

IS-IS Spine-Leaf

IETF 104 Prague

LSR

draft-ietf-isis-spine-leaf-ext-01

March, 2019

Naiming S., Les G., Sanjay T.


DC IS-IS Spine-Leaf

- Incremental protocol improvement
- Reuses well deployed protocol for DC fabric supporting large number of network routing entries and features
- Supports full routing in small/median DC and serves as a DC fabric for overlay routing
- Supports Auto-tier discovery and protocol configuration for Zero Touch Provisioning
- Minimizes IS-IS LSP flooding to leaf nodes in DC spine-leaf topology
- Handles link/node DOWN events in spine-leaf topology to avoid black holes
- Supports other mechanisms/proposals such as Openfabric and Dynamic Flooding

Changes From draft-shen-isis-spine-leaf-ext-07

- Became WG document
- Spine-Leaf TLV simplified and limited to hellos
- New TLV (Leaf-Set TLV) introduced for Circuit Scope LSPs
- (re)introduced limited support for leaf-leaf connections

Spine-Leaf TLV in Hello




- T bit: Tier field is valid
- L: Leaf mode bit – enables reduced flooding
- R: Default Route Gateway bit


blue: sent by leaf nodes

green: sent by spine nodes

Leaf-Set TLV (CS-LSPs)

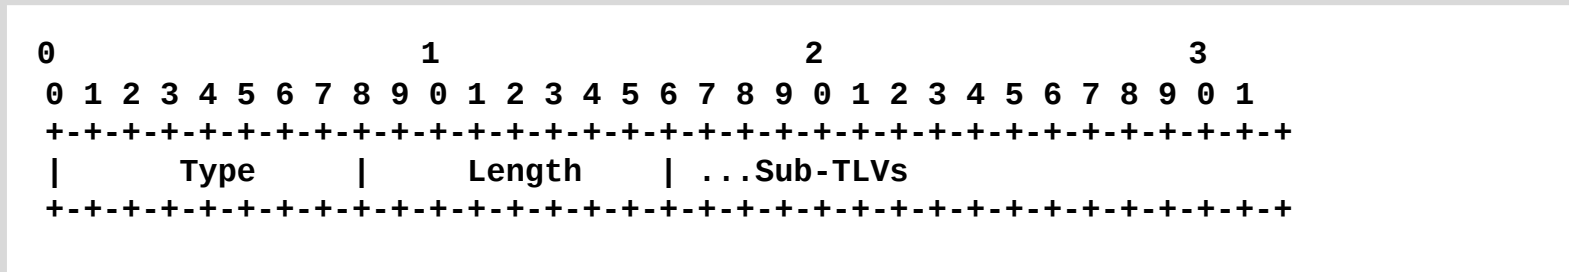


Leaf-Neighbors sub-TLV

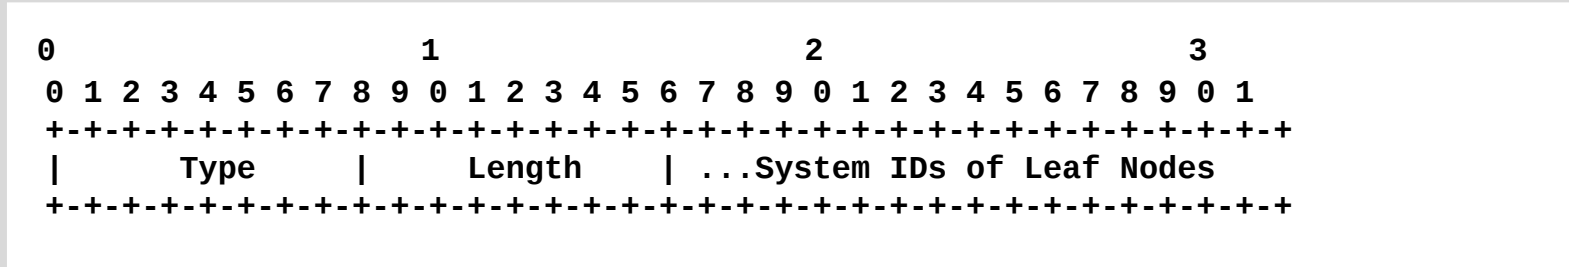


- Sent by Spine Nodes
- Used by Leaf Nodes to detect spine nodes w incomplete reachability to leaf nodes

Leaf-Set TLV (CS-LSPs) (s)




Reachability-Req sub-TLV



- Sent by Leaf Nodes
- Requests Spine Node to send prefix reachability for the specified Leaf Nodes
- Spine Node then sends prefix reachability (TLVs 135/236) in CS-LSPs to the requesting leaf-node


Advertising Connections Between RF Leaf Nodes and Spine Nodes

Link Attribute sub-TLV (RFC 5029)



- Local Protection Available(P) (RFC 5029)
- Excluded from Local Protection(X) (RFC 5029)
- Connect to RF Leaf Node(L) – sent by Spines
 - Signals leaf node can be excluded from flooding
- **Connect to RF Spine(S) – sent by Leaf Nodes**
 - Signals Link is to a Spine Node – possible alternate default path

Use of Leaf-Leaf Connections



- Limited to one leaf neighbor
- Normal flooding
 - Leaf Node LSPs are exchanged as normal and flooded to Spine Nodes
 - Local (to each leaf) Reachability is known
- Local Traffic Only
 - Leaf Node Sets OL bit so spine nodes cannot use the Link for transit
- Transit Traffic Allowed
 - OL bit NOT set
 - High Metric ($2^{24}-2$)
 - Link of last resort – spine nodes can use the link to reach a leaf node which has been disconnected from all spine nodes

Next Steps

Comments welcomed