

MPLS Inter-AS VPNs

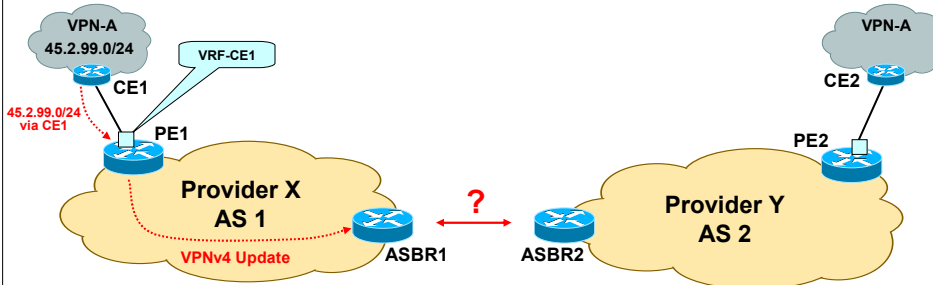
Interconnecting MPLS Networks

(C) Herbert Haas 2005/03/11

“A complex system that works is invariably found to have evolved from a simple system that worked.”

John Gall

Why Inter-AS VPNs?



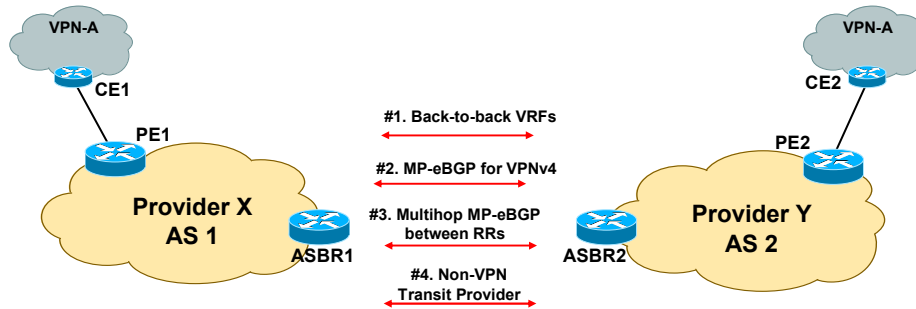
- Same VPN customer may be attached to different providers
 - ♦ Geographically distributed sites
- Service providers must ensure seamless VPN connectivity
 - ♦ Advantage: Quicker geographic service coverage
 - ♦ But how should provider X and Y exchange routes?
- This is called "Inter-AS VPN"
 - ♦ Or "Inter-Provider VPN" or similar...

How to Enable Inter-AS VPN



1. Enable MPLS-VPN in all core networks
2. Exchange IPv4 routes + labels
 - ♦ IGP (or static routes) + LDP between ASBRs
 - or
 - ♦ MP-iBGP + LDP session between ASBRs
3. Exchange VPNv4 routes + labels
 - ♦ MP-eBGP session between ASBRs
 - or
 - ♦ MP-eBGP multihop session between RRs

