



Telemetry Configuration Guide for Cisco NCS 540 Series Routers, IOS XR Release 7.10.x

First Published: 2023-08-16

Americas Headquarters

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS 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.

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)

© 2023 Cisco Systems, Inc. All rights reserved.



CONTENTS

| | | |
|------------------|--|-----------|
| CHAPTER 1 | YANG Data Models for System Management Features | 1 |
| | Using YANG Data Models | 1 |
| CHAPTER 2 | Scale-Up Your Network Monitoring Strategy Using Telemetry | 3 |
| | Benefits of Shifting Network Monitoring from Pull Models to Telemetry Push Model | 5 |
| | Review Mechanisms to Stream Telemetry Data from a Router to a Destination | 5 |
| | Cadence-driven Telemetry | 6 |
| | Event-driven Telemetry | 6 |
| | Learn About the Elements that Enable Streaming Telemetry Data | 7 |
| | Sensor Path | 7 |
| | Subscription | 9 |
| | Encoder | 9 |
| | Transport | 9 |
| | gRPC Network Management Interface | 10 |
| | gRPC Network Operations Interface | 10 |
| | TLS Authentication For Dial-in | 10 |
| | Explore the Methods to Establish a Telemetry Session | 11 |
| | Dial-Out Mode | 11 |
| | Dial-In Mode | 11 |
| | Identify the Telemetry Session Suitable for Your Network | 12 |
| CHAPTER 3 | Establish a Model-Driven Telemetry Session from a Router to a Collector | 13 |
| | Monitor CPU Utilization Using Telemetry Data to Plan Network Infrastructure | 14 |
| | Define a Subscription to Stream Data from Router to Receiver | 15 |
| | Verify Deployment of the Subscription | 18 |
| | Operate on Telemetry Data for In-depth Analysis of the Network | 19 |

| | | |
|------------------|--|-----------|
| CHAPTER 4 | Establish a Model-Driven Telemetry Session from a Collector to a Router | 23 |
| | Monitor Network Parameters Using Telemetry Data for Proactive Analysis | 24 |
| | Define a Subscription to Stream Data from Router to Receiver | 25 |
| | Verify Deployment of the Subscription | 26 |
| | Operate on Telemetry Data for In-depth Analysis of the Network | 27 |

| | | |
|------------------|--|-----------|
| CHAPTER 5 | Build Intelligence on the Router Using AI-Driven Telemetry | 31 |
| | Key Components | 32 |
| | Processing Telemetry Data on the Router to Analyze Traffic Changes | 33 |
| | Define a Subscription to Stream ADT Events from Router to Receiver | 34 |
| | Verify Deployment of the Subscription | 37 |
| | Operate on Telemetry Data for In-Depth Analysis of the Network | 40 |

| | | |
|------------------|--|-----------|
| CHAPTER 6 | Hardware Timestamp | 43 |
| | Stream Telemetry Data about PBR Decapsulation Statistics | 44 |



CHAPTER 1

YANG Data Models for System Management Features

This chapter provides information about the YANG data models for System Management features.

- [Using YANG Data Models, on page 1](#)

Using YANG Data Models

Cisco IOS XR supports a programmatic way of configuring and collecting operational data of a network device using YANG data models. Although configurations using CLIs are easier and human-readable, automating the configuration using model-driven programmability results in scalability.

The data models are available in the release image, and are also published in the [Github](#) repository. Navigate to the release folder of interest to view the list of supported data models and their definitions. Each data model defines a complete and cohesive model, or augments an existing data model with additional XPathS. To view a comprehensive list of the data models supported in a release, navigate to the **Available-Content.md** file in the repository.

You can also view the data model definitions using the [YANG Data Models Navigator](#) tool. This GUI-based and easy-to-use tool helps you explore the nuances of the data model and view the dependencies between various containers in the model. You can view the list of models supported across Cisco IOS XR releases and platforms, locate a specific model, view the containers and their respective lists, leaves, and leaf lists presented visually in a tree structure. This visual tree form helps you get insights into nodes that can help you automate your network.

To get started with using the data models, see the *Programmability Configuration Guide*.



CHAPTER 2

Scale-Up Your Network Monitoring Strategy Using Telemetry

Are you monitoring your network using traditional polling methods such as SNMP, Syslog, and CLI? If yes, does the data that you extract from your network help you answer these questions?

- What percentage of the network bandwidth does the network traffic currently consume?
- Do all the links in the network run at a hundred percent utilization rate?
- If an unmanned router fails, is the network operator notified in real time about the issue and its related consequences?
- Is the CPU over- or under-utilized?
- Can the efficiency of the network be calculated based on traffic and data loss?
- What are the possible performance issues that cause traffic loss or network latency?
- How do you proactively prevent issues that may arise? Does the data support the study of network patterns in real time?


These traditional methods use a *pull* model to request information at regular intervals. The data that you collect may help you to efficiently monitor your network of a manageable size. However, as your network grows in complexity and scale, the data that you poll may be insufficient for efficient and effective monitoring. Additionally, the polling methods are resource-intensive, and network operators face information gaps in the data that they collect. With the pull model, the network device (the server) sends data only when the data collector (the client) requests it. Initiating such requests requires continual manual intervention. This manual intervention makes this model unsuitable, and limits automation and the ability to scale. It inhibits the visibility of the network and therefore provides inefficient control of the network. You need monitoring strategy that adds resiliency and stability to your network.

Telemetry does just that. Telemetry uses a *push* model that automatically streams data from a network device. Instead of a collector requesting data at periodic intervals, the network device streams operational data in real time.

Telemetry focuses on the power of scale, speed, and automation. With the power of flexibility, you can select data of interest from the routers and transmit it in a structured format to remote management stations for monitoring. Using the finer granularity and higher frequency of data available through telemetry, DevOps (development and operations) engineers in your organization can quickly locate and investigate issues as soon as they occur. They can, thus, collaborate to monitor and have better control over the network.

The following image shows the comparative benefits of streaming telemetry data using the telemetry push model over traditional pull models. The pull models create resource bottlenecks that prevent retrieving valuable operational data from the router. On the other hand, the push model is designed to remove such bottlenecks and deliver data efficiently.

Figure 1: Comparison Between Traditional Pull Models and Telemetry Push Model



Watch this [video](#) to see how telemetry data can unlock the intelligence of data in your network to proactively predict and troubleshoot issues.



Note Starting from Cisco IOS XR, Release 7.0.1, Telemetry is part of the base image (<platform>-mini-x.iso). In earlier releases, Telemetry was part of the Manageability package (<platform>-mgbl-3.0.0.0-<release>.x86_64.rpm).

This article describes the benefits of using telemetry data and the various methods to stream meaningful data from your network device:

- [Benefits of Shifting Network Monitoring from Pull Models to Telemetry Push Model](#), on page 5
- [Review Mechanisms to Stream Telemetry Data from a Router to a Destination](#), on page 5
- [Learn About the Elements that Enable Streaming Telemetry Data](#), on page 7
- [Explore the Methods to Establish a Telemetry Session](#), on page 11

Benefits of Shifting Network Monitoring from Pull Models to Telemetry Push Model

Real-time telemetry data is useful in:


- **Managing network remotely:** The primary benefit of telemetry is the ability it offers you as an end user to monitor the state of a network element remotely. After the network is deployed, you cannot be physically present at the network site to find out what works, and what is cumbersome. With telemetry, those insights can be analyzed, leveraged, and acted upon from a remote location.
- **Optimizing traffic:** When link utilization and packet drops in a network are monitored at frequent intervals, 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 traditional SNMP poll interval mechanism. Streaming telemetry data helps in providing quick response time for faster transport of traffic.
- **Preventive troubleshooting:** 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. The finer granularity and higher frequency of data available through telemetry enables better performance monitoring and therefore, better troubleshooting.
- **Visualizing data:** Telemetry data acts as a data lake that analytics toolchains and applications use to visualize valuable insights into your network deployments.
- **Monitoring and controlling distributed devices:** The monitoring function is decoupled from the storage and analysis functions. This decoupling helps to reduce device dependency, while providing flexibility to transform data using [pipelines](#). These pipelines are utilities that consume telemetry data, transform it, and forward the resulting content to a downstream, typically off-the-shelf, consumer. The supported downstream consumers include Apache Kafka, Influxdata, Prometheus, and Grafana.

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.

Review Mechanisms to Stream Telemetry Data from a Router to a Destination

Telemetry data can be streamed using either cadence-driven or event-driven mechanisms.

Figure 2: Cadence-driven and Event-driven Telemetry




Cadence-driven Telemetry

Cadence-driven telemetry continually streams data (operational statistics and state transitions) at a configured cadence. The higher frequency of the data that is continuously streamed helps you closely identify emerging patterns in the network.

