

BIER with RIFT

Zhaohui Zhang (Juniper)
Shaowen Ma (Juniper)
Zheng Zhang (ZTE)

RIFT

- A hybrid routing protocol for CLOS and Fat Tree networks
 - Link State Routing northbound
 - Distance Vector Routing southbound
 - Default route most of the time
 - Some specific disaggregation routes to avoid black-holing or to provide optimal routing in certain situations

BIER-OSPF Signaling

- BIER sub-TLV with MPLS Encap sub-TLV attached to BIER Prefixes, which are flooded throughout an OSPF area
 - This applies to RIFT northbound
- BIER prefixes re-advertised across area boundary, along with the BIER sub-TLV
 - This is extended to readvertise across IGP boundary in draft-zwzw-bier-prefix-redistribute
 - This applies to RIFT southbound

Non-MPLS Encapsulation

- Non-MPLS Encapsulation is important in Data Center, which is the target of RIFT
- The only real difference between non-MPLS and MPLS encapsulation is the BIFT-ID, which is at the same place of a BIER packet regardless of encap type
 - MPLS: BIFT-ID is a 20-bit label
 - Non-MPLS: BIFT-ID is a 20-bit opaque field

BIFT-ID Signaling

- MPLS: a label block in MPLS Encap sub-TLV with BIER-OSPF/ISIS
- Non-MPLS: not currently signaled in OSPF/ISIS
 - There is one proposal to simply construct BIFT-ID as <SD, BSL, SI>
 - Could be signaled just like MPLS case – BIFT-ID block instead of label block
 - Full advantage of MPLS encapsulation w/o requiring MPLS infrastructure
- All the above can be used for BIER-RIFT

Summary

- BIER-OSPF/ISIS-like signaling for RIFT Northbound
- draft-zwzw-bier-prefix-redistribute method for RIFT southbound
- Similar BIFT-ID signaling for both MPLS and non-MPLS encapsulation
- Thrift schema instead of sub-TLV format

Next Steps

- Seeking Comments