

Cisco Nexus 3000 Series NX-OS Programmability Guide, Release 9.3(x)

First Published: 2019-07-20

Last Modified: 2020-09-03

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/c/en/us/about/legal/cloud-and-software/software-terms.html. Cisco product warranty information is available at https://www.cisco.com/c/en/us/products/warranty-listing.html. 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: https://www.cisco.com/c/en/us/about/legal/trademarks.html. 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)

© 2019–2020 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface xix

Audience xix

Document Conventions xix

Related Documentation for Cisco Nexus 3000 Series Switches xx

Documentation Feedback xx

Communications, Services, and Additional Information xx

CHAPTER 1

New and Changed Information 1

New and Changed Information 1

CHAPTER 2

Supported Platforms for Programmability Features 5

Platform Support for Programmability Features 5

CHAPTER 3

Overview 9

Programmability Overview 9

Standard Network Manageability Features 10

Advanced Automation Feature 10

PowerOn Auto Provisioning Support 10

OpenStack Integration 10

Programmability Support 12

NX-API Support 12

Python Scripting 12

Tcl Scripting 12

Broadcom Shell 12

Bash 12

Guest Shell 12

Container Tracker Support 13

PART I Shells and Scripting 15

CHAPTER 4 Bash 17

About Bash 17
Guidelines and Limitations 17
Accessing Bash 17
Escalate Privileges to Root 18
Examples of Bash Commands 21
Displaying System Statistics 21
Running Bash from CLI 21
Managing RPMs 22
Installing RPMs from Bash 22
Upgrading Feature RPMs 23
Downgrading a Feature RPM 23
Erasing a Feature RPM 24
Support for DME Modularity 24
Installing the DME RPMs 24
Verifying the Installed RPM 27
Querying for the RPM in the Local Repo 28
Downgrading Between Versions of DME RPM 28
Downgrading to the Base RPM 30
Persistently Daemonizing an SDK- or ISO-built Third Party Process 32
Persistently Starting Your Application from the Native Bash Shell 33
Copy Through Kstack 33
An Example Application in the Native Bash Shell 34

CHAPTER 5 Guest Shell 37

About the Guest Shell 37
Guidelines and Limitations 38
Accessing the Guest Shell 42
Resources Used for the Guest Shell 43
Capabilities in the Guestshell 43

NX-OS CLI in the Guest Shell 44
Network Access in Guest Shell 44
Access to Bootflash in Guest Shell 46
Python in Guest Shell 47
Python 3 in Guest Shell versions up to 2.10 (CentOS 7) 47
Installing RPMs in the Guest Shell 50
Security Posture for Guest Shell 51
Kernel Vulnerability Patches 51
ASLR and X-Space Support 52
Namespace Isolation 52
Root-User Restrictions 53
Resource Management 54
Guest File System Access Restrictions 54
Managing the Guest Shell 54
Disabling the Guest Shell 57
Destroying the Guest Shell 58
Enabling the Guest Shell 58
Replicating the Guest Shell 60
Exporting Guest Shell rootfs 60
Importing Guest Shell rootfs 60
Importing YAML File 62
show guestshell Command 65
Verifying Virtual Service and Guest Shell Information 66
Persistently Starting Your Application From the Guest Shell 67
Procedure for Persistently Starting Your Application from the Guest Shell 68
An Example Application in the Guest Shell 68
Troubleshooting Guest Shell Issues 69

CHAPTER 6 **Broadcom Shell** 71
About the Broadcom Shell 71
Guidelines and Limitations 71
Accessing the Broadcom Shell (bcm-shell) 71
Accessing bcm-shell with the CLI API 71
Accessing the Native bcm-shell on the Fabric Module 72

Accessing the bcm-shell on the Line Card 73

CHAPTER 7 Barefoot Shell 75

About the Barefoot Shell 75

Guidelines and Limitations 75

Accessing the Barefoot Shell with CLI API 76

CHAPTER 8 Python API 79

About the Python API 79

Using Python 79

Cisco Python Package 79

Using the CLI Command APIs 81

Invoking the Python Interpreter from the CLI 82

Display Formats 83

Non-Interactive Python 84

Running Scripts with Embedded Event Manager 85

Python Integration with Cisco NX-OS Network Interfaces 86

Cisco NX-OS Security with Python 87

Examples of Security and User Authority 87

Example of Running Script with Schedular 88

CHAPTER 9 Scripting with Tcl 91

About Tcl 91

Guidelines and Limitations 91

Tclsh Command Help 91

Tclsh Command History 92

Tclsh Tab Completion 92

Tclsh CLI Command 92

Tclsh Command Separation 93

Tcl Variables 93

Tclquit 93

Tclsh Security 93

Running the Tclsh Command 94

Navigating Cisco NX-OS Modes from the Tclsh Command 95

Tcl References 96

CHAPTER 10 iPXE 97

About iPXE 97

Netboot Requirements 98

Guidelines and Limitations 98

Boot Mode Configuration 98

Verifying the Boot Order Configuration 100

CHAPTER 11 Kernel Stack 101

About Kernel Stack 101

Guidelines and Limitations 101

Changing the Port Range 102

PART II Applications 105

CHAPTER 12 Third-Party Applications 107

About Third-Party Applications 107

Installing Third-Party Native RPMs/Packages 107

Installing Signed RPM 109

Checking a Signed RPM 109

Installing Signed RPMs by Manually Importing Key 109

Installing Signed Third-Party RPMs by Importing Keys Automatically 111

Adding Signed RPM into Repo 113

Persistent Third-Party RPMs 114

Installing RPM from VSH 115

Package Addition 115

Package Activation 116

Deactivating Packages 116

Removing Packages 117

Displaying Installed Packages 117

Displaying Detail Logs 117

Upgrading a Package 118

Downgrading a Package 118

Third-Party Applications 119
NX-OS 119
DevOps Configuration Management Tools 119
collectd 119
Ganglia 119
Iperf 119
LLDP 120
Nagios 120
OpenSSH 120
Quagga 120
Splunk 120
tcollector 121
tcpdump 121
Tshark 121

CHAPTER 13 Ansible 123
Prerequisites 123
About Ansible 123
Cisco Ansible Module 123

CHAPTER 14 SaltStack 125
About SaltStack 125
About NX-OS and SaltStack 126
Guidelines and Limitations 126
Cisco NX-OS Environment for SaltStack 126
Enabling NX-API for SaltStack 127
Installing SaltStack for NX-OS 127

CHAPTER 15 Puppet Agent 129
About Puppet 129
Prerequisites 129
Puppet Agent NX-OS Environment 130
ciscopuppet Module 130

CHAPTER 16 Using Chef Client with Cisco NX-OS 131
About Chef 131
Prerequisites 131
Chef Client NX-OS Environment 132
cisco-cookbook 132

CHAPTER 17 Nexus Application Development - Yocto 135
About Yocto 135
Installing Yocto 135

CHAPTER 18 Nexus Application Development - SDK 139
About the Cisco SDK 139
Installing the SDK 139
Procedure for Installation and Environment Initialization 140
Using the SDK to Build Applications 141
Using RPM to Package an Application 142
Creating an RPM Build Environment 143
Using General RPM Build Procedure 143
Example to Build RPM for collectd with No Optional Plug-Ins 144
Example to Build RPM for collectd with Optional Curl Plug-In 145

CHAPTER 19 NX-SDK 147
About the NX-SDK 147
Considerations for Go Bindings 148
About On-Box (Local) Applications 148
Default Docker Images 148
Guidelines and Limitations for NX-SDK 149
Off-Box (Remote) Applications 149
About NX-SDK 2.0 149
About NX-SDK 2.5 150
About Remote Applications 150
NX-SDK Security 150
Security Profiles for NX SDK 2.0 151

CHAPTER 20 **Using Docker with Cisco NX-OS** 153

About Docker with Cisco NX-OS 153

Prerequisites for Setting Up Docker Containers Within Cisco NX-OS 153

Starting the Docker Daemon 154

Configure Docker to Start Automatically 155

Starting Docker Containers: Host Networking Model 155

Starting Docker Containers: Bridged Networking Model 157

Mounting the bootflash and volatile Partitions in the Docker Container 158

Enabling Docker Daemon Persistence on Enhanced ISSU Switchover 158

Resizing the Docker Storage Backend 159

Stopping the Docker Daemon 161

Docker Container Security 162

Securing Docker Containers With User namespace Isolation 162

Moving the cgroup Partition 163

Docker Troubleshooting 164

Docker Fails to Start 164

Docker Fails to Start Due to Insufficient Storage 164

Failure to Pull Images from Docker Hub (509 Certificate Expiration Error Message) 165

Failure to Pull Images from Docker Hub (Client Timeout Error Message) 165

Docker Daemon or Containers Not Running On Switch Reload or Switchover 166

Resizing of Docker Storage Backend Fails 166

Docker Container Doesn't Receive Incoming Traffic On a Port 166

Unable to See Data Port And/Or Management Interfaces in Docker Container 167

General Troubleshooting Tips 167

PART III **NX-API** 169

CHAPTER 21 **NX-API CLI** 171

About NX-API CLI 171

Transport 171

Message Format 172

Security 172

Using NX-API CLI 172

Escalate Privileges to Root on NX-API 174
NX-API Management Commands 175
Working With Interactive Commands Using NX-API 179
NX-API Request Elements 179
NX-API Response Elements 183
Restricting Access to NX-API 184
Updating an iptable 184
Making an Iptable Persistent Across Reloads 186
Table of NX-API Response Codes 187
XML and JSON Supported Commands 189
About JSON (JavaScript Object Notation) 190
Examples of XML and JSON Output 190

CHAPTER 22 NX-API REST 199
About NX-API REST 199
DME Config Replace Through REST 199
About DME Full Config Replace Through REST Put 199
Guidelines and Limitations 200
Replacing the System-Level Configuration Through REST PUT 200
Replacing Feature-Level Config Through REST PUT 201
Replacing Property-Level Config Through REST POST 202
Troubleshooting Config Replace for REST PUT 203

CHAPTER 23 NX-API Developer Sandbox 205
NX-API Developer Sandbox: NX-OS Releases Prior to 9.2(2) 205
About the NX-API Developer Sandbox 205
Guidelines and Limitations 206
Configuring the Message Format and Command Type 206
Using the Developer Sandbox 208
Using the Developer Sandbox to Convert CLI Commands to Payloads 208
NX-API Developer Sandbox: NX-OS Release 9.2(2) and Later 211
About the NX-API Developer Sandbox 211
Guidelines and Limitations 212
Configuring the Message Format and Input Type 214

Using the Developer Sandbox 217
Using the Developer Sandbox to Convert CLI Commands to REST Payloads 218
Using the Developer Sandbox to Convert from REST Payloads to CLI Commands 220
Using the Developer Sandbox to Convert from RESTCONF to json or XML 225

PART IV Model-Driven Programmability 229

CHAPTER 24 Infrastructure Overview 231
About Model-Driven Programmability 231
About the Programmable Interface Infrastructure 231

CHAPTER 25 Managing Components 235
About the Component RPM Packages 235
Preparing For Installation 237
Downloading Components from the Cisco Artifactory 238
Installing RPM Packages 239
Installing the Programmable Interface Base And Common Model Component RPM Packages 239

CHAPTER 26 NETCONF Agent 241
About the NETCONF Agent 241
Guidelines and Limitations for the NETCONF Agent 242
Configuring the NETCONF Agent 244
Configuring the NETCONF Agent for Cisco NX-OS 9.3(5) and Later 244
Configuring the NETCONF Agent for Cisco NX-OS 9.3(4) and Earlier 245
Establishing a NETCONF Session 245
NETCONF Operations 247
NETCONF Notifications 255
About NETCONF Notifications 255
Capabilities Exchange 256
Event Stream Discovery 256
Creating Subscriptions 257
Receiving Notifications 258
Terminating Subscriptions 258
NETCONF Examples 259

Troubleshooting the NETCONF Agent 263

CHAPTER 27 **Converting CLI Commands to Network Configuration Format** 265

Information About XMLIN 265

Licensing Requirements for XMLIN 265

Installing and Using the XMLIN Tool 266

Converting Show Command Output to XML 266

Configuration Examples for XMLIN 267

CHAPTER 28 **RESTConf Agent** 271

About the RESTCONF Agent 271

Configuring the RESTConf Agent 272

Using the RESTCONF Agent 272

Troubleshooting the RESTCONF Agent 273

CHAPTER 29 **gRPC Agent** 275

About the gRPC Agent 275

Configuring the gRPC Agent for Cisco NX-OS Release 9.3(2) and Earlier 276

Using the gRPC Agent 278

Troubleshooting the gRPC Agent 280

gRPC Protobuf File 280

CHAPTER 30 **gNMI - gRPC Network Management Interface** 289

About gNMI 289

gNMI RPC and SUBSCRIBE 290

Guidelines and Limitations for gNMI 291

Configuring gNMI 293

Configuring Server Certificate 294

Generating Key/Certificate Examples 295

Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(2) and Earlier 295

Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(3) and Later 296

Verifying gNMI 298

Clients 303

Sample DME Subscription - PROTO Encoding 303

Capabilities 304
About Capabilities 304
Guidelines and Limitations for Capabilities 305
Example Client Output for Capabilities 306
Get 308
About Get 308
Guidelines and Limitations for Get 308
Set 309
About Set 309
Guidelines and Limitations for Set 309
Subscribe 310
Guidelines and Limitations for Subscribe 310
gNMI Payload 311
Streaming Syslog 314
About Streaming Syslog for gNMI 314
Guidelines and Limitations for Streaming Syslog - gNMI 314
Syslog Native YANG Model 315
Subscribe Request Example 315
Sample PROTO Output 316
Sample JSON Output 319
Troubleshooting 320
Gathering TM-Trace Logs 320
Gathering MTX-Internal Logs 320

CHAPTER 31 **Dynamic Logger** 325
Dynamic Logger 325
Prerequisites 325
Reference 325

CHAPTER 32 **Model-Driven Telemetry** 333
About Telemetry 333
Telemetry Components and Process 333
High Availability of the Telemetry Process 335
Licensing Requirements for Telemetry 335

Installing and Upgrading Telemetry 335
Guidelines and Limitations for Telemetry 336
Configuring Telemetry Using the CLI 342
Configuring Telemetry Using the NX-OS CLI 342
Configuration Examples for Telemetry Using the CLI 348
Displaying Telemetry Configuration and Statistics 351
Displaying Telemetry Log and Trace Information 360
Configuring Telemetry Using the NX-API 361
Configuring Telemetry Using the NX-API 361
Configuration Example for Telemetry Using the NX-API 370
Telemetry Model in the DME 373
Streaming Syslog 375
About Streaming Syslog for gNMI 375
Configuring the YANG Data Source Path for Syslog Information 375
Telemetry Data Streamed for Syslog Path 376
Sample JSON Output 378
Sample KVGPB Output 378
Telemetry Path Labels 380
About Telemetry Path Labels 380
Polling for Data or Receiving Events 381
Guidelines and Limitations for Path Labels 381
Configuring the Interface Path to Poll for Data or Events 382
Configuring the Interface Path for Non-Zero Counters 383
Configuring the Interface Path for Operational Speeds 385
Configuring the Interface Path with Multiple Queries 386
Configuring the Environment Path to Poll for Data or Events 388
Configuring the Resources Path to Poll for Events or Data 389
Configuring the VXLAN Path to Poll for Events or Data 390
Verifying the Path Label Configuration 392
Displaying Path Label Information 392
About Native Data Source Paths 395
About Native Data Source Paths 395
Telemetry Data Streamed for Native Data Source Paths 395
Guidelines and Limitations 397

Configuring the Native Data Source Path for Routing Information 398
Configuring the Native Data Source Path for MAC Information 399
Configuring the Native Data Path for IP Adjacencies 401
Displaying Native Data Source Path Information 403
Additional References 404
Related Documents 404

CHAPTER 33 **Inband Network Telemetry** 405
About Inband Network Telemetry 405
Licensing Requirements for Inband and Postcard Telemetry 410
INT Collection Parameters 411
Example Inband Network Telemetry Headers 411
Telemetry Report Formats 413
Guidelines and Restrictions 414
Configuring Inband Network Telemetry 415
Flow and Queue Profile Arguments for Report Types 421
Configuring Packet Postcards 421

CHAPTER 34 **Packetized MMU Stats** 427
About Packetized MMU Stats 427
Guidelines and Limitations 427
Using Packetized MMU Stats 427

PART V **XML Management Interface** 437

CHAPTER 35 **XML Management Interface** 439
About the XML Management Interface 439
About the XML Management Interface 439
NETCONF Layers 439
SSH xmlagent 440
Licensing Requirements for the XML Management Interface 440
Prerequisites to Using the XML Management Interface 441
Using the XML Management Interface 441
Configuring SSH and the XML Server Options Through the CLI 441

Starting an SSH Session 442
Sending the Hello Message 443
Obtaining the XSD Files 443
Sending an XML Document to the XML Server 444
Creating NETCONF XML Instances 444
RPC Request Tag rpc 445
NETCONF Operations Tags 446
Device Tags 447
Extended NETCONF Operations 449
NETCONF Replies 452
RPC Response Tag 453
Interpreting Tags Encapsulated in the Data Tag 453
Information About Example XML Instances 454
Example XML Instances 454
NETCONF Close Session Instance 454
NETCONF Kill-session Instance 455
NETCONF copy-config Instance 455
NETCONF edit-config Instance 455
NETCONF get-config Instance 457
NETCONF Lock Instance 457
NETCONF unlock Instance 458
NETCONF Commit Instance - Candidate Configuration Capability 459
NETCONF Confirmed-commit Instance 459
NETCONF rollback-on-error Instance 459
NETCONF validate Capability Instance 460
Additional References 460

APPENDIX A **Streaming Telemetry Sources** 463
About Streaming Telemetry 463
Data Available for Telemetry 463

APPENDIX B **Packet Timestamping** 465
About Packet Timestamping 465
Guidelines and Limitations 467



Preface

This preface includes these sections:

- Audience, on page xix
- Document Conventions, on page xix
- Related Documentation for Cisco Nexus 3000 Series Switches, on page xx
- Documentation Feedback, on page xx
- Communications, Services, and Additional Information, on page xx

Audience

This publication is for network administrators who install, configure, and maintain Cisco Nexus switches.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.

Convention	Description
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Related Documentation for Cisco Nexus 3000 Series Switches

The entire Cisco Nexus 3000 Series switch documentation set is available at the following URL:

https://www.cisco.com/c/en/us/support/switches/nexus-3000-series-switches/tsd-products-support-series-home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus3k-docfeedback@cisco.com. We appreciate your feedback.

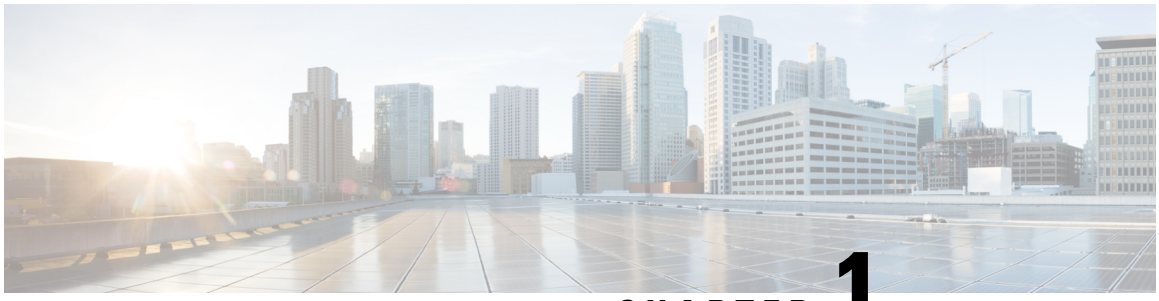
Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco DevNet.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

Preface



CHAPTER 1

New and Changed Information

This chapter provides release-specific information for each new and changed feature in this release of the *Cisco Nexus 3000 Series NX-OS Programmability Guide, 9.3(x)*.

- New and Changed Information, on page 1

New and Changed Information

This table summarizes the new and changed features for the *Cisco Nexus 3000 Series NX-OS Programmability Guide, Release 9.3(x)* and which sections contain them.

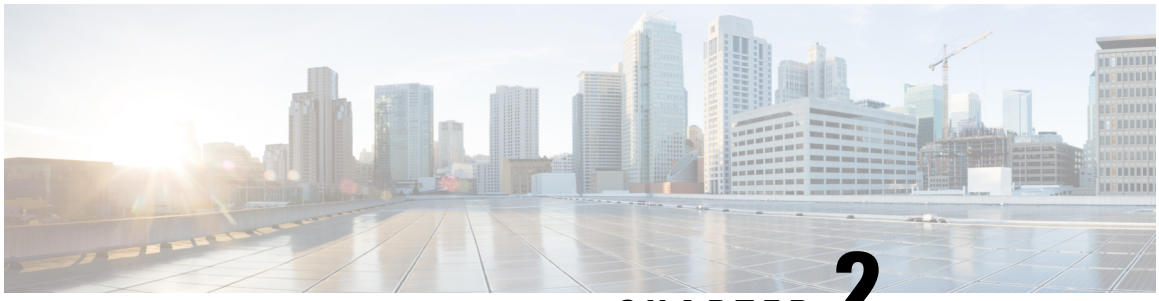
Table 1: New and Changed Features

Feature	Description	Changed in Release	Where Documented
Alias Option for Sensor Path for Model-Driven Telemetry	Added the **alias** keyword for the **path** CLI command.	9.3(5)	Configuring Telemetry Using the NX-OS CLI, on page 342
gNMI Get/Set	Added the Get and Set RPCs.	9.3(5)	gNMI - gRPC Network Management Interface, on page 289
NETCONF gRPC Support	Documents the NETCONF configuration procedure beginning with Cisco NX-OS Release 9.3(5)	9.3(5)	Configuring the NETCONF Agent for Cisco NX-OS 9.3(5) and Later, on page 244
Python 3 on NX-OS	Python 3 support added.	9.3(5)	About the Python API , on page 79
YANG Support for Multiple Keys	Added configuration for a single telemetry DME stream.	9.3(5)	Configuring Telemetry Using the NX-OS CLI, on page 342

Feature	Description	Changed in Release	Where Documented
nxapi idle-timeout <timeout>	Enables you to configure the amount of time before an idle NX-API session is invalidated.	9.3(5)	NX-API Management Commands, on page 175
Streaming Syslog and Filtered Syslog	Support added for Streaming 9.3(3) Syslog and Filtered Syslog in the NX-SDK, gNMI – gRPC Network Management Interface, and Model Driven Telemetry chapters.	9.3(3)	About Streaming Syslog for gNMI, on page 314
NX-SDK	Added support for NX-SDK 2.0, which includes support for remote (off-box) applications and enhanced security	9.3(1)	NX-SDK, on page 147
SaltStack	Added support for SaltStack automation and integration	9.3(1)	SaltStack, on page 125
NX-OS **copy** through kstack	Enabled speed and operability of the switch's **copy** commands	9.3(1)	Copy Through Kstack, on page 33
NX-OS **show** command enhancement	Added JSON Native and JSON Native Pretty support for NX-OS **show** commands	9.3(1)	XML and JSON Supported Commands, on page 189
NX-API Chunking	Enhanced messages and chunking functionality	9.3(1)	NX-API CLI, on page 171
REST API	Added REST support for replacing the switch configuration at the tree, subtree, and leaf level of the DME	9.3(1)	About DME Full Config Replace Through REST Put, on page 199
Modular DME Infrastructure	Added modularity to support non-intrusive upgrade or downgrade of DME RPMs	9.3(1)	Support for DME Modularity, on page 24
Model Driven Telemetry, ease of use enhancements	Added Path Labels, which consolidate multiple queries for telemetry path data. Adds support for setting or changing the node ID string for telemetry messages.	9.3(1)	About Telemetry Path Labels, on page 380 Model-Driven Telemetry, on page 333

Feature	Description	Changed in Release	Where Documented
Model-Driven Telemetry, Native Data Source Paths	Added support for applications to stream telemetry data without the restriction of a specific type of infrastructure or database.	9.3(1)	About Native Data Source Paths, on page 395
No updates since Cisco NX-OS Release 9.2(x)	First 9.3(x) release	Not applicable	Not applicable

New and Changed Information



CHAPTER 2

Supported Platforms for Programmability Features

This chapter contains the following topics:

- 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.

Chef Client

Return to Using Chef Client with Cisco NX-OS, on page 131.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Chef Agent	Cisco Nexus 3100 platform switches with at least 8 GB of system RAM Cisco Nexus 3000 Series switches with at least 8 GB of system RAM	Cisco NX-OS 7.0(3)I2(1)	

Docker

Return to Using Docker with Cisco NX-OS, on page 153.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Docker Container support within NX-OS	Cisco Nexus 3000 Series switches with at least 8 GB of system RAM	Cisco NX-OS 9.2(1)	

gNMI - gRPC Network Management Interface

Return to gNMI - gRPC Network Management Interface, on page 289.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
gNMI Get/Set	Cisco Nexus 3000 Series switches	Cisco NX-OS 9.3(5)	

gRPC Agent

Return to gRPC Agent, on page 275.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Config Replace	Cisco Nexus 3100-V and 3100-XL switches Cisco Nexus 3200 platform switches	Cisco NX-OS 9.3(1)	

Model-Driven Telemetry

Return to Model-Driven Telemetry, on page 333.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Alias Option for Sensor Path	Cisco Nexus 3000 Series switches	Cisco NX-OS 9.3(5)	
YANG Support for Multiple Keys	Cisco Nexus 3000 Series switches	Cisco NX-OS 9.3(5)	

NETCONF Agent

Return to NETCONF Agent, on page 241.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
NETCONF gRPC Support	Cisco Nexus 3000 Series switches	Cisco NX-OS 9.3(5)	

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Config Replace	Cisco Nexus 3000 Series switches with at least 8 GB of system RAM Cisco Nexus 3100-V and 3100-XL platform switches Cisco Nexus 3200 platform switches	Cisco NX-OS 9.3(1)	

NX-API REST

Return to NX-API REST, on page 199.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
DME Config Replace	Cisco Nexus 3000 Series switches with at least 8 GB of system RAM	Cisco NX-OS 9.3(1)	

Packet Timestamping

Return to Packet Timestamping, on page 465.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Packet Timestamping	Cisco Nexus 3132C-Z switch Cisco Nexus 3264C-E switch	-	

Packetized MMU Statistics

Return to Packetized MMU Stats, on page 427.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Packetized MMU Statistics	Cisco Nexus 3132C-Z switch Cisco Nexus 3264C-E switches	-	

Python API

Return to Python API, on page 79.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Python 3 on NX-OS	Cisco Nexus 3000 Series switches	Cisco NX-OS 9.3(5)	

Puppet Agent

Return to Puppet Agent, on page 129.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Puppet Agent	Cisco Nexus 3100 platform switches with at least 8 GB of system RAM Cisco Nexus 3000 Series switches with at least 8 GB of system RAM	Cisco NX-OS 7.0(3)I2(1)	

RESTCONF Agent

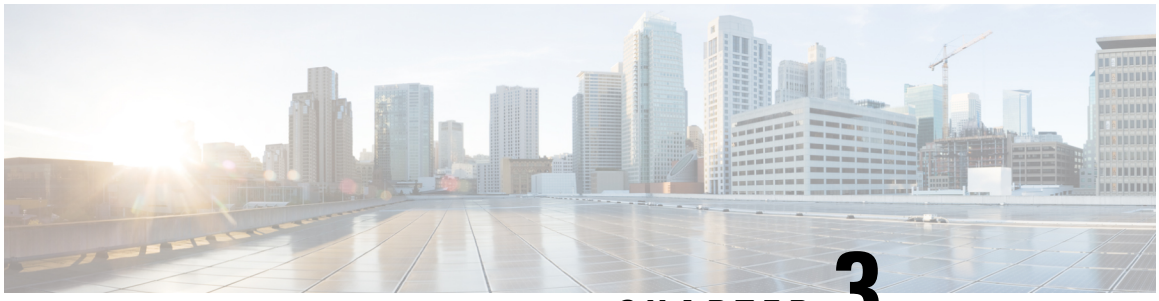
Return to RESTConf Agent, on page 271.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Config Replace	Cisco Nexus 3100-V and 3100-XL switches Cisco Nexus 3200 platform switches	Cisco NX-OS 9.3(1)	

Streaming Telemetry

Return to Streaming Telemetry Sources, on page 463.

Feature	Supported Platforms or Line Cards	First Supported Release	Platform Exceptions
Telemetry sources	Cisco Nexus 3000 switches with at least 8 GB of system RAM	-	



CHAPTER 3

Overview

- Programmability Overview, on page 9
- Standard Network Manageability Features, on page 10
- Advanced Automation Feature, on page 10
- Programmability Support, on page 12

Programmability Overview

The Cisco NX-OS software running on the Cisco Nexus 3000 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/.

Standard Network Manageability Features

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

Advanced Automation Feature

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

PowerOn Auto Provisioning Support

PowerOn Auto Provisioning (POAP) automates the process of installing/upgrading software images and installing configuration files on Cisco Nexus switches that are being deployed in the network for the first time. It reduces the manual tasks required to scale the network capacity.

When a Cisco Nexus switch with the POAP feature boots and does not find the startup configuration, the switch enters POAP mode. It locates a DHCP server and bootstraps itself with its interface IP address, gateway, and DNS server IP addresses. The switch obtains the IP address of a TFTP server or the URL of an HTTP server and downloads a configuration script that enables the switch to download and install the appropriate software image and configuration file.

OpenStack Integration

The Cisco Nexus 3000 Series switches support the Cisco Nexus plugin for OpenStack Networking, also known as Neutron (http://www.cisco.com/web/solutions/openstack/index.html). The plugin allows you to build an infrastructure as a service (IaaS) network and to deploy a cloud network. With OpenStack, you can build an on-demand, self-service, multitenant computing infrastructure. However, implementing OpenStack's VLAN networking model across virtual and physical infrastructures can be difficult.

The OpenStack Networking extensible architecture supports plugins to configure networks directly. However, when you choose a network plugin, only that plugin's target technology is configured. When you are running OpenStack clusters across multiple hosts with VLANs, a typical plugin configures either the virtual network infrastructure or the physical network, but not both.

The Cisco Nexus plugin solves this difficult problem by including support for configuring both the physical and virtual networking infrastructure.

The Cisco Nexus plugin accepts OpenStack Networking API calls and uses the Network Configuration Protocol (NETCONF) by default or the REST API to configure Cisco Nexus switches as well as Open vSwitch (OVS) that runs on the hypervisor. The Cisco Nexus plugin configures VLANs on both the physical and virtual

network. It also allocates scarce VLAN IDs by deprovisioning them when they are no longer needed and reassigning them to new tenants whenever possible. VLANs are configured so that virtual machines that run on different virtualization (compute) hosts that belong to the same tenant network transparently communicate through the physical network. In addition, connectivity from the compute hosts to the physical network is trunked to allow traffic only from the VLANs that are configured on the host by the virtual switch.



Note We recommend configuring the REST API driver. For more configuration details, go to:

http://docwiki.cisco.com/wiki/Neutron_ML2_Driver_For_Cisco_Nexus_Devices_Ocata_Release

The following table lists the features of the Cisco Nexus plugin for OpenStack Networking:

Table 2: Summary of Cisco Nexus Plugin features for OpenStack Networking (Neutron)

Considerations	Description	Cisco Nexus Plugin
Extension of tenant VLANs across virtualization hosts	VLANs must be configured on both physical and virtual networks. OpenStack Networking supports only a single plugin at a time. You must choose which parts of the networks to manually configure.	Accepts networking API calls and configures both physical and virtual switches.
Efficient use of scarce VLAN IDs	Static provisioning of VLAN IDs on every switch rapidly consumes all available VLAN IDs, which limits scalability and makes the network vulnerable to broadcast storms.	Efficiently uses limited VLAN IDs by provisioning and deprovisioning VLANs across switches as tenant networks are created and destroyed.
Easy configuration of tenant VLANs in a top-of-rack (ToR) switch	You must statically provision all available VLANs on all physical switches. This process is manual and error prone.	Dynamically provisions tenant-network-specific VLANs on switch ports connected to virtualization hosts through the Nexus plugin driver.
Intelligent assignment of VLAN IDs	Switch ports connected to virtualization hosts are configured to handle all VLANs. Hardware limits are reached quickly.	Configures switch ports connected to virtualization hosts only for the VLANs that correspond to the networks configured on the host. This feature enables accurate port and VLAN associations.
Aggregation switch VLAN configuration for large multirack deployments.	When compute hosts run in several racks, you must fully mesh top-of-rack switches or manually trunk aggregation switches.	Supports Cisco Nexus 2000 Series Fabric Extenders to enable large, multirack deployments and eliminates the need for an aggregation switch VLAN configuration.

Programmability Support

Cisco NX-OS on Cisco Nexus 3000 Series switches support the following capabilities to aid programmability:

NX-API Support

Cisco NX-API allows for HTTP-based programmatic access to the Cisco Nexus 3000 Series 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 Cisco Nexus 3000 Series switches.

Python Scripting

Cisco Nexus 3000 Series switches support Python v2.7.5 in both interactive and non-interactive (script) modes.

The Python scripting capability on the devices provide 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 3000 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 3000 Series switches 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 3000 Series switches support direct Bourne-Again SHell (Bash) access. With Bash, you can access the underlying Linux system on the device and manage the system.

Guest Shell

The Cisco Nexus 3000 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 NX-OS software. With the guest shell you can add software packages and update libraries as needed without impacting the host system software.

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:

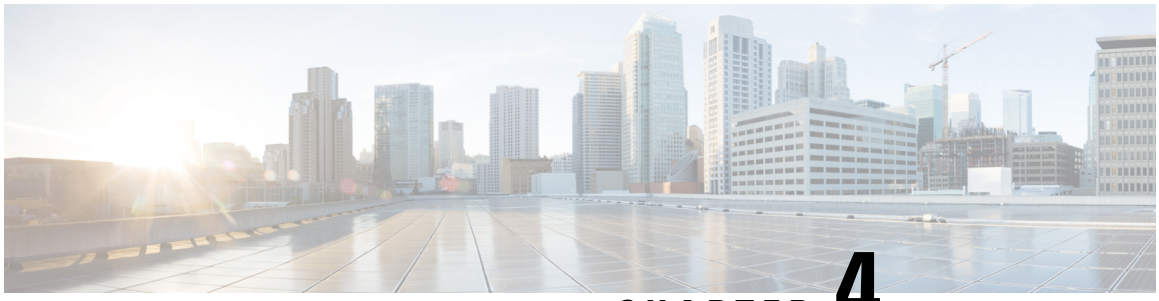
- The **show containers kubernetes** command obtains data from *kube-apiserver* using API calls over HTTP.
- The **kubernetes watch** *resource* command uses a daemon to subscribe to requested resources and process streaming data from *kube-apiserver*.
- The **action** assigned in the **watch** command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



PART I

Shells and Scripting

- Bash, on page 17
- Guest Shell, on page 37
- Broadcom Shell, on page 71
- Barefoot Shell, on page 75
- Python API, on page 79
- Scripting with Tcl, on page 91
- iPXE, on page 97
- Kernel Stack, on page 101



CHAPTER 4

Bash

- About Bash, on page 17
- Guidelines and Limitations, on page 17
- Accessing Bash, on page 17
- Escalate Privileges to Root, on page 18
- Examples of Bash Commands, on page 21
- Managing RPMs, on page 22
- Support for DME Modularity, on page 24
- Persistently Daemonizing an SDK- or ISO-built Third Party Process, on page 32
- Persistently Starting Your Application from the Native Bash Shell, on page 33
- Copy Through Kstack, on page 33
- An Example Application in the Native Bash Shell, on page 34

About Bash

In addition to the Cisco NX-OS CLI, Cisco Nexus 3000 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 importing Cisco Python modules, do not use Python from the Bash shell. Instead use the more recent Python in NX-OS VSH.

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 bash
Linux# whoami
admin
Linux# pwd
/bootflash/home/admin
Linux#
```



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

The following is an example of the **run bash** <*command*> command.

```
run bash whoami
```

Escalate Privileges to Root

The privileges of an admin user can escalate their privileges for root access. Root access is required to pass configuration commands to the NX-OS VSH.

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 *
- 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 sudo prompt-password** command.

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

```
switch# run bash
Linux# sudo su root

We trust you have received the usual lecture from the local System
Administrator. It usually boils down to these three things:

    #1) Respect the privacy of others.
    #2) Think before you type.
    #3) With great power comes great responsibility.

Password:

Linux# whoami
root
Linux# exit
exit
```



Note A user account with network administrator privileges that is configured to login with shell-type Bash must still escalate privileges to root when passing configuration commands to VSH.

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

The example below demonstrates with network administrator user MyUser with a default shelltype 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                                                                    Ch #
--------------------------------------------------------------------------------
Eth1/2        --      eth  routed down    Administratively down      auto(D) --
```

The example below demonstrates with network administrator user MyUser with default shelltype 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                                                                    Ch #
--------------------------------------------------------------------------------
Eth1/2        --      eth  routed down    Administratively down      auto(D) --
```

Do not use **sudo su -** or the system will hang.

Examples of Bash Commands

This section contains examples of Bash commands and output.

Displaying System Statistics

The following example shows how to display system statistics:

```
switch# run bash
Linux# cat /proc/meminfo
MemTotal:      3795100 kB
MemFree:       1472680 kB
Buffers:           136 kB
Cached:        1100116 kB
ShmFS:         1100116 kB
Allowed:        948775 Pages
Free:           368170 Pages
Available:      371677 Pages
SwapCached:          0 kB
Active:        1198872 kB
Inactive:       789764 kB
SwapTotal:           0 kB
SwapFree:            0 kB
Dirty:               0 kB
Writeback:           0 kB
AnonPages:      888272 kB
Mapped:         144044 kB
Slab:           148836 kB
SReclaimable:    13892 kB
SUnreclaim:     134944 kB
PageTables:      28724 kB
NFS_Unstable:        0 kB
Bounce:              0 kB
WritebackTmp:        0 kB
CommitLimit:   1897548 kB
Committed_AS: 19984932 kB
VmallocTotal: 34359738367 kB
VmallocUsed:    215620 kB
VmallocChunk: 34359522555 kB
HugePages_Total:     0
HugePages_Free:      0
HugePages_Rsvd:      0
HugePages_Surp:      0
Hugepagesize:     2048 kB
DirectMap4k:     40960 kB
DirectMap2M:   4190208 kB
Linux#
```

Running Bash from CLI

The following example shows how to run a bash command from the CLI with the **run bash** <*command*> 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 -   497 select ?        00:00:08 init
5 S     0     2     0  0  75  -5 -     0 kthrea ?        00:00:00 kthreadd
1 S     0     3     2  0 -40   - -     0 migrat ?        00:00:00 migration/0
```

```
1 S     0     4     2  0  75  -5 -     0 ksofti ?        00:00:01 ksoftirqd/0
5 S     0     5     2  0  58   - -     0 watchd ?        00:00:00 watchdog/0
1 S     0     6     2  0 -40   - -     0 migrat ?        00:00:00 migration/1
1 S     0     7     2  0  75  -5 -     0 ksofti ?        00:00:00 ksoftirqd/1
5 S     0     8     2  0  58   - -     0 watchd ?        00:00:00 watchdog/1
1 S     0     9     2  0 -40   - -     0 migrat ?        00:00:00 migration/2
1 S     0    10     2  0  75  -5 -     0 ksofti ?        00:00:00 ksoftirqd/2
5 S     0    11     2  0  58   - -     0 watchd ?        00:00:00 watchdog/2
1 S     0    12     2  0 -40   - -     0 migrat ?        00:00:00 migration/3
1 S     0    13     2  0  75  -5 -     0 ksofti ?        00:00:00 ksoftirqd/3
5 S     0    14     2  0  58   - -     0 watchd ?        00:00:00 watchdog/3

...

4 S     0  8864     1  0  80   0 -  2249 wait   ttyS0    00:00:00 login
4 S  2002 28073  8864  0  80   0 - 69158 select ttyS0    00:00:00 vsh
4 R     0 28264  3782  0  80   0 - 54790 select ?        00:00:00 in.dcos-telnet
4 S     0 28265 28264  0  80   0 -  2247 wait   pts/0    00:00:00 login
4 S  2002 28266 28265  0  80   0 - 69175 wait   pts/0    00:00:00 vsh
1 S  2002 28413 28266  0  80   0 - 69175 wait   pts/0    00:00:00 vsh
0 R  2002 28414 28413  0  80   0 -   887 -      pts/0    00:00:00 ps
switch#
```

Managing RPMs

Installing RPMs from Bash

Procedure

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

Example

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

```
bash-4.2$ sudo yum 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$ sudo yum list available
bgp.lib32_n9000                         1.0.0-r0
bash-4.2$ sudo yum -y install bfd
```

Upgrading Feature RPMs

Before you begin

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

Procedure

	Command or Action	Purpose
Step 1	**sudo yum -y upgrade** *rpm*	Upgrades an installed RPM.

Example

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

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

Downgrading a Feature RPM

Procedure

	Command or Action	Purpose
Step 1	**sudo yum -y downgrade** *rpm*	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 yum -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/yum/protected.d/protected_pkgs.conf`.

Procedure

	Command or Action	Purpose
Step 1	**sudo yum -y erase** *rpm*	Erases the RPM.

Example

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

```
bash-4.2$ sudo yum -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.

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 **yum**. If you use **yum**, 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 **yum install**:

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

```
switch-1# yum 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#
```

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

  This option requires user intervention, as shown below.

  Example:

```
switch-1# yum 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 0:2.0.1.0-9.3.1 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 **yum 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 **yum list**, you must log in to the switch's Bash shell (**run bash**) before issuing the **yum** commands.

  yum list --patch-only installed | grep dme

  Example:

```
switch-1# yum 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 **yum**. 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 **yum**, you must run commands in Bash shell as root user (**run bash sudo su**):

- In Bash, run **yum 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# yum 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.

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

  Example:

```
bash-4.3# yum 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 **yum 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 **yum**, you must run commands in Bash shell as root user (**run bash sudo su**):

- In Bash, run **yum 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# yum 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.

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

  Example:

```
bash-4.3# yum 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.

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#

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
```

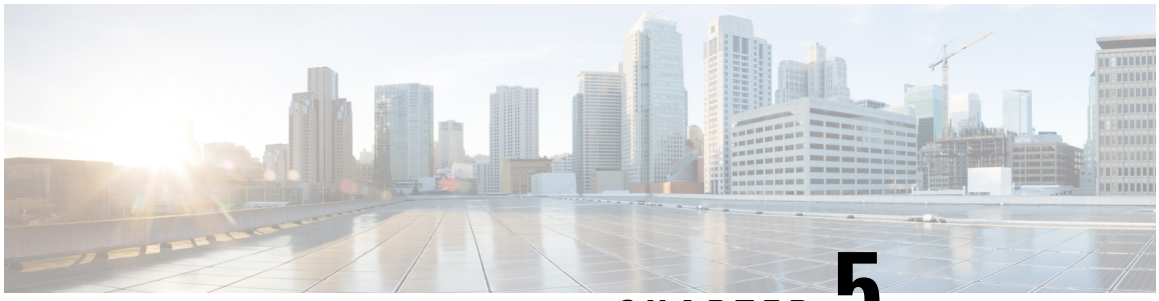
```
PIDFILE=/tmp/hello.pid

# 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 37
- Guidelines and Limitations, on page 38
- Accessing the Guest Shell, on page 42
- Resources Used for the Guest Shell, on page 43
- Capabilities in the Guestshell, on page 43
- Security Posture for Guest Shell, on page 51
- Guest File System Access Restrictions , on page 54
- Managing the Guest Shell, on page 54
- Verifying Virtual Service and Guest Shell Information, on page 66
- Persistently Starting Your Application From the Guest Shell, on page 67
- Procedure for Persistently Starting Your Application from the Guest Shell, on page 68
- An Example Application in the Guest Shell, on page 68
- Troubleshooting Guest Shell Issues, on page 69

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 35 MB of RAM and 350 MB of bootflash when enabled. Use the **guestshell destroy** command to reclaim resources if the Guest Shell is not used.

Guidelines and Limitations

Common Guidelines Across All Releases

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

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

- If you are running a third-party DHCPD server in Guest Shell, 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 Guest Shell 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 Guest Shell and get a Bash prompt or run a command within the context of the Guest Shell. The command uses password-less SSH to an available port on the localhost in the default network namespace.
- The `sshd` utility can secure the pre-configured SSH access into the Guest Shell 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 Guest Shell after Guest Shell 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 Guest Shell. Network-admin users may start another instance of sshd in the Guest Shell to allow remote access directly into the Guest Shell, but any user that logs into the Guest Shell is also given network-admin privilege.

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

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

Guest Shell 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 Guest Shell. Subsequent upgrades to the switch software will not automatically upgrade Guest Shell.
- Guest Shell releases increment the major number when distributions or distribution versions change.
- Guest Shell releases increment the minor number when CVEs have been addressed. The Guest Shell updates CVEs only when CentOS makes them publicly available.
- Cisco recommends using **yum 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 Guest Shell rootfs. Any customizations and software package installations would then need to be performed again within the context of this new Guest Shell rootfs.

Upgrading from Guest Shell 1.0 to Guest Shell 2.x

Guest Shell 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 Guest Shell 1.0, you must repeat the same deployment steps in Guest Shell 2.x. Your deployment script may need to be adjusted to account for the CentOS 7 differences.

Guest Shell 2.x

The Cisco NX-OS automatically installs and enables the Guest Shell by default on systems with sufficient resources. However, if the device is reloaded with a Cisco NX-OS image that does not provide Guest Shell support, the installer will automatically remove the existing Guest Shell and issue a `%VMAN-2-INVALID_PACKAGE`.

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

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: Guest shell 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 Guest Shell with the **guestshell destroy** command before reloading an older Cisco NX-OS image that does not support the Guest Shell.

Pre-Configured SSHD Service

The Guest Shell 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 Guest Shell 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 Guest Shell from the NX-OS CLI might not work.

The following steps instantiate an OpenSSh server within the Guest Shell 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 Guest Shell 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/systm/sshd-mgmt.service` and `/etc/ssh/sshd-mgmt_config`. The files should have the following configurations:

```
-rw-r--r-- 1 root root 394 Apr 7 14:21 /usr/lib/systemd/system/sshd-mgmt.service
-rw------- 1 root root 4478 Apr 7 14:22 /etc/ssh/sshd-mgmt_config
```

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/sshd
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 Guest Shell:

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

8. Save the configuration across multiple Guest Shell 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
```

10. 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/*
```

11. 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 Guest Shell 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 Guest Shell 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 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* | *rootfs*} command changes these limits.

Resource	Default	Minimum/Maximum
CPU	1%	1/6%
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 dnf 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 dnf 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 **vrfinfo**, 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 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 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]# yum 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]# yum 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]# yum install -y rh-python36 | tail
warning: /var/cache/yum/x86_64/7/centos-sclo-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 b1fb 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/23c926cd341ea6b7dd0b2a00aba99ae0f828be89d72b2190f27c11d4b7fb/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/cd551d81dbe15200be1cf41cd03869a46fe7226e7450af7a6545bfc474c9/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/a9/01ffebfb562e4274b6487b4bb1ddec7ca55ec7510b22e4c51f14098443b8/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/df50ac98ea0eb006c55a399c3bf1db9da7b5a24de7890bc9cfd5dd9e99/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/e8/74/6e4f91745020f967d09332bb2b8b9b10090957334692eb88ea4afe91b77f/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 yum install -y rh-python35 | tail
Dependency Installed:
  rh-python35-python.x86_64 0:3.5.1-13.el7
  rh-python35-python-devel.x86_64 0:3.5.1-13.el7
  rh-python35-python-libs.x86_64 0:3.5.1-13.el7
  rh-python35-python-pip.noarch 0:7.1.0-2.el7
  rh-python35-python-setuptools.noarch 0:18.0.1-2.el7
  rh-python35-python-virtualenv.noarch 0:13.1.2-2.el7
  rh-python35-runtime.x86_64 0:2.0-2.el7
```

```
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.

Installing RPMs in the Guest Shell

The /etc/yum.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.

Yum 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 2.x, go to the CentOS 7 repo at http://mirror.centos.org/centos/7/os/x86_64/Packages/.

Yum resolves the dependencies and installs all the required packages.

