

RIFT extensions for SRv6

draft-cheng-rift-srv6-extensions-01

Presenter : [Changwang Lin \(New H3C Technologies\)](#)

Co-authors: Weiqiang Cheng (China Mobile)

Changwang Lin (New H3C Technologies)

Ruixue Wang (China Mobile)

IETF117

Why We Need SRv6 in DataCenter

Using compressed SRv6 in data center networks brings several advantages

- **Simplified end-to-end protocols**

Currently, multiple tunneling technologies such as MPLS and VxLAN are used separately in the data center and inter-data center.

SRv6 can simplify end-to-end protocols.

- **Enhanced TE capabilities**

SRv6 enables easier load balancing and facilitates adaptive routing.

- **Better Service Function Chaining (SFC)**

SRv6 inherently supports SFC and can be considered for simplified service provisioning through end-to-end orchestration with WAN networks.

- **Enable Network Programmability**

Customers have the complete flexibility to program the SID in SRH to enable simplified network programming.

SRv6 in data center networks with RIFT

By using SRv6 END and SRv6 END.X,
It is possible to bypass congested nodes or links.

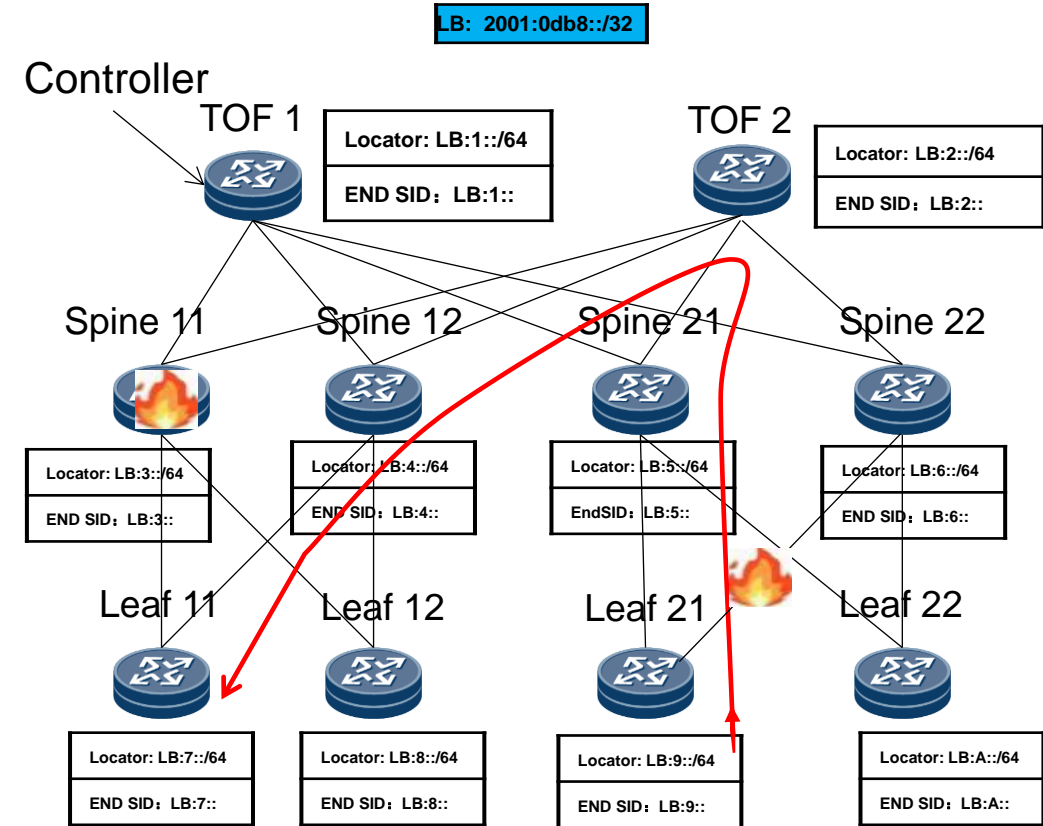
- When network congestion is detected, the controller can optimize network forwarding performance by issuing Segment Routing (SR) paths to avoid the congested paths.
- Either Spine Node 11 is experiencing congestion or the link between Leaf 21 and Spine 22 is congested. In either case, the controller specifies a forwarding path of Leaf21 -> Spine21 -> TOF2 -> Spine12 -> Leaf11 to optimize network performance.
- SRv6 Compress SID Encapsulation offers minimal overhead.

