

TEAS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 6 August 2026

D. Dhody, Ed.
Huawei
D. King
Lancaster University
R. Vilalta
CTTC
I. Busi
Huawei Technologies
D. Ceccarelli
Cisco
2 February 2026

YANG models for Virtual Network (VN)/TE Performance Monitoring Telemetry
and Scaling Intent Autonomics
draft-ietf-teas-actn-pm-telemetry-autonomics-18

Abstract

This document provides YANG data models that describe the performance monitoring parameters and scaling intent mechanisms for Traffic Engineering (TE) tunnels and Virtual Networks (VNs). Their performance monitoring parameters are exposed as the key telemetry data for tunnels and VNs.

The models presented in this document allow customers to subscribe to and monitor the key performance data of the TE-tunnel or the VN. The models also provide customers with the ability to program autonomic scaling intent mechanisms on the level of TE-tunnel as well as VN.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 6 August 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	4
1.2. Tree Diagram	4
1.3. Prefixes in Data Node Names	4
2. Use-Cases	5
3. Design of the Data Models	7
3.1. TE Telemetry Model	7
3.2. VN Telemetry Model	8
3.3. VPN Service Performance Monitoring	9
4. Autonomic Scaling Intent Mechanism	10
5. Performance Monitoring Parameters	12
6. Notification	13
6.1. YANG Push Subscription Examples	13
7. Scaling Examples	15
8. YANG Data Tree	19
9. YANG Data Model	22
9.1. ietf-te-telemetry model	22
9.2. ietf-vn-telemetry model	30
10. Security Considerations	34
11. IANA Considerations	36
12. Acknowledgments	37
13. References	37
13.1. Normative References	37
13.2. Informative References	39
Appendix A. Out of Scope	40
Appendix B. Contributors	41
Authors' Addresses	41

1. Introduction

The YANG [RFC7950] model in [RFC9731] is used to operate customer-driven Virtual Networks (VNs) during the computation of VN, their instantiation, and their life-cycle service management and operations. The YANG model in [I-D.ietf-teas-yang-te] is used to operate TE-tunnels during the tunnel instantiation, and their life-cycle management and operations.

The models presented in this document allow the applications hosted by the customers to subscribe to and monitor the key performance data of their interest on the level of VN [RFC9731] or TE-tunnel [I-D.ietf-teas-yang-te]. The key characteristic of the models presented in this document is their top-down programmability, which enables customer-hosted applications to subscribe to and monitor the key performance data of their interest, as well as utilize the autonomic scaling intent mechanism at both the VN and TE-tunnel levels.

According to the classification of [RFC8309], the YANG data models presented in this document can be classified as customer service models. These can be mapped to the CMI (Customer Network Controller (CNC)- Multi-Domain Service Coordinator (MSDC) interface) of Abstraction and Control of TE Networks (ACTN) [RFC8453].

[RFC8233] describes key network performance data to be considered for end-to-end path computation in TE networks. The services provided can be optimized to meet the requirements (such as traffic patterns, quality, and reliability) of the applications hosted by the customers.

This document provides YANG data models with performance monitoring parameters that can be subscribed to for monitoring and telemetry for any VN/TE-Tunnel via the mechanism specified in [RFC8641] and [RFC8640]. It also provides the ability to program their customized automatic scaling in/out intent. A client network controller can utilize these models and initiate the capabilities via a NETCONF [RFC8341] or a RESTCONF [RFC8040] interface.

The term 'Performance monitoring' in this document refers to the subscription and publication of streaming telemetry data. Subscription is initiated by the client (e.g., CNC) while publication is provided by the network (e.g., MDSC/Provisioning Network Controller (PNC)) based on the client's subscription. As per [RFC7799], this would be classified as a passive method. Note that the actual measurements might be done via any technique, though. As the scope of performance monitoring in this document is to augment the performance monitoring parameters (telemetry data) on the level

of a client's VN or TE-tunnel, the entity interfacing to the client (e.g., MDSC) has to provide VN or TE-tunnel level information. This requires the controller to be able to derive VN or TE-tunnel level performance data based on lower-level data collected via PM counters in the Network Elements (NE). How the controller entity derives such customized level data (i.e., VN or TE-tunnel level) is out of the scope of this document.

The data model includes configuration and state data according to the Network Management Datastore Architecture (NMDA) [RFC8342].

1.1. Terminology

Refer to [RFC8453], [RFC7926], and [RFC8309] for the key terms used in this document.

Scaling: This refers to the network's ability to reshape its own resources. "Scale-out" refers to improving network performance by increasing the allocated resources, while "scale-in" refers to decreasing the allocated resources, typically because the existing resources are unnecessary.

Scaling Intent: Scaling intent is used to declare scaling conditions. Specifically, scaling intent refers to how the client programs or configures conditions that will be applied to their key performance data to trigger either scaling out or scaling in. Various conditions can be set for scaling intent on either the VN or TE-tunnel level.

Network Autonomics: This refers to the network automation capability that allows a client to initiate scaling intent mechanisms and provides the client with the status of the adjusted network resources based on the client's scaling intent in an automated fashion.

1.2. Tree Diagram

A simplified graphical representation of the data model is used in Section 4 and Section 8 of this document. The meaning of the symbols in these diagrams is defined in [RFC8340].

1.3. Prefixes in Data Node Names

In this document, the names of data nodes and other data model objects are prefixed using the standard prefix associated with the corresponding YANG imported modules, as shown in Table 1.

Prefix	YANG module	Reference
te	ietf-te	[I-D.ietf-teas-yang-te]
te-types	ietf-te-types	[RFC8776]
rt-types	ietf-routing-types	[RFC8294]
te-tel	ietf-te-telemetry	[RFCXXXX]
vn	ietf-vn	[RFC9731]
vn-tel	ietf-vn-telemetry	[RFCXXXX]

