

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: April 24, 2014

J. Medved
Cisco
N. Bahadur
Juniper Networks
A. Clemm
Cisco
H. Ananthakrishnan
Juniper Networks
October 21, 2013

An Information Model for Network Topologies
draft-medved-i2rs-topology-im-01.txt

Abstract

This document defines the information model for network topologies.

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 <http://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 April 24, 2014.

Copyright Notice

Copyright (c) 2013 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 (<http://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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	2
2.	Definitions and Acronyms	4
3.	Network Topology Model Overview	4
4.	Network Topology Information Model	5
4.1.	Base Model: the Network-Topology Component	6
4.2.	Layer 3 Unicast Topology (IGP) Extensions	9
4.2.1.	The L3-Unicast-Topology Component	9
4.2.2.	The OSPF-Topology Component	11
4.2.3.	The IS-IS-Topology Component	12
4.2.4.	The TED (Traffic Engineering Data) Component	13
5.	Security Considerations	14
6.	Contributors	15
7.	Acknowledgements	15
8.	References	15
8.1.	Normative References	15
8.2.	Informative References	15

[1.](#) Introduction

This document introduces an information model for network topologies. The model allows applications to have a holistic view of an entire network. [[I-D.amante-i2rs-topology-use-cases](#)] describes an entity - the Topology Manager - that would create a cohesive, abstracted model of the network and expose it to applications via northbound API.

The information model can be related to a corresponding data model, for example, a data model defined in YANG [[RFC6020](#)], defining the actual data that is exchanged across specific interfaces. On the relationship between information and data models, please refer to [[RFC3444](#)].

In order to capture information that is specific to different network topology types, this document defines an abstract (basic) topology model that can be extended and adapted. As a result, the information model is generic in nature and can be applied to many network topologies. Applications can operate on any topology at a generic level where specifics of particular topology types are not required, and at a topology-specific level when those specifics come into play. Specific topology types that are covered in this document include Layer 3 Unicast IGP, IS-IS, and OSPF. We also define the information model for traffic engineering (TE) data. Adaptations and extensions to other types of topologies (such as Layer 2 topology or OpenFlow topology) are possible, using similar model patterns to the ones that are illustrated.

This revision of the document focuses on the "live" topology information model ([[I-D.amante-i2rs-topology-use-cases](#)], Section 3.2). The "inventory" and "statistics collection" information models will be addressed separately.

The information model contains several components:

Network-Topology contains a generic network topology model. It defines a network topology at its most general level of abstraction. It models aspects such as nodes and edges that constitute a topology graph, as well as termination points contained in the nodes that actually terminate edges of the graph. A network can contain multiple topologies, for example topologies at different network layers or overlay topologies. The model therefore allows to show relationships between topologies, as well as dependencies between nodes and termination points across topologies.

L3-Unicast-IGP-Topology applies the general network topology model to Layer 3 Unicast IGP topologies. It extends the general topology with information specific to Layer 3 Unicast IGP. In doing so, it also illustrates the extension patterns associated with extending respectively extending the general topology model to meet the needs of a specific topology.

OSPF-Topology Module "ospf-topology" defines a topology model for OSPF, building on and extending the Layer 3 Unicast IGP topology model. It serves as an example of how the general topology model can be refined across multiple levels.

IS-IS-Topology defines a topology model for IS-IS, again building on and extending the Layer 3 Unicast IGP topology model.

TED defines information kept in the Traffic Engineering Database (TED) that is leveraged by IS-IS and OSPF topologies.

2. Definitions and Acronyms

Data model: An abstract model of a conceptual domain that is intended for implementors and contains enough specifics to result in interoperable implementations and data representations

Datastore: A conceptual store of instantiated management information, with individual data items represented by data nodes which are arranged in hierarchical manner.