The following image shows a continuous stream of data after a configured time interval:

Figure 3: Cadence-driven Telemetry




Event-driven Telemetry

Event-driven telemetry optimizes data that is collected at the receiver and streams data only when a state transition occurs and thus optimizes data that is collected at the receiver. For example, EDT streams data about interface state transitions, IP route updates, and so on.

The following image shows a stream of data after a state change:

Figure 4: Event-driven Telemetry



Learn About the Elements that Enable Streaming Telemetry Data

These elements are the building blocks in enabling telemetry in a network.

Sensor Path

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

A YANG module defines a data model through the data of the router, and the hierarchical organization and constraints on that data.

YANG defines four node types. Each node has a name. Depending on the node type, the node either defines a value or contains a set of child nodes. The nodes types for data modeling are:

- leaf node - contains a single value of a specific type
- leaf-list node - contains a sequence of leaf nodes
- list node - contains a sequence of leaf-list entries, each of which is uniquely identified by one or more key leaves
- container node - contains a grouping of related nodes that have only child nodes, which can be any of the four node types

To get started with using the data models, see the *Programmability Configuration Guide*.

The following table shows few examples of sensor paths:

Table 1: Sensor Paths

| Feature | Sensor Path |
|---------|---|
| CPU | Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization |

| Feature | Sensor Path |
|---|--|
| Memory | Cisco-IOS-XR-nto-misc-oper:memory-summary/nodes/node/summary |
| Interface | Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/latest/generic-counters Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/data-rate openconfig-interfaces:interfaces/interface |
| Node summary | Cisco-IOS-XR-nto-misc-oper:memory-summary/nodes/node/summary |
| Forwarding information base (FIB) | Cisco-IOS-XR-fib-common-oper:fib-statistics/nodes/node/drops Cisco-IOS-XR-fib-common-oper:fib/nodes/node/protocols/protocol/vrfs/vrf/summary |
| MPLS Traffic engineering (MPLS-TE) | Cisco-IOS-XR-mpls-te-oper:mpls-te/tunnels/summary Cisco-IOS-XR-ip-rsvp-oper:rsvp/interface-briefs/interface-brief Cisco-IOS-XR-mpls-te-oper:mpls-te/fast-reroute/protectations/protection Cisco-IOS-XR-mpls-te-oper:mpls-te/signalling-counters/signalling-summary Cisco-IOS-XR-mpls-te-oper:mpls-te/p2p-p2mp-tunnel/tunnel-heads/tunnel-head |
| MPLS Label distribution protocol (MPLS-LDP) | Cisco-IOS-XR-mpls-ldp-oper:mpls-ldp/nodes/node/bindings-summary-all Cisco-IOS-XR-mpls-ldp-oper:mpls-ldp/global/active/default-vrf/summary Cisco-IOS-XR-mpls-ldp-oper:mpls-ldp/nodes/node/default-vrf/neighbors/neighbor |
| Routing | Cisco-IOS-XR-clns-isis-oper:isis/instances/instance/statistics-global Cisco-IOS-XR-clns-isis-oper:isis/instances/instance/neighbors/neighbor Cisco-IOS-XR-ip-rib-ipv4-oper:rib/rib-table-ids/rib-table-id/summary-protos/summary-proto Cisco-IOS-XR-clns-isis-oper:isis/instances/instance/levels/level/adjacencies/adjacency Cisco-IOS-XR-ipv4-bgp-oper:bgp/instances/instance/instance-active/default-vrf/process-info Cisco-IOS-XR-ip-rib-ipv6-oper:ipv6-rib/rib-table-ids/rib-table-id/summary-protos/summary-proto |



Note Use specific paths to avoid streaming data that you may not be interested. For example, if you want to stream information about only the summary of MPLS-TE, use `sensor-path Cisco-IOS-XR-mpls-te-oper:mpls-te/autotunnel/mesh/summary` instead of `sensor-path Cisco-IOS-XR-mpls-te-oper:mpls-te` sensor path.

In the sensor path configuration, the schema node identifier can be configured with or without a leading slash.

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-driven Telemetry, on page 6](#)), or when an event occurs ([Event-driven Telemetry, on page 6](#)), to one or more collectors through subscribed sessions.

Subscription

A subscription binds one or more sensor paths and destinations.

The collector uses the subscription to receive updates about the state of data on the router. A subscription can consist of one or more sensor paths. The data for the paths that you have subscribed starts streaming until the session is terminated by the collector or the telemetry subscription configuration is removed to cancel the subscription.

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
```

Encoder

Data that is streamed from a router can be encoded using one of these formats:

- **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 at Cisco Network Telemetry Proto in Github.
 - **Compact GPB encoding:** Data is streamed in a compressed format and not in a self-descriptive format. A `.proto` file corresponding to each sensor-path must be used by the collector to decode the streamed data.
 - **Self-describing GPB encoding:** Data streamed for each sensor path is in a self-describing and ASCII text format. A single `.proto` file, `telemetry.proto`, is used by the collector to decode any sensor path data. Self-describing GPB encoding is easier to manage because it needs single `.proto` file to decode any sensor path data, even though the message size is large.
- **JSON encoding:** Data is streamed in strings of keys and its values in a human-readable format.

Transport

In the telemetry push model, the router streams telemetry data using a transport protocol. The generated data is encapsulated into the desired format using encoders.

Model-Driven Telemetry (MDT) data is streamed through these supported transport protocols:

- Google Protocol RPC (gRPC): used for both dial-in and dial-out modes.



Note gRPC protocol is not supported over Multiprotocol Label Switching (MPLS) including `explicit-null` label.

- Transmission Control Protocol (TCP): used for only dial-out mode.
- User Datagram Protocol (UDP): used for only dial-out mode. Because UDP is connectionless, the UDP destination is shown as `Active` irrespective of the state of the collector. This is not ideally suitable for a busy network. If a message is dropped by the network before it reaches the collector, the protocol does not resend the data.



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.

gRPC Network Management Interface

gRPC Network Management Interface (gNMI) is a gRPC-based network management protocol used to modify, install or delete configuration from network devices. It is also used to view operational data, control and generate telemetry streams from a target device to a data collection system. It uses a single protocol to manage configurations and stream telemetry data from network devices.

For the list of gNMI RPCs, see the *Programmability Configuration Guide*.

gRPC Network Operations Interface

gRPC Network Operations Interface (gNOI) defines a set of gRPC-based microservices for executing operational commands on network devices. Extensible Manageability Services (EMS) gNOI is the Cisco IOS XR implementation of gNOI. gNOI uses gRPC as the transport protocol and the configuration is same as that of gRPC.

For the list of gNOI RPCs, see the *Programmability Configuration Guide*.

TLS Authentication For Dial-in

The gRPC protocol supports Transport Layer Security (TLS) for encrypting data. By default, model-driven telemetry uses TLS to dial-out.

When TLS is enabled, the server sends a certificate to authenticate it with the collector. The collector validates the certificate verifying which certificate authority has signed it and generates session keys to encrypt the session.

The TLS certificate must be copied at the `/misc/config/grpc/dialout/` path. If only the `protocol grpc` command is configured, by default, TLS is enabled and the hostname defaults to the IP address of the destination. In addition, in the certificate, configure the Common Name (CN) as `protocol grpc tls-hostname` <>.

The following output shows the certificate that gRPC uses to establish a dialout session:

```
Router#run
[node:]$ls -l /misc/config/grpc/dialout/
total 4
-rw-r--r-- 1 root root 4017 dialout.pem
```

To bypass the TLS option, use **grpc no-tls** command.



Note Although TLS provides secure communication between servers and clients, TLS version 1.0 may pose a security threat. You can now disable TLS version 1.0 using the **grpc tlsv1-disable** command.

Explore the Methods to Establish a Telemetry Session

A telemetry session can be initiated using: either the dial-out mode or the dial-in mode. Although the modes to establish a telemetry session are different, both modes use the same data model and stream the same data.

Dial-Out Mode

In a *dial-out* mode, the router dials out to the receiver to establish a subscription-based telemetry session. Because the router initiates the connection, there is no need to manage the ports for inbound traffic. In this default mode of operation, the protocols you use to establish a session gives you the flexibility to choose between simplicity (TCP) and security (gRPC). A simple protocol requires only accessibility to the socket on the collector. A secure protocol, additionally, offers security capabilities to authenticate and encrypt the session. You can, therefore, secure your collector, and establish a much advanced method of communication with the router. If the connection between the router and the destination is lost, the router re-establishes the connection with the destination and continues to push data again. However, data transmitted during the time of reconnection is lost.

To explore the dial-out mode, and to create a dial-out session, see [Establish a Model-Driven Telemetry Session from a Router to a Collector, on page 13](#).

Dial-In Mode

In a *dial-in* mode, a collector dials in to the router, and subscribes dynamically to one or more sensor paths specified in a subscription. The router is open for connections from the collector. This mode is useful to establish a single channel of communication with the router. Because the collector establishes the session, there is no need to create destinations in the configuration. Additionally, the protocol (gRPC) used to establish a session provides advanced security capabilities to authenticate and encrypt the session. If the connection between the router and the collector is lost, the session is cancelled. The collector must reconnect to the router to restart streaming data. Only gRPC supports dial-in session.

To explore the dial-in mode, and to create a dial-in session, see [Establish a Model-Driven Telemetry Session from a Collector to a Router, on page 23](#).

Identify the Telemetry Session Suitable for Your Network

The transport protocols and encoding formats in your network help you determine which mode is suitable for your needs. The encoding efficiency is determined by the space that data occupies on the wire, memory utilization, and the amount of data that you plan to stream from the router.

- Use TCP dial-out mode if you plan to stream telemetry data using a simple setup with a single router and collector. It is simple to configure and does not require extensive knowledge about protocols. It removes the need to manage ports for inbound connections.
- Use gRPC dial-out mode if your setup involves scaling out to many devices or needs encryption of your data. This mode removes the need to manage ports for inbound connections.
- Use gRPC dial-in mode if you are already using gRPC in your network and you want your sessions to be dynamic without having the data streamed to fixed destinations. This mode is convenient if you prefer a centralized way configuring your network and requesting operational data.



CHAPTER 3

Establish a Model-Driven Telemetry Session from a Router to a Collector

Streaming telemetry is a new paradigm in monitoring the health of a network. It provides a mechanism to efficiently stream configuration and operational data of interest from Cisco IOS XR routers. This streamed data is transmitted in a structured format to remote management stations for monitoring and troubleshooting purposes.

With telemetry data, you create a data lake. Analyzing this data, you proactively monitor your network, monitor utilization of CPU and memory, identify patterns, troubleshoot your network in a predictive manner, and devise strategies to create a resilient network using automation.

Telemetry works on a [Subscription](#) model where you subscribe to the data of interest in the form of [Sensor Path](#). The sensor paths describes [OpenConfig data models](#) or native Cisco data models. You can access the [OpenConfig](#) and [Native](#) data models for telemetry from Github, a software development platform that provides hosting services for version control. You choose who initiates the subscription by establishing a telemetry session between the router and the receiver. The session is established using either a [Dial-Out Mode](#) or a [Dial-In Mode](#), described in the [Scale-Up Your Network Monitoring Strategy Using Telemetry](#) article.




Note Watch this [video](#) to discover the power of real-time network management using model-driven telemetry.

This article describes the dial-out mode where the router dials out to the receiver to establish a telemetry session. In this mode, destinations and sensor-paths 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. Even when a session terminates, the router continually attempts to re-establish a new session with the receiver at regular intervals.

The following image shows a high-level overview of the dial-out mode:

Figure 5: Dial-Out Mode



This article describes, with a use case that illustrates the monitoring of CPU utilization, how streaming telemetry data helps you gain better visibility of your network, and make informed decisions to stabilize your network.

**Tip**

You can programmatically configure a dial-out telemetry session using `openconfig-telemetry.yang` OpenConfig data model. To get started with using data models, see the *Programmability Configuration Guide*.

- [Monitor CPU Utilization Using Telemetry Data to Plan Network Infrastructure](#), on page 14

Monitor CPU Utilization Using Telemetry Data to Plan Network Infrastructure

The use case illustrates how, with the [Dial-Out Mode](#), you can use telemetry data to proactively monitor CPU utilization. Monitoring CPU utilization ensures efficient storage capabilities in your network. This use case describes the tools used in the open-sourced collection stack to store and analyse telemetry data.

**Note**

Watch this [video](#) to see how you configure model-driven telemetry to take advantage of data models, open source collectors, encodings and integrate into monitoring tools.

Telemetry involves the following workflow:

- **Define:** You define a subscription to stream data from the router to the receiver. To define a subscription, you create a destination-group and a sensor-group.
- **Deploy:** The router establishes a subscription-based telemetry session and streams data to the receiver. You verify subscription deployment on the router.
- **Operate:** You consume and analyse telemetry data using open-source tools, and take necessary actions based on the analysis.

Before you begin

Make sure you have L3 connectivity between the router and the receiver.

Define a Subscription to Stream Data from Router to Receiver

Create a subscription to define the data of interest to be streamed from the router to the destination.

Step 1

Create one or more destinations to collect telemetry data from a router. Define a destination-group to contain the details about the destinations. Include the destination address (ipv4 or ipv6), port, transport, and encoding format in the destination-group.

Example:**Create a destination-group using data model**

This example uses the native data model `Cisco-IOS-XR-um-telemetry-model-driven-cfg.yang`.

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101">
  <get-config>
    <source>
      <candidate/>
    </source>
    <filter>
      <telemetry-model-driven
xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-um-telemetry-model-driven-cfg">
        <destination-groups>
          <destination-group>
            <destination-id>CPU-Health</destination-id>
            <ipv4-destinations>
              <ipv4-destination>
                <ipv4-address>172.0.0.0</ipv4-address>
                <destination-port>57500</destination-port>
                <encoding>self-describing-gpb</encoding>
                <protocol>
                  <protocol>tcp</protocol>
                </protocol>
              </ipv4-destination>
            </ipv4-destinations>
          </destination-group>
        </destination-groups>
      </telemetry-model-driven>
    </filter>
  </get-config>
</rpc>
```

Create a destination group using CLI

```
##Configuration with tls-hostname##
Router(config)#telemetry model-driven
Router(config-model-driven)#destination-group CPU-Health
Router(config-model-driven-dest)#address family ipv4 172.0.0.0 port 57500
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
```

where -

- CPU-Health is the name of the destination-group

Define a Subscription to Stream Data from Router to Receiver

- 172.0.0.0 is the IP address of the destination where data is to be streamed
- 57500 is the port number of the destination
- self-describing-gpb is the format in which data is encoded and streamed to the destination
- tcp is the protocol through which data is transported to the destination.

Step 2 Specify the subset of the data that you want to stream from the router using sensor paths. The [Sensor Path](#) represents the path in the hierarchy of a YANG data model. Create a sensor-group to contain the sensor paths.

Example:**Create a sensor-group for CPU utilization using data model**

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <telemetry-model-driven
xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-um-telemetry-model-driven-cfg">
        <sensor-groups>
          <sensor-group>
            <sensor-group-identifier>Monitor-CPU</sensor-group-identifier>
            <sensor-paths>
              <sensor-path>

<telemetry-sensor-path>Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization</telemetry-sensor-path>

              </sensor-path>
            </sensor-paths>
          </sensor-group>
        </sensor-groups>
      </telemetry-model-driven>
    </config>
  </edit-config>
</rpc>
```

Create a sensor-group for CPU utilization using CLI

```
Router(config)#telemetry model-driven
Router(config-model-driven)#sensor-group Monitor-CPU
Router(config-model-driven-snsr-grp)# sensor-path
Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization
Router(config-model-driven-snsr-grp)# commit
```

where -

- Monitor-CPU is the name of the sensor-group
- Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization is the sensor path from where data is streamed.

Step 3 Subscribe to telemetry data that is streamed from a router. A [Subscription](#) binds the destination-group with the sensor-group and sets the streaming method. The streaming method can be [Cadence-driven Telemetry](#) or [Event-driven Telemetry](#).

Example:

Note The configuration for event-driven telemetry is similar to cadence-driven telemetry, with only the sample interval as the differentiator. Configuring the sample interval value to 0, zero, sets the subscription for event-driven telemetry, while configuring the interval to any non-zero value sets the subscription for cadence-driven telemetry.

Create a subscription using data model

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="101">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <telemetry-model-driven
xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-um-telemetry-model-driven-cfg">
        <subscriptions>
          <subscription>
            <subscription-identifier>CPU-Utilization</subscription-identifier>
            <sensor-profiles>
              <sensor-profile>
                <sensorgroupid>Monitor-CPU</sensorgroupid>
                <sample-interval>30000</sample-interval>
              </sensor-profile>
            </sensor-profiles>
            <destination-profiles>
              <destination-profile>
                <destination-id>CPU-Health</destination-id>
              </destination-profile>
            </destination-profiles>
            <source-interface>Interface1</source-interface>
          </subscription>
        </subscriptions>
      </telemetry-model-driven>
    </config>
  </edit-config>
</rpc>
```

Create a subscription using CLI

```
Router(config)#telemetry model-driven
Router(config-model-driven)#subscription CPU-Utilization
Router(config-model-driven-subs)#sensor-group-id Monitor-CPU sample-interval 30000
Router(config-model-driven-subs)#destination-id CPU-Health
Router(config-model-driven-subs)#source-interface Interface1
Router(config-model-driven-subs)#commit
```

where -

- CPU-Utilization is the name of the subscription
- Monitor-CPU is the name of the sensor-group
- CPU-Health is the name of the destination-group
- Interface1 is the source interface that is used for establishing the telemetry session. If both the VRF and source interface are configured, the source interface must be in the same VRF as the one specified in the destination group.
- 30000 is the sample interval in milliseconds. The sample interval is the time interval between two streams of data. In this example, the sample interval is 30000 milliseconds or 30 seconds.

