MPLS over Various IP Tunnels

W. Mark Townsley

Generic MPLS over IP

Manual, Point to Point Tunnel

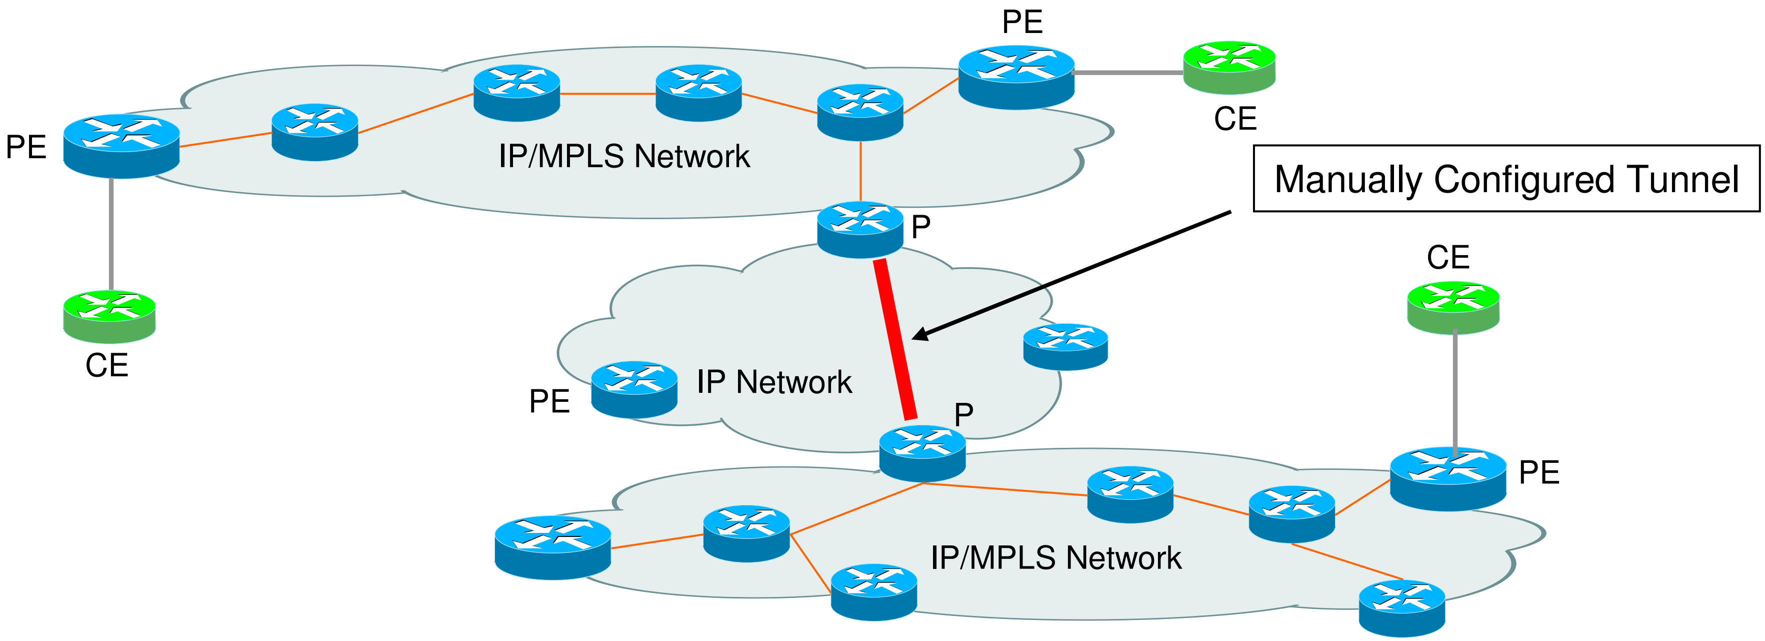

- Typically a GRE tunnel, but may use other encapsulation
- Connects disparate MPLS networks over IP
- Acts as a single MPLS network, so all services enabled by MPLS are available across both clouds