```
[guestshell@guestshell ~]$ sudo chvrf management yum -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 roofs** *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 Guest Shell

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.

ASLR and X-Space Support

Cisco 3000 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 the 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 the Guest Shell follow:

- 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

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 between Guest Shell and services on the host.

Guest File System Access Restrictions

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.

`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** [*guest shell OVA file* | *rootfs-file-URI*]}	• When *guest shell OVA file* 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 *rootfs-file-URI* 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.

Commands	Description
guestshell export rootfs package *destination-file-URI*	Exports a Guest Shell **rootfs** file to a local URI (bootflash, USB1, etc.).
guestshell disable	Shuts down and disables the Guest Shell.
guestshell upgrade {**package** [*guest shell OVA file* | *rootfs-file-URI*]}	• When *guest shell OVA file* is specified: 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. 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. You are prompted for a confirmation prior to carrying out the upgrade command. • When *rootfs-file-URI* is specified: 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.
guestshell reboot	Deactivates the Guest Shell and then reactivates it. You are prompted for a confirmation prior to carrying out the reboot command. **Note** This is the equivalent of a **guestshell disable** command followed by a **guestshell enable** command in exec mode. This is useful when processes inside the Guest Shell have been stopped and need to be restarted. The **run guestshell** command relies on `sshd` running in the Guest Shell. If the command does not work, the `sshd` 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.

Commands	Description
guestshell destroy	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. Issuing this command results in a prompt for a confirmation prior to carrying out the destroy command.
guestshell **run guestshell**	Connects to the Guest Shell that is already running with a shell prompt. No username/password is required.
guestshell run *command* **run guestshell** *command*	Executes a Linux/UNIX command within the context of the Guest Shell environment. After execution of the command you are returned to the switch prompt.
guestshell resize [**cpu** | **memory** | **rootfs**]	Changes the allotted resources available for the Guest Shell. The changes take effect the next time the Guest Shell is enabled or rebooted. **Note** Resize values are cleared when the **guestshell destroy** command is used.
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	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. You are prompted for a confirmation prior to initiating the reset.

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, LXCs 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        guestshe11.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 : 2.3(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
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#`	Displays the global state and limits for virtual services.
show virtual-service list `switch# show virtual-service list *` `Virtual Service List:` `Name Status Package Name` `---------------------------------------------------------------------` `guestshell+ Activated guestshell.ova`	Displays a summary of the virtual services, the status of the virtual services, and installed software packages.

Command	Description
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 : 2.2(0.2)` `      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`	Displays details about the guestshell package (such as version, signing resources, and devices).

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:

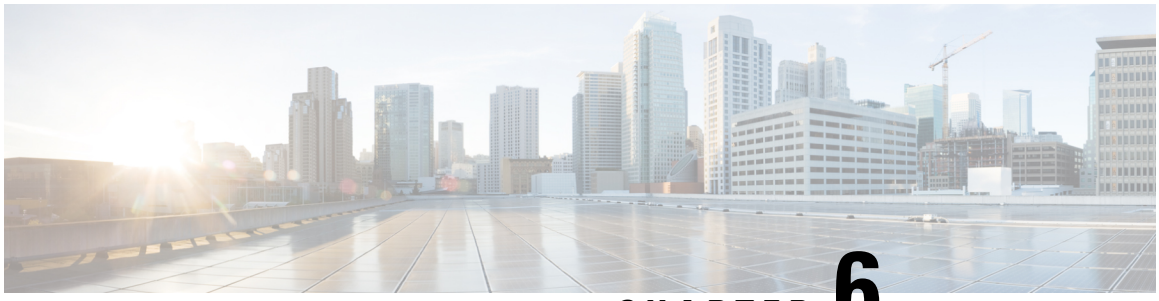
```
root@switch# ls -al /bootflash/try.that
-rw-r--r--  1  root  root  0  Apr 27  20:55  /bootflash/try.that
root@switch#
```

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 71
- Guidelines and Limitations, on page 71
- Accessing the Broadcom Shell (bcm-shell), on page 71

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

module_number	Module number in the chassis.

instance_number	T2 instance number • When not specified, the T2 instance number defaults to 0. • When a wildcard ('*') is specified, all T2 instances are processed.
command	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.

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.

- 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 ttyp0
```

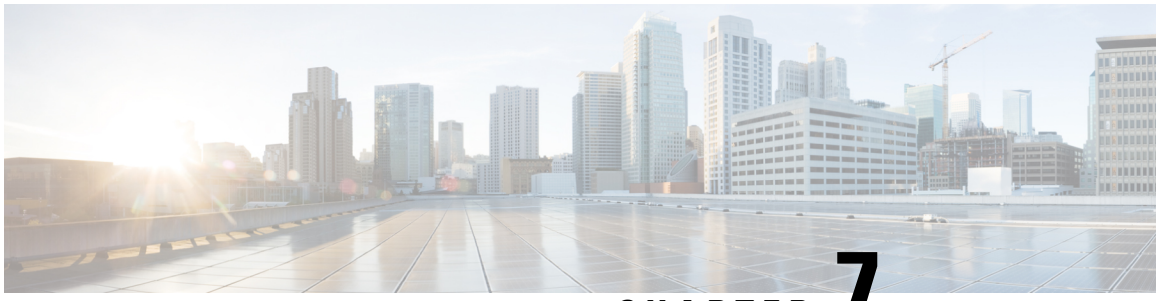
- Enter the command to gain root access to the line card software.

```
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.
```



CHAPTER 7

Barefoot Shell

This section contains the following topics:

- About the Barefoot Shell, on page 75
- Guidelines and Limitations, on page 75
- Accessing the Barefoot Shell with CLI API, on page 76

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
!!  - Subsitute 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.

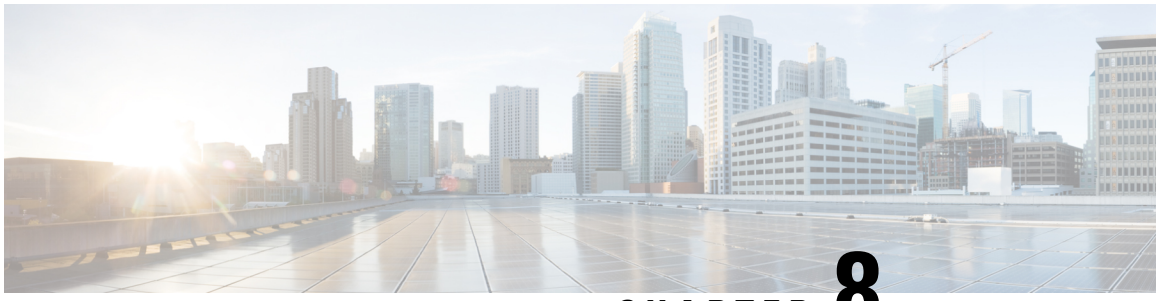
  New commands are available within each subsystem in the hierarchy.

```
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

- **About the Python API** , on page 79
- Using Python, on page 79

About the Python API

Beginning with Cisco NX-OS Release 9.3(5), Python 3 is now supported. Python 2.7 will continue to be supported. We recommend that you use the **python3** command for new scripts.

The Cisco Nexus 3000 Series switches support Python v2.7.11 and v3.7.3 in both interactive and noninteractive (script) modes and are available in the Guest Shell. For a list of switches that support Python scripting, see Supported Platforms for Programmability Features, on page 5.

Python is an easy to learn, powerful programming language. It has efficient high-level data structures and a simple but effective approach to object-oriented programming. Python's elegant syntax and dynamic typing, together with its interpreted nature, make it an ideal language for scripting and rapid application development in many areas on most platforms.

The Python interpreter and the extensive standard library are freely available in source or binary form for all major platforms from the Python website:

http://www.python.org/

The same site also contains distributions of and pointers to many free third-party Python modules, programs and tools, and more documentation.

The Python scripting capability gives programmatic access to the device's command-line interface (CLI) to perform various tasks and Power On Auto Provisioning (POAP) or Embedded Event Manager (EEM) actions. Python can be accessed from the Bash shell.

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

Using Python

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.

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: `string = cli ("cli-command")`	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: `json_string = clid ("cli-command")`	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.

API	Description
clip() Example: `clip ("cli-command")`	Prints the output of the CLI command directly to stdout and returns nothing to Python. **Note** `clip ("cli-command")` is equivalent to `r=cli("cli-command")` `print r`

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.

```
switch# python
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.
>>>
>>> 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
>>>
```

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:

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 Cisco Nexus 3000 Series switches 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 %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 3000 Series 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 3000 Series 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 Schedular

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.

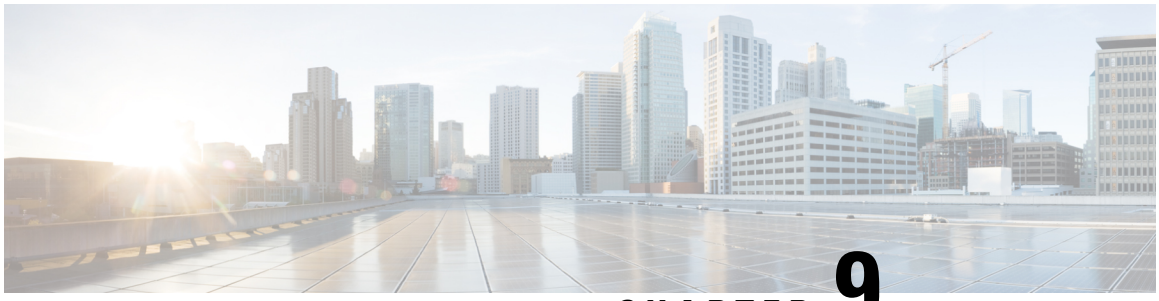
```
#!/bin/env python3
from cli import *
from nxos import *
import os

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#
```



CHAPTER 9

Scripting with Tcl

This chapter contains the following topics:

- About Tcl, on page 91
- Running the Tclsh Command, on page 94
- Navigating Cisco NX-OS Modes from the Tclsh Command, on page 95
- Tcl References, on page 96

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:

- Tcl is supported on Cisco Nexus switches.
- 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.

Procedure

	Command or Action	Purpose
Step 1	**tclsh** [**bootflash:***filename* [*argument* **...**]] **Example:** `switch# tclsh ?` `<CR>` `bootflash: The file to run`	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.

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
```

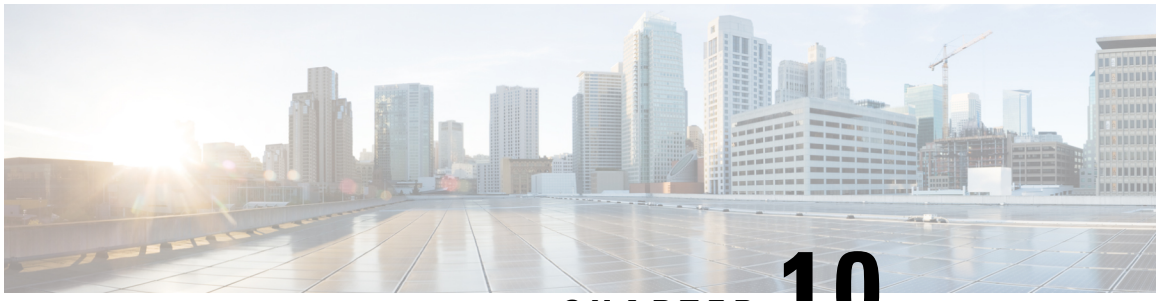
```
  this          Shows info about current object (mode's instance)
  vrf           Configure VRF parameters
  end           Go to exec mode
  exit          Exit from command interpreter
  pop           Pop mode from stack or restore from name
  push          Push current mode to stack or save it under 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

This chapter contains the following sections:

- About iPXE, on page 97
- Netboot Requirements, on page 98
- Guidelines and Limitations, on page 98
- Boot Mode Configuration, on page 98
- Verifying the Boot Order Configuration, on page 100

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

PXE has the following configuration guidelines and limitations:

- While auto-booting 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 UDP based and it can be problematic if packet loss starts appearing. TCP is a window-based protocol and handles bandwidth sharing/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/boots Cisco signed NBI images. Other standard image format support is disabled for security reasons.

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**]

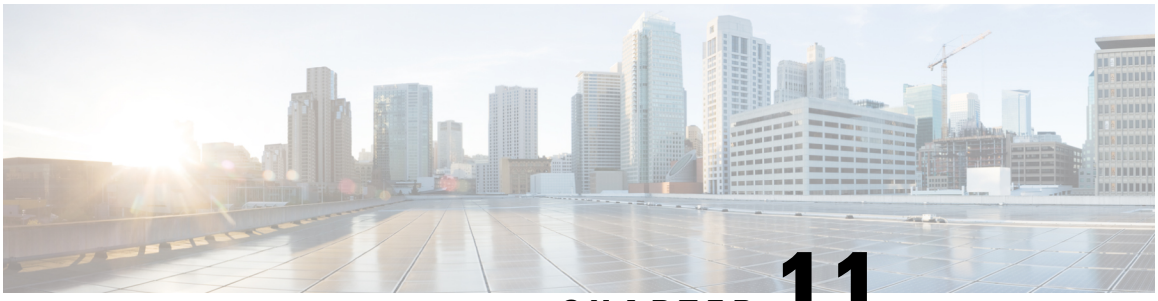
Keyword	Function
– – grub	Grub only
– – pxe	PXE only
– – pxe2grub	PXE first, followed by Grub if PXE failed
– – grub2pxe	Grub first, followed by PXE if Grub failed

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
show boot order	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

This chapter contains the following sections:

- About Kernel Stack, on page 101
- Guidelines and Limitations, on page 101
- Changing the Port Range, on page 102

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.

Guidelines and Limitations

Using the Kernel Stack has the following guidelines and limitations:

- Guest Shell, other open containers, and the host Bash Shell use Kernel Stack (kstack).
- Open containers start in the host default namespace
  - Other network namespaces might 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 PIDs and identify options for the **ip netns** command do not work without modification because of the file system device check. A **vrfinfo** utility is provided to give the network administrator the same information.
- Open containers may read the interface state from `/proc/net/dev` or use other normal Linux utilities such as **netstat** or **ifconfig** without modification. This provides counters for packets that have initiated / terminated on the switch.
- Open containers may use **ethtool –S** to get extended statistics from the net devices. This includes packets switched through the interface.

- Open containers may run packet capture applications like **tcpdump** to capture packets initiated from or terminated on the switch.
- There is no support for networking state changes (interface creation/deletion, IP address configuration, MTU change, etc.) from the Open containers
- IPv4 and IPv6 are supported
- Raw PF_PACKET is supported
- Well-known ports (0-15000) may only be used by one stack (Netstack or kstack) at a time, regardless of the network namespace.
- There is no IP connectivity between Netstack and kstack applications. This is a host limitation which also applies to open containers.
- Open containers are not allowed to send packets directly over an Ethernet out-of-band channel (EOBC) interface to communicate with the linecards or standby Sup.
- From within an open container, direct access to the EOBC interface used for internal communication with linecards or the standby supervisor. The host bash shell should be used if this access is needed.
- 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 3000 Series NX-OS VXLAN Configuration Guide, Release 9.3(x).

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.

Procedure

	Command or Action	Purpose
Step 1	[**no**] **sockets local-port-range** *start-port end-port*	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:

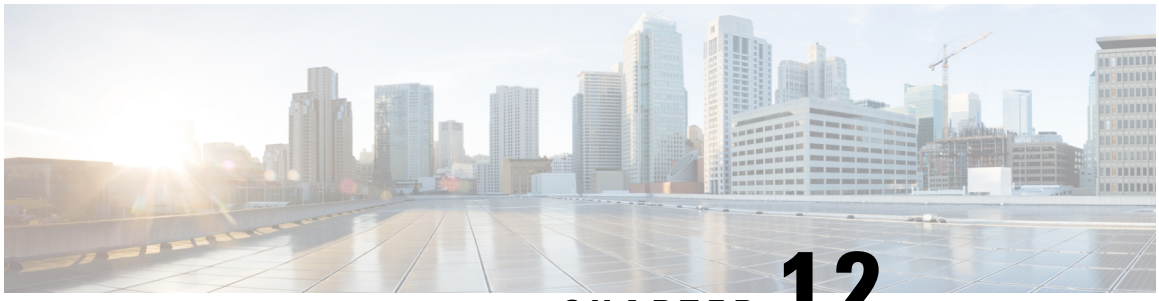
- Reload the switch after entering the command.
- Leave a minimum of 7000 ports unallocated which are used by Netstack.
- Specify the *start-port* as 15001 or the *end-port* as 65535 to avoid holes in the port range.



PART II

Applications

- Third-Party Applications, on page 107
- Ansible, on page 123
- SaltStack, on page 125
- Puppet Agent, on page 129
- Using Chef Client with Cisco NX-OS, on page 131
- Nexus Application Development - Yocto, on page 135
- Nexus Application Development - SDK, on page 139
- NX-SDK, on page 147
- Using Docker with Cisco NX-OS, on page 153



CHAPTER 12

Third-Party Applications

This chapter contains the following sections:

- About Third-Party Applications, on page 107
- Installing Third-Party Native RPMs/Packages, on page 107
- Installing Signed RPM, on page 109
- Persistent Third-Party RPMs, on page 114
- Installing RPM from VSH, on page 115
- Third-Party Applications, on page 119

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-I2-1/x86_64. These applications are installed in the native host by using the **yum** command in the Bash shell or through the NX-OS CLI.

When you enter the **yum install** *rpm* command, a Cisco **YUM** plugin gets executed. This plugin copies the RPM to a hidden location. On switch reload, the system re-installs the RPM.

For configurations in `/etc`, a Linux process, **incrond**, monitors artifacts created in the directory and copies them to a hidden location, which gets copied back to `/etc`.

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/yum/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/
enabled=1
gpgcheck=0
sslverify=0
```

An example of installation of an rpm using *yum*, with full install log.

Example:

```
bash-4.2# yum 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# yum 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** will throw 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 below 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

Setup 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 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 **yum** repository is dedicated to persistent third-party RPMs. The `/etc/yum/repos.d/thirdparty.repo` points to `/bootflash/.rpmstore/thirdparty`.
- Whenever you enter the **yum 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 created in the `/etc` directory persists across reloads. For example, a third-party service script created under `/etc/init.d/` brings up the apps during 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 186.

Installing RPM from VSH

Package Addition

NX-OS feature RPMs can also be installed by using the VSH CLIs.

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.el5.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 116.

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.el5.x86_64.rpm
activate
```

Package Activation

Before you begin

The RPM has to have been previously added.

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** *rpm-packagename*	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.el5 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.el5.x86_64.rpm
[####################] 100%
Install operation completed successfully at Thu Aug  6 12:46:53 2015
```

Deactivating Packages

Procedure

	Command or Action	Purpose
Step 1	**install deactivate** *package-name*	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.

Procedure

	Command or Action	Purpose
Step 1	**install remove** *package-name*	Removes the RPM package.

Example

The following example shows how to remove the Chef RPM package:

```
switch# install remove chef-12.0-1.el5.x86_64.rpm
```

Displaying Installed Packages

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

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

Procedure

	Command or Action	Purpose
Step 1	install add *package-name* activate upgrade	Upgrade a package.

Example

The following example show 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

Procedure

	Command or Action	Purpose
Step 1	install add *package-name* 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 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 Releases
- Ansible NX-OS Sample Modules, NX-OS Ansible
- Puppet, Puppet Forge
- Cisco Puppet Module(Git), Cisco Network Puppet Module
- Chef, Cisco Chef Cookbook
- Cisco Chef Cookbook(Git), Cisco Chef Cookbook(Git)

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 (i.e. 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 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 very 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 prototocol (LLDP) is an industry standard protocol designed to supplant proprietary link layer protocols such as EDP or CDP. The goal of LLDP is to provide an inter-vendor 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 network services (through ICMP, SNMP, SSH, FTP, HTTP etc), host resources (CPU load, disk usage, system logs, etc.), and alert services for servers, switches, applications, and services through the Nagios remote plugin executor (NRPE) and through SSH or SSL tunnels.

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 effectively 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 a search, visualization and pre-packaged 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 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,
- De-duplicates repeated values, and
- 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 out a description of the contents of packets on a network interface that match the boolean expression; the description is preceded by a time stamp, printed, by default, as hours, minutes, seconds, and fractions of a second since midnight. It can also be run with the -w flag, which causes it to save the packet data to a file for later analysis, and/or with the -r flag, which causes it to read from a saved packet file rather than to read packets from a network interface. It can also be run with the -V flag, which causes it to read a list of saved packet files. In all cases, only packets that match expression will be processed by tcpdump.

For more information, see http://www.tcpdump.org/manpages/tcpdump.1.html.

Tshark

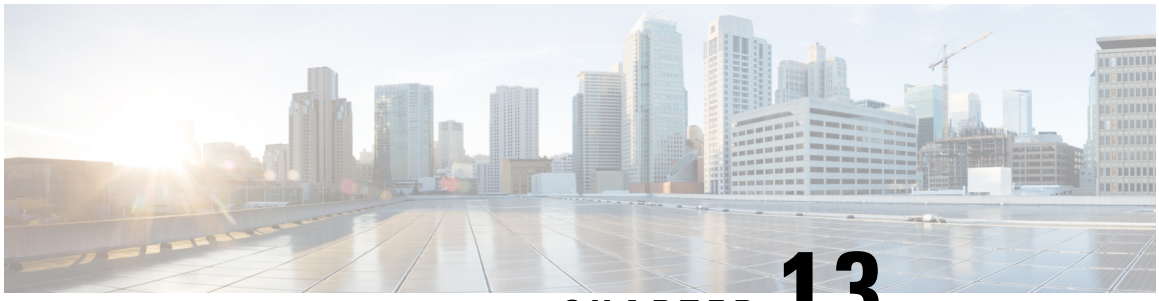
TShark is a network protocol analyzer on the CLI. It lets you capture packet data from a live network, or read packets from a previously saved capture file, You can either print a decoded form of those packets to the standard output or write the packets to a file. TShark's native capture file format is the pcap format, which is also the format used by **tcpdump** and various other tools. Tshark can be used within the Guest Shell 2.1 after removing the cap_net_admin file capability.

```
setcap
 cap_net_raw=ep /sbin/dumpcap
```



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

Ansible

- Prerequisites, on page 123
- About Ansible, on page 123
- Cisco Ansible Module, on page 123

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:

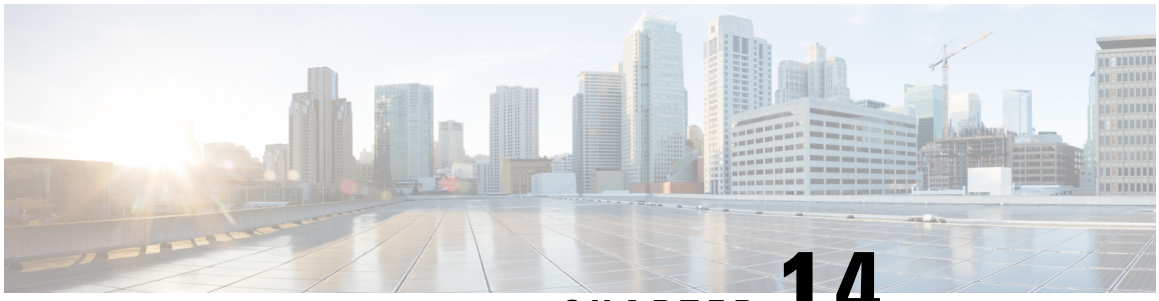
Ansible	https://www.ansible.com/
Ansible Automation Solutions. Includes installation instructions, playbook instructions and examples, module lists, and so on.	https://docs.ansible.com/

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

SaltStack

This chapter contains the following topics:

- About SaltStack, on page 125
- Guidelines and Limitations, on page 126
- Cisco NX-OS Environment for SaltStack, on page 126
- Enabling NX-API for SaltStack, on page 127
- Installing SaltStack for NX-OS, on page 127

About SaltStack

The Cisco Nexus switches support SaltStack through NX-OS. For information about Cisco NX-OS releases that support SaltStack, see https://github.com/saltstack/salt/blob/develop/doc/topics/installation/nxos.rst#step-1-verify-platform-and-software-version-support.

SaltStack is a free and open source automation framework for configuration, management, and remote execution of servers and other network devices. The SaltStack framework consists of a server that is called the Salt primary, and Salt nodes that run client programs, called minions. The Cisco Nexus switch (switch) is a Salt node, not the Salt primary.

SaltStack minions can run either on-box or off-box, respective to the switch, to execute the configuration or management operations:

- On-box, the minions run in the switch's Bash shell. These native minions receive and execute remote commands from the primary, and relay the command's results to the primary. In an on-box deployment, the minions are enabled in the switch's Guest shell.
- Off-box, a different type of minion, a proxy minion, runs over an SSH connection to the switch or through the NX-API. The proxy minion, either the SSH proxy minion or the NX-API proxy minion, receives and executes the commands. The proxy then relays the command's results to the primary.

Keys are used to ensure security between the Salt primary and the minions running on the Cisco Nexus switch. When the Salt primary initiates its connection with a minion running on the Cisco Nexus switch, it first passes a key. The minion receives the key, then computes the correct response, and transmits the key back to the primary. The primary also has computed the correct response value for the key. When the primary receives the key from the minion, if the keys match, the session is open. The Salt primary can then send commands. Sessions are not persistent across power cycles or reboots.

SaltStack manages and configures the switch through execution modules and salt states, which affect the switch's CLI, properties, and features. For example, through the modules, SaltStack can be used to upgrade

the Cisco Nexus switches. The Salt primary sends commands programmatically to leverage automation and scalability.

For more information, consult the following documentation:

SaltStack	https://www.saltstack.com/
SaltStack Documentation	https://docs.saltstack.com/en/latest/
Cisco Nexus Salt Minion Installation and Configuration Guide	https://github.com/saltstack/salt/blob/develop/doc/topics/installation/nxos.rst

About NX-OS and SaltStack

Salt Open is the open source, community edition of the Salt configuration management and distributed remote execution system. Cisco NX-OS provides an intermediate layer between the physical switch and the Salt Open software. Cisco NX-OS and Salt Open interoperate to provide the API and command-execution layer between Salt minions and Cisco Nexus switches. Cisco NX-OS hosts the minions and enables them to run as follows:

- On the switch, the Cisco NX-OS guest shell hosts SaltStack minions and provides automated orchestration of one or more switches through a unified interface. The minion running in the guest shell is a native minion and it connects over the NX-API the UNIX Domain Socket (UDS).
- Off the switch, the Salt primary runs the Salt Open software on a network device and communicates with NX-OS through SSH (the SSH proxy minion) or NX-API over HTTPS (the NX-API proxy minion). Cisco NX-OS interprets the commands, performs required configuration tasks, and reports success or failure back to the appropriate proxy minion. The proxy minion, in turn, transmits this data back to the Salt primary.

Guidelines and Limitations

The following are the guidelines and limitations for implementing SaltStack on Cisco Nexus 3000 Series switches:

- For notes about platform support, see Supported Platforms for Programmability Features, on page 5.
- If you are running SaltStack over SSH or NX-API HTTPS, enable the NX-API feature (**feature nxapi**) before you run Salt.
- The Salt primary listens for minions on port 4506. Make sure that this port is open (unblocked) and not used by another service.

Cisco NX-OS Environment for SaltStack

The Cisco NX-OS environment is different depending on whether you are running Salt on box or off box.

- For on-box management of the switch, you must install the SaltStack minion RPM in the Guest Shell, which is the hosting environment for the minion.
- For off-box management of the switch, SSH or NX-API must be enabled in NX-OS.

For more information, such as which Cisco Nexus switches support SaltStack, go to https://github.com/saltstack/salt/blob/develop/doc/topics/installation/nxos.rst#step-1-verify-platform-and-software-version-support.

Enabling NX-API for SaltStack

Before you begin

For proxy minions running over SSH or NX-API HTTPS, the NX-API feature must be enabled for SaltStack to function. By default, NX-API is enabled. The following instructions are provided in case you need to reenable it.

Procedure

	Command or Action	Purpose
Step 1	**config terminal** **Example:** `switch-1# config terminal` `Enter configuration commands, one per line. End with CNTL/Z.` `switch-1(config)#`	Enters configuration mode.
Step 2	**feature nxapi** **Example:** `switch-1# feature nxapi` `switch-1#(config)#`	Enables NX-API for proxy minions.

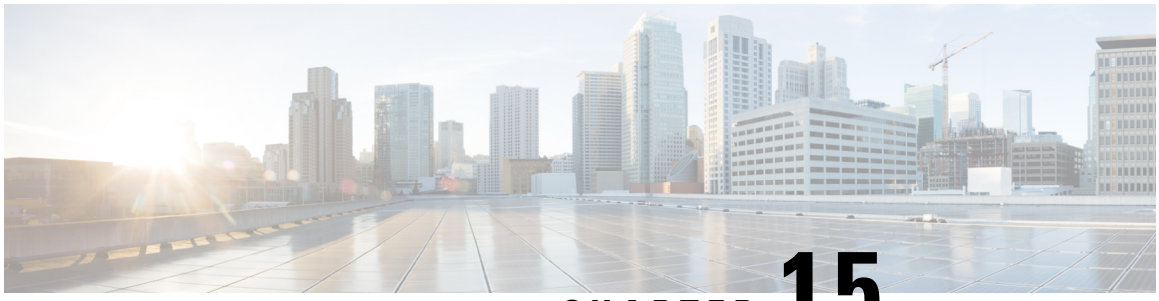
What to do next

Install SaltStack.

Installing SaltStack for NX-OS

Use the following installation guide to install and bring up SaltStack on the Cisco Nexus switches:

https://github.com/saltstack/salt/blob/develop/doc/topics/installation/nxos.rst#cisco-nexus-salt-minion-installation-and-configuration-guide



CHAPTER 15

Puppet Agent

This chapter includes the following sections:

- About Puppet, on page 129
- Prerequisites, on page 129
- Puppet Agent NX-OS Environment, on page 130
- ciscopuppet Module, on page 130

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/
Puppet Labs Documentation	https://puppet.com/docs

Prerequisites

The following are prerequisites for the Puppet Agent:

- For information about supported platforms, see the Nexus Switch Platform Matrix.
- You must have the required disk storage available on the device for virtual services installation and deployment of Puppet Agent.
  - A minimum of 450MB free disk space on bootflash.
- You must have 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

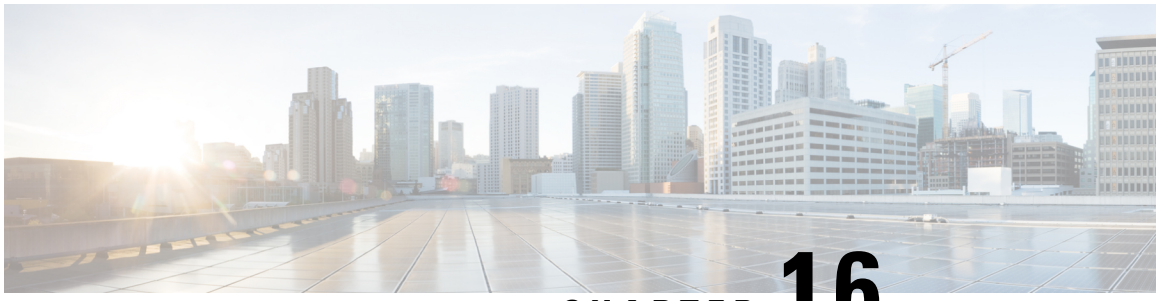
ciscopuppet Module

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
Puppet NX-OS Manifest Examples	Cisco Network Puppet Module Examples
NX-OS developer landing page.	Configuration Management Tools



CHAPTER 16

Using Chef Client with Cisco NX-OS

This chapter includes the following sections:

- About Chef, on page 131
- Prerequisites, on page 131
- Chef Client NX-OS Environment, on page 132
- cisco-cookbook, on page 132

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:

- For information about supported platforms, see Supported Platforms for Programmability Features, on page 5.
- 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: 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-cookbook

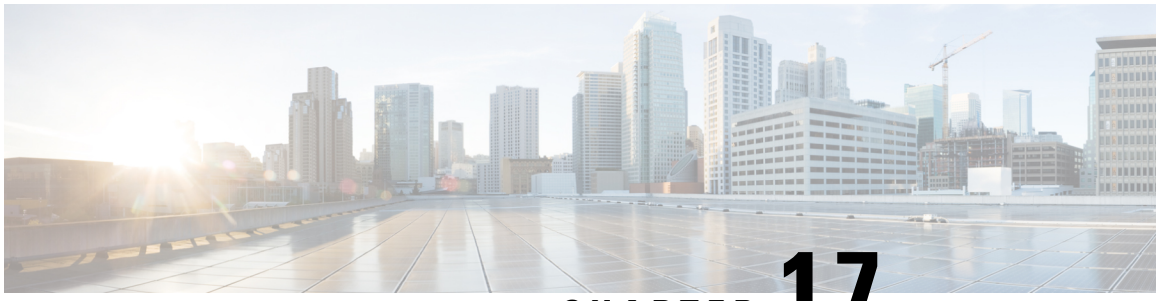
cisco-cookbook is a Cisco-developed open-source interface between the abstract resources configuration in a Chef recipe and the specific implementation details of the switch. This cookbook is installed on the Chef Server and is required for proper Chef Client operation on switches.

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
Chef Supermarket	Chef Supermarket

Topic	Link
Chef NX-OS Manifest Examples	Cisco Network Chef Cookbook Recipes



CHAPTER 17

Nexus Application Development - Yocto

This chapter contains the following sections:

- About Yocto, on page 135
- Installing Yocto, on page 135

About Yocto

The Cisco NX-OS Release 9.2(1) software is based on Yocto 2.0. More applications can be installed by downloading Yocto 2.0, 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 libsdl1.2-dev xterm
```

Step 2 Download Yocto 2.0.

```
wget http://downloads.yoctoproject.org/releases/yocto/yocto-2.0/poky-jethro-14.0.0.tar.bz2
tar -xjfv poky-jethro-14.0.0.tar.bz2
cd poky-jethro-14.0.0
```

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# yum 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# yum 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

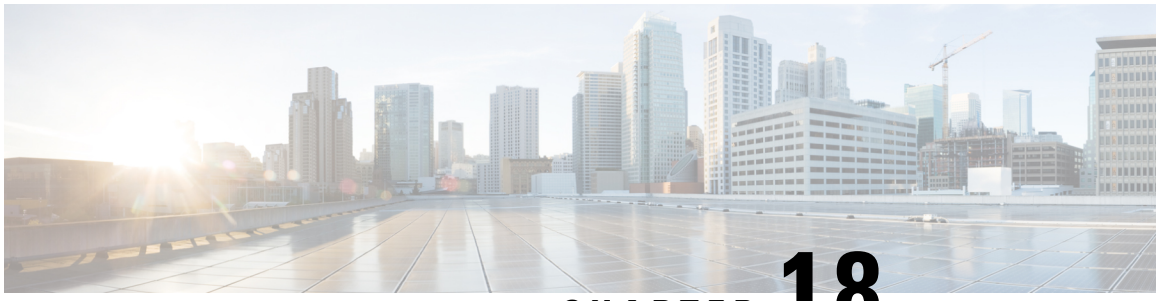
Transaction Summary
================================================================================================
Install       1 Package

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 18

Nexus Application Development - SDK

This chapter contains the following sections:

- About the Cisco SDK, on page 139
- Installing the SDK, on page 139
- Procedure for Installation and Environment Initialization, on page 140
- Using the SDK to Build Applications, on page 141
- Using RPM to Package an Application, on page 142
- Creating an RPM Build Environment, on page 143
- Using General RPM Build Procedure, on page 143
- Example to Build RPM for collectd with No Optional Plug-Ins, on page 144
- Example to Build RPM for collectd with Optional Curl Plug-In, on page 145

About the Cisco SDK

The Cisco SDK is a development kit based on Yocto 1.2. It contains all of the tools needed to build applications for execution on a switch running the Cisco NX-OS Release 7.0(3)I2(1) and later. 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 it might be required that the you download and build any dependencies needed for any particular application. Note that 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 yum 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_64-wrs-linux** command to add the SDK-specific paths to your shell environment. This must be done for each shell you intend to use with the SDK. This is the key to setting up the SDK in order to use the correct versions of the build tools and libraries.

Procedure

Step 1 Browse to the installation directory.

Step 2 Enter the following command at the Bash prompt:

```
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 memcachec 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 plugins.

```
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*

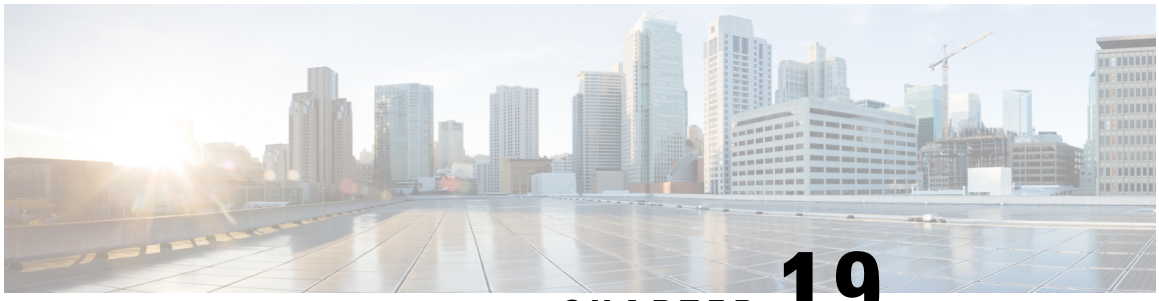
```
bash$ rpmbuild_opts=()
bash$ for rmdep in \
> amqp apache ascent bind curl_xml dbi ipmi java memcachec mysql nginx \
> notify_desktop notify_email nut openldap perl pinba ping postgresql python \
> rrdtool sensors snmp varnish virt write_http write_riemann
> do
>  rpmbuild_opts+=("--without")
>  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 19

NX-SDK

This chapter contains the following topics:

- About the NX-SDK, on page 147
- About On-Box (Local) Applications, on page 148
- Default Docker Images, on page 148
- Guidelines and Limitations for NX-SDK, on page 149
- Off-Box (Remote) Applications, on page 149

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 (with some exceptions) 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 of the switch in the Docker containers that the 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 148
- 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 150.

Related Information

For more information about Cisco NX-SDK, go to:

- Cisco DevNet NX-SDK. Click the `versions.md` link (Cisco DevNet NX-SDK Versions) to get information about features and details on each supported release.
- NX-SDK Readmes

As needed, Cisco will add 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 148.

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.

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 .

Off-Box (Remote) Applications

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.

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	• 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.

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.

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 151.

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.

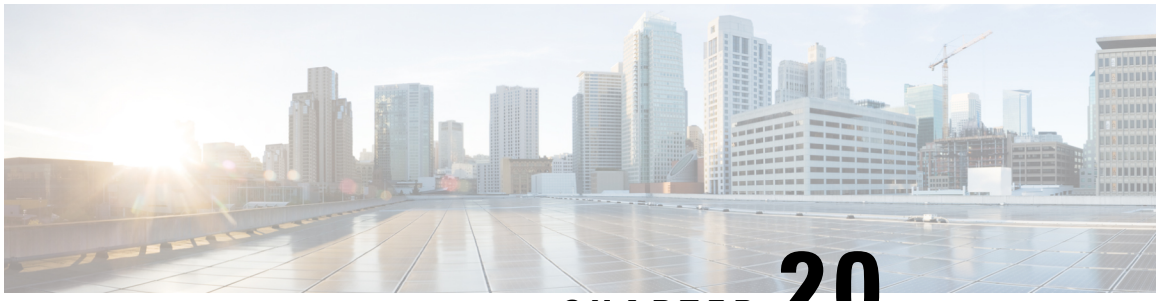
The ability of an app to read or write to the switch can be controlled through a security profile. A security profile is an optional object that is attached to the applications' service running in the switch. Security profiles control an application's ability to write to the switch, and in turn, control 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 20

Using Docker with Cisco NX-OS

This chapter contains the following topics:

- About Docker with Cisco NX-OS, on page 153
- Prerequisites for Setting Up Docker Containers Within Cisco NX-OS, on page 153
- Starting the Docker Daemon, on page 154
- Configure Docker to Start Automatically, on page 155
- Starting Docker Containers: Host Networking Model, on page 155
- Starting Docker Containers: Bridged Networking Model, on page 157
- Mounting the bootflash and volatile Partitions in the Docker Container, on page 158
- Enabling Docker Daemon Persistence on Enhanced ISSU Switchover, on page 158
- Resizing the Docker Storage Backend, on page 159
- Stopping the Docker Daemon, on page 161
- Docker Container Security, on page 162
- Docker Troubleshooting, on page 164

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 1.13.1. 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.

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
2000000000 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 162 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 18887
```

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.

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/userns-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 `--userns-remap=default` option to the `other_args` field in the `/etc/sysconfig/docker` file.

For example:

```
other_args="-debug=true --userns-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/userns-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
/ #
```

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 yum, apk, or something similar.

General Troubleshooting Tips

Problem: You have other issues with Docker containers that were not resolved using other troubleshooting processes.

Solution:

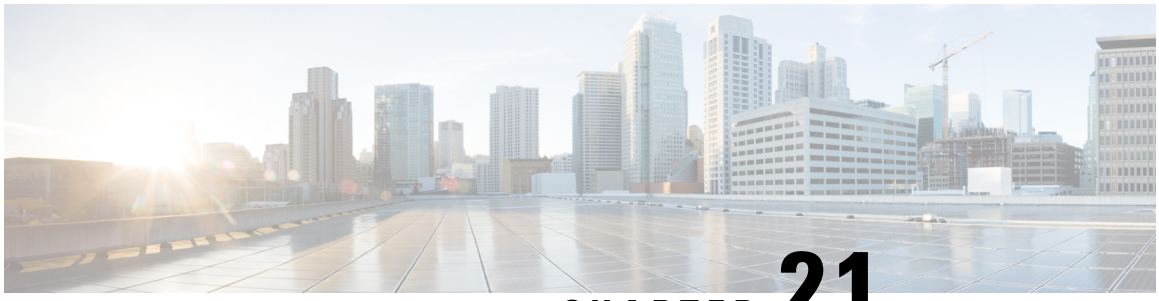
- Look for dockerd debug output in `/var/log/docker` for any clues as to what is wrong.
- Verify that your switch has 8 GB or more of RAM. Docker functionality is not supported on any switch that has less than 8 GB of RAM.



PART III

NX-API

- NX-API CLI, on page 171
- NX-API REST, on page 199
- NX-API Developer Sandbox, on page 205



CHAPTER 21

NX-API CLI

- About NX-API CLI, on page 171
- Using NX-API CLI, on page 172
- XML and JSON Supported Commands, on page 189

About NX-API CLI

On 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 switch. NX-API CLI supports **show** commands, configurations, and Linux Bash.

NX-API CLI supports JSON-RPC.

The NX-API CLI also supports JSON/CLI Execution in Cisco Nexus switches.

Transport

NX-API uses HTTP/HTTPS as its transport. CLIs are encoded into the HTTP/HTTPS POST body.

The NX-API 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. If the Nginx memory usage exceeds the cgroup limitations, the Nginx process is restarted and restored.

Note For the 7.x release, the Nginx process continues to run even after NX-API is disabled using the “no feature NXAPI” command. This is required for other management-related processes. In the 6.x release, all processes were killed when you ran the “no feature NXAPI” command, so this is a change in behavior in the 7.x release.

Message Format

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 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 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 187.

You must enable NX-API with the **feature** manager CLI command on the device. By default, NX-API is disabled.

The following example shows how to configure and launch the NX-API CLI:

- Enable the management interface.

```
switch# conf t
switch(config)# interface mgmt 0
switch(config)# ip address 192.0.20.123/24
switch(config)# vrf context managment
switch(config)# ip route 10.0.113.1/0 1.2.3.1
```

- 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"
            }
        }
    }
}
```

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 6: NX-API Management Commands

NX-API Management Command	Description
feature nxapi	Enables NX-API.
no feature nxapi	Disables NX-API.
nxapi {**http** | **https**} **port** *port*	Specifies a port.
no nxapi {**http** | **https**}	Disables HTTP/HTTPS.
show nxapi	Displays port and certificate information. **Note** The "**show nxapi**" command doesn't display certificate/config information for network-operator role.

NX-API Management Command	Description
nxapi certificate {httpscrt certfile | **httpskey keyfile**} *filename*	Specifies the upload of the following: • HTTPS certificate when httpscrt is specified. • HTTPS key when httpskey is specified. Example of HTTPS certificate: `nxapi certificate httpscrt certfile bootflash:cert.crt` Example of HTTPS key: `nxapi certificate httpskey keyfile bootflash:privkey.key`
nxapi certificatehttpskey keyfile *filename* **password** *passphrase*	Installs NX-API certificates with encrypted private keys: **Note** The passphrase for decrypting the encrypted private key is **pass123!**. Example: `nxapi certificate httpskey keyfile bootflash:encr-cc.pem password pass123!`
nxapi certificate enable	Enables a certificate.

NX-API Management Command	Description
nxapi certificate sudi	This CLI provides a secure way of authenticating to the device by using Secure Unique Device Identifier (SUDI). The SUDI based authentication in nginx will be used by the CISCO SUDI compliant controllers. 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. **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. • "nxapi certificate sudi" will overwrite the custom certificate/key if configured, and there is no way to get the custom certificate/key back. • "nxapi certificate sudi" and "nxapi certificate trustpoint" and "nxapi certificate enable" are mutually exclusive , and configuring one will delete the other configuration. • NX-API do not support SUDI certificate-based client certificate authentication. If client certificate authentication is needed, then Identity certificate need to be used. • 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.
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).
nxapi ssl-protocols {TLSv1.0 TLSv1.1 TLSv1.2}	Beginning with Cisco NX-OS Release 9.2(1), TLS1.0 is disabled by default. Running this command enables the TLS versions specified in the string, including the TLS1.0 that was disabled by default, if necessary. The **no** form of the command changes it to the default (by default, only TLS1.1 and TLS1.2 will be enabled).
nxapi use-vrf *vrf*	Specifies the default VRF, management VRF, or named VRF.

