MPLSCon 2006

New York City, May 25, 2006

Implementing IPv6 and IPv6 VPN in MPLS Networks

Luyuan Fang

AT&T

at&t

Outline

- Benefits and Key Drivers for IPv6 and IPv6 VPN
- General Requirements for IPv6 VPN in MPLS Networks
- Implementing IPv6 and IPv6 VPN over MPLS Backbone
- Design Considerations and Challenges
- Conclusions

Benefits of IPv6

- Key benefit – increasing address space
  - IPv4: 32-bit, ~ 4.3 billion addresses
  - IPv6: 128-bit, ~ 340 undecillion addresses ($3.4x10^{38}$)
- Other advantages
  - Auto-configuration via neighbor discovery
  - Mobility
  - Better routing efficiency and flexibility
    - Six fields removed
    - Extension header added
  - IPSec is mandatory
  - Performance improvement for broadband utilization
    - Jumbograms 4 GB (going to 32) in v6 vs. 64KB in v4
    - use flow label to largely increase the network utilization
  - QoS included in IPv6 headers

at&t

Key Drivers for IPv6 and IPv6 VPN

- Large address space is the most compelling reason for IPv6
  - IPv4 addresses limitation starts to restrict Internet growth and use, especially in China, India, and other heavily populated Asian countries
  - IPv4 address exhaustion has been predicted in less than 8 years
  - Government mandates: Japan, US, China…
  - Explosion of wireless IP devices
    - 2 billion mobile phones by 2006, not enough with what is left today with IPv4 addresses if static addresses are used.
- The urgency of providing IPv6 VPN services in US
  - US Office of Management and Budget (OMB) government Mandate in 8/2005
    - "Federal agencies must use the next-generation Internet service known as Internet protocol version 6 (IPv6) by June 2008…".
  - Following the mandate, government agencies would need to upgrade their IPv4 VPNs to IPv6 VPNs by mid 2008

IPv6 Implementation Options (1)

- Native IPv6 network - All P and PE are IPv6 capable
- IPv6 tunneling through IPv4 network
  - IPv6 manually configured tunnel (RFC 2893)
  - IPv6 over IPv4 GRE Tunnel
  - Tunnel broker
  - Auto 6to4 tunnel

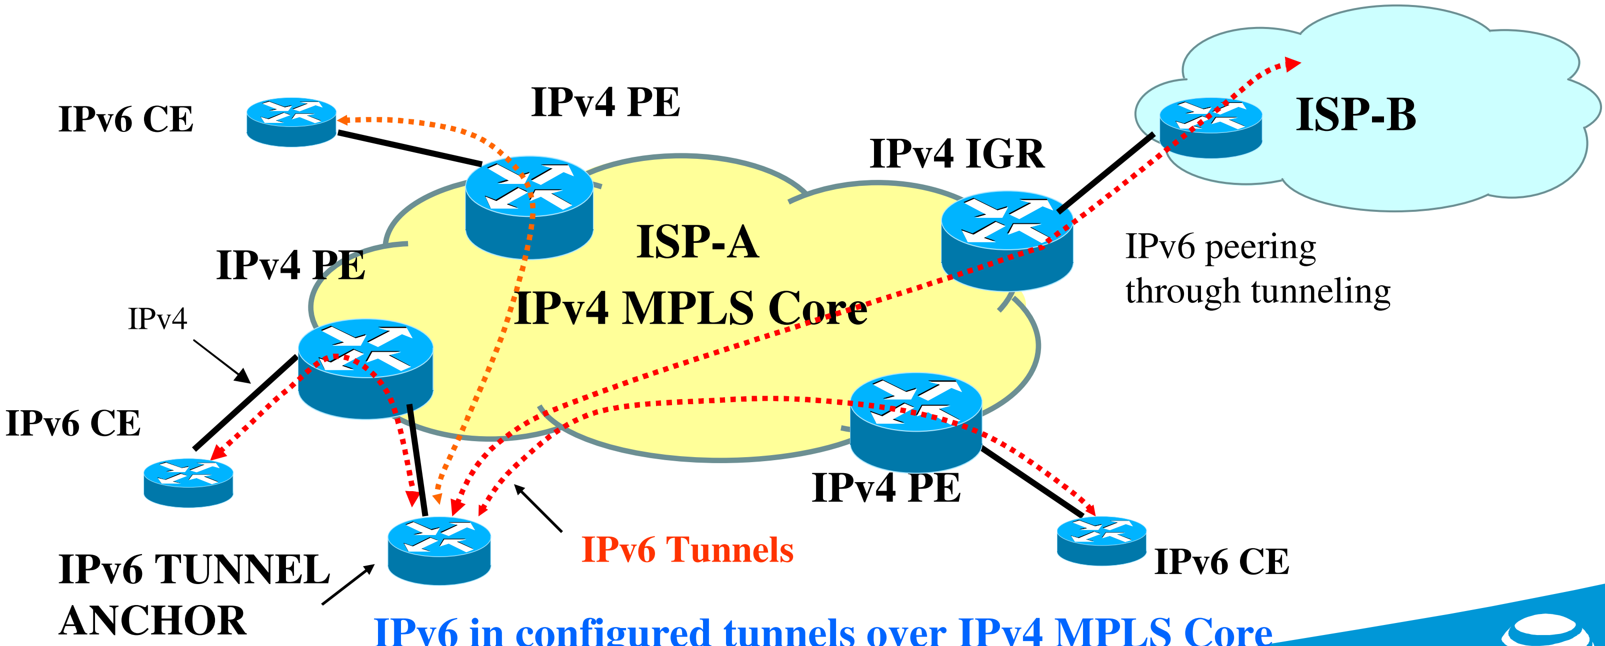

IPv6 in configured tunnels over IPv4 MPLS Core

at&t

IPv6 Implementation Options (2)

- Dual stack 6PE solution
  - Only 6PEs are IPv6 capable, P routers are IPv6 unaware, core remain unchanged
  - Dual stack: Support both IPv4 and IPv6 on the same interface
  - IPv6 reachability established among 6PEs via MP-BGP w v6 support
  - Easy v6 migration or existing IPv4 MPLS backbone
  - Solution is more scalable compared to tunneling solutions.

