server. The NETCONF agent on the switch listens at port 830 of the management port IP address.

Start Session

The client can establish a connection with the NETCONF subsystem over SSH. When a client establishes a session with the NETCONF agent, the server sends a <hello> message to the client. The client likewise must

send its `<hello>` message to the server. The `<hello>` messages are exchanged simultaneously as soon as the connection is open. Each `<hello>` message contains a list of the sending peer's protocol version and capabilities. These messages are used to determine protocol compatibility and capabilities. Both NETCONF client and server must verify that a common protocol version is advertised by the other peer's `<hello>` message. Also, the server's `<hello>` message must include a `<session-id>` whereas the client's `<hello>` message must not include the `<session-id>`.

The following shows an example session establishment using the `ssh` command. The first `<hello>` message is received from the server and the second message is sent from the client. The server's `<hello>` message shows the protocol version "urn:ietf:params:netconf:base:1.1" and includes the supported data models. This below example might not reflect the current Cisco NX-OS release.



Note The server's `<hello>` message has a `<session-id>`, but the client's message does not.

```
client-host % ssh admin@172.19.193.166 -p 830 -s netconf User Access Verification Password:
<?xml version="1.0" encoding="UTF-8"?>
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<capabilities>
<capability>urn:ietf:params:netconf:base:1.0</capability>
<capability>urn:ietf:params:netconf:base:1.1</capability>
<capability>urn:ietf:params:netconf:capability:writable-running:1.0</capability>
<capability>urn:ietf:params:netconf:capability:rollback-on-error:1.0</capability>
<capability>urn:ietf:params:netconf:capability:candidate:1.0</capability>
<capability>urn:ietf:params:netconf:capability:validate:1.1</capability>
<capability>urn:ietf:params:netconf:capability:confirmed-commit:1.1</capability>
<capability>urn:ietf:params:netconf:capability:notification:1.0</capability>
<capability>urn:ietf:params:netconf:capability:interleave:1.0</capability>
<capability>urn:ietf:params:netconf:capability:with-defaults:1.0?basic-mode=report-all</capability>
<capability>http://cisco.com/ns/yang/cisco-nx-os-device?revision=2020-04-20&module=Cisco-NX-OS-device</capability>
<capability>http://openconfig.net/yang/acl?revision=2019-11-27&module=openconfig-acl&path=deviations-cisco-nx-openconfig-acl-deviations</capability>
<capability>http://openconfig.net/yang/bfd?revision=2019-10-25&module=openconfig-bfd&path=deviations-cisco-nx-openconfig-bfd-deviations</capability>
...
</capabilities>
<session-id>1286775422</session-id>
</hello>
]]>]]>

<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<capabilities>
<capability>urn:ietf:params:netconf:base:1.1</capability>
</capabilities>
</hello>
]]>]]>
```

Using NETCONF with the raw `ssh` command is not convenient and is prone to error, as the complexity for message framing can be seen from RFC 6242 (Using the NETCONF Protocol over SSH). The above `ssh` command is used for illustration purposes only. There exist various clients written for NETCONF which are recommended over the `ssh` command. The `ncclient` is one such example. Please refer to the *Examples* section.

Terminate Session

NETCONF supports two operations for terminating a session, namely, `<close-session>` and `<kill-session>`. When the server receives a `<close-session>` request, it gracefully terminates the session by releasing any locks and resources associated with the session and closing the connection with the client.

The following is an example of the `<close-session>` request and response for success:

```
<rpc message-id="1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<close-session/>
</rpc>
<rpc-reply message-id="1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>
```

The `<kill-session>` request forces the termination of another session and requires `<session-id>` in the request message. Upon receiving the `<kill-session>` request, the server terminates current operations, releases locks and resources, and closes the connection associated with the specified session ID.

The following is an example of the `<kill-session>` request and response for success:

```
<rpc message-id="2" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<kill-session>
<session-id>296324181</session-id>
</kill-session>
</rpc>
<rpc-reply message-id="2" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <ok/>
</rpc-reply>
```

Besides the `<close-session>` and `<kill-session>` requests, a session is terminated automatically if the client does not send any request for a certain length of time. The default is five minutes. See [Configuring the NETCONF Agent](#) for configuring the idle timeout.

NETCONF Get / Set

This section describes supported NETCONF operations to manipulate and query datastores. The client can send RPC messages for these operations after establishing a session with the NETCONF agent. Basic usage explanations are given, and [RFC 6242](#) can be referred to for thorough details about these operations.

`<get-config>`

This operation retrieves configuration data from a specified datastore. The supported parameters are `<source>` and `<filter>`. The `<source>` specifies the datastore being queried such as `<running/>`, which holds the currently active configuration. The `<filter>` specifies the portions of the specified datastore to retrieve.

The following are examples of `<get-config>` request and response messages.

- Retrieve the entire `<System>` subtree:

```
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source>
      <running/>
    </source>
    <filter>
      <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device"/>
    </filter>
  </get-config>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101">
  <data>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      ...
    </System>
  </data>
</rpc-reply>
```

- Retrieve a specific list item:

```

<rpc message-id="102" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get-config>
    <source>
      <running/>
    </source>
  </filter>
  <filter>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <bgp-items>
        <inst-items>
          <dom-items>
            <Dom-list>
              <name>default</name>
            </Dom-list>
          </dom-items>
        </inst-items>
      </bgp-items>
    </System>
  </filter>
</get-config>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="102">
  <data>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <bgp-items>
        <inst-items>
          <dom-items>
            <Dom-list> <name>default</name>
            ...
            <rtctrl-items>
              <enforceFirstAs>enabled</enforceFirstAs>
              <fibAccelerate>disabled</fibAccelerate>
              <logNeighborChanges>enabled</logNeighborChanges>
              <supprRt>enabled</supprRt>
            </rtctrl-items>
            <rtrId>1.2.3.4</rtrId>
          </Dom-list>
        </dom-items>
      </inst-items>
    </bgp-items>
  </System>
</data>
</rpc-reply>

```

<edit-config>

This operation writes specified configuration to the target datastore.

The <target> parameter specifies the datastore being edited, such as <running/> or <candidate/>. The <running/> datastore is manipulated to immediately change the switch config while the candidate datastore can be manipulated without impacting the running datastore until its changes are committed. For more information, see the <commit> section.

The <config> parameter specifies the modeled data to be written to the target datastore. The intended model is specified by the “xmlns” attribute. Any number of elements in the <config> subtree may contain an “operation” attribute. The operation of an element is inherited by its descendent elements until it’s overridden by a new “operation” attribute. The supported operations are “merge”, “replace”, “create”, “delete”, and “remove”. If the “operation” attribute is not specified, the “merge” operation is assumed as default; the default operation can be overridden by the optional <default-operation> parameter, which could be “merge”, “replace” or “none”.

- The “merge” operation is different from “create” in that no error is returned if the config data already exists.
- The “remove” operation is different from “delete” in that no error is returned if the config data does not exist.

The following are examples of `<edit-config>` request and response messages.

- Create a port-channel named "po5" with MTU 9216 and the description in the running configuration:

```
<rpc message-id="103" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <running/>
    </target>
    <config
      xmlns:xc="urn:ietf:params:xml:ns:netconf:base:1.0">
      <System
        xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
        <intf-items>
          <aggr-items>
            <AggrIf-list xc:operation="create">
              <id>po5</id>
              <mtu>9216</mtu>
              <descr>port-channel 5</descr>
            </AggrIf-list>
          </aggr-items>
        </intf-items>
      </System>
    </config>
  </edit-config>
</rpc>

<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="103">
  <ok/>
</rpc-reply>
```

- Replace all configurations of a port-channel with new configurations:

```
<rpc message-id="104" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <running/>
    </target>
    <config
      xmlns:xc="urn:ietf:params:xml:ns:netconf:base:1.0">
      <System
        xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
        <intf-items>
          <aggr-items>
            <AggrIf-list xc:operation="replace">
              <id>po5</id>
              <mtu>1500</mtu>
              <adminSt>down</adminSt>
            </AggrIf-list>
          </aggr-items>
        </intf-items>
      </System>
    </config>
  </edit-config>
</rpc>

<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="104">
```

```
<ok/>
</rpc-reply>
```

- Delete a port-channel:

```
<rpc message-id="105" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <running/>
    </target>
    <config xmlns:xc="urn:ietf:params:xml:ns:netconf:base:1.0">
      <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
        <intf-items>
          <aggr-items>
            <AggrIf-list xc:operation="delete"> <id>po5</id>
          </AggrIf-list>
        </aggr-items>
      </intf-items>
    </System>
  </config>
</edit-config>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="105">
  <ok/>
</rpc-reply>
```

<copy-config>

This operation replaces the target configuration datastore with the contents of source configuration datastore. The parameters for source datastore and target datastore are `<source>` and `<target>`, respectively.

The following are examples of `<copy-config>` request and response messages.

- Copy from running configuration to startup configuration:

```
<rpc message-id="106" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <copy-config>
    <target>
      <startup/>
    </target>
    <source>
      <running/>
    </source>
  </copy-config>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="106">
  <ok/>
</rpc-reply>
```

- Copy from running configuration to candidate configuration:

```
<rpc message-id="107" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <copy-config>
    <target>
      <candidate/>
    </target>
    <source>
      <running/>
    </source>
  </copy-config>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="107">
  <ok/>
</rpc-reply>
```

<lock>

The <lock> operation allows a client to lock the configuration datastore, preventing other clients from locking or modifying the datastore. The lock that is held by the client is released with either the <unlock> operation or termination of a session. The <target> parameter is used to specify the datastore to be locked.

The following are examples of <lock> request and response messages.

- A successful acquisition of a lock:

```
<rpc message-id="108" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <lock>
    <target>
      <running/>
    </target>
  </lock>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="108">
  <ok/>
</rpc-reply>
```

- A failed attempt to acquire a lock already in use by another session:

```
<rpc message-id="109" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <lock>
    <target>
      <candidate/>
    </target>
  </lock>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="109">
  <rpc-error>
    <error-type>protocol</error-type>
    <error-tag>lock-denied</error-tag>
    <error-severity>error</error-severity>
    <error-message xml:lang="en">Lock failed, lock is already held</error-message>
    <error-info>
      <session-id>1553704357</session-id>
    </error-info>
  </rpc-error>
</rpc-reply>
```

<unlock>

The <unlock> operation releases a configuration lock, obtained with the <lock> operation. Only the same session that issued the <lock> operation can use the <unlock> operation. The <target> parameter is used to specify the datastore to be unlocked.

The following is an example of <unlock> request and response messages.

- Unlock

```
<rpc message-id="110" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <unlock>
    <target>
      <candidate/>
    </target>
  </unlock>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="110">
  <ok/>
</rpc-reply>
```

<get>

The <get> operation retrieves running configuration and operational state data. The supported parameter is <filter>. The <filter> specifies the portions of the running configuration operational state data to retrieve.

The following is an example of <get> request and response messages.

- Retrieve running configuration and operational state data of a list item:

```
<rpc message-id="111" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter>
      <System
        xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
        <bgp-items>
          <inst-items>
            <dom-items>
              <Dom-list>
                <name>default</name>
              </Dom-list>
            </dom-items>
          </inst-items>
        </bgp-items>
      </System>
    </filter>
  </get>
</rpc>

<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="111">
  <data>
    <System
      xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <bgp-items>
      <inst-items>
        <dom-items>
          <Dom-list>
            <name>default</name>
            <always>disabled</always>
            <bestPathIntvl>300</bestPathIntvl>
            <clusterId>120</clusterId>
            <firstPeerUpTs>2020-04-20T16:19:03.784+00:00</firstPeerUpTs>
            <holdIntvl>180</holdIntvl>
            <id>1</id>
            <kaIntvl>60</kaIntvl>
            <mode>fabric</mode>
            <numEstPeers>0</numEstPeers>
            <numPeers>0</numPeers>
            <numPeersPending>0</numPeersPending>
            <operRtrId>1.2.3.4</operRtrId>
            <operSt>up</operSt>
            <pfxPeerTimeout>90</pfxPeerTimeout>
            <pfxPeerWaitTime>90</pfxPeerWaitTime>
            <reConnIntvl>60</reConnIntvl>
            <rtrId>1.2.3.4</rtrId>
            <vnid>0</vnid> ...

          </Dom-list>
        </dom-items>
      </inst-items>
    </bgp-items>
  </System>
</data>
</rpc-reply>
```


<validate>

This operation validates the configuration contents of the candidate datastore. It is useful for validating the configuration changes made on the candidate datastore before committing them to the running datastore. The `<source>` parameter supports `<candidate/>`.

The following is an example of `<validate>` request and response messages.

- Validate the contents of the candidate datastore:

```
<rpc message-id="112" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <validate>
    <source>
      <candidate/>
    </source>
  </validate>
</rpc>

<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="112">
  <ok/>
</rpc-reply>
```

<commit>

This operation commits the candidate configuration to the running configuration. The operation without any parameter is considered final and cannot be reverted.

If `<commit>` is issued with the `<confirmed/>` parameter, it is considered a confirmed commit, and commit is finalized only if it is followed by another `<commit>` operation without the `<confirmed/>` parameter. That is, the confirming commit. The confirmed commit allows two parameters: `<confirm-timeout>` and `<persist>`. The `<confirm-timeout>` is the period in seconds before the confirmed commit is reverted, restoring the running configuration to its state before the confirmed commit was issued, unless the confirming commit is issued before, or the timeout is reset by another confirmed commit. If the `<confirm-timeout>` is not specified, the default timeout is 600 seconds. Also, the confirmed commit is reverted if the session is terminated. The `<persist>` parameter makes the confirmed commit to persist even if the session is terminated. The value of the `<persist>` parameter is used to identify the confirmed commit from any session and must be used as the value of the `<persist-id>` parameter of subsequent confirmed commit or confirming commit.

The following are examples of `<commit>` request and response messages.

- Commit the contents of the candidate datastore:

```
<rpc message-id="113" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <commit/>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="113">
  <ok/>
</rpc-reply>
```

- Confirmed commit with the timeout:

```
<rpc message-id="114" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <commit>
    <confirmed/>
    <confirm-timeout>120</confirm-timeout>
  </commit>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="114">
  <ok/>
</rpc-reply>
```

- Start a persistent confirmed commit and then confirm the persistent confirmed commit:

```

<rpc message-id="115" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <commit>
    <confirmed/>
    <persist>ID1234</persist>
  </commit>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="115">
  <ok/>
</rpc-reply>
<!-- confirm the persistent confirmed-commit, from the same session or another session
-->
<rpc message-id="116" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <commit>
    <persist-id>ID1234</persist-id>
  </commit>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="116">
  <ok/>
</rpc-reply>

```

<cancel-commit>

This operation cancels an ongoing confirmed commit. If a confirmed commit from a different session needs to be canceled, the `<persist-id>` parameter must be used with the same value that was given in the `<persist>` parameter of the confirmed commit.

- Cancel the confirmed commit from the same sessions:

```

<rpc message-id="117" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <cancel-commit/>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="117">
  <ok/>
</rpc-reply>

```

<discard-changes>

This operation discards any uncommitted changes that are made on the candidate configuration by resetting back to the content of the running configuration. No parameter is required.

The following is an example of `<discard-changes>` request and response messages.

- Discard the changes made on the candidate datastore:

```

<rpc message-id="118" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <discard-changes/>
</rpc>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="118">
  <ok/>
</rpc-reply>

```

NETCONF Notifications

About NETCONF Notifications

NETCONF notification is a mechanism where a NETCONF client can subscribe to system events and then receive notifications to these events from a NETCONF agent. These features are defined in [RFC 5277](#). This is an optional capability that is advertised in the NETCONF hello message.

A NETCONF client can subscribe for notifications using Device YANG or OpenConfig models.

With this support, any NETCONF client can:

- **Subscribe to event notifications**

Each subscription is a one-time request over a session from a NETCONF client. The Cisco NX-OS NETCONF agent responds, and the subscription is active until the session is explicitly closed by the NETCONF client. The subscription can also be closed by an administrative action, such as a switch restart or disabling NETCONF feature on the switch. The subscription is active if the underlying NETCONF session is active. The events that are generated for these subscribed filters are sent as notifications to the client. Clients can subscribe to notifications for config or operational change events. For example, port state change, fan speed change, process memory change, feature enabled, to name a few.

- **Receive event notifications**

An event notification is a well-formed XML document that contains information about the configuration or operational events on the switch. The NETCONF client can send filtering criteria in the subscription request to specify a subset of events instead of all events.

- **Interleave event notifications with other operations**

The Cisco NX-OS NETCONF agent can receive, process, and respond to NETCONF requests on a session with an active notification subscription.

Capabilities Exchange

During the NETCONF handshake, the Cisco NX-OS NETCONF server sends the <capabilities> element to the connecting NETCONF clients to indicate what requests that the server can process. As part of the exchange, the server includes the following identifiers, which inform the client that the Cisco NX-OS NETCONF server supports both notifications and interleave.

Capability identifier for notification:

```
urn:ietf:params:netconf:capability:notification:1.0
```

Capability identifier for interleave:

```
urn:ietf:params:netconf:capability:interleave:1.0
```

Event Stream Discovery

The client can discover the Cisco NX-OS NETCONF server's supported streams by using a NETCONF `<get>` operation for all available `<streams>`. Cisco NX-OS supports the NETCONF stream only. Discovering event streams occurs through a request and reply sequence.

Request to retrieve available streams:

Any NETCONF client can send a NETCONF `<get>` request with filter for `<streams>` to identify all supported streams.

The following example shows the payload of a client request message:

```
<rpc message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <get>
    <filter type="subtree">
      <netconf xmlns="urn:ietf:params:xml:ns:netmod:notification"> <streams/>
    </netconf>
    </filter>
  </get>
</rpc>
```

Reply:

The Cisco NX-OS NETCONF server replies with all the event streams that are available and to which the client can subscribe. Cisco NX-OS supports the NETCONF stream only.

```
<rpc-reply message-id="101" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <netconf xmlns="urn:ietf:params:xml:ns:netmod:notification"> <streams>
      <stream>
        <name>NETCONF</name>
        <description>default NETCONF event stream </description> </stream>
      </streams>
    </netconf>
  </data>
</rpc-reply>
```

Create Subscriptions

NETCONF clients can create subscriptions for events on the switch through an RPC with a `<create-subscription>` protocol operation. When the Cisco NX-OS NETCONF server responds with the `<ok/>` element, the subscription is active.

Unlike synchronous Get and Set operations, a subscription is a persistent, asynchronous operation. The subscription stays active until the client explicitly closes the subscription, or the session goes offline. For example, by a switch restart.

- If a client subscribes to event notifications, but it goes offline, the server terminates the subscription and closes the session.
- If a subscription is closed, the NETCONF client must reconnect and create the subscription again to receive all event notifications.

The server does not initiate subscriptions, so the user must write client programs that contain the send the `<create-subscription>` operation.

The following is an example for `<create-subscription>` sent by a NETCONF client:

```
<create-subscription xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
  <stream>NETCONF</stream>
  <filter
    xmlns:ns1="urn:ietf:params:xml:ns:netconf:base:1.0" type="subtree">
    <System
      xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <intf-items>
        <phys-items>
          <PhysIf-list>
            <id>eth1/54/1</id>
            <phys-items>
              <operSt/>
            </phys-items>
          </PhysIf-list>
        </phys-items>
      </intf-items>
    </System>
  </filter>
</create-subscription>
```

The `<create-subscription>` operation supports the following options:

- `<stream>`, which specifies which stream of events the client wants to subscribe to. If you specify no stream, by default, events in the NETCONF stream are sent to the client.
- `<filter>`, which enables filtering the events to provide a subset of events carried on the stream.

The Cisco NX-OS NETCONF server responds back with an `<ok>` message if the server can create the subscription successfully.

The following is a sample successful response received in the client for the `<create-subscription>`:

```
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
  message-id="urn:uuid:6ff0bda6-d3f1-4288-9a7e-0f30581e4bab">
  <ok/>
</rpc-reply>
```



Note Subscriptions with Replay are not supported, so that “Start Time” and “Stop Time” options cannot be used.

Receive Notifications

When the NETCONF client has successfully created a subscription, the Cisco NX-OS NETCONF server begins sending relevant event notifications, for any events in the switch matching the requested filter. The event notification is its own XML-formatted document that contains the notification element.

The following is a sample notification when an interface “eth1/54/1” went down. In this case, the client has subscribed to interface `operSt` with the example from the above example.

```
<?xml version="1.0" encoding="UTF-8"?>
<notification
  xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
  <eventTime>2020-05-05T10:22:52.260+00:00</eventTime>
  <operation>modified</operation>
  <event>
    <System
      xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <intf-items>
        <phys-items>
          <PhysIf-list>
```

```

    <id>eth1/54/1</id>
    <phys-items>
      <operSt>down</operSt>
    </phys-items>
  </PhysIf-list>
</phys-items>
</intf-items>
</System>
</event>
</notification>

```

The `<notification>` messages contain the following fields:

- `<eventTime>`, the date and timestamp of when the event occurred.
- `<operation>`, the type of event on the model node.
- `<event>`, the model data to which the client is subscribed.

Terminate Subscriptions

Subscriptions are terminated when the NETCONF client sends specific request to the Cisco NX-OS NETCONF server in any of the following ways:

- Closing the subscription session, which occurs when the `<close-session>` operation is sent to the NETCONF Server for a specific subscription session.
- Terminating the NETCONF session, which occurs when the `<kill-session>` operation is sent to the NETCONF server.

Every subscription is tied to one NETCONF session. It is a one-to-one relationship.

NETCONF Device YANG RPC

About Model Driven RPC in NETCONF

In addition to manipulate the switch configuration, YANG offers the extensibility to trigger specific actions in the switch via YANG model, which is equivalent to invoke EXEC mode commands through CLI. Cisco NX-OS has defined the following Native Device RPC.

| Operation | Device YANG RPC | CLI |
|------------|-----------------|--|
| Checkpoint | checkpoint | checkpoint <i><name></i> checkpoint <i><file></i> |
| Rollback | rollback | rollback running-config checkpoint <i><name></i> rollback running-config checkpoint <i><file></i> |

| Operation | Device YANG RPC | CLI |
|--------------------------------|---|---|
| Install | install_all_nxos install_add install_activate install_deactivate install_commit install_remove | install all nxos <image> install {add activate deactivate commit remove} <rpm> |
| Import Crypto Certificate | import_ca_certificate | crypto ca import <trustpoint> pkcs12 <file> <passphrase> |
| Switch Reload or Module Reload | reload | reload [timer <seconds>] reload module <module number> |
| Copy File | copy | copy <source> <destination> |

Native Device RPC Examples

- **Creating checkpoint using filename option**

```
<rpc message-id="checkpoint-3" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <checkpoint xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <file>bootflash:my_checkpoint2</file>
  </checkpoint>
</rpc>
```

- **Creating checkpoint using checkpoint name, description**

```
<rpc message-id="checkpoint-1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <checkpoint xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <action>create</action>
    <name>my_checkpoint1</name>
    <description>test checkpoint one</description> </checkpoint>
  </rpc>
```

- **Deleting checkpoint using checkpoint name**

```
<rpc message-id="delatecheckpoint-1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <checkpoint xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <action>delete</action>
    <name>my_checkpoint1</name>
  </checkpoint>
</rpc>
```

- **Rollback**



Note The following options tags can be used as atomic, stop-at-first-failure, best-effort.

```
<rpc message-id="rollback-cfg-option1"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rollback
    xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
```

```

    <name>my_checkpoint1</name>
    <option>atomic</option>
  </rollback>
</rpc>

```

• Rollback using file option

```

<rpc message-id="rollback-cfg1"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rollback
    xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <file>bootflash:my_checkpoint2</file>
  </rollback>
</rpc>

```

• Copy file

Copy any file from remote server to switch storage (example: bootflash)

```

<rpc message-id="copy-file-1"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <copy
    xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <source>tftp://172.27.xxx.xxx//file_location?/tls1-server.pfx</source>
    <destination>bootflash:</destination>
    <vrf>management</vrf>
  </copy>
</rpc>

```

• Import CA Certificate

Pre-requisite: my_truspoint should be already created on the switch:

```

<rpc message-id="import_ca_certificate-1"
  xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <import_ca_certificate
    xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <trustpoint>my_trustpoint</trustpoint>
    <pkcs12>tls1-server.pfx</pkcs12>
    <passphrase>xxxxxxx</passphrase>
  </import_ca_certificate>
</rpc>

```

• Install RPM package EXEC RPC commands

Install <add>

```

<rpc message-id="install-add-1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <install_add xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <add>rpm_packagenamehere_from_bootflash</add>
  </install_add>
</rpc>

```

Install <activate>

```

<rpc message-id="install-activate-1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <install_activate xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <activate> rpm_packagenamehere_from_bootflash</activate>
  </install_activate>
</rpc>

```

• Install <deactivate>

```

<rpc message-id="install-deactivate-1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <install_deactivate xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <deactivate>rpm_packagenamehere_from_bootflash </deactivate> </install_deactivate>
  </rpc>

```


Install <remove>

```
<rpc message-id="rpc-install_remove-1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<install_remove xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
<remove>rpm_packagenamehere_from_bootflash </remove>
</install_remove>
</rpc>
```

Install all nx-os image

```
<rpc message-id="rpc-install_all_nxos-1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<install_all_nxos xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
<nxos>nxos.image.bin.upg</nxos>
</install_all_nxos>
</rpc>
```

Reload module number

```
<rpc message-id="reload-module-pyld1" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <reload xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <module>29</module>
  </reload>
</rpc>
```

Reload

Note When Client requests or sends the following RPC, the exec command executes switch reload and further Netconf client does not receive <ok> response.

```
<rpc message-id="563" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<reload xmlns="http://cisco.com/ns/yang/cisco-nx-os-device"/>
</rpc>
```

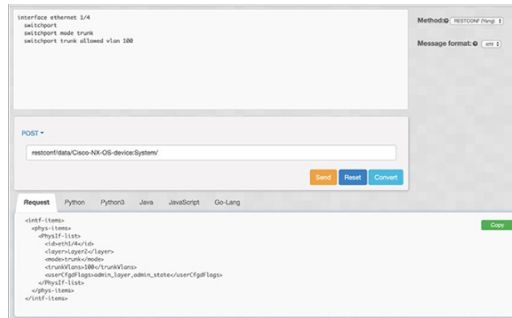
Netconf Client Examples

Using the Sandbox to Generate the NETCONF Payload

Refer to NXAPI Developer Sandbox section to enable it.

To generate a payload for NETCONF, change the method to RESTCONF (Yang) and message format to XML. Enter the command you need to convert in the text window, click **Convert** and the equivalent payload is displayed in the **Request** text box:

Figure 3: NX-OS Developer Sandbox



Connecting Cisco NX-OS with the ncclient



Note There exist various NETCONF clients. All examples in this section use the particular ncclient python library

The ncclient is a Python library for NETCONF clients. The following is an example of how to establish a connection to Cisco NX-OS from the ncclient Manager API:

```
device = {
    "address": "10.10.10.10",
    "netconf_port": 830, "username": "admin",
    "password": "cisco"
}
with manager.connect(host = device["address"], port = device["netconf_port"], username =
device["username"],
    password = device["password"], hostkey_verify = False) as m:
    # do your stuff
```

Getting Configuration Data

Here is an example of how to use the ncclient to get the BGP configuration from Cisco NX-OS:

```
from ncclient import manager
import sys
from lxml import etree

device = {
    "address": "nexus",
    "netconf_port": 830,
    "username": "admin",
    "password": "cisco!"
}

# create a main() method
def main():
    bgp_dom = ""
    <filter type="subtree">
        <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
            <bgp-items>
                <inst-items>
                    <dom-items>
                        <Dom-list/>
```

```

        </dom-items>
    </inst-items>
</bgp-items>
</System>
</filter>
"""

with manager.connect(host=device["address"],
                    port=device["netconf_port"],
                    username=device["username"],
                    password=device["password"],
                    hostkey_verify=False) as m:

    # Collect the NETCONF response
    netconf_response = m.get_config(source='running', filter=bgp_dom)
    # Parse the XML and print the data
    xml_data = netconf_response.data_ele
    print(etree.tostring(xml_data, pretty_print=True).decode("utf-8"))

if __name__ == '__main__':
    sys.exit(main())

```

Getting the Running Configuration and Operational Data

Here is example of getting the interface counters of all the physical interfaces on Cisco NX-OS:

```

from ncclient import manager import sys from lxml import etree
device = {
    "address": "nexus",
    "netconf_port": 830, "username": "admin",
    "password": "cisco"
}
def main():
    intf_ctr_filter = """
<filter>
  <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
<intf-items>
  <phys-items>
    <PhysIf-list>
      <dbgIfIn-items/>
      <dbgIfOut-items/>
    </PhysIf-list>
  </phys-items>
</intf-items>
</System>
</filter>"""
with manager.connect(host=device["address"], port=device["netconf_port"],
                    username=device["username"], password=device["password"],
                    hostkey_verify=False) as m:
    # Collect the NETCONF response
    netconf_response = m.get(filter=intf_ctr_filter)
    # Parse the XML and print the data xml_data =
    netconf_response.data_ele
    print(etree.tostring(xml_data, pretty_print=True).decode("utf-8"))
if __name__ == '__main__': sys.exit(main())

```

Creating a New Configuration

Here is example of how to create VLAN 100 with name using edit config of ncclient:

```

from ncclient import manager import sys from lxml import etree
device = {
    "address": "nexus",
    "netconf_port": 830, "username": "admin",
    "password": "cisco"
}
def main(): add_vlan = """
<config>
  <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <bd-items>
      <bd-items>
        <BD-list>
          <fabEncap>vlan-100</fabEncap>
          <name>inb_mgmt</name>
        </BD-list>
      </bd-items>
    </bd-items>
  </System>
</config>
"""
with manager.connect(host=device["address"], port=device["netconf_port"],
                    username=device["username"],
                    password=device["password"],hostkey_verify=False) as m:
    # create vlan with edit_config
    netconf_response = m.edit_config(target="running", config=add_vlan)
    print(netconf_response)
if __name__ == '__main__': sys.exit(main())

```

Deleting Configuration

Here is example of deleting a loopback interface from Cisco NX-OS:

```

from ncclient import manager import sys from lxml import etree
device = {
    "address": "nexus",
    "netconf_port": 830, "username": "admin",
    "password": "cisco"
}
def main(): remove_loopback = """
<config>
  <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
    <intf-items>
      <lb-items>
        <LbRtdIf-list operation="delete"> <id>lo10</id>
      </LbRtdIf-list>
    </lb-items>
  </intf-items>
</System>
</config>"""
with manager.connect(host=device["address"], port=device["netconf_port"],
                    username=device["username"], password=device["password"],
                    hostkey_verify=False) as m:
    # create vlan with edit_config
    netconf_response = m.edit_config(target="running", config=remove_loopback)
    print(netconf_response)
if __name__ == '__main__':
    sys.exit(main())

```

Troubleshooting the NETCONF Agent

Check Feature Status

- In Cisco NX-OS, enter the **show feature | inc netconf** command to check the agent config.
- To view the status of the NETCONF agent, use the **show feature** command.

```
switch-1# show feature | grep netconf
restconf 1 enabled
switch-1#
```

Check Connectivity

- From a client system, ping the management port of the switch to verify that the switch is reachable.
- There is the XML Management Interface (also known as xmlagent), which is quite different from and often confused as the NETCONF Agent. Please ensure that you connect to the correct port 830 and receive a correct <hello> message (like what is shown in the Establishing a NETCONF Session section) from the server if the server does not respond with the correct NETCONF messages.

Check TLS

If TLS is used, the user can check the TLS status using the **show netconf internal tls service statistics** command.

First check the server start time, which tells whether the TLS server is running.

If the server is not running, then further check the Cert and the Client Root Cert to see whether the certificate is valid.

```
# show netconf internal tls service statistics
=====
TLS Service
=====
Port          : 6513
Cert notBefore : Nov  5 16:48:58 2015 GMT
Cert notAfter  : Nov  5 16:48:58 2035 GMT
Client Root Cert notBefore : Oct  1 18:00:24 2020 GMT
Client Root Cert notAfter  : Sep 26 18:00:24 2040 GMT
Server restarts : 108
Created sessions : 54
Start           : 02/23 01:13:10
Run time        : 142 hour(s) 15 min(s) sec(s)
```

If the server appears to run properly, further check the session status using the **show netconf internal tls session all summary** command:

- If the server does not receive the TLC connection at all, then would recommend debugging the client-side connectivity.
- If the server indeed receives the session but rejects the session immediately, then would recommend checking the user authentication. Check whether the certification is configured properly.

```
# show netconf internal tls session all summary
N9k (config)# show netconf internal tls session all summary
=====
TLS Session
=====
* - history
Client                               Read (KB)   Write (KB)  Status
-----
*10.28.23.116:35490                   0.4         1.1 End
*10.28.23.116:35494                   0.4         1.1 End
--                                     0.0         0.0 Listening
```

Accounting Log for NETCONF Agent

For write operations such as `<edit-config>`, `<commit>` or `<abort>`, NETCONF would emit corresponding accounting log. It would include both the original received request, as well as the eventual changes applied to the switch. This is useful to check the history of config changes via Netconf.

You can see the accounting log using the **show accounting log** command.

Refer the following NETCONF request as an example:

```
---
<edit-config>
<config xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
<System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
  <intf-items>
    <lb-items>
      <LbRtdIf-list>
        <id>lol0</id>
        <descr nc:operation="create">test</descr>
      </LbRtdIf-list>
    </lb-items>
  </intf-items>
</System>
</config>
</edit-config>
---
```

The accounting log shall include the following items:

- Changes applied to the switch:

| Item | Description |
|-----------|---|
| Context | Session ID and user |
| Operation | COMMIT/ABORT |
| Database | Running or Candidate |
| ConfigMO | MO tree's text representation. Up to 3K characters. |
| Status | SUCCESS/FAILED |

Example:

```
Wed Jun 29 13:48:03
2022:type=update:id=2496515744:user=admin:cmd=(COMMIT),database=[running],configMo=[
<topSystem childAction="" dn="sys" status="created,modified"><interfaceEntity
childAction="" rn="intf"
```

```
status="created,modified"><13LbRtdIf childAction="" id="lo10" rn="lb-[lo10]"
status="created,modified"/></interfaceEntity></topSystem>] (SUCCESS)
```

- Original received request

| Item | Description |
|-----------|---|
| Context | Session ID and user |
| Operation | NETCONF:EDIT-CONFIG, NETCONF:COMMIT, NETCONF:DELETE |
| Source IP | NETCONF Client IP |
| Payload | Received XML Request. Up to 3K characters. |
| Status | SUCCESS/FAILED |
| Item | Description |
| Context | Session ID and user |

Example:

```
Wed Jun 29 13:48:03
2022:type=update:id=2496515744:user=admin:cmd=(NETCONF:EDITCONFIG),sourceIp=[192.168.1.2],
payload=[<edit-config><config xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
<System
xmlns="http://cisco.com/ns/yang/cisco-nx-os-device"><intf-items><lb-items><LbRtdIfList><id>lo10</id>
<descr
nc:operation="remove">test</descr></LbRtdIf-list></lb-items></intf-items></System></config></edit-config>]
(SUCCESS)
```

In case of failed request, based on the failed scenarios, a user may not observe both the logs.

- Invalid request:

The invalid request would be rejected without making a configuration change, thus only the original request would be logged.

Example:

```
Wed Jun 29 20:08:36
2022:type=update:id=2517274784:user=admin:cmd=(NETCONF:EDITCONFIG),
sourceIp=[192.168.1.2],payload=[<edit-config><config
xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
<System
xmlns="http://cisco.com/ns/yang/cisco-nx-os-device"><intf-items><lb-items><LbRtdIfList
nc:operation="create"><id>lo10</id>
</LbRtdIf-list></lb-items></intf-items></System></config></edit-config>] (FAILED)
```

- Failures due to various configuration restrictions:

In this case, both the failed configuration attempt and the original request would be logged.

Example:

```
Wed Jun 29 20:11:04
2022:type=update:id=2517274784:user=admin:cmd=(COMMIT),database=[running],
configMo=[<topSystem childAction="" dn="sys"
status="created,modified"><telemetryEntity
childAction=""rn="tm" status="created,modified"><telemetryCertificate childAction=""
filename="foo" hostname="foo" rn="certificate" status="created,modified"
trustpoint="test"/>
```

```

</telemetryEntity></topSystem>] (FAILED)
Wed Jun 29 20:11:04
2022:type=update:id=2517274784:user=admin:cmd=(NETCONF:EDITCONFIG),sourceIp=[192.168.1.2],
payload=[<edit-config><config xmlns:xc="urn:ietf:params:xml:ns:netconf:base:1.0">
<System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
<tm-items><certificateitems><trustpoint>test</trustpoint><hostname>foo</hostname>
<filename>foo</filename></certificateitems></tm-items></System></config></edit-config>]
(FAILED)

```

About NETCONF Explicit Mode

The Network Configuration Protocol Explicit mode is a network management protocol defined by [RFC 6243](#). This protocol defines three standard modes to read the default configuration from the NETCONF server. The standard modes are **report-all**, **trim**, and **explicit**.

Cisco NX-OS already supports **report-all** mode. When configuration is read using **report-all** mode, all the configuration including default configurations are visible to the users. This feature additionally allows NETCONF client to read configuration using the **explicit** mode. In explicit mode, Cisco NX-OS only exposes the configuration that is explicitly done by the user.

Cisco NX-OS only support the explicit mode with NETCONF **<get-config>** and **<edit-config>** RPC:

- **<get-config>**

In explicit mode, **<get-config>** retrieves the data which is explicitly configured by the user irrespective of its value.

- **<edit-config>**

In explicit mode, the operations such as create, delete, merge, replace and remove works slightly different from the **report-all** mode.

- A valid **create** operation attribute for a data node that has been set by a client to its schema default value fails with a **data-exists** error-tag. A valid **create** operation attribute for a data node that has been set by the server to its schema default value succeeds.
- A valid **delete** operation attribute for a data node that has been set by a client to its schema default value succeeds. A valid **delete** operation attribute for a data node that has been set by the server to its schema default value fails with a **data-missing** error-tag.

Guidelines and Limitations

- Cisco NX-OS does not support operations other than `get-config` and `edit-config`.
- As this feature partially supports [RFC 6243](#), switch would not advertise the with-default explicit capability.
- The explicit `get-config` response from clean boot switch would never be empty.
- This feature does not guarantee the expected response when multiple switches are allowed to access the switch at the same time.
- Upgrade from prior 10.3(4) is subject to the below limitations:
 - install ... non-xx

- install ..
- Reload from any release is subject to the below limitation:
- reload ascii
- Upgrade from previous releases
- Downgrade from previous releases

Topic 2.1

NETCONF Explicit Mode Get/Set

To illustrate the explicit mode, consider the below telemetry configuration model.

```
+--rw tm-items
|   +--rw dest-items
|   |   +--rw DestGroup-list* [id]
|   |   |   +--rw id                               telemetry_IDType
|   |   |   +--rw addr-items
|   |   |   |   +--rw Dest-list* [addr port]
|   |   |   |   +--rw addr                       address_Ip
|   |   |   |   +--rw port                       uint16
|   |   |   |   +--rw proto?                    telemetry_Protocol    <===== Default Config
|   |   |   |   +--rw enc?                      telemetry_Encoding    <===== Default Config
|   |   |   |   +--rw nodeid?                  telemetry_NodeIDorZero
```

Add config

- The following edit-config request would configure id along with the destination address and port

Explicit <edit-config> request:

```
-----
<edit-config>
  <with-defaults
  xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-with-defaults">explicit</with-defaults>

  <target>
    <running/>
  </target>
  <config>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
        <dest-items>
          <DestGroup-list>
            <id xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
nc:operation="create">1</id>
              <addr-items>
                <Dest-list>
                  <addr>2.2.2.2</addr>
                  <port>2</port>
                </Dest-list>
              </addr-items>
            </DestGroup-list>
          </dest-items>
        </tm-items>
      </System>
    </config>
```

```
</edit-config>
-----
```

Response:

```
-----

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <ok/>
</rpc-reply>

-----
```

- This following report-all get-config request would only return all the data.

Report-all <get-config>request:

```
-----

<get-config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <with-defaults xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-with-defaults">
    report-all</with-defaults>
  <filter>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
      </tm-items>
    </System>
  </filter>
</get-config>

-----
```

Response:

```
-----

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <data>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
        <adminSt>enabled</adminSt>
        <batchDmeEvt>true</batchDmeEvt>
        <dest-items>
          <DestGroup-list>
            <id>1</id>
            <addr-items>
              <Dest-list>
                <addr>2.2.2.2</addr>
                <port>2</port>
                <enc>GPB</enc>
                <proto>gRPC</proto>
              </Dest-list>
            </addr-items>
          </DestGroup-list>
        </dest-items>
      </tm-items>
    </System>
  </data>
</rpc-reply>

-----
```

- This following explicit get-config request would only return the configured data.

