Inter-domain Traffic Engineering

Retaining System Autonomy With LSP Hierarchies

Oxymoron

- A phrase in which two contradictory terms are combined
  - A deafening silence
  - A mournful optimist
  - A tragic comedy
- To some degree, the term “inter-domain traffic engineering” is also an oxymoron

Inter-Domain Traffic Engineering

- Traffic engineering a flow that originates on one routing domain and terminates in another
- Motivation for separating routing domains
  - Scaling
  - Security considerations
  - Administrative considerations
  - Fate sharing

Decomposition

- In order to understand inter-domain traffic engineering, we must define the following terms
  - Traffic engineering
  - Routing domain

Traffic Engineering

- An aspect of network engineering concerned with the performance optimization
- Typically, traffic engineering tools redirect selected flows from the IGP shortest path onto an alternative path. Redirecting those flows diverts traffic from congested interfaces and distributes traffic across available interfaces
- Most traffic engineering tools are extensions of the routing subsystem
  - RSVP-TE, IGP extensions for traffic engineering

Routing Domain

- A group of routers that enforces a common routing and forwarding policy
  - Determines which subnets are accessible to which gateway interfaces
  - Determines what forwarding behavior is afforded to each flow
  - Determines what path a flow will take through the routing domain
- Typically, all routers within a routing domain are controlled by a single administrative authority

Routing Domain Autonomy

- A peer network can offer traffic, but it cannot determine the following:
  - Whether the traffic will be delivered
  - The QoS that will be afforded to that traffic
  - The route that traffic will traverse within the autonomous routing domain
- Except by mutual agreement
  - For example, using BGP MEDs

Tension Between Inter-domain Traffic Engineering and Routing Domain Autonomy

- Peer networks frequently want to optimize end-to-end performance for flows that originate in one network and terminate in another
- Peer networks rarely want to compromise their autonomy to the degree required to accomplish this goal
  - Don’t want to share IGP information
  - Don’t want to merge RSVP signaling domains

LSP Hierarchy Can Resolve This Conflict

- Each peer network configures MPLS LSPs connecting its Autonomous System Border Routers to its other edge routers
  - PSN LSPs
- Each peer network advertises the PSN LSP as a forwarding adjacency into its IGP
  - Possibly coloring the forwarding adjacency to indicated that it is a PSN LSP

LSP Hierarchy (continued)

- Peer network uses RSVP-TE to signal end-to-end LSP
  - Possibly constraining end-to-end LSP to traverse only links of a specific color
  - Possibly constraining end-to-end LSP with other QoS requirements
- RSVP-TE views each PSN LSP as a hop contain within the end-to-end LSP
- On the forwarding plane, end-to-end LSPs are nested within PSN LSPs

