

Foundry Networks

Principal member of the MPLS Forum

Challenges Facing L2 VPN

MPLS Japan 11/8/2001

Gary Blankenship
Foundry Networks
Systems Engineer - Japan



Agenda

- ➔ L2 Technologies in MAN (Metro)
 - MPLS VLL Draft-Martini
 - MPLS VPLS
 - Interoperability

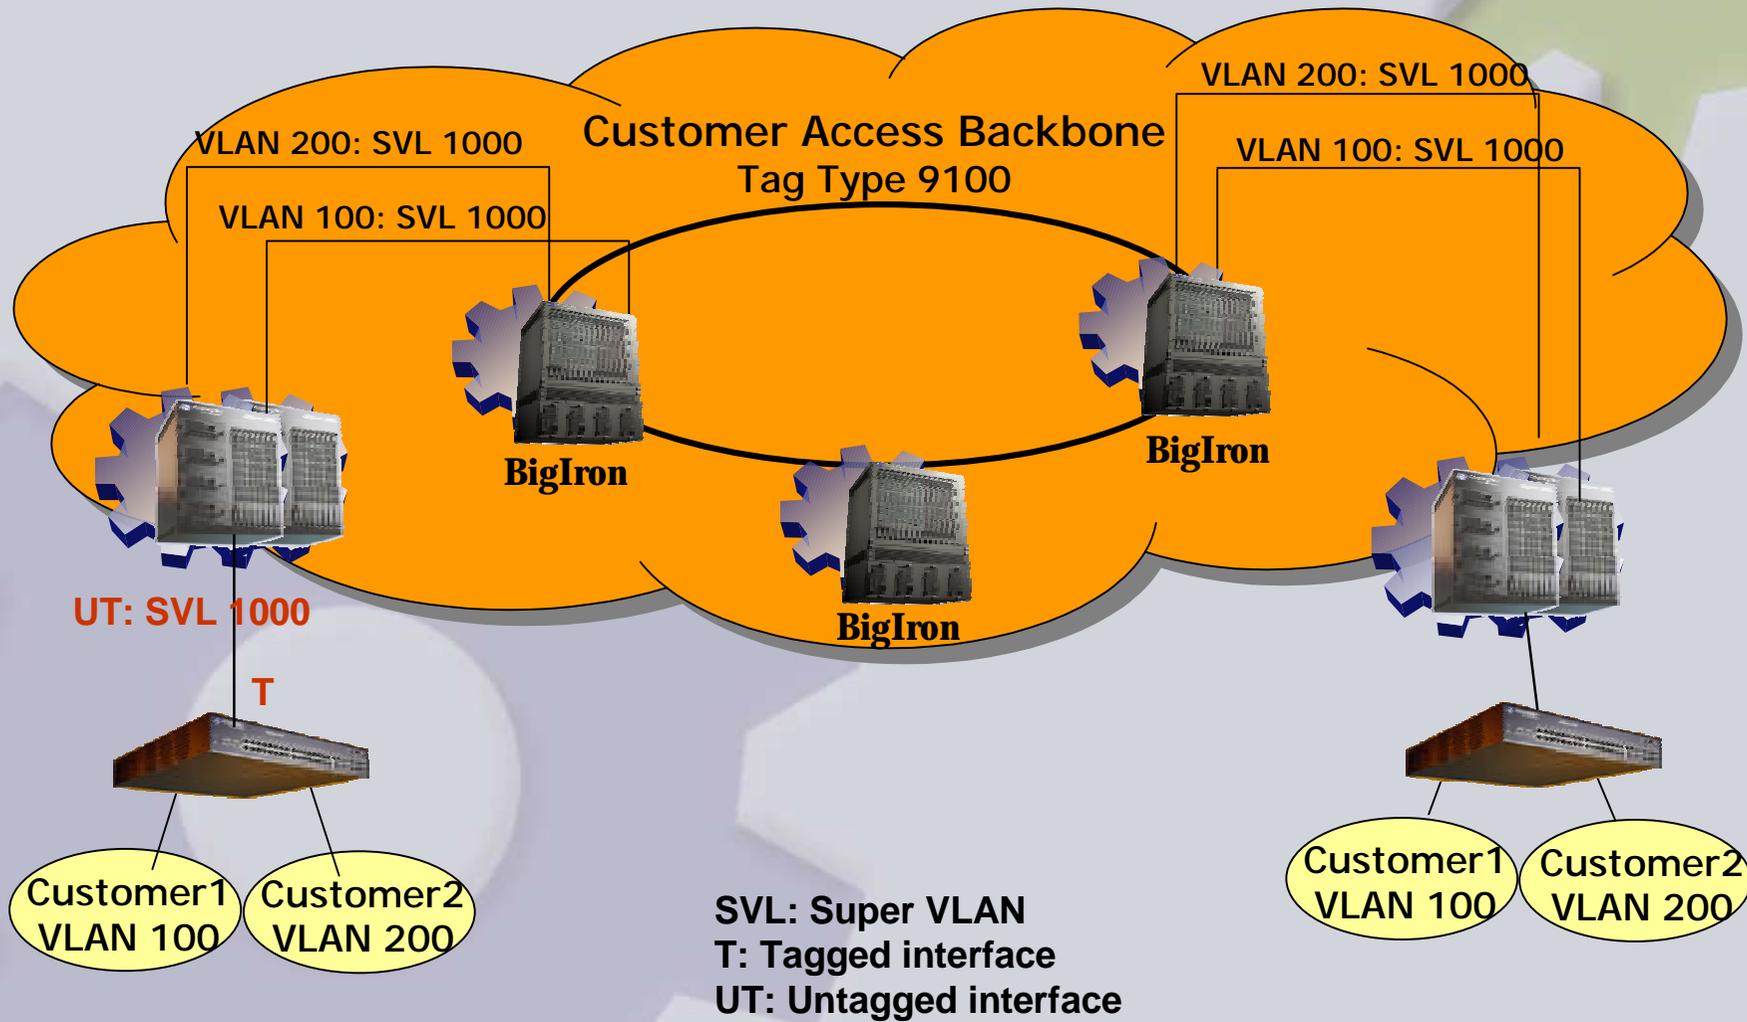
L2 Technologies in Metro Area Networks (MAN)

- 802.1Q allows a maximum of 4,096 unique VLANs
- SuperVLAN. A stack of two 802.1Q headers to be carried in an Ethernet frame. 16 million unique VLANs
 - Potential solution through VLAN tag stacking – 4096x4096 VLANs
 - Can work with Spanning Tree
 - Customers can define their own VLAN ID
 - Core of the network remains limited to 4,096 VLANs
 - Large number of MAC addresses in the core
 - Management of the VLAN hierarchy can become rather complex

L2 Technologies in Metro Area Networks (Cont.)

- SuperVLAN
 - Lacks visibility of individual customer traffic within the network core
 - Customers VLANs within the same Super VLAN must share its resources (Aggregation = Loss of 802.1p information)
 - Broadcast/unknown unicast traffic may be forwarded to more edge switches than necessary
 - Limited means for resources allocation and QoS treatment from 802.1Q
 - STP provides only crude traffic placement ability and is not straightforward to manage
- 802.1w – **Rapid Spanning Tree Protocol** for rapid convergence time of 1 to 5 seconds
- 802.1s – **STP per VLAN Group** provides VLAN and STP scalability, and utilizes dark fiber efficiently

Super VLAN Aggregation Solution to VLAN Growth

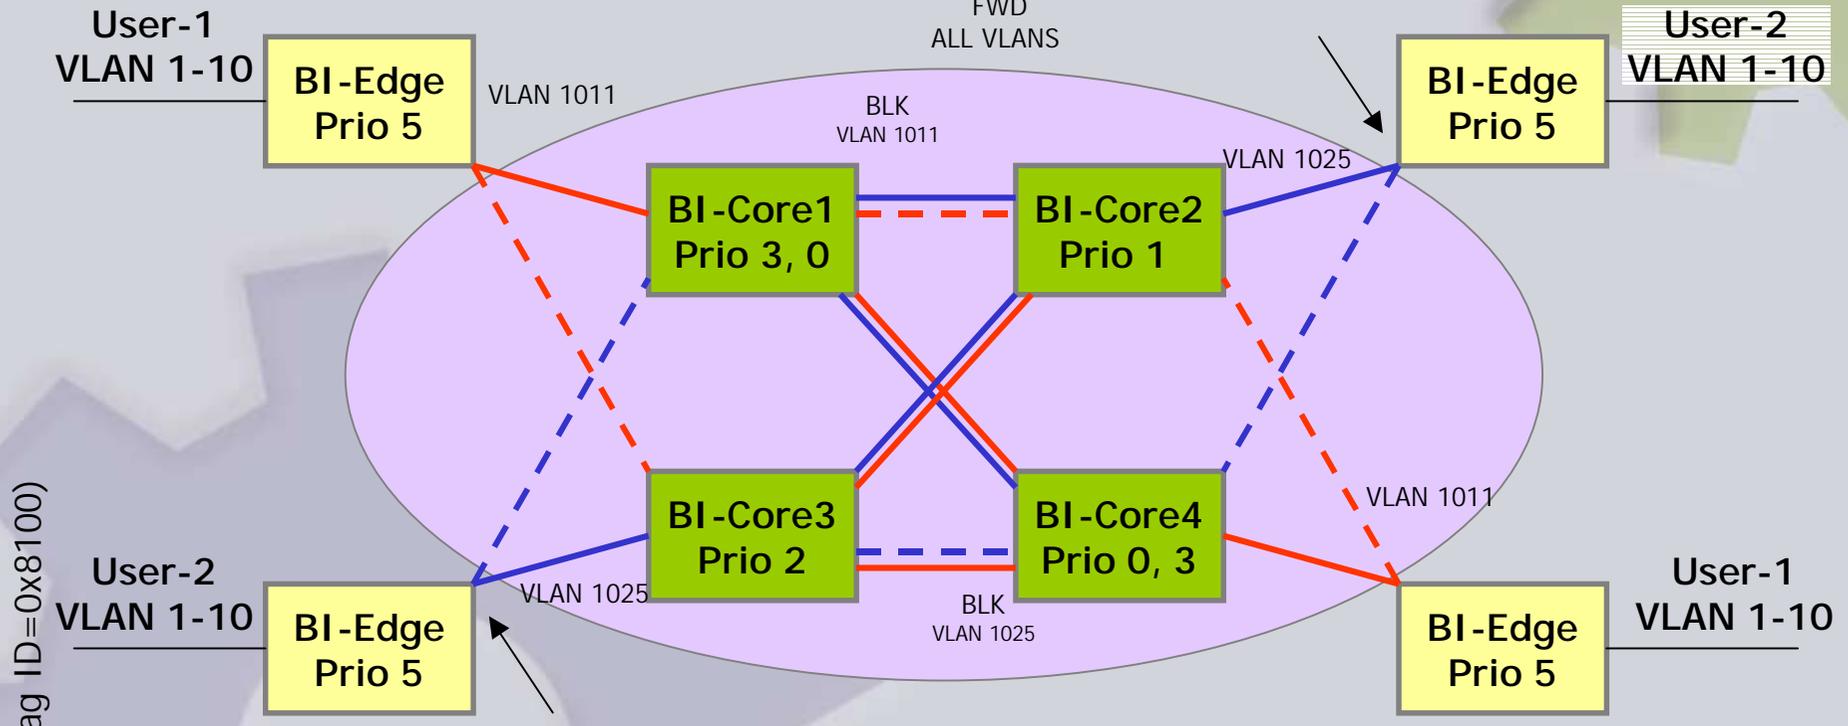


How 802.1s Works

STP with VLAN Grouping

Using 802.1 S
To Support 4 to 8 groups of Spanning Tree

Spanning Tree in Service
Provider's network
(Tag ID=0x9100)



(Tag ID=0x8100)

Spanning Tree in Service
Provider's network
(Tag ID=0x9100)

- VLAN 1-1024 -> VLAN Group 1 = —
 - VLAN 1025-2048 -> VLAN Group 2 = —
 - VLAN 2049-3072 -> VLAN Group 3 = —
 - VLAN 3073-4096 -> VLAN Group 4 = —
- All Core Switches forwarding unless otherwise noted
- Dash line indicates blocked state



Foundry's SuperSpan™

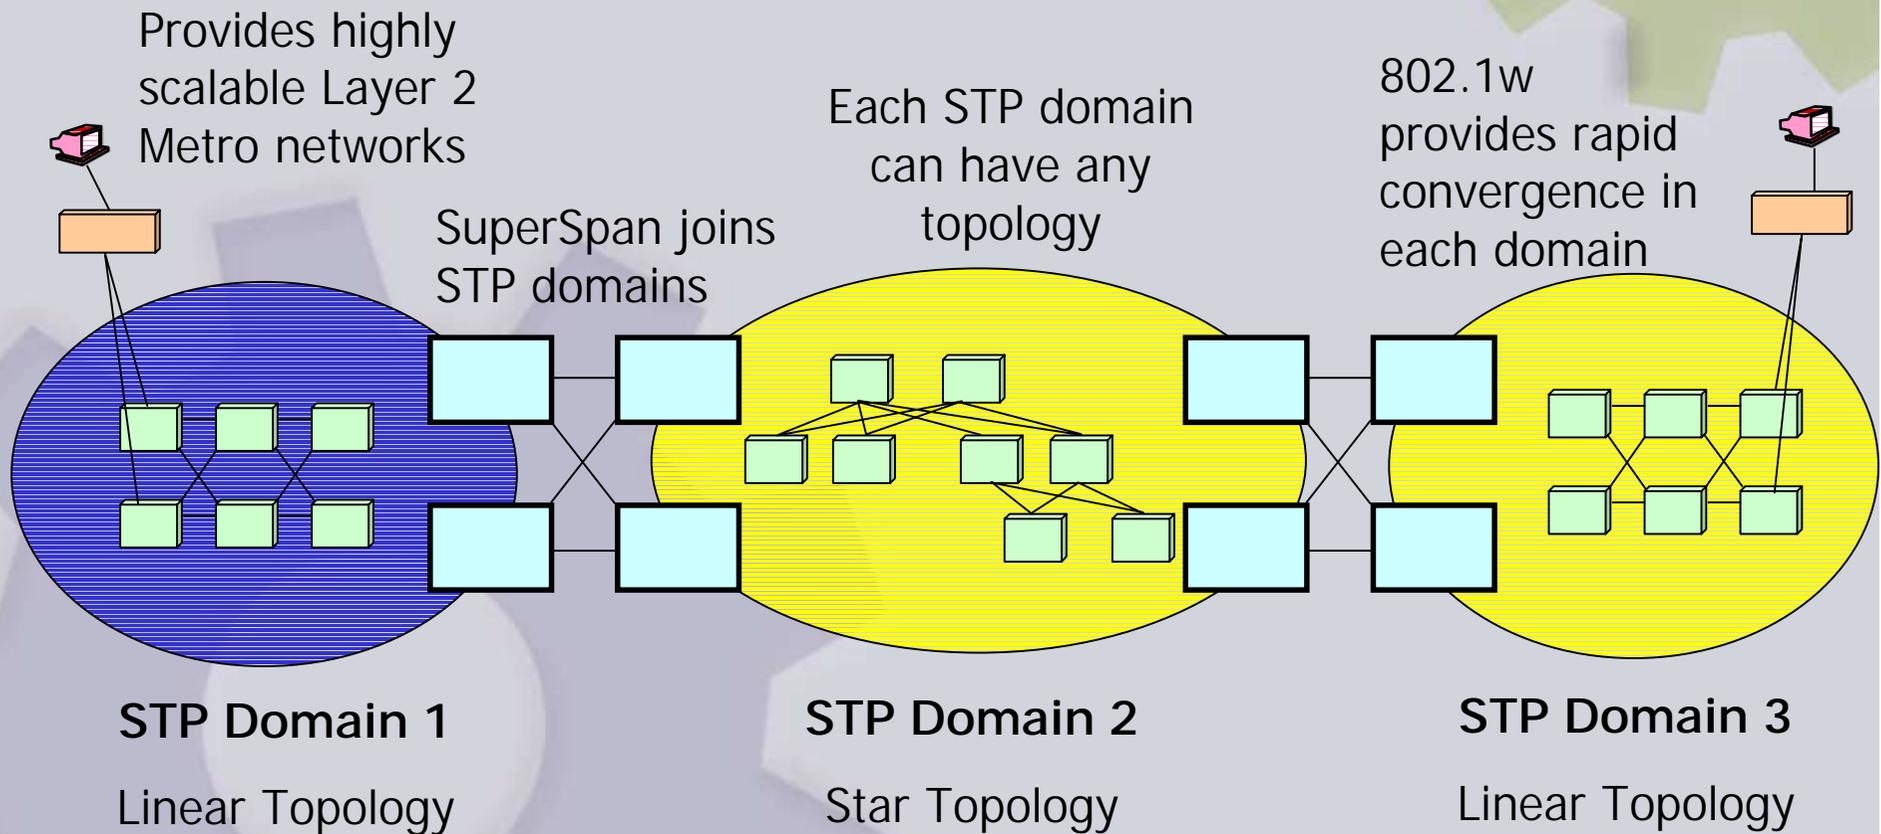
A Breakthrough For Layer 2 Network Scalability and High Availability

- Organizes Layer 2 Metro Core into smaller STP domains that are easy-to-manage and simple to configure
- Localizes STP to the domain – Provides rapid convergence within each domain with 802.1w Rapid STP
- Dramatic improvements in scalability – By joining simple, faster-converging STP domains, SuperSpan allows highly scalable large Layer 2 Metro networks
- Improves High Availability – A link failure is confined to the STP domain as other STP domains will be unaffected
- Allows any flexible topology within each STP domain – star, ring or mesh networks
- Interoperates with Super VLAN aggregation to allow end-to-end VLAN based services



How Foundry's SuperSpan Works

Easy-to-manage, scalable, highly available
Layer 2 Metro Core Networks!