IGP: Interior Gateway Protocol

Information Model: An abstract model of a conceptual domain, independent of a specific implementations or data representation

IS-IS: Intermediate System to Intermediate System protocol

LSP: Label Switched Path

OSPF: Open Shortest Path First, a link state routing protocol

RBNF: Routing Backus-Naur Form

URI: Uniform Resource Identifier

SRLG: Shared Risk Link Group

TED: Traffic Engineering Database

3. Network Topology Model Overview

This section provides an overview of the network topology model. We start with the structure of the foundational model that represents a generic topology. Subsequently, an overview of selected specific topologies is given - Layer 3 Unicast IGP, OSPF, and IS-IS, respectively. Throughout the document, selected design choices are explained and the pattern that should be applied to extend the model to new types of topologies is presented.

The network topology model is defined by the following components, whose relationship is roughly depicted in the figure below.

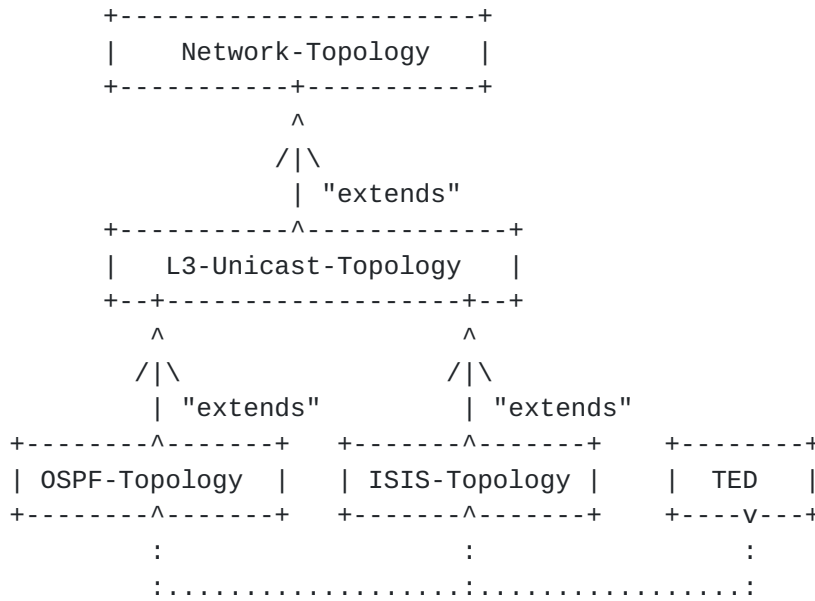


Figure 1: Overall model structure


The Network-Topology component defines the basic network topology model. The L3-Unicast-Topology module extends this model with additional definitions needed to represent Layer 3 Unicast IGP topologies. This component in turn is extended by OSPF-Topology and ISIS-Topology components providing additional definitions for OSPF and IS-IS topologies, respectively. The TED component, used by both OSPF-Topology and ISIS-Topology components, contains a set of auxiliary definitions related to traffic engineering.

4. Network Topology Information Model

This section specifies the network topology information model in Routing Backus-Naur Form (RBNF, [[RFC5511](#)]). It also provides diagrams of the main entities that the information model is comprised of.

4.1. Base Model: the Network-Topology Component

The following diagram contains an informal graphical depiction of the main elements of the information model:



Roughly speaking, the basic information model works as follows: A topology contains nodes and links. Each node in turn contains termination points. A link connects two nodes (a source and a destination), terminating on each at a termination point. Nodes can map onto and be supported by other nodes, while links can map onto and be supported by other links. Topologies can map onto other, underlay topologies.

