

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 10, 2016

A. Clemm
J. Medved
R. Varga
T. Tkacik
Cisco
N. Bahadur
Bracket Computing
H. Ananthakrishnan
Packet Design
December 8, 2015

A Data Model for Network Topologies
draft-ietf-i2rs-yang-network-topo-02.txt

Abstract

This document defines an abstract (generic) YANG data model for network/service topologies and inventories. The model serves as a base model which is augmented with technology-specific details in other, more specific topology and inventory models.

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 June 10, 2016.

Copyright Notice

Copyright (c) 2015 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 3
- 2. Definitions and Acronyms 7
- 3. Model Structure Details 7
 - 3.1. Base Network Model 7
 - 3.2. Base Network Topology Model 10
 - 3.3. Extending the model 11
 - 3.4. Discussion and selected design decisions 12
 - 3.4.1. Container structure 12
 - 3.4.2. Underlay hierarchies and mappings 12
 - 3.4.3. Dealing with changes in underlay networks 13
 - 3.4.4. Use of groupings 13
 - 3.4.5. Cardinality and directionality of links 13
 - 3.4.6. Multihoming and link aggregation 14
 - 3.4.7. Mapping redundancy 14
 - 3.4.8. Typing 14
 - 3.4.9. Representing the same device in multiple networks . . 14
 - 3.5. Open Issues 15
 - 3.5.1. Supporting client as well as server provided network topology 16
 - 3.5.2. Identifiers of string or URI type 16
- 4. YANG Modules 17
 - 4.1. Defining the Abstract Network: network.yang 17
 - 4.2. Creating Abstract Network Topology: network-topology.yang 21
- 5. Security Considerations 28
- 6. Contributors 28
- 7. Acknowledgements 28
- 8. References 28
 - 8.1. Normative References 28

8.2. Informative References 29
Authors' Addresses 29

1. Introduction

This document introduces an abstract (base) YANG [RFC6020] [RFC6021] data model to represent networks and topologies. The data model is divided into two parts. The first part of the model defines a network model that allows to define network hierarchies (i.e. network stacks) and to maintain an inventory of nodes contained in a network. The second part of the model augments the basic network model with information to describe topology information. Specifically, it adds the concepts of links and termination points to describe how nodes in a network are connected to each other. Moreover the model introduces vertical layering relationships between networks that can be augmented to cover both network inventories and network/service topologies.

While it would be possible to combine both parts into a single model, the separation facilitates integration of network topology and network inventory models, by allowing to augment network inventory information separately and without concern for topology into the network model.

The model can be augmented to describe specifics of particular types of networks and topologies. For example, an augmenting model can provide network node information with attributes that are specific to a particular network type. Examples of augmenting models include models for Layer 2 network topologies, Layer 3 network topologies, such as Unicast IGP, IS-IS [RFC1195] and OSPF [RFC2328], traffic engineering (TE) data [RFC3209], or any of the variety of transport and service topologies. Information specific to particular network types will be captured in separate, technology-specific models.

The basic data models introduced in this document are generic in nature and can be applied to many network and service topologies and inventories. The models allow applications to operate on an inventory or topology of any network at a generic level, where specifics of particular inventory/topology types are not required. At the same time, where data specific to a network type does come into play and the model is augmented, the instantiated data still adheres to the same structure and is represented in consistent fashion. This also facilitates the representation of network hierarchies and dependencies between different network components and network types.

The abstract (base) network YANG module introduced in this document, entitled "network.yang", contains a list of abstract network nodes

and defines the concept of network hierarchy (network stack). The abstract network node can be augmented in inventory and topology models with inventory and topology specific attributes. Network hierarchy (stack) allows any given network to have one or more "supporting networks". The relationship of the base network model, the inventory models and the topology models is shown in the following figure (dotted lines in the figure denote possible augmentations to models defined in this document).




Figure 1: The network model structure

The network-topology YANG module introduced in this document, entitled "network-topology.yang", defines a generic topology model at its most general level of abstraction. The module defines a topology graph and components from which it is composed: nodes, edges and termination points. Nodes (from the network.yang module) represent graph vertices and links represent graph edges. Nodes also contain termination points that anchor the links. A network can contain multiple topologies, for example topologies at different layers and overlay topologies. The model therefore allows to capture relationships between topologies, as well as dependencies between nodes and termination points across topologies. An example of a topology stack is shown in the following figure.




