Recent and Future Developments in Profitable MPLS

Matt Kolon

MPLS Japan

October 2003

Juniper NETWORKS™

Virtual Private LAN Service

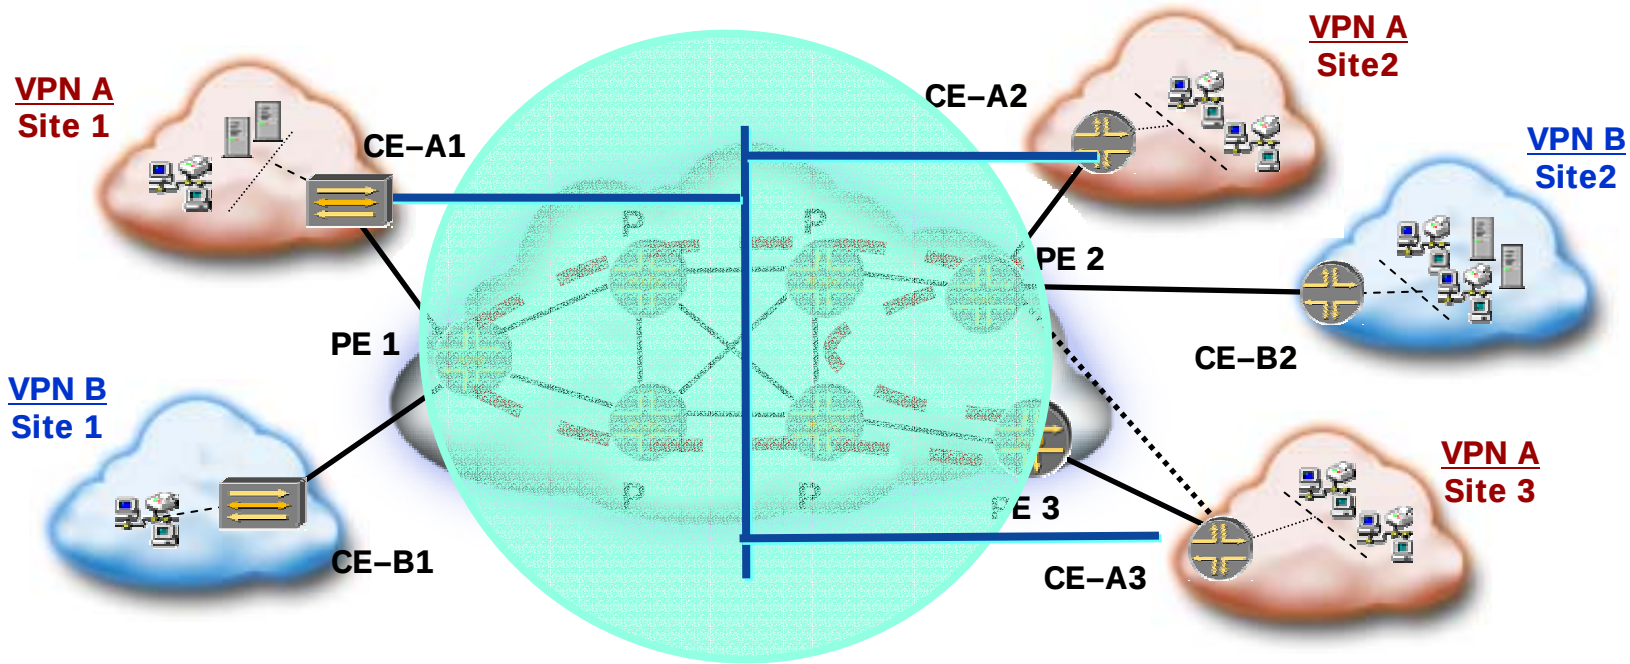

- A private Ethernet network constructed over a ‘shared’ infrastructure which may span several metro areas
- Multipoint to Multipoint Ethernet connectivity where the SP network looks like an Ethernet broadcast domain
- Compliments Layer 3 2547 and Layer 2 VPNs

Another kind of VPN?

- Another choice for customers
  - This is a good thing!
- Another revenue opportunity
  - This is also a good thing!

