



LSRV BGP SPF Applicability IETF-102 Montreal

Keyur Patel, Arrcus
Acee Lindem, Cisco
Shawn Zandi, LinkedIn
Gaurav Dawra, LinkedIn




Status




- Base Applicability draft is now a wg document
 - Many thanks for to wg members for their feedback
- Base Applicability draft went thru 2 version changes ahead of becoming a wg document
- Version 1 changes where presented and discussed at the interim meeting
- Version 2 changes are fairly recent

Version 1 Changes




- Added Sections to Explain
 - Interaction with other BGP AFI/SAFIs
 - Peering models and BFD for liveness support
 - Bi-Connected Graph Heuristic
 - DCI Applicability
 - Non-Clos Applicability
- Security – Pending addition to the draft.

Interaction with Other BGP AFI/SAFIs



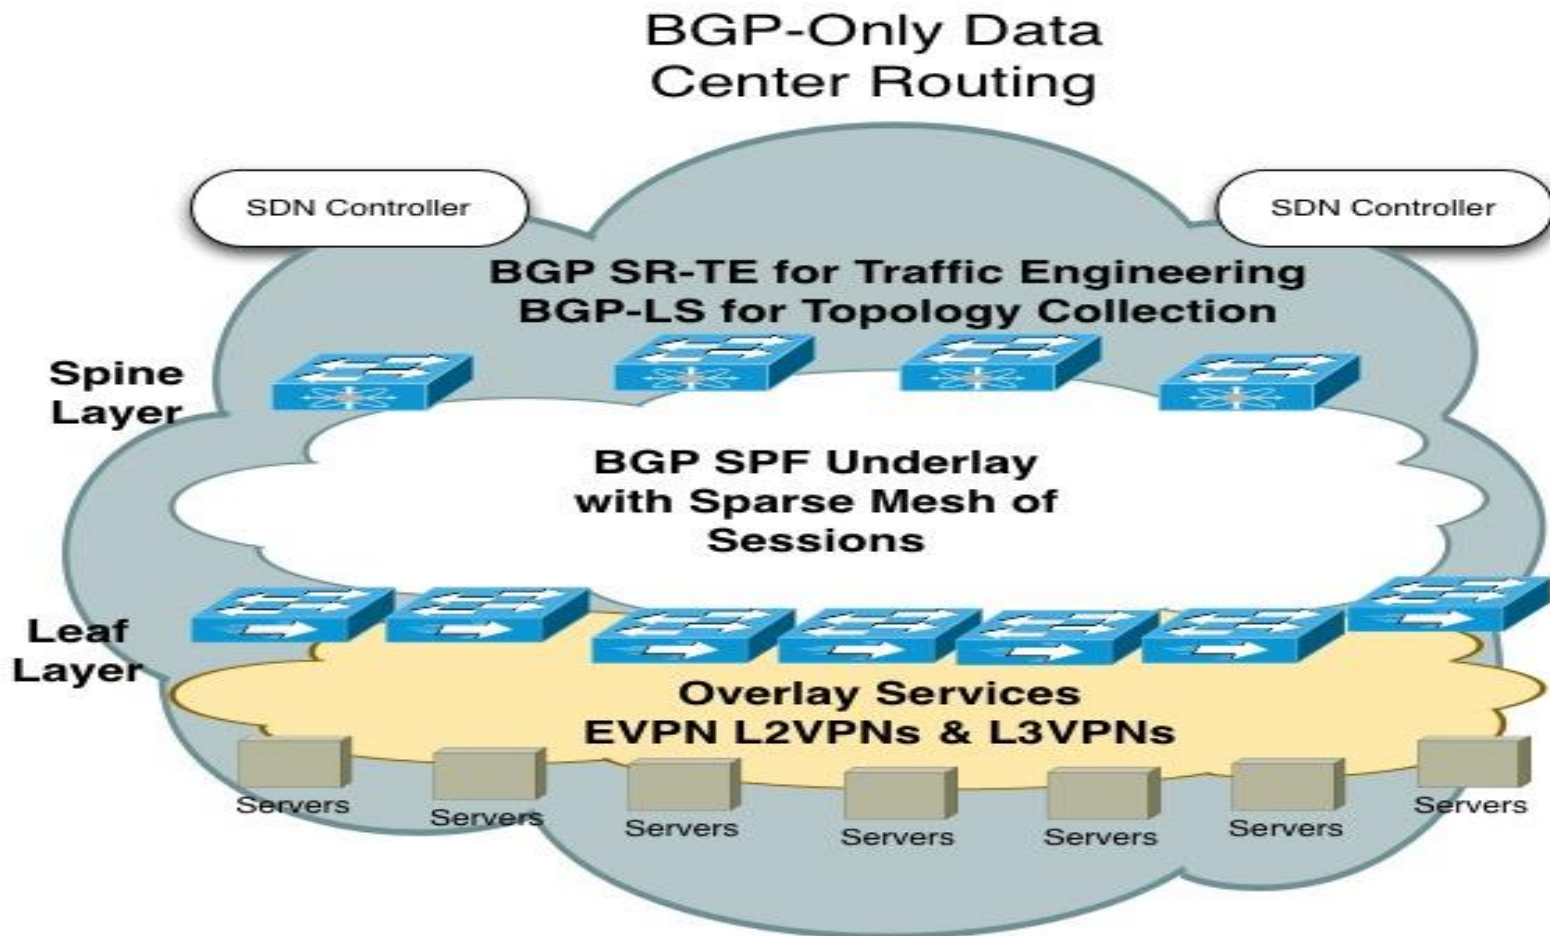
- Intended for underlay where other AFI/SAFIs resolve next-hops using BGP-LS SPF Routes
- Also intended for “flat” data center network
 - RFC 7938 - Use of BGP for Routing in Large-Scale Data Centers
- Interaction with IPv4/IPv6 Unicast
 - Treat as another “Ships in the Night” protocol
 - Recommend BGP-LS SPF routes be given preference
 - No mutual redistribution by default