Figure 3: Topology hierarchy (stack) example

The figure shows two VPN service topologies (VPN1 and VPN2) instantiated over a common L3 topology. Each VPN service topology is mapped onto a subset of nodes from the common L3 topology.

There are multiple applications for such a data model. For example, within the context of I2RS, nodes within the network can use the data model to capture their understanding of the overall network topology and expose it to a network controller. A network controller can then use the instantiated topology data to compare and reconcile its own view of the network topology with that of the network elements that it controls. Alternatively, nodes within the network could propagate this understanding to compare and reconcile this understanding either among themselves or with help of a controller. Beyond the network element and the immediate context of I2RS itself, a network controller might even use the data model to represent its view of the topology that it controls and expose it to applications north of itself. Further use cases that the data model can be applied to are described in [topology-use-cases].

2. Definitions and Acronyms

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

Data subtree: An instantiated data node and the data nodes that are hierarchically contained within it.

HTTP: Hyper-Text Transfer Protocol

IGP: Interior Gateway Protocol

IS-IS: Intermediate System to Intermediate System protocol

NETCONF: Network Configuration Protocol

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

URI: Uniform Resource Identifier

ReST: Representational State Transfer, a style of stateless interface and protocol that is generally carried over HTTP

YANG: A data definition language for NETCONF

3. Model Structure Details

3.1. Base Network Model

The abstract (base) network model is defined in the network.yang module. Its structure is shown in the following figure. Brackets enclose list keys, "rw" means configuration data, "ro" means operational state data, and "?" designates optional nodes. A "+" indicates a line break.

```

module: ietf-network
  +--rw networks
  |   +--rw network* [network-id]
  |   |   +--rw network-types
  |   |   +--rw network-id          network-id
  |   |   +--rw supporting-network* [network-ref]
  |   |   |   +--rw network-ref    -> /networks/network/network-id
  |   |   +--rw node* [node-id]
  |   |   |   +--rw node-id        node-id
  |   |   |   +--rw supporting-node* [network-ref node-ref]
  |   |   |   |   +--rw network-ref -> ../../../../supporting-network/+
  |   |   |   |   |   network-ref
  |   |   |   +--rw node-ref      -> /networks/network/node/node-id
  |   +--ro networks-state
  |   |   +--ro network* [network-ref]
  |   |   +--ro network-ref      -> /networks/network/network-id
  |   +--ro server-provided?    boolean

```

Figure 4: The structure of the abstract (base) network model

The model contains a container with a list of networks, and a container with a list of network state.

Container "networks" contains a list of networks. Each network is captured in its own list entry, distinguished via a network-id.

A network has a certain type, such as L2, L3, OSPF or IS-IS. A network can even have multiple types simultaneously. The type, or types, are captured underneath the container "network-types". In this module it serves merely as an augmentation target; network-specific modules will later introduce new data nodes to represent new network types below this target, i.e. insert them below "network-types" by ways of yang augmentation.

When a network is of a certain type, it will contain a corresponding data node. Network types SHOULD always be represented using presence containers, not leafs of empty type. This allows to represent hierarchies of network subtypes within the instance information. For example, an instance of an OSPF network (which, at the same time, is a layer 3 unicast IGP network) would contain underneath "network-types" another container "l3-unicast-igp-network", which in turn would contain a container "ospf-network".

A network can in turn be part of a hierarchy of networks, building on top of other networks. Any such networks are captured in the list "supporting-network". A supporting network is in effect an underlay network.

Furthermore, a network contains an inventory of nodes that are part of the network. The nodes of a network are captured in their own list. Each node is identified relative to its containing network by a node-id.

It should be noted that a node does not exist independently of a network; instead it is a part of the network that it is contained in. In cases where the same entity takes part in multiple networks, or at multiple layers of a networking stack, the same entity will be represented by multiple nodes, one for each network. In other words, the node represents an abstraction of the device for the particular network that it is part of. To represent that the same entity or same device is part of multiple topologies or networks, it is possible to create one "physical" network with a list of nodes for each of the devices or entities. This (physical) network, respectively the (entities) nodes in that network, can then be referred to as underlay network and nodes from the other (logical) networks and nodes, respectively. Note that the model allows to define more than one underlay network (and node), allowing for simultaneous representation of layered network- and service topologies and physical instantiation.

Similar to a network, a node can be supported by other nodes, and map onto one or more other nodes in an underlay network. This is captured in the list "supporting-node". The resulting hierarchy of nodes allows also to represent device stacks, where a node at one level is supported by a set of nodes at an underlying level. For example, a "router" node might be supported by a node representing a route processor and separate nodes for various line cards and service modules, a virtual router might be supported or hosted on a physical device represented by a separate node, and so on.