Generic MPLS over IP

Manual. Point to Point Tunnels

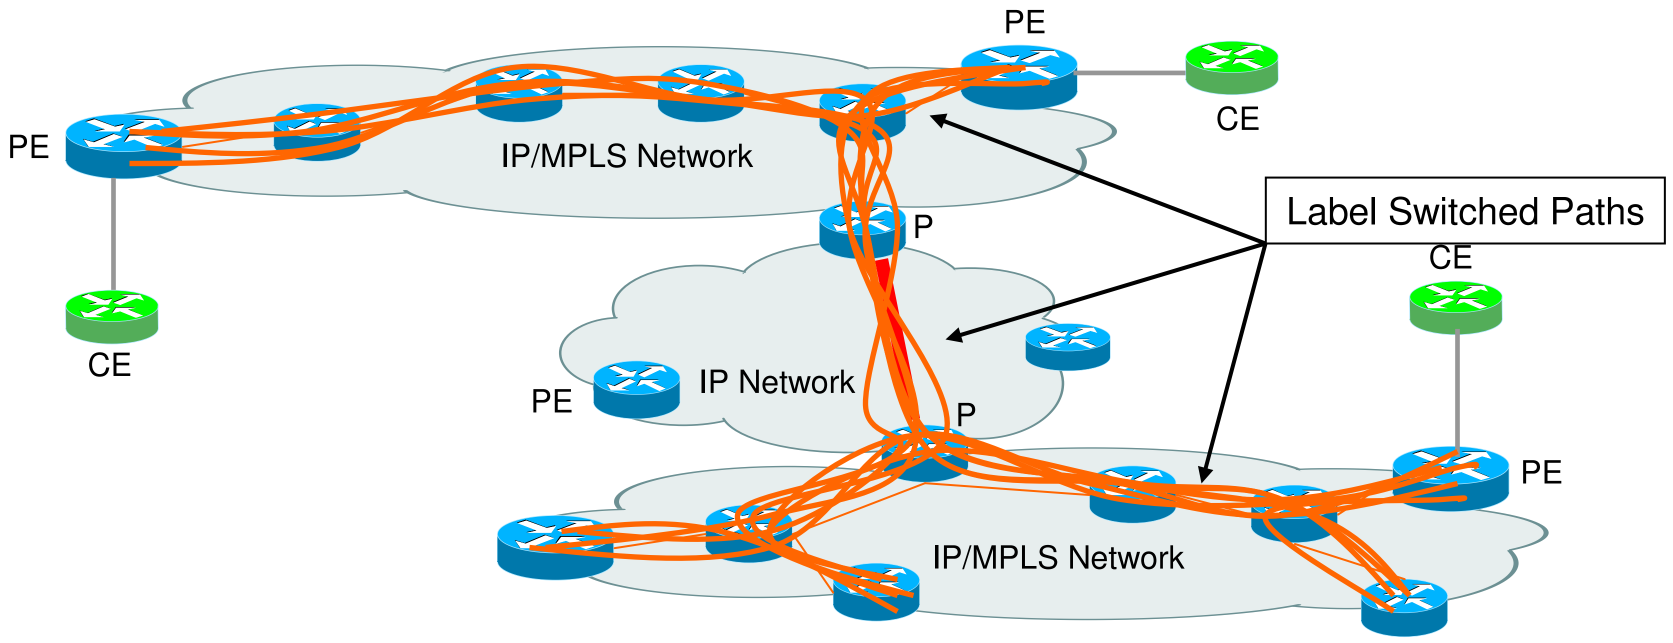

- Tunnel acts as a link layer between MPLS networks
- LSPs are still setup between all nodes as if directly connected on the same MPLS network

Generic MPLS over IP

Manual, Point to Point Tunnels

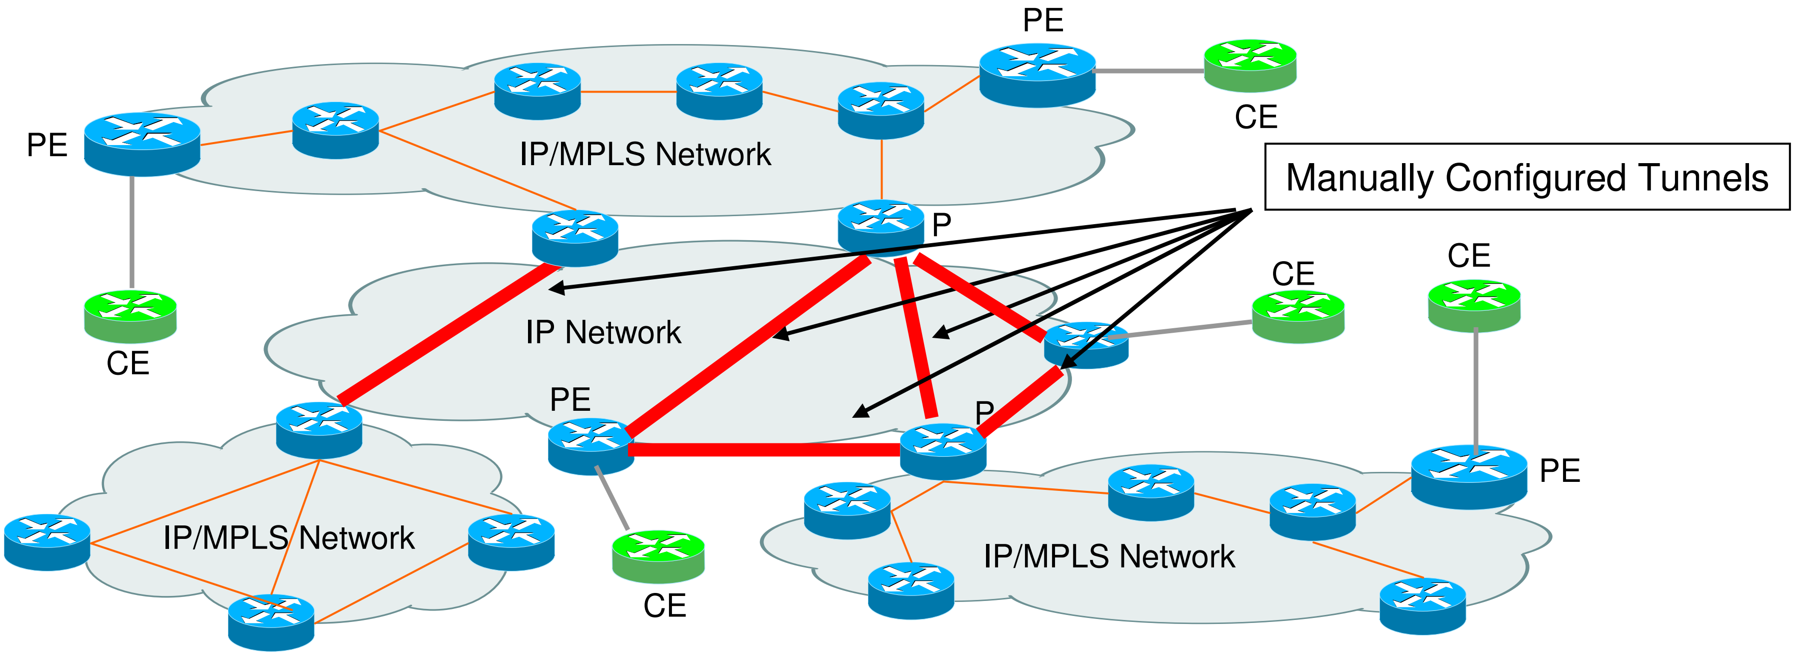

- With multiple MPLS networks and multiple IP-only PEs participating, manual configurations may become cumbersome

MPLS over IP for 2547 VPN Support

- Targeting a specific MPLS application gives us more options.
- Instead of manually configuring tunnels, “Tunnel Reachability Information” is sent via a BGP Next Hop Tunnel SAFI (draft-nalawade-kapoor-tunnel-safi-01.txt)
- This advertises which tunnel method is best to reach a given PE. i.e., MPLS/LSP, MPLS/GRE, MPLS/IP, MPLS/L2TPv3, MPLS/IPsec, etc.
- Includes any parameters necessary to select a given tunnel at a particular PE (IPsec policies, L2TPv3 Session/Cookie, protocol type, etc.)
- No additional configuration necessary beyond locally enabling the encapsulation mode. IPsec is an exception, as it requires IKE for Security Association setup.