Explicit <get-config> request:

```
-----
<get-config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <with-defaults
xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-with-defaults">explicit</with-defaults>

    <filter>
      <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
        <tm-items/>
      </System>
    </filter>
  </get-config>
-----
```

Response:

```
-----
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <data>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
        <dest-items>
          <DestGroup-list>
            <id>1</id>
            <addr-items>
              <Dest-list>
                <addr>2.2.2.2</addr>
                <port>2</port>
              </Dest-list>
            </addr-items>
          </DestGroup-list>
        </dest-items>
      </tm-items>
    </System>
  </data>
</rpc-reply>
-----
```

Add Configuration Using CLI

Below are the example steps to configure NETCONF request/response in explicit mode.

```
switch(config)# telemetry
switch(config-telemetry)# destination-group 1
switch(conf-tm-dest)# ip address 2.2.2.2 port 2 protocol UDP encoding JSON
```

- The above explicit get-config request would return following response

```
-----
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <data>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
        <dest-items>
          <DestGroup-list>
```

```

        <id>1</id>
        <addr-items>
          <Dest-list>
            <addr>2.2.2.2</addr>
            <port>2</port>
            <enc>JSON</enc>
            <proto>UDP</proto>
          </Dest-list>
        </addr-items>
      </DestGroup-list>
    </dest-items>
  </tm-items>
</System>
</data>
</rpc-reply>

```

Remove configuration

Request:

```

<edit-config>
  <with-defaults>
    xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-with-defaults">explicit</with-defaults>

    <target>
      <running/>
    </target>
    <config>
      <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
        <tm-items>
          <dest-items>
            <DestGroup-list>
              <id>1</id>
              <addr-items>
                <Dest-list>
                  <addr>2.2.2.2</addr>
                  <port>2</port>
                  <enc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
nc:operation="delete">JSON</enc>
                </Dest-list>
              </addr-items>
            </DestGroup-list>
          </dest-items>
        </tm-items>
      </System>
    </config>
  </edit-config>

```

Response:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <ok/>
</rpc-reply>

```

- The above explicit get-config request would return following response

```

-----
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <data>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
        <dest-items>
          <DestGroup-list>
            <id>1</id>
            <addr-items>
              <Dest-list>
                <addr>2.2.2.2</addr>
                <port>2</port>
                <proto>UDP</proto>
              </Dest-list>
            </addr-items>
          </DestGroup-list>
        </dest-items>
      </tm-items>
    </System>
  </data>
</rpc-reply>
-----

```

Add Configuration Using CLI

Follow below steps to configure explicit mode:

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **destination-group** *dgrp_id*
4. **ip address** *ip_address* *port port* **protocol** *procedural-protocol* **encoding** *encoding-protocol*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | configure terminal Example: switch# configure terminal | Enter the global configuration mode. |
| Step 2 | telemetry Example: switch# telemetry | Enter configuration mode for streaming telemetry. |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 3 | destination-group <i>dgrp_id</i> Example: switch# destination-group 1 | Create a destination group and enter destination group configuration mode. |
| Step 4 | ip address <i>ip_address</i> port <i>port</i> protocol <i>procedural-protocol</i> encoding <i>encoding-protocol</i> Example: switch# ip address 2.2.2.2 port 2 protocol UDP encoding JSON | Specify an IPv4 IP address and port to receive encoded telemetry data. |

Example

- The above explicit get-config request would return following response

```

-----
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <data>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
        <dest-items>
          <DestGroup-list>
            <id>1</id>
            <addr-items>
              <Dest-list>
                <addr>2.2.2.2</addr>
                <port>2</port>
                <enc>JSON</enc>
                <proto>UDP</proto>
              </Dest-list>
            </addr-items>
          </DestGroup-list>
        </dest-items>
      </tm-items>
    </System>
  </data>
</rpc-reply>

```

Remove configuration

Request:

```

-----
<edit-config>
  <with-defaults
xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-with-defaults">explicit</with-defaults>

  <target>
    <running/>
  </target>
  <config>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
        <dest-items>

```

```

        <DestGroup-list>
          <id>1</id>
          <addr-items>
            <Dest-list>
              <addr>2.2.2.2</addr>
              <port>2</port>
              <enc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
nc:operation="delete">JSON</enc>
            </Dest-list>
          </addr-items>
        </DestGroup-list>
      </dest-items>
    </tm-items>
  </System>
</config>
</edit-config>

```

Response:

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <ok/>
</rpc-reply>

```

- The above explicit get-config request would return following response

```

<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <data>
    <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      <tm-items>
        <dest-items>
          <DestGroup-list>
            <id>1</id>
            <addr-items>
              <Dest-list>
                <addr>2.2.2.2</addr>
                <port>2</port>
                <proto>UDP</proto>
              </Dest-list>
            </addr-items>
          </DestGroup-list>
        </dest-items>
      </tm-items>
    </System>
  </data>
</rpc-reply>

```



CHAPTER 28

RESTCONF Agent

- [About the RESTCONF Agent, on page 371](#)
- [Guidelines and Limitations, on page 372](#)
- [Configuring the RESTCONF Agent, on page 372](#)
- [Manage a RESTCONF Session, on page 373](#)
- [RESTCONF Get / Set, on page 373](#)
- [RESTCONF Device YANG RPC, on page 375](#)
- [Troubleshooting the RESTCONF Agent, on page 377](#)

About the RESTCONF Agent

Cisco NX-OS RESTCONF is a HTTP -based protocol for configuring data that are defined in YANG version 1, using datastores defined in NETCONF.

NETCONF defines configuration datastores and a set of Create, Retrieve, Update, and Delete (CRUD) operations that can be used to access these datastores. The YANG language defines the syntax and semantics of datastore content, operational data, protocol operations, and event notifications.

Cisco NX-OS RESTCONF uses HTTP operations to provide CRUD operations on a conceptual datastore containing YANG-defined data. This data is compatible with a server which implements NETCONF datastores.

The RESTCONF protocol supports both XML and JSON payload encodings. User authentication is done through the HTTP Basic Authentication.

The following table shows the Protocol operations that the Cisco NX-OS RESTCONF Agent supports:

Table 31: Supported Operations

| RESTCONF | NETCONF Equivalent |
|----------|---|
| OPTIONS | NETCONF: none |
| HEAD | NETCONF: none |
| GET | NETCONF: <get-config>, <get> |
| POST | NETCONF: <edit-config> (operation="create") |
| PUT | NETCONF: <edit-config> (operation="create/replace") |
| PATCH | NETCONF: <edit-config> (operation="merge") |

| | |
|--------|---|
| DELETE | NETCONF: <edit-config> (operation="delete") |
|--------|---|

Guidelines and Limitations

The RESTCONF Agent has the following guideline and limitation:

- Cisco NX-OS RESTCONF is based on an RFC draft entitled RESTCONF Protocol. See <https://tools.ietf.org/html/draft-ietf-netconf-restconf-10>.
- Beginning with Cisco NX-OS Release 10.4(3)F, the following items in RFC 8040, <https://datatracker.ietf.org/doc/html/rfc8040>, are supported:
 - 3.2 RESTCONF Media Types
 - 4.2 HEAD
- For TLS, v1.3 is supported, and the minimally supported version is v1.2.
- Beginning with Cisco NX-OS Release 10.4(3)F, RESTCONF is supported on 92348GC-X.

Configuring the RESTCONF Agent

This procedure describes how to enable and configure the RESTCONF.

Before you begin

Before communicating with the switch using RESTCONF, the RESTCONF Agent must be enabled. The RESTCONF Agent is enabled or disabled by entering the **[no] feature restconf** command. **feature nxapi** is needed to run the http server.

SUMMARY STEPS

1. **configure terminal**
2. (Optional) **feature nxapi**
3. (Optional) **nxapi http port port-num**
4. (Optional) **nxapi https port port-num**
5. **feature restconf**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|----------------------------------|
| Step 1 | configure terminal Example: switch# configure terminal | Enter global configuration mode. |

| | Command or Action | Purpose |
|--------|---|---|
| Step 2 | (Optional) feature nxapi Example: switch(config)# feature nxapi | Enable nxapi for the HTTP server. |
| Step 3 | (Optional) nxapi http port port-num Example: switch(conf-tm-sub)# nxapi http port 80 | Specifies the port used for. The default port is 80. |
| Step 4 | (Optional) nxapi https port port-num Example: switch(conf-tm-sub)# nxapi https port 443 | Specifies the port used for https. The default port is 443. |
| Step 5 | feature restconf Example: switch(conf-tm-sensor)# feature restconf | Enable restconf. |

Manage a RESTCONF Session

RESTCONF runs on top of HTTP. The NX-API server on the switch listens on the port 80 / 443 (or configured port) of the management port IP address. The client can establish a connection with the HTTP server, and RESTCONF subsystem hooks to the “/restconf” URL.

The client can maintain a long-lived HTTP connection to interact with the RESTCONF agent, though it is usually easier to send a single request then terminate. Please refer to the below sections for explanations.

RESTCONF Get / Set

The following shows an example request using the linux **curl** command.

<GET>

This operation retrieves configuration data from a specified datastore.

The following is the example of HTTP Get request and the response message.

```
client-host % curl -X GET -H "Authorization: Basic YWRtaW46Y2lzY28=" -H "Accept:
application/yang.data+xml"
"http://192.0.20.123/restconf/data/Cisco-NX-OS-device:System/bgp-items/inst-items/dom-items/Dom-list?content=config"
-i
HTTP/1.1 200 OK
Server: nginx/1.7.10
Date: Tue, 27 Sep 2016 20:26:03 GMT
Content-Type: application/yang.data+xml
Content-Length: 395
Connection: keep-alive
Set-Cookie: nxapi_auth=admin:147500856185650327
Status: 200 OK
  <Dom-list>
    <name>default</name>
```

```

    <always>enabled</always>
    <bestPathIntvl>300</bestPathIntvl>
    <holdIntvl>180</holdIntvl>
    <kaIntvl>60</kaIntvl>
    <maxAsLimit>0</maxAsLimit>
    <pfxPeerTimeout>30</pfxPeerTimeout>
    <pfxPeerWaitTime>90</pfxPeerWaitTime>
    <reConnIntvl>60</reConnIntvl>
    <rtrId>2.2.2.2</rtrId>
  </Dom-list>
client-host %

```

<POST>

This operation writes a specified configuration to the target datastore.

The following is the example of HTTP POST request and the response message.

```

client-host % curl -X POST -H "Authorization: Basic YWRtaW46Y2lzY28=" -H "Content-Type:
application/yang.data+xml" -d '<always>enabled</always><rtrId>2.2.2.2</rtrId>'
"http://192.0.20.123/restconf/data/Cisco-NX-OS-device:System/bgp-items/inst-items/dom-items/Dom-list=default"
-i
HTTP/1.1 201 Created
Server: nginx/1.7.10
Date: Tue, 27 Sep 2016 20:25:31 GMT
Transfer-Encoding: chunked
Connection: keep-alive
Set-Cookie: nxapi_auth=admin:14750085316974134
Status: 201 Created
Location: /System/bgp-items/inst-items/dom-items/Dom-list=default/always/rtrId/

```

Retrieve Ephemeral Data

Ephemeral data is high volume data. DME provides a batching mechanism to retrieve the data so that each batch is of a manageable size in terms of memory usage. The size of the batch is the number of MOs to be retrieved.

You can find information about which data is ephemeral by the comment "Ephemeral data" in the published Cisco-NX-OS-device.yang file.

The output from ephemeral data is returned, if and only if the RESTCONF URI in the request points to:

- A leaf from ephemeral data
- A container or list with ephemeral data children
- An empty container that is used to wrap a list that has direct ephemeral data children
- System level GET queries do not return ephemeral data.

Example

This is an example for retrieving ephemeral data.

The client might send the following GET request message.

```

GET
/restconf/data/Cisco-NX-OS-device:System//urib-items/table4-items/Table4-list=management/route4-items
HTTP/1.1
Host: example.com
Accept: application/yang.data+json

```

The server might respond:

```

HTTP/1.1 200 OK
Date: Fri, 06 Mar 2020 11:10:30 GMT
Server: nginx/1.7.10
Content-Type: application/yang.data+json
{
  "route4-items": {
    "Route4-list": [{
      "prefix": "172.23.167.255/32", "flags": "0", ...

```

RESTCONF Device YANG RPC

About Operational Commands in RESTCONF

This feature provides ways to perform model driven operation commands execution on the switch.

The following is the list of supported execution RPCs. Information about the RPCs can be found in the published Cisco-NX-OS-device.yang file.

Table 32: Device YANG RPC

| Operation | Device YANG RPC | CLI |
|--------------------------------|---|--|
| Checkpoint | checkpoint | checkpoint <name> checkpoint <file> |
| Rollback | rollback | rollback running-config checkpoint <name> rollback running-config checkpoint <file> |
| Install | install_all_nxos install_add install_activate install_deactivate install_commit install_remove | install all nxos <image> install {add activate deactivate commit} <image> |
| Import Crypto Certificate | import_ca_certificate | crypto ca import <trustpoint> pkcs12 <file> |
| Switch Reload or Module Reload | reload | reload [timer <seconds>] reload module <module number> |
| Copy File | copy | copy <source> <destination> |

Device YANG RPC Examples

Creating checkpoint

The client might send the following POST request message:

```
POST /restconf/operations/Cisco-NX-OS-device:checkpoint
Accept: application/yang.operation+json,application/yang.errors+json
Content-type: application/yang.operation+json
Body: {
  "input": {
    "name": "checkpoint-1",
    "description": "testing checkpoint through Restconf" }
}
```

The server might respond:

```
HTTP/1.1 204 No content
```

Rollback

The client might send the following POST request message:

```
POST /restconf/operations/Cisco-NX-OS-device:rollback
Accept: application/yang.operation+json,application/yang.errors+json
Content-type: application/yang.operation+json
Body: {
  "input": {
    "name": "checkpoint-1",
    "action": "create"
  }
}
```

The server might respond:

```
HTTP/1.1 204 No content
```

Install

The client might send the following POST request message:

```
POST /restconf/operations/Cisco-NX-OS-device:install_all_nxos
Accept: application/yang.operation+json,application/yang.errors+json
Content-type: application/yang.operation+json
Body: {
  "input": {
    "nxos": "bootflash:nxos.10.1.1-jcco.bin" }
}
```

The server might respond:

```
HTTP/1.1 204 No content
```

Import ca certificate RPC

The client might send the following POST request message:

```
POST /restconf/operations/Cisco-NX-OS-device:import_ca_certificate
Accept: application/yang.operation+json,application/yang.errors+json
Content-type: application/yang.operation+json
Body: {
  "input": {
    "trustpoint": "mytrustpoint",
    "pkcs12": "bootflash:server.pfx",
    "passphrase": "mypassphrase"
  }
}
```

The server might respond:

```
HTTP/1.1 204 No content
```

Switch reload

The client might send the following POST request message:

```
POST /restconf/operations/Cisco-NX-OS-device:reload
Accept: application/yang.operation+json,application/yang.errors+json
Content-type: application/yang.operation+json
Body: {
  "input": {
  }
}
```

The server might respond:

```
HTTP/1.1 204 No content
```

Module reload

The client might send the following POST request message:

```
The client might send the following POST request message:
POST /restconf/operations/Cisco-NX-OS-device:reload
Accept: application/yang.operation+json,application/yang.errors+json
Content-type: application/yang.operation+json
Body: {
  "input": {
    "module": "31"
  }
}
```

The server might respond:

```
HTTP/1.1 204 No content
```

Copy file RPC

The client might send the following POST request message:

```
POST /restconf/operations/Cisco-NX-OS-device:reload
Accept: application/yang.operation+json,application/yang.errors+json
Content-type: application/yang.operation+json
Body: {
  "input": {
    "source": "tftp://10.1.1.1/users/myname/config1.txt",
    "destination": "bootflash:",
    "vrf": "management"
  }
}
```

The server might respond:

```
HTTP/1.1 204 No content
```

Troubleshooting the RESTCONF Agent

Check Feature Status

- In Cisco NX-OS, enter the `show feature | inc restconf` command to check the agent config.
- To view the status of the RESTCONF agent, use the `show feature` command and include the expression `restconf`.

```
switch-1# show feature | grep restconf
restconf 1 enabled
switch-1#
```

Check NX-API Connectivity

- Enable the web server by issuing the **feature nxapi** command.
- Ensure that the **nxapi http port 80** command is configured to open the port for HTTP.
- Ensure that the **nxapi https port 443** command is configured to open the port for HTTPS.
- Ping the management port of the switch to verify that the switch is reachable.

Check Restconf Errors

The following shows a common error message and offers guidelines for resolving it.

- If you receive this message soon after sending a request (for example, seconds), verify the following:

Error Message: *Sorry, the page you are looking for is currently unavailable*

- The NXAPI feature is enabled as documented in "Troubleshooting Connectivity".
- The RESTCONF feature is enabled (**show feature | grep restconf**). If RESTCONF is not enabled, enable it (**feature restconf**).
- The port is configured for HTTP or HTTPS by NX-API. Use **show nxapi** to verify that the port is configured.

```
switch-1# show nxapi
nxapi enabled
HTTP Listen on port 80
HTTPS Listen on port 443
```

If the port is not configured for HTTP or HTTPS, configure it by issuing **nxapi http port 80** or **nxapi https port 443**.

- If you receive this message long after sending a request (for example, minutes), ensure that the system is not overloaded with excessive concurrent requests to query from the top level of the switch. Excessive top-level queries can create a significant resource burden.

You can ensure the switch is not overloaded by either of the following:

- Throttle back the number of requests that the client is sending.
- On the switch, restart the RESTCONF agent by issuing **no feature restconf**, then **feature restconf**.

Accounting Log for RESTCONF Agent

For write operations such as POST, PUT, PATCH, or DELETE, RESTCONF would emit the relevant accounting log. It would include both the original received request, as well as the eventual changes applied to the switch.

You can see the accounting log using the **show accounting log** command.

Consider the following example request:


```

---
curl -s -L --request POST --user admin: --header 'Content-Type: application/yang.data+json'
  --url `restconf/data/Cisco-NX-OS-device:System/intf-items/lb-items/LbRtdIf-list=lo10`
--data-raw @request.txt
Payload:
<descr>test</descr>
Or
{"descr":"test"}
---

```

The accounting log shall include the following items:

Table 33: Changes applied to the switch

| Item | Description |
|-----------|---|
| Context | Session ID and user |
| Operation | COMMIT/ABORT |
| Database | Running or Candidate |
| ConfigMO | MO tree's text representation. Up to 3K characters. |
| Status | SUCCESS/FAILED |

Example:

```

Wed Jun 29 13:53:37
2022:type=update:id=3180018864:user=admin:cmd=(COMMIT),database=[running],configMo=[
<topSystem childAction="" dn="sys" status="created,modified"><interfaceEntity childAction=""
rn="intf" status="created,modified"><l3LbRtdIf childAction="" descr="test" id="lo10" rn="lb-
[lo10]" status="created,modified"/></interfaceEntity></topSystem>] (SUCCESS)

```

Table 34: Original received request

| Item | Description |
|-----------|--|
| Context | Session ID and user |
| Operation | RESTCONF:POST, RESTCONF:PUT, RESTCONF:PATCH, RESTCONF:DELETE |
| Source IP | RESTCONF Client IP |
| URL | HTTP URL |
| Payload | Received XML/JSON Request. Up to 3K characters. |
| Status | SUCCESS/FAILED |

Example:

```

Wed Jun 29 13:53:37
2022:type=update:id=3180018864:user=admin:cmd=(RESTCONF:POST),sourceIp=[192.168.1.2],
url=[/restconf/data/Cisco-NX-OS-device:System/intf-items/lb-items/LbRtdIfList=lo10],payload=[<descr>test</descr>]
(SUCCESS)

```

In case of failed request, based on the failed scenarios, a user may not observe both the logs.

Invalid request:

The invalid request would be rejected without making a configuration change, thus only the original request would be logged.

Example:

```
Wed Jun 29 20:16:26
2022:type=update:id=3180018864:user=admin:cmd=(RESTCONF:POST),
sourceIp=[192.168.1.2],url=[/restconf/data/Cisco-NX-OS-device:System/intf-itens/lb-itens/lbRtdIfIlist=lol0],payload=[<descr>test</descr>]
(FAILED)
```

Request fails due to various configuration restrictions:

In this case, both the failed configuration attempt and the original request would be logged.

Example:

```
Wed Jun 29 20:32:01
2022:type=update:id=3180018864:user=admin:cmd=(COMMIT),database=[running],
configMo=[<topSystem childAction="" dn="sys" status="created,modified"><telemetryEntity
childAction="" rn="tm" status="created,modified"><telemetryCertificate childAction=""
filename="foo" hostname="foo" rn="certificate" status="created,modified"
trustpoint="test"/></telemetryEntity></topSystem>] (FAILED)
Wed Jun 29 20:32:01
2022:type=update:id=3180018864:user=admin:cmd=(RESTCONF:PATCH),
sourceIp=[192.168.1.2],url=[/restconf/data/Cisco-NX-OS-device:System/intens/certificateitens],payload=[<trustpoint>test/</trustpoint><hostname>foo/</hostname><filename>
foo</filename>] (FAILED)
```



CHAPTER 29

gRPC Agent

- [About gRPC Agent, on page 381](#)
- [Revision History, on page 381](#)
- [Guidelines and Limitations for gRPC Agent , on page 381](#)
- [Troubleshooting, on page 389](#)

About gRPC Agent

gRPC is a modern opensource high performance Remote Procedure Call (RPC) framework. Cisco NX-OS provides a gRPC agent to support gRPC related services including: gNMI and gNOI.

Revision History

| Release | Description |
|---------|---|
| 9.3(3) | Add support for <ul style="list-style-type: none">• gRPC port• gRPC certificate |
| 10.1(1) | Add support for client cert-based authentication <ul style="list-style-type: none">• gRPC client root certificate |
| 10.3(3) | Support NGINX to act as GRPC proxy |

Guidelines and Limitations for gRPC Agent

Following are the guidelines and limitations for gRPC agent:

- When you enable gRPC on both the management VRF and default VRF and later disable on the default VRF, the gNMI operations on the management VRF stop working.

As a workaround, disable gRPC completely by entering the **no feature gRPC** command and reprovision it by entering the **feature gRPC** command or with any existing gRPC configuration commands like, **gRPC**

certificate or **grpc port**. You must also resubscribe to any existing notifications on the management VRF.

- If the gRPC certificate is explicitly configured, after a reload with the saved startup configuration to a prior Cisco NX-OS 9.3(x) image, the gRPC feature does not accept connections.

To confirm this issue, enter the **show grpc gnmi service statistics** command. The following status error message is displayed:

```
Status: Not running - Initializing...Port not available or certificate
invalid.
```

Unconfigure and configure the proper certificate command to restore the service.

- If you have configured a custom gRPC certificate, upon entering the **reload ascii** command the configuration is lost. It reverts to the default day-1 certificate. After entering the **reload ascii** command, the switch reloads. Once the switch is up again, you must reconfigure the gRPC custom certificate.



Note This applies when entering the `grpc certificate` command.

- The reachability in non-default VRF for gRPC is supported only over L3VNI's/EVPN and IP. However, reachability over MPLS in non-default VRF and VXLAN Flood and Learn is not supported.
- For Cisco NX-OS release prior to 9.3(x), information about supported platforms, see *Platform Support for Programmability Features* in the guide for that release. Starting with Cisco NX-OS release 9.3(x), for information about supported platforms, see the [Nexus Switch Platform Matrix](#).
- The gRPC process uses the HIGH_PRIO control group, which limits the CPU usage to 75% of CPU and memory to 4 GB.
- The gRPC agent supports management VRF and one user specified VRF for a total of two gRPC servers on each switch. Supporting a gRPC in the user specified VRF (for example: the default VRF) adds flexibility to offload processing gRPC calls from the management VRF, where significant traffic load is not desirable.
- If two gRPC servers are configured, be aware of the following:
 - VRF boundaries are strictly enforced, so each gRPC server process requests independent of the other. Requests do not cross between VRFs.
 - The two servers are not HA or fault tolerant. One gRPC server does not back up the other, and there is no switchover or switchback between them.
 - Any limits for the gRPC server are per VRF.
- Beginning with Cisco NX-OS Release 10.4(3)F, gRPC is supported on 92348GC-X.

Configuring the gRPC Agent

Configuring gRPC

Configure the gNMI feature through the `grpc` commands.

To import certificates used by the **grpc certificate** command onto the switch, see the [Installing Identity Certificates](#) section of the Cisco Nexus 9000 Series NX-OS Security Configuration Guide.



Note When modifying the installed identity certificates or **grpc port** and **grpc certificate** values, the gRPC server might restart to apply the changes. When the gRPC server restarts, any active subscription is dropped, and you must resubscribe.

Before you begin

Prepare and sign the required certificate files for the server authentication.

You can re-use the existing trustpoint files as this is not specific to gRPC.

SUMMARY STEPS

1. **configure terminal**
2. (Optional) **crypto ca trustpoint** *<server-trustpoint>*
3. **crypto ca import** *<server-trustpoint>* **pkcs12 bootflash:** *<server-ca-file>* *<pkcs-password>*
4. **feature grpc**
5. (Optional) **grpc port** *port-id*
6. **grpc certificate** *certificate-id*
7. (Optional) **use-vrf default**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|---|--|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enters configuration mode. |
| Step 2 | (Optional) crypto ca trustpoint <i><server-trustpoint></i> Example: <pre>switch# crypto ca trustpoint tls_server_trustpoint</pre> | Creates a trustpoint for server authentication. Step 2-3 is optional if there already exist usable server trustpoint. |
| Step 3 | crypto ca import <i><server-trustpoint></i> pkcs12 bootflash: <i><server-ca-file></i> <i><pkcs-password></i> Example: <pre>switch# crypto ca import tls_server_trustpoint pkcs12 bootflash:server.pfx test</pre> | Imports the server pkcs12 file to the trustpoint. |
| Step 4 | feature grpc Example: | Enables the gRPC agent, which supports the gNMI interface for dial-in. |

| | Command or Action | Purpose |
|---------------|---|---|
| | switch# feature grpc switch(config)# | |
| Step 5 | (Optional) grpc port <i>port-id</i> Example: switch(config)# grpc port 50051 | Configures the port number. The range of port-id is from 1024 to 65535. 50051 is the default. |
| Step 6 | grpc certificate <i>certificate-id</i> Example: switch(config)# grpc certificate cert-1 | Specify the certificate trustpoint ID. For more information, see the Installing Identity Certificates section of the Cisco Nexus 9000 Series NX-OS Security Configuration Guide for importing the certificate. |
| Step 7 | (Optional) use-vrf default Example: switch(config)# grpc use-vrf default | Enables the gRPC agent to accept incoming (dial-in) RPC requests from the default VRF. This step enables the default VRF to process incoming RPC requests. By default, the management VRF processes incoming RPC requests when the gRPC feature is enabled. |

Generating Key/Certificate

The following is an example for generating a self-signed key/certificate in the switch bash shell. This is only for experimental usage. For more information on generating identity certificates, see the [Installing Identity Certificates](#) section of the Cisco Nexus 9000 Series NX-OS Security Configuration Guide.



Note This task is an example of how a certificate can be generated on a switch. You can also generate a certificate in any Linux environment. In a production environment, you should consider using a CA signed certificate.

SUMMARY STEPS

1. Generate the self-signed key and pem files.
2. After generating the key and pem files, you must bundle the key and pem files for use in the trustpoint CA Association.
3. Set up the trustpoint CA Association by inputting in the pkcs12 bundle into the trustpoint.
4. Verify the setup.

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | Generate the self-signed key and pem files. | switch# run bash sudo su bash-4.3# openssl req -x509 -newkey rsa:2048 -keyout self_sign2048.key -out self_sign2048.pem -days 365 -nodes |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 2 | After generating the key and pem files, you must bundle the key and pem files for use in the trustpoint CA Association. | After generating the key and pem files, you must bundle the key and pem files for use in the trustpoint CA Association. <pre>switch# run bash sudo su bash-4.3# cd /bootflash/ bash-4.3# openssl pkcs12 -export -out self_sign2048.pfx -inkey self_sign2048.key -in self_sign2048.pem -certfile self_sign2048.pem -password pass:Ciscolab123! bash-4.3# exit</pre> |
| Step 3 | Set up the trustpoint CA Association by inputting in the pkcs12 bundle into the trustpoint. | <pre>switch(config)# crypto ca trustpoint mytrustpoint switch(config-trustpoint)# crypto ca import mytrustpoint pkcs12 self_sign2048.pfx Ciscolab123!</pre> |
| Step 4 | Verify the setup. | <pre>switch(config)# show crypto ca certificates Trustpoint: mytrustpoint certificate: subject= /C=US/O=Cisco Systems, Inc./OU=CSG/L=San Jose/ST=CA/street=3700 Cisco Way/postalCode=95134/CN=ens.cisco.com/serialNumber=FGE18420KOR issuer= /C=US/O=Cisco Systems, Inc./OU=CSG/L=San Jose/ST=CA/street=3700 Cisco Way/postalCode=95134/CN=ens.cisco.com/serialNumber=FGE18420KOR serial=0413 notBefore=Nov 5 16:48:58 2015 GMT notAfter=Nov 5 16:48:58 2035 GMT SHA1 Fingerprint=2E:99:2C:CE:2F:C3:B4:EC:C7:E2:52:3A:19:A2:10:D0:54:CA:79:3E purposes: sslserver sslclient CA certificate 0: subject= /C=US/O=Cisco Systems, Inc./OU=CSG/L=San Jose/ST=CA/street=3700 Cisco Way/postalCode=95134/CN=ens.cisco.com/serialNumber=FGE18420KOR issuer= /C=US/O=Cisco Systems, Inc./OU=CSG/L=San Jose/ST=CA/street=3700 Cisco Way/postalCode=95134/CN=ens.cisco.com/serialNumber=FGE18420KOR serial=0413 notBefore=Nov 5 16:48:58 2015 GMT notAfter=Nov 5 16:48:58 2035 GMT SHA1 Fingerprint=2E:99:2C:CE:2F:C3:B4:EC:C7:E2:52:3A:19:A2:10:D0:54:CA:79:3E purposes: sslserver sslclient</pre> |

Configuring gRPC Client Certificate Authentication

gRPC also allows to authenticate the client based on the cert files (public key). This provides password-less authentication, which is considered more secure than password-based authentication.

Before you begin

Prepare and sign the required certificate files for the server authentication.

You can re-use the existing trustpoint files as this is not specific to gRPC.

SUMMARY STEPS

1. **configure terminal**
2. (Optional) **crypto ca trustpoint** <server-trustpoint>
3. **rsa keypair** <client-key>
4. (Optional) **crypto ca authenticate** <client-root-trustpoint>

5. `grpc client root certificate <client-root-trustpoint>`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enters configuration mode. |
| Step 2 | (Optional) crypto ca trustpoint <server-trustpoint> Example: switch# crypto ca trustpoint tls_server_trustpoint | Creates a trustpoint for server authentication. Step 2-3 is optional if there already exist usable server trustpoint. |
| Step 3 | rsa keypair <client-key> Example: switch# rsa keypar client-key | Generates a rsa key pair for the client trustpoint. |
| Step 4 | (Optional) crypto ca authenticate <client-root-trustpoint> Example: switch# crypto ca authenticate client_trustpoint | Imports the client certificate. This step requires manual copy paste. Please follow the instruction. |
| Step 5 | grpc client root certificate <client-root-trustpoint> Example: switch(config)# grpc client root certificate client_trustpoint | Enters the trustpoint to host the client CA root certificate. |

Example

Config example

This section provides an example config sequence for illustration.

1. Prepare the Client Root CA Certificates.
2. Import the certificate

When you have generated a new certificate to the client root successfully, following are the sample commands to configure them in the switch, and their output.

```
switch(config)# crypto ca trustpoint my_client_trustpoint
switch(config-trustpoint)# crypto ca authenticate my_client_trustpoint
input (cut & paste) CA certificate (chain) in PEM format; end the input with a line
containing only END OF INPUT :
-----BEGIN CERTIFICATE-----
MIIDUDCCAjigAwIBAgIJAJLIsBKCGjQOMA0GCSqGSIb3DQEBCwUAMD0xCzAJBgNV
BAYTA1VTMQswCQYDVQQIDAJDQTERMA8GA1UEBwwIU2FuIEpvc2UxDjAMBGNVBAoM
BUNpc2NvMB4XDTIwMTAxNDIwNTYyN1oXDQwMTAwOTIwNTYyN1owPTELMAkGA1UE
BhMCVVMxCzAJBgNVBAGMAkNBMRERDwYDVQQHDAhTYW4gSm9zZTEOMAwwGA1UECgwF
```



```

Q2lzY28wggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQDEX7qZ2EdogZU4
EW0NSpB3EjY0nSlFLOw/iLKsXfIiQJD0Qhaw16fDnnYZj6vzWEa0ls8canqHCXQl
gUyxFOdGDxa6neQFTqLowSA6UCSQA+eenN2PIpMOjfdFpaPiHu3mmcT1lxP39Ti3
/y548NNORSepApBNkZ1rJSB6Cu9AIFMzgrZXfQDKBGSUOf/CPnvIDZeLcun+zpUu
CxJLA76Et4buPMysuRqMGHIX8CYw8MtjmuCuCTHXNN31ghhgpFxfRw/69pykjU3R
YOrwlSUkvYQhtefHuTHBmqym7MFoBEchwr1C5YTduDzmOvtkhsmogRe3BiIBx45
AnZdtDi1AgMBAAGjUzBRMB0GA1UdDgQWBBSH3IqRrm+mtB5GNsoLXFb3bAVg5Taf
BgNVHSMEGDAWgBSH3IqRrm+mtB5GNsoLXFb3bAVg5TAPBgNVHRMBAf8EBTADAQH/
MA0GCSqGSIb3DQEBCwUAA4IBAQA4Fpc6lRKzBGJQ/7oK1FNcTX/YXkneXDk7Zrj
8W0RS0Kxhgke97d2Cw15P5reXO27kvXsnsz/VZn7JYGUvGS1xTlcCb6x6wNBr4Qr
t9qDBu+LykwqNOFe4VCAv6e4cMXNbH2wHBVS/NSoWnM2FGZ10VppjEGFm6OM+N6z
8n4/rWslfWFbn7T7xHH+N10Ffc+8q8h37opyCnb0ILj+a4rnyus8xXJPQb05DfJe
ahPNfdEsXKDOwkrSDtmKwtWDqdtjSQC4xioKHoshnNgWBjbovPlMQ64UrajBycVv
z9snWBm6p9SdTsV92YwFltRGUqpcI9olsBgH7FUVU1hmdHWE
-----END CERTIFICATE-----END OF INPUT
Fingerprint(s): SHA1
Fingerprint=0A:61:F8:40:A0:1A:C7:AF:F2:F7:D9:C7:12:AE:29:15:52:9D:D2:AE
Do you accept this certificate? [yes/no]:yes switch(config)#
NOTE: Use the CA Certificate from the .pem file content.
switch# show crypto ca certificates Trustpoint: my_client_trustpoint CA certificate 0:
subject=C = US, ST = CA, L = San Jose, O = Cisco
issuer=C = US, ST = CA, L = San Jose, O = Cisco
serial=B7E30B8F4168FB87 notBefore=Oct 1 17:29:47 2020 GMT notAfter=Sep 26 17:29:47 2040
GMT
SHA1 Fingerprint=E4:91:4E:D4:41:D2:7D:C0:5A:E8:F7:2D:32:81:B3:37:94:68:89:10 purposes:
sslserver sslclient

```

3. Associating Trustpoints to gRPC

When you have configured a new certificate to the client root successfully, the following is the output example for associating trustpoints to gRPC server on the switch:

```

switch(config)# feature grpc
switch(config)# grpc client root certificate my_client_trustpoint switch(config)# show
run grpc
!Command: show running-config grpc
!Running configuration last done at: Wed Dec 16 20:18:35 2020
!Time: Wed Dec 16 20:18:40 2020
version 10.1(1) Bios:version N/A feature grpc
grpc gnmi max-concurrent-calls 14 grpc use-vrf default grpc certificate my_trustpoint
grpc client root certificate my_client_trustpoint grpc port 50003

```

4. Validating the Certificate Details

When you have successfully associated the trustpoints to gRPC on the switch, the following is the output example for validating the certificate details:

```

switch# show grpc gnmi service statistics
===== gRPC Endpoint =====
Vrf : management
Server address : [::]:50003
Cert notBefore : Mar 13 19:05:24 2020 GMT
Cert notAfter : Nov 20 19:05:24 2033 GMT
Client Root Cert notBefore : Oct 1 17:29:47 2020 GMT
Client Root Cert notAfter : Sep 26 17:29:47 2040 GMT
...

```

5. Verifying the Connection using Client Certificate Authentication for any gNMI Clients.

The client certificate requests with a private key (pkey) and ca chain (cchain). The password is now optional. Make sure the client needs to supply full chain from to root CA to its client cert.

6. For removing trustpoint reference from gRPC (no command) use the following command:

```

switch(config)# no grpc client root certificate my_client_trustpoint

```

The command will remove the trustpoint reference only from gRPC agent, but the trustpoints CA certificates will NOT be removed. Connections that use client certificate authentication to gRPC server on switch will not establish, but basic authentication with username and password will go through.

Configuring NGINX proxy for GRPC

Like Netconf and Restconf, gRPC agent runs on a dedicate server/port. The gRPC client would need to directly connect to the gRPC agent/server.

Starting from release 10.3(3)F, NX-OS NGINX can also act as GRPC proxy by relaying the gRPC traffic and can be useful for certain use cases.

- GRPC port blocked: The GRPC agent listens on port 50051. If this port is blocked by the firewall, GRPC clients can access the gRPC services indirectly through the NGINX HTTPS port 443.
- Increased VRF support: Currently GRPC services are accessible only via management or one user specified VRFs. NGINX proxy can forward the gRPC requests from any VRF.

This new support does not affect existing behavior. The GRPC client can still connect to the GRPC agent directly. It can also instead connect to the NGINX server, which then proxies the GRPC requests to the GRPC agent. Note that such redirection deems to incur additional request-response latency.

All server and client authentication will be handled by NGINX. Just enable GRPC and configure NGINX server certificate and/or client certificates.

Before you begin

Enable the grpc feature.

Prepare the NX-API certificates. See Using NX-API CLI for detail.

SUMMARY STEPS

1. **configure terminal**
2. **feature nxapi**
3. **nxapi certificate https crt certfile *cert-file***
4. **nxapi certificate https crt keyfile *key-file* password *<password>***
5. **nxapi certificate enable**
6. (Optional) **crypto ca trustpoint *<trustpoint>***
7. (Optional) **crypto ca authenticate *<trustpoint>***
8. (Optional) **nxapi client certificate authentication**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|---|----------------------------|
| Step 1 | configure terminal Example: | Enters configuration mode. |

| | Command or Action | Purpose |
|---------------|---|--|
| | switch# configure terminal switch(config)# | |
| Step 2 | feature nxapi Example: switch# feature nxapi switch(config)# | Enables the feature nxapi. |
| Step 3 | nxapi certificate httpsert certfile <i>cert-file</i> Example: switch# nxapi certificate httpsert certfile bootflash:nxapi.crt | Configures the cert file. |
| Step 4 | nxapi certificate httpsert keyfile <i>key-file</i> password <password> Example: switch# nxapi certificate httpskey keyfile bootflash:nxapi.key password cisco123 | Configures the key file. |
| Step 5 | nxapi certificate enable Example: switch# nxapi certificate enable | Enables the certificate authentication. |
| Step 6 | (Optional) crypto ca trustpoint <i><trustpoint></i> Example: switch# crypto ca trustpoint grpcClientCA | Creates a trustpoint for server authentication. |
| Step 7 | (Optional) crypto ca authenticate <i><trustpoint></i> Example: switch# crypto ca authenticate grpcClientCA | Imports the client root certificate to the trustpoint. |
| Step 8 | (Optional) nxapi client certificate authentication Example: switch# nxapi client certificate authentication | Enables the client certificate authentication. |

Troubleshooting

Check Feature Status

- In Cisco NX-OS device, enter the **show feature grpc** command to check the agent config.
- To view the status of the gRPC agent, use the **show feature** command.

```
switch-1# show feature | grep grpc
restconf 1 enabled
switch-1#
```

Check Connectivity

From a client system, ping the management port of the switch to verify that the switch is reachable.

Gathering gRPC Agent Logs

The /volatile directory houses the grpc agent log.

```
bash-4.3# cd /volatile/ bash-4.3# ls /volatile -al
...
-rw-rw-rw- 1 root root 103412 Jun 21 16:14 grpc-internal-log
...
```

Gathering TM-Trace Logs

```
tmtrace.bin -f gnmi-logs gnmi-events gnmi-errors following are available 2.
Usage:
bash-4.3# tmtrace.bin -d gnmi-events | tail -30 Gives the last 30
...
[06/21/19 15:58:38.969 PDT f8f 3133] [3981658944][tm_transport_internal.c:43] dn:
Cisco-NX-OS-device:System/cdp-items, sub_id: 0,
sub_id_str: 2329, dc_start_time: 0, length: 124, sync_response:1
[06/21/19 15:58:43.210 PDT f90 3133] [3621780288][tm_ec_yang_data_processor.c:93] TM_EC:
[Y] Data received for 2799743488: 49
{
  "cdp-items": {
    "inst-items": {
      "if-items": {
        "If-list": [
          {
            "id": "mgmt0",
            "ifstats-items": {
              "v2Sent": "74",
              "validV2Rcvd": "79"
            }
          }
        ]
      }
    }
  }
}
} [06/21/19 15:58:43.210 PDT f91 3133] [3981658944][tm_transport_internal.c:43] dn:
Cisco-NX-OS-device:System/cdp-items, sub_id: 0,
sub_id_str: 2329, dc_start_time: 0, length: 141, sync_response:1
[06/21/19 15:59:01.341 PDT f92 3133] [3981658944][tm_transport_internal.c:43] dn:
Cisco-NX-OS-device:System/intf-items, sub_id:
4091, sub_id_str: , dc_start_time: 1561157935518, length: 3063619, sync_response:0 [06/21/19
15:59:03.933 PDT f93 3133] [3981658944][tm_transport_internal.c:43] dn:
Cisco-NX-OS-device:System/cdp-items, sub_id:
4091, sub_id_str: , dc_start_time: 1561157940881, length: 6756, sync_response:0 [06/21/19
15:59:03.940 PDT f94 3133] [3981658944][tm_transport_internal.c:43] dn:
Cisco-NX-OS-device:System/lldp-items, sub_id:
```

Gathering MTX-Internal Logs

1. Modify the following file with below /opt/mtx/conf/mtxlogger.cfg

Please refer to the “Diagnose and Serviceability” section to toggle the preferred filter.

2. Disable then enable **feature grpc**.

3. The /volatile directory houses the mtx-internal.log, the log rolls over time so make sure to download the logs before rolled over.

```
bash-4.3# cd /volatile/ bash-4.3# ls /volatile -al
...
-rw-r--r-- 1 root root 24 Jun 21 14:44 mtx-internal-19-06-21-14-46-21.log
-rw-r--r-- 1 root root 24 Jun 21 14:46 mtx-internal-19-06-21-14-46-46.log
-rw-r--r-- 1 root root 175 Jun 21 15:11 mtx-internal-19-06-21-15-11-57.log
-rw-r--r-- 1 root root 175 Jun 21 15:12 mtx-internal-19-06-21-15-12-28.log
-rw-r--r-- 1 root root 175 Jun 21 15:13 mtx-internal-19-06-21-15-13-17.log
-rw-r--r-- 1 root root 175 Jun 21 15:13 mtx-internal-19-06-21-15-13-42.log
-rw-r--r-- 1 root root 24 Jun 21 15:13 mtx-internal-19-06-21-15-14-22.log
-rw-r--r-- 1 root root 24 Jun 21 15:14 mtx-internal-19-06-21-15-19-05.log
-rw-r--r-- 1 root root 24 Jun 21 15:19 mtx-internal-19-06-21-15-47-09.log
-rw-r--r-- 1 root root 24 Jun 21 15:47 mtx-internal.log
```




CHAPTER 30

gNMI - Management Interface

- [About gNMI, on page 393](#)
- [Guidelines and Limitations for gNMI, on page 394](#)
- [Configuring gNMI, on page 397](#)
- [gNMI RPCs, on page 398](#)
- [About Get, on page 400](#)
- [About Set, on page 401](#)
- [About Subscribe, on page 402](#)
- [About Subscribing Custom Syslog Stream, on page 405](#)
- [References, on page 409](#)
- [Troubleshooting gNMI, on page 409](#)
- [Configuration Examples for gRPC Commands, on page 412](#)
- [Gathering Debug Logs, on page 414](#)
- [Accounting Log for gNMI, on page 414](#)

About gNMI

gNMI (gRPC Network Management Interface) is configured on gRPC. Configuring gNMI helps to edit, read the configuration, and operational state of a network device. It also enables the network devices to generate telemetry streams to a designated data collection system.



