



## **Cisco Nexus 3500 Series NX-OS Programmability Guide, Release 10.2(x)**

**First Published:** 2021-08-24

**Last Modified:** 2021-12-16

### **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: <http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

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

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

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

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

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

© 2021 Cisco Systems, Inc. All rights reserved.



## CONTENTS

---

### PREFACE

<b>Preface</b>	<b>xiii</b>
Audience	<b>xiii</b>
Document Conventions	<b>xiii</b>
Related Documentation for Cisco Nexus 3000 Series Switches	<b>xiv</b>
Documentation Feedback	<b>xiv</b>
Communications, Services, and Additional Information	<b>xiv</b>

---

### CHAPTER 1

<b>New and Changed Information</b>	<b>1</b>
New and Changed Information	<b>1</b>

---

### CHAPTER 2

<b>Bash</b>	<b>3</b>
About Bash	<b>3</b>
Accessing Bash	<b>3</b>
Escalate Privileges to Root	<b>4</b>
Examples of Bash Commands	<b>5</b>
Displaying System Statistics	<b>5</b>
Running Bash from CLI	<b>5</b>
Running Python from Bash	<b>6</b>

---

### CHAPTER 3

<b>Guest Shell</b>	<b>7</b>
About the Guest Shell	<b>7</b>
Accessing the Guest Shell	<b>8</b>
Resources Used for the Guest Shell	<b>8</b>
Capabilities in the Guestshell	<b>9</b>
NX-OS CLI in the Guest Shell	<b>9</b>
Network Access in Guest Shell	<b>10</b>

- Access to Bootflash in Guest Shell 12
- Python in Guest Shell 12
- Installing RPMs in the Guest Shell 13
- Security Posture for 14
  - Kernel Vulnerability Patches 14
  - ASLR and X-Space Support 14
  - Root-User Restrictions 14
  - Resource Management 15
- Guest File System Access Restrictions 15
  - Secure IPC 15
- Managing the Guest Shell 15
  - Disabling the Guest Shell 19
  - Destroying the Guest Shell 19
  - Enabling the Guest Shell 20
- Verifying Virtual Service and Guest Shell Information 21
- Persistently Starting Your Application From the Guest Shell 22
- Procedure for Persistently Starting Your Application from the Guest Shell 23
- An Example Application in the Guest Shell 23

---

**CHAPTER 4**

**Python API 25**

- Information About the Python API 25
- Using Python 25
  - Cisco Python Package 25
  - Using the CLI Command APIs 26
  - Invoking the Python Interpreter from the CLI 28
  - Display Formats 28
  - Non-Interactive Python 30
  - Running Scripts with Embedded Event Manager 31
  - Python Integration with Cisco NX-OS Network Interfaces 32
  - Cisco NX-OS Security with Python 32
    - Examples of Security and User Authority 33
    - Example of Running Script with Scheduler 33

---

**CHAPTER 5**

**Scripting with Tcl 35**

About Tel	35
Telsh Command Help	35
Telsh Command History	36
Telsh Tab Completion	36
Telsh CLI Command	36
Telsh Command Separation	36
Tel Variables	37
Telquit	37
Telsh Security	37
Running the Telsh Command	37
Navigating Cisco NX-OS Modes from the Telsh Command	38
Tcl References	40

---

**CHAPTER 6**

<b>Ansible</b>	41
Prerequisites	41
About Ansible	41
Cisco Ansible Module	41

---

**CHAPTER 7**

<b>Puppet Agent</b>	43
About Puppet	43
Prerequisites	43
Puppet Agent NX-OS Environment	44
ciscopuppet Module	44

---

**CHAPTER 8**

<b>Using Chef Client with Cisco NX-OS</b>	47
About Chef	47
Prerequisites	47
Chef Client NX-OS Environment	48
cisco-cookbook	48

---

**CHAPTER 9**

<b>Using Docker with Cisco NX-OS</b>	51
About Docker with Cisco NX-OS	51
Guidelines and Limitations	51
Prerequisites for Setting Up Docker Containers Within Cisco NX-OS	52

Starting the Docker Daemon	52
Configure Docker to Start Automatically	53
Starting Docker Containers: Host Networking Model	53
Starting Docker Containers: Bridged Networking Model	55
Mounting the bootflash and volatile Partitions in the Docker Container	56
Enabling Docker Daemon Persistence on Enhanced ISSU Switchover	56
Enabling Docker Daemon Persistence on the Cisco Nexus Platform Switches Switchover	57
Resizing the Docker Storage Backend	58
Stopping the Docker Daemon	60
Docker Container Security	60
Securing Docker Containers With User namespace Isolation	61
Moving the cgroup Partition	61
Docker Troubleshooting	62
Docker Fails to Start	62
Docker Fails to Start Due to Insufficient Storage	63
Failure to Pull Images from Docker Hub (509 Certificate Expiration Error Message)	63
Failure to Pull Images from Docker Hub (Client Timeout Error Message)	63
Docker Daemon or Containers Not Running On Switch Reload or Switchover	64
Resizing of Docker Storage Backend Fails	64
Docker Container Doesn't Receive Incoming Traffic On a Port	65
Unable to See Data Port And/Or Management Interfaces in Docker Container	65
General Troubleshooting Tips	65

---

**CHAPTER 10****NX-API 67**

About NX-API	67
Feature NX-API	67
Transport	67
Message Format	68
Security	68
Using NX-API	68
NX-API Management Commands	70
Working With Interactive Commands Using NX-API	72
NX-API Request Elements	72
NX-API Response Elements	72

	About JSON (JavaScript Object Notation)	73
	CLI Execution	73
	XML and JSON Supported Commands	73
	Examples of XML and JSON Output	74
<hr/>		
<b>CHAPTER 11</b>	<b>NX-API Response Codes</b>	<b>77</b>
	Table of NX-API Response Codes	77
<hr/>		
<b>CHAPTER 12</b>	<b>NX-API Developer Sandbox</b>	<b>79</b>
	NX-API Developer Sandbox: NX-OS Releases Prior to 9.2(2)	79
	About the NX-API Developer Sandbox	79
	Guidelines and Limitations	80
	Configuring the Message Format and Command Type	80
	Using the Developer Sandbox	82
	Using the Developer Sandbox to Convert CLI Commands to Payloads	82
<hr/>		
<b>CHAPTER 13</b>	<b>XML Support for ABM and LM in N3500</b>	<b>85</b>
	XML Support for ABM and LM in N3500	85
<hr/>		
<b>CHAPTER 14</b>	<b>Converting CLI Commands to Network Configuration Format</b>	<b>93</b>
	Information About XMLIN	93
	Licensing Requirements for XMLIN	93
	Installing and Using the XMLIN Tool	94
	Converting Show Command Output to XML	94
	Configuration Examples for XMLIN	95
<hr/>		
<b>CHAPTER 15</b>	<b>OpenConfig YANG</b>	<b>99</b>
	About OpenConfig YANG	99
	Guidelines and Limitations for OpenConfig YANG	99
	Understanding Deletion of BGP Routing Instance	104
	Verifying YANG	105
	Enabling OpenConfig Support	106

---

<b>CHAPTER 16</b>	<b>XML Management Interface</b>	<b>107</b>
	About the XML Management Interface	107
	About the XML Management Interface	107
	NETCONF Layers	107
	SSH xmlagent	108
	Licensing Requirements for the XML Management Interface	108
	Prerequisites to Using the XML Management Interface	109
	Using the XML Management Interface	109
	Configuring SSH and the XML Server Options	109
	Starting an SSH Session	109
	Sending the Hello Message	110
	Obtaining the XSD Files	110
	Sending an XML Document to the XML Server	111
	Creating NETCONF XML Instances	111
	RPC Request Tag rpc	112
	NETCONF Operations Tags	113
	Device Tags	114
	Extended NETCONF Operations	116
	NETCONF Replies	119
	RPC Response Tag	120
	Interpreting Tags Encapsulated in the Data Tag	120
	Information About Example XML Instances	121
	Example XML Instances	121
	NETCONF Close Session Instance	121
	NETCONF Kill-session Instance	122
	NETCONF copy-config Instance	122
	NETCONF edit-config Instance	122
	NETCONF get-config Instance	124
	NETCONF Lock Instance	124
	NETCONF unlock Instance	125
	NETCONF Commit Instance - Candidate Configuration Capability	126
	NETCONF Confirmed-commit Instance	126
	NETCONF rollback-on-error Instance	126



	NETCONF validate Capability Instance	127
	Additional References	127
<hr/>		
<b>PART I</b>	<b>Model-Driven Programmability</b>	<b>129</b>
<hr/>		
<b>CHAPTER 17</b>	<b>Managing Components</b>	<b>131</b>
	About the Component RPM Packages	131
	Preparing For Installation	133
	Downloading Components from the Cisco Artifactory	134
	Installing RPM Packages	134
	Installing the Programmable Interface Base And Common Model Component RPM Packages	134
<hr/>		
<b>CHAPTER 18</b>	<b>Converting CLI Commands to Network Configuration Format</b>	<b>137</b>
	Information About XMLIN	137
	Licensing Requirements for XMLIN	137
	Installing and Using the XMLIN Tool	138
	Converting Show Command Output to XML	138
	Configuration Examples for XMLIN	139
<hr/>		
<b>CHAPTER 19</b>	<b>gNMI - gRPC Network Management Interface</b>	<b>143</b>
	About gNMI	143
	gNMI RPC and SUBSCRIBE	144
	Guidelines and Limitations for gNMI	145
	Configuring gNMI	146
	Configuring Server Certificate	147
	Generating Key/Certificate Examples	148
	Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3(2) and Earlier	149
	Examples for Generating and Configuring Key/Certificate for Cisco NX-OS Release 9.3(3) and Later	150
	Verifying gNMI	151
	Clients	152
	Sample DME Subscription - PROTO Encoding	152
	Capabilities	154
	About Capabilities	154
	Guidelines and Limitations for Capabilities	154

Example Client Output for Capabilities	155
Get	158
About Get	158
Guidelines and Limitations for Get	158
Set	159
About Set	159
Guidelines and Limitations for Set	159
Subscribe	160
Guidelines and Limitations for Subscribe	160
gNMI Payload	161
Troubleshooting	163
Gathering TM-Trace Logs	163
Gathering MTX-Internal Logs	164

---

<b>CHAPTER 20</b>	<b>gNOI-gRPC Network Operations Interface</b>	<b>169</b>
	About gNOI	169
	Supported gNOI RPCs	169
	System Proto	170
	OS Proto	171
	Cert Proto	172
	File Proto	172
	Guidelines and Limitations	173
	Verifying gNOI	173

---

<b>CHAPTER 21</b>	<b>Model Driven Telemetry</b>	<b>175</b>
	About Telemetry	175
	Telemetry Components and Process	175
	High Availability of the Telemetry Process	176
	Licensing Requirements for Telemetry	177
	Installing and Upgrading Telemetry	177
	Guidelines and Limitations for Model Driven Telemetry	178
	Configuring Telemetry Using the CLI	181
	Configuring Telemetry Using the NX-OS CLI	181
	Configuring Cadence for YANG Paths	183

Configuration Examples for Telemetry Using the CLI	185
Displaying Telemetry Configuration and Statistics	189
Displaying Telemetry Log and Trace Information	194
Configuring Telemetry Using the NX-API	195
Configuring Telemetry Using the NX-API	195
Configuration Example for Telemetry Using the NX-API	203
Telemetry Model in the DME	206
Telemetry Path Labels	208
About Telemetry Path Labels	208
Polling for Data or Receiving Events	208
Guidelines and Limitations for Path Labels	209
Configuring the Interface Path to Poll for Data or Events	209
Configuring the Interface Path for Non-Zero Counters	211
Configuring the Interface Path for Operational Speeds	212
Configuring the Interface Path with Multiple Queries	214
Configuring the Environment Path to Poll for Data or Events	216
Configuring the Resources Path for Poll for Events or Data	217
Configuring the VXLAN Path to Poll for Events or Data	219
Verifying the Path Label Configuration	220
Displaying Path Label Information	221
Native Data Source Paths	224
About Native Data Source Paths	224
Telemetry Data Streamed for Native Data Source Paths	224
Guidelines and Limitations for Native Data Source Path	226
Configuring the Native Data Source Path for Routing Information	227
Configuring the Native Data Source Path for MAC Information	228
Configuring the Native Data Source Path for All MAC Information	230
Configuring the Native Data Path for IP Adjacencies	232
Displaying Native Data Source Path Information	234
Streaming Syslog	235
About Streaming Syslog for Telemetry	235
Configuring the Native Data Source Path for Routing Information	236
Telemetry Data Streamed for Syslog Path	238
Sample JSON Output	239

Sample KVGPB Output 240  
Additional References 242  
Related Documents 242

---

**APPENDIX A**

**Streaming Telemetry Sources 243**  
About Streaming Telemetry 243  
Data Available for Telemetry 243



## Preface

---

This preface includes the following sections:

- [Audience, on page xiii](#)
- [Document Conventions, on page xiii](#)
- [Related Documentation for Cisco Nexus 3000 Series Switches, on page xiv](#)
- [Documentation Feedback, on page xiv](#)
- [Communications, Services, and Additional Information, on page xiv](#)

## 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
<b>bold</b>	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	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
<i>variable</i>	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
<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
<b>boldface screen font</b>	Information you must enter is in boldface screen font.
<i>italic screen font</i>	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](mailto: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 impact 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 Marketplace](#).
- 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.







# CHAPTER 1

## New and Changed Information

---

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

## New and Changed Information

This table summarizes the new and changed features for the *Cisco Nexus 3500 Series NX-OS Programmability Guide, Release 10.2(x)* and where they are documented.

**Table 1: New and Changed Features**

Feature	Description	Changed in Release	Where Documented
No new feature updates	First 10.2(x) release	10.2(1)F	Not applicable
Bundle OC RPM as part of NX-OS	Added a new configuration command "feature openconfig".	10.2(2)F	<a href="#">Enabling OpenConfig Support, on page 106</a>





## CHAPTER 2

# Bash

---

- [About Bash, on page 3](#)
- [Accessing Bash, on page 3](#)
- [Escalate Privileges to Root, on page 4](#)
- [Examples of Bash Commands, on page 5](#)

## About Bash

In addition to the Cisco NX-OS CLI, Cisco Nexus 3500 platform 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.

## 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.

The following are guidelines for escalating privileges:

- Only an admin user can escalate privileges to root.
- Bash must be enabled before escalating privileges.
- Escalation to root is password protected.
- 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.

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
```

## 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#

```

## Running Python from Bash

The following example shows how to load Python and configure a switch using Python objects:

```

switch# run bash
Linux# python
Python 2.7.5 (default, May 16 2014, 10:58:01)
[GCC 4.3.2] on linux2
Type "help", "copyright", "credits" or "license" for more information.
Loaded cisco NxOS lib!
>>>
>>> from cisco import *
>>> from cisco.vrf import *
>>> from cisco.interface import *
>>> vrfobj=VRF('myvrf')
>>> vrfobj.get_name()
'myvrf'
>>> vrfobj.add_interface('Ethernet1/3')
True
>>> intf=Interface('Ethernet1/3')
>>> print intf.config()

!Command: show running-config interface Ethernet1/3
!Time: Thu Aug 21 23:32:25 2014

version 6.0(2)U4(1)

interface Ethernet1/3
 no switchport
 vrf member myvrf

>>>

```



## CHAPTER 3

# Guest Shell

---

- [About the Guest Shell, on page 7](#)
- [Accessing the Guest Shell, on page 8](#)
- [Resources Used for the Guest Shell, on page 8](#)
- [Capabilities in the Guestshell, on page 9](#)
- [Security Posture for , on page 14](#)
- [Guest File System Access Restrictions , on page 15](#)
- [Managing the Guest Shell, on page 15](#)
- [Verifying Virtual Service and Guest Shell Information, on page 21](#)
- [Persistently Starting Your Application From the Guest Shell, on page 22](#)
- [Procedure for Persistently Starting Your Application from the Guest Shell, on page 23](#)
- [An Example Application in the Guest Shell, on page 23](#)

## 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.

## 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/%             |
| 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 **vrinfo**, are provided as a convenience for running in a different namespace and getting information about which namespace/vrf a process is running in.

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

## NX-OS CLI in the Guest Shell

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

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

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

The NX-OS CLI can be chained together:

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

```
[guestshell@guestshell ~]$
```

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

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




---

**Note** 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 10.28.38.48** 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"
```

## 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.

## 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/](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 roots *size-in-MB*** command is used to increase the size of the file system.

---




---

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

---

## Security Posture for

### 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 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.

### 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 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 follow:

## 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 on the host.

## Guest File System Access Restrictions

### Secure IPC

Applications in a guest shell or virtual service can be made more integrated with the host by using Cisco onePK services. The applications communicate with the host network element over TIPC. Applications within various containers are not allowed to communicate with each other over TIPC, they are only allowed to talk to the host. This prevents issues of one container from spoofing that it is where the Cisco onePK services are running. Applications in containers are also not allowed to listen on TIPC ports.

To ensure that only known virtual services can communicate with the host, a unique identifier for each virtual service is created when it is enabled and verified at the time when the onePK communication channel is established.

The system also limits the rate at which an application in an individual virtual service can send messages to the host. This behavior prevents a misbehaving application from sending messages frequently enough to prevent normal operation of the host or to block other virtual services on the same host from communicating with the host.

## Managing the Guest Shell

The following are commands to manage the Guest Shell:

*Table 2: Guest Shell CLI Commands*

Commands	Description

Commands	Description
<b>guestshell enable</b> { <b>package</b> [ <i>guest shell OVA file</i>   <i>rootfs-file-URI</i> ]}	<ul style="list-style-type: none"> <li>• When <i>guest shell OVA file</i> is specified:             <p>Installs and activates the Guest Shell using the OVA that is embedded in the system image.</p> <p>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.</p> <p>When the Guest Shell is already installed, this command enables the installed Guest Shell. Typically this is used after a <b>guestshell disable</b> command.</p> </li> <li>• When <i>rootfs-file-URI</i> is specified:             <p>Imports a Guest Shell <b>rootfs</b> when the Guest Shell is in a destroyed state. This command brings up the Guest Shell with the specified package.</p> </li> </ul>
<b>guestshell export rootfs package</b> <i>destination-file-URI</i>	Exports a Guest Shell <b>rootfs</b> file to a local URI (bootflash, USB1, etc.).
<b>guestshell disable</b>	Shuts down and disables the Guest Shell.



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

Commands	Description
<b>guestshell destroy</b>	<p>Deactivates and uninstalls the Guest Shell. All resources associated with the Guest Shell are returned to the system. The <b>show virtual-service global</b> command indicates when these resources become available.</p> <p>Issuing this command results in a prompt for a confirmation prior to carrying out the destroy command.</p>
<b>guestshell</b> <b>run guestshell</b>	Connects to the Guest Shell that is already running with a shell prompt. No username/password is required.
<b>guestshell run</b> <i>command</i> <b>run guestshell</b> <i>command</i>	<p>Executes a Linux/UNIX command within the context of the Guest Shell environment.</p> <p>After execution of the command you are returned to the switch prompt.</p>
<b>guestshell resize</b> [cpu   memory   rootfs]	<p>Changes the allotted resources available for the Guest Shell. The changes take effect the next time the Guest Shell is enabled or rebooted.</p> <p><b>Note</b>      Resize values are cleared when the <b>guestshell destroy</b> command is used.</p>
<b>guestshell sync</b>	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.
<b>virtual-service reset force</b>	<p>In the event that the guestshell or virtual-services cannot be managed, even after a system reload, the reset command is used to force the removal of the Guest Shell and all virtual-services. The system needs to be reloaded for the cleanup to happen. No Guest Shell or additional virtual-services can be installed or enabled after issuing this command until after the system has been reloaded.</p> <p>You are prompted for a confirmation prior to initiating the reset.</p>




---

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

---




---

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

---

## Disabling the Guest Shell

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

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

Example:

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

2014 Jul 30 19:47:23 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Deactivating virtual
service 'guestshell+'
2014 Jul 30 18:47:29 switch %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE: Successfully deactivated
virtual service 'guestshell+'
switch# show virtual-service list
Virtual Service List:
Name                               Status           Package Name
-----
guestshell+                        Deactivated      guestshell.ova
```




---

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

---

## Destroying the Guest Shell

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

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

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

switch# guestshell destroy
```

```

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

switch# show virtual-service list
Virtual Service List:

```




---

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

---




---

**Note** In the Cisco NX-OS software, the **oneP** feature is automatically enabled for local access when a container is installed. Since the Guest Shell is a container, the **oneP** feature is automatically started.

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

```

## Verifying Virtual Service and Guest Shell Information

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

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

Command	Description
<pre> <b>show guestshell detail</b>  switch# <b>show guestshell detail</b> 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 </pre>	<p>Displays details about the guestshell package (such as version, signing resources, and devices).</p>

## Persistently Starting Your Application From the Guest Shell

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

```

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

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

[Install]
WantedBy=multi-user.target

```



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

# Procedure for Persistently Starting Your Application from the Guest Shell

- 
- 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]#
```





## CHAPTER 4

# Python API

---

- [Information About the Python API, on page 25](#)
- [Using Python, on page 25](#)

## Information About the Python API

The Cisco Nexus 3500 platform switches support Python v2.7.11 in both interactive and noninteractive (script) modes and are available in the Guest Shell.

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

This section describes how to write and execute Python scripts.

## Cisco Python Package

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

you can run the help command for a specific module. For example, `help(cisco.interface)` displays the properties of the `cisco.interface` module.

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
```

## 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 3: CLI Command APIs

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

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

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

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

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

In the following example, the second command is successful because the state persists for the second and third commands:

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



**Note** Commands are separated with ";" as shown in the example. The semicolon (;) must be surrounded with single blank characters.

