

Workgroup: TEAS Working Group

Internet-Draft:

draft-ietf-teas-te-topology-profiles-01

Published: 19 April 2024

Intended Status: Informational

Expires: 21 October 2024

Authors: I. Busi X. Liu I. Bryskin T. Saad
 Huawei Alef Edge Individual Cisco Systems Inc
 O. Gonzalez de Dios
 Telefonica

Profiles for Traffic Engineering (TE) Topology Data Model and Applicability to non-TE Use Cases

Abstract

This document describes how profiles of the Traffic Engineering (TE) Topology Model, defined in RFC8795, can be used to address applications beyond "Traffic Engineering".

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 21 October 2024.

Copyright Notice

Copyright (c) 2024 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](#)
- [2. Examples of non-TE scenarios](#)
 - [2.1. UNI Topology Discovery](#)
 - [2.2. Administrative and Operational status management](#)
 - [2.3. Overlay and Underlay non-TE Topologies](#)
 - [2.4. Nodes with switching limitations](#)
- [3. Technology-specific augmentations](#)
 - [3.1. Multi-inheritance](#)
 - [3.2. Example \(Link augmentation\)](#)
- [4. Open Issues](#)
 - [4.1. Supporting node/link versus overlay/underlay](#)
 - [4.2. Implemented profiles](#)
 - [4.3. Applicability to non-TE use cases](#)
 - [4.3.1. Update UNI topology discovery use case](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [Acknowledgments](#)
- [References](#)
 - [Normative References](#)
 - [Informative References](#)
- [Contributors](#)
- [Authors' Addresses](#)

1. Introduction

There are many network scenarios being discussed in various IETF Working Groups (WGs) that are not classified as "Traffic Engineering" but can be addressed by a sub-set (profile) of the Traffic Engineering (TE) Topology YANG data model, defined in [\[RFC8795\]](#).

Traffic Engineering (TE) is defined in [\[I-D.ietf-teas-rfc3272bis\]](#) as aspects of Internet network engineering that deal with the issues of performance evaluation and performance optimization of operational IP networks. TE encompasses the application of technology and scientific principles to the measurement, characterization, modeling, and control of Internet traffic.

The TE Topology Model is augmenting the Network Topology Model defined in [\[RFC8345\]](#) with generic and technology-agnostic features that some are strictly applicable to TE networks, while others applicable to both TE and non-TE networks.

Examples of such features that are applicable to both TE and non-TE networks are: inter-domain link discovery (plug-id), geo-localization, and admin/operational status.

It is also worth noting that the TE Topology Model is quite an extensive and comprehensive model in which most features are optional. Therefore, even though the full model appears to be complex, at the first glance, a sub-set of the model (profile) can be used to address specific scenarios, e.g. suitable also to non-TE use cases.

The implementation of such TE Topology profiles can simplify and expedite adoption of the full TE topology YANG data model, and allow for its reuse even for non-TE use case. The key question being whether all or some of the attributes defined in the TE Topology Model are needed to address a given network scenario.

[Section 2](#) provides examples where profiles of the TE Topology Model can be used to address some generic use cases applicable to both TE and non-TE technologies.

2. Examples of non-TE scenarios

2.1. UNI Topology Discovery

UNI Topology Discovery is independent from whether the network is TE or non-TE.

The TE Topology Model supports inter-domain link discovery (including but not being limited to UNI link discovery) using the plug-id attribute. This solution is quite generic and does not require the network to be a TE network.

The following profile of the TE Topology model can be used for the UNI Topology Discovery:

```
module: ietf-te-topology
  augment /nw:networks/nw:network/nw:network-types:
    +--rw te-topology!
  augment /nw:networks/nw:network/nw:node/nt:termination-point:
    +--rw te-tp-id?   te-types:te-tp-id
    +--rw te!
      +--rw admin-status?
        |   te-types:te-admin-status
      +--rw inter-domain-plug-id?      binary
      +--ro oper-status?               te-types:te-oper-status
```

Figure 1: UNI Topology

The profile data model shown in [Figure 1](#) can be used to discover TE and non TE UNIs as well as to discover UNIs for TE or non TE networks.