MPLS (RFC2547) VPNs over IP
Extending the reach of MPLS

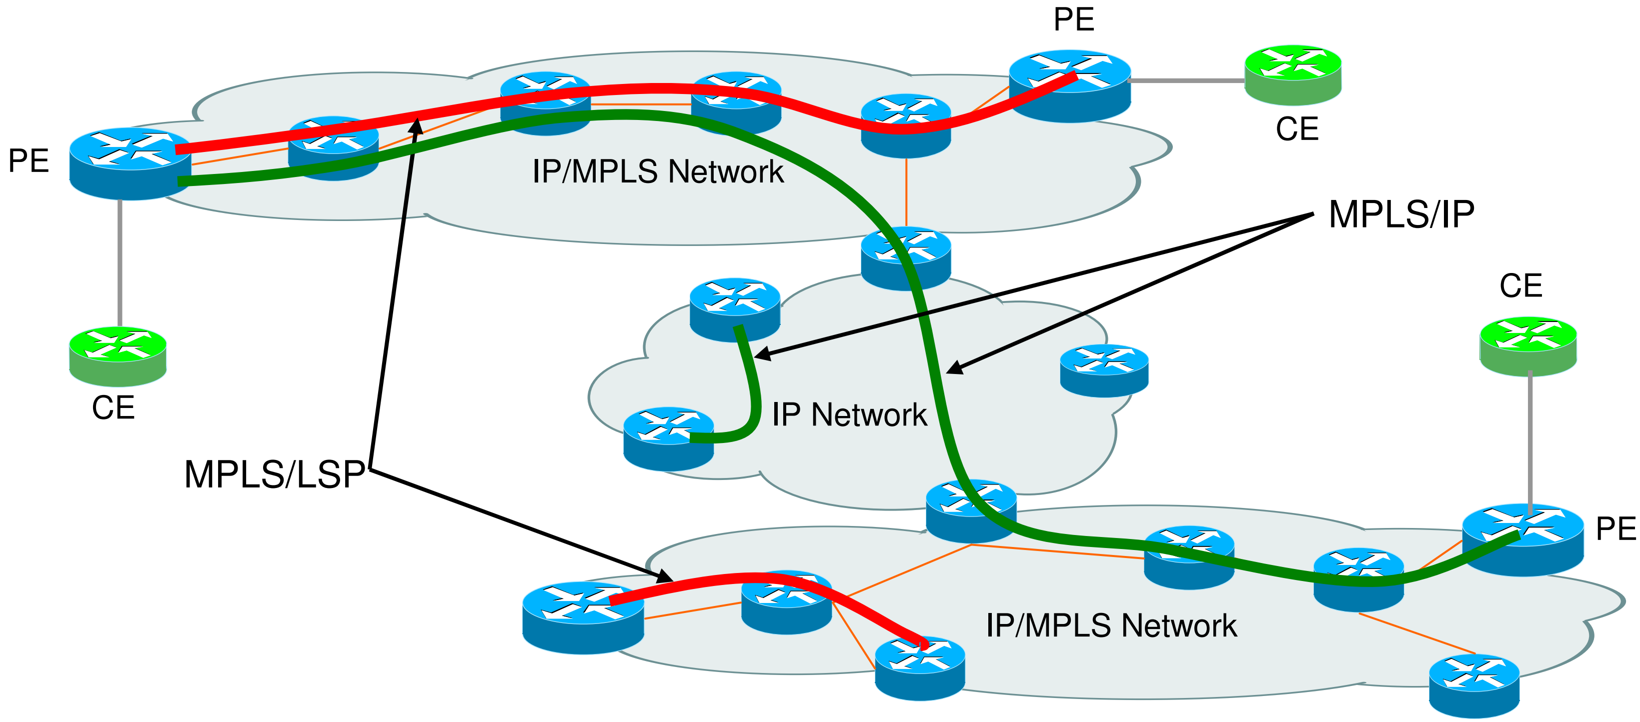

- MPLS/LSP is used when possible, MPLS/IP when not
- MPLS networks need not setup LSPs to reach one another across clouds, only IP reachability between PEs is needed.
- Useful in MPLS migration scenerios

MPLS (RFC2547) VPNs over IP

“Native IP” Core

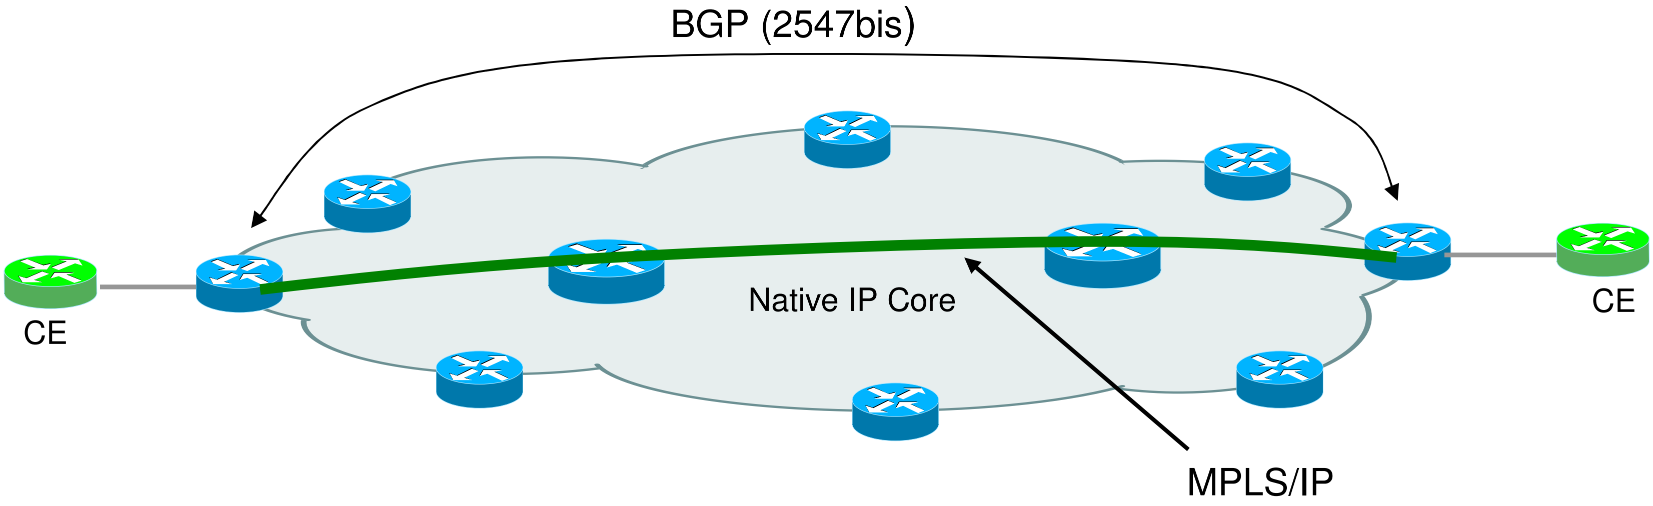

- Core remains IP-only.
- PEs run MPLS only at the edge
- Deploy RFC 2547 service without moving to MPLS core right away

RFC 2547 VPNs:
Cons of MPLS/IP vs. MPLS/LSP

- MTU decreased by at least 16 bytes
- An IP core may be more vulnerable to spoofing attacks vs. an isolated MPLS core
- Potential Interoperability issues due to multiple encapsulation options

Encapsulation Options

- Each of these modes are referred to in one or more IETF drafts
- MPLS over IP
- MPLS over GRE
- MPLS over L2TPv3
- Any of the above with MPLS over IPsec transport mode.
- Which to choose?