NX-API Management Command	Description
ip netns exec management iptables	Implements any access restrictions and can be run in management VRF. **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. *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.* **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 186.
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>**

Following is an example of a successful upload of an HTTPS certificate:

```
switch(config)# nxapi certificate httpscrt 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 httpscrt 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 httpscrt 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 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:

Table 7: 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
type	Specifies the type of command to be executed. The following types of commands are supported: • **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. **Note** • 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.

<table>
<tr><th>NX-API Request Element</th><th>Description</th></tr>
<tr><td>chunk</td><td>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.

Enable or disable chunk with the following settings:
<table><tr><td>0</td><td>Do not chunk output.</td></tr><tr><td>1</td><td>Chunk output.</td></tr></table>Note
• 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 < > is converted into &gt).
You can use XML SAX to parse the chunked output.
• When the output message format is JSON, the chunks are concatenated to create a valid JSON object.

Note
When chunking is enabled, the maximum message size supported is currently 200MB of chunked output.</td></tr>
<tr><td>rollback</td><td>Valid only for configuration CLIs, not for show commands. Specifies the configuration rollback options. Specify one of the following options.
• 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.

Note
The rollback element is available in the cli_conf mode when the input request format is XML or JSON.</td></tr>
</table>

<table>
<tr><th>NX-API Request Element</th><th>Description</th></tr>
<tr><td>sid</td><td>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 sid to match the sid of the previous response message.

NX-OS release 9.3(1) introduces the sid option <code>clear</code>. When a new chunk request is initiated with the sid set to <code>clear</code>, all current chunk requests are discarded or abandoned.

When you receive response code <code>429: Max number of concurrent chunk request is 2</code>, use sid <code>clear</code> to abandon the current chunk requests. After using sid <code>clear</code>, subsequent response codes operate as usual per the rest of the request.</td></tr>
<tr><td>input</td><td>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 cli_show message type and are not supported in cli_conf mode.

Note
Except for bash, multiple commands are separated with " ; ". (The ; must be surrounded with single blank characters.)

For bash, multiple commands are separated with ";". (The ; is not surrounded with single blank characters.)

The following are examples of multiple commands:

<table><tr><td>cli_show</td><td><code>show version ; show interface brief ; show vlan</code></td></tr><tr><td>cli_conf</td><td><code>interface Eth4/1 ; no shut ; switchport</code></td></tr><tr><td>bash</td><td><code>cd /bootflash;mkdir new_dir</code></td></tr></table></td></tr>
<tr><td>output_format</td><td>The available output message formats are the following:

<table><tr><td>xml</td><td>Specifies output in XML format.</td></tr><tr><td>json</td><td>Specifies output in JSON format.</td></tr></table></td></tr>
</table>

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 8: NX-API Request Elements for JSON-RPC Format

NX-API Request Element	Description
jsonrpc	A string specifying the version of the JSON-RPC protocol. Version must be 2.0.

NX-API Request Element	Description
method	A string containing the name of the method to be invoked. NX-API supports either: • **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, [ ].
params	A structured value that holds the parameter values used during the invocation of a method. It must contain the following: • **cmd**–CLI command • **version**–NX-API request version identifier
rollback	Valid only for configuration CLIs, not for show commands. Configuration rollback options. You can specify one of the following options. • 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.
id	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, *id* : 1

NX-API Response Elements

The NX-API elements that respond to a CLI command are listed in the following table:

Table 9: NX-API Response Elements

NX-API Response Element	Description
version	NX-API version.
type	Type of command to be executed.

NX-API Response Element	Description
sid	Session ID of the response. This element is valid only when the response message is chunked.
outputs	Tag that encloses all command outputs. When multiple commands are in cli_show or cli_show_ascii, each command output is enclosed by a single output tag. When the message type is cli_conf or bash, there is a single output tag for all the commands because cli_conf and bash commands require context.
output	Tag that encloses the output of a single command output. For cli_conf and bash 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 184.

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
DROP       tcp  --  anywhere             anywhere             tcp dpt:http
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
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 186.

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_int
```

Step 5 Set the iptables_int 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.

Table of NX-API Response Codes

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 10: 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	Input Bash command error.
CHUNK_ALLOW_ONE_CMD_ERR	400	Chunking only allowed to 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	Request message is invalid.
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.
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.
LIBXML_PARSE_ERR	500	System internal LIBXML parse error.
LIBXML_PATH_CTX_ERR	500	System internal LIBXML path context 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 `show` 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 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 not supported.
PIPE_XML_NOT_ALLOWED_IN_INPUT	501	Pipe XML is not allowed in input.
PIPE_NOT_ALLOWED_IN_INPUT	501	Pipe is not allowed for this input type.
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.
RESP_BIG_JSON_NOT_ALLOWED_ERR	501	Response has large amount of output. JSON not supported.
STRUCT_NOT_SUPPORTED_ERR	501	Structured output unsupported.
ERR_UNDEFINED	600	Undefined.

XML and JSON Supported Commands

The NX-OS supports redirecting the standard output of various **show** commands in the following structured output formats:

- XML
- JSON
- JSON Pretty, which makes the standard block of JSON-formatted output easier to read
- 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 JSON, JSON Pretty, or XML format 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 JSON, JSON Pretty, JSON Native, JSON Native Pretty, or XML, and the NX-OS command output will be 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, and XML output.

Selected examples of this feature follow.

About JSON (JavaScript Object Notation)

JSON is a light-weight text-based open standard 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 Pretty format, as well as JSON Native and JSON Pretty Native, is also supported.

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 /XML output for a show command can also be accessed via sandbox.

CLI Execution

```
BLR-VXLAN-NPT-CR-179# show cdp neighbors  | json
{"TABLE_cdp_neighbor_brief_info": {"ROW_cdp_neighbor_brief_info": [{"ifindex": "
83886080", "device_id": "SW-SPARSHA-SAVBU-F10", "intf_id": "mgmt0", "ttl": "148"
, "capability": ["switch", "IGMP_cnd_filtering"], "platform_id": "cisco WS-C2960
S-48TS-L", "port_id": "GigabitEthernet1/0/24"}, {"ifindex": "436207616", "device
_id": "BLR-VXLAN-NPT-CR-178(FOC1745R01W)", "intf_id": "Ethernet1/1", "ttl": "166
", "capability": ["router", "switch", "IGMP_cnd_filtering", "Supports-STP-Disput
e"], "platform_id": "N3K-C3132Q-40G", "port_id": "Ethernet1/1"}]}}
BLR-VXLAN-NPT-CR-179#
```

Examples of 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", "u
sed_v4_lpm": "6", "used_v6_lpm": "3", "used_v6_lpm_128": "1", "used_host_lpm_tot
al": "0", "used_host_v4_lpm": "0", "used_host_v6_lpm": "0", "used_mcast": "0", "
used_mcast_oifl": "2", "used_host_in_host_total": "13", "used_host4_in_host": "1
2", "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://w
ww.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_oifl>2</used_mcast_oifl>
         <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 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 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://w
ww.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:p
arams: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_summar
y": "0", "acl_name": "acl-test1", "statistics": "enable", "frag_opt_permit_deny
": "permit-all", "TABLE_seqno": {"ROW_seqno": {"seqno": "10", "permitdeny": "pe
rmit", "ip": "ip", "src_ip_prefix": "192.0.2.1/24", "dest_any": "any"}}}}}
switch-1(config-acl)#
```

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 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#
```

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 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_t
os0_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,"m
ax_metric_adver":"false","asext_lsa_cnt":0,"asext_lsa_crc":"0","asopaque_lsa_cn
t":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","ins
tance_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":"PT0
S","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_pe
riod":"PT60S","gr_state":"inactive","gr_last_status":"None","support_tos0_only"
:"true","support_opaque_lsa":"true","is_abr":"false","is_asbr":"false","admin_d
ist":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":"PT
5S","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,"aso
paque_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_d
iscard_rt_ext":"false","no_discard_rt_int":"false"},{"ptag":"112","instance_num
ber":2,"cname":"default","rid":"0.0.0.0","stateful_ha":"true","gr_ha":"true","g
r_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":"PT
0S","spf_hold_time":"PT1S","spf_max_time":"PT5S","lsa_start_time":"PT0S","lsa_h
old_time":"PT5S","lsa_max_time":"PT5S","min_lsa_arr_time":"PT1S","lsa_aging_pac
e":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_norm
al":0,"area_stub":0,"area_nssa":0,"act_area_total":0,"act_area_normal":0,"act_a
rea_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 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#
```

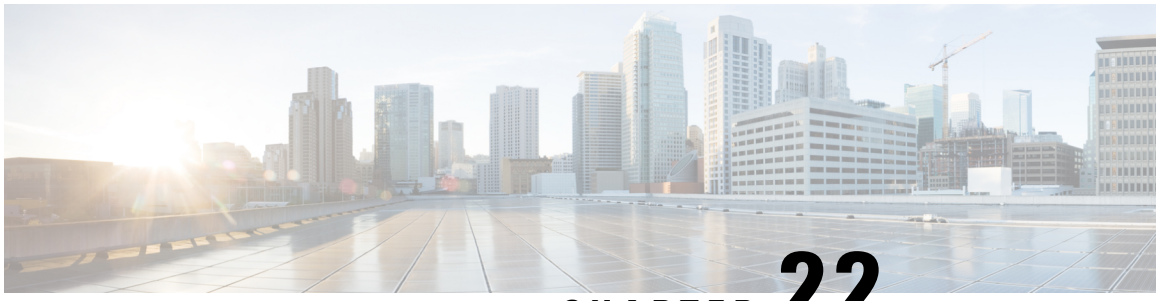
The following example shows how to display the IP route summary in JSON format.

```
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",
                                        "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#
```



CHAPTER 22

NX-API REST

- About NX-API REST, on page 199
- DME Config Replace Through REST, on page 199

About NX-API REST

NX-API REST

NX-API REST is available for use with the Cisco Nexus 3132Q-XL, 3172PQ-XL, and 3172TQ-XL switches [starting with Cisco NX-OS Release 7.0(3)I2(2)], and with the Cisco Nexus 3164Q and 31128PQ switches. NX-API REST is not supported for other Cisco Nexus 3000 Series and 3100 platform switches.

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. If the Nginx resource usage exceeds the cgroup limitations, the Nginx process is restarted and restored.

For more information about the Cisco Nexus 3000 and 9000 Series NX-API REST SDK, see https://developer.cisco.com/site/nx-api/documents/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 Supported Platforms for Programmability Features, on page 5.
- 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 the System-Level Configuration Through REST PUT

You can replace the entire configuration for the switch by sending a REST PUT from the management client.

Use the following procedure:

Procedure

Step 1 From the client, issue a REST PUT operation with the payload as the System level with the URL as `/api/mo/sys.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 `/api/mo/sys.json?rsp-subtree=full&rsp-prop-include=set-config-only`.

Step 2 Send a GET on the DN you used for the config replace by using `/api/mo/sys.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.

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.

Replacing Property-Level Config 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, on page 205.

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.

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:

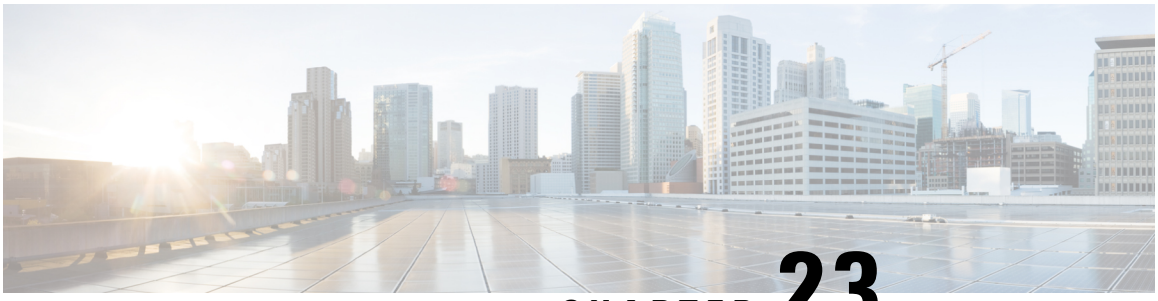
- Successfully set
- Taken from a valid device config

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 23

NX-API Developer Sandbox

This chapter contains the following topics:

- NX-API Developer Sandbox: NX-OS Releases Prior to 9.2(2), on page 205
- NX-API Developer Sandbox: NX-OS Release 9.2(2) and Later, on page 211

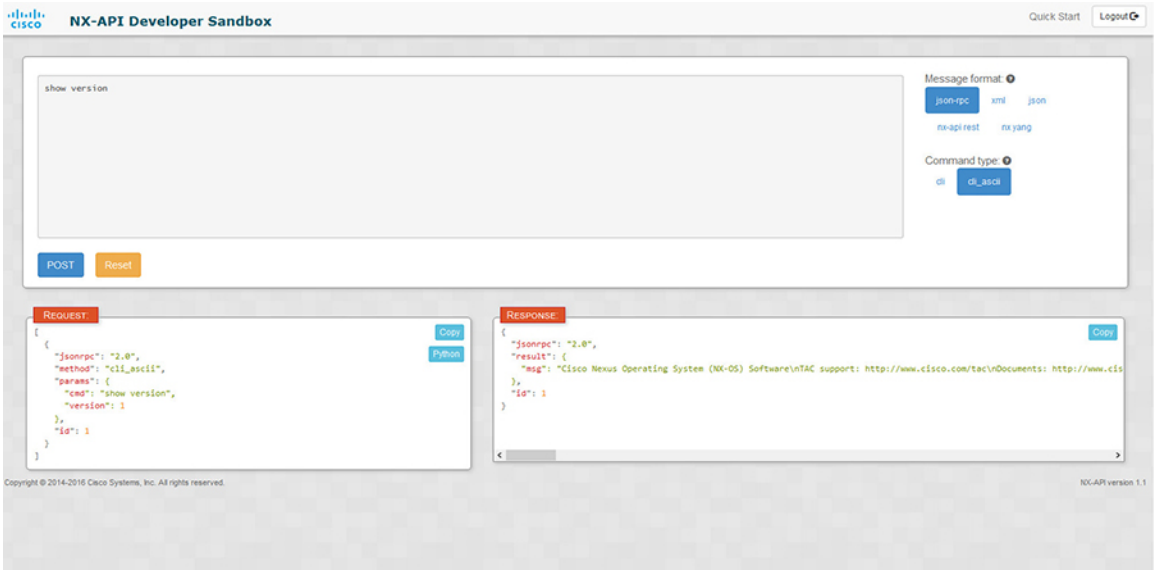
NX-API Developer Sandbox: NX-OS Releases Prior to 9.2(2)

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 **POST** 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.

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 11: 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, see the 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 12: Command Types

Message format	Command type
json-rpc	• cli — show or configuration commands • cli-ascii — show or configuration commands, output without formatting
xml	• 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. **Note** The bash shell must be enabled in the switch.
json	• 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. **Note** The bash shell must be enabled in the switch.
nx-api rest	• cli — configuration commands • model — DN and corresponding payload.
nx yang	• 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 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 NX-API CLI chapter.

Only configuration commands are supported.

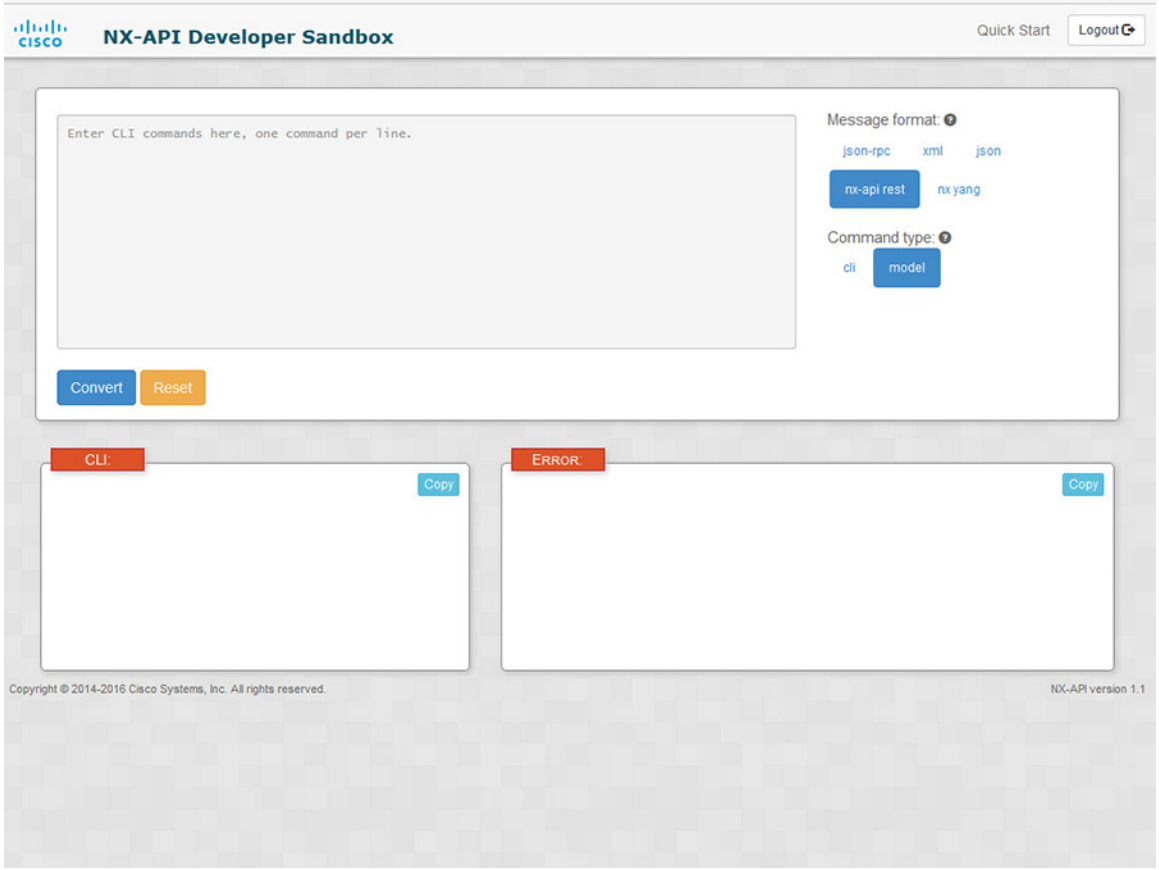
Procedure

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 206.

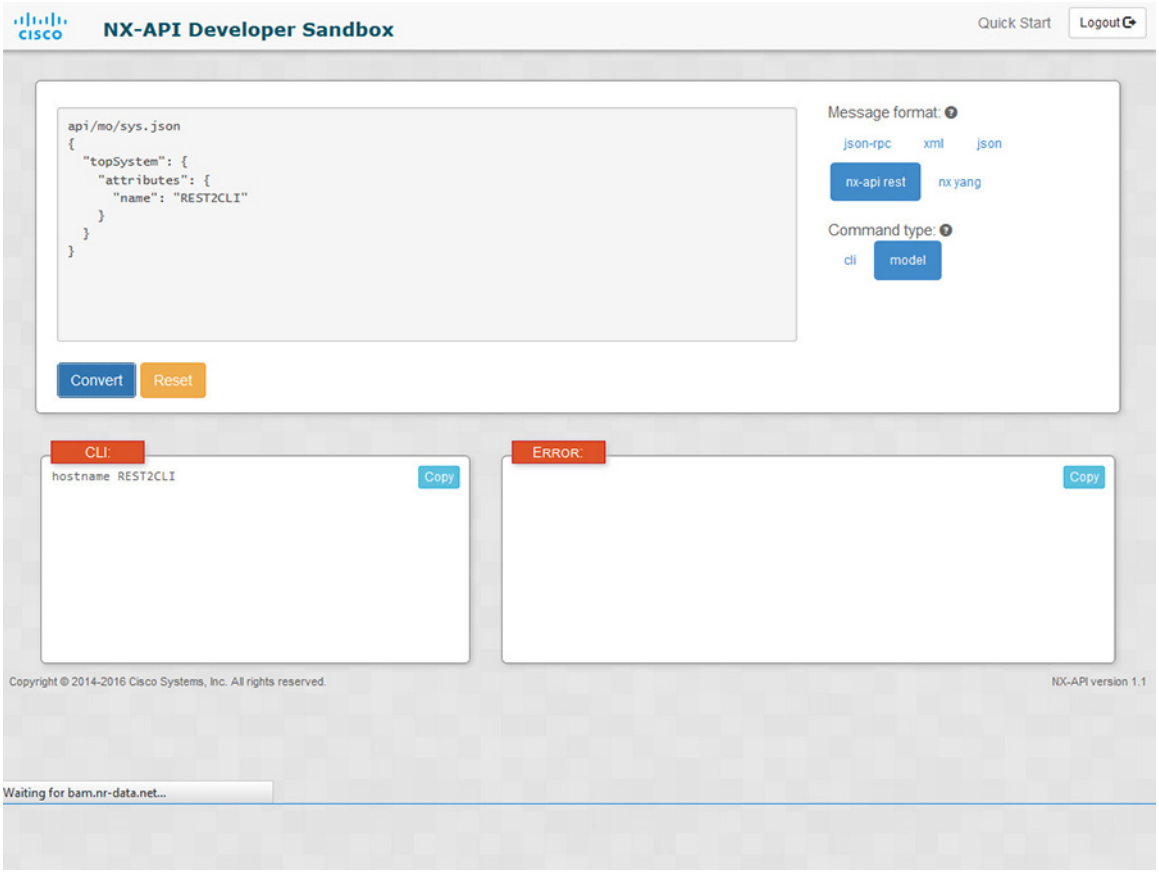
Step 2 Type or paste NX-OS CLI configuration commands, one command per line, into the text entry box in the top pane.

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.

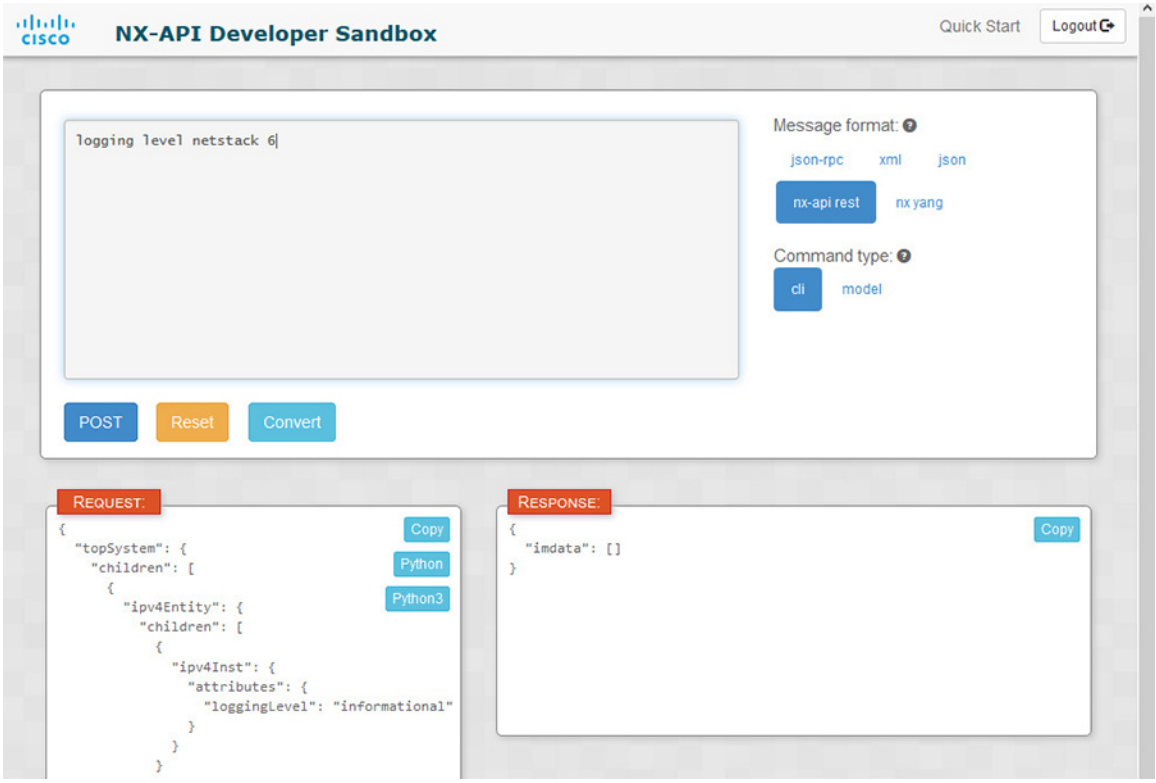
If the CLI commands contain no configuration errors, the payload appears in the **Request** pane. If errors are present, a descriptive error message appears in the **Response** pane.



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.

NX-API Developer Sandbox: NX-OS Release 9.2(2) and Later

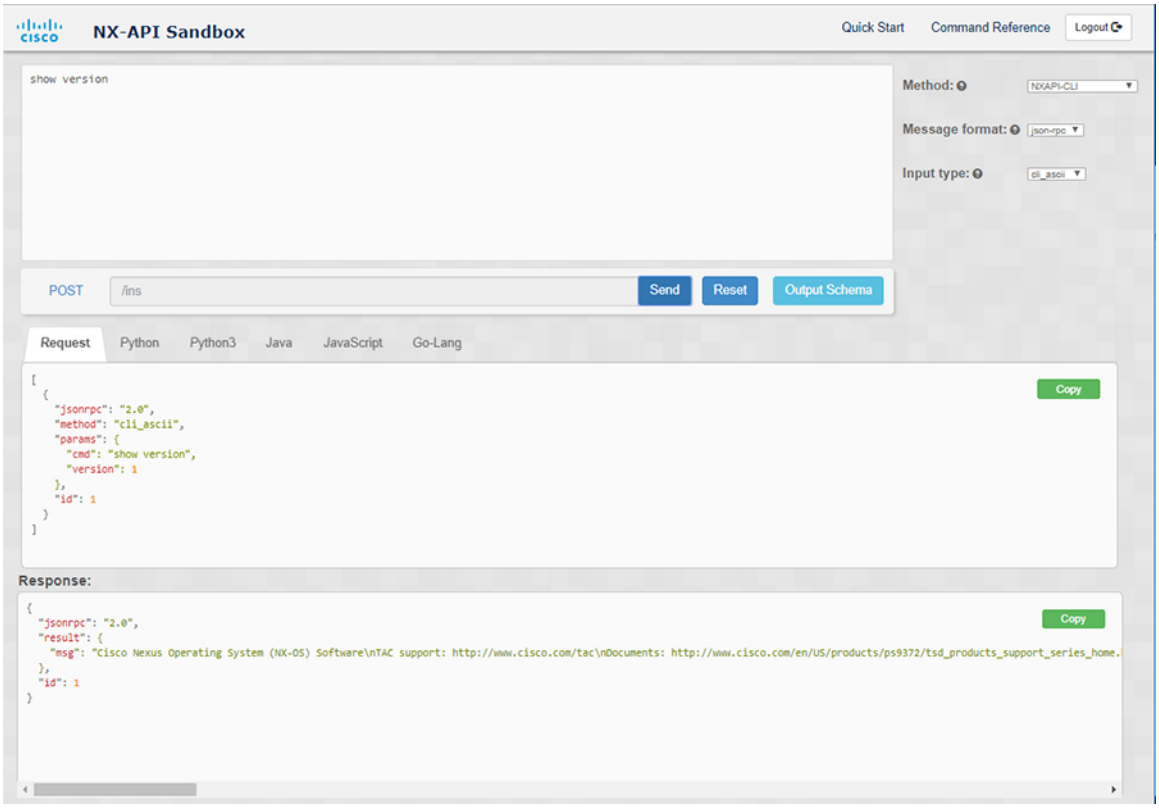
About the NX-API Developer Sandbox

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



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 NXAPI-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 218

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 220

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 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.
- 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 sucessfully.

  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 13: 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.

Protocol	Description
NXAPI-REST (DME)	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: • **POST** • **GET** • **PUT** • **DELETE** 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/.
RESTCONF (Yang)	The YANG ("Yet Another Next Generation") data modeling language for configuration and state data. The RESTCONF (Yang) protocol displays a drop-down list that enables you to choose from the following methods: • **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 14: Command Types

Method	Message format	Input/Command type
NXAPI-CLI	json-rpc	• 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, [ ].

Method	Message format	Input/Command type
NXAPI-CLI	xml	• 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. **Note** The bash shell must be enabled in the switch.
NXAPI-CLI	json	• cli_show — show commands. If the command does not support XML output, an error message will be returned. **Note** Beginning with Cisco NX-OS Release 9.3(3), the cli_show_array command is recommended over the cli_show command. • 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. **Note** The bash shell must be enabled in the switch.
NXAPI-REST (DME)		• cli — CLI to model conversion • model — Model to CLI conversion.
RESTCONF (Yang)	• 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 15: 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

You can use the Cisco NX-API Developer Sandbox to make multiple conversions, including the following:

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

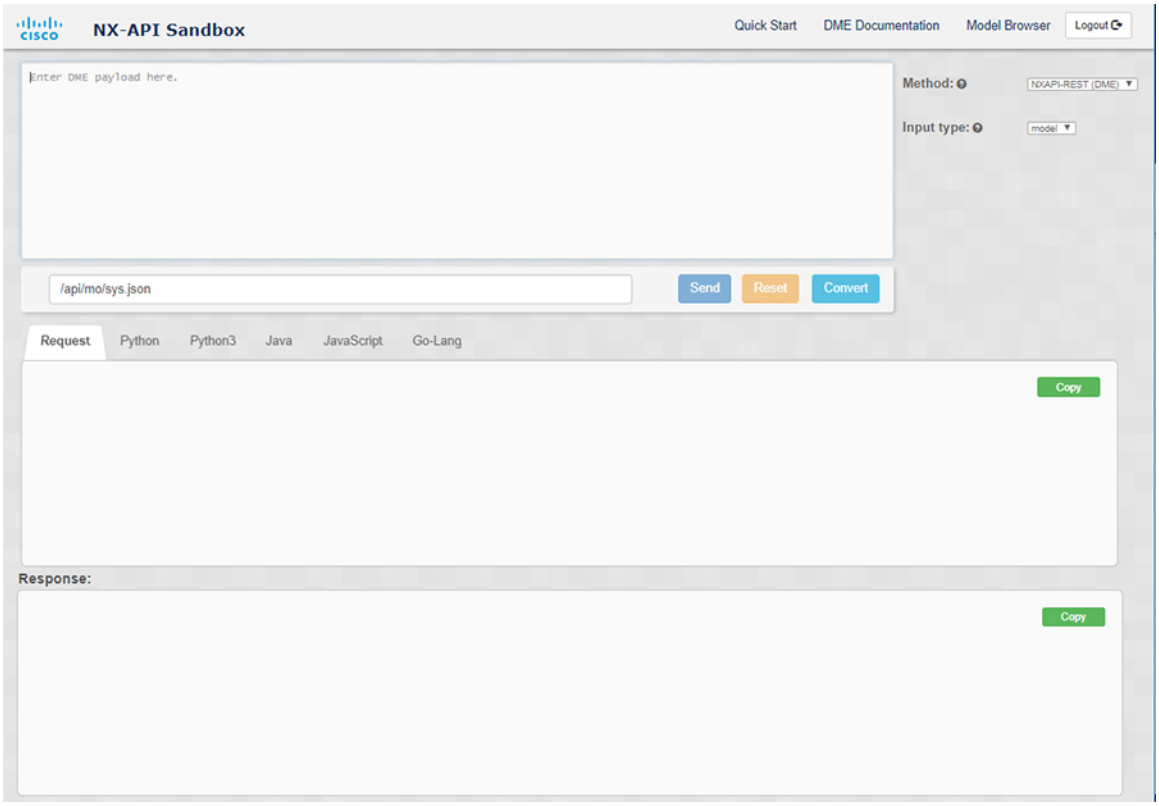
Step 1 Click the **Method** drop-down list and choose **NXAPI-REST (DME)**.

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.



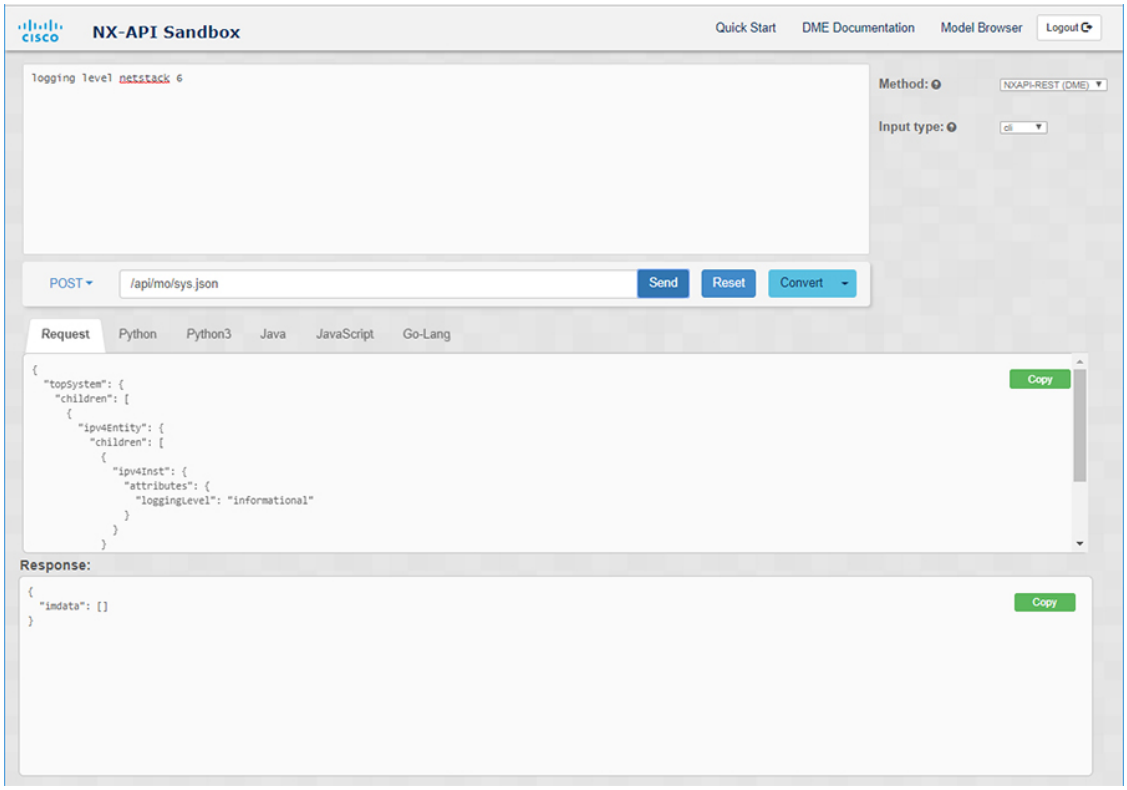
Step 4 Click **Convert**.

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 o the respective pane is copied to the clipboard.

Step 9 (Optional) To convert the request into an 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 get information about the respective field.

  For additional details, such as response codes and security methods, see the chapter *NX-API CLI*.

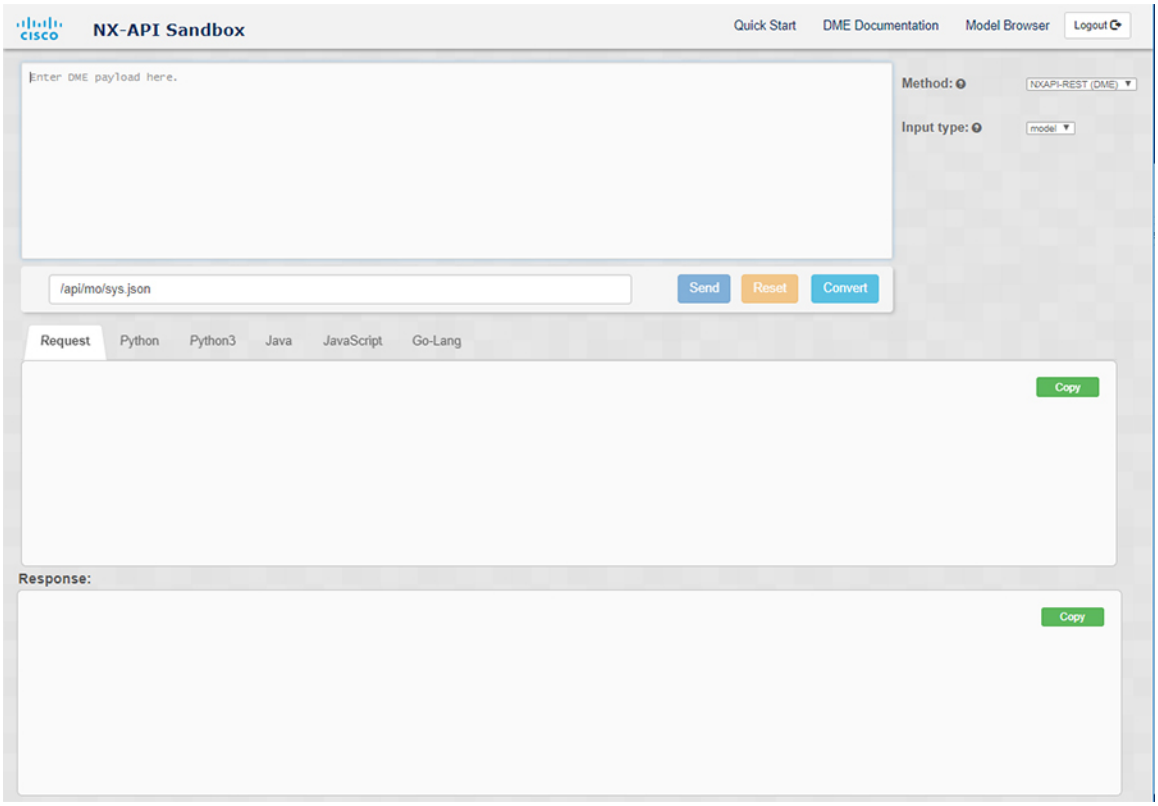
- The top-right corner of the Cisco NX-API Developer Sandbox contains links for additional information. The links that appear depend on the **Method** you choose. The links that appear for the NXAPI-REST (DME) method:
  - **NX-API References**—Enables you to access additional NX-API documentation.
  - **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:



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.

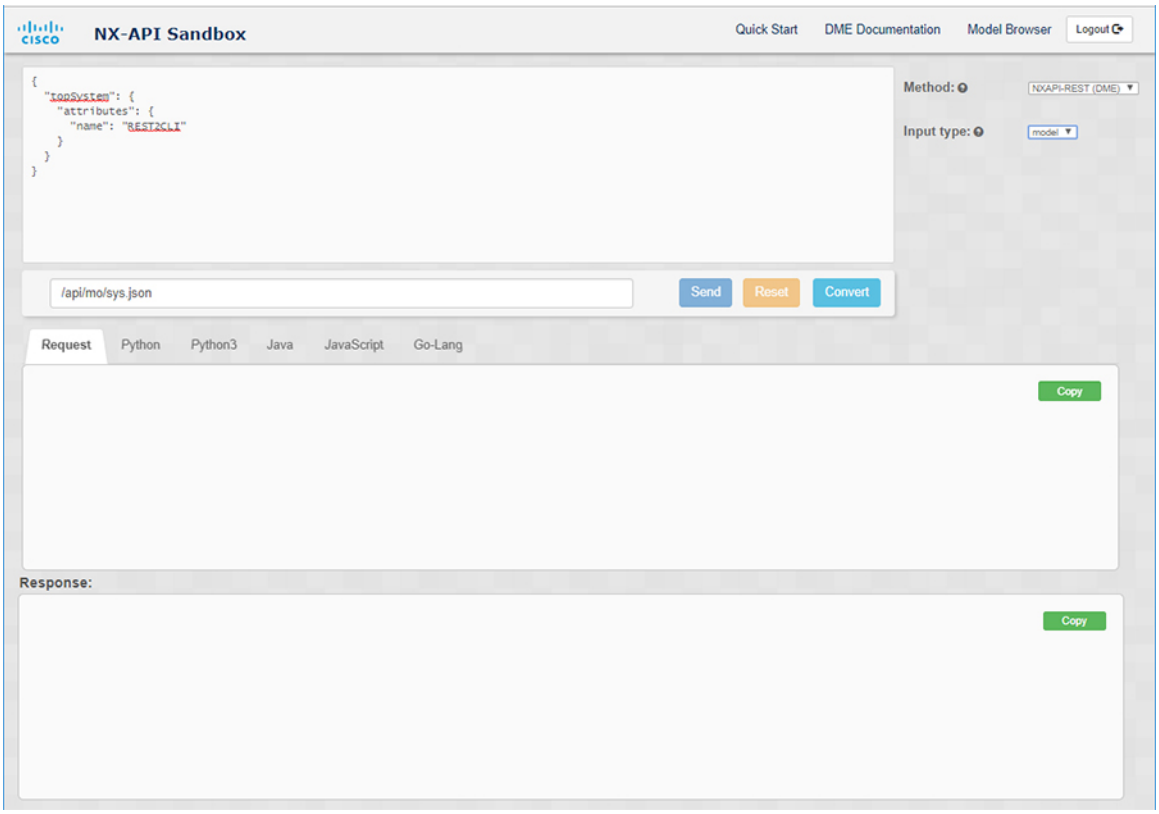
Step 4 Enter the payload in the Command pane.

Step 5 Click **Convert**.

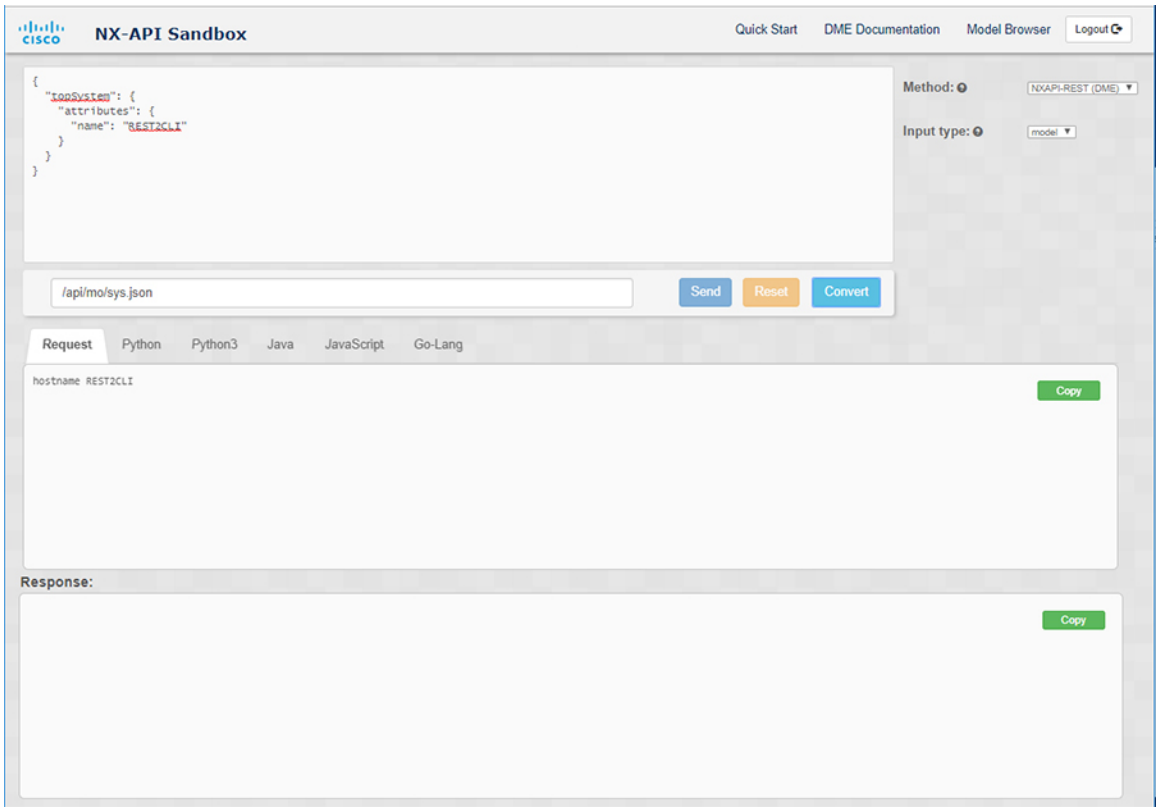
Example:

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.



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 16: Sources of REST2CLI Errors

Payload Issue	Result
The payload contains an attribute that does not exist in the MO. Example: `api/mo/sys.json` `{ "topSystem": { "children": [ { "interfaceEntity": { "children": [ { "l1PhysIf": { "attributes": { "id": "eth1/1", "fakeattribute": "totallyFake" } } } ] } } ] } }`	The **Error** pane will return an error related to the attribute. Example: **CLI** **Error** `unknown attribute 'fakeattribute' in element 'l1PhysIf'`
The payload includes MOs that aren't yet supported for conversion: Example: `api/mo/sys.json` `{ "topSystem": { "children": [ { "dhcpEntity": { "children": [ { "dhcpInst": { "attributes": { "SnoopingEnabled": "yes" } } } ] } } ] } }`	The **Error** Pane will return an error related to the unsupported MO. Example: **CLI** **Error** `The entire subtree of "sys/dhcp" is not converted.`

Using the Developer Sandbox to Convert from RESTCONF to json or XML

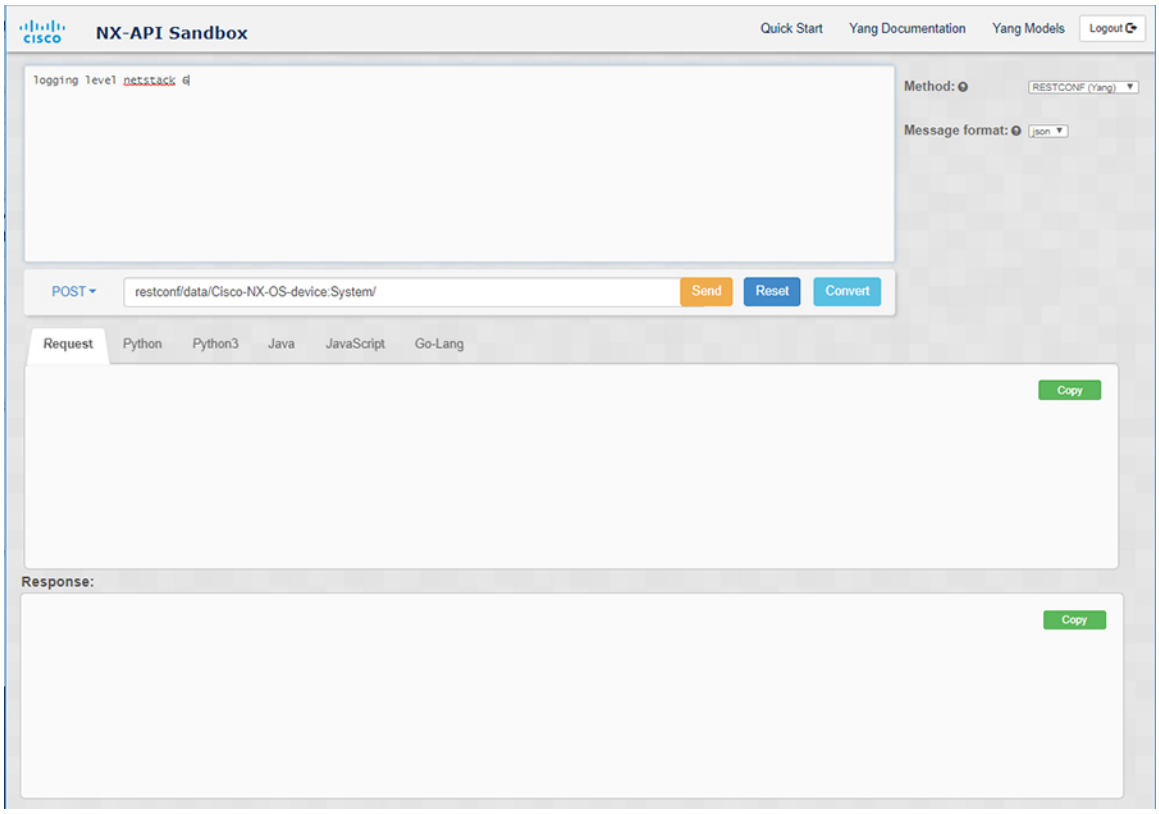
Tip

- 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:



Step 2 Click **Message format** and choose either **json** or **xml**.

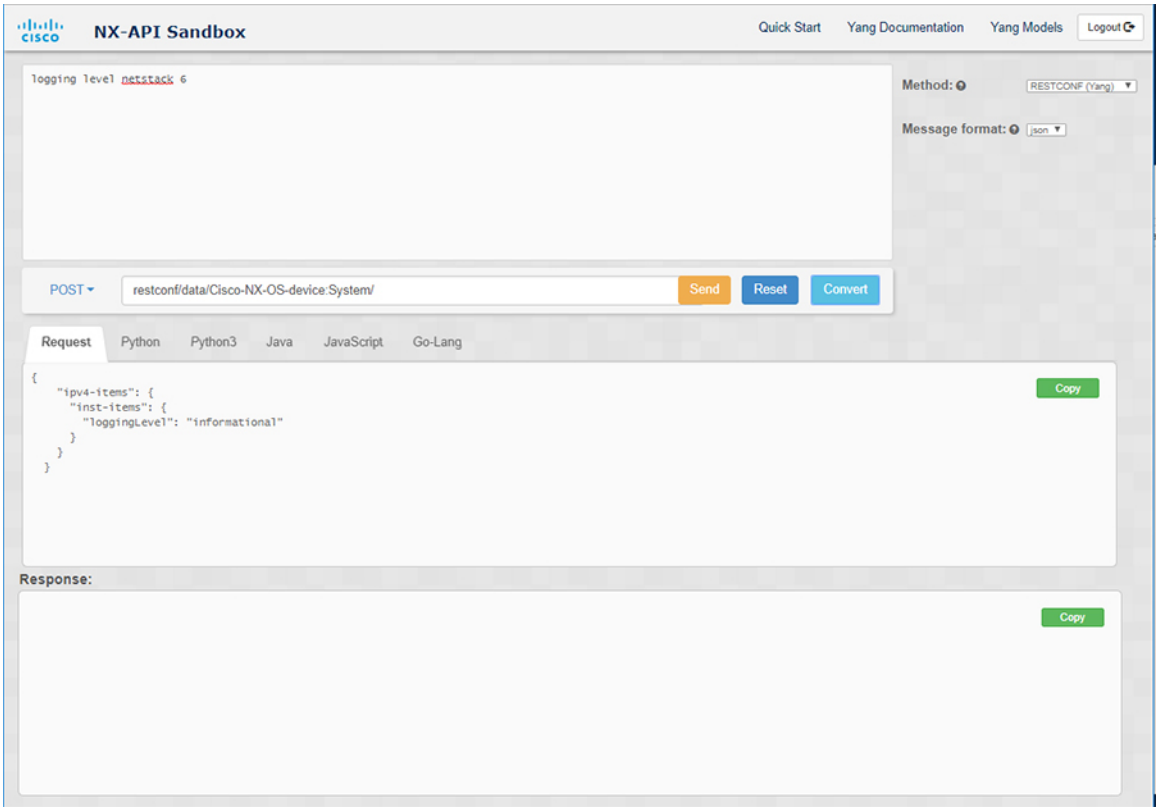
Step 3 Enter a command in the text entry box in the top pane.

Step 4 Choose a message format.

Step 5 Click **Convert**.

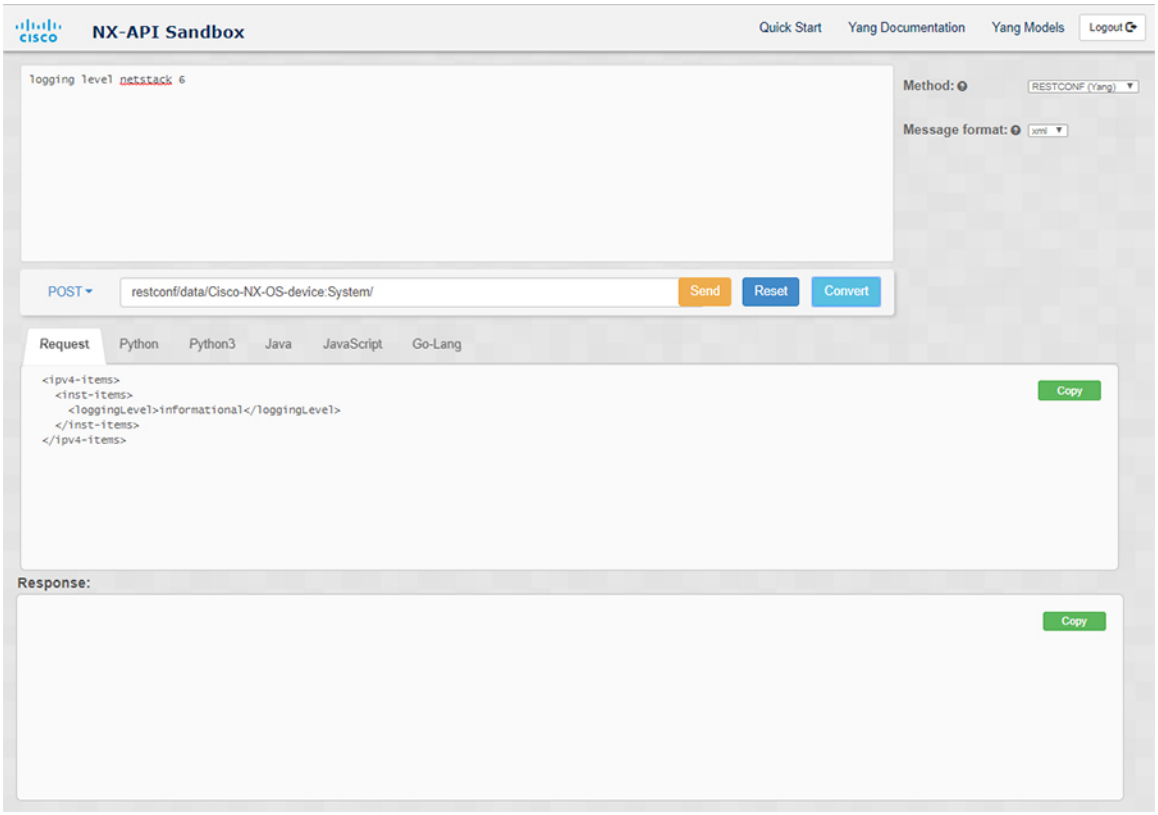
Example:

For this example, the command is **`logging level netstack 6`** and the message format is json:



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

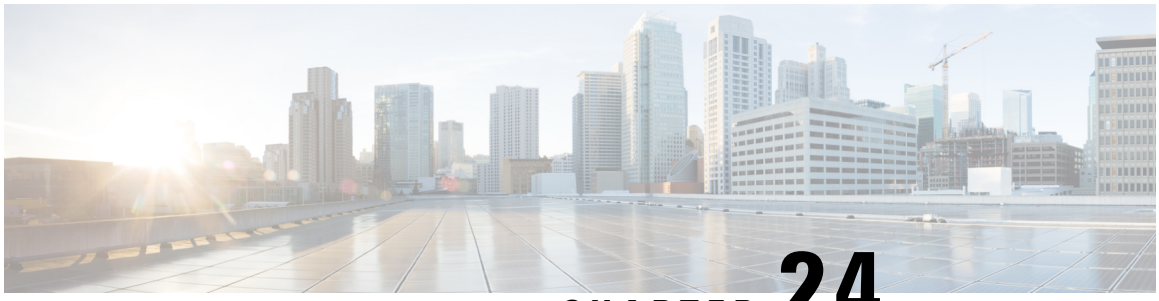
The Java-generated script does not work if you choose the PATCH option from the drop-down menu in the area above the Request tab. This is a known limitation with Java and is expected behavior.



PART IV

Model-Driven Programmability

- Infrastructure Overview, on page 231
- Managing Components, on page 235
- NETCONF Agent, on page 241
- Converting CLI Commands to Network Configuration Format, on page 265
- RESTConf Agent, on page 271
- gRPC Agent, on page 275
- gNMI - gRPC Network Management Interface, on page 289
- Dynamic Logger, on page 325
- Model-Driven Telemetry, on page 333
- Inband Network Telemetry, on page 405
- Packetized MMU Stats, on page 427



CHAPTER 24

Infrastructure Overview

- About Model-Driven Programmability, on page 231
- About the Programmable Interface Infrastructure, on page 231

About Model-Driven Programmability

The model-driven programmability of the NX-OS device allows you to automate the configuration and control of the device.

Data Modeling

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.

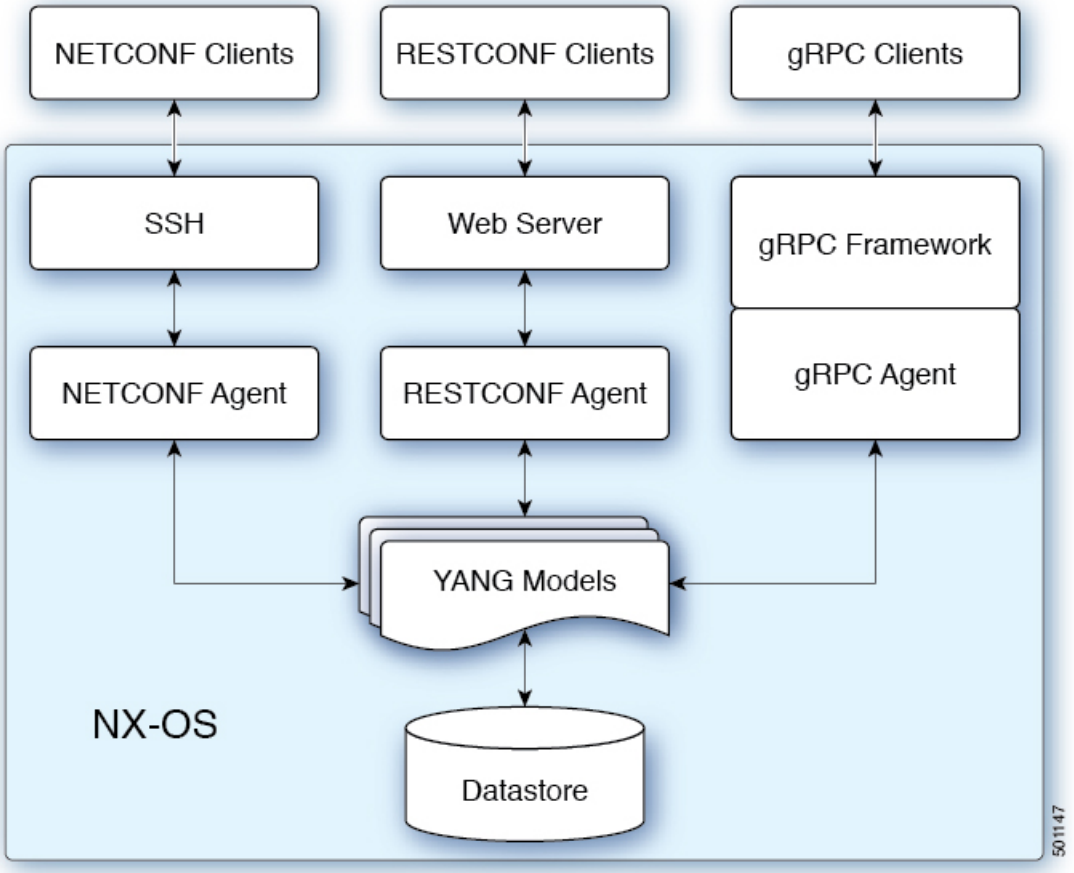
The Cisco NX-OS device supports the YANG data modeling language. YANG is a data modeling language used to describe configuration and operational data, remote procedure calls, and notifications for network devices.

Programmable Interfaces

Three standards-based programmable interfaces are supported by NX-OS for operations on the data model: NETCONF, RESTConf, and gRPC.

About the Programmable Interface Infrastructure

This section provides a brief overview of the NX-OS Programmable Interface infrastructure.



When a request is received whether via NETCONF, RESTConf, or gRPC, the request is converted into an abstract message object that distributed to the underlying model infrastructure based on the 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 device data store returning the results to the agent of origin for response transmission back to the requesting client.

NX-OS Programmable Interface Agents

Agents provide an interface between the Device and clients. They specify the transport, the protocol and the encoding of the communications with the Device. NX-OS Programmable Interfaces supports three agents: NETCONF, RESTConf and gRPC, each providing different interfaces for configuration management of the Device via YANG models.

Note Supported YANG models for each Cisco NX-OS release are provided at https://devhub.cisco.com/artifactory/open-nxos-agents.

Table 17: NX-OS Programmable Interface Agents

Agent	Transport	Protocol	Encoding
NETCONF	SSH	RFC 6241[1]	XML
RESTConf	HTTP	draft-ietf-netconf-restconf-10[2]	XML or JSON
gRPC	HTTP	gRPC Protocol Spec[3]	Google Protobuf

The protocol specifications are described in the following documents:

- [1] R. Enns, Ed., RFC6241: *Network Configuration Protocol (NETCONF)* http://www.rfc-editor.org/rfc/rfc6241.txt
- [2] RESTCONF Protocol draft-ietf-netconf-restconf-10 https://tools.ietf.org/html/draft-ietf-netconf-restconf-10
- [3] Cisco NX-OS gRPC Protocol Specification

Model Infrastructure

The Model Infrastructure takes requests received from the Agent, determines the namespace 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 forth as specified by the Agent protocol.

Device YANG Model

The Device Configuration is described in a YANG model called a Device Model. The Device Model is manifested in the Model Infrastructure as another model component with the Device namespace.

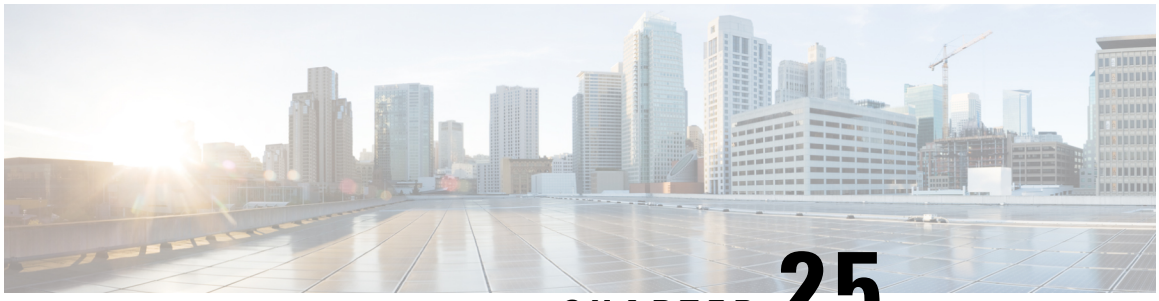
Common YANG Models

A Common Model is another kind of model component that contains within its elements, YANG Paths 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.

Additional YANG References

The YANG modeling language is described in the following document: M. Bjorklund, Ed., RFC6020: *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)* http://www.rfc-editor.org/rfc/rfc6020.txt

Additional information about YANG can be found at the *YANG Central Wiki* http://www.yang-central.org/twiki/bin/view/Main/WebHome (M. Bjorklund, Ed.)



CHAPTER 25

Managing Components

- About the Component RPM Packages, on page 235
- Preparing For Installation, on page 237
- Downloading Components from the Cisco Artifactory, on page 238
- Installing RPM Packages, on page 239

About the Component RPM Packages

Note Beginning with Cisco NX-OS Release 7.0(3)I6(2), the NX-OS Programmable Interface Base Component RPM packages (agents, the Cisco native model, most of the other required models, and infrastructure) are included in the Cisco NX-OS image. As a result, nearly all the required software is installed automatically when the image is loaded. This situation means that there is no need to download and install the bulk of the software from the Cisco Artifactory. The exception is the OpenConfig model, which is required. You must explicitly download the OpenConfig models from the Cisco Artifactory.

But, for Cisco NX-OS Release 7.0(3)I6(1) and earlier releases, if you need to upgrade, the following sections describing downloading and installing the packages are required.

NX-OS Programmable Interface Component RPM packages may be downloaded from the Cisco Artifactory. There are two types of component RPM packages that are needed:

- Base Components (required)
- Common Model Components (OpenConfig models must be explicitly downloaded and installed)

Base Components

The Base Components comprise the following required RPM packages:

- **mtx-infra** — Infrastructure
- **mtx-device** — Cisco native model

At least one of the following agent packages must be installed in order to have access to the modeled NX-OS interface:

- **mtx-netconf-agent** — NETCONF agent

- **mtx-restconf-agent** — RESTCONF agent
- **mtx-grpc-agent** — gRPC agent

Common Model Components

Common Model component RPMs support OpenConfig models. To use the OpenConfig models, you must download and install the OpenConfig RPMs. For convenience, there is a single combined package of all supported OpenConfig models, `mtx-openconfig-all`.

While the single combined package is recommended, an alternative is to download and install RPMs of selected models and their dependencies among the supported models listed in the following table. The `mtx-openconfig-all` RPM is not compatible with the individual model RPMs. You must uninstall the former before installing the latter, and you must unistall the latter before installing the former.

Model Name	Model Rev	Model Ver	Package Name	Dependencies
openconfig-acl	2017-05-26	1.0.0	mtx-openconfig-acl	mtx-openconfig-interfaces
openconfig-bgp-policy	2017-07-30	4.0.1	mtx-openconfig-bgp-policy	mtx-openconfig-interfaces mtx-openconfig-routing-policy
openconfig-if-aggregate	2017-07-14	2.0.0	mtx-openconfig-if-aggregate	mtx-openconfig-if-ethernet mtx-openconfig-interfaces
openconfig-if-ethernet	2017-07-14	2.0.0	mtx-openconfig-if-ethernet	mtx-openconfig-interfaces
openconfig-if-ip	2016-05-26	1.0.2	mtx-openconfig-if-ip	mtx-openconfig-if-aggregate mtx-openconfig-if-ethernet mtx-openconfig-interfaces mtx-openconfig-vlan
openconfig-if-ip-ext	2018-01-05	2.3.0	mtx-openconfig-if-ip-ext	mtx-openconfig-if-aggregate mtx-openconfig-if-ethernet mtx-openconfig-if-ip mtx-openconfig-interfaces mtx-openconfig-vlan
openconfig-interfaces	2017-07-14	2.0.0	mtx-openconfig-interfaces	-

Model Name	Model Rev	Model Ver	Package Name	Dependencies
openconfig-network-instance	2017-08-24	0.8.1	mtx-openconfig-network-instance	mtx-openconfig-bgp-policy mtx-openconfig-if-aggregate mtx-openconfig-if-ethernet mtx-openconfig-interfaces mtx-openconfig-routing-policy mtx-openconfig-vlan
openconfig-network-instance-policy	2017-02-15	0.1.0	mtx-openconfig-network-instance-policy	mtx-openconfig-routing-policy
openconfig-ospf-policy	2017-08-24	0.1.1	mtx-openconfig-ospf-policy	mtx-openconfig-interfaces mtx-openconfig-routing-policy
openconfig-platform	2018-01-16	0.8.0	mtx-openconfig-platform	-
openconfig-platform-linecard	2017-08-03	0.1.0	mtx-openconfig-platform-linecard	mtx-openconfig-platform
openconfig-platform-port	2018-01-20	0.3.0	mtx-openconfig-platform-port	mtx-openconfig-if-ethernet mtx-openconfig-interfaces mtx-openconfig-platform
openconfig-platform-transceiver	2018-01-22	0.4.1	mtx-openconfig-platform-transceiver	mtx-openconfig-if-ethernet mtx-openconfig-interfaces mtx-openconfig-platform
openconfig-relay-agent	2016-05-16	0.1.0	mtx-openconfig-relay-agent	mtx-openconfig-interfaces
openconfig-routing-policy	2016-05-12	2.0.1	mtx-openconfig-routing-policy	-
openconfig-spanning-tree	2017-07-14	0.2.0	mtx-openconfig-spanning-tree	mtx-openconfig-interfaces
openconfig-system	2017-09-18	0.3.0	mtx-openconfig-system	-
openconfig-vlan	2017-07-14	2.0.0	mtx-openconfig-vlan	mtx-openconfig-if-aggregate mtx-openconfig-if-ethernet mtx-openconfig-interfaces

Preparing For Installation

This section contains installation preparation and other useful information for managing NX-OS Programmable Interface components.

Opening the Bash Shell on the Device

RPM installation on the switch is performed in the Bash shell. Make sure that **feature bash** is configured on the device.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# feature bash-shell
Switch(config)# end
Switch# run bash sudo su
bash-4.2#
```

