



Telemetry Configuration Guide for Cisco NCS 540 Series Routers, IOS XR Release 7.0.x

First Published: 2019-08-30

Last Modified: 2023-04-14

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 IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

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 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.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

The documentation set for this product strives to use bias-free language. For 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 standards documentation, or language that is used by a referenced third-party product.

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

© 2019 Cisco Systems, Inc. All rights reserved.

- 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.

© 2019 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1	Stream Telemetry Data	1
	Scope	1
	Need	1
	Benefits	2
	Methods of Telemetry	2

CHAPTER 2	Configure Model-driven Telemetry	3
	Configure Dial-out Mode	4
	Create a Destination Group	4
	Create a Sensor Group	4
	Create a Subscription	5
	Validate Dial-out Configuration	6

CHAPTER 3	Core Components of Model-driven Telemetry Streaming	7
	Session	7
	Dial-in Mode	7
	Dial-out Mode	7
	Sensor Path	8
	Subscription	8
	Transport and Encoding	8

CHAPTER 4	Example: Configure Event-driven Telemetry for LLDP	11
------------------	---	-----------

CHAPTER 5	Configure Policy-based Telemetry	15
	Create Policy File	15
	Copy Policy File	17

Configure Encoder	17
Configure JSON Encoder	18
Configure GPB Encoder	18
Verify Policy Activation	19

CHAPTER 6 **Core Components of Policy-based Telemetry Streaming** **21**

Telemetry Policy File	21
Schema Paths	22
Telemetry Encoder	23
TCP Header	24
JSON Message Format	25
GPB Message Format	27
Telemetry Receiver	30



CHAPTER 1

Stream Telemetry Data

This document will help you understand the process of streaming telemetry data and its core components.

- [Scope, on page 1](#)
- [Need, on page 1](#)
- [Benefits, on page 2](#)
- [Methods of Telemetry, on page 2](#)

Scope

Streaming telemetry lets users direct data to a configured receiver. This data can be used for analysis and troubleshooting purposes to maintain the health of the network. This is achieved by leveraging the capabilities of machine-to-machine communication.

The data is used by development and operations (DevOps) personnel who plan to optimize networks by collecting analytics of the network in real-time, locate where problems occur, and investigate issues in a collaborative manner.

Need

Collecting data for analyzing and troubleshooting has always been an important aspect in monitoring the health of a network.

IOS XR 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.

Network state indicators, network statistics, and critical infrastructure information are exposed to the application layer, where they are used to enhance operational performance and to reduce troubleshooting time. A push model uses this capability to continuously stream data out of the network and notify the client. Telemetry enables the push model, which provides near-real-time access to monitoring data.

Streaming telemetry provides a mechanism to select data of interest from IOS XR routers and to transmit it in a structured format to remote management stations for monitoring. This mechanism enables automatic tuning of the network based on real-time data, which is crucial for its seamless operation. The finer granularity

and higher frequency of data available through telemetry enables better performance monitoring and therefore, better troubleshooting. It helps a more service-efficient bandwidth utilization, link utilization, risk assessment and control, remote monitoring and scalability. Streaming telemetry, thus, converts the monitoring process into a Big Data proposition that enables the rapid extraction and analysis of massive data sets to improve decision-making.

Benefits

Streamed real-time telemetry data is useful in:

- **Traffic optimization:** When link utilization and packet drops in a network are monitored frequently, it is easier to add or remove links, re-direct traffic, modify policing, and so on. With technologies like fast reroute, the network can switch to a new path and re-route faster than the SNMP poll interval mechanism. Streaming telemetry data helps in providing quick response time for faster traffic.
- **Preventive troubleshooting:** Helps to quickly detect and avert failure situations that result after a problematic condition exists for a certain duration.

Methods of Telemetry

Telemetry data can be streamed using these methods:

- **Model-driven telemetry:** provides a mechanism to stream data from an MDT-capable device to a destination. The data to be streamed is driven through subscription. There are two methods of configuration:
 - **Cadence-based telemetry:** Cadence-based Telemetry (CDT) continuously streams data (operational statistics and state transitions) at a configured cadence. The streamed data helps users closely identify patterns in the networks. For example, streaming data about interface counters and so on.
- **Policy-based telemetry:** streams telemetry data to a destination using a policy file. A policy file defines the data to be streamed and the frequency at which the data is to be streamed.



Note Model-driven telemetry supersedes policy-based telemetry.



CHAPTER 2

Configure Model-driven Telemetry

Model-driven Telemetry (MDT) provides a mechanism to stream data from an MDT-capable device to a destination. The data to be streamed is defined through subscription.

The data to be streamed is subscribed from a data set in a YANG model. The data from the subscribed data set is streamed out to the destination either at a configured periodic interval or only when an event occurs. This behavior is based on whether MDT is configured for cadence-based telemetry .

The following YANG models are used to configure and monitor MDT:

- **Cisco-IOS-XR-telemetry-model-driven-cfg.yang** and **openconfig-telemetry.yang**: configure MDT using NETCONF or merge-config over grpc.
- **Cisco-IOS-XR-telemetry-model-driven-oper.yang**: get the operational information about MDT.

The process of streaming MDT data uses these components:

- **Destination**: specifies one or more destinations to collect the streamed data.
- **Sensor path**: specifies the YANG path from which data has to be streamed.
- **Subscription**: binds one or more sensor-paths to destinations, and specifies the criteria to stream data. In cadence-based telemetry, data is streamed continuously at a configured frequency.
- **Transport and encoding**: represents the delivery mechanism of telemetry data.