Table 1: Prefixes and corresponding YANG modules

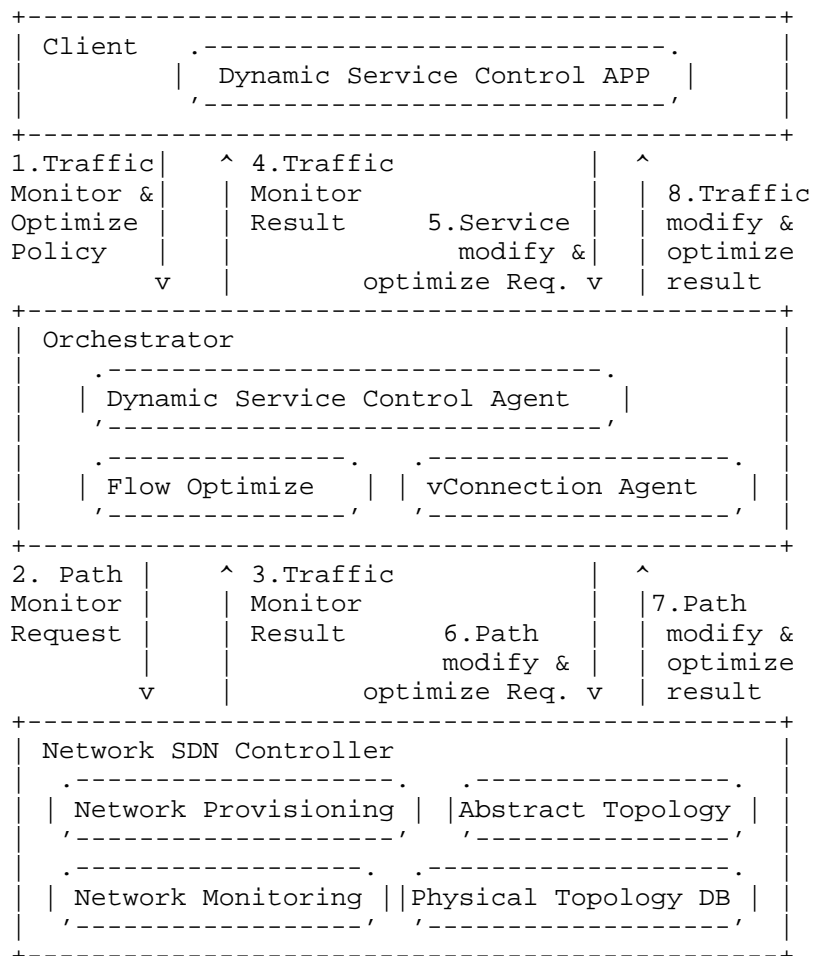
Note: The RFC Editor is requested to replace XXXX with the number assigned to the RFC once this draft becomes an RFC, and to remove this note.

Further, the following additional documents are referenced in the model defined in this document -

- * [RFC7471] - OSPF Traffic Engineering (TE) Metric Extensions.
- * [RFC8570] - IS-IS Traffic Engineering (TE) Metric Extensions.
- * [RFC7823] - Performance-Based Path Selection for Explicitly Routed Label Switched Paths (LSPs) Using TE Metric Extensions.

2. Use-Cases

There is a need for real-time (or semi-real-time) traffic monitoring of the network to optimize the network and the traffic distribution. Figure 1 shows an example of a high-level workflow for dynamic service control based on traffic monitoring that could use the mechanism described in this document.



APP: Application

DB: Database

Req: Request

Figure 1: Workflow for dynamic service control based on traffic monitoring

Some of the key points are as follows:

- * Network traffic monitoring is important to facilitate the automatic discovery of the imbalance of network traffic, and initiate network optimization, thus helping the network operator or the virtual network service provider to use the network more efficiently and save Capital Expenses (CAPEX) and Operating Expenses (OPEX).
- * Customer services have various Service Level Agreements (SLA) requirements, such as service availability, latency, jitter, packet loss rate, Bit Error Rate (BER), etc. The TE network can satisfy service availability and BER requirements by providing different protection and restoration mechanisms. However, for other SLA requirements (like latency), there are no such mechanisms. In order to provide high-quality services according to the customer SLA, one possible solution is to measure the SLA-related performance parameters, and dynamically provision and optimize services based on the performance monitoring results.
- * Performance monitoring in a large-scale network could generate a huge amount of performance information. Therefore, the appropriate way to deliver the information to the client and network interfaces should be carefully considered.

3. Design of the Data Models

This document describes two YANG models:

- (i) TE Telemetry Model, which provides the TE-Tunnel level of performance monitoring mechanism and scaling intent mechanism that allows scale in/out programming by the customer. (See Section 3.1 & Section 9.1 for details).
- (ii) VN Telemetry Model, which provides the VN level of the aggregated performance monitoring mechanism and scaling intent mechanism that allows scale-in/out programming by the customer (See Section 3.2 & Section 9.2 for details).

3.1. TE Telemetry Model

This model describes the performance telemetry for the TE tunnel. The telemetry data is augmented to the TE tunnel. This model also allows autonomic traffic engineering scaling-intent configuration mechanism on the TE-tunnel level. Various conditions can be set for auto-scaling based on the telemetry data (See Section 6 for details)

As shown in Figure 2, the TE Telemetry Model augments the TE-Tunnel Model to enhance TE performance monitoring capability. This monitoring capability will facilitate the re-optimization and reconfiguration of TE tunnels based on the performance monitoring data collected via the TE Telemetry YANG model.




Figure 2: TE Telemetry Model Relationship

3.2. VN Telemetry Model

As shown in Figure 3, the VN Telemetry Model augments the basic VN model to enhance VN monitoring capability. This monitoring capability will facilitate re-optimization and reconfiguration of VNs based on the performance monitoring data collected via the VN Telemetry YANG model. This model also imports the TE telemetry model to reuse the groupings.




Figure 3: VN Telemetry Model Relationships

This model describes the performance telemetry for the VN model. The telemetry data is augmented to the VN model at the VN Level as well as at the individual VN member level. This model also allows autonomic traffic engineering scaling intent configuration mechanism on the VN level. Scale-in/out criteria might be used for network autonomies in order for the controller to react to a certain set of variations in monitored parameters (See Section 4 for illustrations).

Moreover, this model also provides a mechanism to define aggregated VN telemetry parameters as a grouping of underlying VN-member level telemetry parameters. This is unique to the VN model as a VN comprises multiple VN-members, and each VN-member could be further set across multiple TE tunnels. Grouping operations (such as maximum and mean) could be set at the time of configuration. For example, if the "maximum" grouping operation is used for delay at the VN level, the VN telemetry data is reported as the maximum of {delay_vn_member_1, delay_vn_member_2, ... delay_vn_member_N}. Thus, this telemetry aggregation mechanism allows the aggregation (or grouping) of a certain common set of telemetry values under a grouping operation. This can also be done at the VN-member level to suggest how the end-to-end (E2E) telemetry can be inferred from the per-domain tunnels created and monitored by PNCs. Figure 4 provides an example of interactions.

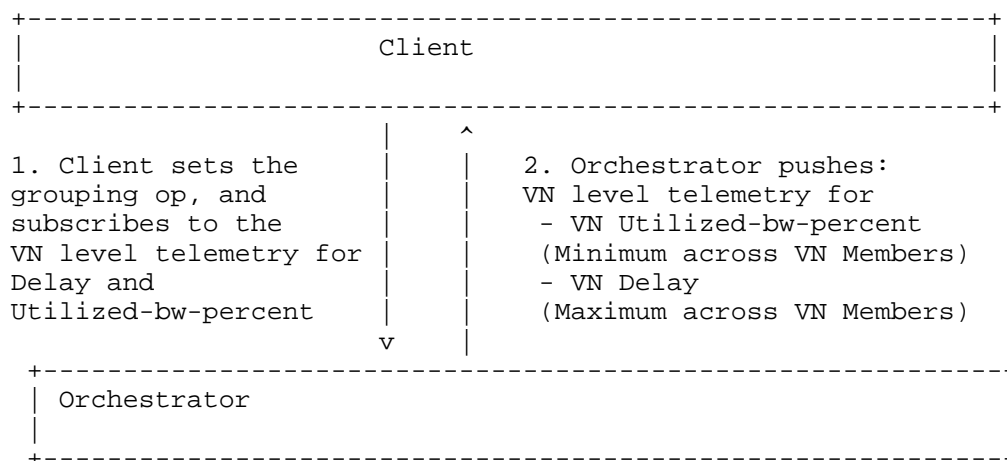


Figure 4: TE Telemetry Model Interactions

3.3. VPN Service Performance Monitoring

The YANG model in [RFC9375] provides network performance monitoring (PM) and VPN service performance monitoring that can be used to monitor and manage network performance on the topology at higher layers or the service topology between VPN sites. Thus the YANG models in this document could be used alongside ietf-network-vpn-pm to understand and correlate the performance monitoring at the VPN service and the underlying TE level.