Verify Deployment of the Subscription

The router dials out to the receiver to establish a session with each destination in the subscription. After the session is established, the router streams data to the receiver to create a data lake.

You can verify the deployment of the subscription on the router.

Step 1 View the model-driven telemetry configuration on the router.

Example:

```
Router#show running-config telemetry model-driven
telemetry model-driven
destination-group CPU-Health
address-family ipv4 172.0.0.0 port 57500
encoding self-describing-gpb
protocol tcp
!
sensor-group Monitor-CPU
sensor-path
Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization
!
subscription CPU-Utilization
sensor-group-id Monitor-CPU sample-interval 30000
destination-id CPU-Health
source-interface GigabitEthernet0/0/0/0
!
!
```

Step 2 Verify the state of the subscription. An **Active** state indicates that the router is ready to stream data to the receiver based on the subscription.

Example:

```
Router# show telemetry model-driven subscription CPU-Utilization

Subscription: CPU-Utilization
-----
State: NA
Source Interface: GigabitEthernet0_0_0_0( 0x0)
Sensor groups:
Id: Monitor-CPU
Sample Interval: 30000 ms
Sensor Path: Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization
Sensor Path State: Resolved

Destination Groups:
Group Id: CPU-Health
Destination IP: 172.0.0.0
Destination Port: 57500
Encoding: self-describing-gpb
Transport: tcp
State: NA
No TLS

Collection Groups:
-----
No active collection groups
```

The router streams data to the receiver using the subscription-based telemetry session and creates a data lake in the receiver.

Operate on Telemetry Data for In-depth Analysis of the Network

You can start consuming and analyzing telemetry data from the data lake using an open-sourced collection stack. This use case uses the following tools from the collection stack:

- Pipeline is a lightweight tool used to collect data. You can download [Network Telemetry Pipeline](#) from Github. You define how you want the collector to interact with routers and where you want to send the processed data using `pipeline.conf` file.
- Telegraph (plugin-driven server agent) and InfluxDB (a time series database (TSDB)) stores telemetry data, which is retrieved by visualization tools. You can download [InfluxDB](#) from Github. You define what data you want to include into your TSDB using the `metrics.json` file.
- [Grafana](#) is a visualization tool that displays graphs and counters for data streamed from the router.

In summary, Pipeline accepts TCP and gRPC telemetry streams, converts data and pushes data to the InfluxDB database. Grafana uses the data from InfluxDB database to build dashboards and graphs. Pipeline and InfluxDB may run on the same server or on different servers.

Consider that the router is streaming data of approximately 350 counters every 5 seconds, and Telegraf requests information from the Pipeline at 1 second intervals. The CPU usage is analysed in three stages using:

- a single router to get initial values
- two routers to find the difference in values and understand the pattern
- five routers to arrive at a proof-based conclusion

This helps you make informed business decisions about deploying the infrastructure; in this case, the CPU.

Step 1 Start Pipeline, and enter your router credentials.

Note The IP address and port that you specify in the destination-group must match the IP address and port on which Pipeline is listening.

Example:

```
$ bin/pipeline -config pipeline.conf

Startup pipeline
Load config from [pipeline.conf], logging in [pipeline.log]

CRYPT Client [grpc_in_mydmtrouter], [http://172.0.0.0:5432]
Enter username: <username>
Enter password: <password>
Wait for ^C to shutdown
```

Step 2 In the Telegraph configuration file, add the following values to read the metrics about CPU usage.

Example:

```
[[inputs.cpu]]
  ## Whether to report per-cpu stats or not
  percpu = true
  ## Whether to report total system cpu stats or not
  totalcpu = true
  ## If true, collect raw CPU time metrics.
  collect_cpu_time = false
  ## If true, compute and report the sum of all non-idle CPU states.
  report_active = false
```


Step 3 Use Grafana to create a dashboard and visualize data about CPU usage.

One router

The router pushes the counters every five seconds.


All CPU cores are loaded equally, and there are spikes up to approximately 10 or 11 percent.

Figure 6: CPU Usage Graph with a Single Router




Two routers

The second router is added at 14:00 in the timeline, and shows an increase in the spikes to around 25 percent with midpoint value at 15 percent.

Figure 7: CPU Usage Graph with Two Routers

Five routers

With five routers, the spikes peak upto approximately 40 percent with midpoint in the range of 22 to 25 percent.

Figure 8: CPU Usage Graph with Five Routers

In conclusion, telemetry data shows that the processes are balanced almost equally across the CPU cores. There is no linear increase on a subset of cores. This analysis helps in planning the CPU utilization based on the number of counters that you stream.



CHAPTER 4

Establish a Model-Driven Telemetry Session from a Collector to a Router

Streaming telemetry is a new paradigm in monitoring the health of a network. It provides a mechanism to efficiently stream configuration and operational data of interest from Cisco IOS XR routers. This streamed data is transmitted in a structured format to remote management stations for monitoring and troubleshooting purposes.

With telemetry data, you create a data lake. Analyzing this data, you proactively monitor your network, monitor utilization of CPU and memory, identify patterns, troubleshoot your network in a predictive manner, and devise strategies to create a resilient network using automation.

Telemetry works on a [Subscription](#) model where you subscribe to the data of interest in the form of [Sensor Path](#). The sensor paths describes [OpenConfig data models](#) or native Cisco data models. You can access the [OpenConfig](#) and [Native](#) data models for telemetry from Github, a software development platform that provides hosting services for version control. You choose who initiates the subscription by establishing a telemetry session between the router and the receiver. The session is established using either a [Dial-Out Mode](#) or [Dial-In Mode](#), described in the [Scale-Up Your Network Monitoring Strategy Using Telemetry](#) article.




Note Watch this [video](#) to discover the power of real-time network management using model-driven telemetry.

This article describes the dial-in mode where a receiver dials in to the router to establish a telemetry session. In this mode, the receiver dials in to the router, and subscribes dynamically to one or more sensor paths specified in a subscription. The router streams telemetry data through the same session that is established by the receiver. The dial-in mode of subscriptions is dynamic. This dynamic subscription terminates when the receiver cancels the subscription or when the session terminates.

The following image shows a high-level overview of the dial-in mode:

Figure 9: Dial-In Mode



This article describes, with a use case that illustrates the simultaneous monitoring of various parameters in the network, how streaming telemetry data helps you gain better visibility of the network, and make informed decisions to stabilize it.



YANG Data Model

You can programmatically configure a dial-in telemetry session using `openconfig-telemetry.yang` OpenConfig data model. To get started with using data models, see the *Programmability Configuration Guide*.

- [Monitor Network Parameters Using Telemetry Data for Proactive Analysis](#), on page 24

Monitor Network Parameters Using Telemetry Data for Proactive Analysis

The use case illustrates how, with the [Dial-In Mode](#), you can use telemetry data to stream various parameters about your network. You use this data for predictive analysis where you monitor patterns, and proactively troubleshoot issues. This use case describes the tools used in the open-sourced collection stack to store and analyse telemetry data.



Note Watch this [video](#) to see how you configure model-driven telemetry to take advantage of data models, open source collectors, encodings and integrate into monitoring tools.

Telemetry involves the following workflow:

- **Define:** You define a subscription to stream data from the router to the receiver. To define a subscription, you create a sensor-group.

- **Deploy:** The receiver initiates a session with the router and establishes a subscription-based telemetry session. The router streams data to the receiver. You verify subscription deployment on the router.
- **Operate:** You consume and analyse telemetry data using open-source tools, and take necessary actions based on the analysis.

Before you begin

Ensure you meet these dependencies:

- Make sure you have L3 connectivity between the router and the receiver.
- Enable gRPC server on the router to accept incoming connections from the receiver.

```
Router#configure
Router(config)#grpc
Router(config-grpc)#port <port-number>
Router(config-grpc)#commit
```

The port-number ranges from 57344 to 57999. If a port number is unavailable, an error is displayed.

Define a Subscription to Stream Data from Router to Receiver

Create a subscription to define the data of interest to be streamed from the router to the destination.

Step 1

Specify the subset of the data that you want to stream from the router using sensor paths. The [Sensor Path](#) represents the path in the hierarchy of a YANG data model. This example uses the native data model

Cisco-IOS-XR-um-telemetry-model-driven-cfg.yang. Create a sensor-group to contain the sensor paths.