Note You can configure gNMI as a child service feature of gRPC feature only.

Cisco NX-OS supports all gNMI RPC as mentioned in the following table:

Table 35: Supported gNMI RPCs

| gNMI RPC | Supported |
|--------------|-----------|
| Capabilities | Yes |
| Get | Yes |
| Set | Yes |
| Subscribe | Yes |

To subscribe RPC, Cisco NX-OS supports the below submodes:

Table 36: Subscribe Options

| Type | Sub Type | Supported | Description |
|--------|-----------|-----------|---|
| Once | | Yes | Switch sends current values only once for all specified paths. |
| Poll | | Yes | Whenever the switch receives a Poll message, the switch sends the current values for all specified paths. |
| Stream | Sample | Yes | Once per stream sample interval, the switch sends the current values for all specified paths. The supported sample interval range is from 1 through 604,800 second. The default sample interval is 10 seconds. |
| | On_Change | Yes | The switch sends current values as its initial state, but updates values only when there are changes such as create, modify, or delete for the specified paths. |

gNMI subscription is commonly referred to as dial-in telemetry. As gNMI requires an external customer to initiate the request in the network device.

Cisco NX-OS supports a separate telemetry feature in which the networking device pushes the telemetry data out of external receivers, it is referred to as dial-out telemetry. For more information, see the Telemetry section.

Beginning with Cisco NX-OS Release 10.4(3)F, OpenConfig path `openconfig-system:/system/processes` supports On-change gnmI subscriptions. This path supports both gnmI-proto and gnmI-json encodings. The subscription will remain active until the client unsubscribes the same path.

It is recommended to make the most strict possible subscription for better efficiency. For example, to monitor `cpu-usage-user`, you can subscribe the data directly rather than using On-target gnmI subscription for general path like `/system/processes`, to avoid event data.

Guidelines and Limitations for gNMI

The following are the guidelines and limitations for gNMI:

- When you subscribe an OpenConfig routing policy with a pre-existing CLI configuration as shown below, it returns an empty value due to current implementation of the OpenConfig model.

```
ip prefix-list bgp_v4_drop seq 5 deny 125.2.0.0/16 le 32
ipv6 prefix-list bgp_v6_drop seq 5 deny cafe:125:2::/48 le 128
```

use below path:

```
openconfig-routing-policy:/routing-policy/defined-sets/prefix-sets/prefix-set[name=bgp_v4_drop]/config
openconfig-routing-policy:/routing-policy/defined-sets/prefix-sets/prefix-set[name=bgp_v6_drop]/config
```

- For gNMI subscriptions, use of `origin`, `use_models` or both these commands are optional.
- Before Cisco NX-OS 9.3(x) Release, information about supported platforms, see *Platform Support for Programmability Features* in the guide for that release. Starting with Cisco NX-OS release 9.3(x), for information about supported platforms, see [Nexus Switch Platform Support Matrix](#).
- The feature supports JSON and gNMI proto encoding.
- The feature does not support protobuf any encoding.
- The feature does not support a path prefix in the subscription request, but the Subscription contains an empty prefix field.

Wildcard Path

- Multilevel wildcard "..." in path is not allowed.
- wildcard '*' in the top of the path is not allowed
- wildcard '*' in the key name is not allowed
- wildcard and value are not compatible in keys

The **show grpc gnmi** command has the following guidelines and limitations:

- This command does not support xml or json output format.
- The gRPC agent retains gNMI call data for maximum of an hour after the call ends.
- If the total number of calls exceeds 2000, the gRPC agent purges ended calls based on the internal cleanup routine.

Table 37: Wildcard Support for gNMI Requests

| Type of Request | Wildcard Support |
|--|------------------|
| gNMI GET | Yes |
| gNMI SET | No |
| gNMI SUBSCRIBE, ONCE | Yes |
| gNMI SUBSCRIBE, POLL | Yes |
| gNMI SUBSCRIBE, STREAM, SAMPLE | Yes |
| gNMI SUBSCRIBE, STREAM, TARGET_DEFINED | Yes |
| gNMI SUBSCRIBE, STREAM, ON_CHANGE | No |

Scale Considerations

- Each gNMI message has a maximum size of 12 MB. If the amount of collected data exceeds the maximum size, the collected data drop. This is applicable for gNMI ON_CHANGE mode only.

You can create focused subscriptions that manage smaller, granular data collection sets to avoid collection data drop. It is recommended to create multiple subscriptions for different, lower-level parts of the path instead of one higher-level path.

Across all subscriptions, there is support of up to 250K aggregate MOs. Subscribing to more MOs can lead to collection data drops

- For all subscriptions, there is support of up to 250K aggregate MOs. If you subscribe to more MOs this impacts collection data drop.

Guidelines and Limitations for gNMI Subscription

- On-change subscriptions work for both gnmi and telemetry, but only one of these may be active at one time. If a subscription is made from one of these agents it will overwrite any existing subscription information.

The following is an example for On-change gnmi subscription for openconfig-system:/system/processes:

Request:

```
./gnmi-console_enhanced_plus --host 172.22.244.142 --port 50051 -u admin -p insieme
--tls --cafile /tmp/grpc.pem --hostnameoverride ems.cisco.com --operation=Subscribe
--submode ON_CHANGE --xpath "openconfig-system:/system/processes/process[pid=1]" -e
JSON
```

Response:

```
///// initial snapshot
```

```
Received response 1 -----
```

```
/system/
{
  "processes": {
    "process": [
      {
        "pid": "1",
        "state": {
          "pid": "1",
          "name": "init",
          "start-time": "1706643350384761800",
          "cpu-usage-user": "14",
          "cpu-usage-system": "12",
          "cpu-utilization": 0,
          "memory-usage": "180224",
          "memory-utilization": 0
        }
      }
    ]
  }
}
```

```
///// first event update
```

```
/system/
{
  "processes": {
    "process": [
      {
```

```

        "pid": "1",
        "state": {
            "start-time": "1706643350384761800",
            "cpu-usage-user": "14",
            "cpu-usage-system": "12"
        }
    }
}
]
}

```

Configuring gNMI

Configure the gNMI feature through the gRPC gNMI commands.

Configuring gNMI Options

Before you begin

gNMI is a child service feature of gRPC feature only.

For more information, see the gRPC Agent documentation to enable the gRPC agent.

SUMMARY STEPS

1. `switch# configure terminal`
2. (Optional) `grpc gnmi max-concurrent-call`
3. `grpc gnmi subscription target-defined min-interval<interval>`
4. `grpc gnmi subscription query-condition keep-data-timestamp`
5. (Optional) `grpc gnmi keepalive-timeout <timeout>`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|--|--|
| Step 1 | <p><code>switch# configure terminal</code></p> <p>Example:</p> <pre>switch# config terminal switch(config)#</pre> | Enters the global configuration mode. |
| Step 2 | <p>(Optional) <code>grpc gnmi max-concurrent-call</code></p> <p>Example:</p> <pre>switch(config)# grpc gnmi max-concurrent-call 16</pre> | <p>Configuring this command sets the limit of simultaneous dial-in calls to the gNMI server on the switch. The default limit is 8. You can configure value with limit range of 1–16.</p> <p>The maximum value that you configure is on each VRF. If you set the limit value of 16 and if gNMI is configured on both management and default VRFs, each VRF supports 16 simultaneous gNMI calls.</p> |

| | Command or Action | Purpose |
|---------------|---|--|
| | | This command has no effect for ongoing or in-progress gNMI calls. gRPC enforces a limit on new calls and the active calls have no impact and you can complete the active call. |
| Step 3 | grpc gnmi subscription target-defined min-interval <interval> Example: <pre>switch(config)# grpc gnmi subscription target-defined min-interval 60</pre> | You can modify the default target-defined sample interval from 30 seconds to other required value. |
| Step 4 | grpc gnmi subscription query-condition keep-data-timestamp Example: <pre>switch(config)# grpc gnmi subscription query-condition keep-data-timestamp</pre> | <p>You can enable sample or once or poll subscriptions to get a timestamp from the database when the data is last updated.</p> <p>Note</p> <ul style="list-style-type: none"> • This command is supported only for PROTO and this is not supported for JSON code. • It supports for once, poll and sample and not supported for not on_change subscriptions. • It supports for DME, YANG, and OpenConfig data sources. • For the properties which are not supported, the command defaults back to the collection time instead of the last database change time. • This command generates verbose responses for each timestamp. If you are not able to collect the messages on time, the switch drops the collection of messages. |
| Step 5 | (Optional) grpc gnmi keepalive-timeout <timeout> Example: <pre>switch(config)# grpc gnmi keepalive-timeout 1200</pre> | <p>Configuring this command, you can delete inactive or unauthorized connections. The gRPC agent periodically sends an empty response to the client. If this response fails to deliver to the client, then the connection stops.</p> <p>The default interval value is 600 seconds. You can configure to change the keepalive interval with the interval range of 600-86400 seconds.</p> |

gNMI RPCs

This section describes each gNMI RPC. You can view use of the reference client *gnmi_cli* in this section.

Capabilities

The capabilities of RPC returns to the list of capabilities of the gNMI service. The response message to an RPC request includes the gNMI service version, version data models, and data encode supported by the server.

Guidelines and Limitations for Capabilities

The following are the guidelines and limitations for capabilities:

Example Client Output

Below mentioned the example of client output for Capabilities. Considering the feature *openconfig* is enabled.

This examples shows the support to YANG model, gNMI version, and the supported encodes.

```
$ ./gnmi_cli -a 172.1.1.1:50051 -ca_crt ./grpc.pem -insecure -capabilities
supported_models: <
  name: "Cisco-NX-OS-device"
  organization: "Cisco Systems, Inc."
  version: "2019-11-13"
>
supported_models: <
  name: "openconfig-acl"
  organization: "OpenConfig working group"
  version: "1.0.0"
>
supported_models: <
  name: "openconfig-bgp-policy"
  organization: "OpenConfig working group"
  version: "4.0.1"
>
.
.
.
supported_models: <
  name: "openconfig-spanning-tree"
  organization: "OpenConfig working group"
  version: "0.2.0"
>
supported_models: <
  name: "openconfig-system"
  organization: "OpenConfig working group"
  version: "0.3.0"
>
supported_models: <
  name: "openconfig-telemetry"
  organization: "OpenConfig working group"
  version: "0.5.1"
>
supported_models: <
  name: "openconfig-vlan"
  organization: "OpenConfig working group"
  version: "3.0.2"
>
supported_models: <
  name: "DME"
  organization: "Cisco Systems, Inc."
>
supported_models: <
  name: "Cisco-NX-OS-Syslog-oper"
  organization: "Cisco Systems, Inc."
  version: "2019-08-15"
>
supported_encodings: JSON
supported_encodings: PROTO
gNMI_version: "0.5.0"
```

About Get

You can use `Get` RPC to retrieve a snapshot of the data tree from the switch. You can request multiple paths in a single request. According to gNMI Path conventions, you can use simple form of XPATH. See [Schema path encoding conventions for gNMI](#). For more information about `GET` operation, see the *Retrieving Snapshots of State Information* section in the [gRPC Network Management Interface](#).

Guidelines and Limitations for Get

The following are guidelines and limitations for `Get`:

- `GetRequest` encode supports only the JSON format.
- For the `GetRequest.type`, only `DataType CONFIG` and `STATE` have correlation and the expression in YANG format. The operation is not supported.
- You cannot have both OpenConfig (OC) YANG and device YANG paths in a single request. A single request has only one path.
- `GetRequest` for root path ("/"): everything from **all** models) is not permitted.
- gNMI `Get` returns all default value. See Report-all mode [RFC 6243 \[4\]](#).
- `Get` does not support the `Cisco-NX-OS-syslog-oper` model.
- To retrieve `openconfig-procmon` data, you can send a query to the path `/system/processes` OR `/system`.



Note Before Cisco NX-OS Release 10.3(x), you cannot retrieve data from the `/system` path.

- The following optionals are not supported:
 - Path prefix
 - Path alias
 - Wildcards in the path
- You can request 10 paths in a single `GetRequest`.
- If the return value of size in `GetResponse` is more than 12 MB, the system displays the error status `grpc::RESOURCE_EXHAUSTED`.
- The maximum receive buffer size of gRPC is 8 MB.
- If you perform an `Get` operation for a switch which has more configuration, the gRPC process consumes all available memory. If the memory exhausts, the below syslog is generated:

```
MTX-API: The memory usage is
reaching the max memory resource limit (3072) MB
```

If the memory exhausts consecutively, the below syslog is generated:

```
The process has become unstable and
the feature should be restarted.
```

Cisco recommends you to restart the gRPC feature for normal functioning of gNMI transactions.

- The maximum number of total concurrent sessions for `Get` is 75% of maximum configured concurrent calls. For an example, if the MTX concurrent calls are configured to 16, the maximum number of total concurrent sessions for `Get` is 12.
- The total number of concurrent sessions for `Get` and `Set` is configured `max gNMI concurrent-1`. For an example, if gNMI concurrent calls are configured to 16, the maximum number of total concurrent sessions for `Get` and `Set` is 15.

About Set

You can use `Set` RPC to change the configuration of the switch. You can delete, replace, and update the switch. All these operations in a single `Set` request that is considered as a transaction which is a successful operation or the switch is set to original state.

The `Set` operations are performed in the order that is specified in the `Set Request` configuration. If a path is requested multiple times, the changes are performed even if the paths overwrite each other path. The final state of the data is achieved with the final operation in the transaction. You can edit the paths that are specified in the `Set Request` such as delete, replace, update fields are configuration data paths.

For more information on `Set` operations, see *Modifying State* section of [gNMI specification](#).

Guidelines and Limitations for Set

The following are guidelines and limitations for `Set`:

- `SetRequest` encode supports only JSON format.
- You cannot have both OpenConfig (OC) YANG and device YANG paths in a single request. A single request has only one path.
- The following optionals are not supported:
 - Path prefix
 - Path alias
 - Wildcards in the path
- You can request 20 paths in a single `SetRequest`.
- The maximum receive buffer size of gRPC is 8 MB.
- The maximum number of total concurrent sessions for `Get` and `Set` is configured maximum gNMI concurrent calls. For an example, if the gNMI concurrent calls are configured to 16, the maximum number of total concurrent sessions for `Get` and `Set` is 15.
- If you perform a `Set::Delete` RPC operation for a switch, if a configuration is bulky, an MTX warning message is generated as shown below:

```
Configuration size for this
namespace exceeds operational limit. Feature may become unstable and require
restart.
```

About Subscribe

You can use `Subscribe` RPC to register a telemetry stream for specific paths. When there is change in the registered paths, the switch pushes the notification to for the change in configuration or the path state.

You can use an optional flag available for `Subscribe` RPC. Beginning with Cisco NX-OS Release 9.3(1), the `UPDATES_ONLY` optional flag is supported and applicable only for `ON_CHANGE` subscriptions. If this optional flag is configured, the switch suppresses the current state and a notification is sent with the first response.

You can use these flags to modify the response to options listed in the following table.

Table 38: Support Metrics for SUBSCRIBE Flags

| Subscription Type | heartbeat_interval | suppress_redundant |
|-------------------|--|---|
| ON_CHANGE | Origin: Device YANG, OpenConfig YANG, DME | N/A |
| SAMPLE | Origin: Device YANG, OpenConfig YANG, DME | Origin: Device YANG, OpenConfig YANG |

Beginning with Cisco NX-OS Release 10.2(3)F, the following optional flags are supported:

- `heartbeat_interval`
- `suppress_redundant`

You can use a `heartbeat_interval` flag to modify the behavior of `suppress_redundant` in a `sample` subscription. In this case, the target generates one telemetry update per `heartbeat_interval`, despite of `suppress_redundant` flag status is set to true. The value is specified as an unsigned 64-bit integer in nano seconds.

You can use the `suppress_redundant` flag for a `sample` subscription. In this case, the set status is true. The target generates a telemetry update message when the value of the path reported has changed the last update is generated. Updates are generated for the individual leaf nodes for which the subscription has changed.

For an example, subscriptions name 'A' and 'B' which has leaf nodes 'C' and 'D' branch from 'B' node. If the value of 'C' is changed, and not the value of 'D'. An update is generated only for 'C' and not for 'D'.

Guidelines and Limitations

The following are guidelines and limitations for Set subscription:

- The following flags are not supported:
 - Aliases
 - `allow_aggregation`
 - `extensions`
 - `prefix`
 - `QoS`

- Make sure that all paths within the same subscription request must have the same `sample` interval. If the same path requires different sample intervals, you must create multiple subscriptions.

gNMI Subscribe Options

The following table lists the subscribe option modes:

Table 39:

| Type | Subscription Type | Description |
|----------------|-------------------|---|
| ONCE | | Subscribe to receive data when and close the session. |
| POLL | | Subscribe to retain an active session. When you raise a poll request, ensure that you enter data for each request. |
| STREAM | SAMPLE | Subscribe to receive data at a specific cadence. The payload value is in nano seconds. 1 second =1000000000. |
| | ON_CHANGE | Subscribe to receive a snapshot and receive data only when there is change in tree. |
| TARGET_DEFINED | | Subscribe to find the best subscription which can be created. |

To Set Modes

Each mode requires 2 settings, such as inside subscription and outside subscription. The following mentioned the combination of subscriptions:

- ONCE: SAMPLE, ONCE
- POLL: SAMPLE, POLL
- STREAM: SAMPLE, STREAM
- ON_CHANGE: ON_CHANGE, STREAM
- TARGET_DEFINED: TARGET_DEFINED, STREAM

Origin

- DME: Subscribing to DME model
- DEVICE: Subscribing to YANG model
- OPENCONFIG: Subscribing to OpenConfig model

Name

- DME: Subscribing to DME model
- Cisco-NX-OS-device: Subscribing to YANG model

Encoding

- JSON: Stream is sent in JSON format.
- PROTO: Stream is sent in original proto format.

Configuring Example of Payload

The following section lists client examples for a payload:

Subscribe to DME Stream/Sample

```
{
  "SubscribeRequest":
  [
    {
      "subscribe":
      {
        "subscription":
        [
          {
            "path":
            {
              "origin": "DME",
              "elem":
              [
                {
                  "name": "sys"
                },
                {
                  "name": "bgp"
                }
              ]
            },
            "mode": "SAMPLE"
          }
        ],
        "mode": "ONCE",
        "allow_aggregation" : false,
        "use_models":
        [
          {
            "name": "DME",
            "organization": "Cisco Systems, Inc.",
            "version": "1.0.0"
          }
        ],
        "encoding": "JSON"
      }
    }
  ]
}
```

Subscribe to OpenConfig YANG/Sample

```

{
  "SubscribeRequest":
  [
    {
      "subscribe":
      {
        "subscription":
        [
          {
            "path":
            {
              "origin": "openconfig",
              "elem":
              [
                {
                  "name": "interfaces"
                }
              ]
            },
            "mode": "SAMPLE",
            "Sample_interval": 10000000000
          }
        ],
        "mode": "ONCE",
        "allow_aggregation" : false,
        "use_models":
        [
          {
            "name": "openconfig-interfaces",
            "organization": "OpenConfig working group",
            "version": "0.8.1"
          }
        ]
      },
      "encoding": "JSON"
    }
  ]
}

```

About Subscribing Custom Syslog Stream

Cisco NX-OS supports the original and OpenConfig model. For gNMI or, subscribe `ON_CHANGE`, a custom YANG model is designed to stream the syslog events. You can configure this feature for the Cisco Nexus 9000 Series switches with minimum memory of 8 GB and above.

Guidelines and Limitations

The following are guidelines and limitations for Syslog Stream:

- An invalid syslog is not supported, such as a syslog with a filter or query condition is not supported.
- Syslog stream supports only below paths:
 - Cisco-NX-OS-Syslog-oper: syslog
 - Cisco-NX-OS-Syslog-oper: syslog messages
- Only stream sample and POLL modes are supported.

- Encoding formats that are supported are JSON and PROTO.

Original YANG Syslog Model

The following example shows the configuration of the YANG Syslog Model:



Note By default, the time-zone field is empty. The time zone field is set when you configure `clock format show-timezone syslog`.

```

PYANG Tree for Syslog Native Yang Model:
>>> pyang -f tree Cisco-NX-OS-infra-syslog-oper.yang module: Cisco-NX-OS-syslog-oper
+--ro syslog
+--ro messages
+--ro message* [message-id]
+--ro message-id int32
+--ro node-name? string
+--ro time-stamp? uint64
+--ro time-of-day? string
+--ro time-zone? string
+--ro category? string
+--ro group? string
+--ro message-name? string
+--ro severity? System-message-severity
+--ro text? string

```

Subscribe Request

The following is an example for the subscribe request:

```

{
  "SubscribeRequest":
  [
    {
      "subscribe":
      {
        "subscription":
        [
          {
            "path":
            {
              "origin": "syslog-oper",
              "elem":
              [
                {
                  "name": "syslog"
                },
                {
                  "name": "messages"
                }
              ]
            },
            "mode": "ON_CHANGE"
          }
        ],
        "mode": "ON_CHANGE",
        "allow_aggregation" : false,
        "use_models":
        [
          {
            "name": "Cisco-NX-OS-Syslog-oper",

```

```

        "organization": "Cisco Systems, Inc.",
        "version": "0.0.0"
    },
    "encoding": "JSON"
}
]
}

```

Example Response

The following is an output response example for PROTO encode:

```

[Subscribe]-----
Sat Aug 24 14:38:06 2019
### Generating request : 1 -----
### Comment : STREAM request
### Delay : 2 sec(s) ...
### Delay : 2 sec(s) DONE

subscribe {
  subscription {
    path {
      origin: "syslog-oper"
      elem {
        name: "syslog"
      }
      elem {
        name: "messages"
      }
    }
    mode: ON_CHANGE
  }

  use_models {
    name: "Cisco-NX-OS-Syslog-oper"
    organization: "Cisco Systems, Inc."
    version: "0.0.0"
  }
  encoding: PROTO
}

Thu Nov 21 14:26:41 2019
Received response 3 -----
update {
  timestamp: 1574375201665688000
  prefix {
    origin: "Syslog-oper"
    elem {
      name: "syslog"
    }
    elem {
      name: "messages"
    }
  }
  ...
  update {
    path {
      elem {
        name: "time-stamp"
      }
    }
    val {
      uint_val: 1574375200000
    }
  }
}

```

```

}
}
update {
  path {
    elem {
      name: "severity"
    }
  }
}
val {
  uint_val: 5
}
}
}

```

/Received -----

The following is an example of output response for JSON encode:

```

[Subscribe]-----
### Reading from file ' testing_b1/stream_on_change/OC_SYSLOG.json '

Tue Nov 26 11:47:00 2019
### Generating request : 1 -----
### Comment : STREAM request
### Delay : 2 sec(s) ...
### Delay : 2 sec(s) DONE
subscribe {
  subscription {
    path {
      origin: "syslog-oper"
      elem {
        name: "syslog"
      }
      elem {
        name: "messages"
      }
    }
    mode: ON_CHANGE
  }
  use_models {
    name: "Cisco-NX-OS-Syslog-oper"
    organization: "Cisco Systems, Inc."
    version: "0.0.0"
  }
}

Tue Nov 26 11:47:15 2019
Received response 5 -----
update {
  timestamp: 1574797636002053000
  prefix {
  }
  update {
    path {
      origin: "Syslog-oper"
      elem {
        name: "syslog"
      }
    }
  }
  val {
    json_val: "[ { \"messages\" : [[
{\"message-id\":657},{\"node-name\": \"task-n9k-1\", \"time-stamp\": \"1574797635000\", \"time-of-day\": \"Nov
26 2019
11:47:15\", \"severity\":3, \"message-name\": \"HDR_L2LEN_ERR\", \"category\": \"ARP\", \"group\": \"ARP\", \"text\": \"arp
[30318] Received packet with incorrect layer 2 address length (8 bytes), Normal pkt with

```

```
S/D MAC: 003a.7d21.d55e ffff.ffff.ffff eff_ifc mgmt0(9), log_ifc mgmt0(9), phy_ifc
mgmt0(9)\",\"time-zone\":\"\"} ]] } ]"
}
}
}
```

References

You can view available clients for gNMI subscriptions. For more information, see https://github.com/influxdata/telegraf/tree/master/plugins/inputs/cisco_telemetry_gnmi.

Troubleshooting gNMI

You can run show commands to view gNMI status. To verify specific gNMI configurations, enter the following commands that are listed in the below table:

Table 40: Show Commands

| Command | Description |
|--|--|
| <code>show grpc gnmi service statistics</code> | Displays the summary of the agent running status, respectively for the management VRF, or the default configured VRF. It also displays: <ul style="list-style-type: none"> • Basic overall counters • Certificate expiration time |
| <code>show grpc gnmi rpc summary</code> | Displays the following: <ul style="list-style-type: none"> • Number of capability RPCs received. • Capability of RPC errors. • Number of Get RPCs received. • Get RPC errors. • Number of Set RPCs received. • Set RPC errors. |

| Command | Description |
|-----------------------------|-------------|
| show grpc gnmi transactions | |

| Command | Description |
|---------|--|
| | <p>The <code>show grpc gnmi transactions</code> command is the densest and contains considerable information. It is history buffer of the most recent 50 gNMI transactions that are received by the switch. As new RPCs come in, the oldest history entry is removed from the end. The following explains what is displayed:</p> <p>This command displays the detailed information. It shows the history of the latest 50 gNMI transactions that are received by the switch. An entry of new RPCs, the history of the oldest entry of RPCs is removed. The following shows the information which is displayed:</p> <ul style="list-style-type: none"> • RPC - This shows the type of RPC that was received (Get, Set, Capabilities) • DataType - For a Get only. It has the values ALL, CONFIG, and STATE. • Session - Displays the unique session-id that is assigned to this transaction. It can be used to correlate data that is found in other log files. • Time In - Displays the timestamp when the RPC was received by the gNMI handler. • Duration - Time delta in milliseconds from receiving the request to giving a response. • Status – Displays the status code of the operation returned to the client (0 = Success, !0 == error) <p>The following shows the data per path within a single gNMI transaction. For a single <code>Get</code> or <code>Set</code>:</p> <ul style="list-style-type: none"> • Subtype – For a <code>Set</code> RPC, it displays the specific operation that is requested per path (Delete, Update, Replace). For <code>Get</code>, there is no subtype. • Dtx – Displays that this path is processed in DTX fast path or not. A dash (-) indicates no, an asterisk (*) indicates yes. • St – Displays the status for this path. The different status that is displayed as shown below: <ul style="list-style-type: none"> • OK: The path is valid and processed by infra successfully. • ERR: The path is either invalid or generated error by infra. • --: The path is not processed yet, or not a |

| Command | Description |
|---------|----------------------------------|
| | valid and not sent to infra yet. |

Configuration Examples for gRPC Commands

gRPC gNMI Service Statistics

The following shows the sample output for the command `show grpc gnmi service statistics`:

```

=====
gRPC Endpoint
=====

Vrf : management
Server address : [::]:50051

Cert notBefore : Mar 13 19:05:24 2020 GMT
Cert notAfter  : Nov 20 19:05:24 2033 GMT

Max concurrent calls : 8
Listen calls : 1
Active calls : 0

Number of created calls : 1
Number of bad calls : 0

Subscription stream/once/poll : 0/0/0

Max gNMI::Get concurrent : 5
Max grpc message size : 8388608
gNMI Synchronous calls : 74
gNMI Synchronous errors : 0
gNMI Adapter errors : 0
gNMI Dtx errors : 0

```

gRPC gNMI rPC Summary

The following shows the sample output for the command `show grpc gnmi service statistics`:

```

=====
gRPC Endpoint
=====

Vrf          : management
Server address : [::]:50051

Cert notBefore : Mar 31 20:55:02 2020 GMT
Cert notAfter  : Apr  1 20:55:02 2020 GMT

Capability rpcs      : 1
Capability errors    : 0
Get rpcs             : 53
Get errors           : 19
Set rpcs             : 23
Set errors           : 8
Resource Exhausted  : 0
Option Unsupported  : 6
Invalid Argument    : 18
Operation Aborted   : 1

```

```
Internal Error      : 2
Unknown Error      : 0
```

| RPC Type | State | Last Activity | Cnt Req | Cnt Resp | Client |
|-----------|--------|----------------|---------|----------|--------|
| Subscribe | Listen | 04/01 07:39:21 | 0 | 0 | |

gRPC gNMI Transactions

The following shows the sample output for the command `show grpc gnmi transactions`:

```
=====
gRPC Endpoint
=====

Vrf          : management
Server address : [::]:50051

Cert notBefore : Mar 31 20:55:02 2020 GMT
Cert notAfter  : Apr  1 20:55:02 2020 GMT

RPC          DataType  Session      Time In          Duration(ms)  Status
-----
Set          -          2361443608   04/01 07:43:49   173           0
subtype: dtx: st: path:
Delete      -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo789]

Set          -          3445444384   04/01 07:43:33   3259          0
subtype: dtx: st: path:
Delete      -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo789]
Delete      -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo790]
...
Delete      -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo807]
Delete      -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo808]

Set          -          2297474560   04/01 07:43:26   186           0
subtype: dtx: st: path:
Update      -          OK /System/ipv4-items/inst-items/dom-items/Dom-list[name=foo]/rt-
items/Route-list[prefix=0.0.0.0/0]/nh-items/NextHop-list[nhAddr=192.168.1.1/32][n
hVrf=foo][nhIf=unspecified]/tag

Set          -          0            04/01 07:43:11   0             3
subtype: dtx: st: path:
Update      -          -- /System/intf-items/lb-items/LbRtdIf-list[id=lo4]/descr
Update      -          ERR /system/processes

Set          -          2464255200   04/01 07:43:05   708           0
subtype: dtx: st: path:
Delete      -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo2]
Replace     -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo3]/descr
Replace     -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo4]/descr
Replace     -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo5]/descr
Update      -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo3]/descr
Update      -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo5]/descr

Set          -          3491213208   04/01 07:42:58   14            0
subtype: dtx: st: path:
Replace     -          OK /System/intf-items/lb-items/LbRtdIf-list[id=lo3]/descr
...

Get          ALL          2293232352   04/01 07:42:35   258           0
```

```

subtype: dtx:  st: path:
-         -      OK  /system

Get          ALL          0          04/01 07:42:33          0          12
subtype: dtx:  st: path:
-         -      --  /intf-items

```

Gathering Debug Logs

gNOI is a child service of the gRPC agent. For more information, see gRPC Agent chapter.

Accounting Log for gNMI

In gNMI, Set RPC changes the configuration on the switch. For the SET requests such as UPDATE, REPLACE, or DELETE configuration, gNMI generates the corresponding accounting logs. These logs include both original request and the changes made on the switch.

You can use `show accounting log` command to view the accounting logs.

The following example shows SetRequest, encoding = JSON to localhost with the following gNMI paths:

```

---
<<<<<< set_delete >>>>>>
[]
<<<<<< set_replace >>>>>>
[] []
<<<<<< set_update >>>>>>
[elem { name: "System"
} elem {
name: "tm-items"
} elem {
name: "certificate-items"
}
] [json_val: "{\"hostname\": \"test\", \"trustpoint\": \"foo\"}"]
]
The SetRequest response is below -----response { path {
elem {
name: "System"
} elem {
name: "tm-items"
} elem {
name: "certificate-items"
}
} op: UPDATE
} timestamp: 1656512303065384369
---

```

The below table shows the options can be configured on the switches. The accounting logs include the following items.

| Item | Description |
|-----------|----------------------|
| Context | Session ID and user |
| Operation | Commit or Abort |
| Database | Running or Candidate |

| Item | Description |
|----------|---|
| ConfigMO | MO tree's text representation. Maximum up to 3000 characters. |
| Status | Success or Failed |

The following shows the sample of accounting logs.

```
Wed Jun 29 14:18:23
2022:type=update:id=1430425712:user=admin:cmd=(COMMIT),database=[candidate],
configMo=[<topSystem childAction="" dn="sys" status="created,modified"><telemetryEntity
childAction="" rn="tm" status="created,modified"><telemetryCertificate childAction=""
hostname="test" rn="certificate" status="created,modified"
trustpoint="foo"/></telemetryEntity></topSystem>] (SUCCESS)
Wed Jun 29 14:18:23
2022:type=update:id=1430425712:user=admin:cmd=(COMMIT:CANDIDATE-TO-RUNNING),
database=[running] (SUCCESS)
```

The below table shows the original received request on the switch.

| Item | Description |
|-----------|--|
| Context | Session ID and user |
| Operation | gNMI:SET:UPDATE, gNMI:SET:REPLACE, gNMI:SET:DELETE, COMMIT:CANDIDATE-TO-RUNNING |
| Source IP | gNMI Client IP |
| Path | gNMI path in the text format |
| Payload | Received JSON Request. Maximum up to 3000 characters. |
| Status | Success or Failed |

```
Wed Jun 29 14:18:23
2022:type=update:id=1430425712:user=admin:cmd=(GNMI:SET:UPDATE),sourceIp=[192.168.1.2],
path=[/System/tm-items/certificate-items],payload=[{"hostname":"test","trustpoint":"foo"}]
(SUCCESS)
```

In case of failed request and based on failed scenario, you cannot view both the logs.

Invalid Requests

If you raise a request which is invalid, this request will be rejected without any configuration changes and only the initial request will be logged. The following shows the example of invalid request logs.

```
Wed Jun 29 14:18:23
2022:type=update:id=1430425712:user=admin:cmd=(GNMI:SET:UPDATE),
sourceIp=[192.168.1.2],path=[/System/tm-items/certificateitems],
payload=[{"hostname":"test","trustpoint":"foo"}] (FAILED)
```

Failed Requests

If you raise a request and if it fails due to various configuration restrictions, in such case both the original and failed configuration request is logged. The following example shows the logs.

```
Wed Jun 29 20:52:15
2022:type=update:id=1429663200:user=admin:cmd=(COMMIT),database=[candidate],
configMo=[<topSystem childAction="" dn="sys"
status="created,modified"><telemetryEntity childAction="" rn="tm"
status="created,modified"><telemetryCertificate childAction="" filename="foo"
hostname="test" rn="certificate" status="created,modified,replaced"
trustpoint="foo"/></telemetryEntity></topSystem>] (FAILED)
```

```
Wed Jun 29 20:52:15
2022:type=update:id=1429663200:user=admin:cmd=(GNMI:SET:REPLACE),
sourceIp=[192.168.1.2],path=[/System/tm-items/certificate-items],
payload=[{"hostname":"test","trustpoint":"foo","filename":"foo"}] (FAILED)
```

If you raise a request and if it fails to commit, in such case the original request is logged with the failed request. The following example shows the logs.

```
Wed Jun 29 14:18:23
2022:type=update:id=1430425712:user=admin:cmd
(COMMIT),database=[candidate],configMo=[<topSystem childAction="" dn="sys"
status="created,modified"><telemetryEntity childAction="" rn="tm"
status="created,modified"><telemetryCertificate childAction="" hostname="test"
rn="certificate" status="created,modified" trustpoint="foo"/></telemetryEntity></topSystem>]
(SUCCESS)
```

```
Wed Jun 29 14:18:23
2022:type=update:id=1430425712:user=admin:cmd=(GNMI:SET:UPDATE),
sourceIp=[192.168.1.2],path=[/System/tm-items/certificate-items],
payload=[{"hostname":"test","trustpoint":"foo"}] (SUCCESS)
```

```
Wed Jun 29 20:34:06
2022:type=update:id=1429665744:user=admin:cmd=(COMMIT:CANDIDATE-TO-RUNNING),
database=[running] (FAILED)
```



CHAPTER 31

gNOI - Operation Interface

- [About gNOI, on page 417](#)
- [Revision History, on page 418](#)
- [Guidelines and Limitations for gNOI, on page 418](#)
- [Configuring gNOI, on page 418](#)
- [System .Proto, on page 419](#)
- [OS .Proto, on page 420](#)
- [Cert .Proto, on page 420](#)
- [File .Proto, on page 421](#)
- [Factory Reset .Proto, on page 421](#)
- [Troubleshooting gNOI, on page 422](#)

About gNOI

gRPC Network Operations Interface (gNOI) defines a set of gRPC-based micro-services for executing operational commands on network devices.

gNOI uses Google Remote Procedure Call (gRPC) as the transport protocol and the configuration is same as that of gNMI. For details on gNMI configuration, see [gRPC Agent, on page 381](#). To send gNOI RPC requests, user needs a client that implements the gNOI client interface for each RPC. In Cisco NX-OS Release 10.1(1) the gNOI defines Remote Procedure Calls (RPCs) for a limited number of components and some of them are related to hardware (like optical interfaces).

Proto files are defined for the gRPC micro-services and are available at [GitHub](#).

Table 41: Supported gNOI RPCs

| Proto | gNOI RPC | Supported |
|--------|-------------------------|-----------|
| System | Ping | Yes |
| | Traceroute | Yes |
| | Time | Yes |
| | SwitchControl Processor | Yes |
| | Reboot | Yes |
| | RebootStatus | Yes |
| | CancelReboot | Yes |
| OS | Activate | Yes |
| | Verify | Yes |
| Cert | LoadCertificate | Yes |
| File | Get | Yes |
| | Stat | Yes |
| | Remove | Yes |

Revision History

| Release | Description |
|---------|--|
| 10.1(1) | gNOI RPCs are implemented with the equivalent CLI. The existing CLI restrictions or valid options remain as applicable |
| 10.2(1) | Add support for file.proto and cert.proto |

Guidelines and Limitations for gNOI

The gNOI feature has the following guidelines and limitations:

- A maximum of 16 active gNOI RPCs are supported.
- The Cisco Nexus 9000 series switches would run one endpoint with one gNMI service and two gNOI microservices.

Configuring gNOI

gNMI is a child functionality of the gRPC agent. See [gRPC Agent, on page 381](#), to enable the gRPC agent. Currently there is no separate configuration for gNOI.

System.Proto

The System proto service is a collection of operational RPCs that allows the management of a target outside the configuration and telemetry pipeline.

The following are the RPC support details for System proto:

| RPC | Support | Description | Limitation |
|-------------------------|------------------------------------|---|--|
| Ping | ping/ping6 cli command | Executes the ping command on the target and streams back the results. Some targets may not stream any results until all results are available. If a packet count is not explicitly provided, ping5 is used. | do_not_resolve option is not supported. |
| Traceroute | traceroute/traceroute6 cli command | Executes the traceroute command on the target and streams back the results. Some targets may not stream any results until all results are available. Max hop count of 30 is used. | itlial_ttl, marx_ttl, wait, do_not_fragment, do_not_resolve and l4protocol options are not supported. |
| Time | local time | Returns the current time on the target. Typically used to test if the target is responding. | - |
| SwitchControl Processor | system switchover cli command | Switches from the current route processor to the provided route processor. Switchover happens instantly and the response may not be guaranteed to return to the client. | Switchover occurs instantly. As a result, the response may not be guaranteed to return to the client. |
| Reboot | cli: reload [module] | Causes the target to reboot. | message option is not supported. Delay option is supported for switch reload, and the path option accepts one module number. |
| RebootStatus | show version [module] cli command | Returns the status of the reboot for the target. | - |
| CancelReboot | reload cancel | Cancel any pending reboot request. | - |



Note The SetPackage RPC is not supported.

OS .Proto

The OS service provides an interface for OS installation on a Target. The OS package file format is platform dependent. The platform must validate that the OS package that is supplied is valid and bootable. This must include a hash check against a known good hash. It is recommended that the hash is embedded in the OS package.

The Target manages its own persistent storage, and OS installation process. It stores a set of distinct OS packages, and always proactively frees up space for incoming new OS packages. It is guaranteed that the Target always has enough space for a valid incoming OS package. The currently running OS packages must never be removed. The Client must expect that the last successfully installed package is available.

The following are the RPC support details for OS proto:

| RPC | Support | Description | Limitation |
|----------|---|--|---|
| Activate | install all nxos bootflash:///img_name | Sets the requested OS version as the version that is used at the next reboot. This RPC reboots the Target. | Cannot rollback or recover if the reboot fails. |
| Verify | show version | Verify checks the running OS version. This RPC may be called multiple times while the Target boots until it is successful. | - |



Note The Install RPC is not supported.

Cert .Proto

The certificate management service is exported by targets. Rotate, Install and other Certificate Proto RPCs are not supported.

| RPC | Support | Description | Limitation |
|-----------------|---|------------------------------------|------------|
| LoadCertificate | crypto ca import <trustpoint> pkcs12 <file> <passphrase> | Loads a bundle of CA certificates. | - |

File .Proto

The file proto streams messages based on the features of the file.proto RPCs. Put and other RPCs that are not listed here are not supported in File Proto.

Get, Stat, and Remove RPCs support file systems such as - bootflash, bootflash://sup-remote, logflash, logflash://sup-remote, usb, volatile, volatile://sup-remote and debug.

The following are the RPC support details for File proto:

| RPC | Support | Description | Limitation |
|--------|---------|--|-----------------------------------|
| Get | | Get reads and streams the contents of a file from the target. The file is streamed by sequential messages, each containing up to 64 KB of data. A final message is sent prior to closing the stream that contains the hash of the data sent. An error is returned if the file does not exist or there was an error reading the file. | Maximum file size limit is 32 MB. |
| Stat | | Stat returns metadata about a file on the target. An error is returned if the file does not exist or if there is an error in accessing the metadata. | - |
| Remove | | Remove removes the specified file from the target. An error is returned if the file does not exist, if it is a directory, or the remove operation encounters an error. | - |

Factory Reset .Proto

This .proto currently defines only one RPC. Refer to https://github.com/openconfig/gnoi/blob/master/factory_reset/factory_%20reset.proto.

| RPC | Support | Description | Limitation |
|--------------|---------|---|---|
| FactoryReset | | Executes the ping command on the target and streams back the results. Some targets may not stream any results until all results are available. If a packet count is not explicitly provided, ping5 is used. | do_not_resolve option is not supported. |

FactoryReset

The gNOI factory reset operation erases all persistent storage on the specified module. This includes configuration, all log data, and the full contents of flash and (Solid State Drives) SSDs. The reset boots to the last boot image, erases all storage including license. gNOI factory reset supports two modes:

- A fast erase which can reformat and repartition only.
- A secure erase which can erase securely and wipe the data which is impossible to recover.

| Option | Description | Values |
|------------|--|---|
| factory_os | Specifies to rollback to the OS version as shipped from factory. | Setting to true on NX-OS is not supported, and it is mandatory to preserve the current boot image. |
| zero_fill | Specifies whether to perform more time consuming and comprehensive secure erase. | zero_fill = true: Specifies factory-reset module all preserve-image force. zero_fill = false: Specifies factory-reset module all bypass-secure-erase preserve-image force. |

Troubleshooting gNOI

Debug gNOI

To verify the gNOI status, enter the following commands.

