

Interconnection of Segment Routing Domains Problem Statement and Solution Landscape

draft-farrel-spring-sr-domain-interconnect-01

Adrian Farrel <afarrel@juniper.net>

John Drake <jdrake@juniper.net>


IETF-100, Singapore, November 2017

Why Did We Write This Document?

- We have a document already adopted in BESS
 - draft-ietf-bess-datacenter-gateway
 - “Gateway Auto-Discovery and Route Advertisement for Segment Routing Enabled Domain Interconnection”
 - Short document to define a new "BGP Tunnel Encapsulation Tunnel Type" for "SR Tunnel"
 - Describes how domain gateways can advertise so that:
 - Path choice is available across a multi-homed or multi-AS backbone
 - End-to-end SR paths (and SID stacks) can be computed over SR and non-SR cores
 - (despite the filename, this is *not* just about DC sites)
 - However, that document does not say “why” or give the full explanation of all the building blocks
- So we wrote this informational document to explain
 - The problems space
 - How many pre-existing protocol elements fit together
 - How to use the small protocol extensions in the BESS draft

Reference Model

- Two SR domains
- Want to steer end-to-end traffic
- Backbone networks might or might not be SR capable



Issues and Challenges


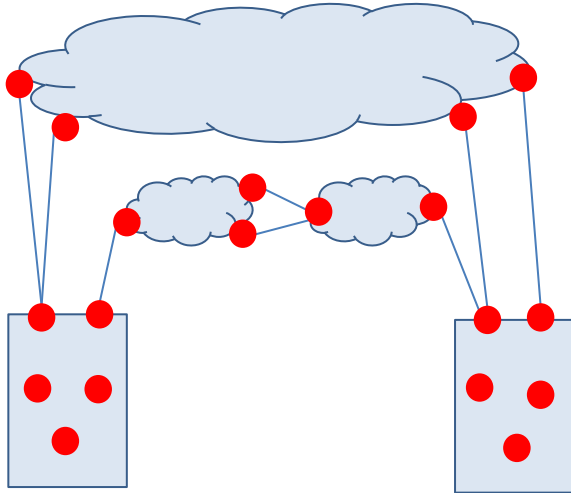
- Gateway selection
 - How to make all gateways to a site visible
- How to select paths (tunnels) over the backbone
 - Don't need to see full paths
 - Do need to see options (reachability)
 - Might or might not be SR-enabled backbone
- How to identify tunnels in the backbone
- How to provide “enough” visibility to a controller

Available Building Blocks

- Learning about the network
 - BGP-LS
 - RFC 7752
 - draft-ietf-idr-bgp-ls-segment-routing-ext
 - Abstraction of topology
 - RFC 7926
- Create tunnels
 - draft-ietf-pce-pce-initiated-lsp
 - draft-ietf-pce-segment-routing
- Bind labels to prefixes
 - draft-ietf-mpls-rfc3107bis
 - draft-ietf-idr-bgp-prefix-sid
- Assign policies for placing traffic on SR paths and advertise them
 - draft-filsfils-spring-segment-routing-policy
 - draft-ietf-idr-segment-routing-te-policy
- Which tunnel type to use to reach a prefix
 - draft-ietf-idr-tunnel-encaps
- Discover gateways and propagate all gateways through all ASBRs
 - draft-ietf-bess-datacenter-gateway