To return to the device CLI prompt from Bash, type **exit** or **Ctrl-D**.

Verify Device Readiness

You can use the following CLI **show** commands to confirm the readiness of the device before installation of an RPM.

- `show module` — Indicates whether all modules are up.

  ```
  Switch# show module
  ```

- `show system redundancy status` — Indicates whether the standby device is up and running and in HA mode. If a standby sync is in progress, the RPM installation may fail.

  ```
  Switch# show system redundancy status
  ```

If the line cards have failed to come up, enter the **createrepo /rpms** command in the Bash shell.

```
bash-4.2# createrepo /rpms
```

Downloading Components from the Cisco Artifactory

The NX-OS Programmable Interface Component RPMs can be downloaded from the Cisco Artifactory at the following URL. The RPMs are organized by NX-OS release-specific directories. Ensure that you are downloading the RPMs from the correct NX-OS release directory.

https://devhub.cisco.com/artifactory/open-nxos-agents

The NX-OS Programmable Interface Component RPMs adhere to the following naming convention:

*<package>-<version>-<NX-OS release>.<architecture>***.rpm**

Select and download the desired NX-OS Programmable Interface Component RPM packages to the device for installation as described in the following sections.

Installing RPM Packages

Installing the Programmable Interface Base And Common Model Component RPM Packages

Before you begin

- From the Cisco Artifactory, download the following packages:
  - mtx-infra
  - mtx-device
  - mtx-netconf-agent/mtx-restconf-agent/mtx-grpc-agent (at least one)
  - mtx-openconfig-all (alternatively, selected individual models)
- Using the CLI commands in Verify Device Readiness, on page 238, confirm that all line cards in the Active and Standby devices are up and ready.

Procedure

Step 1 Copy the downloaded RPMs to the device.

Example:

```
Switch# copy scp://jdoe@192.0.20.123/myrpms/mtx-infra-2.0.0.0-9.2.1.lib32_n9000.rpm bootflash:
 vrf management
Switch# copy scp://jdoe@192.0.20.123/myrpms/mtx-device-2.0.0.0-9.2.1.lib32_n9000.rpm
bootflash: vrf management
Switch# copy scp://jdoe@192.0.20.123/myrpms/mtx-netconf-agent-2.0.0.0-9.2.1.lib32_n9000.rpm
 bootflash: vrf management
Switch# copy scp://jdoe@192.0.20.123/myrpms/mtx-openconfig-all-1.0.0.0-9.2.1.lib32_n9000.rpm
 bootflash: vrf management
```

Step 2 From the Bash shell, install the RPMs.

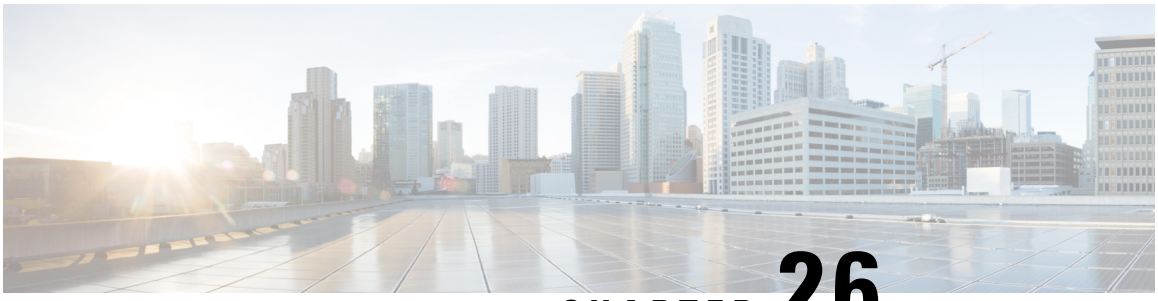
Example:

```
bash-4.2# cd /bootflash
bash-4.2# yum install mtx-infra-2.0.0.0-9.2.1.lib32_n9000.rpm
mtx-device-2.0.0.0-9.2.1.lib32_n9000.rpm mtx-netconf-agent-2.0.0.0-9.2.1.lib32_n9000.rpm
mtx-openconfig-all-1.0.0.0-9.2.1.lib32_n9000.rpm
```

Step 3 From the Bash shell, verify the installation.

Example:

```
bash-4.2# yum list installed | grep mtx
```



CHAPTER 26

NETCONF Agent

This chapter contains the following topics:

- About the NETCONF Agent, on page 241
- Guidelines and Limitations for the NETCONF Agent, on page 242
- Configuring the NETCONF Agent, on page 244
- Establishing a NETCONF Session, on page 245
- NETCONF Operations, on page 247
- NETCONF Notifications, on page 255
- NETCONF Examples, on page 259
- Troubleshooting the NETCONF Agent, on page 263

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 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 descriptions of the operations that are supported:

Table 18: 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
copy-config	Create or replace datastore with the contents of another datastore

Operation	Description
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

Guidelines and Limitations for the NETCONF Agent

The NETCONF Agent has the following guideline and limitation:

- If you configure a routing policy **prefix-list** using the CLI and send a NETCONF **netconf edit-config** or **netconf get-config** or NETCONF Get request for the routing-policy OpenConfig model, it is not supported.

  If you configure a routing policy **prefix-list** using the CLI and request a NETCONF notification for the routing policy OpenConfig model, it is not supported.

- Cisco NX-OS supports both the Cisco Device YANG model and OpenConfig models in NETCONF notifications. Support for OpenConfig models in NETCONF notifications begins with the Cisco NX-OS Release 9.3(5).

- The device YANG model defines ephemeral data and they are marked with a comment "// Ephemeral data". These nonpersistent large-volume data is handled differently from the rest of the model. They are returned only when `<get>` query's `<filter>` parameter points specifically to the particular element marked with the comment. Refer to the ephemeral data support documentation for detailed information on the usage.

- Beginning with Cisco NX-OS Release 9.3(3), 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.

- 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.
  ```

- NETCONF does not support enhanced Role-Based Access Control (RBAC) as specified in RFC 6536. Only users with a "network-admin" role are granted access to the NETCONF agent.

- Beginning with Cisco NX-OS Release 9.3(1), 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 a message that is similar to the following, the requests are not carrying a namespace:

  ```
  Request without namespace and filter is an unsupported operation
  ```

  The following example shows a `get` request and response with the behavior before this change. This example shows the error message that is caused by behavior which is no longer supported.

  Request:

  ```
  <get>
  </get>
  ```

  Response:

  ```
  <rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="1">
    <rpc-error>
        <error-type>protocol</error-type>
        <error-tag>operation-not-supported</error-tag>
        <error-severity>error</error-severity>
        <error-message xml:lang="en">Request without filtering is an unsupported
  operation</error-message>
    </rpc-error>
  </rpc-reply>
  ```

  The following example shows a `get` request and response with the correct behavior in Cisco NX-OS Release 9.3(1) and later.

  Request:

  ```
  <get>
    <filter>
      <System xmlns="http://cisco.com/ns/yang/cisco-nx-os-device">
      </System>
    </filter>
  </get>
  ```

  Response:

  ```
  <rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
      <data>
          <System> …
      </data>
  </rpc-reply>
  ```

- 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.
- The Cisco NX-OS NETCONF server 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.
- Cisco NX-OS supports NETCONF notifications beginning with Cisco NX-OS Release 9.3(1). Cisco NX-OS supports only the Cisco Device YANG model.
- Cisco NX-OS supports both the Cisco Device YANG model and OpenConfig models. Support for OpenConfig models in NETCONF notifications begins with the Cisco NX-OS Release 9.3(5).
- All Cisco Nexus 3100-V, 3100-XL, and the 3200 platform switches support NETCONF.

Configuring the NETCONF Agent

Configuring the NETCONF Agent for Cisco NX-OS 9.3(5) and Later

This procedure describes how to enable and configure the NETCONF Agent.

Note Use this procedure with Cisco NX-OS Release 9.3(5) and later.

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.

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`	Enable NETCONF services.
Step 3	(Optional) **netconf idle-timeout** *it-num* **Example:** `switch(config)# netconf idle-timeout 5`	(Optional) Specifies the timeout in minutes after which idle client sessions are disconnected. The range of*it-num* is 0-1440 minutes. The default timeout is 5 minutes. A value of 0 disables timeout.

	Command or Action	Purpose
Step 4	(Optional) **netconf sessions** *num-sessions* **Example:** `switch(config)# netconf sessions 5`	Specifies the number of maximum simultaneous client sessions. The range of *num-sessions* is 1-10. The default is 5 sessions.

Configuring the NETCONF Agent for Cisco NX-OS 9.3(4) and Earlier

Note Use this procedure with Cisco NX-OS Release 9.3(4) and earlier.

The NETCONF Agent supports the following optional configuration parameters under the `[netconf]` section in the configuration file (`/etc/mtx.conf`).

Parameter	Description
idle_timeout	(Optional) Specifies the timeout in minutes after which idle client sessions are disconnected. The default value is 5 minutes. A value of 0 disables timeout.
limit	(Optional) Specifies the number of maximum simultaneous client sessions. The default value is 5 sessions. The range is 1 to 10.

The following is an example of the `[netconf]` section in the configuration file:

```
[netconf]
mtxadapter=/opt/mtx/lib/libmtxadapternetconf.1.0.1.so
idle_timeout=10
limit=1
```

For the modified configuration file to take effect, you must restart the NETCONF Agent using the CLI command [**no**] **feature netconf** to disable and re-enable.

Establishing 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. 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 when 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 peers 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.

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 NETCONF base capabilities that are supported on Cisco NX-OS Release 9.3(4). Also, the server's <hello> message includes supported data models. They might not match the models that are supported in 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&amp;module=Cisco-NX-OS-device</capability>


<capability>http://openconfig.net/yang/acl?revision=2019-11-27&amp;module=openconfig-acl&amp;deviations=cisco-nx-openconfig-acl-deviations</capability>


<capability>http://openconfig.net/yang/bfd?revision=2019-10-25&amp;module=openconfig-bfd&amp;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 the NETCONF agent with the **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 **ssh** command is used for the example above for illustration purposes only. There are various clients that are written for NETCONF which are recommended over the **ssh** command. The ncclient is one such example and is used in the Usage Examples section.

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 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 Operations

This section describes supported base protocol 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>
            <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 a specified configuration to the target datastore. The <target> parameter specifies the datastore being edited, such as <running/> or <candidate/>. 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 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”. The “remove” operation is different from “delete” in that no error is returned if the configuration data does not exist. 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 has “merge”, “replace” or “none”.

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. Beginning with Cisco NX-OS Release 9.3(1), support for NETCONF notifications began as described in RFC 5277. This is an optional capability that is advertised in the NETCONF hello message.

A NETCONF client can subscribe for notifications using Deviceyang or OpenConfig models. Support for OpenConfig models in NETCONF notifications begins with Cisco NX-OS Release 9.3(5).

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 as long as 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 system events. For example, port state change, fan speed change, and process memory change to name a few. Also, configuration events such as a new feature being enabled.

- 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>
```

Creating 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 you must to write client programs that contain the `<create-subscription>` operation. The following is an example for `<create-subscription>` sent by any 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 any of 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 is able to create the subscription successfully.

The following is a sample successful response received in the client for the `<create-subscription>` request that it sent to the server.

Response for `<create-subscription>`, received in the client:

```
<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 the Start Time and Stop Time options cannot be used.

Receiving 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, for the filter used. The event notification is its own XML-formatted document that contains the notification element.

The following is a sample notification for an Ethernet interface going down, when the client subscribed to interface operSt, from the DeviceYang model. The <create-subscription> is in the Creating Subscriptions section.

```
<?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.

Terminating Subscriptions

Subscriptions are terminated when the NETCONF client sends specific operations to the Cisco NX-OS NETCONF server in the payload of a NETCONF message. Subscription termination occurs 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 Examples



Note All examples in this section use the ncclient python library.

Connecting Cisco NX-OS with the ncclient

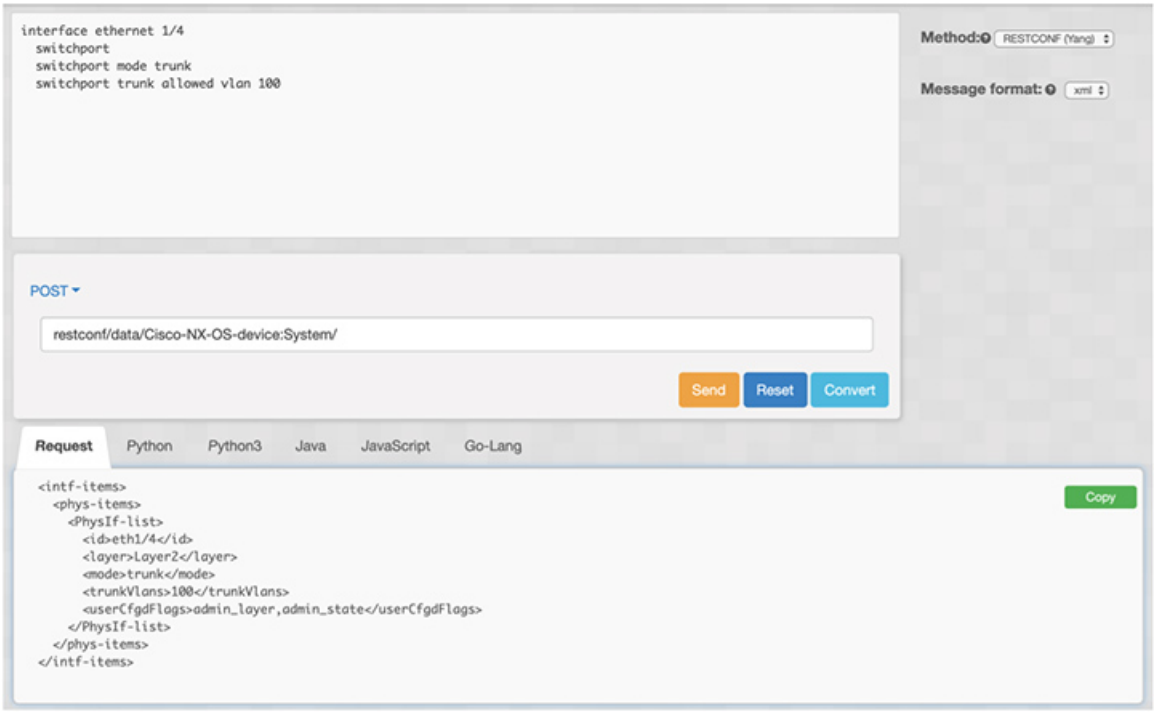
The ncclient is a Python library for NETCONF clients. The following is an example of how to establish a connection to Cisco NX-OS from the ncclient Manager API:

```
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
```

Using the Sandbox to Generate the NETCONF Payload

Refer to NXAPI Developer Sandbox section to enable it. In order 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: NCCLIENT



Getting Configuration Data from Cisco NX-OS

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 from Cisco NX-OS

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 Using NETCONF

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 Using NETCONF

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)

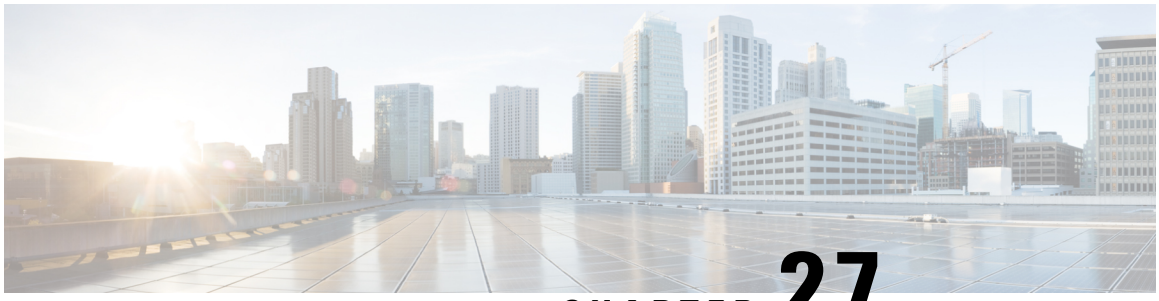
        print(netconf_response)

if __name__ == '__main__':
    sys.exit(main())
```

Troubleshooting the NETCONF Agent

Troubleshooting Connectivity

- From a client system, ping the management port of the switch to verify that the switch is reachable.
- In the bash shell of the switch, execute the **service netconf status** command to check the agent status.
- 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 (similar to what is shown in the Establishing a NETCONF Session section) from the server if the server does not respond with the correct NETCONF messages.



CHAPTER 27

Converting CLI Commands to Network Configuration Format

- Information About XMLIN, on page 265
- Licensing Requirements for XMLIN, on page 265
- Installing and Using the XMLIN Tool, on page 266
- Converting Show Command Output to XML, on page 266
- Configuration Examples for XMLIN, on page 267

Information About XMLIN

The XMLIN tool converts CLI commands to the Network Configuration (NETCONF) protocol format. NETCONF is a network management protocol that provides mechanisms to install, manipulate, and delete the configuration of network devices. It uses XML-based encoding for configuration data and protocol messages. The NX-OS implementation of the NETCONF protocol supports the following protocol operations: <get>, <edit-config>, <close-session>, <kill-session>, and <exec-command>.

The XMLIN tool converts show, EXEC, and configuration commands to corresponding NETCONF <get>, <exec-command>, and <edit-config> requests. You can enter multiple configuration commands into a single NETCONF <edit-config> instance.

The XMLIN tool also converts the output of show commands to XML format.

Licensing Requirements for XMLIN

Table 19: XMLIN Licensing Requirements

Product	License Requirement
Cisco NX-OS	XMLIN requires no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the *Cisco NX-OS Licensing Guide*.

Installing and Using the XMLIN Tool

You can install the XMLIN tool and then use it to convert configuration commands to NETCONF format.

Before you begin

The XMLIN tool can generate NETCONF instances of commands even if the corresponding feature sets or required hardware capabilities are not available on the device. But, you might still need to install some feature sets before entering the **xmlin** command.

Procedure

	Command or Action	Purpose
Step 1	switch# **xmlin**	
Step 2	switch(xmlin)# **configure terminal**	Enters global configuration mode.
Step 3	Configuration commands	Converts configuration commands to NETCONF format.
Step 4	(Optional) switch(config)(xmlin)# **end**	Generates the corresponding <edit-config> request. **Note** Enter the **end** command to finish the current XML configuration before you generate an XML instance for a **show** command.
Step 5	(Optional) switch(config-if-verify)(xmlin)# **show** *commands*	Converts **show** commands to NETCONF format.
Step 6	(Optional) switch(config-if-verify)(xmlin)# **exit**	Returns to EXEC mode.

Converting Show Command Output to XML

You can convert the output of show commands to XML.

Before you begin

Make sure that all features for the commands you want to convert are installed and enabled on the device. Otherwise, the commands fail.

You can use the **terminal verify-only** command to verify that a feature is enabled without entering it on the device.

Make sure that all required hardware for the commands you want to convert are present on the device. Otherwise, the commands fail.

Make sure that the XMLIN tool is installed.

Procedure

	Command or Action	Purpose
Step 1	switch# *show-command* | **xmlin**	Enters global configuration mode. **Note** You cannot use this command with configuration commands.

Configuration Examples for XMLIN

The following example shows how the XMLIN tool is installed on the device and used to convert a set of configuration commands to an <edit-config> instance.

```
switch# xmlin
**************************************
Loading the xmlin tool. Please be patient.
**************************************
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright ©) 2002-2013, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php

switch(xmlin)# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)(xmlin)# interface ethernet 2/1
% Success
switch(config-if-verify)(xmlin)# cdp enable
% Success
switch(config-if-verify)(xmlin)# end
<?xml version="1.0"?>
<nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:6.2.2.:configure_"
xmlns:m="http://www.cisco.com/nxos:6.2.2.:_exec"
xmlns:m1="http://www.cisco.com/nxos:6.2.2.:configure__if-eth-base" message-id="1">
  <nf:edit-config>
     <nf:target>
      <nf:running/>
  </nf:target>
  <nf:config>
    <m:configure>
      <m:terminal>
        <interface>
           <__XML__PARAM__interface>
              <__XML__value>Ethernet2/1</__XML__value>
              <m1:cdp>
                <m1:enable/>
              </m1:cdp>
            </__XML__PARAM__interface>
```

```
          </interface>
         </m:terminal>
        </m:configure>
       </nf:config>
      </nf:edit-config>
     </nf:rpc>
    ]]>]]>
```

The following example shows how to enter the **end** command to finish the current XML configuration before you generate an XML instance for a **show** command.

```
switch(xmlin)# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)(xmlin)# interface ethernet 2/1
switch(config-if-verify)(xmlin)# show interface ethernet 2/1
******************************************************
Please type "end" to finish and output the current XML document before building a new one.
******************************************************
% Command not successful

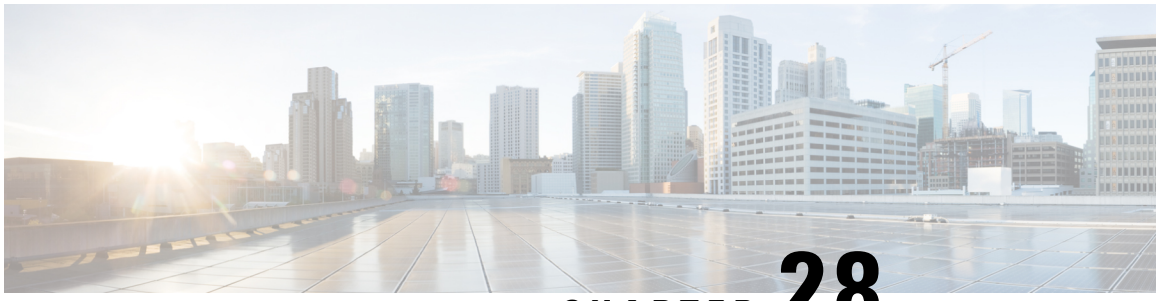
switch(config-if-verify)(xmlin)# end
<?xml version="1.0"?>
<nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:6.2.2.:configure_"
xmlns:m="http://www.cisco.com/nxos:6.2.2.:_exec" message-id="1">
    <nf:edit-config>
      <nf:target>
         <nf:running/>
      </nf:target>
      <nf:config>
         <m:configure>
           <m:terminal>
              <interface>
                 <__XML__PARAM__interface>
                    <__XML__value>Ethernet2/1</__XML__value>
                 </__XML__PARAM__interface>
              </interface>
            </m:terminal>
           </m:configure>
         </nf:config>
      </nf:edit-config>
    </nf:rpc>
   ]]>]]>

switch(xmlin)# show interface ethernet 2/1
<?xml version="1.0"?>
<nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:6.2.2.:if_manager" message-id="1">
  <nf:get>
    <nf:filter type="subtree">
      <show>
      <interface>
        <__XML__PARAM__ifeth>
           <__XML__value>Ethernet2/1</__XML__value>
        </__XML__PARAM__ifeth>
      </interface>
     </show>
   </nf:filter>
  </nf:get>
</nf:rpc>
]]>]]>
switch(xmlin)# exit
switch#
```

The following example shows how you can convert the output of the **show interface brief** command to XML.

```
switch# show interface brief | xmlin
<?xml version="1.0"?>
<nf:rpc xmlns:nf="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns="http://www.cisco.com/nxos:6.2.2.:if_manager"

message-id="1">
  <nf:get>
    <nf:filter type="subtree">
       <show>
          <interface>
              <brief/>
          </interface>
       </show>
     </nf:filter>
   </nf:get>
</nf:rpc>
]]>]]>
```



CHAPTER 28

RESTConf Agent

- About the RESTCONF Agent, on page 271
- Configuring the RESTConf Agent, on page 272
- Using the RESTCONF Agent, on page 272
- Troubleshooting the RESTCONF Agent, on page 273

About the RESTCONF Agent

Cisco NX-OS RESTCONF is an 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:

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")

Configuring the RESTConf Agent

The RESTConf Agent does not require any configuration in the configuration file (**/etc/mtx.conf**).

Using the RESTCONF Agent

General Commands

- Configure the following commands to enable HTTP or HTTPS access:
  - **feature nxapi**
  - **nxapi http port 80**
  - **nxapi https port 443**

General Control Commands

You can enable or disable the RESTCONF Agent [**no**] **feature restconf** command.

Viewing the Agent Status

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#
```

Sending a POST Request to the Server Using Curl

```
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:147500853169574134
Status: 201 Created
Location: /System/bgp-items/inst-items/dom-items/Dom-list=default/always/rtrId/
```

Sending a GET Request to the Server Using Curl

```
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 %
```

Troubleshooting the RESTCONF Agent

Troubleshooting Connectivity

- Enable the web server by issuing the **feature nxapi** command.
- Ensure that the **nxapi http port 80** command is configured to open up the port for HTTP
- Ensure that the **nxapi https port 443** command is configured to open up the port for HTTPS.
- Ping the management port of the switch to verify that the switch is reachable.

Troubleshooting Errors

The following shows a common error message and offers guidelines for resolving it.

Error Message: `Sorry, the page you are looking for is currently unavailable`

- If you receive this message soon after sending a request (for example, seconds), verify the following:
  - 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.

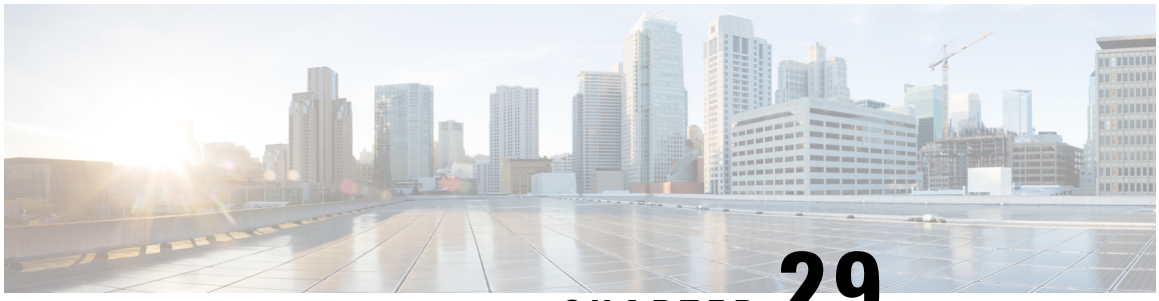
    ```
    switch-1# show nxapi
    nxapi enabled
    HTTP Listen on port 80
    HTTPS Listen on port 443
    ...
    switch-1#
    ```

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**.



CHAPTER 29

gRPC Agent

This chapter contains the following sections:

- About the gRPC Agent, on page 275
- Configuring the gRPC Agent for Cisco NX-OS Release 9.3(2) and Earlier, on page 276
- Using the gRPC Agent, on page 278
- Troubleshooting the gRPC Agent, on page 280
- gRPC Protobuf File, on page 280

About the gRPC Agent

The Cisco NX-OS gRPC protocol defines a mechanism through which a network device can be managed and its configuration data can be retrieved and installed. The protocol exposes a complete and formal Application Programming Interface (API) that clients can use to manage device configurations.

The Cisco NX-OS gRPC protocol uses a remote procedure call (RPC) paradigm where an external client manipulates device configurations utilizing Google Protocol Buffer (GPB)-defined API calls along with their service- specific arguments. These GPB-defined APIs transparently cause an RPC call to the device that return replies in the same GPB-defined API context.

The gRPC Agent provides a secure transport through TLS and user authentication and authorization through AAA.

The functional objective of the Cisco NX-OS gRPC protocol is to mirror that provided by NETCONF, particularly in terms of both stateless and stateful configuration manipulation for maximum operational flexibility.

The Cisco NX-OS gRPC Agent supports the following protocol operations:

- Get
- GetConfig
- GetOper
- EditConfig
- StartSession
- CloseSession
- KillSession

The gRPC Agent supports two types of operations:

- **Stateless operations** are performed entirely within a single message without creating a session.
- **Stateful operations** are performed using multiple messages. The following is the sequence of operations that are performed:
  1. Start the session. This action acquires a unique session ID.
  2. Perform session tasks using the session ID.
  3. Close the session. This action invalidates the session ID.

The following are the supported operations. See the Appendix for their RPC definitions in the **.proto** file that is exported by the gRPC Agent.

Operation	Description
StartSession	Starts a new session between the client and server and acquires a unique session ID.
EditConfig	Writes the specified YANG data subset to the target datastore.
GetConfig	Retrieves the specified YANG configuration data subset from the source datastore.
GetOper	Retrieves the specified YANG operational data from the source datastore.
Get	Retrieves the specified YANG configuration and operational data from the source datastore.
KillSession	Forces the termination of a session.
CloseSession	Requests graceful termination of a session.

GetConfig, GetOper, and Get are stateless operations so don't require a session ID.

EditConfig can be either stateless or stateful. For a stateless operation, specify the SessionID as 0. For a stateful operation, a valid (nonzero) SessionID is required.

The gRPC Agent supports timeout for the sessions. The idle timeout for sessions can be configured on the device, after which idle sessions are closed and deleted.

Configuring the gRPC Agent for Cisco NX-OS Release 9.3(2) and Earlier

The gRPC Agent supports the following configuration parameters under the **[grpc]** section in the configuration file (**/etc/mtx.conf**).