For more information about the core components, see [Core Components of Model-driven Telemetry Streaming, on page 7](#).

The options to initialize a telemetry session between the router and destination is based on two modes:

- **Dial-out mode**: The router initiates a session to the destinations based on the subscription.
- **Dial-in mode**: The destination initiates a session to the router and subscribes to data to be streamed.



Note Dial-in mode is supported only over gRPC.



Important From Release 6.1.1 onwards, Cisco introduces support for the 64-bit Linux-based IOS XR operating system. The 64-bit platforms, such as NCS5500, NCS5000, ASR9000 support gRPC, UDP and TCP protocols. All 32-bit IOS XR platforms, such as CRS and legacy ASR9000, support only TCP protocol.

Streaming model-driven telemetry data to the intended receiver involves these tasks:

- [Configure Dial-out Mode, on page 4](#)

Configure Dial-out Mode

In a dial-out mode, the router initiates a session to the destinations based on the subscription.

All 64-bit IOS XR platforms (except for NCS 6000 series routers) support gRPC and TCP protocols. All 32-bit IOS XR platforms support only TCP.

For more information about the dial-out mode, see [Dial-out Mode, on page 7](#).

The process to configure a dial-out mode involves:

Create a Destination Group

The destination group specifies the destination address, port, encoding and transport that the router uses to send out telemetry data.

1. Identify the destination address, port, transport, and encoding format.
2. Create a destination group.

```
Router(config)#telemetry model-driven
Router(config-model-driven)#destination-group <group-name>

Router(config-model-driven-dest)#address family ipv4 <IP-address> port <port-number>
Router(config-model-driven-dest-addr)#encoding <encoding-format>
Router(config-model-driven-dest-addr)#protocol <transport>
Router(config-model-driven-dest-addr)#commit
```

Example: Destination Group for TCP Dial-out

The following example shows a destination group `DGroup1` created for TCP dial-out configuration with key-value Google Protocol Buffers (also called self-describing-gpb) encoding:

```
Router(config)#telemetry model-driven
Router(config-model-driven)#destination-group DGroup1
Router(config-model-driven-dest)#address family ipv4 172.0.0.0 port 5432
Router(config-model-driven-dest-addr)#encoding self-describing-gpb
Router(config-model-driven-dest-addr)#protocol tcp
Router(config-model-driven-dest-addr)#commit
```

Create a Sensor Group

The sensor-group specifies a list of YANG models that are to be streamed.

1. Identify the sensor path for XR YANG model.
2. Create a sensor group.

```
Router(config)#telemetry model-driven
Router(config-model-driven)#sensor-group <group-name>
Router(config-model-driven-snsr-grp)# sensor-path <XR YANG model>
Router(config-model-driven-snsr-grp)# commit
```

Example: Sensor Group for Dial-out



Note gRPC is supported in only 64-bit platforms.

The following example shows a sensor group `SGroup1` created for dial-out configuration with the YANG model for interface statistics:

```
Router(config)#telemetry model-driven
Router(config-model-driven)#sensor-group SGroup1
Router(config-model-driven-snsr-grp)# sensor-path
Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/latest/generic-counters
Router(config-model-driven-snsr-grp)# commit
```

What to Do Next:

Create a subscription.

Create a Subscription

The subscription associates a destination-group with a sensor-group and sets the streaming method.

A source interface in the subscription group specifies the interface that will be used for establishing the session to stream data to the destination. If both VRF and source interface are configured, the source interface must be in the same VRF as the one specified under destination group for the session to be established.

```
Router(config)#telemetry model-driven
Router(config-model-driven)#subscription <subscription-name>
Router(config-model-driven-subs)#sensor-group-id <sensor-group> sample-interval <interval>

Router(config-model-driven-subs)#destination-id <destination-group>
Router(config-model-driven-subs)#source-interface <source-interface>
Router(config-mdt-subscription)#commit
```

Example: Subscription for Cadence-based Dial-out Configuration

The following example shows a subscription `Sub1` that is created to associate the sensor-group and destination-group, and configure an interval of 30 seconds to stream data:

```
Router(config)#telemetry model-driven
Router(config-model-driven)#subscription Sub1
Router(config-model-driven-subs)#sensor-group-id SGroup1 sample-interval 30000
Router(config-model-driven-subs)#destination-id DGroup1
Router(config-mdt-subscription)# commit
```

Validate Dial-out Configuration

Use the following command to verify that you have correctly configured the router for dial-out.

```
Router#show telemetry model-driven subscription <subscription-group-name>
```

Example: Validation for TCP Dial-out

```
Router#show telemetry model-driven subscription Sub1
Thu Jul 21 15:42:27.751 UTC
Subscription: Sub1                               State: ACTIVE
-----
Sensor groups:
Id          Interval(ms)   State
SGroup1    30000          Resolved

Destination Groups:
Id          Encoding      Transport  State  Port  IP
DGroup1    self-describing-gpb tcp        Active 5432 172.0.0.0
```



CHAPTER 3

Core Components of Model-driven Telemetry Streaming

The core components used in streaming model-driven telemetry data are:

- [Session, on page 7](#)
- [Sensor Path, on page 8](#)
- [Subscription, on page 8](#)
- [Transport and Encoding, on page 8](#)

Session

A telemetry session can be initiated using:

Dial-in Mode

In a dial-in mode, an MDT receiver dials in to the router, and subscribes dynamically to one or more sensor paths or subscriptions. The router acts as the server and the receiver is the client. The router streams telemetry data through the same session. The dial-in mode of subscriptions is dynamic. This dynamic subscription terminates when the receiver cancels the subscription or when the session terminates.

