



Enhancements to Data Models

This section provides an overview of the enhancements made to data models.

- [Install Label in oc-platform Data Model, on page 1](#)
- [OAM for MPLS and SR-MPLS in mpls-ping and mpls-traceroute Data Models, on page 3](#)
- [OpenConfig YANG Model:SR-TE Policies, on page 8](#)
- [Aggregate Prefix SID Counters for OpenConfig SR YANG Module, on page 9](#)
- [OpenConfig YANG Model:AFT, on page 10](#)

Install Label in oc-platform Data Model

Table 1: Feature History Table

Feature Name	Release Information	Description
Enhancements to openconfig-platform YANG Data Model	Release 7.3.2	<p>The openconfig-platform YANG data model provides a structure for querying hardware and software router components via the NETCONF protocol. This release delivers an enhanced openconfig-platform YANG data model to provide information about:</p> <ul style="list-style-type: none">• software version• golden ISO (GISO) label• committed IOS XR packages <p>You can access this data model from the Github repository.</p>

The openconfig-platform (oc-platform.yang) data model is enhanced to provide the following data:

- IOS XR software version (optionally with GISO label)
- Type, description, operational status of the component. For example, a CPU component reports its utilization, temperature or other physical properties.

Install Label in oc-platform Data Model

- List of the committed IOS XR packages

To retrieve oc-platform information from a router via NETCONF, ensure you configured the router with the SH server and management interface:

```
Router#show run
Building configuration...
!! IOS XR Configuration version = 7.3.2
!! Last configuration change at Tue Sep 7 16:18:14 2016 by USER1
!
.....
.....
netconf-yang agent ssh
ssh server netconf vrf default
interface MgmtEth 0/RP0/CPU0/0
    no shut
    ipv4 address dhcp
```

The following example shows the enhanced `OPERATING_SYSTEM` node component (line card or route processor) of the oc-platform data model:

```
<component>
<name>IOSXR-NODE 0/RP0/CPU0</name>
<config>
<name>0/RP0/CPU0</name>
</config>
<state>
<name>0/RP0/CPU0</name>
<type xmlns:idx="http://openconfig.net/yang/platform-types">idx:OPERATING_SYSTEM</type>
<location>0/RP0/CPU0</location>
<description>IOS XR Operating System</description>
<software-version>7.3.2</software-version> -----> Label Info
<removable>true</removable>
<oper-status xmlns:idx="http://openconfig.net/yang/platform-types">idx:ACTIVE</oper-status>
</state>
<subcomponents>
<subcomponent>
<name><platform>-af-ea-7.3.2v1.0.0.1</name>
<config>
<name><platform>-af-ea-7.3.2v1.0.0.1</name>
</config>
<state>
<name><platform>-af-ea-7.3.2v1.0.0.1</name>
</state>
</subcomponent>
...
...
```

The following example shows the enhanced `OPERATING_SYSTEM_UPDATE` package component (RPMs) of the oc-platform data model:

```
<component>
<name>IOSXR-PKG/1 <platform>-isis-2.1.0.0-r732</name>
<config>
<name><platform>-isis-2.1.0.0-r732</name>
</config>
<state>
<name><platform>-isis-2.1.0.0-r732</name>
<type xmlns:idx="http://openconfig.net/yang/platform-types">idx:OPERATING_SYSTEM_UPDATE</type>
<description>IOS XR Operating System Update</description>
<software-version>7.3.2</software-version>-----> Label Info
<removable>true</removable>
<oper-status xmlns:idx="http://openconfig.net/yang/platform-types">idx:ACTIVE</oper-status>
</state>
</component>
```

Associated Commands

- **show install committed**—Shows the committed IOS XR packages.
- **show install committed summary**—Shows a summary of the committed packages along with the committed IOS XR version that is displayed as a label.