Parameter	Description
idle_timeout	(Optional) Specifies the timeout in minutes after which idle client sessions are disconnected. The default timeout is 5 minutes. A value of 0 disables timeout.
limit	(Optional) Specifies the number of maximum simultaneous client sessions. The default limit is 5 sessions. The range is from 1 through 50.
lport	(Optional) Specifies the port number on which the gRPC Agent listens. The default port is 50051.
key	Specifies the key file location for TLS authentication. The default location is **/opt/mtx/etc/grpc.key**.
cert	Specifies the certificate file location for TLS authentication. The default location is **/opt/mtx/etc/grpc.pem**. Beginning with NX-OS release 9.3(1), some changes are made to the certificate for gRPC Agent. See "About Certificates" below.
security	Specifies the type of secure connection. Valid choices are: • **TLS** for TLS • **NONE** for an insecure connection

About Certificates

For new switch deployments with Cisco NX-OS 9.3(1), the switch generates a one-day temporary certificate to allow you enough time to install your custom certificate. This temporary one-day certificate replaces the hard-coded default certificate that was generated in previous versions of Cisco NX-OS. The temporary one-day certificate gets regenerated when a switch reload or system switchover event occurs.

If your switch already has a custom certificate that is installed, for example, if you are upgrading from a previous NX-OS version to NX-OS 9.3(1), your existing certificate is retained and used after upgrade.

The following is an example of the **[grpc]** section in the configuration file:

```
[grpc]
mtxadapter=/opt/mtx/lib/libmtxadaptergrpc.1.0.1.so
idle_timeout=10
limit=1
lport=50051
```

```
security=TLS
cert=/etc/grpc.pem
key=/etc/grpc.key
```



Note For the following, if you have a hard-coded key, you must delete it and install your custom key in a persistent location.

Delete the Hard-Coded Certificate

If your switch was deployed with a version of NX-OS earlier than 9.3(1), you must delete the hard-coded certificate.

Change the Certificate Location

The previous default location for the certificate (`/etc/mtx.conf file`) is no longer persistent.

By default, the `/etc/mtx.conf.user` file is not present, so you will need to create it and specify the new location of the key and certificate so that the change is persistent. For example:

```
[grpc]
cert="new location"
key="new location"
```



Note For switches that have dual supervisor modules, you must upload the certificate files and modify the `/etc/mtx.conf.user` file on both the Active and Standby supervisors. If you do not do this on both supervisors, the certificate files are not persistent after a switchover.

Replace an Expired Certificate

If your temporary, one-day certificate is expired, delete it and replace it with a new certificate.

Reload the Configuration File

Use the appropriate option, depending on whether you want to reload the switch:

- If you are restarting the gRPC agent to use the new certificate without reloading the switch:
  1. Add the certificate and key to `/etc/mtx.conf`.
  2. For the modified configuration file to take effect, issue **no feature grpc** to disable the gRPC Agent, then **feature grpc** to reenable it.
- If you want to reload the switch, the changes to `/etc/mtx.conf.user` are sufficient. You do not need to modify `/etc/mtx.conf`.

Using the gRPC Agent

General Commands

You can enable or disable the gRPC Agent by issuing the [**no**] **feature grpc** command.

Example: A Basic Yang Path in JSON Format

```
client-host % cat payload.json
{
  "namespace": "http://cisco.com/ns/yang/cisco-nx-os-device",
  "System": {
    "bgp-items": {
      "inst-items": {
        "dom-items": {
          "Dom-list": {
            "name": "default",
            "rtrId": "7.7.7.7",
            "holdIntvl": "100"
          }
        }
      }
    }
  }
}
```



Note The JSON structure has been pretty-formatted here for readability.

Sending an EditConfig Request to the Server

```
client-host % ./grpc_client -username=admin -password=cisco -operation=EditConfig
-e_oper=Merge -def_op=Merge -err_op=stop-on-error -infile=payload.json -reqid=1
-source=running -tls=true -serverAdd=192.0.20.123 -lport=50051

#####################################################
Starting the client service
#####################################################
TLS set true for client requests1ems.cisco.com
TLS FLAG:1
192.0.20.123:50051
All the client connections are secured
Sending EditConfig request to the server
sessionid is
0
reqid:1
{"rpc-reply":{"ok":""}}
```

Sending a GetConfig Request to the Server

```
client-host % ./grpc_client -username=admin -password=cisco -operation=GetConfig
-infile=payload.json -reqid=1 -source=running -tls=true -serverAdd=192.0.20.123 -lport=50051

#####################################################
Starting the client service
#####################################################
TLS set true for client requests1ems.cisco.com
TLS FLAG:1
192.0.20.123:50051
All the client connections are secured
Sending GetConfig request to the server
in get config
```

```
Got the response from the server
########################################
Yang Json is:
########################################
{"rpc-reply":{"data":{"System":{"bgp-items":{"inst-items":{"dom-items":{"Dom-list":{"name":"default","rtrId":"7.7.7.7","holdIntvl":"100"}}}}}}}}
########################################
client-host %
```

Troubleshooting the gRPC Agent

Troubleshooting Connectivity

- From a client system, verify that the agent is listening on the port. For example:

```
client-host % nc -z 192.0.20.222 50051
Connection to 192.0.20.222 50051 port [tcp/*] succeeded!
client-host % echo $?
0
client-host %
```

- In the NX-OS, check the gRPC agent status by issuing **show feature | grep grpc**.

gRPC Protobuf File

The gRPC Agent exports the supported operations and data structures in the proto definition file at `/opt/mtx/etc/nxos_grpc.proto`. The file is included in the gRPC Agent RPM. The following shows the definitions:

```
// Copyright 2016, Cisco Systems Inc.
// All rights reserved.

syntax = "proto3";

package NXOSExtensibleManagabilityService;

// Service provided by Cisco NX-OS gRPC Agent
service gRPCConfigOper {

    // Retrieves the specified YANG configuration data subset from the
    // source datastore
    rpc GetConfig(GetConfigArgs) returns(stream GetConfigReply) {};

    // Retrieves the specified YANG operational data from the source datastore
    rpc GetOper(GetOperArgs) returns(stream GetOperReply) {};

    // Retrieves the specified YANG configuration and operational data
    // subset from the source datastore
    rpc Get(GetArgs) returns(stream GetReply){};

    // Writes the specified YANG data subset to the target datastore
    rpc EditConfig(EditConfigArgs) returns(EditConfigReply) {};

    // Starts a new session between the client and server and acquires a
```

```
    // unique session ID
    rpc StartSession(SessionArgs) returns(SessionReply) {};

    // Requests graceful termination of a session
    rpc CloseSession(CloseSessionArgs) returns (CloseSessionReply) {};

    // Forces the termination of a session
    rpc KillSession(KillArgs) returns(KillReply) {};

// Unsupported; reserved for future
    rpc DeleteConfig(DeleteConfigArgs) returns(DeleteConfigReply) {};

    // Unsupported; reserved for future
    rpc CopyConfig(CopyConfigArgs) returns(CopyConfigReply) {};

    // Unsupported; reserved for future
    rpc Lock(LockArgs) returns(LockReply) {};

    // Unsupported; reserved for future
    rpc UnLock(UnLockArgs) returns(UnLockReply) {};

    // Unsupported; reserved for future
    rpc Commit(CommitArgs) returns(CommitReply) {};

    // Unsupported; reserved for future
    rpc Validate(ValidateArgs) returns(ValidateReply) {};

    // Unsupported; reserved for future
    rpc Abort(AbortArgs) returns(AbortReply) {};

}

message GetConfigArgs
{
    // JSON-encoded YANG data to be retrieved
    string YangPath = 1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;

    // (Optional) Specifies the source datastore; only "running" is supported.
    // Default is "running".
    string Source = 3;
}

message GetConfigReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // JSON-encoded YANG data that was retrieved
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message GetOperArgs
{
    // JSON-encoded YANG data to be retrieved
    string YangPath = 1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;
```

```
}

message GetOperReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // JSON-encoded YANG data that was retrieved
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message GetArgs
{
    // JSON-encoded YANG data to be retrieved
    string YangPath=1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;
}

message GetReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // JSON-encoded YANG data that was retrieved
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message EditConfigArgs
{
    // JSON-encoded YANG data to be edited
    string YangPath = 1;

    // Specifies the operation to perform on teh configuration datastore with
    // the YangPath data.  Possible values are:
    //   create
    //   merge
    //   replace
    //   delete
    //   remove
    // If not specified, default value is "merge".
    string Operation = 2;

    // A unique session ID acquired from a call to StartSession().
    // For stateless operation, this value should be set to 0.
    int64 SessionID = 3;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 4;

    // (Optional) Specifies the target datastore; only "running" is supported.
    // Default is "running".
    string Target = 5;

    // Specifies the default operation on the given object while traversing
    // the configuration tree.
    // The following operations are possible:
```

```
    //   merge:     merges the configuration data with the target datastore;
    //              this is the default.
    //   replace:   replaces the configuration data with the target datastore.
    //   none:      target datastore is unaffected during the traversal until
    //              the specified object is reached.
    string DefOp = 6;

    // Specifies the action to be performed in the event of an error during
    // configuration.  Possible values are:
    //   stop
    //   roll-back
    //   continue
    // Default is "roll-back".
    string ErrorOp = 7;
}

message EditConfigReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // If EditConfig is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message DeleteConfigArgs
{
    // A unique session ID acquired from a call to StartSession().
    // For stateless operation, this value should be set to 0.
    int64 SessionID = 1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;

    // (Optional) Specifies the target datastore; only "running" is supported.
    // Default is "running".
    string Target = 3;
}

message DeleteConfigReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // If DeleteConfig is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message CopyConfigArgs
{
    // A unique session ID acquired from a call to StartSession().
    // For stateless operation, this value should be set to 0.
    int64 SessionID = 1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;

    // (Optional) Specifies the source datastore; only "running" is supported.
```

```
    // Default is "running".
    string Source = 3;

    // (Optional) Specifies the target datastore; only "running" is supported.
    // Default is "running".
    string Target = 4;
}

message CopyConfigReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // If CopyConfig is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message LockArgs
{
    // A unique session ID acquired from a call to StartSession().
    int64 SessionID = 1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID=2;

    // (Optional) Specifies the target datastore; only "running" is supported.
    // Default is "running".
    string Target = 3;
}

message LockReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // If Lock is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message UnLockArgs
{
    // A unique session ID acquired from a call to StartSession().
    int64 SessionID = 1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;

    // (Optional) Specifies the target datastore; only "running" is supported.
    // Default is "running".
    string Target = 3;
}

message UnLockReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // If UnLock is successful, YangData contains a JSON-encoded "ok" response.
```

```
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message SessionArgs
{
    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 1;
}

message SessionReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;
    int64 SessionID = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message CloseSessionArgs
{
    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 1;

    // A unique session ID acquired from a call to StartSession().
    int64 SessionID = 2;
}

message CloseSessionReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // If CloseSession is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message KillArgs
{
    // A unique session ID acquired from a call to StartSession().
    int64 SessionID = 1;

    int64 SessionIDToKill = 2;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 3;
}

message KillReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // If Kill is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
```

```
}

message ValidateArgs
{
    // A unique session ID acquired from a call to StartSession().
    int64 SessionID = 1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;
}

message ValidateReply
{
    // The request ID specified in the request.
    int64 ReqID = 1;

    // If Validate is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message CommitArgs
{
    // A unique session ID acquired from a call to StartSession().
    int64 SessionID = 1;

    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;
}

message CommitReply
{
    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 1;

    // If Commit is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
}

message AbortArgs
{
    // A unique session ID acquired from a call to StartSession().
    int64 SessionID = 1;

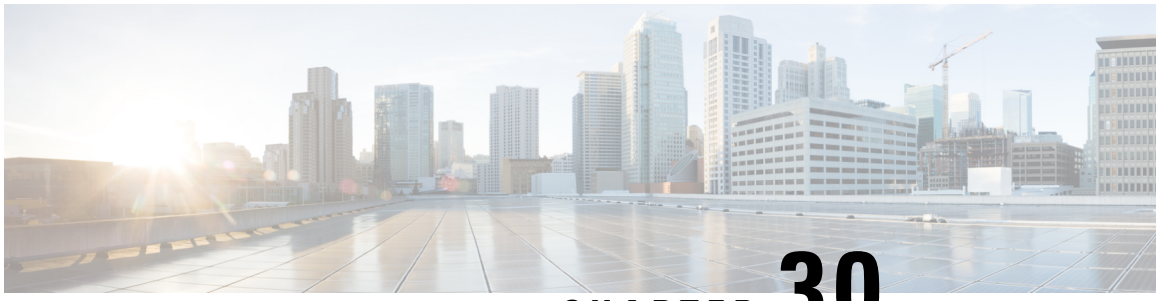
    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 2;
}

message AbortReply
{
    // (Optional) Specifies the request ID.  Default value is 0.
    int64 ReqID = 1;

    // If Abort is successful, YangData contains a JSON-encoded "ok" response.
    string YangData = 2;

    // JSON-encoded error information when request fails
    string Errors = 3;
```

```
}
```



CHAPTER 30

gNMI - gRPC Network Management Interface

This chapter contains the following topics:

- About gNMI, on page 289
- gNMI RPC and SUBSCRIBE, on page 290
- Guidelines and Limitations for gNMI, on page 291
- Configuring gNMI, on page 293
- Configuring Server Certificate, on page 294
- Generating Key/Certificate Examples , on page 295
- Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(2) and Earlier, on page 295
- Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(3) and Later, on page 296
- Verifying gNMI, on page 298
- Clients, on page 303
- Sample DME Subscription - PROTO Encoding, on page 303
- Capabilities, on page 304
- Get, on page 308
- Set, on page 309
- Subscribe, on page 310
- Streaming Syslog, on page 314
- Troubleshooting, on page 320

About gNMI

gNMI uses gRPC (Google Remote Procedure Call) as its transport protocol.

Cisco NX-OS supports gNMI for dial-in subscription to telemetry applications running on the Cisco Nexus 3000 Series switches. Although past release supported telemetry events over gRPC, the switch pushed the telemetry data to the telemetry receivers. This method was called dial out.

With gNMI, applications can pull information from the switch. They subscribe to specific telemetry services by learning the supported telemetry capabilities and subscribing to only the telemetry services that it needs.

Table 20: Supported gNMI RPCs

gNMI RPC	Supported
Capabilities	Yes
Get	Yes
Set	Yes
Subscribe	Yes

gNMI RPC and SUBSCRIBE

The NX-OS 9.3(1) release supports gNMI version 0.5.0. Cisco NX-OS Release 9.3(1) supports the following parts of gNMI version 0.5.0.

Table 21: 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 604800 seconds. The default sample interval is 10 seconds.
	On_Change	Yes	The switch sends current values as its initial state, but then updates the values only when changes, such as create, modify, or delete occur to any of the specified paths.
	Target_Defined	No	

Optional SUBSCRIBE Flags

For the SUBSCRIBE option, some optional flags are available that modify the response to the options listed in the table. In release 9.3(1), the updates_only optional flag is supported, which is applicable to ON_CHANGE subscriptions. If this flag is set, the switch suppresses the initial snapshot data (current state) that is normally sent with the first response.

The following flags are not supported:

- aliases
- allow_aggregation
- extensions
- heart-beat interval
- prefix
- qos
- suppress_redundant

Guidelines and Limitations for gNMI

Following are the guidelines and limitations for gNMI:

- Beginning with Cisco NX-OS Release 9.3(5), Get and Set are supported.
- gNMI queries do not support wildcards in paths.
- If you configure a **prefix-list** in CLI mode and issue an Edit, Get, or Get-Config, for the OpenConfig model, it is not supported.

  If you configure a **prefix-list** in CLI Mode and issue a NETCONF/gNMI Notification or Subscription from the OpenConfig model, it is not supported.
- When you attempt to subscribe an OpenConfig routing policy with a preexisting CLI configuration like the following, it returns empty values due to the 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

using the xpath

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
```

- When you enable gRPC on both the management VRF and default VRF and later disable on the default VRF, the gNMI notifications 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 and any existing gRPC configuration commands. For example, **grpc certificate** or **grpc port**. You must also resubscribe to any existing notifications on the management VRF.

- When you attempt to subscribe an OpenConfig routing policy with a preexisting CLI configuration like the following, it returns empty values due to the 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

using the xpath

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
```

- Beginning with Cisco NX-OS Release 9.3(3), 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 will reload. Once the switch is up again, you must reconfigure the gRPC custom certificate.



Note This applies when entering the **grpc port** and **grpc certificate** commands.

- Use of origin, use_models, or both, is optional for gNMI subscriptions.
- Beginning with Cisco NX-OS Release 9.3(3), Subscribe supports the OpenConfig model.
- The gNMI feature supports Subscribe and Capability as options of the gNMI service.
- The feature supports JSON and gnmi.proto encoding. The feature does not support protobuf.any encoding.
- Each gNMI message has a maximum size of 12 MB. If the amount of collected data exceeds the 12 MB maximum, the collected data is dropped.

  You can avoid this situation by creating more focused subscriptions that handle smaller, more granular data-collection sets. So, instead of subscribing to one higher-level path, create multiple subscriptions for different, lower-level parts of the path.
- Across all subscriptions, there is support of up to 150K aggregate MOs. Subscribing to more MOs can lead to collection data drops.
- All paths within the same subscription request must have the same sample interval. If the same path requires different sample intervals, create multiple subscriptions.
- The gRPC process that supports gNMI uses the HIGH_PRIO cgroup, which limits the CPU usage to 75% of CPU and memory to 1.5 GB.
- The **show grpc gnmi** command has the following considerations:
  - The gRPC agent retains gNMI calls for a maximum of 1 hour after the call has ended.
  - If the total number of calls exceeds 2000, the gRPC agent purges ended calls based an internal cleanup routine.

The gRPC server runs in the management VRF. As a result, the gRPC process communicates only in this VRF forcing the management interface to support all gRPC calls.

gRPC functionality now includes the default VRF for a total of two gRPC servers on each Cisco Nexus 3000 Series switch. You can run one gRPC server in each VRF, or run only one gRPC server in the management

VRF. Supporting a gRPC in the default VRF adds flexibility to offload processing gRPC calls from the management VRF, where significant traffic load might not be desirable.

If two gRPC servers are configured, be aware of the following:

- VRF boundaries are strictly enforced, so each gRPC server processes 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.

Configuring gNMI

Configure the gNMI feature through the **grpc gnmi** commands.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch-1# configure terminal` `switch-1(config)#`	Enters global configuration mode.
Step 2	**feature grpc** **Example:** `switch-1# feature grpc` `switch-1(config)#`	Enables the gRPC agent, which supports the gNMI interface for dial-in.
Step 3	**grpc gnmi max-concurrent-call** *number* **Example:** `switch-1(config)# grpc gnmi max-concurrent-call 16` `switch-1(config)#`	Sets the limit of simultaneous dial-in calls to the gNMI server on the switch. Configure a limit from 1 through 16. The default limit is 8. The maximum value that you configure is for each VRF. If you set a limit of 16 and gNMI is configured for both management and default VRFs, each VRF supports 16 simultaneous gNMI calls. This command does not affect and ongoing or in-progress gNMI calls. Instead, gRPC enforces the limit on new calls, so any in-progress calls are unaffected and allowed to complete. **Note** The configured limit does not affect the gRPCConfigOper service.

Configuring Server Certificate

gNMI communicates over gRPC and uses TLS to secure the channel between the switch and the client. The default hard-coded gRPC certificate is no longer shipped with the switch. The default behavior is a self-signed key and certificate which is generated on the switch as shown below with an expiration date of one day.

```
switch# show grpc gnmi service statistics

=============
gRPC Endpoint
=============

Vrf : management
Server address : [::]:50051

Cert notBefore : Fri Jun 21 10:54:42 PDT 2019
Cert notAfter : Sat Jun 22 10:54:42 PDT 2019

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

The keys are generated in the location below

switch# pwd
/opt/mtx/etc
switch# ls -al
total 20
drwxr-xr-x 4 root floppy 140 Jun 21 10:54 .
drwxr-xr-x 7 root floppy 140 Jun 20 02:43 ..
-rw------- 1 root root 1704 Jun 21 10:54 grpc.key
-rw------- 1 root root 1107 Jun 21 10:54 grpc.pem
drwxr-xr-x 2 root floppy 60 Jun 21 10:54 init.d
-rw-r--r-- 1 root floppy 10712 Jun 20 02:43 nxos_grpc.proto
drwxr-xr-x 2 root floppy 60 Jun 21 10:54 proto-1.0
switch#
```

With an expiration of one day, you can use this temporary certificate for quick testing. For long term a new key/certificate must be generated.



Note After the certificate expires, there are two ways to have the key/certificate to regenerate:

- Reload the switch.
- Manually delete the key/certificate in the `/opt/mtx/etc` folder and enter the **no feature grpc** and **feature grpc** commands.

Generating Key/Certificate Examples

Follow these examples to generate Key/Certificates:

- Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(2) and Earlier, on page 295
- Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(3) and Later, on page 296

Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(2) and Earlier

The following is an example for generating key/certificate:

For more information on generating identify certificates, see the Installing Identity Certificates section of the *Cisco Nexus 3000 Series NX-OS Security Configuration Guide, Release 9.3(x)*.

Procedure

Step 1 Generate the selfsigned key and pem files.

a) switch# openssl req -x509 -newkey rsa:2048 -keyout self_sign2048.key -out self_sign2048.pem -days 365 -nodes

Step 2 After generating the key and pem files, modify the mtx.conf.user files in the Bash shell to have the gRPC service pick up the certificates.

```
[grpc]
key = /bootflash/self-sign2048.key
cert = /bootflash/self-sign2048.pem
```

Step 3 Reload the box to have the gRPC service pick up the certificate.

Step 4 Verify gRPC is now using the certificate.

```
switch# show grpc gnmi service statistics

=============
gRPC Endpoint
=============

Vrf : management
Server address : [::]:50051

Cert notBefore : Nov 5 16:48:58 2015 GMT
Cert notAfter : Nov 5 16:48:58 2035 GMT

Max concurrent calls : 16
Listen calls : 1
Active calls : 0

Number of created calls : 953
```

```
Number of bad calls : 0

Subscription stream/once/poll : 476/238/238

Max gNMI::Get concurrent : 5
Max grpc message size : 8388608
gNMI Synchronous calls : 10
gNMI Synchronous errors : 0
gNMI Adapter errors : 0
gNMI Dtx errors : 0
```

Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(3) and Later

The following is an example for generating key/certificate.



Note We recommend using this method for configuring certificates.

For more information on generating identify certificates, see the Installing Identity Certificates section of the *Cisco Nexus 3000 Series NX-OS Security Configuration Guide, Release 9.3(x)*.

Procedure

Step 1 Generate the selfsigned key and pem files.

a) switch# openssl req -x509 -newkey rsa:2048 -keyout self_sign2048.key -out self_sign2048.pem -days 365 -nodes

Step 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
```

Step 3 Set up the trustpoint CA Association by inputting the pkcs12 bundle into the trustpoint.

```
switch(config)# crypto ca trustpoint mytrustpoint
switch(config-trustpoint)# crypto ca import mytrustpoint pkcs12 self_sign2048.pfx Ciscolab123!
  <- input the pkcs12 bundle into the trustpoint.
```

Step 4 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
```

Step 5 Configure gRPC to use the trustpoint.

```
switch(config)# grpc certificate mytrustpoint
switch(config)# show run grpc

!Command: show running-config grpc
!Running configuration last done at: Thu Jul  2 12:24:02 2020
!Time: Thu Jul  2 12:24:05 2020

version 9.3(5) Bios:version 05.38
feature grpc

grpc gnmi max-concurrent-calls 16
grpc use-vrf default
grpc certificate mytrustpoint
```

Step 6 Verify gRPC is now using the certificate.

```
switch# show grpc gnmi service statistics

=============
gRPC Endpoint
=============

Vrf : management
Server address : [::]:50051

Cert notBefore : Nov 5 16:48:58 2015 GMT
Cert notAfter : Nov 5 16:48:58 2035 GMT

Max concurrent calls : 16
Listen calls : 1
Active calls : 0

Number of created calls : 953
Number of bad calls : 0

Subscription stream/once/poll : 476/238/238

Max gNMI::Get concurrent : 5
Max grpc message size : 8388608
gNMI Synchronous calls : 10
```

```
gNMI Synchronous errors : 0
gNMI Adapter errors : 0
gNMI Dtx errors : 0
```

Verifying gNMI

To verify the gNMI configuration, enter the following command:

Command	Description
show grpc gnmi service statistics	Displays a summary of the agent running status, respectively for the management VRF, or the default VRF (if configured). It also displays: • Basic overall counters • Certificate expiration time **Note** If the certificate is expired, the agent cannot accept requests.
show grpc gnmi rpc summary	Displays the following: • Number of capability RPCs received. • Capability RPC errors. • Number of Get RPCs received. • Get RPC errors. • Number of Set RPCs received. • Set RPC errors. • More error types and counts.

Command	Description
show grpc gnmi transactions	The **show grpc gnmi transactions** command is the most dense and contains considerable information. It is a 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: • RPC – This shows the type of RPC that was received (Get, Set, Capabilities) • DataType – For a Get only. Has values ALL, CONFIG, and STATE. • Session – shows 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 -- shows timestamp of when the RPC was received by the gNMI handler. • Duration – time delta in ms from receiving the request to giving response. • Status – the status code of the operation returned to the client `(0 = Success, !0 == error)` This section is data that is kept per path within a single gNMI transaction. For example, a single Get or Set • subtype – for a Set RPC, shows the specific operation that is requested per path (Delete, Update, Replace). For Get, there is no subtype. • dtx – shows that this path was processed in DTX "fast" path or not. A dash '-' means no, an asterisk '*' means yes. • st – Status for this path. The meaning is as follows: • OK: path is valid and processed by infra successfully. • ERR: path is either invalid or generated error by infra • --: path not processed yet, might or might not be valid and has not been sent to infra yet. • path – the path

show grpc gnmi service statistics Example

```
=============
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
```

show grpc gnmi rpc summary Example

```
=============
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
```

show grpc gnmi transactions Example

```
=============
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          -          2293989720      04/01 07:43:45       183          0
subtype: dtx:  st: path:
Replace  -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo6]

Set          -          2297110560      04/01 07:43:41       184          0
subtype: dtx:  st: path:
Update   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo7]

Set          -          0               04/01 07:43:39       0            10

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=lo791]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo792]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo793]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo794]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo795]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo796]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo797]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo798]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo799]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo800]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo801]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo802]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo803]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo804]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo805]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo806]
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          -          2294408864      04/01 07:43:17       176          13
subtype: dtx:  st: path:
Delete   -     ERR /System/intf-items/lb-items/LbRtdIf-list/descr


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]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo777]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo778]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo779]
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo780]
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=lo4]/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

Set          -          3551604840      04/01 07:42:54       35           0
subtype: dtx:  st: path:
Delete   -     OK  /System/intf-items/lb-items/LbRtdIf-list[id=lo1]

Set          -          2362201592      04/01 07:42:52       13           13
subtype: dtx:  st: path:
Delete   -     ERR /System/intf-items/lb-items/LbRtdIf-list[id=lo3]/lbrtdif-items
/operSt


Set          -          0               04/01 07:42:47       0            3
subtype: dtx:  st: path:
Delete   -     ERR /System/*

Set          -          2464158360      04/01 07:42:46       172          3
subtype: dtx:  st: path:
Delete   -     ERR /system/processes/shabang

Set          -          2295440864      04/01 07:42:46       139          3
subtype: dtx:  st: path:
Delete   -     ERR /System/invalid/path

Set          -          3495739048      04/01 07:42:44       10           0


Get          ALL        3444580832      04/01 07:42:40       3            0
subtype: dtx:  st: path:
-        -     OK  /System/bgp-items/inst-items/disPolBatch

Get          ALL        0               04/01 07:42:36       0            3
subtype: dtx:  st: path:
-        -     --  /system/processes/process[pid=1]

Get          ALL        3495870472      04/01 07:42:36       2            0
subtype: dtx:  st: path:
-        *     OK  /system/processes/process[pid=1]

Get          ALL        2304485008      04/01 07:42:36       33           0
subtype: dtx:  st: path:
-        *     OK  /system/processes

Get          ALL        2464159088      04/01 07:42:36       251          0
subtype: dtx:  st: path:
-        -     OK  /system
```

```
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
```

Clients

There are available clients for gNMI. One such client is located at https://github.com/influxdata/telegraf/tree/master/plugins/inputs/cisco_telemetry_gnmi.

Sample DME Subscription - PROTO Encoding

```
gnmi-console --host >iip> --port 50051 -u <user> -p <pass> --tls --
operation=Subscribe --rpc /root/gnmi-console/testing_bl/once/61_subscribe_bgp_dme_gpb.json

[Subscribe]--------------------------------
### Reading from file ' /root/gnmi-console/testing_bl/once/61_subscribe_bgp_dme_gpb.json '
Wed Jun 26 11:49:17 2019
### Generating request : 1 -----------
### Comment : ONCE request
### Delay : 2 sec(s) ...
### Delay : 2 sec(s) DONE
subscribe {
subscription {
path {
origin: "DME"
elem {
name: "sys"
}
elem {
name: "bgp"
}
}
mode: SAMPLE
}
mode: ONCE
use_models {
name: "DME"
organization: "Cisco Systems, Inc."
version: "1.0.0"
}
encoding: PROTO
}
Wed Jun 26 11:49:19 2019
Received response 1 --------------------------
update {
timestamp: 1561574967761
prefix {
elem {
name: "sys"
}
elem {
name: "bgp"
}
}
```

```
update {
path {
elem {
}
elem {
name: "version_str"
}
}
val {
string_val: "1.0.0"
}
}
update {
path {
elem {
}
elem {
name: "node_id_str"
}
}
val {
string_val: "n9k-tm2"
}
}
update {
path {
elem {
}
elem {
name: "encoding_path"
}
}
val {
string_val: "sys/bgp"
}
}
update {
path {
elem {
}
elem {
/Received -------------------------------------
Wed Jun 26 11:49:19 2019
Received response 2 -------------------------
sync_response: true
/Received -------------------------------------
(_gnmi) [root@tm-ucs-1 gnmi-console]#
```

Capabilities

About Capabilities

The Capabilities RPC returns the list of capabilities of the gNMI service. The response message to the RPC request includes the gNMI service version, the versioned data models, and data encodings supported by the server.

Guidelines and Limitations for Capabilities

Following are the guidelines and limitations for Capabilities:

- Beginning with Cisco NX-OS Release 9.3(3), Capabilities supports the OpenConfig model.
- The gNMI feature supports Subscribe and Capability as options of the gNMI service.
- The feature supports JSON and gnmi.proto encoding. The feature does not support protobuf.any encoding.
- Each gNMI message has a maximum size of 12 MB. If the amount of collected data exceeds the 12-MB maximum, the collected data is dropped.

  You can avoid this situation by creating more focused subscriptions that handle smaller, more granular data-collection sets. So, instead of subscribing to one higher-level path, create multiple subscriptions for different, lower-level parts of the path.
- All paths within the same subscription request must have the same sample interval. If the same path requires different sample intervals, create multiple subscriptions.
- The feature does not support a path prefix in the Subscription request, but the Subscription can contain an empty prefix field.
- The feature supports Cisco DME and Device YANG data models. Openconfig YANG is not supported.
- The gRPC process that supports gNMI uses the HIGH_PRIO cgroup, which limits the CPU usage to 75% of CPU and memory to 1.5 GB.
- The **show grpc gnmi** command has the following considerations:
  - The commands are not XMLized in this release.
  - The gRPC agent retains gNMI calls for a maximum of 1 hour after the call has ended.
  - If the total number of calls exceeds 2000, the gRPC agent purges ended calls based an internal cleanup routine.

The gRPC server runs in the management VRF. As a result, the gRPC process communicates only in this VRF forcing the management interface to support all gRPC calls.

gRPC functionality now includes the default VRF for a total of 2 gRPC servers on each Cisco Nexus 3000 Series switch. You can run one gRPC server in each VRF, or run only one gRPC server in the management VRF. Supporting a gRPC in the default VRF adds flexibility to offload processing gRPC calls from the management VRF, where significant traffic load might not be desirable.

If two gRPC servers are configured, be aware of the following:

- VRF boundaries are strictly enforced, so each gRPC server processes requests independent of the other, and 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.

Example Client Output for Capabilities

In this example, all the OpenConfig model RPMs have been installed on the switch.

The following is an example of client output for Capabilities.

```
hostname user$ ./gnmi_cli -a 172.19.193.166:50051 -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-interfaces"
  organization: "OpenConfig working group"
  version: "2.0.0"
>
supported_models: <
  name: "openconfig-if-aggregate"
  organization: "OpenConfig working group"
  version: "2.0.0"
>
supported_models: <
  name: "openconfig-if-ethernet"
  organization: "OpenConfig working group"
  version: "2.0.0"
>
supported_models: <
  name: "openconfig-if-ip"
  organization: "OpenConfig working group"
  version: "2.3.0"
>
supported_models: <
  name: "openconfig-if-ip-ext"
  organization: "OpenConfig working group"
  version: "2.3.0"
>
supported_models: <
  name: "openconfig-lacp"
  organization: "OpenConfig working group"
  version: "1.0.2"
>
supported_models: <
  name: "openconfig-lldp"
  organization: "OpenConfig working group"
  version: "0.2.1"
>
supported_models: <
  name: "openconfig-network-instance"
  organization: "OpenConfig working group"
  version: "0.11.1"
>
supported_models: <
  name: "openconfig-network-instance-policy"
```

```
  organization: "OpenConfig working group"
  version: "0.1.1"
>
supported_models: <
  name: "openconfig-ospf-policy"
  organization: "OpenConfig working group"
  version: "0.1.1"
>
supported_models: <
  name: "openconfig-platform"
  organization: "OpenConfig working group"
  version: "0.12.2"
>
supported_models: <
  name: "openconfig-platform-cpu"
  organization: "OpenConfig working group"
  version: "0.1.1"
>
supported_models: <
  name: "openconfig-platform-fan"
  organization: "OpenConfig working group"
  version: "0.1.1"
>
supported_models: <
  name: "openconfig-platform-linecard"
  organization: "OpenConfig working group"
  version: "0.1.1"
>
supported_models: <
  name: "openconfig-platform-port"
  organization: "OpenConfig working group"
  version: "0.3.2"
>
supported_models: <
  name: "openconfig-platform-psu"
  organization: "OpenConfig working group"
  version: "0.2.1"
>
supported_models: <
  name: "openconfig-platform-transceiver"
  organization: "OpenConfig working group"
  version: "0.7.0"
>
supported_models: <
  name: "openconfig-relay-agent"
  organization: "OpenConfig working group"
  version: "0.1.0"
>
supported_models: <
  name: "openconfig-routing-policy"
  organization: "OpenConfig working group"
  version: "2.0.1"
>
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"

hostname user$
```

Get

About Get

The purpose of the Get RPC is to allow a client to retrieve a snapshot of the data tree from the device. Multiple paths may be requested in a single request. A simplified form of XPATH according to the gNMI Path Conventions, Schema path encoding conventions for gNMI are used for the path.

For detailed information on the Get operation, refer to the Retrieving Snapshots of State Information section in the gNMI specification: gRPC Network Management Interface (gNMI)

Guidelines and Limitations for Get

The following are guidelines and limitations for Get and Set:

- GetRequest.encoding supports only JSON.
- For GetRequest.type, only DataType CONFIG and STATE have direct correlation and expression in YANG. OPERATIONAL is not supported.
- A single request cannot have both OpenConfig (OC) YANG and device YANG paths. A request must have only OC YANG paths or device YANG paths, but not both.
- GetRequest for root path (“/”: everything from **all** models) is not allowed.
- GetRequest for the top level of the device model (“/System”) is not allowed.
- gNMI Get returns all default values (ref. report-all mode in RFC 6243 [4]).
- Subscribe supports the model `Cisco-NX-OS-syslog-oper`.
- Get does not support the model `Cisco-NX-OS-syslog-oper`.

- Query from the path `/system` does not return data from the path `/system/processes`. The specific path `/system/processes` should be used to query `openconfig-procmon` data.
- The following optional items are not supported:
  - Path prefix
  - Path alias
  - Wildcards in path
- A single GetRequest can have up to 10 paths.
- If the size of value field to be returned in GetResponse is over 12 MB, the system returns error status `grpc::RESOURCE_EXHAUSTED`.
- The maximum gRPC receive buffer size is set to 8 MB.
- The number of total concurrent sessions for Get is limited to five.
- Performing a Get operation when a large configuration is applied to the switch might cause the gRPC process to consume all available memory. If a memory exhaustion condition is hit, the following syslog is generated:

  ```
  MTX-API: The memory usage is reaching the max memory resource limit (3072) MB
  ```

  If this condition is hit several times consecutively, the following syslog is generated:

  ```
  The process has become unstable and the feature should be restarted.
  ```

  We recommend that you restart the gRPC feature at this point to continue normal processing of gNMI transactions.

Set

About Set

The Set RPC is used by a client to change the configuration of the device. The operations, which may be applied to the device data, are (in order) delete, replace, and update. All operations in a single Set request are treated as a transaction, meaning that all operations are successful or the device is rolled-back to the original state. The Set operations are applied in the order that is specified in the SetRequest. If a path is mentioned multiple times, the changes are applied even if they overwrite each other. The final state of the data is achieved with the final operation in the transaction. It is assumed that all paths specified in the SetRequest::delete, replace, update fields are CONFIG data paths and writable by the client.

For detailed information on the Set operation, refer to the Modifying State section of the gNMI Specification https://github.com/openconfig/reference/blob/1cf43d2146f9ba70abb7f04f6b0f6eaa504cef05/rpc/gnmi/gnmi-specification.md.

Guidelines and Limitations for Set

The following are guidelines and limitations for Set:

- SetRequest.encoding supports only JSON.

- A single request cannot have both OpenConfig (OC) YANG and device YANG paths. A request must have only OC YANG paths or device YANG paths, but not both.
- Subscribe supports the model `Cisco-NX-OS-syslog-oper`.
- Query from the path `/system` does not return data from the path `/system/processes`. The specific path `/system/processes` should be used to query `openconfig-procmon` data.
- The following optional items are not supported:
  - Path prefix
  - Path alias
  - Wildcards in path
- A single SetRequest can have up to 20 paths.
- The maximum gRPC receive buffer size is set to 8 MB.
- The number of total concurrent sessions for Get is limited to five.
- Performing a Set operation when a large configuration is applied to the switch might cause the gRPC process to consume all available memory. If a memory exhaustion condition is hit, the following syslog is generated:

  ```
  MTX-API: The memory usage is reaching the max memory resource limit (3072) MB
  ```

  If this condition is hit several times consecutively, the following syslog is generated:

  ```
  The process has become unstable and the feature should be restarted.
  ```

  We recommend that you restart the gRPC feature at this point to continue normal processing of gNMI transactions.
- For the Set::Delete RPC, an MTX log message warns if the configuration being operated on may be too large:

  ```
  Configuration size for this namespace exceeds operational limit. Feature may become
  unstable and require restart.
  ```

Subscribe

Guidelines and Limitations for Subscribe

Following are the guidelines and limitations for Subscribe:

- Beginning with Cisco NX-OS Release 9.3(3), Subscribe supports the OpenConfig model.
- The gNMI feature supports Subscribe and Capability as options of the gNMI service.
- The feature supports JSON and gnmi.proto encoding. The feature does not support protobuf.any encoding.
- Each gNMI message has a maximum size of 12 MB. If the amount of collected data exceeds the 12-MB maximum, the collected data is dropped.

You can avoid this situation by creating more focused subscriptions that handle smaller, more granular data-collection sets. So, instead of subscribing to one higher-level path, create multiple subscriptions for different, lower-level parts of the path.

- All paths within the same subscription request must have the same sample interval. If the same path requires different sample intervals, create multiple subscriptions.
- The feature does not support a path prefix in the Subscription request, but the Subscription can contain an empty prefix field.
- The feature supports Cisco DME and Device YANG data models. Openconfig YANG is not supported.
- The gRPC process that supports gNMI uses the HIGH_PRIO cgroup, which limits the CPU usage to 75% of CPU and memory to 1.5 GB.
- The **show grpc gnmi** command has the following considerations:
  - The commands are not XMLized in this release.
  - The gRPC agent retains gNMI calls for a maximum of 1 hour after the call has ended.
  - If the total number of calls exceeds 2000, the gRPC agent purges ended calls based an internal cleanup routine.

The gRPC server runs in the management VRF. As a result, the gRPC process communicates only in this VRF forcing the management interface to support all gRPC calls.

gRPC functionality now includes the default VRF for a total of 2 gRPC servers on each Cisco Nexus 3000 Series switch. You can run one gRPC server in each VRF, or run only one gRPC server in the management VRF. Supporting a gRPC in the default VRF adds flexibility to offload processing gRPC calls from the management VRF, where significant traffic load might not be desirable.

If two gRPC servers are configured, be aware of the following:

- VRF boundaries are strictly enforced, so each gRPC server processes requests independent of the other, and 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.

gNMI Payload

gNMI uses a specific payload format to subscribe to:

- DME Streams
- YANG Streams

Subscribe operations are supported with the following modes:

- ONCE: Subscribe and receive data once and close session.
- POLL: Subscribe and keep session open, client sends poll request each time data is needed.

- STREAM: Subscribe and receive data at specific cadence. The payload accepts values in nanoseconds 1 second = 1000000000.
- ON_CHANGE: Subscribe, receive a snapshot, and only receive data when something changes in the tree.

Setting modes:

- Each mode requires 2 settings, inside sub and outside sub
- ONCE: SAMPLE, ONCE
- POLL: SAMPLE, POLL
- STREAM: SAMPLE, STREAM
- ON_CHANGE: ON_CHANGE, STREAM

Origin

- DME: Subscribing to DME model
- device: Subscribing to YANG model

Name

- DME = subscribing to DME model
- Cisco-NX-OS-device = subscribing to YANG model

Encoding

- JSON = Stream will be send in JSON format.
- PROTO = Stream will be sent in protobuf.any format.

Sample gNMI Payload for DME Stream



Note Different clients have their own input format.

```
{
    "SubscribeRequest":
    [
        {
            "_comment" : "ONCE request",
            "_delay" : 2,
            "subscribe":
            {
                "subscription":
                [
                    {
                        "_comment" : "1st subscription path",
                        "path":
                        {
                            "origin": "DME",
                            "elem":
                            [
```

```
                                {
                                    "name": "sys"
                                },
                                {
                                    "name": "bgp"
                                 }
                            ]
                        },
                        "mode": "SAMPLE"
                    }
                ],
                "mode": "ONCE",
                "allow_aggregation" : false,
                "use_models":
                [
                    {
                        "_comment" : "1st module",
                        "name": "DME",
                        "organization": "Cisco Systems, Inc.",
                        "version": "1.0.0"
                    }
                ],
                "encoding": "JSON"
            }
        }
    ]
}
```

Sample gNMI Payload YANG Stream

```
{
    "SubscribeRequest":
    [
        {
            "_comment" : "ONCE request",
            "_delay" : 2,
            "subscribe":
            {
                "subscription":
                [
                    {
                        "_comment" : "1st subscription path",
                        "path":
                        {
                            "origin": "device",
                            "elem":
                            [
                                {
                                    "name": "System"
                                 },
                                 {
                                     "name": "bgp-items"
                                  }
                              ]
                          },
                                                    "mode": "SAMPLE"
                      }
                  ],
                  "mode": "ONCE",
                  "allow_aggregation" : false,
                  "use_models":
                  [
                      {
```

```
                        "_comment" : "1st module",
                        "name": "Cisco-NX-OS-device",
                        "organization": "Cisco Systems, Inc.",
                        "version": "0.0.0"
                    }
                ],
                "encoding": "JSON"
            }
        }
    ]
}
```

Streaming Syslog

About Streaming Syslog for gNMI

gNMI Subscribe is a new way of monitoring the network as it provides a real-time view of what's going on in your system by pushing the structured data as per gNMI Subscribe request.

Beginning with the Cisco NX-OS Release 9.3(3), support is added for gNMI Subscribe functionality.

gNMI Subscribe Support Detail

- Syslog-oper model streaming
  - stream_on_change

This feature applies to Cisco Nexus 3000 Series switches with 8 GB or more of memory.

Guidelines and Limitations for Streaming Syslog - gNMI

The following are guidelines and limitations for Streaming Syslog:

- An invalid syslog is not supported. For example, a syslog with a filter or query condition
- Only the following paths are supported:
  - Cisco-NX-OS-Syslog-oper:syslog
  - Cisco-NX-OS-Syslog-oper:syslog/messages
- The following modes are not supported:
  - Stream sample
  - POLL
- A request must be in the YANG model format.
- You can use the internal application or write your own application.
- The payload comes from the controller and gNMI sends a response.
- Encoding formats are JSON and PROTO.

Syslog Native YANG Model

The YangModels are located here.

Note The time-zone field is set only when the **clock format show-timezone syslog** is entered. By default, it's not set, therefore the time-zone field is empty.

```
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 Example

The following is an example of a Subscribe request:

```
{
    "SubscribeRequest":
    [
        {
            "_comment" : "STREAM request",
            "_delay"   : 2,
            "subscribe":
            {
                "subscription":
                [
                    {
                        "_comment" : "1st subscription path",
                        "path":
                        {
                            "origin": "syslog-oper",
                            "elem":
                              [
                                {
                                    "name": "syslog"
                                },
                                {
                                    "name":"messages"
                                }
                              ]
                        },
                        "mode": "ON_CHANGE"
                    }
                ],
                "mode": "ON_CHANGE",
             "allow_aggregation" : false,
```

```
                "use_models":
                [
                    {
                        "_comment" : "1st module",
                        "name": "Cisco-NX-OS-Syslog-oper",
                        "organization": "Cisco Systems, Inc.",
                        "version": "0.0.0"
                    }
                ],
                "encoding":"JSON"
            }
        }
    ]
}
```

Sample PROTO Output

This is a sample of PROTO output.

```
############################

[Subscribe]--------------------------------

### Reading from file ' /root/gnmi-console/testing_bl/stream_on_change/OC_SYSLOG.json '

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: "message-id"
}
}
val {
uint_val: 529
}
}
update {
path {
elem {
name: "node-name"
}
}
val {
string_val: "task-n9k-1"
}
}
update {
path {
elem {
name: "message-name"
}
}
val {
string_val: "VSHD_SYSLOG_CONFIG_I"
}
}
update {
path {
elem {
name: "text"
}
}
val {
string_val: "Configured from vty by admin on console0"
}
```

```
}
update {
path {
elem {
name: "group"
}
}
val {
string_val: "VSHD"
}
}
update {
path {
elem {
name: "category"
}
}
val {
string_val: "VSHD"
}
}
update {
path {
elem {
name: "time-of-day"
}
}
val {
string_val: "Nov 21 2019 14:26:40"
}
}
update {
path {
elem {
name: "time-zone"
}
}
val {
string_val: ""
}
}
update {
path {
elem {
name: "time-stamp"
}
}
val {
uint_val: 1574375200000
}
}
update {
path {
elem {
name: "severity"
}
}
val {
uint_val: 5
}
}
}

/Received ------------------------------------
```

Sample JSON Output

This is a sample JSON output.