There are two methods to request sensor-paths in a dynamic subscription:

- **OpenConfig RPC model:** The `subscribe` RPC defined in the model is used to specify sensor-paths and frequency. In this method, the subscription is not associated with an existing configured subscription. A subsequent `cancel` RPC defined in the model removes an existing dynamic subscription.
- **IOS XR MDT RPC:** IOS XR defines RPCs to subscribe and to cancel one or more configured subscriptions. The sensor-paths and frequency are part of the telemetry configuration on the router. A subscription is identified by its configured subscription name in the RPCs.

Dial-out Mode

In a dial-out mode, the router dials out to the receiver. This is the default mode of operation. The router acts as a client and receiver acts as a server. In this mode, sensor-paths and destinations are configured and bound together into one or more subscriptions. The router continually attempts to establish a session with each destination in the subscription, and streams data to the receiver. The dial-out mode of subscriptions is persistent.

When a session terminates, the router continually attempts to re-establish a new session with the receiver every 30 seconds.

Sensor Path

The sensor path describes a YANG path or a subset of data definitions in a YANG model with a container. In a YANG model, the sensor path can be specified to end at any level in the container hierarchy.

An MDT-capable device, such as a router, associates the sensor path to the nearest container path in the model. The router encodes and streams the container path within a single telemetry message. A receiver receives data about all the containers and leaf nodes at and below this container path.

The router streams telemetry data for one or more sensor-paths, at the configured frequency (cadence-based streaming) to one or more receivers through subscribed sessions.

Subscription

A subscription binds one or more sensor paths and destinations. An MDT-capable device streams data for each sensor path at the configured frequency (cadence-based streaming) to the destination.

The following example shows subscription `SUB1` that associates a sensor-group, sample interval and destination group.

```
Router(config)#telemetry model-driven
Router(config-model-driven)#subscription SUB1
Router(config-model-driven-subs)#sensor-group-id SGROUP1 sample-interval 10000
Router(config-model-driven-subs)#strict-timer
```



Note With a `strict-timer` configured for the sample interval, the data collection starts exactly at the configured time interval allowing a more deterministic behavior to stream data.

In 32-bit platforms, `strict-timer` can be configured only under the subscription. Whereas, 64-bit platforms support configuration at global level in addition to the subscription level. However, configuring at the global level will affect all configured subscriptions.

```
Router(config)#telemetry model-driven
Router(config-model-driven)#strict-timer
```

Transport and Encoding

The router streams telemetry data using a transport mechanism. The generated data is encapsulated into the desired format using encoders.

Model-Driven Telemetry (MDT) data is streamed through :

- **Transmission Control Protocol (TCP):** used for only dial-out mode.
- **User Datagram Protocol (UDP):** used for only dial-out mode.

The data to be streamed can be encoded into Google Protocol Buffers (GPB) encoding. In GPB, the encoding can either be compact GPB (for optimising the network bandwidth usage) or self-describing GPB. The encodings supported are:

- **GPB encoding:** configuring for GPB encoding requires metadata in the form of compiled .proto files. A .proto file describes the GPB message format, which is used to stream data. The .proto files are available in the Github repository.
 - **Compact GPB encoding:** data is streamed in compressed and non self-describing format. A .proto file corresponding to each sensor-path must be used by the receiver to decode the streamed data.
 - **Key-value (KV-GPB) encoding:** data of each sensor path streamed is in a self-describing formatted ASCII text. A single .proto file `telemetry.proto` is used by the receiver to decode any sensor path data. Because the key names are included in the streamed data, the data on the wire is much larger as compared to compact GPB encoding.



Note Telemetry data is streamed out of the router using an Extensible Manageability Services Daemon (emsd) process. The data of interest is subscribed through subscriptions and streamed through gRPC, TCP or UDP sessions. However, a combination of gRPC, TCP and UDP sessions with more than 150 active sessions leads to emsd crash or process restart.



CHAPTER 4

Example: Configure Event-driven Telemetry for LLDP

Telemetry supports NETCONF event notifications where the NETCONF client is configured to receive event notifications from a NETCONF server through a subscription. The NETCONF client must subscribe using a `create-subscription` request. Currently, only the events from Link Layer Discovery Protocol (LLDP) is supported. These event notifications are sent until either the NETCONF session or the subscription is terminated.



Note Configuring a sensor group and a subscription is not required for receiving NETCONF notifications. While sensor path and subscription configurations are required for receiving telemetry events, NETCONF `create-subscription` is required for receiving NETCONF notifications.

To generate NETCONF notifications:

1. Enable NETCONF agent and SSH sub system.

```
ssh server netconf
netconf-yang agent ssh
```

2. Enable model-driven telemetry.

```
telemetry model-driven
```

3. Enable LLDP.

```
lldp
```

This example shows event-driven telemetry for LLDP configuration data.