Spoofing Attack w/Internet Backchannel

2. SP Core enables attack by allowing only a single packet through to an MPLS/GRE/IP router, and to the Customer VPN by guessing one valid MPLS label

3. Host within customer VPN responds, sending packet through firewall over Internet

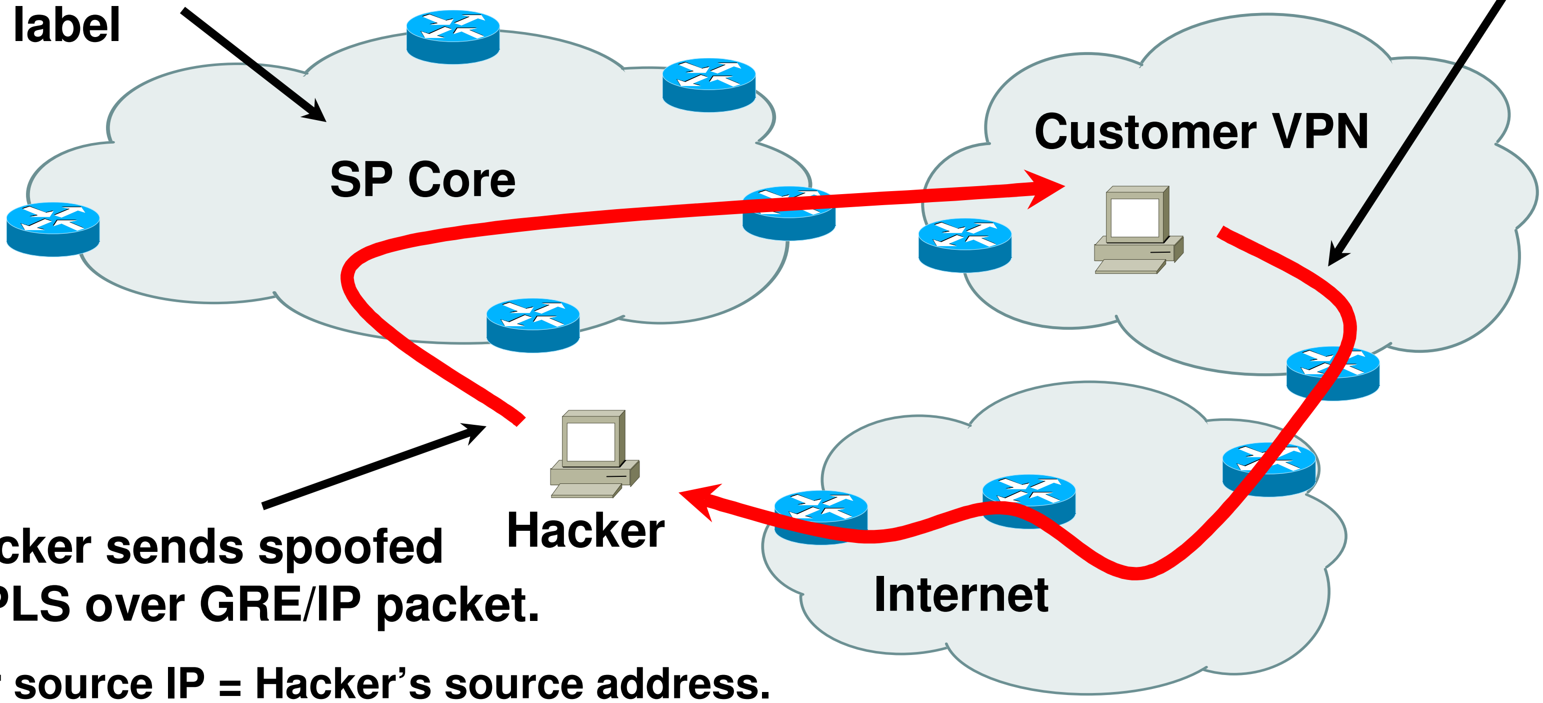

1. Hacker sends spoofed MPLS over GRE/IP packet.

Inner source IP = Hacker's source address.
Inner Dest IP = Enterprise Host
Outer source IP = SP PE source address.
Outer Dest IP = Any SP Router with MPLS/GRE/IP

Packet Spoofing Attacks

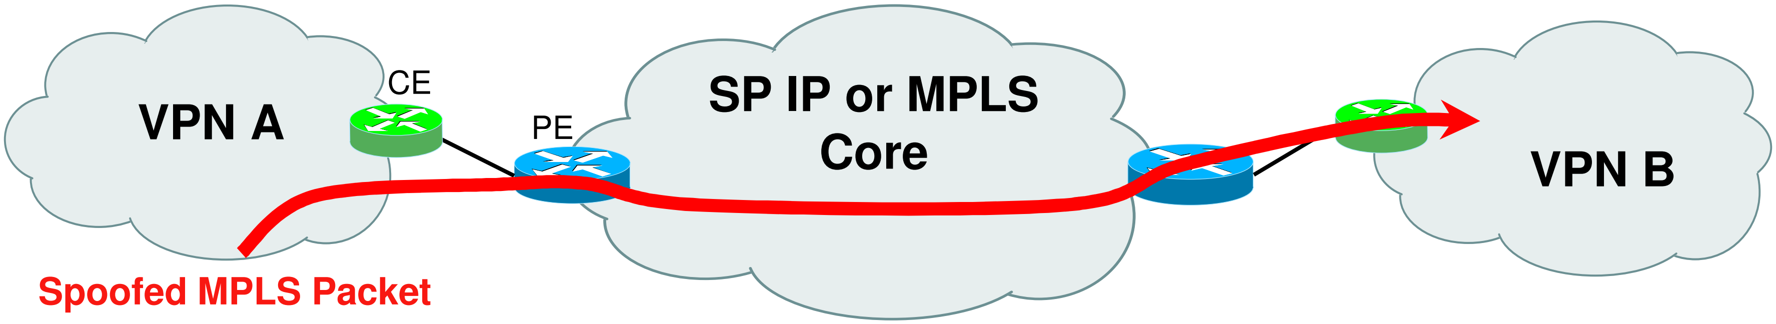

If MPLS VPN packets can make into your core…

20-bit Label	Exp	S	TTL
Rogue PDU (Hacker's choosing)			

Assuming the hacker can send 5000 pps to a PE with 4000 routes, all possible valid labels may be found in 3.5 minutes (an average of 2 discovered per second).