MPLS Layer-3 vs. Layer-2 VPNs

Layer-3

- Provider has to do some routing on the subscriber's behalf on the PE routers – extra burden for the provider
- Each PE router must have multiple routing and forwarding tables, one per customer site
- Introduces operational and management complexity
- Scalability challenges
- IP-only solution as it is a routed service

Layer-2

- Provider doesn't handle any routing on the subscriber's behalf, instead only simple mapping of Layer-2 frames to VCs is done
- No multiple routing and forwarding tables, only the global routing table
- Much simpler to operate and manage
- No scalability challenges in point to point mode
Scalability challenges in VPLS mode
- Subscriber IP addresses are opaque to the provider, that way overlapping address spaces are handling without any extra work on the provider side
- Solution allows for carrying any protocol contained within the Layer-2 frames

Agenda

- L2 Technologies in MAN
- ➔ MPLS VLL Draft-Martini
- MPLS VPLS
- Interoperability

What Is Virtual Leased Line/Draft-Martini???

- “Draft-Martini” refers to two IETF Internet Drafts co-authored by Luca Martini of Level3
- draft-martini-l2circuit-encap-mpls-03.txt
 - Defines encapsulation formats for ATM, FR, Ethernet, Ethernet VLAN, HDLC, and PPP
 - We focus on Ethernet and Ethernet VLAN for now
- draft-martini-l2circuit-trans-mpls-07.txt
 - Proposes an extension to Label Distribution Protocol (LDP)
 - Purpose is to coordinate encapsulations between two routers
 - Introduces a new VC FEC TLV, to be used within Label Mapping messages

VLL Feature Highlights

- Capable of selecting LSP tunnels based on COS criterion
- Load balancing among multiple LSPs, tunnel failover
- The two endpoints of a VLL can have
 - Mixed tagged and untagged
 - Different tagged VLAN IDs
- VLAN ID can be reused
 - Subject to same DMA restriction
- Label withdrawal when local port goes down

The Martini Drafts

- The Martini drafts introduce the concept of **Virtual Circuits (VCs)**
- LSP are used as tunnels, an LSP might carry multiple VCs
- To accomplish this, the drafts leverage the MPLS label stacking ability by adding an extra MPLS label that distinguishes the VC used
- An MPLS frame traversing a service provider network has two labels:

- Tunnel Label
- VC Label



The Martini Drafts

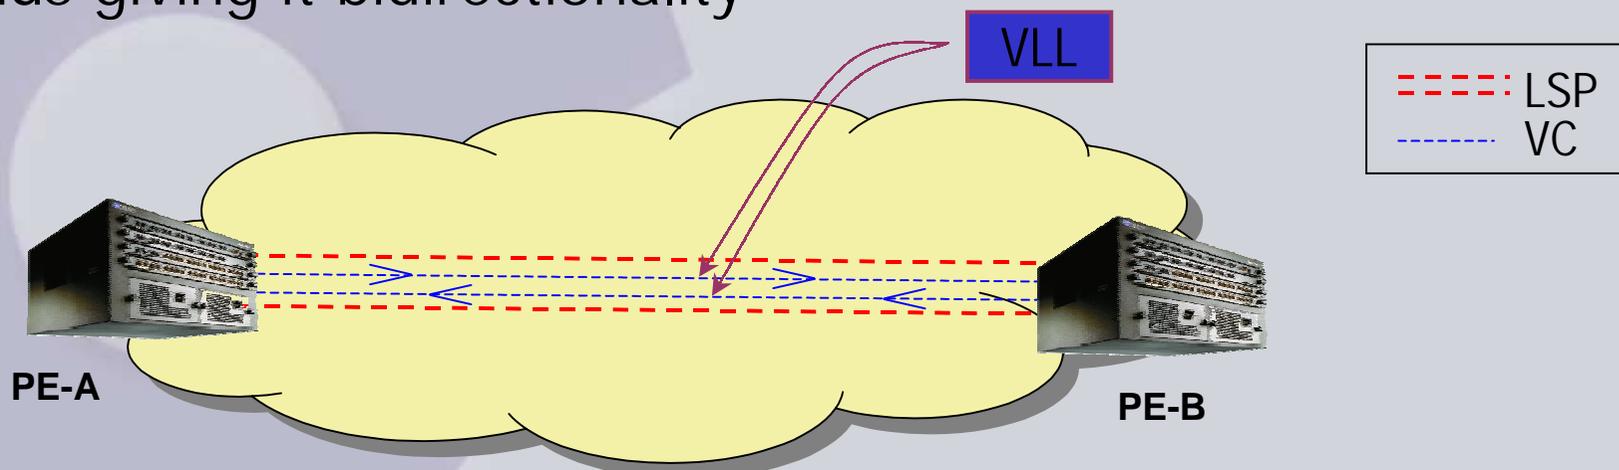
- draft-martini-l2circuit-trans-mpls-07.txt
 - Distribution of VC labels via LDP
 - Introduces a new VC FEC TLV, to be used within Label Mapping messages
- draft-martini-l2circuit-encap-mpls-03.txt
 - Defines encapsulations for ATM, FR, Ethernet, Ethernet VLAN, HDLC, and PPP

Format for Ethernet/VLAN Encapsulation:

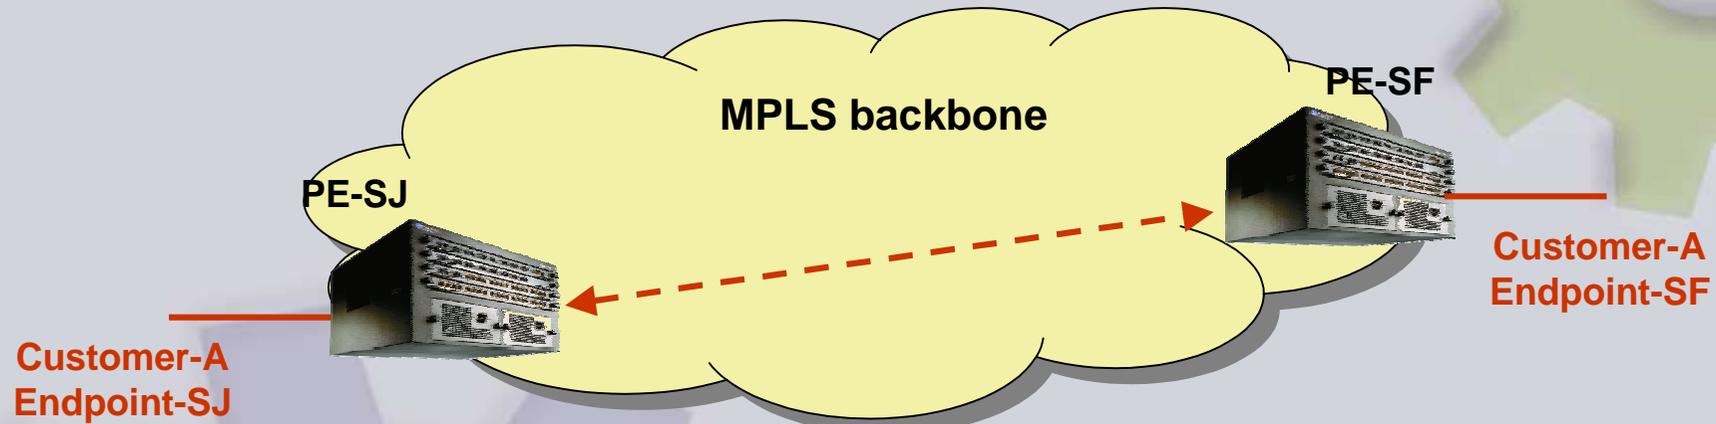
PPP-MPLS or Enet Hdr.	Tunnel Label	VC Label	Optional Control Word	Tagged or untagged Ethernet payload
--------------------------------------	-------------------------	---------------------	--------------------------------------	--

The Virtual Lease Line (VLL) Service

- The mechanisms specified in the Martini drafts could be used to offer a Virtual Leased Line (VLL) service for point to point connectivity
- The Virtual Leased Line is just an abstraction of the Virtual Circuits involved in realizing the point to point connectivity, making the service more manageable
- A VLL is composed of two VCs one in each direction, thus giving it bidirectionality



What Draft-Martini Is and Is NOT



- Layer-2 packet relay between two customer endpoints using MPLS encapsulation
- The VLL is only known to the two Provider Edge (PE) routers involved
- Point-to-Point ONLY
- Draft-Martini does not include Point-to-MultiPoint support – it's not transparent LAN switching (TLS)

What Draft-Martini Is and Is NOT (cont.)

- No switching means
 - PE routers do not have to care and learn about MAC addresses
 - Forwarding simplicity
 - No MAC table scalability concern

Anatomy of a VLL

What Needs to Be Configured?

- VLL name and ID
 - VLL ID must be unique on the two PE routers – it allows the two PEs to know which VLL they are referring to
 - Name is just for user reference and is not signaled
- VLL Peer
 - IP address of the other PE router
- Local Endpoint
 - Customer-facing interface
 - Can be either an untagged port or an tagged port with a specified VLAN-ID
 - Benefit of tagged endpoint: can have multiple endpoints on one physical port

Sample Configurations

- Untagged endpoint

```
router mpls
  vll foundry-sj-to-sf 40000
  vll-peer 192.168.2.100
  untagged e 2/1
```

VLL ID

- Tagged endpoint

```
router mpls
  vll foundry-sj-to-la 50000
  vll-peer 200.100.1.100
  vlan 20
  tagged e 3/1
```

- Options: statically configured label values, vll cos (class of service)

Anatomy of a VLL

How Does a VLL Transport Packets?

- Peers forming a single virtual pipe
- Peers exchange VC-label to VLL-ID binding through LDP session with Draft-Martini extension

“To send packets for VLL-ID 1000 to me, use VC label X”

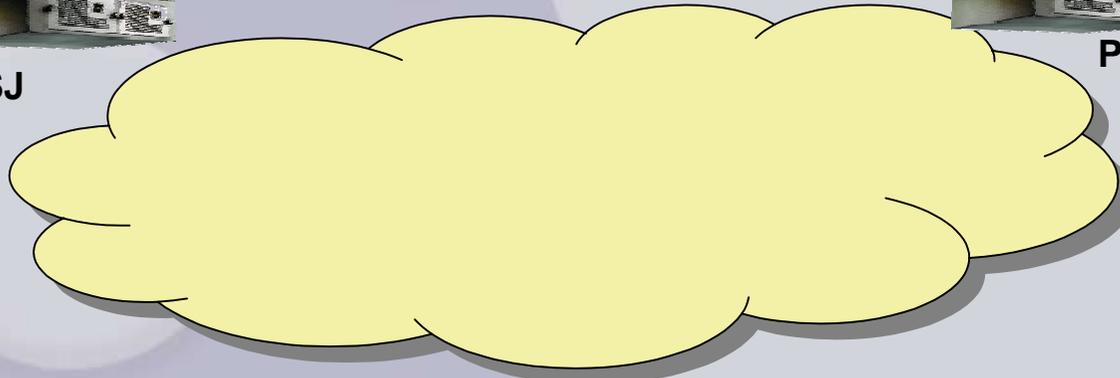


PE-SJ

“To send packets for VLL-ID 1000 to me, use VC label Y”



PE-SF



Anatomy of a VLL

How Does a VLL Transport Packets?

- LSP tunnels, established by RSVP, provide the transport through MPLS backbone

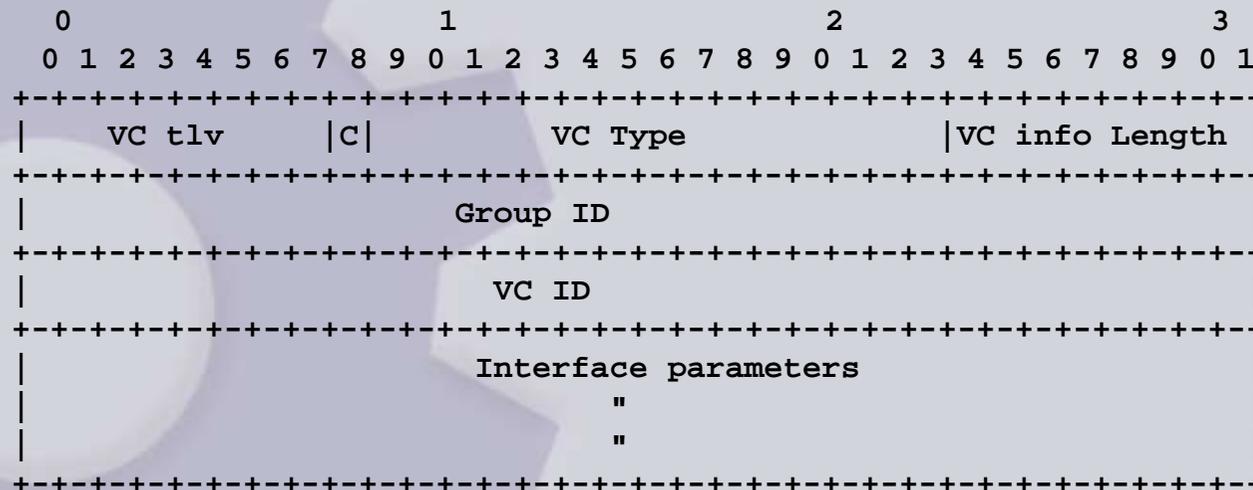


Outgoing MPLS packet generated by PE-SJ:

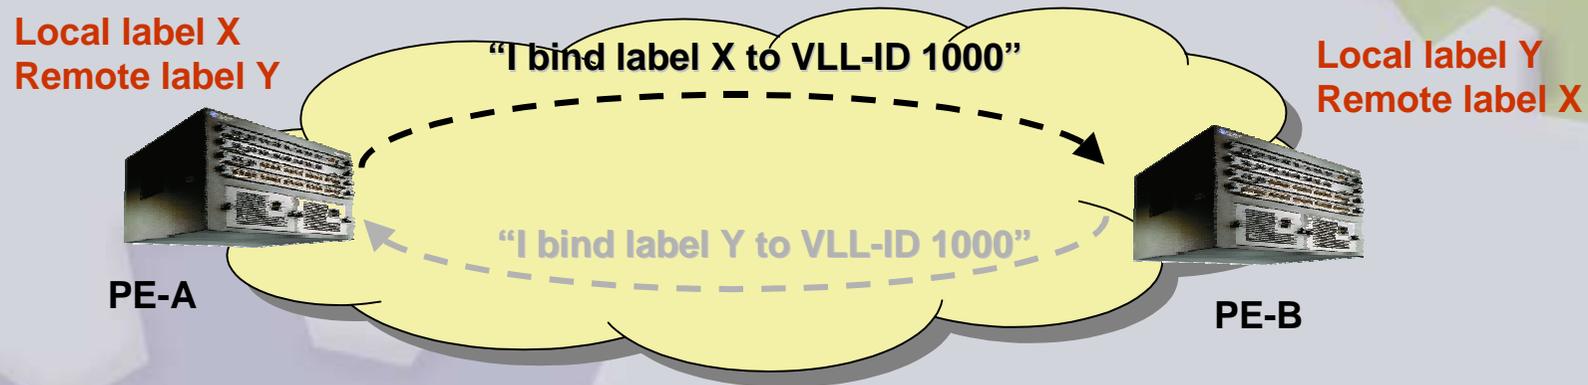
PPP Hdr Type-MPLS	Tunnel Label M	VC Label Y	Tagged or untagged Ethernet payload
----------------------	-----------------------------	-------------------------	-------------------------------------

Anatomy of a VLL VC Label Signaling

- Once all configurations are complete, a PE router sends an LDP Label-Mapping message to the other end (in downstream unsolicited mode)
- The message indicates the binding of a local VC label (in Label TLV) to the VC FEC TLV (introduced by draft-martini)



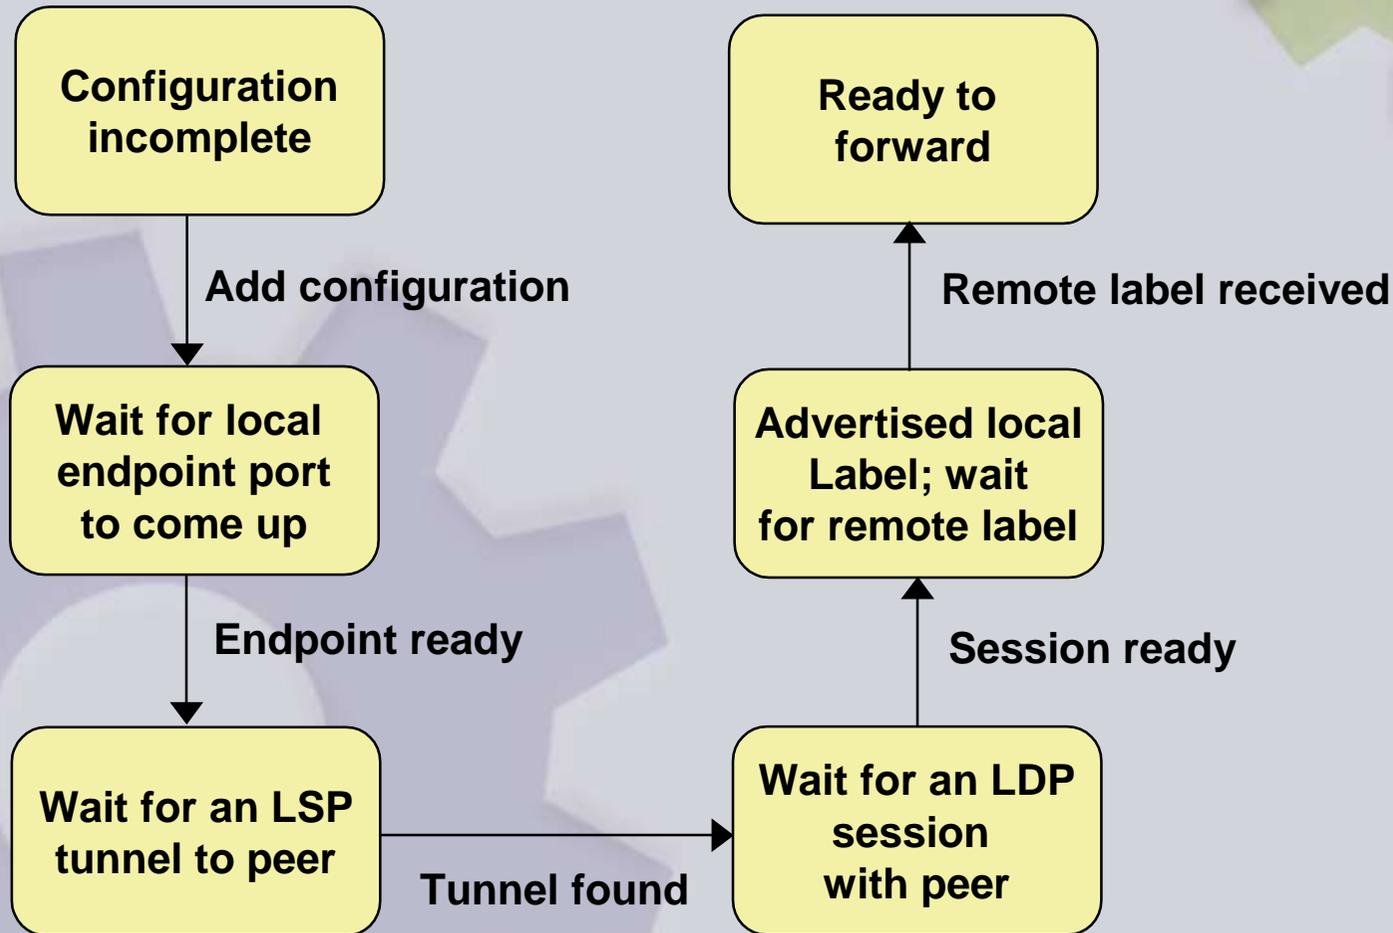
Anatomy of a VLL VC Label Signaling



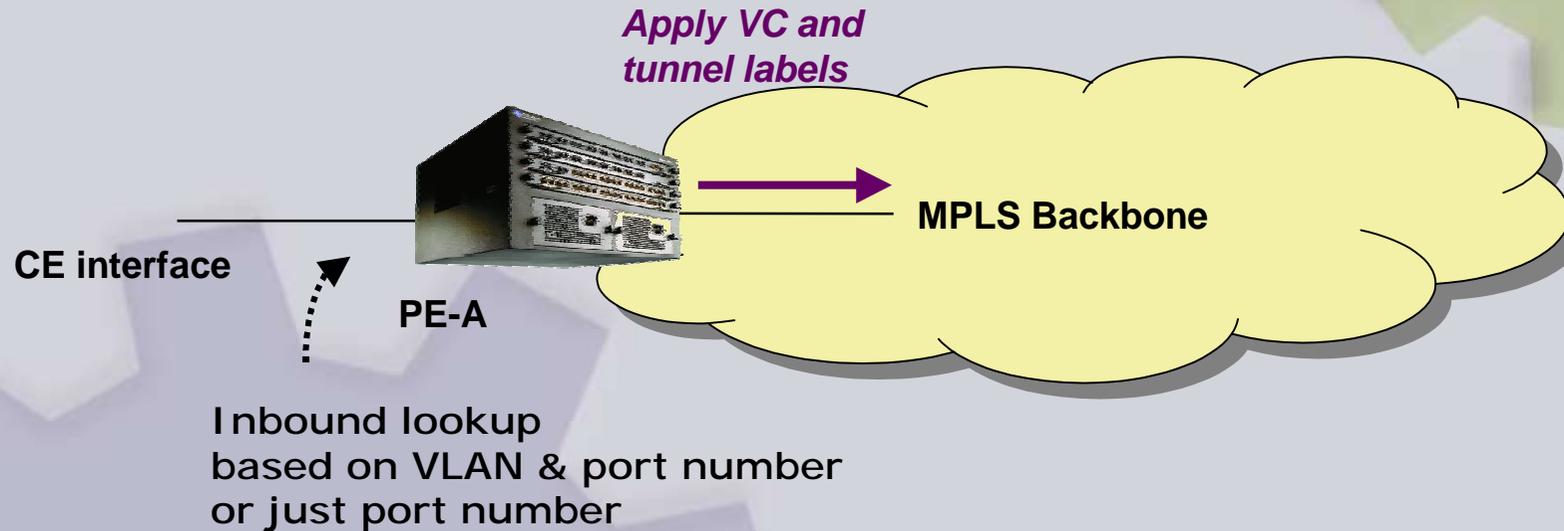
- PE router sends an LDP Label-Mapping message to the other end
- May send Label-Withdraw message when local CE interface goes down

Anatomy of a VLL

How Does a VLL Get Up and Running?



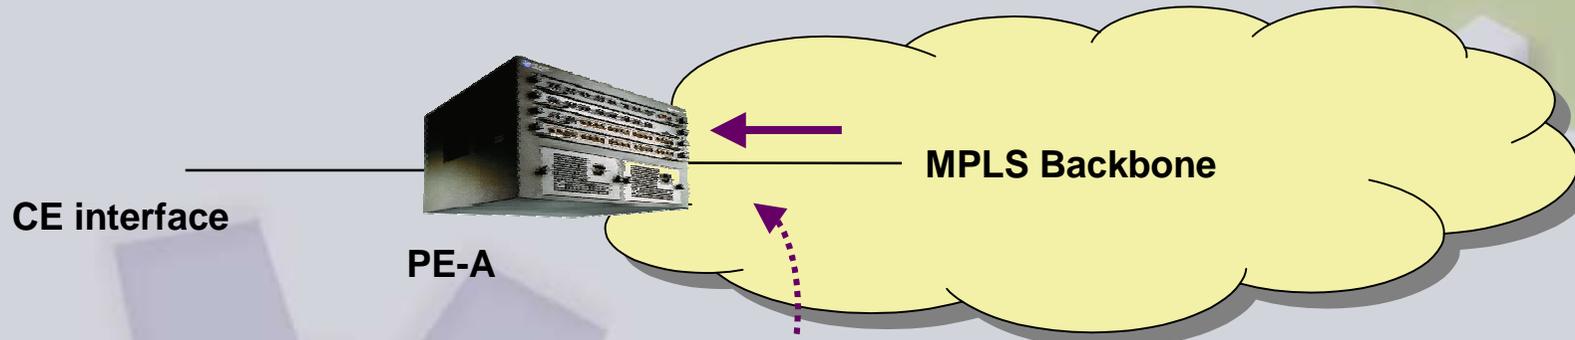
Forwarding Frames: CE to Backbone



Outgoing MPLS Packet Format:

PPP-MPLS or Enet Hdr.	Tunnel Label M	VC Label Y	Optional Control Word	Tagged or untagged Ethernet payload
-----------------------------	-----------------------------	-------------------------	-----------------------------	-------------------------------------

Forwarding Frames: Backbone to CE



Incoming label X indicates VLL payload and which CE interface to send to and VLAN-ID to use

Incoming MPLS Packet Format:

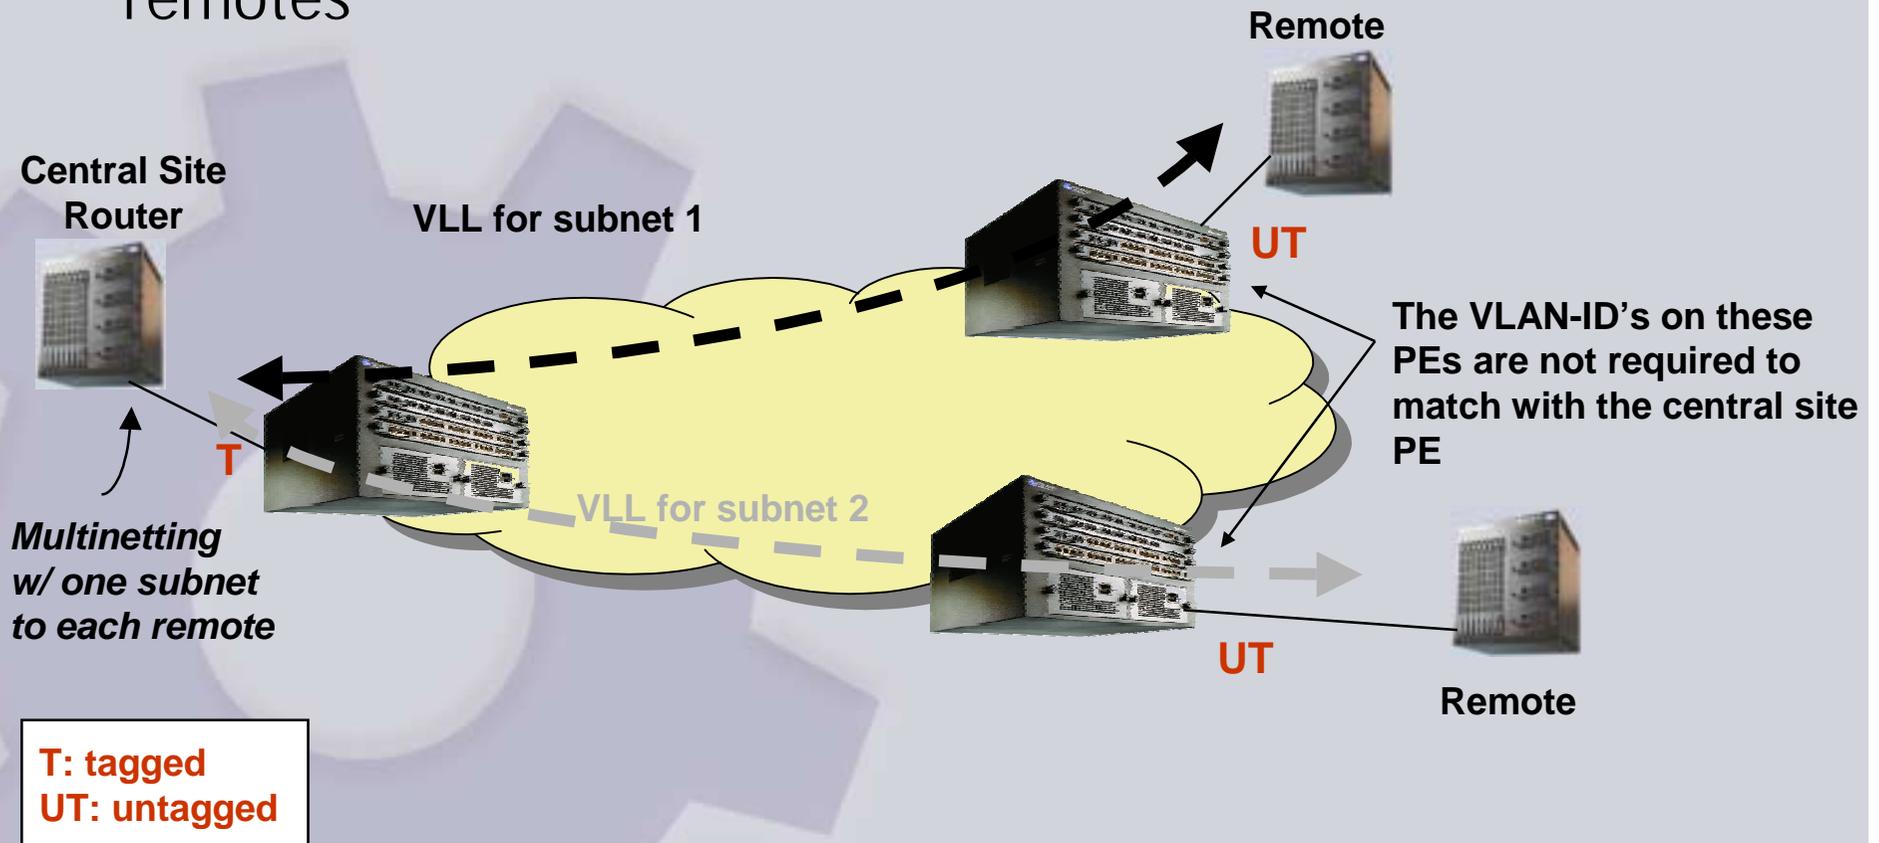
PPP-MPLS or Enet Hdr.	Tunnel Label M	VC Label Y	Optional Control Word	Tagged or untagged Ethernet payload
-----------------------------	-----------------------------	-------------------------	-----------------------------	-------------------------------------