Four Inter-AS Solutions

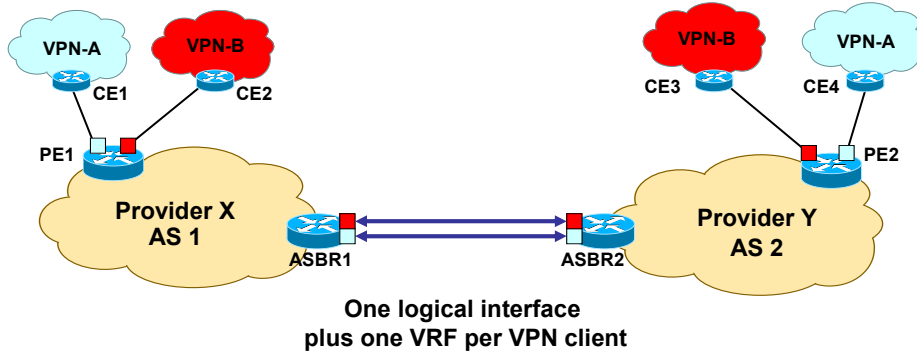


#1: Basic Idea

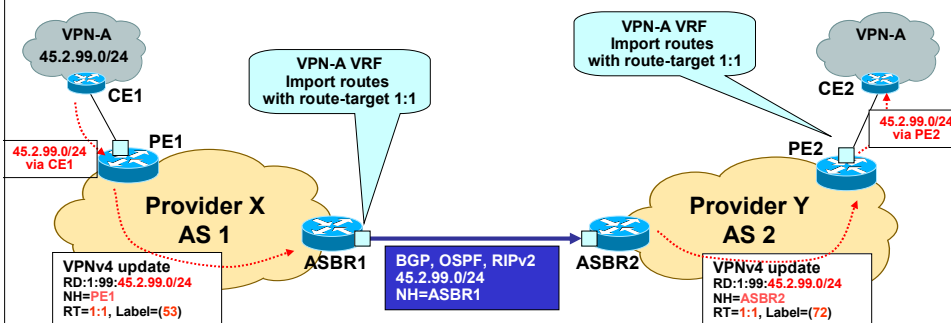


- **MPLS VPN providers exchange routes across VRF interfaces**
 - ♦ Each VRF represents one particular VPN client
- **PE-ASBR routers treat each other as a CE**
 - ♦ Although both provider interfaces associated with a VRF
- **PE routers used as gateways for VPNv4 route exchange**
- **Any PE-CE routing protocol can be used between each two PE-ASBRs**
 - ♦ OSPF, BGP-4, RIPv2 or static routes

#1: Multiple VPNs

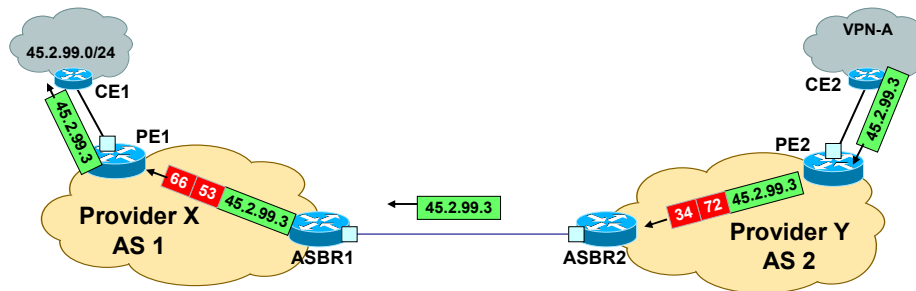


#1: Back-to-back VRF (Control Plane)



- VRF to VRF Connectivity between ASBRs

#1: Back-to-back VRF (Forwarding Plane)



