

VPNs with MPLS

“Traditional” VPN overlay model:

MPLS-based Layer 2 VPNs
draft-kompella-mpls-l2vpn-02.txt

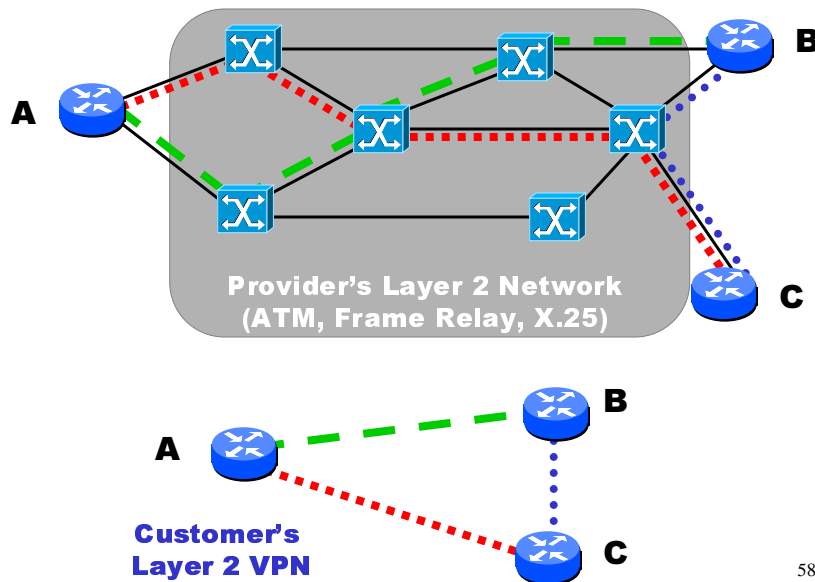
Whither Layer 2 VPNs?
draft-kb-ppvnp-l2vpn-motiv-00.txt

New VPN peering model:

RFC 2547. BGP/MPLS VPNs


57

Traditional Overlay VPNs



58

Why Not Use MPLS Tunnels?



59

Potential Advantages of MPLS Layer 2 VPNs

- **Provider needs only a single network infrastructure to support public IP, and VPN services, traffic engineered services, and differentiated services**
- **Additional routing burden on provider is bounded**
- **Clean separation of administrative responsibilities. Service provider does MPLS connectivity, customer does layer 3 connectivity**
- **Easy transition for customers currently using traditional Layer 2 VPNs**

60


BGP/MPLS VPNs

- RFC 2547
- Is Peer Model of VPN (not Overlay)
- Also draft-rosen-rfc2547bis-02.txt
- Cisco configuration info :
 - <http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/vpn.htm>

AT&T's IPFR service is based on this RFC.

