

RIFT Auto-EVPN

draft-ietf-rift-auto-evpn-01

Jordan Head, Tony Przygienda, Wen Lin

IETF112

What's new in draft-ietf-auto-evpn-01?

- Working Group Adopted
- Data Center Interconnect
- ToF-RR Election Procedures
- Increased Collision-Free Scale
- Operational Section
- Minor Restructuring

DCI Gateways

- ToF-RRs can also act as a DCI gateway.
- Advertised via new bit in Node-TIE.
- ToFs with DCI functionality preferred in RR election.

Single Plane Route Reflector Election

- Previous Election Procedures:

1. Highest SID
2. Lowest SID
3. 2nd highest SID
4. 2nd lowest SID

- New Election Procedures:

1. Highest SID w/ DCI-bit
2. Lowest SID w/ DCI-bit
3. 2nd Highest SID w/ DCI-bit
4. Highest SID w/o DCI-bit ## No more DCI/RRs

Multiplane Route Reflector Election

- Declare a Plane ID
 - Highest SID from each Plane
 - Carried in Node TIE
- Election
 1. Highest Plane ID w/ DCI-bit
 2. Lowest Plane ID w/ DCI-bit
 3. 2nd Highest Plane ID w/ DCI-bit
 4. Highest Plane ID w/o DCI-bit ## No more DCI/RRs
- Distributes ToF-RR/DCI across planes for redundancy.

Scale Considerations

- Previous Maximum
 - 3 Fabrics
 - 3 EVIs
 - 15 VLANs
- New Maximum
 - 6 Fabrics
 - 7 EVIs
 - 30 VLANs
- Will scale higher, but collision-free derivation details **MUST** be defined.

Last but not least...

- Added examples to clarify how Auto-EVPN roles align with underlays.
 - For example:
 - **Physical Underlay:** Superspine
 - **RIFT Underlay:** Top-of-Fabric
 - **Auto-EVPN Overlay:** ToF-RR
- Minor section restructuring.
- “*PICS Proforma*” for scale mentioned in previous slide.

What's Next?

- Co-Authorship and comments are welcome.
- More Data Center Interconnect details/examples.
- More multiplane examples.
- More operational considerations.

Questions?

If you think Auto-EVPN is cool, you'll love this...


First, some background.

- We all know that **flat single area** IGPs can come with some pitfalls.
 - **Flooding** – every node needs to know.
 - **State** – every node needs to remember.
 - **Convergence** – every node needs to compute.
- This gets even worse as the network is scaled.
- However, these deployments may be desirable for things like SR.


How are we addressing the problem?

- Flood Reflection is based on existing LSR work.
 - <https://datatracker.ietf.org/doc/html/draft-ietf-lsr-isis-flood-reflection-04>
- Flood Reflectors are a *bit* like BGP Route Reflectors in that we:
 - Designate one or more Flood Reflectors.
 - Designate one or more Flood Reflector Clients.
- Drastically improves scale for flat single area networks.
 - Less state (adjacencies/LSDB).
 - Less flooding.
 - Less SPF computation.

For the visually inclined.



Level 2 Topology



Flood Reflector Topology

What's this got to do with Auto-EVPN?

- Auto-EVPN-like variable derivation applied to Flood Reflection.
- Derive the Cluster ID.
- Derive Role (Client or Reflector).