



IETF 101 – London
Mar 2018
IDR Working Group

draft-ketant-idr-bgp-ls-bgp-only-fabric-00

Ketan Talaulikar (ketant@cisco.com)

Clarence Filsfils (cf@cisco.com)

Krishna Swamy (kriswamy@cisco.com)

Shawn Zandi (szandi@linkedin.com)

Gaurav Dawra (gdawra.ietf@gmail.com)

What does this draft propose?

- BGP-LS specified originally via RFC 7752 for “north-bound distribution of link-state and TE information”
- BGP-LS focus has been mostly on distribution of information from IGPs
- draft-ietf-idr-bgpls-segment-routing-epe introduced distribution of information from BGP for Segment Routing Egress Peer Engineering

This draft specifies the mechanism for distribution of network topology and TE information in a network running BGP as the only routing protocol

Why is this required?

- Several DCs/MSDCs [RFC7938] run BGP as their only routing protocol today
- In such networks, controllers do not have the underlying link topology and TE information view (as available with IGPs via BGP-LS)
- Having such a topology enables TE use-cases in BGP-only networks using Segment Routing [draft-ietf-spring-segment-routing-msdc]

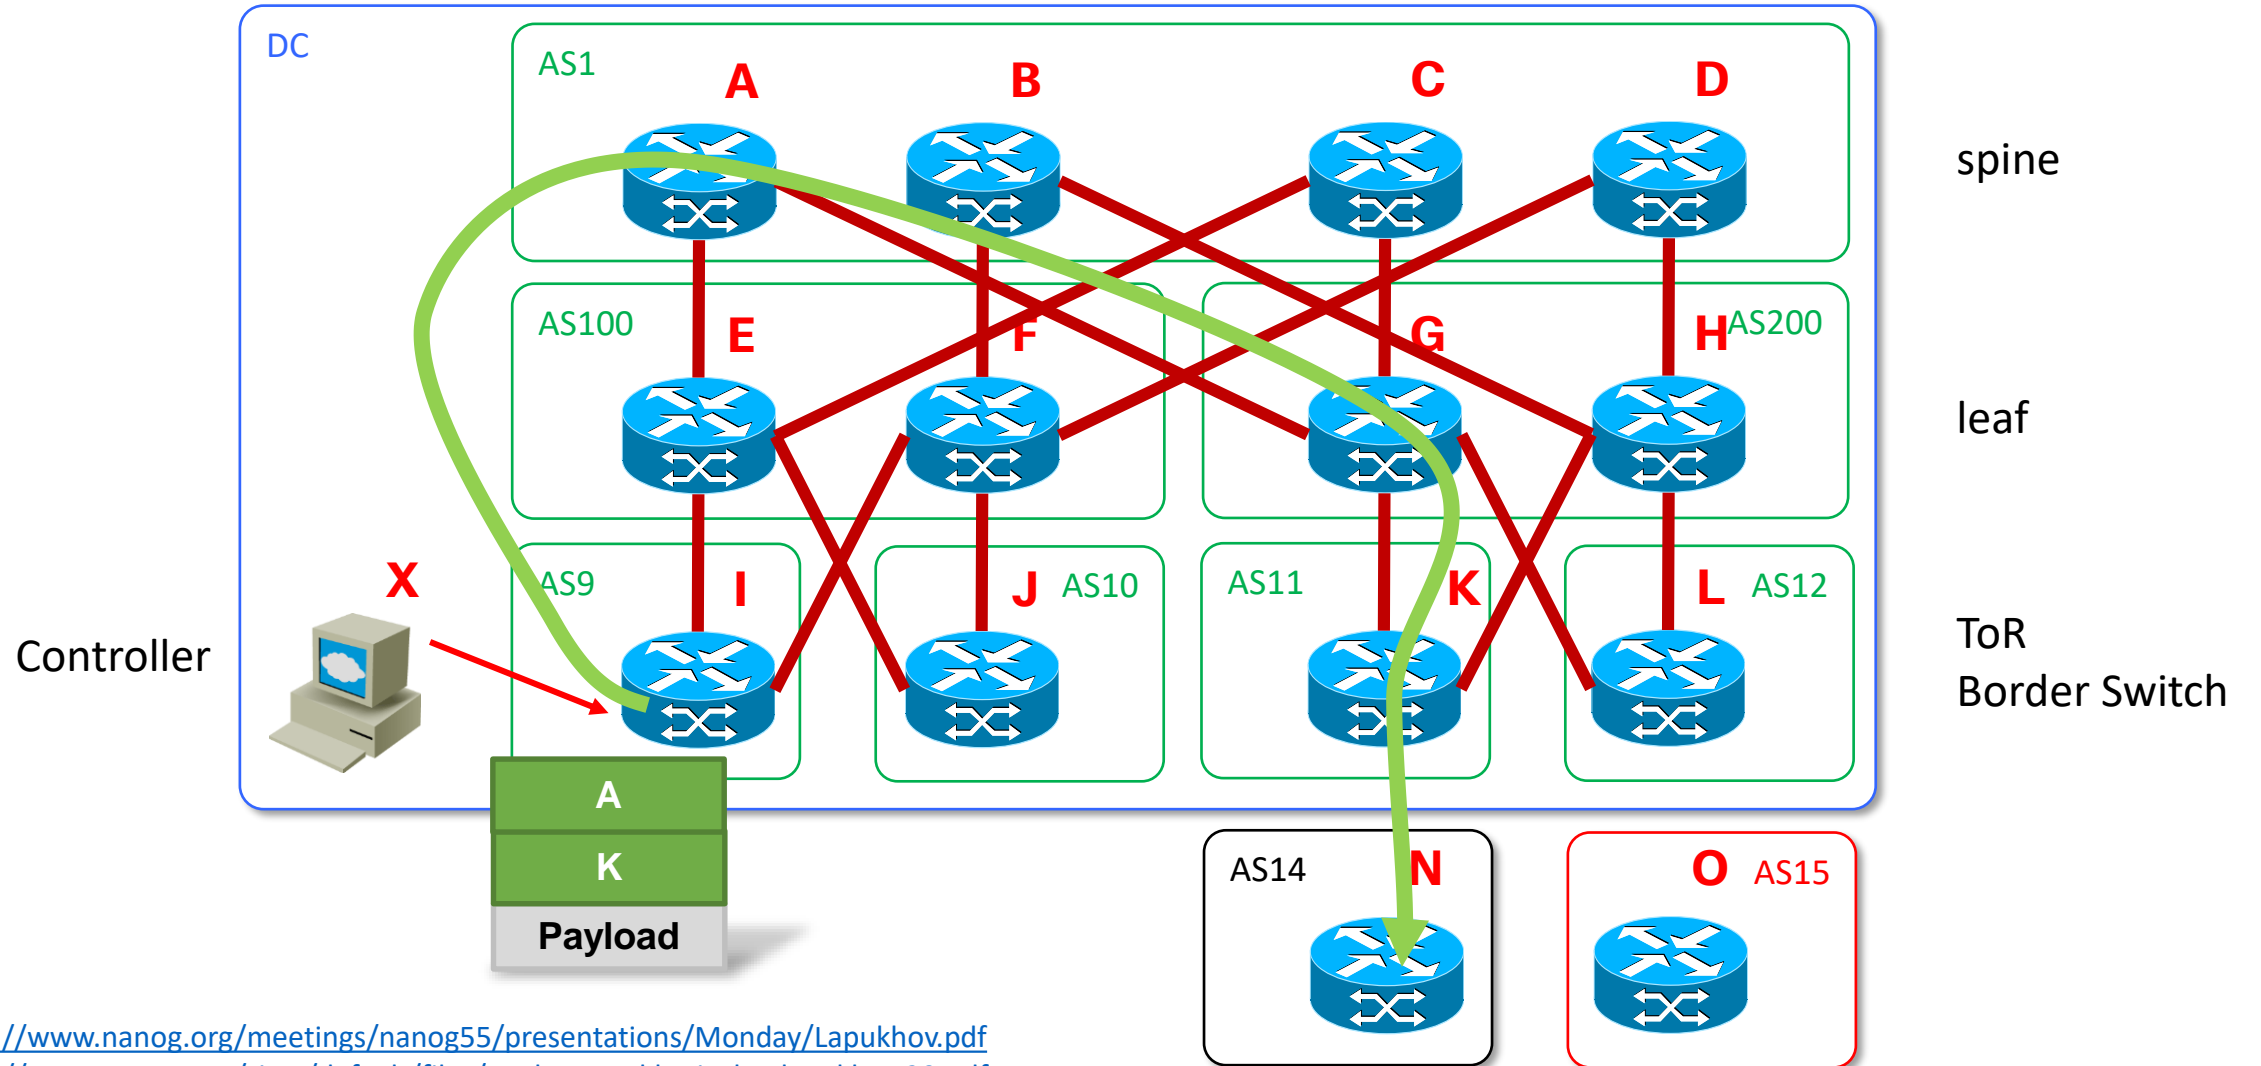
This draft describes the use-case where Segment Routing TE paths can be setup in networks running only BGP by using this topology information