1. Create a destination group.

```
grpc
port 56782
address-family ipv4
!
telemetry model-driven
destination-group <destination-udp>
  address-family ipv4 <client-ip>1 port <udp port num>
  encoding self-describing-gpb
  protocol udp
```

```

!
!
destination-group <destination-tcp>
  address-family ipv4 <client-ip> port <tcp port num>
  encoding gpb
  protocol tcp
!
destination-group <destination-grpc>
  address-family ipv4 <grpc client ip>port <grpc port num>
  encoding self-describing-gpb
  protocol grpc no-tls

```

2. Create a sensor group.

```

sensor-group <sensor-group-name>
  sensor-path Cisco-IOS-XR-ethernet-lldp-oper:lldp/global-lldp/lldp-info
  sensor-path Cisco-IOS-XR-ethernet-lldp-oper:lldp/nodes/node/interfaces/interface
  sensor-path Cisco-IOS-XR-ethernet-lldp-oper:lldp/nodes/node/neighbors/details/detail
!

```

3. Create a subscription.

```

subscription udp-out
  sensor-group-id <sensor-group-name> sample-interval 0
  destination-id <destination-udp>
!

subscription <subscription-name>
  sensor-group-id <sensor-group-name> sample-interval 0
  destination-id <destination-tcp>

subscription <subscription-name>
  sensor-group-id <sensor-group-name> sample-interval 0
!
netconf-yang agent
ssh
!

```

4. Set the notification to stream data when an event occurs.

```

Router(config-lldp)#timer 12
Router(config-lldp)#commit

Router(config-lldp)#holdtime 150
Router (config-lldp)#commit

Router (config-lldp)#exit
#506
<?xml version="1.0"?>
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
  <eventTime>Date-and-Time</eventTime>
  <lldp xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-ethernet-lldp-oper">
    <global-lldp>
      <lldp-info>
        <chassis-id>000b.1bc9.e700</chassis-id>
        <chassis-id-sub-type>4</chassis-id-sub-type>
        <system-name>ios</system-name>
        <timer>12</timer>
        <hold-time>120</hold-time>
        <re-init>2</re-init>
      </lldp-info>
    </global-lldp>
  </lldp>
</notification>

```

```

</lldp>
</notification>
Ready to send a request.
Paste your request or enter 'get', 'get-config', 'create-sub', or 'bye' to quit):

```

5. Validate response received from NETCONF agent.

```

#506
<?xml version="1.0"?>
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
  <eventTime>Date-and-Time</eventTime>
  <lldp xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-ethernet-lldp-oper">
    <global-lldp>
      <lldp-info>
        <chassis-id>000b.1bc9.e700</chassis-id>
        <chassis-id-sub-type>4</chassis-id-sub-type>
        <system-name>ios</system-name>
        <timer>12</timer>
        <hold-time>150</hold-time>
        <re-init>2</re-init>
      </lldp-info>
    </global-lldp>
  </lldp>
</notification>

```




CHAPTER 5

Configure Policy-based Telemetry

Policy-based telemetry (PBT) streams telemetry data to a destination using a policy file. A policy file defines the data to be streamed and the frequency at which the data is to be streamed.

ASR9000 series routers and CRS routers do not support PBT.

The process of streaming telemetry data uses three core components:

- **Telemetry Policy File** specifies the kind of telemetry data to be generated, at a specified frequency.
- **Telemetry Encoder** encapsulates the generated data into the desired format and transmits to the receiver.
- **Telemetry Receiver** is the remote management system that stores the telemetry data.

For more information about the three core components, see [Core Components of Policy-based Telemetry Streaming, on page 21](#).



Note Model-driven telemetry supersedes policy-based telemetry.

Streaming policy-based telemetry data to the intended receiver involves these tasks:

- [Create Policy File, on page 15](#)
- [Copy Policy File, on page 17](#)
- [Configure Encoder, on page 17](#)
- [Verify Policy Activation, on page 19](#)

Create Policy File

You define a telemetry policy file to specify the kind of telemetry data to be generated and pushed to the receiver. Defining the policy files requires a path to stream data. The paths can be schemas, native YANG or allowed list entries.

For more information on the schema paths associated with a corresponding CLI command, see [Schema Paths, on page 22](#).

For more information on policy files, see [Telemetry Policy File, on page 21](#).

1. Determine the schema paths to stream data.

For example, the schema path for interfaces is:

```
RootOper.InfraStatistics.Interface(*).Latest.GenericCounters
```

2. Create a policy file that contains these paths.

Example: Policy File

The following example shows a sample policy file for streaming the generic counters of an interface:

```
{
  "Name": "Test",
  "Metadata": {
    "Version": 25,
    "Description": "This is a sample policy",
    "Comment": "This is the first draft",
    "Identifier": "<data that may be sent by the encoder to the mgmt stn"
  },
  "CollectionGroups": {
    "FirstGroup": {
      "Period": 10,
      "Paths": [
        "RootOper.InfraStatistics.Interface(*).Latest.GenericCounters"
      ]
    }
  }
}
```

The following example shows the paths with allowed list entries in the policy file. Instead of streaming all the data for a particular entry, only specific items can be streamed using allowed list entries. The entries are allowed using `IncludeFields` in the policy file. In the example, the entry within the `IncludeFields` section streams only the latest applied AutoBW value for that TE tunnel, which is nested two levels down from the top level of the path:

```
{
  "Name": "RSVPTEPolicy",
  "Metadata": {
    "Version": 1,
    "Description": "This policy collects auto bw stats",
    "Comment": "This is the first draft"
  },
  "CollectionGroups": {
    "FirstGroup": {
      "Period": 10,
      "Paths": {
        "RootOper.MPLS_TE.P2P_P2MPTunnel.TunnelHead({'TunnelName':
'tunnel-tel0'})": {
          "IncludeFields": [{
            "P2PInfo": [{
              "AutoBandwidthOper": [
                "LastBandwidthApplied"
              ]
            }
          ]
        }
      ]
    }
  }
}
```

```
    }
}
```

The following example shows the paths with native YANG entry in the policy file. This entry will stream the generic counters of the interface:

```
"Paths": [
  "/Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface=*/latest/generic-counters"
]
```

What to Do Next:

Copy the policy file to the router. You may copy the same policy file to multiple routers.