Packet Spoofing Attacks
Isolated MPLS Core

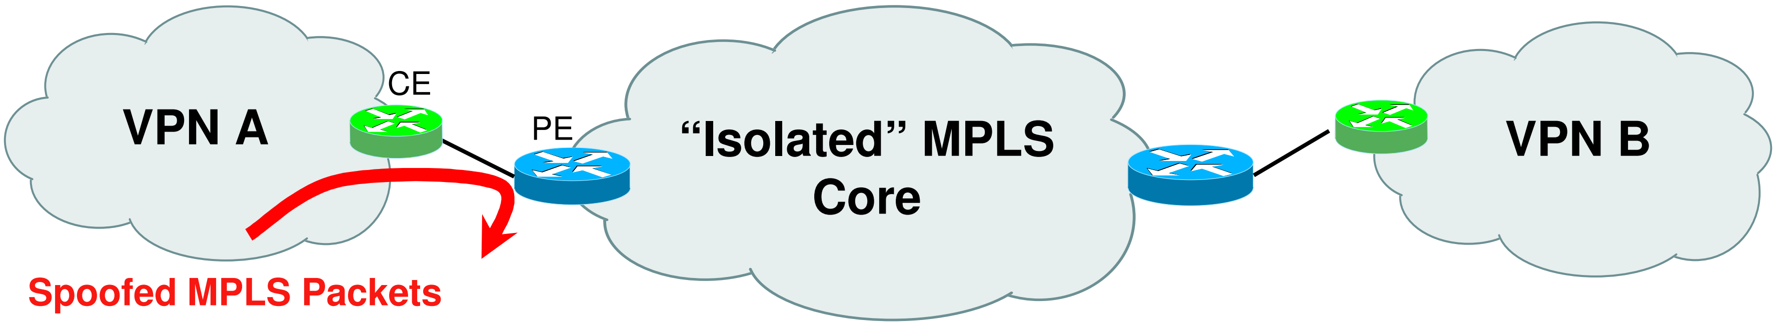

- **“For security reasons a PE router should never accept a packet with a label from a CE router. “**
  - **(draft-behringer-mpls-security-04.txt, section 3.4)**
- **As long as this holds true, all spoofed MPLS packets from the CE are dropped at the customer interface, unable to reach into the MPLS core.**

Packet Spoofing Attacks
MPLS over GRE/IP

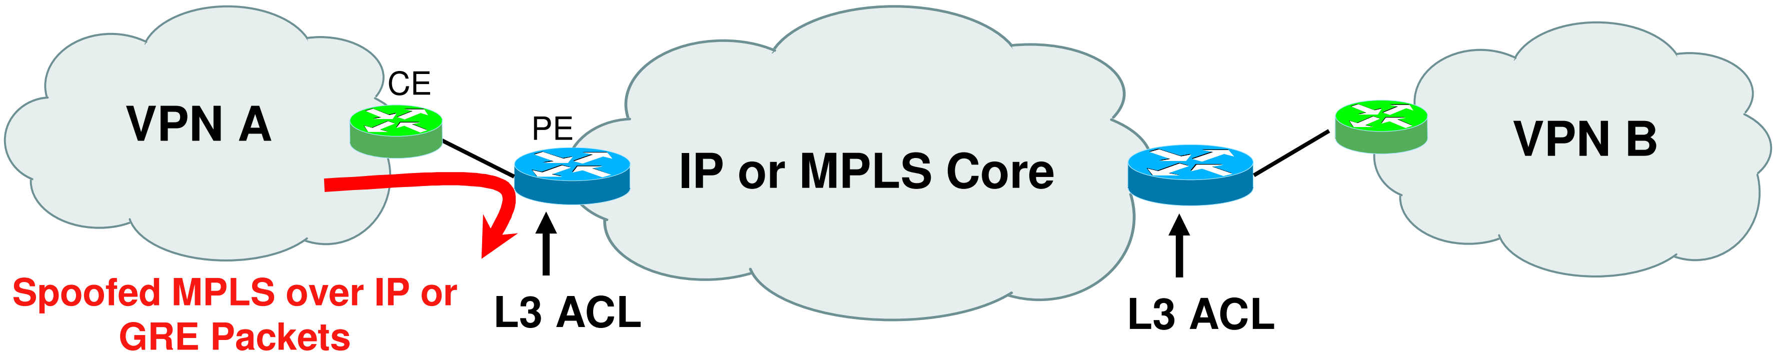

- Enabling MPLS over IP anywhere requires that L3ACLs be maintained across the entire network boundary.
- This may be difficult to maintain, subject to configuration errors, etc.
- Given the ease of spoofing a packet by a *"blind attacker"* it could be dangerous to rely on L3ACLs for MPLS over IP

Blind Insertion Attack

- The aim of the hacker is not to disrupt your core, but to transit the core network to gain access to or disrupt the VPN.
- Hacker can send a packet into your core network and hit a VPN PE (e.g., L3ACLs fail)
- Hacker does *not* have the sophistication to capture and decode packets in the core for use in a orchestrated attack