Incoming MPLS Packet Format in Case of Pen-Ultimate Hop Popping:

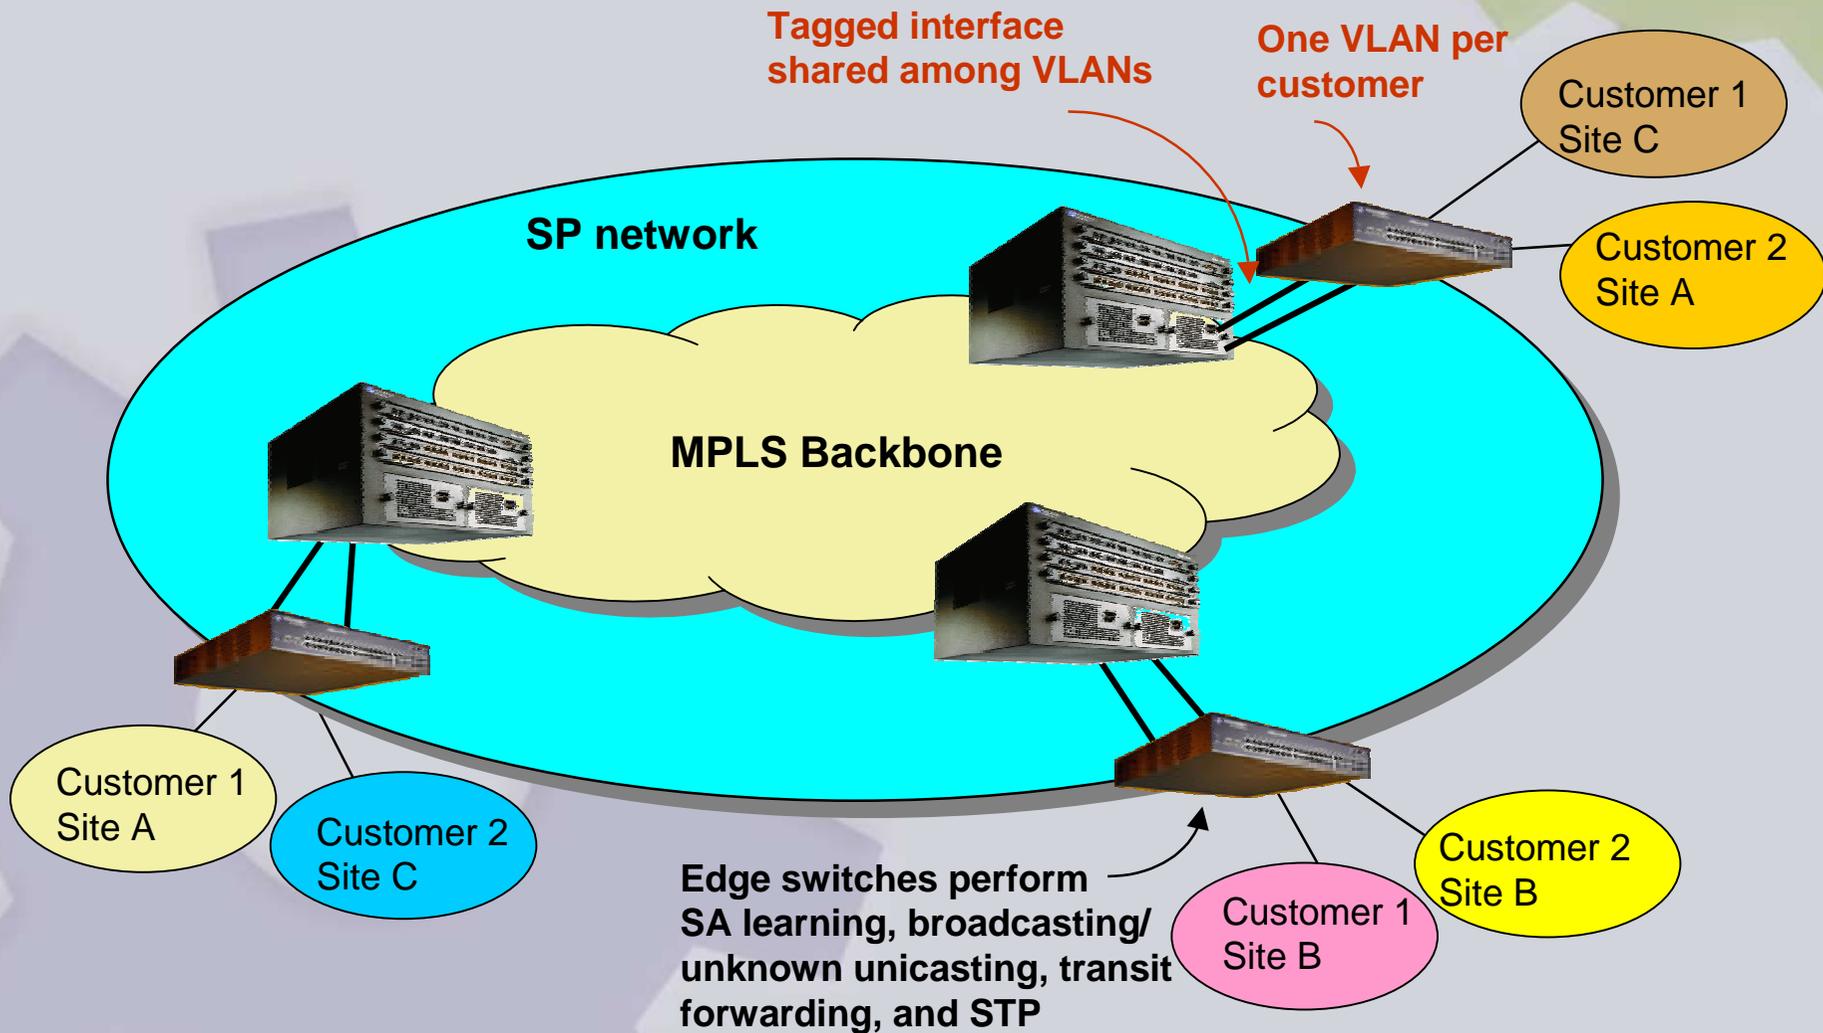
PPP-MPLS or Enet Hdr.	VC Label X	Optional Control Word	Tagged or untagged Ethernet payload
-----------------------------	-------------------------	-----------------------------	-------------------------------------

VLL Example Application: CE Device Handles Switching

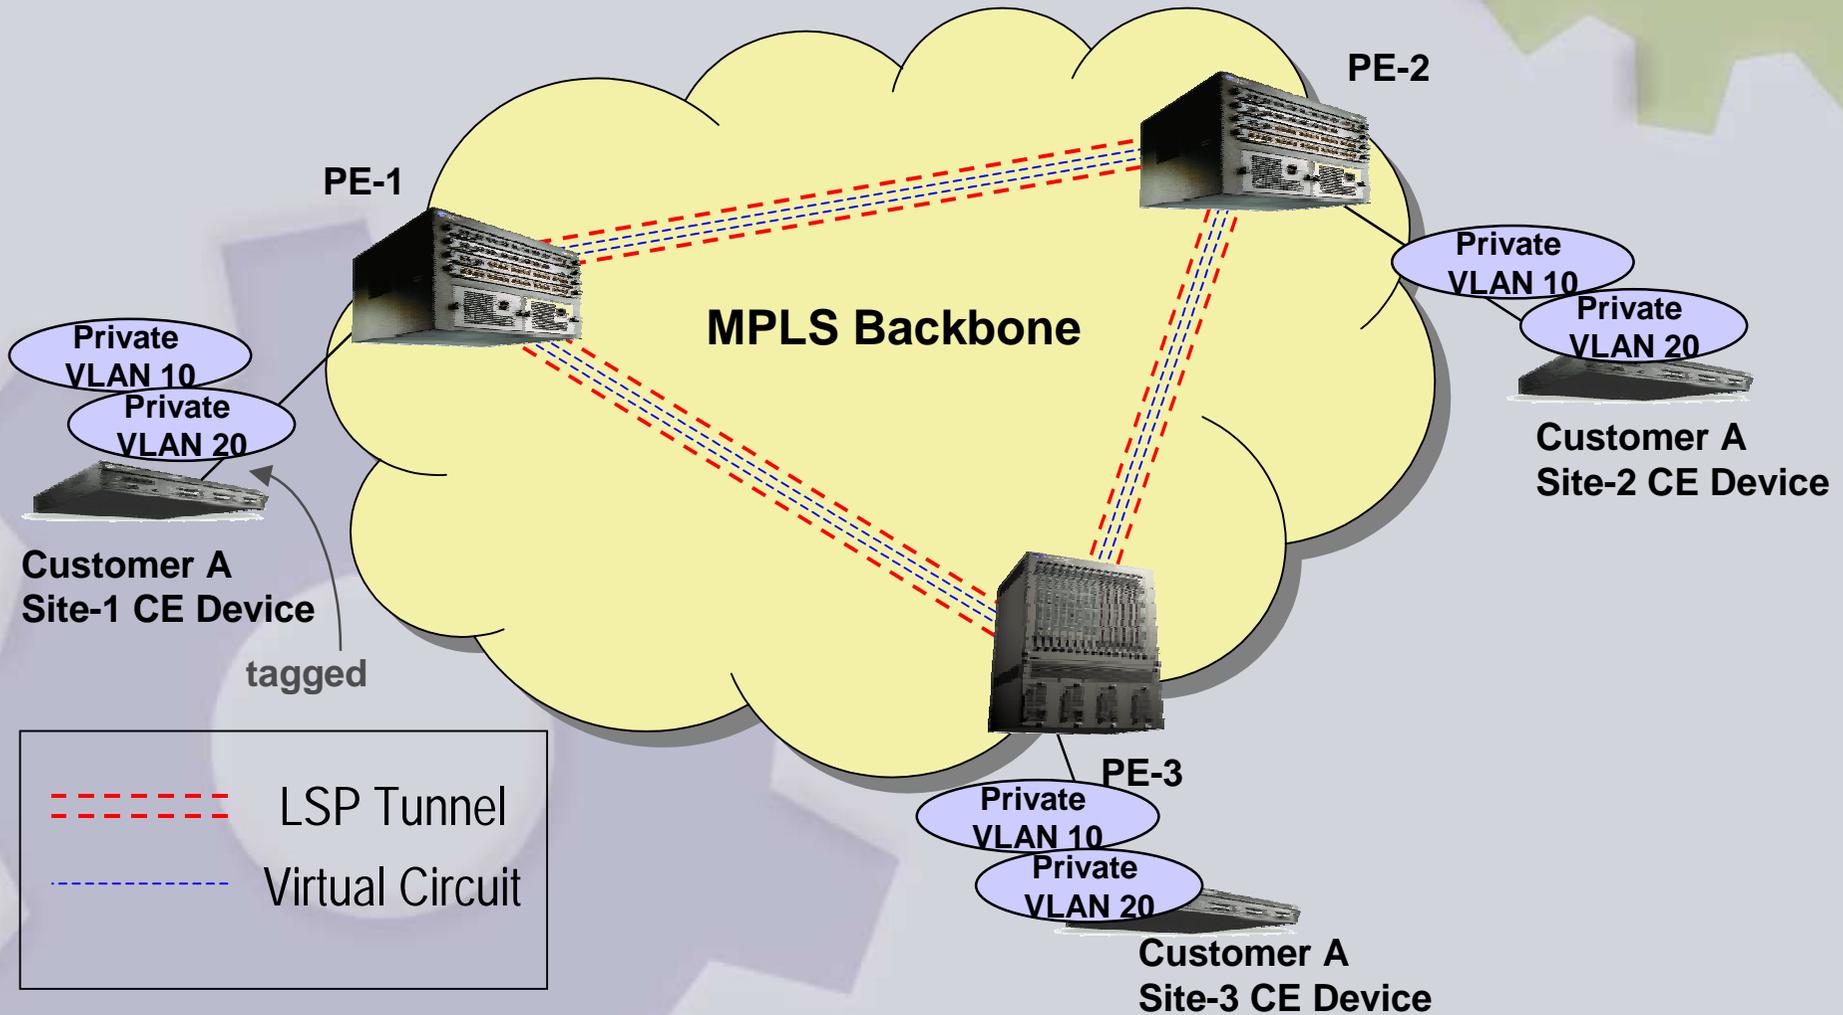
- Router at central location responsible for data forwarding among all remotes; switches or routers at remotes