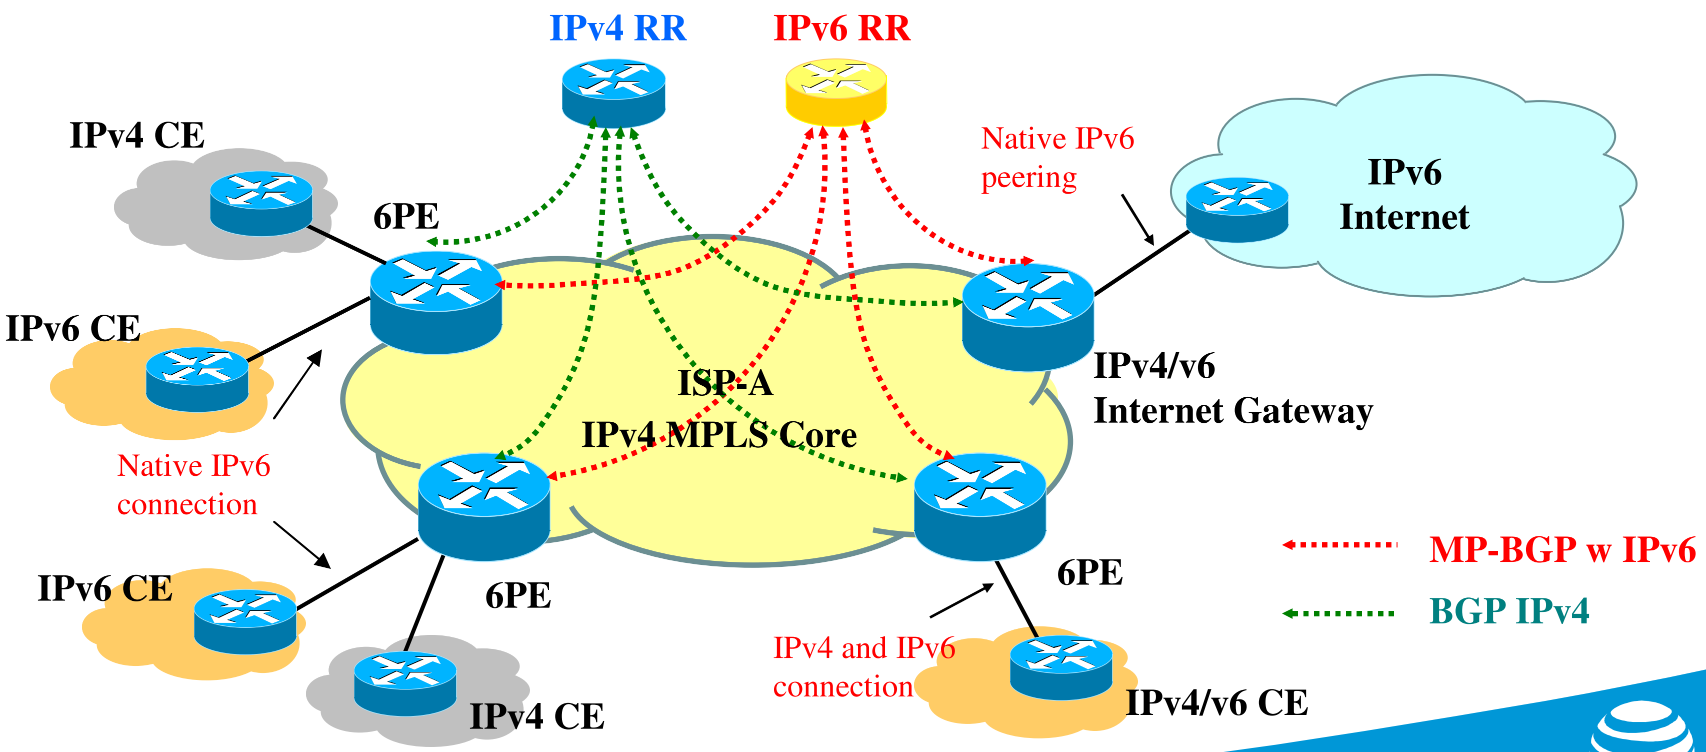

MPLS BGP IPv6 VPN Requirements (1)

- Support the same features set implemented for VPNv4

2547bis VPNv6 features	Access
▪ VPNv6 types ➢ Enterprise VPN ➢ Carrier's Carrier VPN ➢ Inter-AS/Inter-provider VPN	▪ Access types ➢ POS, FR, ATM, PPP, MLPPP, IMA, FR encap over POS, Ethernet
▪ VPNv6 topologies ➢ Any-to-any ➢ Hub-and-spoke ➢ Hybrid	▪ Access speed ➢ FT1/FE1 to T3/E3 ➢ OC-3/STM-1 to OC-192/STM-64 ➢ Ethernet: FE to 10GE
▪ VPNv6 features ➢ QoS per interface / logical interface ➢ Multi-homing load sharing ➢ Multicast VPNv6	▪ PE-CE connections ➢ eBGP, eBGP with labels ➢ Static, Static with labels ➢ OSPF, OSPF with labels ➢ Other protocols supported in VPNv4

MPLS BGP IPv6 VPN Requirements (2)

- No impact to the existing MPLS core
- No performance impact to VPNv4 Services on the same router
- Support of route reflection
- Support of v6 management and monitoring capability
- At least the same level of security as VPNv4
- Scalable in terms of data plane and control plane, e.g.
    - Total bandwidth, # of ports, # of sub-interfaces, # of MLPPP bundle
    - # of BGP sessions, # of routes, # of MVPN routes
- Fast convergence
- Inter-operability among different platforms
- Ease of operation

IPv6 L3 VPN* Technologies
*** Based on rfc2547bis and draft-ietf-l3vpn-bgp-ipv6**

- Reuse existing VPNv4 components:
  - RD, RT, VRF, MPLS
- New components for VPNv6
  - MP-BGP VPNv6 address-family
    - RD [64 bits] + IPv6 prefix [128 bits]
    - Support IPv6 addressing – Global/Unique Local/Link Local
  - Distributing VPNv6 addresses among PEs via MP-iBGP over IPv4
    - RFC 2283 – Multiprotocol extension for BGP4
    - VPN IPv6 NLRI encoding
      - AFI=2 (IPv6); SAFI=128 (MPLS labeled VPNv6)
  - BGP nexthop - IPv4-compatible IPv6 address
    - PE advertises to its peer a Next Hop Network Address field containing a VPN-IPv6 address:
      - RD=0
      - IPv6 address is encoded as an IPv4-mapped IPv6 address (::ffff:IPv4 address)

Implementing IPv6 VPN in MPLS Network

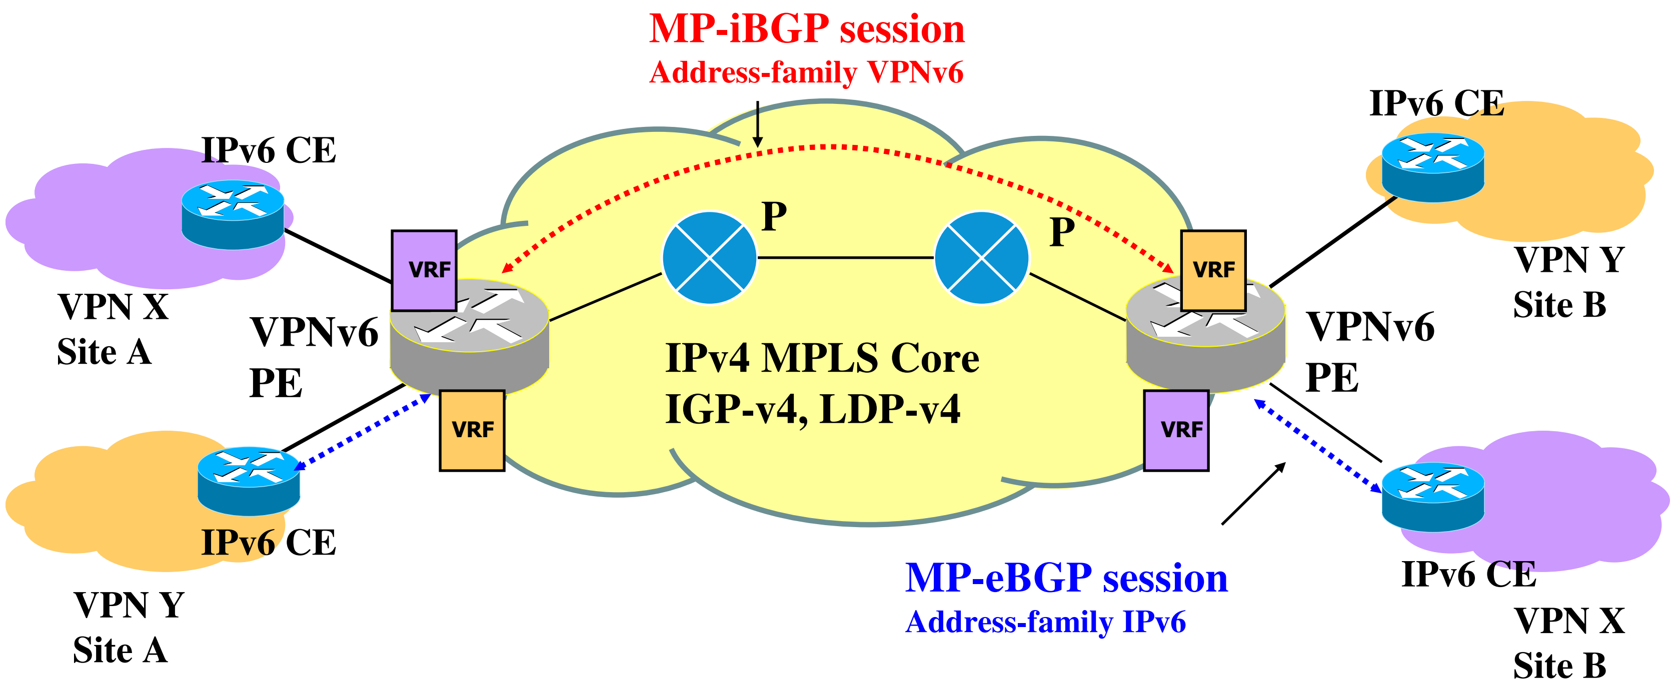

- MPLS IPv4 Backbone unchanged (IGPv4, LDPv4)
  - P - No upgrade required
  - PE - needs support IPv6 BGP VPN extensions
  - LSP – IPv4 signaled
- MP-iBGP with VPNv6 AF peering among VPNv6 PEs
- MP-eBGP with IPv6+VRF AF peering with IPv6 CE