Container "networks-state" contains data about the network state. It contains a list of networks, mirroring the corresponding list under the "networks" container and containing in each data node state information for the corresponding network. Currently, each list element contains a single state, "server-provided". This state indicates how the network came into being. If false, the network was configured by a client application, for example in the case of an overlay network that is configured by a controller application. If true, the network was populated by the server itself, respectively an application on the server that is able to discover the network.

Network data can come into being in one of two ways. In one way, network data is configured by client applications, for example in case of overlay networks that are configured by an SDN Controller application. In another way, it is populated by the server, in case of networks that can be discovered. Client applications SHOULD NOT

modify configurations of networks for which "server-provided" is true. (Please note this is an open issue for further discussion, see section Section 3.5.)

3.2. Base Network Topology Model

The abstract (base) network topology model is defined in the "network-topology.yang" module. It builds on the network model defined in the "network.yang" module, augmenting it with links (defining how nodes are connected) and termination-points (which anchor the links and are contained in nodes). The structure of the network topology module is shown in the following figure. Brackets enclose list keys, "rw" means configuration data, "ro" means operational state data, and "?" designates optional nodes. A "+" indicates a line break.

```

module: ietf-network-topology
augment /nd:networks/nd:network:
  +--rw link* [link-id]
    +--rw source
      | +--rw source-node -> ../../../../nd:node/node-id
      | +--rw source-tp?  -> ../../../../nd:node[nd:node-id=current()/+
      |                   ../../source-node]/termination-point/tp-id
    +--rw destination
      | +--rw dest-node   -> ../../../../nd:node/node-id
      | +--rw dest-tp?   -> ../../../../nd:node[nd:node-id=current()/+
      |                   ../../dest-node]/termination-point/tp-id
    +--rw link-id       link-id
    +--rw supporting-link* [network-ref link-ref]
      +--rw network-ref -> ../../../../nd:supporting-network/
      |                 network-ref
      +--rw link-ref   -> /nd:networks/network[nd:network-id=+
      |                   current()/../../network-ref]/link/link-id
augment /nd:networks/nd:network/nd:node:
  +--rw termination-point* [tp-id]
    +--rw tp-id                tp-id
    +--rw supporting-termination-point* [network-ref node-ref tp-ref]
      +--rw network-ref        -> ../../../../nd:supporting-node/network-ref
      +--rw node-ref           -> ../../../../nd:supporting-node/node-ref
      +--rw tp-ref             -> /nd:networks/network[nd:network-id=+
      |                   current()/../../network-ref]/node+
      |                   [nd:node-id=current()/../../node-ref]/+
      |                   termination-point/tp-id

```

Figure 5: The structure of the abstract (base) network topology model

A node has a list of termination points that are used to terminate links. An example of a termination point might be a physical or logical port or, more generally, an interface.

Like a node, a termination point can in turn be supported by an underlying termination point, contained in the supporting node of the underlay network.

A link is identified by a link-id that uniquely identifies the link within a given 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. Similar to a node, a link can map onto one or more links in an underlay topology (which are terminated by the corresponding underlay termination points). This is captured in the list "supporting-link".

3.3. Extending the model

In order to derive a model for a specific type of network, the base model can be extended. This can be done roughly as follows: for the new network type, a new YANG module is introduced. In this module a number of augmentations are defined against the network and network-topology YANG modules.

We start with augmentations against the network.yang module. First, a new network type needs to be defined. For this, a presence container that resembles the new network type is defined. It is inserted by means of augmentation below the network-types container. Subsequently, data nodes for any network-type specific node parameters are defined and augmented into the node list. The new data nodes can be defined as conditional ("when") on the presence of the corresponding network type in the containing network. In cases where there are any requirements or restrictions in terms of network hierarchies, such as when a network of a new network-type requires a specific type of underlay network, it is possible to define corresponding constraints as well and augment the supporting-network list accordingly. However, care should be taken to avoid excessive definitions of constraints.

Subsequently, augmentations are defined against network-topology.yang. Data nodes are defined both for link parameters, as well as termination point parameters, that are specific to the new network type. Those data nodes are inserted by way of augmentation into the link and termination-point lists, respectively. Again, data nodes can be defined as conditional on the presence of the corresponding network-type in the containing network, by adding a corresponding "when"-statement.

