              Separating Routing Planes using Segment Routing
                 draft-gulkohegde-routing-planes-using-sr-00

Abstract

   Many network deployments arrange the network topologies in two or
   more planes. The traffic generally uses one of the planes and fails
   over to the other plane when there are link or node failure. Certain
   applications require the traffic to be strictly restricted to a
   particular plane and should not failover to the other plane. This
   document proposes a solution for the strict planar routing using
   Segment Routing.

Requirements Language

   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 RFC 2119 [RFC2119].

Table of Contents

1. Introduction

```
              A3---------------------A4
             / |                   / |
            /  |                  /  |
           /   |                 /   |
          A1---|----------------A2   |
          |    B3---------------|---B4
          |   /                 |   /
          |  /                  |  /
          | /                   | /
          B1---------------------B2
```

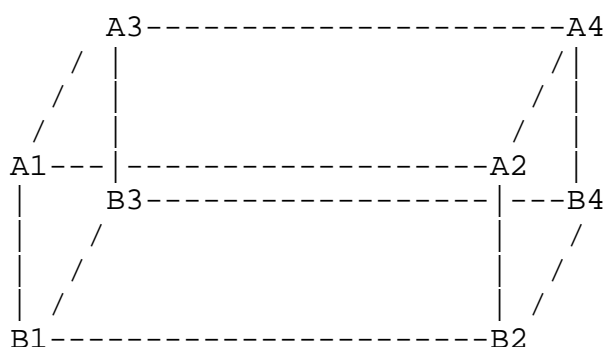   Figure 1: Network Planes

   The above Figure 1 represents a network topology in two planes.Nodes
   A1, A2, A3, A4 are in plane A and B1, B2, B3, B4 are in plane B.
   A1->B1, A2->B2, A3->B3, A4->B4 represent the "shunt links" which
   connect the two planes. Certain applications require that the

traffic follows plane A and remains in plane A in case of failures and does not cross over to plane B. Strict routing plane requirements can be met using multiple techniques. This draft focuses on solutions using Segment Routing technology.

2. Motivation

The motivation of this document is to provide solutions to strict routing plane requirements using Segment Routing. The following objectives help to accomplish this in a range of deployment scenarios.

1. Maintain strict routing within routing planes.

2. Allow traffic to failover within routing plane and do not allow traffic to failover to other planes

3. Achieve ease of configuration and operational management

3. Solutions

There are multiple ways to address the strict routing plane requirements. Section 3.1 describes a mechanism by using diferrent SIDs for each plane.Another option is to use Multi-topology SIDs as defined in Section 3.2.

3.1. Routing-plane SID

A new SID called Routing-Plane SID is defined and associated with new algorithm values. This document proposes 4 new algorithm values which represent SPF in routing-planes. When the network topology is organized into two different planes, each node in plane A associates a new Routing-Plane SID to one of it's loopback addresses and advertises the SID in the segment routing extension defined in [SR-OSPF] section 2.1 and [SR-IS-IS] section 5. The prefix-SID sub-TLV carries the new algorithm values corresponding to the routing-plane. The traffic which needs to be restricted to a certain routing-plane,should use the Routing-Plane SID corresponding to that plane to forward the traffic. The paths through the Routing planes MAY use single Routing Plane SID or a stack of multiple Routing Plane SIDs. Adjacency-SIDs can also be used build paths across routing planes. Adjacency-SIDs are not scoped per-algorithm and it is possible that the protection path for the adjacency SIDs uses a path going over a different routing-plane.It is recommended to use non-protected adjacency-SIDs to avoid backup traffic flowing through different plane.

3.1.1. Elements of procedure

   All the nodes that belong to a certain routing-plane MUST advertise
   the algorithm corresponding to that routing-plane in the algorithm
   sub-TLV as defined in [SR-OSPF] and [SR-IS-IS]. The nodes SHOULD
   also advertise Routing-Plane SID corresponding to that algorithm in
   the prefix-SID Sub-TLV.

   A node that receives the algorithm sub-TLV with new algorithm value
   specified in Section 6 MUST compute SPF with all the nodes that
   advertised the new algorithm. The next-hops and RIB entries for the
   Routing-Plane SIDs MUST be computed from the routing-plane SPF.
   Certain nodes MAY belong to multiple routing-planes. Such nodes MUST
   compute SPF corresponding to each plane and compute the next-hops for
   the SIDs corresponding to each plane.