Show Commands

| Command | Description |
|---------------------|---|
| clear grpc gnoi rpc | Serves to clean up the counters or calls. |

| Command | Description |
|--|---|
| debug grpc events {events errors} show grpc nxsdk event-history {events errors} | Debugs the events and errors from the event history. |
| show grpc internal gnoi service statistics | Display gNOI service statistics |
| show grpc internal gnoi rpc {summary detail} | An internal keyword command added for serviceability. |
| clear grpc gnoi rpc | Serves to clean up the counters or calls. |

Example Output

show grpc gnmi service statistics

```
=====
gRPC Endpoint
=====
```

```
Vrf          : management
Server address : [::]:50051
```

```
Status          : Running - certificate expired
Cert notBefore  : Jun 20 16:43:49 2023 GMT
Cert notAfter   : Jun 21 16:43:49 2023 GMT
Client Root Cert notBefore : n/a
Client Root Cert notAfter  : n/a
```

```
Max concurrent calls      : 16
Active calls              : 0
```

show grpc internal gnoi rpc all summary

```
=====
gRPC Endpoint
=====
```

```
Vrf          : management
Server address : [::]:50051
```

```
RPC Type      State      Last Activity  Cnt Req  Cnt Resp  Client
-----
system.ping   End        01/12 20:22:06      1        6  ipv4:171.68.196.210:53222
system.time   Listen     01/12 20:21:57      0         0
```

Gathering Debug Logs

gNOI is a child service of the gRPC agent. For more information, see [gRPC Agent, on page 381](#) chapter.



CHAPTER 32

gRPC Tunnel

- [About gRPC Tunnel, on page 425](#)
- [Guidelines and Limitations, on page 425](#)
- [Configuring gRPC Tunnel, on page 426](#)
- [Troubleshooting, on page 433](#)

About gRPC Tunnel

Cisco NX-OS supports “gRPC Tunnel”, a specific implementation of application layer tunnel to allow external communication across network segregations like firewall.

In common network deployment, if we use firewall to roughly segregate the network to either "outside" or "inside" of the firewall. Then, for example in gNMI, the "gnmi clients (controllers)" are usually outside, while the "gnmi servers (networking switches)" are inside of the firewall. Therefore, a gNMI connection is referred as an inbound RPC to the switch, and usually referred as "Dial-In".

Such "Dial-In" results in two limitations:

- **Firewall provision:**

"Dial-In" requires punching holes in the firewall to allow gNMI connection to go through by specifying firewall rules for specific hosts, addresses, or ports, just to name a few.

- **Prerequisite network inventory:**

To "Dial-in" to a specific destination, an initial solicitation step is required, either manually or programmatically by the user to collect the host information of gNMI servers.

Only after that, the client can know the address/port to "Dial-In" to those servers.

The grpc-tunnel intends to work around these two limitations, as it established tunnels using the opposite "Dial-out" approach. For more information about grpc tunnel, please refer to the external document at [gRPC Tunnel](#)

Guidelines and Limitations

The gRPC tunnel has the following guidelines and limitations:

- The naming conventions when assigning a target identifier for a tunnel is completely up to the user.

- The user is responsible to make sure the naming convention of the target identifier is unique. It is recommended that an automated deployment workflow should handle the uniqueness of the target identifier.
- Cisco NX-OS supports up to 8 tunnel configurations.

Topic 2.1

Configuring gRPC Tunnel

Configuring gRPC Tunnel without authentication

This procedure describes how to enable and configure the gRPC Tunnel without either server or client authentication. This is mostly for experimental usage.

Before you begin

gRPC Tunnel is an opaque tunnel which supposedly should be able to forward various network traffic. However, in Cisco NX-OS, the primary use case is to proxy the gNMI/gNOI requests. If that is case, it requires to properly config the grpc agent. Please refer to the gRPC Agent programming guide.

SUMMARY STEPS

1. **configure terminal**
2. **[no] feature grpctunnel**
3. **[no] grpctunnel destination**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|---|--|
| Step 1 | configure terminal Example: <pre>switch-1# configure terminal</pre> | Enters global configuration mode. |
| Step 2 | [no] feature grpctunnel Example: <pre>switch-1# feature grpctunnel</pre> | Enables or disables grpc-tunnel feature. |
| Step 3 | [no] grpctunnel destination Example: <pre>switch-1# grpctunnel destination 1.1.1.1 port 8000 target test1 type GNMI_GNOI use-vrf management</pre> | Configure a tunnel connection. The no form of this command removes the tunnel connection. <ul style="list-style-type: none"> • destination - (Type: IPv4/IPv6 address or hostname string) Tunnel server ip address or the hostname. If |

| | Command or Action | Purpose |
|--|-------------------|---|
| | | <p>hostname is given, a valid name-server config is required.</p> <ul style="list-style-type: none"> • port - (Type: tcp port) Tunnel server port number. • target - (Type: string, 64 bytes limit) Target ID is a string. If the user sets the ID as the reserved keyword 'HOSTNAME', the switch will substitute the switch hostname as the target. • type - (Type: string, 64 bytes limit) Type only supports GNMI_GNOI in 10.3.2F release. • use-vrf - (Type: string) The vrf name string that switch will use to dial out for grpc tunnel session. • [Optional] source-interface - (Type: interface name string) source-interface is used to determine the egress source ip address of the tunnel establishment. The switch would select the first ipv6 global unicast address of the interface. Else, it would select the ipv4 unicast address of the interface. This configuration supports loopback and svi interfaces only. The interface must be specified in the short name format such as Lo10, Vlan100. • [Optional] target-vrf - (Type: string) vrf name is used to reach local grpc server target. If not specified, uses the same as the vrf parameter. For example, specifying <code>grpc tunnel ... use-vrf foo ... target-vrf bar</code> means the switch establishes connection to the external tunnel server in vrf foo, but forwards incoming grpc requests to the local switch grpc server residing on vrf bar. |

Configuring gRPC Tunnel with Trustpoints

Note that for gRPC Tunnel, the Cisco NX-OS device initiates a connection to specified external gRPC Tunnel Destination. The user may configure the trustpoints/certificates to secure such outbound connections.

- Server authentication with “cert” option This configures the external destination’s cert into the switch. The switch would refuse to the connection if the configured cert does not match to the remote tunnel
- Client authentication with “client-cert” option This configures the identity cert of the switch. The remote tunnel server would reject the connection from the switch if the switch could not present the matching certificate.
- Mutual authentication is combination of both “cert” and “client-cert” authentication

In the below steps, step 1-3 means to import the “server authentication” while steps 4-5 mean to import the “client authentication”. The user can decide the proper combination to either enable either, or both.



Note Configuring or removing the root certificate for client authentication will cause gRPC tunnel to restart the connection to the remote destination.



Note If the client's certificate is signed by intermediate CAs, but not directly by the root CA that is imported from the above config, the grpc tunnel certification (step 5) needs to supply the full cert chain, including the user, intermediate CA cert, and the root CA cert.

Before you begin

Prepare and sign and the required certificate files for the server authentication. This is not specific to gRPC tunnel, so it is possible to re-use the existing trustpoint files.

This section means to particularly clarify the usage of the 'cert' and 'client-cert' options. The option describes the previous section can freely combine with these two options.

SUMMARY STEPS

1. **configure terminal**
2. (Optional) **crypto ca trustpoint** *<tunnel-trustpoint>*
3. (Optional) **rsa keypair** *<client-key>*
4. (Optional) **crypto ca authenticate** *<tunnel-trustpoint>*
5. (Optional) **crypto ca trustpoint** *<tunnel-client-trustpoint>*
6. (Optional) **crypto ca import** *<tunnel-client-trustpoint>* **pkcs12 bootflash** *:<ca-file> <pkcs-password>*
7. **[no] grpctunnel destination ...**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | configure terminal Example: switch# configure terminal | Enters global configuration mode. |
| Step 2 | (Optional) crypto ca trustpoint <i><tunnel-trustpoint></i> Example: switch# crypto ca trustpoint tunnel_trustpoint | Create a trustpoint for tunnel server authentication. |
| Step 3 | (Optional) rsa keypair <i><client-key></i> Example: switch# rsa keypair key | Generate a rsa key pair for the client trustpoint. |

| | Command or Action | Purpose |
|--------|---|--|
| Step 4 | (Optional) crypto ca authenticate <tunnel-trustpoint> Example: switch# crypto ca authenticate tunnel_trustpoint | Import the tunnel server certificate. This step requires manual copy paste. Please follow the instruction. |
| Step 5 | (Optional) crypto ca trustpoint <tunnel-client-trustpoint> Example: switch# crypto ca trustpoint tunnel_client_trustpoint | Create a trustpoint for tunnel client authentication. |
| Step 6 | (Optional) crypto ca import <tunnel-client-trustpoint> pkcs12 bootflash :<ca-file> <pkcs-password> Example: switch# crypto ca import tunnel_client_trustpoint pkcs12 bootflash:ca.pfx test | Import the pkcs12 file to the trustpoint. |
| Step 7 | [no] grpctunnel destination ... Example: switch(config)# grpctunnel destination 1.1.1.1 port 8000 target test1 type GNMI_GNOI cert tunnel_trustpoint client-cert tunnel_client_trustpoint ... | Configure a tunnel connection. The no form of this command removes the tunnel connection. <ul style="list-style-type: none"> • [Optional] cert - (Type: string) Trustpoint which holds the tunnel server certificate. If not specified, would skip the server verification. • [Optional] client-cert - (Type: string) Trustpoint which holds the client certificate. If specified, would exercise mutual authentication with the tunnel server. |

Configuring gRPC Tunnel with VRF

In the below steps, step 1-3 means to import the “server authentication” while steps 4-5 mean to import the “client authentication”. The user can decide the proper combination to either enable either, or both.

Before you begin

This section means to particularly clarify the usage of the ‘use-vrf and ‘target-vrf’ options. The option describes the previous section can freely combine with these two options.

SUMMARY STEPS

1. **configure terminal**
2. **[no] feature grpctunnel**
3. **[no] grpctunnel destination ...**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | configure terminal Example: switch-1# configure terminal | Enters global configuration mode. |
| Step 2 | [no] feature grpctunnel Example: switch(config)# feature grpctunnel | Enables or disables grpc-tunnel feature. |
| Step 3 | [no] grpctunnel destination ... Example: switch(config)# switch(config)# grpctunnel destination 1.1.1.1 port 8000 target test1 type GNMI_GNOI use-vrf management target-vrf default | Configure a tunnel connection. The no form of this command removes the tunnel connection. <ul style="list-style-type: none"> • use-vrf - (Type: string) The vrf name string that switch will use to dial out for grpc tunnel session. • [Optional] target-vrf - (Type: string) vrf name is used to reach local grpc server target. If not specified, uses the same as the vrf parameter. For example, specifying <code>grpctunnel ... use-vrf foo ... target-vrf bar</code> means the switch establishes connection to the external tunnel server in vrf foo, but forwards incoming grpc requests to the local switch grpc server residing on vrf bar. |

Example

This section lists a few configuration examples to illustrate the tunnel usage.

Without authentication

The following steps describe how to configure the tunnel destinations without server validation.

```
switch(config)# grpctunnel destination 1.1.1.1 port 8000 target test1 type GNMI_GNOI use-vrf
management
switch(config)# grpctunnel destination server.foo.com port 8000 target test2 type GNMI_GNOI
use-vrf management
```

In this example, the user configures two tunnel destinations “1.1.1.1:8000” and “server.foo.com:8000” with target “test1” and “test2” respectively. The connections are initiated over the management namespace.

With server authentication

The following steps describe how to configure the tunnel destinations with server validation.

Execute the following commands to Import server cert to the trustpoint.

```
switch(config)# crypto ca trustpoint tunnel_server_trustpoint switch(config-trustpoint)#
crypto ca authenticate tunnel_server_trustpoint
input (cut & paste) CA certificate (chain) in PEM format; end the input with a line containing
only END OF INPUT :
```

```

-----BEGIN CERTIFICATE-----
MIIC3TCCAcWgAwIBAgIJA04xEeL+IrpMA0GCSqGSIb3DQEBCwUAMBcxFTATBgNV
BAMMDHNqYy1hZHMtNjAxND AeFw0yMjAlMjYwMDE4MzBaFw0zMjAlMjYwMDE4MzBa
MBcxFTATBgNVBAMMDHNqYy1hZHMtNjAxNDCCASIwDQYJKoZIhvcNAQEBBQADggEP
ADCCAQoCggEBALudrG824XmW/4+BNd632CT3x47akV0QfjwAU1xBDScpAw9brERO
YTLp9BxInbA+WAS+zGq16nmBoZxbqZZL/NVD81tLKYYJxtDQHJkqkX21URnMUFr2
9pyJQtuh/udq9hp8zGcEpbPayfIdHCnZqraWMLvk1W0mqAa7ek0iizIZNwKmU3oR
7CGQOxi8aMsAfH5iBsRTNURFdaXdJYTOjry0il+jBKT21F2Z3vGcB7ddTt+I7qrd
GjJs4BI4a22Y3usYb/dnsEa0ZCFTFIq6Y2Pwc3DOuKalUhuJsqisqfMDuqC34ATw
kWwLnHDWVu0iVaWndy3uvQZKDNv/bIuoo8CAwEAAAMsMCowFwYDVR0RBBAwDoIM
c2pjLWFkcy02MDE0MA8GA1UdEwEB/wQFMAMBAf8wDQYJKoZIhvcNAQELBQADggEB
AIjNgq/paYfPtHDe9PlZKzrmGz+U1UAx8saj2WHtrKgBj48J6fYvz1yTPWLKMPct
/5y+nhia6gRlV/navFcpIUUQGpOzQnaa40/nkBMDvVxnTu619UC0WUAYTh217ec
BriY8yq3elpQWHZS4KRnMBH8fuviAv4f0fzOAuNGeIuv7UGnfa8Ed/q/Z3frQxOI
qNxr3vBBTptYTLwdrRM0axagL6waZgZyTffFHpIXBPEtsXKb/5GuP4+nqXvtfkfe
d6P9jA4BKA/e6Gu6NAR0JMOdmJeEFjMbg+uu8jghcRTcwrRsGeb9DqPUL+5IsVg3a
-----END CERTIFICATE-----
END OF INPUT
Fingerprint(s): SHA1 Fingerprint=D4:9D:79:5B:8B:38:D6:50:6D:46:89:A8:C4:41:AB:
C9:D9:9F:D1:66
Do you accept this certificate? [yes/no]:yes

```

Then execute the following command to configure the tunnel destination.

Also use the show command to display the configuration.

```

switch(config)# grpctunnel destination 1.1.1.1 port 8000 target test1 type GNMI_GNOI use-vrf
management cert tunnel_server_trustpoint

```

```

switch(config)# show system internal dme running-config all dn sys/grpctunnel {
"grpctunnelInst": {
"attributes": {
"childAction": "",
"dn": "sys/grpctunnel",
"modTs": "2022-12-02T12:57:37.891+00:00",
"status": ""
},
"children": [
{
"grpctunnelTunnelMgr": {
"attributes": {
"childAction": "",
"dn": "sys/grpctunnel/tunnelmgr",
"modTs": "2022-12-02T12:57:37.891+00:00",
"status": ""
},
"children": [
{
"grpctunnelTunnel": {
"attributes": {
"cert": "tunnel_server_trustpoint",
"certClient": "",
"childAction": "", "dest": "1.1.1.1", "dn":
"sys/grpctunnel/tunnelmgr/tunnel-[1.1.1.1]-port-[8000]-target-[test1]-type-[GNMI_GNOI]-vrf-[management]",
"modTs": "2022-12-05T10:09:45.163+00:00",
"port": "8000",
"srcIf": "unspecified",
"status": "",
"targetId": "test1",
"targetType": "GNMI_GNOI",
"targetVrf": "",
"vrf": "management"
}
}
}
}
}
}
}

```

```

]
}
}
]
}
}
}

```

With client authentication

The following steps describe how to configure the tunnel destinations with client validation.

The following steps describe how to configure the tunnel destinations without server validation.

```

switch(config)# crypto ca trustpoint tunnel_client_trustpoint
switch(config)# crypto ca import tunnel_client_trustpoint pkcs12 bootflash://ca.pfx test

```

Then execute the following command to configure the tunnel destination.

Also use the show command to display the configuration.

```

switch(config)# grpctunnel destination 1.1.1.1 port 8000 target test1 type GNMI_GNOI use-vrf
management client-cert tunnel_client_trustpoint

```

```

switch(config)# show system internal dme running-config all dn sys/grpctunnel {
  "grpctunnelInst": {
    "attributes": {
      "childAction": "",
      "dn": "sys/grpctunnel",
      "modTs": "2022-12-02T12:57:37.891+00:00",
      "status": ""
    },
    "children": [
      {
        "grpctunnelTunnelMgr": {
          "attributes": {
            "childAction": "",
            "dn": "sys/grpctunnel/tunnelmgr",
            "modTs": "2022-12-02T12:57:37.891+00:00",
            "status": ""
          },
          "children": [
            {
              "grpctunnelTunnel": {
                "attributes": {
                  "cert": "",
                  "certClient": "tunnel_client_trustpoint ",
                  "childAction": "", "dest": "1.1.1.1", "dn":
                    "sys/grpctunnel/tunnelmgr/tunnel-[1.1.1.1]-port-[8000]-target-[test1]-type-[GNMI_GNOI]-vrf-[management]",
                  "modTs": "2022-12-05T10:09:45.163+00:00",
                  "port": "8000",
                  "srcIf": "unspecified",
                  "status": "",
                  "targetId": "test1",
                  "targetType": "GNMI_GNOI",
                  "targetVrf": "",
                  "vrf": "management"
                }
              }
            }
          ]
        }
      }
    ]
  }
}

```

With VRF

The combination of 'use-vrf' and 'target-vrf' config offers deployment flexibility but may also incur confusions.

Please note the following difference.

- use-vrf → How to reach the remote tunnel destination.
- target-vrf → How to forward/relay the tunnel traffic to the switch's internal service.

Please refer to the below example scenarios:

- The remote tunnel server is reachable via the 'management' vrf. When the switch received a gNMI connection within the tunnel, the switch would forward to the gnmi 'management' server.

```
grpctunnel destination server1 port 9000 target target2 type GNMI_GNOI vrf management
```

- The remote tunnel server is reachable via the 'management' vrf, while the local grpc agent is running on the default vrf. With the below config, When the switch received a gNMI connection within the tunnel, the switch would stich the gnmi request to the to default vrf.

```
grpc use-vrf default
grpctunnel destination server1 port 9000 target target2 type GNMI_GNOI use-vrf management
target-vrf default
```

- Both the remote tunnel server and the local grpc agent are running on the default vrf.

```
grpc use-vrf default
grpctunnel destination server1 port 9000 target target2 type GNMI_GNOI use-vrf default
```

- The remote tunnel server is reachable via the 'default' vrf, while the local grpc agent is running on the 'test' vrf. With the below config, When the switch received a gNMI connection within the tunnel, the switch would stich the gnmi request to the to test vrf.

```
grpc use-vrf test
grpctunnel destination server1 port 9000 target target2 type GNMI_GNOI vrf default
local-vrf test
```

- In this case, the remote tunnel server is reachable via the 'default' vrf, while the local grpc agent is running on the 'abc' vrf. With the below config, When the switch received a gNMI connection within the tunnel, the switch would stich the gnmi request to the to test vrf, thus the connection would not work. This can be treated as a forward reference. The connection would start to work after changing grpc config to 'grpc use-vrf test'.

```
grpc use-vrf abc
grpctunnel destination server1 port 9000 target target2 type GNMI_GNOI vrf default
local-vrf test
```

Troubleshooting

Check Feature Status

- In Cisco NX-OS, enter the **show feature grpctunnel** command to check the agent config.
- To view the status of the gRPC tunnel, use the **show feature** command.

```
switch-1# show feature | grep grpctunnel
restconf 1 enabled
switch-1#
```

Debug gRPC Tunnel

There exist a series of show commands to display tunnel status.

Show Commands

To verify the tunnel configuration/status, enter the following command:

| Command | Description |
|--|--|
| show grpctunnel internal sessions [all] { summary detail } } | Displays the tunnel status. <ul style="list-style-type: none"> The 'sessions' option would list the tu The 'all' option would list the ended sessio For ended sessions, the switch would retai |
| debug grpctunnel events all | Executed in CLI EXEC mode This allows to display the debug message |

Example Output

show grpctunnel internal sessions summary

```
=====
gRPC Tunnel
=====
Restart Count : 1
* - history
Destination                               Target/Type                               Cnt
  Retry  Cnt Sess  Status/Error
-----
1.1.1.1:8080 (management)                   test/GNMI_GNOI
           0           0 NOT CONNECTED - Dialing [1.1.1.1]:8080
```

Gathering gRPC Tunnel Logs

The /volatile directory houses the grpc tunnel log

```
bash-4.3# cd /volatile/ bash-4.3# ls /volatile -al
...
-rw-rw-rw- 1 root root 103412 Jun 21 16:14 grpc-internal-tunnel-log
...
```




CHAPTER 33

Telemetry

- [About Telemetry, on page 435](#)
- [Telemetry Components and Terminology, on page 435](#)
- [High Availability of the Telemetry Process, on page 436](#)
- [Licensing Requirements for Telemetry, on page 436](#)
- [Guidelines and Limitations for Telemetry, on page 437](#)
- [Configuring Telemetry Using the CLI, on page 443](#)
- [Configuring Telemetry Using the NX-API, on page 462](#)
- [Cloud Scale Software Telemetry, on page 477](#)
- [Telemetry Path Labels, on page 478](#)
- [Native Data Source Paths, on page 497](#)
- [Streaming Syslog, on page 508](#)
- [Troubleshooting Telemetry, on page 513](#)

About Telemetry

Collecting data for analyzing and troubleshooting has always been an important aspect in monitoring the health of a network.

Cisco NX-OS provides several mechanisms such as SNMP, CLI, and Syslog to collect data from a network. These mechanisms have limitations that restrict automation and scale. One limitation is the use of the pull model, where the initial request for data from network elements originates from the client. The pull model does not scale when there is more than one network management station (NMS) in the network. With this model, the server sends data only when clients request it. To initiate such requests, continual manual intervention is required. This continual manual intervention makes the pull model inefficient.

A push model continuously streams data out of the network and notifies the client. Telemetry enables the push model, which provides near-real-time access to monitoring data.

Telemetry Components and Terminology

Telemetry consists of four key elements:

- **Data Collection** — Telemetry data is collected from the Data Management Engine (DME) database in branches of the object model specified using distinguished name (DN) paths or YANG infra in branches

of paths. The data can be retrieved periodically (frequency-based) or only when a change occurs in any object on a specified path (event-based). You can use the NX-API to collect frequency-based data.

- **Data Encoding** — The telemetry encoder encapsulates the collected data into the desired format for transporting. NX-OS encodes telemetry data in the Google Protocol Buffers (GPB) and JSON format. The GPB encoder stores data in a generic key-value format. The encoder requires metadata in the form of a compiled `.proto` file to translate the data into GPB format.
- **Data Transport** — NX-OS transports telemetry data using HTTP for JSON encoding and the Google remote procedure call (gRPC) protocol for GPB encoding. The gRPC receiver supports message sizes greater than 4 MB. (Telemetry data using HTTPS is also supported if a certificate is configured.)
- **Telemetry Receiver** — A telemetry receiver is a remote management system or application that stores the telemetry data. In order to receive and decode the data stream correctly, the receiver requires the `.proto` file that describes the encoding and the transport services. The encoding decodes the binary stream into a key value string pair. A telemetry `.proto` file that describes the GPB encoding and gRPC transport is available on Cisco's GitLab: <https://github.com/CiscoDevNet/nx-telemetry-proto>



Note In the telemetry context, gRPC refers to a specific proprietary `.proto`. Please don't confuse with the "gRPC Agent" which hosts the standard gNMI/gNOI services.

High Availability of the Telemetry Process

High availability of the telemetry process is supported with the following behaviors:

- **System Reload** — During a system reload, any telemetry configuration and streaming services are restored.
- **Supervisor Failover** — Although telemetry is not on hot standby, telemetry configuration and streaming services are restored when the new active supervisor is running.
- **Process Restart** — If the telemetry process freezes or restarts for any reason, configuration and streaming services are restored when telemetry is restarted.

Licensing Requirements for Telemetry

Table 42: Licensing Requirements for Telemetry

| Product | License Requirement |
|-------------|--|
| Cisco NX-OS | Telemetry requires no license. Any feature not included in a license package is bundled with the Cisco NX-OS image and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> . |

Guidelines and Limitations for Telemetry

Telemetry has the following configuration guidelines and limitations:

- For information about supported platforms, see the Nexus Switch Platform Matrix.
- Cisco NX-OS releases that support the data management engine (DME) Native Model support Telemetry.
- Support is in place for the following:
 - DME data collection
 - NX-API data sources
 - Google protocol buffer (GPB) encoding over Google Remote Procedure Call (gRPC) transport
 - JSON encoding over HTTP
- The smallest sending interval (cadence) supported is five seconds for a depth of 0. The minimum cadence values for depth values greater than 0 depends on the size of the data being streamed out. Configuring any cadences below the minimum value may result in undesirable system behavior.
- Telemetry supports up to five remote management receivers (destinations). Configuring more than five remote receivers may result in undesirable system behavior.
- Telemetry can consume up to 20% of the CPU resource.
- Beginning with Cisco NX-OS Release 10.4(2)F, telemetry is supported on the Cisco Nexus 93400LD-H1 platform switches.
- Beginning with Cisco NX-OS Release 10.4(3)F, Telemetry is supported on 92348GC-X.
- Beginning with Cisco NX-OS Release 10.4(2)F, telemetry is supported on the Cisco Nexus N9KC9364C-H1 platform switches.

Configuration Commands After Downgrading to an Older Release

After a downgrade to an older release, some configuration commands or command options can fail because the older release may not support them. When downgrading to an older release, unconfigure and reconfigure the telemetry feature after the new image comes up. This sequence avoids the failure of unsupported commands or command options.

The following example shows this procedure:

- Copy the telemetry configuration to a file:

```
switch# show running-config | section telemetry
feature telemetry
telemetry
destination-group 100
ip address 1.2.3.4 port 50004 protocol gRPC encoding GPB use-chunking size 4096
sensor-group 100 path sys/bgp/inst/dom-default depth 0
subscription 600
dst-grp 100
snsr-grp 100 sample-interval 7000
switch# show running-config | section telemetry > telemetry_running_config
switch# show file bootflash:telemetry_running_config
```

```

feature telemetry
telemetry
destination-group 100
ip address 1.2.3.4 port 50004 protocol gRPC encoding GPB use-chunking size 4096
sensor-group 100 path sys/bgp/inst/dom-default depth 0
subscription 600
dst-grp 100
snsr-grp 100 sample-interval 7000
switch#

```

- Execute the downgrade operation. When the image comes up and the switch is ready, copy the telemetry configurations back to the switch.

```

switch# copy telemetry_running_config running-config echo-commands
`switch# config terminal`
`switch(config)# feature telemetry`
`switch(config)# telemetry`
`switch(config-telemetry)# destination-group 100`
`switch(config-tm-dest)# ip address 1.2.3.4 port 50004 protocol gRPC encoding GPB `
`switch(config-tm-dest)# sensor-group 100`
`switch(config-tm-sensor)# path sys/bgp/inst/dom-default depth 0`
`switch(config-tm-sensor)# subscription 600`
`switch(config-tm-sub)# dst-grp 100`
`switch(config-tm-sub)# snsr-grp 100 sample-interval 7000`
`switch(config-tm-sub)# end`
Copy complete, now saving to disk (please wait)...
Copy complete. switch#

```

gRPC Error Behavior

The switch client disables the connection to the gRPC receiver if the gRPC receiver sends 20 errors. Unconfigure then reconfigure the receiver's IP address under the destination group to enable the gRPC receiver.

Errors include:

- The gRPC client sends the wrong certificate for secure connections.
- The gRPC receiver takes too long to handle client messages and incurs a timeout. Avoid timeouts by processing messages using a separate message processing thread.

Support for gRPC Chunking

Cisco NX-OS supports gRPC chunking. For streaming to occur successfully, you must enable chunking if gRPC has to send an amount of data greater than 12 MB to the receiver.

The gRPC user must do the gRPC chunking. The gRPC client side does the fragmentation, and the gRPC server side does the reassembly. Telemetry is still bound to memory and data can be dropped if the memory size is more than the allowed limit of 12 MB for telemetry. In order to support chunking, use the telemetry .proto file that is available at Cisco's GibLab, which has been updated for gRPC chunking, as described in

The chunking size is from 64 through 4096 bytes.

Following shows a configuration example through the NX-API CLI:

```

feature telemetry !
telemetry
destination-group 1
ip address 171.68.197.40 port 50051 protocol gRPC encoding GPB use-chunking size 4096
destination-group 2
ip address 10.155.0.15 port 50001 protocol gRPC encoding GPB use-chunking size 64

```

```

sensor-group 1 path sys/intf depth unbounded
sensor-group 2 path sys/intf depth unbounded
subscription 1
dst-grp 1
snsr-grp 1 sample-interval 10000
subscription 2
dst-grp 2 snsr-grp 2 sample-interval 15000

```

Following shows a configuration example through the NX-API REST:

```

{
  "telemetryDestGrpOptChunking": {
    "attributes": {
      "chunkSize": "2048",
      "dn": "sys/tm/dest-1/chunking"
    }
  }
}

```

The following error message appears on systems that do not support gRPC chunking, such as the Cisco MDS series switches:

```

switch# use-chunking size 200
ERROR: Operation failed: [chunking support not available]

```

NX-API Sensor Path Limitations

NX-API can collect and stream switch information not yet in the DME using **show** commands. However, using the NX-API instead of streaming data from the DME has inherent scale limitations as outlined:

- The switch backend dynamically processes NX-API calls such as **show** commands,
- NX-API spawns several processes that can consume up to a maximum of 20% of the CPU.
- NX-API data translates from the CLI to XML to JSON.

The following is a suggested user flow to help limit excessive NX-API sensor path bandwidth consumption:

1. Check whether the **show** command has NX-API support. You can confirm whether NX-API supports the command from the VSH with the pipe option:

```

show <command> | json

or show <command> | json pretty.

```



Note Avoid commands that take the switch more than 30 seconds to return JSON output.

2. Refine the show command to include any filters or options.

Avoid enumerating the same command for individual outputs; for example, show vlan id 100 , show vlan id 101 , and so on. Instead, use the CLI range options; for example, show vlan id 100-110,204 , whenever possible to improve performance.

If only the summary or counter is needed, then avoid dumping a whole show command output to limit the bandwidth and data storage that is required for data collection.

3. Configure telemetry with sensor groups that use NX-API as their data sources. Add the show commands as sensor paths

4. Configure telemetry with a cadence of five times the processing time of the respective show command to limit CPI usage.
5. Receive and process the streamed NX-API output as part of the existing DME collection.

Telemetry VRF Support

Telemetry VRF support allows you to specify a transport VRF, which means that the telemetry data stream can egress through front-panel ports and avoid possible competition between SSH or NGINX control sessions.

You can use the **use-vrf** *vrf-name* command to specify the transport VRF.

The following example specifies the transport VRF:

The following is an example of use-vrf as a POST payload:

```
{
  "telemetryDestProfile": {
    "attributes": {
      "adminSt": "enabled"
    },
    "children": [
      {
        "telemetryDestOptVrf": { "attributes": {
          "name": "default"
        }
      }
    ]
  }
}
```

In a given configuration, only one VRF can be used to route data for a specific IP address and port combination. The VRF that will be used is determined by the vrf configuration done at the end.

For example:

```
telemetry
  destination-group 1
    ip address 91.1.1.1 port 50007
    use-vrf default
  destination-group 2
    ip address 91.1.1.1 port 50007
    use-vrf test
  sensor-group 1
    data-source DME
    path sys/fm
  subscription 1
    dst-grp 1
    dst-grp 2
    snsr-grp 1 sample-interval 1000
```

In the above configuration, both destination-group 1 and destination-group 2 have the same destination IP address and port (91.1.1.1 50007), but different VRFs are configured (default and test). In this scenario, if destination-group 2 is configured last, the data sent to 91.1.1.1 50007 will be routed using the **test** VRF. If there are no VRF configured for a destination-group, it will implicitly use **management** VRF.

Certificate Trustpoint Support

Telemetry supports the trustpoint keyword.

The following is the command syntax:

```
switch(config-telemetry)# certificate ?
trustpoint      specify trustpoint label
WORD           .pem certificate filename (Max Size 256)
switch(config-telemetry)# certificate trustpoint
WORD           trustpoint label name (Max Size 256)
switch(config-telemetry)# certificate trustpoint trustpoint1 ?
WORD          Hostname associated with certificate (Max Size 256)
switch(config-telemetry)#certificate trustpoint trustpoint1 foo.test.google.fr
```

Destination Hostname Support

Telemetry supports the **host** keyword in destination-group command.

The following is the example for the destination hostname support:

```
switch(config-telemetry)# destination-group 1
switch(conf-tm-dest)# ?
certificate Specify certificate
host Specify destination host
ip Set destination IPv4 address
ipv6 Set destination IPv6 address
...
switch(conf-tm-dest)# host ?
A.B.C.D|A::B::C:D|WORD IPv4 or IPv6 address or DNS name of destination
switch(conf-tm-dest)#
switch(conf-tm-dest)# host abc port 11111 ?
protocol Set transport protocol
switch(conf-tm-dest)# host abc port 11111 protocol ?
HTTP
UDP
gRPC
switch(conf-tm-dest)# host abc port 11111 protocol gRPC ?
encoding Set encoding format
switch(conf-tm-dest)# host abc port 11111 protocol gRPC encoding ?
Form-data Set encoding to Form-data only
GPB Set encoding to GPB only
GPB-compact Set encoding to Compact-GPB only
JSON Set encoding to JSON
XML Set encoding to XML
switch(conf-tm-dest)# host ip address 1.1.1.1 port 2222 protocol HTTP encoding JSON
<CR>
```

Support for Node ID

Telemetry supports a custom Node ID string for a telemetry receiver through the **use-nodeid** command. By default, the host name is used, but support for a node ID enables you to set or change the identifier for the `node_id_str` of the telemetry receiver data.

You can assign the node ID through the telemetry destination profile, by using the **usenode-id** command. This command is optional.

The following example shows configuring the node ID.

```
switch(config)# telemetry
switch(config-telemetry)# destination-profile
switch(conf-tm-dest-profile)# use-nodeid test-srvr-10
switch(conf-tm-dest-profile)#
```

The following example shows a telemetry notification on the receiver after the node ID is configured.

```

Telemetry receiver:
=====
node_id_str: "test-srvr-10"
subscription_id_str: "1" encoding_path:
"sys/ch/psuslot-1/psu" collection_id:
3896 msg_timestamp: 1559669946501

```

Use the **use-nodeid** sub-command under the **host** command. The destination level **use-nodeid** configuration precedes the global level configuration. The following example shows the command syntax:

```

switch(config-telemetry)# destination-group 1
switch(conf-tm-dest)# host 172.19.216.78 port 18112 protocol http enc json
switch(conf-tm-dest-host)# use-nodeid ?
WORD Node ID (Max Size 128)
switch(conf-tm-dest-host)# use-nodeid session_1:18112

```

The following example shows the output from the Telemetry receiver:

```

>> Message size 923
Telemetry msg received @ 23:41:38 UTC Msg Size: 11
node_id_str : session_1:18112 collection_id : 3118
data_source : DME
encoding_path : sys/ch/psuslot-1/psu collection_
start_time : 1598485314721
collection_end_time : 1598485314721
data :

```

Support for Streaming of YANG Models

Telemetry supports the YANG ("Yet Another Next Generation") data modeling language. Telemetry supports data streaming for both device YANG and OpenConfig YANG.

Support for Proxy

Telemetry supports the **proxy** keyword in the host command. The following is the command syntax:

```

switch(config-telemetry)# destination-group 1
switch(conf-tm-dest)# host 172.19.216.78 port 18112 protocol http enc json
switch(conf-tm-dest-host)# proxy ?
A.B.C.D|A:B::C:D|WORD IPv4 or IPv6 address or DNS name of proxy server
<1-65535> Proxy port number, Default value is 8080 username Set proxy authentication username
password Set proxy authentication password

```

gRPC Asynchronous Mode

The gRPC asynchronous mode is available only under the **host** command. In normal stream condition, this mode allows the receivers to stream data in **mdtDialout** call without exiting or receiving **WriteDone()** call.

The following is the command syntax:

```

nxosv-1(config-telemetry)# destination-group 1
nxosv-1(conf-tm-dest)# host 172.22.244.130 port 50007 ?
nxosv-1(conf-tm-dest-host)# grpc-async ?

```


Configuring Telemetry Using the CLI

Configuring with CLI

The following steps enable streaming telemetry and configuring the source and destination of the data stream.