Interaction with Other BGP AFI/SAFs (Cont'd)




- Interaction with base BGP-LS Address Family
 - BGP-LS SPF Address Family Node, Link, and Prefix NLRI can be used in lieu of base BGP-LS Address when both are required
 - Additional Node, Link, and Prefix NLRI attributes can be piggy-backed on the BGP-LS SPF Address family NLRI

BGP-Only Data Center Routing




Sparse BGP Peering and BFD Peering

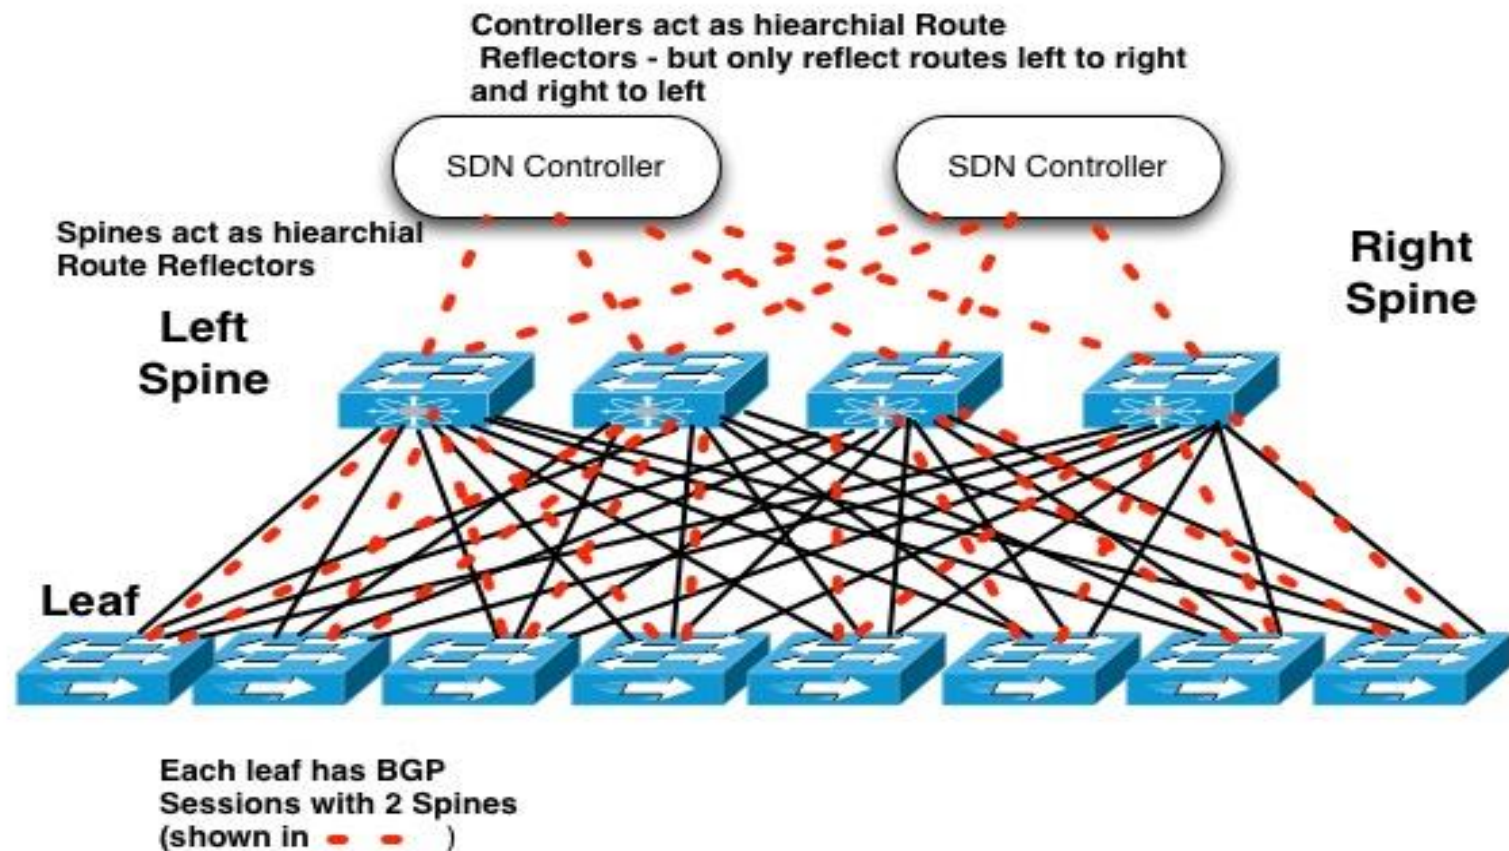


- Liveness detection for links done outside of BGP (i.e., based solely on link status or using BFD)
- Leaves peer with subset of spines (e.g., only 2 to offer redundancy)
 - Spines act as Route Reflector
 - Savings in sessions depends on the number of spines to which leaves are connected
 - Redundancy trade-off versus copies of advertisements
- Spines peer with controllers
 - Controllers reflect between spines that peer with a unique set of leaves


BGP SPF Data Center Sparse Peering Example



BGP SPF Data Center Topology




Bi-Connected Peer Heuristic



- Dependent on BGP routers in fabric knowing southbound (toward servers) and northbound (toward spine) ports
 - Most likely provided by the discovery protocol
- BGP Routers accept connections passively on southbound ports.
- BGP Routers choose a subset of northbound connections (usually 2) to provide “enough” redundancy
 - Selection of northbound sessions is local matter
 - Could use hash or spines with fewest BGP sessions
- BGP Routers attempt to maintain “enough” northbound connections


BGP SPF Bi-Connected Heuristic



BGP SPF Session Heuristic

Spines Passively Accept Connections from Southbound Leaves

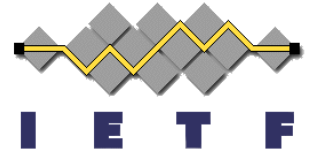
Spines




Leaves

Leaves Determine with which Spines they actively Establish Northbound BGP Sessions

DCI and Non-Data Center Applicability




- Data Center Interconnect (DCI) Interoperability
 - In general, it is expected that individual data centers will act as separate BGP-LS SPF domains