Routing Information Exchange (1)

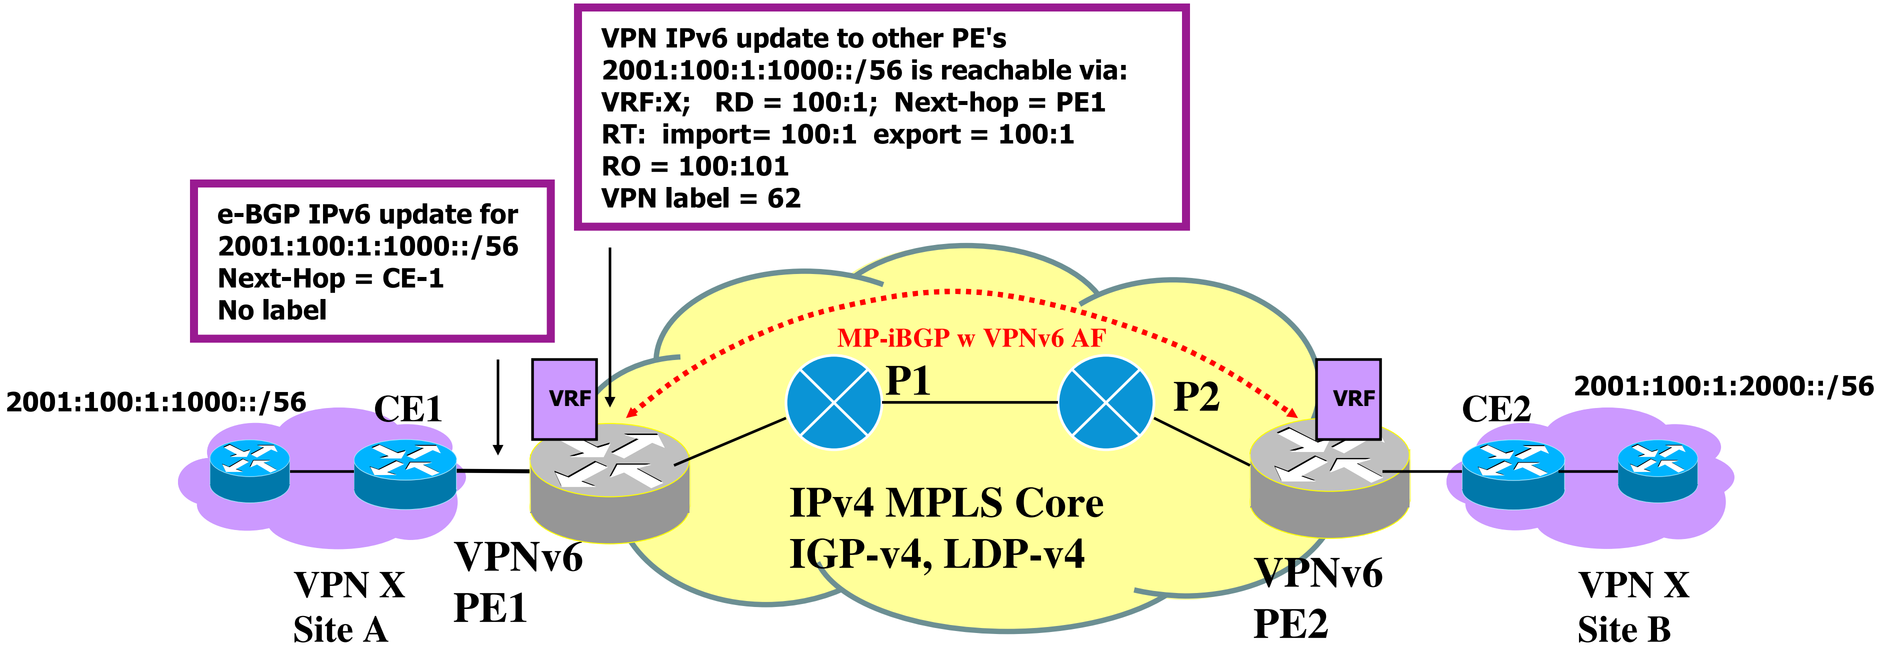

- CE sends IPv6 updates to its peer PE1 (via e-BGP, static, etc.)
- PE1 translates IPv6 into VPN-IPv6
  - Assign RD, RT, RO according to VPN RED configuration
  - Re-write Next-Hop attribute to PE1
  - Assign label for this VRF/ interface
- PE1 sends MP-iBGP updates to all PE neighbors

at&t

Routing Information Exchange (2)