SUMMARY STEPS

1. **configure terminal**
2. **feature telemetry**
3. **feature nxapi**
4. **nxapi use-vrf management**
5. **telemetry**
6. (Optional) **certificate** *certificate_path* *host_URL*
7. **sensor-group** *sgrp_id*
8. **path** *sensor_path* **depth unbounded** [**filter-condition** *filter*] [**alias** *path_alias*]
9. **destination-group** *dgrp_id*
10. (Optional) **ip address** *ip_address* **port** *port* **protocol** *procedural-protocol* **encoding** *encoding-protocol*
11. (Optional) **ipv6 address** *ipv6_address* **port** *port* **protocol** *procedural-protocol* **encoding** *encoding-protocol*
12. (Optional) **source-interface** *interface*
13. **ip_version address ip_address port portnum**
14. (Optional) **use-chunking size** *chunking_size*
15. **subscription sub_id**
16. **snsr-grp** *sgrp_id* **sample-interval** *interval*
17. **dst-grp** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|--------|---|--|
| Step 1 | configure terminal Example: <code>switch# configure terminal</code> | Enter the global configuration mode. |
| Step 2 | feature telemetry | Enable the streaming telemetry feature. |
| Step 3 | feature nxapi | Enable NX-API. |
| Step 4 | nxapi use-vrf management Example: | Enable the VRF management to be used for NX-API communication. |

| | Command or Action | Purpose |
|---------------|--|---|
| | <pre>switch(config)# switch(config)# nxapi use-vrf management switch(config)</pre> | <p>Note The following warnings are seen previous to 10.2(3)F release as ACLs are able to filter only netstack packets:</p> <p>Warning Warning: Management ACLs configured will not be effective for HTTP services. Please use iptables to restrict access."</p> <p>Note Beginning with 10.2(3)F, ACLs are able to filter both netstack and kstack packets which are coming to the management vrf. The following warnings are displayed:</p> <p>Warning Warning: ACLs configured on non-management VRF will not be effective for HTTP services on that VRF."</p> |
| Step 5 | <p>telemetry</p> <p>Example:</p> <pre>switch# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for streaming telemetry. |
| Step 6 | <p>(Optional) certificate <i>certificate_path</i> <i>host_URL</i></p> <p>Example:</p> <pre>switch# certificate /bootflash/server.key localhost</pre> | <p>Use an existing SSL/TLS certificate.</p> <p>For EOR devices, the certificate also has to be copied to the standby SUP.</p> |
| Step 7 | <p>sensor-group <i>sgrp_id</i></p> <p>Example:</p> <pre>switch# sensor-group 100 switch(config-telemetry)#</pre> | <p>Create a sensor group with ID <i>srgp_id</i> and enter sensor group configuration mode.</p> <p>Currently only numeric ID values are supported. The sensor group defines nodes that will be monitored for telemetry reporting.</p> |
| Step 8 | <p>path <i>sensor_path</i> depth unbounded [filter-condition <i>filter</i>] [alias <i>path_alias</i>]</p> <p>Example:</p> <ul style="list-style-type: none"> The following command is applicable for DME, not for NX-API or YANG: <pre>switch(conf-tm-sensor)# path sys/bd/bd-[vlan-100] depth 0 filter-condition eq(l2BD.operSt, "down")</pre> <p>Use the following syntax for state-based filtering to trigger only when operSt changes from up to down, with no notifications of when the MO changes.</p> <pre>switch(conf-tm-sensor)# path sys/bd/bd-[vlan-100] depth 0 filter-condition and(updated(l2BD.operSt),eq(l2BD.operSt,"down"))</pre> | <p>Here unbounded means include child Managed Objects (MO) in the output. So, for POLL telemetry streams, all child MO for that path and EVENT retrieves the changes made in child MO.</p> <p>Note This is applicable for data source DME paths only.</p> <p>Add a sensor path to the sensor group.</p> <ul style="list-style-type: none"> Beginning with the Cisco NX-OS 9.3(5) release, the alias keyword is introduced. The depth setting specifies the retrieval level for the sensor path. Depth settings of 0 - 32 , unbounded are supported. |

| | Command or Action | Purpose |
|----------------|--|---|
| | <p>Use the following syntax to distinguish the path on the UTR side.</p> <pre>switch(conf-tm-sensor)# path sys/ch/fts1ot-1/ft alias ft_1</pre> <ul style="list-style-type: none"> The following command is applicable for NX-API, not for DME or YANG: <pre>switch(conf-tm-sensor)# path "show interface" depth 0</pre> The following command is applicable for device YANG: <pre>switch(conf-tm-sensor)# path Cisco-NX-OS-device:System/bgp-items/inst-items</pre> The following commands are applicable for OpenConfig YANG: <pre>switch(conf-tm-sensor)# path openconfig-bgp:bgp switch(conf-tm-sensor)# path Cisco-NX-OS-device:System/bgp-items/inst-items alias bgp_alias</pre> The following command is applicable for NX-API: <pre>switch(conf-tm-sensor)# path "show interface" depth 0 alias sh_int_alias</pre> The following command is applicable for OpenConfig: <pre>switch(conf-tm-sensor)# path openconfig-bgp:bgp alias oc_bgp_alias</pre> | <p>Note depth 0 is the default depth. NX-API-based sensor paths can only use depth 0 .</p> <p>If a path is subscribed for the event collection, the depth only supports 0 and unbounded. Other values would be treated as 0.</p> <ul style="list-style-type: none"> The optional filter-condition parameter can be specified to create a specific filter for event-based subscriptions. <p>For state-based filtering, the filter returns both when a state has changed and when an event has occurred during the specified state. That is, a filter condition for the DN sys/bd/bd-[vlan] of eq(l2Bd.operSt, "down") triggers when the operSt changes, and when the DN's property changes while the operSt remains down , such as a no shutdown command is issued while the VLAN is operationally down .</p> <ul style="list-style-type: none"> For the YANG model, the sensor path format is as follows: module_name: YANG_path, where module_name is the name of the YANG model file. For example: <ul style="list-style-type: none"> For device YANG: <pre>Cisco-NX-OS-device:System/bgp-items/inst-items</pre> For OpenConfig YANG: openconfig-bgp:bgp <p>Note The depth , filter-condition , and query-condition parameters are not supported for YANG currently.</p> <p>For the openconfig YANG models, go to https://github.com/YangModels/yang/tree/master/vendor/cisco/nx and navigate to the appropriate folder for the latest release.</p> <p>Instead of installing a specific model, you can install the openconfig-all RPM which has all the OpenConfig models.</p> <p>For example:</p> <pre>install add mtx-openconfig-bgp-1.0.0.0-7.0.3.IHD8.1.lib32_n9000.rpm activate</pre> |
| Step 9 | <p>destination-group <i>dgrp_id</i></p> <p>Example:</p> <pre>switch(conf-tm-sensor)# destination-group 100</pre> | <p>Create a destination group and enter destination group configuration mode.</p> <p>Currently dgrp_id only supports numeric ID values.</p> |
| Step 10 | <p>(Optional) ip address <i>ip_address</i> port <i>port</i> protocol <i>procedural-protocol</i> encoding <i>encoding-protocol</i></p> | <p>Specify an IPv4 IP address and port to receive encoded telemetry data.</p> |

| | Command or Action | Purpose |
|----------------|---|---|
| | <p>Example:</p> <pre>switch(conf-tm-sensor)# ip address 171.70.55.69 port 50001 protocol gRPC encoding GPB switch(conf-tm-sensor)# ip address 171.70.55.69 port 50007 protocol HTTP encoding JSON</pre> | <p>Note gRPC is the default transport protocol. GPB is the default encoding.</p> |
| Step 11 | <p>(Optional) ipv6 address <i>ipv6_address</i> port port protocol <i>procedural-protocol</i> encoding <i>encoding-protocol</i></p> <p>Example:</p> <pre>switch(conf-tm-sensor)# ipv6 address 10:10::1 port 8000 protocol gRPC encoding GPB switch(conf-tm-sensor)# ipv6 address 10:10::1 port 8001 protocol HTTP encoding JSON switch(conf-tm-sensor)# ipv6 address 10:10::1 port 8002 protocol UDP encoding JSON</pre> | <p>Specify an IPv6 IP address and port to receive encoded telemetry data.</p> <p>Note gRPC is the default transport protocol. GPB is the default encoding.</p> |
| Step 12 | <p>(Optional) source-interface <i>interface</i></p> <p>Example:</p> <pre>switch(conf-tm-sensor)# source-interface vlan1500</pre> | <p>Specify the telemetry source interface for a destination IP address.</p> <ul style="list-style-type: none"> • Only one source interface will be supported per destination ip:port. For a destination ip:port, last configured source-interface will be used. • The source-interface configuration must be accompanied with respective vrf of that interface using the use-vrf CLI under the destination group. If use-vrf is not specified, management vrf will be used. |
| Step 13 | <p>ip_version address ip_address port <i>portnum</i></p> <p>Example:</p> <p>For IPv4:</p> <pre>switch(conf-tm-dest)# ip address 1.2.3.4 port 50003</pre> <p>For IPv6:</p> <pre>switch(conf-tm-dest)# ipv6 address 10:10::1 port 8000</pre> | <p>Create a destination profile for the outgoing data, where ip_version is either ip (for IPv4) or ipv6 (for IPv6).</p> <p>When the destination group is linked to a subscription, telemetry data is sent to the IP address and port that is specified by this profile.</p> |
| Step 14 | <p>(Optional) use-chunking size <i>chunking_size</i></p> <p>Example:</p> <pre>switch(conf-tm-dest)# use-chunking size 64</pre> | <p>Enable gRPC chunking and set the chunking size, between 64-4096 bytes. See the section "Support for gRPC Chunking" for more information.</p> |
| Step 15 | <p>subscription sub_id</p> <p>Example:</p> <pre>switch(conf-tm-dest)# subscription 100 switch(conf-tm-sub)#</pre> | <p>Create a subscription node with ID and enter the subscription configuration mode.</p> <p>Currently sub_id only supports numeric ID values.</p> <p>Note When subscribing to a DN, check whether the DN is supported by DME using REST to ensure that events will stream.</p> |

| | Command or Action | Purpose |
|----------------|--|--|
| Step 16 | snsr-grp <i>sgrp_id</i> sample-interval <i>interval</i> Example: <pre>switch(conf-tm-sub)# snsr-grp 100 sample-interval 15000</pre> | Link the sensor group with ID <i>sgrp_id</i> to this subscription and set the data sampling interval in milliseconds. An interval value of 0 creates an event-based subscription, in which telemetry data is sent only upon changes under the specified MO. An interval value greater than 0 creates a frequency-based subscription, in which telemetry data is sent periodically at the specified interval. For example, an interval value of 15000 results in the sending of telemetry data every 15 seconds. |
| Step 17 | dst-grp <i>dgrp_id</i> Example: <pre>switch(conf-tm-sub)# dst-grp 100</pre> | Link the destination group with ID <i>dgrp_id</i> to this subscription. |

Configuring Cadence for YANG Paths

The cadence for YANG paths must be greater than the total streaming time. If the total streaming time and cadence are incorrectly configured, gathering telemetry data can take longer than the streaming interval. In this situation, you can see:

- Queues that incrementally fill because telemetry data is accumulating faster than it is streaming to the receiver.
- Stale telemetry data which is not from the current interval.

Configure the cadence to a value greater than the total streaming time.

SUMMARY STEPS

1. **show telemetry control database sensor-groups**
2. **sensor group** *number*
3. **subscription** *number*
4. **snsr-grp** *number* **sample-interval** *milliseconds*
5. **show system resources**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | show telemetry control database sensor-groups Example: <pre>switch# show telemetry control database sensor-groups Sensor Group Database size = 2</pre> <hr/> <pre>Row ID Sensor Group ID Sensor Group type</pre> | Calculate the total streaming time. The total streaming time is the sum of the individual current streaming times of each sensor group. Individual streaming times are displayed in Streaming time in ms (Cur) . In this example, total streaming time is 2.664 seconds (2515 milliseconds plus 149 milliseconds). |

| | Command or Action | Purpose |
|---------------|---|--|
| | <pre> Sampling interval(ms) Linked subscriptions SubID ----- 1 2 Timer /YANG 5000 /Running 1 Collection Time in ms (Cur/Min/Max): 2444/2294/2460 Encoding Time in ms (Cur/Min/Max): 56/55/57 Transport Time in ms (Cur/Min/Max): 0/0/1 Streaming Time in ms (Cur/Min/Max): 2515/2356/28403 Collection Statistics: collection_id_dropped = 0 last_collection_id_dropped = 0 drop_count = 0 2 1 Timer /YANG 5000 /Running 1 Collection Time in ms (Cur/Min/Max): 144/142/1471 Encoding Time in ms (Cur/Min/Max): 0/0/1 Transport Time in ms (Cur/Min/Max): 0/0/0 Streaming Time in ms (Cur/Min/Max): 149/147/23548 Collection Statistics: collection_id_dropped = 0 last_collection_id_dropped = 0 drop_count = 0 switch# telemetry destination-group 1 ip address 192.0.2.1 port 9000 protocol HTTP encoding JSON sensor-group 1 data-source YANG path /Cisco-NX-OS-device:System/procsys-items depth unbounded sensor-group 2 data-source YANG path /Cisco-NX-OS-device:System/intf-items/phys-items depth unbounded subscription 1 dst-grp 1 snsr-grp 1 sample-interval 5000 snsr-grp 2 sample-interval 5000 </pre> | <p>Compare the configured cadence to the total streaming time for the sensor group.</p> <p>The cadence is displayed in sample-interval . In this example, the cadence is correctly configured because the total streaming time (2.664 seconds) is less than the cadence (5.000 seconds, which is the default).</p> |
| Step 2 | <p>sensor group number</p> <p>Example:</p> <pre>switch(config-telemetry)# sensor group1</pre> | <p>If the total streaming time is not less than the cadence, enter the sensor group for which you want to set the interval.</p> |
| Step 3 | <p>subscription number</p> <p>Example:</p> <pre>switch(conf-tm-sensor)# subscription 100</pre> | <p>Edit the subscription for the sensor group.</p> |
| Step 4 | <p>snsr-grp number sample-interval milliseconds</p> <p>Example:</p> | <p>For the appropriate sensor group, set the sample interval to a value greater than the total streaming time.</p> |

| | Command or Action | Purpose |
|---------------|--|--|
| | <pre>switch(conf-tm-sub)# snsr-grp number sample-interval 5000</pre> | In this example, the sample interval is set to 5.000 seconds, which is valid because it is larger than the total streaming time of 2.664 seconds. |
| Step 5 | <p>show system resources</p> <p>Example:</p> <pre>switch# show system resources Load average: 1 minute: 0.38 5 minutes: 0.43 15 minutes: 0.43 Processes: 555 total, 3 running CPU states : 24.17% user, 4.32% kernel, 71.50% idle CPU0 states: 0.00% user, 2.12% kernel, 97.87% idle CPU1 states: 86.00% user, 11.00% kernel, 3.00% idle CPU2 states: 8.08% user, 3.03% kernel, 88.88% idle CPU3 states: 0.00% user, 1.02% kernel, 98.97% idle Memory usage: 16400084K total, 5861652K used, 10538432K free Current memory status: OK</pre> | Check the CPU usage. If the CPU user state shows high usage, as shown in this example, your cadence and streaming value are not configured correctly. Repeat this procedure to properly configure the cadence. |

Configuration Examples for Telemetry Using the CLI

Configure a single telemetry DME stream

with a ten second cadence with GPB encoding.

```
switch# configure terminal
switch(config)# feature telemetry
switch(config)# telemetry
switch(config-telemetry)# destination-group 1
switch(config-tm-dest)# ip address 171.70.59.62 port 50051 protocol gRPC encoding GPB
switch(config-tm-dest)# source-interface vlan1500
switch(config-tm-dest)# exit
switch(config-telemetry)# sensor group sg1
switch(config-tm-sensor)# data-source DME
switch(config-tm-dest)# path interface depth unbounded query-condition keep-data-type
switch(config-tm-dest)# subscription 1
switch(config-tm-dest)# dst-grp 1
switch(config-tm-dest)# snsr grp 1 sample interval 10000
```

Subscribe sys/bgp root MO

every 5 seconds to the destination IP 1.2.3.4 port 50003.

```
switch(config)# telemetry
switch(config-telemetry)# sensor-group 100
switch(conf-tm-sensor)# path sys/bgp depth 0
switch(conf-tm-sensor)# destination-group 100
switch(conf-tm-dest)# ip address 1.2.3.4 port 50003
switch(conf-tm-dest)# subscription 100
switch(conf-tm-sub)# snsr-grp 100 sample-interval 5000
switch(conf-tm-sub)# dst-grp 100
```

Subscribe sys/intf MO

- every 5 seconds to destination IP 1.2.3.4 port 50003
- encrypts the stream using GPB encoding verified using the test.pem.

```
switch(config)# telemetry
switch(config-telemetry)# certificate /bootflash/test.pem foo.test.google.fr
switch(conf-tm-telemetry)# destination-group 100
switch(conf-tm-dest)# ip address 1.2.3.4 port 50003 protocol gRPC encoding GPB
switch(config-dest)# sensor-group 100
switch(conf-tm-sensor)# path sys/bgp depth 0
switch(conf-tm-sensor)# subscription 100
switch(conf-tm-sub)# snsr-grp 100 sample-interval 5000
switch(conf-tm-sub)# dst-grp 100
```

Subscribe sys/cdp MO

every 15 seconds to destination IP 1.2.3.4 port 50004

```
switch(config)# telemetry
switch(config-telemetry)# sensor-group 100
switch(conf-tm-sensor)# path sys/cdp depth 0
switch(conf-tm-sensor)# destination-group 100
switch(conf-tm-dest)# ip address 1.2.3.4 port 50004
switch(conf-tm-dest)# subscription 100
switch(conf-tm-sub)# snsr-grp 100 sample-interval 15000
switch(conf-tm-sub)# dst-grp 100
```

Subscribe cadence-based collection of show command

every 750 seconds.

```
switch(config)# telemetry
switch(config-telemetry)# destination-group 1
switch(conf-tm-dest)# ip address 172.27.247.72 port 60001 protocol gRPC encoding GPB
switch(conf-tm-dest)# sensor-group 1
switch(conf-tm-sensor)# data-source NX-API
switch(conf-tm-sensor)# path "show system resources" depth 0
switch(conf-tm-sensor)# path "show version" depth 0
switch(conf-tm-sensor)# path "show environment power" depth 0
switch(conf-tm-sensor)# path "show environment fan" depth 0
switch(conf-tm-sensor)# path "show environment temperature" depth 0
switch(conf-tm-sensor)# path "show process cpu" depth 0
switch(conf-tm-sensor)# path "show nve peers" depth 0
switch(conf-tm-sensor)# path "show nve vni" depth 0
switch(conf-tm-sensor)# path "show nve vni 4002 counters" depth 0
switch(conf-tm-sensor)# path "show int nve 1 counters" depth 0
switch(conf-tm-sensor)# path "show policy-map vlan" depth 0
switch(conf-tm-sensor)# path "show ip access-list test" depth 0
switch(conf-tm-sensor)# path "show system internal access-list resource utilization" depth
0
switch(conf-tm-sensor)# subscription 1
switch(conf-tm-sub)# dst-grp 1
switch(conf-tm-dest)# snsr-grp 1 sample-interval 750000
```

Subscribe event-based subscription for sys/fm

streamed only if there is a change under the sys/fm MO

```
switch(config)# telemetry
switch(config-telemetry)# sensor-group 100
switch(conf-tm-sensor)# path sys/fm depth 0
```



```
switch(conf-tm-sensor)# destination-group 100
switch(conf-tm-dest)# ip address 1.2.3.4 port 50005
switch(conf-tm-dest)# subscription 100
switch(conf-tm-sub)# snsr-grp 100 sample-interval 0
switch(conf-tm-sub)# dst-grp 100
```

During operation, you can change a sensor group from frequency-based to event-based, and change event-based to frequency-based by changing the sample-interval. This example changes the sensor-group from the previous example to frequency-based. After the following commands, the telemetry application will begin streaming the sys/fm data to the destination every 7 seconds.

```
switch(config)# telemetry
switch(config-telemetry)# subscription 100
switch(conf-tm-sub)# snsr-grp 100 sample-interval 7000
```

Subscribe to multiple sensor and destination groups

Multiple sensor groups and destinations can be linked to a single subscription. The subscription in this example streams the data for Ethernet port 1/1 to four different destinations every 10 seconds.

```
switch(config)# telemetry
switch(config-telemetry)# sensor-group 100
switch(conf-tm-sensor)# path sys/intf/phys-[eth1/1] depth 0
switch(conf-tm-sensor)# destination-group 100
switch(conf-tm-dest)# ip address 1.2.3.4 port 50004
switch(conf-tm-dest)# ip address 1.2.3.4 port 50005
switch(conf-tm-sensor)# destination-group 200
switch(conf-tm-dest)# ip address 5.6.7.8 port 50001 protocol HTTP encoding JSON
switch(conf-tm-dest)# ip address 1.4.8.2 port 60003
switch(conf-tm-dest)# subscription 100
switch(conf-tm-sub)# snsr-grp 100 sample-interval 10000
switch(conf-tm-sub)# dst-grp 100
switch(conf-tm-sub)# dst-grp 200
```

A sensor group can contain multiple paths, a destination group can contain multiple destination profiles, and a subscription can be linked to multiple sensor groups and destination groups, as shown in this example.

```
switch(config)# telemetry
switch(config-telemetry)# sensor-group 100
switch(conf-tm-sensor)# path sys/intf/phys-[eth1/1] depth 0
switch(conf-tm-sensor)# path sys/epId-1 depth 0
switch(conf-tm-sensor)# path sys/bgp/inst/dom-default depth 0

switch(config-telemetry)# sensor-group 200
switch(conf-tm-sensor)# path sys/cdp depth 0
switch(conf-tm-sensor)# path sys/ipv4 depth 0

switch(config-telemetry)# sensor-group 300
switch(conf-tm-sensor)# path sys/fm depth 0
switch(conf-tm-sensor)# path sys/bgp depth 0

switch(conf-tm-sensor)# destination-group 100
switch(conf-tm-dest)# ip address 1.2.3.4 port 50004
switch(conf-tm-dest)# ip address 4.3.2.5 port 50005

switch(conf-tm-dest)# destination-group 200
switch(conf-tm-dest)# ip address 5.6.7.8 port 50001

switch(conf-tm-dest)# destination-group 300
switch(conf-tm-dest)# ip address 1.2.3.4 port 60003

switch(conf-tm-dest)# subscription 600
switch(conf-tm-sub)# snsr-grp 100 sample-interval 7000
switch(conf-tm-sub)# snsr-grp 200 sample-interval 20000
```

```

switch(conf-tm-sub)# dst-grp 100
switch(conf-tm-sub)# dst-grp 200

switch(conf-tm-dest)# subscription 900
switch(conf-tm-sub)# snsr-grp 200 sample-interval 7000
switch(conf-tm-sub)# snsr-grp 300 sample-interval 0
switch(conf-tm-sub)# dst-grp 100
switch(conf-tm-sub)# dst-grp 300

```

Subscribe to UDP transport

Use the following command to configure the UDP transport to stream data using a datagram socket either in JSON or GPB:

```

destination-group num
    ip address xxx.xxx.xxx.xxx port xxxx protocol UDP encoding {JSON | GPB }
Example for an IPv6 destination:
destination-group 100 ipv6 address 10:10::1 port 8000 protocol gRPC encoding GPB

```

The UDP telemetry is with the following header:

```

typedef enum tm_encode_ {
    TM_ENCODE_DUMMY,
    TM_ENCODE_GPB,
    TM_ENCODE_JSON,
    TM_ENCODE_XML,
    TM_ENCODE_MAX, } tm_encode_type_t;
typedef struct tm_pak_hdr_ {
    uint8_t version; /* 1 */ uint8_t encoding; uint16_t msg_size; uint8_t secure; uint8_t
padding;
}__attribute__((packed, aligned (1))) tm_pak_hdr_t;

```

Use the first 6 bytes in the payload to process telemetry data using UDP, using one of the following methods:

- Read the information in the header to determine which decoder to use to decode the data, JSON or GPB, if the receiver is meant to receive different types of data from multiple endpoints.
- Remove the header if you are expecting one decoder (JSON or GPB) but not the other.

Verify telemetry configuration

You can verify the telemetry configuration using the show running-config telemetry command, as shown in this example.

```

switch(config)# telemetry
switch(config-telemetry)# destination-group 100
switch(conf-tm-dest)# ip address 1.2.3.4 port 50003
switch(conf-tm-dest)# ip address 1.2.3.4 port 50004
switch(conf-tm-dest)# end
switch# show run telemetry

!Command: show running-config telemetry
!Time: Thu Oct 13 21:10:12 2016

version 7.0(3)I5(1)
feature telemetry

telemetry
destination-group 100
ip address 1.2.3.4 port 50003 protocol gRPC encoding GPB
ip address 1.2.3.4 port 50004 protocol gRPC encoding GPB

```

Displaying Telemetry Configuration and Statistics

Use the following NX-OS CLI **show** commands to display telemetry configuration, statistics, errors, and session information.

show telemetry yang direct-path cisco-nxos-device

This command displays YANG paths that are directly encoded to perform better than other paths.

```
switch# show telemetry yang direct-path cisco-nxos-device
) Cisco-NX-OS-device:System/lldp-items
2) Cisco-NX-OS-device:System/acl-items
3) Cisco-NX-OS-device:System/mac-items
4) Cisco-NX-OS-device:System/intf-items
5) Cisco-NX-OS-device:System/procsys-items/sysload-items
6) Cisco-NX-OS-device:System/ospf-items
7) Cisco-NX-OS-device:System/procsys-items
8) Cisco-NX-OS-device:System/ipqos-items/queuing-items/policy-items/out-items
9) Cisco-NX-OS-device:System/mac-items/static-items
10) Cisco-NX-OS-device:System/ch-items
11) Cisco-NX-OS-device:System/cdp-items
12) Cisco-NX-OS-device:System/bd-items
13) Cisco-NX-OS-device:System/eps-items
14) Cisco-NX-OS-device:System/ipv6-items
```

show telemetry control database

This command displays YANG paths that are directly encoded to perform better than other paths.

```
switch# show telemetry control database ?
<CR>
>                                Redirect it to a file
>>                               Redirect it to a file in append mode
destination-groups              Show destination-groups
destinations                     Show destinations
sensor-groups                   Show sensor-groups
sensor-paths                    Show sensor-paths
subscriptions                   Show subscriptions
|                                Pipe command output to filter

switch# show telemetry control database

Subscription Database size = 1

-----
Subscription ID      Data Collector Type
-----
100                  DME NX-API

Sensor Group Database size = 1

-----
Sensor Group ID  Sensor Group type  Sampling interval(ms)  Linked subscriptions
-----
100              Timer                10000 (Running)        1

Sensor Path Database size = 1

-----
Subscribed Query Filter  Linked Groups  Sec Groups  Retrieve level  Sensor Path
-----
```

```
No                1                0                Full                sys/fm
```

```
Destination group Database size = 2
```

```
-----
Destination Group ID  Refcount
-----
```

```
100                1
```

```
Destination Database size = 2
```

```
-----
Dst IP Addr        Dst Port   Encoding   Transport  Count
-----
```

```
192.168.20.111    12345     JSON      HTTP      1
```

```
192.168.20.123  50001     GPB       gRPC      1
```

show telemetry control database sensor-paths

This command displays sensor path details for telemetry configuration, including counters for encoding, collection, transport, and streaming.

```
switch(conf-tm-sub)# show telemetry control database sensor-paths
Sensor Path Database size = 4
```

```
-----
Row ID Subscribed Linked Groups Sec Groups Retrieve level Path(GroupId) : Query :
Filter
-----
```

```
1 No 1 0 Full sys/cdp(1) : NA : NA
```

```
GPB Encoded Data size in bytes (Cur/Min/Max): 0/0/0
```

```
JSON Encoded Data size in bytes (Cur/Min/Max): 65785/65785/65785
```

```
Collection Time in ms (Cur/Min/Max): 10/10/55
```

```
Encoding Time in ms (Cur/Min/Max): 8/8/9
```

```
Transport Time in ms (Cur/Min/Max): 0/0/0
```

```
Streaming Time in ms (Cur/Min/Max): 18/18/65
```

```
2 No 1 0 Self show module(2) : NA : NA
```

```
GPB Encoded Data size in bytes (Cur/Min/Max): 0/0/0
```

```
JSON Encoded Data size in bytes (Cur/Min/Max): 1107/1106/1107
```

```
Collection Time in ms (Cur/Min/Max): 603/603/802
```

```
Encoding Time in ms (Cur/Min/Max): 0/0/0
```

```
Transport Time in ms (Cur/Min/Max): 0/0/1
```

```
Streaming Time in ms (Cur/Min/Max): 605/605/803
```

```
3 No 1 0 Full
```

```
GPB Encoded Data size in bytes (Cur/Min/Max): 0/0/0
```

```
JSON Encoded Data size in bytes (Cur/Min/Max): 0/0/0
```

```
Collection Time in ms (Cur/Min/Max): 0/0/44
```

```
Encoding Time in ms (Cur/Min/Max): 0/0/0
```

```
Transport Time in ms (Cur/Min/Max): 0/0/0
```

```
Streaming Time in ms (Cur/Min/Max): 1/1/44 sys/bgp(1) : NA : NA
```

```
4 No 1 0 Self
```

```
GPB Encoded Data size in bytes (Cur/Min/Max): 0/0/0
```

```
JSON Encoded Data size in bytes (Cur/Min/Max): 2442/2441/2442
```

```
Collection Time in ms (Cur/Min/Max): 1703/1703/1903
```

```
Encoding Time in ms (Cur/Min/Max): 0/0/0
```

```
Transport Time in ms (Cur/Min/Max): 0/0/0
```

```
Streaming Time in ms (Cur/Min/Max): 1703/1703/1904
```

```
switch(conf-tm-sub)#
```

```
show version(2) : NA : NA
```

show telemetry control stats

This command displays the statistics about the internal databases about configuration of telemetry.

```
switch# show telemetry control stats
show telemetry control stats entered
```

```
-----
Error Description                                     Error Count
-----
Chunk allocation failures                             0
Sensor path Database chunk creation failures         0
Sensor Group Database chunk creation failures        0
Destination Database chunk creation failures         0
Destination Group Database chunk creation failures   0
Subscription Database chunk creation failures        0
Sensor path Database creation failures               0
Sensor Group Database creation failures              0
Destination Database creation failures               0
Destination Group Database creation failures         0
Subscription Database creation failures              0
Sensor path Database insert failures                 0
Sensor Group Database insert failures                0
Destination Database insert failures                0
Destination Group Database insert failures           0
Subscription insert to Subscription Database failures 0
Sensor path Database delete failures                0
Sensor Group Database delete failures                0
Destination Database delete failures                0
Destination Group Database delete failures           0
Delete Subscription from Subscription Database failures 0
Sensor path delete in use                            0
Sensor Group delete in use                           0
Destination delete in use                            0
Destination Group delete in use                      0
Delete destination(in use) failure count             0
Failed to get encode callback                        0
Sensor path Sensor Group list creation failures     0
Sensor path prop list creation failures              0
Sensor path sec Sensor path list creation failures  0
Sensor path sec Sensor Group list creation failures 0
Sensor Group Sensor path list creation failures     0
Sensor Group Sensor subs list creation failures     0
Destination Group subs list creation failures       0
Destination Group Destinations list creation failures 0
Destination Destination Groups list creation failures 0
Subscription Sensor Group list creation failures   0
Subscription Destination Groups list creation failures 0
Sensor Group Sensor path list delete failures       0
Sensor Group Subscriptions list delete failures     0
Destination Group Subscriptions list delete failures 0
Destination Group Destinations list delete failures 0
Subscription Sensor Groups list delete failures     0
Subscription Destination Groups list delete failures 0
Destination Destination Groups list delete failures 0
Failed to delete Destination from Destination Group 0
Failed to delete Destination Group from Subscription 0
Failed to delete Sensor Group from Subscription     0
Failed to delete Sensor path from Sensor Group     0
Failed to get encode callback                       0
Failed to get transport callback                    0
switch# Destination Database size = 1
```

```
-----
Dst IP Addr      Dst Port  Encoding  Transport  Count
-----
192.168.20.123 50001    GPB       gRPC       1
```

show telemetry data collector brief

This command displays the brief statistics about the data collection.

```
switch# show telemetry data collector brief
-----
Collector Type Successful Collections Failed Collections
-----
DME 143 0
```

show telemetry data collector details

This command displays detailed statistics about the data collection which includes breakdown of all sensor paths.

```
switch# show telemetry data collector details
-----
Succ Collections Failed Collections Sensor Path
-----
150 0 sys/fm
```

show telemetry event collector errors

This command displays the errors statistic about the event collection.

```
switch# show telemetry event collector errors
-----
Error Description Error Count
-----
APIC-Cookie Generation Failures - 0
Authentication Failures - 0
Authentication Refresh Failures - 0
Authentication Refresh Timer Start Failures - 0
Connection Timer Start Failures - 0
Connection Attempts - 3
Dme Event Subscription Init Failures - 0
Event Data Enqueue Failures - 0
Event Subscription Failures - 0
Event Subscription Refresh Failures - 0
Pending Subscription List Create Failures - 0
Subscription Hash Table Create Failures - 0
Subscription Hash Table Destroy Failures - 0
Subscription Hash Table Insert Failures - 0
Subscription Hash Table Remove Failures - 0
Subscription Refresh Timer Start Failures - 0
Websocket Connect Failures - 0
```

show telemetry event collector stats

This command displays the statistics about the event collection which includes breakdown of all sensor paths.

```
switch# show telemetry event collector stats
-----
Collection Count Latest Collection Time Sensor Path
-----
```

show telemetry control pipeline stats

This command displays the statistics for the telemetry pipeline.

```
switch# show telemetry pipeline stats
Main Statistics:
  Timers:
```

```

Errors:
  Start Fail          =    0

Data Collector:
  Errors:
    Node Create Fail =    0

Event Collector:
  Errors:
    Node Create Fail =    0    Node Add Fail    =    0
    Invalid Data     =    0

Queue Statistics:
  Request Queue:
    High Priority Queue:
      Info:
        Actual Size   =   50    Current Size   =    0
        Max Size      =    0    Full Count    =    0

      Errors:
        Enqueue Error =    0    Dequeue Error =    0

    Low Priority Queue:
      Info:
        Actual Size   =   50    Current Size   =    0
        Max Size      =    0    Full Count    =    0

      Errors:
        Enqueue Error =    0    Dequeue Error =    0

  Data Queue:
    High Priority Queue:
      Info:
        Actual Size   =   50    Current Size   =    0
        Max Size      =    0    Full Count    =    0

      Errors:
        Enqueue Error =    0    Dequeue Error =    0

    Low Priority Queue:
      Info:
        Actual Size   =   50    Current Size   =    0
        Max Size      =    0    Full Count    =    0

      Errors:
        Enqueue Error =    0    Dequeue Error =    0

```

show telemetry transport

This command displays all configured transport sessions.

```
switch# show telemetry transport
```

| Session Id | IP Address | Port | Encoding | Transport | Status |
|------------|----------------|-------|----------|-----------|-----------|
| 0 | 192.168.20.123 | 50001 | GPB | gRPC | Connected |

Table 43: Syntax Description for show telemetry transport

| Syntax | Description |
|--------|-------------|
|--------|-------------|

| | |
|--------------------------------------|---|
| show | Shows running system information |
| telemetry | Shows telemetry information |
| transport | Shows telemetry transport information |
| <i>session_id</i> | (Optional) Session id |
| stats | (Optional) Shows all telemetry statistics information |
| errors | (Optional) Show all telemetry error information |
| readonly | (Optional) |
| TABLE_transport_info | (Optional) Transport information |
| <i>session_idx</i> | (Optional) Session Id |
| <i>ip_address</i> | (Optional) Transport IP address |
| <i>port</i> | (Optional) Transport port |
| <i>dest_info</i> | (Optional) Destination information |
| <i>encoding_type</i> | (Optional) Encoding type |
| <i>transport_type</i> | (Optional) Transport type |
| <i>transport_status</i> | (Optional) Transport status |
| <i>transport_security_cert_fname</i> | (Optional) Transport security file name |
| <i>transport_last_connected</i> | (Optional) Transport last connected |
| <i>transport_last_disconnected</i> | (Optional) Last time this destination configuration was |
| <i>transport_errors_count</i> | (Optional) Transport errors count |
| <i>transport_last_tx_error</i> | (Optional) Transport last tx error |
| transport_statistics | (Optional) Transport statistics |
| <i>t_session_id</i> | (Optional) Transport Session id |
| connect_statistics | (Optional) Connection statistics |
| <i>connect_count</i> | (Optional) Connection count |
| <i>last_connected</i> | (Optional) Last connected timestamp |
| <i>disconnect_count</i> | (Optional) Disconnect count |
| <i>last_disconnected</i> | (Optional) Last time this destination configuration was |
| trans_statistics | (Optional) Transport statistics |
| <i>compression</i> | (Optional) Compression status |
| <i>source_interface_name</i> | (Optional) Source interface name |
| <i>source_interface_ip</i> | (Optional) Source interface IP |
| <i>transmit_count</i> | (Optional) Transmission count |
| <i>last_tx_time</i> | (Optional) Last Transmission time |

| | |
|------------------------------|---|
| <i>min_tx_time</i> | (Optional) Minimum transmission time |
| <i>max_tx_time</i> | (Optional) Maximum transmission time |
| <i>avg_tx_time</i> | (Optional) Average transmission time |
| <i>cur_tx_time</i> | (Optional) Current transmission time |
| <i>transport_errors</i> | (Optional) Transport errors |
| <i>connect_errors</i> | (Optional) Connection errors |
| <i>connect_errors_count</i> | (Optional) Connection error count |
| <i>trans_errors</i> | (Optional) Transport errors |
| <i>trans_errors_count</i> | (Optional) Transport error count |
| <i>last_tx_error</i> | (Optional) Last transport error |
| <i>last_tx_return_code</i> | (Optional) Last transport return code |
| <i>transport_retry_stats</i> | (Optional) Retry Statistics |
| <i>ts_event_retry_bytes</i> | (Optional) Event Retry buffer size |
| <i>ts_timer_retry_bytes</i> | (Optional) Timer Retry buffer size |
| <i>ts_event_retry_size</i> | (Optional) Event Retry number of messages |
| <i>ts_timer_retry_size</i> | (Optional) Timer Retry number of messages |
| <i>ts_retries_sent</i> | (Optional) Number of retries sent |
| <i>ts_retries_dropped</i> | (Optional) Number of retries dropped |
| <i>event_retry_bytes</i> | (Optional) Event Retry buffer size |
| <i>timer_retry_bytes</i> | (Optional) Timer Retry buffer size |
| <i>retries_sent</i> | (Optional) Number of retries sent |
| <i>retries_dropped</i> | (Optional) Number of retries dropped |
| <i>retry_buffer_size</i> | (Optional) Retry buffer size |

show telemetry transport <session-id>

This command displays detailed session information for a specific transport session.

```
switch# show telemetry transport 0

Session Id:          0
IP Address:Port     192.168.20.123:50001
Encoding:           GPB
Transport:          gRPC
Status:             Disconnected
Last Connected:     Fri Sep 02 11:45:57.505 UTC

Tx Error Count:     224
Last Tx Error:      Fri Sep 02 12:23:49.555 UTC

switch# show telemetry transport 1

Session Id:          1
```

```

IP Address:Port      10.30.218.56:51235
Encoding:           JSON
Transport:          HTTP
Status:             Disconnected
Last Connected:     Never

Tx Error Count:     3
Last Tx Error:      Wed Apr 19 15:56:51.617 PDT
The following example shows output from an IPv6 entry.
switch# show telemetry transport 0
Session Id: 0
IP Address:Port [10:10::1]:8000
Transport: GRPC
Status: Idle
Last Connected: Never
Last Disconnected: Never
Tx Error Count: 0
Last Tx Error: None
Event Retry Queue Bytes: 0
Event Retry Queue Size: 0
Timer Retry Queue Bytes: 0
Timer Retry Queue Size: 0
Sent Retry Messages: 0
Dropped Retry Messages: 0

```

show telemetry transport <session-id> stats

This command displays details of a specific transport session.

```

switch# show telemetry transport 0 stats

Session Id:          0
IP Address:Port      192.168.20.123:50001
Encoding:            GPB
Transport:           GRPC
Status:              Connected
Last Connected:      Mon May 01 11:29:46.912 PST
Last Disconnected:   Never
Tx Error Count:      0
Last Tx Error:       None

```

show telemetry transport <session-id> errors

This command displays detailed error statistics for a specific transport session.

```

switch# show telemetry transport 0 errors

Session Id:          0
Connection Stats
  Connection Count    1
  Last Connected:     Mon May 01 11:29:46.912 PST
  Disconnect Count    0
  Last Disconnected:  Never
Transmission Stats
  Transmit Count:     1225
  Last TX time:       Tue May 02 11:40:03.531 PST
  Min Tx Time:        7 ms
  Max Tx Time:        1760 ms
  Avg Tx Time:        500 ms

```

show telemetry control databases sensor-paths

These following configuration steps result in the **show telemetry control databases sensor-paths** command output

```

below. feature telemetry
telemetry
destination-group 1
ip address 172.25.238.13 port 50600 protocol gRPC encoding GPB
sensor-group 1
path sys/cdp depth unbounded path sys/intf depth unbounded path sys/mac depth 0
subscription 1
dst-grp 1 snsr-grp 1 sample-interval 1000 Command output. switch# show telemetry control
databases sensor-paths
Sensor Path Database size = 3
-----
-----
Row ID Subscribed Linked Groups Sec Groups Retrieve level Path(GroupId) : Query : Filter
-----
-----
1 No 1 0 Full sys/cdp(1) : NA
: NA
GPB Encoded Data size in bytes (Cur/Min/Max): 30489/30489/30489
JSON Encoded Data size in bytes (Cur/Min/Max): 0/0/0
CGPB Encoded Data size in bytes (Cur/Min/Max): 0/0/0
Collection Time in ms (Cur/Min/Max): 6/5/54
Encoding Time in ms (Cur/Min/Max): 5/5/6
Transport Time in ms (Cur/Min/Max): 1027/55/1045
Streaming Time in ms (Cur/Min/Max): 48402/5/48402
2 No 1 0 Full sys/intf(1) : N
A : NA
GPB Encoded Data size in bytes (Cur/Min/Max): 539466/539466/539466
JSON Encoded Data size in bytes (Cur/Min/Max): 0/0/0
CGPB Encoded Data size in bytes (Cur/Min/Max): 0/0/0
Collection Time in ms (Cur/Min/Max): 66/64/114
Encoding Time in ms (Cur/Min/Max): 91/90/92
Transport Time in ms (Cur/Min/Max): 4065/4014/5334
Streaming Time in ms (Cur/Min/Max): 48365/64/48365
3 No 1 0 Self sys/mac(1) : NA
: NA
GPB Encoded Data size in bytes (Cur/Min/Max): 247/247/247
JSON Encoded Data size in bytes (Cur/Min/Max): 0/0/0
CGPB Encoded Data size in bytes (Cur/Min/Max): 0/0/0
Collection Time in ms (Cur/Min/Max): 1/1/47
Encoding Time in ms (Cur/Min/Max): 1/1/1
Transport Time in ms (Cur/Min/Max): 4/1/6
Streaming Time in ms (Cur/Min/Max): 47369/1/47369

```

show telemetry transport sessions

The following commands loop through all the transport sessions and prints the information in one command:

```

switch# show telemetry transport sessions
switch# show telemetry transport stats
switch# show telemetry transport errors
switch# show telemetry transport all
The following is an example for telemetry transport session:
switch# show telemetry transport sessions
Session Id:          0
IP Address:Port     172.27.254.13:50004
Transport:          GRPC
Status:             Transmit Error
SSL Certificate:    trustpoint1
Last Connected:     Never

```

```

Last Disconnected:  Never
Tx Error Count:    2
Last Tx Error:    Wed Aug 19 23:32:21.749 UTC
...
Session Id:       4
IP Address:Port   172.27.2

```

Telemetry Ephemeral Event

To support ephemeral event, a new sensor path query-condition is added. To enable accounting log ephemeral event streaming, use the following query condition:

```

sensor-group 1 path sys/accounting/log query-condition
query-target=subtree&complete-mo=yes&notify-interval=1

```

The following are the other sensor paths that support ephemeral event:
 sys/pim/inst/routedb-route, sys/pim/pimifdb-adj, sys/pim/pimifdb-prop
 sys/igmp/igmpifdb-prop, sys/igmp/inst/routedb, sys/igmpsnoop/inst/dom/db-extrack,
 sys/igmpsnoop/inst/dom/db-group, sys/igmpsnoop/inst/dom/db-mrouter
 sys/igmpsnoop/inst/dom/db-querier, sys/igmpsnoop/inst/dom/db-snoop

Configuring Telemetry Using the NX-API

Telemetry Model in the DME

The telemetry application is modeled in the DME with the following structure:

```

model
|----package [name:telemetry]
|   @name:telemetry
|----objects
|   |----mo [name:Entity]
|   |   @name:Entity
|   |   @label:Telemetry System
|   |--property
|   |   @name:adminSt
|   |   @type:AdminState
|   |
|   |----mo [name:SensorGroup]
|   |   @name:SensorGroup
|   |   @label:Sensor Group
|   |--property
|   |   @name:id [key]
|   |   @type:string:Basic
|   |
|   |----mo [name:SensorPath]
|   |   @name:SensorPath
|   |   @label:Sensor Path
|   |--property
|   |   @name:path [key]
|   |   @type:string:Basic
|   |   @name:filterCondition
|   |   @type:string:Basic
|   |   @name:excludeFilter
|   |   @type:string:Basic
|   |   @name:depth
|   |   @type:RetrieveDepth
|   |
|   |----mo [name:DestGroup]

```

```

|   |   @name:DestGroup
|   |   @label:Destination Group
|   |--property
|   |   @name:id
|   |   @type:string:Basic
|   |
|   |----mo [name:Dest]
|   |   @name:Dest
|   |   @label:Destination
|   |   |--property
|   |   |   @name:addr [key]
|   |   |   @type:address:Ip
|   |   |   @name:port [key]
|   |   |   @type:scalar:Uint16
|   |   |   @name:proto
|   |   |   @type:Protocol
|   |   |   @name:enc
|   |   |   @type:Encoding
|   |
|   |----mo [name:Subscription]
|   |   @name:Subscription
|   |   @label:Subscription
|   |   |--property
|   |   |   @name:id
|   |   |   @type:scalar:Uint64
|   |   |----reldef
|   |   |   @name:SensorGroupRel
|   |   |   @to:SensorGroup
|   |   |   @cardinality:ntom
|   |   |   @label:Link to sensorGroup entry
|   |   |----property
|   |   |   @name:sampleIntvl
|   |   |   @type:scalar:Uint64
|   |   |----reldef
|   |   |   @name:DestGroupRel
|   |   |   @to:DestGroup
|   |   |   @cardinality:ntom
|   |   |   @label:Link to destGroup entry

```

Configuring with NX-API

In the object model of the switch DME, the configuration of the telemetry feature is defined in a hierarchical structure of objects as shown in the section "Telemetry Model in the DME." Following are the main objects to be configured:

- **fmEntity** — Contains the NX-API and Telemetry feature states.
- **fmNxapi** — Contains the NX-API state.
- **fmTelemetry** — Contains the Telemetry feature state.
- **telemetryEntity** — Contains the telemetry feature configuration.
- **telemetrySensorGroup** — Contains the definitions of one or more sensor paths or nodes to be monitored for telemetry. The telemetry entity can contain one or more sensor groups.
- **telemetryRtSensorGroupRel** — Associates the sensor group with a telemetry subscription.