4. Autonomic Scaling Intent Mechanism

The scaling intent configuration mechanism allows the client to configure automatic scale-in and scale-out mechanisms on both the TE-tunnel and the VN level. Various conditions can be set for auto-scaling based on the PM telemetry data.

There are several parameters involved in the mechanism:

- * Scale-out-intent or Scale-in-intent: whether to scale-out or scale-in.
- * Performance-type: performance metric type (e.g., one-way-delay, one-way-delay-min, one-way-delay-max, two-way-delay, two-way-delay-min, two-way-delay-max, utilized bandwidth, etc.)
- * Threshold-value: the threshold value for a certain Performance-type that triggers scale-in or scale-out.
- * scaling-operation-type: in the case where the scaling condition can be set with one or more performance types, then scaling-operation-type (AND, OR, MIN, MAX, etc.) is applied to these selected performance types and their threshold values.
- * Threshold-time: the duration for which the criteria need to hold true.
- * Cooldown-time: the duration after a scaling action has been triggered, for which there will be no further operation.

The tree in Figure 5 is a part of ietf-te-telemetry tree whose model is presented in full detail in Sections 6 & 7.

```

module: ietf-te-telemetry

augment /te:te/te:tunnels/te:tunnel:
  +--rw te-scaling-intent
  |   +--rw scale-in-intent
  |   |   +--rw threshold-time?      uint32
  |   |   +--rw cooldown-time?      uint32
  |   |   +--rw scaling-condition* [performance-type]
  |   |   |   +--rw performance-type      identityref
  |   |   |   +--rw threshold-value?      scale-value
  |   |   |   +--rw scale-in-operation-type?
  |   |   |       scaling-criteria-operation
  |   |   +--rw scale-in-op?          scale-op
  |   |   +--rw scale?                scale-value
  |   +--rw scale-out-intent
  |   |   +--rw threshold-time?      uint32
  |   |   +--rw cooldown-time?      uint32
  |   |   +--rw scaling-condition* [performance-type]
  |   |   |   +--rw performance-type      identityref
  |   |   |   +--rw threshold-value?      scale-value
  |   |   |   +--rw scale-out-operation-type?
  |   |   |       scaling-criteria-operation
  |   |   +--rw scale-out-op?          scale-op
  |   |   +--rw scale?                scale-value

```

Figure 5: The scaling intent

Let's say the client wants to set the scaling out operation based on two performance-types (e.g., two-way-delay and utilized-bandwidth for a te-tunnel), it can be done as follows:

- * Set Threshold-time: x (sec) (duration for which the criteria must hold true)
- * Set Cooldown-time: y (sec) (the duration after a scaling action has been triggered, for which there will be no further operation)
- * Set AND for the scale-out-operation-type

In the scaling condition's list, the following two components can be set:

List 1: Scaling Condition for Two-way-delay

- * performance type: Two-way-delay
- * threshold-value: z milli-seconds

List 2: Scaling Condition for Utilized bandwidth

- * performance type: Utilized bandwidth
- * threshold-value: w megabytes

Refer to Section 7 for some examples of scaling intent.

5. Performance Monitoring Parameters

This model augments the Tunnel model to include performance parameters from the grouping performance-metrics-attributes from te-types [RFC8776]:

- * one-way-delay
- * one-way-delay-normality
- * one-way-residual-bandwidth
- * one-way-residual-bandwidth-normality
- * one-way-available-bandwidth
- * one-way-available-bandwidth-normality
- * one-way-utilized-bandwidth
- * one-way-utilized-bandwidth-normality
- * two-way-delay
- * two-way-delay-normality

```

+--ro te-telemetry
  +--ro performance-metrics-one-way
    |   +--ro one-way-delay?                               uint32
    |   +--ro one-way-delay-normality?
    |   |       te-types:performance-metrics-normality
    |   +--ro one-way-residual-bandwidth?
    |   |       rt-types:bandwidth-ieee-float32
    |   +--ro one-way-residual-bandwidth-normality?
    |   |       te-types:performance-metrics-normality
    |   +--ro one-way-available-bandwidth?
    |   |       rt-types:bandwidth-ieee-float32
    |   +--ro one-way-available-bandwidth-normality?
    |   |       te-types:performance-metrics-normality
    |   +--ro one-way-utilized-bandwidth?
    |   |       rt-types:bandwidth-ieee-float32
    |   +--ro one-way-utilized-bandwidth-normality?
    |       te-types:performance-metrics-normality
  +--ro performance-metrics-two-way
    +--ro two-way-delay?                               uint32
    +--ro two-way-delay-normality?
        te-types:performance-metrics-normality

```

Figure 6: Performance Monitoring Parameters

6. Notification

This model does not define specific notifications. To enable notifications, the mechanism defined in [RFC8641] and [RFC8640] can be used. This mechanism currently allows the user to:

- * Subscribe to notifications on a per-client basis.
- * Specify subtree filters or xpath filters so that only interested contents will be sent.
- * Specify either periodic or on-demand notifications.

6.1. YANG Push Subscription Examples

[RFC8641] allows subscriber applications to request a continuous, customized stream of updates from a YANG datastore.

The example in Figure 7 shows the way for a client to subscribe to the telemetry information for a particular tunnel (Tunnell). The telemetry parameter that the client is interested in is one-way-delay.

```
<netconf:rpc netconf:message-id="101"
  xmlns:netconf="urn:ietf:params:xml:ns:netconf:base:1.0">
  <establish-subscription
    xmlns="urn:ietf:params:xml:ns:yang:ietf-yang-push:1.0">
    <filter netconf:type="subtree">
      <te xmlns="urn:ietf:params:xml:ns:yang:ietf-te">
        <tunnels>
          <tunnel>
            <name>Tunnell</name>
            <te-telemetry xmlns="urn:ietf:params:xml:ns:yang:
              ietf-te-telemetry">
              <performance-metrics-one-way>
                <one-way-delay/>
              </performance-metrics-one-way>
            </te-telemetry>
          </tunnel>
        </tunnels>
      </te>
    </filter>
    <period>500</period>
    <encoding>encode-xml</encoding>
  </establish-subscription>
</netconf:rpc>
```

Figure 7: TE Tunnel Subscription Example

The example in Figure 8 shows the way for a client to subscribe to the telemetry information for all VNs. The telemetry parameter that the client is interested in is one-way-delay and one-way-utilized-bandwidth.

```
<netconf:rpc netconf:message-id="101"
  xmlns:netconf="urn:ietf:params:xml:ns:netconf:base:1.0">
  <establish-subscription
    xmlns="urn:ietf:params:xml:ns:yang:ietf-yang-push:1.0">
    <filter netconf:type="subtree">
      <virtual-network xmlns="urn:ietf:params:xml:ns:yang:ietf-vn">
        <vn>
          <vn-id/>
          <vn-telemetry xmlns="urn:ietf:params:xml:ns:yang:
            ietf-vn-telemetry">
            <params>
              <performance-metrics-one-way>
                <one-way-delay/>
                <one-way-utilized-bandwidth/>
              </performance-metrics-one-way>
            </params>
          </vn-telemetry>
        </vn>
      </virtual-network>
    </filter>
    <period>500</period>
  </establish-subscription>
</netconf:rpc>
```