Worked Example: Step 1

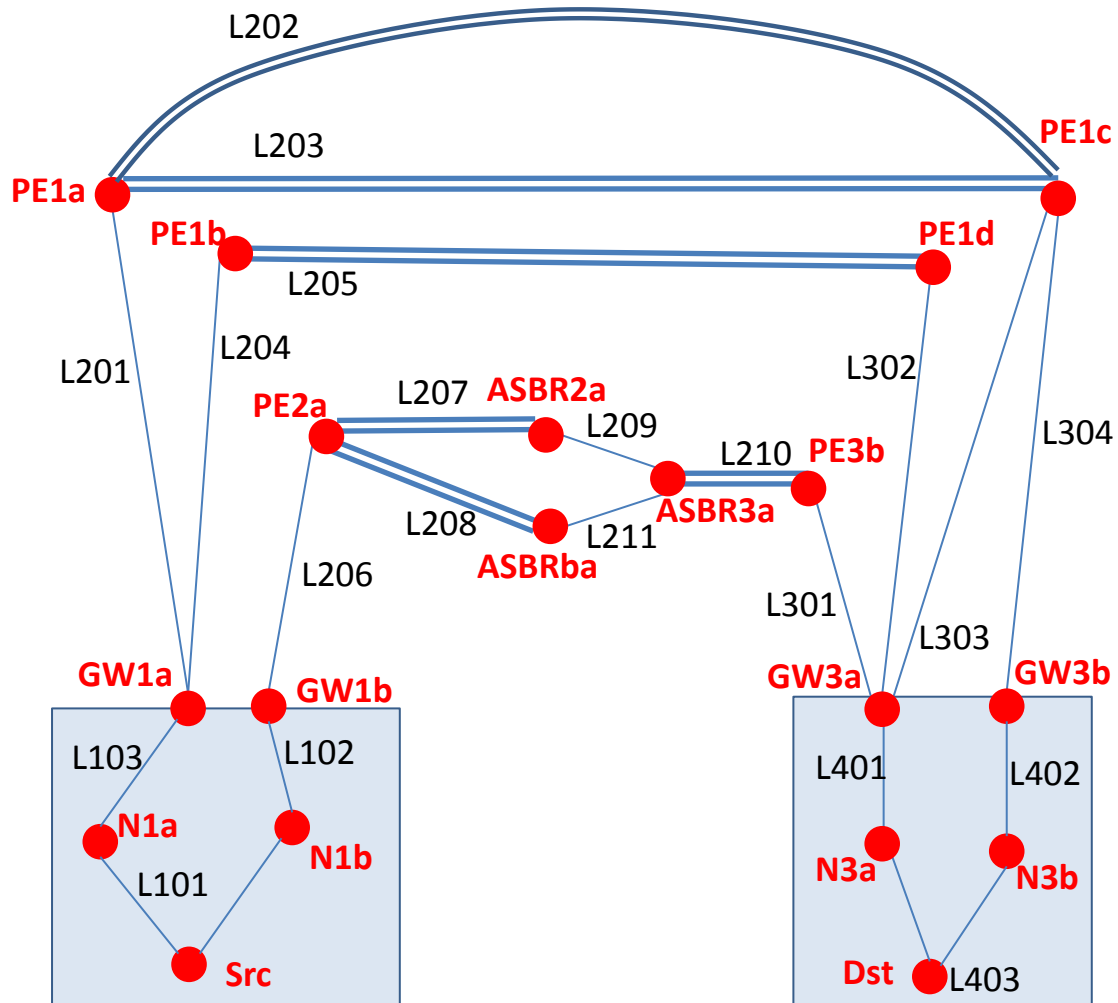
Abstraction, Build Tunnels, and Assign SIDs



Step 2

Select a Path and Build a SID Stack

- Selected path
Src-N1a-GW1a-PE1a-
(RSVP-TE LSP)-PE1c-
GW3b-N3b-Dst
- Fully explicit label stack
L103, L201, L202, L304,
L402, L403
- Egress domain doesn't
share topology
L103, L201, L202, L304,
L403
- Source domain only has
reachability info
L103, L304, L403