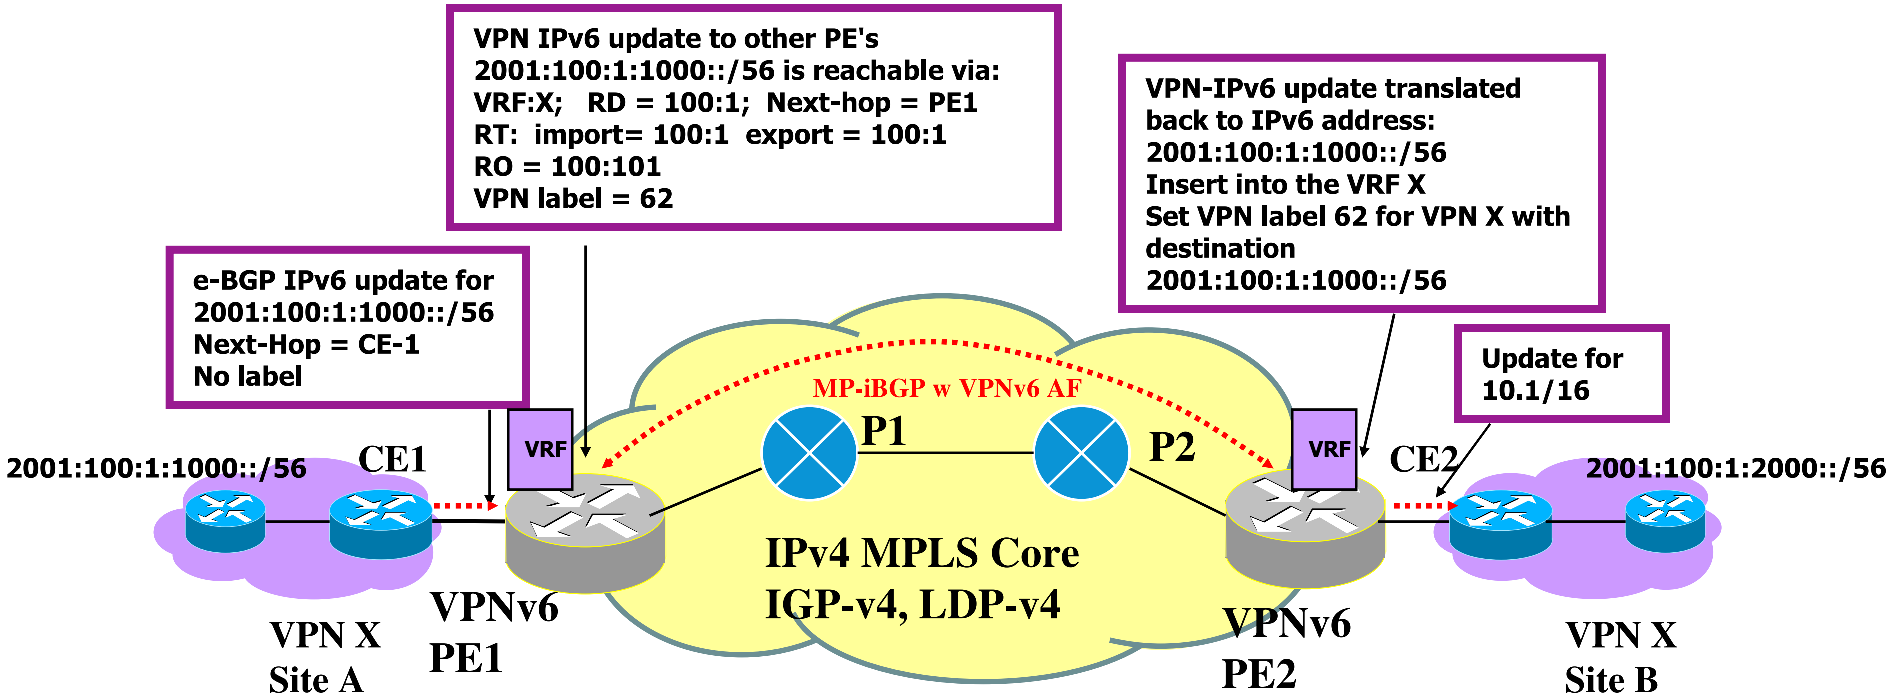

- PE2 receives VPN-IPv6 update, translates it to IPv6
  - Insert the VPN route into VRF X as indicated by RT, per PE2 configuration
- PE2 sets the label associated to VPN-IPv6 address for VPN X, and uses it for forwarding packets to the VPN destination

MPLS VPNv6 Forwarding (1)

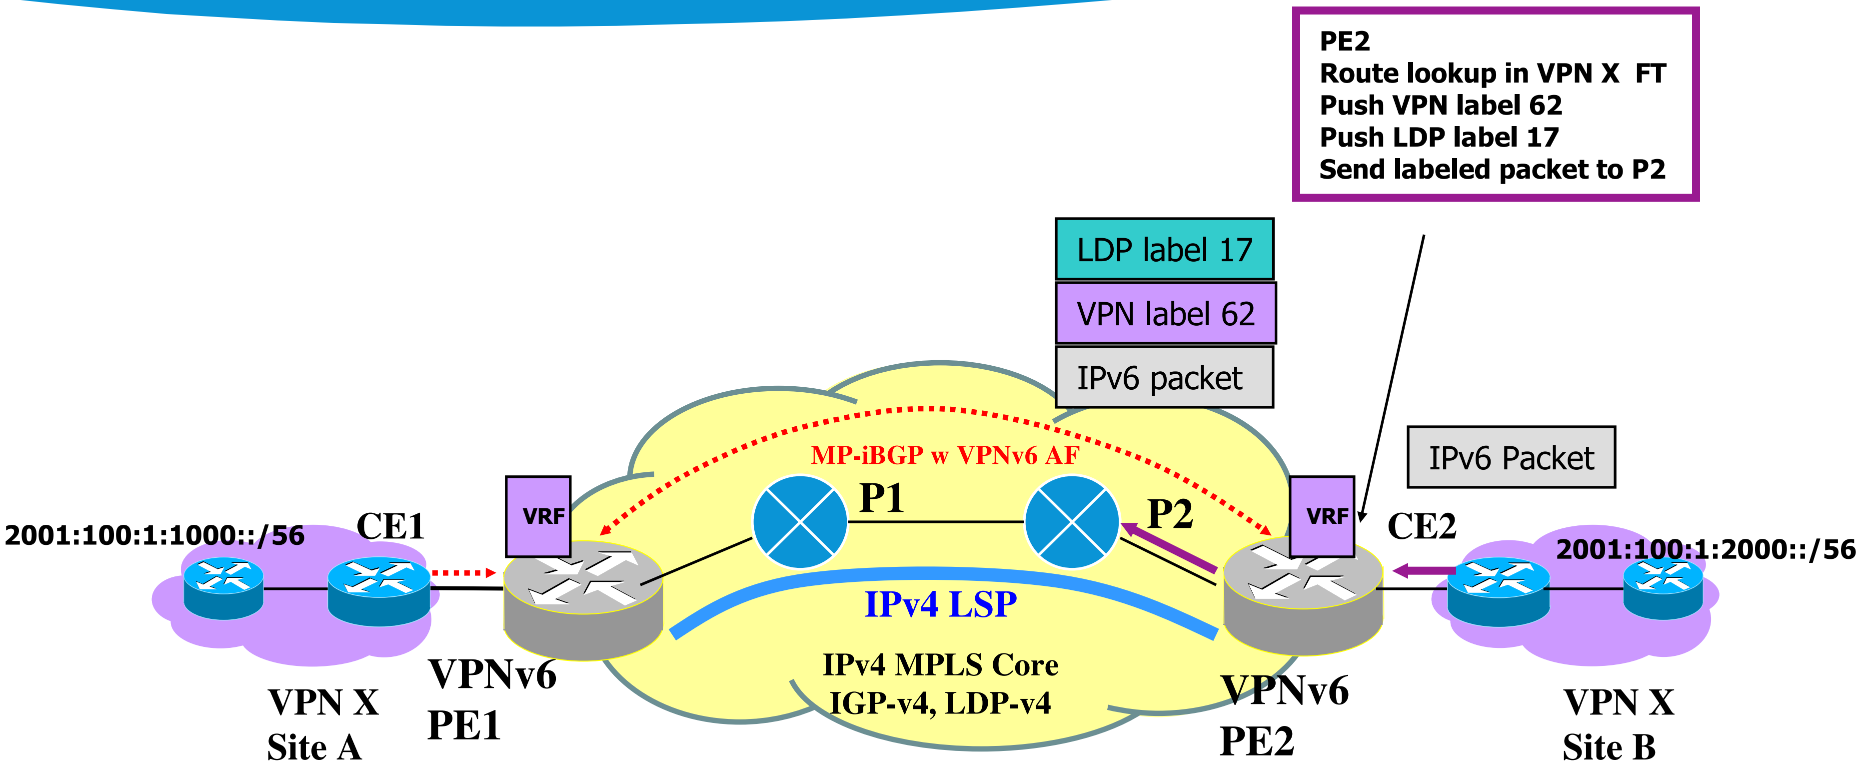

- CE2 sends normal IP packets to PE2, destination 2001:100:1:1000::/56
- PE2 performs "longest match" from VRF, find iBGP next hop PE1, then attaches two MPLS labels on the packet:
    - VPN label as the inner label
    - LDP label as the outer label

MPLS VPNv6 Forwarding (2)