Such a UNI TE Topology profile model can also be used with technology-specific UNI augmentations, as described in section 3.

For example, in [[I-D.ietf-ccamp-eth-client-te-topo-yang](#)], the eth-svc container is defined to represent the capabilities of the Termination Point (TP) to be configured as an Ethernet client UNI, together with the Ethernet classification and VLAN operations supported by that TP.

The [[I-D.ietf-ccamp-otn-topo-yang](#)] provides another example, where:

- *the client-svc container is defined to represent the capabilities of the TP to be configured as an transparent client UNI (e.g., STM-N, Fiber Channel or transparent Ethernet);

- *the OTN technology-specific Link Termination Point (LTP) augmentations are defined to represent the capabilities of the TP to be configured as an OTN UNI, together with the information about OTN label and bandwidth availability at the OTN UNI.

For example, the UNI TE Topology profile can be used to model features defined in [[I-D.ogondio-opsawg-uni-topology](#)]:

- *The inter-domain-plug-id attribute would provide the same information as the attachment-id attribute defined in [[I-D.ogondio-opsawg-uni-topology](#)];

- *The admin-status and oper-status that exists in this TE topology profile can provide the same information as the admin-status and oper-status attributes defined in [[I-D.ogondio-opsawg-uni-topology](#)].

Following the same approach in [[I-D.ietf-ccamp-eth-client-te-topo-yang](#)] and [[I-D.ietf-ccamp-otn-topo-yang](#)], the type and encapsulation-type attributes can be defined by technology-specific UNI augmentations to represent the capability of a TP to be configured as a L2VPN/L3VPN UNI Service Attachment Point (SAP).

The advantages of using a TE Topology profile would be having common solutions for:

- *discovering UNIs as well as inter-domain NNI links, which is applicable to any technology (TE or non TE) used at the UNI or within the network;

*modelling non TE UNIs such as Ethernet, and TE UNIs such as OTN, as well as UNIs which can be configured as TE or non-TE (e.g., being configured as either Ethernet or OTN UNI).

2.2. Administrative and Operational status management

The TE Topology Model supports the management of administrative and operational state, including also the possibility to associate some administrative names, for nodes, termination points and links. This solution is generic and also does not require the network to be a TE network.

The following profile of the TE Topology Model can be used for administrative and operational state management:

```
module: ietf-te-topology
  augment /nw:networks/nw:network/nw:network-types:
    +--rw te-topology!
  augment /nw:networks/nw:network:
    +--rw te-topology-identifier
    | +--rw provider-id?   te-global-id
    | +--rw client-id?    te-global-id
    | +--rw topology-id?  te-topology-id
    +--rw te!
    +--rw name?           string
  augment /nw:networks/nw:network/nw:node:
    +--rw te-node-id?    te-types:te-node-id
    +--rw te!
    +--rw te-node-attributes
    | +--rw admin-status? te-types:te-admin-status
    | +--rw name?        string
    +--ro oper-status?   te-types:te-oper-status
  augment /nw:networks/nw:network/nt:link:
    +--rw te!
    +--rw te-link-attributes
    | +--rw name?        string
    | +--rw admin-status? te-types:te-admin-status
    +--ro oper-status?   te-types:te-oper-status
  augment /nw:networks/nw:network/nw:node/nt:termination-point:
    +--rw te-tp-id?    te-types:te-tp-id
    +--rw te!
    +--rw admin-status? te-types:te-admin-status
    +--rw name?        string
    +--ro oper-status? te-types:te-oper-status
```

Figure 2: Generic Topology with admin and operational state

The TE topology data model profile shown in [Figure 2](#) is applicable to any technology (TE or non-TE) that requires management of the administrative and operational state and administrative names for nodes, termination points and links.