Figure 8: VN Subscription Example

7. Scaling Examples

The example in Figure 9 shows the way to configure a TE tunnel with the scaling-out intent to re-optimize when the scaling condition of two-way-delay crossing 100 milliseconds (100000 microseconds) for a threshold of 1 min (60 seconds).

```
<edit-config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <target>
    <running/>
  </target>
  <config>
    <te xmlns="urn:ietf:params:xml:ns:yang:ietf-te">
      <tunnels>
        <tunnel>
          <name>Tunnell</name>
          <te-scaling-intent
            xmlns="urn:ietf:params:xml:ns:yang:
              ietf-te-telemetry">
            <scale-out-intent>
              <threshold-time>
                60
              </threshold-time>
              <scaling-condition>
                <performance-type>
                  two-way-delay
                </performance-type>
                <threshold-value>
                  100000
                </threshold-value>
              </scaling-condition>
            </scale-out-intent>
          </te-scaling-intent>
        </tunnel>
      </tunnels>
    </te>
  </config>
</edit-config>
```

Figure 9: TE Tunnel Scaling Example

The example in Figure 10 shows the way to configure a VN with the scaling-in intent to reduce bandwidth when the scaling condition of utilized-percentage crossing 50 percent for a threshold of 5 minutes (300 seconds).


```
<edit-config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <target>
    <running/>
  </target>
  <config>
    <virtual-network xmlns="urn:ietf:params:xml:ns:yang:ietf-vn">
      <vn>
        <vn-id>VN1</vn-id>
        <vn-scaling-intent
          xmlns="urn:ietf:params:xml:ns:yang:
            ietf-vn-telemetry">
          <scale-in-intent>
            <threshold-time>300</threshold-time>
            <scaling-condition>
              <performance-type>
                utilized-percentage
              </performance-type>
              <threshold-value>
                50
              </threshold-value>
            </scaling-condition>
          </scale-in-intent>
        </vn-scaling-intent>
      </vn>
    </virtual-network>
  </config>
</edit-config>
```

Figure 10: VN Scaling Example

The example in Figure 11 shows the way to configure a VN with the scaling-in when the scaling condition of one-way-delay-variation crossing 100 milliseconds (100000 microseconds) OR one-way-delay crossing 50 milliseconds (50000 microseconds) for a threshold of 2 minutes (120 seconds).

```

<edit-config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <target>
    <running/>
  </target>
  <config>
    <virtual-network xmlns="urn:ietf:params:xml:ns:yang:ietf-vn">
      <vn>
        <vn-id>VN2</vn-id>
        <vn-scaling-intent
          xmlns="urn:ietf:params:xml:ns:yang:
            ietf-vn-telemetry">
          <scale-in-intent>
            <threshold-time>120</threshold-time>
            <scaling-condition>
              <performance-type>
                one-way-delay-variation
              </performance-type>
              <threshold-value>
                100000
              </threshold-value>
              <scale-in-operation-type>
                OR
              </scale-in-operation-type>
            </scaling-condition>
            <scaling-condition>
              <performance-type>
                one-way-delay
              </performance-type>
              <threshold-value>
                50000
              </threshold-value>
              <scale-in-operation-type>
                OR
              </scale-in-operation-type>
            </scaling-condition>
          </scale-in-intent>
        </vn-scaling-intent>
      </vn>
    </virtual-network>
  </config>
</edit-config>

```

Figure 11: VN Scaling Example with OR condition

The example in Figure 12 shows the way to configure a grouping operation at the VN level to require that the VN level one-way-delay needs to be reported as the max of the one-way-delay at the VN-member level, whereas the utilized-percentage is the mean.

```
<edit-config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <target>
    <running/>
  </target>
  <config>
    <virtual-network xmlns="urn:ietf:params:xml:ns:yang:ietf-vn">
      <vn>
        <vn-id>VN1</vn-id>
        <vn-telemetry
          xmlns="urn:ietf:params:xml:ns:yang:
            ietf-vn-telemetry">
          <operation>
            <performance-type>
              one-way-delay
            </performance-type>
            <grouping-operation>
              maximum
            </grouping-operation>
          </operation>
          <operation>
            <performance-type>
              utilized-percentage
            </performance-type>
            <grouping-operation>
              mean
            </grouping-operation>
          </operation>
        </vn-telemetry>
      </vn>
    </virtual-network>
  </config>
</edit-config>
```

Figure 12: VN Grouping Operation Example

8. YANG Data Tree

```

module: ietf-te-telemetry

augment /te:te/te:tunnels/te:tunnel:
  +--rw te-scaling-intent
  |   +--rw scale-in-intent
  |   |   +--rw threshold-time?          uint32
  |   |   +--rw cooldown-time?          uint32
  |   |   +--rw scaling-condition* [performance-type]
  |   |   |   +--rw performance-type          identityref
  |   |   |   +--rw threshold-value?         scale-value
  |   |   |   +--rw scale-in-operation-type?
  |   |   |       scaling-criteria-operation
  |   |   +--rw scale-in-op?             scale-op
  |   |   +--rw scale?                   scale-value
  |   +--rw scale-out-intent
  |   |   +--rw threshold-time?          uint32
  |   |   +--rw cooldown-time?          uint32
  |   |   +--rw scaling-condition* [performance-type]
  |   |   |   +--rw performance-type          identityref
  |   |   |   +--rw threshold-value?         scale-value
  |   |   |   +--rw scale-out-operation-type?
  |   |   |       scaling-criteria-operation
  |   |   +--rw scale-out-op?            scale-op
  |   |   +--rw scale?                   scale-value
  +--ro te-telemetry
  |   +--ro performance-metrics-one-way
  |   |   +--ro one-way-delay?            uint32
  |   |   +--ro one-way-delay-normality?
  |   |   |       te-types:performance-metrics-normality
  |   |   +--ro one-way-residual-bandwidth?
  |   |   |       rt-types:bandwidth-ieee-float32
  |   |   +--ro one-way-residual-bandwidth-normality?
  |   |   |       te-types:performance-metrics-normality
  |   |   +--ro one-way-available-bandwidth?
  |   |   |       rt-types:bandwidth-ieee-float32
  |   |   +--ro one-way-available-bandwidth-normality?
  |   |   |       te-types:performance-metrics-normality
  |   |   +--ro one-way-utilized-bandwidth?
  |   |   |       rt-types:bandwidth-ieee-float32
  |   |   +--ro one-way-utilized-bandwidth-normality?
  |   |   |       te-types:performance-metrics-normality
  |   +--ro performance-metrics-two-way
  |   |   +--ro two-way-delay?            uint32
  |   |   +--ro two-way-delay-normality?
  |   |   |       te-types:performance-metrics-normality