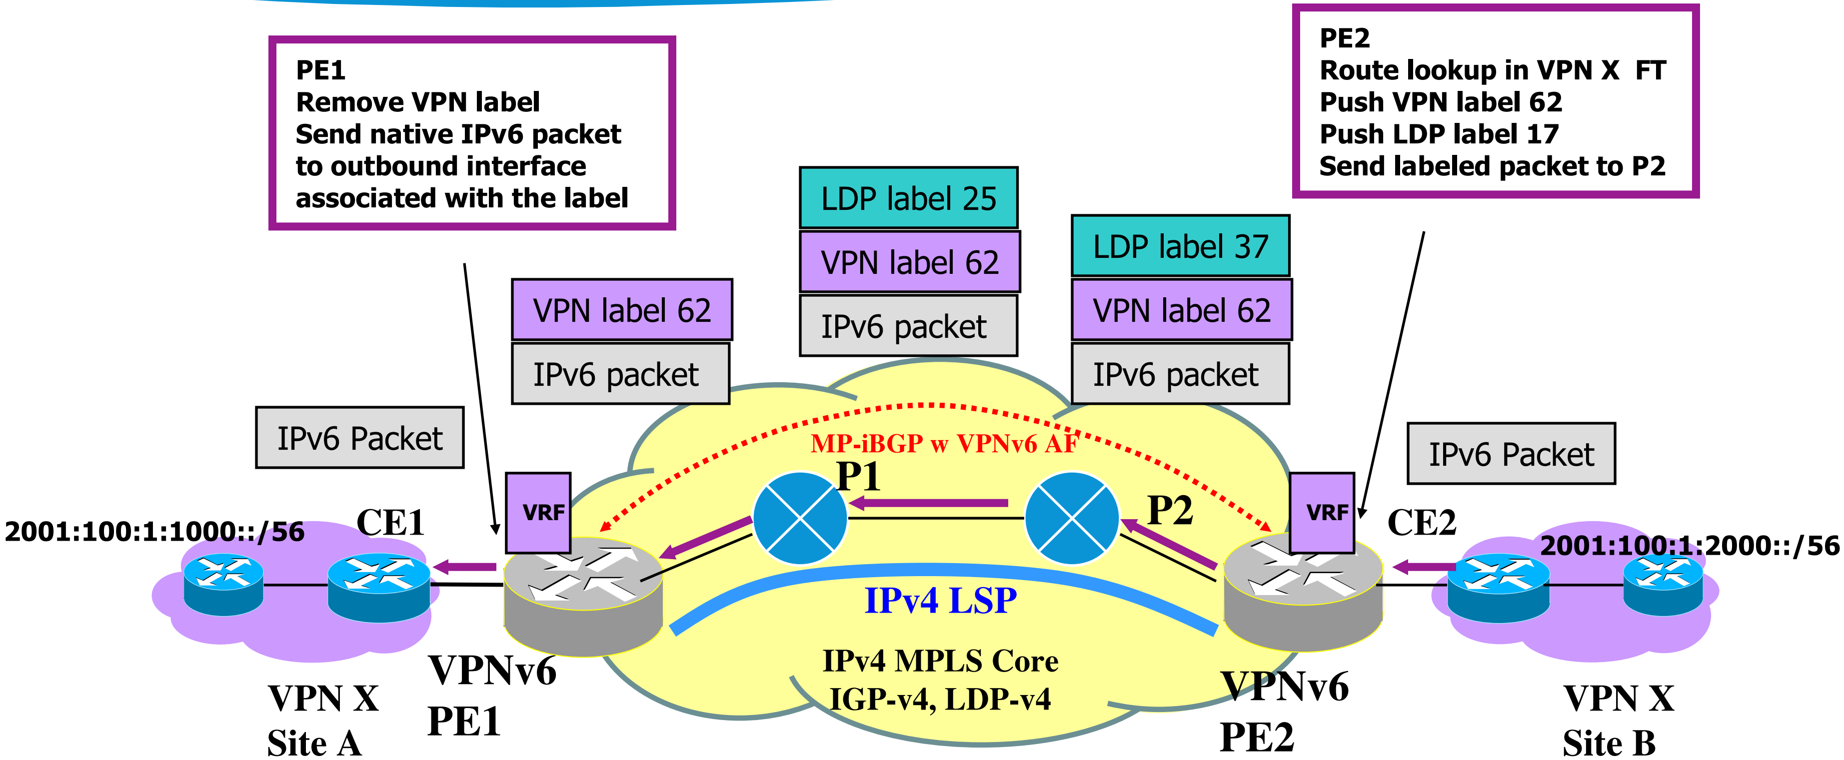

- P2 label switches the packet to P1
  - LDP label (outer) is swapped based the in/out interfaces with pre-assigned label
  - VPN label (inner) remain untouched - *P routers are not VPN aware*
- P1 performs Penultimate Hop Popping
  - Remove the top label and forward the rest to PE1
- PE1 aggregates VPN traffic
  - Use the VPN label to identify outgoing interface (VRF)
  - Remove the VPN label and forward the IP packet to its IP neighbor CE1

at&t

Design/Deployment Considerations

- Meeting customer requirements
  - Feature check list – no less than VPNv4
  - Time constrains – US government mandate: IPv6 compliance by June 2008
- Minimize network impact at initial deployment
  - No backbone changes except P routers may need to support IPv6 for
    - ECMP
    - Traffic flow accounting
  - Expanding VPNv6 footprint with time and experience
  - Use dedicated VPNv6 RR can be a clean start
  - RD, RT, VRF assignment
    - Same RD, RT can be used for VPNv4 and VPNv6 in the same VPN
    - Single VRF support both VPNv4 and VPNv6 in the same VPN
- Dual stack support
    - VPNv4 must be consistent with existing VPNv4 services regardless of the platforms
    - VPNv6 is green filed, address assignment can take advantage of IPv6
- OSS development for IPv6 and VPNv6 support is a major task

Dual Stack IPv4/v6 VPN Deployment Scenario

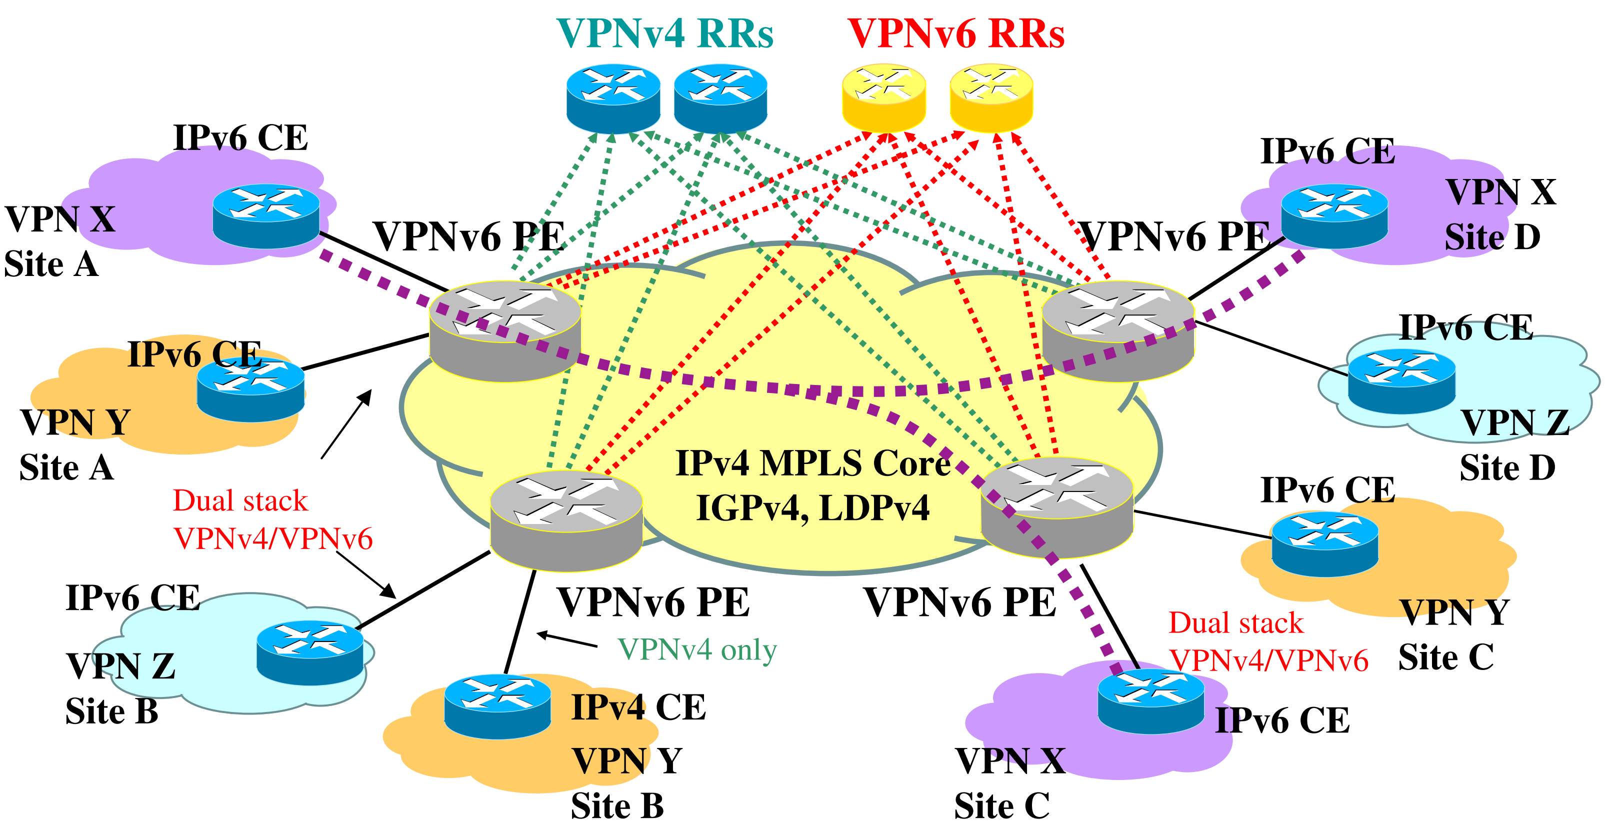

MP-BGP VPNv6
MP-BGP VPNv4

- All VPNv6 PEs are dual stack IPv4/IPv6 capable
  - Single BGP session per PE-CE link
  - Both IPv4 and IPv6 address assigned
  - No V4/V6 translation function on PE initially
- Route reflection supports VPNv4 and VPNv6 separate
- Scale VPNv6 routes support by using RT-filter feature to avoid adv. not needed routes to peers