Spoofing MPLS over IP

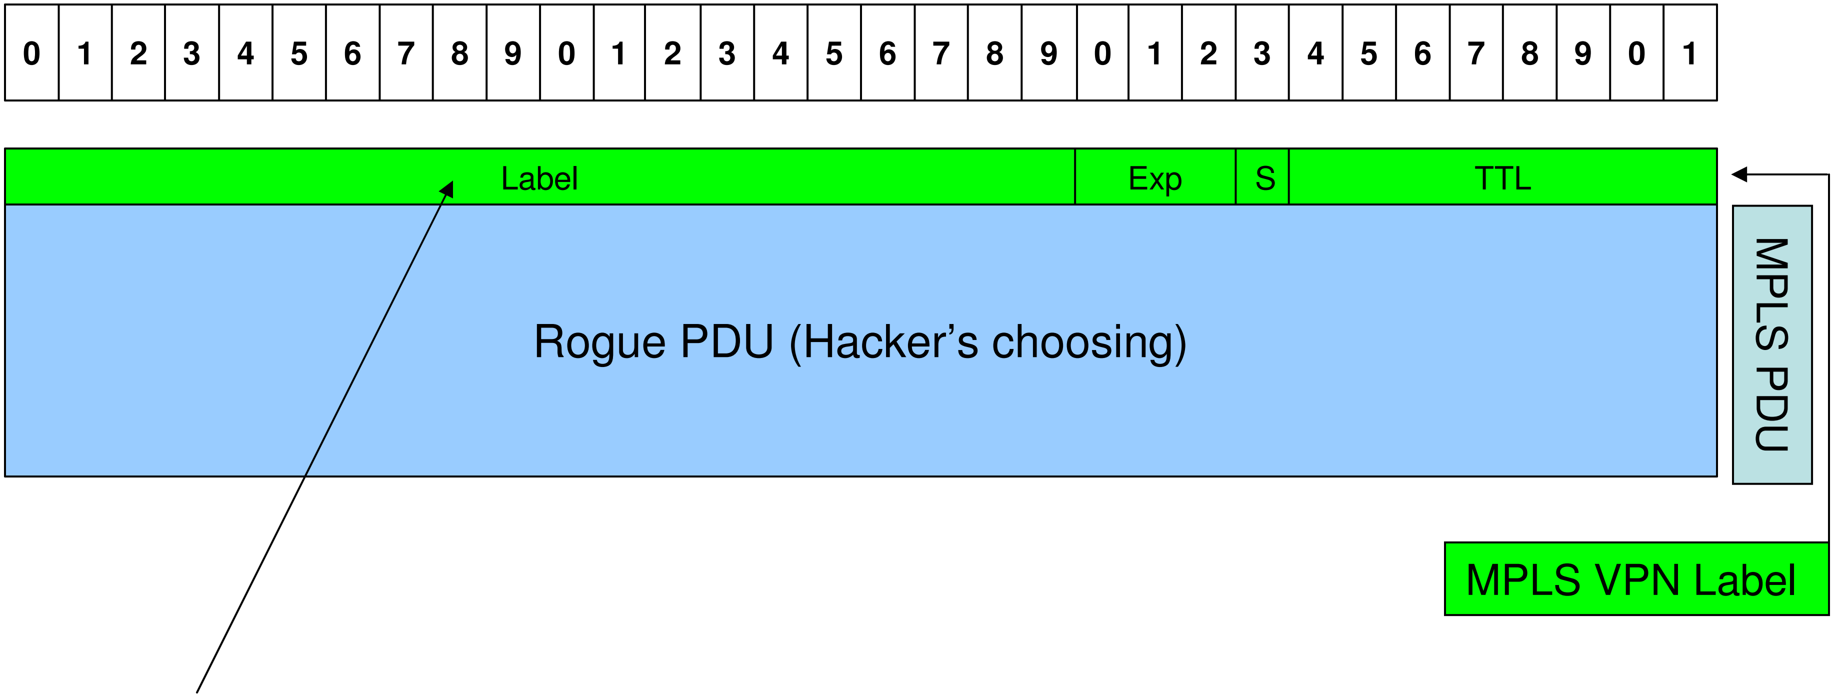

One correct guess at the 20-bit MPLS label, and the Hacker wins

Spoofing MPLS over GRE

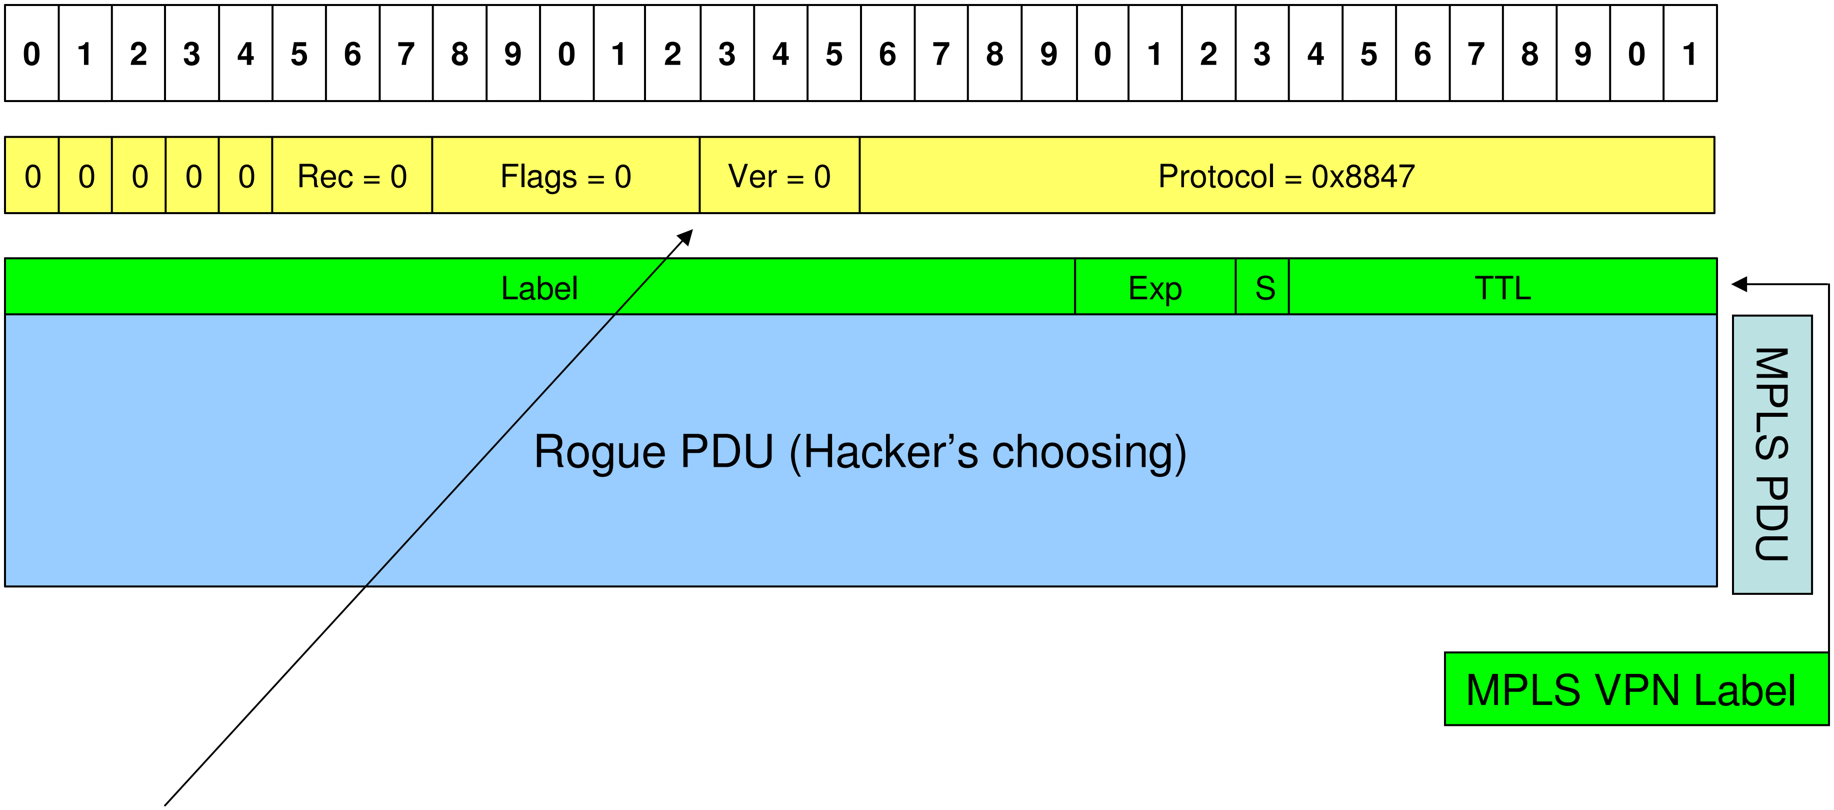

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1

0 0 0 0 0 Rec = 0 Flags = 0 Ver = 0 Protocol = 0x8847

Label Exp S TTL

Rogue PDU (Hacker’s choosing)

MPLS PDU

MPLS VPN Label