SRv6 Full SID Encapsulation

SA: 2001::21
 DA:2001:0db8:0005:0000:0000:0000:0000
 SL = 3
 Seglist [0]:DA:2001:0db8:0007:0000:0000:0000:0000
 Seglist [1]:DA:2001:0db8:0004:0000:0000:0000:0000
 Seglist [2]:DA:2001:0db8:0002:0000:0000:0000:0000
 Seglist [3]:DA:2001:0db8:0005:0000:0000:0000:0000

SRv6 Compress SID Encapsulation

SA: 2001::21
 DA:2001:0db8:0005:0002:0004:0007:0000:0000



ISIS/OSPF for SRv6 Extension

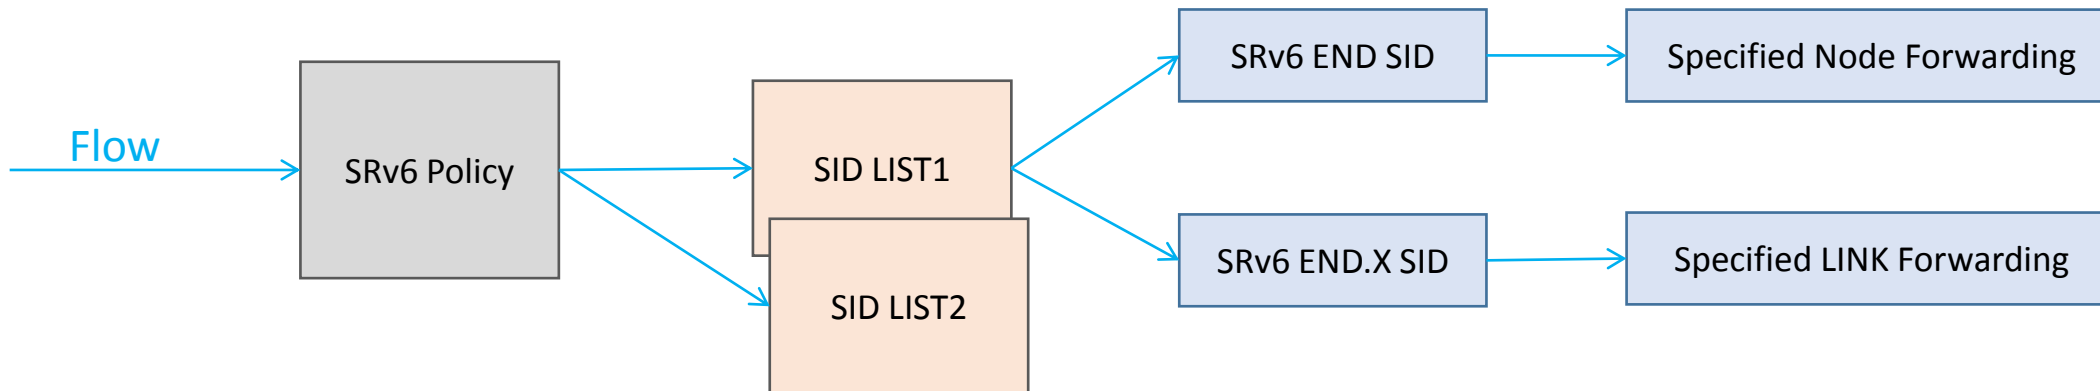
| ISIS/OSPF | Usage | RIFT |
|---|--|-----------------|
| SRv6 Capbility TLV | An SRv6 Capabilities TLV to advertise the SRv6 features and SRH operations | No need |
| Maxmum SID Depths(MSD) | Carried in SRv6 Capbility TLV. Several sub-TLVs to advertise various SRv6 Maximum SID Depths | No need |
| SRv6 Locator | Carried in SRv6 Locator TLV. Advertise the SRv6 Locator information with SRv6 END SIDs and LB/LN. | Need |
| SRv6 End SID | See SRv6 Locator | Need |
| Local Block Length(LB) Local Node Length(LN) | See SRv6 Locator | Need |
| SRv6 End.X SID | Advertising SRv6 Adjacency SIDs(SRv6 End.X SID and SRv6 LAN End.X SID) | Need |
| BGP-LS | Used to advertise link-state information including SRv6 link-state information through the BGP protocol. | To be discussed |

RIFT Extension for SRv6

- SRv6 Locator KV TIE for Zero Touch Provisioning (ZTP)
 - ✓ Distribute SRv6 Locator information with SRv6 END SIDs
 - ✓ Distribute SRv6 Adjacency SIDs
- SRv6 Locator Prefix TIE
 - ✓ Northbound: The Locator Prefix and End SID information
 - ✓ Southbound: Only the Locator Prefix of the local device is advertised
- SRv6 Adjacency SIDs
 - ✓ End.X SID is advertised via NeighborsTIEElement in Node TIE

