Intelligent Information Network

MPLS VPN Security

Klaudia Bakšová

Systems Engineer, Cisco Systems

kbaksova@cisco.com

Agenda

Cisco.com

- **Analysis of MPLS/VPN Security**
  - **Inter-AS VPNs**
  - **Provider Edge DoS possibility**
- **Secure MPLS VPN Design**
  - **Internet Access**
- **Security Recommendations**
- **Summary**

The Principle: A "Virtual Router"

Virtual Routing and Forwarding Instance

Route Distinguisher: Makes VPN routes unique

```
!
ip vrf Customer_A
  rd 100:110
  route-target export 100:1000
  route-target import 100:1000
!
interface Serial0/1
  ip vrf forwarding Customer_A
!
```

Export this VRF with community 100:1000

Import routes from other VRFs with community 100:1000

Assign Interface to "Virtual Router"

General VPN Security Requirements

- Address Space and Routing Separation
- Hiding of the MPLS Core Structure
- Resistance to Attacks
- Impossibility of VPN Spoofing

Working assumption: The core (PE+P) is secure

Address Planes: True Separation!

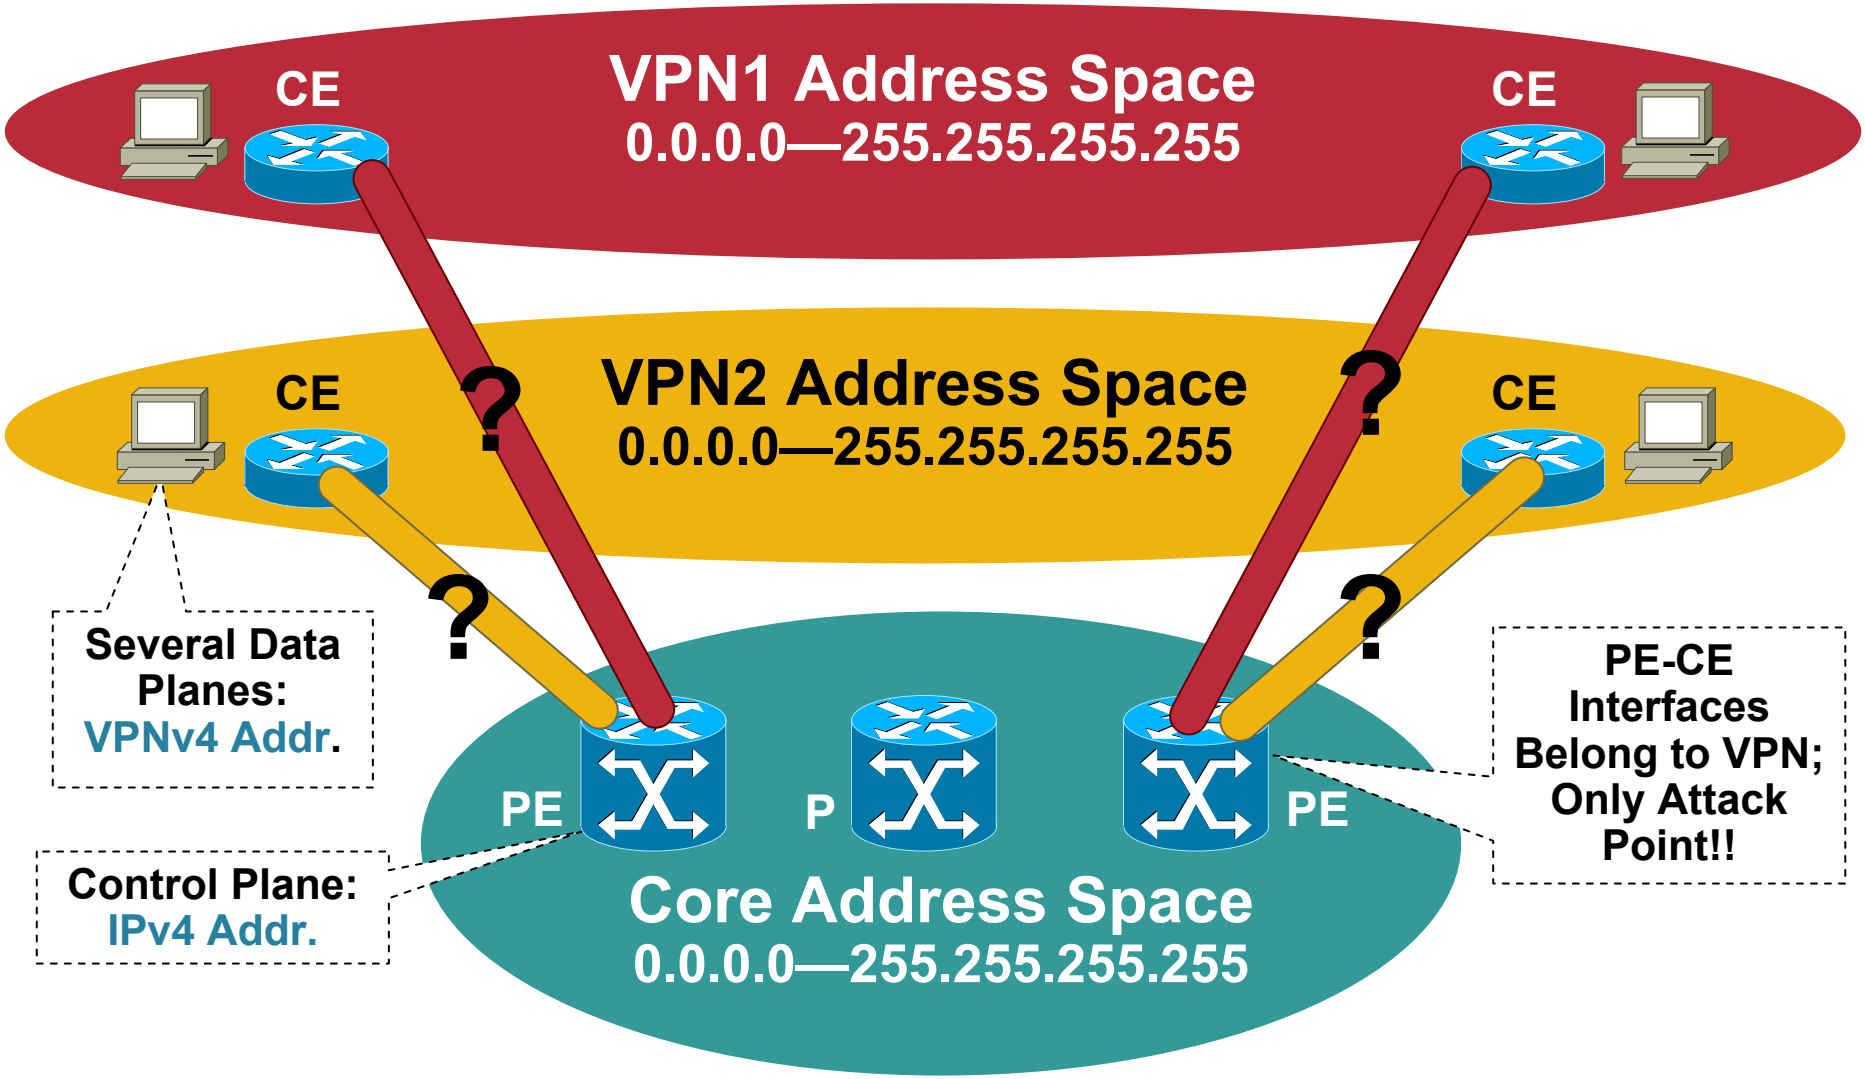

Hiding of the MPLS Core Structure

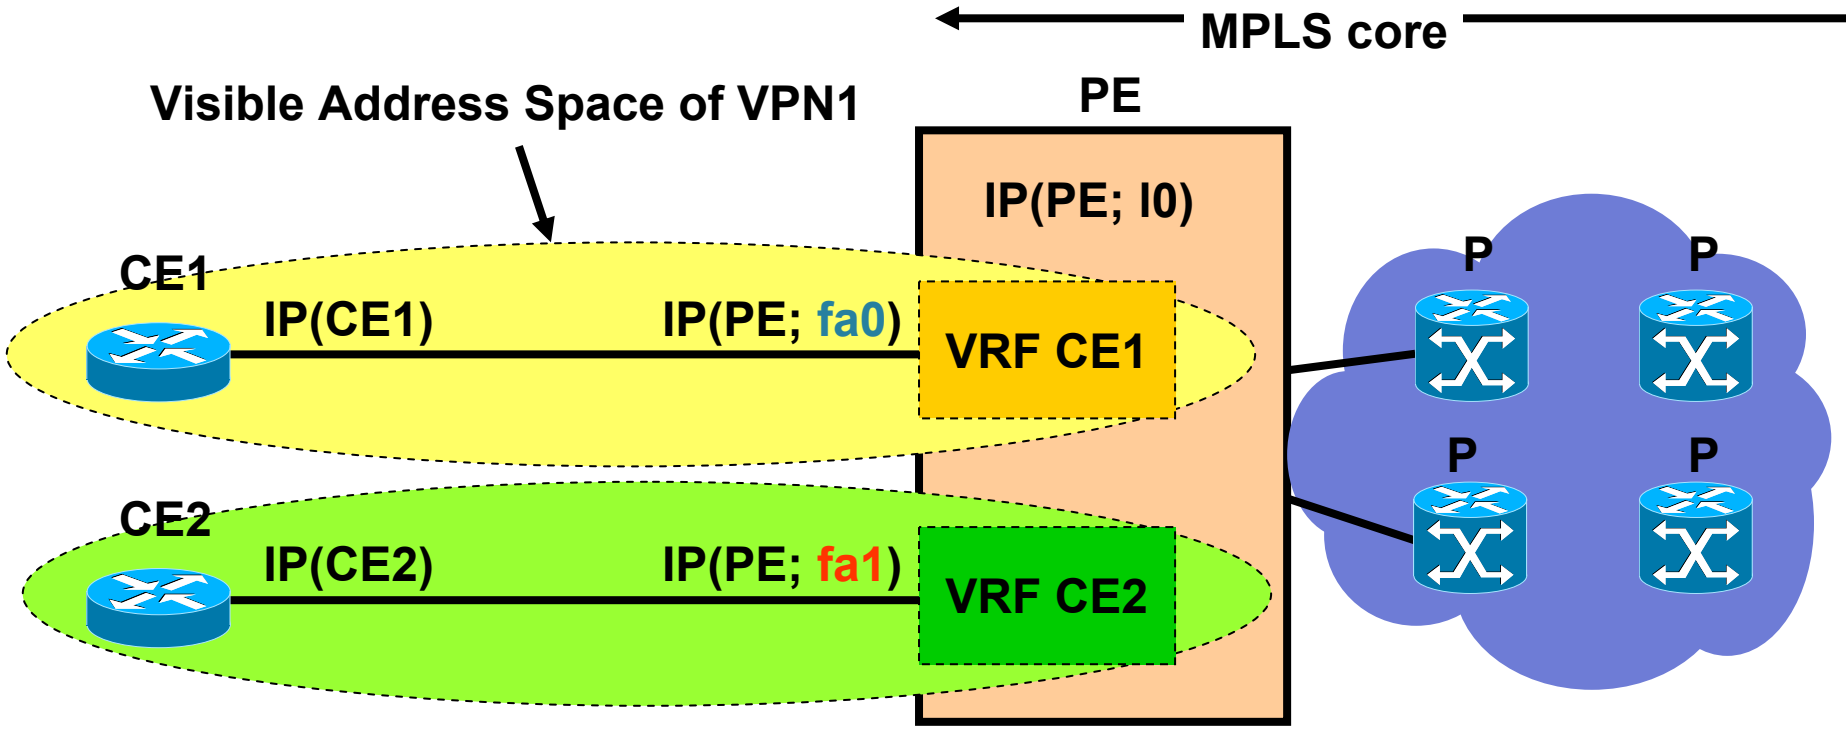

