



gNMI - gRPC Network Management Interface

This chapter contains the following topics:

- [About gNMI, on page 1](#)
- [gNMI RPC and SUBSCRIBE, on page 2](#)
- [Guidelines and Limitations for gNMI, on page 3](#)
- [Configuring gNMI, on page 4](#)
- [Configuring Server Certificate, on page 5](#)
- [Generating Key/Certificate Examples , on page 6](#)
- [Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3\(2\) and Earlier, on page 7](#)
- [Generating and Configuring Key/Certificate Examples for Cisco NX-OS Release 9.3\(3\) and Later, on page 8](#)
- [Verifying gNMI, on page 9](#)
- [Clients, on page 10](#)
- [Sample DME Subscription - PROTO Encoding, on page 10](#)
- [Capabilities, on page 12](#)
- [Get, on page 16](#)
- [Set, on page 17](#)
- [Subscribe, on page 18](#)
- [Troubleshooting, on page 21](#)

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 1: 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 2: 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:

- 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 3: 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	configure terminal Example: <pre>switch-1# configure terminal switch-1(config)#</pre>	Enters global configuration mode.
Step 2	feature grpc Example: <pre>switch-1# feature grpc switch-1(config)#</pre>	Enables the gRPC agent, which supports the gNMI interface for dial-in.
Step 3	grpc gnmi max-concurrent-call <i>number</i> Example: <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>Note 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 7](#)

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
```

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

The following is an example for generating key/certificate.



Note 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
show grpc gnmi service statistics	<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"> • Basic overall counters • Certificate expiration time <p>Note 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.

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: <
```

Example Client Output for Capabilities

```

    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>

```

```

        <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 /volaiflcls -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#.
```