It is possible, but not required, to group data nodes for a given network-type under a dedicated container. Doing so introduces further structure, but lengthens data node path names.

In cases where a hierarchy of network types is defined, augmentations can in turn augmenting modules, with the module of a network "sub-type" augmenting the module of a network "super-type".

3.4. Discussion and selected design decisions

3.4.1. Container structure

Rather than maintaining lists in separate containers, the model is kept relatively flat in terms of its containment structure. Lists of nodes, links, termination-points, and supporting-nodes, supporting-links, and supporting-termination-points are not kept in separate containers. Therefore, path specifiers used to refer to specific nodes, be it in management operations or in specifications of constraints, can remain relatively compact. Of course, this means there is no separate structure in instance information that separates elements of different lists from one another. Such structure is semantically not required, although it might enhance human readability in some cases.

3.4.2. Underlay hierarchies and mappings

To minimize assumptions of what a particular entity might actually represent, mappings between networks, 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.

Where additional specifics about mappings between upper and lower layers are required, those can be captured in augmenting modules. For example, to express that a termination point in a particular network type maps to an interface, an augmenting module can introduce an augmentation to the termination point which introduces a leaf of type `ifref` that references the corresponding interface [RFC7223]. Similarly, if a node maps to a particular device or network element, an augmenting module can augment the node data with a leaf that references the network element.

It is possible for links at one level of a hierarchy to map to multiple links at another level of the hierarchy. For example, a VPN topology might model VPN tunnels as links. Where a VPN tunnel maps to a path that is composed of a chain of several links, the link will contain a list of those supporting links. Likewise, it is possible

for a link at one level of a hierarchy to aggregate a bundle of links at another level of the hierarchy.

3.4.3. Dealing with changes in underlay networks

It is possible for a network to undergo churn even as other networks are layered on top of it. When a supporting node, link, or termination point is deleted, the supporting leafrefs in the overlay will be left dangling. To allow for this possibility, the model makes use of the "require-instance" construct of YANG 1.1 [I.D.draft-ietf-netmod-rfc6020bis].

It is the responsibility of the application maintaining the overlay to deal with the possibility of churn in the underlay network. When a server receives a request to configure an overlay network, it SHOULD validate whether supporting nodes/links/tps refer to nodes in the underlay are actually in existence. Configuration requests in which supporting nodes/links/tps refer to objects currently not in existence SHOULD be rejected. It is the responsibility of the application to update the overlay when a supporting node/link/tp is deleted at a later point in time. For this purpose, an application might subscribe to updates when changes to the underlay occur, for example using mechanisms defined in [I-D.draft-ietf-netconf-yang-push].

3.4.4. Use of groupings

The model makes use of groupings, instead of simply defining data nodes "in-line". This allows to more easily include the corresponding data nodes in notifications, which then do not need to respecify each data node that is to be included. The tradeoff for this is that it makes the specification of constraints more complex, because constraints involving data nodes outside the grouping need to be specified in conjunction with a "uses" statement where the grouping is applied. This also means that constraints and XPath-statements need to be specified in such a way that they navigate "down" first and select entire sets of nodes, as opposed to being able to simply specify them against individual data nodes.

3.4.5. Cardinality and directionality of links

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 to applications that make use of graph algorithms. Bidirectional connections can be represented through pairs of unidirectional links. Multipoint networks can be represented through

pseudo-nodes (similar to IS-IS, for example). By introducing hierarchies of nodes, with nodes at one level mapping onto a set of other nodes at another level, and introducing new links for nodes at that level, topologies with connections representing non-point-to-point communication patterns can be represented.

3.4.6. Multihoming and link aggregation

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

3.4.7. Mapping redundancy

In a hierarchy of networks, there are nodes mapping to nodes, links mapping to links, and termination points mapping to termination points. Some of this information is redundant. Specifically, if the link-to-links mapping is known, and the termination points of each link are known, termination point mapping information can be derived via transitive closure and does not have to be explicitly configured. Nonetheless, in order to not constrain applications regarding which mappings they want to configure and which should be derived, the model does provide for the option to configure this information explicitly. The model includes integrity constraints to allow for validating for consistency.

3.4.8. Typing

A network's network types are represented using a container which contains a data node for each of its network types. A network can encompass several types of network simultaneously, hence a container is used instead of a case construct, with each network type in turn represented by a dedicated presence container itself. The reason for not simply using an empty leaf, or even simpler, do away even with the network container and just use a leaf-list of network-type instead, is to be able to represent "class hierarchies" of network types, with one network type refining the other. Network-type specific containers are to be defined in the network-specific modules, augmenting the network-types container.