```

Figure 13: ietf-te-telemetry YANG model tree

```
module: ietf-vn-telemetry
```

```
augment /vn:virtual-network/vn:vn:
```

```

  +--rw vn-scaling-intent
  |   +--rw scale-in-intent
  |   |   +--rw threshold-time?          uint32
  |   |   +--rw cooldown-time?          uint32
  |   |   +--rw scaling-condition* [performance-type]
  |   |   |   +--rw performance-type          identityref
  |   |   |   +--rw threshold-value?         scale-value
  |   |   |   +--rw scale-in-operation-type?
  |   |   |       scaling-criteria-operation
  |   |   +--rw scale-in-op?             scale-op
  |   |   +--rw scale?                   scale-value
  |   +--rw scale-out-intent
  |   |   +--rw threshold-time?          uint32
  |   |   +--rw cooldown-time?          uint32
  |   |   +--rw scaling-condition* [performance-type]
  |   |   |   +--rw performance-type          identityref
  |   |   |   +--rw threshold-value?         scale-value
  |   |   |   +--rw scale-out-operation-type?
  |   |   |       scaling-criteria-operation
  |   |   +--rw scale-out-op?            scale-op
  |   |   +--rw scale?                   scale-value
  +--rw vn-telemetry
  |   +--ro params
  |   |   +--ro performance-metrics-one-way
  |   |   |   +--ro one-way-delay?          uint32
  |   |   |   +--ro one-way-delay-normality?
  |   |   |       te-types:performance-metrics-normality
  |   |   +--ro one-way-residual-bandwidth?
  |   |   |   rt-types:bandwidth-ieee-float32
  |   |   +--ro one-way-residual-bandwidth-normality?
  |   |   |   te-types:performance-metrics-normality
  |   |   +--ro one-way-available-bandwidth?
  |   |   |   rt-types:bandwidth-ieee-float32
  |   |   +--ro one-way-available-bandwidth-normality?
  |   |   |   te-types:performance-metrics-normality
  |   |   +--ro one-way-utilized-bandwidth?
  |   |   |   rt-types:bandwidth-ieee-float32
  |   |   +--ro one-way-utilized-bandwidth-normality?
  |   |   |   te-types:performance-metrics-normality
  |   |   +--ro performance-metrics-two-way
  |   |   |   +--ro two-way-delay?          uint32
  |   |   |   +--ro two-way-delay-normality?
  |   |   |       te-types:performance-metrics-normality
  +--rw operation* [performance-type]
  |   +--rw performance-type          identityref

```

```

      +--rw grouping-operation?  identityref
augment /vn:virtual-network/vn:vn/vn:vn-member:
  +--rw vn-member-telemetry
    +--ro params
      +--ro performance-metrics-one-way
        +--ro one-way-delay?          uint32
        +--ro one-way-delay-normality?
          | te-types:performance-metrics-normality
        +--ro one-way-residual-bandwidth?
          | rt-types:bandwidth-ieee-float32
        +--ro one-way-residual-bandwidth-normality?
          | te-types:performance-metrics-normality
        +--ro one-way-available-bandwidth?
          | rt-types:bandwidth-ieee-float32
        +--ro one-way-available-bandwidth-normality?
          | te-types:performance-metrics-normality
        +--ro one-way-utilized-bandwidth?
          | rt-types:bandwidth-ieee-float32
        +--ro one-way-utilized-bandwidth-normality?
          | te-types:performance-metrics-normality
      +--ro performance-metrics-two-way
        +--ro two-way-delay?          uint32
        +--ro two-way-delay-normality?
          | te-types:performance-metrics-normality
      +--ro te-tunnel-ref*
        -> /te:te/tunnels/tunnel/name
    +--rw operation* [performance-type]
      +--rw performance-type          identityref
      +--rw grouping-operation?      identityref

```

Figure 14: ietf-vn-telemetry YANG model tree

9. YANG Data Model

9.1. ietf-te-telemetry model

The YANG code is as follows:

```

<CODE BEGINS> file "ietf-te-telemetry@2025-10-13.yang"
module ietf-te-telemetry {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-te-telemetry";
  prefix te-tel;

  /* Import TE */

  import ietf-te {
    prefix te;