- But on to the details:
  - How can we implement this new service in a way that allows excellent performance but low operational cost?
  - How can we build an economical, practical service?
  - How can we best capitalize on our current skills and infrastructure?

VPLS Operations

- Control Plane
  - VPN Discovery
    - Discover who are the PE members of a given VPN
  - VPN Signaling
    - Setup and teardown of the pseudo-wires between VPLS instances that constitute the VPLS Domain
- Forwarding Plane (not talking about this today)
  - MAC Learning and Packet Forwarding
  - MAC Aging
  - MAC Flushing

VPLS Control Plane

- Control Plane
  - VPN Auto-Discovery
    - Auto-discovery can be done by BGP
    - IETF proposals to extend DNS or RADIUS for auto-discovery
  - VPN Signaling
    - Demultiplexors can be signaled
      - by *targeted* LDP (draft-lasserre-vkompella-ppvpn-vpls)
        - O(N^2) LDP sessions operational challenge
      - by BGP (draft-kompella-ppvpn-vpls)
        - O(N) BGP sessions operational challenge
  - A single MP-BGP NLRI supports both Auto-Discovery and Signaling
  - Using two different protocols for Auto-discovery & Signaling
    - More complexity and inter-protocol interactions
    - More protocol state in the network

BGP does both Auto-discovery & Signaling

- IP VPN services (aka RFC2547 VPN)
  - RFC2547, draft-ietf-ppvnp-2547bis
- BGP for VPN auto-discovery
  - draft-ietf-ppvpn-bgpvpn-auto
- IPv6 VPN
  - draft-ietf-ppvpn-bgp-ipv6-vpn-03.txt
  - Extensions to RFC 2547bis to support IPv6 VPNs
- Virtual Private LAN Service (VPLS)
  - draft-kompella-ppvpn-vpls
- BGP is a proven, multi-vendor solution deployed in production networks today

VPLS Control Plane functionality with MP-BGP

- Using BGP for VPN Auto-discovery and Signaling provides the following benefits:
  - A single MP-BGP NLRI for most efficient Auto-Discovery and Signaling
    - No additional overhead
    - No need for complex inter-protocol interactions
  - Same framework as IP-VPNs (2547bis), and others
  - Takes advantage of all the scalability, redundancy and operational simplicity features available in BGP:
    - Route Reflectors, Refresh, ORF, etc…
  - Supports Multi-AS/Multi-provider operations