- **telemetrySensorPath** — A path to be monitored. The sensor group can contain multiple objects of this type.
- **telemetryDestGroup** — Contains the definitions of one or more destinations to receive telemetry data. The telemetry entity can contain one or more destination groups.
- **telemetryRtDestGroupRel** — Associates the destination group with a telemetry subscription.
- **telemetryDest** — A destination address. The destination group can contain multiple objects of this type.
- **telemetrySubscription** — Specifies how and when the telemetry data from one or more sensor groups is sent to one or more destination groups.
- **telemetryRsDestGroupRel** — Associates the telemetry subscription with a destination group.
- **telemetryRsSensorGroupRel** — Associates the telemetry subscription with a sensor group.
- **telemetryCertificate** — Associates the telemetry subscription with a certificate and hostname.

To configure the telemetry feature using the NX-API, you must construct a JSON representation of the telemetry object structure and push it to the DME with an HTTP or HTTPS POST operation.



Note For detailed instructions on using the NX-API, see the *Cisco Nexus 3000 and 9000 Series NX-API REST SDK User Guide and API Reference*

Before you begin

Your switch must be configured to run the NX-API from the CLI:

```
switch(config)# feature nxapi
nxapi use-vrf vrf_name nxapi http port port_number
```

SUMMARY STEPS

1. Enable the telemetry feature.
2. Create the root level of the JSON payload to describe the telemetry configuration.
3. Create a sensor group to contain the defined sensor paths.
4. (Optional) Add an SSL/TLS certificate and a host.
5. Define a telemetry destination group.
6. Define a telemetry destination profile.
7. Define one or more telemetry destinations, consisting of an IP address and port number to which telemetry data will be sent.
8. Enable gRPC chunking and set the chunking size, between 64 and 4096 bytes.
9. Create a telemetry subscription to configure the telemetry behavior
10. Add the sensor group object as a child object to the telemetrySubscription element under the root element (telemetryEntity).
11. Create a relation object as a child object of the subscription to associate the subscription to the telemetry sensor group and to specify the data sampling behavior.
12. Define one or more sensor paths or nodes to be monitored for telemetry.
13. Add sensor paths as child objects to the sensor group object (telemetrySensorGroup).

14. Add destinations as child objects to the destination group object (telemetryDestGroup).
15. Add the destination group object as a child object to the root element (telemetryEntity).
16. Create a relation object as a child object of the telemetry sensor group to associate the sensor group to the subscription.
17. Create a relation object as a child object of the telemetry destination group to associate the destination group to the subscription.
18. Create a relation object as a child object of the subscription to associate the subscription to the telemetry destination group.
19. Send the resulting JSON structure as an HTTP/HTTPS POST payload to the NX-API endpoint for telemetry configuration.

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | Enable the telemetry feature. Example: <pre> { "fmEntity" : { "children" : [{ "fmTelemetry" : { "attributes" : { "adminSt" : "enabled" } }] } } </pre> | The root element is fmTelemetry and the base path for this element is <code>sys/fm</code> . Configure the adminSt attribute as enabled. |
| Step 2 | Create the root level of the JSON payload to describe the telemetry configuration. Example: <pre> { "telemetryEntity": { "attributes": { "dn": "sys/tm" }, } } </pre> | The root element is telemetryEntity and the base path for this element is <code>sys/tm</code> . Configure the dn attribute as <code>sys/tm</code> . |
| Step 3 | Create a sensor group to contain the defined sensor paths. Example: <pre> "telemetrySensorGroup": { "attributes": { "id": "10", "rn": "sensor-10" }, "children": [{ }] } </pre> | A telemetry sensor group is defined in an object of class <code>telemetrySensorGroup</code> . Configure the following attributes of the object: <ul style="list-style-type: none"> • <code>id</code> — An identifier for the sensor group. Currently only numeric ID values are supported. • <code>rn</code> — The relative name of the sensor group object in the format: <code>sensor-id</code>. |

| | Command or Action | Purpose |
|---------------|--|---|
| | | <ul style="list-style-type: none"> dataSrc — Selects the data source from DEFAULT , DME , YANG , or NX-API . <p>Children of the sensor group object include sensor paths and one or more relation objects (telemetryRtSensorGroupRel) to associate the sensor group with a telemetry subscription.</p> |
| Step 4 | (Optional) Add an SSL/TLS certificate and a host. Example: <pre>{ "telemetryCertificate": { "attributes": { "filename": "root.pem" "hostname": "c.com" } } }</pre> | The telemetryCertificate defines the location of the SSL/TLS certificate with the telemetry subscription/destination. |
| Step 5 | Define a telemetry destination group. Example: <pre>{ "telemetryDestGroup": { "attributes": { "id": "20" } } }</pre> | A telemetry destination group is defined in telemetryEntity . Configure the id attribute. |
| Step 6 | Define a telemetry destination profile. Example: <pre>{ "telemetryDestProfile": { "attributes": { "adminSt": "enabled" }, "children": [{ "telemetryDestOptSourceInterface": { "attributes": { "name": "lo0" } } }] } }</pre> | A telemetry destination profile is defined in telemetryDestProfile . <ul style="list-style-type: none"> Configure the adminSt attribute as enabled. <p>Under telemetryDestOptSourceInterface , configure the name attribute with an interface name to stream data from the configured interface to a destination with the source IP address.</p> |
| Step 7 | Define one or more telemetry destinations, consisting of an IP address and port number to which telemetry data will be sent. Example: <pre>{</pre> | A telemetry destination is defined in an object of class telemetryDest . Configure the following attributes of the object: <ul style="list-style-type: none"> addr — The IP address of the destination. port — The port number of the destination. |

| | Command or Action | Purpose |
|---------------|--|--|
| | <pre> "telemetryDest": { "attributes": { "addr": "1.2.3.4", "enc": "GPB", "port": "50001", "proto": "gRPC", "rn": "addr-[1.2.3.4]-port-50001" } } </pre> | <ul style="list-style-type: none"> • rn — The relative name of the destination object in the format: path-[path] . • enc — The encoding type of the telemetry data to be sent. NX-OS supports: <ul style="list-style-type: none"> • Google protocol buffers (GPB) for gRPC. • JSON for C. • proto — The transport protocol type of the telemetry data to be sent. NX-OS supports: <ul style="list-style-type: none"> • gRPC • HTTP • Supported encoded types are: <ul style="list-style-type: none"> • HTTP/JSON YES • HTTP/Form-data YES Only supported for Bin Logging. • GRPC/GPB-Compact YES Native Data Source Only. • GRPC/GPB YES • UDP/GPB YES <p>UDP/JSON YES</p> |
| Step 8 | <p>Enable gRPC chunking and set the chunking size, between 64 and 4096 bytes.</p> <p>Example:</p> <pre> { "telemetryDestGrpOptChunking": { "attributes": { "chunkSize": "2048", "dn": "sys/tm/dest-1/chunking" } } } </pre> | See gRPC Chunking section for more information. |
| Step 9 | <p>Create a telemetry subscription to configure the telemetry behavior</p> <p>Example:</p> <pre> "telemetrySubscription": { "attributes": { "id": "30", "rn": "subs-30" }, "children": [{ </pre> | <p>A telemetry subscription is defined in an object of class <code>telemetrySubscription</code> . Configure the following attributes of the object:</p> <ul style="list-style-type: none"> • id — An identifier for the subscription. Currently only numeric ID values are supported. • rn — The relative name of the subscription object in the format: subs-id . |

| | Command or Action | Purpose |
|----------------|---|--|
| | <pre>]] } } } } } </pre> | Children of the subscription object include relation objects for sensor groups (telemetryRsSensorGroupRel) and destination groups (telemetryRsDestGroupRel). |
| Step 10 | <p>Add the sensor group object as a child object to the telemetrySubscription element under the root element (telemetryEntity).</p> <p>Example:</p> <pre> { "telemetrySubscription": { "attributes": { "id": "30" } } "children": [{ "telemetryRsSensorGroupRel": { "attributes": { "sampleIntvl": "5000", "tDn": "sys/tm/sensor-10" } }] } </pre> | |
| Step 11 | <p>Create a relation object as a child object of the subscription to associate the subscription to the telemetry sensor group and to specify the data sampling behavior.</p> <p>Example:</p> <pre> "telemetryRsSensorGroupRel": { "attributes": { "rType": "mo", "rn": "rSensorGroupRel-[sys/tm/sensor-10]", "sampleIntvl": "5000", "tCl": "telemetrySensorGroup", "tDn": "sys/tm/sensor-10", "tType": "mo" } } </pre> | <p>The relation object is of class telemetryRsSensorGroupRel and is a child object of telemetrySubscription . Configure the following attributes of the relation object:</p> <ul style="list-style-type: none"> • rn — The relative name of the relation object in the format: rSensorGroupRel-[sys/tm/sensor-group-id] . • sampleIntvl — The data sampling period in milliseconds. An interval value of 0 creates an event-based subscription, in which telemetry data is sent only upon changes under the specified MO. An interval value greater than 0 creates a frequency-based subscription, in which telemetry data is sent periodically at the specified interval. For example, an interval value of 15000 results in the sending of telemetry data every 15 seconds. • tCl — The class of the target (sensor group) object, which is telemetrySensorGroup . • tDn — The distinguished name of the target (sensor group) object, which is sys/tm/sensor-group-id . • rType — The relation type, which is mo for managed object. <p>tType — The target type, which is mo for managed object.</p> |

| | Command or Action | Purpose |
|-----------------------|--|--|
| <p>Step 12</p> | <p>Define one or more sensor paths or nodes to be monitored for telemetry.</p> <p>Example: Single sensor path</p> <pre data-bbox="289 506 732 856"> { "telemetrySensorPath": { "attributes": { "path": "sys/cdp", "rn": "path-[sys/cdp]", "excludeFilter": "", "filterCondition": "", "path": "sys/fm/bgp", "secondaryGroup": "0", "secondaryPath": "", "depth": "0" } } } </pre> <p>Example: Multiple sensor paths</p> <pre data-bbox="289 1020 732 1675"> { "telemetrySensorPath": { "attributes": { "path": "sys/cdp", "rn": "path-[sys/cdp]", "excludeFilter": "", "filterCondition": "", "path": "sys/fm/bgp", "secondaryGroup": "0", "secondaryPath": "", "depth": "0" } } }, { "telemetrySensorPath": { "attributes": { "excludeFilter": "", "filterCondition": "", "path": "sys/fm/dhcp", "secondaryGroup": "0", "secondaryPath": "", "depth": "0" } } } </pre> <p>Example: Single sensor path filtering for BGP disable events:</p> <pre data-bbox="289 1835 305 1856"> { </pre> | <p>A sensor path is defined in an object of class <code>telemetrySensorPath</code>. Configure the following attributes of the object:</p> <ul style="list-style-type: none"> • <code>path</code> — The path to be monitored. • <code>rn</code> — The relative name of the path object in the format: <code>path-[path]</code> • <code>depth</code> — The retrieval level for the sensor path. A depth setting of 0 retrieves only the root MO properties. <p><code>filterCondition</code> — (Optional) Creates a specific filter for event-based subscriptions. The DME provides the filter expressions. For more information about filtering, see the Cisco APIC REST API Usage Guidelines on composing queries. You can find it at the following Cisco APIC documents landing page:</p> |

| | Command or Action | Purpose |
|----------------|--|--|
| | <pre> "telemetrySensorPath": { "attributes": { "path": "sys/cdp", "rn": "path-[sys/cdp]", "excludeFilter": "", "filterCondition": "eq(fmBgp.operSt.\"disabled\")", "path": "sys/fm/bgp", "secondaryGroup": "0", "secondaryPath": "", "depth": "0" } } </pre> | |
| Step 13 | Add sensor paths as child objects to the sensor group object (telemetrySensorGroup). | |
| Step 14 | Add destinations as child objects to the destination group object (telemetryDestGroup). | |
| Step 15 | Add the destination group object as a child object to the root element (telemetryEntity). | |
| Step 16 | <p>Create a relation object as a child object of the telemetry sensor group to associate the sensor group to the subscription.</p> <p>Example:</p> <pre> "telemetryRtSensorGroupRel": { "attributes": { "rn": "rtsensorGroupRel-[sys/tm/subs-30]", "tCl": "telemetrySubscription", "tDn": "sys/tm/subs-30" } } </pre> | <p>The relation object is of class telemetryRtSensorGroupRel and is a child object of telemetrySensorGroup. Configure the following attributes of the relation object:</p> <ul style="list-style-type: none"> • rn — The relative name of the relation object in the format: rtsensorGroupRel-[sys/tm/subscription-id]. • tCl — The target class of the subscription object, which is telemetrySubscription. <p>tDn — The target distinguished name of the subscription object, which is sys/tm/subscription-id.</p> |
| Step 17 | <p>Create a relation object as a child object of the telemetry destination group to associate the destination group to the subscription.</p> <p>Example:</p> <pre> "telemetryRtDestGroupRel": { "attributes": { "rn": "rtdestGroupRel-[sys/tm/subs-30]", "tCl": "telemetrySubscription", "tDn": "sys/tm/subs-30" } } </pre> | <p>The relation object is of class telemetryRtDestGroupRel and is a child object of telemetryDestGroup. Configure the following attributes of the relation object:</p> <ul style="list-style-type: none"> • rn — The relative name of the relation object in the format: rtdestGroupRel-[sys/tm/subscription-id]. • tCl — The target class of the subscription object, which is telemetrySubscription. <p>tDn — The target distinguished name of the subscription object, which is sys/tm/subscription-id.</p> |
| Step 18 | <p>Create a relation object as a child object of the subscription to associate the subscription to the telemetry destination group.</p> <p>Example:</p> <pre> "telemetryRsDestGroupRel": { </pre> | <p>The relation object is of class telemetryRsDestGroupRel and is a child object of telemetrySubscription. Configure the following attributes of the relation object:</p> <ul style="list-style-type: none"> • rn — The relative name of the relation object in the format: rsdestGroupRel-[sys/tm/destination-group-id]. |

| | Command or Action | Purpose |
|----------------|--|--|
| | <pre> "attributes": { "rType": "mo", "rn": "rsdestGroupRel-[sys/tm/dest-20]", "tCl": "telemetryDestGroup", "tDn": "sys/tm/dest-20", "tType": "mo" } </pre> | <ul style="list-style-type: none"> • tCl — The class of the target (destination group) object, which is telemetryDestGroup . • tDn — The distinguished name of the target (destination group) object, which is sys/tm/destination-group-id . • rType — The relation type, which is mo for managed object. <p>tType — The target type, which is mo for managed object.</p> |
| Step 19 | <p>Send the resulting JSON structure as an HTTP/HTTPS POST payload to the NX-API endpoint for telemetry configuration.</p> <p>Example:</p> | <p>The base path for the telemetry entity is <code>sys/tm</code> and the NX-API endpoint is:</p> <pre> {{URL}}/api/node/mo/sys/tm.json </pre> |

Example

The following is an example of all the previous steps that are collected into one POST payload (note that some attributes may not match):

```

{
  "telemetryEntity": {
    "children": [{
      "telemetrySensorGroup": {
        "attributes": {
          "id": "10"
        }
        "children": [{
          "telemetrySensorPath": { "attributes": {
            "excludeFilter": "",
            "filterCondition": "", "path": "sys/fm/bgp",
            "secondaryGroup": "0",
            "secondaryPath": "",
            "depth": "0"
          }
        }
      }
    ]
  },
  {
    "telemetryDestGroup": {
      "attributes": {
        "id": "20"
      }
      "children": [{
        "telemetryDest": {
          "addr": "10.30.217.80",
          "attributes": {
            "port": "50051",
            "enc": "GPB",
            "proto": "gRPC"
          }
        }
      }
    }
  }
}

```



```

    }
  }, {
    "telemetrySensorPath": {
      "attributes": {
        "path": "sys/cdp",
        "rn": "path-[sys/cdp]",
        "excludeFilter": "",
        "filterCondition": "",
        "secondaryGroup": "0",
        "secondaryPath": "",
        "depth": "0"
      }
    }
  }, {
    "telemetrySensorPath": {
      "attributes": {
        "path": "sys/ipv4",
        "rn": "path-[sys/ipv4]",
        "excludeFilter": "",
        "filterCondition": "",
        "secondaryGroup": "0",
        "secondaryPath": "",
        "depth": "0"
      }
    }
  ]
}
}, {
  "telemetryDestGroup": {
    "attributes": {
      "id": "20",
      "rn": "dest-20"
    },
    "children": [{
      "telemetryRtDestGroupRel": {
        "attributes": {
          "rn": "rtdestGroupRel-[sys/tm/subs-30]",
          "tCl": "telemetrySubscription",
          "tDn": "sys/tm/subs-30"
        }
      }
    }
  ], {
    "telemetryDest": {
      "attributes": {
        "addr": "1.2.3.4",
        "enc": "GPB",
        "port": "50001",
        "proto": "gRPC",
        "rn": "addr-[1.2.3.4]-port-50001"
      }
    }
  ]
}
}, {
  "telemetrySubscription": {
    "attributes": {
      "id": "30",
      "rn": "subs-30"
    },
    "children": [{
      "telemetryRsDestGroupRel": {
        "attributes": {
          "rType": "mo",
          "rn": "rsdestGroupRel-[sys/tm/dest-20]",

```

```

        "tCl": "telemetryDestGroup",
        "tDn": "sys/tm/dest-20",
        "tType": "mo"
      }
    }, {
      "telemetryRsSensorGroupRel": {
        "attributes": {
          "rType": "mo",
          "rn": "rssensorGroupRel-[sys/tm/sensor-10]",
          "sampleIntvl": "5000",
          "tCl": "telemetrySensorGroup",
          "tDn": "sys/tm/sensor-10",
          "tType": "mo"
        }
      }
    }
  ]
}

```

Filter Conditions on BGP Notifications

The following example payload enables notifications that trigger when the BFP feature is disabled as per the `filterCondition` attribute in the `telemetrySensorPath` MO. The data is streamed to `10.30.217.80` port `50055`.

POST `https://192.168.20.123/api/node/mo/sys/tm.json`

Payload:

```

{
  "telemetryEntity": {
    "children": [{
      "telemetrySensorGroup": {
        "attributes": {
          "id": "10"
        }
      }
    ]
  },
  "telemetrySensorPath": {
    "attributes": {
      "excludeFilter": "",
      "filterCondition": "eq(fmBgp.operSt,\"disabled\")",
      "path": "sys/fm/bgp",
      "secondaryGroup": "0",
      "secondaryPath": "",
      "depth": "0"
    }
  }
},
{
  "telemetryDestGroup": {
    "attributes": {
      "id": "20"
    }
  },
  "children": [{
    "telemetryDest": {
      "attributes": {
        "addr": "10.30.217.80",
        "port": "50055",
        "enc": "GPB",

```


MFDM is a Multicast FIB distribution management which consumes the information from the upper-level component, builds an intelligence for each multicast feature, and then propagates the information to the consumer. This is the core component where the feature is implemented along with DME. It is responsible for publishing all the multicast states to DME, based on the information provided by MRIB and the statistics collected by MFIB.

DME is used to store all the information that needs to be made available to the consumer/controller. It will also be responsible of generating the appropriate notifications to telemetry whenever an object is created or deleted or modified to support event-based notifications.

Telemetry process is responsible for streaming out all the data stored in DME to the consumers and format the data in proper form.

CLIs for Multicast Flow Path Visibility

The following are the CLIs that are introduced to verify the accurate functionality of the Multicast Flow Path Visibility:

- A configuration command to enable the export of information to DME. This CLI enables the feature for every route present in the system.

```
switch(config)# multicast flow-path export

switch(config)# sh system internal dme run all dn sys/mca/config
```

- A consistency checker show command to perform consistency checks between states present in MFDM and DME. This command allows you to catch inconsistencies quickly, especially on high scale setups.

```
switch# show forwarding distribution internal multicast
consistency-checker flow-path route
```

```
Starting flow-path DME consistency-check for VRF:
default
```

```
(0.0.0.0/0, 230.0.0.1/32). Result:
PASS
```

```
(10.0.0.10/32, 230.0.0.1/32). Result:
PASS
```

```
(0.0.0.0/0, 232.0.0.0/8). Result: PASS
```

- A global show command is used to check if the feature is enabled in the system or not.

```
switch(config)# show forwarding distribution internal
multicast global_state
```

```
**** MFDM Flow PATH VISIBILITY INFO
****
```

```
Multicast flow-path info export enabled:
Y
```

```
BE DME Handler: 0x117c3e6c
```

```
PE DME Handler: 0x117b955c
```

```
switch(config)# show forwarding distribution internal
multicast fpv CC
```

```
PASS/FAIL (In case of fail, it will highlight the
inconsistencies)
```

Cloud Scale Software Telemetry

About Cloud Scale Software Telemetry

Cisco NX-OS supports Cisco Nexus Cloud Scale switches that use the Tahoe ASIC. In such platform, supported Cloud Scale switches host a TCP/IP server that is tightly integrated with the ASICs, which expedites reporting telemetry data from the switch. The server runs on TCP port 7891, and telemetry clients can connect to the server on this port to retrieve hardware-counter data in a maximum of 10 milliseconds.

Cloud Scale software telemetry offers the users the flexibility of creating your own client programs or using the default client program that is bundled into NX-OS release 9.3.1 and later. The user can write client programs in any programming language that supports TCP/IP, such as Python 2.7 or higher, C, or PHP. Client programs must be constructed with the correct message formatting.

Beginning with NX-OS release 9.3(1), the Cloud Scale software telemetry feature is available in NX-OS. The feature is enabled by default, so supported switches running NX-OS 9.3(1) or later can use this feature.

Cloud Scale Software Telemetry Message Formats

Cloud Scale telemetry begins with a handshake between the client and TCP/IP server on the switch, during which the client initiates the connection over the TCP socket. The client message is a 32-bit integer set to zero. The switch responds with a message that contains the counter data in a specific format.

In NX-OS release 9.3(1), the following message format is supported. If you create your own client programs, make sure that the messages that your clients initiate conform to this format

| Length | Specifies |
|----------|---|
| 4 bytes | The number of ports, N |
| 56 bytes | The data for each port, for a total of $56 * N$ bytes. Each 56-byte chunk of data consists of the following: <ul style="list-style-type: none"> • 24 bytes of interface name • 8 bytes of the transmitted (TX) packets • 8 bytes of transmitted (TX) bytes • 8 bytes of received (RX) packets 8 bytes of received (RX) bytes |

Guidelines and Limitations for Cloud Scale Software Telemetry

The following are the guidelines and limitations for the Cloud Scale software telemetry feature:

- For information about supported platforms for Cisco NX-OS prior to release 9.3(x), see the section for <https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/platform/platform.html> in that guide. Starting with Cisco NX-OS release 9.3(x) for information about supported platforms, see <https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/platform/platform.html>.

- For custom client telemetry programs, one message format is supported. Your client programs must comply with this format.
- Beginning with Cisco NX-OS Release 10.3(1)F, software telemetry is supported on the Cisco Nexus 9800 platform switches.

Telemetry Path Labels

About Telemetry Path Labels

Beginning with NX-OS release 9.3(1), model-driven telemetry supports path labels. Path labels provide an easy way to gather telemetry data from multiple sources at once. With this feature, you specify the type of telemetry data you want collected, and the telemetry feature gathers that data from multiple paths. The feature then returns the information to one consolidated place, the path label. This feature simplifies using telemetry because you no longer must:

- Have a deep and comprehensive knowledge of the Cisco DME model.
- Create multiple queries and add multiple paths to the subscription, while balancing the number of collected events and the cadence.
- Collect multiple chunks of telemetry information from the switch, which simplifies serviceability.

Path labels span across multiple instances of the same object type in the model, then gather and return counters or events. Path labels support the following telemetry groups:

- **Environment**, which monitors chassis information, including fan, temperature, power, storage, supervisors, and line cards.
- **Interface**, which monitors all the interface counters and status changes.

This label supports predefined keyword filters that can refine the returned data by using the **query-condition** command.

- **Resources**, which monitors system resources such as CPU utilization and memory utilization.
- **VXLAN**, which monitors VXLAN EVPNs including VXLAN peers, VXLAN counters, VLAN counters, and BGP Peer data.

Polling for Data or Receiving Events

The sample interval for a sensor group determines how and when telemetry data is transmitted to a path label. The sample interval can be configured either to periodically poll for telemetry data or gather telemetry data when events occur.

- When the sample interval for telemetry is configured as a non-zero value, telemetry periodically sends the data for the environment, interfaces, resources, and vxlan labels during each sample interval.
- When the sample interval is set to zero, telemetry sends event notifications when the environment, interfaces, resources, and vxlan labels experience operational state updates, as well as creation and deletion of MOs.

Polling for data or receiving events are mutually exclusive. You can configure polling or event-driven telemetry for each path label.

Guidelines and Limitations for Path Labels

The telemetry path labels feature has the following guidelines and limitations:

- The feature supports only Cisco DME data source only.
- You cannot mix and match usability paths with regular DME paths in the same sensor group. For example, you cannot configure `sys/intf` and `interface` in the same sensor group. Also, you cannot configure the same sensor group with `sys/intf` and `interface`. If this situation occurs, NX-OS rejects the configuration.
- User filter keywords, such as `oper-speed` and `counters=[detailed]`, are supported only for the `interface` path.
- The feature does not support other sensor path options, such as `depth` or `filter-condition`.
- The telemetry path labels has the following restrictions in using path labels:
 - Must start with prefix **show** in lowercase, as it is case sensitive.

For example: **show version** is allowed. However, **SHOW version** or `version` is not allowed.

- Cannot include following characters:
 - ;
 - |
- " " or ''
 - Cannot include following words:
 - telemetry
 - conf t

`configure`

Configuring the Interface Path to Poll for Data or Events

The interface path label monitors all the interface counters and status changes. It supports the following interface types:

- Physical
- Subinterface
- Management
- Loopback
- VLAN
- Port Channel

You can configure the interface path label to either periodically poll for data or receive events. See Polling for Data or Receiving Events, on page 45.



Note The model does not support counters for subinterface, loopback, or VLAN, so they are not streamed out.

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*
4. **path interface**
5. **destination-group** *grp_id*
6. **ip address** *ip_addr* **port** *port*
7. **subscription** *sub_id*
8. **snsr-group** *sgrp_id* **sample-interval** *interval*
9. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enter configuration mode. |
| Step 2 | telemetry Example: <pre>switch(config)# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: <pre>switch(config-telemetry)# sensor-group 6 switch(conf-tm-sensor)#</pre> | Create a sensor group for telemetry data. |
| Step 4 | path interface Example: <pre>switch(conf-tm-sensor)# path interface switch(conf-tm-sensor)#</pre> | <p>Configure the interface path label, which enables sending one telemetry data query for multiple individual interfaces. The label consolidates the queries for multiple interfaces into one. Telemetry then telemetry gathers the data and returns it to the label.</p> <p>Depending on how the polling interval is configured, interface data is sent based on a periodic basis or whenever the interface state changes.</p> |

| | Command or Action | Purpose |
|--------|--|--|
| Step 5 | destination-group <i>grp_id</i> Example: <pre>switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)#</pre> | Enter telemetry destination group submode and configure the destination group. |
| Step 6 | ip address <i>ip_addr</i> port <i>port</i> Example: <pre>switch(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch(conf-tm-dest)#</pre> | Configure the telemetry data for the subscription to stream to the specified IP address and port |
| Step 7 | subscription <i>sub_id</i> Example: <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 8 | snsr-group <i>sgrp_id</i> sample-interval <i>interval</i> Example: <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur. |
| Step 9 | dst-group <i>dgrp_id</i> Example: <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Configuring the Interface Path for Non-Zero Counters

You can configure the interface path label with a pre-defined keyword filter that returns only counters that have non-zero values. The filter is `counters=[detailed]`.

By using this filter, the interface path gathers all the available interface counters, filters the collected data, then forwards the results to the receiver. The filter is optional, and if you do not use it, all counters, including zero-value counters, are displayed for the interface path.



Note Using the filter is conceptually similar to issuing `show interface mgmt0 counters detailed`

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*
4. **path interface query-condition** `counters=[detailed]`
5. **destination-group** *grp_id*
6. **ip address** *ip_addr* **port** *port*

7. **subscription** *sub_id*
8. **snsr-group** *sgrp_id* **sample-interval** *interval*
9. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enter configuration mode. |
| Step 2 | telemetry Example: <pre>switch(config)# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: <pre>switch(config-telemetry)# sensor-group 6 switch(conf-tm-sensor)#</pre> | Create a sensor group for telemetry data. |
| Step 4 | path interface query-condition counters=[detailed] Example: <pre>switch(conf-tm-sensor)# path interface query-condition counters=[detailed] switch(conf-tm-sensor)#</pre> | Configure the interface path label and query for only the non-zero counters from all interfaces. |
| Step 5 | destination-group <i>grp_id</i> Example: <pre>switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)#</pre> | Enter telemetry destination group submode and configure the destination group. |
| Step 6 | ip address <i>ip_addr</i> port <i>port</i> Example: <pre>switch(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch(conf-tm-dest)#</pre> | Configure the telemetry data for the subscription to stream to the specified IP address and port. |
| Step 7 | subscription <i>sub_id</i> Example: <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 8 | snsr-group <i>sgrp_id</i> sample-interval <i>interval</i> Example: | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling |

| | Command or Action | Purpose |
|---------------|---|---|
| | <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | interval determines whether the switch sends telemetry data periodically, or when interface events occur. |
| Step 9 | <p>dst-group <i>dgrp_id</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Configuring the Interface Path for Operational Speeds

You can configure the interface path label with a pre-defined keyword filter that returns counters for interfaces of specified operational speeds. The filter is `oper-speed=[]`. The following operational speeds are supported: auto, 10M, 100M, 1G, 10G, 40G, 200G, and 400G.

By using this filter, the interface path gathers the telemetry data for interfaces of the specified speed, then forwards the results to the receiver. The filter is optional. If you do not use it, counters for all interfaces are displayed, regardless of their operational speed.

The filter can accept multiple speeds as a comma-separated list, for example `oper-speed=[1G,10G]` to retrieve counters for interfaces that operate at 1 and 10 Gbps. Do not use a blank space as a delimiter.



Note Interface types subinterface, loopback, and VLAN do not have operational speed properties, so the filter does not support these interface types.

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **snsr-group** *sgrp_id* **sample-interval** *interval*
4. **path interface query-condition** **oper-speed**=[*speed*]
5. **destination-group** *grp_id*
6. **ip address** *ip_addr* *port* *port*
7. **subscription** *sub_id*
8. **snsr-group** *sgrp_id* **sample-interval** *interval*
9. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---------------------------|
| Step 1 | <p>configure terminal</p> <p>Example:</p> | Enter configuration mode. |

| | Command or Action | Purpose |
|---------------|--|--|
| | switch# configure terminal switch(config)# | |
| Step 2 | telemetry Example: switch(config)# telemetry switch(config-telemetry)# | Enter configuration mode for the telemetry features. |
| Step 3 | snsr-group <i>sgrp_id</i> sample-interval <i>interval</i> Example: switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)# | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur. |
| Step 4 | path interface query-condition oper-speed=[<i>speed</i>] Example: switch(conf-tm-sensor)# path interface query-condition oper-speed=[1G,40G] switch(conf-tm-sensor)# | Configure the interface path label and query for counters from interfaces running the specified speed, which in this example, is 1 and 40 Gbps only |
| Step 5 | destination-group <i>grp_id</i> Example: switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)# | Enter telemetry destination group submode and configure the destination group. |
| Step 6 | ip address <i>ip_addr</i> port <i>port</i> Example: switch(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch(conf-tm-dest)# | Configure the telemetry data for the subscription to stream to the specified IP address and port. |
| Step 7 | subscription <i>sub_id</i> Example: switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)# | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 8 | snsr-group <i>sgrp_id</i> sample-interval <i>interval</i> Example: switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)# | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur. |
| Step 9 | dst-group <i>dgrp_id</i> Example: switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)# | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Configuring the Interface Path with Multiple Queries

You can configure multiple filters for the same query condition in the interface path label. When you do so, the individual filters you use are ANDed.

Separate each filter in the query condition by using a comma. You can specify any number of filters for the query-condition, but be aware that the more filters you add, the more focused the results become.

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*
4. **path interface query-condition** **counters=[detailed],oper-speed=[1G,40G]**
5. **destination-group** *grp_id*
6. **ip address** *ip_addr* **port** *port*
7. **subscription** *sub_id*
8. **snsr-group** *sgrp_id* **sample-interval** *interval*
9. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enter configuration mode. |
| Step 2 | telemetry Example: <pre>switch(config)# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: <pre>switch(config-telemetry)# sensor-group 6 switch(conf-tm-sensor)#</pre> | Create a sensor group for telemetry data. |
| Step 4 | path interface query-condition counters=[detailed],oper-speed=[1G,40G] Example: <pre>switch(conf-tm-sensor)# path interface query-condition counters=[detailed],oper-speed=[1G,40G] switch(conf-tm-sensor)#</pre> | Configures multiple conditions in the same query. In this example, the query does both of the following: <ul style="list-style-type: none"> • Gathers and returns non-zero counters on interfaces running at 1 Gbps. Gathers and returns non-zero counters on interfaces running at 40 Gbps. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 5 | destination-group <i>grp_id</i> Example: <pre>switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)#</pre> | Enter telemetry destination group submode and configure the destination group. |
| Step 6 | ip address <i>ip_addr</i> port <i>port</i> Example: <pre>switch(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch(conf-tm-dest)#</pre> | Configure the telemetry data for the subscription to stream to the specified IP address and port. |
| Step 7 | subscription <i>sub_id</i> Example: <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 8 | snsr-group <i>sgrp_id</i> sample-interval <i>interval</i> Example: <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur. |
| Step 9 | dst-group <i>dgrp_id</i> Example: <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Configuring the Environment Path to Poll for Data or Events

The environment path label monitors chassis information, including fan, temperature, power, storage, supervisors, and line cards. You can configure the environment path to either periodically poll for telemetry data or get the data when events occur. For information, see *Polling for Data or Receiving Events*, on page 45.

You can set the resources path to return system resource information through either periodic polling or based on events. This path does not support filtering.

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*
4. **path environment**
5. **destination-group** *grp_id*
6. **ip address** *ip_addr* **port** *port*
7. **subscription** *sub_id*
8. **snsr-group** *sgrp_id* **sample-interval** *interval*
9. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enter configuration mode. |
| Step 2 | telemetry Example: <pre>switch(config)# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for the telemetry features |
| Step 3 | sensor-group <i>sgrp_id</i> Example: <pre>switch(config-telemetry)# sensor-group 6 switch(conf-tm-sensor)#</pre> | Create a sensor group for telemetry data. |
| Step 4 | path environment Example: <pre>switch(conf-tm-sensor)# path environment switch(conf-tm-sensor)#</pre> | <p>Configures the environment path label, which enables telemetry data for multiple individual environment objects to be sent to the label. The label consolidates the multiple data inputs into one output.</p> <p>Depending on the sample interval, the environment data is either streaming based on the polling interval, or sent when events occur.</p> |
| Step 5 | destination-group <i>grp_id</i> Example: <pre>switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)#</pre> | Enter telemetry destination group submode and configure the destination group. |
| Step 6 | ip address <i>ip_addr</i> port <i>port</i> Example: <pre>switch(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch(conf-tm-dest)#</pre> | Configure the telemetry data for the subscription to stream to the specified IP address and port. |
| Step 7 | subscription <i>sub_id</i> Example: <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 8 | snsr-group <i>sgrp_id</i> sample-interval <i>interval</i> Example: <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when environment events occur. |

| | Command or Action | Purpose |
|---------------|---|---|
| Step 9 | dst-group <i>dgrp_id</i> Example: <pre>switch(conf-tm-sub) # dst-grp 33 switch(conf-tm-sub) #</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Enabling Power Usage Tracking Functionality

Starting from NX-OS Release 10.4.1(F), the **power usage-history** command is supported on Cisco Nexus 9336C-FX2 and 9332D-GX2B switches to track power consumption. By default this feature is disabled.

Beginning with Cisco NX-OS Release 10.4(2)F, the **power usage-history** command is supported on Cisco Nexus 9000 Series platform switches.

Follow the steps to enable the feature.

SUMMARY STEPS

1. **configure terminal**
2. **[no] power usage-history**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config) #</pre> | Enters configuration mode. |
| Step 2 | [no] power usage-history Example: <pre>switch# power usage-history switch(config) #</pre> | Enables power usage tracking feature. Use the no form of this command to disable this feature. |

Displaying Power Consumption History

Power Usage Tracking Show Command

See [Enabling Power Usage Tracking Functionality, on page 488](#) to enable power usage tracking feature. After enabling, use **show environment power history** to display power usage statistics for various targets.

| Command | Shows |
|--|--|
| show environment power history { peak target 1min target 1hr target 14 days target 14days day _day_no } | Power usage information for various targets. |

Command Examples

The following shows an example of the **show environment power history peak** command.

```
switch# show environment power history peak
Power                               Output Power      Peak Time          Input
Power                               (peak)           (Output Power)    (peak)
Supply    Model                      Status
(Input Power)
-----
1         N9K-PAC-3000W-B             Ok                334.30 W          06/07/2023 10:54:29  362.45 W
06/07/2023 10:54:59
2         N9K-PAC-3000W-B             Ok                362.45 W          06/07/2023 10:53:29  425.80 W
06/07/2023 10:51:44
switch#
```

Last 1min usage data would contain average usage in last 15secs, 30secs and 60 secs.

```
module-4# show environment power history target 1min

Power                               Output/Input      Output/Input
Output/Input                         (15 sec)         (30 sec)
Supply    Model                      Status
(60 sec)
-----
1         N9K-PAC-3000W-B             Ok                330.78 W / 362.45 W  330.00 W / 362.00 W
330.00 W / 362.00 W
2         N9K-PAC-3000W-B             Ok                358.94 W / 415.24 W  358.00 W / 415.00 W
358.00 W / 415.00 W
switch#
```

The following shows the output of the **show environment power history target 1hr** command.

```
switch# show environment power history target 1hr
1 min avg data for 1 Hr for
slot: 1 Product Name: N9K-PAC-3000W-B status: Ok
Output Power      Input Power      Time
-----
331.00 W          362.00 W        06/07/2023 11:34:44
330.00 W          362.00 W        06/07/2023 11:33:44
333.00 W          362.00 W        06/07/2023 11:32:44
333.00 W          362.00 W        06/07/2023 11:31:44
331.00 W          362.00 W        06/07/2023 11:30:44

1 min avg data for 1 Hr for
slot: 2 Product Name: N9K-PAC-3000W-B status: Ok
Output Power      Input Power      Time
-----
358.00 W          417.00 W        06/07/2023 11:34:44
358.00 W          417.00 W        06/07/2023 11:33:44
358.00 W          417.00 W        06/07/2023 11:32:44
358.00 W          417.00 W        06/07/2023 11:31:44
357.00 W          415.00 W        06/07/2023 11:30:44
```

The following shows an example of the **show environment power history target 24hr** command.

```
switch# show environment power history target 24hr
1HR avg data for 24 Hr for
slot: 1 Product Name: N9K-PAC-3000W-B status: Ok
Output Power      Input Power      Time
```

```

-----
332.15 W          363.56 W          06/07/2023 12:50:44
332.13 W          363.66 W          06/07/2023 11:50:44

1HR avg data for 24 Hr for
slot: 2 Product Name: N9K-PAC-3000W-B status: Ok
  Output Power      Input Power          Time
-----
358.23 W          416.68 W          06/07/2023 12:50:44
358.35 W          417.05 W          06/07/2023 11:50:44
switch#

```

The following shows an example of the **show environment power history target 14days** command.

```

switch# show environment power history target 14days
 1 Day avg data over a period of 14 days
slot: 1 Product Name: N9K-PAC-3000W-B status: Ok
Day  Output Power      Input Power          Date
-----
 1      332.17 W          363.61 W          06/07/23

 1 Day avg data over a period of 14 days
slot: 2 Product Name: N9K-PAC-3000W-B status: Ok
Day  Output Power      Input Power          Date
-----
 1      358.23 W          416.81 W          06/07/23
switch#

```

This CLI displays the average usage throughout the day for each day in last 14days. For each PSU 14 days average usage is displayed. A detailed per hour usage for each day is displayed when day number is given. Output for that is given in next slide.

The following shows an example of the **show environment power history target 14days day 1** command.

```

switch# show environment power history target 14days day 1
 1 HR avg data for 1 Day
slot: 1 Product Name: N9K-PAC-3000W-B status: Ok
Day 1
  Output Power      Input Power          Time
-----
 332.23 W          363.61 W          06/07/2023 13:50:44
 332.15 W          363.56 W          06/07/2023 12:50:44
 332.13 W          363.66 W          06/07/2023 11:50:44

 1 HR avg data for 1 Day
slot: 2 Product Name: N9K-PAC-3000W-B status: Ok
Day 1
  Output Power      Input Power          Time
-----
 358.11 W          416.71 W          06/07/2023 13:50:44
 358.23 W          416.68 W          06/07/2023 12:50:44
 358.35 W          417.05 W          06/07/2023 11:50:44
switch#
switch#

```

Configuring the Resources Path to Poll for Events or Data

The resources path monitors system resources such as CPU utilization and memory utilization. You can configure this path to either periodically gather telemetry data, or when events occur. See [Polling for Data or Receiving Events](#), on page 45.

This path does not support filtering.

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*
4. **path resources**
5. **destination-group** *grp_id*
6. **ip address** *ip_addr* **port** *port*
7. **subscription** *sub_id*
8. **snsr-group** *sgrp_id* **sample-interval** *interval*
9. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enter configuration mode. |
| Step 2 | telemetry Example: <pre>switch(config)# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: <pre>switch(config-telemetry)# sensor-group 6 switch(conf-tm-sensor)#</pre> | Create a sensor group for telemetry data. |
| Step 4 | path resources Example: <pre>switch(conf-tm-sensor)# path resources switch(conf-tm-sensor)#</pre> | <p>Configure the resources path label, which enables telemetry data for multiple individual system resources to be sent to the label. The label consolidates the multiple data inputs into one output.</p> <p>Depending on the sample interval, the resource data is either streaming based on the polling interval, or sent when system memory changes to Not OK.</p> |
| Step 5 | destination-group <i>grp_id</i> Example: <pre>switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)#</pre> | Enter telemetry destination group submode and configure the destination group. |
| Step 6 | ip address <i>ip_addr</i> port <i>port</i> Example: | Configure the telemetry data for the subscription to stream to the specified IP address and port. |

| | Command or Action | Purpose |
|---------------|--|---|
| | <pre>switch(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch(conf-tm-dest)#</pre> | |
| Step 7 | <p>subscription <i>sub_id</i></p> <p>Example:</p> <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 8 | <p>snsr-group <i>sgrp_id</i> sample-interval <i>interval</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when resource events occur. |
| Step 9 | <p>dst-group <i>dgrp_id</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Configuring the VXLAN Path to Poll for Events or Data

The vxlan path label provides information about the switch's Virtual Extensible LAN EVPNs, including VXLAN peers, VXLAN counters, VLAN counters, and BGP Peer data. You can configure this path label to gather telemetry information either periodically, or when events occur. See [Telemetry Data Streamed for Native Data Source Paths, on page 497](#).

