

# Segment Routing Conflict Resolution

## draft-ietf-spring-conflict-resolution-01

Les Ginsberg ([ginsberg@cisco.com](mailto:ginsberg@cisco.com))

Stefano Previdi ([sprevidi@cisco.com](mailto:sprevidi@cisco.com))

Peter Psenak ([ppsenak@cisco.com](mailto:ppsenak@cisco.com))

Martin Pilka ([martin@infobox.sk](mailto:martin@infobox.sk))

# 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.**

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

**Multiple options currently in the draft – one (and only one) will be selected and specified as the normative behavior.**

# Since Buenos Aires...

**Draft has become a WG document**

**V0 incorporated agreement on handling SRGB misconfigurations**

**V1 updated options for SID conflict resolution**

**Options are still being actively discussed**

# 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:  $[\text{SRGB}(\text{next\_hop}) + \text{index}]$ . 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 configuration of SRMS advertisements**

**Local configuration conflicts can be prevented before they are advertised but...**

**Conflicts between advertisements from different nodes cannot be prevented before they are advertised**

**Receivers must apply consistent conflict resolution policy on a consistent Database**

# Mapping Entry

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

Src – PFX or SRMS

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

T – Topology

A - Algorithm

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

PFX, 1.1.1.1/32, 100, 1, 0,0

SRMS, 1,1,1,1/32, 100, 100, 0, 0



# Terminology: Conflict Types

## **PREFIX CONFLICT**

When different SIDs are assigned to the same prefix

(PFX, 192.0.2.120/32, 200, 1, 0, 0)

(PFX, 192.0.2.120/32, 30, 1, 0, 0)

Same topology and algorithm

## **SID CONFLICT:**

When the same SID has been assigned to multiple prefixes

(PFX, 192.0.2.1/32, 200, 1, 0, 0)

(PFX, 192.0.2.222/32, 200,1, 0, 1)

Across all topologies and all algorithms

# 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**

# Sources of Data

Type	Sources
Protocol prefix reachability advertisements	Local: Use only what is advertised Remote: Use what is received and reachable regardless of whether the prefix will be installed in the RIB
SRMS Advertisements	Local: Use only if advertised Remote: Use from sources which are reachable

# Preference Rule

1. **PFX source wins over SRMS source**
  2. **Smaller range wins**
  3. **IPv6 entry wins over IPv4 entry**
  4. **Longer prefix length wins**
  5. **Smaller algorithm wins**
  6. **Smaller starting address (considered as an unsigned integer value) wins**
  7. **Smaller starting SID wins**
  8. **If topology IDs are NOT identical, ignore BOTH entries**
- Otherwise - Identical entries from different sources – does not matter which is used**

# Prefix Sid Conflict Resolution Policies

Policy	Description	Comment
Ignore	<del>Ignore all entries which have a conflict</del>	<del>Maximum negative impact on traffic.</del>
Quarantine (AKA “per advertisement”)	Use preference rule to compare advertisements, quarantine the losing advertisement	Moderate impact on traffic with moderate complexity
Ignore Overlap Only (AKA “per FEC”)	Use preference rule on a per FEC basis to determine a winner. Ignore only the losing conflicting entries – not the entire advertisement	Maximizes traffic delivery – most complex to implement.

# Working Example

## Partial SR Deployment:

100 nodes

50 SR capable

50 non-SR capable

## SR capable nodes – advertise SIDs in prefix reachability:

(PFX, 192.0.2.1/32, 1, 1)

...

(PFX, 192.0.2.50/32, 50, 1)

## SR incapable nodes – advertise SIDs in SRMS from one or more SR capable nodes

(SRMS, 192.0.2.1/32, 1, 100)

Active Policy (nn) – “nn” the number of prefixes with a SID in the Active database

Excluded Policy (nn) – “nn” the number of prefixes without a SID in the Active database

## Example 1a: Local Prefix SID Config Error (smaller range preferred)

Quarantine		
Advertisements (100)	Active Policy (50)	Excluded Entries (50)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	
...	...	
<b>(PFX, 192.0.2.12/32,52,1)</b>	(192.0.2.12,52,1)	
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, 1,100)		(192.0.2.1/32, 1, 100)

Ignore Overlap Only		
Advertisements (100)	Active Policy (99)	Excluded Entries (1)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	
...	...	
<b>(PFX, 192.0.2.12/32,52,1)</b>	(192.0.2.12,52,1)	
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, 1,100)	<b>(192.0.2.1/32,1,11)</b> <b>(192.0.2.13,13,39)</b> <b>(192.0.2.53/32,53,48)</b>	<b>(192.0.2.12/32, 12, 1)</b> <b>(192.0.2.52/32,52,1)</b>

**Yellow => derived entry**

## Example 2a: SRMS Prefix error (smaller range preferred)

Quarantine		
Advertisements (100)	Active Policy (50)	Excluded Entries (50)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, <b>192.0.2.0</b> /32, 1,100)		(192.0.2.0/32, 1, 100)

Ignore Overlap Only		
Advertisements (100)	Active Policy (98)	Excluded Entries (2)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, <b>192.0.2.0</b> /32, 1,100)	<b>(192.0.2.51/32,52,48)</b>	<b>(192.0.2.0/32, 1, 52)</b>

**Yellow => derived entry**



Example 3a: SRMS SID error (smaller range preferred)

<b>Quarantine</b>		
Advertisements (100)	Active Policy (50)	Excluded Entries (50)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, <b>101</b> ,100)		(192.0.2.0/32, 101, 100)

