TEAS Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 14, 2021

Y. Lee, Ed.
Samsung Electronics
D. Dhody, Ed.
S. Karunanithi
Huawei Technologies
R. Vilalta
CTTC
D. King
Lancaster University
D. Ceccarelli
Ericsson
July 13, 2020

# YANG models for VN/TE Performance Monitoring Telemetry and Scaling Intent Autonomics
## draft-ietf-teas-actn-pm-telemetry-autonomics-03

Abstract

This document provides YANG data models that describe performance monitoring telemetry and scaling intent mechanism for TE-tunnels and Virtual Networks (VN).

The models presented in this draft allow customers to subscribe to and monitor their key performance data of their interest on the level of TE-tunnel or VN. The models also provide customers with the ability to program autonomic scaling intent mechanism 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 January 14, 2021.

Copyright Notice

Copyright (c) 2020 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction . . . . . 2
  1.1. Terminology . . . . . 4
    1.1.1. Requirements Language . . . . . 4
  1.2. Tree diagram . . . . . 4
  1.3. Prefixes in Data Node Names . . . . . 5
2. Use-Cases . . . . . 5
3. Design of the Data Models . . . . . 7
  3.1. TE KPI Telemetry Model . . . . . 7
  3.2. VN KPI Telemetry Model . . . . . 8
4. Autonomic Scaling Intent Mechanism . . . . . 9
5. Notification . . . . . 11
  5.1. YANG Push Subscription Examples . . . . . 11
6. YANG Data Tree . . . . . 12
7. YANG Data Model . . . . . 15
  7.1. ietf-te-kpi-telemetry model . . . . . 15
  7.2. ietf-vn-kpi-telemetry model . . . . . 21
8. Security Considerations . . . . . 25
9. IANA Considerations . . . . . 25
10. Acknowledgements . . . . . 26
11. References . . . . . 26
  11.1. Normative References . . . . . 26
  11.2. Informative References . . . . . 28
Authors' Addresses . . . . . 29

## 1. Introduction

The YANG [RFC7950] model discussed in [I-D.ietf-teas-actn-vn-yang] is used to operate customer-driven Virtual Networks (VNs) during the VN instantiation, VN computation, and its life-cycle service management and operations. YANG model discussed 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, which is mapped to 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. Key performance indicator (KPI) is a term that describes critical performance data that may affect VN/TE-tunnel service. 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 generically applicable to any VN/TE-Tunnel service clients to provide an ability to program their customized performance monitoring subscription and publication data models and automatic scaling in/out intent data models. These models can be utilized by a client network controller to initiate these capability to a transport network controller communicating with the client controller via a NETCONF [RFC8341] or a RESTCONF [RFC8040] interface.

The term performance monitoring being used in this document is different from the term that has been used in transport networks for many years. 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., MDSC/PNC) based on the client's subscription. As the scope of performance monitoring in this document is telemetry data on the level of client's VN or TE- tunnel, the entity interfacing the client (e.g., MDSC) has to provide VN or TE-tunnel level information. This would require controller 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 new Network Management Datastore Architecture [RFC8342].

## 1.1. Terminology

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

Key Performance Data: This refers to a set of data the customer is interested in monitoring for their instantiated VNs or TE-tunnels. Key performance data and key performance indicators are inter-exchangeable in this draft.

Scaling: This refers to the network ability to re-shape its own resources. Scale out refers to improve network performance by increasing the allocated resources, while scale in refers to decrease the allocated resources, typically because the existing resources are unnecessary.

Scaling Intent: To declare scaling conditions, scaling intent is used. Specifically, scaling intent refers to the intent expressed by the client that allows the client to program/configure conditions of their key performance data either for 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 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.

### 1.1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 1.2. Tree diagram

A simplified graphical representation of the data model is used in Section 5 of this 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 |
|---|---|---|
| inet | ietf-inet-types | [RFC6991] |
| te | ietf-te | [I-D.ietf-teas-yang-te] |
| te-types | ietf-te-types | [RFC8776] |
| te-tel | ietf-te-kpi-telemetry | [RFCXXXX] |
| vn | ietf-vn | [I-D.ietf-teas-actn-vn-yang] |
| vn-tel | ietf-vn-kpi-telemetry | [RFCXXXX] |

Table 1: Prefixes and corresponding YANG modules

Note: The RFC Editor will replace XXXX with the number assigned to the RFC once this draft becomes an RFC.