```
[Subscribe]--------------------------------
### Reading from file ' testing_bl/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\":\"\"} ]] } ]"
}
}
}

/Received -----------------------------------
```

Troubleshooting

Gathering TM-Trace Logs

```
1. 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:
4091, sub_id_str: , dc_start_time: 1561157940912, length: 8466, sync_response:1
bash-4.3#
```

Gathering MTX-Internal Logs

```
1. Modify the following file with below /opt/mtx/conf/mtxlogger.cfg

<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">true<
/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="facility">
          <leaf name="info" type="boolean" default="true">true</leaf>
          <leaf name="warning" type="boolean" default="true">true<
/leaf>
          <leaf name="error" type="boolean" default="true">true</leaf>
          <leaf name="debug" type="boolean" default="false">true<
/leaf>
        </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>

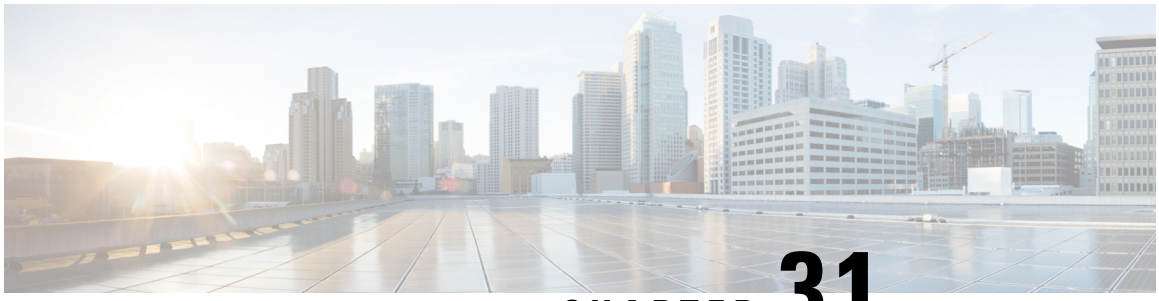

        <leaf name="location" type="string" default="./mtxlogs">
/volatile</leaf>
              <leaf name="mbytes-rollover" type="uint32" default="10"
>50</leaf>
              <leaf name="hours-rollover" type="uint32" default="24"
>24</leaf>
              <leaf name="startup-rollover" type="boolean" default="
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">MTX-EvtMgr</leaf>
                     <leaf name="active" type="boolean" default="true"
>true</leaf>
        </listitem>
        <listitem>
               <leaf name="id" type="string">TM-ADPT</leaf>
                   <leaf name="active" type="boolean" default="true"
>false</leaf>
```

```
        </listitem>
        <listitem>
              <leaf name="id" type="string">TM-ADPT-JSON</leaf>
                 <leaf name="active" type="boolean" default="true"
>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"
>true</leaf>
        </listitem>
        <listitem>
                 <leaf name="id" type="string">INST-ADAPTER-NC</leaf>
                    <leaf name="active" type="boolean" default="true"
>true</leaf>
        </listitem>
        <listitem>
               <leaf name="id" type="string">INST-ADAPTER-RC</leaf>
                  <leaf name="active" type="boolean" default="true"
>true</leaf>
        </listitem>
        <listitem>
               <leaf name="id" type="string">INST-ADAPTER-GRPC</leaf>
                  <leaf name="active" type="boolean" default="true"
>true</leaf>
         </listitem>
      </list>
    </container>
  </container>
</config>

2. Run "no feature grpc" / "feature grpc"
```

3. The /volataile directory houses the mtx-internal.log, the log rolls over over time so be sure to grab what you need before thenbash-4.3# cd /volatile/

```
bash-4.3# cd /volaiflels -al
total 148
drwxrwxrwx 4 root root 340 Jun 21 15:47 .
drwxrwxr-t 64 root network-admin 1600 Jun 21 14:45 ..
-rw-rw-rw- 1 root root 103412 Jun 21 16:14 grpc-internal-log
-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
-rw-rw-rw- 1 root root 355 Jun 21 14:44 netconf-internal-log
-rw-rw-rw- 1 root root 0 Jun 21 14:45 nginx_logflag
drwxrwxrwx 3 root root 60 Jun 21 14:45 uwsgipy
drwxrwxrwx 2 root root 40 Jun 21 14:43 virtual-instance
bash-4.3#.
```



CHAPTER 31

Dynamic Logger

- Dynamic Logger, on page 325
- Prerequisites, on page 325
- Reference, on page 325

Dynamic Logger

Prerequisites

Before using dynamic logging, confirm that the following are on your switch:

- The `libmtxlogmgr*.so` library is installed `/opt/mtx/lib/`. The `libmtxlogmgr*.so` library is part of the `mtx_infra` RPM.
- The `mtx.conf` file that is located in `/etc/` contains:

```
[mtxlogger]
config=/opt/mtx/conf/mtxlogger.cfg
```

- The `mtxlogger.cfg` file is in `/opt/mtx/conf/`.

Reference

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="$DATETIME$ $COMPONENTID$ $TYPE$:
$MSG$"></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"></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 22: Container and Leaf Descriptions

Container	Container Description	Contained Containers	Contained Leaf and Description
logging	Contains all logging data types	format dest file **Note** Also contains list tag "logitems"	enabled: Boolean that determines whether logging is on or off. Default off. allActive: Boolean that activates all defined logging items for logging. Default off

Container	Container Description	Contained Containers	Contained Leaf and Description
format	Contains the log message format information	componentID dateTime type fcn	content: String listing data types included in log messages. Includes: • $DATETIME$: Include date or time in log message • $COMPONENTID$: Include component name in log message. • $TYPE$: Includes message type ("", INFO, WARNING, ERROR) • $SRCFILE$: Includes name of source file. • $SRCLINE$: Include line number of source file • $FCNINFO$ Include class::function name from the source file. • MSG: Include actual log message text.
componentID	Name of logged component.	NA	enabled: Boolean that determines if the log message includes the component ID. Default to "true." Value of "false" returns a "" string in log message.
dateTime	Date or time of log message	NA	enabled: Boolean whether to include date or time information in log message. Default is enabled. format: String of values to include in log message. Format of %y%m%d.%H%M%S.

Container	Container Description	Contained Containers	Contained Leaf and 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 and Description
file	Determines the settings of the destination file.	NA	enabled: Boolean that determines whether the destination is enabled. Default is "false." name: String of the destination log file. Default of "mtx-internal.log" location: String of destination file path. Default at "./mtxlogs." mbytes-rollover: uint32 that determines the length of the log file before the system overwrites the oldest data. Default is 10 Mbytes. hours-rollover: uint32 that determines the length of the log file in terms of hours. Default is 24 hours. startup-rollover: Boolean that determines if the log file is rolled over upon agent start or restart. Default value of "false." 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.

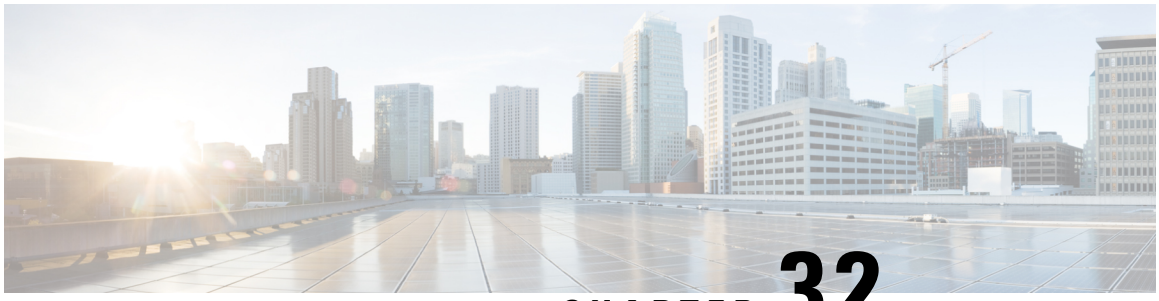
Example

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="./mtxlogs">/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>
         <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>
```



CHAPTER 32

Model-Driven Telemetry

This chapter contains the following topics:

- About Telemetry, on page 333
- Licensing Requirements for Telemetry, on page 335
- Installing and Upgrading Telemetry, on page 335
- Guidelines and Limitations for Telemetry, on page 336
- Configuring Telemetry Using the CLI, on page 342
- Configuring Telemetry Using the NX-API, on page 361
- Streaming Syslog, on page 375
- Telemetry Path Labels, on page 380
- About Native Data Source Paths, on page 395
- Additional References, on page 404

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 Process

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. 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.

- **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.)

  Starting with Cisco NX-OS Release 7.0(3)I7(1), UDP and secure UDP (DTLS) are supported as telemetry transport protocols. You can add destinations that receive UDP. The encoding for UDP and secure UDP can be GPB or JSON.

  Starting with Cisco NX-OS Release 9.2(1), telemetry now supports streaming to IPv6 destinations and IPv4 destinations.

  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 IPv4 destination:

```
destination-group 100
  ip address 171.70.55.69 port 50001 protocol UDP encoding 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.

Note Depending on the receiving operation system and the network load, using the UDP protocol may result in packet drops.

- **Telemetry Receiver** — A telemetry receiver is a remote management system or application that stores the telemetry data.

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.

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

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

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 *Cisco NX-OS Licensing Guide*.

Installing and Upgrading Telemetry

Installing the Application

The telemetry application is packaged as a feature RPM and included with the NX-OS release. The RPM is installed by default as part of the image bootup. After installation, you can start the application using the **feature telemetry** command. The RPM file is located in the `/rpms` directory and is named as follows:

telemetry-*version-build_ID***.libn32_n3000.rpm**

as in the following example:

```
telemetry-2.0.0-7.0.3.I5.1.lib32_n3000.rpm
```

Installing Incremental Updates and Fixes

Copy the RPM to the device bootflash and use the following commands from the `bash` prompt:

```
feature bash
run bash sudo su
```

Then copy the RPM to the device bootflash. Use the following commands from the `bash` prompt:

yum upgrade *telemetry_new_version*.**rpm**

The application is upgraded and the change will appear when the application is started again.

Downgrading to a Previous Version

To downgrade the telemetry application to a previous version, use the following command from the `bash` prompt:

```
yum downgrade telemetry
```

Verifying the Active Version

To verify the active version, run the following command from the switch `exec` prompt:

```
show install active
```



Note The `show install active` command will only show the active installed RPM after an upgrade has occurred. The default RPM that comes bundled with the NX-OS will not be displayed.

Guidelines and Limitations for Telemetry

Telemetry has the following configuration guidelines and limitations:

- Telemetry is supported in Cisco NX-OS releases that support the data management engine (DME) Native Model.
- Support is in place for DME data collection, NX-API data sources, Google protocol buffer (GPB) encoding over Google Remote Procedure Call (gRPC) transport, and 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 cadences below the minimum value may result in undesirable system behavior.
- Up to five remote management receivers (destinations) are supported. Configuring more than five remote receivers may result in undesirable system behavior.
- In the event that a telemetry receiver goes down, other receivers will see data flow interrupted. The failed receiver must be restarted. Then start a new connection with the switch by unconfiguring then reconfiguring the failer receiver's IP address under the destination group.
- Telemetry can consume up to 20% of the CPU resource.

- To configure SSL certificate based authentication and the encryption of streamed data, you can provide a self signed SSL certificate with **certificate** *ssl cert path* **hostname "CN"** command.
- QoS Explicit Congestion Notification (ECN) statistics are supported only on Cisco Nexus 9364C, 9336C-FX, and 93240YC-FX switches.

Configuration Commands After Downgrading to an Older Release

After a downgrade to an older release, some configuration commands or command options might fail because the older release may not support them. As a best practice when downgrading to an older release, unconfigure and reconfigure the telemetry feature after the new image comes up to avoid 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(conf-tm-dest)# ip address 1.2.3.4 port 50004 protocol gRPC encoding GPB `
`switch(conf-tm-dest)# sensor-group 100`
`switch(conf-tm-sensor)# path sys/bgp/inst/dom-default depth 0`
`switch(conf-tm-sensor)# subscription 600`
`switch(conf-tm-sub)# dst-grp 100`
`switch(conf-tm-sub)# snsr-grp 100 sample-interval 7000`
`switch(conf-tm-sub)# end`
Copy complete, now saving to disk (please wait)...
Copy complete.
switch#
```

gRPC Error Behavior

The switch client will disable the connection to the gRPC receiver if the gRPC receiver sends 20 errors. You will then need to 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.

Telemetry Compression for gRPC Transport

Telemetry compression support is available for gRPC transport. You can use the **use-compression gzip** command to enable compression. (Disable compression with the **no use-compression gzip** command.)

The following example enables compression:

```
switch(config)# telemetry
switch(config-telemetry)# destination-profile
switch(config-tm-dest-profile)# use-compression gzip
```

The following example shows compression is enabled:

```
switch(conf-tm-dest)# show telemetry transport 0 stats

Session Id:                    0
Connection Stats
   Connection Count            0
   Last Connected:             Never
   Disconnect Count            0
   Last Disconnected:          Never
Transmission Stats
   Compression:                gzip
   Source Interface:           loopback1(1.1.3.4)
   Transmit Count:             0
   Last TX time:               None
   Min Tx Time:                0                   ms
   Max Tx Time:                0                   ms
   Avg Tx Time:                0                   ms
   Cur Tx Time:                0                   ms


switch2(config-if)# show telemetry transport 0 stats

Session Id: 0
Connection Stats
Connection Count 0
Last Connected: Never
Disconnect Count 0
Last Disconnected: Never
Transmission Stats
Compression: disabled
Source Interface: loopback1(1.1.3.4)
Transmit Count: 0
Last TX time: None
Min Tx Time: 0 ms
Max Tx Time: 0 ms
Avg Tx Time: 0 ms
```

```
Cur Tx Time: 0 ms
switch2(config-if)#
```

The following is an example of use-compression as a POST payload:

```
{
              "telemetryDestProfile": {
                "attributes": {
                  "adminSt": "enabled"
                },
                "children": [
                  {
                    "telemetryDestOptCompression": {
                      "attributes": {
                        "name": "gzip"
                      }
                    }
                  }
                ]
              }
            }
```

Support for gRPC Chunking

Starting with Release 9.2(1), support for gRPC chunking has been added. For streaming to occur successfully, you must enable chunking if gRPC has to send an amount of data greater than 12MB to the receiver.

gRPC chunking has to be done by the gRPC user. Fragmentation has to be done on the gRPC client side and reassembly has to be done on the gRPC server side. Telemetry is still bound to memory and data can be dropped if the memory size is more than the allowed limit of 12MB 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 Telemetry Components and Process, on page 333.

The chunking size is between 64 and 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 will appear on systems that do not support gRPC chunking, such as the Cisco MDS series switches:

```
MDS-9706-86(conf-tm-dest)# 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; i.e., show vlan id 100, show vlan id 101, etc.. Instead, use the CLI range options; i.e., show vlan id 100-110,204, whenever possible to improve performance.

     If only the summary/counter is needed, then avoid dumping a whole show command output to limit the bandwidth and data storage 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 5 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. This means that the telemetry data stream can egress via front-panel ports and avoid possible competition between SSH/NGINX control sessions.

You can use the **use-vrf** *vrf-name* command to specify the transport VRF.

The following example specifies the transport VRF:

```
switch(config)# telemetry
switch(config-telemetry)# destination-profile
switch(config-tm-dest-profile)# use-vrf test_vrf
```

The following is an example of use-vrf as a POST payload:

```
{
              "telemetryDestProfile": {
                "attributes": {
                  "adminSt": "enabled"
                },
                "children": [
                  {
                    "telemetryDestOptVrf": {
                      "attributes": {
                        "name": "default"
                      }
                    }
                  }
                ]
              }
            }
```

Support for Node ID

Beginning in NX-OS release 9.3.1, you can configure 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-1(config)# telemetry
switch-1(config-telemetry)# destination-profile
switch-1(conf-tm-dest-profile)# use-nodeid test-srvr-10
switch-1(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
```

Support for Streaming of YANG Models

Starting with Release 9.2(1), the YANG ("Yet Another Next Generation") data modeling language is now supported as part of telemetry. Both device YANG and open config YANG model data streaming are supported.

Configuring Telemetry Using the CLI

Configuring Telemetry Using the NX-OS CLI

The following steps enable streaming telemetry and configuring the source and destination of the data stream. These steps also include optional steps to enable and configure SSL/TLS certificates and GPB encoding.

Procedure

	Command or Action	Purpose
Step 1	(Optional) **openssl** *argument* **Example:** Generate an SSL/TLS certificate using a specific argument, such as the following: • To generate a private RSA key: **openssl genrsa** *-cipher* **-out** *filename.key cipher-bit-length* For example: `switch# openssl genrsa -des3 -out server.key 2048` • To write the RSA key: **openssl rsa -in** *filename.key* **-out** *filename.key* For example: `switch# openssl rsa -in server.key -out server.key` • To create a certificate that contains the public or private key: **openssl req** *-encoding-standard* **-new** **-new** *filename.key* **-out** *filename.csr* **-subj '/CN=localhost'** For example: `switch# openssl req -sha256 -new -key server.key -out server.csr -subj '/CN=localhost'` • To create a public key: **openssl x509 -req** *-encoding-standard* **-days** *timeframe* **-in** *filename.csr* **-signkey** *filename.key* **-out** *filename.csr* For example: `switch# openssl x509 -req -sha256 -days 365 -in server.csr -signkey`	Create an SSL or TLS certificate on the server that receives the data, where the *private.key* file is the private key and the *public.crt* is the public key.

	Command or Action	Purpose
	`server.key` `-out server.crt`	
Step 2	**configure terminal** **Example:** `switch# configure terminal` `switch(config)#`	Enter the global configuration mode.
Step 3	**feature telemetry**	Enable the streaming telemetry feature.
Step 4	**feature nxapi**	Enable NX-API.
Step 5	**nxapi use-vrf management** **Example:** `switch(config)#` `switch(config)# nxapi use-vrf management` `switch(config)#`	Enable the VRF management to be used for NX-API communication. **Note** The following warnings are seen previous to 10.2(3)F release as ACLs are able to filter only netstack packets: "Warning: Management ACLs configured will not be effective for HTTP services. Please use iptables to restrict access." **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: "Warning: ACLs configured on non-management VRF will not be effective for HTTP services on that VRF."
Step 6	**telemetry** **Example:** `switch(config)# telemetry` `switch(config-telemetry)#`	Enter configuration mode for streaming telemetry.
Step 7	(Optional) **certificate** *certificate_path host_URL* **Example:** `switch(config-telemetry)# certificate /bootflash/server.key localhost`	Use an existing SSL/TLS certificate. For EOR devices, the certificate also has to be copied to the standby SUP.
Step 8	(Optional) Specify a transport VRF or enable telemetry compression for gRPC transport. **Example:** `switch(config-telemetry)#`	• Enter the **destination-profile** command to specify the default destination profile. • Enter any of the following commands: • **use-vrf** *vrf* to specify the destination VRF.

	Command or Action	Purpose
	destination-profile switch(conf-tm-dest-profile)# **use-vrf default** switch(conf-tm-dest-profile)# **use-compression gzip** switch(conf-tm-dest-profile)# **use-retry size 10** switch(conf-tm-dest-profile)# **source-interface loopback1**	• **use-compression gzip** to specify the destination compression method. • **use-retry size** *size* to specify the send retry details, with a retry buffer size between 10 - 1500 megabytes. • **source-interface** *interface-name* to stream data from the configured interface to a destination with the source IP address. **Note** After configuring the **use-vrf** command, you must configure a new destination IP address within the new VRF. However, you may re-use the same destination IP address by unconfiguring and reconfiguring the destination. This action ensures that the telemetry data streams to the same destination IP address in the new VRF.
Step 9	**sensor-group** *sgrp_id* **Example:** switch(config-telemetry)# **sensor-group 100** switch(conf-tm-sensor)#	Create a sensor group with ID *srgp_id* and enter sensor group configuration mode. Alphanumeric ID values are supported. The sensor group defines nodes that will be monitored for telemetry reporting.
Step 10	(Optional) **data-source** *data-source-type* **Example:** switch(config-telemetry)# **data-source NX-API**	Select a data source. Select from either YANG, DME or NX-API as the data source. **Note** DME is the default data source.
Step 11	**path** *sensor_path* **depth unbounded** [**filter-condition** *filter*] [**alias** *path_alias*] **Example:** • The following command is applicable for DME, not for NX-API or YANG: switch(conf-tm-sensor)# **path sys/bd/bd-[vlan-100] depth 0 filter-condition eq(l2BD.operSt, "down")** 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.	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. **Note** This is applicable for data source DME paths only. Add a sensor path to the sensor group. • Beginning with the Cisco NX-OS 9.3(5) release, the **alias** keyword is introduced.

	Command or Action	Purpose
	`switch(conf-tm-sensor)# path sys/bd/bd-[vlan-100] depth 0 filter-condition and(updated(l2BD.operSt),eq(l2BD.operSt,"down"))` Use the following syntax to distinguish the path on the UTR side. `switch(conf-tm-sensor)# path sys/ch/ftslot-1/ft alias ft_1` • The following command is applicable for NX-API, not for DME or YANG: `switch(conf-tm-sensor)# path "show interface" depth 0` • The following command is applicable for device YANG: `switch(conf-tm-sensor)# path Cisco-NX-OS-device:System/bgp-items/inst-items` • The following commands are applicable for OpenConfig YANG: `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` • The following command is applicable for NX-API: `switch(conf-tm-sensor)# path "show interface" depth 0 alias sh_int_alias` • The following command is applicable for OpenConfig: `switch(conf-tm-sensor)# path openconfig-bgp:bgp alias oc_bgp_alias`	• The **depth** setting specifies the retrieval level for the sensor path. Depth settings of **0** - **32**, **unbounded** are supported. **Note** **depth 0** is the default depth. NX-API-based sensor paths can only use **depth 0**. If a path is subscribed for the event collection, the depth only supports 0 and unbounded. Other values would be treated as 0. • The optional **filter-condition** parameter can be specified to create a specific filter for event-based subscriptions. 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**. **Note** query-condition parameter — For DME, based on the DN, the query-condition parameter can be specified to fetch MOTL and ephemeral data with the following syntax: query-condition "rsp-foreign-subtree=applied-config"; query-condition "rsp-foreign-subtree=ephemeral". • 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: • For device YANG: **Cisco-NX-OS-device:System/bgp-items/inst-items** • For OpenConfig YANG: **openconfig-bgp:bgp**

	Command or Action	Purpose
		Note The **depth**, **filter-condition**, and **query-condition** parameters are not supported for YANG currently. 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. Instead of installing a specific model, you can install the openconfig-all RPM which has all the OpenConfig models. For example: **install add mtx-openconfig-bgp-1.0.0.0-7.0.3.I7.1.HDB.1.0.632_n9000.rpm activate**
Step 12	**destination-group** *dgrp_id* **Example:** `switch(conf-tm-sensor)# destination-group 100` `switch(conf-tm-dest)#`	Create a destination group and enter destination group configuration mode. The *dgrp_id* supports alphanumeric ID values..
Step 13	(Optional) **ip address** *ip_address* **port** *port* **protocol** *procedural-protocol* **encoding** *encoding-protocol* **Example:** `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` `switch(conf-tm-sensor)# ip address 171.70.55.69 port 50009 protocol UDP encoding JSON`	Specify an IPv4 IP address and port to receive encoded telemetry data. **Note** gRPC is the default transport protocol. GPB is the default encoding.
Step 14	(Optional) **ipv6 address** *ipv6_address* **port** *port* **protocol** *procedural-protocol* **encoding** *encoding-protocol* **Example:** `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`	Specify an IPv6 IP address and port to receive encoded telemetry data. **Note** gRPC is the default transport protocol. GPB is the default encoding.

	Command or Action	Purpose
	`10:10::1 port 8002 protocol UDP encoding JSON`	
Step 15	*ip_version* **address** *ip_address* **port** *portnum* **Example:** • For IPv4: `switch(conf-tm-dest)# ip address 1.2.3.4 port 50003` • For IPv6: `switch(conf-tm-dest)# ipv6 address 10:10::1 port 8000`	Create a destination profile for the outgoing data, where *ip_version* is either ip (for IPv4) or ipv6 (for IPv6). 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.
Step 16	(Optional) **use-chunking size** *chunking_size* **Example:** `switch(conf-tm-dest)# use-chunking size 64`	Enable gRPC chunking and set the chunking size, between 64-4096 bytes. See the section "Support for gRPC Chunking" for more information.
Step 17	**subscription** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 100` `switch(conf-tm-sub)#`	Create a subscription node with ID and enter the subscription configuration mode. The *sub_id* supports alphanumeric ID values.. **Note** When subscribing to a DN, check whether the DN is supported by DME using REST to ensure that events will stream.
Step 18	**snsr-grp** *sgrp_id* **sample-interval** *interval* **Example:** `switch(conf-tm-sub)# snsr-grp 100 sample-interval 15000`	Link the sensor group with ID *sgrp_id* 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 19	**dst-grp** *dgrp_id* **Example:** `switch(conf-tm-sub)# dst-grp 100`	Link the destination group with ID *dgrp_id* to this subscription.

Configuration Examples for Telemetry Using the CLI

The following steps describe how to 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)# 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
```

This example creates a subscription that streams data for the `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
```

This example creates a subscription that streams data for `sys/intf` every 5 seconds to destination IP 1.2.3.4 port 50003, and 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
```

This example creates a subscription that streams data for `sys/cdp` 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
```

This example creates a cadence-based collection of **show** command data 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
```

This example creates an event-based subscription for `sys/fm`. Data is streamed to the destination 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
```

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
```

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
```

You can specify transport VRF and telemetry data compression for gRPC using the **use-vrf** and **use-compression gzip** commands, as shown in this example.

```
switch(config)# telemetry
switch(config-telemetry)# destination-profile
switch(conf-tm-dest-profile)# use-vrf default
switch(conf-tm-dest-profile)# use-compression gzip
switch(conf-tm-dest-profile)# sensor-group 1
switch(conf-tm-sensor)# path sys/bgp depth unbounded
switch(conf-tm-sensor)# destination-group 1
switch(conf-tm-dest)# ip address 1.2.3.4 port 50004
switch(conf-tm-dest)# subscription 1
switch(conf-tm-sub)# dst-grp 1
switch(conf-tm-sub)# snsr-grp 1 sample-interval 10000
```

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 the internal databases that reflect the configuration of telemetry.

```
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            sys/bgp(1) : NA : NA
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

4          No         1              0           Self            show version(2) : NA : NA
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 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

    Memory:
            Allowed Memory Limit               = 1181116006 bytes
            Occupied Memory                    = 93265920 bytes

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 23: Syntax Description for show telemetry transport

Syntax	Description
show	Shows running system information
telemetry	Shows telemetry information
transport	Shows telemetry transport information
session_id	(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

Syntax	Description
session_idx	(Optional) Session Id
ip_address	(Optional) Transport IP address
port	(Optional) Transport port
dest_info	(Optional) Destination information
encoding_type	(Optional) Encoding type
transport_type	(Optional) Transport type
transport_status	(Optional) Transport status
transport_security_cert_fname	(Optional) Transport security file name
transport_last_connected	(Optional) Transport last connected
transport_last_disconnected	(Optional) Last time this destination configuration was removed
transport_errors_count	(Optional) Transport errors count
transport_last_tx_error	(Optional) Transport last tx error
transport_statistics	(Optional) Transport statistics
t_session_id	(Optional) Transport Session id
connect_statistics	(Optional) Connection statistics
connect_count	(Optional) Connection count
last_connected	(Optional) Last connected timestamp
disconnect_count	(Optional) Disconnect count
last_disconnected	(Optional) Last time this destination configuration was removed
trans_statistics	(Optional) Transport statistics
compression	(Optional) Compression status
source_interface_name	(Optional) Source interface name
source_interface_ip	(Optional) Source interface IP
transmit_count	(Optional) Transmission count
last_tx_time	(Optional) Last Transmission time
min_tx_time	(Optional) Minimum transmission time
max_tx_time	(Optional) Maximum transmission time

Syntax	Description
avg_tx_time	(Optional) Average transmission time
cur_tx_time	(Optional) Current transmission time
transport_errors	(Optional) Transport errors
connect_errors	(Optional) Connection errors
connect_errors_count	(Optional) Connection error count
trans_errors	(Optional) Transport errors
trans_errors_count	(Optional) Transport error count
last_tx_error	(Optional) Last transport error
last_tx_return_code	(Optional) Last transport return code
transport_retry_stats	(Optional) Retry Statistics
ts_event_retry_bytes	(Optional) Event Retry buffer size
ts_timer_retry_bytes	(Optional) Timer Retry buffer size
ts_event_retry_size	(Optional) Event Retry number of messages
ts_timer_retry_size	(Optional) Timer Retry number of messages
ts_retries_sent	(Optional) Number of retries sent
ts_retries_dropped	(Optional) Number of retries dropped
event_retry_bytes	(Optional) Event Retry buffer size
timer_retry_bytes	(Optional) Timer Retry buffer size
retries_sent	(Optional) Number of retries sent
retries_dropped	(Optional) Number of retries dropped
retry_buffer_size	(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
Last Disconnected:   Never
```

```
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
Transport:           HTTP
Status:              Disconnected
Last Connected:      Never
Last Disconnected:   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> stats

This command displays details of a specific transport session.

```
Session Id:                   0
Transmission Stats
   Compression:               disabled
   Source Interface:          not set()
   Transmit Count:            319297
   Last TX time:              Fri Aug 02 03:51:15.287 UTC
   Min Tx Time:               1                  ms
   Max Tx Time:               3117               ms
```

```
   Avg Tx Time:                3                   ms
   Cur Tx Time:                1                   ms
```

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 Errors
Connection Error Count:     0
Transmission Errors
Tx Error Count:             30
Last Tx Error:              Thu Aug 01 04:39:47.083 UTC
Last Tx Return Code:        No error
```

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
```

show system internal telemetry trace

The **show system internal telemetry trace** [**tm-events** | **tm-errors** |**tm-logs** | **all**] command displays system internal telemetry trace information.

```
switch# show system internal telemetry trace all
Telemetry All Traces:
Telemetry Error Traces:
[07/26/17 15:22:29.156 UTC 1 28577] [3960399872][tm_cfg_api.c:367] Not able to destroy dest
 profile list for config node rc:-1610612714 reason:Invalid argument
[07/26/17 15:22:44.972 UTC 2 28577] [3960399872][tm_stream.c:248] No subscriptions for
destination group 1
[07/26/17 15:22:49.463 UTC 3 28577] [3960399872][tm_stream.c:576] TM_STREAM: Subscriptoin
1 does not have any sensor groups

3 entries printed
Telemetry Event Traces:
[07/26/17 15:19:40.610 UTC 1 28577] [3960399872][tm_debug.c:41] Telemetry xostrace buffers
 initialized successfully!
[07/26/17 15:19:40.610 UTC 2 28577] [3960399872][tm.c:744] Telemetry statistics created
successfully!
[07/26/17 15:19:40.610 UTC 3 28577] [3960399872][tm_init_n9k.c:97] Platform intf:
grpc_traces:compression,channel
switch#

switch# show system internal telemetry trace tm-logs
Telemetry Log Traces:
0 entries printed
switch#
switch# show system internal telemetry trace tm-events
Telemetry Event Traces:
```

```
[07/26/17 15:19:40.610 UTC 1 28577] [3960399872][tm_debug.c:41] Telemetry xostrace buffers
 initialized successfully!
[07/26/17 15:19:40.610 UTC 2 28577] [3960399872][tm.c:744] Telemetry statistics created
successfully!
[07/26/17 15:19:40.610 UTC 3 28577] [3960399872][tm_init_n9k.c:97] Platform intf:
grpc_traces:compression,channel
[07/26/17 15:19:40.610 UTC 4 28577] [3960399872][tm_init_n9k.c:207] Adding telemetry to
cgroup
[07/26/17 15:19:40.670 UTC 5 28577] [3960399872][tm_init_n9k.c:215] Added telemetry to
cgroup successfully!

switch# show system internal telemetry trace tm-errors
Telemetry Error Traces:
0 entries printed
switch#
```

Configuring Telemetry Using the NX-API

Configuring Telemetry Using the 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 Telemetry Model in the DME, on page 373. 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
```

Procedure

	Command or Action	Purpose
Step 1	Enable the telemetry feature. **Example:** `{ "fmEntity" : { "children" : [{ "fmTelemetry" : { "attributes" : { "adminSt" : "enabled" } } } } ] } }`	The root element is **fmTelemetry** and the base path for this element is `sys/fm`. Configure the **adminSt** attribute as `enabled`.
Step 2	Create the root level of the JSON payload to describe the telemetry configuration. **Example:** `{ "telemetryEntity": { "attributes": { "dn": "sys/tm" }, } }`	The root element is **telemetryEntity** and the base path for this element is `sys/tm`. Configure the **dn** attribute as `sys/tm`.
Step 3	Create a sensor group to contain the defined sensor paths. **Example:**	A telemetry sensor group is defined in an object of class **telemetrySensorGroup**. Configure the following attributes of the object:

	Command or Action	Purpose
	<pre>"telemetrySensorGroup": { "attributes": { "id": "10", "rn": "sensor-10" "dataSrc": "NX-API" }, "children": [{ }] }</pre>	• **id** — An identifier for the sensor group. Currently only numeric ID values are supported. • **rn** — The relative name of the sensor group object in the format: **sensor-***id*. • **dataSrc** — Selects the data source from **DEFAULT**, **DME**, **YANG**, or **NX-API**. Children of the sensor group object will include sensor paths and one or more relation objects (**telemetryRtSensorGroupRel**) to associate the sensor group with a telemetry subscription.
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**. • Configure the **adminSt** attribute as `enabled`. • 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.

	Command or Action	Purpose
	<pre>] } }</pre>	
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>{ "telemetryDest": { "attributes": { "addr": "1.2.3.4", "enc": "GPB", "port": "50001", "proto": "gRPC", "rn": "addr-[1.2.3.4]-port-50001" } } }</pre>	A telemetry destination is defined in an object of class **telemetryDest**. Configure the following attributes of the object: • **addr** — The IP address of the destination. • **port** — The port number of the destination. • **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: • Google protocol buffers (GPB) for gRPC. • JSON for C. • GPB or JSON for UDP and secure UDP (DTLS). • **proto** — The transport protocol type of the telemetry data to be sent. NX-OS supports: • gRPC • HTTP • VUDP and secure UDP (DTLS)
Step 8	Enable gRPC chunking and set the chunking size, between 64 and 4096 bytes. **Example:** <pre>{ "telemetryDestGrpOptChunking": { "attributes": { "chunkSize": "2048", "dn": "sys/tm/dest-1/chunking" } } }</pre>	See Support for gRPC Chunking, on page 339 for more information.

	Command or Action	Purpose
Step 9	Create a telemetry subscription to configure the telemetry behavior. **Example:** `"telemetrySubscription": { "attributes": { "id": "30", "rn": "subs-30" }, "children": [{ }] }`	A telemetry subscription is defined in an object of class **telemetrySubscription**. Configure the following attributes of the object: • **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*. Children of the subscription object will include relation objects for sensor groups (**telemetryRsSensorGroupRel**) and destination groups (**telemetryRsDestGroupRel**).
Step 10	Add the sensor group object as a child object to the **telemetrySubscription** element under the root element (**telemetryEntity**). **Example:** `{ "telemetrySubscription": { "attributes": { "id": "30" } "children": [{ "telemetryRsSensorGroupRel": { "attributes": { "sampleIntvl": "5000", "tDn": "sys/tm/sensor-10" } } } ] } }`	
Step 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. **Example:** `"telemetryRsSensorGroupRel": { "attributes": { "rType": "mo", "rn": "rssensorGroupRel-[sys/tm/sensor-10]", "sampleIntvl": "5000", "tCl": "telemetrySensorGroup", "tDn": "sys/tm/sensor-10", "tType": "mo"`	The relation object is of class **telemetryRsSensorGroupRel** and is a child object of **telemetrySubscription**. Configure the following attributes of the relation object: • **rn** — The relative name of the relation object in the format: **rssensorGroupRel-[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

	Command or Action	Purpose
	<pre> } }</pre>	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. • **tType** — The target type, which is **mo** for managed object.
Step 12	Define one or more sensor paths or nodes to be monitored for telemetry. **Example:** Single sensor path <pre>{ "telemetrySensorPath": { "attributes": { "path": "sys/cdp", "rn": "path-[sys/cdp]", "excludeFilter": "", "filterCondition": "", "path": "sys/fm/bgp", "secondaryGroup": "0", "secondaryPath": "", "depth": "0" "alias": "cdp_alias", } } }</pre>**Example:** Single sensor path for NX-API <pre>{ "telemetrySensorPath": { "attributes": { "path": "show interface", "path": "show bgp", "rn": "path-[sys/cdp]", "excludeFilter": "",</pre>	A sensor path is defined in an object of class **telemetrySensorPath**. Configure the following attributes of the object: • **path** — The path to be monitored. • **rn** — The relative name of the path object in the format: **path-[***path***]** • **depth** — The retrieval level for the sensor path. A depth setting of **0** retrieves only the root MO properties. • **filterCondition** — (Optional) Creates a specific filter for event-based subscriptions. The DME provides the filter expressions. For more information regarding filtering, see the Cisco APIC REST API Usage Guidelines on composing queries: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/rest_cfg/2_1_x/b_Cisco_APIC_REST_API_Configuration_Guide/b_Cisco_APIC_REST_API_Configuration_Guide_chapter_01.html#d25e1534a1635 • **alias** - Specify an alias for this path.

	Command or Action	Purpose

```
            "filterCondition": "",
            "path": "sys/fm/bgp",
            "secondaryGroup": "0",
            "secondaryPath": "",
            "depth": "0"
        }
    }
}
```

Example:

Multiple sensor paths

```
{
    "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"
        }
    }
}
```

Example:

Single sensor path filtering for BGP disable events:

```
{
    "telemetrySensorPath": {
        "attributes": {
            "path": "sys/cdp",
            "rn": "path-[sys/cdp]",
            "excludeFilter": "",
            "filterCondition":
"eq(fmBgp.operSt.\"disabled\")",
            "path": "sys/fm/bgp",
            "secondaryGroup": "0",
            "secondaryPath": "",
            "depth": "0"
        }
    }
```

	Command or Action	Purpose
	`}`	
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	Create a relation object as a child object of the telemetry sensor group to associate the sensor group to the subscription. **Example:** `"telemetryRtSensorGroupRel": { "attributes": { "rn": "rtsensorGroupRel-[sys/tm/subs-30]", "tCl": "telemetrySubscription", "tDn": "sys/tm/subs-30" } }`	The relation object is of class **telemetryRtSensorGroupRel** and is a child object of **telemetrySensorGroup**. Configure the following attributes of the relation object: • **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**. • **tDn** — The target distinguished name of the subscription object, which is **sys/tm/***subscription-id*.
Step 17	Create a relation object as a child object of the telemetry destination group to associate the destination group to the subscription. **Example:** `"telemetryRtDestGroupRel": { "attributes": { "rn": "rtdestGroupRel-[sys/tm/subs-30]", "tCl": "telemetrySubscription", "tDn": "sys/tm/subs-30" } }`	The relation object is of class **telemetryRtDestGroupRel** and is a child object of **telemetryDestGroup**. Configure the following attributes of the relation object: • **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**. • **tDn** — The target distinguished name of the subscription object, which is **sys/tm/***subscription-id*.
Step 18	Create a relation object as a child object of the subscription to associate the subscription to the telemetry destination group. **Example:** `"telemetryRsDestGroupRel": { "attributes": { "rType": "mo", "rn":`	The relation object is of class **telemetryRsDestGroupRel** and is a child object of **telemetrySubscription**. Configure the following attributes of the relation object: • **rn** — The relative name of the relation object in the format: **rsdestGroupRel-[sys/tm/***destination-group-id***]**.

	Command or Action	Purpose
	<pre>"rsdestGroupRel-[sys/tm/dest-20]", "tCl": "telemetryDestGroup", "tDn": "sys/tm/dest-20", "tType": "mo" } }</pre>	• **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. • **tType** — The target type, which is **mo** for managed object.
Step 19	Send the resulting JSON structure as an HTTP/HTTPS POST payload to the NX-API endpoint for telemetry configuration.	The base path for the telemetry entity is `sys/tm` and the NX-API endpoint is: `{{URL}}/api/node/mo/sys/tm.json`

Example

The following is an example of all the previous steps 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": {
            "attributes": {
              "addr": "10.30.217.80",
              "port": "50051",
              "enc": "GPB",
```

```
            "proto": "gRPC"
          }
        }
      }
      ]
    }
  },
  {
    "telemetrySubscription": {
      "attributes": {
        "id": "30"
      }
      "children": [{
        "telemetryRsSensorGroupRel": {
          "attributes": {
            "sampleIntvl": "5000",
            "tDn": "sys/tm/sensor-10"
          }
        }
      },
      {
        "telemetryRsDestGroupRel": {
          "attributes": {
            "tDn": "sys/tm/dest-20"
          }
        }
      }
      ]
    }
  }
  ]
 }
}
```

Configuration Example for Telemetry Using the NX-API

Streaming Paths to a Destination

This example creates a subscription that streams paths `sys/cdp` and `sys/ipv4` to a destination `1.2.3.4 port 50001` every five seconds.

```
POST https://192.168.20.123/api/node/mo/sys/tm.json

Payload:
{
    "telemetryEntity": {
        "attributes": {
            "dn": "sys/tm"
        },
        "children": [{
            "telemetrySensorGroup": {
                "attributes": {
                    "id": "10",
                    "rn": "sensor-10"
                },  "children": [{
                    "telemetryRtSensorGroupRel": {
                        "attributes": {
                            "rn": "rtsensorGroupRel-[sys/tm/subs-30]",
                            "tCl": "telemetrySubscription",
                            "tDn": "sys/tm/subs-30"
```

```
                    }
                }
            }, {
                "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"
        }
        "children": [{
          "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",
                "proto": "gRPC"
              }
            }
          }
          ]
        }
      },
      {
        "telemetrySubscription": {
          "attributes": {
            "id": "30"
          }
          "children": [{
            "telemetryRsSensorGroupRel": {
              "attributes": {
                "sampleIntvl": "0",
                "tDn": "sys/tm/sensor-10"
              }
            }
          },
          {
            "telemetryRsDestGroupRel": {
              "attributes": {
                "tDn": "sys/tm/dest-20"
              }
            }
          }
          ]
        }
      }
      ]
    }
}
```

Using Postman Collection for Telemetry Configuration

An example Postman collection is an easy way to start configuring the telemetry feature, and can run all telemetry CLI equivalents in a single payload. Modify the file in the preceding link using your preferred text editor to update the payload to your needs, then open the collection in Postman and run the collection.

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
|     |    @name:dataSrc
|     |      @type:DataSource
|     |
|     |----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
```

DNs Available to Telemetry

For a list of DNs available to the telemetry feature, see Streaming Telemetry Sources, on page 463.

Streaming Syslog

About Streaming Syslog for gNMI

gNMI Subscribe is a new way of monitoring the network as it provides a real-time view of what's going on in your system by pushing the structured data as per gNMI Subscribe request.

Beginning with the Cisco NX-OS Release 9.3(3), support is added for gNMI Subscribe functionality.

gNMI Subscribe Support Detail

- Syslog-oper model streaming
  - stream_on_change

This feature applies to Cisco Nexus 3000 Series switches with 8 GB or more of memory.

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.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enter global configuration mode.
Step 2	**telemetry** **Example:** `switch(config)# telemetry`	Enter configuration mode for telemetry.
Step 3	**sensor-group** *sgrp_id* **Example:** `switch(config-telemetry)# sensor-group 6`	Creates a sensor group.
Step 4	**data source** *data-source-type* **Example:** `switch(config-tm-sensor)# data source YANG`	Set the data source to YANG, so that it uses the native YANG streaming model to stream syslogs

	Command or Action	Purpose
Step 5	**path Cisco-NX-OS-Syslog-oper:syslog/messages** **Example:** `switch(config-tm-sensor)# path Cisco-NX-OS-Syslog-oper:syslog/messages`	Configure the syslog path which streams syslog generated on the switch.
Step 6	**destination-group** *grp_id* **Example:** `switch(config-tm-sensor)# destination-group 33`	Enter telemetry destination group sub-mode and configure the destination group.
Step 7	ip address *ip_addr* **port** *port* **protocol** {**HTTP** | **gRPC** } **encoding** { **JSON** | **GPB** | **GPB-compact** } **Example:** `switch(config-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json` **Example:** `switch(config-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb`	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** *sub-id* **Example:** `switch(config-tm-dest)# subscription 33`	Enter telemetry subscription submode and configure the telemetry subscription.
Step 9	**snsr-group** *sgrp_id* **sample-interval** *interval* **Example:** `switch(config-tm-sub)# snsr-group 6 sample-interval 0`	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** *dgrp_id* **Example:** `switch(config-tm-sub)# dst-grp 33`	Link the destination group to the current subscription. The destination group that you specify must match the destination group thtat 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 Notification
Cisco-NX-OS-Syslog-oper:syslog/messages	Stream all the existing syslogs from the switch.	Sends event notification for syslog occurred on the switch: • message-id • node-name • time-stamp • time-of-day • time-zone • category • message-name • 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:

```
switch# show telemetry event collector stats

------------------------------------------------------------------------------
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.

```
KVGPB Output:
---Telemetry msg received @ 18:22:04 UTC

Read frag:1 size:339 continue to block on read..