No help here as the GRE header is set with constant, well-known values. The same 20 bits must be guessed as with MPLS over IP

Spoofing MPLS over L2TPv3

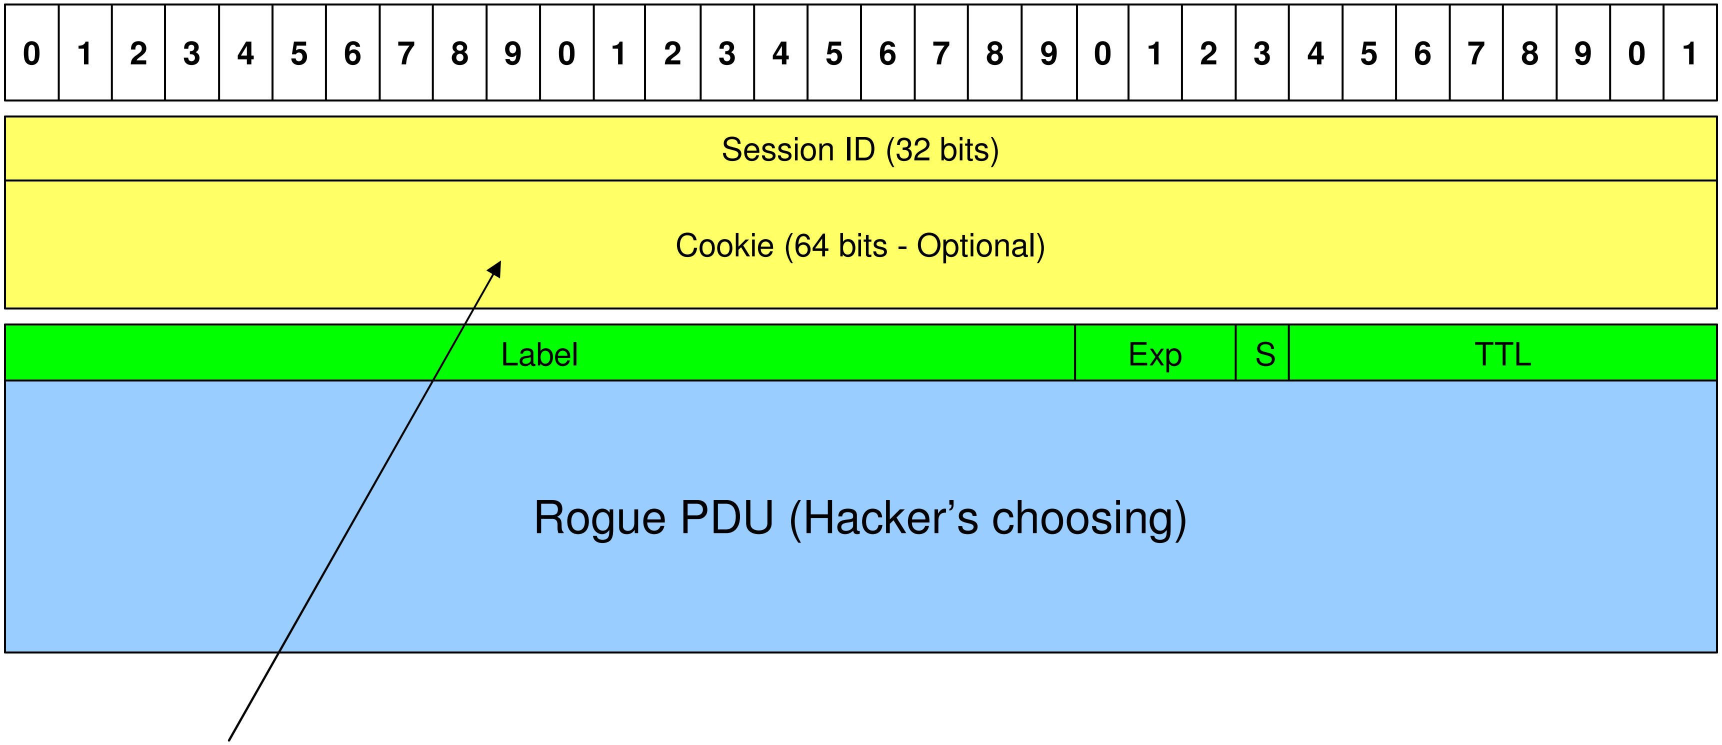

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1

Session ID (32 bits)

Cookie (64 bits - Optional)

Label | Exp | S | TTL

Rogue PDU (Hacker's choosing)

Hacker must guess 64 cryptographically random bits, in addition to the MPLS label.

