

Support for RSVP in Layer 3 VPNs



draft-davie-tsvwg-rsvp-l3vpn-02.txt

Bruce Davie
François le Faucheur
Ashok Narayanan
Cisco Systems



Problem Overview (1)

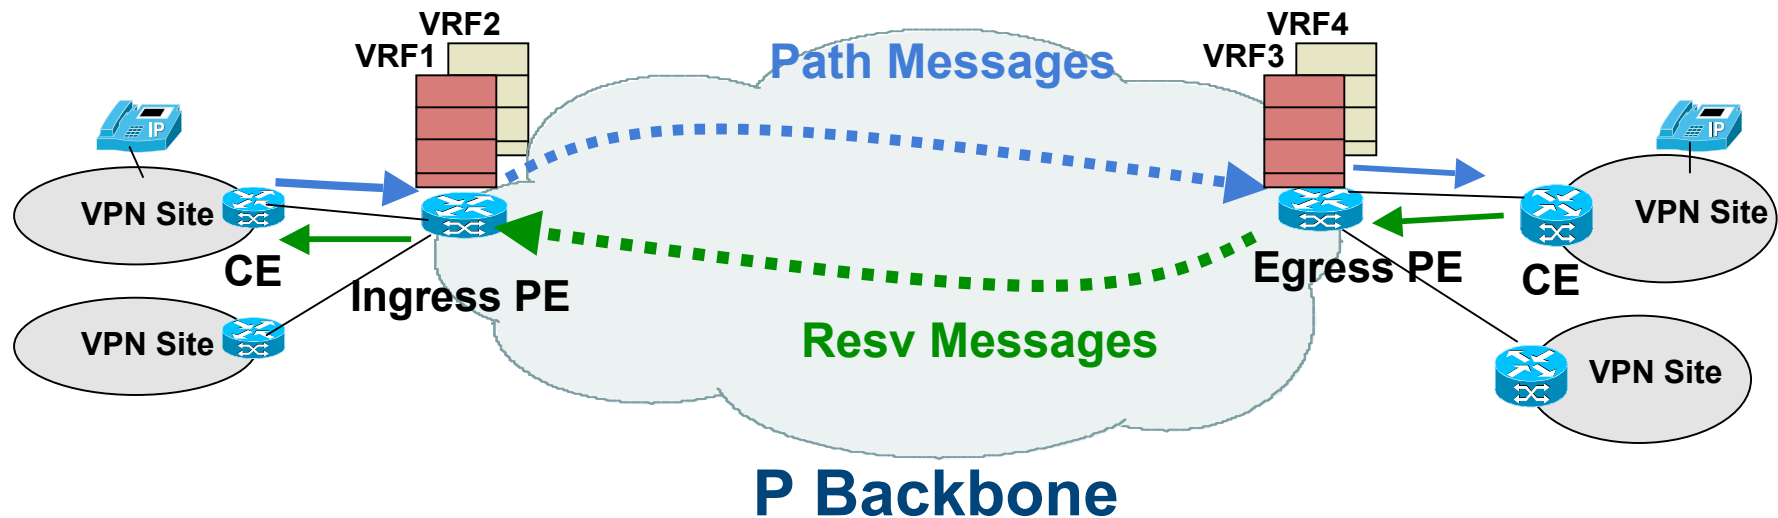
- Admission control may be desired on CE ↔ PE links of layer 3 VPNs (RFC4364)
- Running RSVP across these links presents several issues:
 - Need to associate RSVP messages (which contain C addresses) with appropriate VRF context when they arrive at PE across backbone
 - customer address spaces may overlap
 - Need to intercept Path messages at egress PE but Router Alert IP option may not be visible/accessible
- NB: Focus on admission control, not TE
 - TE has enough differences to warrant new draft



Problem Overview (2)

- May also wish to perform admission control for e2e flows in backbone
 - Clearly need some sort of aggregation for scalability and to avoid installation of per-customer state in P routers
 - Similar to other RSVP aggregation scenarios (e.g. RFC 3175, RFC 4804)
- Need to support Inter-AS operation

Model of operation



- Path messages sent by data senders
- Receivers send Resv messages
 - forwarded back up the path to senders
- Neither Paths nor Resvs are processed in P routers



Changes from -01 to -02

- Main change: introduce VPN-IPv4 and VPN-IPv6 as proper RSVP address families
 - SESSION, SENDER_TEMPLATE and FILTER_SPEC objects modified to carry VPN-IPv4 (or VPN-IPv6) addresses within the MPLS VPN.
 - The VRF_ID and VPN_LABEL objects removed - their function provided by above objects
- Can now support Option B without CAC on ASBRs (or with CAC)
 - New VPN-IPv4 HOP object for this purpose
- Aggregate RSVP (RFC3175) sessions across MPLS VPNs added.
- Explicit support for IPv6 VPNs has been added