3.4.9. Representing the same device in multiple networks

One common requirement concerns the ability to represent that the same device can be part of multiple networks and topologies. However, the model defines a node as relative to the network that it is contained in. The same node cannot be part of multiple

topologies. In many cases, a node will be the abstraction of a particular device in a network. To reflect that the same device is part of multiple topologies, the following approach might be chosen: A new type of network to represent a "physical" (or "device") network is introduced, with nodes representing devices. This network forms an underlay network for logical networks above it, with nodes of the logical network mapping onto nodes in the physical network.

This scenario is depicted in the following figure. It depicts three networks with two nodes each. A physical network P consists of an inventory of two nodes, D1 and D2, each representing a device. A second network, X, has a third network, Y, as its underlay. Both X and Y also have the physical network P as underlay. X1 has both Y1 and D1 as underlay nodes, while Y1 has D1 as underlay node. Likewise, X2 has both Y2 and D2 as underlay nodes, while Y2 has D2 as underlay node. The fact that X1 and Y1 are both instantiated on the same physical node D1 can be easily derived.




Figure 6: Topology hierarchy example - multiple underlays

In the case of a physical network, nodes represent physical devices and termination points physical ports. It should be noted that it is also conceivable to augment the model for a physical network-type, defining augmentations that have nodes reference system information and termination points reference physical interfaces, in order to provide a bridge between network and device models.

3.5. Open Issues

3.5.1. Supporting client as well as server provided network topology

YANG requires data needs to be designated as either configuration or operational data, but not both, yet it is important to have all network information, including vertical cross-network dependencies, captured in one coherent model. In most cases network topology information is discovered about a network; the topology is considered a property of the network that is reflected in the model. That said, it is conceivable that certain types of topology need to also be configurable by an application.

There are several alternatives in which this can be addressed. The alternative chosen in this draft does not restrict network topology information as read-only, but includes a state that indicates for each network whether it is populated by the server or by a client application. (The drawback of this solution is that it stretches its use of the configuration concept. Configuration information is subject to backup and restore, which is not applicable to server-provided information. Perhaps more serious is the ability of a client to lock a configuration and thus prevent changes to server-provided network topology while the lock is in effect. Preventing this requires special treatment of network topology related configuration information.)

In another alternative, all information about network topology that is in effect is represented as network state, i.e. as read-only information, regardless of how it came into being. For cases where network topology needs to be configured, a second branch for configurable topology information is introduced. Any network topology configuration is mirrored by network state information. A configurable network will thus be represented twice: once in the read-only list of all networks, a second time in a configuration sandbox. (The drawback of this solution is slightly increased complexity of augmentations due to two target branches.)

3.5.2. Identifiers of string or URI type

The current model defines identifiers of nodes, networks, links, and termination points as URIs. An alternative would define them as string.

The case for strings is that they will be easier to implement. The reason for choosing URIs is that the topology/node/tp exists in a larger context, hence it is useful to be able to correlate identifiers across systems. While strings, being the universal data type, are easier for human beings (a string is a string is a string), they also muddle things. What typically happens is that strings have some structure which is magically assigned and the knowledge of this

structure has to be communicated to each system working with the data. A URI makes the structure explicit and also attaches additional semantics: the URI, unlike a free-form string, can be fed into a URI resolver, which can point to additional resources associated with the URI. This property is important when the topology data is integrated into a larger, more complex system.

4. YANG Modules

4.1. Defining the Abstract Network: network.yang

```
<CODE BEGINS> file "ietf-network@2015-12-08.yang"
module ietf-network {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-network";
  prefix nd;

  import ietf-inet-types {
    prefix inet;
  }

  organization
    "IETF I2RS (Interface to the Routing System) Working Group";

  contact
    "WG Web:      <http://tools.ietf.org/wg/i2rs/>
     WG List:     <mailto:i2rs@ietf.org>

     WG Chair:    Susan Hares
                 <mailto:shares@ndzh.com>

     WG Chair:    Jeffrey Haas
                 <mailto:jhaas@pfrc.org>

     Editor:      Alexander Clemm
                 <mailto:alex@cisco.com>

     Editor:      Jan Medved
                 <mailto:jmedved@cisco.com>

     Editor:      Robert Varga
                 <mailto:rovarga@cisco.com>

     Editor:      Tony Tkacik
                 <mailto:ttkacik@cisco.com>

     Editor:      Nitin Bahadur
                 <mailto:nitin_bahadur@yahoo.com>
```

Editor: Hariharan Ananthakrishnan
<mailto:hari@packetdesign.com>;

description

"This module defines a common base model for a collection of nodes in a network. Node definitions are further used in network topologies and inventories.