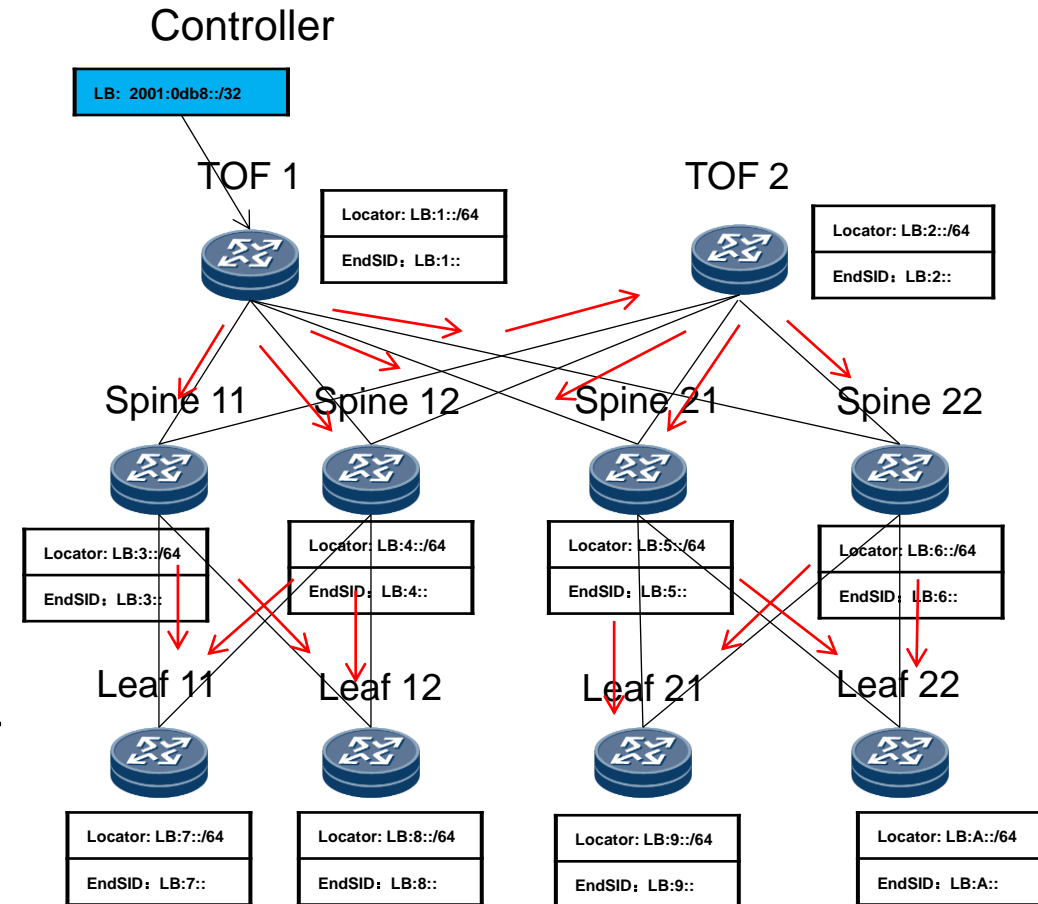
RIFT Extension for SRv6(Traffic Forwarding)

- Traffic steering based on SRv6 Policy
- SRv6 END SIDs: lookup default route or SRv6 locator route
- SRv6 Adjacency-SIDs(END.X) : only makes sense with parallel links between nodes
- SRv6 Local SIDs: END、 END.X、 BindingSID、 L3/L2 ServiceSID



SRv6 Locator KV TIE for ZTP

- The controller is linked to TOF and distributes the configuration of all routers to TOF, including the SRv6 Locator and End SID of each router. TOF then sends this information to all routers in the network through the **new defined SRv6 Locator KV TIE**.
- After the device has established neighbors, the neighbor information is transmitted to TOF through the northbound Node TIE and then passed to the controller. If the device enables the ZTP feature, **the End.X SID will also be configured by the controller**.
- After receiving the neighbor information of the router, the controller configures the End.X SID for each neighbor of the router, and **passes the corresponding End.X SID information to the router through Locator KV TIE**.
- **Through SRv6 Locator KV TIE, all routers can obtain their own Locator, END SID and End.x SID configurations, achieving zero configuration.**



Defines of SRv6 Locator KV TIE(#1)

Publish SRv6 related configurations for all routers via the Locator KV TIE.