Further, the following additional documents are refrenced 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

[I-D.xu-actn-perf-dynamic-service-control] describes use-cases relevant to this draft. It introduces the dynamic creation, modification and optimization of services based on the performance monitoring. Figure 1 shows a high-level workflows for dynamic service control based on traffic monitoring.

```
      +-----------------------------------------------+
      | Client   +-----------------------------+     |
      |          | Dynamic Service Control APP |     |
      |          +-----------------------------+     |
      +-----------------------------------------------+
      1.Traffic|  /|\4.Traffic            | /|\
      Monitor& |   | Monitor              |  | 8.Traffic
      Optimize |   | Result     5.Service |  | modify &
      Policy   |   |              modify& |  | optimize
              \|/  |        optimize Req.\|/ | result
      +-----------------------------------------------+
      | Orchestrator                                  |
      |    +-------------------------------+         |
      |    |Dynamic Service Control Agent  |         |
      |    +-------------------------------+         |
      |    +---------------+ +--------------------+  |
      |    | Flow Optimize | | vConnection Agent |   |
      |    +---------------+ +--------------------+  |
      +-----------------------------------------------+
      2. Path |   /|\3.Traffic            | /|\
      Monitor |    | Monitor              |  |7.Path
      Request |    | Result      6.Path   |  | modify &
              |    |             modify&  |  | optimize
             \|/   |        optimize Req.\|/ | result
      +-----------------------------------------------+
      | Network SDN Controller                        |
      |  +----------------------+ +-----------------+|
      |  | Network Provisioning | |Abstract Topology||
      |  +----------------------+ +-----------------+|
      |  +------------------+ +--------------------+ |
      |  |Network Monitoring| |Physical Topology DB| |
      |  +------------------+ +--------------------+ |
      +-----------------------------------------------+
```

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

Some of the key points from [I-D.xu-actn-perf-dynamic-service-control] are as follows:

o  Network traffic monitoring is important to facilitate automatic discovery of the imbalance of network traffic, and initiate the network optimization, thus helping the network operator or the virtual network service provider to use the network more efficiently and save the Capital Expense (CAPEX) and the Operating Expense (OPEX).

- o Customer services have various Service Level Agreement (SLA) requirements, such as service availability, latency, latency jitter, packet loss rate, Bit Error Rate (BER), etc. The transport network can satisfy service availability and BER requirements by providing different protection and restoration mechanisms. However, for other performance parameters, there are no such mechanisms. In order to provide high quality services according to 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.

- o Performance monitoring in a large scale network could generate a huge amount of performance information. Therefore, the appropriate way to deliver the information in the client and network interfaces should be carefully considered.

## 3. Design of the Data Models

The YANG models developed in this document describe two models:

(i) TE KPI 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 7.1 for details).

(ii) VN KPI 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 7.2 for details).

### 3.1. TE KPI Telemetry Model

This module describes performance telemetry for TE-tunnel model. The telemetry data is augmented to tunnel state. This module 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 5 for details)

The TE KPI Telemetry Model augments the TE-Tunnel Model to enhance TE performance monitoring capability. This monitoring capability will facilitate proactive re-optimization and reconfiguration of TEs based on the performance monitoring data collected via the TE KPI Telemetry YANG model.

```
                +------------+          +--------------+
                |  TE-Tunnel |          |    TE KPI    |
                |   Model    |<---------|  Telemetry   |
                +------------+ augments |     Model    |
                                        +--------------+
```

### 3.2. VN KPI Telemetry Model

This module describes performance telemetry for VN model. The telemetry data is augmented both at the VN Level as well as individual VN member level. This module 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 the controller to react to a certain set of variations in monitored parameters (See Section 4 for illustrations).

