

Workgroup: TEAS Working Group
Internet-Draft:
draft-ietf-teas-actn-pm-telemetry-
autonomics-09

Published: 11 July 2022

Intended Status: Standards Track

Expires: 12 January 2023

Authors: Y. Lee, Ed. D. Dhody, Ed.
 Samsung Electronics Huawei Technologies
 R. Vilalta D. King D. Ceccarelli
 CTTC Lancaster University Ericsson

**YANG models for Virtual Network (VN)/TE Performance Monitoring
Telemetry and Scaling Intent Autonomics**

Abstract

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

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 12 January 2023.

Copyright Notice

Copyright (c) 2022 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](#)
 - [1.1. Terminology](#)
 - [1.2. Tree Diagram](#)
 - [1.3. Prefixes in Data Node Names](#)
- [2. Use-Cases](#)
- [3. Design of the Data Models](#)
 - [3.1. TE Telemetry Model](#)
 - [3.2. VN Telemetry Model](#)
 - [3.3. VPN Service Performance Monitoring](#)
- [4. Autonomic Scaling Intent Mechanism](#)
- [5. Performance Monitoring Parameters](#)
- [6. Notification](#)
 - [6.1. YANG Push Subscription Examples](#)
- [7. Scaling Examples](#)
- [8. YANG Data Tree](#)
- [9. YANG Data Model](#)
 - [9.1. ietf-te-telemetry model](#)
 - [9.2. ietf-vn-telemetry model](#)
- [10. Security Considerations](#)
- [11. IANA Considerations](#)
- [12. Acknowledgments](#)
- [13. References](#)
 - [13.1. Normative References](#)
 - [13.2. Informative References](#)
- [Appendix A. Out of Scope](#)
- [Appendix B. Contributors](#)
- [Authors' Addresses](#)

1. Introduction

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

The models presented in this draft allow the applications hosted by the customers to subscribe to and monitor their key performance data of their interest on the level of VN [[I-D.ietf-teas-actn-vn-yang](#)] or TE-tunnel [[I-D.ietf-teas-yang-te](#)]. The key characteristic of the models presented in this document is a top-down programmability that allows the applications hosted by the customers to subscribe to and monitor key performance data of their interest and autonomic scaling intent mechanism on the level of VN as well as TE-tunnel.

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 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 mechanism specified in [[RFC8641](#)] and [[RFC8640](#)]. It also provides an ability to program their customized automatic scaling in/out intent as well. These models can be utilized by a client network controller to initiate the capabilities via a NETCONF [[RFC8341](#)] or a RESTCONF [[RFC8040](#)] interface.

The term 'Performance monitoring' in this document refers to 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., MSDC/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 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., MSDC) has to provide VN or TE-tunnel level information. This requires the controller to have the capability 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 re-shape its own resources. "Scale out" refers to improve 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 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, 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 | [I-D.ietf-teas-actn-vn-yang] |
| vn-tel | ietf-vn-telemetry | [RFCXXXX] |

Table 1: Prefixes and corresponding YANG modules

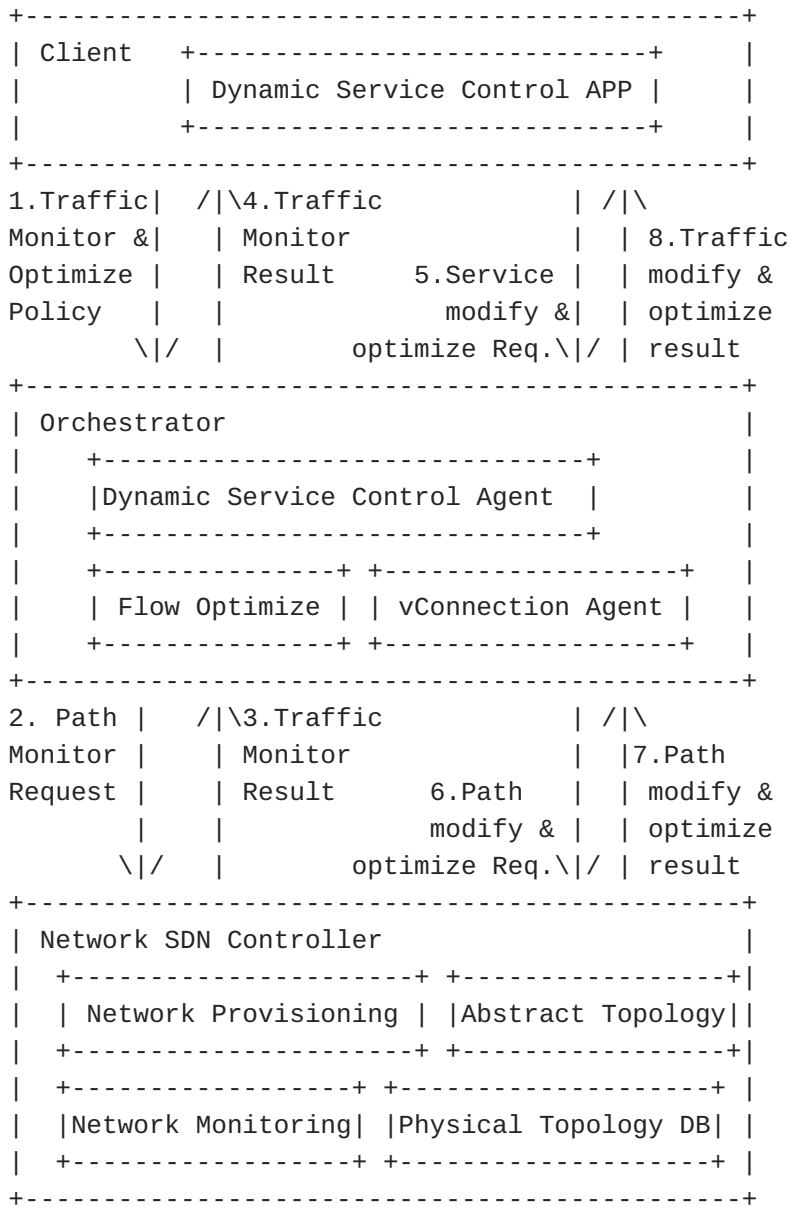
Note: The RFC Editor is requested to replace XXXX with the number assigned to the RFC once this draft becomes an RFC, and to remove this note.

