

Gateway Auto-Discovery and Route Advertisement for Segment Routing Enabled Domain Interconnection

draft-drake-bess-datacenter-gateway-04

John Drake <jdrake@juniper.net>

Adrian Farrel <afarrel@juniper.net>

Eric Rosen <erosen@juniper.net>

Keyur Patel <keyur@arccus.com>

Luay Jalil <luay.jalil@verizon.com>


IETF-99, Prague, July 2017

Document History and Status

- Originally written for connecting “segment routing enabled data center sites”
- This revision generalizes the concept to interconnect and SR domains
- The work is stable and referenced by draft-farrel-spring-sr-domain-interconnect
 - This document provides protocol features
 - That document combines those features and other protocol components to describe a full solution

Reference Model

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



Gateway Discovery

- All GWs advertize themselves as having routes to all prefixes in the domain
- Each GW is configured with an identifier for the SR domain
 - Common across all GWs to the domain
 - Unique across all SR domains that are connected
- A route target [RFC4360] is attached to each GW's auto-discovery route and has its value set to the SR domain identifier
- Each GW constructs an import filtering rule to import any route that carries a route target with the same SR domain identifier that it uses
 - Only GWs to the domain will import those routes
 - All GWs to the same SR domain will import each other's routes and will learn (auto-discover) the current set of active GWs for the SR domain

Tunnel Usage

- When a GW advertizes it includes a Tunnel Encapsulation attribute [draft-ietf-idr-tunnel-encaps]
 - Contains an SR Tunnel TLV (a new TLV)
 - One per tunnel to the gateway
 - Includes a Remote Endpoint sub-TLV
 - Can contain a Prefix SID sub-TLV or an MPLS label stack sub-TLV to help build an end-to-end SID stack or to swap out a binding SID
- Tunnel head ends (i.e., remote GWs) can assign binding SIDs to identify the tunnels
 - But that is out of scope for this document

And the magic is...

- When a GW to a domain discovers other GWs to the same domain
 - It also adds those SR Tunnel TLVs to its advertisement
- This means that even when a GW is accessed through multiple ASBRs the full choice of access points is known
 - Combined with BGP-LS and EPE this gives full end-to-end traffic engineering

What Next

- Document still needs
 - Management and security sections
- Planning FCFS codepoint request as implementation progresses
- Would like more reviews
 - Especially in the context of draft-farrel-spring-sr-domain-interconnect
- It's a simple draft, could we consider adopting it?