<b>Ignore Overlap Only</b>		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, <b>101</b> ,100)	<b>(192.0.2.51/32,151,50)</b>	<b>(192.0.2.1/32, 101, 51)</b>

**Yellow => derived entry**

## Example 4a: SRMS SID error – Redundant SRMS (smaller range preferred)

Quarantine		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	(192.0.2.0/32, 1, 100)
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, <b>101</b> ,100)		
(SRMS, 192.0.2.1/32, 1,100)	(192.0.2.1/32,1,100)	

Ignore Overlap Only		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	(192.0.2.1/32,101, 100)
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, 101,100)		
(SRMS, 192.0.2.1/32, 1,100)	(192.0.2.1/32,1,100)	

Yellow => derived entry

# Revised Preference Rule

Smaller range preferred (in draft)

1. **PFX source wins over SRMS source**
2. **Smaller range wins**
3. IPv6 entry wins over IPv4 entry
4. Longer prefix length wins
5. Smaller algorithm wins
6. Smaller starting address (considered as an unsigned integer value) wins
7. Smaller starting SID wins
8. If topology IDs are NOT identical, ignore BOTH entries

Otherwise - Identical entries from different sources – does not matter which is used

Larger range preferred (New rule)

1. **Larger range wins**
2. IPv6 entry wins over IPv4 entry
3. Longer prefix length wins
4. Smaller algorithm wins
5. Smaller starting address (considered as an unsigned integer value) wins
6. Smaller starting SID wins
7. If topology IDs are NOT identical, ignore BOTH entries

Otherwise - Identical entries from different sources – does not matter which is used

## Example 1b: Local Prefix SID Config Error (prefer larger range)

Quarantine		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	
...	...	
<b>(PFX, 192.0.2.12/32,52,1)</b>		(192.0.2.12,52,1)
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, 1,100)	(192.0.2.1/32, 1, 100)	

Ignore Overlap Only		
Advertisements	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	
...	...	
<b>(PFX, 192.0.2.12/32,52,1)</b>		(192.0.2.12,52,1)
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, 1,100)	(192.0.2.1/32,1,100)	

**Yellow => derived entry**

## Example 2b: SRMS Prefix error (larger range preferred)

Quarantine		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)		(192.0.2.1/32,1,1)
...		...
(PFX,192.0.2.50/32,50,1)		(192.0.2.50/32,50,1)
(SRMS, <b>192.0.2.0</b> /32, 1,100)	(192.0.2.0/32, 1, 100)	

Ignore Overlap Only		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)		(192.0.2.1/32,1,1)
...		...
(PFX,192.0.2.50/32,50,1)		(192.0.2.50/32,50,1)
(SRMS, <b>192.0.2.0</b> /32, 1,100)	(192.0.2.0/32,1,100)	

Yellow => derived entry

Example 3b: SRMS SID error (larger range preferred)

Quarantine		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)		(192.0.2.1/32,1,1)
...		...
(PFX,192.0.2.50/32,50,1)		(192.0.2.50/32,50,1)
(SRMS, 192.0.2.1/32, <b>101</b> ,100)	(192.0.2.1/32,101,100)	

Ignore Overlap Only		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)		(192.0.2.1/32,1,1)
...		...
(PFX,192.0.2.50/32,50,1)		(192.0.2.50/32,50,1)
(SRMS, 192.0.2.1/32, <b>101</b> ,100)	(192.0.2.1/32,101,100)	

Yellow => derived entry

## Example 4b: SRMS SID error – Redundant SRMS (larger range preferred)

Quarantine		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	(192.0.2.0/32, 1, 100)
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, <b>101</b> ,100)	(192.0.2.1/32,1,100)	
(SRMS, 192.0.2.1/32, 1,100)	(192.0.2.1/32,1,100)	

Ignore Overlap Only		
Advertisements (100)	Active Policy (100)	Excluded Entries (0)
(PFX,192.0.2.1/32, 1, 1)	(192.0.2.1/32,1,1)	(192.0.2.1/32,101, 100)
...	...	
(PFX,192.0.2.50/32,50,1)	(192.0.2.50/32,50,1)	
(SRMS, 192.0.2.1/32, 101,100)	(192.0.2.1/32,1,100)	
(SRMS, 192.0.2.1/32, 1,100)	(192.0.2.1/32,1,100)	

Yellow => derived entry

# Summary of Test Results

Range Preference	Quarantine		Ignore Overlap		
	Test	Active	Excluded	Active	Excluded
Small	1a	50	50	99	1
	2a	50	50	98	2
	3a	50	50	100	0
	4a	100	0	100	0
Large	1b	100	0	100	0
	2b	100	0	100	0
	3b	100	0	100	0
	4b	100	0	100	0



# Goals and Strategies

Goal	Quarantine	Ignore Overlap Only	Comments
Minimize Complexity	Less complex	More complex	Time to market Interoperability
Maximize SR Forwarding	May ignore non-conflicting SIDs	Best	Traffic loss

Strategy	Quarantine	Ignore Overlap Only	Comments
Use redundant SRMS	Improves coverage	Little impact	Provides redundancy
Prefer larger range	Improves coverage	No impact	Promotes use of SRMS as global provisioning tool

# Next Steps

Input from protocol vendors:

*Provide feedback on relative complexity between Quarantine and Ignore Overlap*

Input from network operators:

*Is quarantine + revised preference rule a desirable solution?*

Decide how “weight” will be advertised/used for SRMS entries (per node vs per advertisement)

New version of the draft reflecting input