Overview of Proposed Solution

- New SESSION, SENDER_TEMPLATE, FILTER_SPEC types in Path, Resv etc. use VPN-IPv4/6 addresses
 - enable PEs to identify appropriate VRF context during RSVP processing
 - appear only in PE-PE messages, not outside provider's backbone (except inter-AS options B and C)
- Control-plane approach to direct Path messages to egress PE for processing, avoiding need for Router Alert handling in data plane
- RSVP over TE tunnels as per RFC 4804 if admission control over provider backbone required



Why VPN-IPv4/6 in RSVP SESSION etc.?


- Responding to feedback from WG
- Simpler than previous approach in Option B ASBRs
 - no per-VPN state as ASBRs
 - allows for CAC-less option B ASBRs
- Preserves the RSVP SESSION as complete and unique identifier of a session (unlike approach in -01 draft)



Summary

- Admission control on PE-CE links would be useful
- Small set of new mechanisms makes RSVP work in VRF context and solves router alert issue
 - Put VPN-IPv4 addresses in Path and Resv messages to enable correct VRF to be identified
 - Address Path messages directly to egress PE or ASBR
- Admission control over backbone is optional, leverages existing techniques (RFC 4804)
- No change to RFC4364 (MPLS/BGP VPN) protocols or operations
- Solution now close to complete, IOHO

Backup





Details

- Path message at ingress PE
 - Find the RD for the prefix that matches dest, append it to dest, use RD:dest in SESSION
 - Find the RD for the source, use RD:src in SENDER_TEMPLATE
 - Set PHOP to an address of the PE
 - Set IP dest of datagram to address of remote PE/ASBR (BGP next hop from lookup in VRF)
 - Forward the message to egress PE
 - Router Alert not required



Details (2)

- Path message at egress PE
 - Use RD:dest from Session to find egress VRF
 - Store Path state, including VRF info
 - Strip RDs from message
 - Set PHOP to outgoing interface address
 - Forward the message to CE, with Router Alert option (as normal)



Details (3)

- Resv message at egress PE
 - Process in appropriate VRF to find the Path state
 - Find SESSION and SENDER_TEMPLATE (with VPN-IPv4/6 addresses) in Path State
 - Use them to create SESSION, FILTER_SPEC in Resv
 - Do admission control on PE-CE link
 - Send to ingress PE



Details (4)

- Resv message at ingress PE
 - Use RD and src addr from FILTER_SPEC to find correct VRF context
 - Find Path state
 - Strip RDs from SESSION and FILTER_SPEC
 - Optional - do admission control on PE-PE tunnel as per RFC 4804
 - Send message to CE (found in Path state)



New PHOP

- New PHOP will contain RD:IPv4_addr
 - RD:IPv4_addr is a VPNv4 route, advertised in BGP with a label
 - IPv4_addr could be almost anything, as long as RD:IPv4_addr is unique - address of the PE-CE link a fine choice
 - May wish to prevent advertisement of this route outside provider's backbone
 - An LSP will exist to this VPNv4 route
 - Resv can be sent along that LSP



Why not use RA label?

- Doesn't provide any obvious benefit
- Requires PEs and ASBRs to look at ALL messages with RA label to find the ones they care about



Can we label switch the RSVP messages?

- What does it take to label switch the Path and Resv messages to the right VRF, rather than using new RSVP objects?
 - Need a per-VRF label
 - ...and a way to advertise it
 - ...and a way to find the right label when sending Path or Resv



Label switching Path msgs

- Can advertise a route to each VRF
- Need some way to identify the VRF (e.g. RD+loopback address)
- Need some way to distinguish VRF advertisement from a customer route advertisement
- Need to identify the correct egress PE (or ASBR) and VRF given a customer address in the Path message



Label Switching Resv Msgs

- Similar issues to Path, but
 - Resv does not contain a customer address as its destination - contains address of PHOP found in Path State
 - That PHOP must contain enough information to tell a PE which VRF label to use (so it can't just be a PE loopback)
 - Probably need a new PHOP type in RSVP
 - Need some means to associate PHOP with correct VRF label advertisement



Summary of label-switch approach

- Yes, we think it can be made to work
- Requires extensions to RFC4364 to ensure VRF labels are advertised and identifiable
- Requires RSVP extension to support new PHOP type
- Not obviously better than documented approach