- **PE interface to CE** – the only point where a VPN can 'see' the core and send packets to the core device; seen and accessible from VPN1 space **only**, **VPN1 cannot see any other interface on the PE**
- **Only PE peer addresses of VPN1 exposed** (-> CE)!
  -> ACL for PE interfaces – for 'receive traffic'
- **IP unnumbered** for PE interfaces – complete hiding of the core from that VPN!
- **P routers** – not reachable from VPN

Protection Against Spoofing

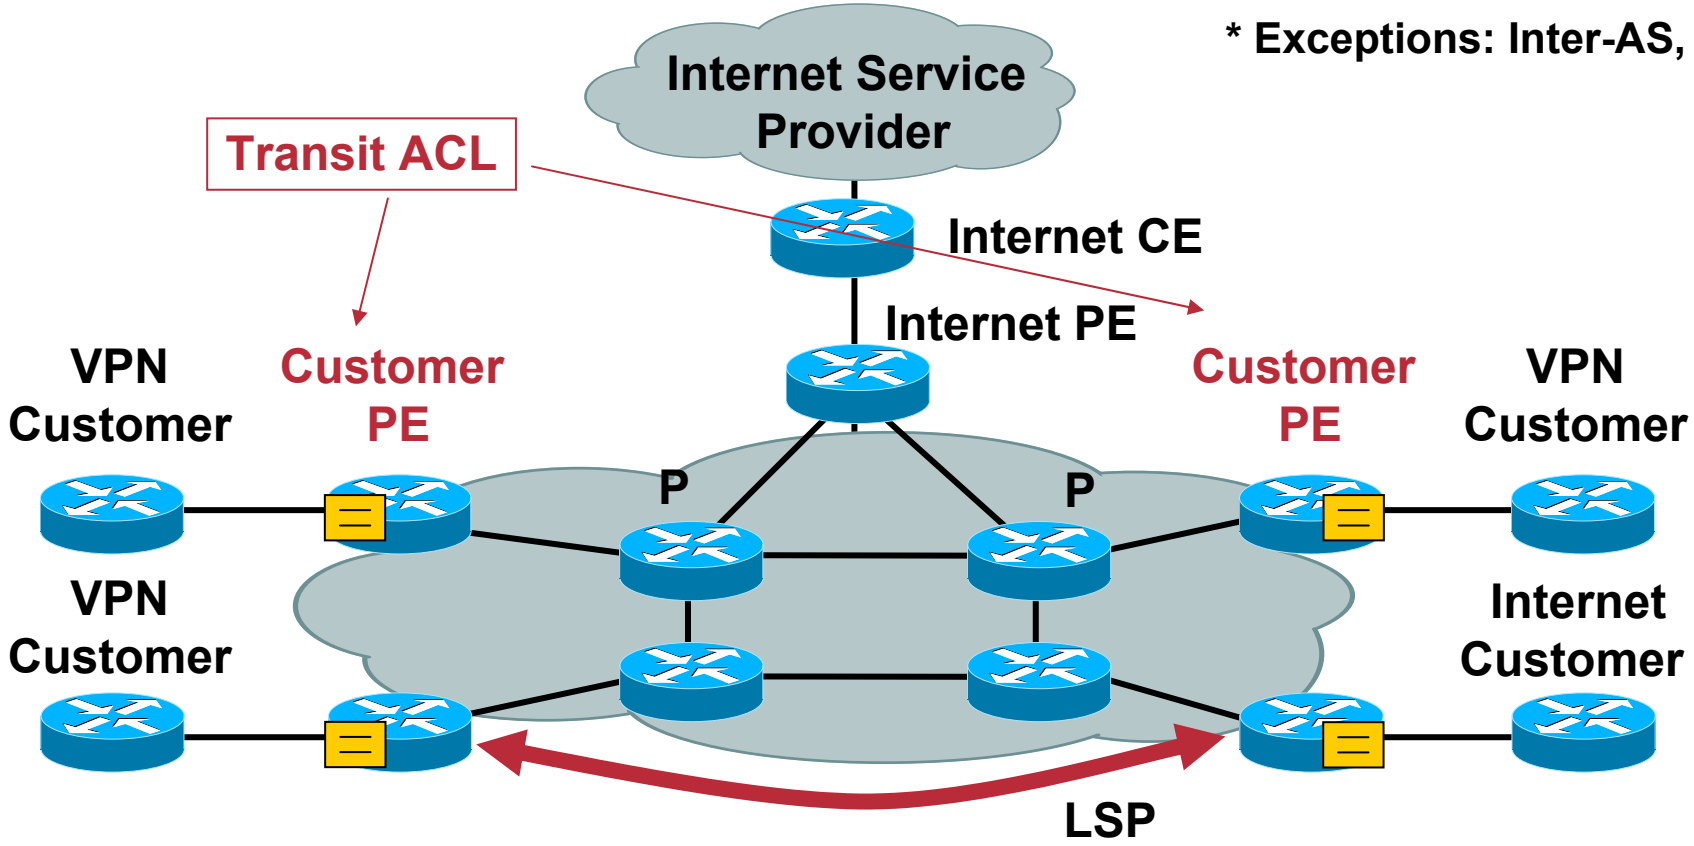

- **Label Spoofing** - Interface between PE and CE – pure IP without labels → labeled packet received from CE, **PE automatically drops it**

  → **Cannot spoof labels from outside!**

- **IP spoofing** – possible, **remains within the originating VPN – RFC2827**

Inter-AS: What are we trying to achieve?

- **An SP should have:**

  100% (full) reachability to all Inter-AS VPNs shared between them (control plane and data plane)

  0% (no) reachability to VPNs that are not shared (control plane and data plane)

- **SP networks should be independent:**

  Must be secured against each other

  Not attackable from outside (other SP, customer, Internet)

Inter-AS:
What Are We NOT Trying to Achieve?