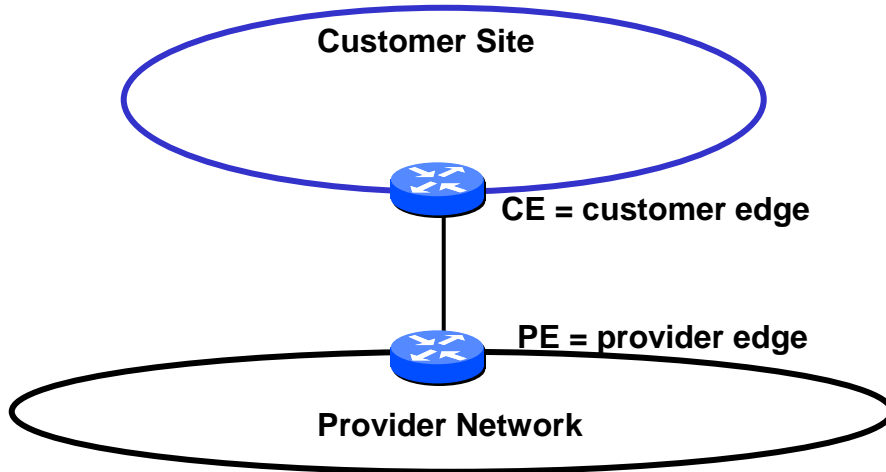
61

RFC 2547 Model



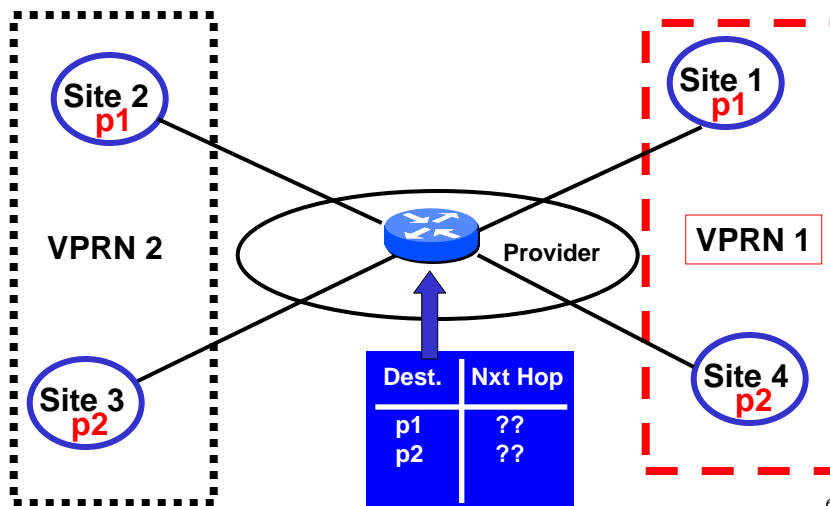
62

CEs and PEs



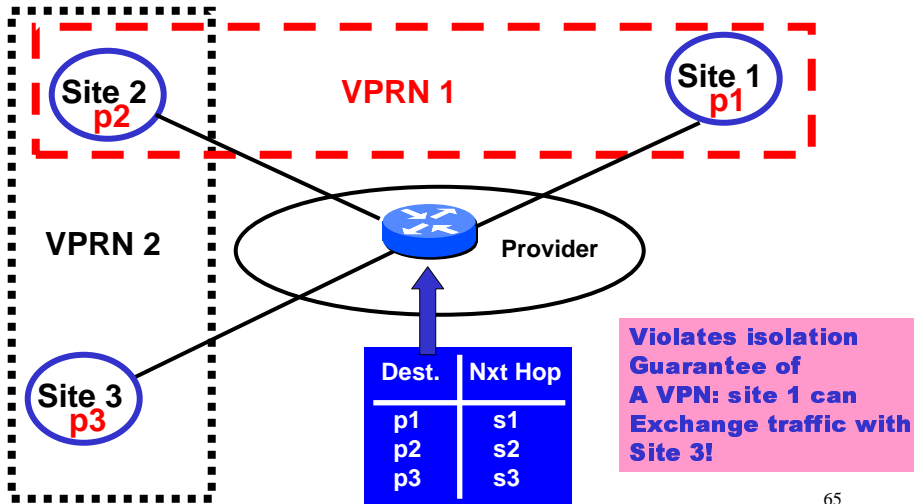
63

VPN Address Overlap Means Vanilla Forwarding Tables Can't Work



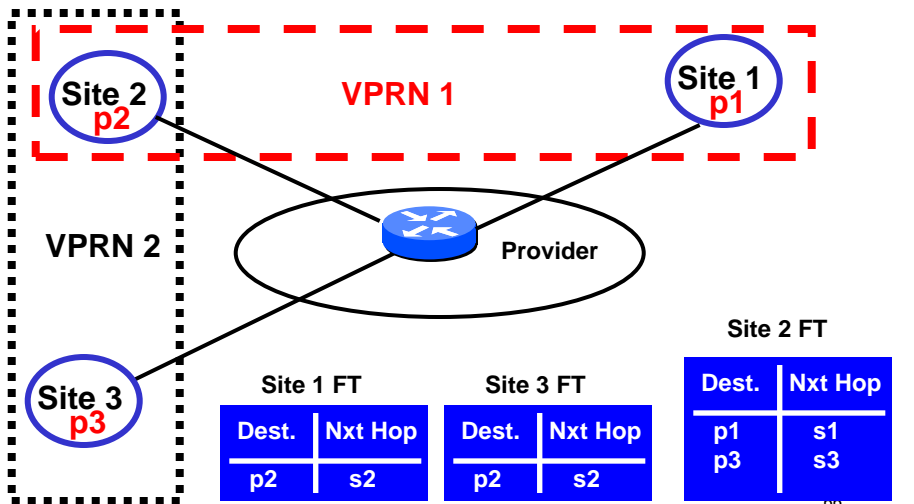
64

VPN Overlap Means Vanilla Forwarding Tables Can't Work

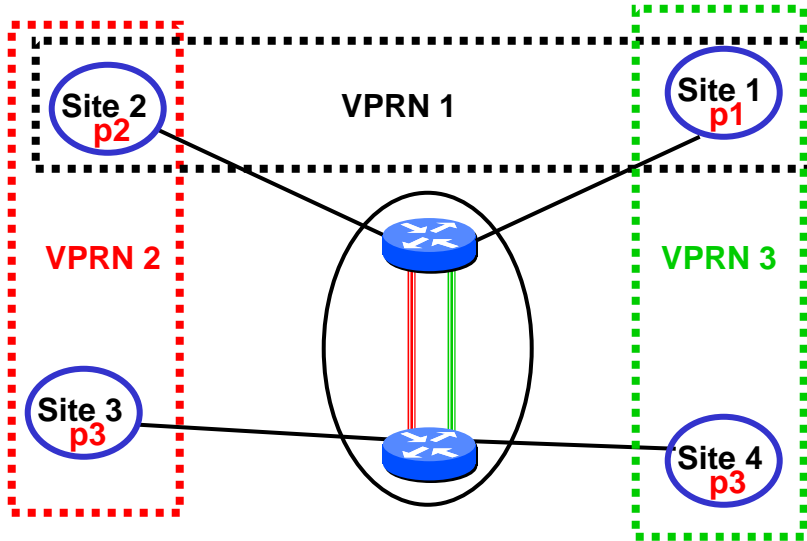


