

I2RS working group
Internet-Draft
Intended status: Standards Track
Expires: August 1, 2015

S. Hares
Q. Wu
X. Guan
Huawei
January 28, 2015

An Information model for service topology
draft-hares-i2rs-info-model-service-topo-02

Abstract

This service model is a protocol independent topology for the interface to routing system (I2RS).

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 August 1, 2015.

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.

Table of Contents

1.	Introduction	2
2.	Conventions used in this document	2
3.	Service Topology Information Model	3
3.1.	Base Model: the Service-Topology Component	3
3.2.	The TED (Traffic Engineering Data) Component	6
3.3.	Inventory datastore Component	7
4.	Security Considerations	8
5.	IANA Considerations	8
6.	References	8
6.1.	Normative References	8
6.2.	Informative References	9
	Authors' Addresses	9

1. Introduction

A service specific overlay utilized by Service chaining creates the service topology. The overlay creates a path between service function(SF) nodes. Service functions can be co-located on one SF Node or physically separated across several SF Nodes with each having one or more Service Functions. In either case, a service function may be running in its own virtualized system space or natively on the hosting system.

Within the service topology, an ordered set of Service functions will be invoked for each packet that belongs to a given flow for which a SFC will be applied. Adding new service function to SF Node in the topology is easily accomplished, and no underlying network changes are required. Furthermore, additional service Functions or Service Function instances, for redundancy or load distribution purpose, can be added or removed to the service topology as required.

As stated in [I-D.ietf-sfc-problem-statement] the service overlay is independent of the network topology and allows operators to use whatever overlay or underlay they prefer and to locate service nodes in the network as needed.

2. Conventions used in this document

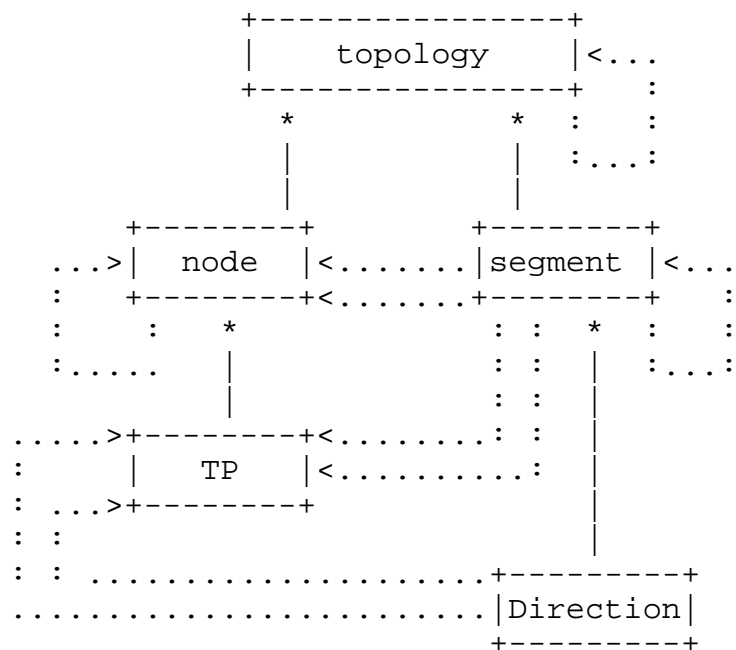
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [RFC2119].

3. Service Topology Information Model

This section specifies the service 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.

3.1. Base Model: the Service-Topology Component

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



The basic information model works as follows: A service topology contains service nodes and segments. A segment connects two nodes (a source and a destination) and have direction, may be unidirectional or bidirectional. unidirectional is one where traffic is passed through a set of SF nodes in one forwarding direction only. Bidirectional is one where traffic is passed through a set of SF nodes in both forwarding directions. Each SF node contains termination points. It occurs before or after other service node, therefore each node may have its upstream SF node and/or downstream SF node.

