

Egress Peer Engineering using BGP-LU

draft-gredler-idr-bgplu-epe-01

Hannes Gredler hannes@juniper.net

Balaji Rajagopalan balajir@juniper.net

Kaliraj Vairavakkalai kaliraj@juniper.net

Chandra Ramachandran csekar@juniper.net

Luyuan Fang lufang@microsoft.com

Egress TE – Feedback Cycle



The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

Rationale

- Content Providers want to balance their outgoing load *without* hacking BGP import policy at ASBRs
- Deliver Interface stats using *PUSH* (rather than PULL) model
- Use Egress assigned, *1-hop MPLS LSPs* (strict ERO) MPLS LSPs for traffic path control
- This can be done using *existing* Protocols
 - Gang of four

Gang of Four

- ***RFC 3107***
 - Advertising the egress assigned 1-hop LSP
- ***Best-external***
 - Path diversity, FRR
- ***Add-Path***
 - Path diversity, FRR
- ***Link-Bandwidth community***
 - Report dynamic per-link utilization

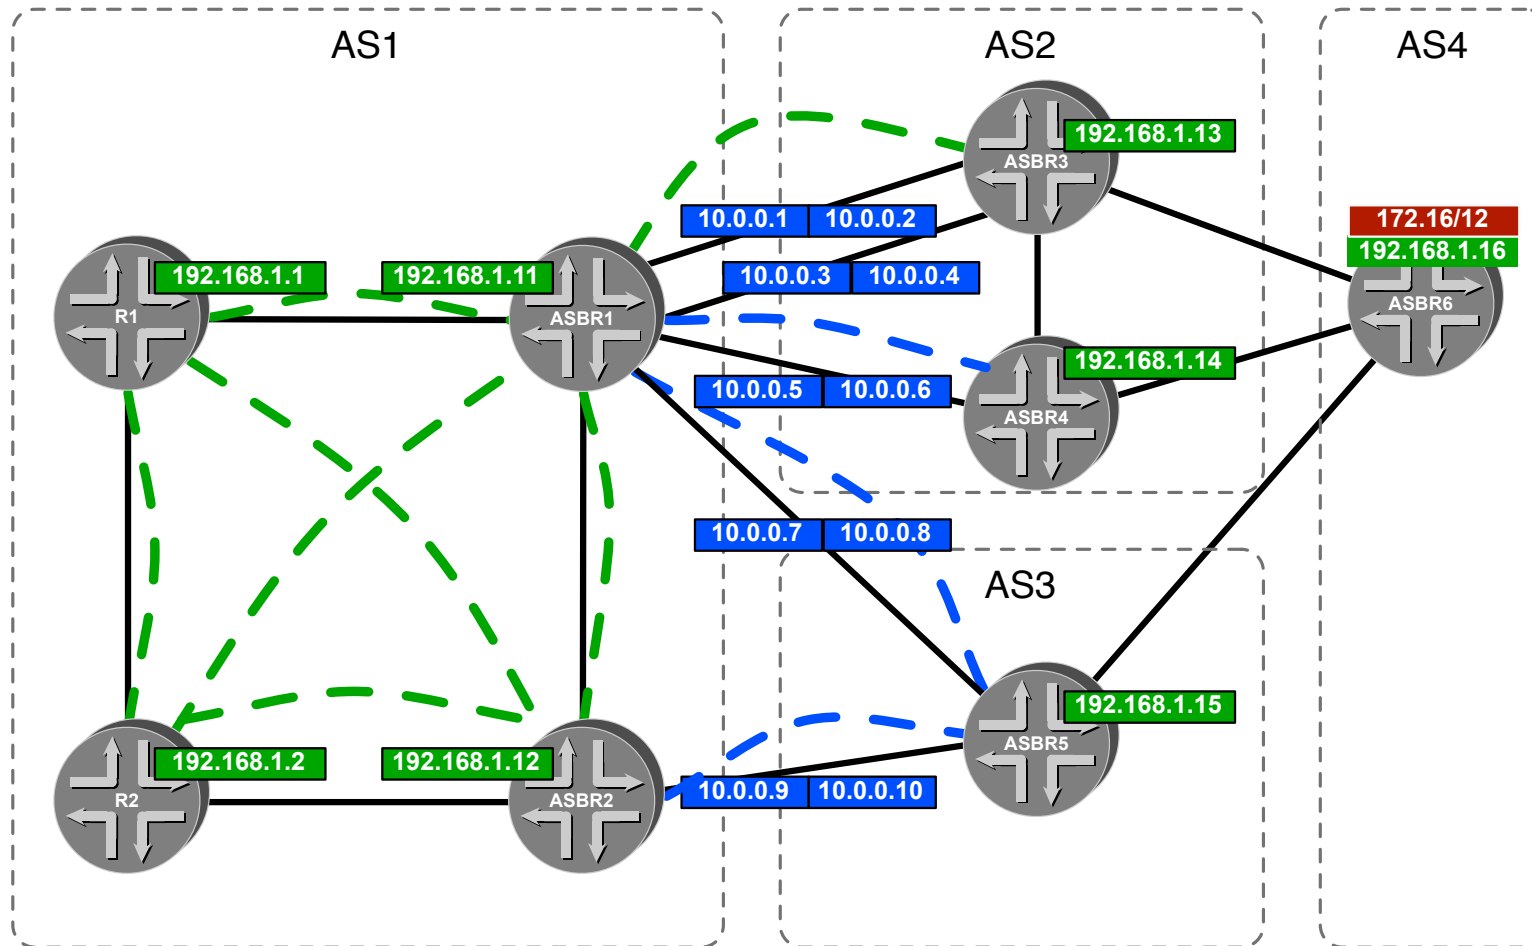
Link-Bandwidth community

- Implemented as
 - “Static BW of an inter-AS link”
 - Affects load-balancing of ECMP paths
- Used here as
 - “Available BW (as seen from the data plane) of an inter-AS link”
 - **Does not break** above assumptions
 - The higher the (available) bandwidth, the higher the traffic portion,

Fast Re-route

- Backup path candidates
 - Local Links
 - Remote non-best paths
 - Re-use some of the machinery described in *[draft-minto-2547-egress-node-fast-protection-03](#)*
 - Local IP forwarding tables

Sample Topology in the I.D.



Next Steps

- Add support for “Aggregate Labels” = “PeerSet”
- Add support for BMP as “Controller Service Protocol”
- Describe IPv6 (dual-stack) operation
- Questions ?
- Please provide Feedback on the list
 - especially if you’re a content-provider ;-)