Any Form of Separation Between Inter-AS VPNs (Control or Data Plane) -

- **Interconnection of VPNs is 100%**
- **No firewalling, no limitations, no sanity checks within an Inter-AS VPN**

If an SP Holds VPN Sites in an Inter-AS Set-Up, He Has Full Access to *All* VPN Sites, Also on Other ASes

Inter-AS: Case A
VRF-VRF Back-to-Back

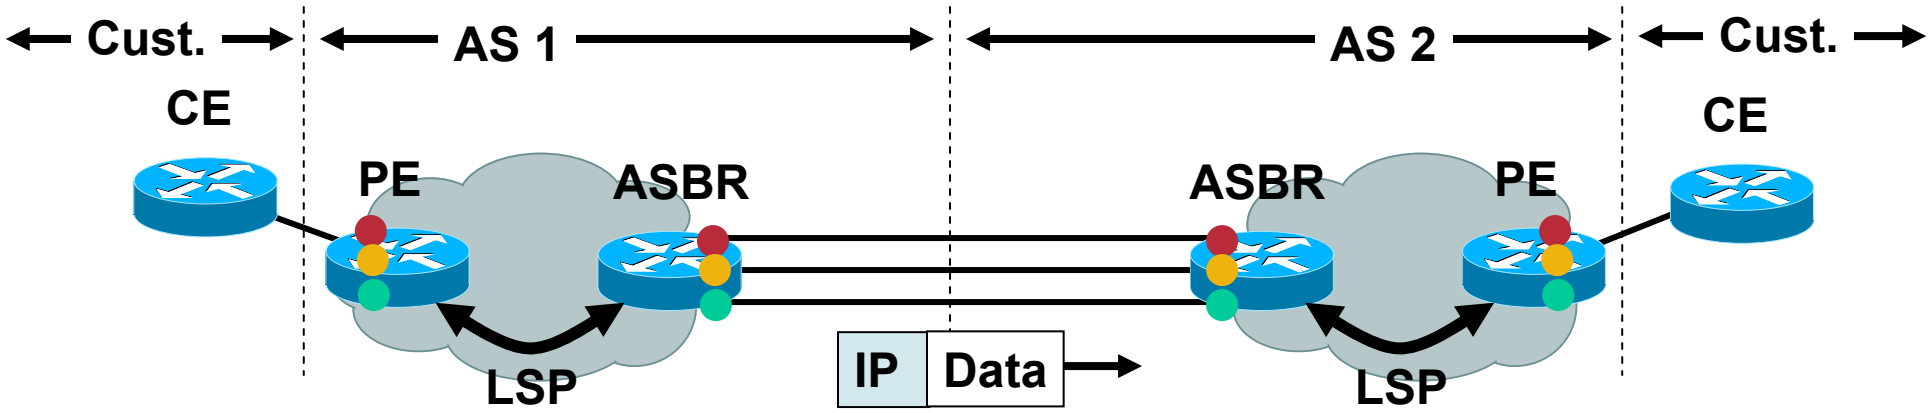

- **Control plane:** No signalling, no labels – interfaces external to AS are pure IP, each ASBR holds its own VRF for the shared VPN

  ASBR - as if a single PE router connecting a CE router (the other ASBR)

- **Data plane:** IPv4 only, no labels accepted
- Not very scalable

Inter-AS: Case A
Potential Security Issues

- **Accidental misconnection at the ASBR** – SPs have to make sure they are clear about which interface/subinterface connects which VPN
- **Routing issues** –VRFs on both ASBRs will exchange routing for a given Inter-AS VPN
  - **Routing security**
  - **Prefix number limited to avoid memory overflow**

- **Security:** as in RFC2547; most secure interconnection model – no labels accepted due to **'PE-CE' analogy**, neighbouring AS cannot see the AS core
- SPs are completely separated, VRF-to-VRF connection, no global routing table connection
- Neighboring ASBR - just an IP interface to MPLS core – **no label spoofing**

Inter-AS: Case B
ASBR exchange labelled VPNv4 routes

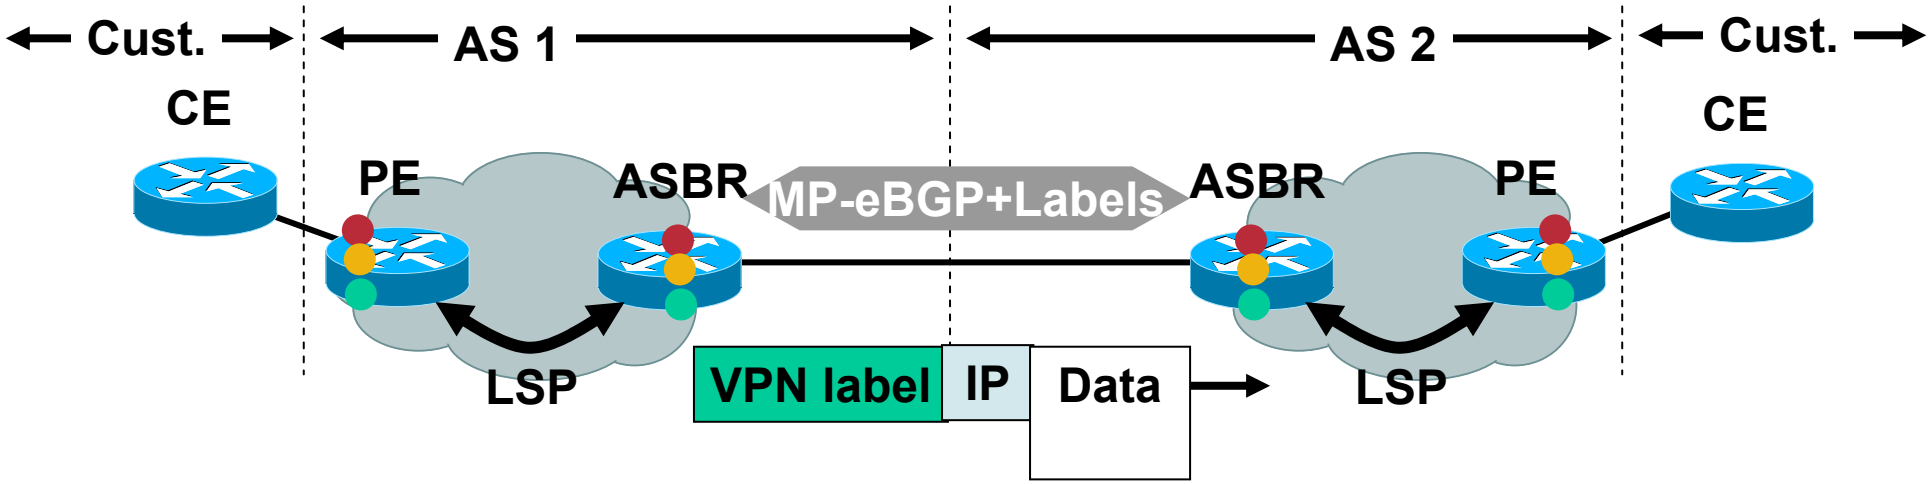

- **Control plane:** MP-eBGP between ASBRs, no IGP or TDP/LDP
- Inter-AS VPNv4 routes held in BGP table, not in VRFs
- **Data plane:** one connection between ASBRs – data plane traffic for different VPNs must be kept separate – labelling packets before sending them to the other ASBR (label stack swapped for ASBR VPN label)

➢ inherent behaviour to MP-eBGP

➢ Better scalability, BGP table size might be an issue

Inter-AS: Case B
Potential Security Issues

- **No AS VPN label is checked on ASBRs when forwarding, => possible label spoofing** => data plane not possible to secure completely
  - External interfaces accept labelled packets instead of just IP packets
  - No way for ASBR to check on the VPN membership of the packet, as there is no VRF on ASBR
- **Control plane**: ingress ASBR interfaces – **ACL to filter any IP accept BGP**
- SPs are completely separate
- Visibility – only the neighbouring ASBR, via eBGP

Inter-AS Case C: ASBRs Exchange PE loopbacks

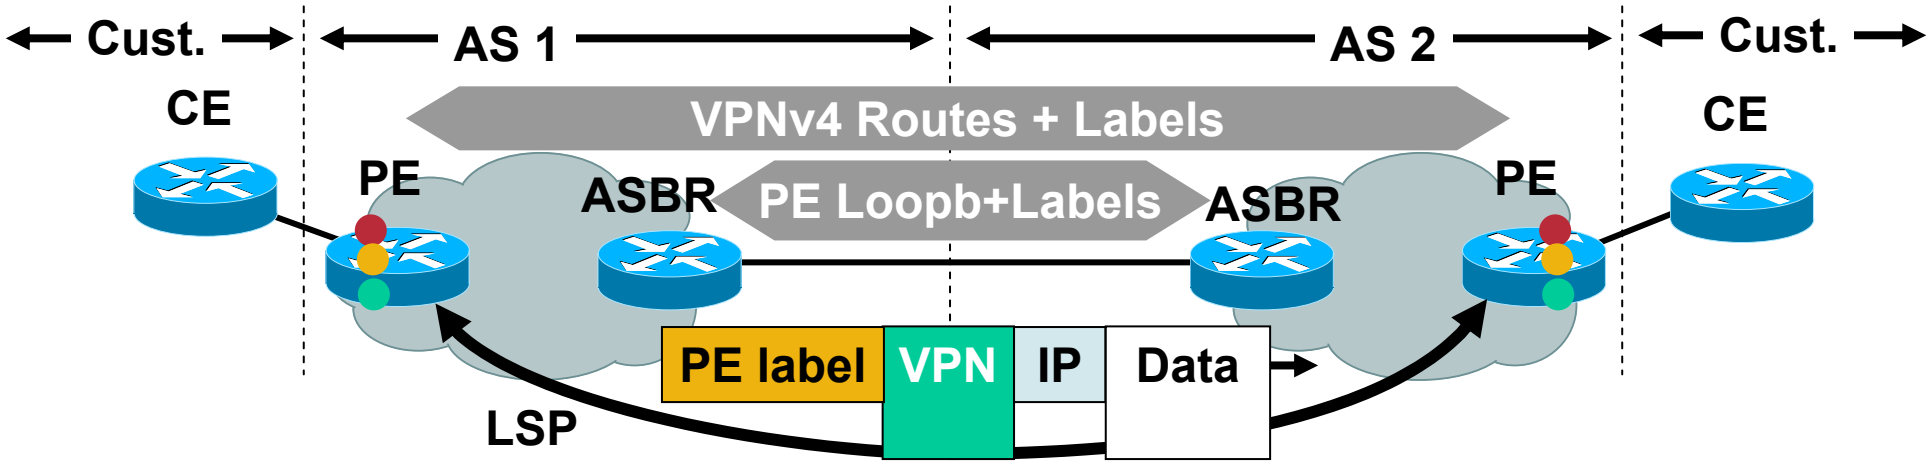

- **Control plane:** PE visibility of both SPs – through Multihop MP-BGP
- ASBR exchange just PE loopback vie eBPG + labels; PEs exchange VPNv4 routes + labels end to end without involving ASBRs => no need to hold VPN specific information, only PE loopbacks and their labels => very scalable
- **Data plane:** PE label + VPN label, ASBRs only as P routers, LSP built from PE in AS1 to PE in AS2

Inter-AS: Case C
Potential Security Issues

- **Security:** SP must be able to reach all PEs of neighbouring AS which hold connections of shared VPNs, issue: ASBR cannot check VPN label, sees only egress PE label, possible VPN label spoofing => probability of mis-insertion
- **Control plane**: ingress ASBR interfaces – ACL to filter any IP accept BGP
- ASBR – no VRF, no VPN routing information => VPN label below egress PE label cannot be checked (e.g. intrusion – no VPN label appended, PHP pops egress PE label at P router, PE receives a pure IP packet – gets routed into SP core)

All these label spoofing attacks carried out by SP, not by customer VPN, as data can be injected at ASBR only!