- Transport labels (66 and 34) are learned via LDP/TDP

#1: Considerations



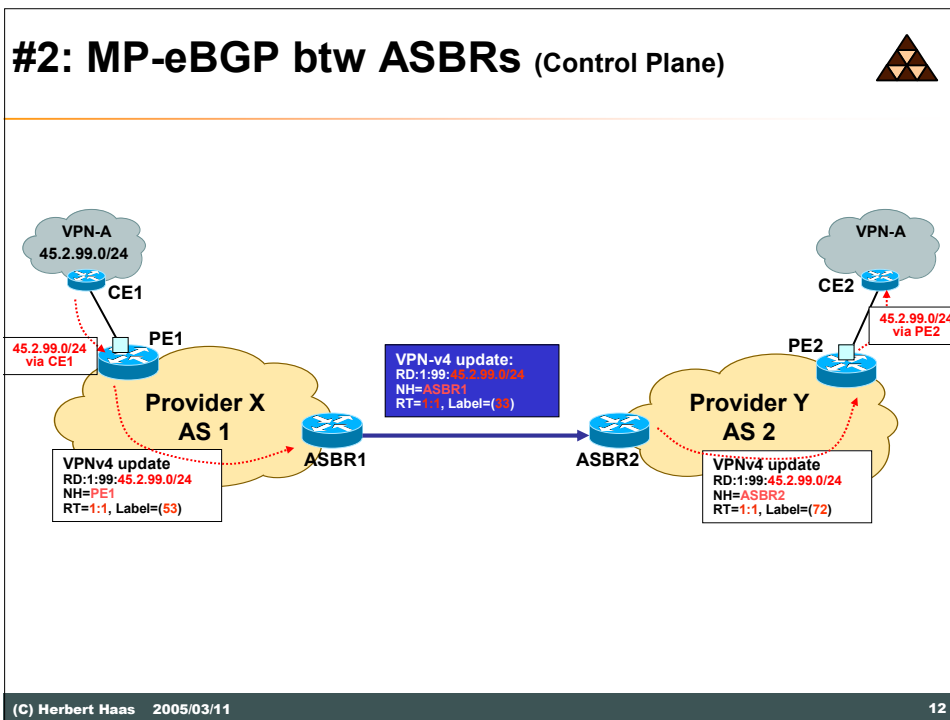
- Bad scalability
 - ◆ 1000 VRFs (for Inter-AS) would need 1000 interfaces (physical or logical) on both ASBRs
- No end-to-end MPLS
- Unnecessary memory overhead in RIB/(L)FIB
- But simple and backwards compatible
 - ◆ No need to load the Inter-AS code
 - ◆ Still not widely deployed

#2: MP-eBGP between ASBRs

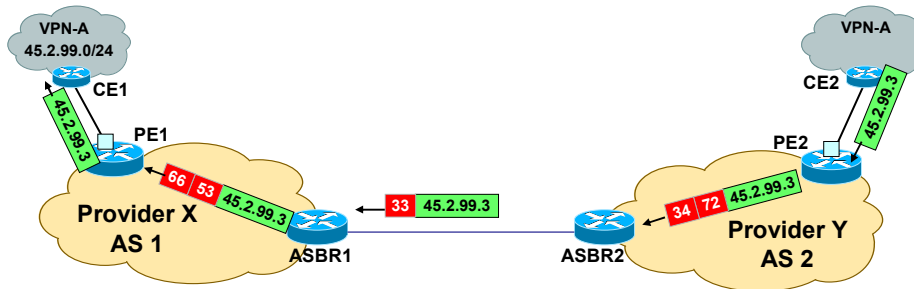


- ASBRs exchange VPN routes using eBGP (VPNv4 af)
 - ◆ No VRFs on ASBRs
 - ◆ Only LDP
 - ◆ New IOS option on ASBRs:
`no bgp default route-target filter`
- ASBRs store all VPN routes
 - ◆ But only in BGP table and LFIB table
 - ◆ No entry in routing table/CEF table

#2: MP-eBGP btw ASBRs (Control Plane)



#2: MP-eBGP btw ASBRs (Forwarding Plane)



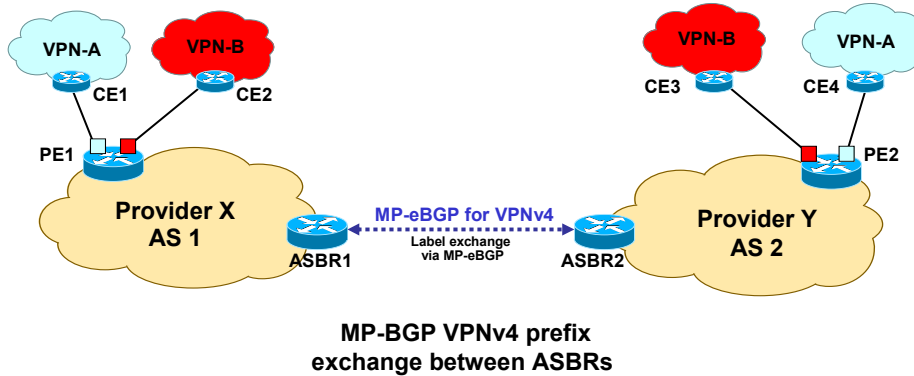
- Transport labels (66 and 34) are learned via LDP/TDP