Copyright (c) 2015 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 (<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of draft-ietf-i2rs-yang-network-topo-02; see the RFC itself for full legal notices.

NOTE TO RFC EDITOR: Please replace above reference to draft-ietf-i2rs-yang-network-topo-02 with RFC number when published (i.e. RFC xxxx).";

```
revision 2015-12-08 {
  description
    "Initial revision.
    NOTE TO RFC EDITOR: Please replace the following reference
    to draft-ietf-i2rs-yang-network-topo-02 with
    RFC number when published (i.e. RFC xxxx).";
  reference
    "draft-ietf-i2rs-yang-network-topo-02";
}

typedef node-id {
  type inet:uri;
  description
    "Identifier for a node.";
}

typedef network-id {
  type inet:uri;
  description
    "Identifier for a network.";
}
```

```

grouping network-ref {
  description
    "Contains the information necessary to reference a network,
    for example an underlay network.";
  leaf network-ref {
    type leafref {
      path "/nd:networks/nd:network/nd:network-id";
      require-instance false;
    }
    description
      "Used to reference a network, for example an underlay
      network.";
  }
}

grouping node-ref {
  description
    "Contains the information necessary to reference a node.";
  leaf node-ref {
    type leafref {
      path "/nd:networks/nd:network[nd:network-id=current()/../"+
        "network-ref]/nd:node/nd:node-id";
      require-instance false;
    }
    description
      "Used to reference a node.
      Nodes are identified relative to the network they are
      contained in.";
  }
  uses network-ref;
}

container networks {
  description
    "Serves as top-level container for a list of networks.";
  list network {
    key "network-id";
    description
      "Describes a network.
      A network typically contains an inventory of nodes,
      topological information (augmented through
      network-topology model), as well as layering
      information.";
    container network-types {
      description
        "Serves as an augmentation target.
        The network type is indicated through corresponding
        presence containers augmented into this container.";
    }
  }
}

```

```

}
leaf network-id {
  type network-id;
  description
    "Identifies a network.";
}
list supporting-network {
  key "network-ref";
  description
    "An underlay network, used to represent layered network
    topologies.";
  leaf network-ref {
    type leafref {
      path "/networks/network/network-id";
      require-instance false;
    }
    description
      "References the underlay network.";
  }
}
}
list node {
  key "node-id";
  description
    "The inventory of nodes of this network.";
  leaf node-id {
    type node-id;
    description
      "Identifies a node uniquely within the containing
      network.";
  }
  list supporting-node {
    key "network-ref node-ref";
    description
      "Represents another node, in an underlay network, that
      this node is supported by. Used to represent layering
      structure.";
    leaf network-ref {
      type leafref {
        path "../..../supporting-network/network-ref";
        require-instance false;
      }
      description
        "References the underlay network that the
        underlay node is part of.";
    }
  }
  leaf node-ref {
    type leafref {
      path "/networks/network/node/node-id";
    }
  }
}

```



```
    prefix inet;
  }
import ietf-network {
  prefix nd;
}
```

```
organization
  "IETF I2RS (Interface to the Routing System) Working Group";
```

```
contact
```

```
"WG Web:    <http://tools.ietf.org/wg/i2rs/>
WG List:    <mailto:i2rs@ietf.org>
```

```
WG Chair:   Susan Hares
            <mailto:shares@ndzh.com>
```

```
WG Chair:   Jeffrey Haas
            <mailto:jhaas@pfr.org>
```

```
Editor:     Alexander Clemm
            <mailto:alex@cisco.com>
```

```
Editor:     Jan Medved
            <mailto:jmedved@cisco.com>
```

```
Editor:     Robert Varga
            <mailto:rovarga@cisco.com>
```

```
Editor:     Tony Tkacik
            <mailto:ttkacik@cisco.com>
```

```
Editor:     Nitin Bahadur
            <mailto:nitin_bahadur@yahoo.com>
```

```
Editor:     Hariharan Ananthakrishnan
            <mailto:hari@packetdesign.com>";
```

```
description
```

```
"This module defines a common base model for network topology,
augmenting the base network model with links to connect nodes,
as well as termination points to terminate links on nodes.
```

```
Copyright (c) 2015 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
(<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of
draft-ietf-i2rs-yang-network-topo-02;
see the RFC itself for full legal notices.

NOTE TO RFC EDITOR: Please replace above reference to
draft-ietf-i2rs-yang-network-topo-02 with RFC
number when published (i.e. RFC xxxx).";