Example:

```
sensor-group health
  sensor-path Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization
  sensor-path Cisco-IOS-XR-nto-misc-oper:memory-summary/nodes/node/summary
  sensor-path Cisco-IOS-XR-shellutil-oper:system-time/uptime
!
sensor-group interfaces
  sensor-path
Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/latest/generic-counters
  sensor-path Cisco-IOS-XR-pfi-im-cmd-oper:interfaces/interface-summary
!
sensor-group optics
  sensor-path Cisco-IOS-XR-controller-optics-oper:optics-oper/optics-ports/optics-port/optics-info
!
sensor-group routing
  sensor-path Cisco-IOS-XR-clns-isis-oper:isis/instances/instance/levels/level/adjacencies/adjacency

  sensor-path Cisco-IOS-XR-clns-isis-oper:isis/instances/instance/statistics-global
  sensor-path
Cisco-IOS-XR-ip-rib-ipv4-oper:rib/vrfs/vrf/afs/af/safs/saf/ip-rib-route-table-names/ip-rib-route-table-name/protocol/isis/as/information

  sensor-path Cisco-IOS-XR-ipv4-bgp-oper:bgp/instances/instance/instance-active/default-vrf/process-info
!
sensor-group mpls-te
  sensor-path Cisco-IOS-XR-mpls-te-oper:mpls-te/tunnels/summary
  sensor-path Cisco-IOS-XR-ip-rsvp-oper:rsvp/interface-briefs/interface-brief
```

```
sensor-path Cisco-IOS-XR-ip-rsvp-oper:rsvp/counters/interface-messages/interface-message
!
```

Step 2

Subscribe to telemetry data that is streamed from a router. A [Subscription](#) binds the sensor-group, and sets the streaming method. The streaming method can be [Cadence-driven Telemetry](#) or [Event-driven Telemetry](#). Separating the sensor-paths into different subscriptions enhances the efficiency of the router to retrieve operational data at scale.

Example:

Note The configuration for event-driven telemetry is similar to cadence-driven telemetry, with only the sample interval as the differentiator. Configuring the sample interval value to 0 (zero), sets the subscription for event-driven telemetry, while configuring the interval to any non-zero value sets the subscription for cadence-driven telemetry.

```
subscription health
  sensor-group-id health strict-timer
  sensor-group-id health sample-interval 30000
!
subscription interfaces
  sensor-group-id interfaces strict-timer
  sensor-group-id interfaces sample-interval 30000
!
subscription optics
  sensor-group-id optics strict-timer
  sensor-group-id optics sample-interval 30000
!
subscription routing
  sensor-group-id routing strict-timer
  sensor-group-id routing sample-interval 30000
!
subscription mpls-te
  sensor-group-id mpls-te strict-timer
  sensor-group-id mpls-te sample-interval 30000
!
```

Verify Deployment of the Subscription

The receiver dials into the router to establish a dynamic session based on the subscription. After the session is established, the router streams data to the receiver to create a data lake.

You can verify the deployment of the subscription on the router.

Verify the state of the subscription. An `Active` state indicates that the router is ready to stream data to the receiver based on the subscription.

Example:

```
Router#show telemetry model-driven subscription
Thu Jan 16 09:48:14.293 UTC
Subscription: health                               State: Active
-----
Sensor groups:
Id              Interval (ms)      State
health          30000              Resolved

Subscription: optics                               State: NA
-----
```

```

Sensor groups:
Id          Interval(ms)  State
optics      30000         Resolved

Subscription: mpls-te          State: NA
-----
Sensor groups:
Id          Interval(ms)  State
mpls-te     30000         Resolved

Subscription: routing          State: NA
-----
Sensor groups:
Id          Interval(ms)  State
routing     30000         Resolved

Subscription: interfaces       State: NA
-----
Sensor groups:
Id          Interval(ms)  State
interfaces  30000         Resolved

Subscription: CPU-Utilization  State: NA
-----
Sensor groups:
Id          Interval(ms)  State
Monitor-CPU 30000         Resolved

Destination Groups:
Id          Encoding      Transport  State  Port  Vrf  IP
CPU-Health  self-describing-gpb tcp        NA     57500 Vrf  172.0.0.0
No TLS

```

The router streams data to the receiver using the subscription-based telemetry session and creates a data lake in the receiver.

Operate on Telemetry Data for In-depth Analysis of the Network

You can start consuming and analyzing telemetry data from the data lake using an open-sourced collection stack. This use case uses the following tools from the collection stack:

- Pipeline is a lightweight tool used to collect data. You can download [Network Telemetry Pipeline](#) from Github. You define how you want the collector to interact with routers, and where you want to send the processed data using `pipeline.conf` file.
- Telegraph or InfluxDB is a time series database (TSDB) that stores telemetry data, which is retrieved by visualization tools. You can download [InfluxDB](#) from Github. You define what data you want to include into your TSDB using the `metrics.json` file.
- [Grafana](#) is a visualization tool that displays graphs and counters for data streamed from the router.

In summary, Pipeline accepts TCP and gRPC telemetry streams, converts data and pushes data to the InfluxDB database. Grafana uses the data from InfluxDB database to build dashboards and graphs. Pipeline and InfluxDB may run on the same server or on different servers.

Consider that the router is monitored for the following parameters:

- Memory and CPU utilization

- Interface counters and interface summary
- Transmitter and receiver power levels from optic controllers
- ISIS route counts and ISIS interfaces
- BGP neighbours, path count, and prefix count
- MPLS-TE tunnel summary
- RSVP control messages and bandwidth allocation for each interface

Step 1 Start Pipeline from the shell, and enter your router credentials.

Example:

```
$ bin/pipeline -config pipeline.conf


Startup pipeline
Load config from [pipeline.conf], logging in [pipeline.log]

CRYPT Client [grpc_in_myndtrouter], [http://172.0.0.0:5432]
Enter username: <username>
Enter password: <password>
Wait for ^C to shutdown
```

The streamed telemetry data is stored in InfluxDB.


Step 2 Use Grafana to create a dashboard and visualize the streamed data.

Figure 10: Visual Analysis of Network Health using Telemetry Data



369765

Figure 11: Visual Analysis of System Monitoring using Telemetry Data



In conclusion, telemetry data shows that various parameters of the network can be monitored simultaneously. This data is streamed in near real-time without affecting the performance of the network. With this data, you gain better visibility into your network.



CHAPTER 5

Build Intelligence on the Router Using AI-Driven Telemetry

Table 2: Feature History Table

| Feature Name | Release Information | Description |
|---------------------------|---------------------|---|
| AI-driven telemetry (ADT) | Release 7.3.1 | <p>This feature leverages machine learning to detect and retrieve important network-state changes on the router. Relevant data is filtered and exported to the network management system for analysis or troubleshooting purposes.</p> <p>ADT significantly simplifies the configuration of streaming telemetry, and you are no longer required to manually choose sensor paths or tune the cadence at which counters have to be collected.</p> |

Cisco IOS XR offers a rich set of sensor paths that help you stream telemetry data about your network using [Establish a Model-Driven Telemetry Session from a Router to a Collector](#). However, this rich offering of sensor paths leads to the following questions:

- Which sensor paths do I have to subscribe to?
- When do I know that the state of the system has changed, and which parameters do I look for troubleshooting purposes?
- How do I map between my router configuration and the sensor paths that I must monitor?
- Can I get the state changes in an automated, unsupervised, and data-driven way?
- Can I automatically filter a small set of sensor paths that properly portray an event?

AI-Driven Telemetry (ADT) is the answer to all these questions.

This article describes ADT, a component of the Cisco IOS XR operating system. ADT leverages machine learning (ML) and artificial intelligence (AI) to detect and describe important state changes on the router, and export only the relevant data using telemetry.

ADT eases streaming telemetry data and does not require the complex configurations of MDT. MDT supports streaming telemetry data at a configured [Cadence-driven Telemetry](#), or when an [Event-driven Telemetry](#) occurs in the network. However, with MDT, not all telemetry data can be exported by the router at high frequencies. You must choose and subscribe to only a subset of the available information. Data availability does not mean that the data is easily consumable. You may miss some events generated within a cadence duration, either because you did not subscribe to the appropriate sensor paths or data models, or because the sampling rate was too coarse. On the other hand, you may be overwhelmed with redundant information, because a single event is represented through a large set of correlated counters.

ADT provides data-correlation and filtering on the router so that the router exports only the relevant data to network management systems (NMS). ADT correlates data by detecting important state changes on the router using a multivariate, unsupervised machine learning approach. ADT significantly simplifies the configuration of streaming telemetry. You are not required to choose the set of sensor paths (counters) to collect from the system, nor tune the cadence. ADT describes a change in the state of the router by choosing a small set of representative counters that best portray the change. The results of ADT change detection and description are provided as event-driven streaming telemetry using YANG model.