#2: Basics



- ASBRs must provide “no bgp default route-target filter”
- ASBRs exchange routes directly using BGP
 - ♦ VPNv4 prefixes are exchanged via MP-eBGP
 - ♦ No LDP or IGP needed
- MP-BGP session has next-hop set to advertising ASBR
 - ♦ Next-hop and labels are rewritten when advertised across the Inter-Provider MP-BGP session
- ASBR stores all VPN routes in the BGP table that need to be exchanged
 - ♦ No entries in the routing/CEF table
 - ♦ Labels are populated into the LFIB of the ASBR
- No VRF is needed on the ASBRs

#2: MP-eBGP for VPNv4



#2: Considerations



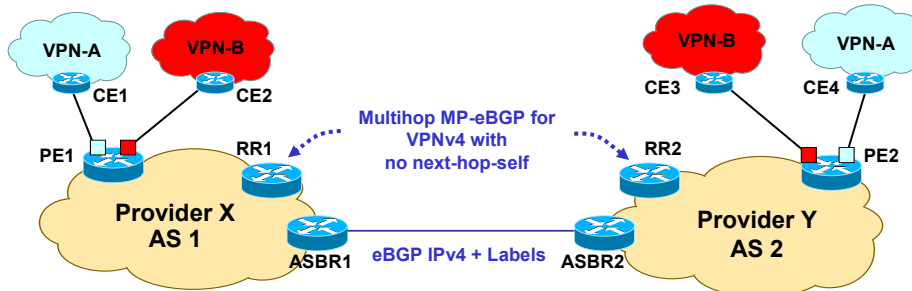
- **More scalable than solution #1**
 - ◆ No VRF configuration needed on ASBR
 - ◆ Less memory consumption (no RIB/FIB memory)
 - ◆ MPLS between providers
- **ASBRs are still required to hold VPN routes**
- **Automatic Route Filtering must be disabled**
 - ◆ However BGP filtering is allowed
- **Simple and working solution**

#3: Multihop MP-eBGP between RRs



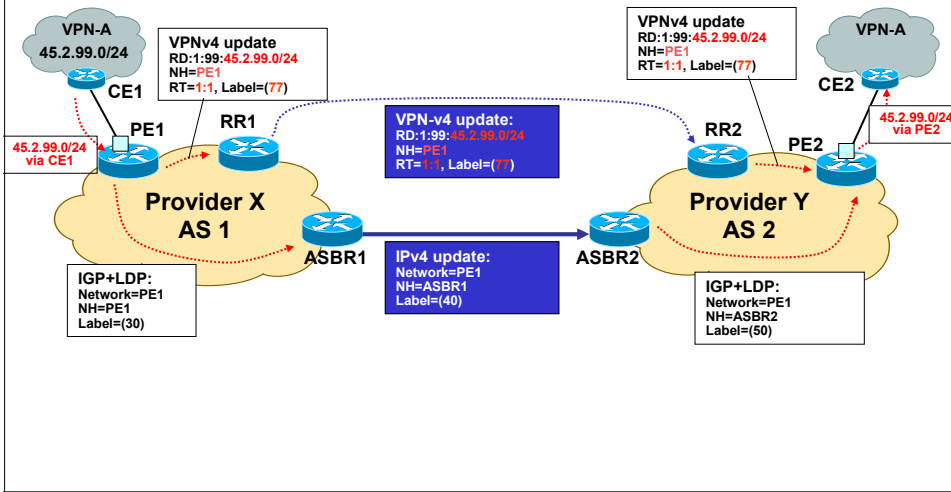
- Multihop MP-eBGP between RRs to exchange VPNv4 routes
 - ♦ Using next-hop-self option
- Normal eBGP between directly connected ASBRs to exchange IPv4 routes + labels
 - ♦ Only exchanges PE loopback addresses which are BGP next-hops for the VPN routes