The Key Issue: Designing a DoS Resistant Provider Edge

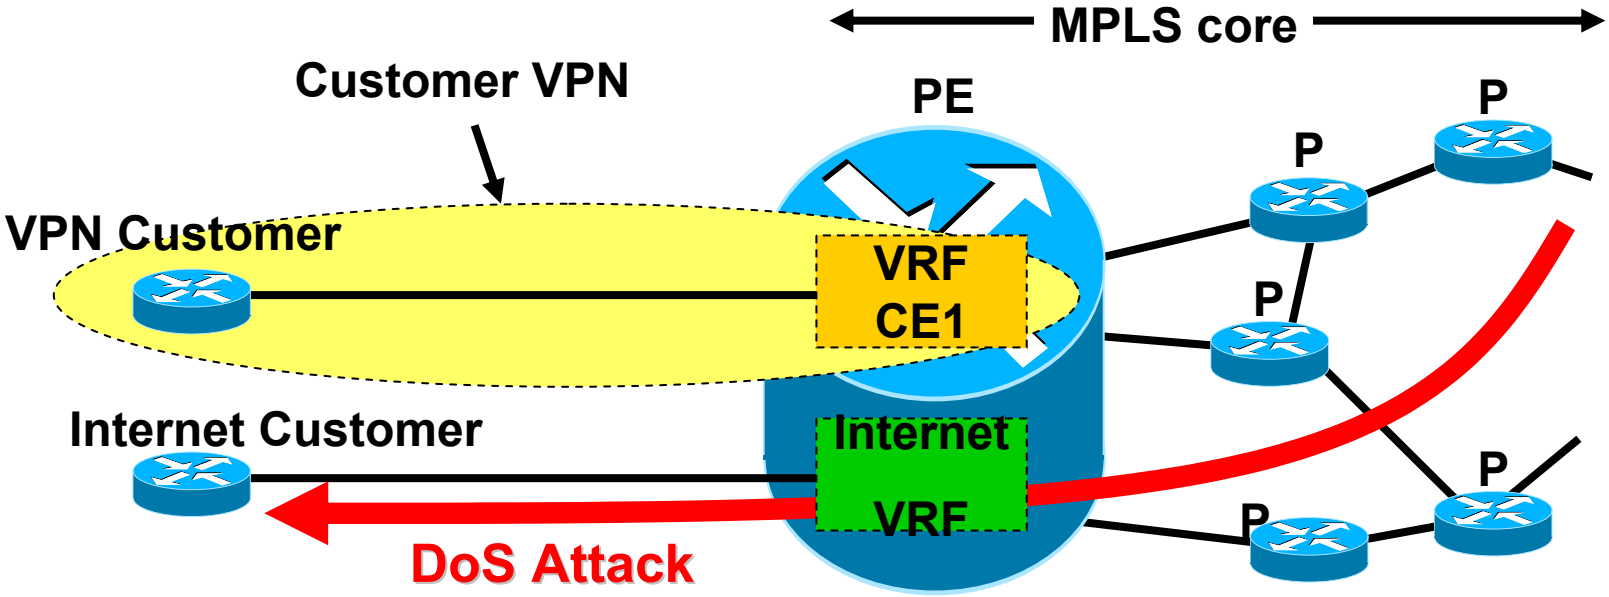

- **Primary prerequisite – IP address visibility**
- **PE has shared CPU / memory / bandwidth resources for different VRFs:**

  → Traffic can affect VPN customer(s) via performance degradation up to complete loss of connectivity
- DoS attacks usually perceived as coming from Internet, however also coming from customer VPNs
- A way to compromise MPLS core – thorough security of PEs crucial to avoid the threat

Today's Best Practice: DoS Through a Shared PE Solved by Using a different design

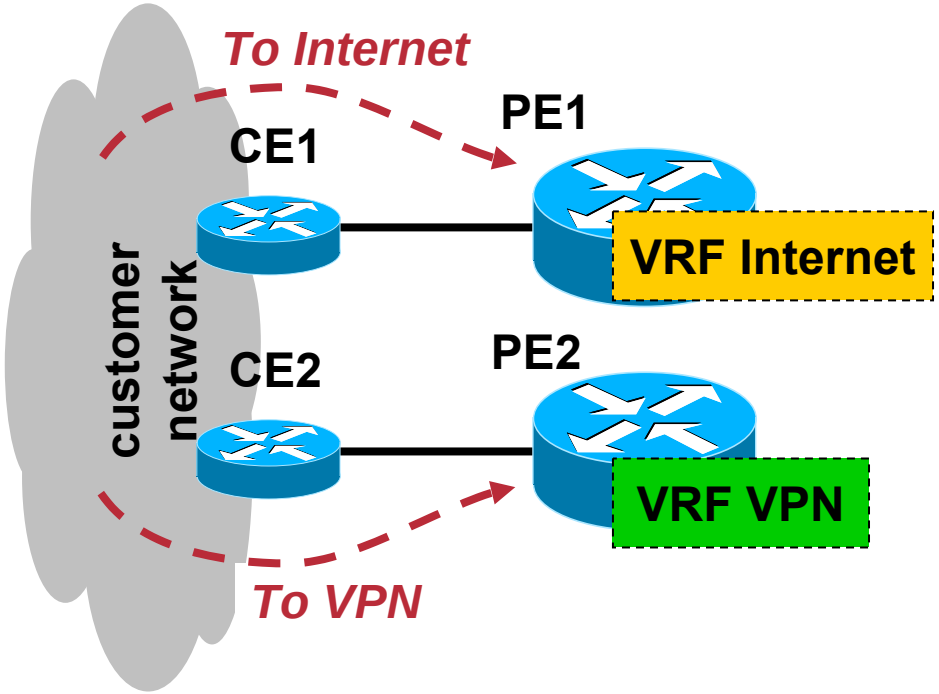

- **Separate VPN and Internet traffic on physically different PE routers**
- **PE routers should contain only VRFs of the same security level. Example:**

  Level 0: Internet

  Level 1: VPN customers

- Internet VPN subject to DoS attack in no different way than other network technologies, i.e. this is not an MPLS-specific issue
- **DO NOT expose PE addresses to Internet at all, or with dynamic routing use limit to routing reachability only – Infrastructure ACL!**

