Segment Routing Conflict Resolution

draft-ginsberg-spring-conflict-resolution-00

Les Ginsberg (ginsberg@cisco.com)
Stefano Previdi (sprevidi@cisco.com)
Peter Psenak (ppsenak@cisco.com)

WHY?

For identifiers w global scope/usage conflicts may occur due to misconfiguration. This will cause forwarding issues (drops, loops).

Consistent (network-wide) and deterministic conflict resolution policy is needed to minimize the damage.

Prior discussions have not reached consensus.

Draft is the vehicle to drive discussions to consensus and document the agreed upon policies.

Problem is cross-protocol – therefore SPRING is the right working group.

Handling Invalid SRGB Entries

Example:

Range 1: (100, 199]

Range 2: (1000, 1099)

Range 3: (100, 599) !Overlaps w Range #1

Range 4: (2000, 2099)

As this is local configuration burden should be on the local node to detect and prevent misconfiguration BEFORE it is advertised.

New text agreed upon – aided by clarification added to draft-ietf-spring-segment-routing-mpls

Update to conflict resolution draft

For the set of ranges to be usable the ranges MUST be disjoint. Sender behavior is defined in various SR protocol drafts such as [SR-IS-IS] which specify that senders MUST NOT advertise overlapping ranges.

Receivers of SRGB ranges MUST validate the SRGB ranges advertised by other nodes. If overlapping ranges are detected receivers MUST ignore all advertised SRGB ranges from that node. Operationally the node is treated as though it did not advertise any SRGB ranges. When the procedures defined in [SR-MPLS] for mapping global SIDs to outgoing labels are followed the advertising node is determined to be incapable of supporting all global SIDs.

Note that utilization of local SIDs (e.g. adjacency SIDs) advertised by a node is not affected by the state of the advertised SRGB.

Update to sr-mpls draft

When different SRGBs are used, the outgoing label value is set as: [SRGB(next_hop)+index]. If the index can't be applied to the SRGB (i.e.: if the index points outside the SRGB of the next-hop or the next-hop has not advertised a valid SRGB), then no outgoing label value can be computed and the next-hop MUST be considered as not supporting the MPLS operations for that particular SID.

Prefix SID Conflict Resolution

No consensus reached yet.

Configuration is distributed –

- **Local configuration of SIDs for local prefixes**
- **Local configuratyion of SRMS advertisements**

Conflicts cannot be prevented before they are advertised

Receivers must apply consistent conflict resolution policy on a consistent Database

Prefix Sid Conflict ResolutionPolicy

Policy	Advantages	Disadvantages
Ignore	Easy to diagnose No unintended traffic flow	**Delivery to all destinations in conflict is compromised**
Preference Rule	***Traffic to some of the destinations in conflict may continue to be forwarded successfully***	Harder to diagnose based on forwarding behavior Introduction of new conflicts may cause other entries in conflict to be used

(Draft is currently agnostic)

Mapping Entry

A generalized mapping entry can be represented using the following definitions:

Pi - Initial prefix
Pe - End prefix
L - Prefix length
Lx - Maximum prefix length (32 for IPv4, 128 for IPv6)
Si - Initial SID value
Se - End SID value
R - Range value

Mapping Entry is then the tuple: (Pi/L, Si, R)

1.1.1.1/32 100 range 1
or
(1.1.1.1/32, 100, 1)

Terminology: Conflict Types

PREFIX CONFLICT

When different SIDs are assigned to the same prefix

(192.0.2.120/32, 200, 1)

(192.0.2.120/32, 30, 1)

Intra-topology

SID CONFLICT:

When the same SID has been assigned to multiple prefixes

(192.0.2.1/32, 200, 1)

(192.0.2.222/32, 200,1)

Inter-topology

How to achieve consistency

All routers MUST have the same database.

Local configuration does not matter unless it is also advertised.

Priority is based on the content of the advertisement – NOT the source of the advertisement

- **Local vs remote does NOT matter**

Consistency: Databases

Two categories of databases:

1)SIDs associated with prefix reachability advertisements

2)SRMS advertisements

Local configuration

Use only what is advertised

Advertisements received from protocol peers

Preference Rule

1. **Smaller range wins**
2. **IPv4 entry wins over IPv6 entry**
3. **Smaller prefix length wins**
4. **Smaller starting address (considered as an unsigned integer value) wins**
5. **Smaller starting SID wins**

(Identical entries from different sources – does not matter which is used)

#1 is key –

biased towards prefix advertisements over SRMS

minimizes impact of conflicts – one misconfig cannot override many valid entries

Smaller range wins

1.1.1.1/32 10 range 1
2.2.2.2/32 20 range 1
3.3.3.3/32 1 range 50

Larger range wins

1.1.1.1/32 10 range 1
2.2.2.2/32 20 range 1
3.3.3.3/32 1 range 50

Implementing Policy: Quarantine

Advertisements	Active Policy	Excluded Entries
(1.1.1.1/32, 10, 1) (2.2.2.2/32, 20, 1) (3.3.3.3/32, 1, 50)	(1.1.1.1/32, 10, 1) (2.2.2.2/32, 20, 1)	(3.3.3.3/32, 1, 50)

Implementing Policy: Overlap Only

Advertisements	Active Policy	Excluded Entries
(1.1.1.1/32, 10, 1) (2.2.2.2/32, 20, 1) (3.3.3.3/32, 1, 50)	(1.1.1.1/32, 10, 1) (2.2.2.2/32, 20, 1) *(3.3.3.3/32, 1, 9)* *(3.3.3.13/32, 11, 9)* *(3.3.3.23/32, 21, 30)*	*(3.3.3.12/32, 10, 1)* *(3.3.3.22/32, 20, 1)*

Yellow => derived entry

Implementing Policy: Ignore

Advertisements	Active Policy	Excluded Entries
(1.1.1.1/32, 10, 1) (2.2.2.2/32, 20, 1) (3.3.3.3/32, 1, 50)		(1.1.1.1/32, 10, 1) (2.2.2.2/32, 20, 1) (3.3.3.3/32, 1, 50)

Next Steps

Some Consensus on conflict resolution policy

Issue V1 of the draft

WG adoption