#3: MP-eBGP btw RRs

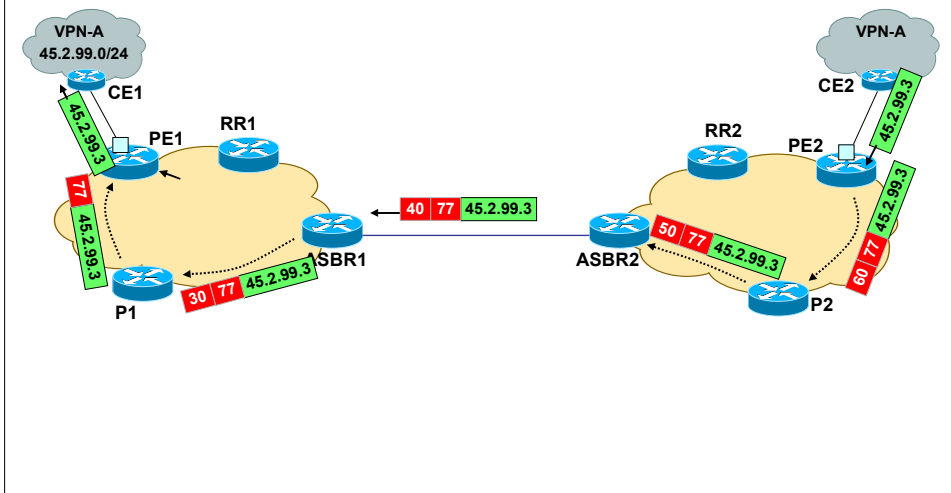


Multihop MP-eBGP VPNv4 prefix exchange between Route Reflectors

#3: MP-eBGP btw RRs (Control Plane)



#3: MP-eBGP btw RRs (Forwarding Plane)

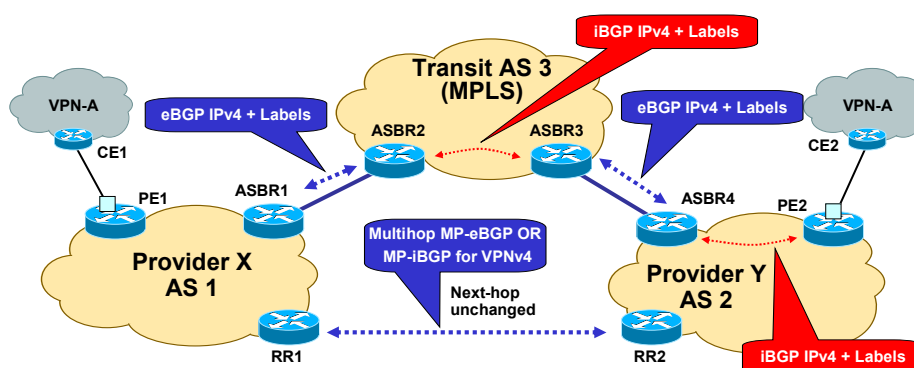


#3: Considerations



- RRs exchange VPNv4 routes plus associated labels
 - ◆ Since RRs know the VPNv4 information anyway
- ASBRs only exchange IPv4 routes plus forwarding labels
 - ◆ Only care about forwarding of MPLS packets
- Good scalability
 - ◆ Better than solutions 1 and 2
 - ◆ Allows separation of control and forwarding planes
- But advertising PE addresses to another provider might not be achievable !!!