Copy Policy File

Run the Secure Copy Protocol (SCP) command to securely copy the policy file from the server where it is created. For example:

```
$ scp Test.policy <ip-address-of-router>:/telemetry/policies
```

For example, to copy the `Test.policy` file to the `/telemetry/policies` folder of a router with IP address 10.0.0.1:

```
$ scp Test.policy cisco@10.0.0.1:/telemetry/policies
cisco@10.0.0.1's password:
Test.policy
100% 779    0.8KB/s   00:00
Connection to 10.0.0.1 closed by remote host.
```

Verify Policy Installation

In this example, the policy is installed in the `/telemetry/policies/` folder in the router file system. Run the **show telemetry policies brief** command to verify that the policy is successfully copied to the router.

```
Router#show telemetry policy-driven policies brief
Wed Aug 26 02:24:40.556 PDT

Name                               |Active?| Version | Description
-----|-----|-----|-----
Test                               |N      | 1       | This is a sample policy
```

What to Do Next:

Configure the telemetry encoder to activate and stream data.

Configure Encoder

An encoder calls the streaming Telemetry API to:

- Specify policies to be explicitly defined
- Register all policies of interest

Configure the encoder to activate the policy and stream data. More than one policy and destination can be specified. Multiple policy groups can be specified under each encoder and each group can be streamed to multiple destinations. When multiple destinations are specified, the data is streamed to all destinations.

Configure an encoder based on the requirement.

Configure JSON Encoder

The JavaScript Object Notation (JSON) encoder is packaged with the IOS XR software and provides the default format for streaming telemetry data.

To stream data in JavaScript Object Notation (JSON) format, specify the encoder, policies, policy group, destination, and port:

```
Router# configure
Router(config)#telemetry policy-driven encoder json
Router(config-telemetry-json)#policy group FirstGroup
Router(config-policy-group)#policy Test
Router(config-policy-group)#destination ipv4 10.0.0.1 port 5555
Router(config-policy-group)#commit
```

The names of the policy and the policy group must be identical to the policy and its definition that you create. For more information on policy files, see [Create Policy File, on page 15](#).

For more information about the message format of JSON encoder, see [JSON Message Format, on page 25](#)

Configure GPB Encoder

Configuring the GPB (Google Protocol Buffer) encoder requires metadata in the form of compiled `.proto` files. A `.proto` file describes the GPB message format, which is used to stream data.

Two encoding formats are supported:

- **Compact encoding** stores data in a compressed and non-self-describing format. A `.proto` file must be generated for each path in the policy file to be used by the receiver to decode the resulting data.
- **Key-value encoding** uses a single `.proto` file to encode data in a self-describing format. This encoding does not require a `.proto` file for each path. The data on the wire is much larger because key names are included.

To stream GPB data, complete these steps:

1. For compact encoding, create `.proto` files for all paths that are to be streamed using the following tool:

```
telemetry generate gpb-encoding path <path> [file <output_file>]
```

or

```
telemetry generate gpb-encoding policy <policy_file> directory <output_dir>
```



Attention A parser limitation does not support the use of quotes within paths in the tool. For example, for use in the tool, change this policy path,

```
RootOper.InfraStatistics.Interface(*) .Latest.Protocol(['IPV4_UNICAST']) to
RootOper.InfraStatistics.Interface(*) .Latest.Protocol.
```


2. Copy the policy file to the router.
3. Configure the telemetry policy specifying the encoder, policies, policy group, destination, and port:

```
Router# configure
Router(config)#telemetry policy-driven encoder gpb
Router(config-telemetry-json)#policy group FirstGroup
Router(config-policy-group)#policy Test
Router(config-policy-group)#destination ipv4 10.0.0.1 port 5555
Router(config-policy-group)#commit
```

For more information about the message format of GPB encoder, see [GPB Message Format, on page 27](#)