The information model for the Network-Topology component is more formally shown in the following diagram.

```

<network-topology> ::= (<topology>...)

<topology> ::= <TOPOLOGY_IDENTIFIER>
              (<node>
               (<link>...)
               [<topology-type>]
               [<underlay-topologies>]
               [<topology-extension>])

<topology-type> ::= (<IGP> [<igp-topology-type>]) |
                   (<BGP> [<bgp-topology-type>])
<igp-topology-type> ::= <OSPF> | <ISIS>

<bgp-topology-type> ::= <>
    
```



```
<underlay-topologies> ::= (<TOPOLOGY_IDENTIFIER>...)  
  
<topology-extension> ::= <igp-topology-extension> |  
                        <bgp-topology-extension> |  
                        ...  
  
<node> ::= <NODE_IDENTIFIER>  
          (<termination-point>...)  
          [<supporting-nodes>]  
          [<node-extension>]  
  
<termination-point> ::= <TERMINATION_POINT_IDENTIFIER>  
                        [<supporting-termination-points>]  
                        [<termination-point-extension>]  
  
<supporting-termination-points> ::=  
    (<TERMINATION_POINT_IDENTIFIER>...)  
  
<termination-point-extension> ::=  
    <igp-termination-point-extension> |  
    <bgp-termination-point-extension>  
    ...  
  
<supporting-nodes> ::= (<NODE_IDENTIFIER>...)  
  
<node-extension> ::= <igp-node-extension> |  
                    <bgp-node-extension> |  
                    ...  
  
<link> ::= <LINK_IDENTIFIER>  
          <source>  
          <destination>  
          [<supporting-links>]  
          [<link-extension> ]  
  
<source> ::= <termination-point-reference>  
<destination> ::= <termination-point-reference>  
  
<termination-point-reference> ::= <NODE_IDENTIFIER>  
                                  <TERMINATION_POINT_IDENTIFIER>  
  
<supporting-links> ::= (<LINK_IDENTIFIER>...)  
  
<link-extension> ::= <igp-link-extension> |  
                    <bgp-link-extension> |  
                    ...
```


The elements of the Network-Topology information model are as follows:

- o A network can contain multiple topologies. Each topology is captured in its own list element, distinguished via a topology-id.
- o A topology has a certain type, such as OSPF or IS-IS. A topology can even have multiple types simultaneously. The type, or types, are captured in the list of "topology-type" components.
- o A topology can in turn be part of a hierarchy of topologies, building on top of other topologies. Any such topologies are captured in list "underlay-topology".
- o Furthermore, a topology contains nodes and links, each captured in their own list.
- o A node has a node-id. This distinguishes the node from other nodes in the list. In addition, a node has a list of termination points, used to terminate links. An examples of a termination point might be a physical or logical port or, more generally, an interface. Also, a node can in turn map onto other nodes in an underlay topology. This is captured in list "supporting-node".
- o A link is identified by a link-id, uniquely identifying the link within the topology. Links are point-to-point and unidirectional. Accordingly, a link contains a source and a destination. Both source and destination reference a corresponding node, as well as a termination point on that node. Analogous to a node, a link can in turn map onto other links an underlay topology. This is captured in list "supporting-link".
- o The topology, node, link and termination-point elements can be extended with topology-specific components (topology-extensions, node-extension, link-extension and termination-point-extension respectively). This document defines extensions for the L3 unicast topology.

The topology model includes links that are point-to-point and unidirectional. It does not directly support multipoint and bidirectional links. While this may appear as a limitation, it does keep the model simple, generic, and allows it to very easily be subjected applications that make use of graph algorithms. Bi-directional connections can be represented through pairs of unidirectional links. By introducing hierarchies of nodes, with nodes at one level mapping onto a set of other nodes at another level, and the introducing new links for nodes at that level, topologies with connections representing non-point-to-point communication patterns can be represented.

To minimize assumptions of what a topology might actually represent, mappings between topologies, nodes, links, and termination points are kept strictly generic. For example, no assumptions are made whether a termination point actually refers to an interface, or whether a node refers to a specific "system" or device; the model at this generic level makes no provisions for that. Any greater specifics about mappings between upper and lower layers can be captured in extending modules.


Links are terminated by a single termination point, not sets of termination points. Connections involving multihoming or link aggregation schemes need to be modeled using multiple point-to-point links, then define a link at a higher layer that is supported by those individual links.