## Invoking the Python Interpreter from the CLI

The following example shows how to invoke Python 2 from the CLI:



**Note** The Python interpreter is designated with the ">>>" or "... " prompt.

```
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
>>>
```

## 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:

```
>>> from cli import *
>>> cli("conf ; interface loopback 1")
''
>>> r = cli('where detail') ; print r
mode:
username:          admin
vdc:               EOR-1
routing-context vrf: default
>>>
```

### Example 4:

```
>>> from cli import *
>>> import json
>>> out=json.loads(clid('show version'))
>>> for k in out.keys():
...     print "%30s = %s" % (k, out[k])
...
kern_uptm_secs = 21
kick_file_name = bootflash:///nxos.9.2.1.bin.S246
rr_service = None
module_id = Supervisor Module
kick_tmstamp = 07/11/2018 00:01:44
bios_cmpl_time = 05/17/2018
bootflash_size = 20971520
kickstart_ver_str = 9.2(1)
kick_cmpl_time = 7/9/2018 9:00:00
chassis_id = Nexus9000 C9504 (4 Slot) Chassis
proc_board_id = SAL171211LX
memory = 16077872
manufacturer = Cisco Systems, Inc.
kern_uptm_mins = 26
bios_ver_str = 05.31
cpu_name = Intel(R) Xeon(R) CPU D-1528 @ 1.90GHz
kern_uptm_hrs = 2
rr_usecs = 816550
rr_sys_ver = 9.2(1)
rr_reason = Reset Requested by CLI command reload
rr_ctime = Wed Jul 11 20:44:39 2018
header_str = Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2018, Cisco and/or its affiliates.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their own
licenses, such as open source. This software is provided "as is," and unless
otherwise stated, there is no warranty, express or implied, including but not
limited to warranties of merchantability and fitness for a particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0 or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
http://www.gnu.org/licenses/old-licenses/library.txt.
host_name = switch
```

```

        mem_type = kB
        kern_uptm_days = 0
    >>>

```

## 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 3500 platform switches also support 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:deltaCounters.py
#!/isan/bin/python

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 '=====
      %8d %8d %8d %8d %8d %8d' % (rxuc, rxmc, rxbc, txuc, txmc, txbc)
print '=====

i = 0
while (i < count):
    time.sleep(delay)
    out = json.loads(clid(cmd))
    rxucNew = int(out['TABLE_rx_counters']['ROW_rx_counters'][0]['eth_inucast'])
    rxmcNew = int(out['TABLE_rx_counters']['ROW_rx_counters'][1]['eth_inmcast'])
    rxbcNew = int(out['TABLE_rx_counters']['ROW_rx_counters'][1]['eth_inbcast'])
    txucNew = int(out['TABLE_tx_counters']['ROW_tx_counters'][0]['eth_outucast'])
    txmcNew = int(out['TABLE_tx_counters']['ROW_tx_counters'][1]['eth_outmcast'])
    txbcNew = int(out['TABLE_tx_counters']['ROW_tx_counters'][1]['eth_outbcast'])
    i += 1
    print '%-3d %8d %8d %8d %8d %8d' % \
        (i, rxucNew - rxuc, rxmcNew - rxmc, rxbcNew - rxbc, txucNew - txuc, txmcNew - txmc,
        txbcNew - txbc)

switch# python bootflash:deltaCounters.py Ethernet1/1 1 5
row rx_ucast rx_mcast rx_bcast tx_ucast tx_mcast tx_bcast
=====
          0      791         1         0    212739         0
=====
1          0         0         0         0         26         0
2          0         0         0         0         27         0

```

```

3          0          1          0          0          54          0
4          0          1          0          0          55          0
5          0          1          0          0          81          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 3500 platform 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 3500 platform 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
Python 2.7.5 (default, Oct 8 2013, 23:59:43)
[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

- 

## Example of Running Script with Scheduler

-





## CHAPTER 5

# Scripting with Tcl

---

- [About Tcl, on page 35](#)
- [Running the Tclsh Command, on page 37](#)
- [Navigating Cisco NX-OS Modes from the Tclsh Command, on page 38](#)
- [Tcl References, on page 40](#)

## 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.

## Tclsh Command Help

Command help is not available for Tcl commands. You can still access the help functions of Cisco NX-OS commands from within an interactive Tcl shell.

This example shows the lack of Tcl command help in an interactive Tcl shell:

```
switch# tclsh
switch-tcl# set x 1
switch-tcl# puts ?
           ^
% Invalid command at '^' marker.
switch-tcl# configure ?
<CR>
  session  Configure the system in a session
  terminal  Configure the system from terminal input

switch-tcl#
```



---

**Note** In the preceding example, the Cisco NX-OS command help function is still available but the Tcl **puts** command returns an error from the help function.

---

## Tclsh Command History

You can use the arrow keys on your terminal to access commands you previously entered in the interactive Tcl shell.




---

**Note** The **tclsh** command history is not saved when you exit the interactive Tcl shell.

---

## Tclsh Tab Completion

You can use tab completion for Cisco NX-OS commands when you are running an interactive Tcl shell. Tab completion is not available for Tcl commands.

## Tclsh CLI Command

Although you can directly access Cisco NX-OS commands from within an interactive Tcl shell, you can only execute Cisco NX-OS commands in a Tcl script if they are prepended with the Tcl **cli** command.

In an interactive Tcl shell, the following commands are identical and execute properly:

```
switch-tcl# cli show module 1 | incl Mod
switch-tcl# cli "show module 1 | incl Mod"
switch-tcl# show module 1 | incl Mod
```

In a Tcl script, you must prepend Cisco NX-OS commands with the Tcl **cli** command as shown in the following example:

```
set x 1
cli show module $x | incl Mod
cli "show module $x | incl Mod"
```

If you use the following commands in your script, the script fails and the Tcl shell displays an error:

```
show module $x | incl Mod
"show module $x | incl Mod"
```

## Tclsh Command Separation

The semicolon (;) is the command separator in both Cisco NX-OS and Tcl. To execute multiple Cisco NX-OS commands in a Tcl command, you must enclose the Cisco NX-OS commands in quotes ("").

In an interactive Tcl shell, the following commands are identical and execute properly:

```
switch-tcl# cli "configure terminal ; interface loopback 10 ; description loop10"
switch-tcl# cli configure terminal ; cli interface loopback 10 ; cli description loop10
switch-tcl# cli configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

switch(config-tcl)# cli interface loopback 10
switch(config-if-tcl)# cli description loop10
switch(config-if-tcl)#
```

In an interactive Tcl shell, you can also execute Cisco NX-OS commands directly without prepending the Tcl **cli** command:

```
switch-tcl# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.

switch(config-tcl)# interface loopback 10
switch(config-if-tcl)# description loop10
switch(config-if-tcl)#
```

## Tcl Variables

You can use Tcl variables as arguments to the Cisco NX-OS commands. You can also pass arguments into Tcl scripts. Tcl variables are not persistent.

The following example shows how to use a Tcl variable as an argument to a Cisco NX-OS command:

```
switch# tclsh
switch-tcl# set x loop10
switch-tcl# cli "configure terminal ; interface loopback 10 ; description $x"
switch(config-if-tcl)#
```

## Tclquit

The **tclquit** command exits the Tcl shell regardless of which Cisco NX-OS command mode is currently active. You can also press **Ctrl-C** to exit the Tcl shell. The **exit** and **end** commands change Cisco NX-OS command modes. The **exit** command terminates the Tcl shell only from the EXEC command mode.

## Tclsh Security

The Tcl shell is executed in a sandbox to prevent unauthorized access to certain parts of the Cisco NX-OS system. The system monitors CPU, memory, and file system resources being used by the Tcl shell to detect events such as infinite loops, excessive memory utilization, and so on.

You configure the initial Tcl environment with the **scripting tcl init** *init-file* command.

You can define the looping limits for the Tcl environment with the **scripting tcl recursion-limit** *iterations* command. The default recursion limit is 1000 iterations.

## Running the Tclsh Command

You can run Tcl commands from either a script or on the command line using the **tclsh** command.



---

**Note** You cannot create a Tcl script file at the CLI prompt. You can create the script file on a remote device and copy it to the bootflash: directory on the Cisco NX-OS device.

---

### SUMMARY STEPS

1. **tclsh** [**bootflash:***filename* [*argument* ... ]]

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>tclsh</b> [ <b>bootflash:filename</b> [ <i>argument . . .</i> ]]  <b>Example:</b> <pre>switch# tclsh ? &lt;CR&gt; bootflash: The file to run</pre>	<p>Starts a Tcl shell.</p> <p>If you run the <b>tclsh</b> 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 <b>tclquit</b> or <b>Ctrl-C</b>.</p> <p>If you run the <b>tclsh</b> 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.</p>

**Example**

The following example shows an interactive Tcl shell:

```
switch# tclsh
switch-tcl# set x 1
switch-tcl# cli show module $x | incl Mod
Mod Ports Module-Type Model Status
1 36 36p 40G Ethernet Module N9k-X9636PQ ok
Mod Sw Hw
Mod MAC-Address(es) Serial-Num

switch-tcl# exit
switch#
```

The following example shows how to run a Tcl script:

```
switch# show file bootflash:showmodule.tcl
set x 1
while {$x < 19} {
cli show module $x | incl Mod
set x [expr {$x + 1}]
}

switch# tclsh bootflash:showmodule.tcl
Mod Ports Module-Type Model Status
1 36 36p 40G Ethernet Module N9k-X9636PQ ok
Mod Sw Hw
Mod MAC-Address(es) Serial-Num

switch#
```

## Navigating Cisco NX-OS Modes from the Tclsh Command

You can change modes in Cisco NX-OS while you are running an interactive Tcl shell.

## SUMMARY STEPS

1. `tclsh`
2. `configure terminal`
3. `tclquit`

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>tclsh</b> <b>Example:</b> <pre>switch# tclsh switch-tcl#</pre>	Starts an interactive Tcl shell.
Step 2	<b>configure terminal</b> <b>Example:</b> <pre>switch-tcl# configure terminal switch(config-tcl)#</pre>	Runs a Cisco NX-OS command in the Tcl shell, changing modes.  <b>Note</b> The Tcl prompt changes to indicate the Cisco NX-OS command mode.
Step 3	<b>tclquit</b> <b>Example:</b> <pre>switch-tcl# tclquit switch#</pre>	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 cli 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 6

# Ansible

- [Prerequisites](#), on page 41
- [About Ansible](#), on page 41
- [Cisco Ansible Module](#), on page 41

## Prerequisites

Go to [https://docs.ansible.com/ansible/latest/getting\\_started/index.html](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	<a href="https://www.ansible.com/">https://www.ansible.com/</a>
Ansible Automation Solutions. Includes installation instructions, playbook instructions and examples, module lists, and so on.	<a href="https://docs.ansible.com/">https://docs.ansible.com/</a>

## 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.	<a href="#">Configuration Management Tools</a>
-------------------------------	------------------------------------------------

Ansible NX-OS playbook examples	<a href="#">Repo for ansible nxos playbooks</a>
Ansible NX-OS network modules	<a href="#">nxos network modules</a>



## CHAPTER 7

# Puppet Agent

This chapter includes the following sections:

- [About Puppet, on page 43](#)
- [Prerequisites, on page 43](#)
- [Puppet Agent NX-OS Environment, on page 44](#)
- [ciscopuppet Module, on page 44](#)

## 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	<a href="https://puppetlabs.com">https://puppetlabs.com</a>
Puppet Labs FAQ	<a href="https://puppet.com/blog/how-get-started-puppet-enterprise-faq/">https://puppet.com/blog/how-get-started-puppet-enterprise-faq/</a>
Puppet Labs Documentation	<a href="https://puppet.com/docs">https://puppet.com/docs</a>

## Prerequisites

The following are prerequisites for the Puppet Agent:

- You must have a switch and operating system software release that supports the installation.
  - Cisco Nexus 3600 platform switches.
  - Cisco Nexus 3100 platform switches.
  - Cisco Nexus 3000 Series switches.
  - Cisco NX-OS Release 7.0(3)I2(1) or later.
- 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)	<a href="https://github.com/cisco/cisco-network-puppet-module/blob/develop/docs/README-agent-install.md">https://github.com/cisco/cisco-network-puppet-module/blob/develop/docs/README-agent-install.md</a>
---------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## 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)	<a href="#">Puppet Forge</a>
Resource Type Catalog	<a href="#">Cisco Puppet Resource Reference</a>
ciscopuppet Module: Source Code Repository	<a href="#">Cisco Network Puppet Module</a>

ciscopuppet Module: Setup & Usage	<a href="#">Cisco Puppet Module::README.md</a>
Puppet Labs: Installing Modules	<a href="https://puppet.com/docs/puppet/7/modules_installing.html">https://puppet.com/docs/puppet/7/modules_installing.html</a>
Puppet NX-OS Manifest Examples	<a href="#">Cisco Network Puppet Module Examples</a>
NX-OS developer landing page.	<a href="#">Configuration Management Tools</a>





## CHAPTER 8

# Using Chef Client with Cisco NX-OS

This chapter includes the following sections:

- [About Chef, on page 47](#)
- [Prerequisites, on page 47](#)
- [Chef Client NX-OS Environment, on page 48](#)
- [cisco-cookbook, on page 48](#)

## About Chef

Chef is an open-source software package developed by Chef Software, Inc. It 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 is comprised of one or more workstations, a single server, and every node that will be configured and maintained by the chef-client. 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 can also be used to install software packages, copy files, start services, and more.

The following references provide more information from Chef:

Topic	Link
Chef home	<a href="https://www.chef.io">https://www.chef.io</a>
Chef overview	<a href="https://docs.chef.io/chef_overview.html">https://docs.chef.io/chef_overview.html</a>
Chef documentation (all)	<a href="https://docs.chef.io/">https://docs.chef.io/</a>

## Prerequisites

The following are prerequisites for Chef:

- You must have a Cisco switch and operating system software release that supports the installation:
  - Cisco NX-OS Release 6.1(2)I3(4) or higher
- You must have the required disk storage available on the device for Chef deployment:
  - A minimum of 500 MB 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 Cisco Nexus switches. Customers can install chef-client in one of the Linux environments provided by the Cisco Nexus switch:

- Bash Shell — This is the native WindRiver Linux environment underlying Cisco NX-OS.
- Guest Shell — This is a secure Linux container environment running CentOS. Its advantage is a secure, open execution environment that is decoupled from the host.

The workflow for both use cases is similar.

The following documents provide step-by-step guidance on agent software download, installation, and setup:

Topic	Link
Chef Client (Native)	Latest information on Client RPM is available <a href="#">here</a> .
Chef Client (Guest Shell, CentOS7)	Latest information on Client RPM is available <a href="#">here</a> .
Chef Client: Installation and setup on Cisco Nexus platform (manual setup)	<a href="#">cisco-cookbook::README-install-agent.md</a>
Chef Client: Installation and setup on Cisco Nexus platform (automated installation using the Chef provisioner)	<a href="#">cisco-cookbook::README-chef-provisioning.md</a>

## 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 Cisco NX-OS operating system and Cisco Nexus switches. This cookbook is installed on the Chef Server and is required for proper Chef Client operation on Cisco Nexus switches.

cisco-cookbook can be found on Chef Supermarket.

The following documents provide additional detail for cisco-cookbook and generic cookbook installation procedures:



Topic	Link
cisco-cookbook location	<a href="https://supermarket.chef.io/cookbooks/cisco-cookbook">https://supermarket.chef.io/cookbooks/cisco-cookbook</a>
Resource Type Catalog	<a href="https://github.com/cisco/cisco-network-chef-cookbook#resource-by-tech">https://github.com/cisco/cisco-network-chef-cookbook#resource-by-tech</a>
cisco-cookbook: Source Code Repository	<a href="https://github.com/cisco/cisco-network-chef-cookbook">https://github.com/cisco/cisco-network-chef-cookbook</a>
cisco-cookbook: Setup and usage	<a href="#">cisco-cookbook::README.md</a>
Chef Supermarket	<a href="https://supermarket.chef.io">https://supermarket.chef.io</a>
NX-OS developer landing page.	<a href="#">Configuration Management Tools</a>





## CHAPTER 9

# Using Docker with Cisco NX-OS

---

This chapter contains the following topics:

- [About Docker with Cisco NX-OS, on page 51](#)
- [Guidelines and Limitations, on page 51](#)
- [Prerequisites for Setting Up Docker Containers Within Cisco NX-OS, on page 52](#)
- [Starting the Docker Daemon, on page 52](#)
- [Configure Docker to Start Automatically, on page 53](#)
- [Starting Docker Containers: Host Networking Model, on page 53](#)
- [Starting Docker Containers: Bridged Networking Model, on page 55](#)
- [Mounting the bootflash and volatile Partitions in the Docker Container, on page 56](#)
- [Enabling Docker Daemon Persistence on Enhanced ISSU Switchover, on page 56](#)
- [Enabling Docker Daemon Persistence on the Cisco Nexus Platform Switches Switchover, on page 57](#)
- [Resizing the Docker Storage Backend, on page 58](#)
- [Stopping the Docker Daemon, on page 60](#)
- [Docker Container Security, on page 60](#)
- [Docker Troubleshooting, on page 62](#)

## 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.

## Guidelines and Limitations

Following are the guidelines and limitations for using Docker on Cisco NX-OS on a switch:

- Docker functionality is supported on the switches with at least 8 GB of system RAM.

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

Following are the prerequisites for using Docker on Cisco NX-OS on a switch:

- Enable the host Bash shell. To use Docker on Cisco NX-OS on a switch, you must be the root user on the host Bash shell:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# feature bash-shell
```

- If the switch is in a network that uses an HTTP proxy server, the `http_proxy` and `https_proxy` environment variables must be set up in `/etc/sysconfig/docker`. For example:

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

- Verify that the switch clock is set correctly, or you might see the following error message:

```
x509: certificate has expired or is not yet valid
```

- Verify that the domain name and name servers are configured appropriately for the network and that it is reflected in the `/etc/resolv.conf` file:

```
switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# vrf context management
switch(config-vrf)# ip domain-name ?
WORD Enter the default domain (Max Size 64)

switch(config-vrf)# ip name-server ?
A.B.C.D Enter an IPv4 address
A:B::C:D Enter an IPv6 address

root@switch# cat /etc/resolv.conf
domain cisco.com #bleed
nameserver 171.70.168.183 #bleed
root@switch#
```

## Starting the Docker Daemon

When you start the Docker daemon for the first time, a fixed-size backend storage space is carved out in a file called `dockerpart` on the bootflash, which is then mounted to `/var/lib/docker`. If necessary, you can adjust the default size of this space by editing `/etc/sysconfig/docker` before you start the Docker daemon for the first time. You can also resize this storage space if necessary as described later on.

To start the Docker daemon:

---

**Step 1** Load Bash and become superuser.

```
switch# run bash sudo su -
```

**Step 2** Start the Docker daemon.

```
root@switch# service docker start
```

**Step 3** Check the status.

```
root@switch# service docker status
dockerd (pid 3597) is running...
root@switch#
```

**Note** Once you start the Docker daemon, do not delete or tamper with the `dockerpart` file on the bootflash since it is critical to the docker functionality.

```
switch# dir bootflash:dockerpart
20000000000 Mar 14 12:50:14 2018 dockerpart
```

---

## Configure Docker to Start Automatically

You can configure the Docker daemon to always start up automatically when the switch boots up.

---

**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>`.

**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.

**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 61](#) 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.

**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.

**Step 1** Load Bash and become superuser.

```
switch# run bash sudo su -
```

**Step 2** Before starting the switchover, use the `chkconfig` utility to make the Docker service persistent.



```
root@switch# chkconfig --add docker
root@n9k-2#
```

**Step 3** Start any containers using the `--restart unless-stopped` option so that they will be restarted automatically after the switchover.

The following example starts an Alpine container and configures it to always restart unless it is explicitly stopped or Docker is restarted:

```
root@switch# docker run -dit --restart unless-stopped alpine
root@n9k-2#
```

The Docker containers are disrupted (restarted) during the switchover, so they will not be running continuously.

---

## Enabling Docker Daemon Persistence on the Cisco Nexus Platform Switches Switchover

You can have both the Docker daemon and any running containers persist on a switchover between two separate physical supervisors with distinct bootflash partitions. However, for the Cisco Nexus switches, the bootflash partitions on both supervisors are physically separate. You will therefore need to copy the `dockerpart` file manually to the standby supervisor before performing the switchover.

---

**Step 1** Load Bash and become superuser.

```
switch# run bash sudo su -
```

**Step 2** Start any containers using the `--restart unless-stopped` option so that they will be restarted automatically after the switchover.

The following example starts an Alpine container and configures it to always restart unless it is explicitly stopped or Docker is restarted:

```
root@switch# docker run -dit --restart unless-stopped alpine
root@n9k-2#
```

Note that the Docker containers will be disrupted (restarted) during the switchover, so they will not be running continuously.

**Step 3** Before starting the switchover, use the `chkconfig` utility to make the Docker service persistent.

```
root@switch# chkconfig --add docker
root@n9k-2#
```

**Step 4** Copy the Docker backend storage partition from the active to the standby supervisor bootflash:

```
root@switch# service docker stop
Stopping dockerd: dockerd shutdown

root@switch# cp /bootflash/dockerpart /bootflash_sup-remote/
```

```
root@switch# service docker start
```

---

## Resizing the Docker Storage Backend

After starting or using the Docker daemon, you can grow the size of the Docker backend storage space according to your needs.

---

### 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.

---

**Step 1** Load Bash and become superuser.

```
switch# run bash sudo su -
```

**Step 2** Stop the Docker daemon.

```
root@switch# service docker stop
Stopping dockerd: dockerd shutdown
```

**Step 3** Verify that the Docker daemon is stopped.