A SF node may be dedicated to a tenant(e.g., an IPVPN customer), globally shared among tenants, or available to be assigned in whole or in part to a tenant or a set of tenants. Therefore SF Nodes can map onto and be supported by other SF nodes, while Segment can map onto and be supported by other segments,e.g., one segment can be mapped to two consecutive segments stitching together. Service Topologies can map onto other, underlay topologies. However in some

cases when some services are dedicated to a tenant or topology information are not gathered using IGP or BGP, Service Topologies should be independent from network topology and therefore should not map onto other, underlay topologies.

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

```

/* exterior definitions for service topology */

<service-topology> ::= (<topology>...)
<STopo> ::= (<topology> ...)

/* Topology definitions */
<topology> ::= <TOPOLOGY_IDENTIFIER>
              [<node-count>]
              (<segment>...)
              (<node>...)
              [<topology-type>]
              [<underlay-topologies>]
              [<topology-extension>]

<node-count> ::= INTEGER-32;
<topology-type> ::= (<snmp> [<snmp-topology-type>]) |
                   (<ipfix> [<ipfix-topology-type>]) |
                   (<i2rs> [<i2rs-topology-type>])

<underlay-topologies> ::= (<TOPOLOGY_IDENTIFIER>...)

<topology-extension> ::= <snmp-topology-extension> |
                        <ipfix-topology-extension> |
                        <igp-topology-extension> |
                        <bgp-topology-extension> |
                        ...

<segment> ::= <Segment_IDENTIFIER>
             <source>
             <destination>
             [<direction>]
             [<segment-extension> ]

<source> ::= <termination-point-reference>
<destination> ::= <termination-point-reference>

<termination-point-reference> ::= <SF_NODE_IDENTIFIER>

<direction> ::= (<Unidirection>)|
                (<Bidirection>)

```

```

<segment-extension> ::= <snmp-segment-extension> |
                        <ipfix-segment-extension> |
                        ...
<node> ::= <SF_NODE_IDENTIFIER>
           (<termination-point>...)
           [<SF-NODE_TYPE>]
           [<NEXT-HOP>]
           [<SF-node-extension>]

<termination-point> ::= <TERMINATION_POINT_IDENTIFIER>
                        [<supporting-termination-points>]
                        [<termination-point-extension>]
<supporting-termination-points> ::=
                        (<TERMINATION_POINT_IDENTIFIER>...)
<SF_NODE-TYPE> ::=
                        (<SF-Ingress-Node>)|
                        (<SF-Enabled-Node>)|
                        (<SF-Egress-Node>)
                        ...
<NEXT-HOP> ::= (<SF_NODE_IDENTIFIER>...)
<node-extension> ::= <SF-Ingress-Node-extension> |
                     <SF-Enable-Node-extension> |
                     <SF-Egress-Node-extension>
                     ...

```

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

- o A service overlay 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 SNMP or IPFIX. A topology can even have multiple types simultaneously. The type, or types, are captured in the list of "topology-type" components.
- o A topology contains segments and nodes, 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 segment. An examples of a termination point might be a physical or logical port or, more generally, an interface.
- o A segment is identified by a segment-identifier, uniquely identifying the segment within the topology. segment are point-to-point and has direction. The direction can be unidirectional or

bidirectional. Accordingly, a segment contains a source and a destination. Both source and destination reference a corresponding node, as well as a termination point on that node.

- o The topology, node, segment and direction elements can be extended with topology-specific components (topology-extensions, node-extension, segment-extension and direction-extension respectively).

The topology model includes segment that are either bidirectional unidirectional. Service function chain path is analogue to linked list data structure and can be represented through a set of Ordered segments from source to destination. Each node in the service overlay may be located at different layer. The segment can be setup to steer traffic through these specific service nodes at different layers or bypass some specific service nodes at different layers.

The topology model only supports point to point and does not support multipoint. Therefore Segments are terminated by a single termination point, not sets of termination points. Connections involving multihoming or segment aggregation schemes need to be modeled using multiple point-to-point segment, e.g., connection from service node A at lower layer to service node D at higher layer can comprise a segment 1 from service node A to service node B and segment 2 from service node B to service node C and segment 3 from service node C to service node D. By using segment aggregation, we can define a new segment from service A to service node D which is supported by segment 1, 2 and 3.