OAM for MPLS and SR-MPLS in mpls-ping and mpls-traceroute Data Models

Table 2: Feature History Table

Feature Name	Release Information	Description
YANG Data Models for MPLS OAM RPCs	Release 7.3.2	<p>This feature introduces the <code>Cisco-IOS-XR-mpls-ping-act</code> and <code>Cisco-IOS-XR-mpls-traceroute-act</code> YANG data models to accommodate operations, administration and maintenance (OAM) RPCs for MPLS and SR-MPLS.</p> <p>You can access these Cisco IOS XR native data models from the Github repository.</p>

The Cisco-IOS-XR-mpls-ping-act and Cisco-IOS-XR-mpls-traceroute-act YANG data models are introduced to provide the following options:

- Ping for MPLS:
 - MPLS IPv4 address
 - MPLS TE
 - FEC-129 Pseudowire
 - FEC-128 Pseudowire
 - Multisegment Pseudowire
- Ping for SR-MPLS:
 - SR policy name or BSID with LSP end-point
 - SR MPLS IPv4 address
 - SR Nil-FEC labels
 - SR Flexible Algorithm
- Traceroute for MPLS:

- MPLS IPv4 address
- MPLS TE
- Traceroute for SR-MPLS:
 - SR policy name or BSID with LSP end-point
 - SR MPLS IPv4 address
 - SR Nil-FEC labels
 - SR Flexible Algorithm

The following example shows the ping operation for an SR policy and LSP end-point:

```
<mpls-ping xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-mpls-ping-act">
  <sr-mpls>
    <policy>
      <name>srtc_c_10_ep_10.10.10.1</name>
      <lsp-endpoint>10.10.10.4</lsp-endpoint>
    </policy>
  </sr-mpls>
  <request-options-parameters>
    <brief>true</brief>
  </request-options-parameters>
</mpls-ping>
```

Response:

```
<?xml version="1.0"?>
<mpls-ping-response xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-mpls-ping-act">
  <request-options-parameters>
    <exp>0</exp>
    <fec>false</fec>
    <interval>0</interval>
    <ddmap>false</ddmap>
    <force-explicit-null>false</force-explicit-null>
    <packet-output>
      <interface-name>None</interface-name>
      <next-hop>0.0.0.0</next-hop>
    </packet-output>
    <pad>abcd</pad>
    <repeat>5</repeat>
    <reply>
      <dscp>255</dscp>
      <reply-mode>default</reply-mode>
      <pad-tlv>false</pad-tlv>
    </reply>
    <size>100</size>
    <source>0.0.0.0</source>
    <destination>127.0.0.1</destination>
    <sweep>
      <minimum>100</minimum>
      <maximum>100</maximum>
      <increment>1</increment>
    </sweep>
    <brief>true</brief>
    <timeout>2</timeout>
    <ttl>255</ttl>
  </request-options-parameters>
  <replies>
```

```

<reply>
<reply-index>1</reply-index>
<return-code>3</return-code>
<return-char>!</return-char>
<reply-addr>14.14.14.3</reply-addr>
<size>100</size>
</reply>
<reply>
<reply-index>2</reply-index>
<return-code>3</return-code>
<return-char>!</return-char>
<reply-addr>14.14.14.3</reply-addr>
<size>100</size>
</reply>
<reply>
<reply-index>3</reply-index>
<return-code>3</return-code>
<return-char>!</return-char>
<reply-addr>14.14.14.3</reply-addr>
<size>100</size>
</reply>
<reply>
<reply-index>4</reply-index>
<return-code>3</return-code>
<return-char>!</return-char>
<reply-addr>14.14.14.3</reply-addr>
<size>100</size>
</reply>
<reply>
<reply-index>5</reply-index>
<return-code>3</return-code>
<return-char>!</return-char>
<reply-addr>14.14.14.3</reply-addr>
<size>100</size>
</reply>
</replies>
</mpls-ping-response>

```

The following example shows the ping operation for an SR policy BSID and LSP end-point:

```

<mpls-ping xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-mpls-ping-act">
<sr-mpls>
<policy>
<bsid>1000</bsid>
<lsp-endpoint>10.10.10.4</lsp-endpoint>
</policy>
</sr-mpls>
<request-options-parameters>
<brief>true</brief>
</request-options-parameters>
</mpls-ping>