Scaling VPN routes

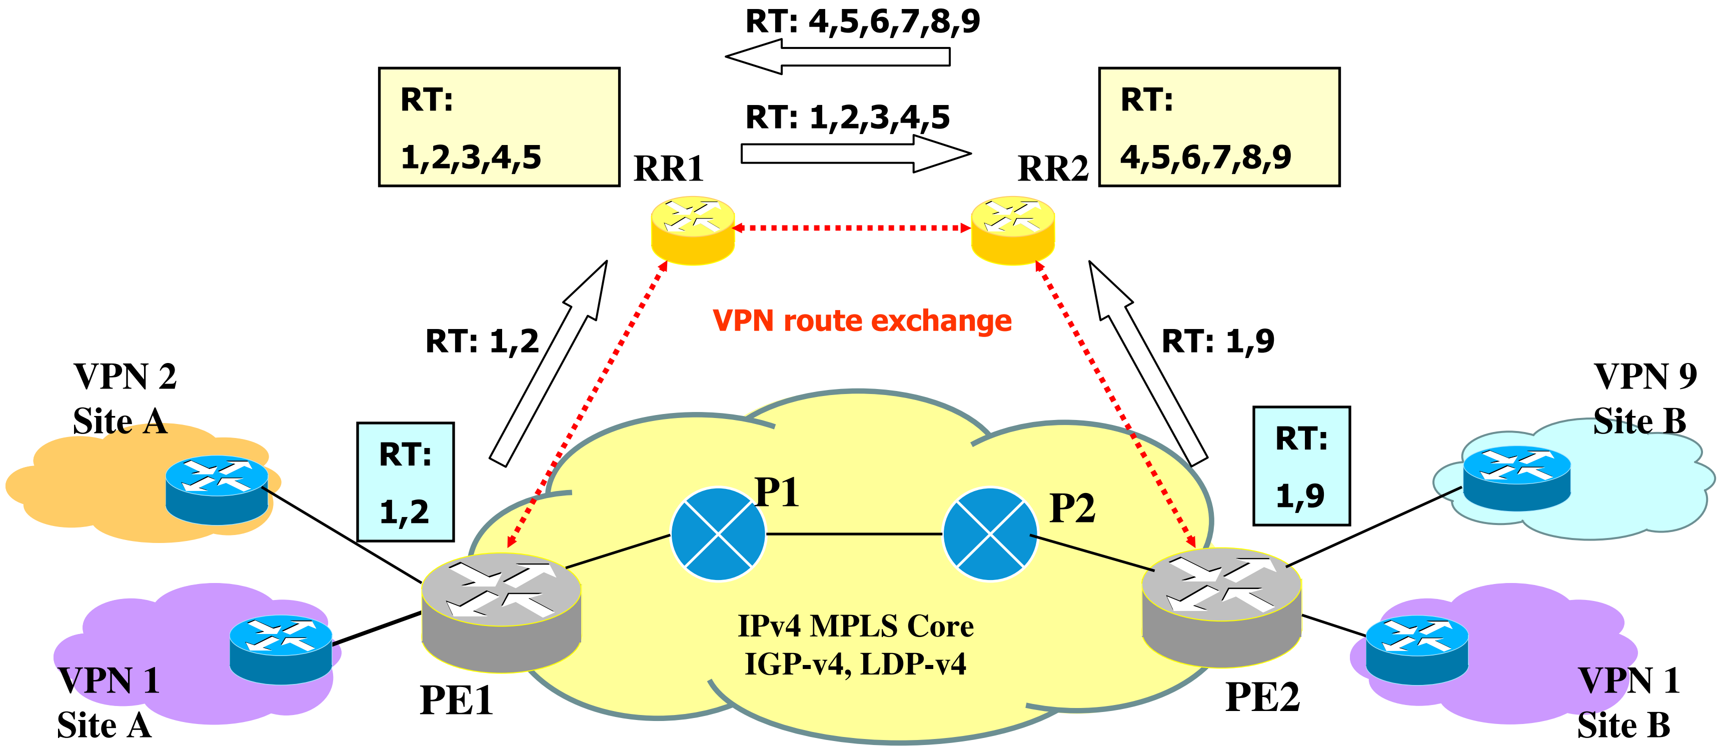

- RT-filter between MP-BGP peers for constrained VPN routes exchange
  - Based on <draft-ietf-l3vpn-rt-constrain-02>
  - Advertise import RTs (RR->RR, PE->RR), not all VPN routes
  - Advertise VPN routes on inverse direction of RT advertisement
  - BGP best path selection selects 1 path per NLRI. NLRI: <as#:RT>
  - Applicable to intra-AS and inter-AS (option B, C)

Security Considerations - v4 and v6

- Alter/disable TTL propagation (core protection - RFC 4111) at the PE
  - Make the backbone look like one hop from the outside, Prevent the backbone addresses from being exposed through trace route
  - Prevent TTL expiry packets cause ICMP time exceeded replies to consume line-card CPU
- Using ACL against infrastructure attacks
  - Control plane protection/policing to protect route processor/routing engine
  - Distributed line-card protection
- QoS pollution control
  - QoS (MPLS EXP) re-coloring to prevent illegitimate traffic from impacting high priority traffic within the backbone
- eBGP security
  - Protect against disruption, redirection of traffic flow
  - Route filtering, dampening, maxas-limit, and MD5
  - Route limit for VPN
  - Control Plane TTL Sanity Check (RFC 3682, GTSM) **-** TTL check on BGP peering packets can effectively block all non-directed BGP spoofing

Security Considerations - v6 Specific

- Routing headers filtering
  - Extension Headers (EHs) filtering and limiting to protect network resources
    - EHs can be manipulated with context causing intensive processing by network elements
    - Header chain can be unlimited (per spec) – a large number of EHs can drain the resources of the routers/devices
  - Filter main header field including Flow Label
- ICMPv6 filtering
  - A large number of functions, message types, and options
  - Security considerations
    - Denial of Service attacks
    - Probing
    - Redirection attacks
    - Renumbering attacks
    - Problems due to ICMPv6 transparency
  - ICMP filtering using IPv6 ACLs, e.g.
    - Rate-limit the number of ICMP error messages generated
    - Re-direct ACL to disable sending redirect packet
  - Best practice guidelines for Filtering ICMPv6 Messages:
    - <draft-ietf-v6ops-icmpv6-filtering-bcp-00.txt>

