

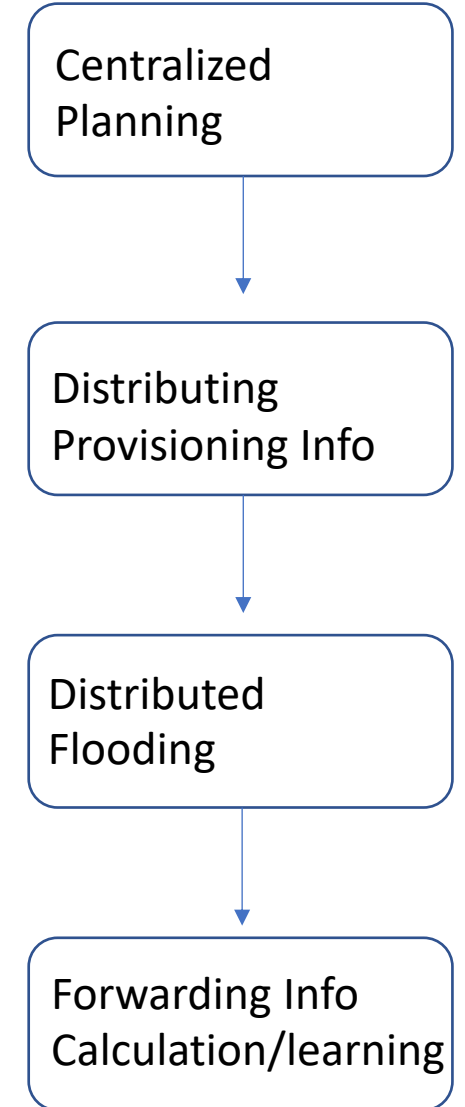
SR with RIFT

Jeffrey Zhang, Jeff Tantsura, Jordan Head, Don Fedyk

IETF 110

SR with OSPF/ISIS/BGP (1/2)

- Each router is provisioned with SRGB and Prefix/Adjacency SIDs, and then floods the provisioned information
 - The provisioning is done to each router after centralized planning
 - Provisioning info sent to each router by:
 - Primitive means of Pen & paper, or
 - Modern Netconf/YANG/other orchestration means



SR with OSPF/ISIS/BGP (2/2)

- Two traffic forwarding models
 - Forwarding based on IGP/BGP routing, using a single prefix SID
 - Replacing LDP with IGP/BGP based label distribution
 - Steering based on a SID-list in packets
 - Route to the segment represented by the top SID
 - then route to the next segment in the list
 - SID-list determined on an ingress node based on:
 - Local calculation, or
 - Programmed by SR controllers

SR with RIFT: SRGB/SID Distribution

- No provisioning of SRGB/SID on non-ToF nodes
 - No flooding from those either
- Result of central planning flooded south by RIFT KV mechanism from ToF nodes
 - ToF nodes learn all the SRGB/SID information
- This combines the two-step process in IGP/BGP SR into a single step
 - Perfectly matching the ZTP principle of RIFT

SR with RIFT: Traffic Forwarding

- Only cares about traffic steering based on SID-list
 - i.e. Label stack (SRv6 not in scope for now)
 - SID-list determined only by SR policies programmed by controllers
- In absence of parallel links or multi-planar topologies
 - Leaves only have northbound default route – so they can't determine SID lists
 - Only cares about Node-SIDs
 - Adjacency-SIDs only makes sense with parallel links between nodes
 - No benefit for non-Node prefix SIDs
- In presence of parallel links or multi-planar topologies
 - Need adj-SID to choose the right link/plane(spine), SRLB used (how do derive)?
 - Will be addressed in the upcoming draft version

Loopback Addresses Preferred

- RIFT does not require loopback addresses in principle
- SR traffic steering is based on routing to node segments in SID-list
 - i.e. must have individual routes to those nodes
 - It's easier to have those routes based on loopback (IPv4/6) addresses
 - Vs. based on System-IDs
- An operator may want to be able SSH into any node using loopback address anyway
- Or should we simply require loopback addresses?
 - Adj-SID to be addressed in the upcoming version of the draft

KV Registry for SR RIFT

| Name | Value | Description |
|------------|-------|---|
| SRIFT Node | TBD | Key-Type describing a SRIFT node |
| TBD2 | | SRIFT Node (System ID, Loopback Address, SRGB Label Base, SRGB Label Range, Node-SID,) |

Next Steps

- Seeking comments
- Seeking WG adoption after revision