RFC 2547 : Per site forwarding tables

Called VRFs, for "VPN Routing and Forwarding" tables.

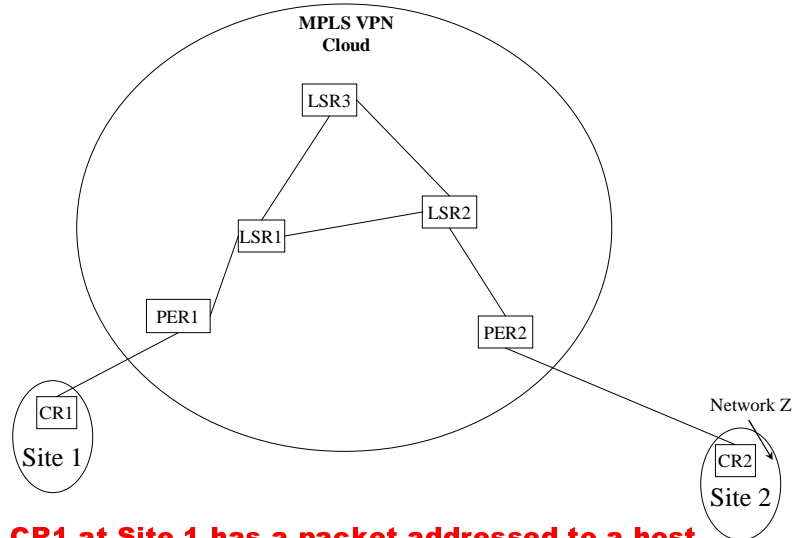


Tunnels required across backbone



67

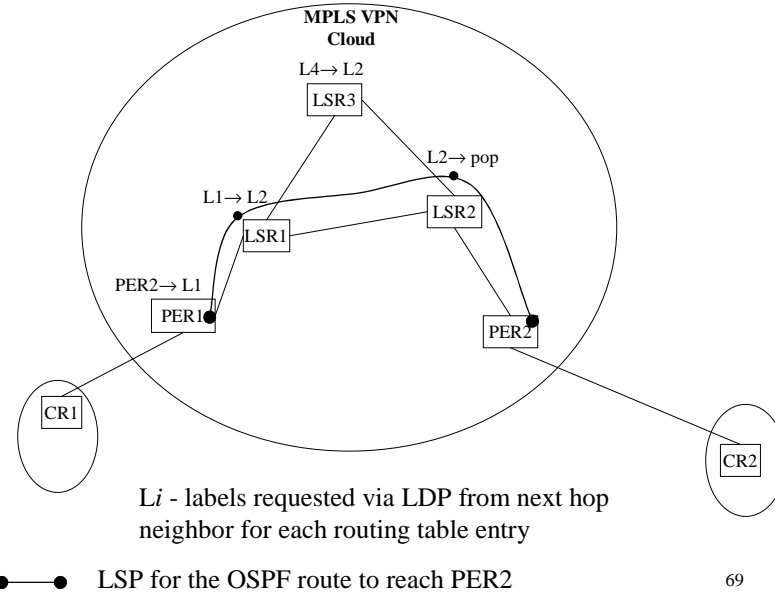
Follow the Route and Follow the Packet



CR1 at Site 1 has a packet addressed to a host in network Z at Site 2. How does it get there?