Alternate Step 2

Select a Path and Build a SID Stack

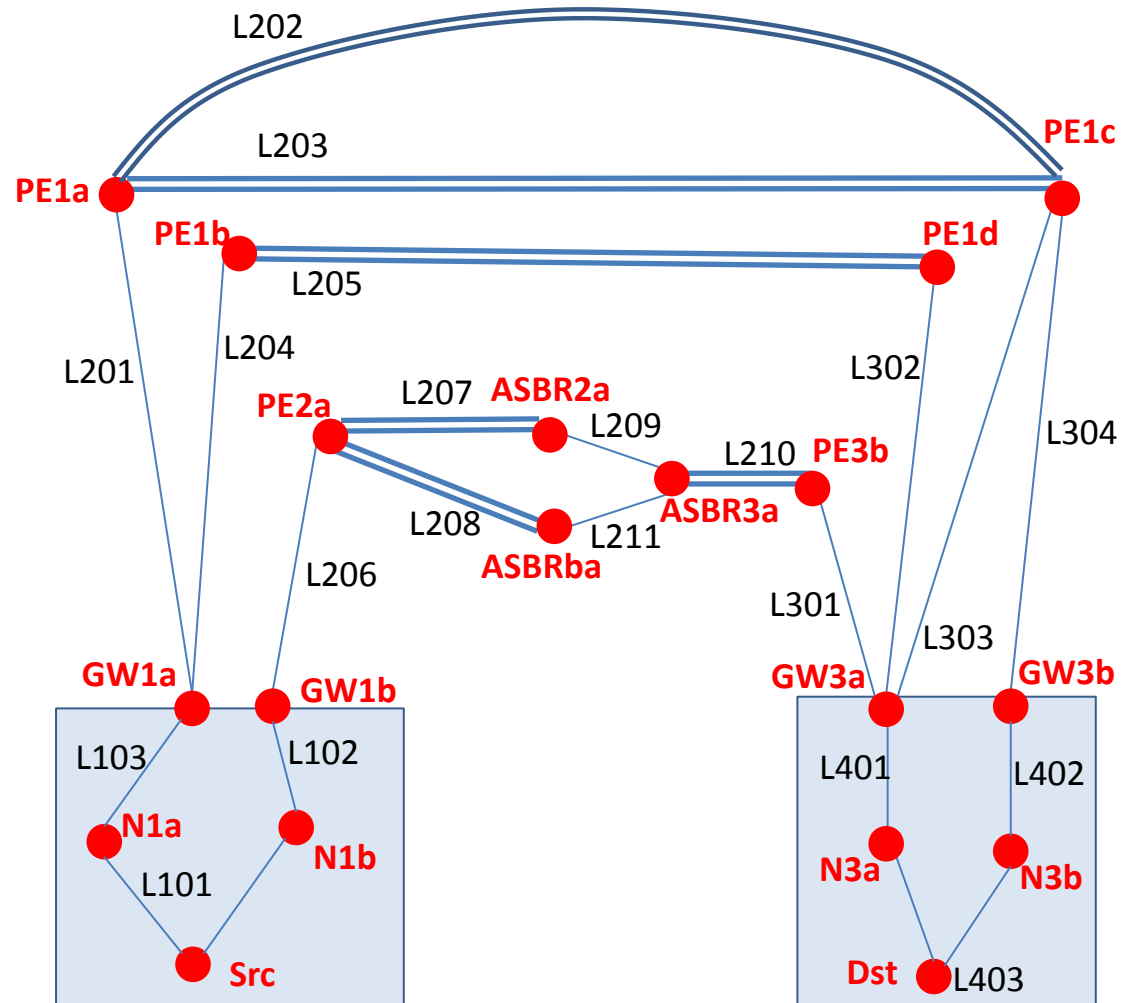
- Selected path spans ASes

Src-N1b-GW1b-PE2a-(RSVP-TE LSP)-ASBR2a-ASBR3a-(RSVP-TE LSP)-PE3b-GW3a-N3a-Dst

- Fully explicit label stack

L102, L206, L207, L209, L210, L301, L401, L403

- Source domain only has reachability info
L102, L301, L403



Next Steps

- Is it helpful to have this type of document?
 - Some people say “It’s just a white paper. Publish it elsewhere.”
 - Some people say “This is really helpful to understand how all of the pieces of IETF work fit together for a real deployment.”
- The authors plan to keep this draft alive as a framework for discussion
 - It seems to us (the authors) that SPRING is a good place to have those discussions