```
root@switch# service docker status
dockerd is stopped
root@switch#
```

**Note** You can also delete the `dockerpart` file on the bootflash at this point, if necessary:

```
switch# delete bootflash:dockerpart
Do you want to delete "/dockerpart" ? (yes/no/abort) y
switch#
```

---

## Docker Container Security

Following are the Docker container security recommendations:

- Run in a separate user `namespace` if possible.
- Run in a separate network `namespace` if possible.
- Use `cgroups` to limit resources. An existing `cgroup` (`ext_ser`) is created to limit hosted applications to what the platform team has deemed reasonable for extra software running on the switch. Docker allows use of this and limiting per-container resources.
- Do not add unnecessary POSIX capabilities.

## Securing Docker Containers With User namespace Isolation

For containers using the bridged networking model, you can also set up user namespace isolation to further improve security. See <https://docs.docker.com/engine/security/usersns-remap/> for more information.

**Step 1** Determine if a `dockremap` group already exists on your system.

A `dockremap` user must already be set up on your system by default. If the `dockremap` group doesn't already exist, follow these steps to create it.

a) Enter the following command to create the `dockremap` group:

```
root@switch# groupadd dockremap -r
```

b) Create the `dockremap` user, unless it already exists:

```
root@switch# useradd dockremap -r -g dockremap
```

c) Verify that the `dockremap` group and the `dockremap` user were created successfully:

```
root@switch# id dockremap
uid=999(dockremap) gid=498(dockremap) groups=498(dockremap)
root@switch#
```

**Step 2** Add the desired re-mapped ID and range to the `/etc/subuid` and `/etc/subgid`.

For example:

```
root@switch# echo "dockremap:123000:65536" >> /etc/subuid
root@switch# echo "dockremap:123000:65536" >> /etc/subgid
```

**Step 3** Using a text editor, add the `--usersns-remap=default` option to the `other_args` field in the `/etc/sysconfig/docker` file.

For example:

```
other_args="--debug=true --usersns-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/usersns-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.

**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=alpinerrun -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.





# CHAPTER 10

## NX-API

---

- [About NX-API, on page 67](#)
- [Using NX-API, on page 68](#)
- [XML and JSON Supported Commands, on page 73](#)

## About NX-API

### Feature NX-API

- Feature NX-API is required to be enabled for access the device through sandbox.
- | json on the device internally uses python script to generate output.
- NX-API can be enabled either on http/https via ipv4:

```
BLR-VXLAN-NPT-CR-179# show nxapi
nxapi enabled
HTTP Listen on port 80
HTTPS Listen on port 443
BLR-VXLAN-NPT-CR-179#
```
- NX-API is internally spawning third-party NGINX process, which handler receive/send/processing of http requests/response:

```
nxapi certificate {httpsCRT |httpskey}
nxapi certificate enable
```
- NX-API Certificates can be enabled for https
- Default port for nginx to operate is 80/443 for http/https respectively. It can also be changed using the following CLI command:

```
nxapi {http|https} port port-number
```

## 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.

## 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 or JSON-RPC.

## 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.

## Using NX-API

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 Sandbox:

- Enable the management interface.

```
switch# conf t
switch(config)# interface mgmt 0
switch(config)# ip address 198.51.100.1/24
switch(config)# vrf context management
switch(config)# ip route 203.0.113.1/0 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"
      }
    }
  }
}

```

### Using the Management Interface for NX-API calls

It is recommended to use the management interface for NX-API calls.

When using non-management interface and a custom port for NX-API an entry should be made in the CoPP policy to prevent NX-API traffic from hitting the default copp entry which could unfavorably treat API traffic.



**Note** It is recommended to use the management interface for NX-API traffic. If that is not possible and a custom port is used, the "copp-http" class should be updated to include the custom NX-API port.

The following example port 9443 is being used for NX-API traffic.

This port is added to the `copp-system-acl-http` ACL to allow it to be matched under the `copp-http` class resulting on 100 pps policing. (This may need to be increased in certain environments.)

```
!
ip access-list copp-system-acl-http
 10 permit tcp any any eq www
 20 permit tcp any any eq 443
 30 permit tcp any any eq 9443 <-----
!
class-map type control-plane match-any copp-http
 match access-group name copp-system-acl-http
!
!
policy-map type control-plane copp-system-policy
 class copp-http
  police pps 100
!
```

## NX-API Management Commands

You can enable and manage NX-API with the CLI commands listed in the following table.

**Table 4: NX-API Management Commands**

NX-API Management Command	Description
<b>feature nxapi</b>	Enables NX-API.
<b>no feature nxapi</b>	Disables NX-API.
<b>nxapi {http   https} port <i>port</i></b>	Specifies a port.
<b>no nxapi {http   https}</b>	Disables HTTP/HTTPS.
<b>show nxapi</b>	Displays port information.

NX-API Management Command	Description
<b>nxapi certificate {httpsrct certfile   httpskey keyfile} filename</b>	Specifies the upload of the following: <ul style="list-style-type: none"> <li>• HTTPS certificate when httpsrct is specified.</li> <li>• HTTPS key when httpskey is specified.</li> </ul> Example of HTTPS certificate: <b>nxapi certificate httpsrct certfile bootflash:cert.crt</b> Example of HTTPS key: <b>nxapi certificate httpskey keyfile bootflash:privkey.key</b>
<b>nxapi certificate enable</b>	Enables a certificate.

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

```
switch(config)# nxapi certificate httpsrct certfile certificate.crt
Upload done. Please enable. Note cert and key must match.
switch(config)# nxapi certificate enable
switch(config)#
```

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)#
```

In some situations, you might get an error message saying that the certificate is invalid:

```
switch(config)# nxapi certificate httpskey keyfile bootflash:privkey.key
Upload done. Please enable. Note cert and key must match.
switch(config)# nxapi certificate enable
Nginx certificate invalid.
switch(config)#
```

This might occur if the key file is encrypted. In that case, the key file must be decrypted before you can install it. You might have to go into Guest Shell to decrypt the key file, as shown in the following example:

```
switch(config)# guestshell
[b3456@guestshell ~]$
[b3456@guestshell bootflash]$ /bin/openssl rsa -in certfilename.net.pem -out clearkey.pem

Enter pass phrase for certfilename.net.pem:
writing RSA key
[b3456@guestshell bootflash]$
[b3456@guestshell bootflash]$ exit
switch(config)#
```

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
Upload done. Please enable. Note cert and key must match.
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 Response Elements

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

**Table 5: NX-API Response Elements**

NX-API Response Element	Description
version	NX-API version.
type	Type of command to be executed.
sid	Session ID of the response. This element is valid only when the response message is chunked.
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 ( <a href="http://www.iana.org/assignments/http-status-codes/http-status-codes.xhtml">http://www.iana.org/assignments/http-status-codes/http-status-codes.xhtml</a> ).



NX-API Response Element	Description
msg	Error message associated with the returned error code.

## 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. The JSON/CLI Execution is currently supported in Cisco Nexus 3500 platform switches.



**Note** The NX-API/JSON functionality is now available on the Cisco Nexus 3500 platform switches.

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

## CLI Execution

### Show\_Command | json

#### Example Code

```
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-Dispute"], "platform_id": "N3K-C3132Q-40G", "port_id": "Ethernet1/1"}]}}
```

## 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

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, 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.

## 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>
```

```

        </__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://www.cisco.com/nxos:1.0:lldp">
  <nf:data>
    <show>
      <lldp>
        <timers>
          <__XML__OPT_Cmd_lldp_show_timers__readonly__>
            <__readonly__>
              <ttl>120</ttl>
              <reinit>2</reinit>
              <tx_interval>30</tx_interval>
              <tx_delay>2</tx_delay>
              <hold_mplier>4</hold_mplier>
              <notification_interval>5</notification_interval>
            </__readonly__>
          </__XML__OPT_Cmd_lldp_show_timers__readonly__>
        </timers>
      </lldp>
    </show>
  </nf:data>
</nf:rpc-reply>
]]>]]>
switch(config)#

```





# CHAPTER 11

## NX-API Response Codes

- [Table of NX-API Response Codes, on page 77](#)

### 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 6: 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.
IN_MSG_ERR	400	Request message is invalid.
NO_INPUT_CMD_ERR	400	No input command.
PERM_DENY_ERR	401	Permission denied.
CONF_NOT_ALLOW_SHOW_ERR	405	Configuration mode does not allow <b>show</b> .
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.
BACKEND_ERR	500	Backend processing error.
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.
USER_NOT_FOUND_ERR	500	User not found from input or cache.
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.
JSON_NOT_SUPPORTED_ERR	501	JSON not supported due to large amount of output.
MSG_TYPE_UNSUPPORTED_ERR	501	Message type 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.
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.



# CHAPTER 12

## NX-API Developer Sandbox

- [NX-API Developer Sandbox: NX-OS Releases Prior to 9.2\(2\)](#), on page 79


### 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.

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.

## 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 7: 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 <a href="http://jsonrpc.org">jsonrpc.org</a> .
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 <a href="#">Cisco Nexus NX-API References</a> .
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 8: Command Types

Message format	Command type
json-rpc	<ul style="list-style-type: none"> <li>cli — show or configuration commands</li> <li>cli-ascii — show or configuration commands, output without formatting</li> </ul>
xml	<ul style="list-style-type: none"> <li>cli_show — show commands. If the command does not support XML output, an error message will be returned.</li> <li>cli_show_ascii — show commands, output without formatting</li> <li>cli_conf — configuration commands. Interactive configuration commands are not supported.</li> <li>bash — bash commands. Most non-interactive bash commands are supported.</li> </ul> <p><b>Note</b> The bash shell must be enabled in the switch.</p>
json	<ul style="list-style-type: none"> <li>cli_show — show commands. If the command does not support XML output, an error message will be returned.</li> <li>cli_show_ascii — show commands, output without formatting</li> <li>cli_conf — configuration commands. Interactive configuration commands are not supported.</li> <li>bash — bash commands. Most non-interactive bash commands are supported.</li> </ul> <p><b>Note</b> The bash shell must be enabled in the switch.</p>
nx-api rest	<ul style="list-style-type: none"> <li>cli — configuration commands</li> </ul>
nx yang	<ul style="list-style-type: none"> <li>json — JSON structure is used for payload</li> <li>xml — XML structure is used for payload</li> </ul>

### 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.

**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 80](#).

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

You can erase the contents of the text entry box (and the **Request** and **Response** panes) by clicking **Reset** at the bottom of the top pane.



**Step 3** Click the **Convert** at the bottom of the top pane.

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

The screenshot displays the NX-API Developer Sandbox interface. At the top left is the Cisco logo and the title "NX-API Developer Sandbox". On the top right are "Quick Start" and "Logout" links. The main area is divided into several sections:

- Request Pane:** Contains a text area with a JSON payload: 

```
api/mo/sys.json
{
  "topSystem": {
    "attributes": {
      "name": "REST2CLI"
    }
  }
}
```

 Below this are "Convert" and "Reset" buttons.
- Message format:** A dropdown menu with options "json-rpc", "xml", and "json". The "json" option is selected.
- Command type:** A dropdown menu with options "cli" and "model". The "model" option is selected.
- CLI Pane:** A text area with the text "hostname REST2CLI" and a "Copy" button.
- ERROR Pane:** An empty text area with a "Copy" button.

At the bottom left, there is a status bar that says "Waiting for bam.nr-data.net...". At the bottom right, there is a copyright notice: "Copyright © 2014-2016 Cisco Systems, Inc. All rights reserved." and the version "NX-API version 1.1".

**Step 4** When a valid payload is present in the **Request** pane, you can click **POST** to send the payload as an API call to the switch.

The response from the switch appears in the **Response** pane.

**Warning** Clicking **POST** commits the command to the switch, which can result in a configuration or state change.

The screenshot shows the NX-API Developer Sandbox interface. At the top, there is a header with the Cisco logo, the title "NX-API Developer Sandbox", and links for "Quick Start" and "Logout". The main area contains a text input field with the command "Logging level netstack 6". To the right of the input field are two dropdown menus: "Message format:" with options "json-rpc", "xml", "json", "nx-api rest", and "nx yang"; and "Command type:" with options "cli" and "model". Below the input field are three buttons: "POST", "Reset", and "Convert".

Below the main area, there are two panels: "REQUEST:" and "RESPONSE:". The "REQUEST:" panel shows a JSON payload with a "loggingLevel" of "informational". The "RESPONSE:" panel shows a JSON payload with an empty "imdata" array. Both panels have "Copy" buttons. The "REQUEST:" panel also has "Python" and "Python3" buttons.

```

REQUEST:
{
  "topSystem": {
    "children": [
      {
        "ipv4Entity": {
          "children": [
            {
              "ipv4Inst": {
                "attributes": {
                  "loggingLevel": "informational"
                }
              }
            }
          ]
        }
      }
    ]
  }
}
RESPONSE:
{
  "imdata": []
}

```

**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.



## CHAPTER 13

# XML Support for ABM and LM in N3500

- [XML Support for ABM and LM in N3500](#) , on page 85

## XML Support for ABM and LM in N3500

The following commands show XML Output for ABM and LM:

### show hardware profile buffer monitor sampling

**CLI :**

```
MTC-8(config)# show hardware profile buffer monitor sampling
```

```
Sampling CLI issued at: 05/25/2016 04:18:56
```

```
Sampling interval: 200
```

**XML :**

```
MTC-8(config)# show hardware profile buffer monitor sampling | 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:mtc_usd_cli">
```

```
<nf:data>
```

```
<show>
```

```
<hardware>
```

```
<profile>
```

```
<buffer>
```

```
<monitor>
```

```
<__XML__BLK_Cmd_show_hardware_profile_buffer_monitor_summary>
```

```
<__XML__OPT_Cmd_show_hardware_profile_buffer_monitor__readonly__>
```

```
<__readonly__>
```

```
<cmd_name>Sampling CLI</cmd_name>
```

```

<cmd_issue_time>05/25/2016 04:19:12</cmd_issue_time>

<TABLE_sampling>

  <ROW_sampling>

    <sampling_interval>200</sampling_interval>

  </ROW_sampling>

</TABLE_sampling>

</__readonly__>

</__XML__OPT_Cmd_show_hardware_profile_buffer_monitor__readonly__>

</__XML__BLK_Cmd_show_hardware_profile_buffer_monitor_summary>

</monitor>

</buffer>

</profile>

</hardware>

</show>

</nf:data>

</nf:rpc-reply>

]]>]]>

```

### show hardware profile buffer monitor detail | xml

**XML :**

```

<show>
  <hardware>
    <profile>
      <buffer>
        <monitor>
          <__XML__BLK_Cmd_show_hardware_profile_buffer_monitor_summary>
            <__XML__OPT_Cmd_show_hardware_profile_buffer_monitor__readonly__>
              <__readonly__>
                <cmd_name>Detail CLI</cmd_name>
                <cmd_issue_time>10/02/2001 10:58:58</cmd_issue_time>
                <TABLE_detail_entry>
                  <ROW_detail_entry>
                    <detail_util_name>Ethernet1/1</detail_util_name>
                    <detail_util_state>Active</detail_util_state>
                  </ROW_detail_entry>
                <ROW_detail_entry>
                  <time_stamp>10/02/2001 10:58:58</time_stamp>
                  <__XML__DIGIT384k_util>0</__XML__DIGIT384k_util>
                  <__XML__DIGIT768k_util>0</__XML__DIGIT768k_util>
                  <__XML__DIGIT1152k_util>0</__XML__DIGIT1152k_util>
                  <__XML__DIGIT1536k_util>0</__XML__DIGIT1536k_util>
                  <__XML__DIGIT1920k_util>0</__XML__DIGIT1920k_util>
                  <__XML__DIGIT2304k_util>0</__XML__DIGIT2304k_util>
                  <__XML__DIGIT2688k_util>0</__XML__DIGIT2688k_util>
                  <__XML__DIGIT3072k_util>0</__XML__DIGIT3072k_util>
                  <__XML__DIGIT3456k_util>0</__XML__DIGIT3456k_util>
                </ROW_detail_entry>
              </__readonly__>
            </__XML__OPT_Cmd_show_hardware_profile_buffer_monitor__readonly__>
          </__XML__BLK_Cmd_show_hardware_profile_buffer_monitor_summary>
        </monitor>
      </buffer>
    </profile>
  </hardware>
</show>

```



```

</ROW_detail_entry>
<ROW_detail_entry>
  <time_stamp>10/02/2001 10:58:54</time_stamp>
  <__XML__DIGIT384k_util>0</__XML__DIGIT384k_util>
  <__XML__DIGIT768k_util>0</__XML__DIGIT768k_util>
  <__XML__DIGIT1152k_util>0</__XML__DIGIT1152k_util>
  <__XML__DIGIT1536k_util>0</__XML__DIGIT1536k_util>
  <__XML__DIGIT1920k_util>0</__XML__DIGIT1920k_util>
  <__XML__DIGIT2304k_util>0</__XML__DIGIT2304k_util>
  <__XML__DIGIT2688k_util>0</__XML__DIGIT2688k_util>
  <__XML__DIGIT3072k_util>0</__XML__DIGIT3072k_util>
  <__XML__DIGIT3456k_util>0</__XML__DIGIT3456k_util>
  <__XML__DIGIT3840k_util>0</__XML__DIGIT3840k_util>
  <__XML__DIGIT4224k_util>0</__XML__DIGIT4224k_util>
  <__XML__DIGIT4608k_util>0</__XML__DIGIT4608k_util>
  <__XML__DIGIT4992k_util>0</__XML__DIGIT4992k_util>
  <__XML__DIGIT5376k_util>0</__XML__DIGIT5376k_util>
  <__XML__DIGIT5760k_util>0</__XML__DIGIT5760k_util>
  <__XML__DIGIT6144k_util>0</__XML__DIGIT6144k_util>
</ROW_detail_entry>

```

### show hardware profile buffer monitor brief

XML :

```

show hardware profile buffer monitor brief | 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:mtc_usd_cli">
  <nf:data>
    <show>
      <hardware>
        <profile>
          <buffer>
            <monitor>
              <__XML__BLK_Cmd_show_hardware_profile_buffer_monitor_summary>
                <__XML__OPT_Cmd_show_hardware_profile_buffer_monitor__readonly__>
                  <__readonly__>
                    <cmd_name>Brief CLI</cmd_name>
                    <cmd_issue_time>03/21/2016 09:06:38</cmd_issue_time>
                    <TABLE_ucst_hdr>
                      <ROW_ucst_hdr>
                        <ucst_hdr_util_name>Buffer Block 1</ucst_hdr_util_name>
                        <ucst_hdr_1sec_util>0KB</ucst_hdr_1sec_util>
                        <ucst_hdr_5sec_util>0KB</ucst_hdr_5sec_util>
                        <ucst_hdr_60sec_util>N/A</ucst_hdr_60sec_util>
                        <ucst_hdr_5min_util>N/A</ucst_hdr_5min_util>
                        <ucst_hdr_1hr_util>N/A</ucst_hdr_1hr_util>
                        <ucst_hdr_total_buffer>Total Shared Buffer Available = 5397 Kbytes
                      </ucst_hdr_total_buffer>
                        <ucst_hdr_class_threshold>Class Threshold Limit = 5130 Kbytes
                      </ucst_hdr_class_threshold>
                    </ROW_ucst_hdr>
                  </TABLE_ucst_hdr>
                  <TABLE_brief_entry>
                    <ROW_brief_entry>
                      <brief_util_name>Ethernet1/45</brief_util_name>
                      <brief_1sec_util>0KB</brief_1sec_util>
                      <brief_5sec_util>0KB</brief_5sec_util>
                      <brief_60sec_util>N/A</brief_60sec_util>
                      <brief_5min_util>N/A</brief_5min_util>
                      <brief_1hr_util>N/A</brief_1hr_util>
                      <brief_util_name>Ethernet1/46</brief_util_name>
                      <brief_1sec_util>0KB</brief_1sec_util>

```





```
<brief_5min_util>N/A</brief_5min_util>
<brief_1hr_util>N/A</brief_1hr_util>
```

### show hardware profile latency monitor sampling

#### CLI

```
MTC-8(config)# show hardware profile latency monitor sampling

Sampling CLI issued at: 05/25/2016 04:19:54

Sampling interval: 20
```

#### XML

```
MTC-8(config)# show hardware profile latency monitor sampling | 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:mtc_usd_cli">

  <nf:data>

    <show>

      <hardware>

        <profile>

          <latency>

            <monitor>

              <__XML__BLK_Cmd_show_hardware_profile_latency_monitor_summary>

                <__XML__OPT_Cmd_show_hardware_profile_latency_monitor__readonly__>

                  <__readonly__>

                    <cmd_issue_time>05/25/2016 04:20:06</cmd_issue_time>

                    <device_instance>0</device_instance>

                    <TABLE_sampling>

                      <ROW_sampling>

                        <sampling_interval>20</sampling_interval>

                      </ROW_sampling>

                    </TABLE_sampling>

                  </__readonly__>

                </__XML__OPT_Cmd_show_hardware_profile_latency_monitor__readonly__>

              </__XML__BLK_Cmd_show_hardware_profile_latency_monitor_summary>

            </monitor>

          </latency>
```

```

    </profile>
  </hardware>
</show>
</nf:data>
</nf:rpc-reply>
]]>]]>

```

### show hardware profile latency monitor threshold

#### CLI

```
MTC-8(config)# show hardware profile latency monitor threshold
```

```
Sampling CLI issued at: 05/25/2016 04:20:53
```

```
Threshold Avg: 3000
```

```
Threshold Max: 300000
```

#### XML

```
MTC-8(config)# show hardware profile latency monitor threshold | 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:mtc_usd_cli">
  <nf:data>
    <show>
      <hardware>
        <profile>
          <latency>
            <monitor>
              <__XML__BLK_Cmd_show_hardware_profile_latency_monitor_summary>
                <__XML__OPT_Cmd_show_hardware_profile_latency_monitor__readonly__>
                  <__readonly__>
                    <cmd_issue_time>05/25/2016 04:21:04</cmd_issue_time>
                    <device_instance>0</device_instance>
                    <TABLE_threshold>
                      <ROW_threshold>
                        <threshold_avg>3000</threshold_avg>
                        <threshold_max>300000</threshold_max>
                      </ROW_threshold>
                    </TABLE_threshold>
                  </__readonly__>
                </__XML__OPT_Cmd_show_hardware_profile_latency_monitor__readonly__>
              </__XML__BLK_Cmd_show_hardware_profile_latency_monitor_summary>
            </monitor>
          </latency>
        </profile>
      </hardware>
    </show>
  </nf:data>
</nf:rpc-reply>

```

```
        </TABLE_threshold>
    </__readonly__>
    </__XML__OPT_Cmd_show_hardware_profile_latency_monitor__readonly__>
    </__XML__BLK_Cmd_show_hardware_profile_latency_monitor_summary>
</monitor>
</latency>
</profile>
</hardware>
</show>
</nf:data>
</nf:rpc-reply>
]]>]]>
```



## CHAPTER 14

# Converting CLI Commands to Network Configuration Format

- [Information About XMLIN, on page 93](#)
- [Licensing Requirements for XMLIN, on page 93](#)
- [Installing and Using the XMLIN Tool, on page 94](#)
- [Converting Show Command Output to XML, on page 94](#)
- [Configuration Examples for XMLIN, on page 95](#)

## 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 9: 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 <i>Cisco NX-OS Licensing Guide</i> .