68

LSP Setup for OSPF Route to PER2



69

How MPLS VPNs work

1) Follow the routes

- Each VPN on a PER has a private routing table
 - Called a Virtual Routing Forwarding (vrf) table
 - vrf is assigned attributes that are unique to the VPN
 - Route Targets (RT) - attached to VPN routes.
 - » only vrfs with common RTs share routes
 - Route Distinguishers (RD) - appended to routes to ensure uniqueness even if VPNs have overlapping address spaces
 - » Creates a new address family called vpv4 = RD+ipv4
- NOTE: RTs and RDs are applied to routes, **NOT packets**

2) Follow the packet

- A stack of two labels is used to forward the packet on the interior LSP and then external interface

70

VPN extensions

- **Route Target (RT)**
 - **BGP 64 bit extended community value**
 - **First 16bit identify as RT type.**
 - **Other 48 bit is variable**
 - **Conventional format – ASN:X, i.e., 16b:32b**

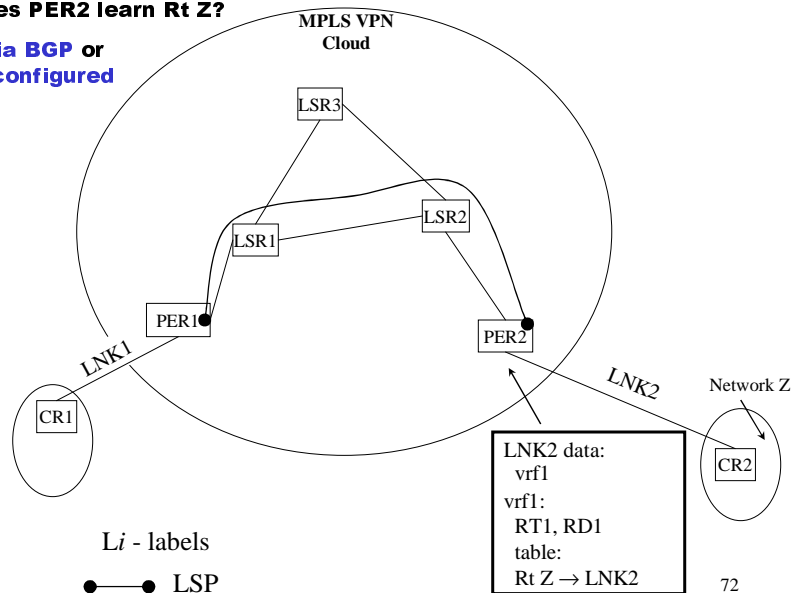
- **Route Distinguisher (RD)**
 - **BGP 64 bit extended community value**
 - **First 16bit identify as RD type.**
 - **Other 48 bit is variable**
 - **Conventional format – ASN:X, i.e., 16b:32b**