```

Response:

```

<?xml version="1.0"?>
<mpls-ping-response xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-mpls-ping-act">
<request-options-parameters>
<exp>0</exp>
<fec>false</fec>
<interval>0</interval>
<ddmap>false</ddmap>
<force-explicit-null>false</force-explicit-null>
<packet-output>

```

```

<interface-name>None</interface-name>
<next-hop>0.0.0.0</next-hop>
</packet-output>
<pad>abcd</pad>
<repeat>5</repeat>
<reply>
  <dscp>255</dscp>
  <reply-mode>default</reply-mode>
  <pad-tlv>false</pad-tlv>
</reply>
<size>100</size>
<source>0.0.0.0</source>
<destination>127.0.0.1</destination>
<sweep>
  <minimum>100</minimum>
  <maximum>100</maximum>
  <increment>1</increment>
</sweep>
<brief>true</brief>
<timeout>2</timeout>
<ttl>255</ttl>
</request-options-parameters>
<replies>
  <reply>
    <reply-index>1</reply-index>
    <return-code>3</return-code>
    <return-char>!</return-char>
    <reply-addr>14.14.14.3</reply-addr>
    <size>100</size>
  </reply>
  <reply>
    <reply-index>2</reply-index>
    <return-code>3</return-code>
    <return-char>!</return-char>
    <reply-addr>14.14.14.3</reply-addr>
    <size>100</size>
  </reply>
  <reply>
    <reply-index>3</reply-index>
    <return-code>3</return-code>
    <return-char>!</return-char>
    <reply-addr>14.14.14.3</reply-addr>
    <size>100</size>
  </reply>
  <reply>
    <reply-index>4</reply-index>
    <return-code>3</return-code>
    <return-char>!</return-char>
    <reply-addr>14.14.14.3</reply-addr>
    <size>100</size>
  </reply>
  <reply>
    <reply-index>5</reply-index>
    <return-code>3</return-code>
    <return-char>!</return-char>
    <reply-addr>14.14.14.3</reply-addr>
    <size>100</size>
  </reply>
</replies>
</mpls-ping-response>

```

The following example shows the traceroute operation for an SR policy and LSP end-point:

```
<mpls-traceroute xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-mpls-traceroute-act">
<sr-mpls>
<policy>
  <name>srtc_10_ep_10.10.10.1</name>
  <lsp-endpoint>10.10.10.4</lsp-endpoint>
</policy>
</sr-mpls>
<request-options-parameters>
  <brief>true</brief>
</request-options-parameters>
</mpls-traceroute>
```

Response:

```
<?xml version="1.0"?>
<mpls-traceroute-response xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-mpls-traceroute-act">

<request-options-parameters>
  <exp>0</exp>
  <fec>false</fec>
  <ddmap>false</ddmap>
  <force-explicit-null>false</force-explicit-null>
  <packet-output>
    <interface-name>None</interface-name>
    <next-hop>0.0.0.0</next-hop>
  </packet-output>
  <reply>
    <dscp>255</dscp>
    <reply-mode>default</reply-mode>
  </reply>
  <source>0.0.0.0</source>
  <destination>127.0.0.1</destination>
  <brief>true</brief>
  <timeout>2</timeout>
  <ttl>30</ttl>
</request-options-parameters>
<paths>
  <path>
    <path-index>0</path-index>
    <hops>
      <hop>
        <hop-index>0</hop-index>
        <hop-origin-ip>11.11.11.1</hop-origin-ip>
        <hop-destination-ip>11.11.11.2</hop-destination-ip>
        <mtu>1500</mtu>
        <dsmap-label-stack>
          <dsmap-label>
            <label>16003</label>
          </dsmap-label>
        </dsmap-label-stack>
        <return-code>0</return-code>
        <return-char> </return-char>
      </hop>
      <hop>
        <hop-index>1</hop-index>
        <hop-origin-ip>11.11.11.2</hop-origin-ip>
        <hop-destination-ip>14.14.14.3</hop-destination-ip>
        <mtu>1500</mtu>
        <dsmap-label-stack>
          <dsmap-label>
            <label>3</label>
          </dsmap-label>
        </dsmap-label-stack>
        <return-code>8</return-code>
      </hop>
    </hops>
  </path>
</paths>
</mpls-traceroute-response>
```

```

<return-char>L</return-char>
</hop>
<hop>
  <hop-index>2</hop-index>
  <hop-origin-ip>14.14.14.3</hop-origin-ip>
  <hop-destination-ip></hop-destination-ip>
  <mtu>0</mtu>
  <dsmpl-label-stack/>
  <return-code>3</return-code>
  <return-char>!</return-char>
</hop>
</hops>
</path>
</paths>
</mpls-traceroute-response>