2.3. Overlay and Underlay non-TE Topologies

The TE Topology model supports the management of overlay/underlay relationship for nodes and links, as described in section 5.8 of [\[RFC8795\]](#). This solution is generic and does not require the network to be a TE network.

The following TE topology data model profile can be used to manage overlay/underlay network data:

```
module: ietf-te-topology
  augment /nw:networks/nw:network/nw:network-types:
    +--rw te-topology!
  augment /nw:networks/nw:network/nw:node:
    +--rw te-node-id?   te-types:te-node-id
    +--rw te!
      +--rw te-node-attributes
        +--rw underlay-topology {te-topology-hierarchy}?
          +--rw network-ref? -> /nw:networks/network/network-id
  augment /nw:networks/nw:network/nt:link:
    +--rw te!
      +--rw te-link-attributes
        +--rw underlay {te-topology-hierarchy}?
          +--rw enabled?                boolean
          +--rw primary-path
            +--rw network-ref?
              | -> /nw:networks/network/network-id
            +--rw path-element* [path-element-id]
              +--rw path-element-id      uint32
              +--rw (type)?
                +--:(numbered-link-hop)
                  | +--rw numbered-link-hop
                  |   +--rw link-tp-id    te-tp-id
                  |   +--rw hop-type?     te-hop-type
                  |   +--rw direction?    te-link-direction
                +--:(unnumbered-link-hop)
                  +--rw unnumbered-link-hop
                    +--rw link-tp-id     te-tp-id
                    +--rw node-id        te-node-id
                    +--rw hop-type?      te-hop-type
                    +--rw direction?     te-link-direction
```

Figure 3: Generic Topology with overlay/underlay information

This profile is applicable to any technology (TE or non-TE) when it is needed to manage the overlay/underlay information. It is also allows a TE underlay network to support a non-TE overlay network and, vice versa, a non-TE underlay network to support a TE overlay network.

2.4. Nodes with switching limitations

A node can have some switching limitations where connectivity is not possible between all its TP pairs, for example when:

- *the node represents a physical device with switching limitations;

- *the node represents an abstraction of a network topology.

This scenario is generic and applies to both TE and non-TE technologies.

A connectivity TE Topology profile data model supports the management of the node connectivity matrix to represent feasible connections between termination points across the nodes. This solution is generic and does not necessarily require a TE enabled network.

The following profile of the TE Topology model can be used for nodes with connectivity constraints:

```
module: ietf-te-topology
  augment /nw:networks/nw:network/nw:network-types:
    +--rw te-topology!
  augment /nw:networks/nw:network/nw:node:
    +--rw te-node-id?   te-types:te-node-id
    +--rw te!
      +--rw te-node-attributes
        +--rw connectivity-matrices
          +--rw number-of-entries?      uint16
          +--rw is-allowed?             boolean
          +--rw connectivity-matrix* [id]
            +--rw id                    uint32
            +--rw from
              | +--rw tp-ref?           leafref
            +--rw to
              | +--rw tp-ref?           leafref
            +--rw is-allowed?           boolean
```

Figure 4: Generic Topology with connectivity constraints

The TE topology data model profile shown in [Figure 4](#) is applicable to any technology (TE or non-TE) networks that requires managing

nodes with certain connectivity constraints. When used with TE technologies, additional TE attributes, as defined in [RFC8795], can also be provided.

3. Technology-specific augmentations

There are two main options to define technology-specific Topology Models which can use the attributes defined in the TE Topology Model [RFC8795].

Both options are applicable to any possible profile of the TE Topology Model, such as those defined in Section 2.