# 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.

## SUMMARY STEPS

1. switch# **xmlin**
2. switch(xmlin)# **configure terminal**
3. Configuration commands
4. (Optional) switch(config)(xmlin)# **end**
5. (Optional) switch(config-if-verify)(xmlin)# **show commands**
6. (Optional) switch(config-if-verify)(xmlin)# **exit**

## DETAILED STEPS

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

**SUMMARY STEPS**

1. switch# *show-command* | **xmlin**

**DETAILED STEPS**

	Command or Action	Purpose
Step 1	switch# <i>show-command</i>   <b>xmlin</b>	Enters global configuration mode.  <b>Note</b> 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>
      <ml:cdp>
        <ml:enable/>
      </ml: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>
    
```



```
        </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 15

# OpenConfig YANG

This section contains the following topics:

- [About OpenConfig YANG, on page 99](#)
- [Guidelines and Limitations for OpenConfig YANG, on page 99](#)
- [Understanding Deletion of BGP Routing Instance, on page 104](#)
- [Verifying YANG, on page 105](#)
- [Enabling OpenConfig Support, on page 106](#)

## About OpenConfig YANG

OpenConfig YANG supports modern networking principles, such as declarative configuration and model-driven management and operations. OpenConfig provides vendor-neutral data models for configuration and monitoring of the network. And, helping with moving from a pull model to a push model, with subscriptions and event update streaming.

Beginning with Cisco NX-OS Release 9.2(1), support is added across a broad range of functional areas. Those include BGP, OSPF, Interface L2 and L3, VRFs, VLANs, and TACACs.

CPU Utilization is the sum of user and kernel utilization values. This value is approximately equal to the sum of user and kernel percentages that are displayed in “show system resources” output.

The min, max, average, min time and max time values are calculated over a 30 second moving window. In this 30-second window, CPU utilization is calculated every 5 seconds. Min, max and average values are calculated over these six samples.

The time instant at which the minimum and maximum values occur within the last 30-second window, are displayed in epoch time format.

For additional information about OpenConfig YANG, see [About OpenConfig YANG](#).

For the OpenConfig models for Cisco NX-OS 9.2(1), see [YANG Models 9.2\(1\)](#). OpenConfig YANG models are grouped by Cisco NX-OS release, so when the Cisco NX-OS release number changes, the last digits in the URL change.

## Guidelines and Limitations for OpenConfig YANG

OpenConfig YANG has the following guidelines and limitations:

- The following OpenConfig YANG limitations exist for OC-BGP-POLICY:
  - Action type is always permit for `community-set` and `as-path-set`, which applies to the following containers:
    - `/bgp-defined-sets/community-sets/community-set/`
    - `/bgp-defined-sets/as-path-sets/as-path-set/`

In OpenConfig YANG, there is no action type concept as there is in the CLI for `community-set` and `as-path-set`. Therefore, the action type is always permit for `community-set` and `as-path-set`.

- The following OpenConfig YANG limitation applies to this container:  
`/bgp-defined-sets/community-sets/community-set/`

In the CLI, `community-list` can have two different types: standard and expanded. However, in the OpenConfig YANG model, `community-set-name` has no such differentiation.

When you create the `community-set-name` through OpenConfig YANG, the following things happen internally:

- The `_std` suffix will be appended after `community-set-name` if `community-member` is in the standard form (AS:NN).
- The `_exp` suffix will be appended after `community-set-name` if `community-member` is in the expanded form (regex):

```
<community-set>
  <community-set-name>oc_commset1d</community-set-name>
  <config>
    <community-set-name>oc_commset1d</community-set-name>
    <community-member>0:1</community-member>
    <community-member>_1_</community-member>
  </config>
</community-set>
```

The preceding OpenConfig YANG configuration is mapped to the following CLI:

```
ip community-list expanded oc_commset1d_exp seq 5 permit "_1_"
ip community-list standard oc_commset1d_std seq 5 permit 0:1
```

- The following OpenConfig YANG limitation applies to this container:  
`/bgp-conditions/match-community-set/config/community-set/`

OpenConfig YANG can only map to one `community-set`, while the CLI can match to multiple instances of the `community-set`:

- In the CLI:

```
ip community-list standard 1-1 seq 1 permit 1:1
ip community-list standard 1-2 seq 1 permit 1:2
ip community-list standard 1-3 seq 1 permit 1:3
route-map To_LC permit 10
  match community 1-1 1-2 1-3
```

- The corresponding OpenConfig YANG payload follows:

```

<config>
  <routing-policy xmlns="http://openconfig.net/yang/routing-policy">
    <defined-sets>
      <bgp-defined-sets xmlns="http://openconfig.net/yang/bgp-policy">
        <community-sets>
          <community-set>
            <community-set-name>cs</community-set-name>
            <config>
              <community-set-name>cs</community-set-name>
              <community-member>1:1</community-member>
              <community-member>1:2</community-member>
              <community-member>1:3</community-member>
            </config>
          </community-set>
        </community-sets>
      </bgp-defined-sets>
    </defined-sets>
    <policy-definitions>
      <policy-definition>
        <name>To_LC</name>
        <statements>
          <statement>
            <name>10</name>
            <conditions>
              <bgp-conditions xmlns="http://openconfig.net/yang/bgp-policy">
                <match-community-set>
                  <config>
                    <community-set>cs</community-set>
                  </config>
                </match-community-set>
              </bgp-conditions>
            </conditions>
          </statement>
        </statements>
      </policy-definition>
    </policy-definitions>
  </routing-policy>
</config>

```

As a workaround, create one community with multiple statements through OpenConfig YANG:

```

ip community-list standard cs_std seq 5 permit 1:1
  ip community-list standard cs_std seq 10 permit 1:2
  ip community-list standard cs_std seq 15 permit 1:3
route-map To_LC permit 10
  match community cs_std

```

- The following OpenConfig YANG limitation applies to this container:

/bgp-conditions/state/next-hop-in

In OpenConfig YANG, the `next-hop-in` type is an IP address, but in the CLI, it is an IP prefix.

While creating the `next-hop-in` through OpenConfig YANG, the IP address is converted to a "/32" mask prefix in the CLI configuration. For example:

- Following is an example of `next-hop-in` in the OpenConfig YANG payload:

```

<policy-definition>
  <name>sc0</name>
  <statements>
    <statement>
      <name>5</name>

```

```

        <conditions>
          <bgp-conditions xmlns="http://openconfig.net/yang/bgp-policy">
            <config>
              <next-hop-in>2.3.4.5</next-hop-in>
            </config>
          </bgp-conditions>
        </conditions>
      </statement>
    </statements>
  </policy-definition>

```

- Following is an example of the same information in the CLI:

```

ip prefix-list IPV4_PFX_LIST_OPENCONFIG_sc0_5 seq 5 permit 2.3.4.5/32
route-map sc0 permit 5
  match ip next-hop prefix-list IPV4_PFX_LIST_OPENCONFIG_sc0_5

```

- The following NX-OS limitations exist for OC-BGP-POLICY:

- /bgp-actions/set-community/config/method enum "REFERENCE" is not supported.
- enum "SELF", which is supported in the OpenConfig YANG model for /bgp-actions/config/set-next-hop, is not supported.

- For OC-BGP-POLICY, /bgp-conditions/match-community-set/config/community-set get mapped only to match community <community-set>\_std, so only standard community is supported. Match to expanded community set is not supported.
- There is a limitation in replacing match-tag-set because defined sets for tag-sets are not currently implemented.

Currently, replacing match-tag-set appends the values. To replace match-tag-set, delete it, then create it again.

- The following guidelines and limitations apply to OSPF OpenConfig YANG:

- If you configure and remove an area configuration in OSPF, the deleted areas (stale entries) are still shown in DME. Those stale area entries are shown in the GETCONFIG/GET output in OpenConfig YANG.
- Only one area is supported in OpenConfig YANG in the OSPF policy match ospf-area configuration. In the CLI, you can configure to match multiple areas, such as match ospf-area 100 101. However, in OpenConfig YANG, you can configure only one area (for example, match ospf-area 100).
- The area virtual-link and area interface configurations payload cannot go under the same area list. Split the area container payload as a Virtual link area and interface area in the same payload.
- The MD5 authentication string cannot be configured in OSPF OpenConfig YANG.

In the OSPF model, Authentication-type is defined for the Authentication:

```

leaf authentication-type {
  type string;
  description
    "The type of authentication that should be used on this

```

```
    interface";
  }

```

OSPF OpenConfig YANG does not support an option for authentication password.

- The OSPF area authentication configuration is not supported. For example, `area 0.0.0.200 authentication message-digest` cannot be configured from OpenConfig YANG.
  - The OSPF/BGP instance configuration that falls under default VRF (for example, **router ospf 1/router bgp 1**) is not deleted when you delete the Protocols container with the default network instance.
- The following are guidelines and limitations for VLAN configuration between the OpenConfig payload and the Cisco Nexus 9000 interfaces:
    - When you attempt to simultaneously configure a trunk-mode interface and trunk VLANs in the same OpenConfig payload, the configuration does not complete successfully. However, when you split the payload so that the trunk-mode interface is sent first, then the trunk VLANs are sent, the configuration completes successfully.

On Cisco NX-OS interfaces, the default interface mode is **access**. To implement any trunk-related configurations, you must first change the interface mode to **trunk**, then configure the trunk VLAN ranges. Do these configurations in separate payloads.

The following examples show the separate payloads for the configuring trunk mode and VLAN ranges.

Example 1, payload configuring the interface to trunk mode.

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101">
  <edit-config>
    <target>
      <running/>
    </target>
    <config>
      <interfaces xmlns="http://openconfig.net/yang/interfaces">
        <interface>
          <name>eth1/47</name>
          <subinterfaces>
            <subinterface>
              <index>0</index>
              <config>
                <index>0</index>
              </config>
            </subinterface>
          </subinterfaces>
          <ethernet xmlns="http://openconfig.net/yang/interfaces/ethernet">
            <switched-vlan xmlns="http://openconfig.net/yang/vlan">
              <config>
                <interface-mode>TRUNK</interface-mode>
              </config>
            </switched-vlan>
          </ethernet>
        </interface>
      </interfaces>
    </config>
  </edit-config>
</rpc>
```

Example 2, payload configuring the VLAN ranges.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101">
  <edit-config>
    <target>
      <running/>
    </target>
    <config>
      <interfaces xmlns="http://openconfig.net/yang/interfaces">
        <interface>
          <name>eth1/47</name>
          <subinterfaces>
            <subinterface>
              <index>0</index>
              <config>
                <index>0</index>
              </config>
            </subinterface>
          </subinterfaces>
          <ethernet xmlns="http://openconfig.net/yang/interfaces/ethernet">
            <switched-vlan xmlns="http://openconfig.net/yang/vlan">
              <config>
                <native-vlan>999</native-vlan>
                <trunk-vlans xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
nc:operation="delete">1..4094</trunk-vlans>
                <trunk-vlans>401</trunk-vlans>
                <trunk-vlans>999</trunk-vlans>
              </config>
            </switched-vlan>
          </ethernet>
        </interface>
      </interfaces>
    </config>
  </edit-config>
</rpc>

```

- Because of the design of OpenConfig YANG, when you configure VLANs, there must be no overlap between the VLANs in the payload and the VLANs already configured on an interface. If an overlap exists, the configuration through OpenConfig is not successful. Make sure that the VLANs configured on an interface are different from the VLANs in the OpenConfig payload. Pay particular attention to the starting and ending VLANs in a range.

## Understanding Deletion of BGP Routing Instance

With OpenConfig YANG network-instance (OCNI), when attempting to delete only the BGP configuration of the default VRF instead of deleting the entire BGP routing instance, BGP information might not be deleted at the protocols/BGP level. In this situation, when the delete is at the protocols or BGP level with the autonomous system number in the payload, only the configuration of the default VRF is deleted instead of removing the entire BGP routing instance.

Following is an example payload that would be used to delete the configuration under the default VRF in BGP.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101">
  <edit-config>
    <target>
      <running/>
    </target>
    <config>
      <network-instances xmlns="http://openconfig.net/yang/network-instance">

```



```

<network-instance>
  <name>default</name>
  <protocols>
    <protocol>
      <identifier>BGP</identifier>
      <name>bgp</name>
      <bgp xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" nc:operation="delete">
        <global>
          <config>
            <as>100</as>
          </config>
        </global>
      </bgp>
    </protocol>
  </protocols>
</network-instance>
</network-instances>
</config>
</edit-config>
</rpc>

```

**Expected Behavior:** The BGP routing instance itself should be deleted, which is the equivalent to **no router bgp 100**.

**Actual Behavior:** Only the BGP configuration under the default VRF is deleted, and there is no equivalent single CLI configuration.

Following is the running configuration before the delete operation:

```

router bgp 100
  router-id 1.2.3.4
  address-family ipv4 unicast
  vrf abc
    address-family ipv4 unicast
      maximum-paths 2

```

And following is the running configuration after the delete operation:

```

router bgp 100
  vrf abc
    address-family ipv4 unicast
      maximum-paths 2

```

## Verifying YANG

Use the following commands to verify YANG settings: :

**Table 10: YANG Verification**

Command	Description
<code>show telemetry yang direct-path cisco-nxos-device</code>	Displays the paths which are supported.

## Enabling OpenConfig Support

To enable or disable OpenConfig support on the programmability agents (NETCONF, RESTCONF and gRPC), configure "[no] feature openconfig". For example:

```
switch(config)# feature netconf
switch(config)# feature restconf
switch(config)# feature grpc
switch(config)# feature openconfig
```



---

**Note** In previous releases, mt-x-openconfig-all RPM was downloaded separately and installed. This method is deprecated in 10.2(2) release.

---



## CHAPTER 16

# XML Management Interface

---

This section contains the following topics:

- [About the XML Management Interface, on page 107](#)
- [Licensing Requirements for the XML Management Interface, on page 108](#)
- [Prerequisites to Using the XML Management Interface, on page 109](#)
- [Using the XML Management Interface, on page 109](#)
- [Information About Example XML Instances, on page 121](#)
- [Additional References, on page 127](#)

## 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 111](#) and [RFC 4741](#).

For more information about using NETCONF over SSH, see [RFC 4742](#).

This section includes the following topics:

- [NETCONF Layers, on page 107](#)
- [SSH xmlagent, on page 108](#)

### NETCONF Layers

The following are the NETCONF layers:

Table 11: 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 111](#) 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 <i>Cisco NX-OS Licensing Guide</i> .

# 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

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.

## 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

- 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"
  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 12: NETCONF Operations in Cisco NX-OS**

NETCONF Operation	Description	Example
close-session	Closes the current XML server session.	<a href="#">NETCONF Close Session Instance, on page 121</a>
commit	Sets the running configuration to the current contents of the candidate configuration.	<a href="#">NETCONF Commit Instance - Candidate Configuration Capability, on page 126</a>
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.	<a href="#">NETCONF Confirmed-commit Instance , on page 126</a>
copy-config	Copies the content of source configuration datastore to the target datastore.	<a href="#">NETCONF copy-config Instance, on page 122</a>
delete-config	Operation not supported.	—
edit-config	Configures features in the running configuration of the device. You use this operation for configuration commands.	<a href="#">NETCONF edit-config Instance, on page 122</a> <a href="#">NETCONF rollback-on-error Instance , on page 126</a>
get	Receives configuration information from the device. You use this operation for <b>show</b> commands. The source of the data is the running configuration.	<a href="#">Creating NETCONF XML Instances, on page 111</a>
get-config	Retrieves all or part of a configuration	<a href="#">NETCONF get-config Instance, on page 124</a>
kill-session	Closes the specified XML server session. You cannot close your own session. See the close-session NETCONF operation.	<a href="#">NETCONF Kill-session Instance, on page 122</a>

NETCONF Operation	Description	Example
lock	Allows the client to lock the configuration system of a device.	<a href="#">NETCONF Lock Instance, on page 124</a>
unlock	Releases the configuration lock that the session issued.	<a href="#">NETCONF unlock Instance, on page 125</a>
validate	Checks a candidate configuration for syntactical and semantic errors before applying the configuration to the device.	<a href="#">NETCONF validate Capability Instance , on page 127</a>

## 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 110](#) section.

Using this schema, it is possible to build an XML instance. In the following examples, the relevant portions of the ncli.xsd schema file that was used to build [Creating NETCONF XML Instances, on page 111](#) 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 120](#) 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 13: 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 <i>Configuration CLI Commands Sent Through &lt;exec-command&gt;</i> .

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 120](#)
- [Interpreting Tags Encapsulated in the Data Tag, on page 120](#)

## 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 14: RPC Response Elements**

Element	Description
<code>&lt;ok&gt;</code>	The RPC request completed successfully. This element is used when no data is returned in the response.
<code>&lt;data&gt;</code>	The RPC request completed successfully. The data associated with the RPC request is enclosed in the <code>&lt;data&gt;</code> element.
<code>&lt;rpc-error&gt;</code>	The RPC request failed. Error information is enclosed in the <code>&lt;rpc-error&gt;</code> 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 121](#)
- [NETCONF Kill-session Instance, on page 122](#)
- [NETCONF copy-config Instance, on page 122](#)
- [NETCONF edit-config Instance, on page 122](#)
- [NETCONF get-config Instance, on page 124](#)
- [NETCONF Lock Instance, on page 124](#)
- [NETCONF unlock Instance, on page 125](#)
- [NETCONF Commit Instance - Candidate Configuration Capability, on page 126](#)
- [NETCONF Confirmed-commit Instance , on page 126](#)
- [NETCONF rollback-on-error Instance , on page 126](#)
- [NETCONF validate Capability Instance , on page 127](#)

### 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**

<b>RFCs</b>	<b>Title</b>
<a href="#">RFC 4741</a>	NETCONF Configuration Protocol
<a href="#">RFC 4742</a>	Using the NETCONF Configuration Protocol over Secure Shell (SSH)





## PART I

# Model-Driven Programmability

- [Managing Components, on page 131](#)
- [Converting CLI Commands to Network Configuration Format, on page 137](#)
- [gNMI - gRPC Network Management Interface, on page 143](#)
- [gNOI-gRPC Network Operations Interface, on page 169](#)
- [Model Driven Telemetry, on page 175](#)





## CHAPTER 17

# Managing Components

- [About the Component RPM Packages, on page 131](#)
- [Preparing For Installation, on page 133](#)
- [Downloading Components from the Cisco Artifactory, on page 134](#)
- [Installing RPM Packages, on page 134](#)

## About the Component RPM Packages

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 uninstall 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	-
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

Model Name	Model Rev	Model Ver	Package Name	Dependencies
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 Artifacts

The NX-OS Programmable Interface Component RPMs can be downloaded from the Cisco Artifacts 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/artifacts/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 Artifacts, 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 133](#), confirm that all line cards in the Active and Standby devices are up and ready.

---

**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 18

# Converting CLI Commands to Network Configuration Format

- [Information About XMLIN, on page 137](#)
- [Licensing Requirements for XMLIN, on page 137](#)
- [Installing and Using the XMLIN Tool, on page 138](#)
- [Converting Show Command Output to XML, on page 138](#)
- [Configuration Examples for XMLIN, on page 139](#)

## 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 15: 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 <i>Cisco NX-OS Licensing Guide</i> .

# 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.

## SUMMARY STEPS

1. switch# **xmlin**
2. switch(xmlin)# **configure terminal**
3. Configuration commands
4. (Optional) switch(config)(xmlin)# **end**
5. (Optional) switch(config-if-verify)(xmlin)# **show commands**
6. (Optional) switch(config-if-verify)(xmlin)# **exit**

## DETAILED STEPS

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

## SUMMARY STEPS

1. switch# *show-command* | **xmlin**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# <i>show-command</i>   <b>xmlin</b>	Enters global configuration mode.  <b>Note</b> 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>
      <ml:cdp>
        <ml:enable/>
      </ml: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 19

## gNMI - gRPC Network Management Interface

This chapter contains the following topics:

- [About gNMI, on page 143](#)
- [gNMI RPC and SUBSCRIBE, on page 144](#)
- [Guidelines and Limitations for gNMI, on page 145](#)
- [Configuring gNMI, on page 146](#)
- [Configuring Server Certificate, on page 147](#)
- [Generating Key/Certificate Examples , on page 148](#)
- [Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3\(2\) and Earlier, on page 149](#)
- [Examples for Generating and Configuring Key/Certificate for Cisco NX-OS Release 9.3\(3\) and Later, on page 150](#)
- [Verifying gNMI, on page 151](#)
- [Clients, on page 152](#)
- [Sample DME Subscription - PROTO Encoding, on page 152](#)
- [Capabilities, on page 154](#)
- [Get, on page 158](#)
- [Set, on page 159](#)
- [Subscribe, on page 160](#)
- [Troubleshooting, on page 163](#)

### About gNMI

Cisco NX-OS supports gNMI for dial-in subscription to telemetry applications running on switches. Although the 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 16: Supported gNMI RPCs**

gNMI RPC	Supported?
Get	No

gNMI RPC	Supported?
Set	No
Capabilities	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 17: 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:

- gRPC traffic destined for a Nexus device will hit the control-plane policer (CoPP) in the default class. To limit the possibility of gRPC drops, configure a custom CoPP policy using the gRPC configured port in the management class.
- For Cisco NX-OS prior to 9.3(x), information about supported platforms, see *Platform Support for Programmability Features* in the guide for that release.
- The gNMI feature supports the Subscribe and Capability gNMI RPCs.
- 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. Applies to gNMI ON\_CHANGE mode only.  
  
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.
- The feature does not support a path prefix in the Subscription request, but the Subscription can contain an empty prefix field.
- The gRPC process that supports gNMI uses the HIGH\_PRIO control group, 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 one hour after the call has ended.
  - If the total number of calls exceeds 2000, the gRPC agent purges ended calls based on the internal cleanup routine.

gRPC functionality now includes the default VRF for a total of two gRPC servers on each 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 is not desirable.

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

- VRF boundaries are strictly enforced, so each gRPC server process requests independent of the other. Requests do not cross between VRFs.
- The two servers are not HA or fault tolerant. One gRPC server does not back up the other, and there is no switchover or switchback between them.
- Any limits for the gRPC server are per VRF.

The following are the limitations for gNMI:

- multi-level wildcard "..." in path is not allowed
- wildcard '\*' in the top of the path is not allowed
- wildcard '\*' in key name is not allowed
- wildcard and value cannot be mixed in keys

The following table shows the wildcard support details for gNMI:

**Table 18: Wildcard Support for gNMI Requests**

Type of Request	Wildcard Support
gNMI GET	YES
gNMI SET	NO
gNMI SUBSCRIBE, ONCE	YES
gNMI SUBSCRIBE, POLL	YES
gNMI SUBSCRIBE, STREAM, SAMPLE	YES
gNMI SUBSCRIBE, STREAM, TARGET_DEFINED	YES
gNMI SUBSCRIBE, STREAM, ON_CHANGE	NO

## Configuring gNMI

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

### SUMMARY STEPS

1. **configure terminal**
2. **feature grpc**
3. **grpc gnmi max-concurrent-call *number***

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enters global configuration mode.
Step 2	<b>feature grpc</b> <b>Example:</b> <pre>switch-1# feature grpc switch-1(config)#</pre>	Enables the gRPC agent, which supports the gNMI interface for dial-in.
Step 3	<b>grpc gnmi max-concurrent-call <i>number</i></b> <b>Example:</b> <pre>switch-1(config)# grpc gnmi max-concurrent-call 16 switch-1(config)#</pre>	<p>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.</p> <p>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.</p> <p>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.</p> <p><b>Note</b> The configured limit does not affect the gRPCConfigOper service.</p>

## Configuring Server Certificate

When you configured a TLS certificate and imported successfully onto the switch, the following is an example of the **show grpc gnmi service statistics** command output.