```

```
reference
  "I-D.ietf-teas-yang-te: A YANG Data Model for Traffic
    Engineering Tunnels and Interfaces";
}

/* Import TE Common types */

import ietf-te-types {
  prefix te-types;
  reference
    "RFC 8776: Common YANG Data Types for Traffic Engineering";
}

/* Import Routing Common types */

import ietf-routing-types {
  prefix rt-types;
  reference
    "RFC 8294: Common YANG Data Types for the Routing Area";
}

organization
  "IETF Traffic Engineering Architecture and Signaling (TEAS)
    Working Group";
contact
  "WG Web:  <https://datatracker.ietf.org/wg/teas/>
  WG List:  <mailto:teas@ietf.org>
  Editor:   Young Lee <younglee.tx@gmail.com>
           Dhruv Dhody <dhruv.ietf@gmail.com>";
description
  "This module describes the YANG data model for performance
    monitoring parameters (telemetry data) for TE tunnels.

    Copyright (c) 2025 IETF Trust and the persons identified as
    authors of the code.  All rights reserved.

    Redistribution and use in source and binary forms, with or
    without modification, is permitted pursuant to, and subject to
    the license terms contained in, the Revised BSD License set
    forth in Section 4.c of the IETF Trust's Legal Provisions
    Relating to IETF Documents
    (https://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX; see the
    RFC itself for full legal notices.";

/* Note: The RFC Editor will replace XXXX with the number
   assigned to the RFC once draft-ietf-teas-pm-telemetry-
```

```
autonomics becomes an RFC.*/  
  
revision 2025-10-13 {  
  description  
    "Initial revision.";   
  reference  
    "RFC XXXX: YANG models for VN/TE Performance Monitoring  
    Telemetry and Scaling Intent Autonomics";  
}  
  
identity telemetry-param-type {  
  description  
    "Base identity for telemetry parameter types";  
}  
  
identity one-way-delay {  
  base telemetry-param-type;  
  description  
    "To specify average Delay in one (forward) direction  
    in microseconds.  
  
    At the VN level, it is the maximum delay of the  
    VN-members.  
  
    The threshold-value for this type is interpreted as  
    microseconds.";  
  reference  
    "RFC 7471: OSPF Traffic Engineering (TE) Metric Extensions.  
    RFC 8570: IS-IS Traffic Engineering (TE) Metric Extensions.  
    RFC 7823: Performance-Based Path Selection for Explicitly  
    Routed Label Switched Paths (LSPs) Using TE Metric  
    Extensions";  
}  
  
identity two-way-delay {  
  base telemetry-param-type;  
  description  
    "To specify average Delay in both (forward and reverse)  
    directions in microseconds.  
  
    At the VN level, it is the maximum delay of the  
    VN-members.  
  
    The threshold-value for this type is interpreted as  
    microseconds.";  
  reference  
    "RFC 7471: OSPF Traffic Engineering (TE) Metric Extensions.  
    RFC 8570: IS-IS Traffic Engineering (TE) Metric Extensions.
```



```
    RFC 7823: Performance-Based Path Selection for Explicitly
    Routed Label Switched Paths (LSPs) Using TE Metric
    Extensions";
}

identity one-way-delay-variation {
  base telemetry-param-type;
  description
    "To specify average Delay Variation in one (forward) direction
    in microseconds.

    At the VN level, it is the max delay variation of the
    VN-members.

    The threshold-value for this type is interpreted as
    microseconds.";
  reference
    "RFC 7471: OSPF Traffic Engineering (TE) Metric Extensions.
    RFC 8570: IS-IS Traffic Engineering (TE) Metric Extensions.
    RFC 7823: Performance-Based Path Selection for Explicitly
    Routed Label Switched Paths (LSPs) Using TE Metric
    Extensions";
}

identity two-way-delay-variation {
  base telemetry-param-type;
  description
    "To specify average Delay Variation in both (forward and
    reverse) directions in microseconds.

    At the VN level, it is the max delay variation of the
    VN-members.

    The threshold-value for this type is interpreted as
    microseconds.";
  reference
    "RFC 7471: OSPF Traffic Engineering (TE) Metric Extensions.
    RFC 8570: IS-IS Traffic Engineering (TE) Metric Extensions.
    RFC 7823: Performance-Based Path Selection for Explicitly
    Routed Label Switched Paths (LSPs) Using TE Metric
    Extensions";
}

identity utilized-bandwidth {
  base telemetry-param-type;
  description
    "To specify utilized bandwidth over the specified source
    and destination in bytes per second.
```

```
    The threshold-value for this type is interpreted as
    bytes per second.";
reference
    "RFC 7471: OSPF Traffic Engineering (TE) Metric Extensions.
    RFC 8570: IS-IS Traffic Engineering (TE) Metric Extensions.
    RFC 7823: Performance-Based Path Selection for Explicitly
    Routed Label Switched Paths (LSPs) Using TE Metric
    Extensions";
}

identity utilized-percentage {
    base telemetry-param-type;
    description
        "To specify utilization percentage of the entity
        (e.g., tunnel, link, etc.)";
}

/* Typedef */

typedef scale-op {
    type enumeration {
        enum UP {
            description
                "Scale up the bandwidth capacity";
        }
        enum DOWN {
            description
                "Scale down the bandwidth capacity";
        }
    }
    description
        "Scaling operation";
}

typedef scaling-criteria-operation {
    type enumeration {
        enum AND {
            description
                "AND operation";
        }
        enum OR {
            description
                "OR operation";
        }
    }
    description
        "Operations to analyze the list of scaling criteria.";
}
```

```
typedef scale-value {
  type union {
    type uint32;
    type rt-types:bandwidth-ieee-float32;
    type rt-types:percentage;
    type te-types:te-bandwidth;
  }
  description
    "Union of scale values of various types";
}

grouping scaling-duration {
  description
    "Base scaling criteria durations";
  leaf threshold-time {
    type uint32;
    units "seconds";
    description
      "The duration for which the criteria must hold true. The
       value of '0' indicates an immediate scaling with no
       duration to wait.";
  }
  leaf cooldown-time {
    type uint32;
    units "seconds";
    description
      "The duration after a scaling-in/scaling-out action has been
       triggered, for which there will be no further operation.
       The value of '0' indicates an immediate scaling action with
       no duration to wait.";
  }
}

grouping scaling-criteria {
  description
    "Grouping for scaling criteria";
  leaf performance-type {
    type identityref {
      base telemetry-param-type;
    }
    description
      "Reference to the tunnel level telemetry type";
  }
  leaf threshold-value {
    type scale-value;
    description
      "Scaling threshold for the telemetry parameter type. The
       value is to be interpreted as per the type.";
  }
}
```

```
    }  
  }  
  
  grouping scaling-in-intent {  
    description  
      "Basic scaling in intent";  
    uses scaling-duration;  
    list scaling-condition {  
      key "performance-type";  
      description  
        "Scaling conditions";  
      uses scaling-criteria;  
      leaf scale-in-operation-type {  
        type scaling-criteria-operation;  
        default "AND";  
        description  
          "Operation to be applied to check between scaling criteria  
          if the scale-in threshold condition has been met.  
          Defaults to AND.";  
      }  
    }  
    leaf scale-in-op {  
      type scale-op;  
      default "DOWN";  
      description  
        "The scaling operation to be performed when scaling condition  
        is met";  
    }  
    leaf scale {  
      type scale-value;  
      description  
        "Additional scaling-by information to be interpreted as per  
        the scale-in-op.";  
    }  
  }  
}  
  
grouping scaling-out-intent {  
  description  
    "Basic scaling out intent";  
  uses scaling-duration;  
  list scaling-condition {  
    key "performance-type";  
    description  
      "Scaling conditions";  
    uses scaling-criteria;  
    leaf scale-out-operation-type {  
      type scaling-criteria-operation;  
      default "OR";  
    }  
  }  
}
```

```
        description
          "Operation to be applied to check between scaling criteria
           if the scale-out threshold condition has been met.
           Defaults to OR.";
      }
    }
  leaf scale-out-op {
    type scale-op;
    default "UP";
    description
      "The scaling operation to be performed when scaling condition
       is met.";
  }
  leaf scale {
    type scale-value;
    description
      "Additional scaling-by information to be interpreted as per
       the scale-out-op.";
  }
}

augment "/te:te/te:tunnels/te:tunnel" {
  description
    "Augmentation parameters for config scaling-criteria TE
     tunnel topologies. Scale in/out criteria might be used
     for network autonomies in order for the controller to
     react to a certain set of monitored parameters.";
  container te-scaling-intent {
    description
      "The scaling intent";
    container scale-in-intent {
      description
        "scale-in";
      uses scaling-in-intent;
    }
    container scale-out-intent {
      description
        "scale-out";
      uses scaling-out-intent;
    }
  }
  container te-telemetry {
    config false;
    description
      "Telemetry Data";
    uses te-types:performance-metrics-attributes;
  }
}
```

```
}  
<CODE ENDS>
```

9.2. ietf-vn-telemetry model

The YANG code is as follows:

```
<CODE BEGINS> file "ietf-vn-telemetry@2025-10-13.yang"  
module ietf-vn-telemetry {  
  yang-version 1.1;  
  namespace "urn:ietf:params:xml:ns:yang:ietf-vn-telemetry";  
  prefix vn-tel;  
  
  /* Import VN */  
  
  import ietf-vn {  
    prefix vn;  
    reference  
      "RFC 9731: A YANG Data Model for Virtual Network (VN)  
      Operations";  
  }  
  
  /* Import TE */  
  
  import ietf-te {  
    prefix te;  
    reference  
      "I-D.ietf-teas-yang-te: A YANG Data Model for Traffic  
      Engineering Tunnels and Interfaces";  
  }  
  
  /* Import TE Common types */  
  
  import ietf-te-types {  
    prefix te-types;  
    reference  
      "RFC 8776: Common YANG Data Types for Traffic Engineering";  
  }  
  
  /* Import TE Telemetry */  
  
  import ietf-te-telemetry {  
    prefix te-tel;  
    reference  
      "RFC XXXX: YANG models for VN/TE Performance Monitoring  
      Telemetry and Scaling Intent Autonomics";  
  }  
}
```

```
/* Note: The RFC Editor will replace XXXX with the number
   assigned to this draft.*/
```

```
organization
```

```
"IETF Traffic Engineering Architecture and Signaling (TEAS)
  Working Group";
```

```
contact
```

```
"WG Web:  <https://datatracker.ietf.org/wg/teas/>
WG List:  <mailto:teas@ietf.org>
Editor:   Young Lee <younglee.tx@gmail.com>
          Dhruv Dhody <dhruv.ietf@gmail.com>";
```

```
description
```

```
"This module describes YANG data models for performance
  monitoring parameters (telemetry data) for Virtual Network
  (VN)."
```