The first option is to define a technology-specific TE Topology Model which augments the TE Topology Model, as shown in Figure 5:




Figure 5: Augmenting the TE Topology Model

This approach is more suitable for cases when the technology-specific TE topology model provides augmentations to the TE Topology constructs, such as bandwidth information (e.g., link bandwidth), tunnel termination points (TTPs) or connectivity matrices. It also allows providing augmentations to the Network Topology constructs, such as nodes, links, and termination points (TPs).

This is the approach currently used in [I-D.ietf-ccamp-eth-client-te-topo-yang] and [I-D.ietf-ccamp-otn-topo-yang].

It is worth noting that a profile of the technology-specific TE Topology model not using any TE topology attribute or constructs can be used to address any use case that do not require these attributes. In this case, only the te-topology presence container of the TE Topology Model needs to be implemented.

The second option is to define a technology-specific Network Topology Model which augments the Network Topology Model and to rely on the multiple inheritance capability, which is implicit in the network- types definition of [RFC8345], to allow using also the generic attributes defined in the TE Topology model:




Figure 6: Augmenting the Network Topology Model with multi-inheritance

This approach is more suitable in cases where the technology-specific Network Topology Model provides augmentation only to the constructs defined in the Network Topology Model, such as nodes, links, and termination points (TPs). Therefore, with this approach, only the generic attributes defined in the TE Topology Model could be used.

It is also worth noting that in this case, technology-specific augmentations for the bandwidth information could not be defined.

In principle, it would be also possible to define both a technology specific TE Topology Model which augments the TE Topology Model, and a technology-specific Network Topology Model which augments the Network Topology Model and to rely on the multiple inheritance capability, as shown in [Figure 7](#):

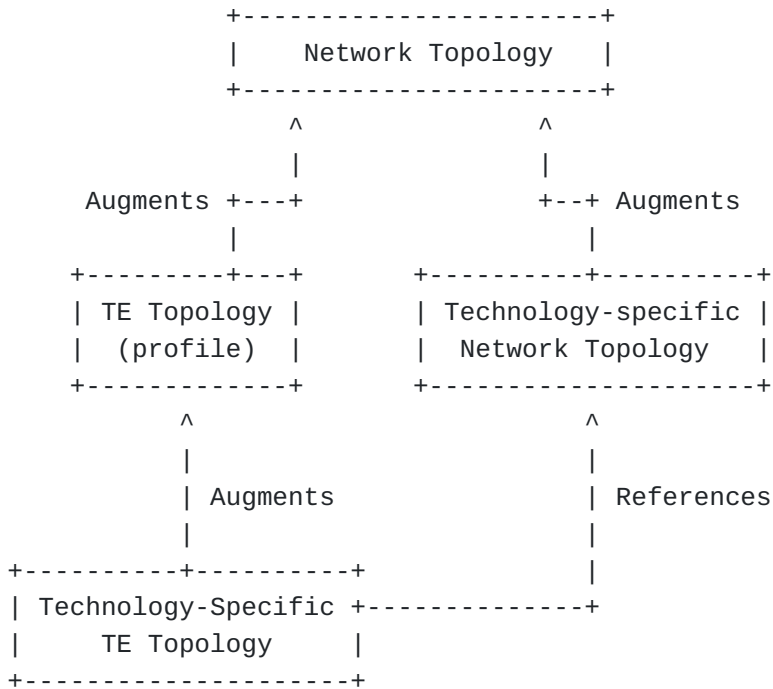


Figure 7: Augmenting both the Network and TE Topology Models

This option does not provide any technical advantage with respect to the first option, shown in [Figure 5](#), but could be useful to add augmentations to the TE Topology constructs and to re-use an already existing technology-specific Network Topology Model.

It is worth noting that the technology-specific TE Topology model can reference constructs defined by the technology-specific Network Topology model but it could not augment constructs defined by the technology-specific Network Topology model.

3.1. Multi-inheritance

As described in section 4.1 of [\[RFC8345\]](#), the network types should be defined using presence containers to allow the representation of network subtypes.