SRv6 Locator Entry is shown below:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Well-Known      |      Key-Type describing a SRv6 Locator Info
|
+-----+-----+-----+-----+
|                  (System ID,
|
|                  SRv6 Locator Entry,
|                  SRv6 End.x Entry)
|
+-----+-----+-----+-----+
SRv6 Locator Key/Value Pair
  
```

```

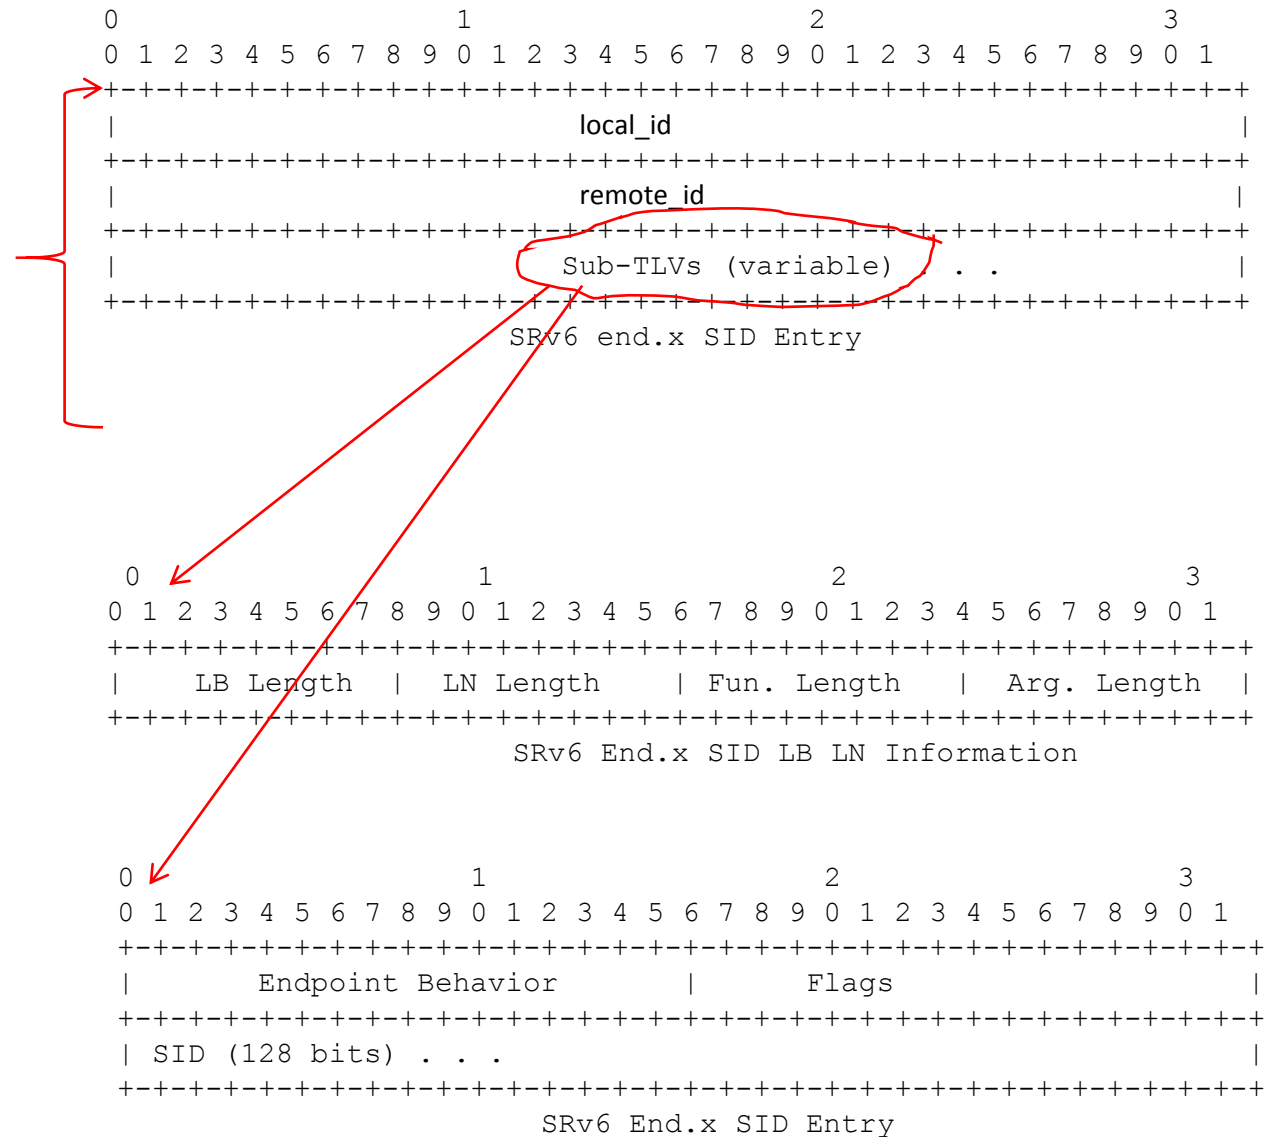
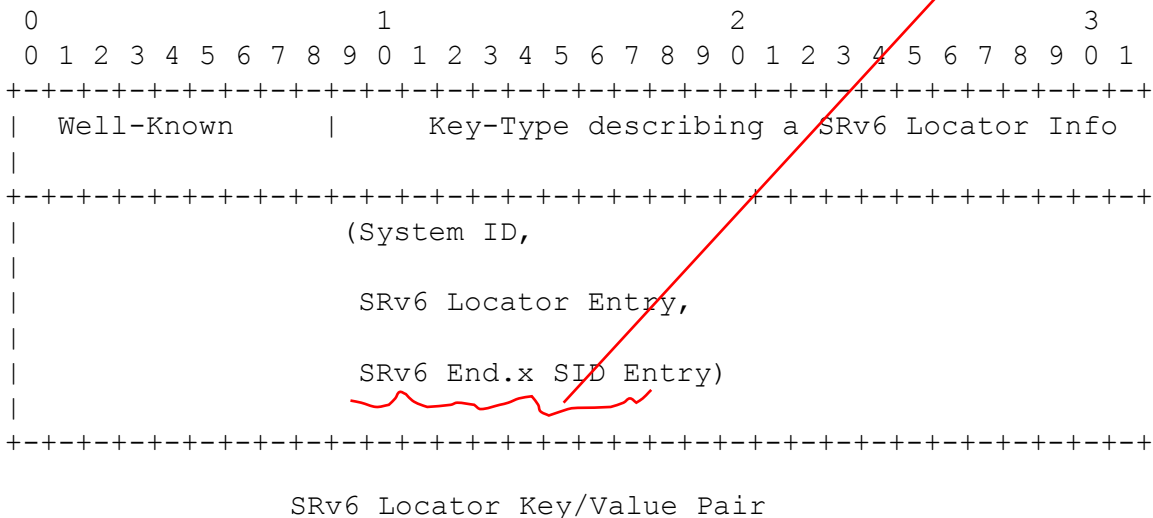
0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|                                     Metric
+-----+-----+-----+-----+
|  Loc Size      |                                     Resv
+-----+-----+-----+-----+
//                                     Locator (continued, variable)                                     //
|                                     Sub-TLVs (variable) . . .
+-----+-----+-----+-----+
SRv6 Locator Entry