 - Initially, no intention to cover DCI for BGP-LS SPF
- Non-Data Center Applicability
 - BGP-LS SPF could be applicable to other use cases including Service Provider (SP) backbone underlays.
 - Dependent on how successful we are with the standards and, more importantly, the implementations.



BGP SPF Security


- Really no different than classic BGP underlay security
 - Simple for both full and sparse peering
 - Tolerance required for alternate sparse peering model
- Use of TTL security on intra-fabric BGP sessions (RFC 5082)
- If BGP fabric is not isolated, recommend control plane protection as well (RFC 6192)
- If BGP fabric may be subverted, TCP-AO (RFC 5925) is recommended (MD5 - RFC 2385 if unavailable)
 - Keys should support key-chain rollover via the YANG model as described in RFC 8177 and be changed periodic or when there is potential for a breach.

Version 2 Changes




- Added Sections to explain BGP Peer Discovery requirements
 - Laid out high level discovery parameters for BGP peer discovery
 - Discovery requirement specifically targets CLOS based networks using BGP but doesn't limit the solution to such networks

BGP Discovery Mechanisms



- BGP Peer Discover using LLDP - draft-acee-idr-lldp-peer-discovery-02
 - IEEE Protocol used today for layer 2 discovery
 - Somewhat limited given based on single Protocol Data Unit
- BGP Peer Discovery - draft-xu-idr-neighbor-autodiscovery-06
 - LDP-like discovery using multicast UDP
 - Part of BGP protocol as option
- Link State Over Ethernet - draft-ymbk-lsvr-lsoe-00
 - New protocol – information could be used for other purposes
- Should use "at least" one of the above.
- Where is work done? Link-State over Ethernet could/should be done in LSVR.

BGP-LS SP Applicability Discussion Points



- What more needs to go into an applicability draft?
- Routing Policy – discussion of aggregation policies and what of an implementation should provide.
- Partitioning of BGP-LS SPF Domains?