Reference Model

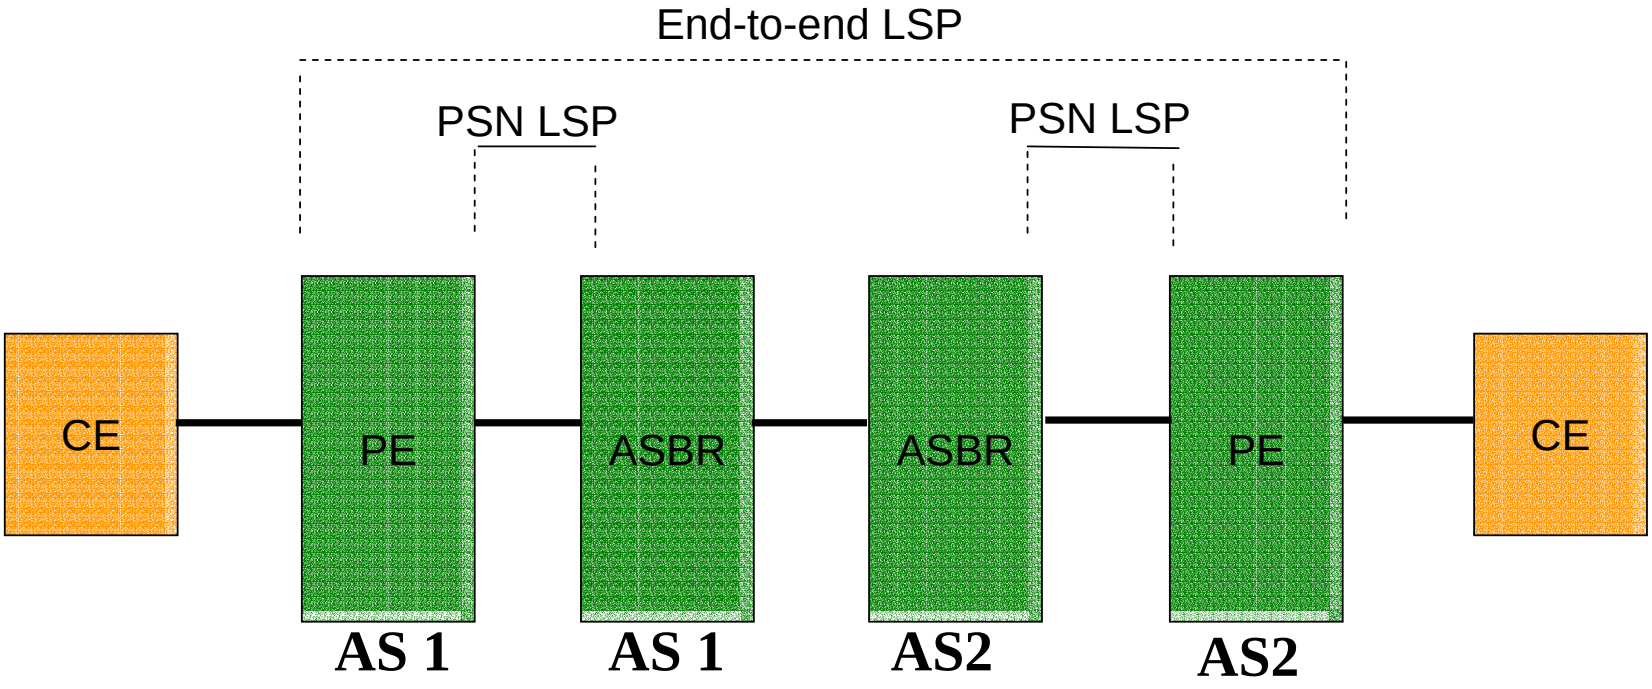

Inter-Domain TE is Achieved

- RSVP-TE used to signal end-to-end LSP
- RSVP-TE treats each PSN LSP as a single hop in the end-to-end LSP
- RSVP-TE views the QoS characteristics of each PSN LSP as the QoS characteristics of a hop within the end-to-end LSP
- RSVP-TE determines which ASBRs and PSN LSP's the end-to-end will traverse

Autonomy is maintained

- One peer cannot influence the path that a PSN LSP will traverse in another peer's network
- No need to merge IGP domains
- There is a need to signal end-to-end LSPs across routing domain boundaries, but this can be constrained with policy
  - For example, policy may require end-to-end LSP to traverse PSN LSPs

Other Benefits

- Only PEs and ASBRs need to maintain state regarding end-to-end LSP
  - Interior routers in either peer network need not maintain state regarding end-to-end LSP
- Network events that cause the PSN LSP to reroute need not propagate to the end-to-end LSP
  - So long as PSN LSP restores within required period of time

Application

- IP Traffic Engineering
  - Managing micro-flows between MAN and WAN
- Multi-segment pseudo-wires

Conclusions

- LSP Hierarchy should be considered as a tool for achieving inter-domain traffic engineering
- RSVP-TE should be considered as a tool for signaling LSP hierarchies.

Juniper™
NETWORKS