Copyright (c) 2025 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Revised BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.";

```
/* Note: The RFC Editor will replace XXXX with the number
   assigned to the RFC once draft-lee-teas-pm-telemetry-
   autonomics becomes an RFC.*/
```

```
revision 2025-10-13 {
```

```
  description
```

```
    "Initial revision.";
```

```
  reference
```

```
    "RFC XXXX: YANG models for VN/TE Performance Monitoring
      Telemetry and Scaling Intent Autonomics";
```

```
}
```

```
identity grouping-op {
```

```
  description
```

```
    "Base identity for grouping-operation";
```

```
}
```

```
identity minimum {
```

```
  base grouping-op;
```

```
    description
      "Select the minimum of the monitored parameters";
  }

  identity maximum {
    base grouping-op;
    description
      "The maximum of the monitored parameters";
  }

  identity mean {
    base grouping-op;
    description
      "The mean of the monitored parameters";
  }

  identity standard-deviation {
    base grouping-op;
    description
      "The standard deviation of the monitored parameters";
  }

  identity sum {
    base grouping-op;
    description
      "The sum of the monitored parameters";
  }

  identity and {
    base grouping-op;
    description
      "Logical AND operation";
  }

  identity or {
    base grouping-op;
    description
      "Logical OR operation";
  }

  grouping grouping-operation {
    list operation {
      key "performance-type";
      leaf performance-type {
        type identityref {
          base te-tel:telemetry-param-type;
        }
      }
      description

```



```
        "Reference to the tunnel level telemetry type";
    }
    leaf grouping-operation {
        type identityref {
            base grouping-op;
        }
        description
            "Describes the operation to apply to the underlying
             TE tunnels";
    }
    description
        "Grouping operation for each performance-type";
}
description
    "Grouping operation for each performance-type";
}

augment "/vn:virtual-network/vn:vn" {
    description
        "Augmentation parameters for state TE VN topologies.";
    container vn-scaling-intent {
        description
            "scaling intent";
        container scale-in-intent {
            description
                "VN scale-in";
            uses te-tel:scaling-in-intent;
        }
        container scale-out-intent {
            description
                "VN scale-out";
            uses te-tel:scaling-out-intent;
        }
    }
    container vn-telemetry {
        description
            "VN telemetry params";
        container params {
            config false;
            description
                "Read-only telemetry parameters";
            uses te-types:performance-metrics-attributes;
        }
        uses grouping-operation;
    }
}

augment "/vn:virtual-network/vn:vn/vn:vn-member" {
```

```
description
  "Augmentation parameters for state TE vn member topologies.";
container vn-member-telemetry {
  description
    "VN member telemetry params";
  container params {
    config false;
    description
      "Read-only telemetry parameters";
    uses te-types:performance-metrics-attributes;
    leaf-list te-tunnel-ref {
      type leafref {
        path "/te:te/te:tunnels/te:tunnel/te:name";
      }
      description
        "A list of underlying TE tunnels that form the
        VN-member";
    }
  }
  uses grouping-operation;
}
}
}
<CODE ENDS>
```

10. Security Considerations

This section is modeled after the template described in Section 3.7.1 of [I-D.ietf-netmod-rfc8407bis].

The "ietf-te-telemetry" and "ietf-vn-telemetry" YANG modules define data models that are designed to be accessed via YANG-based management protocols, such as NETCONF [RFC6241] and RESTCONF [RFC8040]. These protocols have to use a secure transport layer (e.g., SSH [RFC4252], TLS [RFC8446], and QUIC [RFC9000]) and have to use mutual authentication.

The Network Configuration Access Control Model (NACM) [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a preconfigured subset of all available NETCONF or RESTCONF protocol operations and content.

There are a number of data nodes defined in this YANG module that are writable/creatable/deletable (i.e., "config true", which is the default). All writable data nodes are likely to be reasonably sensitive or vulnerable in some network environments. Write operations (e.g., edit-config) and delete operations to these data nodes without proper protection or authentication can have a negative

effect on network operations. The following subtrees and data nodes have particular sensitivities/vulnerabilities and are potentially disruptive if abused:

- * /te:te/te:tunnels/te:tunnel/te-scaling-intent/scale-in-intent
- * /te:te/te:tunnels/te:tunnel/te-scaling-intent/scale-out-intent
- * /vn:virtual-network/vn:vn/vn-scaling-intent/scale-in-intent
- * /vn:virtual-network/vn:vn/vn-scaling-intent/scale-out-intent

Malicious or incorrect modification of these scaling intent parameters could trigger unintended network behavior, including:

- * Inappropriate scaling-out leading to unnecessary resource consumption and bandwidth allocation;
- * Inappropriate scaling-in causing service degradation or capacity shortages;
- * Setting overly aggressive scaling thresholds that create network instability through rapid oscillation between scaled states;
- * Denial of service by continuously triggering scaling operations that overwhelm network controllers;

Unauthorized modification of the following subtrees could compromise network visibility and monitoring accuracy:

- * /vn:virtual-network/vn:vn/vn-telemetry/operation
- * /vn:virtual-network/vn:vn/vn:vn-member/vn-member-telemetry/operation

An attacker with write access to these operation nodes could manipulate aggregation functions (e.g., changing 'max' to 'mean' or 'min') to mask performance degradation, hide capacity issues, or misrepresent network state. This could lead to incorrect operational decisions, delayed incident response, or failure to detect service-level agreement violations.

Some of the readable data nodes in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control read access (e.g., via get, get-config, or notification) to these data nodes. The following subtrees and data nodes have particular sensitivities and raise privacy concerns:

- * /te:te/te:tunnels/te:tunnel/te-telemetry
- * /vn:virtual-network/vn:vn/vn-telemetry
- * /vn:virtual-network/vn:vn/vn:vn-member/vn-member-telemetry

The notifications defined in this document may also expose telemetry state; therefore access to notifications must be controlled (e.g., via NACM) similarly to read access.

Unauthorized access to these telemetry subtrees could expose sensitive operational information, including: network performance characteristics (delay, jitter, packet loss), bandwidth utilization patterns, and tunnel/VN topology. This information could be exploited to:

- * Identify high-value targets for attacks;
- * Determine optimal timing for denial-of-service attacks when resources are constrained;
- * Infer business relationships and traffic patterns;
- * Gain competitive intelligence about network capacity and service quality.

Such exposure raises both security and privacy concerns for network operators and their customers.

There are no particularly sensitive RPC or action operations.

This YANG module uses groupings from other YANG modules that define nodes that may be considered sensitive or vulnerable in network environments. Refer to the Security Considerations of [RFC8776] for information as to which nodes may be considered sensitive or vulnerable in network environments.