Attacking at 10 Mpps, a 64-bit cookie will average on the order of 15,000 years to guess one correct value.

Packet Spoofing Attacks

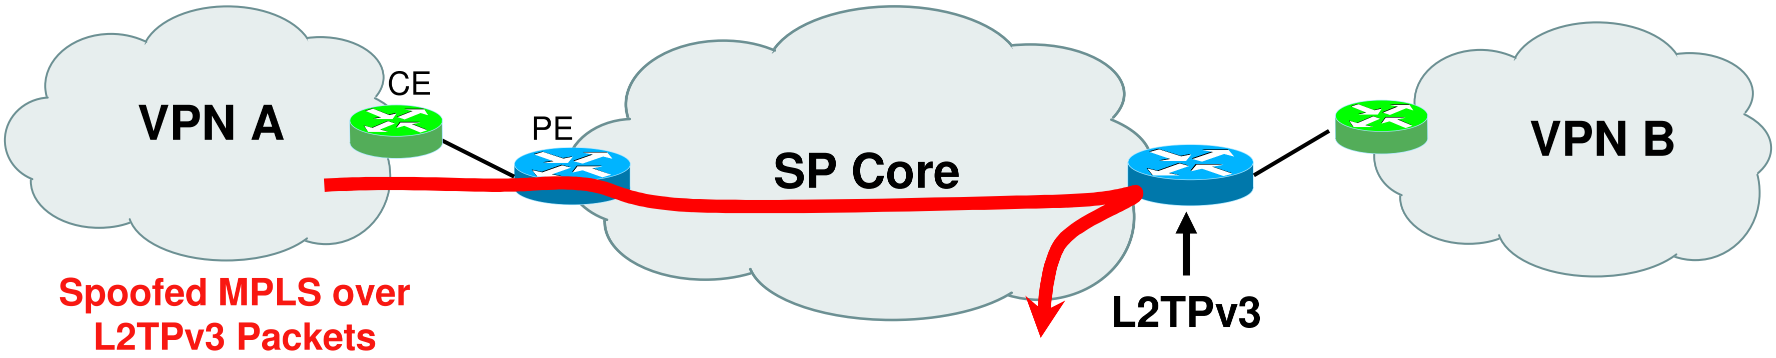

- **If boundary protections fail and MPLS packets can enter your core, L2TPv3 offers a second layer of spoofing protection**
- **Very lightweight: No additional configuration necessary vs. MPLS over IP or MPLS over GRE**

What about IPsec?

- **All** MPLS over IP encapsulations may be protected by IPsec transport mode (GRE, IP or L2TPv3).
- To IPsec, this looks like “host to host” security. There is no “IPsec tunneling” involved.
- Only packets from authenticated sources are processed, so the VPN is protected from packet spoofing attacks, including ones where the hacker can sniff the core

Packet Spoofing Attacks: IPsec

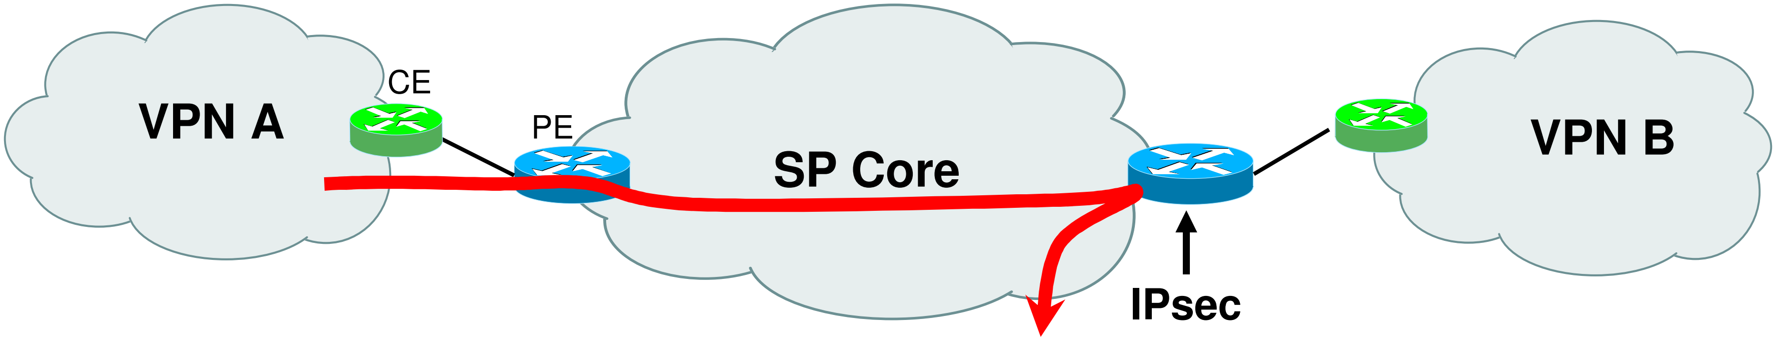

- IPsec provides full cryptographic protection of each packet traversing the SP Core, certainly protecting against packet spoofing
- Heavyweight solution: Requires provisioning a full mesh of p2p IKE (Internet Key Exchange) sessions, cryptographic acceleration, synchronization of IPSec state with other control planes (PE Reachability w/MP-BGP, IGP, LSP), etc.

Summary

- **MPLS over IP may be leveraged for**
  - Migrating to MPLS
  - Enabling MPLS applications across multiple, disparate MPLS networks
  - Enabling MPLS applications over a “Native IP” core network, using MPLS only at the edge
- **IP Tunnels may be configured manually to carry MPLS, or dynamically for certain MPLS applications**
  - Manually configured tunnels link disparate MPLS networks or IP-only PEs into one larger MPLS network
  - MPLS “edge applications” such as RFC 2547 VPNs may be operated over IP without manually configuring IP Tunnels.
- **There are a variety of MPLS over IP encapsulations to choose from**
  - MPLS directly over IP is the most efficient encapsulation, but the easiest to spoof.
  - MPLS over GRE has effectively the same properties as MPLS over IP, but with a 4-byte larger header
  - MPLS over L2TPv3 has an even larger encapsulation (8 additional bytes), but protects against blind packet spoofing attacks with very little additional operational overhead.
  - MPLS over IPsec is the most secure encapsulation, but has the most operational and encapsulation overhead