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|  LB Length    |  LN Length    |  Fun. Length    |  Arg. Length    |
+-----+-----+-----+-----+
SRv6 SID LB LN Information

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|  Endpoint Behavior  |  Flags
+-----+-----+-----+-----+
|  SID (128 bits) . . .
+-----+-----+-----+-----+
SRv6 End SID Entry
  
```

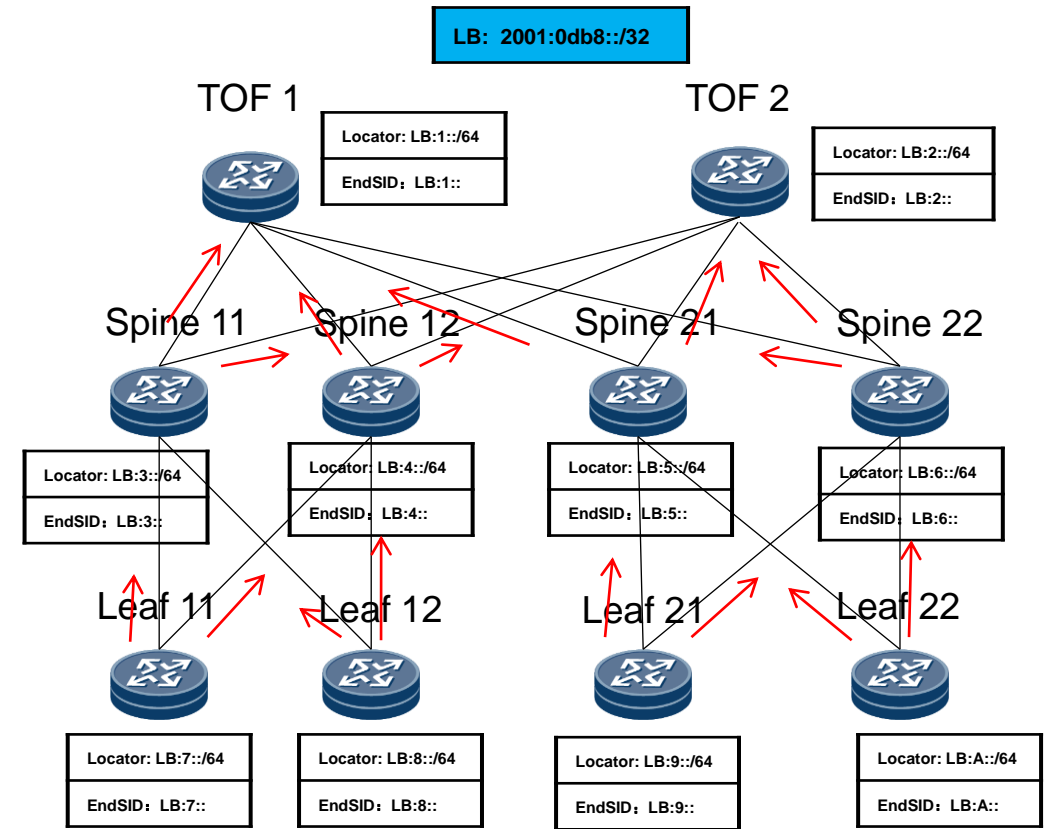
Defines of SRv6 Locator KV TIE(#2)

Publish SRv6 related configurations for all routers via the SRv6 Locator KV TIE.
SRv6 End.x SID Entry is shown below:



SRv6 Locator Prefix TIE

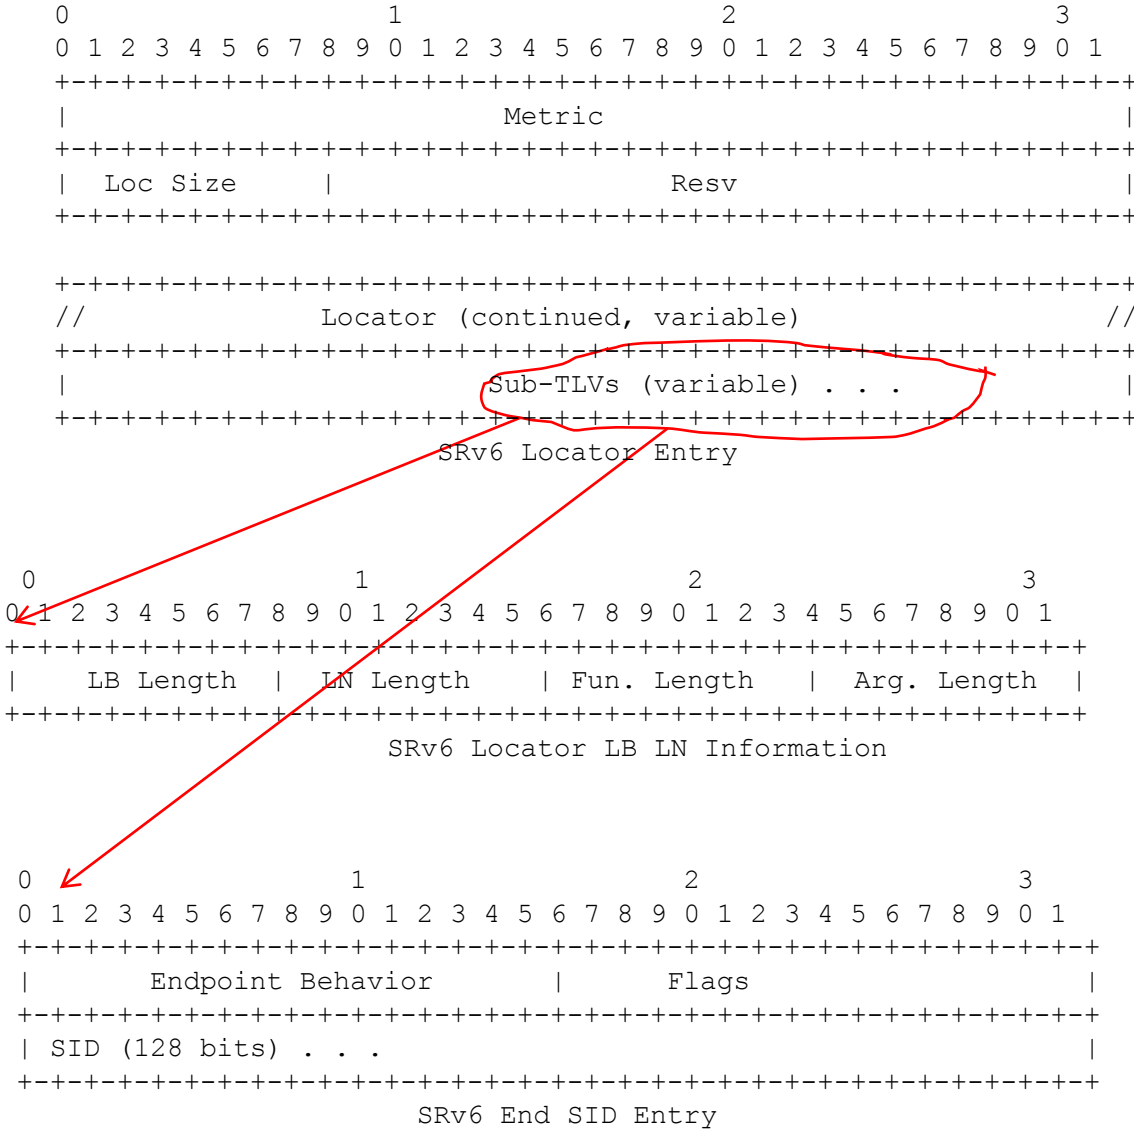
- The treatment of Locator prefixes is the same as that of regular route prefixes.