Moreover, this module also provides mechanism to define aggregated telemetry parameters as a grouping of underlying VN level telemetry parameters. 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 {delay_vn_member_1, delay_vn_member_2,.. delay_vn_member_N}. Thus, this telemetry abstraction mechanism allows the grouping of a certain common set of telemetry values under a grouping operation. This can be done at the VN-member level to suggest how the E2E telemetry be inferred from the per domain tunnel created and monitored by PNCs. One proposed example is the following:

```
  +-------------------------------------------------------------+
  |                      Client                                 |
  |                                                             |
  +-------------------------------------------------------------+
  1.Client sets the      |   /|\   2. Orchestrator pushes:
  grouping op, and       |    |
  subscribes to the      |    |    VN level telemetry for
  VN level telemetry for |    |    - VN Utilized-bw-percentage
  Delay and              |    |       (Minimum across VN Members)
  Utilized-bw-pecentage  |    |    - VN Delay (Maximum across VN
                        \|/   |     Members)
   +-------------------------------------------------------------+
   | Orchestrator                                                |
   |                                                             |
   +-------------------------------------------------------------+
```

The VN Telemetry Model augments the basic VN model to enhance VN monitoring capability. This monitoring capability will facilitate proactive re-optimization and reconfiguration of VNs based on the

performance monitoring data collected via the VN Telemetry YANG model.

```
+----------+          +--------------+
|    VN    | augments |      VN      |
|   Model  |<---------|   Telemetry  |
+----------+          |     Model    |
                      +--------------+
```

## 4. Autonomic Scaling Intent Mechanism

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:

- o scale-out-intent or scale-in-intent: whether to scale-out or scale-in.

- o 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.)

- o threshold-value: the threshold value for a certain performance-type that triggers scale-in or scale-out.

- o 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.

- o Threshold-time: the duration for which the criteria MUST hold true.

- o Cooldown-time: the duration after a scaling action has been triggered, for which there will be no further operation.

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

```
module: ietf-te-kpi-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?           string
    |  |     +--rw scale-in-operation-type?
    |  |             scaling-criteria-operation
    |  +--rw scale-out-intent
    |     +--rw threshold-time?      uint32
    |     +--rw cooldown-time?       uint32
    |     +--rw scaling-condition* [performance-type]
    |        +--rw performance-type            identityref
    |        +--rw threshold-value?            string
    |        +--rw scale-out-operation-type?
    |                scaling-criteria-operation
```

Let 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:

o  Set Threshold-time: x (sec) (duration for which the criteria must hold true)

o  Set Cooldown-time: y (sec) (the duration after a scaling action has been triggered, for which there will be no further operation)

o  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

o  performance type: Two-way-delay

o  threshold-value: z milli-seconds

List 2: Scaling Condition for Utilized bandwidth

o  performance type: Utilized bandwidth

o  threshold-value: w megabytes

## 5. 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:

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

### 5.1. YANG Push Subscription Examples

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

Below example 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>
                  <identifier/>
                  <state>
                    <te-telemetry xmlns="urn:ietf:params:xml:ns:yang:
                                         ietf-te-kpi-telemetry">
                        <one-way-delay/>
                     </te-telemetry>
                  </state>
                 </tunnel>
              </tunnels>
          </te>
       </filter>
       <period>500</period>
       <encoding>encode-xml</encoding>
    </establish-subscription>
 </netconf:rpc>
```

This example 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">
          <vn-state xmlns="urn:ietf:params:xml:ns:yang:ietf-vn">
             <vn>
                <vn-list>
                  <vn-id/>
                  <vn-name/>
                  <vn-telemetry xmlns="urn:ietf:params:xml:ns:yang:
                                       ietf-vn-kpi-telemetry">
                      <one-way-delay/>
                      <one-way-utilized-bandwidth/>
                  </vn-telemetry >
                </vn-list>
              </vn>
          </vn-state>
       </filter>
       <period>500</period>
    </establish-subscription>
 </netconf:rpc>
```

## 6. YANG Data Tree

```
module: ietf-te-kpi-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?           string
    |  |     +--rw scale-in-operation-type?
    |  |             scaling-criteria-operation
    |  +--rw scale-out-intent
    |     +--rw threshold-time?      uint32
    |     +--rw cooldown-time?       uint32
    |     +--rw scaling-condition* [performance-type]
    |        +--rw performance-type            identityref
```

```
       |        +--rw threshold-value?            string
       |        +--rw scale-out-operation-type?
       |                scaling-criteria-operation
       +--ro te-telemetry
          +--ro id?                           telemetry-id
          +--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
```

```
   module: ietf-vn-kpi-telemetry
     augment /vn:vn/vn:vn-list:
       +--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?           string
       |  |     +--rw scale-in-operation-type?
       |  |             scaling-criteria-operation
       |  +--rw scale-out-intent
       |     +--rw threshold-time?      uint32
       |     +--rw cooldown-time?       uint32
       |     +--rw scaling-condition* [performance-type]
       |        +--rw performance-type            identityref
       |        +--rw threshold-value?            string
       |        +--rw scale-out-operation-type?
       |                scaling-criteria-operation
```

```
       +--ro vn-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
          +--ro grouping-operation?            grouping-operation
     augment /vn:vn/vn:vn-list/vn:vn-member-list:
       +--ro vn-member-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
          +--ro te-grouped-params*
          |       -> /te:te/tunnels/tunnel/te-kpi:te-telemetry/id
          +--ro grouping-operation?            grouping-operation
```