```

OpenConfig YANG Model:SR-TE Policies

Table 3: Feature History Table

Feature Name	Release Information	Description
OpenConfig YANG Model:SR-TE Policies	Release 7.3.4	<p>This release supports the OpenConfig (OC) Segment Routing-Traffic Engineering (SR-TE) YANG data model that provides data definitions for SR-TE policy configuration and associated signaling and traffic engineering protocols. Using the model, you can stream a collection of SR-TE operational statistics, such as color, endpoint, and state.</p> <p>You can access the OC data model from the Github repository.</p>

The OC SR-TE policies YANG Data Model supports Version 0.22. Subscribe to the following sensor path to send a pull request to the YANG leaf, list, or container:

`openconfig-network-instance:network-instances/network-instance/segment-routing/te-policies`

The response from the router is a collection of SR-TE operational statistics, such as color, endpoint, and state.

Limitations

- Segment-list ID
 - All locally-configured segment-lists have a unique segment-list ID except for the BGP TE controller. Instead, the BGP TE controller uses the index of the segment-list as the segment-list ID. This ID depends on the local position of the segment-list and can change over time. Therefore for BGP TE controller, you must stream the entire table of the segment-list to ensure that the segment-list ID is always up-to-date.
- Next-hop index

- The Next-hop container is imported from the `openconfig-aft-common.yang` module where the next-hop index is defined as `Uint64`. However, the AFT OC in the FIB uses a positional value of the index and does not identify the next-hop entry separately. Similarly, the next-hop container for OC-SRTE also implements as a positional value of the entry in the list. Ensure that you stream the entire table of the next-hop to get an updated index along with the next-hop entry.

Aggregate Prefix SID Counters for OpenConfig SR YANG Module

Table 4: Feature History Table

Feature Name	Release Information	Description
Aggregate Prefix SID Counters for OpenConfig SR YANG Module	Release 7.3.4	<p>The following components are now available in the OpenConfig (OC) Segment-Routing (SR) YANG model:</p> <ul style="list-style-type: none"> The aggregate-sid-counters container in the sr-mpls-top group to aggregate the prefix segment identifier (SID) counters across the router interfaces. The aggregate-sid-counter and the mpls-label key to aggregate counters across all the router interfaces corresponding to traffic forwarded with a particular prefix-SID. <p>You can access the OC data model from the Github repository.</p>

The OpenConfig SR YANG model supports Version 0.3. Subscribe to the following sensor path:

`openconfig-mpls/mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter/mpls-label/state`

When a receiver subscribes to the sensor path, the router periodically streams the statistics to telemetry for each SR-label. The default collection interval is 30 seconds.

OpenConfig YANG Model:AFT

Table 5: Feature History Table

Feature Name	Release Information	Description
OpenConfig YANG Model:AFT	Release 7.3.4	<p>This release supports the OpenConfig Abstract Forwarding Table (AFT) containers, such as IPv4, IPv6, Network Instance, and MPLS. With this support, the AFT sends only essential interface forwarding entries, such as the next-hop, next-hop group, and RSVP-TE for an IP prefix, to the Network Management System (NMS). Since the NMS receives only essential entries, the forwarding process is simplified.</p> <p>You can access the OC data model from the Github repository.</p>

Supported Agents

The following agents are supported in the SAMPLE and ON-CHANGE modes:

- gNMI
- IOS-XR proprietary telemetry dial-in and dial-out

Limitations

- The Netconf agent is not supported on configuration and operation data.
- The ON-CHANGE mode is supported only at the path level as shown below:
 - /network-instances/network-instance/afts/ipv4-unicast/ipv4-entry
 - /network-instances/network-instance/afts/ipv6-unicast/ipv6-entry
 - /network-instances/network-instance/afts/mpls/label-entry
 - /network-instances/network-instance/afts/next-hop-groups/next-hop-group/state
 - /network-instances/network-instance/afts/next-hop-groups/next-hop-group/next-hops/next-hop
 - /network-instances/network-instance/afts/next-hops/next-hop
- The current implementation of the OC-AFT model, version 0.6.0 does not set the atomic flag for atomic updates for gNMI.