11. IANA Considerations

This document registers the following namespace URIs in the IETF XML registry [RFC3688]:

```
-----  
URI: urn:ietf:params:xml:ns:yang:ietf-te-telemetry  
Registrant Contact: The IESG.  
XML: N/A, the requested URI is an XML namespace.  
-----
```

```
-----  
URI: urn:ietf:params:xml:ns:yang:ietf-vn-telemetry  
Registrant Contact: The IESG.  
XML: N/A, the requested URI is an XML namespace.  
-----
```

This document registers the following YANG modules in the YANG Module Names registry [RFC6020]:

```
-----  
name:          ietf-te-telemetry  
namespace:     urn:ietf:params:xml:ns:yang:ietf-te-telemetry  
prefix:        te-tel  
reference:     RFC XXXX  
-----
```

```
-----  
name:          ietf-vn-telemetry  
namespace:     urn:ietf:params:xml:ns:yang:ietf-vn-telemetry  
prefix:        vn-tel  
reference:     RFC XXXX  
-----
```

12. Acknowledgments

Special thanks to Young Lee, who has provided support during the development of this document.

We also thank Adrian Farrel, Rakesh Gandhi, Tarek Saad, Igor Bryskin, Kenichi Ogaki, and Greg Mirsky for useful discussions and their suggestions for this work.

Thanks to Reshad Rahman for an excellent YANGDOCTOR review.

13. References

13.1. Normative References

[I-D.ietf-netmod-rfc8407bis]
Bierman, A., Boucadair, M., and Q. Wu, "Guidelines for Authors and Reviewers of Documents Containing YANG Data Models", Work in Progress, Internet-Draft, draft-ietf-

netmod-rfc8407bis-28, 5 June 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-netmod-rfc8407bis-28>>.

[I-D.ietf-teas-yang-te]

Saad, T., Gandhi, R., Liu, X., Beeram, V. P., and I. Bryskin, "A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths, and Interfaces", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-41, 20 January 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-41>>.

[RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.

[RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.

[RFC7926] Farrel, A., Ed., Drake, J., Bitar, N., Swallow, G., Ceccarelli, D., and X. Zhang, "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", BCP 206, RFC 7926, DOI 10.17487/RFC7926, July 2016, <<https://www.rfc-editor.org/info/rfc7926>>.

[RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.

[RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.

[RFC8233] Dhody, D., Wu, Q., Manral, V., Ali, Z., and K. Kumaki, "Extensions to the Path Computation Element Communication Protocol (PCEP) to Compute Service-Aware Label Switched Paths (LSPs)", RFC 8233, DOI 10.17487/RFC8233, September 2017, <<https://www.rfc-editor.org/info/rfc8233>>.

- [RFC8294] Liu, X., Qu, Y., Lindem, A., Hopps, C., and L. Berger, "Common YANG Data Types for the Routing Area", RFC 8294, DOI 10.17487/RFC8294, December 2017, <<https://www.rfc-editor.org/info/rfc8294>>.
- [RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", BCP 215, RFC 8340, DOI 10.17487/RFC8340, March 2018, <<https://www.rfc-editor.org/info/rfc8340>>.
- [RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <<https://www.rfc-editor.org/info/rfc8341>>.
- [RFC8342] Bjorklund, M., Schoenwaelder, J., Shafer, P., Watsen, K., and R. Wilton, "Network Management Datastore Architecture (NMDA)", RFC 8342, DOI 10.17487/RFC8342, March 2018, <<https://www.rfc-editor.org/info/rfc8342>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8640] Voit, E., Clemm, A., Gonzalez Prieto, A., Nilsen-Nygaard, E., and A. Tripathy, "Dynamic Subscription to YANG Events and Datastores over NETCONF", RFC 8640, DOI 10.17487/RFC8640, September 2019, <<https://www.rfc-editor.org/info/rfc8640>>.
- [RFC8641] Clemm, A. and E. Voit, "Subscription to YANG Notifications for Datastore Updates", RFC 8641, DOI 10.17487/RFC8641, September 2019, <<https://www.rfc-editor.org/info/rfc8641>>.
- [RFC8776] Saad, T., Gandhi, R., Liu, X., Beeram, V., and I. Bryskin, "Common YANG Data Types for Traffic Engineering", RFC 8776, DOI 10.17487/RFC8776, June 2020, <<https://www.rfc-editor.org/info/rfc8776>>.
- [RFC9731] Lee, Y., Ed., Dhody, D., Ed., Ceccarelli, D., Bryskin, I., and B. Yoon, "A YANG Data Model for Virtual Network (VN) Operations", RFC 9731, DOI 10.17487/RFC9731, March 2025, <<https://www.rfc-editor.org/info/rfc9731>>.

13.2. Informative References

- [RFC4252] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Authentication Protocol", RFC 4252, DOI 10.17487/RFC4252, January 2006, <<https://www.rfc-editor.org/info/rfc4252>>.

- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", RFC 7471, DOI 10.17487/RFC7471, March 2015, <<https://www.rfc-editor.org/info/rfc7471>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [RFC7823] Atlas, A., Drake, J., Giacalone, S., and S. Previdi, "Performance-Based Path Selection for Explicitly Routed Label Switched Paths (LSPs) Using TE Metric Extensions", RFC 7823, DOI 10.17487/RFC7823, May 2016, <<https://www.rfc-editor.org/info/rfc7823>>.
- [RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", RFC 8309, DOI 10.17487/RFC8309, January 2018, <<https://www.rfc-editor.org/info/rfc8309>>.
- [RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <<https://www.rfc-editor.org/info/rfc8453>>.
- [RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <<https://www.rfc-editor.org/info/rfc8570>>.
- [RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/info/rfc9000>>.
- [RFC9375] Wu, B., Ed., Wu, Q., Ed., Boucadair, M., Ed., Gonzalez de Dios, O., and B. Wen, "A YANG Data Model for Network and VPN Service Performance Monitoring", RFC 9375, DOI 10.17487/RFC9375, April 2023, <<https://www.rfc-editor.org/info/rfc9375>>.

Appendix A. Out of Scope

This document exclusively focuses on performance monitoring telemetry and scaling intent mechanisms of the underlying transport (TE-tunnels and Virtual Networks (VNs)). The performance monitoring of the services is out of scope. See Section 3.3 for details about VPN performance monitoring. Similarly, performance monitoring of IETF network slices could be developed, and it is clearly out of the scope

of this document.

Appendix B. Contributors

The following have contributed significantly and should be considered as co-author:

Satish Karunanithi
Kochava
India
Email: satish.karunanithi@gmail.com

Authors' Addresses

Dhruv Dhody (editor)
Huawei
India
Email: dhruv.ietf@gmail.com

Daniel King
Lancaster University
Email: d.king@lancaster.ac.uk

Ricard Vilalta
CTTC
Centre Tecnologic de Telecomunicacions de Catalunya (CTTC/CERCA)
Barcelona
Spain
Email: ricard.vilalta@cttc.es

Italo Busi
Huawei Technologies
Email: italo.busi@huawei.com

Daniele Ceccarelli
Cisco
Email: daniele.ietf@gmail.com