## 7. YANG Data Model

### 7.1. ietf-te-kpi-telemetry model

The YANG code is as follows:

```
 <CODE BEGINS> file "ietf-te-kpi-telemetry@2020-07-13.yang"
module ietf-te-kpi-telemetry {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-te-kpi-telemetry";
  prefix te-tel;

  /* Import inet-types */

  import ietf-inet-types {
    prefix inet;
    reference
      "RFC 6991: Common YANG Data Types";
  }

  /* 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
      "I-D.ietf-teas-yang-te-types: Traffic Engineering Common
       YANG Types";
  }

  organization
    "IETF Traffic Engineering Architecture and Signaling (TEAS)
     Working Group";
  contact
    "WG Web:  <https://tools.ietf.org/wg/teas/>
     WG List: <mailto:teas@ietf.org>
     Editor:  Young Lee <leeyoung@huawei.com>
              Dhruv Dhody <dhruv.ietf@gmail.com>";
  description
    "This module describes YANG data model for performance
```

```
      monitoring telemetry for te tunnels.

      Copyright (c) 2020 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 Simplified 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.

      The key words 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL
      NOT', 'SHOULD', 'SHOULD NOT', 'RECOMMENDED', 'NOT RECOMMENDED',
      'MAY', and 'OPTIONAL' in this document are to be interpreted as
      described in BCP 14 (RFC 2119) (RFC 8174) when, and only when,
      they appear in all capitals, as shown here.";

   /* 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 2020-03-08 {
     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";
     reference
       "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";
    }

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

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

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

    identity utilized-bandwidth {
      base telemetry-param-type;
      description
        "To specify utilized bandwidth over the specified source
         and destination.";
      reference
        "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";
   }

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

   /* Typedef */

   typedef telemetry-id {
     type inet:uri;
     description
       "Identifier for telemetry data.  The precise
        structure of the telemetry-id will be up to the
        implementation. The identifier SHOULD be chosen
        such that the same telemetry data will always be
        identified through the same identifier, even if
        the data model is instantiated in separate
        datastores.";
   }

   typedef scaling-criteria-operation {
     type enumeration {
       enum AND {
         description
           "AND operation";
       }
       enum OR {
         description
           "OR operation";
       }
     }
     description
       "Operations to analize list of scaling criterias";
   }

   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";
    }
    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";
    }
  }

  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 string;
      description
        "Scaling threshold for the telemetry parameter 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 criterias
           to check if the scale in threshold condition has been met.
           Defaults to AND";
      }
    }
  }
```

```
   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 criterias
            to check if the scale out threshold condition has been met.
            Defauls to OR";
       }
     }
   }

   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";
       leaf id {
         type telemetry-id;
         description
           "ID of telemetry data used for easy reference";
```

```
      }
      uses te-types:performance-metrics-attributes;
    }
  }
}

<CODE ENDS>
```