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"
      }
    }
  }
}
```

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`.

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 381.

Note The model does not support counters for subinterface, loopback, or VLAN, so they are not streamed out.

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** *sgrp_id* **Example:** `switch(config-telemetry)# sensor-group 6` `switch(conf-tm-sensor)#`	Create a sensor group for telemetry data.
Step 4	**path interface** **Example:** `switch(conf-tm-sensor)# path interface` `switch(conf-tm-sensor)#`	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. Depending on how the polling interval is configured, interface data is sent based on a

	Command or Action	Purpose
		periodic basis or whenever the interface state changes.
Step 5	**destination-group** *grp_id* **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** *ip_addr* **port** *port* **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** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	**snsr-group** *sgrp_id* **sample-interval** *interval* **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** *dgrp_id* **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 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**

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** *sgrp_id* **Example:** `switch(config-telemetry)# sensor-group 6` `switch(conf-tm-sensor)#`	Create a sensor group for telemetry data.
Step 4	**path interface query-condition counters=[detailed]** **Example:** `switch(conf-tm-sensor)# path interface query-condition counters=[detailed]` `switch(conf-tm-sensor)#`	Configure the interface path label and query for only the non-zero counters from all interfaces.
Step 5	**destination-group** *grp_id* **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** *ip_addr* **port** *port* **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** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	**snsr-group** *sgrp_id* **sample-interval** *interval* **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.

	Command or Action	Purpose
Step 9	**dst-group** *dgrp_id* **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 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.

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	**snsr-group** *sgrp_id* **sample-interval** *interval* **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=[***speed***]** **Example:**	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.

	Command or Action	Purpose
	`switch(conf-tm-sensor)# path interface query-condition oper-speed=[1G,40G]` `switch(conf-tm-sensor)#`	
Step 5	**destination-group** *grp_id* **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** *ip_addr* **port** *port* **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** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	**snsr-group** *sgrp_id* **sample-interval** *interval* **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** *dgrp_id* **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.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal` `switch(config)#`	Enter configuration mode.

	Command or Action	Purpose
Step 2	**telemetry** **Example:** `switch(config)# telemetry` `switch(config-telemetry)#`	Enter configuration mode for the telemetry features.
Step 3	**sensor-group** *sgrp_id* **Example:** `switch(config-telemetry)# sensor-group 6` `switch(conf-tm-sensor)#`	Create a sensor group for telemetry data.
Step 4	**path interface query-condition counters=[detailed],oper-speed=[1G,40G]** **Example:** `switch(conf-tm-sensor)# path interface query-condition counters=[detailed],oper-speed=[1G,40G]` `switch(conf-tm-sensor)#`	Configures multiple conditions in the same query. In this example, the query does both of the following: • Gathers and returns non-zero counters on interfaces running at 1 Gbps. • Gathers and returns non-zero counters on interfaces running at 40 Gbps.
Step 5	**destination-group** *grp_id* **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** *ip_addr* **port** *port* **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** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	**snsr-group** *sgrp_id* **sample-interval** *interval* **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** *dgrp_id* **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 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 381.

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.

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** *sgrp_id* **Example:** `switch(config-telemetry)# sensor-group 6` `switch(conf-tm-sensor)#`	Create a sensor group for telemetry data.
Step 4	**path environment** **Example:** `switch(conf-tm-sensor)# path environment` `switch(conf-tm-sensor)#`	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. Depending on the sample interval, the environment data is either streaming based on the polling interval, or sent when events occur.
Step 5	**destination-group** *grp_id* **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** *ip_addr* **port** *port* **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.

	Command or Action	Purpose
Step 7	**subscription** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	**snsr-group** *sgrp_id* **sample-interval** *interval* **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 environment events occur.
Step 9	**dst-group** *dgrp_id* **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 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 381.

This path does not support filtering.

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** *sgrp_id* **Example:** `switch(config-telemetry)# sensor-group 6` `switch(conf-tm-sensor)#`	Create a sensor group for telemetry data.
Step 4	**path resources** **Example:**	Configure the resources path label, which enables telemetry data for multiple individual system resources to be sent to the label. The

	Command or Action	Purpose
	`switch(conf-tm-sensor)# path resources` `switch(conf-tm-sensor)#`	label consolidates the multiple data inputs into one output. 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.
Step 5	**destination-group** *grp_id* **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** *ip_addr* **port** *port* **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** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	**snsr-group** *sgrp_id* **sample-interval** *interval* **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 resource events occur.
Step 9	**dst-group** *dgrp_id* **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 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 Polling for Data or Receiving Events, on page 381.

This path does not support filtering.

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** *sgrp_id* **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** *grp_id* **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** *ip_addr* **port** *port* **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** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	**snsr-group** *sgrp_id* **sample-interval** *interval* **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** *dgrp_id* **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.

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.

Procedure

	Command or Action	Purpose
Step 1	**show running-config-telemetry** **Example:** `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)#`	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
show telemetry usability {all | environment | interface | resources | vxlan}	Either all telemetry paths for all path labels, or all telemetry paths for a specified path label. Also, the output shows whether each path reports telemetry data based on periodic polling or events. For the interfaces path label, also any keyword filters or query conditions you configured.
show running-config telemetry	The running configuration for telemetry and selected path information.

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     :
[illegible]
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(l1PhysIf.adminSt,"up")&rsp-subtree=children&rsp-subtree-class=rmonEtherStats,rmonIfIn,rmonIfOut,rmonIfHCIn,rmonIfHCOut


  2) label_name          : interface

     path_name           : sys/mgmt-[mgmt0]
     query_type          : poll
     query_condition     :
query-target=subtree&query-target-filter=eq(mgmtMgmtIf.adminSt,"up")&rsp-subtree=full&rsp-subtree-class=rmonEtherStats,rmonIfIn,rmonIfOut,rmonIfHCIn,rmonIfHCOut


  3) label_name          : interface
```

```
     path_name           : sys/intf
     query_type          : event
     query_condition     :
<illegible>
ethpmEncRtdIf.operSt,"down")),and(updated(ethpmEncRtdIf.operSt),eq(ethpmEncRtdIf.operSt,"up"))))


  4) label_name          : interface

     path_name           : sys/mgmt-[mgmt0]
     query_type          : event
     query_condition     :
<illegible>
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     :
<illegible>


  3) label_name          : resources

     path_name           : sys/procsys/sysmem
     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&target-subtree-class=bgpDom,bgpPeer,bgpPeerAf,bgpDomAf,bgpPeerAfEntry,bgpOperRtctrlL3,bgpOperRtt,bgpOperRttEntry,bgpOperAfCtrl

switch#
```

About 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 Notifications
RIB	Sends all routes	Sends event notifications for create, update, and delete events. The following values are exported through telemetry for the RIB path: • Next-hop routing information: • Address of the next hop • Outgoing interface for the next hop • VRF name for the next hop • Owner of the next hop • Preference for the next hop • Metric for the next hop • Tag for the next hop • Segment ID for the next hop • Tunnel ID for the next hop • Encapsulation type for the next hop • Bitwise OR of flags for the Next Hop Type • For Layer-3 routing information: • VRF name of the route • Route prefix address • Mask length for the route • Number of next hops for the route • Event type • Next hops

Path Type	Subscription Baseline	Event Notifications
MAC	Executes a `GETALL` from DME for static and dynamic MAC entries	Sends event notifications for add, update, and delete events. The following values are exported through telemetry for the MAC path: • MAC address • MAC address type • VLAN number • Interface name • Event types Both static and dynamic entries are supported in event notifications.
Adjacency	Sends the IPv4 and IPv6 adjacencies	Sends event notifications for add, update, and delete events. The following values are exported through telemetry for the Adjacency path: • IP address • MAC address • Interface name • Physical interface name • VRF name • Preference • Source for the adjacency • Address family for the adjacency • Adjacency event type

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**.

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 395.

Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

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** *sgrp_id* **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 rib** **Example:** `nxosv2(conf-tm-sensor)# path rib` `nxosv2(conf-tm-sensor)#`	Configure the RIB path which streams routes and route update information.
Step 6	**destination-group** *grp_id* **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	**ip address** *ip_addr* **port** *port* **protocol** { **HTTP** | **gRPC** } **encoding** { **JSON** | **GPB** | **GPB-compact** } **Example:** `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)#`	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** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 9	**snsr-group** *sgrp_id* **sample-interval** *interval* **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 10	**dst-group** *dgrp_id* **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 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 395.

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**).

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** *sgrp_id* **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** *grp_id* **Example:** `switch(conf-tm-sensor)# destination-group 33` `switch(conf-tm-dest)#`	Enter telemetry destination group submode and configure the destination group.
Step 7	**ip address** *ip_addr* **port** *port* **protocol** { **HTTP** | **gRPC** } **encoding** { **JSON** | **GPB** | **GPB-compact** } **Example:** `switch(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json` `switch(conf-tm-dest)#` **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
	`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)#`	
Step 8	**subscription** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 9	**snsr-group** *sgrp_id* **sample-interval** *interval* **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 10	**dst-group** *dgrp_id* **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 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 395.

Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal` `switch(config)#`	Enter configuration mode.

	Command or Action	Purpose
Step 2	**telemetry** **Example:** `switch(config)# telemetry` `switch(config-telemetry)#`	Enter configuration mode for the telemetry features.
Step 3	**sensor-group** *sgrp_id* **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.
Step 5	**path adjacency** **Example:** `nxosv2(conf-tm-sensor)# path adjacency` `nxosv2(conf-tm-sensor)#`	Configure the Adjacency path which streams information about the IPv4 and IPv6 adjacencies.
Step 6	**destination-group** *grp_id* **Example:** `switch(conf-tm-sensor)# destination-group 33` `switch(conf-tm-dest)#`	Enter telemetry destination group submode and configure the destination group.
Step 7	**ip address** *ip_addr* **port** *port* **protocol** { **HTTP** | **gRPC** } **encoding** { **JSON** | **GPB** | **GPB-compact** } **Example:** `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)#`	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
Step 8	**subscription** *sub_id* **Example:** `switch(conf-tm-dest)# subscription 33` `switch(conf-tm-sub)#`	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 9	**snsr-group** *sgrp_id* **sample-interval** *interval* **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 10	**dst-group** *dgrp_id* **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.

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.

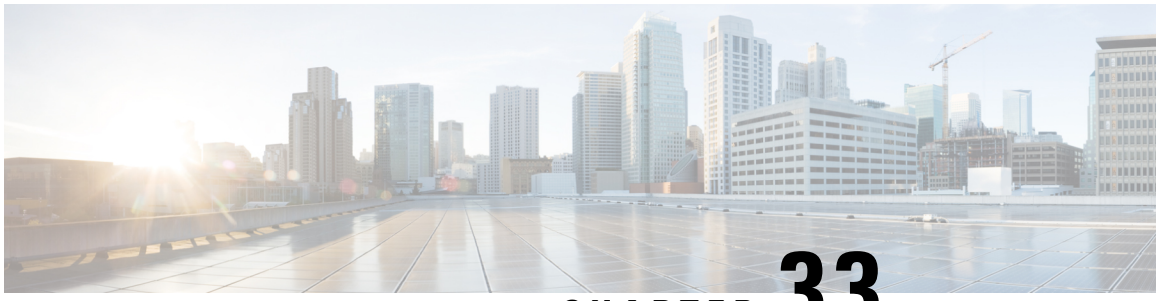
```
switch# show telemetry event collector errors

------------------------------------------------------------------------------
-
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#
```

Additional References

Related Documents

Related Topic	Document Title
Example configurations of telemetry deployment for VXLAN EVPN.	*Telemetry Deployment for VXLAN EVPN Solution*



CHAPTER 33

Inband Network Telemetry

This chapter contains the following topics:

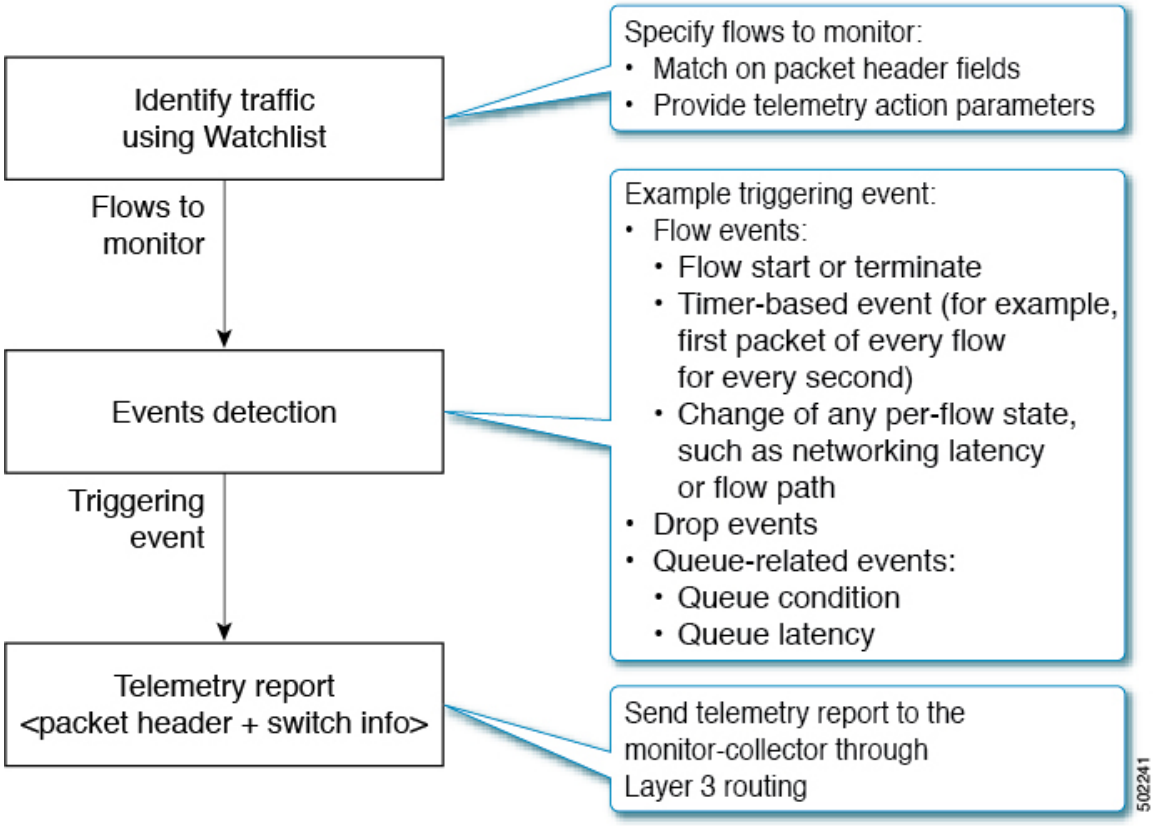
- About Inband Network Telemetry, on page 405
- Licensing Requirements for Inband and Postcard Telemetry, on page 410
- INT Collection Parameters, on page 411
- Example Inband Network Telemetry Headers, on page 411
- Telemetry Report Formats, on page 413
- Guidelines and Restrictions, on page 414
- Configuring Inband Network Telemetry, on page 415
- Flow and Queue Profile Arguments for Report Types, on page 421
- Configuring Packet Postcards, on page 421

About Inband Network Telemetry

Inband Network Telemetry (INT) is a framework that is designed to monitor, collect, and report flows and network states, by the data plane, without requiring intervention or work by the control plane. INT sources (applications, end-host networking stacks, hypervisors, NICs, send-side TORs, and so on) embed the monitored INT flow information in normal data packets, and all the downstream devices add the same information with their details that are based on the source information received.

Inband Network Telemetry is used to achieve per-packet network visibility with low overheads. The following figure shows the typical workflow for Inband Network Telemetry.

Figure 4: Example of Inband Network Telemetry Workflow



As part of the Inband Network Telemetry process, every data packet is inspected without interfering with other packet-processing logic in the switch data plane. Not all the components that are shown in the preceding figure are necessarily located in each INT switch. For example, the watchlist component can be in one switch, and the event detection and telemetry report components can be located in another switch.

A telemetry watchlist table specifies the flows to monitor. It performs a match on the packet headers and switch ports, and provides telemetry action parameters. Packets that match the watchlist entries with telemetry actions are processed by the event detection logic. If a triggering event is detected, the switch generates a telemetry report to the monitor-collector. The report message includes packet header and switch metadata associated with the packet (for example, timestamp, ingress or egress ports, and queue depth or latency).

Terminology

Following are terms used for this feature:

- **INT Header** — A packet header that carries INT information.
- **INT Packet** — A packet containing an INT Header.
- **INT Instruction** — Instructions that are embedded in the INT header, indicating which INT Metadata (defined below) to collect at each INT switch. The collected data is written into the INT Header.
- **INT Source** — A trusted entity that creates and inserts INT Headers into the packets it sends.

- **INT Sink** — A trusted entity that extracts the INT Headers and collects the path state that is contained in the INT Headers. The INT Sink is responsible for removing INT Headers so that INT is transparent to upper layers.
- **INT Metadata** — Information that an INT Source inserts into the INT Header. Examples of metadata are described in INT Collection Parameters, on page 411.
- **INT Domain** — A group of interconnected INT switches under the same administration.

INT Packet Flow

At INT, source flows to be monitored are filtered using the flow watchlist. Various pieces of information on the filtered flows are extracted, based on the collect parameters enabled, such as ingress or egress port ID, ingress or egress timestamp, switch ID, and queue occupancy. The data packets are encapsulated with the INT information and sent toward the destination. The telemetry INT reports are then sent to the monitor-collector, based on whether the events are flow events, packet-drop events, or queue-congestion events.

Report Types

Following are the different types of reports used by INT:

- **Local flow reports** — A general telemetry report, which is generated from flow events. Sent from the source or sink for host-to-host data flows matching the watchlist.
- **Drop reports** — A general telemetry report, which is generated from drop events. Sent for certain supported drops. Every INT-enabled switch sends these reports to the monitor-collector.
- **Queue Congestion reports** — A general telemetry report, which is generated from queue-related events. Sent for packets exceeding the queue depth or latency. Every INT-enabled switch sends these reports to the monitor-collector.
- **INT reports** — A report that is specific to INT. Sent by the sink. When INT-encapsulated data packets are received on the sink fabric port, two reports are generated by the sink:
  - Local report for traffic arriving on a fabric port
  - INT report for data that is received from the source

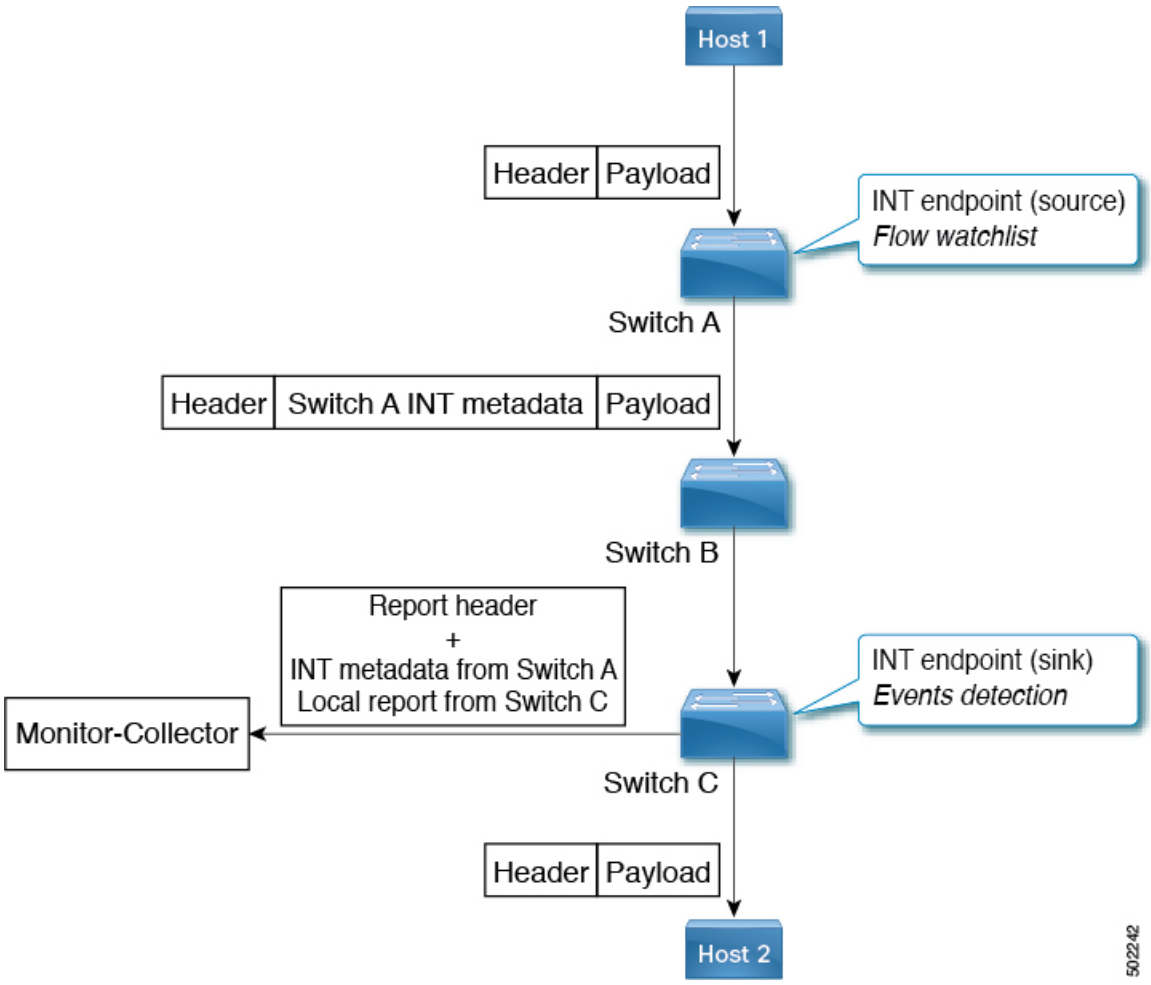
About Flow Events

Each switch can play the role of *endpoint* for INT-enabled packets. The endpoint acts both as *source* and *sink*, where:

- *Source* initiates INT operations by inserting a telemetry header into a packet and then instructing downstream network devices along the routing path to add desired telemetry information into the packet.
- *Sink* retrieves the embedded or collected INT information from the received data packets, extracts the telemetry information from the incoming packets, and sends telemetry reports to the monitor-collector if triggering flow events are detected.

The following figure shows an example flow event for Inband Network Telemetry. Switches along the route path add switch metadata into the packet header based on the telemetry instructions that are carried in the telemetry header.

Figure 5: Example Flow Event, Source, and Sink



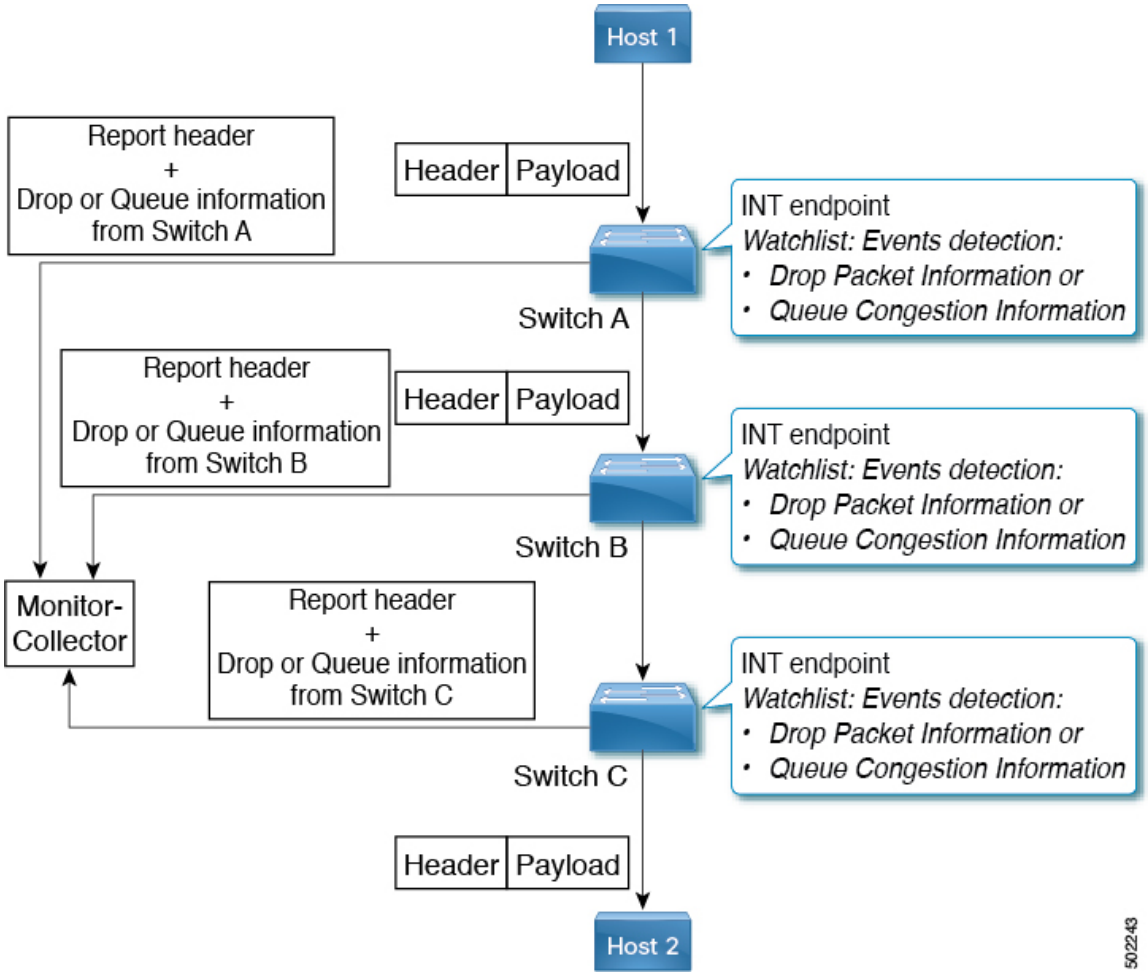
About Drop Packet and Queue Congestion Events

Each switch provides a report to the monitor-collector, based on the type of triggering event that occurred:

- **Dropped Packets** — Drop reports are generated for most dropped packets, except for STP-blocked ports and ACL deny intentional drop packets. This provides visibility into the impact of packet drops on user traffic.
- **Congested Queues** — Queue congestion reports are generated for traffic entering a specific queue during a period of queue congestion. This provides visibility into the traffic causing and prolonging queue congestion.

The following figure shows an example of drop and queue congestion reports.

Figure 6: Example Drop and Queue Congestion Reports



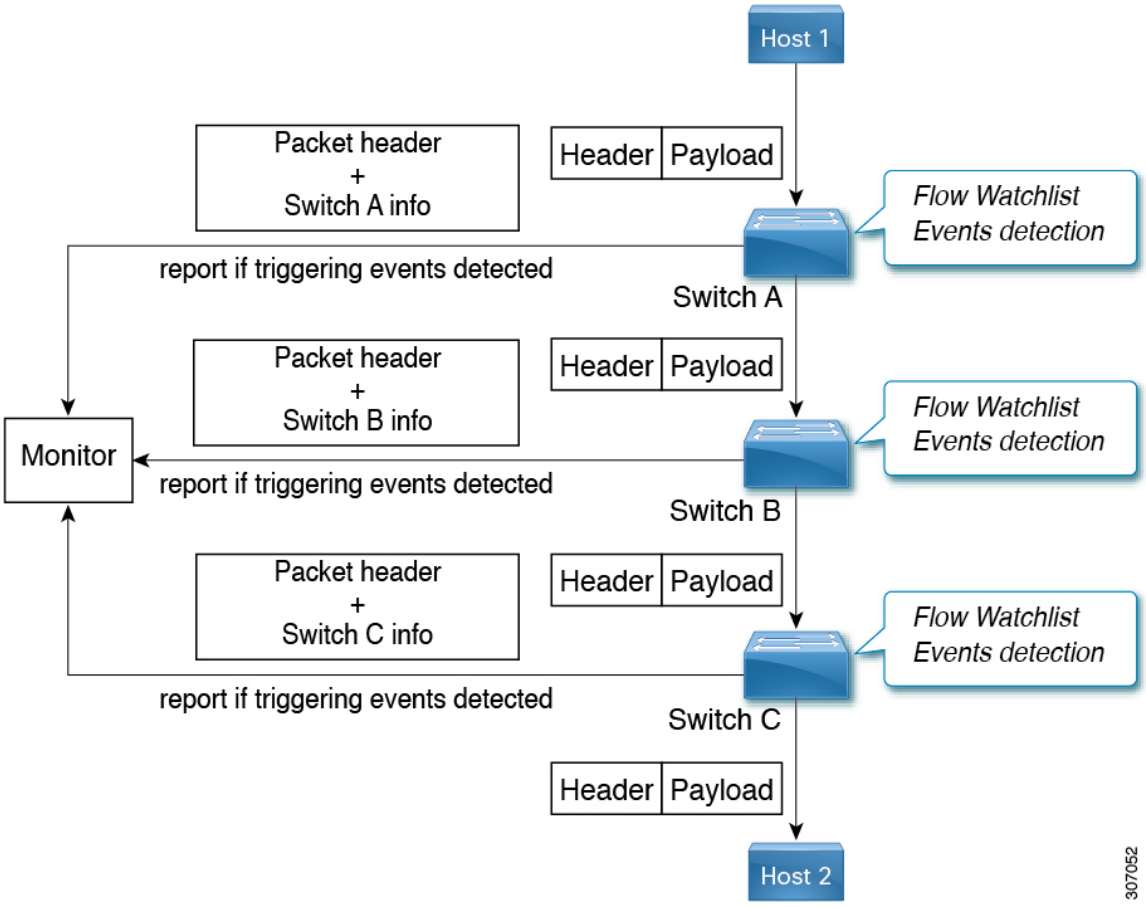
About Packet Postcards

Beginning in NX-OS release 9.2.(3), packet postcards are a different way of transmitting the telemetry data to a monitor. With packet postcards, the host's data packet is sent to another host, just like with the traditional model of INT, which uses source and sink switches as INT endpoints. However, the way that telemetry data is processed is different.

- In the traditional model, a host sends a flow, an INT source switch adds telemetry data to the host's packet, then the sink switch extracts all telemetry from each hop in the flow. The INT sink switch sends all telemetry data for the flow at once, which results in a larger amount of data that is transmitted to the monitor at one time.
- With packet postcards, each switch in the INT domain can extract telemetry data and send it to the collector. The result is faster convergence, distributed packet processing, and less overhead than with a traditional model.

The following figure shows an example of packet postcards on a flow.

Figure 7: Example of Packet Postcards



With packet postcards, the defined roles of source and sink switches become irrelevant because telemetry data is communicated at each switch.

- If the packet or flow matches the watchlist conditions, the comparison is true, and event-detection logic generates the report and sends it to the monitor. Triggering events are detected based on the switch local information such as ingress or egress ports and queueing latency for the monitored flow. Unlike standard Inband Network Telemetry, a postcard switch never modifies the original data packets by adding INT metadata. After telemetry data is sent to the collector, the packet is then forwarded onto the network to either its next hop or destination.
- If the packet or flow does not match the watchlist conditions, the comparison is false, and the event-detection logic does not process the telemetry data. The packet is forwarded by the switch, continuing along the route to its next hop. If it ingresses another switch in the INT domain, the packet is compared against that switch's watchlist, if necessary a packet postcard is sent to the monitor-collector, and the packet is forwarded until it reaches its eventual destination.

Licensing Requirements for Inband and Postcard Telemetry

The following table shows the licensing requirements for this feature:

Product	License Requirement
Cisco NX-OS	Inband Telemetry or Packet Postcards requires a N3K-STR1K9 license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the Cisco NX-OS Licensing Guide.

INT Collection Parameters

The following information can be collected for a given flow for INT encapsulation.

- **Switch ID** — The unique ID of a switch. Switch IDs must be unique within an INT domain. The switch ID is generated automatically from bytes [0][1][4][5] of the VDC MAC address. The switch ID is collected by default.

  For example, using the **show vdc** command to get the VDC MAC address on this device:

```
switch(config-int)# show vdc

vdc_id  vdc_name   state   mac                type        lc
------  --------   -----   ----------         ---------   ------
1       bod-int2   active  b0:12:80:8f:e3:45  Ethernet    None
```

  In this case, the switch ID for this switch is `0xB012E345`.

- **Ingress port identifier** — The port on which the INT packet was received.
- **Ingress timestamp** — The device local time when the INT packet was received on the ingress physical or logical port.
- **Egress port identifier** — The port on which the INT packet was sent out.
- **Egress timestamp** — The device local time when the physical or logical port processed the INT packet.
- **Queue occupancy** — The build-up of traffic in the queue (in bytes, cells, or packets) that the INT packet observes in the device while being forwarded.

Example Inband Network Telemetry Headers

This section shows example INT headers with two hosts (Host 1 and Host 2), communicating over a network path composed of three switches (Switch 1, Switch 2 and Switch 3) as shown in the following example.

```
==> packet P travels from Host 1 to Host 2 ==>

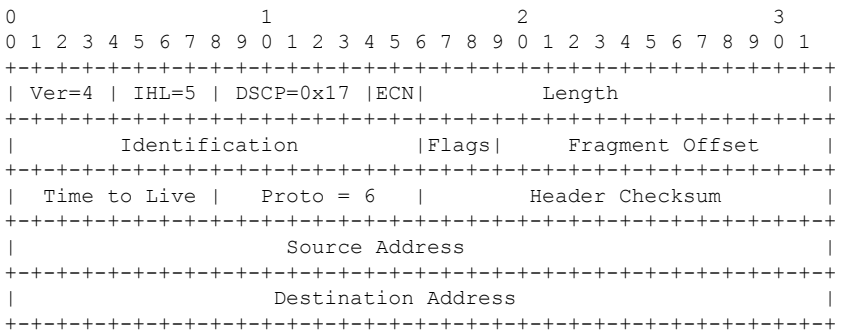
Host 1 --------> Switch 1 ---------> Switch 2 ---------> Switch 3 --------> Host 2
```

In this scenario, Host 1 sends a TCP packet to Host 2. The ToR switch of Host 1 (Switch 1) acts as the INT source. It adds INT headers and its own metadata in the packet.

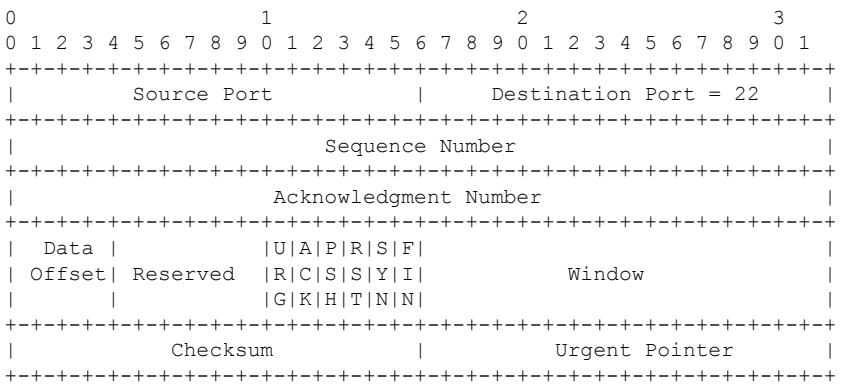
Finally, the ToR switch of Host 2 (Switch 3) acts as the INT sink and removes INT headers before forwarding the packet to Host 2.

The following is the packet that INT sink Switch 3 receives, starting from the IPv4 header.

IP Header

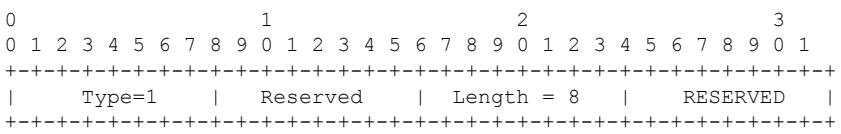


TCP Header



INT Shim Header for TCP/UDP

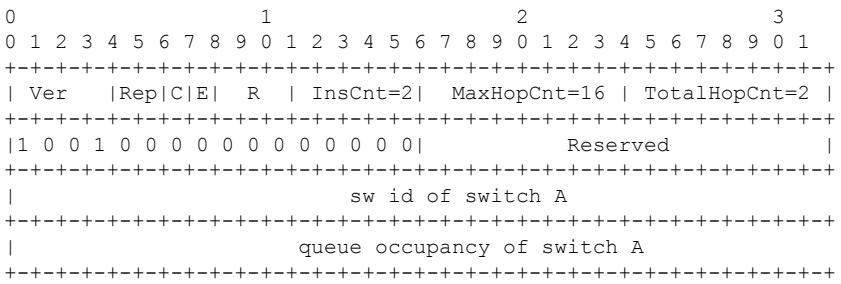
A shim header is a header that is inserted to allow for space in different packet protocol headers. In a TCP/UDP packet, following a TCP/UDP header, the shim header creates an encapsulation into which the INT headers are inserted. One boundary of the encapsulation is the TCP/UDP payload, while the other boundary of the INT header is the shim header.



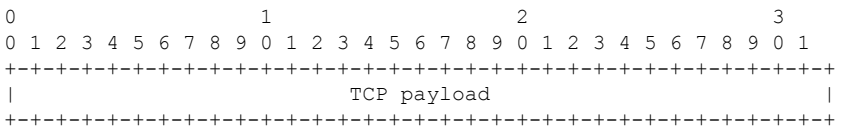
INT Metadata Header and Metadata Stack

The following shows a basic example of one set of instruction bit values. These instruction bit values correspond to one record consisting of one collect queue occupancy command that is configured through **inband telemetry record** command. No other collect commands are configured in the record.

The instruction set values in the INT Metadata header will differ depending on how a record is configured.



TCP Payload



Telemetry Report Formats

This section specifies the packet formats for telemetry reports.

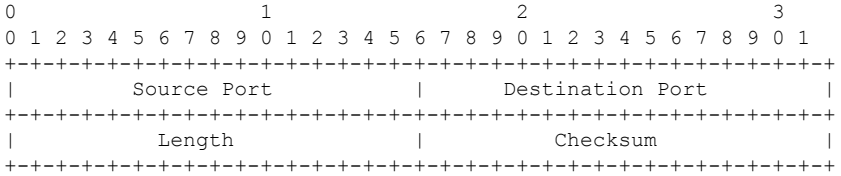
Outer Encapsulation

There are two types of reports:

- **Upstream reports** — The sink generates these reports only if there is a violation of a threshold for the flows being monitored. These reports contain the whole INT stack of embedded metadata up until the sink.
- **Local reports** — Each device generates these reports in cases of queue and drop violations. These reports contain information and metadata of the current device only.

Telemetry reports are defined using a UDP-based encapsulation. The source IP address identifies the switch that generates the telemetry report. The Destination IP address identifies a location in the distributed telemetry monitoring system that receives the telemetry report.

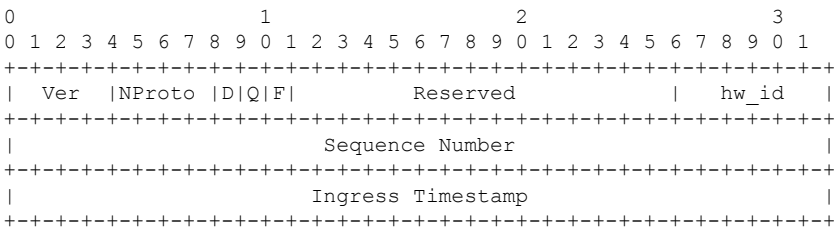
UDP header (8 Octets)



The Source Port can be used to carry flow entropy, for example based on a hash of the inner 5-tuple. Otherwise, it should be set to 0.

The Destination Port is user-configurable. Use the same Destination Port value for all telemetry reports in a particular deployment.

Telemetry Report Header (16+ Octets)



- **Ver** — Version 0.5.
- **NProt** — Next Protocol. Indicates:
  - 0: Ethernet
  - 1: Telemetry Drop Header, followed by Ethernet
  - 2: Telemetry Switch Local Header, followed by Ethernet
- **D|Q|F** — An enumeration that indicates the reason why a packet was dropped:
  - **D: Dropped** — Indicates that at least one packet matching a drop watchlist was dropped.
  - **Q: Congested Queue Association** — Indicates the presence of congestion on a monitored queue.
  - **F: Tracked Flow Association** — Indicates that this telemetry report is for a tracked flow (for example, if the packet matched a flow watchlist somewhere, for INT, or locally, for postcard). The report can include INT metadata beyond the inner Ethernet header. Other telemetry reports are likely to be received for the same tracked flow, from the same switch, and from other switches, for drop reports, postcards, or path changes.
- **hw_id** — Identifies the hardware subsystem within the switch that generated this report. For example, in a chassis with multiple line cards, the hw_id could identify a specific line card, or a subsystem within a line card.
- **Sequence Number** — Reflects the sequence of reports from a specific hw_id to a particular telemetry report destination. The sequence number can be used to detect loss of telemetry reports before they reach their intended destination.
- **Ingress Timestamp** — The device local time when the packet was first received on the ingress physical or logical port, in nanoseconds.

Guidelines and Restrictions

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

Following are the guidelines and restrictions for the INT feature:

- For information about supported platforms, see Supported Platforms for Programmability Features, on page 5.

  Ensure the data that you are using for the DSCP value is different from the DSCP value that the INT flow-profile configured.

- Only flows with the Layer 4 header are exported by INT. Layer 2-only or IP-only flows with no Layer 4 header are not supported. Support exists only for IP TCP, UDP, and ICMP flows.
- INT is supported only for IPv4. IPv6 and multicast is not supported.
- When INT is enabled, it adds bytes to the data packet. Verify that the required MTU is configured end-to-end throughout the network to accommodate for the increased overhead for the INT-encapsulated packet.
- Only Default VRF is supported for exporter reachability.
- Only IPv4 Destination is supported under the exporter.
- Upstream reports are not exported on Sink when SPAN is enabled. Only local reports are allowed.
- Drop reports that are exported are not throttled.
- By default, drop reporting is suppressed for both user ACL deny drops and drops on spanning-tree (STP) blocked ports.
- Ingress MTU failure drops are not reported, as the drops happen before the reporting occurs.
- The current scale limit is 20,000 flows per pipeline.
- As a best practice, Cisco recommends using specific flow watchlists to filter flows for reporting.
- INT reports timestamp is 32 bit, which will not reflect the exact time when the flow comes to the switch. You will need a mechanism on the monitor-collector to correlate and derive the time for when the flow is received.
- Only one monitor-collector is currently supported.
- The programming of the watchlist hardware entries may take a minute or two if the number of watchlist entries is greater than 345.

Configuring Inband Network Telemetry

Telemetry flow reports are generated for flows matching the watchlist. For drop and queue reports, the watchlist is not required.

Before you begin

Make sure that standard packet postcard mode is not enabled. Traditional INT and packet postcards are mutually exclusive. If standard packet postcards are configured, and you try to configure INT, the following error is displayed.

```
Error: Cannot configure inband-telemetry when postcard is configured.
```

If you see this error, you must disable standard INT:

```
switch(config)# no feature postcard-telemetry
```

Also, INT and packet postcard mode are licensed features. Make sure that you have enabled the correct license. See Licensing Requirements for Inband and Postcard Telemetry, on page 410.

Procedure

Step 1 **For all three types of INT reports:** Enable the **hardware-telemetry** feature on all switches in the INT domain.

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# feature hardware-telemetry
HW Telemetry enabled
switch(config)#
switch(config)# hardware-telemetry inband-telemetry
```

Step 2 **For flow reports:** For each switch in the INT domain, configure an INT record.

```
switch(config)# inband-telemetry record record-ID
switch(config-int-record)# collect port-id
switch(config-int-record)# collect queue-occupancy
switch(config-int-record)# collect ingress-timestamp
switch(config-int-record)# collect egress-timestamp
switch(config-int-record)# exit
```

Where *record-ID* is the name of the record. The default switch-id option is present when a new record is created.

This command affects only INT reports. Drop reports, queue congestion reports, and local flow reports are not affected.

For example:

```
switch(config)# inband-telemetry record r1
switch(config-int-record)# collect port-id
switch(config-int-record)# collect queue-occupancy
switch(config-int-record)# collect ingress-timestamp
switch(config-int-record)# collect egress-timestamp
switch(config-int-record)# exit
```

Step 3 **For all three types of INT reports:** For each switch in the INT domain, configure an INT exporter.

```
switch(config)# inband-telemetry exporter exporter-ID
switch(config-int-exporter)# destination destination-IP-address
switch(config-int-exporter)# source loopback loopback-interface-number
switch(config-int-exporter)# transport udp udp-port
switch(config-int-exporter)# exit
```

Where:

- *exporter-ID* is the name that you want to give to the exporter on this particular switch.
- *destination-IP-address* is a destination reachable through the default VRF.
- *loopback-interface-number* is the number for the loopback interface.

- *udp-port* is the destination UDP port. Acceptable values are from 1 to 65535. The default UDP port where exports are sent is 31337.

For example:

```
switch(config)# inband-telemetry exporter e1
switch(config-int-exporter)# destination 17.18.19.20
switch(config-int-exporter)# source loopback 36
switch(config-int-exporter)# transport udp 1000
switch(config-int-exporter)# exit
```

Step 4 **For all three types of INT reports:** For each switch in the INT domain, configure an INT watchlist.

A watchlist is mandatory when the monitor is applied at the system level. A watchlist is applicable to flow monitoring only. A watchlist is not effective for drop or queue congestion reports.

Be careful when configuring watchlists. Try to avoid configuring watchlists that match on control traffic going to the local CPU and generated from the local CPU.

```
switch(config)# inband-telemetry watchlist ip watchlist-ID
switch(config-int-watchlist)# 20 permit ip permit-IP-address-subnet-1
permit-IP-address-subnet-2
switch(config-int-watchlist)# 10 deny ip deny-IP-address-subnet-1 deny-IP-address-subnet-2
switch(config-int-watchlist)# exit
```

Where:

- *watchlist-ID* is the name of the watchlist. A maximum of 512 entries are supported under one watchlist. Watchlist entries are needed on the source to match specific flows that must be exported.
- *permit-IP-address-subnet-#* are the permit IP addresses and subnets.
- *deny-IP-address-subnet-#* are the deny IP addresses and subnets.

For example:

```
switch(config)# inband-telemetry watchlist ip wl1
switch(config-int-watchlist)# 20 permit ip 1.2.3.4/24 5.6.7.8/32
switch(config-int-watchlist)# 10 deny ip 20.30.40.50/24 60.70.80.90/16
switch(config-int-watchlist)# exit
```

Step 5 **For all three types of INT reports:** For each switch in the INT domain, configure an INT monitor.

```
switch(config)# inband-telemetry monitor monitor-ID
switch(config-int-monitor)# record record-ID
switch(config-int-monitor)# exporter exporter-ID
switch(config-int-monitor)# watchlist watchlist-ID
switch(config-int-monitor)# exit
```

Where:

- *monitor-ID* is the name of this INT monitor.
- *record-ID* is the name of the record that you entered in Step 2, on page 416.
- *exporter-ID* is the name of the exporter that you entered in Step 3, on page 416.
- *watchlist-ID* is the name of the watchlist that you entered in Step 4, on page 417.

For example:

```
switch(config)# inband-telemetry monitor m1
switch(config-int-monitor)# record r1
switch(config-int-monitor)# exporter e1
switch(config-int-monitor)# watchlist wl1
switch(config-int-monitor)# exit
```

Step 6 **For flow reports:** Configure a port-type fabric to encapsulate the INT domain.

The port-type fabric must be enabled on all links connecting leaf and spine.

```
switch(config)# interface ethernet interface1/interface2
switch(config-if)# port-type fabric
switch(config-if)# exit
```

Where *interface#* is the Ethernet interface that you are configuring.

For example:

```
switch(config)# interface ethernet 1/1
switch(config-if)# port-type fabric
switch(config-if)# exit
```

Step 7 **For all three types of INT reports:** Apply an INT monitor at both the source and the sink.

```
switch(config-int-record)# inband-telemetry system monitor monitor-ID
```

Where *monitor-ID* is the name of the INT monitor.

For example:

```
switch(config-int-record)# inband-telemetry system monitor m1
```

Step 8 **For flow reports (optional):** Configure a flow profile and queue profile.