VLL Example Application: SP Handles Switching

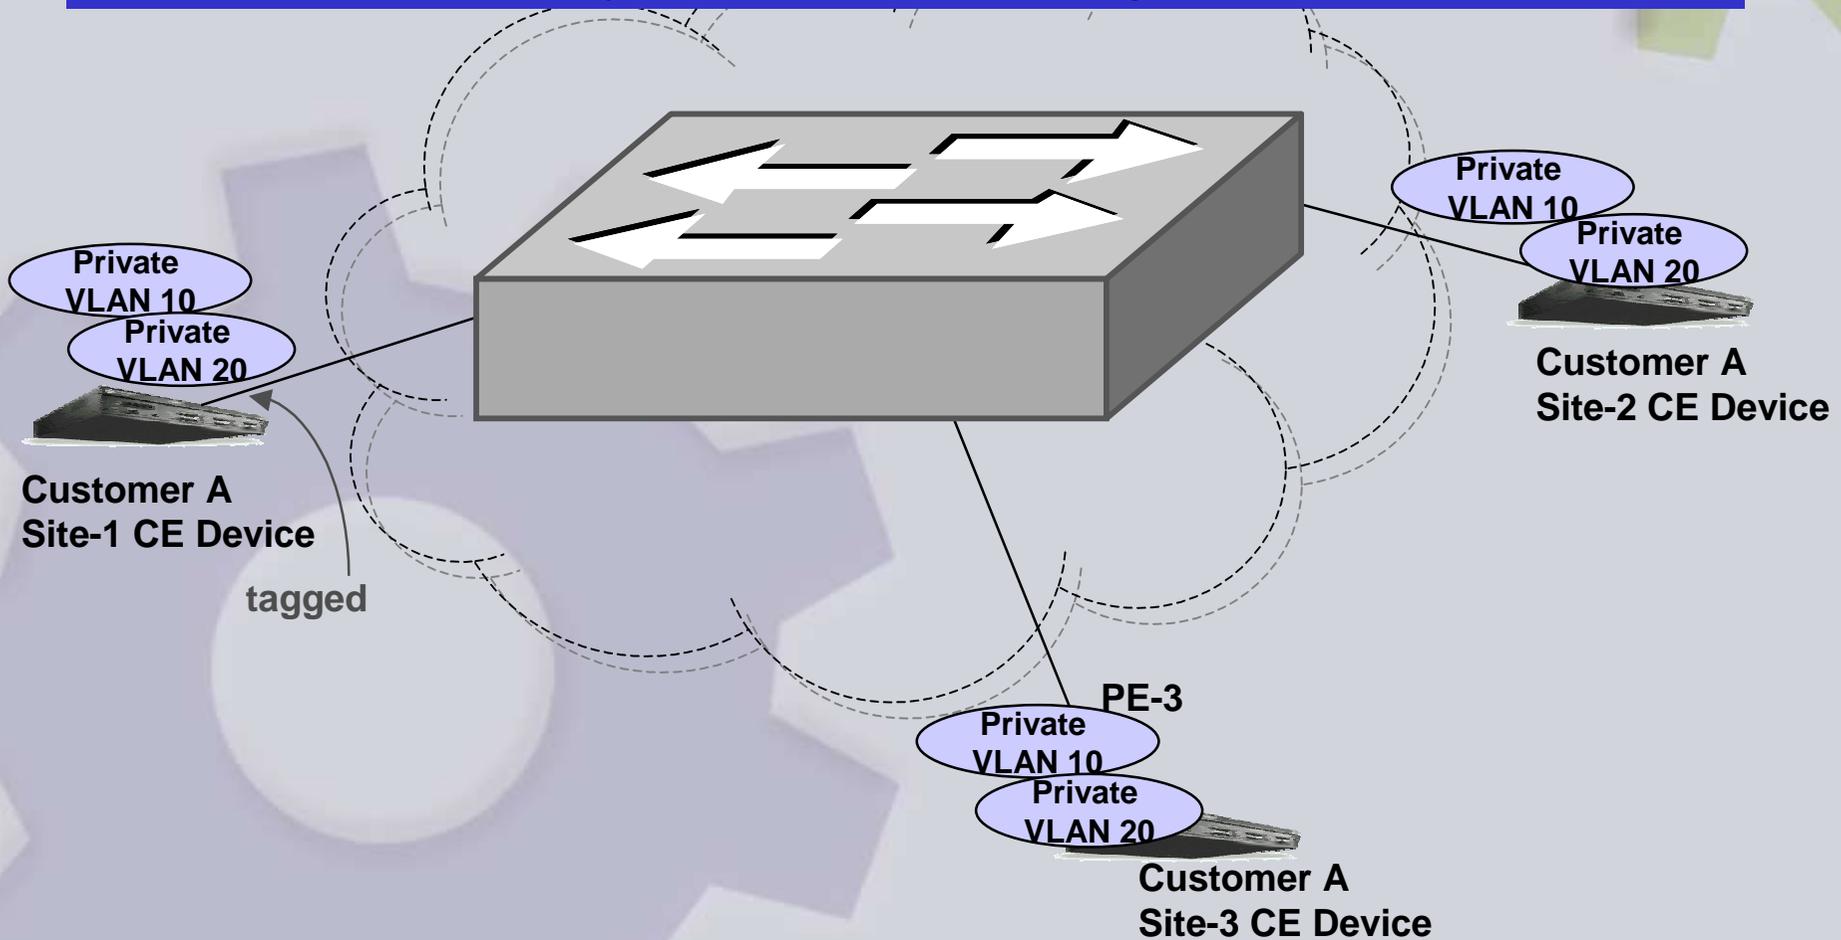


Service Provider View of the Service



Customer's View of the Service

Customer's CE devices think that they are all in the same subnet and that they are connected through a switch



Agenda

- L2 Technologies in MAN
- MPLS VLL Draft-Martini
- ➔ MPLS VPLS
- Interoperability

Service Terminology

- In this scenario, the provider is said to be offering:

A Virtual Private LAN Segment (VPLS)

or (alternative name):

A Transparent LAN Service

Not to be confused with "Virtual Private LAN Services" which is an IP only solution described in "draft-tsenevir-vpl-ip-00.txt"

VPLS Defining Documents

- RFC 2764 – “A Framework for IP Based Virtual Private Networks”:
 - Provides the general framework for IP Based VPNs – MPLS included
- draft-lasserre-tls-mpls-00.txt:
 - Describes the forwarding of Ethernet/802.3 frames between multiple customer sites, as if they were in the same layer-2 broadcast domain
 - Defines a new parameter “VPN ID” to identify PE routers participating in the same customer VPN
- draft-vkompella-ppvvpn-vpsn-mpls-00.txt:
 - Describes MAC address learning and aging
 - Describes MAC address signaling
- A single WG document will be formulated

Functional Requirements

- As frames are transported over a provider backbone, efficient use of resources is crucial
- Frames should be forwarded only over the needed VCs
- PE routers should be capable of learning Ethernet source MAC addresses, just like a regular switch
- A PE router keeps a VPLS Forwarding Table (quite similar to a bridge's MAC table) that distinguishes the entries belonging to the different VPNs
- Before forwarding a frame, the PE router checks its VPLS forwarding table to determine which VCs the frame should go over

Functional Requirements (Cont.)

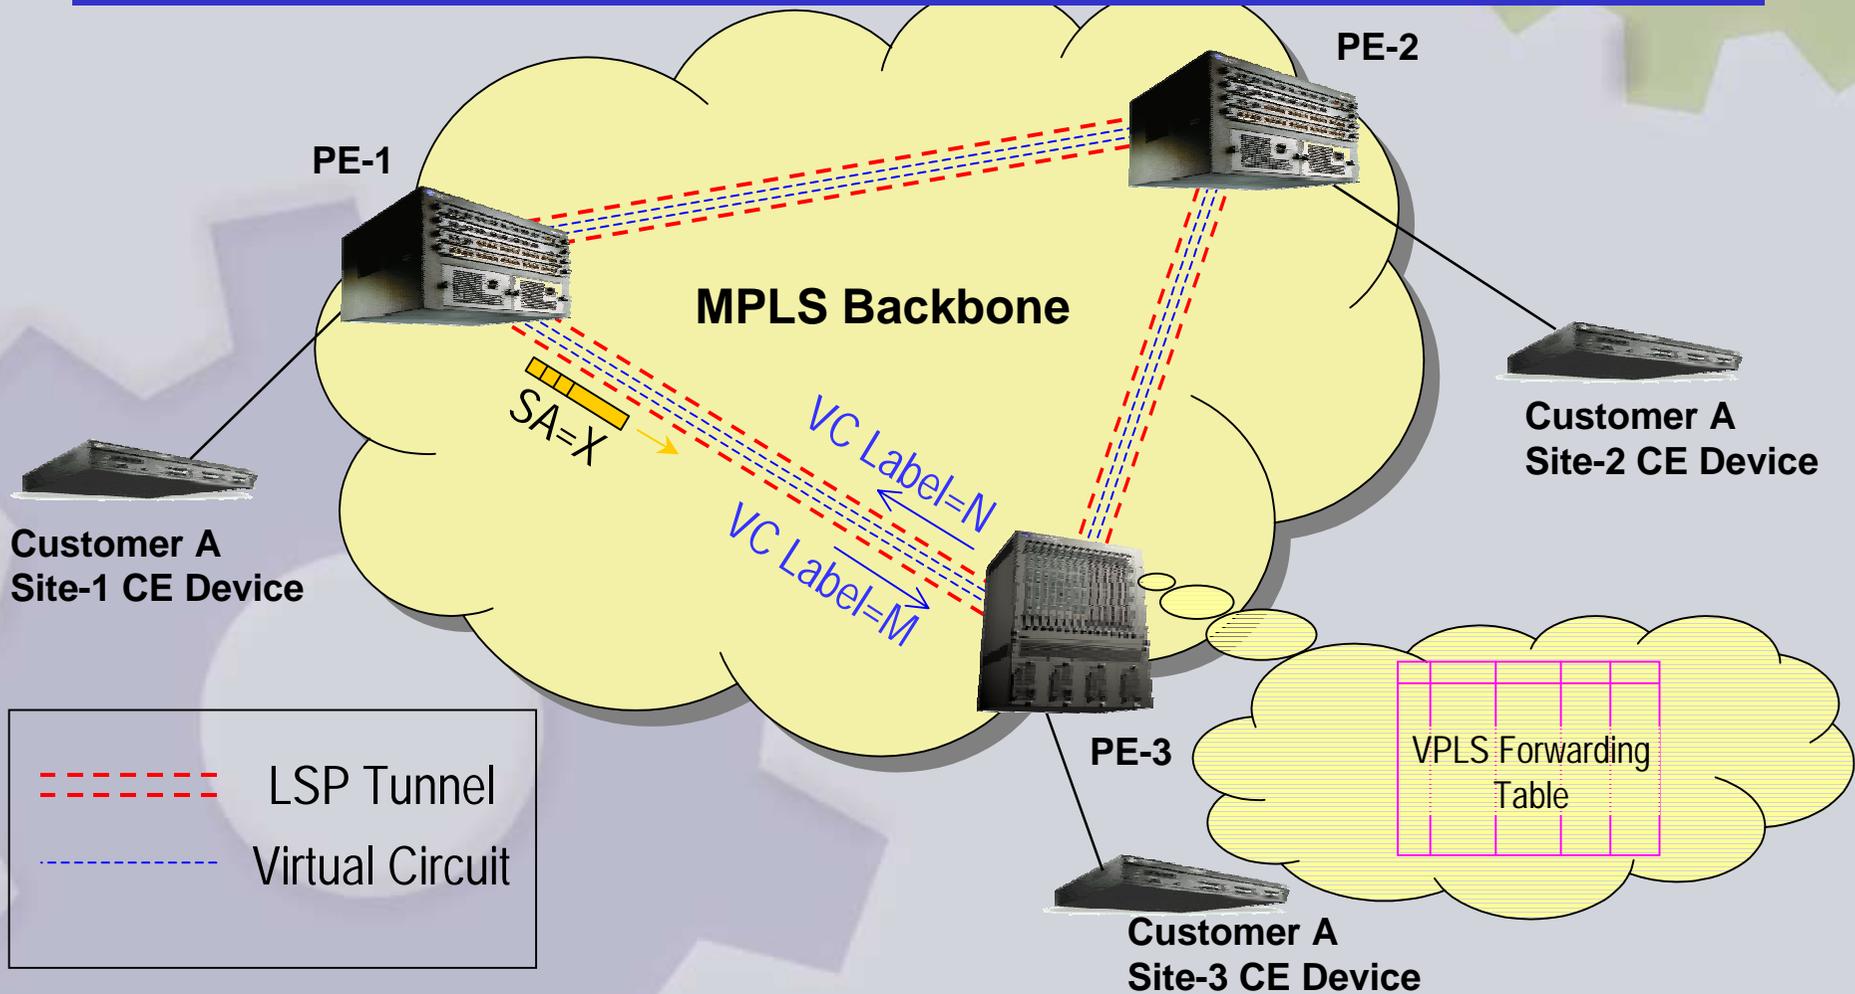
- Broadcast or unknown unicast frames are flooded over all outgoing VCs belonging to the VPN
- The PE routers have to be connected via VCs forming a full mesh topology
- For loop prevention, running an instance of STP per VPN in the provider network would not scale:
 - Solution:
When forwarding frames a PE follows a split-horizon rule: it never forwards a frame from one VC to another in the same VPN

Handling Multicast Frames

- Multicast frames are currently handled like broadcasts
- Extensions for performing smarter mutlicasts are to be developed in the future
- Possible scenarios:
 - Interaction with 802.1 GMRP
 - IGMP snooping
 - Static MAC multicast filters

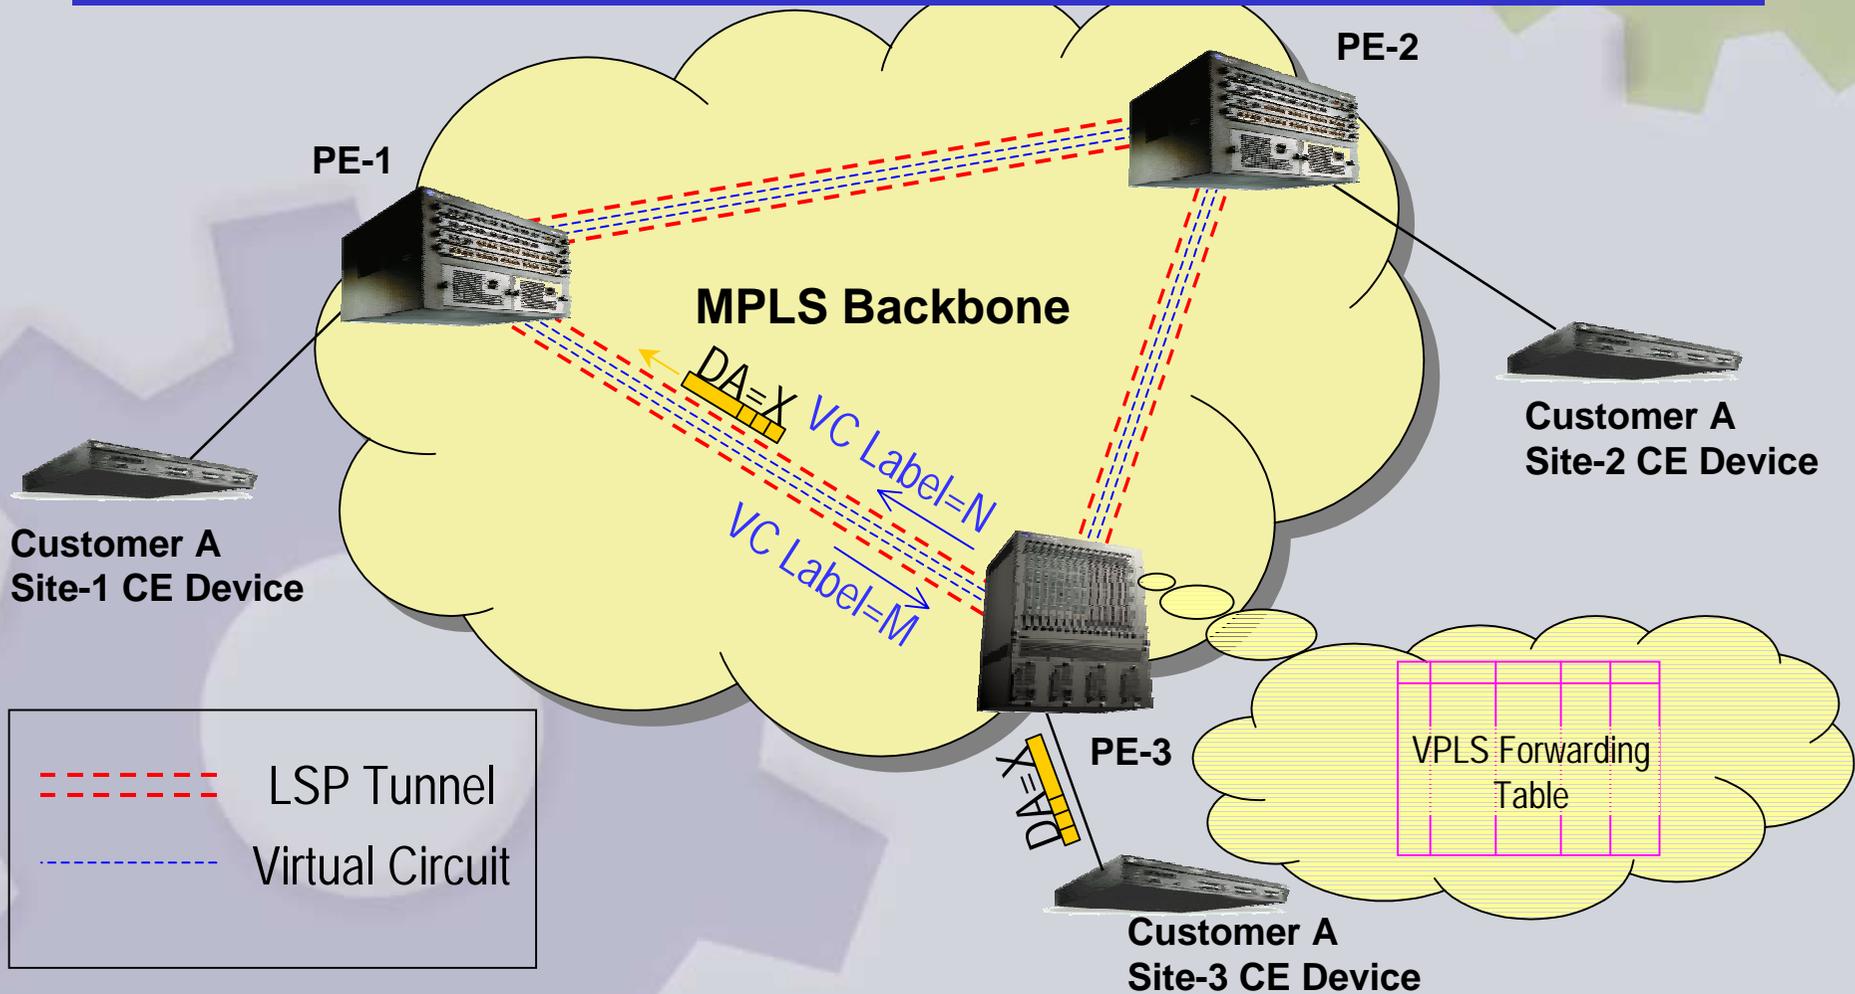
PE Router Performs Source Address Learning