VPLS Auto-discovery & Signaling

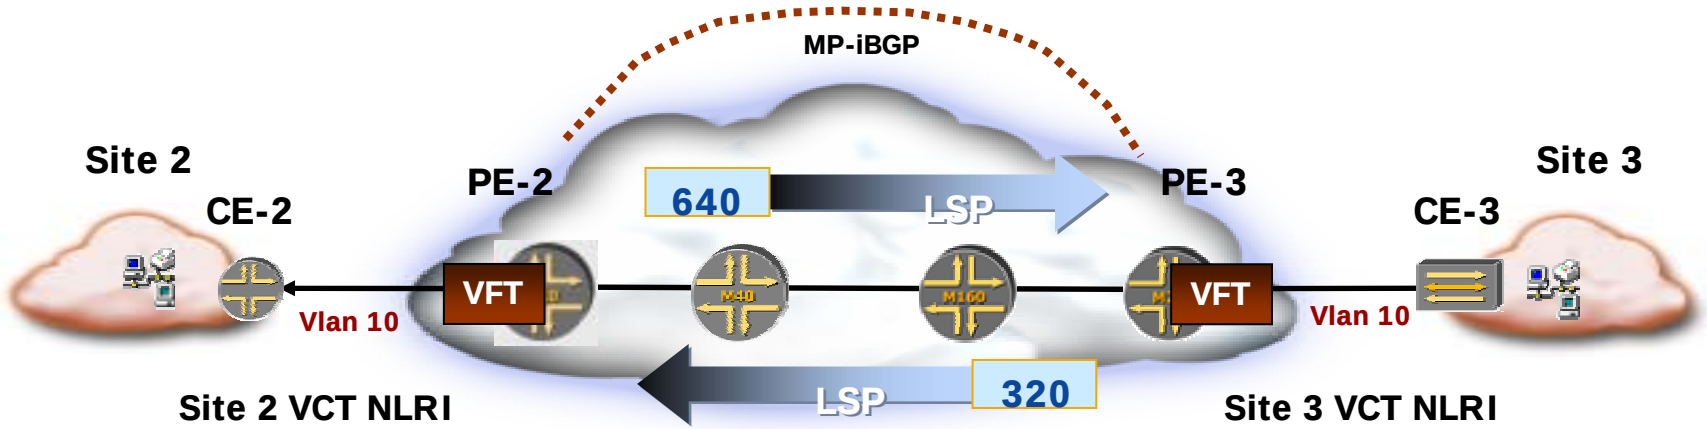

Site 2 VCT NLRI

Route Dist	100:1.2.3.2
VE ID	2
Sites	20
Label base	2000
Route Target	RED
Next Hop	PE-2

Text

Site 3 VCT NLRI

Route Dist	100:1.2.3.3
VE ID	3
Sites	20
Label base	3000
Route Target	RED
Next Hop	PE-3

PE-PE VCT distribution using Multi-Protocol BGP (RFC 2858)

Requires full-mesh MP-iBGP or Route Reflectors

Route Distinguisher: disambiguates VCT information

Route Target: determines VPN topology

Analogous to CE-PE routes advertisements in RFC2547 VPNs

VPLS requires one single NLRI advertisement per VPLS instance per PE

BGP Delivers Scalable Route Reflector Architecture

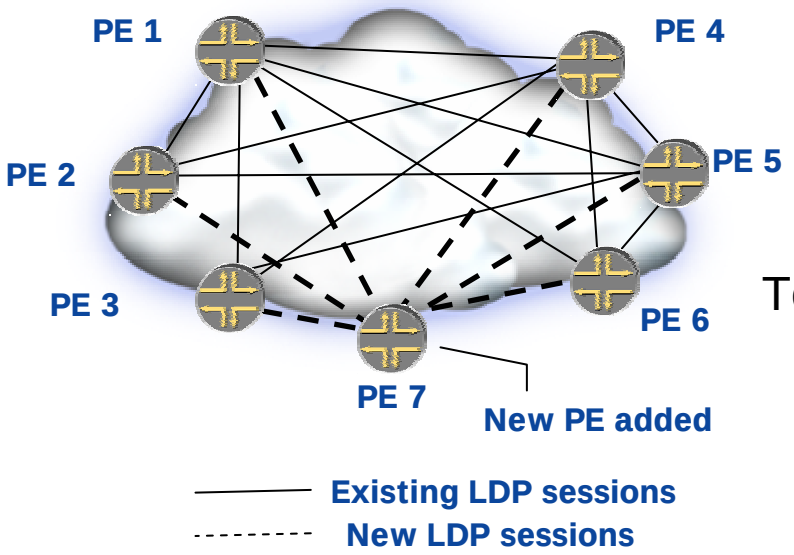

Text

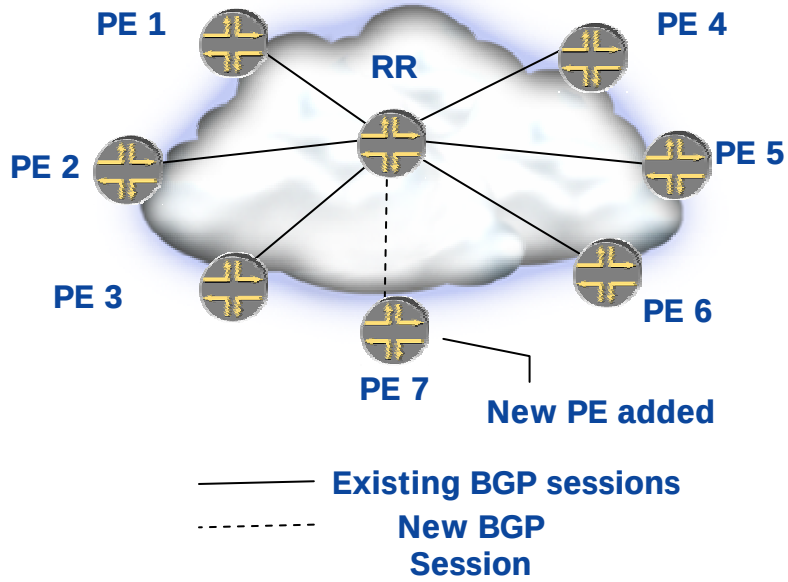

Targeted LDP: New PE Added

7 PEs need to be configured to establish LDP sessions between PEs. All 7 need to be configured again if MD5 authentication desired.