Verify Policy Activation

Verify that the policy is activated using the `show telemetry policies` command.

```
Router#show telemetry policy-driven policies
Wed Aug 26 02:24:40.556 PDT

Filename:           Test.policy
Version:            25
Description:        This is a sample policy to demonstrate the syntax
Status:             Active
CollectionGroup:   FirstGroup
  Cadence:          10s
  Total collections: 2766
  Latest collection: 2015-08-26 02:25:07
  Min collection time: 0.000s
  Max collection time: 0.095s
  Avg collection time: 0.000s
  Min total time:   0.022s
  Max total time:   0.903s
  Avg total time:   0.161s
  Collection errors: 0
  Missed collections: 0
```

```
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Path                                                                                                                                           | Avg (s) |
| Max (s) | Err |                                                                                                                                           |          |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| RootOper.InfraStatistics.Interface(*) .Latest.GenericCounters                                     | 0.000 |
0.000 |    0 |                                                                                                                                           |          |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

After the policy is validated, the telemetry encoder starts streaming data to the receiver. For more information on the receiver, see [Telemetry Receiver, on page 30](#).



CHAPTER 6

Core Components of Policy-based Telemetry Streaming

The core components used in streaming policy-based telemetry data are:

- [Telemetry Policy File, on page 21](#)
- [Telemetry Encoder, on page 23](#)
- [Telemetry Receiver, on page 30](#)

Telemetry Policy File

A telemetry policy file is defined by the user to specify the kind of telemetry data that is generated and pushed to the receiver. The policy must be stored in a text file with a `.policy` extension. Multiple policy files can be defined and installed in the `/telemetry/policies/` folder in the router file system.

A policy file:

- Contains one or more collection groups; a collection group includes different types of data to be streamed at different intervals
- Includes a period in seconds for each group
- Contains one or more paths for each group
- Includes metadata that contains version, description, and other details about the policy

Policy file syntax

The following example shows a sample policy file:

```
{
  "Name": "NameOfPolicy",
  "Metadata": {
    "Version": 25,
    "Description": "This is a sample policy to demonstrate the syntax",
    "Comment": "This is the first draft",
    "Identifier": "<data that may be sent by the encoder to the mgmt stn"
  },
  "CollectionGroups": {
    "FirstGroup": {
      "Period": 10,
      "Paths": [
```

```

        "RootOper.MemorySummary.Node",
        "RootOper.RIB.VRF",
        "...",
    ]
},
"SecondGroup": {
    "Period": 300,
    "Paths": [
        "RootOper.Interfaces.Interface"
    ]
}
}
}

```

The syntax of the policy file includes:

- **Name** the name of the policy. In the previous example, the policy is stored in a file named `NameOfPolicy.policy`. The name of the policy must match the filename (without the `.policy` extension). It can contain uppercase alphabets, lower-case alphabets, and numbers. The policy name is case sensitive.
- **Metadata** information about the policy. The metadata can include the version number, date, description, author, copyright information, and other details that identify the policy. The following fields have significance in identifying the policy:
 - Description is displayed in the **show policies** command.
 - Version and Identifier are sent to the receiver as part of the message header of the telemetry messages.
- **CollectionGroups** an encoder object that maps the group names to information about them. The name of the collection group can contain uppercase alphabets, lowercase alphabets, and numbers. The group name is case sensitive.
- **Period** the cadence for each collection group. The period specifies the frequency in seconds at which data is queried and sent to the receiver. The value must be within the range of 5 and 86400 seconds.
- **Paths** one or more schema paths, allowed list entries or native YANG paths (for a container) for the data to be streamed and sent to the receiver. For example,

Schema path:

```
RootOper.InfraStatistics.Interface(*).Latest.GenericCounters
```

YANG path:

```
/Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface=*/latest/generic-counters
```

Allowed list entry:

```

"RootOper.Interfaces.Interface(*)":
{
    "IncludeFields": ["State"]
}

```

Schema Paths

A schema path is used to specify where the telemetry data is collected. A few paths are listed in the following table for your reference:

Table 1: Schema Paths

Operation	Path
Interface Operational data	RootOper.Interfaces.Interface(*)
Packet/byte counters	RootOper.InfraStatistics.Interface(*).Latest.GenericCounters
Packet/byte rates	RootOper.InfraStatistics.Interface(*).Latest.DataRate
IPv4 packet/byte counters	RootOper.InfraStatistics.Interface(*).Latest.Protocol(['IPV4_UNICAST'])
MPLS stats	<ul style="list-style-type: none"> • RootOper.MPLS_TE.Tunnels.TunnelAutoBandwidth • RootOper.MPLS_TE.P2P_P2MPTunnel.TunnelHead • RootOper.MPLS_TE.SignallingCounters.HeadSignallingCounters
QOS Stats	<ul style="list-style-type: none"> • RootOper.QOS.Interface(*).Input.Statistics • RootOper.QOS.Interface(*).Output.Statistics
BGP Data	RootOper.BGP.Instance({'InstanceName': 'default'}).InstanceActive.DefaultVRF.Neighbor([*])
Inventory data	RootOper.PlatformInventory.Rack(*).Attributes.BasicInfo RootOper.PlatformInventory.Rack(*).Slot(*).Card(*).Sensor(*).Attributes.BasicInfo

Telemetry Encoder

The telemetry encoder encapsulates the generated data into the desired format and transmits to the receiver.

An encoder calls the streaming Telemetry API to:

- Specify policies to be explicitly defined
- Register all policies of interest

Telemetry supports two types of encoders:

- **JavaScript Object Notation (JSON) encoder**

This encoder is packaged with the IOS XR software and provides the default method of streaming telemetry data. It can be configured by CLI and XML to register for specific policies. Configuration is grouped into policy groups, with each policy group containing one or more policies and one or more destinations. JSON encoding is supported over only TCP transport service.

JSON encoder supports two encoding formats:

- **Restconf-style encoding** is the default JSON encoding format.
- **Embedded-keys encoding** treats naming information in the path as keys.

- **Google Protocol Buffers (GPB) encoder**

This encoder provides an alternative encoding mechanism, streaming the data in GPB format over UDP or TCP. It can be configured by CLI and XML and uses the same policy files as those of JSON.

Additionally, a GPB encoder requires metadata in the form of compiled .proto files to translate the data into GPB format.

GPB encoder supports two encoding formats:

- **Compact encoding** stores data in a compact GPB structure that is specific to the policy that is streamed. This format is available over both UDP and TCP transport services. A .proto file must be generated for each path in the policy file to be used by the receiver to decode the resulting data.
- **Key-value encoding** stores data in a generic key-value format using a single .proto file. The encoding is self-describing as the keys are contained in the message. This format is available over TCP transport service. A .proto file is not required for each policy file because the receiver can interpret the data.

TCP Header

Streaming data over a TCP connection either with a JSON or a GPB encoder and having it optionally compressed by zlib ensures that the stream is flushed at the end of each batch of data. This helps the receiver to decompress the data received. If data is compressed using zlib, the compression is done at the policy group level. The compressor resets when a new connection is established from the receiver because the decompressor at the receiver has an empty initial state.

Header of each TCP message:

Type	Flags	Length	Message
4 bytes	4 bytes <ul style="list-style-type: none"> • default - Use 0x0 value to set no flags. • zlib compression - Use 0x1 value to set zlib compression on the message. 	4 bytes	Variable

where:

- The Type is encoded as a big-endian value.
- The Length (in bytes) is encoded as a big-endian value.
- The flags indicates modifiers (such as compression) in big-endian format.
- The message contains the streamed data in either JSON or GPB object.

Type of messages:

Type	Name	Length	Value
1	Reset Compressor	0	No value
2	JSON Message	Variable	JSON message (any format)
3	GPB compact	Variable	GPB message in compact format

Type	Name	Length	Value
4	GPB key-value	Variable	GPB message in key-value format

JSON Message Format

JSON messages are sent over TCP and use the header message described in [TCP Header, on page 24](#).

The message consists of the following JSON objects:

```
{
  "Policy": "<name-of-policy>",
  "Version": "<policy-version>",
  "Identifier": "<data from policy file>"
  "CollectionID": <id>,
  "Path": <Policy Path>,
  "CollectionStartTime": <timestamp>,
  "Data": { ... object as above ... },
  "CollectionEndTime": <timestamp>,
}
```

where:

- `Policy`, `Version` and `Identifier` are specified in the policy file.
- `CollectionID` is an integer that allows messages to be grouped together if data for a single path is split over multiple messages.
- `Path` is the base path of the corresponding data as specified in the policy file.
- `CollectionStartTime` and `CollectionEndTime` are the timestamps that indicate when the data was collected

The JSON message reflects the hierarchy of the router's data model. The hierarchy consists of:

- containers: a container has nodes that can be of different types.
- tables: a table also contains nodes, but the number of child nodes may vary, and they must be of the same type.
- leaf node: a leaf contains a data value, such as integer or string.

The schema objects are mapped to JSON are in this manner:

- Each container maps to a JSON object. The keys are strings that represent the schema names of the nodes; the values represent the values of the nodes.
- JSON objects are also used to represent tables. In this case, the keys are based on naming information that is converted to string format. Two options are provided for encoding the naming information:
 - The default is restconf-style encoding, where naming parameters are contained within the child node to which it refers.
 - The embedded-keys option uses the naming information as keys in a JSON dictionary, with the corresponding child node forming the value.
- Leaf data types are mapped in this manner:

- Simple strings, integers, and booleans are mapped directly.
- Enumeration values are stored as the string representation of the value.
- Other simple data types, such as IP addresses, are mapped as strings.

Example: Rest-conf Encoding

For example, consider the path -

```
Interfaces(*).Counters.Protocols("IPv4")
```

This has two naming parameters - the interface name and the protocol name - and represents a container holding leaf nodes which are packet and byte counters. This would be represented as follows:

```
{
  "Interfaces": [
    {
      "Name": "GigabitEthernet0/0/0/1"
      "Counters": {
        "Protocols": [
          {
            "ProtoName": "IPv4",
            "CollectionTime": 12345678,
            "InputPkts": 100,
            "InputBytes": 200,
          }
        ]
      }
    }, {
      "Name": "GigabitEthernet0/0/0/2"
      "Counters": {
        "Protocols": [
          {
            "ProtoName": "IPv4",
            "CollectionTime": 12345678,
            "InputPkts": 400,
            "InputBytes": 500,
          }
        ]
      }
    }
  ]
}
```

A naming parameter with multiple keys, for example `Foo.Destination(IPAddress=10.1.1.1, Port=2000)` would be represented as follows:

```
{
  "Foo":
  {
    "Destination": [
      {
        "IPAddress": 10.1.1.1,
        "Port": 2000,
        "CollectionTime": 12345678,
        "Leaf1": 100,
      }
    ]
  }
}
```


Example: Embedded Keys Encoding

The embedded-keys encoding treats naming information in the path as keys in the JSON dictionary. The key name information is lost and there are extra levels in the hierarchy but it is clearer which data constitutes the key which may aid collectors when parsing it. This option is provided primarily for backwards-compatibility with 6.0.