```
revision 2015-12-08 {
  description
    "Initial revision.
    NOTE TO RFC EDITOR: Please replace the following reference
    to draft-ietf-i2rs-yang-network-topo-02 with
    RFC number when published (i.e. RFC xxxx).";
  reference
    "draft-ietf-i2rs-yang-network-topo-02.";
}

typedef link-id {
  type inet:uri;
  description
    "An identifier for a link in a topology.
    The identifier SHOULD be chosen such that the same link in a
    real network topology will always be identified through the
    same identifier, even if the model is instantiated in
    separate datastores. An implementation MAY choose to capture
    semantics in the identifier, for example to indicate the type
    of link and/or the type of topology that the link is a part
    of.";
}

typedef tp-id {
  type inet:uri;
  description
    "An identifier for termination points on a node.
    The identifier SHOULD be chosen such that the same TP in a
    real network topology will always be identified through the
    same identifier, even if the model is instantiated in
    separate datastores. An implementation MAY choose to capture
    semantics in the identifier, for example to indicate the type
    of TP and/or the type of node and topology that the TP is a
    part of.";
}
```

```

grouping link-ref {
  description
    "References a link in a specific network.";
  leaf link-ref {
    type leafref {
      path "/nd:networks/nd:network[nd:network-id=current()/../"+
        "network-ref]/lnk:link/lnk:link-id";
      require-instance false;
    }
    description
      "A type for an absolute reference a link instance.
      (This type should not be used for relative references.
      In such a case, a relative path should be used instead.)";
  }
  uses nd:network-ref;
}

grouping tp-ref {
  description
    "References a termination point in a specific node.";
  leaf tp-ref {
    type leafref {
      path "/nd:networks/nd:network[nd:network-id=current()/../"+
        "network-ref]/nd:node[nd:node-id=current()/../"+
        "node-ref]/lnk:termination-point/lnk:tp-id";
      require-instance false;
    }
    description
      "A type for an absolute reference to a termination point.
      (This type should not be used for relative references.
      In such a case, a relative path should be used instead.)";
  }
  uses nd:node-ref;
}

augment "/nd:networks/nd:network" {
  description
    "Add links to the network model.";
  list link {
    key "link-id";
    description
      "A Network Link connects a by Local (Source) node and
      a Remote (Destination) Network Nodes via a set of the
      nodes' termination points.
      As it is possible to have several links between the same
      source and destination nodes, and as a link could
      potentially be re-homed between termination points, to
      ensure that we would always know to distinguish between

```

```

links, every link is identified by a dedicated link
identifier.
Note that a link models a point-to-point link, not a
multipoint link.
Layering dependencies on links in underlay topologies are
not represented as the layering information of nodes and of
termination points is sufficient.";
container source {
  description
    "This container holds the logical source of a particular
    link.";
  leaf source-node {
    type leafref {
      path "../..../nd:node/nd:node-id";
    }
    mandatory true;
    description
      "Source node identifier, must be in same topology.";
  }
  leaf source-tp {
    type leafref {
      path "../..../nd:node[nd:node-id=current()/../"+
        "source-node]/termination-point/tp-id";
    }
    description
      "Termination point within source node that terminates
      the link.";
  }
}
container destination {
  description
    "This container holds the logical destination of a
    particular link.";
  leaf dest-node {
    type leafref {
      path "../..../nd:node/nd:node-id";
    }
    mandatory true;
    description
      "Destination node identifier, must be in the same
      network.";
  }
  leaf dest-tp {
    type leafref {
      path "../..../nd:node[nd:node-id=current()/../"+
        "dest-node]/termination-point/tp-id";
    }
    description

```