```
#show grpc gnmi service statistics

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

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

Cert notBefore : Mon Jan 27 15:34:08 PDT 2020
Cert notAfter  : Tue Jan 26 15:34:08 PDT 2021

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
```

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.

When the certificate is expired or failed to install successfully, you will see the 1-D default certificate. The following is an example of the **show grpc gnmi service statistics** command output.

```
#show grpc gnmi service statistics

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

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

Cert notBefore : Wed Mar 11 19:43:01 PDT 2020
Cert notAfter  : Thu Mar 12 19:43:01 PDT 2020

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
```

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 149](#)

# 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 9000 Series NX-OS Security Configuration Guide, Release 9.3(x)*.

**Step 1** Generate the selfsigned key and pem files.

```
switch# run bash sudo su
bash-4.3# openssl req -x509 -newkey rsa:2048 -keyout self_sign2048.key -out self_sign2048.pem -days
365 -nodes
```

**Step 2** After generating the key and pem files, modify the mtz.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
```

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

The following is an example for generating key/certificate.



**Note** This task is an example of how a certificate can be generated on a switch. You can also generate a certificate in any Linux environment. In a production environment, you should consider using a CA signed certificate.

For more information on generating identity certificates, see the [Installing Identity Certificates](#) section of the *Cisco Nexus 9000 Series NX-OS Security Configuration Guide, Release 9.3(x)*.

**Step 1** Generate the selfsigned key and pem files.

```
switch# run bash sudo su
bash-4.3# openssl req -x509 -newkey rsa:2048 -keyout self_sign2048.key -out self_sign2048.pem -days
365 -nodes
```

**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** 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 4** 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 5 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
<b>show grpc gnmi service statistics</b>	<p>Displays a summary of the agent running status, respectively for the management VRF, or the default VRF (if configured). It also displays:</p> <ul style="list-style-type: none"> <li>• Basic overall counters</li> <li>• Certificate expiration time</li> </ul> <p><b>Note</b> If the certificate is expired, the agent cannot accept requests.</p>

### 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

```

## 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](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:

- 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 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_cert ./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:

- 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 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"
      }
    }
  ]
}

```

## 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="$DATETIME$
$COMPONENTID$ $TYPE$: $MSG$">$DATETIME$ $COMPONENTID$ $TYPE$
$SRCLINE$ @ $SRCLINE$ $FCNINFO$: $MSG$</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>
    </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 /volaiflrels -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 uwsgi.py
drwxrwxrwx 2 root root 40 Jun 21 14:43 virtual-instance
bash-4.3#.
```







## CHAPTER 20

# gNOI-gRPC Network Operations Interface

---

- [About gNOI, on page 169](#)
- [Supported gNOI RPCs, on page 169](#)
- [System Proto, on page 170](#)
- [OS Proto, on page 171](#)
- [Cert Proto, on page 172](#)
- [File Proto, on page 172](#)
- [Guidelines and Limitations, on page 173](#)
- [Verifying gNOI, on page 173](#)

## About gNOI

gRPC Network Operations Interface (gNOI) defines a set of gRPC-based micro-services for executing operational commands on network devices. The operational commands supported are Ping, Traceroute, Time, SwitchControlProcessor, Reboot, RebootStatus, CancelReboot, Activate and Verify.

gNOI uses gRPC as the transport protocol and the configuration is same as that of gNMI. For details on configuration, please refer to [Configuring gNMI](#).

To send gNOI RPC requests, user needs a client that implements the gNOI client interface for each RPC.

In Cisco NX-OS Release 10.1(1) the gNOI defines Remote Procedure Calls (RPCs) for a limited number of components and some of them related to hardware (like optical interfaces).

Proto files are defined for the gRPC micro-services and are available at [GitHub](#).

## Supported gNOI RPCs

The following are the supported gNOI RPCs:

Table 19:

Proto	gNOI RPC	Supported
System	Ping	Yes
	Traceroute	Yes
	Time	Yes
	SwitchControl Processor	Yes
	Reboot	Yes
	RebootStatus	Yes
	CancelReboot	Yes
OS	Activate	Yes
	Verify	Yes
Cert	LoadCertificate	Yes
File	Get	Yes
	Stat	Yes
	Remove	Yes

## System Proto

The System proto service is a collection of operational RPCs that allows the management of a target outside the configuration and telemetry pipeline.

The following are the RPC support details for System proto:

RPC	Support	Description	Limitation
Ping	ping/ping6 cli command	Executes the ping command on the target and streams back the results. Some targets may not stream any results until all results are available. If a packet count is not explicitly provided, 5 is used.	do_not_resolve option is not supported.

RPC	Support	Description	Limitation
Traceroute	traceroute/traceroute6 cli command	Executes the traceroute command on the target and streams back the results. Some targets may not stream any results until all results are available. Max hop count of 30 is used.	initial_ttl, marx_ttl, wait, do_not_fragment, do_not_resolve and l4protocol options are not supported.
Time	local time	Returns the current time on the target. Typically used to test if the target is responding.	-
SwitchControl Processor	system switchover cli command	Switches from the current route processor to the provided route processor. Switchover happens instantly and the response may not be guaranteed to return to the client.	Switchover occurs instantly. As a result, the response may not be guaranteed to return to the client.
Reboot	cli: reload [module]	Causes the target to reboot.	message option is not supported, delay option is supported for switch reload, and the path option accepts one module number.
RebootStatus	show version [module] cli command	Returns the status of the reboot for the target.	-
CancelReboot	reload cancel	Cancels any pending reboot request.	-



**Note** The SetPackage RPC is not supported.

## OS Proto

The OS service provides an interface for OS installation on a Target. The OS package file format is platform dependent. The platform must validate that the OS package that is supplied is valid and bootable. This must include a hash check against a known good hash. It is recommended that the hash is embedded in the OS package.

The Target manages its own persistent storage, and OS installation process. It stores a set of distinct OS packages, and always proactively frees up space for incoming new OS packages. It is guaranteed that the

Target always has enough space for a valid incoming OS package. The currently running OS packages must never be removed. The Client must expect that the last successfully installed package is available.

The following are the RPC support details for OS proto:

RPC	Support	Description	Limitation
Activate	install all nxos bootflash:///img_name	Sets the requested OS version as the version that is used at the next reboot. This RPC reboots the Target.	Cannot rollback or recover if the reboot fails.
Verify	show version	Verify checks the running OS version. This RPC may be called multiple times while the Target boots until it is successful.	-



**Note** The Install RPC is not supported.

## Cert Proto

The certificate management service is exported by targets. Rotate, Install and other Cert Proto RPCs are not supported.

The following are the RPC support details for Cert proto:

RPC	Support	Description	Limitation
LoadCertificate	crypto ca import <trustpoint> pkcs12 <file> <passphrase>	Loads a bundle of CA certificates.	-

## File Proto

The file proto streams messages based on the features of the file.proto RPCs. Put and other RPCs that are not listed here are not supported in File Proto.

Get, Stat, and Remove RPCs support file systems - bootflash, bootflash://sup-remote, logflash, logflash://sup-remote, usb, volatile, volatile://sup-remote and debug.

The following are the RPC support details for File proto:

RPC	Description	Limitation
Get	Get reads and streams the contents of a file from the target. The file is streamed by sequential messages, each containing up to 64 KB of data. A final message is sent prior to closing the stream that contains the hash of the data sent. An error is returned if the file does not exist or there was an error reading the file.	Maximum file size limit is 32 MB.
Stat	Stat returns metadata about a file on the target. An error is returned if the file does not exist or if there is an error in accessing the metadata.	-
Remove	Remove removes the specified file from the target. An error is returned if the file does not exist, is a directory, or the remove operation encounters an error.	-

## Guidelines and Limitations

The gNOI feature has the following guidelines and limitations:

- A maximum of 16 active gNOI RPCs are supported.
- The Cisco Nexus 9000 series switches would run one endpoint with one gNMI service and two gNOI microservices.
- In 10.1(1) release, the gNOI RPCs are implemented with the equivalent CLI. The existing CLI restrictions or valid options remain as applicable.
- gRPC traffic destined for a Nexus device will hit the control-plane policer (CoPP) in the default class. To limit the possibility of gRPC drops, configure a custom CoPP policy using the gRPC configured port in the management class.

## Verifying gNOI

To verify the gNOI configuration, enter the following commands:

Command	Description
clear grpc gnoi rpc	Serves to clean up the counters or calls.

Command	Description
debug grpc events {events errors} show grpc nxsdk event-history {events errors}	Debugs the events and errors from the event history.
show grpc internal gnoi rpc {summary detail}	An internal keyword command added for serviceability.



## CHAPTER 21

# Model Driven Telemetry

- [About Telemetry, on page 175](#)
- [Licensing Requirements for Telemetry, on page 177](#)
- [Installing and Upgrading Telemetry, on page 177](#)
- [Guidelines and Limitations for Model Driven Telemetry, on page 178](#)
- [Configuring Telemetry Using the CLI, on page 181](#)
- [Configuring Telemetry Using the NX-API, on page 195](#)
- [Telemetry Path Labels, on page 208](#)
- [Native Data Source Paths, on page 224](#)
- [Streaming Syslog, on page 235](#)
- [Additional References, on page 242](#)

## 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.)

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 }
```

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.
- **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 <i>Cisco NX-OS Licensing Guide</i> .

## 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:

As in the following example:

### 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 appears 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 Model Driven Telemetry

Telemetry has the following configuration guidelines and limitations:

- Cisco NX-OS releases that support the data management engine (DME) Native Model support Telemetry.
- Support is in place for the following:
  - DME data collection
  - NX-API data sources
  - Google protocol buffer (GPB) encoding over Google Remote Procedure Call (gRPC) transport
  - JSON encoding over HTTP
- The smallest sending interval (cadence) supported is five seconds for a depth of 0. The minimum cadence values for depth values greater than 0 depends on the size of the data being streamed out. Configuring any cadences below the minimum value may result in undesirable system behavior.
- Telemetry supports up to five remote management receivers (destinations). Configuring more than five remote receivers may result in undesirable system behavior.
- Telemetry can consume up to 20% of the CPU resource.

### Configuration Commands After Downgrading to an Older Release

After a downgrade to an older release, some configuration commands or command options can fail because the older release may not support them. When downgrading to an older release, unconfigure and reconfigure the telemetry feature after the new image comes up. This sequence avoids the failure of unsupported commands or command options.

The following example shows this procedure:

- Copy the telemetry configuration to a file:

```
switch# show running-config | section telemetry
feature telemetry
telemetry
  destination-group 100
    ip address 1.2.3.4 port 50004 protocol gRPC encoding GPB
    use-chunking size 4096
  sensor-group 100
    path sys/bgp/inst/dom-default depth 0
  subscription 600
    dst-grp 100
    snsr-grp 100 sample-interval 7000
switch# show running-config | section telemetry > telemetry_running_config
switch# show file bootflash:telemetry_running_config
feature telemetry
telemetry
```

```

destination-group 100
  ip address 1.2.3.4 port 50004 protocol gRPC encoding GPB
  use-chunking size 4096
sensor-group 100
  path sys/bgp/inst/dom-default depth 0
subscription 600
  dst-grp 100
  snsr-grp 100 sample-interval 7000
switch#

```

- Execute the downgrade operation. When the image comes up and the switch is ready, copy the telemetry configurations back to the switch.

```

switch# copy telemetry_running_config running-config echo-commands
`switch# config terminal`
`switch(config)# feature telemetry`
`switch(config)# telemetry`
`switch(config-telemetry)# destination-group 100`
`switch(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 disables the connection to the gRPC receiver if the gRPC receiver sends 20 errors. Unconfigure then reconfigure the receiver's IP address under the destination group to enable the gRPC receiver. Errors include:

- The gRPC client sends the wrong certificate for secure connections.
- The gRPC receiver takes too long to handle client messages and incurs a timeout. Avoid timeouts by processing messages using a separate message processing thread.

### NX-API Sensor Path Limitations

NX-API can collect and stream switch information not yet in the DME using **show** commands. However, using the NX-API instead of streaming data from the DME has inherent scale limitations as outlined:

- The switch backend dynamically processes NX-API calls such as **show** commands,
- NX-API spawns several processes that can consume up to a maximum of 20% of the CPU.
- NX-API data translates from the CLI to XML to JSON.

The following is a suggested user flow to help limit excessive NX-API sensor path bandwidth consumption:

1. Check whether the **show** command has NX-API support. You can confirm whether NX-API supports the command from the VSH with the pipe option: `show <command> | json` or `show <command> | json pretty`.



**Note** Avoid commands that take the switch more than 30 seconds to return JSON output.

2. Refine the **show** command to include any filters or options.
  - Avoid enumerating the same command for individual outputs; for example, **show vlan id 100** , **show vlan id 101** , and so on. Instead, use the CLI range options; for example, **show vlan id 100-110,204** , whenever possible to improve performance.

If only the summary or counter is needed, then avoid dumping a whole show command output to limit the bandwidth and data storage that is required for data collection.
3. Configure telemetry with sensor groups that use NX-API as their data sources. Add the **show** commands as sensor paths
4. Configure telemetry with a cadence of five times the processing time of the respective **show** command to limit CPI usage.
5. Receive and process the streamed NX-API output as part of the existing DME collection.

### Telemetry VRF Support

Telemetry VRF support allows you to specify a transport VRF, which means that the telemetry data stream can egress through front-panel ports and avoid possible competition between SSH or NGINX control sessions.

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

The following example specifies the transport VRF:

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

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

### 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 hostname 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
```

## 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.

#### SUMMARY STEPS

1. **configure terminal**
2. **feature telemetry**
3. **feature nxapi**
4. **nxapi use-vrf management**
5. **telemetry**
6. (Optional) **certificate** *certificate\_path* *host\_URL*
7. **sensor-group** *sgrp\_id*
8. **path** *sensor\_path* **depth** **0** [**filter-condition** *filter*] [**alias** *path\_alias*]
9. **destination-group** *dgrp\_id*
10. (Optional) **ip address** *ip\_address* **port** *port* **protocol** *procedural-protocol* **encoding** *encoding-protocol*
11. **ip\_version address** *ip\_address* **port** *portnum*
12. **subscription** *sub\_id*
13. **snsr-grp** *sgrp\_id* **sample-interval** *interval*
14. **dst-grp** *dgrp\_id*

#### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b>  <b>Example:</b> <pre>switch# <b>configure terminal</b> switch(config)#</pre>	Enter the global configuration mode.
Step 2	<b>feature telemetry</b>	Enable the streaming telemetry feature.

	Command or Action	Purpose
<b>Step 3</b>	<b>feature nxapi</b>	Enable NX-API.
<b>Step 4</b>	<b>nxapi use-vrf management</b>	Enable the VRF management to be used for NX-API communication.
<b>Step 5</b>	<b>telemetry</b> <b>Example:</b> switch(config)# <b>telemetry</b> switch(config-telemetry)#	Enter configuration mode for streaming telemetry.
<b>Step 6</b>	(Optional) <b>certificate</b> <i>certificate_path</i> <i>host_URL</i> <b>Example:</b> switch(config-telemetry)# <b>certificate</b> /bootflash/server.key localhost	Use an existing SSL/TLS certificate.
<b>Step 7</b>	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> switch(config-telemetry)# <b>sensor-group</b> 100 switch(conf-tm-sensor)#	Create a sensor group with ID <i>sgrp_id</i> and enter sensor group configuration mode.  Currently only numeric ID values are supported. The sensor group defines nodes that will be monitored for telemetry reporting.
<b>Step 8</b>	<b>path</b> <i>sensor_path</i> <b>depth 0</b> [ <b>filter-condition</b> <i>filter</i> ] [ <b>alias</b> <i>path_alias</i> ] <b>Example:</b> <ul style="list-style-type: none"> <li>The following command is applicable for DME, not for NX-API: <pre>switch(conf-tm-sensor)# path sys/bd/bd-[vlan-100] depth 0 filter-condition eq(12BD.operSt, "down")</pre>Use the following syntax for state-based filtering to trigger only when <b>operSt</b> changes from <b>up</b> to <b>down</b>, with no notifications of when the MO changes. <pre>switch(conf-tm-sensor)# path sys/bd/bd-[vlan-100] depth 0 filter-condition and(updated(12BD.operSt),eq(12BD.operSt,"down"))</pre></li> <li>The following command is applicable for NX-API, not for DME: <pre>switch(conf-tm-sensor)# path "show interface" depth 0</pre></li> </ul>	Add a sensor path to the sensor group.  <ul style="list-style-type: none"> <li>The <b>depth</b> setting specifies the retrieval level for the sensor path. Depth settings of <b>0 - 32</b>, <b>unbounded</b> are supported. <b>Note</b> <b>depth 0</b> is the default depth. NX-API-based sensor paths can only use <b>depth 0</b>. If a path is subscribed for the event collection, the depth only supports 0 and unbounded. Other values would be treated as 0.</li> <li>The optional <b>filter-condition</b> 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 <b>sys/bd/bd-[vlan]</b> of <b>eq(12Bd.operSt, "down")</b> triggers when the operSt changes, and when the DN's property changes while the operSt remains <b>down</b>, such as a <b>no shutdown</b> command is issued while the VLAN is operationally <b>down</b>.</li> </ul>

	Command or Action	Purpose
Step 9	<b>destination-group</b> <i>dgrp_id</i> <b>Example:</b> <pre>switch(conf-tm-sensor)# destination-group 100 switch(conf-tm-dest)#</pre>	Create a destination group and enter destination group configuration mode. Currently <i>dgrp_id</i> only supports numeric ID values.
Step 10	(Optional) <b>ip address</b> <i>ip_address</i> <b>port</b> <i>port</i> <b>protocol</b> <i>procedural-protocol</i> <b>encoding</b> <i>encoding-protocol</i> <b>Example:</b> <pre>switch(conf-tm-sensor)# ip address 171.70.55.69 port 50001 protocol gRPC encoding GPB switch(conf-tm-sensor)# ip address 171.70.55.69 port 50007 protocol HTTP encoding JSON</pre>	Specify an IPv4 IP address and port to receive encoded telemetry data. <b>Note</b> gRPC is the default transport protocol. GPB is the default encoding.
Step 11	<b>ip_version</b> <b>address</b> <i>ip_address</i> <b>port</b> <i>portnum</i> <b>Example:</b> <ul style="list-style-type: none"> <li>For IPv4:</li> </ul> <pre>switch(conf-tm-dest)# ip address 1.2.3.4 port 50003</pre>	Create a destination profile for the outgoing data. 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 12	<b>subscription</b> <i>sub_id</i> <b>Example:</b> <pre>switch(conf-tm-dest)# subscription 100 switch(conf-tm-sub)#</pre>	Create a subscription node with ID and enter the subscription configuration mode. Currently <i>sub_id</i> only supports numeric ID values. <b>Note</b> When subscribing to a DN, check whether the DN is supported by DME using REST to ensure that events will stream.
Step 13	<b>snsr-grp</b> <i>sgrp_id</i> <b>sample-interval</b> <i>interval</i> <b>Example:</b> <pre>switch(conf-tm-sub)# snsr-grp 100 sample-interval 15000</pre>	Link the sensor group with ID <i>sgrp_id</i> to this subscription and set the data sampling interval in milliseconds. An interval value of 0 creates an event-based subscription, in which telemetry data is sent only upon changes under the specified MO. An interval value greater than 0 creates a frequency-based subscription, in which telemetry data is sent periodically at the specified interval. For example, an interval value of 15000 results in the sending of telemetry data every 15 seconds.
Step 14	<b>dst-grp</b> <i>dgrp_id</i> <b>Example:</b> <pre>switch(conf-tm-sub)# dst-grp 100</pre>	Link the destination group with ID <i>dgrp_id</i> to this subscription.