This path does not support filtering.

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*
4. **vxlan environment**
5. **destination-group** *grp_id*
6. **ip address** *ip_addr* **port** *port*
7. **subscription** *sub_id*
8. **snsr-group** *sgrp_id* **sample-interval** *interval*
9. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enter configuration mode. |
| Step 2 | telemetry Example: switch(config)# telemetry switch(config-telemetry)# | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: switch(config-telemetry)# sensor-group 6 switch(conf-tm-sensor)# | Create a sensor group for telemetry data. |
| Step 4 | vxlan environment Example: switch(conf-tm-sensor)# vxlan environment switch(conf-tm-sensor)# | Configure the vxlan path label, which enables telemetry data for multiple individual VXLAN objects to be sent to the label. The label consolidates the multiple data inputs into one output. Depending on the sample interval, the VXLAN data is either streaming based on the polling interval, or sent when events occur. |
| Step 5 | destination-group <i>grp_id</i> Example: switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)# | Enter telemetry destination group submode and configure the destination group. |
| Step 6 | ip address <i>ip_addr</i> port <i>port</i> Example: switch(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch(conf-tm-dest)# | Configure the telemetry data for the subscription to stream to the specified IP address and port. |
| Step 7 | subscription <i>sub_id</i> Example: switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)# | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 8 | snsr-group <i>sgrp_id</i> sample-interval <i>interval</i> Example: switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)# | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when VXLAN events occur. |

| | Command or Action | Purpose |
|---------------|---|---|
| Step 9 | dst-group <i>dgrp_id</i> Example: <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Verifying the Path Label Configuration

At any time, you can verify that path labels are configured, and check their values by displaying the running telemetry configuration.

SUMMARY STEPS

1. **show running-config-telemetry**

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | show running-config-telemetry Example: <pre>switch(conf-tm-sensor)# show running-config telemetry !Command: show running-config telemetry !Running configuration last done at: Mon Jun 10 08:10:17 2019 !Time: Mon Jun 10 08:10:17 2019 version 9.3(1) Bios:version feature telemetry telemetry destination-profile use-nodeid tester sensor-group 4 path interface query-condition and(counters=[detailed],oper-speed=[1G,10G]) sensor-group 6 path interface query-condition oper-speed=[1G,40G] subscription 6 snsr-grp 6 sample-interval 6000 nxosv2(conf-tm-sensor)#</pre> | Displays the current running config for telemetry. In this example, sensor group 4 is configured to gather non-zero counters from interfaces running at 1 and 10 Gbps. Sensor group 6 is configured to gather all counters from interfaces running at 1 and 40 Gbps. |

Displaying Path Label Information

Path Label Show Commands

Through the **show telemetry usability** commands, you can display the individual paths that the path label walks when you issue a query.

| Command | Shows |
|---|--|
| <code>show telemetry usability {all environment interface resources vxlan}</code> | Either all telemetry paths for all path labels or the paths for the specified path label. Also, the output shows the configuration data based on periodic polling or event-driven polling. For the interfaces path label, also any keywords that are configured. |
| <code>show running-config telemetry</code> | The running configuration for telemetry. |

Command Examples



Note The `show telemetry usability all` command is a concatenation of all the individual commands that are shown in this section.

The following shows an example of the **show telemetry usability environment** command.

```
switch# show telemetry usability environment
1) label_name : environment
path_name : sys/ch query_type : poll query_condition :
rsp-subtree-full&query-target-subtree&target-subtree-class=eqptPsuSlot,eqptFtSlot,eqptSupCSlot,eqptPsu,eqptFt,eqptSensor,eqptLCSlot
2) label_name : environment
path_name : sys/ch query_type : event query_condition :
&query-target-subtree=eq(LBsysIf.adminSt,"up")&rsp-subtree-children&rsp-subtree-class=monEthStats,monIfIn,monIfOut,monIfCIn,monIfCOut
switch#
```

The following shows the output of the **show telemetry usability interface** command.

```
switch# show telemetry usability interface
1) label_name : interface
path_name : sys/intf query_type : poll query_condition :
&query-target-children&query-target-filter=eq(LBsysIf.adminSt,"up")&rsp-subtree-children&rsp-subtree-class=monEthStats,monIfIn,monIfOut,monIfCIn,monIfCOut
2) label_name : interface
path_name : sys/mgmt-[mgmt0] query_type : poll query_condition :
&query-target-subtree&query-target-filter=eq(mgmtMgtIf.adminSt,"up")&rsp-subtree-full&rsp-subtree-class=monEthStats,monIfIn,monIfOut,monIfCIn,monIfCOut
3) label_name : interface
path_name : sys/intf query_type : event query_condition :
&query-target-children&query-target-filter=eq(LBsysIf.adminSt,"down"),and(updated(ethpmEncRtdIf.operSt,"down")),and(updated(ethpmEncRtdIf.operSt,"up")))
4) label_name : interface
path_name : sys/mgmt-[mgmt0] query_type : event query_condition :
&query-target-subtree&query-target-filter=or(blead(),cstat()),and(updated(mgmtIf.operSt,"down")),and(updated(mgmtIf.operSt,"up")))
switch#
```

The following shows an example of the **show telemetry usability resources** command.

```
switch# show telemetry usability resources
1) label_name : resources
path_name : sys/proc
```

```

query_type          : poll
query_condition     : rsp-subtree=full&rsp-foreign-subtree=ephemeral

2) label_name      : resources

path_name          : sys/procsys
query_type         : poll
query_condition    :
query-target-filter=and(updated(procSysMem.memstatus),ne(procSysMem.memstatus,"OK"))

3) label_name      : resources

path_name          : sys/procsys/systemem
query_type         : event
query_condition    :
query-target-filter=and(updated(procSysMem.memstatus),ne(procSysMem.memstatus,"OK"))

switch#

```

The following shows an example of the **show telemetry usability vxlan** command.

```

switch# show telemetry usability vxlan
1) label_name      : vxlan

path_name          : sys/bd
query_type         : poll
query_condition    : query-target=subtree&target-subtree-class=l2VlanStats

2) label_name      : vxlan

path_name          : sys/eps
query_type         : poll
query_condition    : rsp-subtree=full&rsp-foreign-subtree=ephemeral

3) label_name      : vxlan

path_name          : sys/eps
query_type         : event
query_condition    : query-target=subtree&target-subtree-class=nvoDyPeer

4) label_name      : vxlan

path_name          : sys/bgp
query_type         : event
query_condition    : query-target=subtree&query-target-filter=or(deleted(),created())

5) label_name      : vxlan

path_name          : sys/bgp
query_type         : event
query_condition    :
query-target-subtree=and(updated(procSysMem.memstatus),ne(procSysMem.memstatus,"OK"))
query-target-subtree-class=bgpDm,bgpPeer,bgpPeerAf,bgpDmAf,bgpPeerAfEntry,bgpOperPctrlI3,bgpOperPttP,bgpOperPttEntry,bgpOperAfCtrl

switch#

```

Native Data Source Paths

About Native Data Source Paths

NX-OS Telemetry supports the native data source, which is a neutral data source that is not restricted to a specific infrastructure or database. Instead, the native data source enables components or applications to hook into and inject relevant information into the outgoing telemetry stream. This feature provides flexibility because the path for the native data source does not belong to any infrastructure, so any native applications can interact with NX-OS Telemetry.

The native data source path enables you to subscribe to specific sensor paths to receive selected telemetry data. The feature works with the NX-SDK to support streaming telemetry data from the following paths:

- RIB path, which sends telemetry data for the IP routes.
- MAC path, which sends telemetry data for static and dynamic MAC entries.
- Adjacency path, which sends telemetry data for IPv4 and IPv6 adjacencies.

When you create a subscription, all telemetry data for the selected path streams to the receiver as a baseline. After the baseline, only event notifications stream to the receiver.

Streaming of native data source paths supports the following encoding types:

- Google Protobuf (GPB)
- JavaScript Object Notation (JSON)
- Compact Google Protobuf (compact GPB)

Telemetry Data Streamed for Native Data Source Paths

For each source path, the following table shows the information that is streamed when the subscription is first created (the baseline) and when event notifications occur.

| Path Type | Subscription Baseline | Event Notif |
|-----------|-----------------------|-------------|
|-----------|-----------------------|-------------|

| | | |
|------------|---|--|
| <p>RIB</p> | <p>Sends all routes</p> | <p>Sends event notifications for create and delete events. The following fields are included in the telemetry for the event:</p> <ul style="list-style-type: none"> • Next-hop name • Address • Outgoing interface • VRF name • Owner • Preference • Metric • Tag • Segment • Tunnel • Encapsulation • Bitwise • Type <p>• For Layer-3:</p> <ul style="list-style-type: none"> • VRF name • Route • Mask • Number • Event <p>Next hops</p> |
| <p>MAC</p> | <p>Executes a <code>GETALL</code> from DME for static and dynamic MAC entries</p> | <p>Sends event notifications for create and delete events. The following fields are included in the telemetry for the event:</p> <ul style="list-style-type: none"> • MAC address • MAC address • VLAN number • Interface name • Event type <p>Both static and dynamic MAC entries are included in the event notification.</p> |

| | | |
|-----------|-------------------------------------|---|
| Adjacency | Sends the IPv4 and IPv6 adjacencies | Sends event events. The telemetry for <ul style="list-style-type: none"> • IP address • MAC address • Interface • Physical • VRF name • Preference • Source • Address Adjacency e |
|-----------|-------------------------------------|---|

For additional information, refer to Github <https://github.com/CiscoDevNet/nx-telemetry-proto>.

Guidelines and Limitations

The native data source path feature has the following guidelines and limitations:

- For streaming from the RIB, MAC, and Adjacency native data source paths, sensor-path property updates do not support custom criteria like **depth**, **query-condition**, or **filter-condition**.
- Beginning with Cisco NX-OS Release 10.4(3)F, new query conditions are introduced to support sample-based subscription or updates-only support for RIB native path.

Configuring the Native Data Source Path for Routing Information

You can configure the native data source path for routing information, which sends information about all routes that are contained in the URIB. When you subscribe, the baseline sends all the route information. After the baseline, notifications are sent for route update and delete operations for the routing protocols that the switch supports. For the data sent in the RIB notifications, see Telemetry Data Streamed for Native Data Source Paths, on page 61.

Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group *sgrp_id***
4. **data-source native**
5. **path rib query-condition [data=ephemeral | updates_only]**

6. **destination-group** *grp_id*
7. **ip address** *ip_addr* **port** *port* **protocol** { HTTP | gRPC } **encoding** { JSON | GPB | GPB-compact }
8. **subscription** *sub_id*
9. **snsr-group** *sgrp_id* **sample-interval** *interval*
10. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enter configuration mode. |
| Step 2 | telemetry Example: <pre>switch(config)# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: <pre>switch(conf-tm-sub)# sensor-grp 6 switch(conf-tm-sub)#</pre> | Create a sensor group. |
| Step 4 | data-source native Example: <pre>switch(conf-tm-sensor)# data-source native switch(conf-tm-sensor)#</pre> | Set the data source to native so that any native application can use the streamed data without requiring a specific model or database. |
| Step 5 | path rib query-condition [data=ephemeral updates_only] Example: <pre>nxosv2(conf-tm-sensor)# path rib nxosv2(conf-tm-sensor)#</pre> Example: <pre>nxosv2(conf-tm-sensor)# path rib query condition data=ephemeral nxosv2(conf-tm-sensor)#</pre> Example: <pre>nxosv2(conf-tm-sensor)# path rib query condition updates_only nxosv2(conf-tm-sensor)#</pre> | Configure the RIB path which streams routes and route update information. query condition data=ephemeral (optional) - You can configure sample interval 0 or other than 0. This sample interval will determine how frequently the route information is sent to the destination periodically (at the configured sample interval). query condition updates-only (optional) - Supported only for sample interval 0. With this query condition, the initial snapshot data will not be sent, only the route information updates will be sent to the destination. |
| Step 6 | destination-group <i>grp_id</i> Example: | Enter telemetry destination group submode and configure the destination group. |

| | Command or Action | Purpose |
|----------------|---|---|
| | <pre>switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)#</pre> | |
| Step 7 | <p>ip address <i>ip_addr</i> port <i>port</i> protocol { HTTP gRPC } encoding { JSON GPB GPB-compact }</p> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch(conf-tm-dest)# Example: switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch(conf-tm-dest)# Example: switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch(conf-tm-dest)#</pre> | Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream. |
| Step 8 | <p>subscription <i>sub_id</i></p> <p>Example:</p> <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 9 | <p>snsr-group <i>sgrp_id</i> sample-interval <i>interval</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | <p>Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when rib events occur.</p> <p>Note Depending on the sample interval, the rib sensor path streams based on the polling interval.</p> |
| Step 10 | <p>dst-group <i>dgrp_id</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Configuring the Native Data Source Path for MAC Information

You can configure the native data source path for MAC information, which sends information about all entries in the MAC table. When you subscribe, the baseline sends all the MAC information. After the baseline, notifications are sent for add, update, and delete MAC address operations. For the data sent in the MAC notifications, see [Telemetry Data Streamed for Native Data Source Paths, on page 497](#).



Note For update or delete events, MAC notifications are sent only for the MAC addresses that have IP adjacencies.

Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*
4. **data-source native**
5. **path mac**
6. **destination-group** *grp_id*
7. **ip address** *ip_addr* **port** *port* **protocol** { HTTP | gRPC } **encoding** { JSON | GPB | GPB-compact }
8. **subscription** *sub_id*
9. **snsr-group** *sgrp_id* **sample-interval** *interval*
10. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enter configuration mode. |
| Step 2 | telemetry Example: switch(config)# telemetry switch(config-telemetry)# | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: switch(conf-tm-sub)# sensor-grp 6 switch(conf-tm-sub)# | Create a sensor group. |
| Step 4 | data-source native Example: switch(conf-tm-sensor)# data-source native switch(conf-tm-sensor)# | Set the data source to native so that any native application can use the streamed data without requiring a specific model or database. |
| Step 5 | path mac Example: nxosv2(conf-tm-sensor)# path mac nxosv2(conf-tm-sensor)# | Configure the MAC path which streams information about MAC entries and MAC notifications. |
| Step 6 | destination-group <i>grp_id</i> Example: switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)# | Enter telemetry destination group submode and configure the destination group. |

| | Command or Action | Purpose |
|---------|---|--|
| Step 7 | <p>ip address <i>ip_addr</i> port <i>port</i> protocol { HTTP gRPC } encoding { JSON GPB GPB-compact }</p> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch(conf-tm-dest)#</pre> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch(conf-tm-dest)#</pre> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch(conf-tm-dest)#</pre> | Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream. |
| Step 8 | <p>subscription <i>sub_id</i></p> <p>Example:</p> <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 9 | <p>snsr-group <i>sgrp_id</i> sample-interval <i>interval</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur. |
| Step 10 | <p>dst-group <i>dgrp_id</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Configuring the Native Data Source Path for all MAC Information

You can configure the native data source path for MAC information, which sends information about all entries in the MAC table from Layer 3 and Layer 2. When you subscribe, the baseline sends all the MAC information.

After the baseline, notifications are sent for add, update, and delete MAC address operations. For the data sent in the MAC notifications, see [Telemetry Data Streamed for Native Data Source Paths, on page 497](#).



Note For update or delete events, MAC notifications are sent only for the MAC addresses that have IP adjacencies.

Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*
4. **data-source** *native*
5. **path mac-all**
6. **destination-group** *grp_id*
7. **ip address** *ip_addr* **port** *port* **protocol** { HTTP | gRPC } **encoding** { JSON | GPB | GPB-compact }
8. **subscription** *sub_id*
9. **snsr-group** *sgrp_id* **sample-interval** *interval*
10. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | configure terminal Example: switch# configure terminal switch(config)# | Enter configuration mode. |
| Step 2 | telemetry Example: switch(config)# telemetry switch(config-telemetry)# | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: switch(conf-tm-sub)# sensor-grp 6 switch(conf-tm-sub)# | Create a sensor group. |
| Step 4 | data-source <i>native</i> Example: switch(conf-tm-sensor)# data-source native switch(conf-tm-sensor)# | Set the data source to native so that any native application can use the streamed data without requiring a specific model or database. |
| Step 5 | path mac-all Example: nxosv2(conf-tm-sensor)# path mac-all nxosv2(conf-tm-sensor)# | Configure the MAC path which streams information about all MAC entries and MAC notifications. |
| Step 6 | destination-group <i>grp_id</i> Example: switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)# | Enter telemetry destination group submode and configure the destination group. |

| | Command or Action | Purpose |
|---------|---|--|
| Step 7 | <p>ip address <i>ip_addr</i> port <i>port</i> protocol { HTTP gRPC } encoding { JSON GPB GPB-compact }</p> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch(conf-tm-dest)#</pre> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch(conf-tm-dest)#</pre> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch(conf-tm-dest)#</pre> | Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream. |
| Step 8 | <p>subscription <i>sub_id</i></p> <p>Example:</p> <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 9 | <p>snsr-group <i>sgrp_id</i> sample-interval <i>interval</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur. |
| Step 10 | <p>dst-group <i>dgrp_id</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Configuring the Native Data Path for IP Adjacencies

You can configure the native data source path for IP adjacency information, which sends information about all IPv4 and IPv6 adjacencies for the switch. When you subscribe, the baseline sends all the adjacencies. After the baseline, notifications are sent for add, update, and delete adjacency operations. For the data sent in the adjacency notifications, see Telemetry Data Streamed for Native Data Source Paths, on page 61.

Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp_id*

4. **data-source** *native*
5. **path adjacency**
6. **destination-group** *grp_id*
7. **ip address** *ip_addr* **port** *port* **protocol** { HTTP | gRPC } **encoding** { JSON | GPB | GPB-compact }
8. **subscription** *sub_id*
9. **snsr-group** *sgrp_id* **sample-interval** *interval*
10. **dst-group** *dgrp_id*

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enter configuration mode. |
| Step 2 | telemetry Example: <pre>switch(config)# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: <pre>switch(conf-tm-sub)# sensor-grp 6 switch(conf-tm-sub)#</pre> | Create a sensor group. |
| Step 4 | data-source <i>native</i> Example: <pre>switch(conf-tm-sensor)# data-source native switch(conf-tm-sensor)#</pre> | Set the data source to native so that any native application can use the streamed data. |
| Step 5 | path adjacency Example: <pre>nxosv2(conf-tm-sensor)# path adjacency nxosv2(conf-tm-sensor)#</pre> | Configure the Adjacency path which streams information about the IPv4 and IPv6 adjacencies. |
| Step 6 | destination-group <i>grp_id</i> Example: <pre>switch(conf-tm-sensor)# destination-group 33 switch(conf-tm-dest)#</pre> | Enter telemetry destination group submode and configure the destination group. |
| Step 7 | ip address <i>ip_addr</i> port <i>port</i> protocol { HTTP gRPC } encoding { JSON GPB GPB-compact } Example: | Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream. |

| | Command or Action | Purpose |
|----------------|---|--|
| | <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch(conf-tm-dest)#</pre> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch(conf-tm-dest)#</pre> <p>Example:</p> <pre>switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch(conf-tm-dest)#</pre> | |
| Step 8 | <p>subscription <i>sub_id</i></p> <p>Example:</p> <pre>switch(conf-tm-dest)# subscription 33 switch(conf-tm-sub)#</pre> | Enter telemetry subscription submode, and configure the telemetry subscription. |
| Step 9 | <p>snsr-group <i>sgrp_id</i> sample-interval <i>interval</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch(conf-tm-sub)#</pre> | Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur. |
| Step 10 | <p>dst-group <i>dgrp_id</i></p> <p>Example:</p> <pre>switch(conf-tm-sub)# dst-grp 33 switch(conf-tm-sub)#</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Displaying Native Data Source Path Information

Use the NX-OS **show telemetry event collector** commands to display statistics and counters, or errors for the native data source path.

Displaying Statistics

You can issue **show telemetry event collector stats** command to display the statistics and counters for each native data source path.

An example of statistics for the RIB path:

```
switch# show telemetry event collector stats
-----
Row ID Collection Count Latest Collection Time Sensor Path(GroupId)
-----
```

```
1 4 Mon Jul 01 13:53:42.384 PST rib(1) switch#
```

An example of the statistics for the MAC path:

```
switch# show telemetry event collector stats
-----
Row ID Collection Count Latest Collection Time Sensor Path(GroupId)
-----
```

```
1 3 Mon Jul 01 14:01:32.161 PST mac(1) switch#
```

An example of the statistics for the Adjacency path:

```
switch# show telemetry event collector stats
-----
Row ID Collection Count Latest Collection Time Sensor Path(GroupId)
-----
1 7 Mon Jul 01 14:47:32.260 PST adjacency(1)
switch#
```

Displaying Error Counters

You can use the **show telemetry event collector stats** command to display the error totals for all the native data source paths.

```
switch# show telemetry event collector errors
-----
Error Description Error Count
-----
Dme Event Subscription Init Failures - 0
Event Data Enqueue Failures - 0
Event Subscription Failures - 0
Pending Subscription List Create Failures - 0
Subscription Hash Table Create Failures - 0
Subscription Hash Table Destroy Failures - 0
Subscription Hash Table Insert Failures - 0
Subscription Hash Table Remove Failures - 0
switch#
```

Streaming Syslog

About Streaming Syslog for Telemetry

Beginning with Cisco NX-OS release 9.3(3), model-driven telemetry supports streaming of syslogs using YANG as a data source. When you create a subscription, all the syslogs are streamed to the receiver as a baseline. This feature works with the NX-SDK to support streaming syslog data from the following syslog paths:

- Cisco-NX-OS-Syslog-oper:syslog
- Cisco-NX-OS-Syslog-oper:syslog/messages

After the baseline, only syslog event notifications stream to the receiver. Streaming of syslog paths supports the following encoding types:

- Google Protobuf (GPB)
- JavaScript Object Notation (JSON)

Configuring the YANG Data Source Path for Syslog Information

You can configure the syslog path for syslogs, which sends information about all syslogs that are generated on the switch. When you subscribe, the baseline sends all the existing syslog information. After the baseline, notifications are sent for only for new syslogs that are generated on the switch.

Before you begin

If you have not enabled the telemetry feature, enable it now with the **feature telemetry** command.

SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group *sgrp_id***
4. **data source *data-source-type***
5. **path Cisco-NX-OS-Syslog-oper:syslog/messages**
6. **destination-group *grp_id***
7. **ip address *ip_addr* port *port* protocol { HTTP | gRPC } encoding { JSON | GPB | GPB-compact }**
8. **subscription *sub-id***
9. **snsr-group *sgrp_id* sample-interval *interval***
10. **dst-group *dgrp_id***

DETAILED STEPS**Procedure**

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | configure terminal Example: <pre>switch# configure terminal switch(config)#</pre> | Enter configuration mode. |
| Step 2 | telemetry Example: <pre>switch(config)# telemetry switch(config-telemetry)#</pre> | Enter configuration mode for the telemetry features. |
| Step 3 | sensor-group <i>sgrp_id</i> Example: <pre>switch(conf-tm-sub)# sensor-grp 6 switch(conf-tm-sub)#</pre> | Create a sensor group. |
| Step 4 | data source <i>data-source-type</i> Example: <pre>switch(config-tm-sensor)# data source YANG</pre> | Set the data source to YANG, so that it uses the native YANG streaming model to stream syslogs |
| Step 5 | path Cisco-NX-OS-Syslog-oper:syslog/messages Example: <pre>switch(config-tm-sensor)# path Cisco-NX-OS-Syslog-oper:syslog/messages</pre> | Configure the syslog path which streams syslog generated on the switch. |

| | Command or Action | Purpose |
|----------------|---|--|
| Step 6 | destination-group <i>grp_id</i> Example: <pre>switch(config-tm-sensor)# destination-group 33</pre> | Enter telemetry destination group sub-mode and configure the destination group. |
| Step 7 | ip address <i>ip_addr</i> port <i>port</i> protocol { HTTP gRPC } encoding { JSON GPB GPB-compact } Example: <pre>switch(config-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json</pre> Example: <pre>switch(config-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb</pre> | Configure the telemetry data for the subscription to stream to the specified IP address and port, and set the protocol and encoding for the data stream. |
| Step 8 | subscription <i>sub-id</i> Example: <pre>switch(config-tm-dest)# subscription 33</pre> | Enter telemetry subscription submode and configure the telemetry subscription. |
| Step 9 | snsr-group <i>sgrp_id</i> sample-interval <i>interval</i> Example: <pre>switch(config-tm-sub)# snsr-group 6 sample-interval 0</pre> | Link the sensor group to the current subscription and set the data sampling to 0 so that the switch sends telemetry data when syslog events occur. For interval, 0 is the only acceptable value. |
| Step 10 | dst-group <i>dgrp_id</i> Example: <pre>switch(config-tm-sub)# dst-grp 33</pre> | Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the destination-group command. |

Telemetry Data Streamed for Syslog Path

For each source path, the following table shows the information that is streamed when the subscription is first created "the baseline" and when event notifications occur.

| Path | Subscription Baseline | Event Notificat |
|------|-----------------------|-----------------|
|------|-----------------------|-----------------|

| | | |
|---|--|---|
| Cisco-NX-OS-Syslog-oper:syslog/messages | Stream all the existing syslogs from the switch. | Sends event switch: <ul style="list-style-type: none"> • message • node-name • time-stamp • time-of-day • time-zone • category • message-id • severity • text |
|---|--|---|

Displaying Syslog Path Information

Use the Cisco NX-OS **show telemetry event collector** commands to display statistics and counters, or errors for the syslog path.

Displaying Statistics

You can enter the **show telemetry event collector stats** command to display the statistics and counters for each syslog path.

The following is an example of statistics for the syslog path:

```

-----
Row ID           Collection Count  Latest Collection Time      Sensor Path(GroupID)
-----
1                138              Tue Dec 03 11:20:08.200 PST Cisco-NX-OS-Syslog-oper:syslog(1)
2                138              Tue Dec 03 11:20:08.200 PST Cisco-NX-OS-Syslog-oper:syslog/messages(1)

```

Displaying Error Counters

You can use the **show telemetry event collector errors** command to display the error totals for all the syslog paths.

```
switch(config-if)# show telemetry event collector errors
```

```

-----
Error Description                               Error Count
-----
Dme Event Subscription Init Failures            - 0
Event Data Enqueue Failures                     - 0
Event Subscription Failures                     - 0
Pending Subscription List Create Failures        - 0
Subscription Hash Table Create Failures         - 0
Subscription Hash Table Destroy Failures        - 0
Subscription Hash Table Insert Failures         - 0
Subscription Hash Table Remove Failures         - 0

```

Sample JSON Output

The following is a sample of JSON output:

```
172.19.216.13 - - [03/Dec/2019 19:38:50] "POST
/network/Cisco-NX-OS-Syslog-oper%3Asyslog%2Fmessages HTTP/1.0" 200 -
172.19.216.13 - - [03/Dec/2019 19:38:50] "POST
/network/Cisco-NX-OS-Syslog-oper%3Asyslog%2Fmessages HTTP/1.0" 200 -
>>> URL           : /network/Cisco-NX-OS-Syslog-oper%3Asyslog%2Fmessages
>>> TM-HTTP-VER    : 1.0.0
>>> TM-HTTP-CNT    : 1
>>> Content-Type   : application/json
>>> Content-Length : 578
      Path => Cisco-NX-OS-Syslog-oper:syslog/messages
              node_id_str   : task-n9k-1
              collection_id : 40
              data_source    : YANG
              data           :

[
  [
    {
      "message-id": 420
    },
    {
      "category": "ETHPORT",
      "group": "ETHPORT",
      "message-name": "IF_UP",
      "node-name": "task-n9k-1",
      "severity": 5,
      "text": "Interface loopback10 is up ",
      "time-of-day": "Dec 3 2019 11:38:51",
      "time-stamp": "1575401931000",
      "time-zone": ""
    }
  ]
]
```

Sample KVGPB Output

The following is a sample KVGPB output

```
---Telemetry msg received @ 18:22:04 UTC
```

```
All the fragments:1 read successfully total size read:339
```

```
node_id_str: "task-n9k-1"
subscription_id_str: "1"
collection_id: 374
data_gpbkv {
  fields {
    name: "keys"
    fields {
      name: "message-id"
      uint32_value: 374
    }
  }

  fields {
    name: "content"
    fields {
      fields {
        name: "node-name"
        string_value: "task-n9k-1"
      }
    }
  }
}
```

```
    }

    fields {
      name: "time-of-day"
      string_value: "Jun 26 2019 18:20:21"
    }

    fields {
      name: "time-stamp"
      uint64_value: 1574293838000
    }

    fields {
      name: "time-zone"
      string_value: "UTC"
    }

    fields {
      name: "process-name"
      string_value: ""
    }

    fields {
      name: "category"
      string_value: "VSHD"
    }

    fields {
      name: "group"
      string_value: "VSHD"
    }

    fields {
      name: "message-name"
      string_value: "VSHD_SYSLOG_CONFIG_I"
    }

    fields {
      name: "severity"
      uint32_value: 5
    }

    fields {
      name: "text"
      string_value: "Configured from vty by admin on console0"
    }
  }
}
}
```

Troubleshooting Telemetry

.

Displaying Telemetry Log and Trace Information

Use the following NX-OS CLI commands to display the log and trace information.

show tech-support telemetry

This NX-OS CLI command collects the telemetry log contents from the tech-support log. In this example, the command output is redirected into a file in bootflash.

```
switch# show tech-support telemetry > bootflash:tmst.log
```

Additional References

Related Documents

| Related Topic | Document Title |
|--|---|
| Example configurations of telemetry deployment for VXLAN EVPN. | Telemetry Deployment for VXLAN EVPN |



CHAPTER 34

Diagnosis and Serviceability

This chapter contains the following topics:

- [About Diagnosis and Serviceability, on page 515](#)
- [Show Commands, on page 515](#)
- [Debug Logs, on page 516](#)
- [Diagnosis Suggestions, on page 524](#)

About Diagnosis and Serviceability

Cisco NX-OS supports Model-Driven Programmability (MDP) through a range of different protocol interfaces, such as Netconf, Restconf, gNMI/gNOI, and Telemetry. In fact, these interfaces operate around the common underlying YANG and DME/CLI infrastructure. The user can diagnose the behavior through a common collection of utilities.

Show Commands

This section lists the commonly used show commands that you can use to verify the running state of the switch.

Table 44: Show Commands - Diagnosis and Serviceability

| Item | Command | Usage |
|---------|--|-----------------------------------|
| netconf | show running-config netconf | Display netconf config |
| | show netconf nxsdk event-history {events errors} | Display event history |
| | show tech-support netconf | Collect netconf tech-support |
| | show netconf internal details | Verify internal state |
| | show netconf internal tls service | Verify TLS server state |
| | show netconf internal tls session [all] { summary detail } | List current/history TLS sessions |

| Item | Command | Usage |
|------------|--|--|
| restconf | show running-config grep restconf | Display restconf config |
| | show netconf nxsdk event-history {events errors} | Display event history |
| grpc | show running-config grpc | Display grpc config |
| | show grpc nxsdk event-history {events errors} | Display event history |
| | show tech-support grpc | Collect grpc tech-support |
| gnmi | show grpc gnmi service statistics | Verify grpc server state |
| | show grpc gnmi rpc [all] {summary detail } | List current/history gNMI subscription |
| | show grpc gnmi transactions | List gNMI Get/Set |
| | show grpc internal gnmi subscription... | Display internal subscription data |
| | show grpc internal gnmi mtx {sessions statistics subscriptions} | Display internal infra logs |
| gnoi | show grpc gnoi service statistics | Verify grpc server state |
| | show grpc internal gnoi rpc [all] {summary detail} | List current/history gNOI connections |
| openconfig | show running-config openconfig | Display openconfig config |
| | show openconfig nxsdk event-history {events errors} | Display event history |
| dme | show system internal dme transaction history | Verify the DME transaction |
| | show tech-support dme | Collect DME tech-support |

Debug Logs

This section describes how to enable and collect the debug logs.

Programmability Agent Logs

For Netconf, Restconf, and gRPC agents, you can collect the logs in the following ways:

- **Show commands**

This is a straight-forward way to view/check the agent event. These commands are useful to see how the agents interact with the client connections. This log is in-memory log, and thus it could only keep a relatively short history.

```
show netconf nxsdk event-history {events | errors}
show restconf nxsdk event-history {events | errors}
show grpc nxsdk event-history {events | errors}
```

- **Log files**

If you prefer to check the longer history, or even the logs after disabling the agents, then see the log files stored under the **/volatile** directory. The user needs the permission to access the switch bash shell.

```
/volatile/netconf-internal-log
      grpc-internal-log
      restconf-internal-log
```

YANG Infra Logs

The YANG infra logs are saved in the **/volatile** directory. The user needs the permission to access the switch bash shell. In Cisco NX-OS, Bash is accessible from user accounts that are associated with the Cisco NX-OS dev-ops role or the Cisco NX-OS network-admin role.

```
/volatile/mtx-internal.netconf.log
      mtx-internal.grpc.log
      mtx-internal.restconf.log
```

DME Logs

The DME infra logs are saved in the **/nxos/dme_logs** directory. The user needs the permission to access the switch bash shell. See <https://developer.cisco.com/site/cisco-nexus-nx-api-references/>.

```
/nxos/dme_logs/svc_ifc_policyelem.<pid>.log
```

Change the Log Configuration

Cisco NX-OS enables very limited logs by default due to the performance consideration.

The user can change the verbosity by editing **/opt/mtx/conf/mtxlogger.cfg**.

The configuration file has the following structure:

```
<config name="nxos-device-mgmt">
  <container name="mgmtConf">
    <container name="logging">
      <leaf name="enabled" type="boolean" default="false"></leaf>
      <leaf name="allActive" type="boolean" default="false"></leaf>
      <container name="format">
        <leaf name="content" type="string" default="$DATETIME$ $COMPONENTID$ $TYPES$ :
$MSG$"></leaf>
        <container name="componentID">
          <leaf name="enabled" type="boolean" default="true"></leaf>
        </container>
        <container name="dateTime">
          <leaf name="enabled" type="boolean" default="true"></leaf>
          <leaf name="format" type="string" default="%y%m%d.%H%M%S"></leaf>
        </container>
      </container>
    </container>
  </container>
</config>
```

```

    <container name="fcn">
      <leaf name="enabled" type="boolean" default="true"></leaf>
      <leaf name="format" type="string"
default="`${CLASS}$:${FCNNAME}$($ARCS$)@$LINE$"></leaf>
    </container>
  </container>
  <container name="dest">
    <container name="console">
      <leaf name="enabled" type="boolean" default="false"></leaf>
    </container>
    <container name="file">
      <leaf name="enabled" type="boolean" default="false"></leaf>
      <leaf name="name" type="string" default="mtx-internal.log"></leaf>
      <leaf name="location" type="string" default=".mtxlogs"></leaf>
      <leaf name="mbytes-rollover" type="uint32" default="10"></leaf>
      <leaf name="hours-rollover" type="uint32" default="24"></leaf>
      <leaf name="startup-rollover" type="boolean" default="false"></leaf>
      <leaf name="max-rollover-files" type="uint32" default="10"></leaf>
    </container>
  </container>
  <list name="logitems" key="id">
    <listitem>
      <leaf name="id" type="string"></leaf>
      <leaf name="active" type="boolean" default="true"></leaf>
    </listitem>
  </list>
</container>
</container>
</config>

```

The `<list>` tag defines the log filters by `<componentID>`.

The following table describes some of the containers and their leaves.

Table 45: Containers and Leaves

| Container | Container Description | Contained Containers | Contained Leaf Description |
|-----------|----------------------------------|--|--|
| logging | Contains all logging data types. | <ul style="list-style-type: none"> format dest file <p>Note Also contains list tag logitems</p> | <p>enabled: Boolean that determines whether logging is on or off. Default off.</p> <p>allActive: Boolean that activates all defined logging items for logging. Default off</p> |

| Container | Container Description | Contained Containers | Contained Leaf Description |
|-------------|--|--|---|
| format | Contains the log message format information. | <ul style="list-style-type: none"> • componentID • dateTime • type • fcn | <p>content: String listing data types included in log messages. Includes:</p> <ul style="list-style-type: none"> • \$DATETIMES\$: Include date or time in log message. • \$COMPONENTID\$: Include component name in log message. • \$TYPES\$: Includes message type ("", INFO, WARNING, ERROR) • \$SRCFILES\$: Includes name of source file. • \$SRCLINES\$: Include line number of source file. • \$FCNINFOS\$ Include class::function name from the source file. <p>\$MSG\$: Include actual log message text.</p> |
| componentID | Name of logged component. | NA | <p>enabled: Boolean that determines if the log message includes the component ID. Default to "true." Value of "false" returns a "" string in log message.</p> |
| dateTime | Date or time of log message. | NA | <p>enabled: Boolean whether to include date or time information in log message. Default is enabled.</p> <p>format: String of values to include in log message. Format of %y%m%d.%H%M%S.</p> |

| Container | Container Description | Contained Containers | Contained Leaf Description |
|-----------|--|--|--|
| dest | Holds destination logger's configuration settings. | console: Destination console. Only one allowed. file: destination file. Multiple allowed. | NA |
| console | Destination console. | NA | enabled: Boolean that determines whether the console is enabled for logging. Default of "false." |

| Container | Container Description | Contained Containers | Contained Leaf Description |
|-----------|--|----------------------|---|
| file | Determines the settings of the destination file. | NA | <p>enabled: Boolean that determines whether the destination is enabled. Default is "false."</p> <p>name: String of the destination log file. Default of "mtx-internal.log"</p> <p>location: String of destination file path. Default at "./mtxlogs."</p> <p>mbytes-rollover: uint32 that determines the length of the log file before the system overwrites the oldest data. Default is 10 Mbytes.</p> <p>hours-rollover: uint32 that determines the length of the log file in terms of hours. Default is 24 hours.</p> <p>startup-rollover: Boolean that determines if the log file is rolled over upon agent start or restart. Default value of "false."</p> <p>max-rollover-files: uint32 that determines the maximum number of rollover files; deletes the oldest file when the max-rollover-files value exceeded. Default value of 10.</p> |

Default Config Example

The following is the configuration file with the default installed configuration.

```
<config name="nxos-device-mgmt">
  <container name="mgmtConf">
    <container name="logging">
      <leaf name="enabled" type="boolean" default="false">true</leaf>
      <leaf name="allActive" type="boolean" default="false">>false</leaf>
    <container name="format">
      <leaf name="content" type="string" default="$DATETIME$ $COMPONENTID$ $TYPE$:
```

```

$MSG$">$DATETIME$ $COMPONENTID$ $TYPE$ $SRCFILE$ @ $SRCLINE$ $FCNINFO$: $MSG$</leaf>
  <container name="componentID">
    <leaf name="enabled" type="boolean" default="true"></leaf>
  </container>
  <container name="dateTime">
    <leaf name="enabled" type="boolean" default="true"></leaf>
    <leaf name="format" type="string" default="%y%m%d.%H%M%S"></leaf>
  </container>
  <container name="fcn">
    <leaf name="enabled" type="boolean" default="true"></leaf>
    <leaf name="format" type="string"
default="$CLASS$: :$FCNNAME$($ARGS$)@$LINE$"></leaf>
  </container>
</container>
<container name="dest">
  <container name="console">
    <leaf name="enabled" type="boolean" default="false">true</leaf>
  </container>
  <container name="file">
    <leaf name="enabled" type="boolean" default="false">true</leaf>
    <leaf name="name" type="string" default="mtx-internal.log"></leaf>
    <leaf name="location" type="string" default="/volatile"></leaf>
    <leaf name="mbytes-rollover" type="uint32" default="10">50</leaf>
    <leaf name="hours-rollover" type="uint32" default="24">24</leaf>
    <leaf name="startup-rollover" type="boolean" default="false">true</leaf>
    <leaf name="max-rollover-files" type="uint32" default="10">10</leaf>
  </container>
</container>
<list name="logitems" key="id">
  <listitem>
    <leaf name="id" type="string">*</leaf>
    <leaf name="active" type="boolean" default="false">>false</leaf>
  </listitem>
  <listitem>
    <leaf name="id" type="string">SYSTEM</leaf>
    <leaf name="active" type="boolean" default="true">true</leaf>
  </listitem>
  <listitem>
    <leaf name="id" type="string">LIBUTILS</leaf>
    <leaf name="active" type="boolean" default="true">true</leaf>
  </listitem>
  <listitem>
    <leaf name="id" type="string">MTX-API</leaf>
    <leaf name="active" type="boolean" default="true">true</leaf>
  </listitem>
  <listitem>
    <leaf name="id" type="string">Model-*</leaf>
    <leaf name="active" type="boolean" default="true">true</leaf>
  </listitem>
  <listitem>
    <leaf name="id" type="string">Model-Cisco-NX-OS-device</leaf>
    <leaf name="active" type="boolean" default="true">>false</leaf>
  </listitem>
  <listitem>
    <leaf name="id" type="string">Model-openconfig-bgp</leaf>
    <leaf name="active" type="boolean" default="true">>false</leaf>
  </listitem>
  <listitem>
    <leaf name="id" type="string">INST-MTX-API</leaf>
    <leaf name="active" type="boolean" default="true">>false</leaf>
  </listitem>
  <listitem>
    <leaf name="id" type="string">INST-ADAPTER-NC</leaf>
    <leaf name="active" type="boolean" default="true">>false</leaf>
  </listitem>

```



```

</listitem>
<listitem>
  <leaf name="id" type="string">INST-ADAPTER-RC</leaf>
  <leaf name="active" type="boolean" default="true">>false</leaf>
</listitem>
<listitem>
  <leaf name="id" type="string">INST-ADAPTER-GRPC</leaf>
  <leaf name="active" type="boolean" default="true">>false</leaf>
</listitem>
</list>
</container>
</container>
</config>

```

Change the Log Configuration Using CLI

Since 10.4(2)F, CLIs are available to change the above logging configuration dynamically without restarting the process. These are per agent EXEC. There are not configuration, and thus can be changed without impacting the current operations.