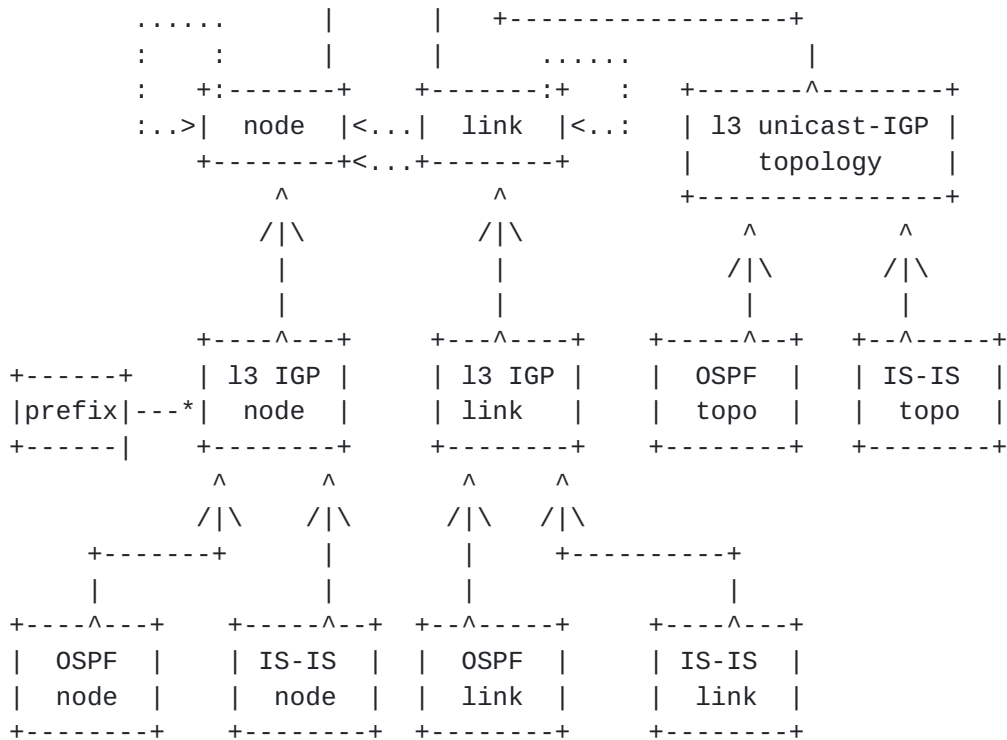
In a hierarchy of topologies, there are nodes mapping to nodes, links mapping to links, and termination points mapping to termination points. Some of this information is redundant. Specifically, with the link-to-links mapping known, and the termination points of each link known, maintaining separate termination point mapping information is not needed but can be derived via transitive closure.

4.2. Layer 3 Unicast Topology (IGP) Extensions

4.2.1. The L3-Unicast-Topology Component

In order to represent a general Layer 3 Unicast IGP topology, the basic network topology model needs to be extended. The corresponding extensions are introduced in a component, whose structure is informally depicted in the following diagram.





Roughly speaking, layer 3 IGP topology refines the generic topology, and layer 3 IGP nodes, IGP links, and IGP termination points (not depicted) refine the generic nodes, links, and termination points. In addition, layer 3 IGP nodes can contain prefixes. The pattern recurses with OSPF and IS-IS topologies, which are in turn derived from the corresponding layer 3 IGP entities.

A more formal depiction in RBNF format follows below:

```

<igp-topology-extension> ::= <TOPOLOGY_NAME>
                             [(<FLAGS>...)]
                             [(<ospf-topology> | <isis-topology>)]

<igp-node-extension> ::= <NODE_NAME>
                          (<router-id>...)
                          [(<prefix>...)]
                          [(<FLAGS>...)]
                          [(<isis-node> | <ospf-node>)]

<router-id> ::= <IPV4><IPV4_ADDRESS> | <NUMBER>

<ip-address> ::= (<IPV4><IPV4_ADDRESS>) | (<IPV6><IPV6_ADDRESS>)

<prefix> ::= <ip-route>
            <METRIC>