Inter-As IPv6 VPN

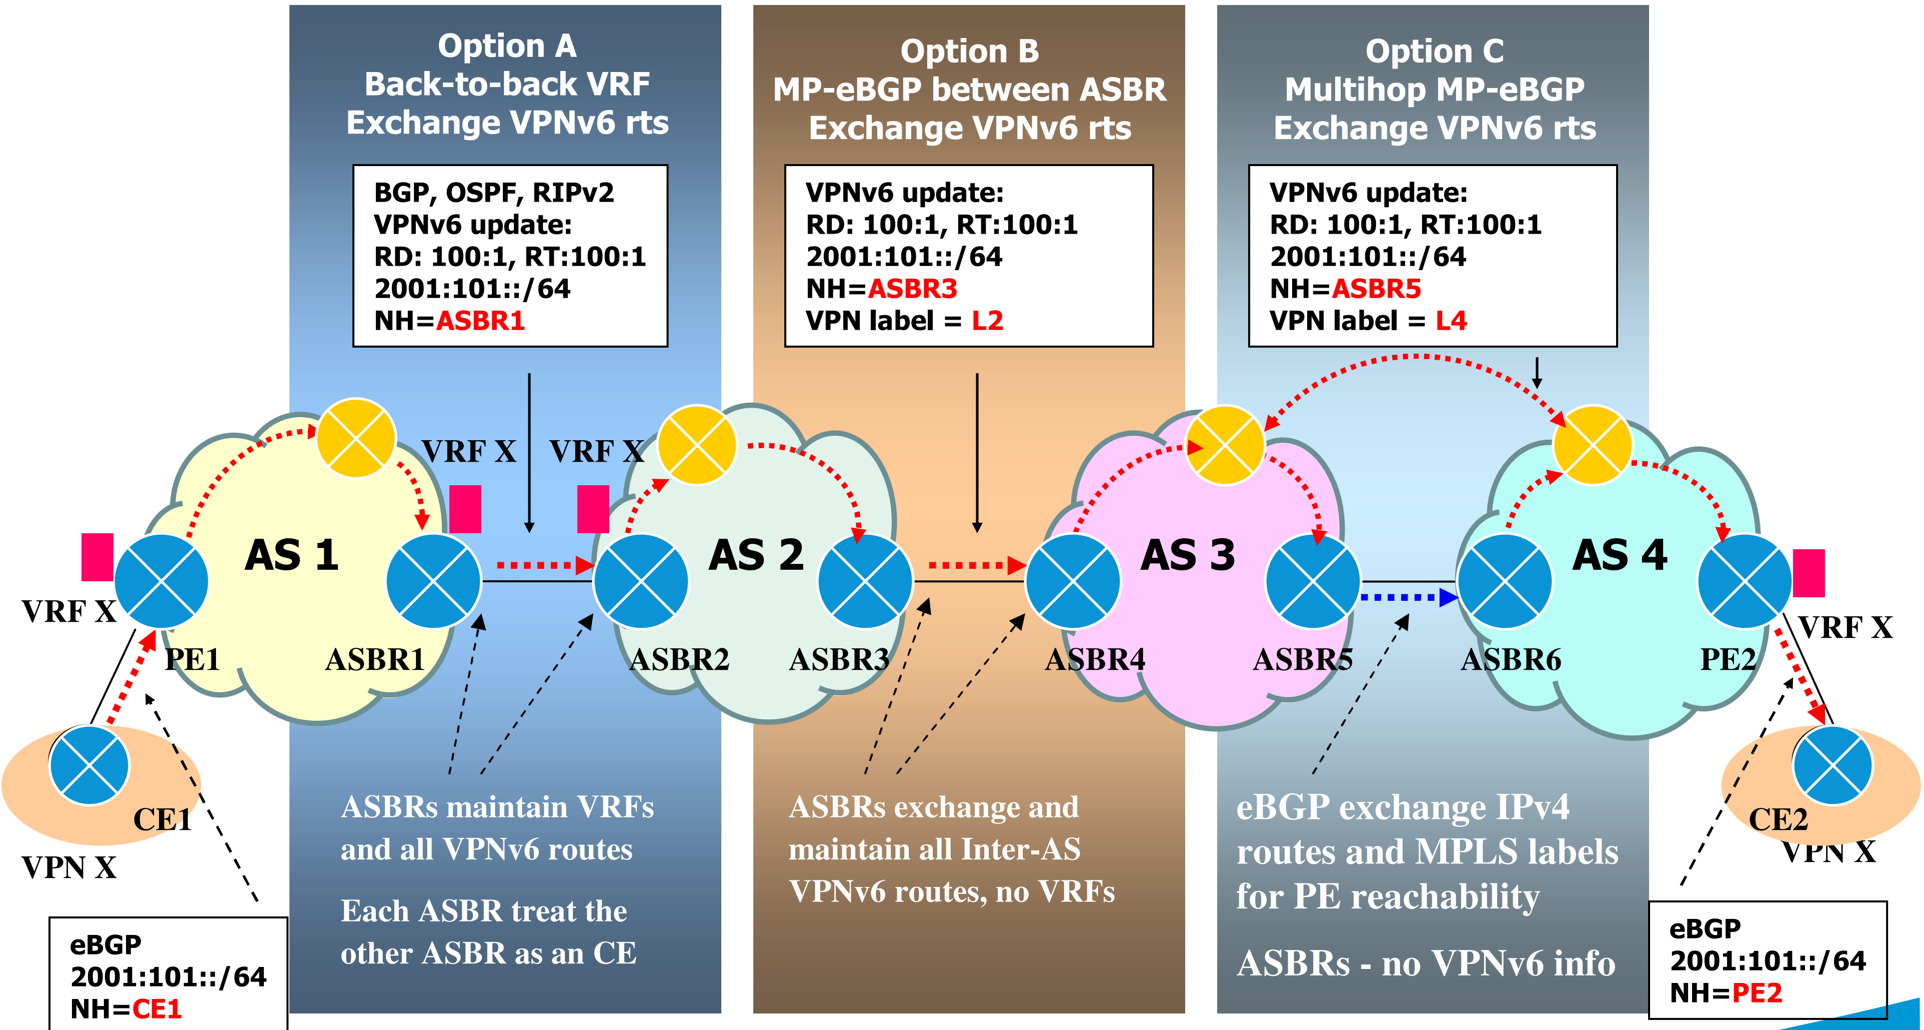

VPNv6 Inter-AS - same options as VPNv4 inter-AS

Multicast IPv6 VPN (1)

- MVPNv6 support challenges – still waiting for VPNv4 scalable design to settle in IETF
  - ➢ Work in progress <draft-ietf-l3vpn-2547bis-mcast-01.txt>