SUMMARY STEPS

1. `[no] debug grpc mtx enable-all`
2. `[no] debug grpc mtx level <level>`
3. `[no] debug grpc mtx item <item>`

DETAILED STEPS

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | <code>[no] debug grpc mtx enable-all</code> | This is a convenient cli to enable all logs. |
| Step 2 | <code>[no] debug grpc mtx level <level></code> Example: switch# debug grpc mtx level info | Toggle the logging level: error, warning, info, debug. The default level is info . |
| Step 3 | <code>[no] debug grpc mtx item <item></code> Example: switch# debug grpc mtx item MTX-EvtMgr | Toggle the logging for specific item. This is a free form string, and use show grpc internal mtx debug to see the available items. |

Example

The below show cli would display the current logging configuration.

```
show grpc internal mtx debug
```

Example:

```

Log enabled : 1
All active  : 0
Log Level   : Debug

```

```

Log items :
* : 0
DtxUserFunc : 0
INST-ADAPTER : 0
INST-ADAPTER-GNMI : 0
INST-ADAPTER-GNOI : 0
INST-ADAPTER-GRPC : 1
INST-ADAPTER-NC : 1
INST-ADAPTER-RC : 1
INST-ADAPTER-TM : 0
INST-MTX-API : 1
LIBUTILS : 1
MTX-API : 1
MTX-ActionMgr : 0
MTX-Coder : 0
MTX-Dy-EvtMgr : 1
MTX-EvtMgr : 1
MTX-RbacMgr : 0
MTXEXPR : 0
MTXItem : 0
MTXNetConfMessage : 0
MTXOperation : 0
MTXRestConfMessage : 0
MTXgNMIMessage : 0
Model-* : 1
Model-Cisco-NX-OS-device : 1
Model-openconfig-bgp : 0
RPC : 0
SYSTEM : 1
TM-ADPT : 0
TM-ADPT-JSON : 0

```

Diagnosis Suggestions

This section provides a few steps to triage frequently seen issues.

Connection Issues

If the user's programming client cannot connect to the switch, then check the following:

- Check whether the feature is enabled by checking the running configuration.
- Check individual agent's show command to confirm that the server is running.
- Check the ip / port to confirm the connectivity is not restricted by firewall, etc.
- Check the client sends the correct user/password.
- If cert-based authentication is used, check that the trustpoint has been properly configured to the switch, and the client certification matches and has not expired.

Native Device Yang

If there is an issue with the native openconfig YANG related to read/write operations, then check the following:

- For "write" operations, check the DME transaction to see the failure details.
- Send equivalent DME REST request, to confirm whether it has the same issue.

OpenConfig Yang

If there is issue read/write the native openconfig YANG, then check the following:

- Check whether **feature openconfig** is enabled.
- Check the published YANG and deviation to confirm the support status.
- For **write** operations, check the DME transaction to see the failure details.

Telemetry

Telemetry is used to collect YANG and other data sources through the “feature telemetry” configuration. Telemetry is also used for gNMI subscribe via “feature grpc”. Troubleshooting steps are different depending on the usage scenario.

Debug Logs

Debug logs can be viewed through:

- **show telemetry internal event history { errors | events }**
- **show grpc nxsdk event-history { events | errors }**

Data / Event Collection Issues:

Check show command for failed or skipped collections.

- **show telemetry data collector detail**
- **show telemetry event collector {errors | stats}**
- **show grpc internal gnmi subscription statistics**

Collection time or size issues:

Check collection sizes and times via following show commands:

- **show telemetry control database**
- **show grpc internal gnmi rpc subscription-data**

Transport Issues:

Check for transport issues with following show command. Note that transport issues only impact **feature telemetry** scenario.

- **show telemetry transport <num> stats | errors**



PART VI

XML Management Interface

- [XML Management Interface, on page 529](#)



CHAPTER 35

XML Management Interface

- [About the XML Management Interface, on page 529](#)
- [Licensing Requirements for the XML Management Interface, on page 530](#)
- [Prerequisites to Using the XML Management Interface, on page 530](#)
- [Using the XML Management Interface, on page 531](#)
- [Information About Example XML Instances, on page 544](#)
- [Additional References, on page 551](#)

About the XML Management Interface

Information About the XML Management Interface

You can use the XML management interface to configure a device. The interface uses the XML-based Network Configuration Protocol (NETCONF), which allows you to manage devices and communicate over the interface with an XML management tool or program. The Cisco NX-OS implementation of NETCONF requires you to use a Secure Shell (SSH) session for communication with a device.

NETCONF is implemented with an XML Schema (XSD) that allows you to enclose device configuration elements within a remote procedure call (RPC) message. From within an RPC message, select one of the NETCONF operations that matches the type of command that you want the device to execute. You can configure the entire set of CLI commands on the device with NETCONF. For information about using NETCONF, see the [Creating NETCONF XML Instances, on page 534](#) and [RFC 4741](#).

For more information about using NETCONF over SSH, see [RFC 4742](#).

This section includes the following topics:

NETCONF Layers

The following table lists the NETCONF layers:

Table 46: NETCONF Layers

| Layer | Example |
|--------------------|----------------|
| Transport protocol | SSHv2 |
| RPC | RPC, RPC-reply |

| Layer | Example |
|------------|---|
| Operations | get-config, edit-config |
| Content | show or configuration command |

The following is a description of the four NETCONF layers:

- SSH transport protocol—Provides an encrypted connection between a client and the server.
- RPC tag—Introduces a configuration command from the requestor and the corresponding reply from the XML server.
- NETCONF operation tag—Indicates the type of configuration command.
- Content—Indicates the XML representation of the feature that you want to configure.

SSH xmlagent

The device software provides an SSH service that is called xmlagent that supports NETCONF over SSH Version 2.



Note The xmlagent service is referred to as the XML server in Cisco NX-OS software.

NETCONF over SSH starts with the exchange of a Hello message between the client and the XML server. After the initial exchange, the client sends XML requests, which the server responds to with XML responses. The client and server terminate requests and responses with the character sequence >. Because this character sequence is not valid in XML, the client and the server can interpret when messages end, which keeps communication in sync.

The XML schemas that define the XML configuration instances that you can use are described in [Creating NETCONF XML Instances](#), on page 534.

Licensing Requirements for the XML Management Interface

| Product | License Requirement |
|-------------|---|
| Cisco NX-OS | The XML management interface requires no license. Any feature that is not included in a license package is bundled with the Cisco NX-OS image and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> . |

Prerequisites to Using the XML Management Interface

Using the XML management interface has the following prerequisites:

- You must install SSHv2 on the client PC.
- You must install an XML management tool that supports NETCONF over SSH on the client PC.
- You must set the appropriate options for the XML server on the device.

Using the XML Management Interface

This section describes how to manually configure and use the XML management interface.



Note Use the XML management interface with the default settings on the device.

Configuring the SSH and the XML Server Options Through the CLI

By default, the SSH server is enabled on your device. If you disable SSH, you must enable it before you start an SSH session on the client PC.

You can configure the XML server options to control the number of concurrent sessions and the timeout for active sessions. You can also enable XML document validation and terminate XML sessions.



Note The XML server timeout applies only to active sessions.

For more information about configuring SSH, see the Cisco NX-OS security configuration guide for your platform.

For more information about the XML commands, see the Cisco NX-OS system management configuration guide for your platform.

Procedure

- Step 1** Enter global configuration mode.
- ```
configure terminal
```
- Step 2** (Optional) Display information about XML server settings and active XML server sessions. You can find session numbers in the command output.
- ```
show xml server status
```
- Step 3** Validate XML documents for the specified server session.
- ```
xml server validate all
```
- Step 4** Terminate the specified XML server session.
- ```
xml server terminate session
```

- Step 5** (Optional) Disable the SSH server so that you can generate keys.
no feature ssh
- Step 6** Enable the SSH server. (The default is enabled.)
feature ssh
- Step 7** (Optional) Display the status of the SSH server.
show ssh server
- Step 8** Set the number of XML server sessions allowed.
xml server max-session sessions
The range is from 1 to 8. The default is 8.
- Step 9** Set the number of seconds after which an XML server session is terminated.
xml server timeout seconds
The range is from 1 to 1200. The default is 1200 seconds.
- Step 10** (Optional) Display information about the XML server settings and active XML server sessions.
show xml server status
- Step 11** (Optional) Saves the running configuration to the startup configuration.
copy running-config startup-config

Example

The following example shows how to configure SSH and XML server options through the CLI:

```
switch# configure terminal
switch(config)# xml server validate all
switch(config)# xml server terminate 8665
switch(config)# no feature ssh
switch(config)# feature ssh server
switch(config)# xml server max-session 6
switch(config)# xml server timeout 24001200
switch(config)# copy running-config startup-config
```

Starting an SSHv2 Session

You can start an SSHv2 session on a client PC with the **ssh2** command that is similar to the following:

```
ssh2 username@ip-address -s xmlagent
```

Enter the login username, the IP address of the device, and the service to connect to. The `xmlagent` service is referred to as the XML server in the device software.



Note The SSH command syntax can differ based on the SSH software on the client PC.

If you do not receive a Hello message from the XML server, verify the following conditions:

- The SSH server is enabled on the device.
- The *max-sessions* option of the XML server is adequate to support the number of SSH connections to the device.
- The active XML server sessions on the device are not all in use.

Sending a Hello Message

You must advertise your capabilities to the server with a Hello message before the server processes any other requests. When you start an SSH session to the XML server, the server responds immediately with a Hello message. This message informs the client of the capabilities of the server. The XML server supports only base capabilities and, in turn, expects that the client supports only these base capabilities.

The following are sample Hello messages from the server and the client:



Note You must end all XML documents with `]]>]]>` to support synchronization in NETCONF over SSH.

Hello Message from a Server

```
<?xml version="1.0"?>
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <capabilities>
    <capability>urn:ietf:params:xml:ns:netconf:base:1.0</capability>
  </capabilities>
  <session-id>25241</session-id>
</hello>]]>]]>
```

Hello Message from a Client

```
<?xml version="1.0"?>
<nc:hello xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
  <nc:capabilities>
    <nc:capability>urn:ietf:params:xml:ns:netconf:base:1.0</nc:capability>
  </nc:capabilities>
</nc:hello>]]>]]>
```

Obtaining XML Schema Definition (XSD) Files

Procedure

-
- Step 1** switch# feature bash shell
Step 2 switch# run bash
Step 3 bash-3.2\$ cd /isan/etc/schema

Step 4 Obtain the necessary schema.

Sending an XML Document to the XML Server

To send an XML document to the XML server through an SSH session that you opened in a command shell, copy the XML text from an editor and paste it into the SSH session. Although typically you use an automated method to send XML documents to the XML server, you can verify the SSH connection to the XML server through this copy-paste method.

The following are the guidelines to follow when sending an XML document to the XML server:

- Verify that the XML server has sent the Hello message immediately after you started the SSH session, by looking for the Hello message text in the command shell output.
- Send the client Hello message before you send XML requests. Note that the XML server sends the Hello response immediately, and no additional response is sent after you send the client Hello message.
- Always terminate the XML document with the character sequence `]]>]]>`.

Creating NETCONF XML Instances

You can create NETCONF XML instances by enclosing the XML device elements within an RPC tag and NETCONF operation tags. The XML device elements are defined in feature-based XML schema definition (XSD) files, which enclose available CLI commands in an XML format.

The following are the tags that are used in the NETCONF XML request in a framework context. Tag lines are marked with the following letter codes:

- X—XML declaration
- R—RPC request tag
- N—NETCONF operation tags
- D—Device tags

NETCONF XML Framework Context

```
X <?xml version="1.0"?>
R <nc:rpc message-id="1" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
R xmlns="http://www.cisco.com/nxos:1.0:nfcli">
N <nc:get>
N <nc:filter type="subtree">
D <show>
D <xml>
D <server>
D <status/>
D </server>
D </xml>
D </show>
N </nc:filter>
N </nc:get>
R </nc:rpc>]]>]]>
```



Note You must use your own XML editor or XML management interface tool to create XML instances.

RPC Request Tag

All NETCONF XML instances must begin with the RPC request tag `<rpc>`. The `<rpc>` element has a message ID (message-id) attribute. This message-id attribute is replicated in the `<rpc-reply>` and can be used to correlate requests and replies. The `<rpc>` node also contains the following XML namespace declarations:

- NETCONF namespace declaration—The `<rpc>` and NETCONF tags that are defined in the `urn:ietf:params:xml:ns:netconf:base:1.0` namespace, are present in the `netconf.xsd` schema file.
- Device namespace declaration—Device tags encapsulated by the `<rpc>` and NETCONF tags are defined in other namespaces. Device namespaces are feature-oriented. Cisco NX-OS feature tags are defined in different namespaces. RPC Request Tag `<rpc>` is an example that uses the NFCLI feature. It declares that the device namespace is `xmlns=http://www.cisco.com/nxos:1.0:nfcli`. `nfcli.xsd` contains this namespace definition. For more information, see [Obtaining XML Schema Definition \(XSD\) Files](#), on page 533.

Examples

RPC Request Tag `<rpc>`

```
<nc:rpc message-id="315" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns=http://www.cisco.com/nxos:1.0:nfcli">
...
</nc:rpc>]]>]]>
```

Configuration Request

```
<?xml version="1.0"?>
<nc:rpc message-id="16" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0:if_manager">
  <nc:edit-config>
    <nc:target>
      <nc:running/>
    </nc:target>
    <nc:config>
      <configure>
        <__XML_MODE__exec_configure>
          <interface>
            <ethernet>
              <interface>2/30</interface>
              <__XML_MODE__if-ethernet>
                <__XML_MODE__if-eth-base>
                  <description>
                    <desc_line>Marketing Network</desc_line>
                  </description>
                </__XML_MODE__if-eth-base>
              </__XML_MODE__if-ethernet>
            </ethernet>
          </interface>
        </__XML_MODE__exec_configure>
```

```

    </configure>
  </nc:config>
</nc:edit-config>
</nc:rpc>]]>]]>

```



Note `__XML__MODE` tags are used internally by the NETCONF agent. Some tags are present only as children of a certain `__XML__MODE`. By examining the schema file, you can find the correct mode tag that leads to the tags representing the CLI command in XML.

NETCONF Operations Tags

NETCONF provides the following configuration operations:

Table 47: NETCONF Operations in Cisco NX-OS

| NETCONF Operation | Description | Example |
|-------------------|---|--|
| close-session | Closes the current XML server session. | NETCONF Close Session Instance, on page 545 |
| commit | Sets the running configuration to the current contents of candidate configuration. | NETCONF Commit Instance: Candidate Configuration Capability, on page 549 |
| confirmed-commit | Provides the parameters to commit the configuration for a specified time. If a commit operation does not follow this operation within the confirm-timeout period, the configuration is reverted to the state before the confirmed-commit operation. | NETCONF Confirmed Commit Instance, on page 550 |
| copy-config | Copies the contents of the source configuration datastore to the target datastore. | NETCONF Copy Config Instance, on page 545 |
| delete-config | Operation not supported. | — |
| edit-config | Configures the features in the running configuration of the device. You use this operation for configuration commands. | NETCONF Edit Config Instance, on page 546 NETCONF Rollback-On-Error Instance, on page 550 |
| get | Receives configuration information from a device. You use this operation for show commands. The source of the data is the running configuration. | Creating NETCONF XML Instances, on page 534 |

| NETCONF Operation | Description | Example |
|-------------------|--|---|
| get-config | Retrieves all or part of a configuration. | Creating NETCONF XML Instances, on page 534 |
| kill-session | Closes the specified XML server session. You cannot close your own session. | NETCONF Kill Session Instance, on page 545 |
| lock | Allows a client to lock the configuration system of a device. | NETCONF Lock Instance, on page 548 |
| unlock | Releases the configuration lock that the session issued. | NETCONF Unlock Instance, on page 549 |
| validate | Checks the configuration of a candidate for syntactical and semantic errors before applying the configuration to a device. | NETCONF Validate Capability Instance, on page 551 |

Device Tags

The XML device elements represent the available CLI commands in XML format. The feature-specific schema files contain the XML tags for CLI commands of that particular feature. See [Obtaining XML Schema Definition \(XSD\) Files, on page 533](#).

Using this schema, it is possible to build an XML instance. The relevant portions of the nfcli.xsd schema file that was used to build the NETCONF instances. See [\(Creating NETCONF XML Instances, on page 534\)](#).

show xml Device Tags

```
<xs:element name="show" type="show_type_Cmd_show_xml"/>
<xs:complexType name="show_type_Cmd_show_xml">
  <xs:annotation>
    <xs:documentation>to display xml agent information</xs:documentation>
  </xs:annotation>
  <xs:sequence>
    <xs:choice maxOccurs="1">
      <xs:element name="xml" minOccurs="1" type="xml_type_Cmd_show_xml"/>
      <xs:element name="debug" minOccurs="1" type="debug_type_Cmd_show_debug"/>
    </xs:choice>
  </xs:sequence>
  <xs:attribute name="xpath-filter" type="xs:string"/>
  <xs:attribute name="uses-namespace" type="nxos:bool_true"/>
</xs:complexType>
```

Server Status Device Tags

```
<xs:complexType name="xml_type_Cmd_show_xml">
  <xs:annotation>
    <xs:documentation>xml agent</xs:documentation>
  </xs:annotation>
  <xs:sequence>
    <xs:element name="server" minOccurs="1" type="server_type_Cmd_show_xml"/>
  </xs:sequence>
</xs:complexType>
<xs:complexType name="server_type_Cmd_show_xml">
  <xs:annotation>
```

```

<xs:documentation>xml agent server</xs:documentation>
</xs:annotation>
<xs:sequence>
<xs:choice maxOccurs="1">
<xs:element name="status" minOccurs="1" type="status_type_Cmd_show_xml"/>
<xs:element name="logging" minOccurs="1" type="logging_type_Cmd_show_logging_facility"/>
</xs:choice>
</xs:sequence>
</xs:complexType>

```

Device Tag Response

```

<xs:complexType name="status_type_Cmd_show_xml">
<xs:annotation>
<xs:documentation>display xml agent information</xs:documentation>
</xs:annotation>
<xs:sequence>
<xs:element name="__XML__OPT_Cmd_show_xml__readonly__" minOccurs="0">
<xs:complexType>
<xs:sequence>
<xs:group ref="og_Cmd_show_xml__readonly__" minOccurs="0" maxOccurs="1"/>
</xs:sequence>
</xs:complexType>
</xs:element>
</xs:sequence>
</xs:complexType>
<xs:group name="og_Cmd_show_xml__readonly__">
<xs:sequence>
<xs:element name="__readonly__" minOccurs="1" type="__readonly__type_Cmd_show_xml"/>
</xs:sequence>
</xs:group>
<xs:complexType name="__readonly__type_Cmd_show_xml">
<xs:sequence>
<xs:group ref="bg_Cmd_show_xml_operational_status" maxOccurs="1"/>
<xs:group ref="bg_Cmd_show_xml_maximum_sessions_configured" maxOccurs="1"/>
<xs:group ref="og_Cmd_show_xml_TABLE_sessions" minOccurs="0" maxOccurs="1"/>
</xs:sequence>
</xs:complexType>

```



Note The `__XML__OPT_Cmd_show_xml__readonly__` tag is optional. This tag represents the response. For more information on responses, see [RPC Response Tag, on page 543](#).

You can use the | XML option to find the tags that you can use to execute a <get> operation. The following is an example of the | XML option. This example shows you that the namespace-defining tag to execute operations on this device is `http://www.cisco.com/nxos:1.0:nfcli`, and that the `nfcli.xsd` file can be used to build requests.

You can enclose the NETCONF operation tags and the device tags within the RPC tag. The </rpc> end tag is followed by the XML termination character sequence.

XML Example

```

Switch#> show xml server status | xml
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0:nfcli">
<nf:data>
<show>
<xml>

```



```

<server>
<status>
<__XML_OPT_Cmd_show_xml__readonly__>
<__readonly__>
<operational_status>
<o_status>enabled</o_status>
</operational_status>
<maximum_sessions_configured>
<max_session>8</max_session>
</maximum_sessions_configured>
</__readonly__>
</__XML_OPT_Cmd_show_xml__readonly__>
</status>
</server>
</xml>
</show>
</nf:data>
</nf:rpc-reply>
]]>]]>

```

Extended NETCONF Operations

Cisco NX-OS supports an <rpc> operation named <exec-command>. The operation allows client applications to send CLI **configuration** and **show** commands and to receive responses to those commands as XML tags.

The following is an example of the tags that are used to configure an interface. Tag lines are marked with the following letter codes:

- X—XML declaration
- R—RPC request tag
- EO—Extended operation

The following table provides a detailed explanation of the operation tags:

Table 48: Operation Tags

| Tag | Description |
|----------------|---|
| <exec-command> | Executes a CLI command. |
| <cmd> | Contains the CLI command. A command can be a show command or configuration command. Separate multiple configuration commands by using a semicolon (;). Although multiple show commands are not supported, you can send multiple configuration commands in different <cmd> tags as part of the same request. For more information, see the Example on <i>Configuration CLI Commands Sent Through <exec-command></i> . |

Replies to configuration commands that are sent through the <cmd> tag are as follows:

- <nf:ok>—All **configuration** commands are executed successfully.

- `<nf:rpc-error>`—Some commands have failed. The operation stops at the first error, and the `<nf:rpc-error>` subtree provides more information about which configuration has failed. Configurations that are executed before the failed command would have been applied to the running configuration.

Configuration CLI Commands Sent Through the `<exec-command>`

The `show` command must be sent in its own `<exec-command>` instance as shown in the following example:

```
X <?xml version="1.0"?>
R <nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="3">
EO <nxos:exec-command>
EO <nxos:cmd>conf t ; interface ethernet 2/1 </nxos:cmd>
EO <nxos:cmd>channel-group 2000 ; no shut; </nxos:cmd>
EO </nxos:exec-command>
R </nf:rpc>]]]]>
```

Response to CLI Commands Sent Through the `<exec-command>`

The following is the response to a send operation:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="3">
<nf:ok/>
</nf:rpc-reply>
]]]]>
```

Show CLI Commands Sent Through the `<exec-command>`

The following example shows how the `show` CLI commands that are sent through the `<exec-command>` can be used to retrieve data:

```
<?xml version="1.0"?>
<nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="110">
<nxos:exec-command>
<nxos:cmd>show interface brief</nxos:cmd>
</nxos:exec-command>
</nf:rpc>]]]]>
```

Response to the show CLI Commands Sent Through the `<exec-command>`

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:nxos="http://www.cisco.com/nxos:1.0"
  xmlns:mod="http://www.cisco.com/nxos:1.0:if_manager" message-id="110">
<nf:data>
<mod:show>
<mod:interface>
<mod:__XML__OPT_Cmd_show_interface_brief__readonly__>
```

```

<mod:__readonly__>
<mod:TABLE_interface>
<mod:ROW_interface>
<mod:interface>mgmt0</mod:interface>
<mod:state>up</mod:state>
<mod:ip_addr>192.0.2.20</mod:ip_addr>
<mod:speed>1000</mod:speed>
<mod:mtu>1500</mod:mtu>
</mod:ROW_interface>
<mod:ROW_interface>
<mod:interface>Ethernet2/1</mod:interface>
<mod:vlan>--</mod:vlan>
<mod:type>eth</mod:type>
<mod:portmode>routed</mod:portmode>
<mod:state>down</mod:state>
<mod:state_rsn_desc>Administratively down</mod:state_rsn_desc>
<mod:speed>auto</mod:speed>
<mod:ratemode>D</mod:ratemode>
</mod:ROW_interface>
</mod:TABLE_interface>
</mod:__readonly__>
</mod:__XML_OPT_Cmd_show_interface_brief__readonly__>
</mod:interface>
</mod:show>
</nf:data>
</nf:rpc-reply>
]]>]]>

```

Failed Configuration

```

<?xml version="1.0"?>
<nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="3">
<nxos:exec-command>
<nxos:cmd>configure terminal ; interface ethernet2/1 </nxos:cmd>
<nxos:cmd>ip address 192.0.2.2/24 </nxos:cmd>
<nxos:cmd>no channel-group 2000 ; no shut; </nxos:cmd>
</nxos:exec-command>
</nf:rpc>]]>]]>
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="3">
<nf:rpc-error>
<nf:error-type>application</nf:error-type>
<nf:error-tag>invalid-value</nf:error-tag>
<nf:error-severity>error</nf:error-severity>
<nf:error-message>Ethernet2/1: not part of port-channel 2000
</nf:error-message>
<nf:error-info>
<nf:bad-element>cmd</nf:bad-element>
</nf:error-info>
</nf:rpc-error>
</nf:rpc-reply>
]]>]]>

```

After a command is executed, the interface IP address is set, but the administrative state is not modified (the **no shut** command is not executed. The administrative state is not modified because the **no port-channel 2000** command results in an error.

The <rpc-reply> is due to a **show** command that is sent through the <cmd> tag that contains the XML output of the **show** command.

You cannot combine configuration and show commands on the same `<exec-command>` instance. The following example shows **config** and **show** commands that are combined in the same instance.

Combination of configure and show Commands

```
<?xml version="1.0"?>
<nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="110">
<nxos:exec-command>
<nxos:cmd>conf t ; interface ethernet 2/1 ; ip address 1.1.1.4/24 ; show xml
server status </nxos:cmd>
</nxos:exec-command>
</nf:rpc>]]>]]>
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="110">
<nf:rpc-error>
<nf:error-type>application</nf:error-type>
<nf:error-tag>invalid-value</nf:error-tag>
<nf:error-severity>error</nf:error-severity>
<nf:error-message>Error: cannot mix config and show in exec-command. Config cmds
before the show were executed.
Cmd:show xml server status</nf:error-message>
<nf:error-info>
<nf:bad-element>cmd</nf:bad-element>
</nf:error-info>
</nf:rpc-error>
</nf:rpc-reply>
]]>]]>
```

show CLI Commands Sent Through the `<exec-command>`

```
<?xml version="1.0"?>
<nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="110">
<nxos:exec-command>
<nxos:cmd>show xml server status ; show xml server status </nxos:cmd>
</nxos:exec-command>
</nf:rpc>]]>]]>
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:nxos="http://www.cisco.com/nxos:1.0" message-id="110">
<nf:rpc-error>
<nf:error-type>application</nf:error-type>
<nf:error-tag>invalid-value</nf:error-tag>
<nf:error-severity>error</nf:error-severity>
<nf:error-message>Error: show cmds in exec-command shouldn't be followed by anything
</nf:error-message>
<nf:error-info>
<nf:bad-element>&lt;cmd&gt;</nf:bad-element>
</nf:error-info>
</nf:rpc-error>
</nf:rpc-reply>
]]>]]>
```

NETCONF Replies

For every XML request sent by a client, the XML server sends an XML response that is enclosed in the RPC response tag `<rpc-reply>`.

RPC Response Tag

The following example shows the RPC response tag `<rpc-reply>`:

RPC Response Tag `<rpc-reply>`

```
<nc:rpc-reply message-id="315" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns=http://www.cisco.com/nxos:1.0:nfcli">
<ok/>
</nc:rpc-reply>]]>]]>
```

RPC Response Elements

The elements `<ok>`, `<data>`, and `<rpc-error>` can appear in the RPC response. The following table describes the RPC response elements that can appear in the `<rpc-reply>` tag:

Table 49: RPC Response Elements

| Element | Description |
|--------------------------------|---|
| <code><ok></code> | The RPC request completed successfully. This element is used when no data is returned in the response. |
| <code><data></code> | The RPC request completed successfully. The data that are associated with the RPC request is enclosed in the <code><data></code> element. |
| <code><rpc-error></code> | The RPC request failed. Error information is enclosed in the <code><rpc-error></code> element. |

Interpreting the Tags Encapsulated in the data Tag

The device tags encapsulated in the `<data>` tag contain the request, followed by the response. A client application can safely ignore all the tags before the `<readonly>` tag, as show in the following example:

RPC Reply Data

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<nf:rpc-reply xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0:if_manager">
<nf:data>
<show>
<interface>
<_XML_OPT_Cmd_show_interface_brief__readonly_>
<_readonly_>
<TABLE_interface>
<ROW_interface>
<interface>mgmt0</interface>
```

```

<state>up</state>
<ip_addr>xx.xx.xx.xx</ip_addr>
<speed>1000</speed>
<mtu>1500</mtu>
</ROW_interface>
<ROW_interface>
<interface>Ethernet2/1</interface>
<vlan>--</vlan>
<type>eth</type>
<portmode>routed</portmode>
<state>down</state>
<state_rsn_desc>Administratively down</state_rsn_desc>
<speed>auto</speed>
<ratemode>D</ratemode>
</ROW_interface>
</TABLE_interface>
</__readonly__>
</__XML_OPT_Cmd_show_interface_brief__readonly__>
</interface>
</show>
</nf:data>
</nf:rpc-reply>
]]>]]>

```



Note <__XML_OPT.*> and <__XML_BLK.*> appear in responses and are sometimes used in requests. These tags are used by the NETCONF agent and are present in responses after the <__readonly__> tag. They are necessary in requests, and should be added according to the schema file to reach the XML tag that represents the CLI command.

Information About Example XML Instances

Example XML Instances

This section provides examples of the following XML instances:

- [NETCONF Close Session Instance, on page 545](#)
- [NETCONF Kill Session Instance, on page 545](#)
- [NETCONF Copy Config Instance, on page 545](#)
- [NETCONF Edit Config Instance, on page 546](#)
- [NETCONF Get Config Instance, on page 548](#)
- [NETCONF Lock Instance, on page 548](#)
- [NETCONF Unlock Instance, on page 549](#)
- [NETCONF Commit Instance: Candidate Configuration Capability, on page 549](#)
- [NETCONF Confirmed Commit Instance, on page 550](#)
- [NETCONF Rollback-On-Error Instance, on page 550](#)

- [NETCONF Validate Capability Instance, on page 551](#)

NETCONF Close Session Instance

The following examples show the close-session request, followed by the close-session response:

Close Session Request

```
<?xml version="1.0"?>
<nc:rpc message-id="101" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0">
<nc:close-session/>
</nc:rpc>]]>]]>
```

Close Session Response

```
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0" message-id="101">
<nc:ok/>
</nc:rpc-reply>]]>]]>
```

NETCONF Kill Session Instance

The following examples show the kill session request, followed by the kill session response:

Kill Session Request

```
<nc:rpc message-id="101" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0">
<nc:kill-session>
<nc:session-id>25241</nc:session-id>
</nc:kill-session>
</nc:rpc>]]>]]>
```

Kill Session Response

```
<?xml version="1.0"?>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0" message-id="101">
<nc:ok/>
</nc:rpc-reply>]]>]]>
```

NETCONF Copy Config Instance



Note <startup/> is not supported as a source or target datastore. To perform any copy operation on **startup-config** like entering the **copy running-config startup-config** command, you need to fallback to the <exec-command> method.

The following examples show the copy config request, followed by the copy config response:

Copy Config Request

```
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<copy-config>
<target>
<running/>
</target>
<source>
<url>https://user@example.com:passphrase/cfg/new.txt</url>
</source>
</copy-config>
</rpc>
```

Copy Config Response

```
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>
```

NETCONF Edit Config Instance



Note XML edit-config with candidate datastore is not supported with 1.0 version XML request. It is supported only with the newer version which can be generated using xml in tool.

The following examples show the use of NETCONF edit config:

Edit Config Request

```
<?xml version="1.0"?>
<nc:rpc message-id="16" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0:if_manager">
<nc:edit-config>
<nc:target>
<nc:running/>
</nc:target>
<nc:config>
<configure>
<__XML_MODE__exec_configure>
<interface>
<ethernet>
<interface>2/30</interface>
<__XML_MODE_if-ethernet>
<__XML_MODE_if-eth-base>
<description>
<desc_line>Marketing Network</desc_line>
</description>
</__XML_MODE_if-eth-base>
</__XML_MODE_if-ethernet>
</ethernet>
</interface>
</__XML_MODE__exec_configure>
</configure>
```



```
</nc:config>
</nc:edit-config>
</nc:rpc>]]]]>
```

Edit Config Response

```
<?xml version="1.0"?>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:1.0:if_manager" message-id="16">
<nc:ok/>
</nc:rpc-reply>]]]]>
```

The operation attribute in edit config identifies the point in configuration where the specified operation is performed. If the operation attribute is not specified, the configuration is merged into the existing configuration data store. The operation attribute can have the following values:

- create
- merge
- delete

Edit Config: Delete Operation Request

The following example shows how to delete the configuration of interface Ethernet 0/0 from the running configuration:

```
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<edit-config>
<target>
<running/>
</target>
<default-operation>none</default-operation>
<config xmlns:xc="urn:ietf:params:xml:ns:netconf:base:1.0">
<top xmlns="http://example.com/schema/1.2/config">
<interface xc:operation="delete">
<name>Ethernet0/0</name>
</interface>
</top>
</config>
</edit-config>
</rpc>]]]]>
```

Response to Edit Config: Delete Operation

The following example shows how to edit the configuration of interface Ethernet 0/0 from the running configuration:

```
<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>]]]]>
```

NETCONF Get Config Instance

The following examples show the use of NETCONF get config:

Get Config Request to Retrieve the Entire Subtree

```
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<get-config>
<source>
<running/>
</source>
<filter type="subtree">
<top xmlns="http://example.com/schema/1.2/config">
<users/>
</top>
</filter>
</get-config>
</rpc>]]>]]>
```

Get Config Response with Results of a Query

```
<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<data>
<top xmlns="http://example.com/schema/1.2/config">
<users>
<user>
<name>root</name>
<type>superuser</type>
<full-name>Charlie Root</full-name>
<company-info>
<dept>1</dept>
<id>1</id>
</company-info>
</user>
<!-- additional <user> elements appear here... -->
</users>
</top>
</data>
</rpc-reply>]]>]]>
```

NETCONF Lock Instance

The following examples show a lock request, a success response, and a response to an unsuccessful attempt:

Lock Request

```
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<lock>
<target>
<running/>
</target>
</lock>
</rpc>]]>]]>
```

Response to a Successful Acquisition of Lock

```
<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/> <!-- lock succeeded -->
</rpc-reply>]]>]]>
```

Response to an Unsuccessful Attempt to Acquire Lock

```
<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<rpc-error> <!-- lock failed -->
<error-type>protocol</error-type>
<error-tag>lock-denied</error-tag>
<error-severity>error</error-severity>
<error-message>
Lock failed, lock is already held
</error-message>
<error-info>
<session-id>454</session-id>
<!-- lock is held by NETCONF session 454 -->
</error-info>
</rpc-error>
</rpc-reply>]]>]]>
```

NETCONF Unlock Instance

The following examples show the use of NETCONF unlock:

Unlock Request

```
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<unlock>
<target>
<running/>
</target>
</unlock>
</rpc>
```

Response to an Unlock Request

```
<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>
```

NETCONF Commit Instance: Candidate Configuration Capability

The following examples show a commit operation and a commit reply:

Commit Operation

```
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<commit/>
</rpc>
```

Commit Reply

```
<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>
```

NETCONF Confirmed Commit Instance

The following examples show a confirmed commit operation and a confirmed commit reply:

Confirmed Commit Request

```
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<commit>
<confirmed/>
<confirm-timeout>120</confirm-timeout>
</commit>
</rpc>]]]]>
```

Confirmed Commit Response

```
<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>]]]]>
```

NETCONF Rollback-On-Error Instance

The following examples show how to configure rollback on error and the response to this request:

Rollback-On-Error Capability

```
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<edit-config>
<target>
<running/>
</target>
<error-option>rollback-on-error</error-option>
<config>
<top xmlns="http://example.com/schema/1.2/config">
<interface>
<name>Ethernet0/0</name>
<mtu>100000</mtu>
</interface>
```

```

</top>
</config>
</edit-config>
</rpc>]]>]]>

```

Rollback-On-Error Response

```

<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>]]>]]>

```

NETCONF Validate Capability Instance

The following examples show the use of NETCONF validate capability. The string `urn:ietf:params:netconf:capability:validate:1.0` identifies the NETCONF validate capability.

Validate Request

```

xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<validate>
<source>
<candidate/>
</source>
</validate>
</rpc>]]>]]>

```

Response to Validate Request

```

<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>]]>]]>

```

Additional References

This section provides additional information that is related to implementing the XML management interface.

RFCs

| RFCs | Title |
|--------------------------|--|
| RFC 4741 | NETCONF Configuration Protocol |
| RFC 4742 | Using the NETCONF Configuration Protocol over Secure Shell (SSH) |



APPENDIX **A**

Streaming Telemetry Sources

- [About Streaming Telemetry, on page 553](#)
- [Guidelines and Limitations, on page 553](#)
- [Data Available for Telemetry, on page 553](#)

About Streaming Telemetry

The streaming telemetry feature of Cisco Nexus switches continuously streams data out of the network and notifies the client, providing near-real-time access to monitoring data.

Guidelines and Limitations

Following are the guideline and limitations for the streaming telemetry:

- For information about supported platforms, see the [Nexus Switch Platform Matrix](#).
- Cisco Nexus switches with less than 8 GB of memory do not support telemetry.

Data Available for Telemetry

For each component group, the distinguished names (DNs) in the appendix of the [NX-API DME Model Reference](#) can provide the listed properties as data for telemetry.



APPENDIX **B**

WebSocket Subscription

- [WebSocket Subscription, on page 555](#)

WebSocket Subscription

Cisco NX-OS provides an interface capability to enable the switch to push notifications to interested subscribers. Through the NX-API WebSocket interface, programs and end-users can receive notifications about various state changes on the switch, eliminating the need for periodic polling.

When you perform an API query using the Cisco NX-API REST interface, you have the option to create a subscription to any future changes in the results of a given query. When any management object (MO) is created, changed, or deleted, because of a user-initiated or system-initiated action, an event is generated. If the received event changes the results of a subscribed query, the switch generates a push notification to the API client that created the subscription.

•

Opening a WebSocket

The API subscription feature uses the WebSocket protocol (RFC 6455) to implement a two-way connection with the API client. This way, the API can send unsolicited notification messages to the client itself. To establish the notification channel, you must first open a WebSocket connection with the respective API. Only a single WebSocket connection is needed to support multiple query subscriptions within each switch. The WebSocket connection is dependent on your API session connection (via token validation), and closes when your API session ends.

There are many ways to open a WebSocket connection. You can write python client as following:

```
from websocket import create_connection

connection_string = "ws:// 10.1.2.3/socket{0}".format(token)

ws = create_connection(connection_string, sslopt={"check_hostname": False})
```

In the URI, the token is the current API session token (cookie). This example shows the URI with a token:

```
ws://10.1.2.3/socketGkZ15NLRZJ15+jqChouaZ9CYjgE58W/pMccR+LeXmd00obG9NB
Iwo1VBo7+YC1oiJL9mS6I9qh62BkX+Xddhe0JYrTmSG4JcKZ4t3bcP2Mxy3VBmgoJjwZ76Zouf9V9AD6X
1831yoR4bLBzqbSSU1R2NIgUotCGWjZt5JX6CJF0=
```

Creating a Subscription

To create a subscription to a query, perform the query with the option “?subscription=yes”. This example creates a subscription to a query of the `sys/intf/phys-[eth1/1]` in the JSON format:

```
GET http://10.1.1.1/api/mo/sys/intf/phys-[eth1/1].json?subscription=yes
```

The query response contains a subscription identifier, `subscriptionId`, that you can use to refresh the subscription and identify future notifications from the given subscription.

```
{"totalCount":"0","subscriptionId":"18374686685813276673","imdata":[]}
```

Receiving Notifications

An event notification from the subscription delivers a data structure that contains the subscription ID and the MO description. In this JSON example, `sys/intf/phys-[eth1/1]` description is changed to “test”.

```
{"subscriptionId":["18374686685813276673"],"imdata":[{"11PhysIf": {"attributes": {"childAction": "", "descr": "test", "dn": "sys/intf/phys-[eth1/1]", "modTs": "2019-10-18T19:42:29.446+00:00", "rn": "", "status": "modified"}}}]}
```

As multiple active subscriptions can exist for a given query, a notification can contain multiple subscription IDs; similar as shown in the example above. Notifications are supported in either JSON or XML format.

Refreshing the Subscription

In order to continue receiving event notifications, you must periodically refresh each subscription during your API session. To refresh a subscription, send an HTTP GET message to the API method `subscriptionRefresh` with the parameter `id` equal to the `subscriptionId` shown in the example:

```
GET http://10.1.1.1/api/subscriptionRefresh.json?id=18374686685813276673
```

The API returns an empty response to the refresh message unless the subscription has expired.



Note The timeout period for a WebSocket subscription is 90 seconds by default. To prevent loss of notifications, you must send a subscription refresh message at least once every 90 seconds.

In summary, WebSocket provides a powerful tool for allowing publisher-subscriber communication for event subscription within the NX-OS REST API.



APPENDIX **C**

Programmability RFCs

- [Programmability RFCs, on page 557](#)

Programmability RFCs

This table lists the RFC compliance standards. For information on each RFC, see www.ietf.org.

Table 50: RFC Compliance Standards

| RFCs | Title |
|----------|---|
| RFC 5277 | NETCONF Event Notifications |
| RFC 6241 | Network Configuration Protocol (NETCONF) |
| RFC 6243 | With-defaults Capability for NETCONF (Supported for report-all only) |



INDEX

B

- Bash [21, 24](#)
 - accessing [21](#)
 - examples [24](#)
 - feature bash-shell [21](#)
- bcm-shell [87, 89–91](#)
 - CLI [87](#)
 - examples [91](#)
 - fabric module [89](#)
 - line card [90](#)

N

- NX-API [221–223, 226, 239, 244, 249, 261, 269, 281](#)
 - CLI [223](#)
 - cookie [223](#)
 - management commands [226](#)
 - message format [222](#)
 - request elements [239](#)
 - response codes [249](#)
 - response elements [244](#)
 - sample scripts [261](#)
 - security [223](#)
 - transport [222](#)
 - user interface [269, 281](#)

P

- Python [97, 99–100, 102–104](#)
 - Cisco package [97](#)
 - CLI command API [99](#)
 - display format examples [100](#)
 - embedded event manager (EEM) [103](#)
 - invoking [100](#)
 - non-interactive [102](#)
 - NX-OS network interfaces [104](#)
 - NX-OS security [104](#)
 - scripts [97](#)

T

- tcl [109–111, 114](#)
 - cli commands [110](#)
 - command separation [110](#)
 - history [110](#)
 - no interactive help [109](#)
 - options [111](#)
 - references [114](#)
 - sandbox [111](#)
 - security [111](#)
 - tab completion [110](#)
 - telquit command [111](#)
 - variables [111](#)
- Tool Command Language, *See* tcl