## Configuring Cadence for YANG Paths

The cadence for YANG paths must be greater than the total streaming time. If the total streaming time and cadence are incorrectly configured, gathering telemetry data can take longer than the streaming interval. In this situation, you can see:

- Queues that incrementally fill because telemetry data is accumulating faster than it is streaming to the receiver.
- Stale telemetry data which is not from the current interval.

Configure the cadence to a value greater than the total streaming time.

## SUMMARY STEPS

1. **show telemetry control database sensor-groups**
2. **sensor group** *number*
3. **subscription** *number*
4. **snsr-grp** *number sample-interval milliseconds*
5. **show system resources**

## DETAILED STEPS

	Command or Action	Purpose																																																																																				
Step 1	<b>show telemetry control database sensor-groups</b> <b>Example:</b> <pre>switch-1# show telemetry control database sensor-groups Sensor Group Database size = 2</pre> <table border="1"> <thead> <tr> <th>Row ID</th> <th>Sensor Group ID</th> <th>Sensor Group type</th> <th>SubID</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>2</td> <td>Timer /YANG</td> <td>5000</td> </tr> <tr> <td colspan="4">/Running 1 1</td> </tr> <tr> <td colspan="4">Collection Time in ms (Cur/Min/Max): 2444/2294/2460</td> </tr> <tr> <td colspan="4">Encoding Time in ms (Cur/Min/Max): 56/55/57</td> </tr> <tr> <td colspan="4">Transport Time in ms (Cur/Min/Max): 0/0/1</td> </tr> <tr> <td colspan="4"><b>Streaming Time in ms (Cur/Min/Max): 2515/2356/28403</b></td> </tr> <tr> <td colspan="4">Collection Statistics:</td> </tr> <tr> <td colspan="4">collection_id_dropped = 0</td> </tr> <tr> <td colspan="4">last_collection_id_dropped = 0</td> </tr> <tr> <td colspan="4">drop_count = 0</td> </tr> <tr> <td>2</td> <td>1</td> <td>Timer /YANG</td> <td>5000</td> </tr> <tr> <td colspan="4">/Running 1 1</td> </tr> <tr> <td colspan="4">Collection Time in ms (Cur/Min/Max): 144/142/1471</td> </tr> <tr> <td colspan="4">Encoding Time in ms (Cur/Min/Max): 0/0/1</td> </tr> <tr> <td colspan="4">Transport Time in ms (Cur/Min/Max): 0/0/0</td> </tr> <tr> <td colspan="4"><b>Streaming Time in ms (Cur/Min/Max): 149/147/23548</b></td> </tr> <tr> <td colspan="4">Collection Statistics:</td> </tr> <tr> <td colspan="4">collection_id_dropped = 0</td> </tr> <tr> <td colspan="4">last_collection_id_dropped = 0</td> </tr> <tr> <td colspan="4">drop_count = 0</td> </tr> </tbody> </table> <pre>switch-1# telemetry destination-group 1 ip address 192.0.2.1 port 9000 protocol HTTP</pre>	Row ID	Sensor Group ID	Sensor Group type	SubID	1	2	Timer /YANG	5000	/Running 1 1				Collection Time in ms (Cur/Min/Max): 2444/2294/2460				Encoding Time in ms (Cur/Min/Max): 56/55/57				Transport Time in ms (Cur/Min/Max): 0/0/1				<b>Streaming Time in ms (Cur/Min/Max): 2515/2356/28403</b>				Collection Statistics:				collection_id_dropped = 0				last_collection_id_dropped = 0				drop_count = 0				2	1	Timer /YANG	5000	/Running 1 1				Collection Time in ms (Cur/Min/Max): 144/142/1471				Encoding Time in ms (Cur/Min/Max): 0/0/1				Transport Time in ms (Cur/Min/Max): 0/0/0				<b>Streaming Time in ms (Cur/Min/Max): 149/147/23548</b>				Collection Statistics:				collection_id_dropped = 0				last_collection_id_dropped = 0				drop_count = 0				<p>Calculate the total streaming time.</p> <p>The total streaming time is the sum of the individual current streaming times of each sensor group. Individual streaming times are displayed in Streaming time in ms (Cur). In this example, total streaming time is 2.664 seconds (2515 milliseconds plus 149 milliseconds).</p> <p>Compare the configured cadence to the total streaming time for the sensor group.</p> <p>The cadence is displayed in sample-interval. In this example, the cadence is correctly configured because the total streaming time (2.664 seconds) is less than the cadence (5.000 seconds, which is the default).</p>
Row ID	Sensor Group ID	Sensor Group type	SubID																																																																																			
1	2	Timer /YANG	5000																																																																																			
/Running 1 1																																																																																						
Collection Time in ms (Cur/Min/Max): 2444/2294/2460																																																																																						
Encoding Time in ms (Cur/Min/Max): 56/55/57																																																																																						
Transport Time in ms (Cur/Min/Max): 0/0/1																																																																																						
<b>Streaming Time in ms (Cur/Min/Max): 2515/2356/28403</b>																																																																																						
Collection Statistics:																																																																																						
collection_id_dropped = 0																																																																																						
last_collection_id_dropped = 0																																																																																						
drop_count = 0																																																																																						
2	1	Timer /YANG	5000																																																																																			
/Running 1 1																																																																																						
Collection Time in ms (Cur/Min/Max): 144/142/1471																																																																																						
Encoding Time in ms (Cur/Min/Max): 0/0/1																																																																																						
Transport Time in ms (Cur/Min/Max): 0/0/0																																																																																						
<b>Streaming Time in ms (Cur/Min/Max): 149/147/23548</b>																																																																																						
Collection Statistics:																																																																																						
collection_id_dropped = 0																																																																																						
last_collection_id_dropped = 0																																																																																						
drop_count = 0																																																																																						



	Command or Action	Purpose
	<pre>encoding JSON sensor-group 1   data-source YANG   path /Cisco-NX-OS-device:System/procsys-items depth unbounded sensor-group 2   data-source YANG   path /Cisco-NX-OS-device:System/intf-items/phys-items depth unbounded subscription 1   dst-grp 1   snsr-grp 1 sample-interval 5000   <b>snsr-grp 2 sample-interval 5000</b></pre>	
<b>Step 2</b>	<p><b>sensor group number</b></p> <p><b>Example:</b></p> <pre>switch-1(config-telemetry)# <b>sensor group1</b></pre>	If the total streaming time is not less than the cadence, enter the sensor group for which you want to set the interval.
<b>Step 3</b>	<p><b>subscription number</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-sensor)# <b>subscription 100</b></pre>	Edit the subscription for the sensor group.
<b>Step 4</b>	<p><b>snsr-grp number sample-interval milliseconds</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-sub)# <b>snsr-grp number</b> <b>sample-interval 5000</b></pre>	<p>For the appropriate sensor group, set the sample interval to a value greater than the total streaming time.</p> <p>In this example, the sample interval is set to 5.000 seconds, which is valid because it is larger than the total streaming time of 2.664 seconds.</p>
<b>Step 5</b>	<p><b>show system resources</b></p> <p><b>Example:</b></p> <pre>switch-1# <b>show system resources</b> Load average:  1 minute: 0.38  5 minutes: 0.43                15 minutes: 0.43 Processes:    555 total, 3 running CPU states   :  <b>24.17% user,</b>  4.32% kernel, 71.50% idle   CPU0 states:  0.00% user,  2.12% kernel, 97.87% idle   CPU1 states: 86.00% user, 11.00% kernel, 3.00% idle   CPU2 states:  8.08% user,  3.03% kernel, 88.88% idle   CPU3 states:  0.00% user,  1.02% kernel, 98.97% idle Memory usage: 16400084K total, 5861652K used, 10538432K free Current memory status: OK</pre>	<p>Check the CPU usage.</p> <p>If the CPU user state shows high usage, as shown in this example, your cadence and streaming value are not configured correctly. Repeat this procedure to properly configure the cadence.</p>

## 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 that is 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
```

## 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 stats

This command displays the statistics about the internal databases about configuration of telemetry.

```
switch# show telemetry control stats
show telemetry control stats entered
```

```
-----
Error Description                                     Error Count
-----
```

```
Chunk allocation failures                               0
Sensor path Database chunk creation failures           0
Sensor Group Database chunk creation failures          0
Destination Database chunk creation failures           0
Destination Group Database chunk creation failures     0
Subscription Database chunk creation failures          0
Sensor path Database creation failures                 0
Sensor Group Database creation failures                0
Destination Database creation failures                 0
Destination Group Database creation failures           0
Subscription Database creation failures                0
Sensor path Database insert failures                  0
Sensor Group Database insert failures                 0
Destination Database insert failures                  0
Destination Group Database insert failures             0
Subscription insert to Subscription Database failures 0
Sensor path Database delete failures                  0
Sensor Group Database delete failures                 0
Destination Database delete failures                  0
Destination Group Database delete failures             0
Delete Subscription from Subscription Database failures 0
Sensor path delete in use                             0
Sensor Group delete in use                            0
Destination delete in use                             0
Destination Group delete in use                       0
Delete destination(in use) failure count               0
Failed to get encode callback                         0
Sensor path Sensor Group list creation failures       0
Sensor path prop list creation failures                0
Sensor path sec Sensor path list creation failures    0
Sensor path sec Sensor Group list creation failures   0
Sensor Group Sensor path list creation failures       0
Sensor Group Sensor subs list creation failures       0
Destination Group subs list creation failures         0
Destination Group Destinations list creation failures 0
```

```

Destination Destination Groups list creation failures      0
Subscription Sensor Group list creation failures          0
Subscription Destination Groups list creation failures    0
Sensor Group Sensor path list delete failures             0
Sensor Group Subscriptions list delete failures           0
Destination Group Subscriptions list delete failures      0
Destination Group Destinations list delete failures      0
Subscription Sensor Groups list delete failures          0
Subscription Destination Groups list delete failures      0
Destination Destination Groups list delete failures      0
Failed to delete Destination from Destination Group      0
Failed to delete Destination Group from Subscription     0
Failed to delete Sensor Group from Subscription          0
Failed to delete Sensor path from Sensor Group           0
Failed to get encode callback                            0
Failed to get transport callback                         0
switch# Destination Database size = 1

```

```

-----
Dst IP Addr      Dst Port   Encoding   Transport   Count
-----
192.168.20.123 50001      GPB        gRPC        1

```

### show telemetry data collector brief

This command displays the brief statistics about the data collection.

```
switch# show telemetry data collector brief
```

```

-----
Collector Type      Successful Collections   Failed Collections
-----
DME                  143                       0

```

### show telemetry data collector details

This command displays detailed statistics about the data collection which includes breakdown of all sensor paths.

```
switch# show telemetry data collector details
```

```

-----
Succ Collections   Failed Collections   Sensor Path
-----
150                0                    sys/fm

```

### show telemetry event collector errors

This command displays the errors statistic about the event collection.

```
switch# show telemetry event collector errors
```

```

-----
Error Description                                     Error Count
-----
APIC-Cookie Generation Failures                       - 0
Authentication Failures                               - 0

```

```

Authentication Refresh Failures           - 0
Authentication Refresh Timer Start Failures - 0
Connection Timer Start Failures           - 0
Connection Attempts                       - 3
Dme Event Subscription Init Failures      - 0
Event Data Enqueue Failures               - 0
Event Subscription Failures               - 0
Event Subscription Refresh Failures       - 0
Pending Subscription List Create Failures - 0
Subscription Hash Table Create Failures   - 0
Subscription Hash Table Destroy Failures  - 0
Subscription Hash Table Insert Failures   - 0
Subscription Hash Table Remove Failures   - 0
Subscription Refresh Timer Start Failures - 0
Websocket Connect Failures                - 0

```

### show telemetry event collector stats

This command displays the statistics about the event collection which includes breakdown of all sensor paths.

```
switch# show telemetry event collector stats
```

```

-----
Collection Count  Latest Collection Time  Sensor Path
-----

```

### show telemetry control pipeline stats

This command displays the statistics for the telemetry pipeline.

```
switch# show telemetry pipeline stats
```

```
Main Statistics:
```

```
  Timers:
```

```
    Errors:
```

```
      Start Fail      =      0
```

```
  Data Collector:
```

```
    Errors:
```

```
      Node Create Fail =      0
```

```
  Event Collector:
```

```
    Errors:
```

```
      Node Create Fail =      0   Node Add Fail      =      0
```

```
      Invalid Data    =      0
```

```
Queue Statistics:
```

```
  Request Queue:
```

```
    High Priority Queue:
```

```
      Info:
```

```
        Actual Size    =    50   Current Size    =    0
```

```
        Max Size      =    0     Full Count     =    0
```

```
      Errors:
```

```
        Enqueue Error  =    0     Dequeue Error  =    0
```

```
    Low Priority Queue:
```

```
      Info:
```

```
        Actual Size    =    50   Current Size    =    0
```



```

          Max Size           = 0      Full Count       = 0
Errors:
  Enqueue Error             = 0      Dequeue Error    = 0
Data Queue:
  High Priority Queue:
    Info:
      Actual Size           = 50     Current Size     = 0
      Max Size              = 0      Full Count       = 0
    Errors:
      Enqueue Error         = 0      Dequeue Error    = 0
  Low Priority Queue:
    Info:
      Actual Size           = 50     Current Size     = 0
      Max Size              = 0      Full Count       = 0
    Errors:
      Enqueue Error         = 0      Dequeue Error    = 0

```

### show telemetry transport

This command displays all configured transport sessions.

```

switch# show telemetry transport

Session Id      IP Address      Port      Encoding  Transport  Status
-----
0               192.168.20.123 50001     GPB       gRPC       Connected

```

### show telemetry transport <session-id>

This command displays detailed session information for a specific transport session.

```

switch# show telemetry transport 0

Session Id:      0
IP Address:Port  192.168.20.123:50001
Encoding:        GPB
Transport:       gRPC
Status:          Disconnected
Last Connected:  Fri Sep 02 11:45:57.505 UTC

Tx Error Count:  224
Last Tx Error:   Fri Sep 02 12:23:49.555 UTC

switch# show telemetry transport 1

Session Id:      1
IP Address:Port  10.30.218.56:51235 Encoding:        JSON
Transport:       HTTP
Status:          Disconnected
Last Connected:  Never

Tx Error Count:  3
Last Tx Error:   Wed Apr 19 15:56:51.617 PDT

```

**show telemetry transport <session-id> stats**

This command displays details of a specific transport session.

```
switch# show telemetry transport 0 stats

Session Id:          0
IP Address:Port     192.168.20.123:50001
Encoding:           GPB
Transport:          GRPC
Status:             Connected
Last Connected:     Mon May 01 11:29:46.912 PST
Last Disconnected:  Never
Tx Error Count:     0
Last Tx Error:      None
```

**show telemetry transport <session-id> errors**

This command displays detailed error statistics for a specific transport session.

```
switch# show telemetry transport 0 errors

Session Id:          0
Connection Stats
  Connection Count   1
  Last Connected:   Mon May 01 11:29:46.912 PST
  Disconnect Count   0
  Last Disconnected: Never
Transmission Stats
  Transmit Count:   1225
  Last TX time:     Tue May 02 11:40:03.531 PST
  Min Tx Time:      7 ms
  Max Tx Time:      1760 ms
  Avg Tx Time:      500 ms
```

## 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
```

# 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 the section "Telemetry Model in the DME." Following are the main objects to be configured:

- **fmEntity** — Contains the NX-API and Telemetry feature states.
  - **fmNxapi** — Contains the NX-API state.
  - **fmTelemetry** — Contains the Telemetry feature state.
- **telemetryEntity** — Contains the telemetry feature configuration.
  - **telemetrySensorGroup** — Contains the definitions of one or more sensor paths or nodes to be monitored for telemetry. The telemetry entity can contain one or more sensor groups.
    - **telemetryRtSensorGroupRel** — Associates the sensor group with a telemetry subscription.
    - **telemetrySensorPath** — A path to be monitored. The sensor group can contain multiple objects of this type.
  - **telemetryDestGroup** — Contains the definitions of one or more destinations to receive telemetry data. The telemetry entity can contain one or more destination groups.
    - **telemetryRtDestGroupRel** — Associates the destination group with a telemetry subscription.
    - **telemetryDest** — A destination address. The destination group can contain multiple objects of this type.
  - **telemetrySubscription** — Specifies how and when the telemetry data from one or more sensor groups is sent to one or more destination groups.
    - **telemetryRsDestGroupRel** — Associates the telemetry subscription with a destination group.
    - **telemetryRsSensorGroupRel** — Associates the telemetry subscription with a sensor group.