Unlike network topology collection, the service topology information may be not available from each SF by using IGP advertisement or BGP-LS northbound distribution since SF may be not located at network layer. However these SF at different layer may have affinity with one SF node (e.g., SF egress node or SF ingress node or SF enabled node), therefore service topology information associated with SF nodes between SFC ingress node and SFC egress node can be collected using IGP or BGP-LS from egress network node or ingress network node. Alternatively, the service node may rely on SNMP or IPFIX interface for interrogation of a virtual device's state, statistics and configuration.

3.2. The TED (Traffic Engineering Data) Component

Traffic Engineering Data for service overlay can be built or supplemented from other sources inventory management system and share to PCE, ALTO server or other topology manager defined in [I.D-ietf-i2rs-architecture]. Information shared by them is defined as the

component, "TED". This component defines a set of groupings with auxiliary information required and shared by those other components.

```
<SF-TED_Ext> ::=      <SF-Enable-Node-Extension>
                      <SF-Node-Locator>
                      <Supported-Context-Type>
                      [<FIB-Size>]
                      [<RIB-Size>]
                      [<MAC-Forwarding-Table-Size>]
                      <SF-Chain-Index>
                      [( <SF-Identifier>
                        <SF-Type>
                        <Customer-ID>...
                        <SF-inventory-data>)... ]

<SF-type> ::=
  <firewall> |
  <loadbalancer> |
  <NAT44> |
  <NAT64> |
  <DPI>
```

This module details traffic-engineering node attributes:

- o TED node attributes include SF-Node-Locator, SF-Type and SF-Identifier, SF-Chain-Index and inventory-data information.

3.3. Inventory datastore Component

Inventory Data for service overlay can be obtained by using SNMP or IPFIX and share to PCE, ALTO server or other topology manager defined in [I-D.ietf-i2rs-architecture]. Information shared by them is defined as the component, "inventory database". This component defines a set of groupings with auxiliary information required and shared by those other components.

```
<SF-inventory-data> ::=
    <SF-capabilities>
    <SF-administrative-info>

<SF-capabilities> ::=
    (<supported-ACL-number >)|
    (<virtual-context-number >)|
    (<supported-packet-rate>)|
    (<supported-bandwidth>)

<SF-administrative-info> ::= =
    (<Packet-rate-utilization>)|
    (<Bandwidth-utilization-per-CoS>)|
    (<Packet-rate-utilization-per-Cos>)|
    (<Memory-utilization>)|
    (<available-memory>)|
    (<RIB-utilization-per-address-family>)|
    (<FIB-utilization-per-address-family>)|
    (<CPU-utilization>)|
    (<Available storage>)|
    (<Bandwidth-utilization>)|
    (<Flow-resource-utilization-per-flow-type>)
```

This module details inventory node attributes:

- o Inventory node attributes include SF-capabilities and SF-administrative-info.

4. Security Considerations

This document does not introduce any new security issues above those identified in [RFC5511].

5. IANA Considerations

This draft includes no request to IANA.

6. References

6.1. Normative References

[I-D.ietf-i2rs-architecture]

Atlas, A., Halpern, J., Hares, S., Ward, D., and T. Nadeau, "An Architecture for the Interface to the Routing System", draft-ietf-i2rs-architecture-08 (work in progress), January 2015.

[I-D.ietf-i2rs-usecase-reqs-summary]

Hares, S. and M. Chen, "Summary of I2RS Use Case Requirements", draft-ietf-i2rs-usecase-reqs-summary-00 (work in progress), November 2014.

[I-D.ietf-sfc-problem-statement]

Quinn, P. and T. Nadeau, "Service Function Chaining Problem Statement", draft-ietf-sfc-problem-statement-10 (work in progress), August 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, April 2009.

6.2. Informative References

[I-D.ietf-idr-ls-distribution]

Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution-10 (work in progress), January 2015.

Authors' Addresses

Susan Hares
Huawei
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: sunseawq@huawei.com

Xiaoran Guan
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: guanxiaoran@huawei.com