Consider a scenario where a bidirectional forwarding detection (BFD) session is enabled. Ideally, the system must filter a set of counters that are related to BFD. ADT uses automated feature selection process, wherein, in addition to streaming BFD-related counters, ADT might also highlight counters related to CPU. This is because BFD changes can trigger route re-calculation and eventually CPU occupancy. This shows that ADT learns the associated configuration in a system, and streams a representative set of sensor paths that get impacted with the state change. So, the use cases for ADT falls under a broad spectrum where any event that causes fluctuation at the interface-level traffic is detected in an unsupervised manner.

With this streamed telemetry data, a data lake is created at the collector. Analyzing this data, you proactively monitor your network, identify patterns, troubleshoot your network in a predictive manner, and devise strategies to create a resilient network using automation.

- [Key Components, on page 32](#)
- [Processing Telemetry Data on the Router to Analyze Traffic Changes, on page 33](#)

Key Components

ADT comprises of the following key components:

- **Collector:** Collects all data from the router to provide a comprehensive view of the system. This function is synonymous to MDT.
- **Detector:** Prepares the collected data, remove redundancies, and detect changes to provide a macroscopic view of the system unlike Event Driven Telemetry (EDT). The detector learns how the sensor paths evolve to depend on each other and monitors changes at the system level. This learning about dependencies is unsupervised, multi-variate and adaptive to the data stream.
- **Selector:** Chooses a small, representative subset of the set of counters that best portrays events that are detected by the Detector using AI and unsupervised ML.
- **Exporter:** Streams the selected sensor paths using the ADT YANG data model to the MDT infrastructure. The ADT data model supports EDT where data is streamed only when an event in the system is detected. With sample interval 0, there is no overhead due to high frequency push.

Processing Telemetry Data on the Router to Analyze Traffic Changes

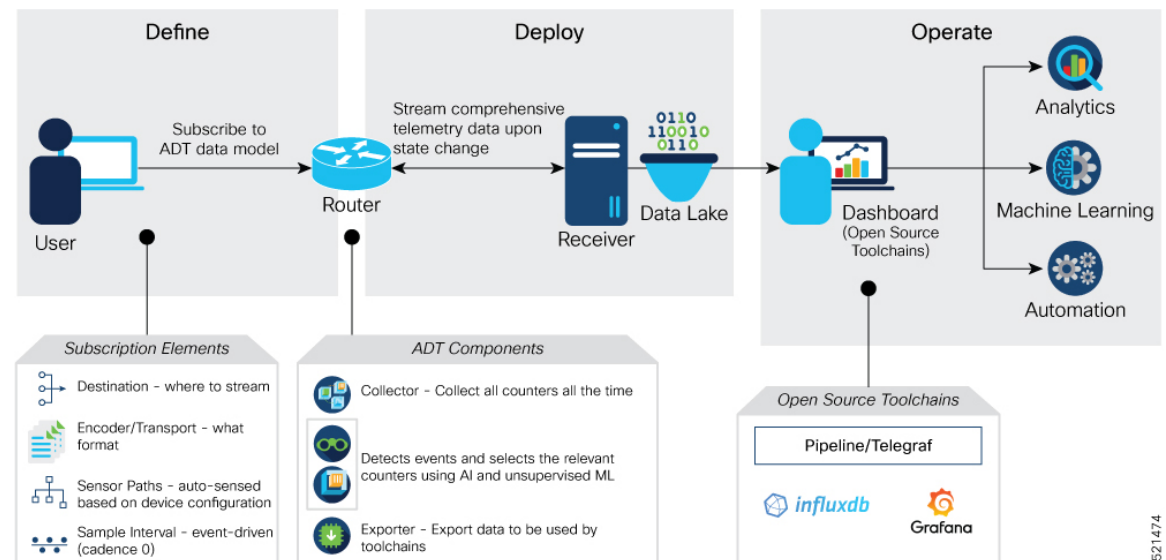
The use case illustrates how, with the AI-driven telemetry (ADT), you can use telemetry data to monitor events in the network. Monitoring the events ensures efficient network management. This use case describes the tools that are used in the open-sourced collection stack to store and analyse telemetry data.



Note Watch this [video](#) to see how you configure model-driven telemetry to take advantage of data models, Open Source collectors, encodings, and integrate into monitoring tools.

Telemetry involves the following workflow:

Figure 12: AI-Driven Telemetry



- **Define:** You define a subscription to stream data from the router to the receiver. To define a subscription, you create a destination-group and a sensor-group.
- **Deploy:** The router establishes a subscription-based telemetry session and streams data to the receiver. You verify subscription deployment on the router.
- **Operate:** You consume and analyze telemetry data using Open Source tools, and take necessary actions based on analysis.

Before you begin

Make sure you have L3 connectivity between the router and the receiver for external traffic.

Define a Subscription to Stream ADT Events from Router to Receiver

Create a subscription to define the data of interest to be streamed from the router to the destination.

Step 1 Enable telemetry on the router.

Example:

```
Router(config)#telemetry model-driven
Router(config)#commit
```

Step 2 Enable ADT on the router.

Example:

```
Router(config)#adt enable
Router(config)#commit
```

Once ADT is enabled, it monitors a set of sensor paths that are derived from the configuration on the router. ADT learns these sensor paths in an unsupervised machine learning approach.

Note The `Cisco-IOS-XR-adt-config-cfg.yang` YANG data model can be used to configure ADT using NETCONF protocol.

```

container adt-config {
  description "Container Schema adt configuration";

  container enable-feature {
    presence "CLI submode compatibility.";
    description "Enable adt feature";

    container input {
      description "Sources for ADT collector";

      container mdt {
        description "MDT config for collector";

        container mdt-sensor-group-ids {
          description "MDT sensor groups";

          list mdt-sensor-group-id {
            key "groupname";
            description "MDT sensor group id";
            leaf groupname {
              type xr:Cisco-ios-xr-string;
              description "Mdt Sensor Group name";
            }
          }
        }
      }
    }
  }
}

```

The following example shows the basic configuration to enable ADT using data model.

```

<config>
  <adt-config xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-adt-config-cfg">
    <enable-feature/>
  </adt-config>
</config>

```

To disable ADT, use the following operation:

```

<config>
  <adt-config xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-adt-config-cfg">
    <enable-feature xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
nc:operation="delete"/>
  </adt-config>
</config>

```

Step 3 Create one or more destinations to collect telemetry data from a router. Define a destination-group to contain the details about the destinations. Include the destination address (IPv4 or IPv6) or fully qualified domain name (FQDN), port, transport, and encoding format in the destination-group.

Example:

```

Router(config)#telemetry model-driven
Router(config-model-driven)#destination-group ADT-DESTINATION
Router(config-model-driven-dest)#address family ipv4 172.0.0.0 port 57500
Router(config-model-driven-dest-addr)#encoding self-describing-gpb
Router(config-model-driven-dest-addr)#protocol tcp

```

Where -

- ADT-DESTINATION is the name of the destination-group
- 172.0.0.0 is the IP address of the destination where data is to be streamed
- 57500 is the port number of the destination
- self-describing-gpb is the format in which data is encoded and streamed to the destination. For the supported formats, see [Encoder, on page 9](#).
- tcp is the protocol through which data is transported to the destination. For the supported protocols, see [Transport, on page 9](#).

Step 4 Specify the subset of the data that you want to stream from the router using sensor paths. The [Sensor Path](#) represents the path in the hierarchy of a YANG data model. Create a sensor-group to contain the sensor paths.

Example:

```
Router(config)#telemetry model-driven
Router(config-model-driven)#sensor-group ADT-EVENT
Router(config-model-driven-snsr-grp)#sensor-path Cisco-IOS-XR-adt-oper:adt/adt-output
```

Where -

- ADT-EVENT is the name of the sensor-group
- Cisco-IOS-XR-adt-oper:adt/adt-output is the sensor path from where data is streamed.

Note The following example showss how to add a custom sensor group using data model.

```
<config>
  <adt-config xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-adt-config-cfg">
    <enable-feature>
      <input>
        <mdt>
          <mdt-sensor-group-ids>
            <mdt-sensor-group-id>
              <groupname> QOS_VOQ </groupname>
            </mdt-sensor-group-id>
          </mdt-sensor-group-ids>
        </mdt>
      </input>
    </enable-feature>
  </adt-config>
</config>
```

Step 5 Subscribe to telemetry data that is streamed from a router. A [Subscription](#) binds the destination-group with the sensor-group and sets the streaming method. The streaming method can be [Cadence-driven Telemetry](#) or [Event-driven Telemetry](#). ADT is event-driven; subscriptions are defined with 0 (zero) cadence.