- Under normal circumstances, **SRv6 Locator Prefix TIEs are flooded only northbound**. However, if a node detects that its default IP prefix covers one or more locator prefixes that are reachable through it but not through one or more other nodes at the same level, then it **MUST** explicitly advertise those SRv6 locators in an South TIE.
- **Positive Disaggregation Locator TIE** and **Negative Disaggregation Locator TIE** are used to advertise Locator prefix-related de-aggregation.



Defines of SRV6 Locator Prefix TIE

A new type of TIE called **SRv6 Locator Prefix TIE** is introduced. The Locator address and End-SID is advertised in Locator TIE.

| | | | |
|--|-----|-----|---------------|
| LocatorTIEType | TBD | 6.1 | This Document |
| PositiveDisaggregation LocatorTIEType | TBD | 6.1 | This Document |
| NegativeDisaggregation LocatorTIEType | TBD | 6.1 | This Document |



Advertise End.X SID

End.X SID is advertised via NeighborsTIEElement in Node TIE

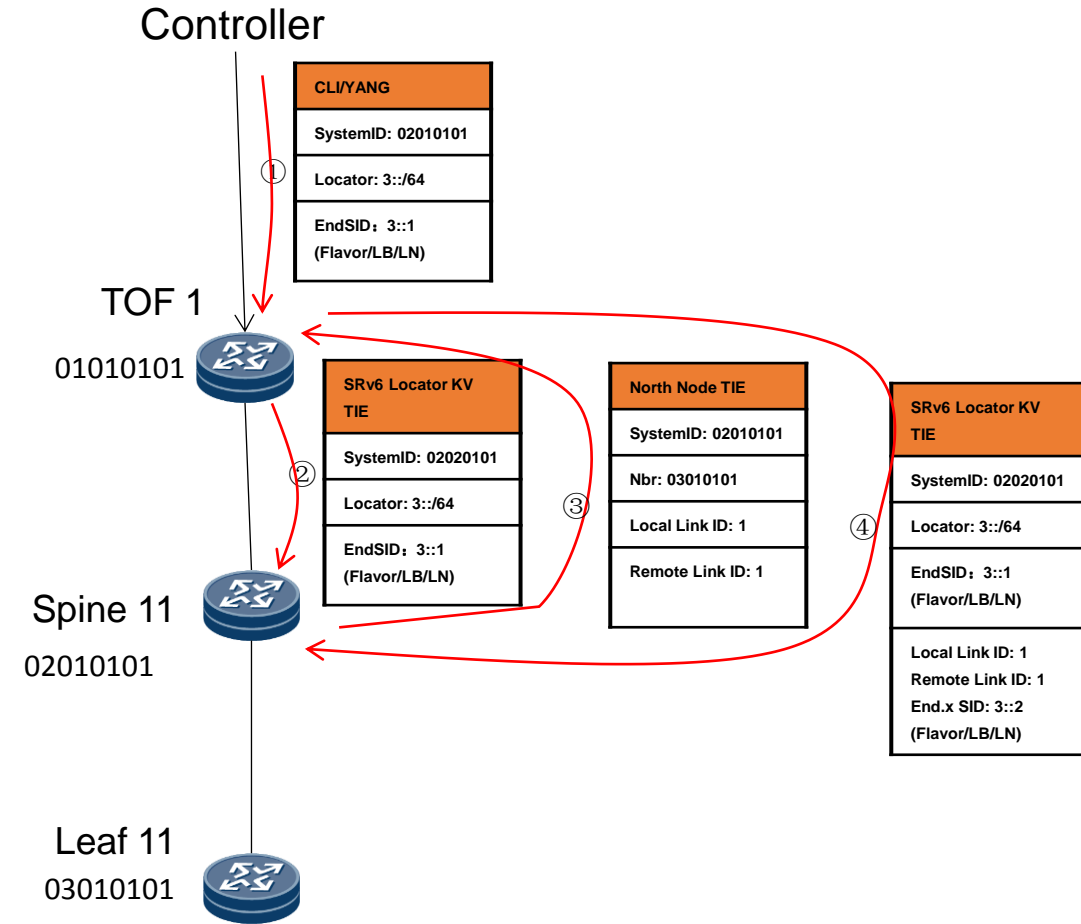
```
Registry RIFT_v6/encoding/NodeNeighborsTIEElement
|link_ids | 4 | 6.1 | Can carry description of multiple |
| | | | parallel links in a TIE |
+-----+-----+-----+-----+
|bandwidth| 5 | 6.1 | Total bandwidth to neighbor as sum |
| | | | of all parallel links |
+-----+-----+-----+-----+
| End.x | TBD | 6.1 | SRv6 End.X SID |
+-----+-----+-----+-----+
```

 Content of End.x