PE-3 associates MAC X with the outgoing VC (VC Label N)

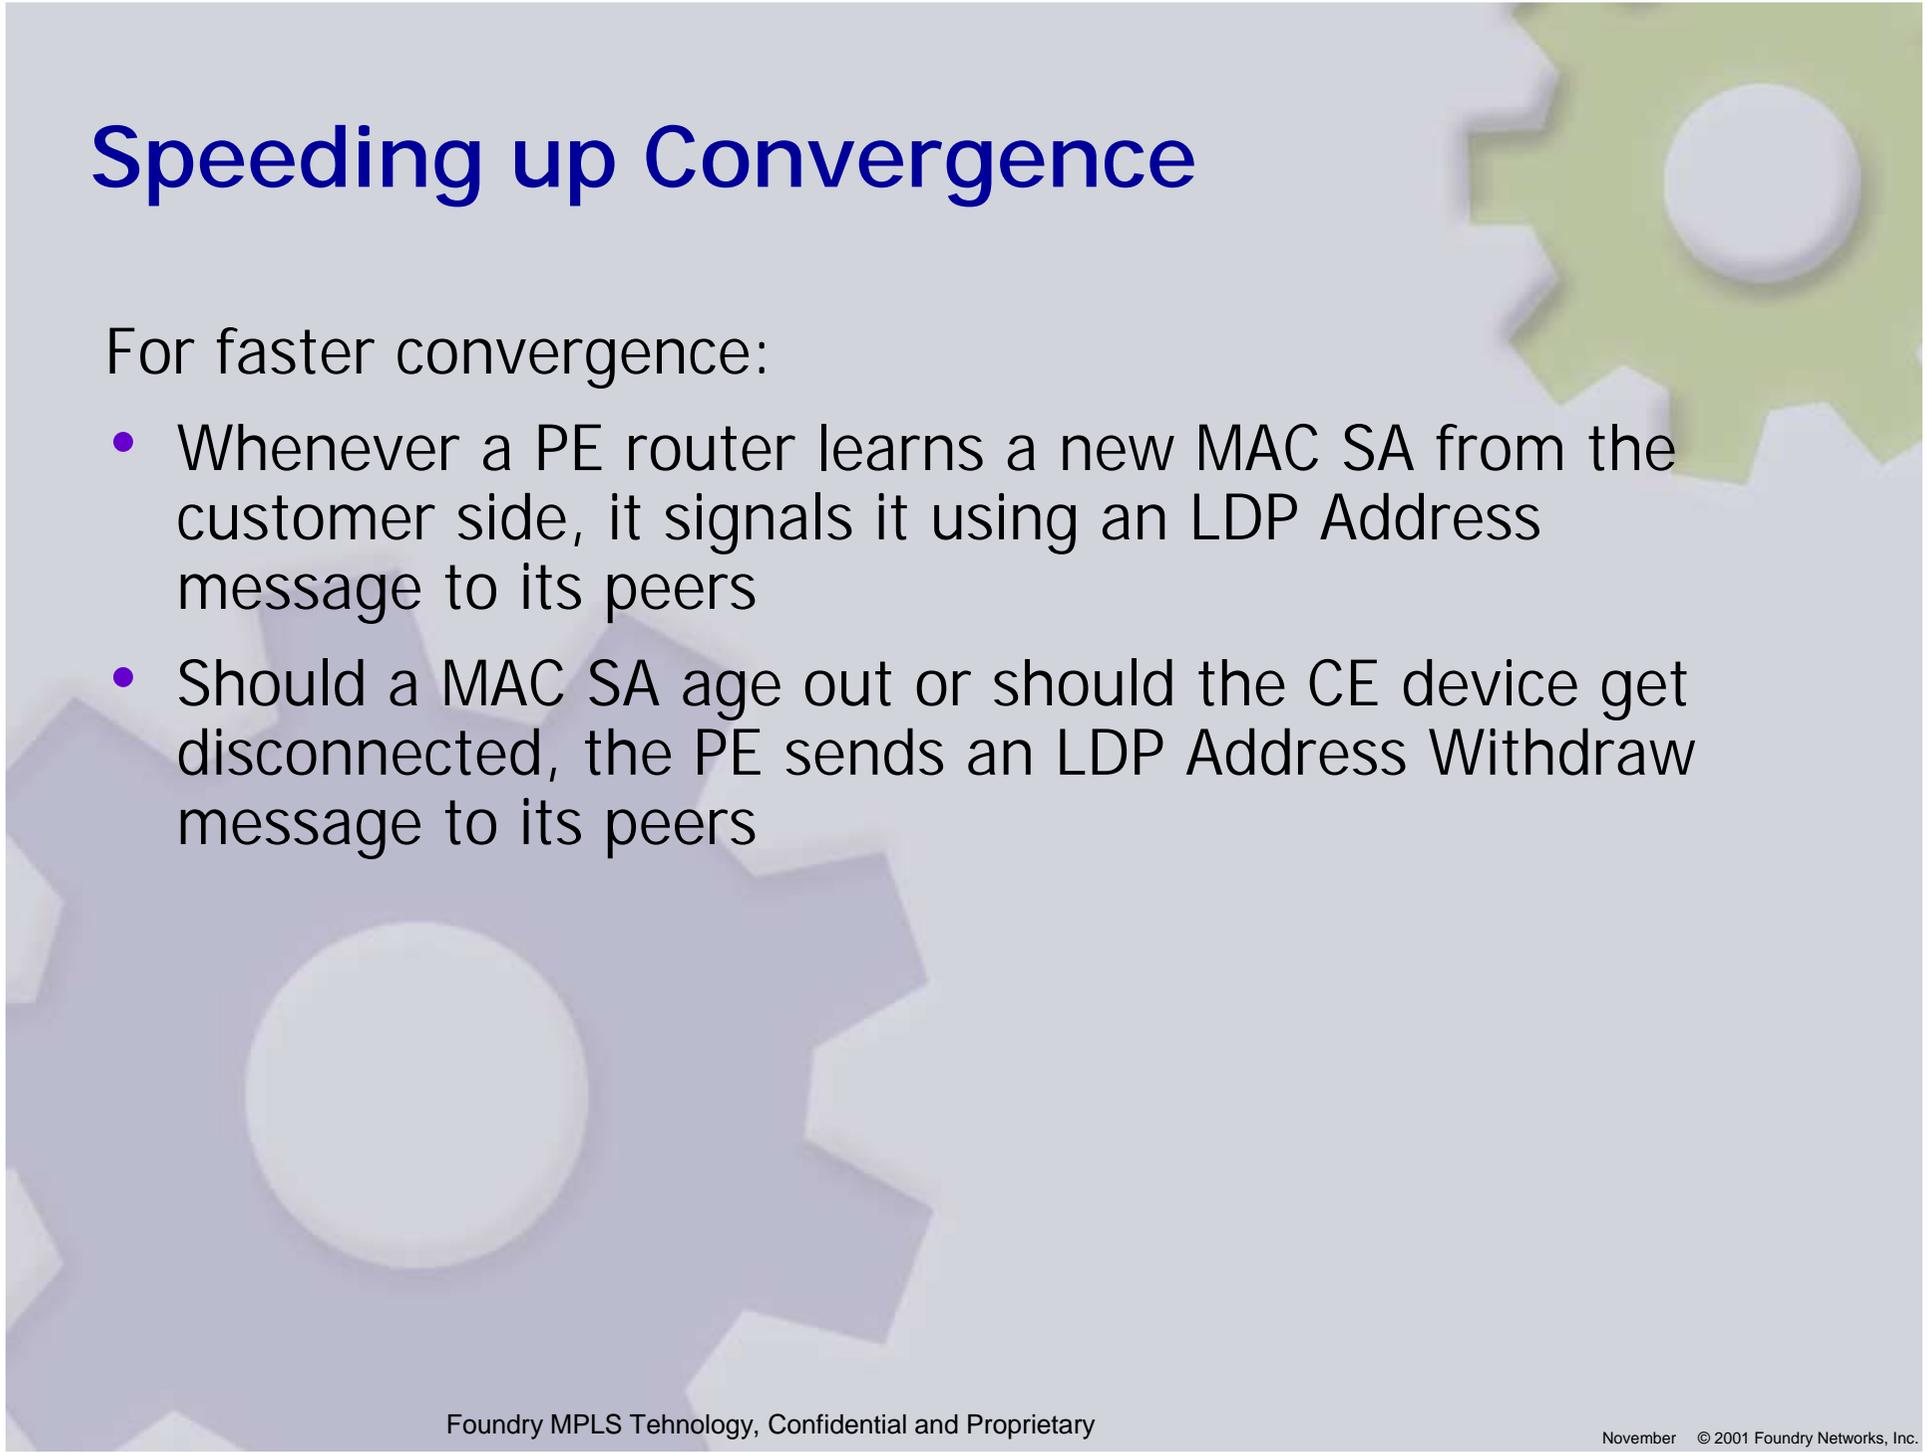


PE Router Performs Forwarding

PE-3 Checks its VPLS forwarding table and selects VC Label N



Speeding up Convergence

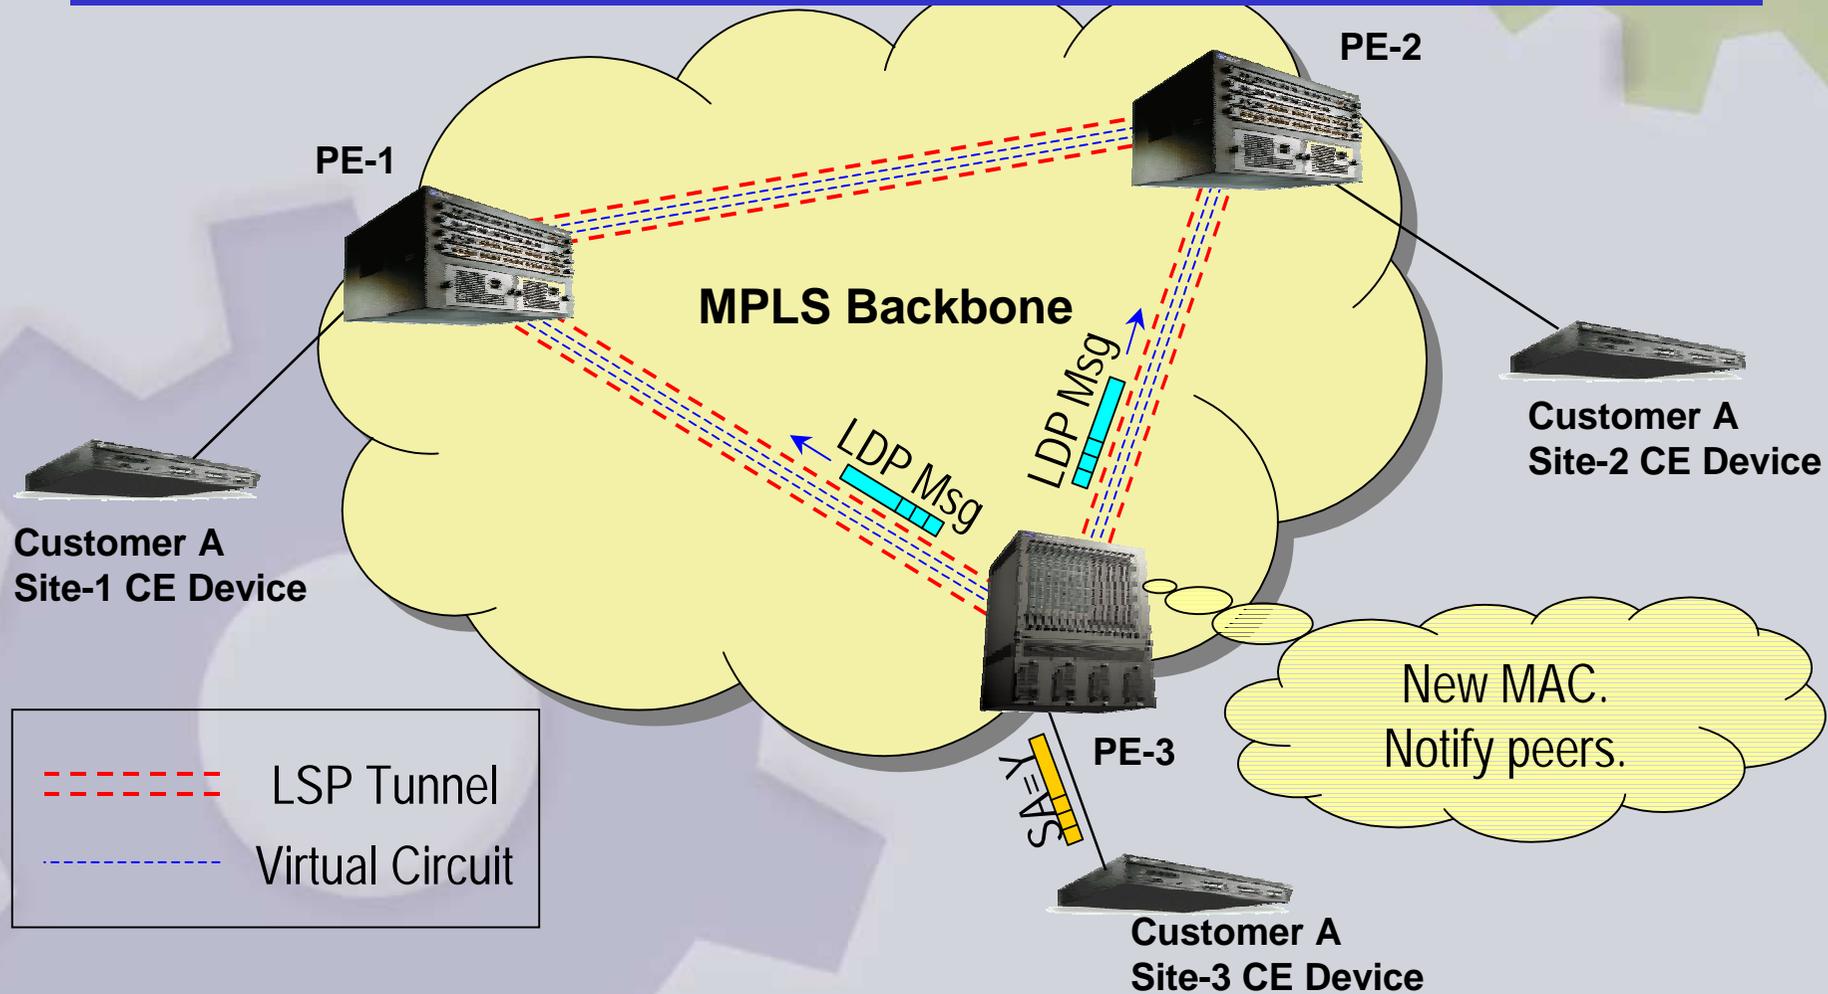


For faster convergence:

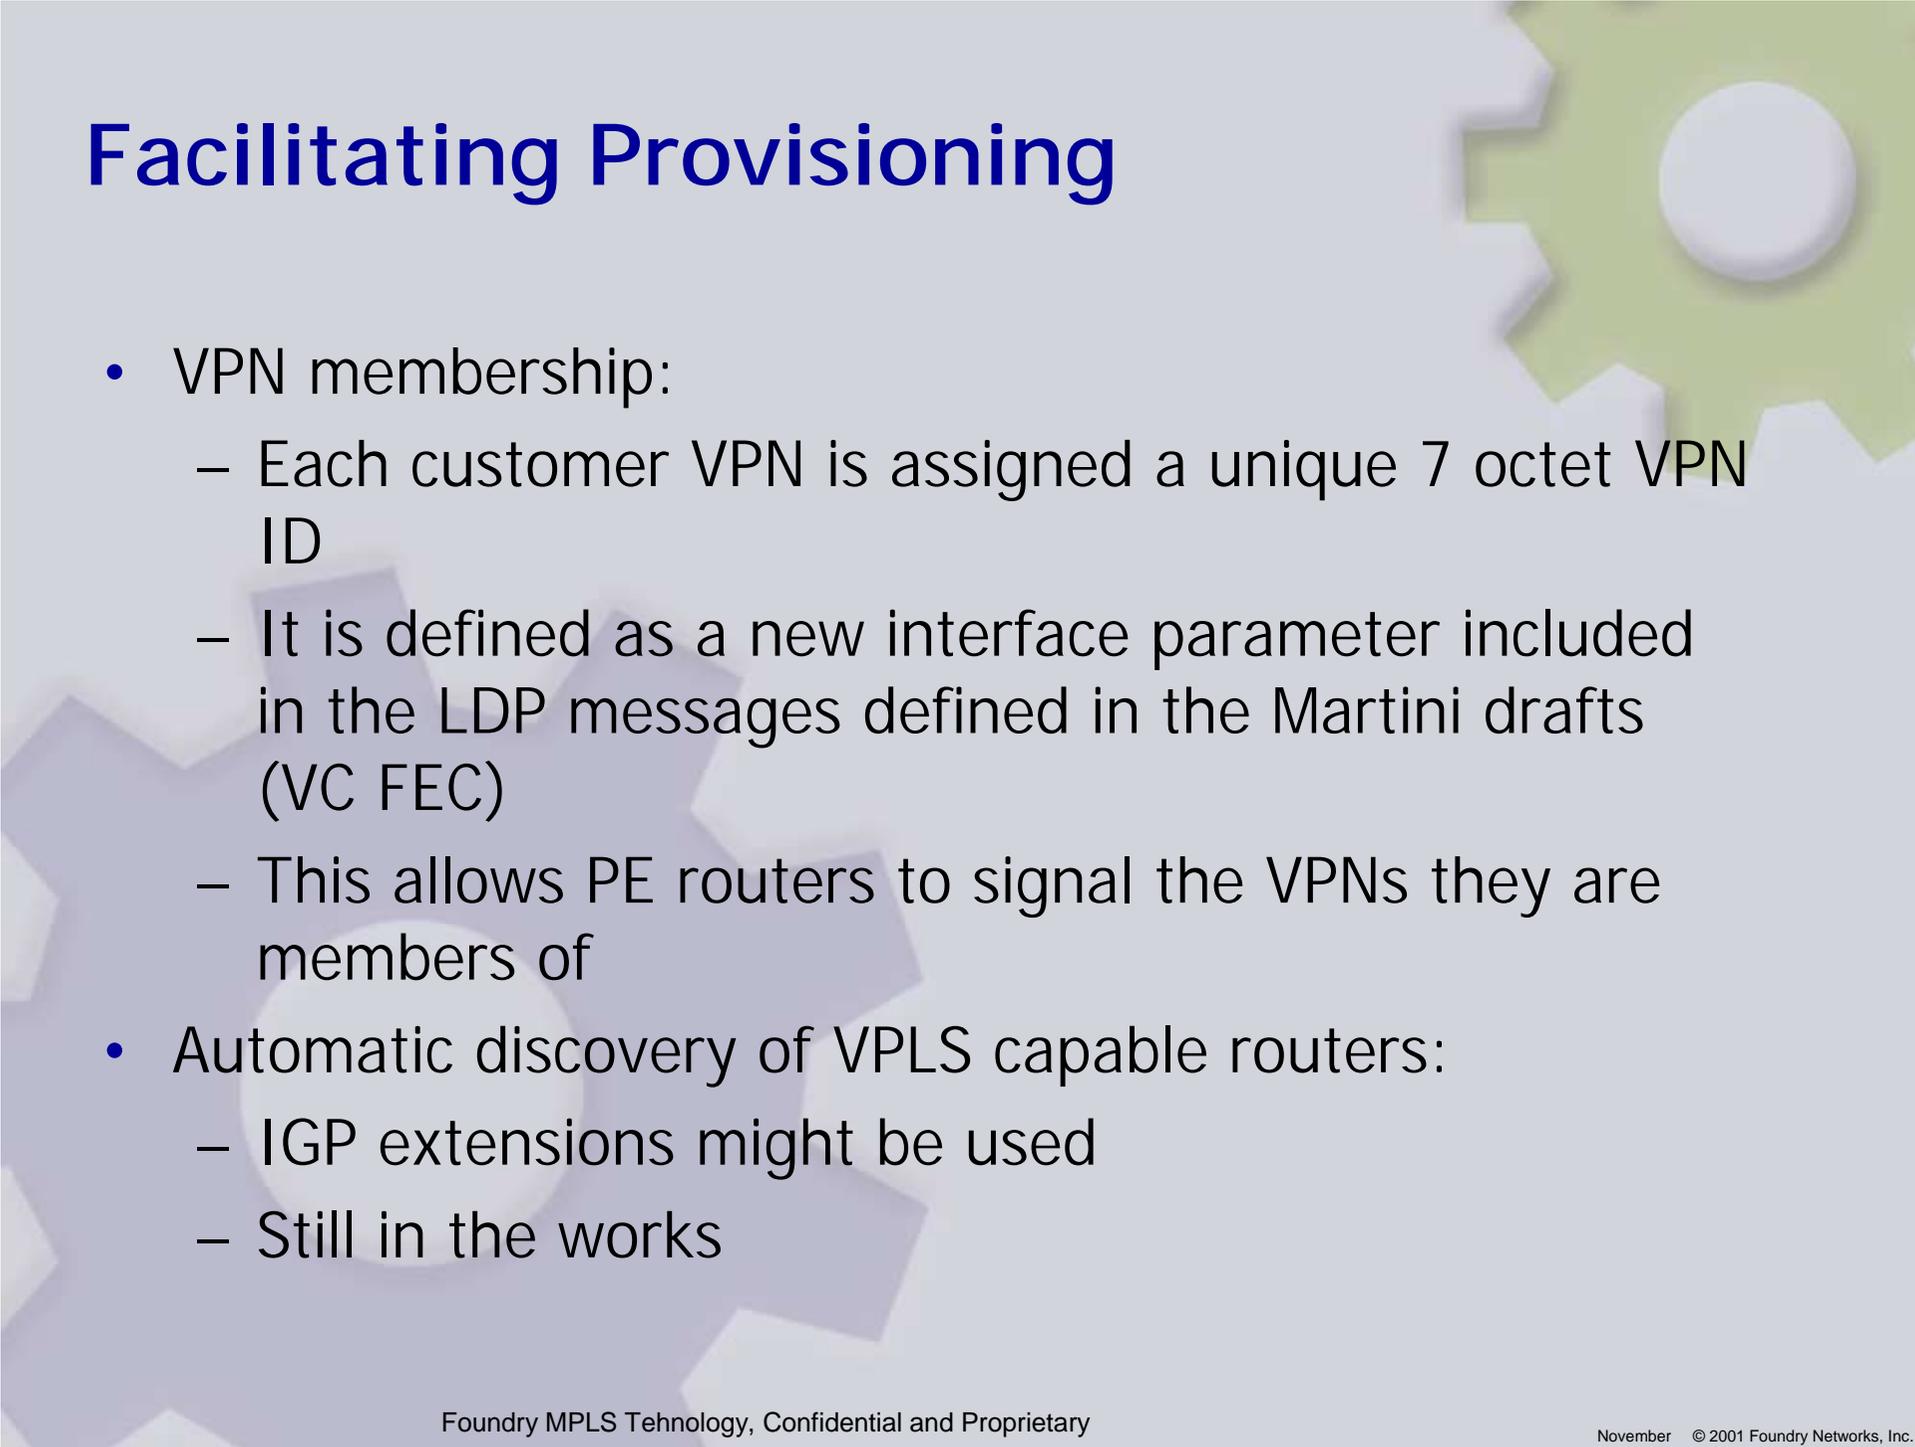
- Whenever a PE router learns a new MAC SA from the customer side, it signals it using an LDP Address message to its peers
- Should a MAC SA age out or should the CE device get disconnected, the PE sends an LDP Address Withdraw message to its peers

MAC Signaling

PE-3 signals the new MAC to speed up convergence

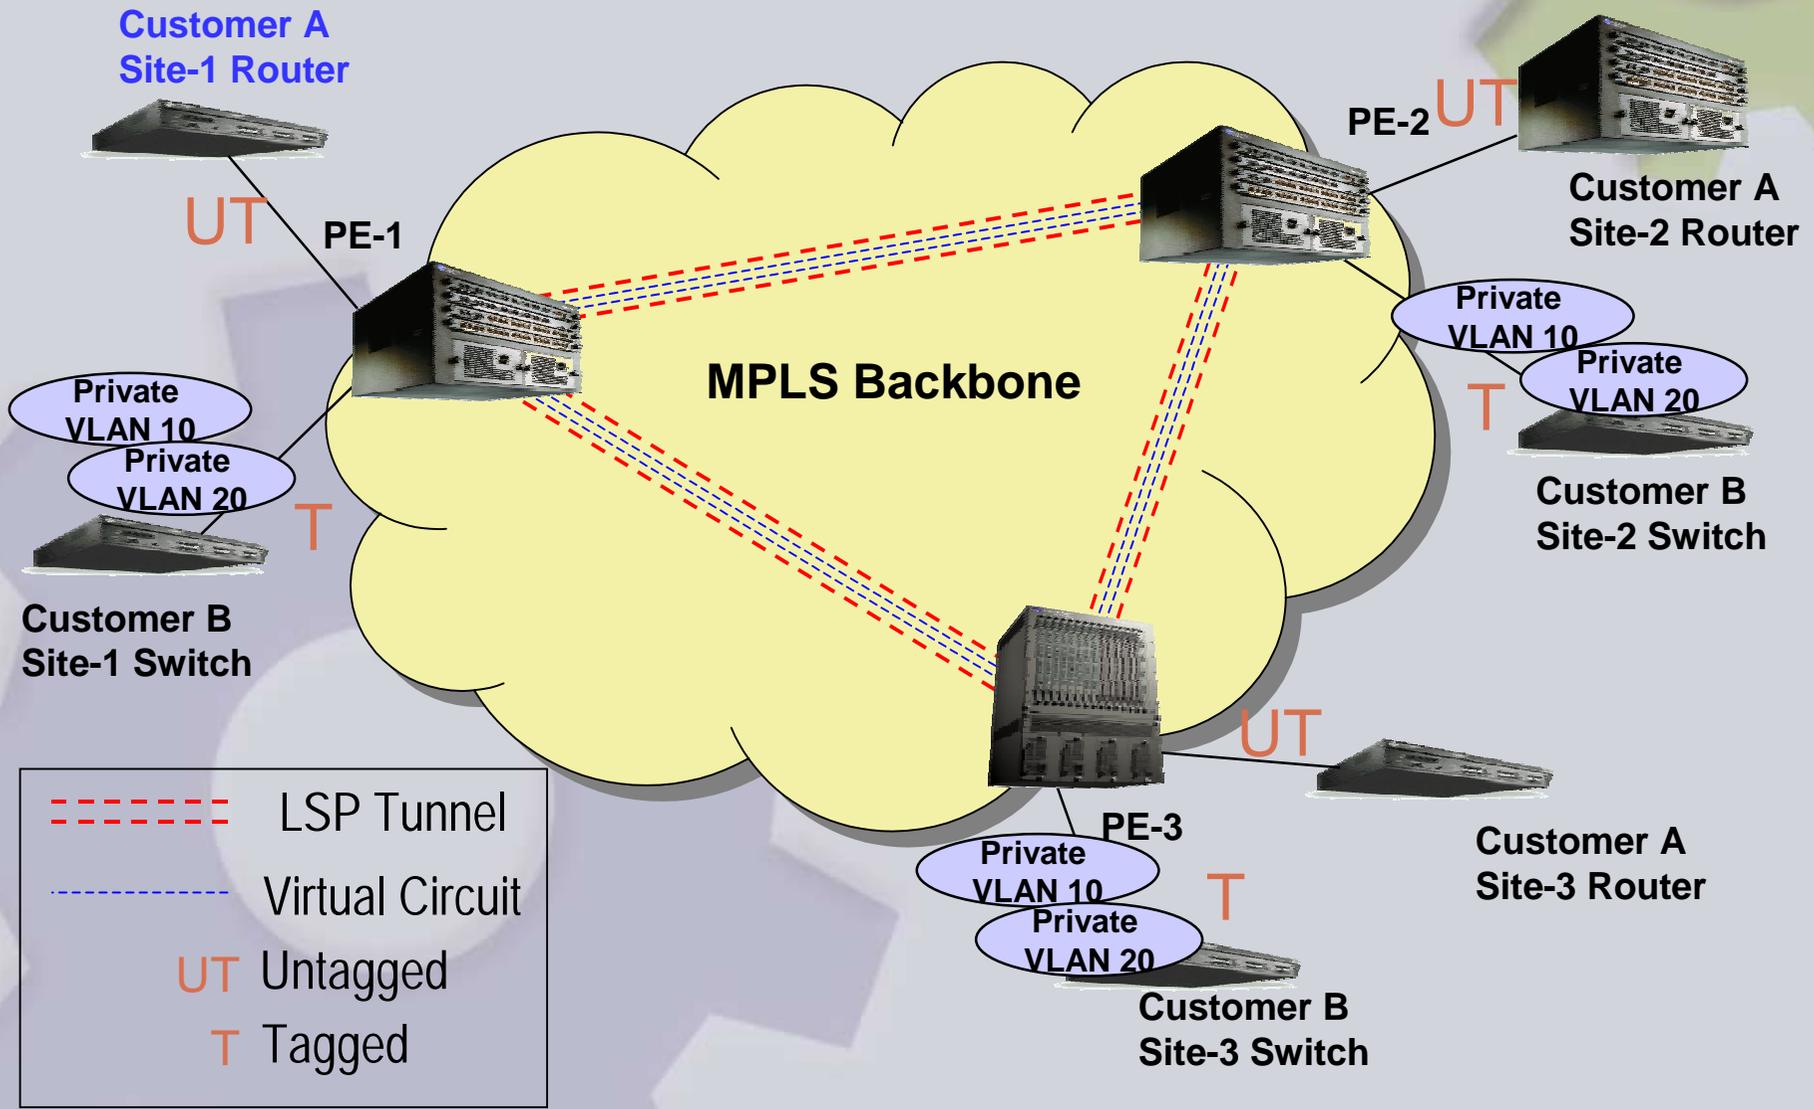


Facilitating Provisioning



- VPN membership:
 - Each customer VPN is assigned a unique 7 octet VPN ID
 - It is defined as a new interface parameter included in the LDP messages defined in the Martini drafts (VC FEC)
 - This allows PE routers to signal the VPNs they are members of
- Automatic discovery of VPLS capable routers:
 - IGP extensions might be used
 - Still in the works