```



```
(<FLAGS>...)
[(<ospf-prefix> | <isis-prefix>)]

<ip-route> ::= <ip-address>
             <PREFIX_LENGTH>

<igp-link-extension> ::= <LINK_NAME>
                        <METRIC>
                        [(<FLAGS>...)]
                        [ (<isis-link> |<ospf-link>) ]

<igp-termination-point-extension> ::= (<ip-address>...) | <NUMBER>
```

The model extends the original network-topology model as follows:

- o Additional topology attributes are introduced. This allows to introduce additional flags in extending modules that are associated with specific IGP topologies, without needing to revise this component.
- o Additional node attributes are introduced. New attributes include router-id, a set of flags, as well a list of prefixes. Each prefix in turn includes an ip prefix, a metric and a prefix-specific set of flags. The prefix is also extensible to any OSPF/ISIS specific attributes.
- o Links are extended as well with a set of parameters, allowing to associate a link with an IGP name, another set of flags, and a link metric.

4.2.2. The OSPF-Topology Component

OSPF is the next type of topology represented in the model. OSPF represents a particular type of Layer 3 Unicast IGP. Accordingly, the Layer 3 Unicast IGP topology model needs to be extended. The corresponding extensions are introduced in a separate component "ospf-topology", whose structure is depicted in the following diagram. For the most part, this module extends the IGP topology info model.

```
<ospf-topology> ::= <AREA_IDENTIFIER>

<ospf-node> ::= <ospf-router-type>
               [<DR_INTERFACE_IDENTIFIER>]
               [(<MULTI_TOPOLOGY_IDENTIFIER>...)]
               [<ospf-node-capabilities>]
               [<ted-node>]
```



```
<ospf-router-type> ::= <ABR> | <ASBR> | <INTERNAL> | <PSEUDONODE>

<ospf-node-capabilities> ::= <>

<ospf-link> ::= [<MULTI_TOPOLOGY_IDENTIFIER>]
                [<ted-link>]

<ospf-prefix> ::= [<forwarding-address>]

<forwarding-address> ::= <IPV4><IPV4_ADDRESS>
```

The module extends the IGP topology component as follows:

- o Additional topology attributes are introduced, which extend the `igp-topology-extension` of the L3 Unicast topology component. The attributes include an OSPF area-id identifying the area.
- o OSPF node extends IGP node extension with additional OSPF attributes. New attributes include `router-type`, `DR_INTERFACE_IDENTIFIER` for pseudonodes, list of multi-topology-ids, OSPF node capabilities and traffic engineering attributes.
- o Links are extended with multi-topology-id and traffic engineering link attributes.
- o IGP prefixes are extended with OSPF specific forwarding address.

4.2.3. The IS-IS-Topology Component

IS-IS is another type of Layer 3 Unicast IGP. Like OSPF topology, IS-IS topology is defined in a separate component, "isis-topology", which extends IGP topology.


```
<isis-topology> ::= <NET_IDENTIFIER>

<isis-node> ::= <isis-router-type> <iso>
               (<NET_IDENTIFIER>...)
               [( <MULTI_TOPOLOGY_IDENTIFIER>...)]
               [<ted-node>]

<iso> ::= <ISO_SYSTEM_ID> <ISO_PSEUDONODE_ID>
<isis-router-type> ::= <LEVEL_2> | <LEVEL_1> | <LEVEL_2_1>

<isis-link> ::= [ <MULTI_TOPOLOGY_IDENTIFIER> ]
                [<ted-link>]

<isis-prefix> ::= <>
```

The module extends the IGP topology component as follows:

- o Additional topology attributes are introduced, which extend the `igp-topology-attributes` of the L3 Unicast topology component. The attributes include an ISIS NET-id identifying the area.
- o ISIS node extends IGP node with the following ISIS attributes `router-type`, `iso-system-id` to identify the router, list of `multi-topology-id`, list of NET ids and traffic engineering attributes.
- o Links are extended with `multi-topology-id` and traffic engineering link attributes.