71

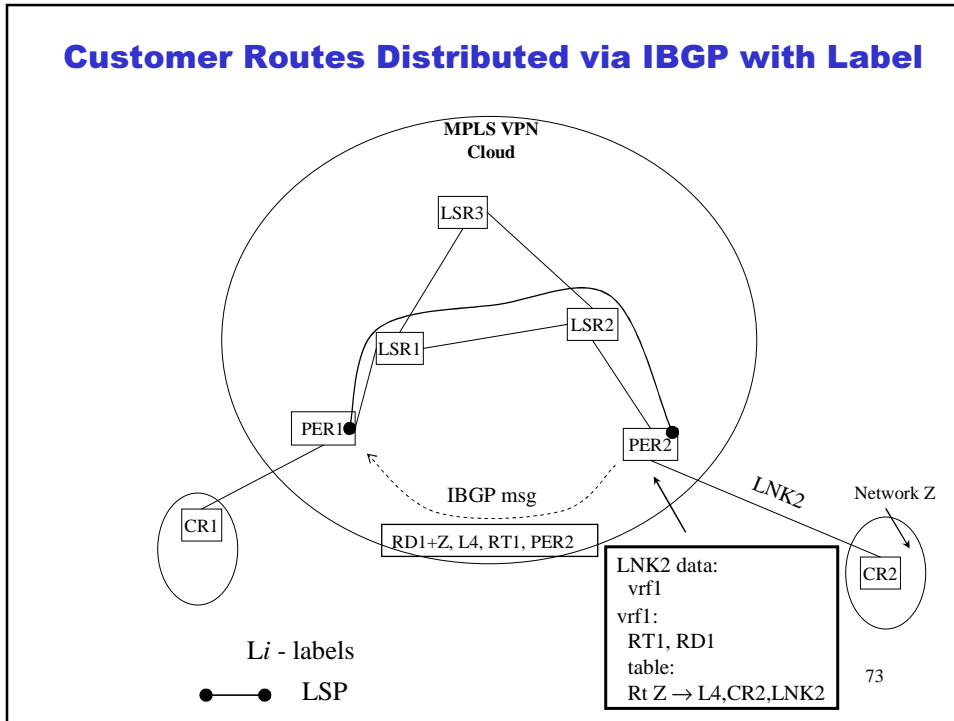
Distributing Customer Routes

Q: How does PER2 learn Rt Z?

A: Either via BGP or statically configured

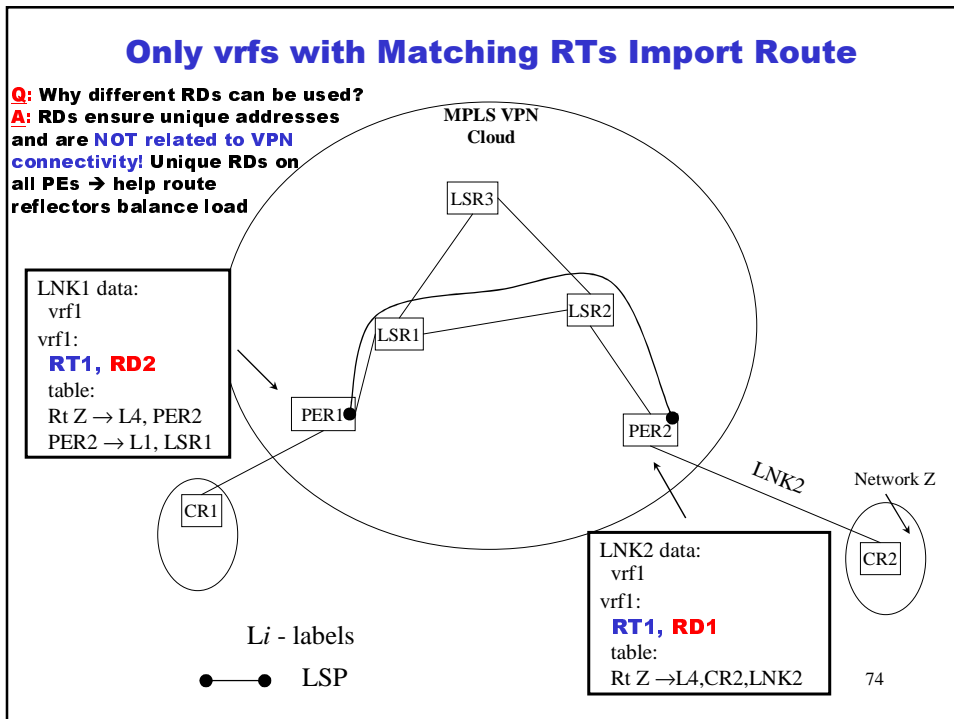


Customer Routes Distributed via IBGP with Label



Only vrfs with Matching RTs Import Route

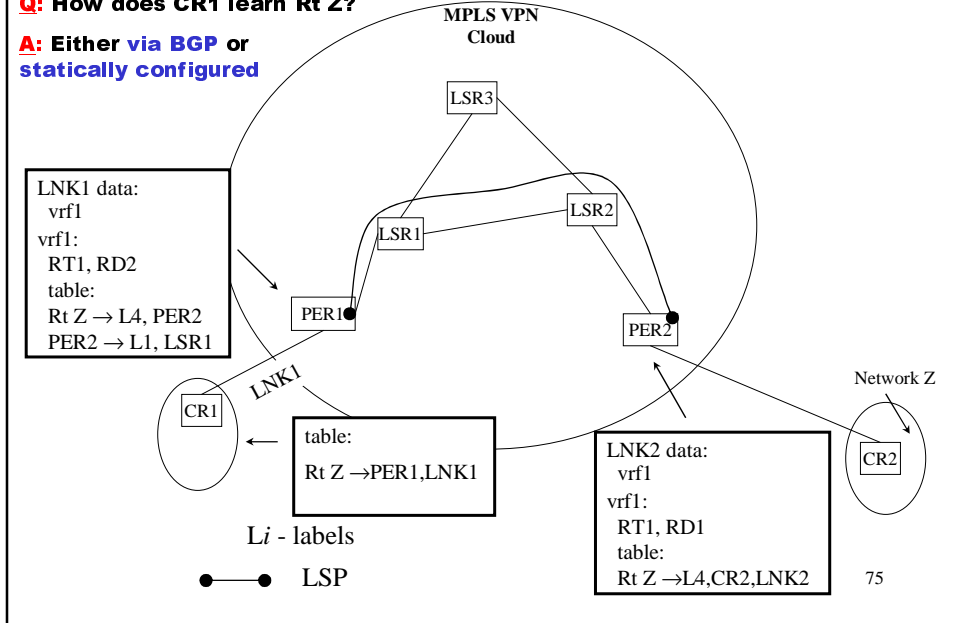
Q: Why different RDs can be used?
A: RDs ensure unique addresses and are **NOT** related to VPN connectivity! Unique RDs on all PEs → help route reflectors balance load