Separate VPN and Internet Access

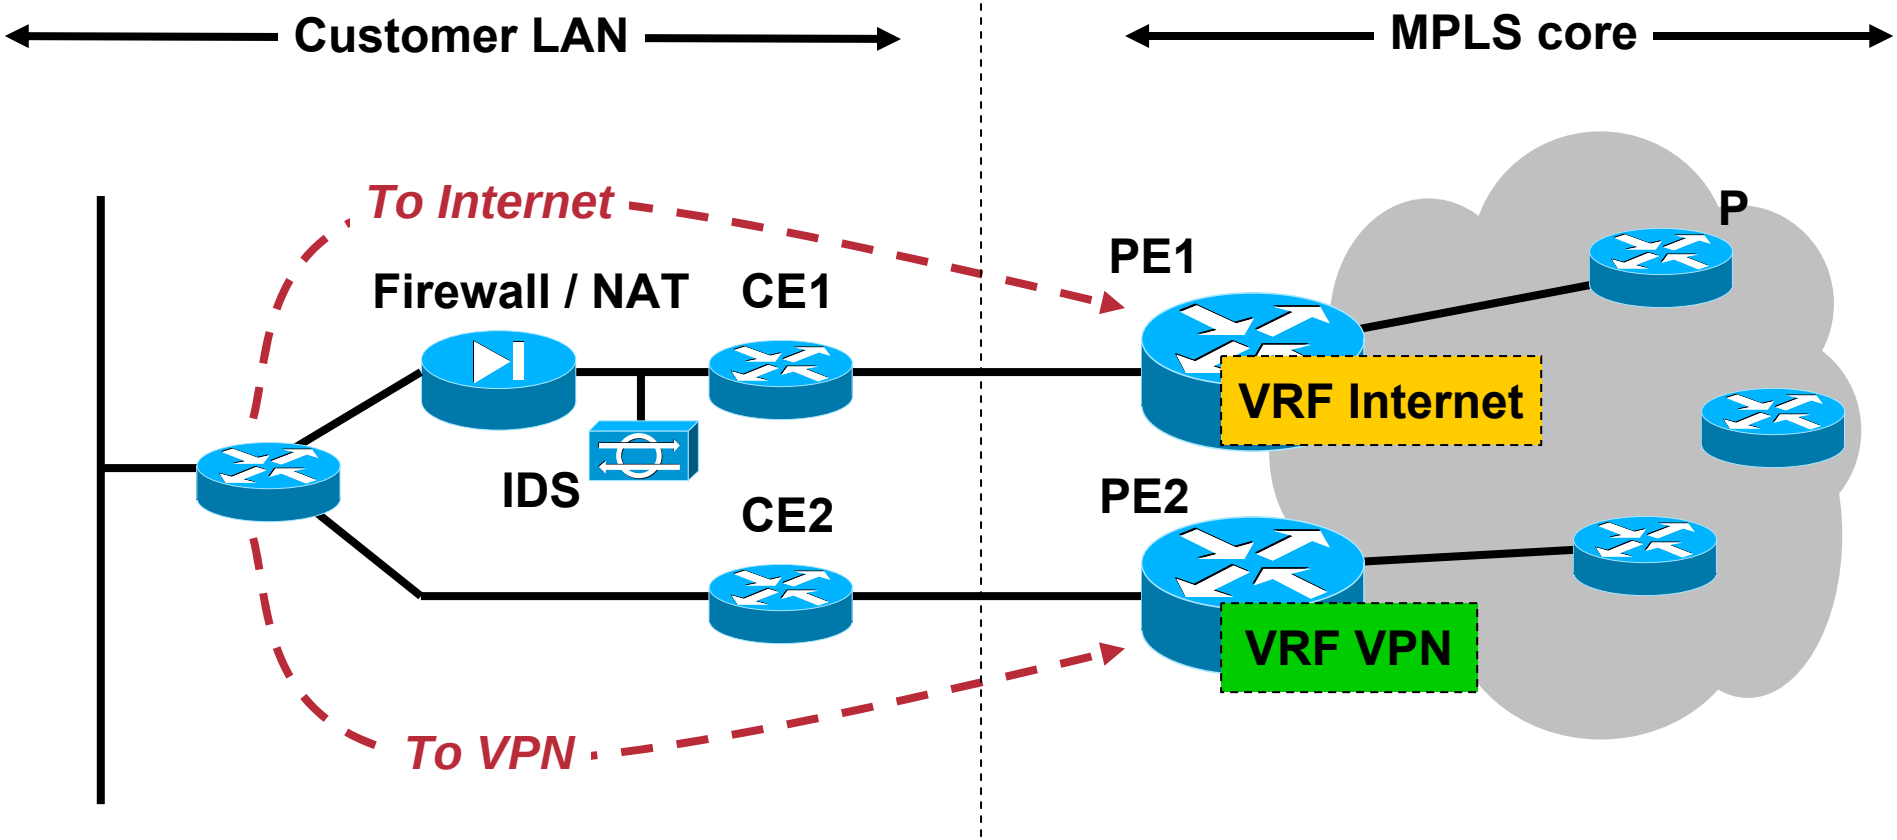

- Separation
- DoS resistance

Agenda

Cisco.com

- **Analysis of MPLS/VPN Security**
  - **Inter-AS VPNs**
  - **Provider Edge DoS possibility**
- **Secure MPLS VPN Design**
  - **Internet Access**
- **Security Recommendations**
- **Summary**

Internet Provisioning on an MPLS Core

Most common VPN user requirement – SP to provide Internet access in addition to VPN connectivity

Two basic possibilities:

1. Internet in global table, either:

   1a) Internet-free MPLS core (using LSPs between PEs)

   1b) Internet routing held by the entire MPLS core (PE and P)

2. Internet in VRF

   Internet carried as a VPN on the core

➢ **Issue – how to design an MPLS core for Internet access such that VPNs remain secure**

MPLS Core Without Internet Connectivity

- MPLS Core – no connection to the Internet; only VPNs connect to the core, **P not reachable, also PE** (except in case seen below)
- Pure MPLS VPN service considered “most secure” – well secured against intrusions and DoS attacks from the outside (core invisible from the outside)
- VPN Spoofing impossible, **VPNs not reachable from the outside**
- **But what about:**

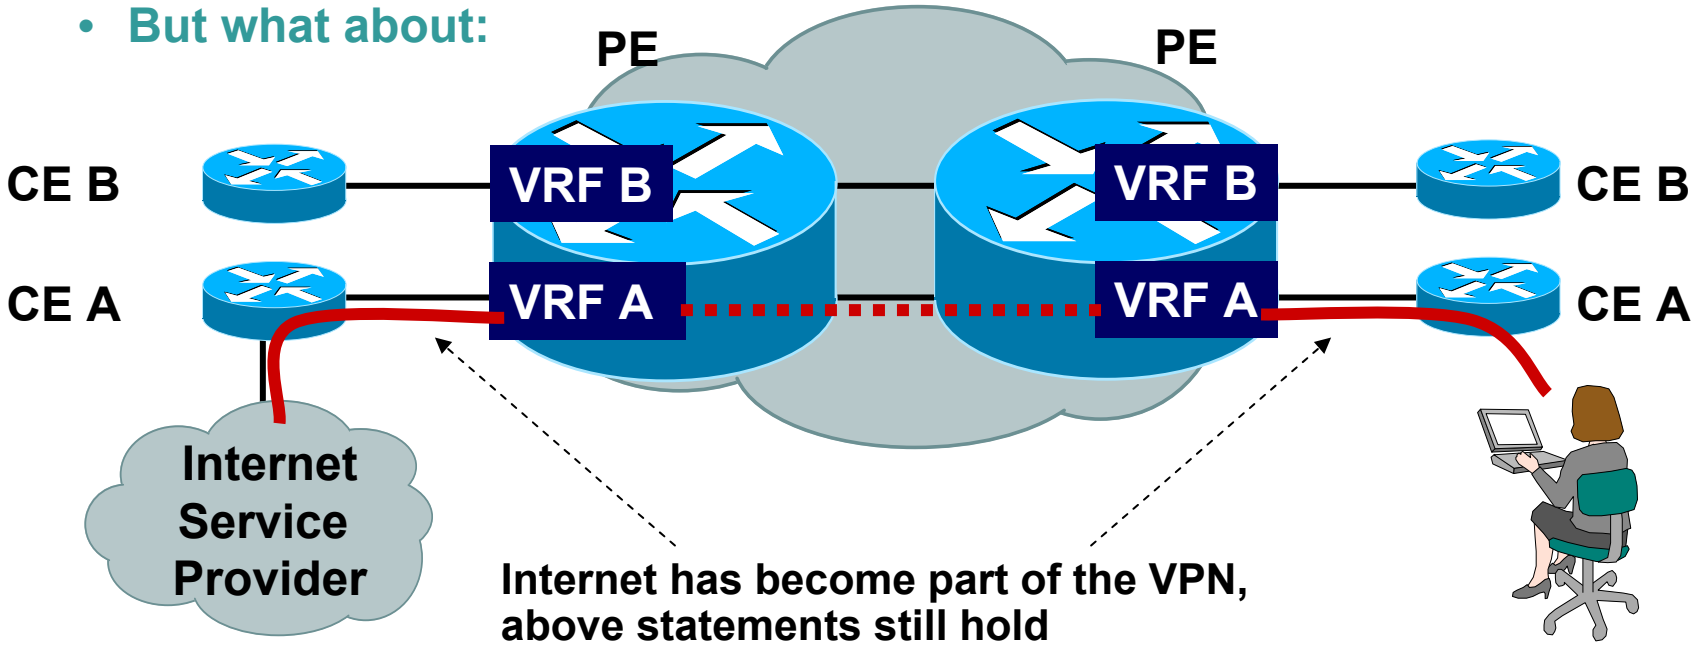

Internet has become part of the VPN, above statements still hold

- DoS attack within such VPN – no immense threat as **access capacity of VPN A can be limited by configuration**