```
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Flags | Algorithm | Weight |
+-----+-----+-----+
| Endpoint Behavior |
+-----+-----+
| SID (128 bits) . . . |
+-----+-----+
| SRv6 SID LB LN Information (optional) |
+-----+-----+
```

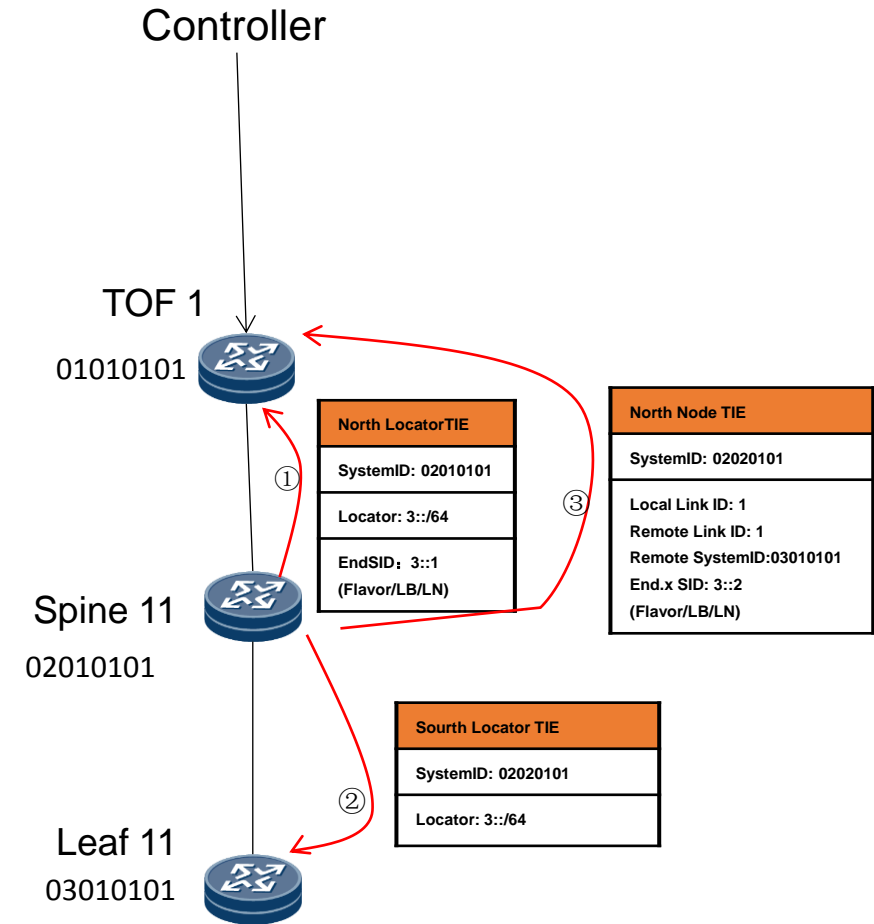
Example #1: SRv6 with ZTP

1. Controller distributes the Locator configuration of Spine11 to TOF1, including the Locator Prefix(3::/64), End SID(3::1).
2. TOF1 generates the SRv6 Locator KV TIE and propagates it to Spine11. Spine11 get the SRv6 Locator configure.
3. Spine11 and Leaf11 establish neighbors, and transmit the neighbor information to TOF1 through the northbound Node TIE. Finally, the neighbor information is passed to the controller.
4. The controller assigns an End.X SID for this neighbor and passes it to Spine11 through the updated SRv6 Locator KV TIE.



Example #2: SRv6 without ZTP

1. Spine11 manually configures SRv6-related settings, including SRv6 Locator Prefix, End SID. These information are propagated to TOF via the North SRv6 Locator TIE.
2. In the southbound SRv6 Locator TIE, only the Locator prefix information is advertised.
3. Spine11 and Leaf11 establish neighbors, then allocate the corresponding End.X SID according to the SRv6 configuration, and advertise the End.X information by inserting it into the neighbor information in the Node SID.



Next Step

- Seeking comments
- Seeking WG adoption after revision

THANKS