The hierarchy of network subtypes can be single hierarchy, as shown in [Figure 5](#). In this case, each presence container contains at most one child presence container, as shows in the JSON code below:

```

{
  "ietf-network:ietf-network": {
    "ietf-te-topology:te-topology": {
      "example-te-topology": {}
    }
  }
}

```

The hierarchy of network subtypes can also be multi-hierarchy, as shown in [Figure 6](#) and [Figure 7](#). In this case, one presence container can contain more than one child presence containers, as show in the JSON codes below:

```
{
  "ietf-network:ietf-network": {
    "ietf-te-topology:te-topology": {}
    "example-network-topology": {}
  }
}

{
  "ietf-network:ietf-network": {
    "ietf-te-topology:te-topology": {
      "example-te-topology": {}
    }
    "example-network-topology": {}
  }
}
```

Other examples of multi-hierarchy topologies are described in [\[I-D.ietf-teas-yang-sr-te-topo\]](#).

3.2. Example (Link augmentation)

This section provides an example on how technology-specific attributes can be added to the Link construct:

```

+--rw link* [link-id]
  +--rw link-id          link-id
  +--rw source
  | +--rw source-node?  -> ../../../../nw:node/node-id
  | +--rw source-tp?   leafref
+--rw destination
  | +--rw dest-node?   -> ../../../../nw:node/node-id
  | +--rw dest-tp?    leafref
+--rw supporting-link* [network-ref link-ref]
  | +--rw network-ref
  | |           -> ../../../../nw:supporting-network/network-ref
  | +--rw link-ref    leafref
+--rw example-link-attributes
  | <...>
+--rw te!
  +--rw te-link-attributes
    +--rw name?                string
    +--rw example-te-link-attributes
    | <...>
    +--rw max-link-bandwidth
      +--rw te-bandwidth
        +--rw (technology)?
          +--:(generic)
          | +--rw generic?   te-bandwidth
          +--:(example)
          | +--rw example?   example-bandwidth

```

Figure 8: Augmenting the Link with technology-specific attributes

The technology-specific attributes within the example-link-attributes container can be defined either in the technology-specific TE Topology Model (Option 1) or in the technology-specific Network Topology Model (Option 2 or Option 3). These attributes can only be non-TE and do not require the implementation of the te container.

The technology-specific attributes within the example-te-link-attributes container as well as the example max-link-bandwidth can only be defined in the technology-specific TE Topology Model (Option 1 or Option 3). These attributes can be TE or non-TE and require the implementation of the te container.

4. Open Issues

4.1. Supporting node/link versus overlay/underlay

Some more explanation of the difference between supporting-node/supporting-link and overlay/underlay has been requested.

Note: that this issue is also tracked in github as issue #167.

4.2. Implemented profiles

When a server implements a profile of the TE topology model, there is no standardized mechanism for the server to report to the client the subset of the model being implemented.

This might not be an issue in case the TE topology profile is read by the the client because the server reports in the operational datastore only the leaves which have been implemented, as described in section 5.3 of [[RFC8342](#)].

More investigation is instead required in case the TE topology profile is configured by the client, to avoid that the client tries to write an attribute not used in the TE Topology profile implemented by the server.

It is also worth noting that the supported profile may also depend on other attributes (for example the network type), so the YANG deviation mechanism is not applicable to this scenario.

Note: that this issue is also tracked in github as issue #161.

4.3. Applicability to non-TE use cases

Extending the applicability of RFC8795 to non-TE use cases is important. However, it is desirable to avoid any debate about whether these use cases in section 2 are or are not TE.

Note: that this issue is also tracked in github as issue #276.

4.3.1. Update UNI topology discovery use case

[Section 2.1](#) points to individual drafts and does reflect the progress made since then. For example, the UNI draft was replaced by other drafts that then led to the SAP model [[RFC9408](#)], which covers both UNI and NNI.

Note: that this issue is also tracked in github as issue #275.