Internet in a VRF

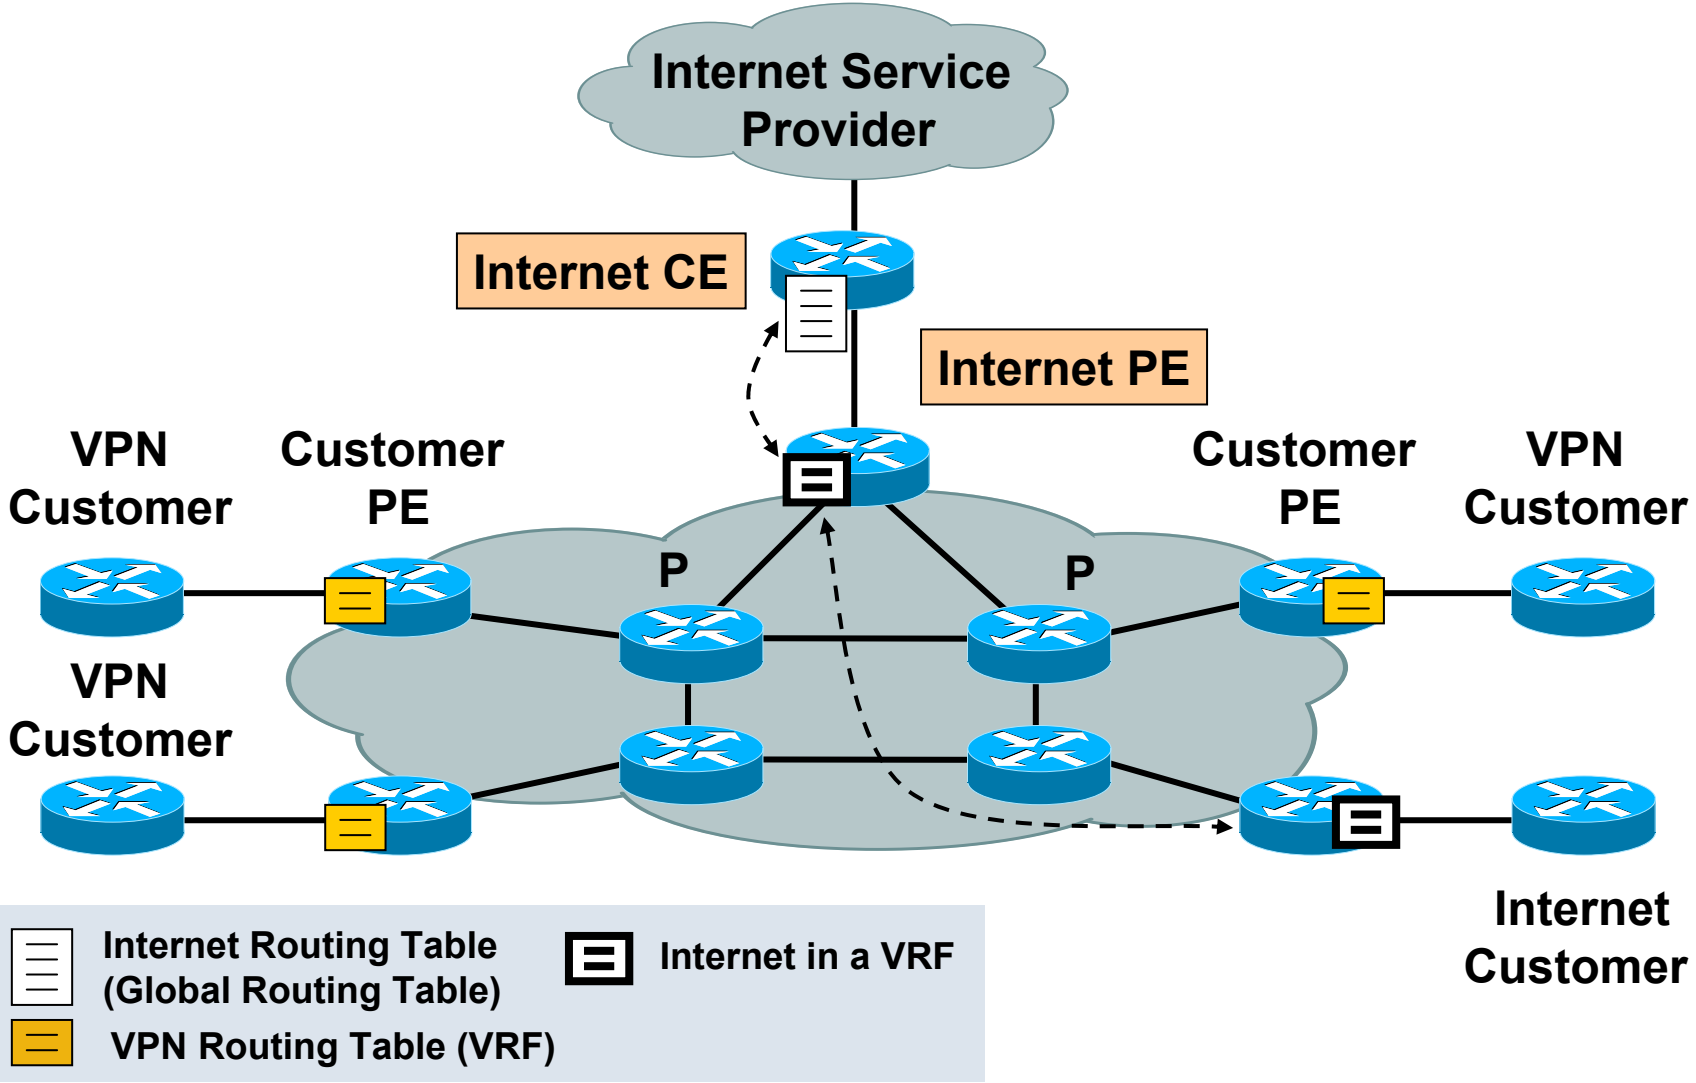

Internet in a VRF – Security Features

- Internet is handled just the same as a VPN, **Customer VPNs not reachable from Internet VPN**
- The core is secure against attacks from the outside as the Internet has no access to the core – **P not reachable**
- Spoofing is impossible between VPNs and Internet in a VPN
- Internet VPN – possibility of DoS of higher magnitude – **PE can be reachable from Internet if not secured properly**

➢ Customer VPNs must not be affected -> provide sufficient capacity in the core OR use QoS to prioritize VPN traffic over Internet traffic

➢ **Scalability Issue** – a prefix held in a VRF requires about three times as much memory as a prefix held in the global table => additional memory required

Internet in the Global Routing Table Using LSPs Between PEs

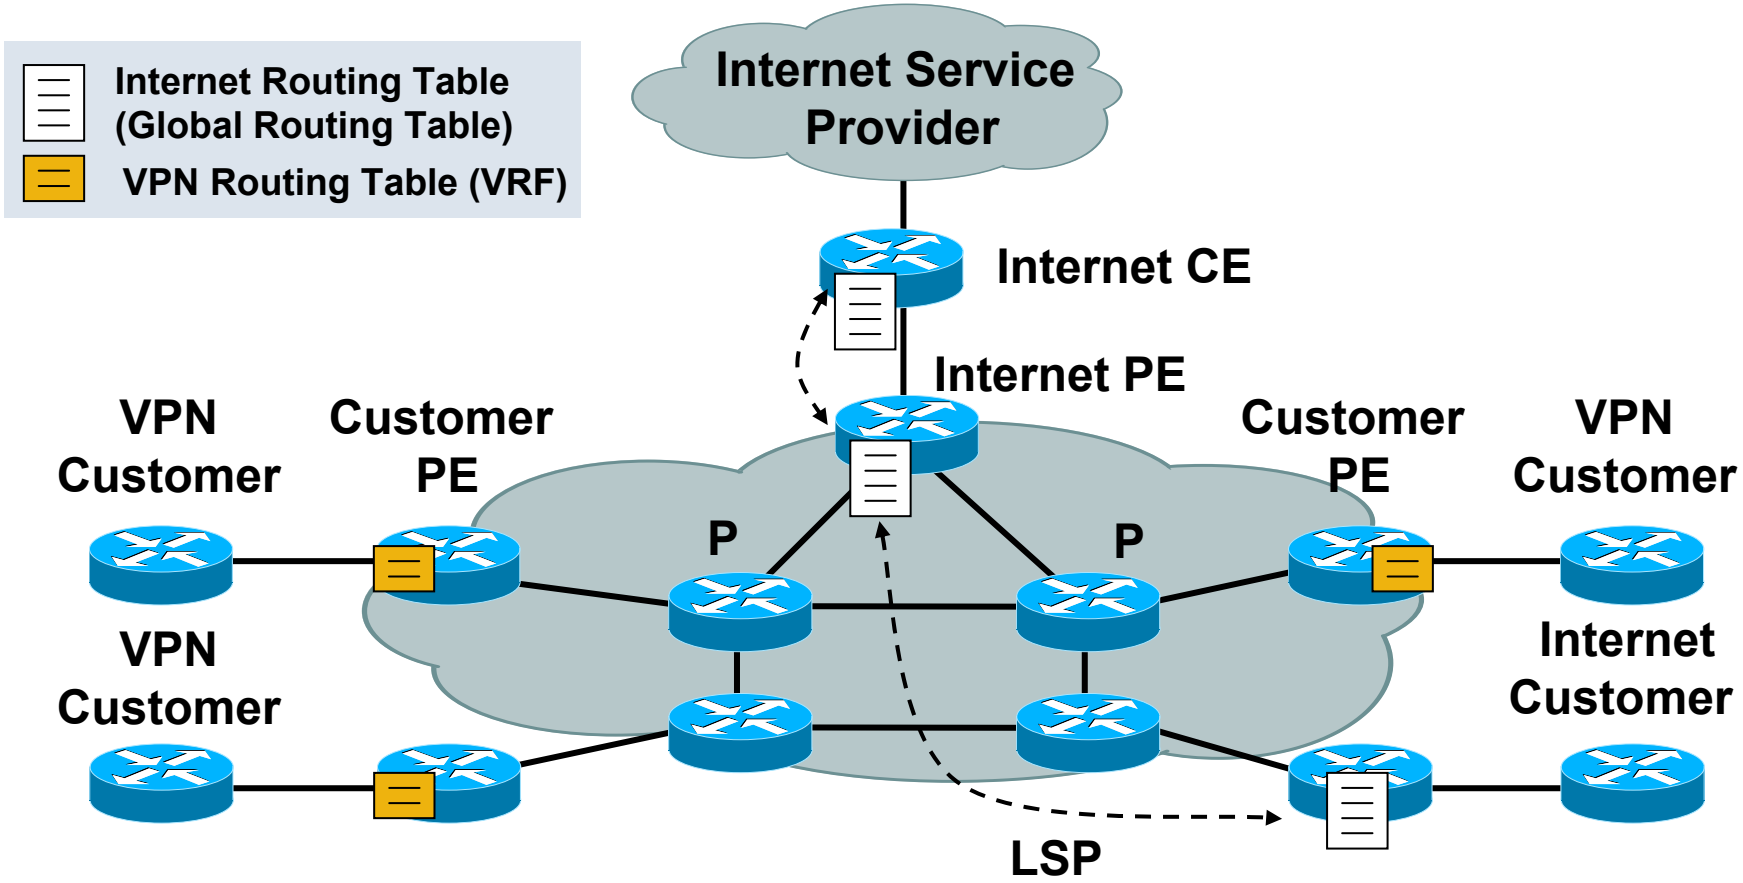

- **Ingress PE - iBGP next hop - Egress PE loopback**

  Next hop to egress usually has label, LSP is used to reach egress PE

  P routers do not need to know Internet routes (nor run BGP, only IGP and LDP)

Internet in the Global Routing Table Using LSPs Between PEs - Recommendations

- In this model PE routers have to carry routes for P routers in their IGP
- Traffic coming from the outside into a PE router's global routing table will have normally a route to the P routers (P reachable unidirectionally)
- LDP and iBGP threatened via attacks against TCP – **usage of MD5 authentication as a solution**

➢ use **Infrastructure ACLs** to prevent packets from outside reach the inside of the core

➢ use **Receive ACLs** and **Control Plane Policing** to protect the control plane of a single platform

➢ **Consider using NSAP addresses in core – IS-IS**

Agenda

Cisco.com

- **Analysis of MPLS/VPN Security**
  - **Inter-AS VPNs**
  - **Provider Edge DoS possibility**
- **Secure MPLS VPN Design**
  - **Internet Access**
- **Security Recommendations**
- **Summary**

Securing the Core: Infrastructure ACLs

In MPLS: PE address belongs to VRF!