Response

A SubscribeRequest message is sent by a gNMI client to request updates from the router for a specified set of paths. The following SubscriptionResponse messages are sent by the router:

AFT IPv4 unicast

```
SubscribeResponse.update: <
timestamp: 1647978999525525791
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts ipv4-unicast ipv4-entry[prefix=10.0.0.1/32] > < json_ietf_val:"{
"state": {
"prefix": "10.0.0.1/32",
"next-hop-group": "1152921642045939938"
}
}" > >

SubscribeResponse.update: <
timestamp: 1647978999341662576
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts ipv4-unicast ipv4-entry[prefix=10.1.1.1/32] > < json_ietf_val:"{
"state": {
"prefix": "10.1.1.1/32",
"next-hop-group": "1152921779484853982"
}
}" > >
```

AFT IPv6 unicast

```
SubscribeResponse.update: <
timestamp: 1647984444644492536
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts ipv6-unicast ipv6-entry[prefix=50:50:58::331/128] > < json_ietf_val:"{
"state": {
"prefix": "50:50:58::331/128",
"next-hop-group": "1153062379534237025"
}
}" > >
```

List of MPLS entries within the AFT

```
SubscribeResponse.update: <
timestamp: 1648009876493069763
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts mpls label-entry[label=12000] > < json_ietf_val:"{
"state": {
"label": 12000,
"next-hop-group": "1152921642046007012"
}
}" > >

SubscribeResponse.update: <
timestamp: 1648011005293000000
prefix: <
```

```

origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts mpls label-entry[label=12000] > < json_ietf_val:"{
"state": {
"label": 12000,
"packets-forwarded": "0",
"octets-forwarded": "0"
}
}" > >
```

AFT next-hop-group

```

SubscribeResponse.update: <
timestamp: 1648011006899606800
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts next-hop-groups next-hop-group[id=1152921642045939938] >
< json_ietf_val:"{
"next-hops": {
"next-hop": {
"index": "1152921642045903362",
"state": {
"index": "1152921642045903362",
"weight": "0"
}
}
}
}" > >

>
SubscribeResponse.update: <
timestamp: 1648011006899606800
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts next-hop-groups next-hop-group[id=1152921642045939938] >
< json_ietf_val:"{
"next-hops": {
"next-hop": {
"index": "1152921642045903355",
"state": {
"index": "1152921642045903355",
"weight": "0"
}
}
}
}" > >

SubscribeResponse.update: <
timestamp: 1648011006899606800
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts next-hop-groups next-hop-group[id=1152921642045939938] >
< json_ietf_val:"{
"next-hops": {
"next-hop": {
"index": "1152921642045903348",
"state": {
```

```
"index": "1152921642045903348",
"weight": "0"
}
}
}
}" > >

AFT next-hops next-hop

SubscribeResponse.update: <
timestamp: 1648011006713962739
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts next-hops next-hop[index=1152921642045903362] > < json_ietf_val:"{
"state": {
"index": "1152921642045903362",
"ip-address": "13.1.1.1"
},
"interface-ref": {
"state": {
"interface": "tunnel-ip2",
"subinterface": 0
}
}
}" > >

SubscribeResponse.update: <
timestamp: 1648011006713954259
prefix: <
origin: openconfig-network-instance
>
update: < path: < element: network-instances network-instance[name=default]
afts next-hops next-hop[index=1152921642045903355] > < json_ietf_val:"{
"state": {
"index": "1152921642045903355",
"ip-address": "13.1.1.2"
},
"interface-ref": {
"state": {
"interface": "tunnel-ip3",
"subinterface": 0
}
}
}" > >
```