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
```

## SUMMARY STEPS

1. Enable the telemetry feature.
2. Create the root level of the JSON payload to describe the telemetry configuration.
3. Create a sensor group to contain the defined sensor paths.
4. Define a telemetry destination group.
5. Define a telemetry destination profile.
6. Define one or more telemetry destinations, consisting of an IP address and port number to which telemetry data will be sent.
7. Create a telemetry subscription to configure the telemetry behavior.
8. Add the sensor group object as a child object to the **telemetrySubscription** element under the root element (**telemetryEntity**).
9. 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.
10. Define one or more sensor paths or nodes to be monitored for telemetry.
11. Add sensor paths as child objects to the sensor group object (**telemetrySensorGroup**).
12. Add destinations as child objects to the destination group object (**telemetryDestGroup**).
13. Add the destination group object as a child object to the root element (**telemetryEntity**).
14. Create a relation object as a child object of the telemetry sensor group to associate the sensor group to the subscription.
15. Create a relation object as a child object of the telemetry destination group to associate the destination group to the subscription.
16. Create a relation object as a child object of the subscription to associate the subscription to the telemetry destination group.
17. Send the resulting JSON structure as an HTTP/HTTPS POST payload to the NX-API endpoint for telemetry configuration.

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	Enable the telemetry feature. <b>Example:</b> <pre>{   "fmEntity" : {     "children" : [{       "fmTelemetry" : {         "attributes" : {           "adminSt" : "enabled"         }       }     ]   } }</pre>	The root element is <b>fmTelemetry</b> and the base path for this element is <code>sys/fm</code> . Configure the <b>adminSt</b> attribute as <code>enabled</code> .

	Command or Action	Purpose
<b>Step 2</b>	<p>Create the root level of the JSON payload to describe the telemetry configuration.</p> <p><b>Example:</b></p> <pre>{   "telemetryEntity": {     "attributes": {       "dn": "sys/tm"     },   } }</pre>	<p>The root element is <b>telemetryEntity</b> and the base path for this element is <code>sys/tm</code>. Configure the <b>dn</b> attribute as <code>sys/tm</code>.</p>
<b>Step 3</b>	<p>Create a sensor group to contain the defined sensor paths.</p> <p><b>Example:</b></p> <pre>"telemetrySensorGroup": {   "attributes": {     "id": "10",     "rn": "sensor-10"   },   "children": [{   }] }</pre>	<p>A telemetry sensor group is defined in an object of class <b>telemetrySensorGroup</b>. Configure the following attributes of the object:</p> <ul style="list-style-type: none"> <li>• <b>id</b> — An identifier for the sensor group. Currently only numeric ID values are supported.</li> <li>• <b>rn</b> — The relative name of the sensor group object in the format: <b>sensor-id</b>.</li> </ul> <p>Children of the sensor group object include sensor paths and one or more relation objects (<b>telemetryRtSensorGroupRel</b>) to associate the sensor group with a telemetry subscription.</p>
<b>Step 4</b>	<p>Define a telemetry destination group.</p> <p><b>Example:</b></p> <pre>{   "telemetryDestGroup": {     "attributes": {       "id": "20"     }   } }</pre>	<p>A telemetry destination group is defined in <b>telemetryEntity</b>. Configure the <b>id</b> attribute.</p>
<b>Step 5</b>	<p>Define a telemetry destination profile.</p> <p><b>Example:</b></p> <pre>{   "telemetryDestProfile": {     "attributes": {       "adminSt": "enabled"     },     "children": [       {         "telemetryDestOptSourceInterface": {           "attributes": {             "name": "lo0"           }         }       }     ]   } }</pre>	<p>A telemetry destination profile is defined in <b>telemetryDestProfile</b>.</p> <ul style="list-style-type: none"> <li>• Configure the <b>adminSt</b> attribute as <code>enabled</code>.</li> <li>• Under <b>telemetryDestOptSourceInterface</b>, configure the <b>name</b> attribute with an interface name to stream data from the configured interface to a destination with the source IP address.</li> </ul>

	Command or Action	Purpose
	<pre> } } </pre>	
<b>Step 6</b>	<p>Define one or more telemetry destinations, consisting of an IP address and port number to which telemetry data will be sent.</p> <p><b>Example:</b></p> <pre> {   "telemetryDest": {     "attributes": {       "addr": "1.2.3.4",       "enc": "GPB",       "port": "50001",       "proto": "gRPC",       "rn": "addr-[1.2.3.4]-port-50001"     }   } } </pre>	<p>A telemetry destination is defined in an object of class <b>telemetryDest</b>. Configure the following attributes of the object:</p> <ul style="list-style-type: none"> <li>• <b>addr</b> — The IP address of the destination.</li> <li>• <b>port</b> — The port number of the destination.</li> <li>• <b>rn</b> — The relative name of the destination object in the format: <b>path-[path]</b>.</li> <li>• <b>enc</b> — The encoding type of the telemetry data to be sent. NX-OS supports: <ul style="list-style-type: none"> <li>• Google protocol buffers (GPB) for gRPC.</li> <li>• JSON for C.</li> </ul> </li> <li>• <b>proto</b> — The transport protocol type of the telemetry data to be sent. NX-OS supports: <ul style="list-style-type: none"> <li>• gRPC</li> <li>• HTTP</li> </ul> </li> <li>• Supported encoded types are: <ul style="list-style-type: none"> <li>• HTTP/JSON YES</li> <li>• HTTP/Form-data YES Only supported for Bin Logging.</li> <li>• GRPC/GPB-Compact YES Native Data Source Only.</li> <li>• GRPC/GPB YES</li> <li>• UDP/GPB YES</li> <li>• UDP/JSON YES</li> </ul> </li> </ul>
<b>Step 7</b>	<p>Create a telemetry subscription to configure the telemetry behavior.</p> <p><b>Example:</b></p> <pre> "telemetrySubscription": {   "attributes": {     "id": "30",     "rn": "subs-30"   },   "children": [{   }] } </pre>	<p>A telemetry subscription is defined in an object of class <b>telemetrySubscription</b>. Configure the following attributes of the object:</p> <ul style="list-style-type: none"> <li>• <b>id</b> — An identifier for the subscription. Currently only numeric ID values are supported.</li> <li>• <b>rn</b> — The relative name of the subscription object in the format: <b>subs-id</b>.</li> </ul>

	Command or Action	Purpose
		Children of the subscription object include relation objects for sensor groups ( <b>telemetryRsSensorGroupRel</b> ) and destination groups ( <b>telemetryRsDestGroupRel</b> ).
<b>Step 8</b>	<p>Add the sensor group object as a child object to the <b>telemetrySubscription</b> element under the root element (<b>telemetryEntity</b>).</p> <p><b>Example:</b></p> <pre> {   "telemetrySubscription": {     "attributes": {       "id": "30"     }     "children": [{       "telemetryRsSensorGroupRel": {         "attributes": {           "sampleIntvl": "5000",           "tDn": "sys/tm/sensor-10"         }       }     ]   } } </pre>	
<b>Step 9</b>	<p>Create a relation object as a child object of the subscription to associate the subscription to the telemetry sensor group and to specify the data sampling behavior.</p> <p><b>Example:</b></p> <pre> "telemetryRsSensorGroupRel": {   "attributes": {     "rType": "mo",     "rn": "rssensorGroupRel-[sys/tm/sensor-10]",     "sampleIntvl": "5000",     "tCl": "telemetrySensorGroup",     "tDn": "sys/tm/sensor-10",     "tType": "mo"   } } </pre>	<p>The relation object is of class <b>telemetryRsSensorGroupRel</b> and is a child object of <b>telemetrySubscription</b>. Configure the following attributes of the relation object:</p> <ul style="list-style-type: none"> <li>• <b>rn</b> — The relative name of the relation object in the format: <b>rssensorGroupRel-[sys/tm/sensor-group-id]</b>.</li> <li>• <b>sampleIntvl</b> — The data sampling period in milliseconds. An interval value of 0 creates an event-based subscription, in which telemetry data is sent only upon changes under the specified MO. An interval value greater than 0 creates a frequency-based subscription, in which telemetry data is sent periodically at the specified interval. For example, an interval value of 15000 results in the sending of telemetry data every 15 seconds.</li> <li>• <b>tCl</b> — The class of the target (sensor group) object, which is <b>telemetrySensorGroup</b>.</li> <li>• <b>tDn</b> — The distinguished name of the target (sensor group) object, which is <b>sys/tm/sensor-group-id</b>.</li> <li>• <b>rType</b> — The relation type, which is <b>mo</b> for managed object.</li> <li>• <b>tType</b> — The target type, which is <b>mo</b> for managed object.</li> </ul>

	Command or Action	Purpose
<p><b>Step 10</b></p>	<p>Define one or more sensor paths or nodes to be monitored for telemetry.</p> <p><b>Example:</b></p> <p>Single sensor path</p> <pre data-bbox="251 478 698 829"> {   "telemetrySensorPath": {     "attributes": {       "path": "sys/cdp",       "rn": "path-[sys/cdp]",       "excludeFilter": "",       "filterCondition": "",       "path": "sys/fm/bgp",       "secondaryGroup": "0",       "secondaryPath": "",       "depth": "0"     }   } } </pre> <p><b>Example:</b></p> <p>Multiple sensor paths</p> <pre data-bbox="251 993 698 1648"> {   "telemetrySensorPath": {     "attributes": {       "path": "sys/cdp",       "rn": "path-[sys/cdp]",       "excludeFilter": "",       "filterCondition": "",       "path": "sys/fm/bgp",       "secondaryGroup": "0",       "secondaryPath": "",       "depth": "0"     }   } }, {   "telemetrySensorPath": {     "attributes": {       "excludeFilter": "",       "filterCondition": "",       "path": "sys/fm/dhcp",       "secondaryGroup": "0",       "secondaryPath": "",       "depth": "0"     }   } } </pre> <p><b>Example:</b></p> <p>Single sensor path filtering for BGP disable events:</p> <pre data-bbox="251 1812 609 1850"> {   "telemetrySensorPath": { </pre>	<p>A sensor path is defined in an object of class <b>telemetrySensorPath</b>. Configure the following attributes of the object:</p> <ul style="list-style-type: none"> <li>• <b>path</b> — The path to be monitored.</li> <li>• <b>rn</b> — The relative name of the path object in the format: <b>path-[path]</b></li> <li>• <b>depth</b> — The retrieval level for the sensor path. A depth setting of <b>0</b> retrieves only the root MO properties.</li> <li>• <b>filterCondition</b> — (Optional) Creates a specific filter for event-based subscriptions. The DME provides the filter expressions. For more information about filtering, see the Cisco APIC REST API Usage Guidelines on composing queries: <a href="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">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</a></li> </ul>



	Command or Action	Purpose
	<pre>       "attributes": {         "path": "sys/cdp",         "rn": "path-[sys/cdp]",         "excludeFilter": "",         "filterCondition": "eq(fmBgp.operSt.\"disabled\")",         "path": "sys/fm/bgp",         "secondaryGroup": "0",         "secondaryPath": "",         "depth": "0"       }     }   } </pre>	
<b>Step 11</b>	Add sensor paths as child objects to the sensor group object ( <b>telemetrySensorGroup</b> ).	
<b>Step 12</b>	Add destinations as child objects to the destination group object ( <b>telemetryDestGroup</b> ).	
<b>Step 13</b>	Add the destination group object as a child object to the root element ( <b>telemetryEntity</b> ).	
<b>Step 14</b>	<p>Create a relation object as a child object of the telemetry sensor group to associate the sensor group to the subscription.</p> <p><b>Example:</b></p> <pre> "telemetryRtSensorGroupRel": {   "attributes": {     "rn": "rtsensorGroupRel-[sys/tm/subs-30]",      "tCl": "telemetrySubscription",     "tDn": "sys/tm/subs-30"   } } </pre>	<p>The relation object is of class <b>telemetryRtSensorGroupRel</b> and is a child object of <b>telemetrySensorGroup</b>. Configure the following attributes of the relation object:</p> <ul style="list-style-type: none"> <li>• <b>rn</b> — The relative name of the relation object in the format: <b>rtsensorGroupRel-[sys/tm/subscription-id]</b>.</li> <li>• <b>tCl</b> — The target class of the subscription object, which is <b>telemetrySubscription</b>.</li> <li>• <b>tDn</b> — The target distinguished name of the subscription object, which is <b>sys/tm/subscription-id</b>.</li> </ul>
<b>Step 15</b>	<p>Create a relation object as a child object of the telemetry destination group to associate the destination group to the subscription.</p> <p><b>Example:</b></p> <pre> "telemetryRtDestGroupRel": {   "attributes": {     "rn": "rtdestGroupRel-[sys/tm/subs-30]",     "tCl": "telemetrySubscription",     "tDn": "sys/tm/subs-30"   } } </pre>	<p>The relation object is of class <b>telemetryRtDestGroupRel</b> and is a child object of <b>telemetryDestGroup</b>. Configure the following attributes of the relation object:</p> <ul style="list-style-type: none"> <li>• <b>rn</b> — The relative name of the relation object in the format: <b>rtdestGroupRel-[sys/tm/subscription-id]</b>.</li> <li>• <b>tCl</b> — The target class of the subscription object, which is <b>telemetrySubscription</b>.</li> <li>• <b>tDn</b> — The target distinguished name of the subscription object, which is <b>sys/tm/subscription-id</b>.</li> </ul>
<b>Step 16</b>	<p>Create a relation object as a child object of the subscription to associate the subscription to the telemetry destination group.</p> <p><b>Example:</b></p>	<p>The relation object is of class <b>telemetryRsDestGroupRel</b> and is a child object of <b>telemetrySubscription</b>. Configure the following attributes of the relation object:</p>

	Command or Action	Purpose
	<pre> "telemetryRsDestGroupRel": {   "attributes": {     "rType": "mo",     "rn": "rsdestGroupRel-[sys/tm/dest-20]",     "tCl": "telemetryDestGroup",     "tDn": "sys/tm/dest-20",     "tType": "mo"   } } </pre>	<ul style="list-style-type: none"> <li>• <b>rn</b> — The relative name of the relation object in the format: <b>rsdestGroupRel-[sys/tm/destination-group-id]</b>.</li> <li>• <b>tCl</b> — The class of the target (destination group) object, which is <b>telemetryDestGroup</b>.</li> <li>• <b>tDn</b> — The distinguished name of the target (destination group) object, which is <b>sys/tm/destination-group-id</b>.</li> <li>• <b>rType</b> — The relation type, which is <b>mo</b> for managed object.</li> <li>• <b>tType</b> — The target type, which is <b>mo</b> for managed object.</li> </ul>
<b>Step 17</b>	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 <code>sys/tm</code> and the NX-API endpoint is:  <code>{{URL}}/api/node/mo/sys/tm.json</code>

### Example

The following is an example of all the previous steps that are collected into one POST payload (note that some attributes may not match):

```

{
  "telemetryEntity": {
    "children": [{
      "telemetrySensorGroup": {
        "attributes": {
          "id": "10"
        }
      }
    ]
  },
  "telemetryDestGroup": {
    "attributes": {
      "id": "20"
    }
  }
}

```

```

        "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"
        }
      }
    ]
  },
  "telemetrySensorPath": {
    "attributes": {
      "excludeFilter": "",
      "filterCondition": "eq(fmBgp.operSt,\"disabled\")",
      "path": "sys/fm/bgp",
      "secondaryGroup": "0",
      "secondaryPath": "",
      "depth": "0"
    }
  }
}

```



```

|----mo [name:SensorGroup]
|   |   @name:SensorGroup
|   |   @label:Sensor Group
|   |--property
|   |   @name:id [key]
|   |   @type:string:Basic
|   |
|   |----mo [name:SensorPath]
|   |   |   @name:SensorPath
|   |   |   @label:Sensor Path
|   |   |--property
|   |   |   @name:path [key]
|   |   |   @type:string:Basic
|   |   |   @name:filterCondition
|   |   |   @type:string:Basic
|   |   |   @name:excludeFilter
|   |   |   @type:string:Basic
|   |   |   @name:depth
|   |   |   @type:RetrieveDepth
|   |
|----mo [name:DestGroup]
|   |   @name:DestGroup
|   |   @label:Destination Group
|   |--property
|   |   @name:id
|   |   @type:string:Basic
|   |
|   |----mo [name:Dest]
|   |   |   @name:Dest
|   |   |   @label:Destination
|   |   |--property
|   |   |   @name:addr [key]
|   |   |   @type:address:Ip
|   |   |   @name:port [key]
|   |   |   @type:scalar:Uint16
|   |   |   @name:proto
|   |   |   @type:Protocol
|   |   |   @name:enc
|   |   |   @type:Encoding
|   |
|----mo [name:Subscription]
|   |   @name:Subscription
|   |   @label:Subscription
|   |--property
|   |   @name:id
|   |   @type:scalar:Uint64
|   |----reldef
|   |   |   @name:SensorGroupRel
|   |   |   @to:SensorGroup
|   |   |   @cardinality:ntom
|   |   |   @label:Link to sensorGroup entry
|   |   |--property
|   |   |   @name:sampleIntvl
|   |   |   @type:scalar:Uint64
|   |
|   |----reldef
|   |   |   @name:DestGroupRel
|   |   |   @to:DestGroup
|   |   |   @cardinality:ntom
|   |   |   @label:Link to destGroup entry

```

# 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 208](#).



---

**Note** The model does not support counters for subinterface, loopback, or VLAN, so they are not streamed out.

---

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **path interface**
5. **destination-group** *grp\_id*
6. **ip address** *ip\_addr* **port** *port*
7. **subscription** *sub\_id*
8. **snsr-group** *sgrp\_id* **sample-interval** *interval*
9. **dst-group** *dgrp\_id*

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
<b>Step 2</b>	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
<b>Step 3</b>	<b>sensor-group <i>sgrp_id</i></b> <b>Example:</b> <pre>switch-1(config-telemetry)# sensor-group 6 switch-1(conf-tm-sensor)#</pre>	Create a sensor group for telemetry data.
<b>Step 4</b>	<b>path interface</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# path interface switch-1(conf-tm-sensor)#</pre>	<p>Configure the interface path label, which enables sending one telemetry data query for multiple individual interfaces. The label consolidates the queries for multiple interfaces into one. Telemetry then gathers the data and returns it to the label.</p> <p>Depending on how the polling interval is configured, interface data is sent based on a periodic basis or whenever the interface state changes.</p>
<b>Step 5</b>	<b>destination-group <i>grp_id</i></b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
<b>Step 6</b>	<b>ip address <i>ip_addr</i> port <i>port</i></b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port.
<b>Step 7</b>	<b>subscription <i>sub_id</i></b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
<b>Step 8</b>	<b>snsr-group <i>sgrp_id</i> sample-interval <i>interval</i></b> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur.

	Command or Action	Purpose
Step 9	<b>dst-group</b> <i>dgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# <b>dst-grp</b> 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> command.

## Configuring the Interface Path for Non-Zero Counters

You can configure the interface path label with a predefined keyword filter that returns only counters that have nonzero values. The filter is `counters=[detailed]`.

By using this filter, the interface path gathers all the available interface counters, filters the collected data, then forwards the results to the receiver. The filter is optional, and if you do not use it, all counters, including zero-value counters, are displayed for the interface path.



**Note** Using the filter is conceptually similar to issuing **show interface mgmt0 counters detailed**

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **path interface query-condition** `counters=[detailed]`
5. **destination-group** *grp\_id*
6. **ip address** *ip\_addr* **port** *port*
7. **subscription** *sub\_id*
8. **snsr-group** *sgrp\_id* **sample-interval** *interval*
9. **dst-group** *dgrp\_id*

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# <b>configure terminal</b> switch-1(config)#</pre>	Enter configuration mode.
Step 2	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# <b>telemetry</b> switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
Step 3	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b>	Create a sensor group for telemetry data.

	Command or Action	Purpose
	<pre>switch-1(config-telemetry)# sensor-group 6 switch-1(conf-tm-sensor)#</pre>	
<b>Step 4</b>	<p><b>path interface query-condition counters=[detailed]</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-sensor)# path interface query-condition counters=[detailed] switch-1(conf-tm-sensor)#</pre>	Configure the interface path label and query for only the nonzero counters from all interfaces.
<b>Step 5</b>	<p><b>destination-group grp_id</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
<b>Step 6</b>	<p><b>ip address ip_addr port port</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port.
<b>Step 7</b>	<p><b>subscription sub_id</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
<b>Step 8</b>	<p><b>snsr-group sgrp_id sample-interval interval</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur.
<b>Step 9</b>	<p><b>dst-group dgrp_id</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-sub)# dst-grp 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> command.

## Configuring the Interface Path for Operational Speeds

You can configure the interface path label with a pre-defined keyword filter that returns counters for interfaces of specified operational speeds. The filter is `oper-speed=[]`. The following operational speeds are supported: auto, 10M, 100M, 1G, 10G, 40G, 200G, and 400G.

By using this filter, the interface path gathers the telemetry data for interfaces of the specified speed, then forwards the results to the receiver. The filter is optional. If you do not use it, counters for all interfaces are displayed, regardless of their operational speed.

The filter can accept multiple speeds as a comma-separated list, for example `oper-speed=[1G,10G]` to retrieve counters for interfaces that operate at 1 and 10 Gbps. Do not use a blank space as a delimiter.



**Note** Interface types subinterface, loopback, and VLAN do not have operational speed properties, so the filter does not support these interface types.

## SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **snsr-group** *sgrp\_id* **sample-interval** *interval*
4. **path interface query-condition oper-speed**=[*speed*]
5. **destination-group** *grp\_id*
6. **ip address** *ip\_addr* **port** *port*
7. **subscription** *sub\_id*
8. **snsr-group** *sgrp\_id* **sample-interval** *interval*
9. **dst-group** *dgrp\_id*

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
Step 2	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
Step 3	<b>snsr-group</b> <i>sgrp_id</i> <b>sample-interval</b> <i>interval</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur.
Step 4	<b>path interface query-condition oper-speed</b> =[ <i>speed</i> ] <b>Example:</b> <pre>switch-1(conf-tm-sensor)# path interface query-condition oper-speed=[1G,40G] switch-1(conf-tm-sensor)#</pre>	Configure the interface path label and query for counters from interfaces running the specified speed, which in this example, is 1 and 40 Gbps only.
Step 5	<b>destination-group</b> <i>grp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.

	Command or Action	Purpose
<b>Step 6</b>	<b>ip address</b> <i>ip_addr</i> <b>port</b> <i>port</i> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port.
<b>Step 7</b>	<b>subscription</b> <i>sub_id</i> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
<b>Step 8</b>	<b>snsr-group</b> <i>sgrp_id</i> <b>sample-interval</b> <i>interval</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur.
<b>Step 9</b>	<b>dst-group</b> <i>dgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# dst-grp 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> 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 the more filters you add, the more focused the results become.

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **path interface query-condition** counters=[detailed],oper-speed=[1G,40G]
5. **destination-group** *grp\_id*
6. **ip address** *ip\_addr* **port** *port*
7. **subscription** *sub\_id*
8. **snsr-group** *sgrp\_id* **sample-interval** *interval*
9. **dst-group** *dgrp\_id*

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
Step 2	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
Step 3	<b>sensor-group <i>sgrp_id</i></b> <b>Example:</b> <pre>switch-1(config-telemetry)# sensor-group 6 switch-1(conf-tm-sensor)#</pre>	Create a sensor group for telemetry data.
Step 4	<b>path interface query-condition</b> <b>counters=[detailed],oper-speed=[1G,40G]</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# path interface query-condition counters=[detailed],oper-speed=[1G,40G] switch-1(conf-tm-sensor)#</pre>	Configures multiple conditions in the same query. In this example, the query does both of the following: <ul style="list-style-type: none"> <li>• Gathers and returns non-zero counters on interfaces running at 1 Gbps.</li> <li>• Gathers and returns non-zero counters on interfaces running at 40 Gbps.</li> </ul>
Step 5	<b>destination-group <i>grp_id</i></b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
Step 6	<b>ip address <i>ip_addr</i> port <i>port</i></b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port.
Step 7	<b>subscription <i>sub_id</i></b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	<b>snsr-group <i>sgrp_id</i> sample-interval <i>interval</i></b> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur.

	Command or Action	Purpose
<b>Step 9</b>	<b>dst-group</b> <i>dgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# <b>dst-grp</b> 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> 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 208](#).

You can set the resources path to return system resource information through either periodic polling or based on events. This path does not support filtering.

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **path environment**
5. **destination-group** *grp\_id*
6. **ip address** *ip\_addr* **port** *port*
7. **subscription** *sub\_id*
8. **snsr-group** *sgrp\_id* **sample-interval** *interval*
9. **dst-group** *dgrp\_id*

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# <b>configure terminal</b> switch-1(config)#</pre>	Enter configuration mode.
<b>Step 2</b>	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# <b>telemetry</b> switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
<b>Step 3</b>	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> <pre>switch-1(config-telemetry)# <b>sensor-group</b> 6 switch-1(conf-tm-sensor)#</pre>	Create a sensor group for telemetry data.



	Command or Action	Purpose
Step 4	<b>path environment</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# path environment switch-1(conf-tm-sensor)#</pre>	<p>Configures the environment path label, which enables telemetry data for multiple individual environment objects to be sent to the label. The label consolidates the multiple data inputs into one output.</p> <p>Depending on the sample interval, the environment data is either streaming based on the polling interval, or sent when events occur.</p>
Step 5	<b>destination-group grp_id</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
Step 6	<b>ip address ip_addr port port</b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port.
Step 7	<b>subscription sub_id</b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 8	<b>snsr-group sgrp_id sample-interval interval</b> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when environment events occur.
Step 9	<b>dst-group dgrp_id</b> <b>Example:</b> <pre>switch-1(conf-tm-sub)# dst-grp 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> command.

## Configuring the Resources Path for 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 208](#).

This path does not support filtering.

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**

3. **sensor-group** *sgrp\_id*
4. **path resources**
5. **destination-group** *grp\_id*
6. **ip address** *ip\_addr* **port** *port*
7. **subscription** *sub\_id*
8. **snsr-group** *sgrp\_id* **sample-interval** *interval*
9. **dst-group** *dgrp\_id*

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
<b>Step 2</b>	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
<b>Step 3</b>	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> <pre>switch-1(config-telemetry)# sensor-group 6 switch-1(conf-tm-sensor)#</pre>	Create a sensor group for telemetry data.
<b>Step 4</b>	<b>path resources</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# path resources switch-1(conf-tm-sensor)#</pre>	<p>Configure the resources path label, which enables telemetry data for multiple individual system resources to be sent to the label. The label consolidates the multiple data inputs into one output.</p> <p>Depending on the sample interval, the resource data is either streaming based on the polling interval, or sent when system memory changes to Not OK.</p>
<b>Step 5</b>	<b>destination-group</b> <i>grp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
<b>Step 6</b>	<b>ip address</b> <i>ip_addr</i> <b>port</b> <i>port</i> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port.
<b>Step 7</b>	<b>subscription</b> <i>sub_id</i> <b>Example:</b>	Enter telemetry subscription submode, and configure the telemetry subscription.

	Command or Action	Purpose
	switch-1 (conf-tm-dest) # <b>subscription 33</b> switch-1 (conf-tm-sub) #	
<b>Step 8</b>	<b>snsr-group</b> <i>sgrp_id</i> <b>sample-interval</b> <i>interval</i>  <b>Example:</b> switch-1 (conf-tm-sub) # <b>snsr-grp 6 sample-interval 5000</b> switch-1 (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.
<b>Step 9</b>	<b>dst-group</b> <i>dgrp_id</i>  <b>Example:</b> switch-1 (conf-tm-sub) # <b>dst-grp 33</b> switch-1 (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 <b>destination-group</b> command.

## Configuring the VXLAN Path to Poll for Events or Data

The VXLAN path label provides information about the switch's Virtual Extensible LAN (VXLAN) 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 208](#).

This path does not support filtering.

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **vxlan environment**
5. **destination-group** *dgrp\_id*
6. **ip address** *ip\_addr* **port** *port*
7. **subscription** *sub\_id*
8. **snsr-group** *sgrp\_id* **sample-interval** *interval*
9. **dst-group** *dgrp\_id*

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch-1# <b>configure terminal</b> switch-1 (config) #	Enter configuration mode.
<b>Step 2</b>	<b>telemetry</b>  <b>Example:</b> switch-1 (config) # <b>telemetry</b> switch-1 (config-telemetry) #	Enter configuration mode for the telemetry features.

	Command or Action	Purpose
<b>Step 3</b>	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> <pre>switch-1(config-telemetry)# sensor-group 6 switch-1(conf-tm-sensor)#</pre>	Create a sensor group for telemetry data.
<b>Step 4</b>	<b>vxlan environment</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# vxlan environment switch-1(conf-tm-sensor)#</pre>	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.
<b>Step 5</b>	<b>destination-group</b> <i>grp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
<b>Step 6</b>	<b>ip address</b> <i>ip_addr port port</i> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 1.2.3.4 port 50004 switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port.
<b>Step 7</b>	<b>subscription</b> <i>sub_id</i> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
<b>Step 8</b>	<b>snsr-group</b> <i>sgrp_id sample-interval interval</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when VXLAN events occur.
<b>Step 9</b>	<b>dst-group</b> <i>dgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# dst-grp 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> command.

## Verifying the Path Label Configuration

At any time, you can verify that path labels are configured, and check their values by displaying the running telemetry configuration.

### SUMMARY STEPS

1. `show running-config-telemetry`

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<p><b>show running-config-telemetry</b></p> <p><b>Example:</b></p> <pre>switch-1(conf-tm-sensor)# show running-config telemetry  !Command: show running-config telemetry !Running configuration last done at: Mon Jun 10 08:10:17 2019 !Time: Mon Jun 10 08:10:17 2019  version 9.3(1) Bios:version feature telemetry  telemetry   destination-profile     use-nodeid tester   sensor-group 4     path interface query-condition and(counters=[detailed],oper-speed=[1G,10G])   sensor-group 6     path interface query-condition oper-speed=[1G,40G]   subscription 6     snsr-grp 6 sample-interval 6000 nxosv2(conf-tm-sensor)#</pre>	<p>Displays the current running config for telemetry,</p> <p>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.</p>

## 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
<b>show telemetry usability {all   environment   interface   resources   vxlan}</b>	<p>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.</p> <p>For the interfaces path label, also any keyword filters or query conditions you configured.</p>
<b>show running-config telemetry</b>	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-1# 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=egptPsuSlot, egptFtSlot, egptSupCSlot, egptPsu, egptFt, egptSensor, egptICSlot

  2) label_name      : environment

      path_name      : sys/ch
      query_type     : event
      query_condition :
      rsp-subtree-full&query-target= subtree&target-subtree-class=egptPsuSlot, egptFtSlot, egptSupCSlot, egptPsu, egptFt, egptSensor, egptICSlot
switch-1#
```