Example:

```
Router(config)#telemetry model-driven
Router(config-model-driven)#subscription ADT-SUBSCRIPTION
Router(config-model-driven-subs)#sensor-group-id ADT-EVENT sample-interval 0
Router(config-model-driven-subs)#destination-id ADT-DESTINATION
Router(config-model-driven-subs)#source-interface Interfacel
```

Where -

- ADT-SUBSCRIPTION is the name of the subscription
- ADT-EVENT is the name of the sensor-group
- ADT-DESTINATION is the name of the destination-group
- Interface1 is the source interface that is used for establishing the telemetry session. If both the VRF and source interface are configured, the source interface must be in the same VRF as the one specified in the destination-group.
- 0 is the sample interval.

Running Configuration

```

adt enable
!
telemetry model-driven
  destination-group ADT-DESTINATION
    destination 172.0.0.0 port 57500
    encoding self-describing-gpb
    protocol tcp
  !
!
  sensor-group ADT-EVENT
    sensor-path Cisco-IOS-XR-adt-oper:adt/adt-output
  !
  subscription ADT-SUBSCRIPTION
    sensor-group-id ADT-EVENT sample-interval 0
    destination-id ADT-DESTINATION
  !
!

```

Alternate Method: Configure ADT Using YANG Data Model

The above ADT configuration using CLI can also be configured using YANG data model Cisco-IOS-XR-adt-config-cfg.yang. You can obtain the data model from [Github](#).

```

<config>
  <adt-config xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-adt-config-cfg">
    <enable-feature/>
  </adt-config>
</config>

```

Verify Deployment of the Subscription

The router dials out to the receiver to establish a session with each destination in the subscription. After the session is established, the router streams data to the receiver to create a data lake.

ADT operational models are classified into three categories:

- **ADT events:** This model defines how ADT generates output to describe an event using Cisco-IOS-XR-adt-oper:adt/adt-output sensor path.
- **ADT subscriptions:** This model exports the sensor paths that are monitored by ADT and the cadence at which each path is monitored using Cisco-IOS-XR-adt-oper:adt/subscription-info sensor path.

```

Router#show adt subscription details
Wed Feb  3 14:51:39.444 IST

```

```

ADT SUBSCRIPTION Details
[Subscription ID, Cadence(in seconds), (E)xplicit/(I)mPLICIT] Sensor Path
*Subscription ID = 0: Not enough system resources to subscribe

Active Groups : 2

    Group: QOS_VOQ
    -----
    [300000028, 20, E]
    Cisco-IOS-XR-qos-ma-oper:qos/interface-table/interface/output/service-policy-names/service-policy-instance/statistics

    [300000027, 20, E]
    Cisco-IOS-XR-freta-bcm-dpa-npu-stats-oper:dpa/stats/nodes/node/npu-numbers/npu-number/display/base-numbers/base-number

    Group: implicit
    -----
    [300000001, 60, I]
    Cisco-IOS-XR-telemetry-model-driven-oper:telemetry-model-driven/subscriptions/subscription

    [300000002, 40, I]
    Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization
    [300000003, 40, I] Cisco-IOS-XR-nto-misc-oper:memory-summary/nodes/node/summary

    [300000004, 40, I]
    Cisco-IOS-XR-procmem-oper:processes-memory/nodes/node[node-name="0/RP0/CPU0"]/process-ids/process-id[process-id="8893"]

    [300000005, 20, I]
    Cisco-IOS-XR-pfi-im-cmd-oper:interfaces/node-type-sets/node-type-set/interface-summary

    [300000006, 20, I]
    Cisco-IOS-XR-infra-stats-oper:infra-statistics/interfaces/interface[interface-name=".*(GigE|Ethernet).*"]/latest/generic-counters

    [300000007, 20, I] Cisco-IOS-XR-ipv6-io-oper:ipv6-io
    [300000008, 20, I] Cisco-IOS-XR-ip-iarm-v6-oper:ipv6arm/router-id

    ----- Truncated for Brevity -----

```

- **ADT statistics:** This model reports health statistics of ADT system using `Cisco-IOS-XR-adt-oper:adt/statistics` sensor path.

In this example, the output is shown for ADT events. ADT event can have one or more events reported.

- Each event contains:
 - An event identifier
 - The time-stamp of the event
 - A short description of the event
 - List of sensors that changed their behavior during the event
- Each sensor path exported contains:
 - Tags which define the scope of the sensor path
 - List of value, timestamp pair containing samples of the sensor output before and after the event.

Step 1 View the generated events.

Example:

Following is a snippet of ADT event output, generated by traffic change.

```
"node_id_str": "PE4",
"subscription_id_str": "app_1887_75f00000001",
"encoding_path": "Cisco-IOS-XR-adt-oper:adt/adt-output",
"collection_id": "9569581",
"collection_start_time": "1607525488535",
"msg_timestamp": "1607525488556",
"data_json": [
  {
    "timestamp": "1607525488552",
    "keys": [],
    "content": {
      "adt-event": [
        {
          "event-id": ,
          "change-description": "Traffic",
          "timestamp": "1607431905419",
          "change": [
            {
              "sensor-path":
"Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/latest/generic-counters/bytes-received",
              "sensor-path-tags": "interface-name=GigabitEthernet0/3/0/19",
              "data": [
                {
                  "value": {
                    "value-type": 8,
                    "val-counter64": "62808023132655"
                  },
                  "timestamp": "1607431545418"
                },
                ...
                {
                  "value": {
                    "value-type": 8,
                    "val-counter64": "62869633436614"
                  },
                  "timestamp": "1607432235421"
                },
                {
                  "value": {
                    "value-type": 8,
                    "val-counter64": "62872314602090"
                  },
                  "timestamp": "1607432265421"
                }
              ]
            }
          ]
        }
      ]
    }
  ],
  "collection_end_time": "1607525488556"
```

The router streams data to the receiver upon state change using the subscription-based telemetry session. A data lake is created at the receiver.

See the ADT events.

Step 2 See the details of events.

Example:

In this example, the details of event are displayed.

```
Router#show adt events id detail
Event Id      :
Timestamp     : Tue 2020-12-08 18:21:45:419
Description   : Traffic
Number of Sensor paths : 1
Sensor Path   :
Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/latest/generic-counters/bytes-received

Sensor Path Tags : interface-name=GigabitEthernet0/3/0/19
Message          :
  Number of entries in list: 25
  Value : [ 62808023132655, 62810701566605, 62813380497605, 62816056633805,
            62818733833405, 62821412539395, 62824092600551, 62826773125641,
            62829461449196, 62832132687840, 62834805072011, 62837462182289,
            62840165873427, 62842846468785, 62845523517431, 62848199893619,
            62852053505122, 62853557475735, 62856237814551, 62858917998694,
            62861594969436, 62864275943572, 62866948687442, 62869633436614,
            62872314602090, ]
  First Timestamp : Tue 2020-12-08 18:15:45:418 [1607431545418]
  Last Timestamp  : Tue 2020-12-08 18:27:45:421 [1607432265421]
```

ADT event lists 25 sample values of data that are reported by sensor paths before and after the event. So, 25 data points help us describe a particular event and its associated sensor path.

Operate on Telemetry Data for In-Depth Analysis of the Network

You can start consuming and analysing telemetry data from the data lake using an open-sourced collection stack. This use case uses the following tools from the collection stack:

- Pipeline is a lightweight tool that is used to collect data. You can download [Network Telemetry Pipeline](#) from Github. You define how you want the collector to interact with routers and where you want to send the processed data using `pipeline.conf` file.
- Telegraph (plugin-driven server agent) and InfluxDB (time series database (TSDB)) stores telemetry data, which is retrieved by visualization tools. You can download [InfluxDB](#) from Github. You define what data you want to include into your TSDB using the `metrics.json` file.
- [Grafana](#) is a visualization tool that displays graphs and counters for data that is streamed from the router.

In summary, Pipeline accepts TCP and gRPC telemetry streams, converts data and pushes data to the InfluxDB database. Grafana uses the data from the InfluxDB database to build dashboards and graphs. Pipeline and InfluxDB may run on the same server or on different servers.

Consider that the router is streaming data on an event change, and Telegraf requests information from the Pipeline at 1-second intervals.

Step 1 Start Pipeline, and enter your router credentials.

Note The IP address and port that you specify in the destination-group must match the IP address and port on which Pipeline is listening.

Example:

```
$ bin/pipeline -config pipeline.conf

Startup pipeline
Load config from [pipeline.conf], logging in [pipeline.log]

CRYPT Client [grpc_in_myndtrouter], [http://172.0.0.0:5432]
Enter username: <username>
Enter password: <password>
Wait for ^C to shutdown
```


Step 2 In the Telegraph configuration file, add the following values to read the metrics about CPU usage.

Example:

```
[[inputs.cpu]]
  ## Whether to report per-cpu stats or not
  percpu = true
  ## Whether to report total system cpu stats or not
  totalcpu = true
  ## If true, collect raw CPU time metrics.
  collect_cpu_time = false
  ## If true, compute and report the sum of all non-idle CPU states.
  report_active = false
```


Step 3 Use Grafana to create a dashboard and visualize data.