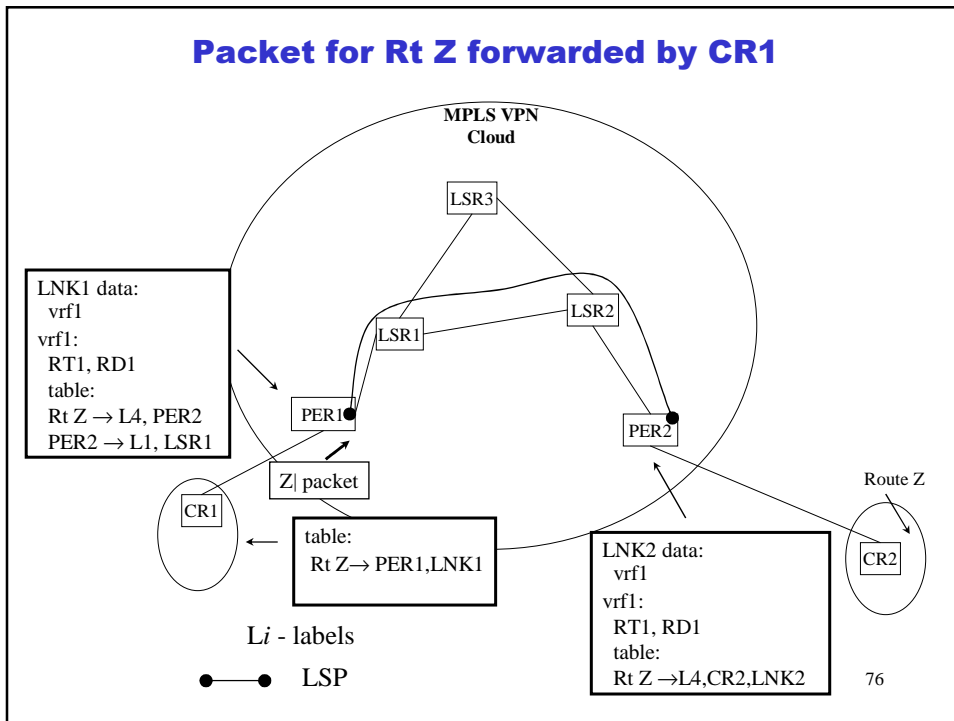
CR1 learns RT Z

Q: How does CR1 learn Rt Z?