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 the an example of a high-level work-flow 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: Work flow for dynamic service control based on traffic monitoring

Some of the key points are as follows:

*Network traffic monitoring is important to facilitate 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 Expense (CAPEX) and Operating Expense (OPEX).

*Customer services have various Service Level Agreement (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 at 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 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 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 autonomics 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 is made up of multiple VN-members and further each VN-member could be set across multiple TE tunnels. Grouping operation (such as maximum, mean) could be set at the time of configuration. For example, if "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 be inferred from the per domain tunnels created and monitored by PNCs. The [Figure 4](#) provides an example interaction.




Figure 4: TE Telemetry Model Interactions

3.3. VPN Service Performance Monitoring

The YANG model in [[I-D.ietf-opsawg-yang-vpn-service-pm](#)] 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 layer or the service topology between VPN sites. Thus the YANG models in this document could be used along side with 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 a number of 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 case where 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 its threshold values.

*Threshold-time: the duration for which the criteria needs 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 [Section 7](#) for some examples of scaling intent.

5. Performance Monitoring Parameters

This model augments 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 (Tunnel1). 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>Tunnel1</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 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>Tunnel1</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 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 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 the reported as the max of the one-way-delay at the VN-member level, where as 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@2022-07-11.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 YANG data model for performance
    monitoring parameters (telemetry data) for TE tunnels.

    Copyright (c) 2022 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 2022-07-11 {
  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 param 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 max 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 max 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 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 it 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
                to check 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
        to check 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 autonomics in order the controller to react
        to a certain set of monitored params.";
    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@2022-07-11.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
      "I-D.ietf-teas-actn-vn-yang: A YANG Data Model for VN
      Operation";
  }

  /* 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) 2022 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 2022-07-11 {

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

The YANG modules specified in this document defines a schema for data that is designed to be accessed via network management protocols such as NETCONF [[RFC6241](#)] or RESTCONF [[RFC8040](#)]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [[RFC6242](#)]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [[RFC8446](#)].

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). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g., edit-config) to these data nodes without proper protection can have a negative effect on network operations. These are the subtrees with the write operation that can be exploited to impact the network monitoring. An incorrect condition could cause frequent scaling operation to be executed causing harm to the network:

```

*/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

```

Further, following are the subtrees with the write operation that can be exploited by setting an incorrect grouping operation for the VN operation impacting the network monitoring:

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

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

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. These are the subtrees with the read operations that can be exploited to learn real-time (and sensitive) telemetry information about the TE tunnels and VN:

```
*/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
```

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

We 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-teas-actn-vn-yang] Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., and B. Y. Yoon, "A YANG Data Model for Virtual Network (VN) Operations", Work in Progress, Internet-Draft, draft-ietf-teas-actn-vn-yang-15, 11 July 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-vn-yang-15>>.

[I-D.ietf-teas-yang-te] Saad, T., Gandhi, R., Liu, X., Beeram, V. P., Bryskin, I., and O. G. D. Dios, "A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths and Interfaces", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-29, 7 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-29>>.

[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>>.

[RFC6242]

Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011, <<https://www.rfc-editor.org/info/rfc6242>>.

[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>>.

13.2. Informative References

[I-D.ietf-opsawg-yang-vpn-service-pm]

Wu, B., Wu, Q., Boucadair, M., Dios, O. G. D., and B. Wen, "A YANG Model for Network and VPN Service Performance Monitoring", Work in Progress, Internet-Draft, draft-ietf-opsawg-yang-vpn-service-pm-09, 17 May 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-yang-vpn-service-pm-09>>.

[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>>.

Appendix A. Out of Scope

This document exclusively focus 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 scope of this document.

Appendix B. Contributors

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

Satish Karunanithi
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore 560066
Karnataka
India
Email: satish.karunanithi@gmail.com

Authors' Addresses

Young Lee (editor)
Samsung Electronics

Email: younglee.tx@gmail.com

Dhruv Dhody (editor)
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore 560066
Karnataka
India

Email: dhruv.ietf@gmail.com

Ricard Vilalta
CTTC
Centre Tecnologic de Telecomunicacions de Catalunya (CTTC/CERCA)
Barcelona
Spain

Email: ricard.vilalta@cttc.es

Daniel King
Lancaster University

Email: d.king@lancaster.ac.uk

Daniele Ceccarelli
Ericsson
Torshamnsgatan, 48
Stockholm, Sweden

Email: daniele.ceccarelli@ericsson.com