The following shows the output of the **show telemetry usability interface** command.

```
switch-1# show telemetry usability interface
  1) label_name      : interface

      path_name      : sys/intf
      query_type     : poll
      query_condition :
      query-target=child&query-target-filter=eq(l1PhysIf.adminSt, "up")&rsp-subtree=child&rsp-subtree-class=monEthStats, monIfIn, monIfOut, monIfCIn, monIfCOut

  2) label_name      : interface

      path_name      : sys/mgmt-[mgmt0]
      query_type     : poll
      query_condition :
      query-target=tree&query-target-filter=eq(mgmtIf.adminSt, "up")&rsp-subtree=full&rsp-subtree-class=monEthStats, monIfIn, monIfOut, monIfCIn, monIfCOut

  3) label_name      : interface

      path_name      : sys/intf
      query_type     : event
      query_condition :
      opState=update&eq(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 :
      query-target=tree&query-target-filter=or((l1Eth0, oas0), or(and(not(immIf.operSt), eq(immIf.operSt, "down")), and(not(immIf.operSt), eq(immIf.operSt, "up"))))
switch-1#
```

The following shows an example of the **show telemetry usability resources** command.



# 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	<p>Sends event notifications for create, update, and delete events. The following values are exported through telemetry for the RIB path:</p> <ul style="list-style-type: none"> <li>• Next-hop routing information: <ul style="list-style-type: none"> <li>• Address of the next hop</li> <li>• Outgoing interface for the next hop</li> <li>• VRF name for the next hop</li> <li>• Owner of the next hop</li> <li>• Preference for the next hop</li> <li>• Metric for the next hop</li> <li>• Tag for the next hop</li> <li>• Segment ID for the next hop</li> <li>• Tunnel ID for the next hop</li> <li>• Encapsulation type for the next hop</li> <li>• Bitwise OR of flags for the Next Hop Type</li> </ul> </li> <li>• For Layer-3 routing information: <ul style="list-style-type: none"> <li>• VRF name of the route</li> <li>• Route prefix address</li> <li>• Mask length for the route</li> <li>• Number of next hops for the route</li> <li>• Event type</li> <li>• Next hops</li> </ul> </li> </ul>

Path Type	Subscription Baseline	Event Notifications
MAC	Executes a <code>GETALL</code> from DME for static and dynamic MAC entries	<p>Sends event notifications for add, update, and delete events. The following values are exported through telemetry for the MAC path:</p> <ul style="list-style-type: none"> <li>• MAC address</li> <li>• MAC address type</li> <li>• VLAN number</li> <li>• Interface name</li> <li>• Event types</li> </ul> <p>Both static and dynamic entries are supported in event notifications.</p>
Adjacency	Sends the IPv4 and IPv6 adjacencies	<p>Sends event notifications for add, update, and delete events. The following values are exported through telemetry for the Adjacency path:</p> <ul style="list-style-type: none"> <li>• IP address</li> <li>• MAC address</li> <li>• Interface name</li> <li>• Physical interface name</li> <li>• VRF name</li> <li>• Preference</li> <li>• Source for the adjacency</li> <li>• Address family for the adjacency</li> <li>• Adjacency event type</li> </ul>

For additional information, refer to Github <https://github.com/CiscoDevNet/nx-telemetry-PROTO>.

## Guidelines and Limitations for Native Data Source Path

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 224](#).

### Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **data-source native**
5. **path rib query-condition** [**data=ephemeral** | **updates\_only**]
6. **destination-group** *grp\_id*
7. **ip address** *ip\_addr* **port** *port* **protocol** { **HTTP** | **gRPC** } **encoding** { **JSON** | **GPB** | **GPB-compact** }
8. **subscription** *sub\_id*
9. **snsr-group** *sgrp\_id* **sample-interval** *interval*
10. **dst-group** *dgrp\_id*

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
Step 2	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
Step 3	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# sensor-grp 6 switch-1(conf-tm-sub)#</pre>	Create a sensor group.
Step 4	<b>data-source native</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# data-source native switch-1(conf-tm-sensor)#</pre>	Set the data source to native so that any native application can use the streamed data without requiring a specific model or database.

	Command or Action	Purpose
<b>Step 5</b>	<b>path rib query-condition</b> [ <b>data=ephemeral</b>   <b>updates_only</b> ]  <b>Example:</b> <pre>nxosv2(conf-tm-sensor)# path rib nxosv2(conf-tm-sensor)#</pre>	Configure the RIB path which streams routes and route update information.
<b>Step 6</b>	<b>destination-group</b> <i>grp_id</i>  <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
<b>Step 7</b>	<b>ip address</b> <i>ip_addr</i> <i>port</i> <i>port</i> <b>protocol</b> { <b>HTTP</b>   <b>gRPC</b> } <b>encoding</b> { <b>JSON</b>   <b>GPB</b>   <b>GPB-compact</b> }  <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream.
<b>Step 8</b>	<b>subscription</b> <i>sub_id</i>  <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
<b>Step 9</b>	<b>snsr-group</b> <i>sgrp_id</i> <b>sample-interval</b> <i>interval</i>  <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when rib events occur.
<b>Step 10</b>	<b>dst-group</b> <i>dgrp_id</i>  <b>Example:</b> <pre>switch-1(conf-tm-sub)# dst-grp 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> 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 224.



**Note** For update or delete events, MAC notifications are sent only for the MAC addresses that have IP adjacencies.

### Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

## SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **data-source native**
5. **path mac**
6. **destination-group** *grp\_id*
7. **ip address** *ip\_addr* **port** *port* **protocol** { HTTP | gRPC } **encoding** { JSON | GPB | GPB-compact }
8. **subscription** *sub\_id*
9. **snsr-group** *sgrp\_id* **sample-interval** *interval*
10. **dst-group** *dgrp\_id*

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
Step 2	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
Step 3	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# sensor-grp 6 switch-1(conf-tm-sub)#</pre>	Create a sensor group.
Step 4	<b>data-source native</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# data-source native switch-1(conf-tm-sensor)#</pre>	Set the data source to native so that any native application can use the streamed data without requiring a specific model or database.

	Command or Action	Purpose
<b>Step 5</b>	<b>path mac</b> <b>Example:</b> <pre>nxosv2(conf-tm-sensor)# path mac nxosv2(conf-tm-sensor)#</pre>	Configure the MAC path which streams information about MAC entries and MAC notifications.
<b>Step 6</b>	<b>destination-group grp_id</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
<b>Step 7</b>	<b>ip address ip_addr port port protocol { HTTP   gRPC } encoding { JSON   GPB   GPB-compact }</b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream.
<b>Step 8</b>	<b>subscription sub_id</b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
<b>Step 9</b>	<b>snsr-group sgrp_id sample-interval interval</b> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur.
<b>Step 10</b>	<b>dst-group dgrp_id</b> <b>Example:</b> <pre>switch-1(conf-tm-sub)# dst-grp 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> command.

## Configuring the Native Data Source Path for All MAC Information

You can configure the native data source path for MAC information, which sends information about all entries in the MAC table from Layer 3 and Layer 2. When you subscribe, the baseline sends all the MAC information.

After the baseline, notifications are sent for add, update, and delete MAC address operations. For the data sent in the MAC notifications, see [Telemetry Data Streamed for Native Data Source Paths](#), on page 224.



**Note** For update or delete events, MAC notifications are sent only for the MAC addresses that have IP adjacencies.

### Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **data-source native**
5. **path mac-all**
6. **destination-group** *grp\_id*
7. **ip address** *ip\_addr* **port** *port* **protocol** { HTTP | gRPC } **encoding** { JSON | GPB | GPB-compact }
8. **subscription** *sub\_id*
9. **snsr-group** *sgrp\_id* **sample-interval** *interval*
10. **dst-group** *dgrp\_id*

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
Step 2	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
Step 3	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# sensor-grp 6 switch-1(conf-tm-sub)#</pre>	Create a sensor group.
Step 4	<b>data-source native</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# data-source native switch-1(conf-tm-sensor)#</pre>	Set the data source to native so that any native application can use the streamed data without requiring a specific model or database.

	Command or Action	Purpose
<b>Step 5</b>	<b>path mac-all</b> <b>Example:</b> <pre>nxosv2(conf-tm-sensor)# path mac-all nxosv2(conf-tm-sensor)#</pre>	Configure the MAC path which streams information about all MAC entries and MAC notifications.
<b>Step 6</b>	<b>destination-group grp_id</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
<b>Step 7</b>	<b>ip address ip_addr port port protocol { HTTP   gRPC } encoding { JSON   GPB   GPB-compact }</b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream.
<b>Step 8</b>	<b>subscription sub_id</b> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
<b>Step 9</b>	<b>snsr-group sgrp_id sample-interval interval</b> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur.
<b>Step 10</b>	<b>dst-group dgrp_id</b> <b>Example:</b> <pre>switch-1(conf-tm-sub)# dst-grp 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> 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 224](#).

### Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **data-source native**
5. **path adjacency**
6. **destination-group** *grp\_id*
7. **ip address** *ip\_addr* **port** *port* **protocol** { HTTP | gRPC } **encoding** { JSON | GPB | GPB-compact }
8. **subscription** *sub\_id*
9. **snsr-group** *sgrp\_id* **sample-interval** *interval*
10. **dst-group** *dgrp\_id*

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
Step 2	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.
Step 3	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# sensor-grp 6 switch-1(conf-tm-sub)#</pre>	Create a sensor group.
Step 4	<b>data-source native</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# data-source native switch-1(conf-tm-sensor)#</pre>	Set the data source to native so that any native application can use the streamed data.
Step 5	<b>path adjacency</b> <b>Example:</b> <pre>nxosv2(conf-tm-sensor)# path adjacency nxosv2(conf-tm-sensor)#</pre>	Configure the Adjacency path which streams information about the IPv4 and IPv6 adjacencies.

	Command or Action	Purpose
<b>Step 6</b>	<b>destination-group</b> <i>grp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
<b>Step 7</b>	<b>ip address</b> <i>ip_addr</i> <b>port</b> <i>port</i> <b>protocol</b> { HTTP   gRPC } <b>encoding</b> { JSON   GPB   GPB-compact } <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream.
<b>Step 8</b>	<b>subscription</b> <i>sub_id</i> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
<b>Step 9</b>	<b>snsr-group</b> <i>sgrp_id</i> <b>sample-interval</b> <i>interval</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when interface events occur.
<b>Step 10</b>	<b>dst-group</b> <i>dgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# dst-grp 33 switch-1(conf-tm-sub)#</pre>	Link the destination group to the current subscription. The destination group that you specify must match the destination group that you configured in the <b>destination-group</b> 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-1# 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-1#
```

An example of the statistics for the MAC path:

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

An example of the statistics for the Adjacency path:

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

### 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-1# 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-1#
```

## Streaming Syslog

### About Streaming Syslog for Telemetry

Beginning with Cisco NX-OS release 9.3(3), model-driven telemetry supports streaming of syslogs using YANG as a data source. When you create a subscription, all the syslogs are streamed to the receiver as a baseline. This feature works with the NX-SDK to support streaming syslog data from the following syslog paths:

- `Cisco-NX-OS-Syslog-oper:syslog`
- `Cisco-NX-OS-Syslog-oper:syslog/messages`

After the baseline, only syslog event notifications stream to the receiver. Streaming of syslog paths supports the following encoding types:

- Google Protobuf (GPB)
- JavaScript Object Notation (JSON)

## Configuring the 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 224](#).

### Before you begin

If you have not enabled the telemetry feature, enable it now (**feature telemetry**).

### SUMMARY STEPS

1. **configure terminal**
2. **telemetry**
3. **sensor-group** *sgrp\_id*
4. **data-source native**
5. **path rib query-condition** [**data=ephemeral** | **updates\_only**]
6. **destination-group** *grp\_id*
7. **ip address** *ip\_addr* **port** *port* **protocol** { **HTTP** | **gRPC** } **encoding** { **JSON** | **GPB** | **GPB-compact** }
8. **subscription** *sub\_id*
9. **snsr-group** *sgrp\_id* **sample-interval** *interval*
10. **dst-group** *dgrp\_id*

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure terminal</b> <b>Example:</b> <pre>switch-1# configure terminal switch-1(config)#</pre>	Enter configuration mode.
Step 2	<b>telemetry</b> <b>Example:</b> <pre>switch-1(config)# telemetry switch-1(config-telemetry)#</pre>	Enter configuration mode for the telemetry features.

	Command or Action	Purpose
Step 3	<b>sensor-group</b> <i>sgrp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# sensor-grp 6 switch-1(conf-tm-sub)#</pre>	Create a sensor group.
Step 4	<b>data-source native</b> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# data-source native switch-1(conf-tm-sensor)#</pre>	Set the data source to native so that any native application can use the streamed data without requiring a specific model or database.
Step 5	<b>path rib query-condition</b> [ <b>data=ephemeral</b>   <b>updates_only</b> ] <b>Example:</b> <pre>nxosv2(conf-tm-sensor)# path rib nxosv2(conf-tm-sensor)#</pre>	Configure the RIB path which streams routes and route update information.
Step 6	<b>destination-group</b> <i>grp_id</i> <b>Example:</b> <pre>switch-1(conf-tm-sensor)# destination-group 33 switch-1(conf-tm-dest)#</pre>	Enter telemetry destination group submode and configure the destination group.
Step 7	<b>ip address</b> <i>ip_addr</i> <b>port</b> <i>port</i> <b>protocol</b> { <b>HTTP</b>   <b>gRPC</b> } <b>encoding</b> { <b>JSON</b>   <b>GPB</b>   <b>GPB-compact</b> } <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol http encoding json switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb switch-1(conf-tm-dest)#</pre> <b>Example:</b> <pre>switch-1(conf-tm-dest)# ip address 192.0.2.11 port 50001 protocol grpc encoding gpb-compact switch-1(conf-tm-dest)#</pre>	Configure the telemetry data for the subscription to stream to the specified IP address and port and set the protocol and encoding for the data stream.
Step 8	<b>subscription</b> <i>sub_id</i> <b>Example:</b> <pre>switch-1(conf-tm-dest)# subscription 33 switch-1(conf-tm-sub)#</pre>	Enter telemetry subscription submode, and configure the telemetry subscription.
Step 9	<b>snsr-group</b> <i>sgrp_id</i> <b>sample-interval</b> <i>interval</i> <b>Example:</b> <pre>switch-1(conf-tm-sub)# snsr-grp 6 sample-interval 5000 switch-1(conf-tm-sub)#</pre>	Link the sensor group to the current subscription and set the data sampling interval in milliseconds. The sampling interval determines whether the switch sends telemetry data periodically, or when rib events occur.

	Command or Action	Purpose
<b>Step 10</b>	<b>dst-group</b> <i>dgrp_id</i> <b>Example:</b> switch-1 (conf-tm-sub) # <b>dst-grp 33</b> switch-1 (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 <b>destination-group</b> 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: <ul style="list-style-type: none"> <li>• message-id</li> <li>• node-name</li> <li>• time-stamp</li> <li>• time-of-day</li> <li>• time-zone</li> <li>• category</li> <li>• message-name</li> <li>• severity</li> <li>• text</li> </ul>

### 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"
  }
}
}
```

## Additional References

### Related Documents

Related Topic	Document Title
Example configurations of telemetry deployment for VXLAN EVPN.	<a href="#">Telemetry Deployment for VXLAN EVPN Solution</a>



## APPENDIX **A**

# Streaming Telemetry Sources

---

- [About Streaming Telemetry, on page 243](#)
- [Data Available for Telemetry, on page 243](#)

## About Streaming Telemetry

The streaming telemetry feature of Cisco Nexus switches continuously streams data out of the network and notifies the client, providing near-real-time access to monitoring data.

## 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.





## INDEX

### B

- Bash [3, 5](#)
  - accessing [3](#)
  - examples [5](#)
  - feature bash-shell [3](#)
- Bourne-Again SHell, *See* Bash

### F

- feature grpc [146–147](#)

### G

- grpc gnmi max-concurrent-call [146–147](#)

### N

- NX-API [67–68, 70, 72, 77, 79](#)
  - CLI [68](#)
  - cookie [68](#)
  - management commands [70](#)
  - message format [68](#)

### NX-API (*continued*)

- request elements [72](#)
- response codes [77](#)
- response elements [72](#)
- security [68](#)
- transport [67](#)
- user interface [79](#)

### T

- tcl [35–37, 40](#)
  - cli commands [36](#)
  - command separation [36](#)
  - history [36](#)
  - no interactive help [35](#)
  - options [37](#)
  - references [40](#)
  - sandbox [37](#)
  - security [37](#)
  - tab completion [36](#)
  - telquit command [37](#)
  - variables [37](#)
- Tool Command Language, *See* tcl