## 7.2. ietf-vn-kpi-telemetry model

The YANG code is as follows:

```
<CODE BEGINS> file "ietf-vn-kpi-telemetry@2020-07-13.yang"
module ietf-vn-kpi-telemetry {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-vn-kpi-telemetry";
  prefix vn-kpi;

  /* 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 KPI */

  import ietf-te-kpi-telemetry {
    prefix te-kpi;
```

```
      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://tools.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 telemetry for Virtual Network (VN).

       Copyright (c) 2020 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 Simplified 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.

       The key words 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL
       NOT', 'SHOULD', 'SHOULD NOT', 'RECOMMENDED', 'NOT RECOMMENDED',
       'MAY', and 'OPTIONAL' in this document are to be interpreted as
       described in BCP 14 (RFC 2119) (RFC 8174) when, and only when,
       they appear in all capitals, as shown here.";

    /* 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 2020-07-13 {
      description
        "Initial revision.";
      reference
        "RFC XXXX: YANG models for VN/TE Performance Monitoring
```

```
         Telemetry and Scaling Intent Autonomics";
    }

    typedef grouping-operation {
      type enumeration {
        enum MINIMUM {
          description
            "Select the minimum param";
        }
        enum MAXIMUM {
          description
            "Select the maximum param";
        }
        enum MEAN {
          description
            "Select the MEAN of the params";
        }
        enum STD_DEV {
          description
            "Select the standard deviation of the monitored params";
        }
        enum AND {
          description
            "Select the AND of the params";
        }
        enum OR {
          description
            "Select the OR of the params";
        }
      }
      description
        "Operations to analize list of monitored params";
    }

    grouping vn-telemetry-param {
      description
        "augment of te-kpi:telemetry-param for VN specific params";
      leaf-list te-grouped-params {
        type leafref {
          path
            "/te:te/te:tunnels/te:tunnel/te-kpi:te-telemetry/te-kpi:id";
        }
        description
          "Allows the definition of a vn-telemetry param
           as a grouping of underlying TE params";
      }
      leaf grouping-operation {
        type grouping-operation;
```

```
        description
          "describes the operation to apply to
           te-grouped-params";
      }
    }

    augment "/vn:vn/vn:vn-list" {
      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-kpi:scaling-in-intent;
        }
        container scale-out-intent {
          description
            "VN scale-out";
          uses te-kpi:scaling-out-intent;
        }
      }
      container vn-telemetry {
        config false;
        description
          "VN telemetry params";
        uses te-types:performance-metrics-attributes;
        leaf grouping-operation {
          type grouping-operation;
          description
            "describes the operation to apply to the VN-members";
        }
      }
    }

    augment "/vn:vn/vn:vn-list/vn:vn-member-list" {
      description
        "Augmentation parameters for state TE vn member topologies.";
      container vn-member-telemetry {
        config false;
        description
          "VN member telemetry params";
        uses te-types:performance-metrics-attributes;
        uses vn-telemetry-param;
      }
    }
  }
```

<CODE ENDS>

## 8. Security Considerations

The YANG module 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 NETCONF access control model [RFC8341] provides the means to restrict access for particular NETCONF users to a preconfigured subset of all available NETCONF protocol operations and content. The NETCONF Protocol over Secure Shell (SSH) [RFC6242] describes a method for invoking and running NETCONF within a Secure Shell (SSH) session as an SSH subsystem. 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.

A number of configuration data nodes defined in this document are writable/deletable (i.e., "config true"). These data nodes may be considered sensitive or vulnerable in some network environments.

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 and data nodes and their sensitivity/vulnerability:

- /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:vn/vn:vn-list/vn-scaling-intent/scale-in-intent

- /vn:vn/vn:vn-list/vn-scaling-intent/scale-out-intent

## 9. IANA Considerations

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

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

--------------------------------------------------------------------
URI: urn:ietf:params:xml:ns:yang:ietf-vn-kpi-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 [RFC7950]:

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

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

## 10. Acknowledgements

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

## 11. References

### 11.1. Normative References

[I-D.ietf-teas-actn-vn-yang]
Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., and B. Yoon, "A Yang Data Model for VN Operation", draft-ietf-teas-actn-vn-yang-08 (work in progress), March 2020.

[I-D.ietf-teas-yang-te]
Saad, T., Gandhi, R., Liu, X., Beeram, V., and I. Bryskin, "A YANG Data Model for Traffic Engineering Tunnels and Interfaces", draft-ietf-teas-yang-te-23 (work in progress), March 2020.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

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

[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>.

[RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <https://www.rfc-editor.org/info/rfc6991>.

[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>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[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>.

[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>.

## 11.2. Informative References

[I-D.xu-actn-perf-dynamic-service-control]
Xu, Y., Zhang, G., Cheng, W., and z. zhenghaomian@huawei.com, "Use Cases and Requirements of Dynamic Service Control based on Performance Monitoring in ACTN Architecture", draft-xu-actn-perf-dynamic-service-control-03 (work in progress), April 2015.

[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>.

[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>.

## Authors' Addresses

Young Lee (editor)
Samsung Electronics

Email: younglee.tx@gmail.com

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

Email: dhruv.ietf@gmail.com

Satish Karunanithi
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

Email: satish.karunanithi@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