#4: Non-VPN Transit Provider



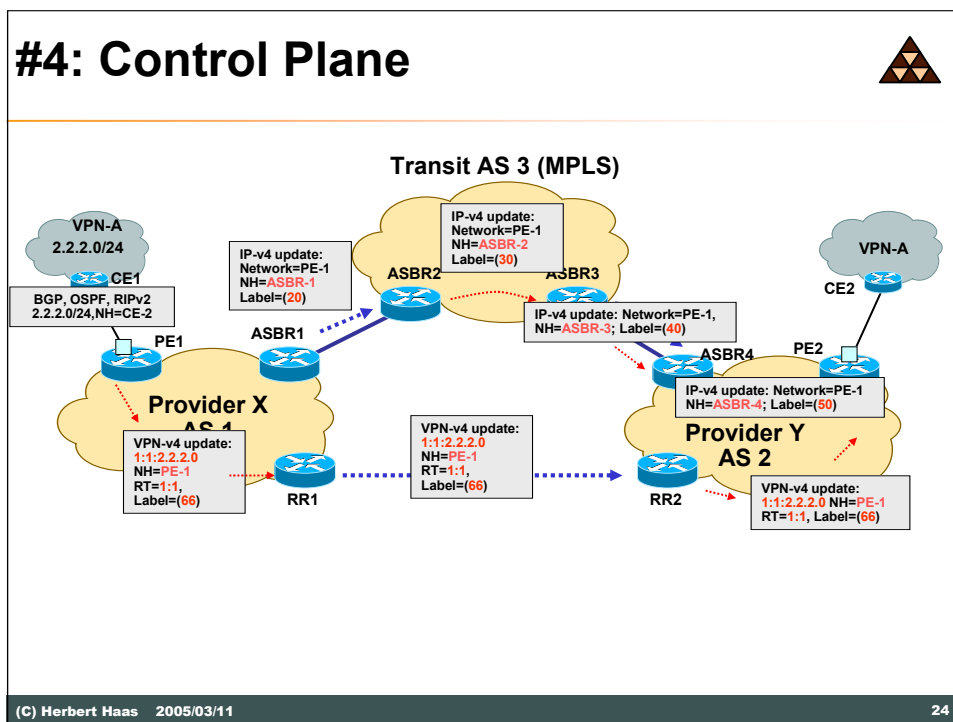
- Transit providers: Only need to run MPLS
- Edge providers run Multihop MP-eBGP
 - ◆ Exchange BGP next-hops via transit provider

#4: Non-VPN Transit Provider

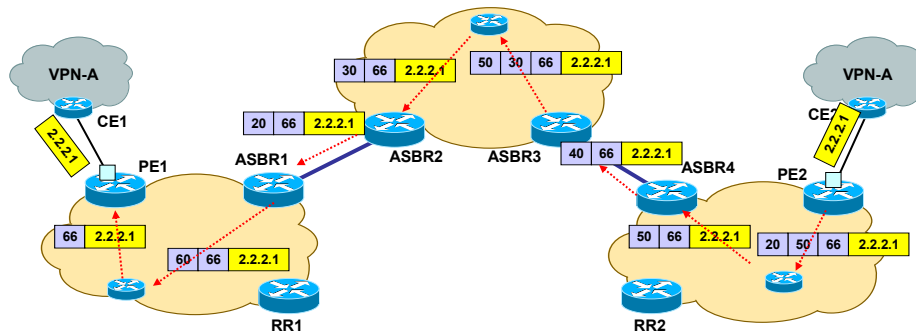


- Requirement to propagate BGP next-hops and also build end-to-end LSPs
- Four options for end-to-end LSP creation:
 1. Merge IGP of all AS's including the transit network
 2. Redistribute PE host routes between AS's
 3. Use static routes across boundaries and redistribution into IGP
 4. Use IPv4 + labels

#4: Control Plane



#4: Forwarding Plane



Inter-AS Deployment Guidelines



- Carry AS-number in the RT
- Set max-prefix limit (both BGP and VRF) on the PEs
- Strengthen security on ASBRs
 - ◆ Enable BGP MD5, BGP filtering ...
- Specify end-to-end QoS on ASBRs
- Rewrite RT on ASBRs (see next slide)
- Consider Internet connectivity on same ASBR