A: Either via BGP or statically configured

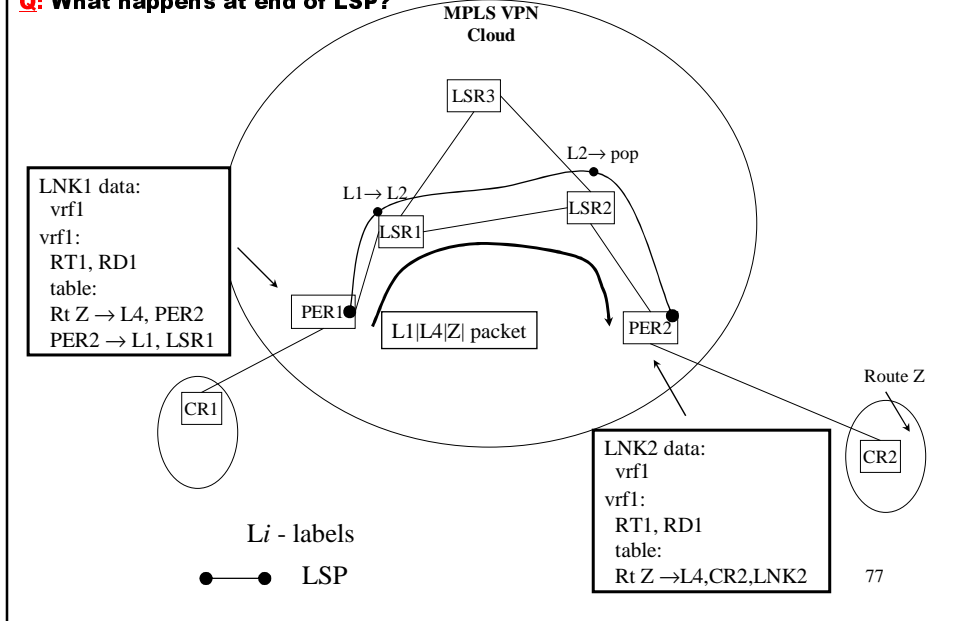


Packet for Rt Z forwarded by CR1



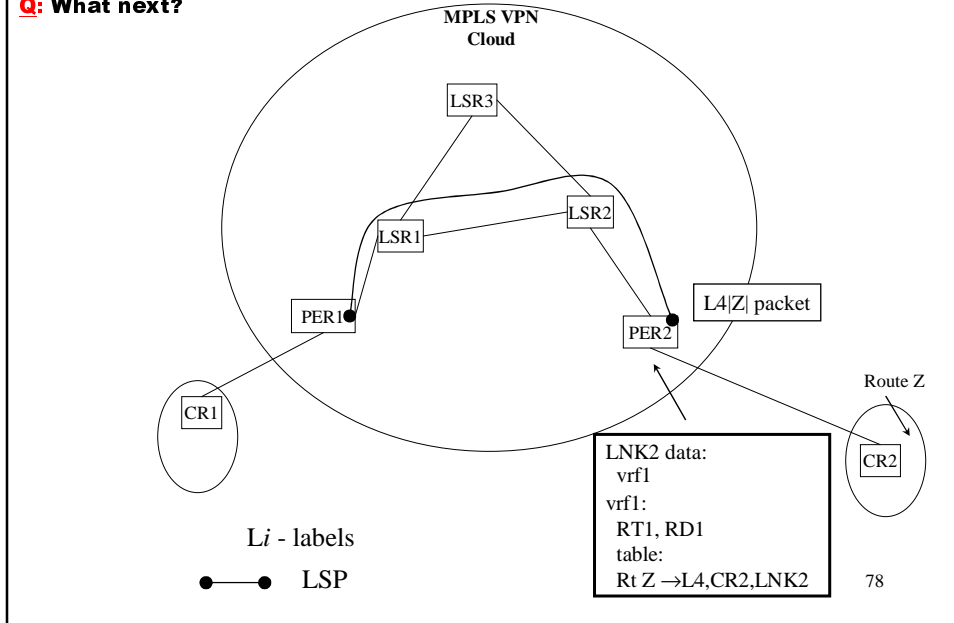
Top label is label-switched through interior

Q: What happens at end of LSP?



Top label popped at end of LSP

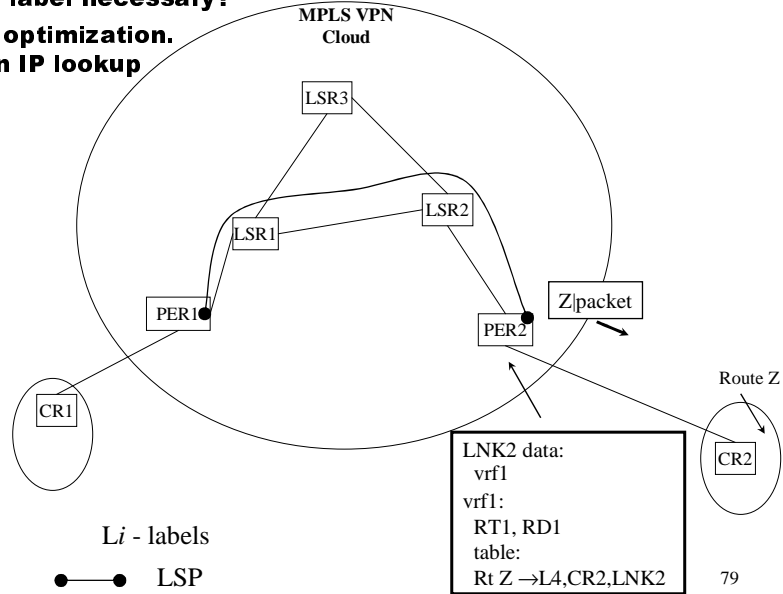
Q: What next?