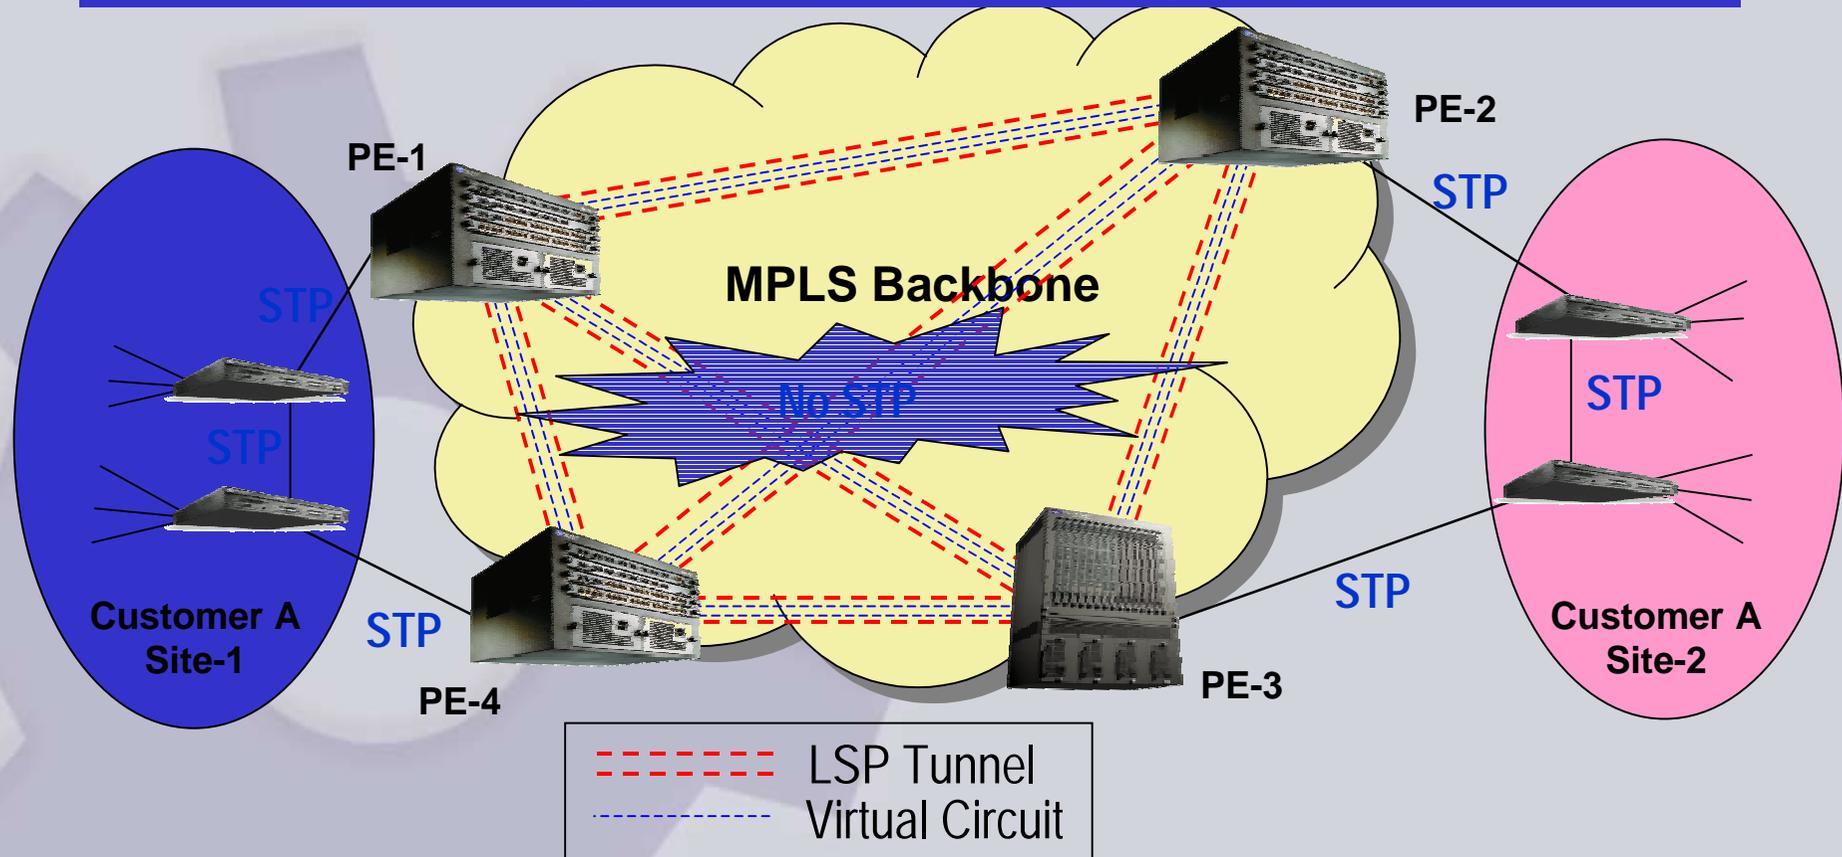
VPLS Example Application: Providing TLS

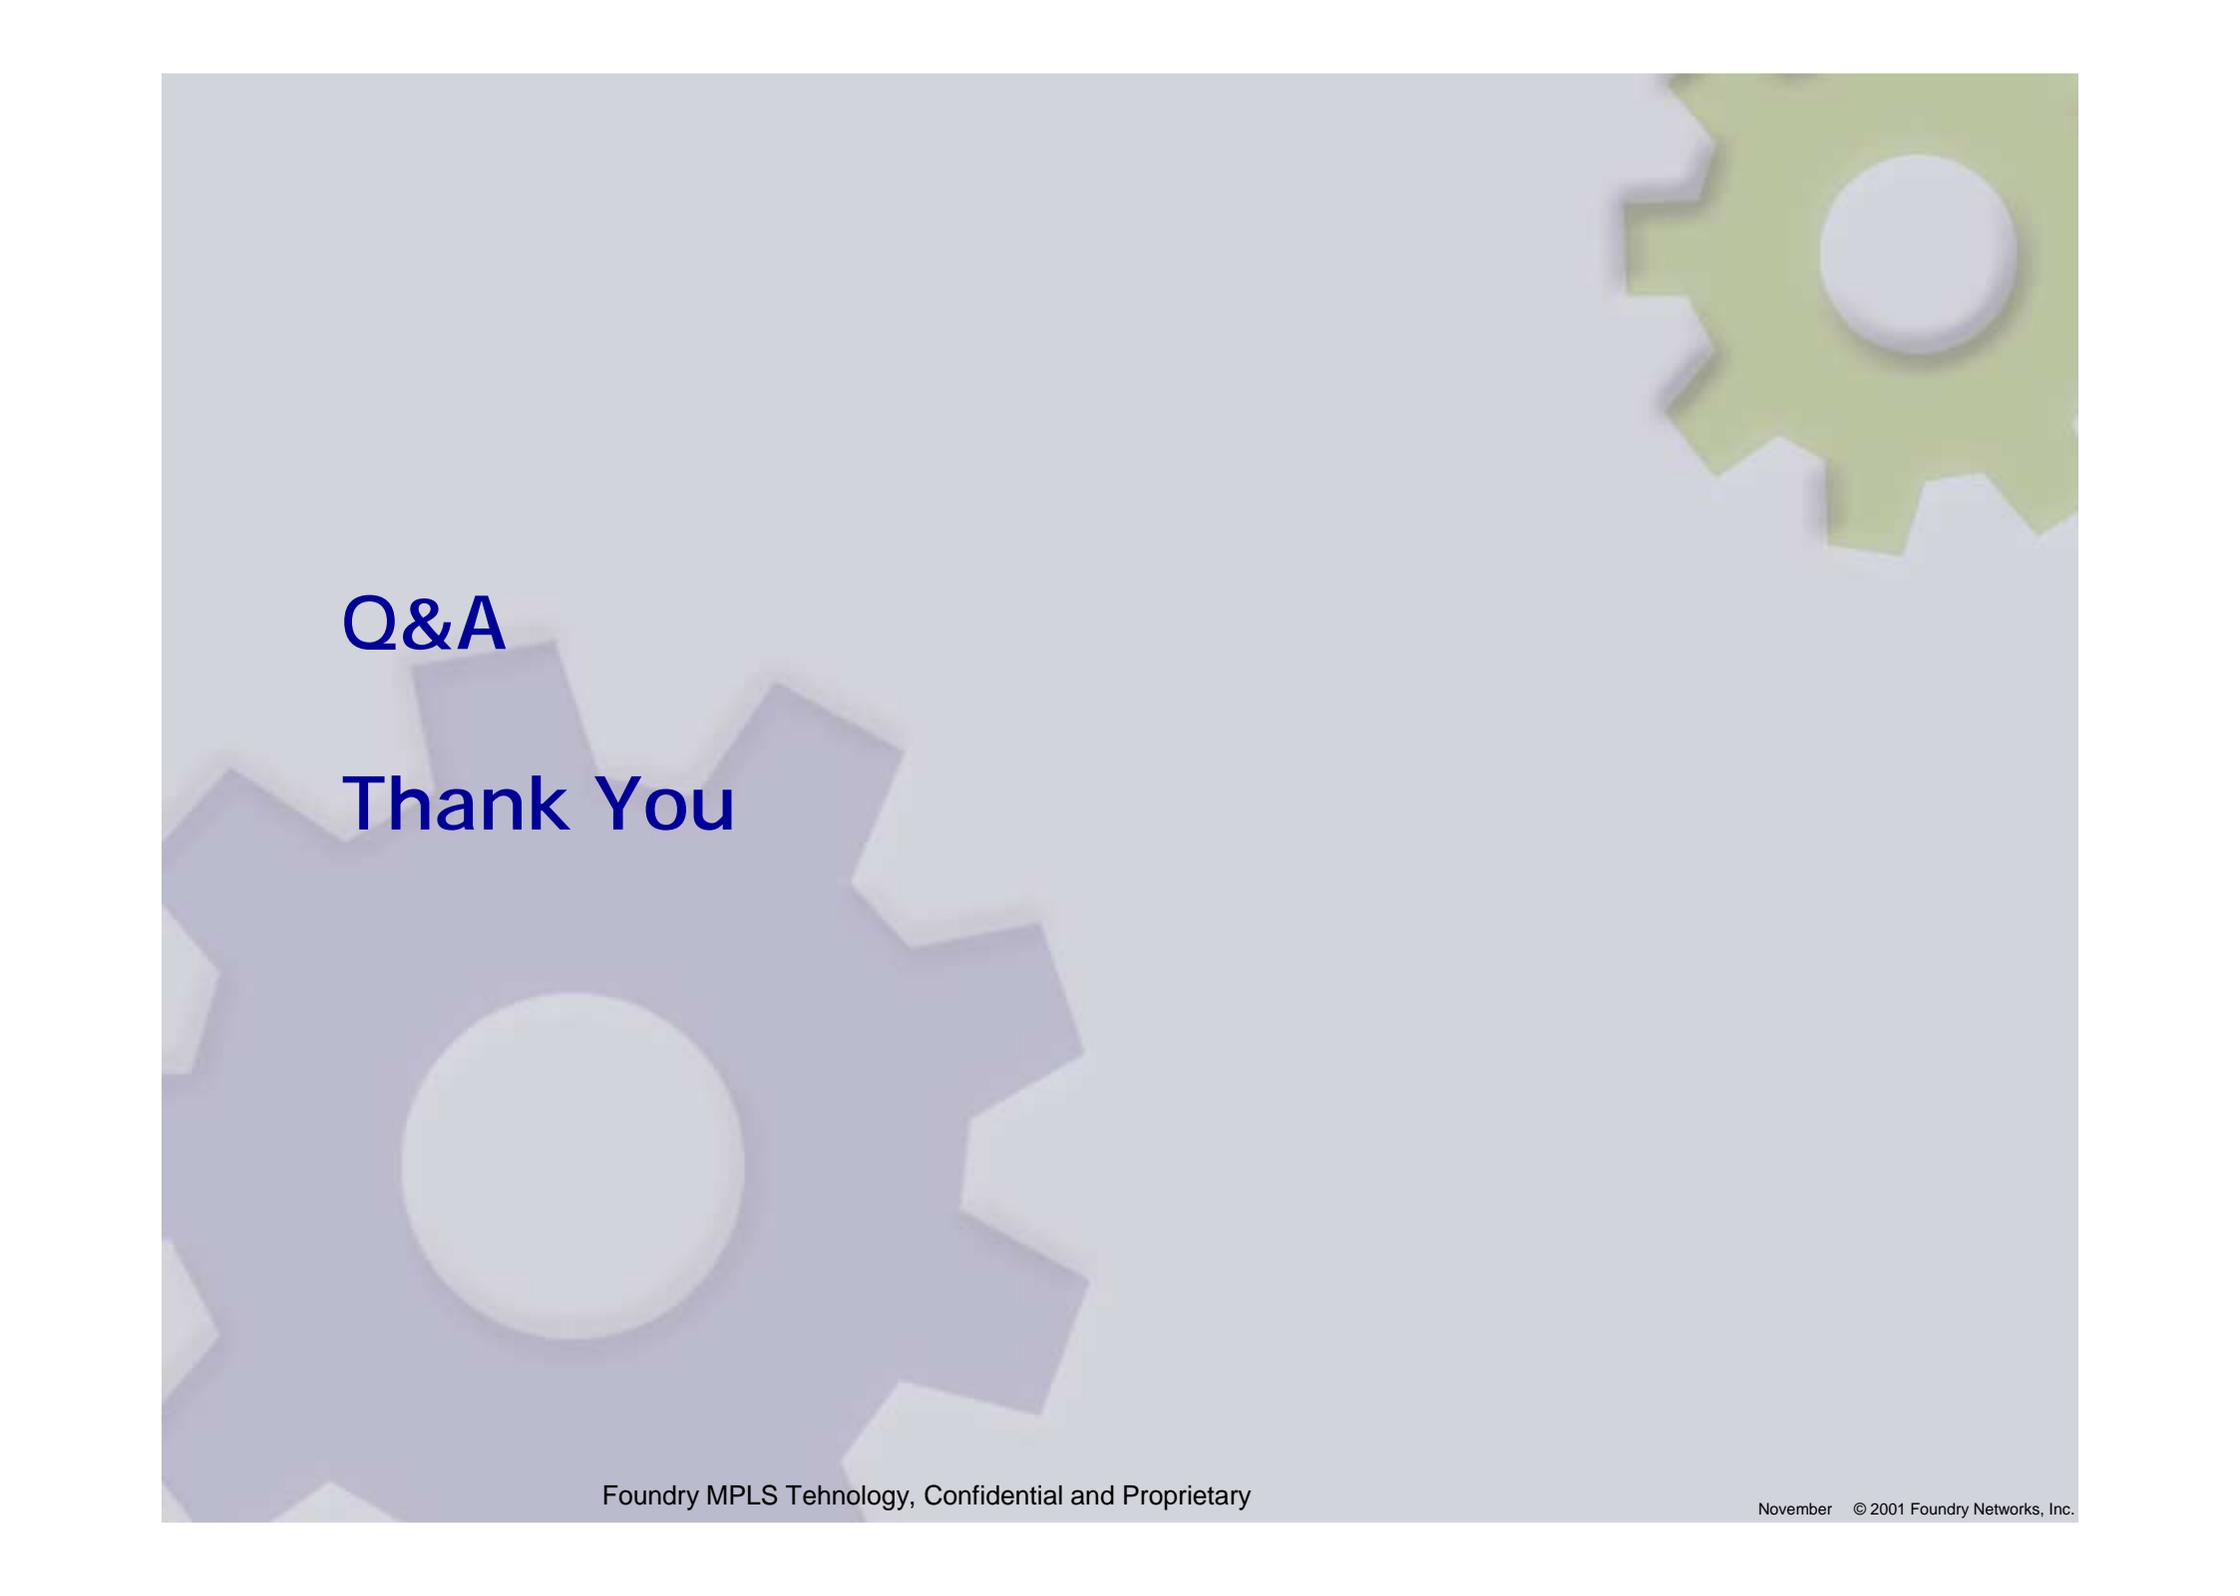


VPLS Example Application: Dual Homing the Customer

- The customer might run and manage their own instance of STP
- The provider does not run STP on their backbone, they just carry the customer's BPDUs

STP used by the customer but not by the provider



The background of the slide features a light gray gradient. On the left side, there is a large, semi-transparent purple gear. On the right side, there is a smaller, semi-transparent green gear. Both gears have a white circular center and a drop shadow effect, giving them a 3D appearance.

Q&A

Thank You