BGP: New PE Added

PE and Route Reflector need to be configured for 1 BGP session established with route reflector

Why not BGP + LDP ?

- BGP or LDP for signaling (distributing VPN labels) is a wrong comparison; ignores autodiscovery
- Compare BGP for both autodiscovery and signaling with BGP for autodiscovery ***and*** LDP for signaling
  - With BGP + LDP approach, signaling of VPN labels requires a completely separate protocol (LDP)
  - With BGP approach signaling, of VPN labels is a “side effect” of (BGP-based) autodiscovery
- The overhead of using BGP for both autodiscovery and signaling is about the same as using BGP just for autodiscovery

Overloading BGP ?

- Common concern:
  {public Internet + RFC2547 VPNs + V6 VPNs + L2 VPNs + VPLS}
  will overload BGP, causing it to crash

- Real question:
  should a single PE provide all of the above services ?
  - If so, it doesn't matter which protocol –
    either the PE device can take it, or not
  - Putting these services in different protocols doesn't reduce overall stress on the PE!
  - Existence proof that some PEs *can* take it

Overloading BGP ? (cont.)

- Other concern: BGP is complex to implement
  - The (perceived) vendor complexity of implementing Multi-Protocol BGP is well worth the greatly simplified operations for the Service Providers
    - Using Multi-Protocol BGP for both auto-discovery and distributing VPN labels means fewer protocols to operate, manage and debug
    - Therefore it means cost savings in terms of:
      - Educating the NOC staff
      - Building tools/expertise to monitor VPNs
      - Building tools/expertise to debug and manage the VPNs

BGP/MPLS VPLS Scalability “divide and conquer”

- Two levels of labels to keep P routers free of all the VPN routing information (v4, v6, L2 VPNs & VPLS)
- PE router has to maintain VPN information only for VPNs whose sites are directly connected to the PE router
- Partition BGP Route Reflectors within the VPN Service Provider among VPNs served by the Provider

- No single component within the system is required to maintain information for all the VPNs
- Routing/forwarding capacity of the system isn’t bounded by the capacity of an individual component

Applies to RFC2547 VPN, L2 VPN and VPLS !

Inter-AS/Inter-provider operations

- Exchange VPN information + VPN labels across AS/provider boundary by using BGP between BGP Route Reflectors in each AS/provider
  - Route Reflectors preserve the next hop information and the VPN label across the AS/provider
- PEs learn routes and label information of the PEs in the neighboring ASs through ASBRs
  - Using labeled IPv4 routes
- No VPN information (e.g., VRF, VFT) on ASBRs

Perhaps most important of all!

Applies to RFC2547 VPN, L2 VPN and VPLS !

Summary

Customers want:

- IP VPNs (RFC 2547 VPN)
- Point-to-point Layer 2 VPNs
- Virtual Private LAN Service (VPLS)
- Service Providers can offer all of the above:
- Over a common infrastructure (MPLS)
- With a common BGP framework
  - Auto-discovery and Signaling
  - Product proven, multivendor
  - Leveraging BGP scalability
  - Supporting multi-AS/multi-provider

A single operational infrastructure and a small set of basic mechanisms means considerable savings for providers, and maximum flexibility and choice for your customers.

Thank you!

Matt Kolon

matt@juniper.net

References

draft-augustyn-vpls-arch

draft-ietf-ppvpn-vpls-requirements

draft-kompella-ppvpn-vpls

draft-kompella-ppvpn-l2vpn

draft-lasserre-vkompella-ppvpn-vpls

draft-martini-l2circuit-encap-mpls

draft-martini-l2circuit-trans-mpls