- Intended to filter data **destined** for network infrastructure equipment, i.e. what protocols and addresses **can access** critical infrastructure equipment
- On all reachable PE VRF interfaces:

  deny ip any <PE – CE address space>

  permit ip any any

  exception: routing protocol from CE only and all transit traffic
- Idea: Protecting the Core
- DoS: traffic over router theoretically enables DoS, **primary threat – traffic destined for RP**
- **iACLs also to deny** source private address space, reserved addresses, SPs own address space - **antispoofing**

Securing the Core: Infrastructure ACLs

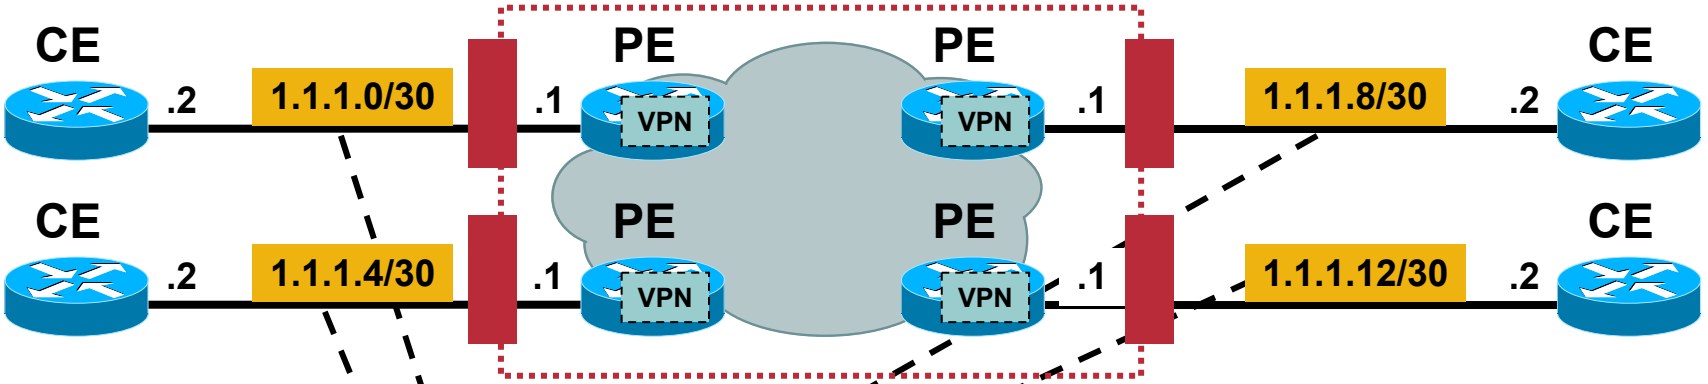

- **Example:**

  deny ip any 1.1.1.0 0.0.0.255

  permit ip any any

This Is VPN Address Space, Not Core!

- **Caution:** This also blocks packets to the CE's!

  Alternatives: List all PE i/f in ACL, or use secondary i/f on CE

Securing the Core: PE-CE routing protocol security

In order of security preference:

1. **Static: If no dynamic routing required**
   (no security implications – no fabricated routing updates, less CPU impact, possible sniffing not revealing routes due to no updates**)**
2. **BGP: For redundancy and dynamic updates**
   (many security features – prefix filtering, route dampening, one BGP process, multiple address-families (per customer/VRF), redistribution at PE not necessary into iBGP**)**
3. **IGPs: If BGP not supported**
   (limited security features – PE peering address known, no 'neighbor' definition, use iACLs**)**

Routing Security: Neighbor Authentication and BGP TTL

Cisco.com

- **Use static routing between CE and PE where possible**

  no errant routes announced, no routing data crossing the ‘wire’, no CPU impact

- Routers authenticate each time a routing update is exchange between them – reliable information received from a trusted source

  Verification through MD5 hash

- Supported: BGP, ISIS, OSPF, EIGRP, RIPv2, LDP
- **MD5 for LDP –** label spoofing protection, enable also on MP-iBGP

Control of Routes from a BGP Peer

- **Injection of too many routes – possible attack at routing table stability, CPU and memory:**

  Potential DoS attack, leading e.g. to CEF disabling or reload

- **Control with "maximum prefix" command**

  After exceeding the number – BGP peering disabled, neighbor down

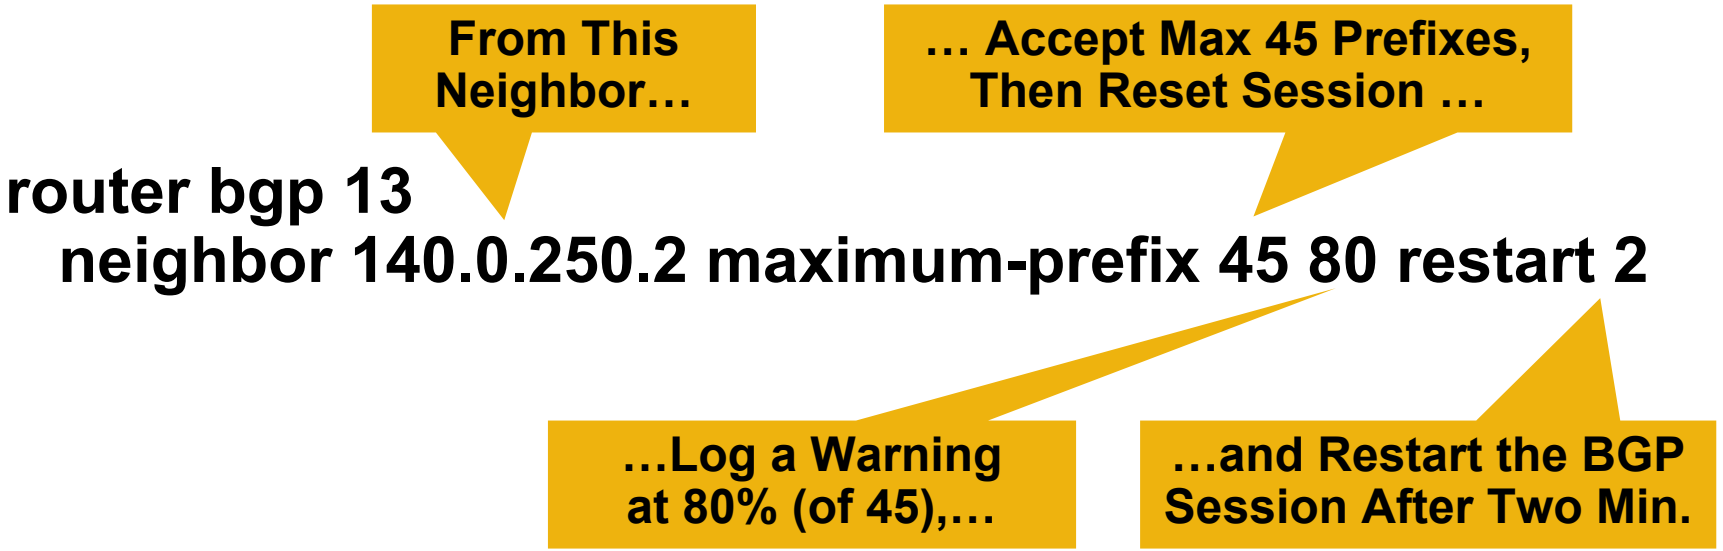

Control of Routes from a BGP Peer: Logging

Cisco.com

6d22h: %BGP-4-MAXPFX: No. of prefix received from 140.0.250.2 (afi 2) reaches 37, max 45

6d22h: %BGP-3-MAXPFXEXCEED: No. of prefix received from 140.0.250.2 (afi 2): 46 exceed limit 45

6d22h: %BGP-5-ADJCHANGE: neighbor 140.0.250.2 vpn vrf VPN_20499 Down BGP Notification sent

6d22h: %BGP-3-NOTIFICATION: sent to neighbor 140.0.250.2 3/1 (update malformed) 0 bytes FFFF FFFF FF

VRF Maximum Prefix Number

- **Injection of too many routes:**
  - Potential memory overflow
  - Potential DoS attack
- **For a VRF: Specify the maximum number of routes allowed**