4.2.4. The TED (Traffic Engineering Data) Component

Traffic Engineering Data is required both by OSPF and IS-IS, which are defined in separate components. Information shared by both is defined in another component, "TED". This component defines a set of groupings with auxiliary information required and shared by those other components.

```
<ted-node> ::= <te-router-id>
               <local-address>
               <pcc-capabilities>

<te-router-id> ::= [ <IPV4><IPV4_ADDRESS> ]
                  [ <IPV6><IPV6_ADDRESS> ]

<local-address> ::= [ (( <IPV4><IPV4_ADDRESS> )... ) ] |
                   [ (( <IPV6><IPV6_ADDRESS><PREFIX_OPTION> )... ) ]

<pcc-capabilities> ::= <>
```



```
<ted-link> ::= <COLOR>
             <MAX_LINK_BANDWIDTH>
             <MAX_RESV_LINK_BANDWIDTH>
             (<UNRESERVED_BANDWIDTH>...)
             <TE_DEFAULT_METRIC>
             [<srlg-attributes>]

<srlg-attributes> ::= ( <interface-switching-capabilities>...)
                    (<SRLG_VALUE>...)
                    <LINK_PROTECTION_TYPE>

<interface-switching-capabilities> ::= <switching-capabilities>
                                       <ENCODING>
                                       (<MAX_LSP_BANDWIDTH>...)
                                       [<packet-switch-capable>]
                                       [<tdm-capable>]

<packet-switch-capable> ::= <MINIMUM_LSP_BANDWIDTH>
                           <INTERFACE_MTU>

<time-division-switch-capable> ::= <MINIMUM_LSP_BANDWIDTH>
                                   <INDICATION>

<switching-capabilities> ::= <>
```

This module details traffic-engineering node and link attributes:

- o TED node attributes include te-router-id for IPv4 and IPv6, local IPv4 and IPv6 addresses and path computation client capabilities.
- o TED link attributes comprise COLOR, MAX_LINK_BANDWIDTH, MAX_RESV_LINK_BANDWIDTH, a list of UNRESERVED_BANDWIDTH and TE_DEFAULT_METRIC. They also include SRLG attributes which contains interface switching capabilities, a list of SRLG_VALUE and LINK_PROTECTION_TYPE. The interface switching capabilities in turn contains switching capability, ENCODING, MAX_LSP_BANDWIDTH and interface switching specific attributes.

5. Security Considerations

TBD

6. Contributors

The model presented in this paper was contributed to by more people than can be listed on the author list. Additional contributors include:

- o Ken Gray, Juniper Networks
- o Tom Nadeau, Juniper Networks
- o Aleksandr Zhdankin, Cisco
- o Tony Tkacik, Cisco
- o Robert Varga, Pantheon Technologies

7. Acknowledgements

We wish to acknowledge the helpful contributions, comments, and suggestions that were received from many people. We'd also like to thank the people that contributed to the topology information model during the I2RS Interim meeting in April 2013 - Shane Amante, Andy Bierman, Ed Crabbe, Adrian Farrel, and Joel Halpern.

8. References

8.1. Normative References

- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", [RFC 5511](#), April 2009.
- [RFC6020] Bjorklund, M., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", [RFC 6020](#), October 2010.

8.2. Informative References

- [I-D.amante-i2rs-topology-use-cases] Medved, J., Previdi, S., Lopez, V., and S. Amante, "Topology API Use Cases", [draft-amante-i2rs-topology-use-cases-01](#) (work in progress), October 2013.
- [RFC3444] Pras, A. and J. Schoenwaelder, "On the Difference between Information Models and Data Models", [RFC 3444](#), January 2003.

Authors' Addresses

Jan Medved
Cisco

E-Mail: jmedved@cisco.com

Nitin Bahadur
Juniper Networks

E-Mail: nitinb@juniper.net>

Alexander Clemm
Cisco

E-Mail: alex@cisco.com

Hariharan Ananthakrishnan
Juniper Networks

E-Mail: hanantha@juniper.net