5. Security Considerations

This document provides only information about how the TE Topology Model, as defined in [[RFC8795](#)], can be profiled to address some scenarios which are not considered as TE.

As such, this document does not introduce any additional security considerations besides those already defined in [[RFC8795](#)].

6. IANA Considerations

This document requires no IANA actions.

Acknowledgments

The authors would like to thank Vishnu Pavan Beeram, Daniele Ceccarelli, Jonas Ahlberg and Scott Mansfield for providing useful suggestions for this draft.

This document was prepared using kramdown.

Initial versions of this document were prepared using 2-Word-v2.0.template.dot.

References

Normative References

- [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>>.
- [RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", RFC 8345, DOI 10.17487/RFC8345, March 2018, <<https://www.rfc-editor.org/info/rfc8345>>.
- [RFC8795] Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and O. Gonzalez de Dios, "YANG Data Model for Traffic Engineering (TE) Topologies", RFC 8795, DOI 10.17487/RFC8795, August 2020, <<https://www.rfc-editor.org/info/rfc8795>>.

Informative References

[I-D.ietf-ccamp-eth-client-te-topo-yang]

Yu, C., Zheng, H., Guo, A., Busi, I., Xu, Y., Zhao, Y., and X. Liu, "A YANG Data Model for Ethernet TE Topology", Work in Progress, Internet-Draft, draft-ietf-ccamp-eth-client-te-topo-yang-06, 12 April 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-eth-client-te-topo-yang-06>>.

- [I-D.ietf-ccamp-otn-topo-yang] Zheng, H., Busi, I., Liu, X., Belotti, S., and O. G. de Dios, "A YANG Data Model for Optical Transport Network Topology", Work in Progress, Internet-Draft, draft-ietf-ccamp-otn-topo-yang-18, 19

April 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-otn-topo-yang-18>>.

[I-D.ietf-teas-rfc3272bis]

Farrel, A., "Overview and Principles of Internet Traffic Engineering", Work in Progress, Internet-Draft, draft-ietf-teas-rfc3272bis-27, 12 August 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-rfc3272bis-27>>.

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

Liu, X., Bryskin, I., Beeram, V. P., Saad, T., Shah, H. C., and S. Litkowski, "YANG Data Model for SR and SR TE Topologies on MPLS Data Plane", Work in Progress, Internet-Draft, draft-ietf-teas-yang-sr-te-topo-18, 21 October 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-sr-te-topo-18>>.

[I-D.ogondio-opsawg-uni-topology]

de Dios, O. G., Barguil, S., Wu, Q., and M. Boucadair, "A YANG Model for User-Network Interface (UNI) Topologies", Work in Progress, Internet-Draft, draft-ogondio-opsawg-uni-topology-01, 2 April 2020, <<https://datatracker.ietf.org/doc/html/draft-ogondio-opsawg-uni-topology-01>>.

[RFC9408]

Boucadair, M., Ed., Gonzalez de Dios, O., Barguil, S., Wu, Q., and V. Lopez, "A YANG Network Data Model for Service Attachment Points (SAPs)", RFC 9408, DOI 10.17487/RFC9408, June 2023, <<https://www.rfc-editor.org/info/rfc9408>>.

Contributors

Aihua Guo
Futurewei Inc.

Email: aihuaguo.ietf@gmail.com

Haomian Zheng
Huawei

Email: zhenghaomian@huawei.com

Sergio Belotti
Nokia

Email: sergio.belotti@nokia.com

Authors' Addresses

Italo Busi

Huawei

Email: italo.busi@huawei.com

Xufeng Liu

Alef Edge

Email: xufeng.liu.ietf@gmail.com

Igor Bryskin

Individual

Email: i_bryskin@yahoo.com

Tarek Saad

Cisco Systems Inc

Email: tсад.net@gmail.com

Oscar Gonzalez de Dios

Telefonica

Email: oscar.gonzalezdedios@telefonica.com