   Each router MAY assign different IP address corresponding to each
   plane or MAY use the same IP address to assign the node-SIDs and
   Routing-Plane SIDs. The ingress routers MAY advertise binding-SIDs
   as defined in [SR-ARCH]section 3.5.2, for the label stacks that are
   constructed using routing-Plane SIDs. The ingress routers MAY map
   the incoming IP traffic onto the Routing-Plane SIDs, the mechanisms
   to do so is implementation dependant and outside the scope of this
   document.

   When the network topology is organized into multiple IGP levels or
   areas, the Routing Plane SIDs MAY be re-originated from one IGP
   domain into the other domain by the border routers. The border IGP
   routers MUST re-advertise the Routing-Plane SIDs if they belong to
   the corresponding Routing plane and has advertised the algorithm
   corresponding to the routing-plane.

3.2. Multi-topology SID

   Multi topology Routing defines mechanisms to support multiple
   topologies in a single physical network. ISIS and OSPF extensions to
   support multi-topology routing is defined in [RFC5120] and [RFC4915]
   respectively. Multi-topology routing defined in [RFC5120] and
   [RFC4915] describes mechanisms to separate topologies in ISIS and
   OSPF by advertising separate MT-TLVs in ISIS and multi-topology
   scoped Router LSA in OSPF. The different routing planes in customer
   network can be assigned with different MT-ID for each routing-plane
   and the multi-topology SIDs can be advertised for each MT-ID as
   described in [SR-OSPF] and [SR-IS-IS]. Multi-topology SIDs are
   associated with algorithm 0 and no new algorithm definition is
   required. All the nodes in the network MUST also support multi-
   topology routing as defined in [RFC5120] and [RFC4915]. All the
   nodes in the network compute separate SPF per MT-ID and program the

forwarding planes with MT-SIDs accordingly. Multi-topology SIDs are used to build the explicit paths through the network. Multi-topology based solution has an advantage of possibility of assigning different IGP costs to links for different MTs. For deployments that do not need separate IGP costs and topologies for each routing plane, it comes with an additional operational overhead of maintaining multi-topology configurations.

4. Backward compatibility

The mechanism described in the document is fully backward compatible. If a node does not support the extensions defined in this document, it will not advertise the additional algorithm values in the algorithm sub-TLV. All the computing nodes will not consider the node in the SPF computation if it has not advertised the specific algorithm. For the multi-topology based solution backward compatibility mechanism described in [RFC5120] and [RFC4915] are applicable.

5. Security Considerations

This document does not introduce any further security issues other than those discussed in [SR-OSPF] and [SR-IS-IS].

6. IANA Considerations

This specification updates OSPF and ISIS registry:

OSPF Router Information (RI) TLVs Registry

8 (IANA Preallocated) - SR-Algorithm TLV

Algorithm 2 -5 : SPF in routing plane

ISIS Sub TLVs for Type 242

Type: TBD (suggested value 19)

Description: Segment Routing Algorithm

Algorithm 2-5 : SPF in Routing Plane

7. Acknowledgements

8. References

8.1. Normative References

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

   [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P.
              Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF",
              RFC 4915, DOI 10.17487/RFC4915, June 2007,
              <http://www.rfc-editor.org/info/rfc4915>.

   [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi
              Topology (MT) Routing in Intermediate System to
              Intermediate Systems (IS-ISs)", RFC 5120,
              DOI 10.17487/RFC5120, February 2008,
              <http://www.rfc-editor.org/info/rfc5120>.

   [SR-IS-IS]
              "IS-IS Extensions for Segment Routing, draft-ietf-isis-
              segment-routing-extensions-11(work in progress)", October
              2016.

   [SR-OSPF] "OSPF Extensions for Segment Routing, draft-ietf-ospf-
              segment-routing-extensions-11(work in progress)", July
              2016.

   [SR-OSPFv3]
              "OSPFv3 Extensions for Segment Routing, draft-ietf-ospf-
              ospfv3-segment-routing-extensions-06(work in progress)",
              July 2016.

8.2. Informative References

   [RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W.,
              Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute
              Advertisement", RFC 7684, DOI 10.17487/RFC7684, November
              2015, <http://www.rfc-editor.org/info/rfc7684>.

   [SR-ARCH] "Segment Routing Architecture, draft-ietf-spring-segment-
              routing-09(work in progress)", July 2016.

Authors' Addresses

    Shraddha Hegde
    Juniper Networks, Inc.
    Embassy Business Park
    Bangalore, KA 560093
    India

    Email: shraddha@juniper.net

    Arkadiy Gulko
    Thomson Reuters

    Email: arkadiy.gulko@thomsonreuters.com