Inner label determines egress interface and then is popped

Q: Is inner label necessary?

**A: No. An optimization.
It saves an IP lookup**



Purpose of BGP Label

- Indicates which vrf and optionally which interface on the egress PER
- Locally, the egress PER will treat labels in two possible ways:
 - Non-aggregate label is associated with an external route
 - Will be switched directly to an outgoing interface
 - IP header is not examined
 - Aggregate label is associated with a locally originated or directly connected route
 - Packet will be looked up in the vrf context

MPLS in Core Not Needed

- **MPLS for IGP domain serves as a tunneling method among PERs**
- **Could use other tunneling methods**
- **Advantages to MPLS:**
 - **Full mesh of LSP tunnels automatically created**
 - **Can use MPLS TE**
- **Internet draft to use IP or GRE tunneling**
 - **Automatically (treat vpnv4 BGP next hop as a recursive encapsulation)**
 - **BGP/IPsec VPN**
<draft-declercq-bgp-ipsec-vpn-00.txt>

81

RFC 2547 Summary

- **Piggyback VPN information on BGP**
 - **New address family**
 - **New attributes for membership**
- **New Per-site forwarding tables (VRFs)**
- **Use MPLS Tunnels between PEs**
 - **No need for VPN routes on backbone LSRs, only on PEs**

82

MPLS VPN Security

- **Private routing table for each VPN (vrf)**
- **VPN membership identity associated with each access connection**
 - **VPN membership is not determined by IP header, only by interface (e.g., DLCI, VPI/VCI, PPP, VLAN tag).**
 - **Label and RT for VPN attached to routes advertised for interface.**
 - **Route and its matching label are only imported by routing tables that match the VPN RT.**
 - **Impossible for a packet on a PVC in one vrf to spoof its way or jump into another vrf**

83

Layer 2 VPNs vs. BGP/MPLS VPNs

- | | |
|---|---|
| <ul style="list-style-type: none">• Customer routing stays with customer• May allow an easier transition for customers currently using Frame/ATM circuits• Familiar paradigm• Easier to extend to multiple providers | <ul style="list-style-type: none">• Customer routing is “outsourced” to provider• Transition may be complicated if customer has many extranets or multiple providers• New “peering” paradigm• Not clear how multiple provider will work (IMHO) |
|---|---|

84

Summary

MPLS is an interesting and potentially valuable technology because it

- **provides an efficient and scalable tunneling mechanism**
- **provides an efficient and scalable mechanism for extending IP routing with explicit routes**

85

More info on MPLS

- **MPLS working group**
 - <http://www.ietf.org/html.charters/mpls-charter.html>
- **MPLS email list archive**
 - <http://cell.onecall.net/cell-relay/archives/mpls/mpls.index.html>
- **MPLS Resource Center**
 - <http://www.mplsrc.com>
- **Peter Ashwood-Smith's NANOG Tutorial**
 - <http://www.nanog.org/mtg-9910/mpls.html>
- **MPLS: Technology and Applications. By Bruce Davie and Yakov Rekhter. Morgan Kaufmann. 2000.**
- **MPLS: Is it all it's cracked up to be? Talk by Pravin K. Johri**
 - <http://buckaroo.mt.att.com/~pravin/docs/mpls.pdf>

86

More info on MPLS TE

- **tewg working group**
 - <http://www.ietf.org/html.charters/tewg-charter.html>
- **NANOG Tutorial by Jeff Doyle and Chris Summers**
 - <http://www.nanog.org/mtg-0006/mpls.html>
- **NANOG Tutorial by Robert Raszuk**
 - <http://www.nanog.org/mtg-0002/robert.html>

87

More info on MPLS VPNs

- **PPVPN working group**
 - <http://www.ietf.org/html.charters/ppvpn-charter.html>
- **PPVPN Archive**
 - <http://nbvpn.francetelecom.com>
- **NANOG Panel: Provider-Provisioned VPNs**
 - <http://www.nanog.org/mtg-0102/jessica.html>
- **MPLS and VPN Architectures. By Ivan Pepelnjak and Jim Guichard. Cisco Press. 2001**

88