```
switch(config-int-record)# inband-telemetry flow-profile
switch(config-int-flow-prof)# dscp dscp-value
switch(config-int-flow-prof)# age age-value
switch(config-int-flow-prof)# latency quantization latency-quant-value
exit

switch(config-int-record)# inband-telemetry queue-profile queue-profile-default
switch(config-int-q-prof)# depth depth-value
switch(config-int-q-prof)# latency latency-value
exit
```

Where the following arguments are applicable to different types of reports. See Flow and Queue Profile Arguments for Report Types, on page 421.

- *dscp-value* is the flow-profile DSCP value that is used for INT. The value must match on both the source and the sink for INT encapsulation and decapsulation to happen. This value must not overlap with the DSCP values that are used in the data traffic. Allowed values: 1, 2, 4, 8, 16, 32, 63. Default is 1.
- *age-value* is the flow-profile age, in seconds, when the exports are sent out. Used to determine the frequency at which a given flow is reported. Allowed values: 0 to 7200. Default is 30.
- *latency-quant-value* is the flow-profile latency quantization value. Flow reports are generated when the flow is exceeding the configured latency value. Only flows that match the watchlist are generated. This

value is converted to a power of 2, in nanoseconds. For example, a value of 11 is converted to 2 power 11, or 2048. This calculation means that if the difference between the ingress timestamp and the egress timestamp is more than 2048 nanoseconds at any point in the path, one upstream or local report is generated per flow. Allowed values are 0 to 28. Default is 11.

- *depth-value* is the queue-profile depth value, in number of cells. Congestion queue reports are generated for flows exceeding the configured queue depth. Allowed values: 1 to 204800. Default is 300.

  Adjust the depth value as needed for the number of reports that are generated. For example, if your deployment is generating too many reports, increase the depth value to reduce the number of reports. Conversely, if you need to generate more reports, lower the depth value.

- *latency-value* is the queue-profile latency, in nanoseconds. The latency value is the time from when the packet is put into the queue to the time when it is dequeued from the queue. Allowed values: 30 to 8192000. Default is 2048. If a violation happens, a queue congestion report is generated from the local switch when the congestion occurs. If needed, adjust the latency value to increase or decrease the number of reports you want in your deployment.

  Note
  The *depth-value* and *latency-value* arguments are compared, and the lowest value is used, so you do not need to configure both. If you configure one argument, the default value for the other argument is used for the comparison.

  Note
  The *depth-value* and *latency-value* are used to derive the queue quota which controls creation of queue congestion reports. Queue congestion reports are sent when either the depth or latency that is defined for the queue profile is hit. The quota controls the number of congestion queue reports sent. You can display the queue quota value through **show inband-telemetry queue-profile**.

For example:

```
switch(config-int-record)# inband-telemetry flow-profile
switch(config-int-flow-prof)# dscp 1
switch(config-int-flow-prof)# age 10
switch(config-int-flow-prof)# latency quantization 28
exit

switch(config-int-record)# inband-telemetry queue-profile queue-profile-default
switch(config-int-q-prof)# depth 300
switch(config-int-q-prof)# latency 8192000
exit
```

Step 9 Verify that the Inband Network Telemetry was configured correctly.

- To view the INT exporter configuration information:

  ```
  switch(config-int-watchlist)# show inband-telemetry exporter exporter-ID
  ```

  For example:

  ```
  switch(config-int-watchlist)# show inband-telemetry exporter e1
  INT exporter e1:
      Destination: 10.1.1.2
      VRF: default
      Destination UDP Port 1000
      Source Interface mgmt0 (209.165.202.129)
  ```

- To view the INT watchlist configuration information:

```
switch(config-int-watchlist)# show inband-telemetry watchlist watchlist-ID
```

For example:

```
switch(config-int-watchlist)# show inband-telemetry watchlist wl1
INT watchlist wl1:
    No. of users: 1
    No. of aces: 1
        20 permit ip 1.2.3.4/24 5.6.7.8/32
```

The **No. of users:** field displays the number of ACE entries that are used in the watchlist when the monitor is applied.

- To view the INT monitor configuration information:

```
switch(config-int-monitor)#  show inband-telemetry monitor monitor-ID
```

For example:

```
switch(config-int-monitor)#  show inband-telemetry monitor m1
INT Monitor m1:
    INT Record: r1
    INT Exporter: e1
    INT Watchlist: wl1
switch(config-int-monitor)#
```

- To display the monitor session that is active at the system level:

```
switch(config-int-watchlist)# show inband-telemetry sessions
INT Sessions:
     system monitor:m1
```

- To view the running configuration of the Inband Network Telemetry:

```
switch(config)# show running-config hardware-telemetry
!Running configuration last done at: Tue Oct 30 15:24:08 2018
!Time: Tue Oct 30 15:24:13 2018

version 9.2(2) Bios:version 05.25
feature hardware-telemetry

hardware-telemetry inband-telemetry
  inband-telemetry exporter e1
    destination 10.1.1.2
    transport udp 1000
  inband-telemetry record r1
    collect switch-id
    collect port-id
    collect queue-occupancy
    collect ingress-timestamp
    collect egress-timestamp
  inband-telemetry watchlist ip wl1
    20 permit ip 1.2.3.4/24 5.6.7.8/32
  inband-telemetry monitor m1
    record r1
    exporter e1
    watchlist wl1
  inband-telemetry queue-profile queue-profile-default
    depth 300
```

```
  latency 8192000
inband-telemetry flow-profile
  dscp 1
  age 10
  latency quantization 28
inband-telemetry system monitor m1
```

Flow and Queue Profile Arguments for Report Types

The following table shows the arguments for the **inband-telemetry flow-profile** and **inband-telemetry queue-profile** commands and how they apply to the different report types. Unless documented, the arguments apply to both INT and Packet Postcards.

Table 24: Flow Reports for Specific Arguments

Report Type	Command and Arguments
Only flow reports **Note** The DSCP value is applicable to INT, but not Packet Postcard mode.	**inband-telemetry flow-profile dscp** *dscp-value* **inband-telemetry flow-profile age** *age-value*
Both flow and queue congestion reports	**inband-telemetry flow-profile latency-quantization** *latency-quant-value*
Only queue congestion reports	**inband-telemetry queue-profile depth** *depth-value* **inband-telemetry queue-profile latency** *latency-value*

Configuring Packet Postcards

With packet postcards, each switch in the INT domain performs a watchlist comparison and events detection. When an event is triggered, each switch sends its own telemetry data to the monitor-collector independent of the other switches. Each switch requires a watchlist and events detection logic.

Before you begin

Make sure that standard INT is not enabled. Packet postcard mode and the traditional INT model are mutually exclusive. If standard INT is configured, and you try to configure packet-postcard mode, the following error is displayed.

```
Error: Cannot configure postcard when inband-telemetry is configured.
```

If you see this error, you must disable standard INT:

```
switch(config)# no feature hardware-telemetry
```

Also, INT and packet postcard mode are licensed features. Make sure that you have enabled the correct license. See Licensing Requirements for Inband and Postcard Telemetry, on page 410.

Procedure

Step 1 Enable the **postcard-telemetry** feature on all switches in the INT domain.

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# feature hardware-telemetry
HW Telemetry enabled
switch(config)#
switch(config)# hardware-telemetry postcard-telemetry
switch(config-postcard)#
```

Step 2 On each switch in the INT domain, configure the appropriate profile:

a) **For a flow profile:** Configure the flow profile parameters.

```
switch(config-postcard)# postcard-telemetry flow-profile
switch(config-postcard-flow-prof)# age age-value
switch(config-postcard-flow-prof)# latency quantization latency-quant-value
exit
```

Where the following arguments are applicable to different types of reports. See Flow and Queue Profile Arguments for Report Types, on page 421.

- *age-value* is the flow-profile age, in seconds, when the exports are sent out. Used to determine the frequency at which a given flow is reported. Allowed values: 0 to 7200. Default is 30.
- *latency-quant-value* is the flow-profile latency quantization value. Flow reports are generated when the flow exceeds the configured latency value. Only flows that match the watchlist are generated. This value is converted to a power of 2, in nanoseconds. For example, a value of 11 is converted to 2 power 11, or 2048. This calculation creates the difference between the ingress timestamp and the egress timestamp. If the difference is more than 2048 nanoseconds at any point in the path, one upstream or local report is generated per flow. When the difference between the ingress and egress timestamp goes below the configured value, a recovery report is sent. Allowed values are 0 to 28. Default is 11.

For example:

```
switch(config-postcard)# postcard-telemetry flow-profile
switch(config-postcard-flow-prof)# age 45
switch(config-postcard-flow-prof)# latency quantization 4
switch(config-postcard-flow-prof)# exit
switch(config-postcard)#
```

b) **For a queue congestion profile:** Configure the watchlist parameters for the default queue profile.

```
switch(config-postcard)# postcard-telemetry queue-profile queue-profile-default
switch(config-postcard-q-prof)# depth depth-value
switch(config-postcard-q-prof)# latency latency-value
exit
```

Where the following arguments are applicable to different types of reports. See Flow and Queue Profile Arguments for Report Types, on page 421.

- *depth-value* is the queue-profile depth value, in number of cells. Congestion queue reports are generated for flows exceeding the configured queue depth. Allowed values: 1 to 204800. Default is 300.

  Adjust the depth value as needed for the number of reports that are generated. For example, if your deployment is generating too many reports, increase the depth value to reduce the number of reports. Conversely, if you need to generate more reports, lower the depth value.

- *latency-value* is the queue-profile latency, in nanoseconds. The latency value is the time from when the packet is put into the queue to the time when it is dequeued from the queue. Allowed values: 30 to 8192000. Default is 2048. If a violation happens, a queue congestion report is generated from the local switch when the congestion occurs. If needed, adjust the latency value to increase or decrease the number of reports you want in your deployment.

  Note
  The *depth-value* and *latency-value* arguments are compared, and the lowest value is used, so you do not need to configure both. If you configure one argument, the default value for the other argument is used for the comparison.

  Note
  The *depth-value* and *latency-value* are used to derive the queue quota which controls creation of queue congestion reports. Queue congestion reports are sent when either the depth or latency that is defined for the queue profile is hit. The quota controls the number of congestion queue reports that are sent. You can display the queue quota value through **show postcard-telemetry queue-profile**.

For example:

```
switch(config-postcard)# postcard-telemetry queue-profile queue-profile-default
switch(config-postcard-q-prof)# depth 200
switch(config-postcard-q-prof)# latency 2048
exit
```

Step 3 On each switch, configure a postcard exporter.

```
switch(config)# postcard-telemetry exporter exporter-ID
switch(config-postcard-exporter)# destination destination-IP-address
switch(config-postcard-exporter)# source interface-type
switch(config-postcard-exporter)# transport udp udp-port
switch(config-postcard-exporter)# exit
```

Where:

- *exporter-ID* is the name that you want to give to the exporter on this particular switch.
- *destination-IP-address* is a destination of the monitor-collector reachable through the default VRF. The destination can be an IPv4 address.
- *source* is any of the following interface types:
  - *ethernet* plus an *ethernet-interface* specifies an Ethernet IEEE 802.3z interface.
  - *loopback* plus a *loopback-interface* specifies a loopback interface.
  - *mgmt* plus a *management-interface* specifies a management interface.
  - *port-channel* plus a *port-channel-interface* specifies a port-channel interface.
  - *vlan* plus a *vlan-id* specifies a VLAN interface.

- *udp-port* is the destination UDP port. Acceptable values are from 1 to 65535. The default UDP port where exports are sent is 31337.

For example:

```
switch(config)# postcard-telemetry exporter ex4
switch(config-postcard-exporter)# destination 192.0.2.12
switch(config-postcard-exporter)# source loopback 40
switch(config-postcard-exporter)# transport udp 999
switch(config-postcard-exporter)# exit
```

Step 4 On each switch in the INT domain, configure a watchlist.

A watchlist is mandatory when the monitor is applied at the system level. However, a watchlist is not effective for drop or queue congestion reports.

Be careful when configuring a watchlist. Try to avoid configuring watchlists that match on control traffic going to the local CPU and generated from the local CPU.

```
switch(config)# postcard-telemetry watchlist ip watchlist-ID
switch(config-postcard-watchlist)# sequence-number permit ip permit-IP-address-subnet-1
permit-IP-address-subnet-2
switch(config-postcard-watchlist)# sequence-number deny ip deny-IP-address-subnet-1
deny-IP-address-subnet-2
switch(config-postcard-watchlist)# exit
```

Where:

- *watchlist-ID* is the name of the watchlist. A maximum of 512 entries are supported under one watchlist. Watchlist entries are needed on the source to match specific flows that must be exported.
- *permit-IP-address-subnet-#* are the permit IP addresses and subnets. When you set permit, flows on the specified IP addresses and subnets are monitored.
- *deny-IP-address-subnet-#* are the deny IP addresses and subnets. When you set deny, flows on the specified IP addresses and subnets are not monitored.

For example:

```
switch(config)# postcard-telemetry watchlist ip watch1
switch(config-postcard-watchlist)# 25 permit ip 10.10.20.30/24 192.0.2.1/32
switch(config-postcard-watchlist)# 30 deny ip 192.0.2.50/24 192.168.70.80/24
switch(config-postcard-watchlist)# exit
```

Step 5 For each switch in the INT domain, configure a monitor.

```
switch(config)# postcard-telemetry monitor monitor-ID
switch(config-postcard-monitor)# exporter exporter-ID
switch(config-postcard-monitor)# watchlist watchlist-ID
switch(config-postcard-monitor)# exit
```

Where:

- *monitor-ID* is the name of this INT monitor.
- *exporter-ID* is the name of the exporter that you entered in Step 3, on page 423.
- *watchlist-ID* is the name of the watchlist that you entered in Step 4, on page 424.

For example:

```
switch(config)# postcard-telemetry monitor mon10
switch(config-postcard-monitor)# exporter ex4
switch(config-postcard-monitor)# watchlist watch10
switch(config-postcard-monitor)# exit
```

Step 6 Apply a monitor at all switches in the INT domain.

```
switch(config-postcard)# postcard-telemetry system monitor monitor-ID
```

Where *monitor-ID* is the name of the INT monitor.

For example:

```
switch(config-postcard)# postcard-telemetry system monitor mon10
```

Step 7 Verify that packet postcard telemetry is configured correctly.

- To view the postcard exporter configuration information:

  ```
  switch(config-postcard-exporter)# show postcard-telemetry exporter exporter-ID
  ```

  For example:

  ```
  switch(config-postcard-exporter)# show postcard-telemetry exporter name ex4
  POSTCARD exporter ex4:
      Destination: 192.0.2.12
      VRF: default
      Destination UDP Port 999
      Source Interface loopback40
  switch(config-postcard-exporter)#
  ```

- To view the postcard telemetry-watchlist configuration information:

  ```
  switch(config-postcard-watchlist)# show postcard-telemetry watchlist watchlist-ID
  ```

  For example:

  ```
  switch(config-postcard-watchlist)# show postcard-telemetry watchlist 27
  No. of users: 1
      No. of aces: 2
          30 permit ip 192.0.2.50/24 192.168.70.80/24
          100 deny ip 10.10.10.10/32 1.2.3.4/24

  switch(config-postcard-watchlist)#
  ```

  The **No. of users:** field displays the number of ACE entries in the watchlist that are used when the monitor is applied.

- To view the postcard monitor configuration information for packet postcards:

  ```
  switch(config-postcard-monitor)#  show postcard-telemetry monitor monitor-ID
  ```

  For example:

  ```
  switch(config-postcard-monitor)# show postcard-telemetry monitor mon10
  POSTCARD Monitor mon10:
      POSTCARD Exporter: ex4
      POSTCARD Watchlist: 27
  ```

```
switch(config-postcard-monitor)#
```

- To display the monitor session that is active at the system level:

```
switch(config-postcard-monitor)# show postcard-telemetry sessions
POSTCARD Sessions:
     system monitor:mon10
switch(config-postcard-monitor)#
```

- To display the queue profile configuration:

```
switch(config-postcard-q-prof)# show postcard-telemetry queue-profile
POSTCARD Queue Profile queue-profile-default:
    Fields:
        Depth: 200
        Latency: 444
        Quota configured: 34
```

- To display the flow profile configuration:

```
switch(config-postcard-flow-prof)# # show postcard-telemetry flow-profile
POSTCARD flow Profile flow-profile-default:
    Fields:
        AGE: 45
        Latency Quantization: 16 nanosec
                (User configured quantization value is 4)
```

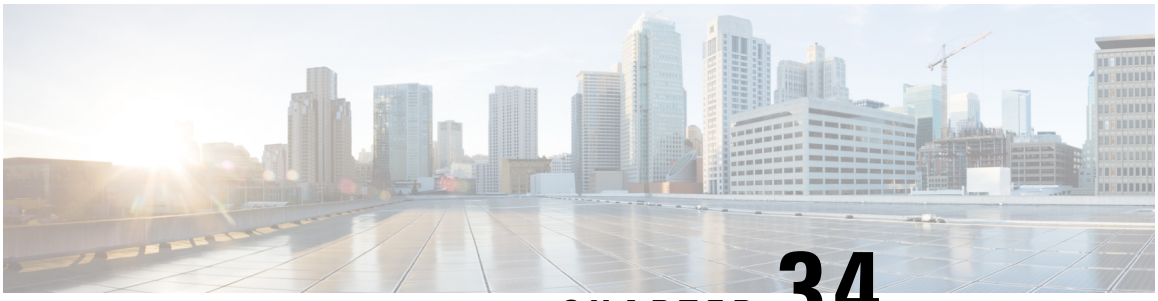
- To view the running configuration of telemetry for packet postcards:

```
switch(config)# show running-config hardware-telemetry

!Command: show running-config hardware-telemetry
!Running configuration last done at: Fri Nov 30 18:59:24 2018
!Time: Fri Nov 30 18:59:47 2018

version 9.2(3) Bios:version 05.25
feature hardware-telemetry

hardware-telemetry postcard-telemetry
  postcard-telemetry exporter ex4
    destination 192.02.12
    transport udp 999
    source loopback40
  postcard-telemetry watchlist ip 27
    30 permit ip 192.0.2.50/24 192.168.70.80/24
    100 deny ip 10.10.10.10/32 1.2.3.4/24
  postcard-telemetry monitor mon10
    exporter ex4
    watchlist 27
  postcard-telemetry queue-profile queue-profile-default
    depth 200
    latency 444
  postcard-telemetry flow-profile
    age 45
    latency quantization 4
  postcard-telemetry system monitor mon10
```



CHAPTER 34

Packetized MMU Stats

This chapter contains the following sections:

- About Packetized MMU Stats, on page 427
- Guidelines and Limitations, on page 427
- Using Packetized MMU Stats, on page 427

About Packetized MMU Stats

Beginning with Cisco NX-OS release 9.2(2), the Packetized MMU Stats (PSTATS) is now supported. PSTATS uses the ASIC capability to read stats using direct memory access (DMA) instead of reading port-by-port or queue-by-queue. It uses the streaming telemetry framework to push data to a receiver.

Guidelines and Limitations

The PSTATS feature has the following guidelines and limitations:

- For information about supported platforms, see the Nexus Switch Platform Matrix.
- Because of a hardware limitation, either Out-of-Band or PSTATS can be enabled at one time.
- You must change the Broadcom configuration variable to use this feature, which also requires a reload.

Using Packetized MMU Stats

Either the Out-of-Band or the PSTATS feature can be enabled at one time. By default, the Out-of-Band feature will be enabled.

Following are the commands to use for the PSTATS feature:

- To disable the default Out-of-Band feature and enable PSTATS (to begin fetching data using the PSTAT DMA):

  ```
  hardware profile statistics pstat [peak]
  ```

  You must reload to allow the enabling of the PSTATS to take affect:

  ```
  switch# reload
  ```

The following displays all the commands entered by you that require a reload, and what is displayed with relation to PSTATS:

```
switch(config)# show system config reload-pending

Following config commands require copy r s + reload :
=====================================================
0       hardware profile statistics pstat
=====================================================
```



Note If you want to disable PSTATS and re-enable the default Out-of-Band feature, enter:

```
no hardware profile statistics pstat [peak]
```

You must reload to allow the enabling of the default Out-of-Band feature to take affect:

```
switch# reload
```

- To get the PSTATS statistics:

  show hardware internal buffer info pstats [interface <if-name-range>][service-pool][detail]

  The telemetry framework uses the command to read the data from the hardware, registers all ports one by one, and pushes the data to the collector.

  For example:

```
telemetry  destination-group 100
    ip address 171.68.248.76 port 50001 protocol gRPC encoding GPB
    sensor-group 100    data-source NX-API
    path "show hardware internal buffer info pkt-stats pstats" depth unbounded
    subscription 100    dst-grp 100    snsr-grp 100 sample-interval 5000
```

- To get debugging information:
  - At the DME level:

```
switch# show system internal dme running-config all dn sys/pltfm | grep pstat

      "pstatCfg": "PSTAT_ENABLE_PEAK",
```

  - At the `pltfm_config` level:

```
switch# show system internal pltfm_config info all | grep -A 10 "STAT"
Buffer stat collect type and mode OOBSTA/PSTAT
Buf stat collect type OOBSTAT/PSTAT:1 [OOBSTAT]
Buf stat collect mode Instantaneous/Peak: 0 [Instantaneous]
```

  - At the BCM level:

```
switch # bcm-shell module 1 "config show" | grep -i buffer
buffer_stats_collect_mode=0
buffer_stats_collect_type=1
switch#
```

- Verify the telemetry configuration using the **show run telemetry** command, as shown in this example:

```
switch# show run telemetry

!Command: show running-config telemetry
!Time: Tue Mar  6 18:36:38 2018

version 9.2(2)
feature telemetry

telemetry
  destination-group 100
    ip address 171.68.248.76 port 50001 protocol gRPC encoding GPB
  sensor-group 100
    data-source NX-API
    path "show hardware internal buffer info pkt-stats" depth unbounded
  subscription 100
    dst-grp 100
    snsr-grp 100 sample-interval 5000
```

Following is example file output when running telemetry on an ADS server:

```
D0308 16:38:23.229714038    2519 env_linux.c:77]             Warning: insecure environment
 read function 'getenv' used
E0308 16:38:23.230022365    2519 tcp_server_posix.c:148]     check for SO_REUSEPORT:
{"created":"@1520555903.230006737","description":"SO_REUSEPORT unavailable on compiling
system","file":"src/core/lib/iomgr/socket_utils_common_posix.c","file_line":175}
Server listening on 0.0.0.0:50001
Received GBP RPC at: Thu Mar  8 16:38:25 2018
Data size is: 722
header {
  node_id_str: C14-HAV-2029
  encoding_path: show hardware internal buffer info pkt-stats pstats
  collection_id: 30
  data_gpbkv {
    timestamp: 0
    name:
    fields {
      timestamp: 0
      name: keys
      value (invalid):
        fields {
          timestamp: 0
          name: show hardware internal buffer info pkt-stats pstats
          value: show hardware internal buffer info pkt-stats pstats
        }
    }
    fields {
      timestamp: 0
      name: content
      value (invalid):
        fields {
          timestamp: 0
          name:
          value:
        subfields {
          name: TABLE_module
            fields {
              timestamp: 0
              name:
              value:
```

```
subfields {
  name: ROW_module
    fields {
      timestamp: 0
      name:
      value:
    subfields {
      name: module_number
      timestamp: 0
      value (uint64): 1
      }
    subfields {
      name: TABLE_instance
        fields {
          timestamp: 0
          name:
          value:
        subfields {
          name: ROW_instance
            fields {
              timestamp: 0
              name:
              value:
            subfields {
              name: instance
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: supports_8q
              timestamp: 0
              value (uint64): 1
              }
            subfields {
              name: total_instant_usage_1
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: total_instant_usage_2
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: total_instant_usage_3
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: total_instant_usage_4
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: rem_instant_usage_1
              timestamp: 0
              value (uint64): 63764
              }
            subfields {
              name: rem_instant_usage_2
              timestamp: 0
              value (uint64): 0
              }
            subfields {
```

```
                                name: rem_instant_usage_3
                                timestamp: 0
                                value (uint64): 0
                                }
                              subfields {
                                name: rem_instant_usage_4
                                timestamp: 0
                                value (uint64): 6960
                                }
                              subfields {
                                name: max_cell_usage_1
                                timestamp: 0
                                value (uint64): 0
                                }
                              subfields {
                                name: max_cell_usage_2
                                timestamp: 0
                                value (uint64): 0
                                }
                              subfields {
                                name: max_cell_usage_3
                                timestamp: 0
                                value (uint64): 0
                                }
                              subfields {
                                name: max_cell_usage_4
                                timestamp: 0
                                value (uint64): 212
                                }
                              subfields {
                                name: switch_cell_count_1
                                timestamp: 0
                                value (uint64): 63764
                                }
                              subfields {
                                name: switch_cell_count_2
                                timestamp: 0
                                value (uint64): 0
                                }
                              subfields {
                                name: switch_cell_count_3
                                timestamp: 0
                                value (uint64): 0
                                }
                              subfields {
                                name: switch_cell_count_4
                                timestamp: 0
                                value (uint64): 6960
                                }
                              }
                            }
                          }
                        }
                      }
                    }
                  }
                }
              }
            }
        }
    }
}
------------------------------------------------------------------------------

Recevied GBP RPC at: Thu Mar  8 16:38:29 2018
Data size is: 722
```

```
header {
  node_id_str: C14-HAV-2029
  encoding_path: show hardware internal buffer info pkt-stats
  collection_id: 31
  data_gpbkv {
    timestamp: 0
    name:
    fields {
      timestamp: 0
      name: keys
      value (invalid):
        fields {
          timestamp: 0
          name: show hardware internal buffer info pkt-stats
          value: show hardware internal buffer info pkt-stats
        }
    }
    fields {
      timestamp: 0
      name: content
      value (invalid):
        fields {
          timestamp: 0
          name:
          value:
        subfields {
          name: TABLE_module
            fields {
              timestamp: 0
              name:
              value:
            subfields {
              name: ROW_module
                fields {
                  timestamp: 0
                  name:
                  value:
                subfields {
                  name: module_number
                  timestamp: 0
                  value (uint64): 1
                  }
                subfields {
                  name: TABLE_instance
                    fields {
                      timestamp: 0
                      name:
                      value:
                    subfields {
                      name: ROW_instance
                        fields {
                          timestamp: 0
                          name:
                          value:
                        subfields {
                          name: instance
                          timestamp: 0
                          value (uint64): 0
                          }
                        subfields {
                          name: supports_8q
                          timestamp: 0
                          value (uint64): 1
                          }
```

```
subfields {
  name: total_instant_usage_1
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: total_instant_usage_2
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: total_instant_usage_3
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: total_instant_usage_4
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: rem_instant_usage_1
  timestamp: 0
  value (uint64): 63764
  }
subfields {
  name: rem_instant_usage_2
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: rem_instant_usage_3
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: rem_instant_usage_4
  timestamp: 0
  value (uint64): 6960
  }
subfields {
  name: max_cell_usage_1
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: max_cell_usage_2
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: max_cell_usage_3
  timestamp: 0
  value (uint64): 0
  }
subfields {
  name: max_cell_usage_4
  timestamp: 0
  value (uint64): 212
  }
subfields {
  name: switch_cell_count_1
  timestamp: 0
  value (uint64): 63764
```

```
                        }
                      subfields {
                        name: switch_cell_count_2
                        timestamp: 0
                        value (uint64): 0
                        }
                      subfields {
                        name: switch_cell_count_3
                        timestamp: 0
                        value (uint64): 0
                        }
                      subfields {
                        name: switch_cell_count_4
                        timestamp: 0
                        value (uint64): 6960
                        }
                      }
                    }
                  }
                }
              }
            }
          }
        }
    }
}
--------------------------------------------------------------------------------


Recevied GBP RPC at: Thu Mar  8 16:38:35 2018
Data size is: 722
header {
  node_id_str: C14-HAV-2029
  encoding_path: show hardware internal buffer info pkt-stats
  collection_id: 32
  data_gpbkv {
    timestamp: 0
    name:
    fields {
      timestamp: 0
      name: keys
      value (invalid):
        fields {
          timestamp: 0
          name: show hardware internal buffer info pkt-stats
          value: show hardware internal buffer info pkt-stats
        }
    }
    fields {
      timestamp: 0
      name: content
      value (invalid):
        fields {
          timestamp: 0
          name:
          value:
        subfields {
          name: TABLE_module
            fields {
              timestamp: 0
              name:
              value:
```

```
subfields {
  name: ROW_module
    fields {
      timestamp: 0
      name:
      value:
    subfields {
      name: module_number
      timestamp: 0
      value (uint64): 1
      }
    subfields {
      name: TABLE_instance
        fields {
          timestamp: 0
          name:
          value:
        subfields {
          name: ROW_instance
            fields {
              timestamp: 0
              name:
              value:
            subfields {
              name: instance
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: supports_8q
              timestamp: 0
              value (uint64): 1
              }
            subfields {
              name: total_instant_usage_1
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: total_instant_usage_2
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: total_instant_usage_3
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: total_instant_usage_4
              timestamp: 0
              value (uint64): 0
              }
            subfields {
              name: rem_instant_usage_1
              timestamp: 0
              value (uint64): 63764
              }
            subfields {
              name: rem_instant_usage_2
              timestamp: 0
              value (uint64): 0
              }
            subfields {
```

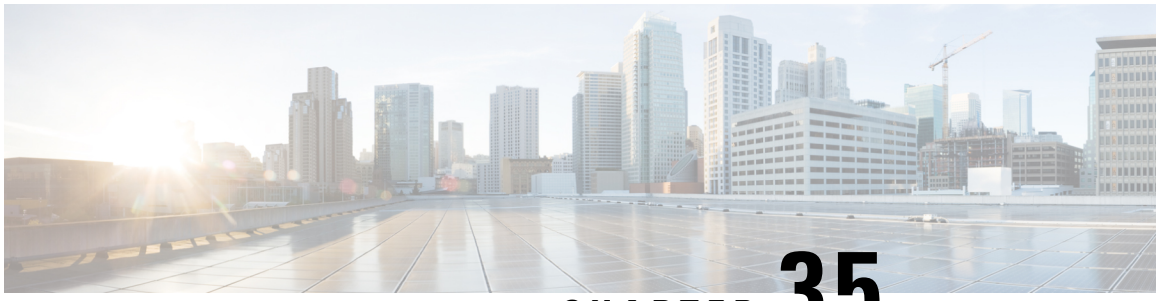
```
          name: rem_instant_usage_3
          timestamp: 0
          value (uint64): 0
          }
        subfields {
          name: rem_instant_usage_4
          timestamp: 0
          value (uint64): 6960
          }
        subfields {
          name: max_cell_usage_1
          timestamp: 0
          value (uint64): 0
          }
        subfields {
          name: max_cell_usage_2
          timestamp: 0
          value (uint64): 0
          }
        subfields {
          name: max_cell_usage_3
          timestamp: 0
          value (uint64): 0
          }
        subfields {
          name: max_cell_usage_4
          timestamp: 0
          value (uint64): 212
          }
        subfields {
          name: switch_cell_count_1
          timestamp: 0
          value (uint64): 63764
          }
        subfields {
          name: switch_cell_count_2
          timestamp: 0
          value (uint64): 0
          }
        subfields {
          name: switch_cell_count_3
          timestamp: 0
          value (uint64): 0
          }
        subfields {
          name: switch_cell_count_4
          timestamp: 0
          value (uint64): 6960
          }
        }
      }
    }
  }
 }
}
}
}
}
}
}
```



PART V

XML Management Interface

- XML Management Interface, on page 439



CHAPTER 35

XML Management Interface

This section contains the following topics:

- About the XML Management Interface, on page 439
- Licensing Requirements for the XML Management Interface, on page 440
- Prerequisites to Using the XML Management Interface, on page 441
- Using the XML Management Interface, on page 441
- Information About Example XML Instances, on page 454
- Additional References, on page 460

About the XML Management Interface

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 the 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, you 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 444 and RFC 4741.

For more information about using NETCONF over SSH, see RFC 4742.

This section includes the following topics:

- NETCONF Layers, on page 439
- SSH xmlagent, on page 440

NETCONF Layers

The following are the NETCONF layers:

Table 25: NETCONF Layers

Layer	Example
Transport protocol	SSHv2
RPC	<rpc>, <rpc-reply>
Operations	<get-config>, <edit-config>
Content	show or configuration command

The following is a description of the four NETCONF layers:

- SSH transport protocol—Provides a secure, 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 the 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 the messages end, which keeps communication in sync.

The XML schemas that define XML configuration instances that you can use are described in the Creating NETCONF XML Instances, on page 444 section.

Licensing Requirements for the XML Management Interface

Product	Product
Cisco NX-OS	The XML management interface 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 *Cisco NX-OS Licensing Guide*.

Prerequisites to Using the XML Management Interface

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. Use the XML management interface with the default settings on the device.

Configuring SSH and the XML Server Options Through the CLI

By default, the SSH server is enabled on the device. If you disable SSH, you must enable it before you start an SSH session on the client PC.

You can configure 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

	Command or Action	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	show **xml server status**	(Optional) Displays information about XML server settings and active XML server sessions. You can find session numbers in the command output.
Step 3	**xml server validate all**	Causes validation of XML documents for the specified server session.
Step 4	**xml server terminate** *session*	Terminates the specified XML server session.
Step 5	**no feature ssh**	(Optional) Disables the SSH server so that you can generate keys.

	Command or Action	Purpose
Step 6	**feature ssh**	Enables the SSH server. The default is enabled.
Step 7	**show ssh server**	(Optional) Displays the status of the SSH server.
Step 8	**xml server max-session** *sessions*	Sets the number of allowed XML server sessions. The default is 8. The range is from 1 to 8.
Step 9	**xml server timeout** *seconds*	Sets the number of seconds after which the XML server session is terminated. The default is 1200 seconds. The range is from 1 to 1200.
Step 10	show **xml server status**	(Optional) Displays information about the XML server settings and active XML server sessions.
Step 11	copy running-config startup-config	(Optional) Saves the running configuration to the startup configuration.

Example

The following example shows how to configure SSH and XML server options through 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 2400
switch(config)# copy running-config startup-config
```

Starting an SSH Session

You can start an SSHv2 session on the client PC with a command 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 from 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 XML server max-sessions option 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 the Hello Message

When you start an SSH session to the XML server, the server responds immediately with a hello message that informs the client of the server's capabilities. You must advertise your capabilities to the server with a hello message before the server processes any other requests. The XML server supports only base capabilities and expects support only for the base capabilities from the client.

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 the 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 the 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 the XSD Files

Procedure

Step 1 From your browser, navigate to the Cisco software download site at the following URL:

http://software.cisco.com/download/navigator.html

The Download Software page opens.

Step 2 In the Select a Product list, choose **Switches > Data Center Switches >** *platform* > *model*.

Step 3 If you are not already logged in as a registered Cisco user, you are prompted to log in now.

Step 4 From the Select a Software Type list, choose **NX-OS XML Schema Definition.**

Step 5 Find the desired release and click **Download.**

Step 6 If you are requested, follow the instructions to apply for eligibility to download strong encryption software images.

The Cisco End User License Agreement opens.

Step 7 Click **Agree** and follow the instructions to download the file to your PC.

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, you can 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 with this method.

Follow these guidelines for this method:

- Verify that the XML server 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 any XML requests. Because the XML server sends the hello response immediately, 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 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 rpc

All NETCONF XML instances must begin with the RPC request tag <rpc>. The example *RPC Request Tag <rpc>* shows the <rpc> element with its required **message-id** attribute. The 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 section on *Obtaining the XSD Files*.

RPC Tag Request

```
<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

The following is an example of a 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>]]>]]>
```

__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 26: NETCONF Operations in Cisco NX-OS

NETCONF Operation	Description	Example
close-session	Closes the current XML server session.	NETCONF Close Session Instance, on page 454
commit	Sets the running configuration to the current contents of the candidate configuration.	NETCONF Commit Instance - Candidate Configuration Capability, on page 459
confirmed-commit	Provides parameters to commit the configuration for a specified time. If this operation is not followed by a commit operation within the confirm-timeout period, the configuration is reverted to the state before the confirmed-commit operation.	NETCONF Confirmed-commit Instance , on page 459
copy-config	Copies the content of source configuration datastore to the target datastore.	NETCONF copy-config Instance, on page 455
delete-config	Operation not supported.	—
edit-config	Configures features in the running configuration of the device. You use this operation for configuration commands.	NETCONF edit-config Instance, on page 455 NETCONF rollback-on-error Instance , on page 459
get	Receives configuration information from the device. You use this operation for **show** commands. The source of the data is the running configuration.	Creating NETCONF XML Instances, on page 444
get-config	Retrieves all or part of a configuration	NETCONF get-config Instance, on page 457

NETCONF Operation	Description	Example
kill-session	Closes the specified XML server session. You cannot close your own session. See the close-session NETCONF operation.	NETCONF Kill-session Instance, on page 455
lock	Allows the client to lock the configuration system of a device.	NETCONF Lock Instance, on page 457
unlock	Releases the configuration lock that the session issued.	NETCONF unlock Instance, on page 458
validate	Checks a candidate configuration for syntactical and semantic errors before applying the configuration to the device.	NETCONF validate Capability Instance , on page 460

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 the Obtaining the XSD Files, on page 443 section.

Using this schema, it is possible to build an XML instance. In the following examples, the relevant portions of the nfcli.xsd schema file that was used to build Creating NETCONF XML Instances, on page 444 is shown.

The following example shows XML device tags.

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>
```

The following example shows the server status device tags.

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>
```

The following example shows the device tag response.

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 “__XML__OPT_Cmd_show_xml___readonly__” is optional. This tag represents the response. For more information on responses, see the RPC Response Tag, on page 453 section.

You can use the | XML option to find the tags you can use to execute a <get>. The following is an example of the | XML option.

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>
]]>]]>
```

From this response, you can see that the namespace defining tag to execute operations on this component is http://www.cisco.com/nxos:1.0:nfcli and the nfcli.xsd file can be used to build requests for this feature.

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.

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

Configuration CLI Commands Sent Through <exec-command>

```
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>]]>]]>
```

The following is the response to the operation:

Response to CLI Commands Sent Through <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" message-id="3">
<nf:ok/>
</nf:rpc-reply>
]]>]]>
```

The following example shows how the show CLI commands that are sent through the <exec-command> can be used to retrieve data.

show CLI Commands Sent Through <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 interface brief</nxos:cmd>
</nxos:exec-command>
</nf:rpc>]]>]]>
```

The following is the response to the operation.

Response to the show CLI commands Sent Through <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>172.23.152.20</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>
]]>]]>
```

The following table provides a detailed explanation of the operation tags:

Table 27: Tags

Tag	Description
<exec-command>	Executes a CLI command.

Tag	Description
<cmd>	Contains the CLI command. A command can be a show or configuration command. Separate multiple configuration commands by using a semicolon ";". 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 in *Configuration CLI Commands Sent Through <exec-command>*.

Replies to configuration commands that are sent through the <cmd> tag are as follows:

- <nf:ok>: All configure commands are executed successfully.
- <nf:rpc-error>: Some commands have failed. The operation stops on the first error, and the <nf:rpc-error> subtree provides more information on what configuration failed. Notice that any configuration that is executed before the failed command would have been applied to the running configuration.

The following example shows a failed configuration:

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 1.1.1.2/24 </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>
]]>]]>
```

Because of a command execution, the interface IP address is set, but the administrative state is not modified (the no shut command is not executed). The reason the administrative state is not modified is because the no port-channel 2000 command results in an error.

The <rpc-reply> results from 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 a configuration and **show** command that are combined in the same instance.

Combination of Configuration 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>
]]>]]>
```

The show command must be sent in its own <exec-command> instance as shown in the following example:

Show CLI Commands Sent Through <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 the client, the XML server sends an XML response enclosed in the RPC response tag <rpc-reply>.

This section contains the following topics:

• RPC Response Tag, on page 453
• Interpreting Tags Encapsulated in the Data Tag, on page 453

RPC Response Tag

The following example shows the RPC response tag <rpc-reply>.

RPC Response Elements

```
<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>]]>]]>
```

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 28: RPC Response Elements

Element	Description
<ok>	The RPC request completed successfully. This element is used when no data is returned in the response.
<data>	The RPC request completed successfully. The data associated with the RPC request is enclosed in the <data> element.
<rpc-error>	The RPC request failed. Error information is enclosed in the <rpc-error> element.

Interpreting Tags Encapsulated in the Data Tag

The device tags encapsulated by the <data> tag contain the request followed by the response. A client application can safely ignore all tags before the <readonly> tag. The following is an 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>
]]>]]>
```

<__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 the examples of the following XML instances:

- NETCONF Close Session Instance, on page 454
- NETCONF Kill-session Instance, on page 455
- NETCONF copy-config Instance, on page 455
- NETCONF edit-config Instance, on page 455
- NETCONF get-config Instance, on page 457
- NETCONF Lock Instance, on page 457
- NETCONF unlock Instance, on page 458
- NETCONF Commit Instance - Candidate Configuration Capability, on page 459
- NETCONF Confirmed-commit Instance , on page 459
- NETCONF rollback-on-error Instance , on page 459
- NETCONF validate Capability Instance , on page 460

NETCONF Close Session Instance

The following example shows 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 example shows 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 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>]]>]]>
```

NETCONF copy-config Instance

The following example shows 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

The following example shows 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. Operation attribute can have the following values:

- create
- merge
- delete

The following example shows how to delete the configuration of interface Ethernet 0/0 from the running configuration.

Edit-config: Delete Operation Request

```
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

```
<rpc-reply message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<ok/>
</rpc-reply>]]>]]>
```

NETCONF get-config Instance

The following example shows 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 the 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 example shows the use of NETCONF lock operation.

The following examples show the 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 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 Unsuccessful Attempt to Acquire the 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 example shows the use of the NETCONF unlock operation.

unlock request

```
<rpc message-id="101"
xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<unlock>
<target>
<running/>
</target>
</unlock>
</rpc>
```

response to 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 example shows the commit operation and the 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 example shows the confirmed-commit operation and the 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 example shows the use of NETCONF rollback on error capability. The string urn:ietf:params:netconf:capability:rollback-on-error:1.0 identifies the capability.

The following example shows 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 example shows the use of the NETCONF validate capability. The string **urn:ietf:params:netconf:capability:validate:1.0** identifies the 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

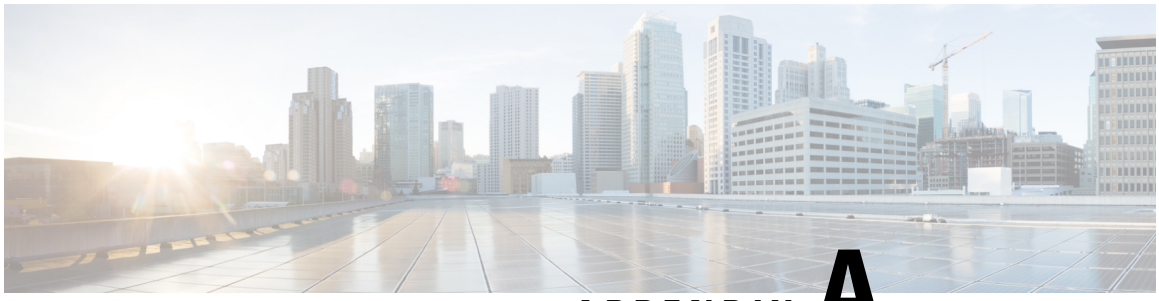
This section provides additional information that is related to implementing the XML management interface.

Standards

Standards	Title
No new or modified standards are supported by this feature. Support for existing standards has not been modified by this feature.	—

RFCs

RFCs	Title
RFC 4741	NETCONF Configuration Protocol
RFC 4742	Using the NETCONF Configuration Protocol over Secure Shell (SSH)



APPENDIX A

Streaming Telemetry Sources

This appendix contains the following topics:

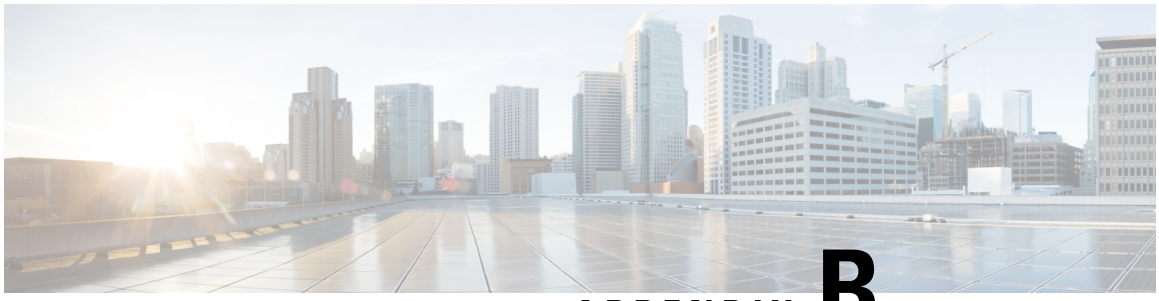
- About Streaming Telemetry, on page 463
- Data Available for Telemetry, on page 463

About Streaming Telemetry

The streaming telemetry feature of the switch continuously streams data out of the network and notifies the client, providing near-real-time access to monitoring data.

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

Packet Timestamping

This chapter contains the following sections:

- About Packet Timestamping, on page 465
- Guidelines and Limitations, on page 467

About Packet Timestamping

Packet timestamping enables precise, scalable traffic monitoring. It helps to detect congestion spots on routers or devices in the network.

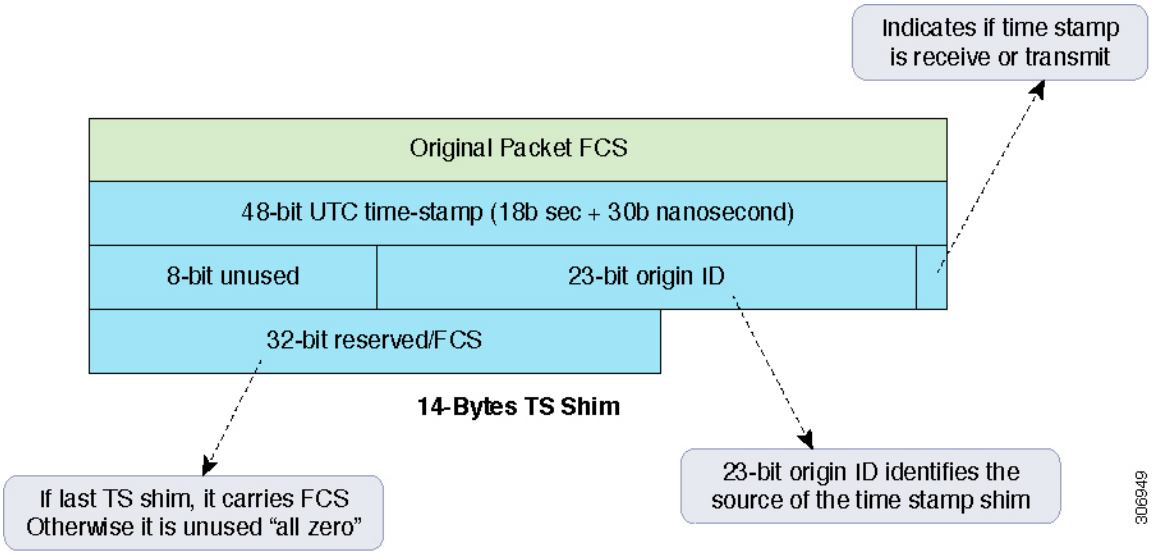
Every participating switch can add one or more timestamp shims, and the decision is based on local configuration.

Timestamping consists of:

- Per-port or Per-flow timestamping
- Insert up to two timestamps at the end of the frame (pre-enqueue and post-dequeue)
- Convey a notion of source identifier that accompany every timestamp record (path topology)

The following figure provides a graphical representation of packet timestamping.

Figure 8: Packet Timestamping



Per-Port Timestamping

An advantage of per-port timestamping is that you can save IFP entries and all packets get timestamped. Each port can be configured to enable timestamping in this way:

- Packets entering ports with timestamping enabled get an ingress timestamp.
- Packets that leave timestamp-enabled ports get an egress timestamp.

To enable per-port timestamping on ingress and egress of port `ethernet1/1` using the CLI in NX-OS:

```
configure terminal
interface ethernet1/1
timestamp ingress id source_id egress id source_id
```

To disable per-port timestamping on port `e1/1` using the CLI in NX-OS:

```
no timestamp
```

Per-Flow Timestamping

Per-flow granularity is achieved in timestamping by defining new action fields for the IFP policy table.

Captured Data

Following is an example of captured date without packet timestamping:

```
0000    00 00 01 00 00 01 00 10   94 00 00 02 08 00 45 00
0010    00 52 d2 ee 00 00 ff fd   66 67 c0 55 01 02 c0 00
```

```
0020    00 01 00 00 00 00 00 00   00 00 00 00 00 00 00 00
0030    00 01 00 00 00 00 00 00   00 00 00 00 00 00 00 00
0040    00 01 00 00 00 00 00 00   00 00 00 00 11 48 f9 8d
0050    8c 9d 07 50 50 c7 10 dc   2c 8d 18 f3 d8 85 e1 94
```

Following is an example of captured date with packet timestamping in ingress and egress enabled:

```
0000    00 00 01 00 00 01 00 10   94 00 00 02 08 00 45 00
0010    00 52 81 ef 00 00 ff fd   b7 66 c0 55 01 02 c0 00
0020    00 01 00 00 00 00 00 00   00 00 00 00 00 00 00 00
0030    00 01 00 00 00 00 00 00   00 00 00 00 00 00 00 00
0040    00 01 00 00 00 00 00 00   00 00 00 00 10 39 27 a7
0050    54 df 93 d4 cd 46 97 80   2c 89 f1 0e 50 33 0c d9
0060    9b e9 fc 50 00 00 b1 cc   e0 6e 00 00 01 54 00 00
0070    00 00 00 00 b1 cc e2 24   00 00 01 77
```

Guidelines and Limitations

Following are the guidelines and limitations for the timestamping feature:

- For information about supported platforms, see the Nexus Switch Platform Matrix.
- Timestamp is not part of the L3 packet. Any checks that assume that the L3+ packet length field represents the total frame length will not be accurate. Systems that need to subject packets to such checks must disable timestamping for the corresponding system or port of flow.
- Header length fields or checksum fields (for example, UDP checksum) will not be updated with the insertion of the timestamp.
- IEEE 802.3 frames (for example, SNAP LLC) are not supported.
- Features that rely on the I2E_CLASSID and HG_CLASSID extended header will not co-exist with packet timestamping.
- No switches across the timestamping path should do pad-stripping or otherwise adjust frame content based on the IP header `payload_len/total_len` field for Ethernet II frames.
- Timestamping is not available for:
  - Mirrored copy
  - SOBMH packets
  - Truncated packets
  - Ingress of HiGig port
  - RCPU