```

        "Termination point within destination node that
        terminates the link.";
    }
}
leaf link-id {
    type link-id;
    description
        "The identifier of a link in the topology.
        A link is specific to a topology to which it belongs.";
}
list supporting-link {
    key "network-ref link-ref";
    description
        "Identifies the link, or links, that this link
        is dependent on.";
    leaf network-ref {
        type leafref {
            path "../..../nd:supporting-network/nd:network-ref";
            require-instance false;
        }
        description
            "This leaf identifies in which underlay topology
            supporting link is present.";
    }
    leaf link-ref {
        type leafref {
            path "/nd:networks/nd:network[nd:network-id=current()]/"+
                "../network-ref]/link/link-id";
            require-instance false;
        }
        description
            "This leaf identifies a link which is a part
            of this link's underlay. Reference loops, in which
            a link identifies itself as its underlay, either
            directly or transitively, are not allowed.";
    }
}
}
}
}
augment "/nd:networks/nd:network/nd:node" {
    description
        "Augment termination points which terminate links.
        Termination points can ultimately be mapped to interfaces.";
    list termination-point {
        key "tp-id";
        description
            "A termination point can terminate a link.
            Depending on the type of topology, a termination point

```

```

    could, for example, refer to a port or an interface.";
leaf tp-id {
  type tp-id;
  description
    "Termination point identifier.";
}
list supporting-termination-point {
  key "network-ref node-ref tp-ref";
  description
    "The leaf list identifies any termination points that
    the termination point is dependent on, or maps onto.
    Those termination points will themselves be contained
    in a supporting node.
    This dependency information can be inferred from
    the dependencies between links. For this reason,
    this item is not separately configurable. Hence no
    corresponding constraint needs to be articulated.
    The corresponding information is simply provided by the
    implementing system.";
leaf network-ref {
  type leafref {
    path "../..../nd:supporting-node/nd:network-ref";
    require-instance false;
  }
  description
    "This leaf identifies in which topology the
    supporting termination point is present.";
}
leaf node-ref {
  type leafref {
    path "../..../nd:supporting-node/nd:node-ref";
    require-instance false;
  }
  description
    "This leaf identifies in which node the supporting
    termination point is present.";
}
leaf tp-ref {
  type leafref {
    path "/nd:networks/nd:network[nd:network-id=current()/" +
      "../network-ref]/nd:node[nd:node-id=current()/" +
      "node-ref]/termination-point/tp-id";
    require-instance false;
  }
  description
    "Reference to the underlay node, must be in a
    different topology";
}
}

```

```
    }  
  }  
}
```

<CODE ENDS>

5. Security Considerations

The transport protocol used for sending the topology data **MUST** support authentication and **SHOULD** support encryption. The data-model by itself does not create any security implications.

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, Cisco Systems
- o Xufeng Liu, Ericsson
- o Tom Nadeau, Brocade
- o Aleksandr Zhdankin, Cisco

7. Acknowledgements

We wish to acknowledge the helpful contributions, comments, and suggestions that were received from Alia Atlas, Vishna Pavan Beeram, Andy Bierman, Martin Bjorklund, Igor Bryskin, Benoit Claise, Susan Hares, Ladislav Lhotka, Carlos Pignataro, Juergen Schoenwaelder, and Xian Zhang.

8. References

8.1. Normative References

- [RFC1195] Callon, R., "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", RFC 1195, December 1990.
- [RFC2328] Moy, J., "OSPF Version 2", RFC 2328, April 1998.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

- [RFC6020] Bjorklund, M., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, October 2010.
- [RFC6021] Schoenwaelder, J., "Common YANG Data Types", RFC 6021, October 2010.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.
- [RFC7223] Bjorklund, M., "A YANG Data Model for Interface Management", RFC 7223, May 2014.

8.2. Informative References

- [I-D.draft-ietf-netconf-yang-push]
Clemm, A., Voit, E., and A. Gonzalez Prieto, "Subscribing to YANG datastore push updates", I-D draft-ietf-netconf-yang-push-00, October 2015.
- [I.D.draft-ietf-netmod-rfc6020bis]
Bjorklund, M., "The YANG 1.1 Data Modeling Language", I-D draft-ietf-netmod-rfc6020bis-08, October 2015.
- [topology-use-cases]
Medved, J., Previdi, S., Lopez, V., and S. Amante, "Topology API Use Cases", I-D draft-amante-i2rs-topology-use-cases-01, October 2013.

Authors' Addresses

Alexander Clemm
Cisco

E-Mail: alex@cisco.com

Jan Medved
Cisco

E-Mail: jmedved@cisco.com

Robert Varga
Cisco

E-Mail: rovarga@cisco.com

Tony Tkacik
Cisco

EMail: ttkacik@cisco.com

Nitin Bahadur
Bracket Computing

EMail: nitin_bahadur@yahoo.com

Hariharan Ananthakrishnan
Packet Design

EMail: hari@packetdesign.com