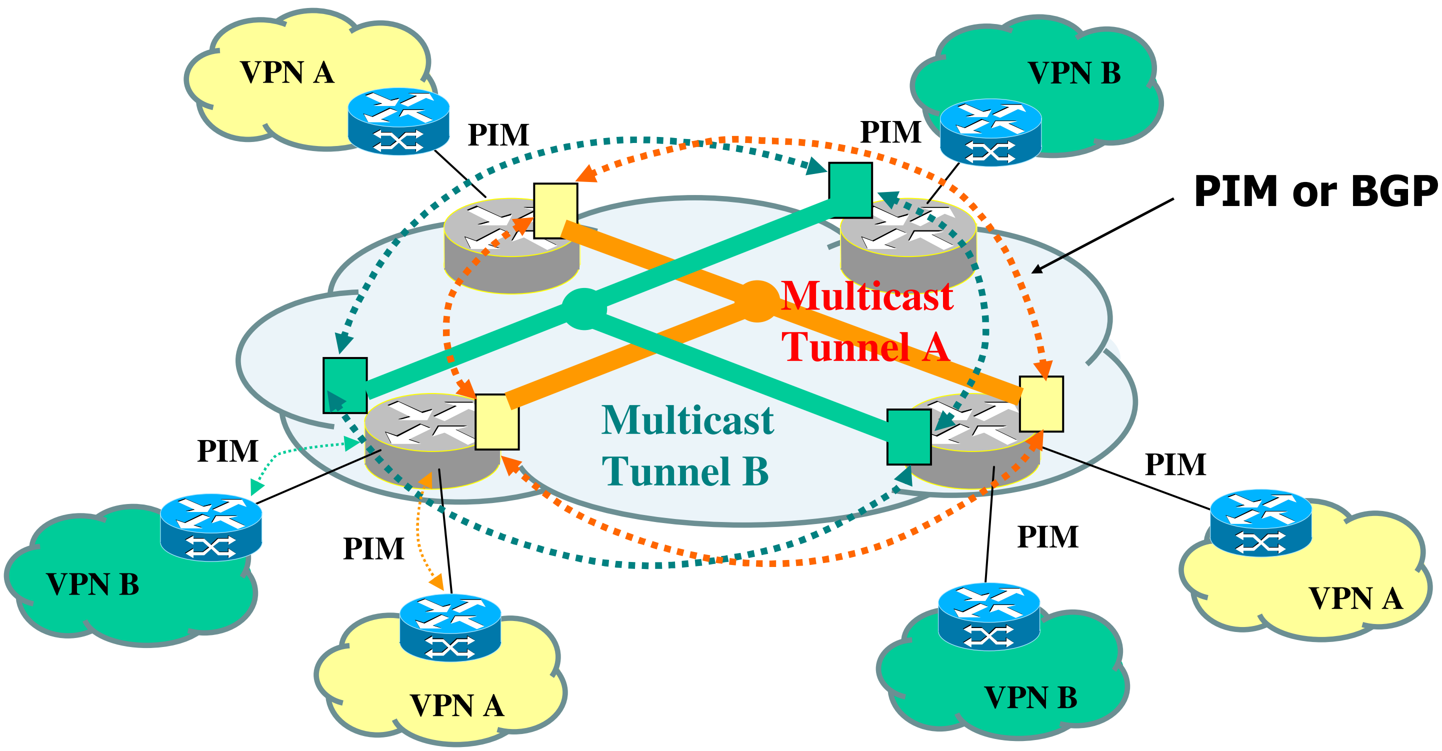

adjacency among VRFs for IPv4/v6 VPN R
adjacency among VRFs for IPv4/v6 VPN B

Multicast IPv6 VPN (2)

- Follow the similar requirements as for IPv4 MVPN
  - MVPN Requirements: <draft-ietf-l3vpn-ppvpn-mcast-reqts-04>
    - e.g. Avoid sending MVPN traffic to non-receivers

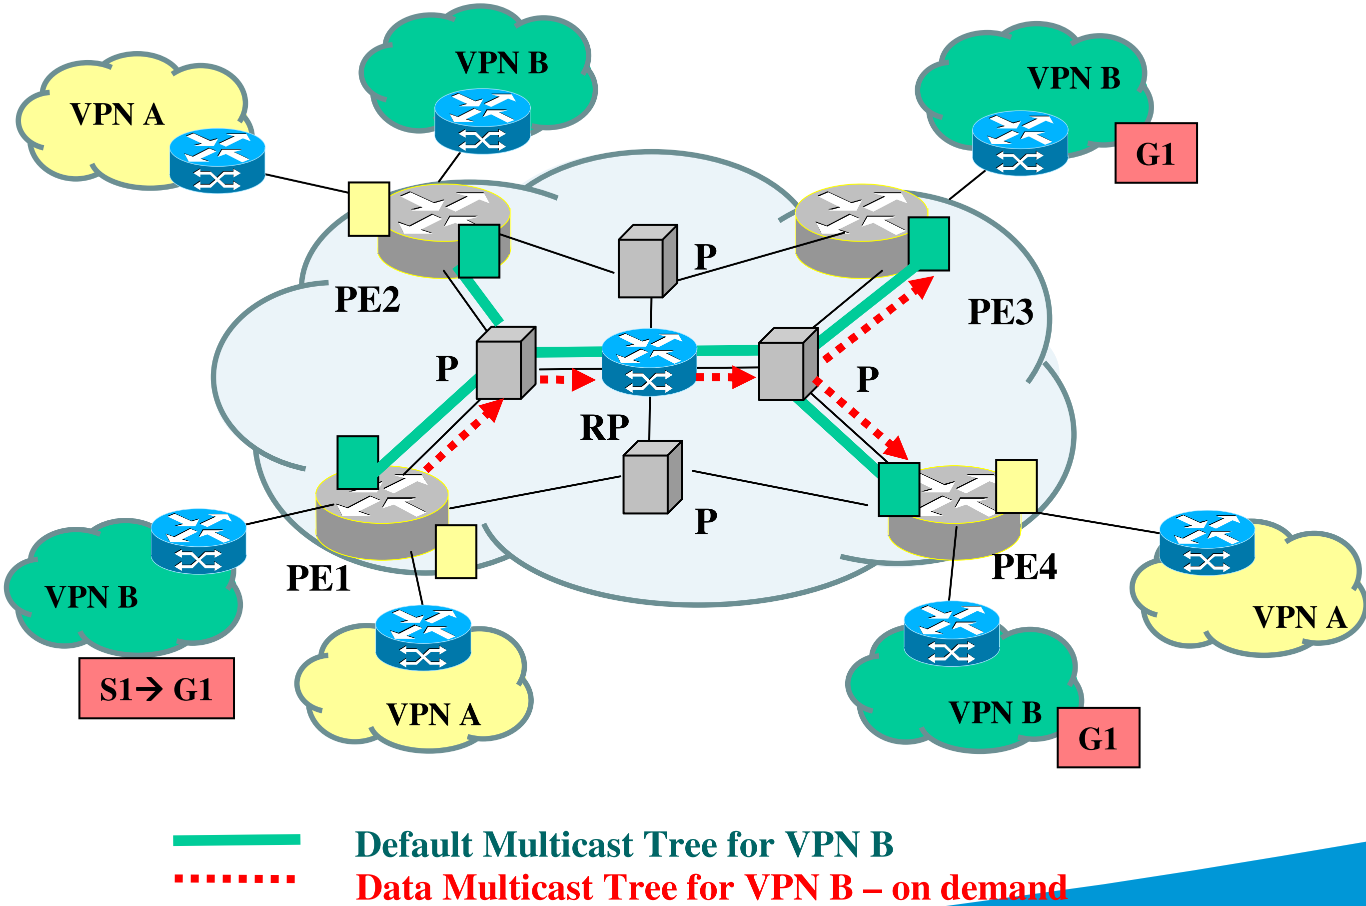

MPLS Multicast Development

- Current Issues
  - No native MPLS (LDP or RSVP-TE) support for multicast
    - Relying on other tunneling mechanisms
    - Piggybacking MPLS labels distribution on PIM
- IETF recent development for MPLS multicast
  - Building p2mp (Point-to-multipoint) LSP with RSVP and LDP
  - MPLS extension to support upstream labels allocation
  - RSVP-TE and LDP extension to support upstream label distribution

MPLS 2547 VPN Multicast Development (1)

- Current Issues - Scalability
  - Control plane
    - Multicast: PE maintain PIM peering with all other PEs which have common VPN(s), vs. Unicast: PE BGP peering with limited number of VPN RRs
    - Per VPN PIM peering
  - Data plane
    - Multicast VPN States Grow with the number of VPNs vs. Unicast: no VPN states on P and only connected VPN states on each PE
      - PIM-SM: # of multicast trees = # of MVPNs
      - PIM-SSM: # of multicast trees = # of MVPNs x average # of PE per VPN
  - Multicast VPN forwarding
    - limited to PIM based IP/GRE tunnels

MPLS 2547 VPN Multicast Development (2)

- IETF recent development on 2547 MVPN
    - <draft-ietf-l3vpn-ppvpn-mcast-reqts-05.txt>
    - <draft-ietf-l3vpn-2547bis-mcast-01.txt>
  - Routing exchange
    - Using BGP for VPN multicast routing info exchange to reduce control plane overhead
      - Similar approach as 2547 VPN: CE <-> PE and PE<->PE
    - Using route reflection to scale as in Unicast
  - Forwarding state aggregation
    - Inter-AS tunnels, <source AS, MVPN> vs. <source PE, MVPN>
    - Using p2mp LSP hierarchy
  - Forwarding MVPN traffic
    - Using p2mp LSP (RSVP-TE or LDP)

Conclusions

- Where are we with v6?
  - IPv6 work started in IETF more than 10 years ago
  - Many networks deployed IPv6 in recent years, more to follow
  - IPv6 VPN deployment is in progress
- Requirements for IPv6 VPN
  - Support all features/capabilities as in IPv4 VPN
  - Ease of migration
- Technologies for IPv6 VPN in MPLS network
  - Same principle as IPv4 VPN
  - MP-BGP extension - VPNv6 Address Family
- Challenges
  - Migration toward IPv6 VPN support
    - Network upgrade
    - OSS development is a major task
  - Additional security mitigation for IPv6 and IPv6 VPN
  - Multicast for IPv6 VPN support

MPLSCon 2006

New York City, May 25, 2006

THANK YOU!

luyuanfang@att.com

at&t