How is this done?

- BGP-LS is enabled on each BGP node to advertise the local topology of that node via different NLRI types (node, links, prefixes and TE policies)
- BGP-LS descriptors to be used for the NLRI types are specified (no new additions)
- BGP-LS attributes for the NLRI types are specified (no new additions as yet)
- Procedures for the BGP-LS advertisements are specified

SR-TE in BGP-only Networks

— eBGP ipv4 labeled-unicast with SR
— BGP SR-TE



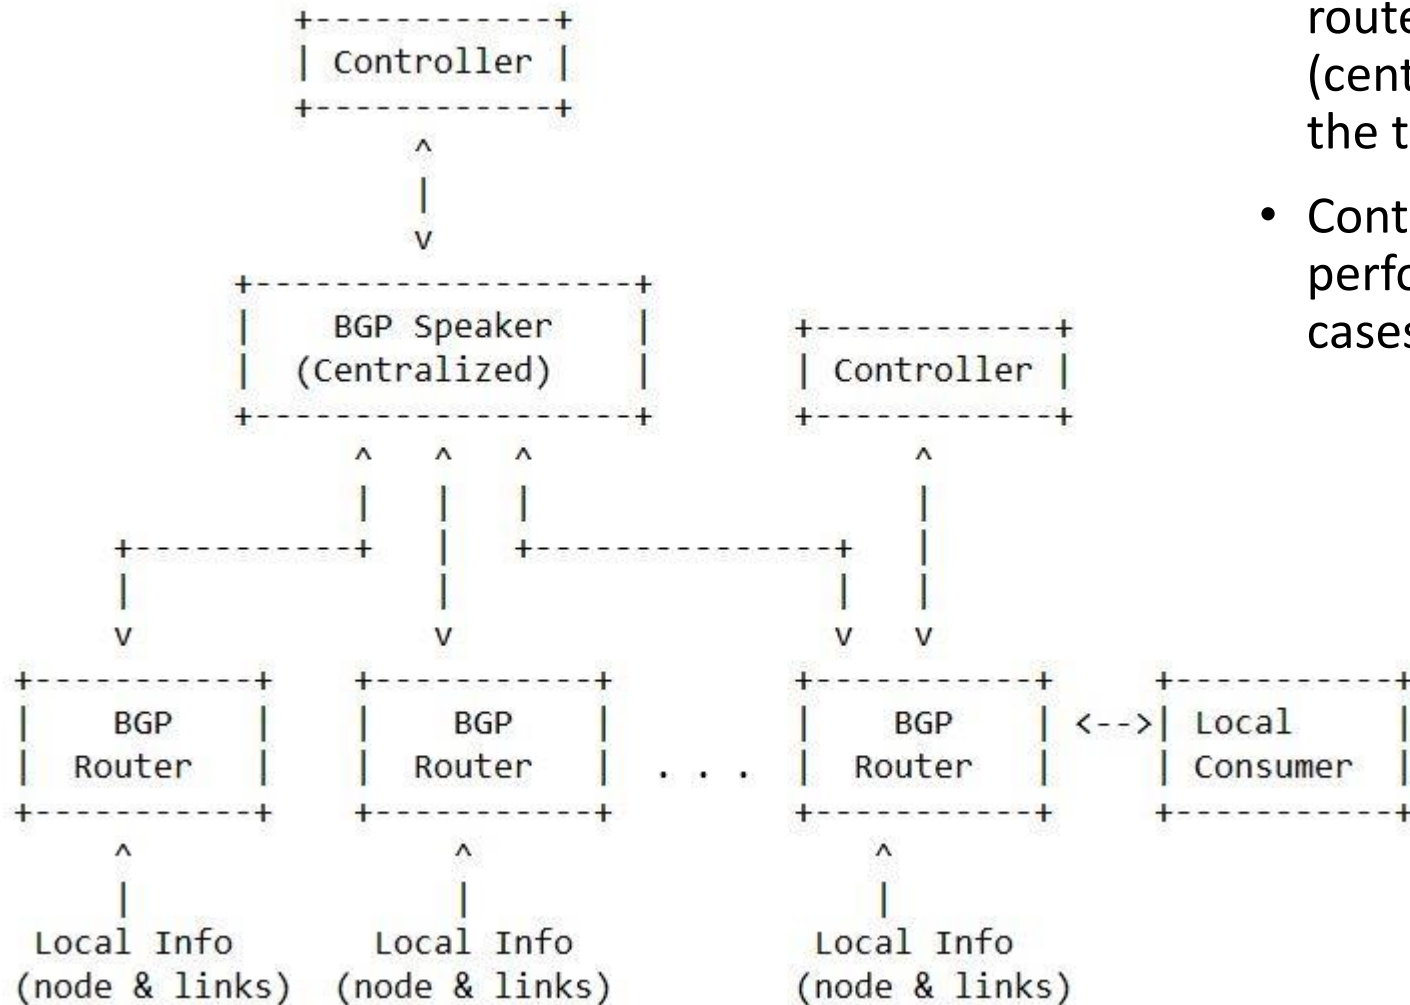
<https://www.nanog.org/meetings/nanog55/presentations/Monday/Lapukhov.pdf>
<https://www.nanog.org/sites/default/files/wed.general.brainslug.lapukhov.20.pdf>

What is NOT changed ...

- Base BGP underlay using IPv4/IPv6/LU/SR is unchanged and follows the BGP decision process
- Existing BGP Peering models are used (current draft version focuses on eBGP single hop as in RFC7938; other models are not precluded)
- BGP-LS sessions to controller(s) may be formed using the reachability provided by this underlay routing

In the Segment Routing TE use-case, controller provides specific TE paths for the overlay services that is instantiated via SR Policies on specific BGP head-end nodes

Topology Collection



- BGP-LS Sessions are setup from the BGP routers to the controller(s) or via one or more (centralized) BGP speakers which consolidate the topology feed from the network
- Controller has the network topology and perform centralized computation for use-cases like TE

- BGP Routers run IPv4/IPv4/LU sessions in the underlay and build their forwarding with the usual BGP decision process

Topology Information

- Node Level Information
 - Node Name, SRGB, SRLB, Node MSDs
- Link Level Information
 - Link Name, Bandwidth, Adjacency/Peering SIDs, Link MSDs
 - Affinities, Delay & Loss metrics, TE metric
- Prefix Information
 - Node SID
- SR-TE Policy Information
 - SR Policies along with their BSID & Segment List instantiated on the node

Next Steps ...

- Solicit WG review and comments/inputs/feedback