Figure 13: Visual Analysis of the Interfaces Impacted Due to Traffic Change Using Telemetry Data



Further, narrow down to view the impact on individual interfaces. In this example, the bytes received for interfaces HundredGigE0/0/30 and HundredGigE0/0/31 are visualized and analysed.

Figure 14: Visual Analysis of System Monitoring on Two Interfaces Using Telemetry Data



In conclusion, until the point where an event occurred, there is relatively no change. When an event is detected, there is significant drop in traffic on few interfaces and a peak on few other interfaces. ADT predicted these changes in the interfaces accurately using the associated sensor paths that it learned using AI and unsupervised ML from the router's configuration.



CHAPTER 6

Hardware Timestamp

Table 3: Feature History Table

Whenever periodic stats are streamed, the collector reads the stats from its internal cache, instead of fetching the stats from the hardware. When the data is read from the sensor paths of Stats manager cache, the rate calculation shows spikes. This behavior is due to the timestamp from the collector that is off by several seconds. Therefore, timestamp of some other collector takes precedence because timestamps of collectors are not in synchronization with the current timestamp. This is observed when there are multiple collectors providing stats updates for the same interface.

The YANG data model for Stats manager `Cisco-IOS-XR-infra-statsd-oper.yang` is enhanced to enable the collector to read periodic stats data from the router using hardware timestamp.

The hardware timestamp is taken into account when a primary collector (for generic or proto stats) provides stats updates from the hardware to the Stats manager. With hardware timestamping in rate computation while streaming periodic stats, the spikes due to the timestamp issue is resolved.

The hardware timestamp is updated only when the collector attempts to read the counters from hardware. Else, the value remains 0. The latest stats can be streamed at a minimum cadence of 10 seconds and periodic stats at a cadence of 30 seconds. The support is available only for physical interfaces and subinterfaces, and bundle interface and subinterfaces.

When there is no traffic flow on protocols for an interface, the hardware timestamp for the protocols is published as 0. This is due to non-synchronized timestamps sent by the collector for protocols in traffic as compared to non-traffic scenarios.

A non-zero value is published for protocols that have stats published by a primary collector for both traffic and non-traffic scenarios.



Note The hardware timestamp is supported only for primary collectors. When the hardware has no update, the timestamp will be same. However generic counters are computed for primary and non-primary collectors. The non-primary collectors show the latest stats, but not the timestamp.

When the counters are cleared for an interface using **clear counters interface** command, all counter-related data including the timestamps for the interface is cleared. After all counter values are cleared and set to 0, the last data time is updated only when there is a request for it from a collector. For example, last data time gets updated from a collector:

```
Router# :Aug 7 09:01:08.471 UTC: statsd_manager_1[168]: Updated last data time for ifhandle
0x02000408,
```

```
stats type 2 from collector with node 0x100, JID 250, last data time 1596790868.
INPUT: last 4294967295 updated 1596469986. OUTPUT: last 4294967295 updated 1596469986
```

All other counter values and hardware timestamp are updated when the counters are fetched from the hardware. In this case, all counters including the hardware timestamp is 0:

```
{"node_id_str":"MGBL_MTB_5504","subscription_id_str":"app_TEST_200000001",
"encoding_path":"Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/cache/generic-counters",
"collection_id":"7848",
"collection_start_time":"1596790879567",
"msg_timestamp":"1596790879571","data_json":
[{"timestamp":"1596790879570","keys":[{"interface-name":"FortyGigE0/1/0/11"}],
"content":{"packets-received":"0","bytes-received":"0","packets-sent":"0",
"bytes-sent":"0","multicast-packets-received":"0","broadcast-packets-received":"0",
"multicast-packets-sent":"0","broadcast-packets-sent":"0","output-drops":0,"output-queue-drops":0,
"input-drops":0,"input-queue-drops":0,"runt-packets-received":0,"giant-packets-received":0,
"throttled-packets-received":0,"parity-packets-received":0,"unknown-protocol-packets-received":0,
"input-errors":0,"crc-errors":0,"input-overruns":0,"framing-errors-received":0,"input-ignored-packets":0,
"input-aborts":0,"output-errors":0,"output-underruns":0,"output-buffer-failures":0,"output-buffers-swapped-out":0,
"applique":0,"resets":0,"carrier-transitions":0,"availability-flag":0,
"last-data-time":"1596790868","hardware-timestamp":"0",
"seconds-since-last-clear-counters":15,"last-discontinuity-time":1596469946,"seconds-since-packet-received":0,
"seconds-since-packet-sent":0}}],"collection_end_time":"1596790879571"}
```

- [Stream Telemetry Data about PBR Decapsulation Statistics, on page 44](#)

Stream Telemetry Data about PBR Decapsulation Statistics

You can stream telemetry data about PBR decapsulation statistics for GRE and GUE encapsulation protocols that deliver packets using IPv4 or IPv6. The encapsulated data has source and destination address that must match with the source and destination address in the classmap. Both encapsulation and decapsulation interfaces collect statistics periodically. The statistics can be displayed on demand using **show policy-map type pbr [vrf vrf-name] address-family ipv4/ipv6 statistics** command. For more information on PBR-based decapsulation, see .

With this release, the decapsulation statistics can be displayed using

Cisco-IOS-XR-infra-policymgr-oper.yang data model and telemetry data. You can stream telemetry data from the sensor path:

```
Cisco-IOS-XR-infra-policymgr-oper:policy-manager/global/policy-map/policy-map-types/policy-map-type/vrf-table/vrf/afi-table/afi/stats
```

The following steps show the PBR configuration and the decapsulation statistics that is streamed as telemetry data to the collector.

Step 1 Check the running configuration to view the configured PBR per VRF.

Example:

```
Router#show running-config
Building configuration...
!! IOS XR Configuration 0.0.0
!!
vrf vrf1
 address-family ipv4 unicast
 !
 address-family ipv6 multicast
 !
netconf-yang agent
```

```

ssh
!
!
class-map type traffic match-all cmap1
  match protocol gre
  match source-address ipv4 161.0.1.1 255.255.255.255
  match destination-address ipv4 161.2.1.1 255.255.255.255
end-class-map
!
policy-map type pbr gre-policy
  class type traffic cmap1
    decapsulate gre
  !
  class type traffic class-default
  !
end-policy-map
!
interface GigabitEthernet0/0/0/1
  vrf vrf1
  ipv4 address 2.2.2.2 255.255.255.0
  shutdown
!
vrf-policy
  vrf vrf1 address-family ipv4 policy type pbr input gre-policy
!
end

```

Step 2 View the output of the VRF statistics.

Example:

Router#**show policy-map type pbr vrf vrf1 addr-family ipv4 statistics**

```

VRF Name:      vrf1
Policy-Name:   gre-policy
Policy Type:   pbr
Addr Family:   IPv4

Class:      cmap1
  Classification statistics      (packets/bytes)
    Matched      :      13387587/1713611136
  Transmitted statistics      (packets/bytes)
    Total Transmitted      :      13387587/1713611136

Class:      class-default
  Classification statistics      (packets/bytes)
    Matched      :      0/0
  Transmitted statistics      (packets/bytes)
    Total Transmitted      :      0/0

```

After you have verified that the statistics are displayed correctly, stream telemetry data and check the streamed data at the collector. For more information about collectors, see *Operate on Telemetry Data for In-depth Analysis of the Network* section in the [Monitor CPU Utilization Using Telemetry Data to Plan Network Infrastructure, on page 14](#) chapter.

```

ios.0/0/CPU0/ $ mdt_exec -s Cisco-IOS-XR-infra-policymgr-oper:policy-manager
/global/policy-map/policy-map-types/policy-map-type/vrf-table/vrf/afi-table/afi/stats -c 100
{"node_id_str":"ios","subscription_id_str":"app_TEST_200000001","encoding_path":
"Cisco-IOS-XR-infra-policymgr-oper:policy-manager/global/policy-map/policy-map-types/policy-map-type
/vrf-table/vrf/afi-table/afi/stats","collection_id":"1","collection_start_time":"1601361558157",
"msg_timestamp":"1601361559179","data_json":[{"timestamp":"1601361559178","keys":[{"type":"ipv6"},
{"vrf_name":"vrf_gue_ipv4"}, {"type":"ipv4"}]}, {"content":{"pmap_name":"gre-policy","vrf_name":
"vrf1","appln_type":2,"addr_family":1,"rc":0,"plmgr_vrf_stats":{"pmap_name":"gre-policy",
"cmmap_stats_arr":[{"cmmap_name":"cmap1","matched_bytes":"1713611136","matched_packets":"13387587",
"transmit_bytes":"1713611136","transmit_packets":"13387587"}]}]}]}],

```

```
"collection_end_time":"1601361559183"}
```

```
----- snipped for brevity -----
```