```
{
  "Interfaces": {
    "GigabitEthernet0/0/0/1": {
      "Counters": {
        "Protocols": {
          "IPv4": {
            "CollectionTime": 12345678,
            "InputPkts": 100,
            "InputBytes": 200,
          }
        }
      }
    },
    "GigabitEthernet0/0/0/2": {
      "Counters": {
        "Protocols": {
          "IPv4": {
            "CollectionTime": 12345678,
            "InputPkts": 400,
            "InputBytes": 500,
          }
        }
      }
    }
  }
}
```

A naming parameter with multiple keys, for example `Foo.Destination(IPAddress=10.1.1.1, Port=2000)`, would be represented by nesting each key in order:

```
{
  "Foo": [
    {
      "Destination": {
        10.1.1.1: {
          2000: {
            Leaf1": 100,
          }
        }
      }
    }
  ]
}
```

GPB Message Format

The output of the GPB encoder consists entirely of GPBs and allows multiple tables in a single packet for scalability.

GPB (Google Protocol Buffer) encoder requires metadata in the form of compiled `.proto` files. A `.proto` file describes the GPB message format, which is used to stream data.

For TCP, the message body is either a `Telemetry` message or a `TelemetryHeader` message, depending on which of the following encoding types is configured:

- **Compact GPB format** stores data in a compressed and non-self-describing format. A `.proto` file must be generated for each path in the policy file to be used by the receiver to decode the resulting data.
- **Key-value GPB format** uses a single `.proto` file to encode data in a self-describing format. This encoding does not require a `.proto` file for each path. The data on the wire is much larger because key names are included.

In the following example, the policy group, *alpha* uses the default configuration of compact encoding and UDP transport. The policy group, *beta* uses compressed TCP and key-value encoding. The policy group, *gamma* uses compact encoding over uncompressed TCP.

```
telemetry policy-driven encoder gpb
  policy group alpha
    policy foo
    destination ipv4 192.168.1.1 port 1234
    destination ipv4 10.0.0.1 port 9876
  policy group beta
    policy bar
    policy whizz
    destination ipv4 10.20.30.40 port 3333
    transport tcp
    compression zlib
  policy group gamma
    policy bang
    destination ipv4 10.11.1.1 port 4444
    transport tcp
    encoding-format gpb-compact
```

Compact GPB Format

The compact GPB format is intended for streaming large volumes of data at frequent intervals. The format minimizes the size of the message on the wire. Multiple tables can be sent in in a single packet for scalability.



Note The tables can be split over multiple packets but fragmenting a row is not supported. If a row in the table is too large to fit in a single UDP frame, it cannot be streamed. Instead either switch to TCP, increase the MTU, or modify the `.proto` file.

The following `.proto` file shows the header, which is common to all packets sent by the encoder:

```
message TelemetryHeader {
  optional uint32 encoding = 1;

  optional string policy_name = 2;
  optional string version = 3;
  optional string identifier = 4;

  optional uint64 start_time = 5;
  optional uint64 end_time = 6;

  repeated TelemetryTable tables = 7;
}

message TelemetryTable {
  optional string policy_path = 1;
  repeated bytes row = 2;
}
```

where:

- encoding is used by receivers to verify that the packet is valid.
- policy name, version and identifier are metadata taken from the policy file.
- start time and end time indicate the duration when the data is collected.
- tables is a list of tables within the packet. This format indicates that it is possible to receive results for multiple schema paths in a single packet.
- For each table:
 - policy path is the schema path.
 - row is one or more byte arrays that represents an encoded GPB.

Key-value GPB Format

The self-describing key-value GPB format uses a generic .proto file. This file encodes data as a sequence of key-value pairs. The field names are included in the output for the receiver to interpret the data.

The following .proto file shows the field containing the key-value pairs:

```
message Telemetry {
  uint64   collection_id = 1;
  string   base_path = 2;
  string   subscription_identifier = 3;
  string   model_version = 4;
  uint64   collection_start_time = 5;
  uint64   msg_timestamp = 6;
  repeated TelemetryField fields = 14;
  uint64   collection_end_time = 15;
}

message TelemetryField {
  uint64   timestamp = 1;
  string   name = 2;
  bool     augment_data = 3;
  oneof value_by_type {
    bytes   bytes_value = 4;
    string  string_value = 5;
    bool    bool_value = 6;
    uint32  uint32_value = 7;
    uint64  uint64_value = 8;
    sint32  sint32_value = 9;
    sint64  sint64_value = 10;
    double  double_value = 11;
    float   float_value = 12;
  }
  repeated TelemetryField fields = 15;
}
```

where:

- collection_id, base_path, collection_start_time and collection_end_time provide streaming details.
- subscription_identifier is a fixed value for cadence-driven telemetry. This is used to distinguish from event-driven data.
- model_version contains a string used for the version of the data model, as applicable.

Telemetry Receiver

A telemetry receiver is used as a destination to store streamed data.

A sample receiver that handles both JSON and GPB encodings is available in the Github repository.

A copy of the `cisco.proto` file is required to compile code for a GPB receiver. The `cisco.proto` file is available in the Github repository.

If you are building your own collector, use the standard protoc compiler. For example, for the GPB compact encoding:

```
protoc --python_out . -I=/sw/packages/protoc/current/google/include/.. generic_counters.proto  
  ipv4_counters.proto
```

where:

- `--python_out <out_dir>` specifies the location of the resulting generated files. These files are of the form `<name>_pb2.py`.
- `-I <import_path>` specifies the path to look for imports. This must include the location of `descriptor.proto` from Google. (in `/sw/packages`) and `cisco.proto` and the `.proto` files that are compiled.

All files shown in the above example are located in the local directory.