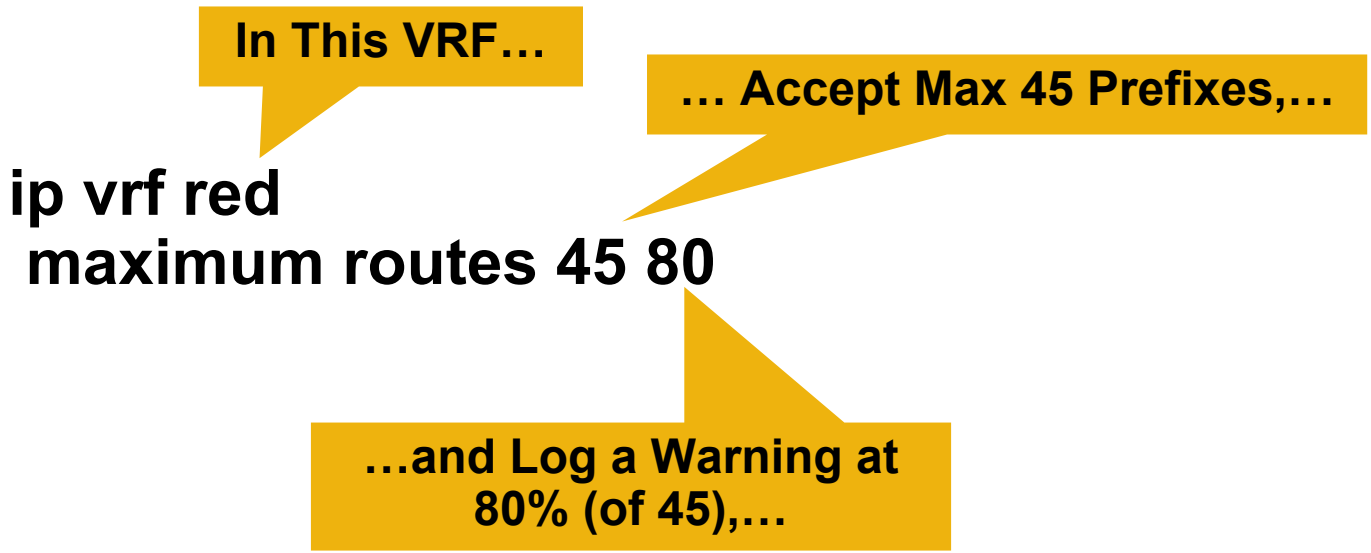

PE-Specific Router Security

- **PE Control Plane hardening – Receive traffic**
  - L3 routing environment (authentication, max number of prefixes…)
  - Infrastructure ACLs
  - Protection ACLs (anti-spoofing, etc.)
- **PE Data Plane Hardening**
  - Use **uRPF Strict mode** on each interface of the PE routers' CE-facing interfaces and on the CE routers' PE-facing interfaces

Attacking a CE from MPLS (other VPN)

- **Is the CE reachable from the MPLS side?**

  -> only if this is an Internet CE, otherwise not!
  (CE-PE addressing is part of VPN!)

- **For Internet CEs:**

  Same security rules apply as for any other access router.

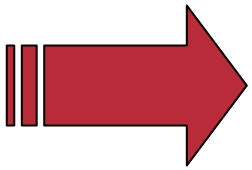

MPLS hides VPN-CEs: Secure!
Internet CEs: Same as in other networks

Agenda

Cisco.com

- **Analysis of MPLS/VPN Security**
  - **Inter-AS VPNs**
  - **Provider Edge DoS possibility**
- **Secure MPLS VPN Design**
  - **Internet Access**
- **Security Recommendations**
- **Summary**

Securing the MPLS Core: Wrap-Up

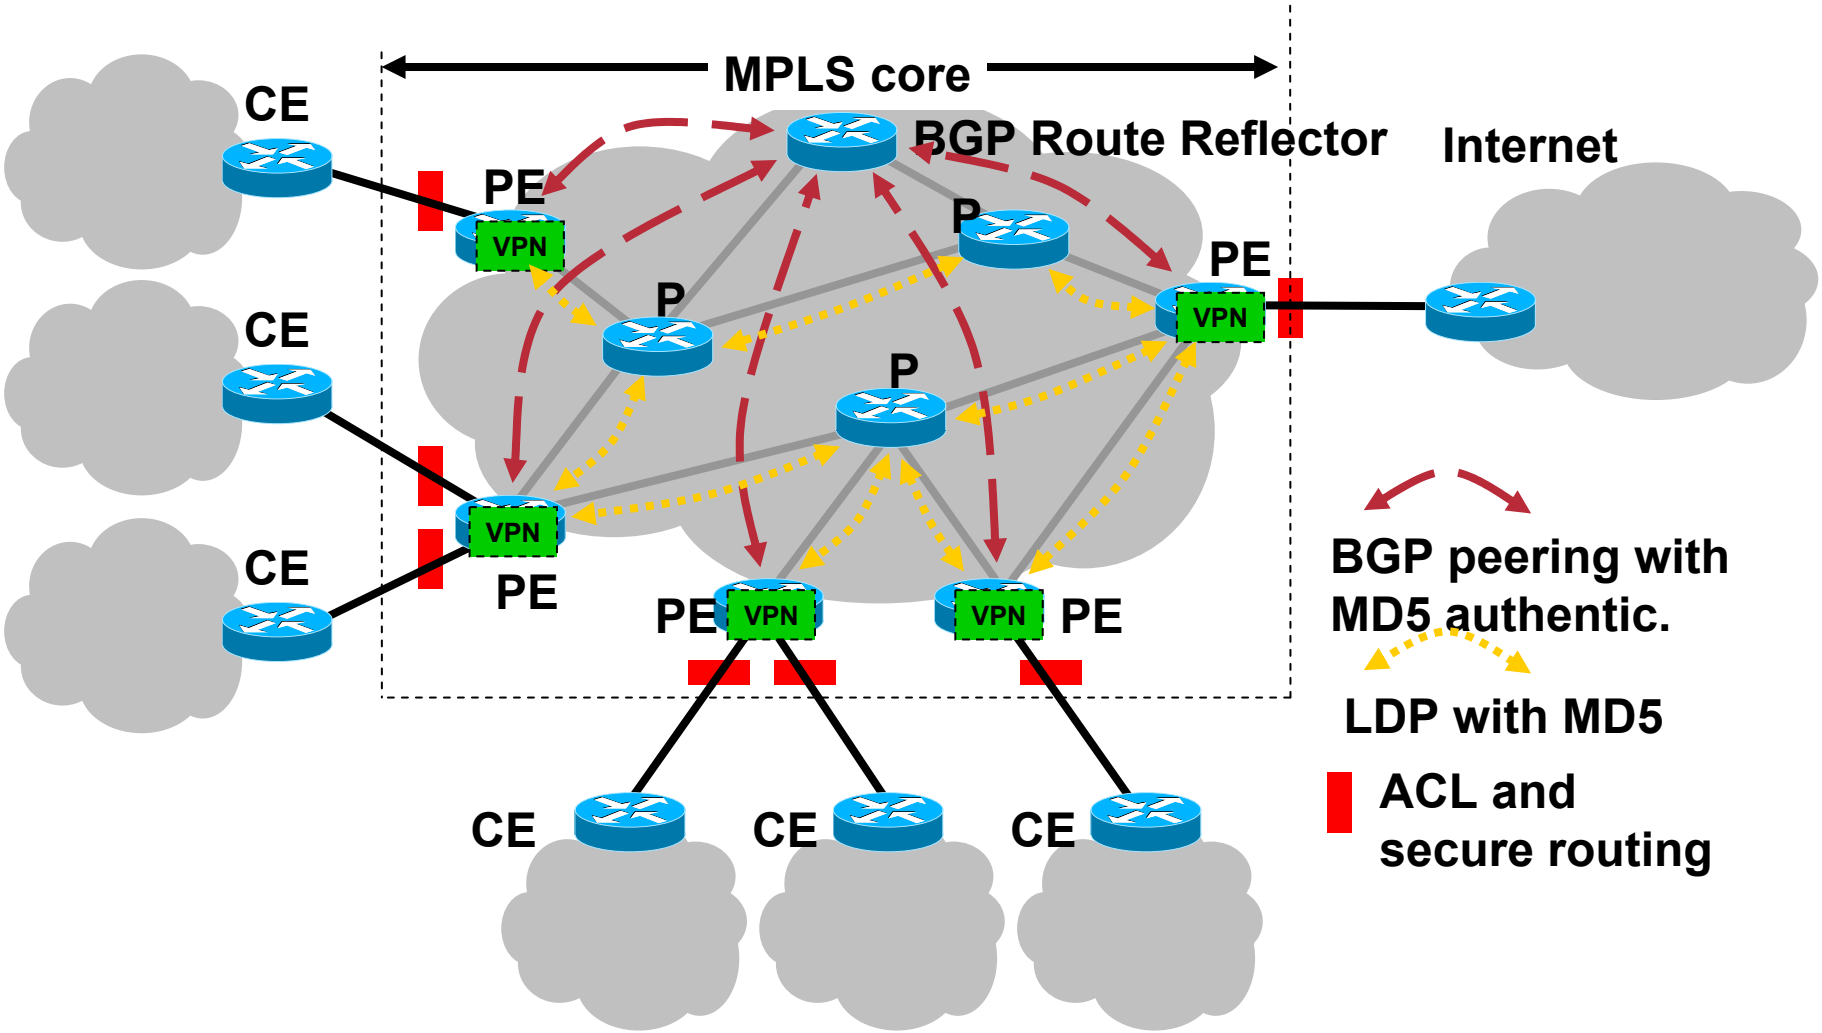

MPLS Security Overview

1. **Don't let packets into (!) the core**

   → No way to attack core, except through routing, thus:

   → **Still "open": routing protocol**

2. **Secure the routing protocol**

   Neighbor authentication, maximum routes, dampening, …

   → **Only attack vector: Transit traffic**

3. **Design for transit traffic**

   QoS to give VPN priority over Internet

   Choose correct router for bandwidth

   Separate PEs where necessary

   → **Now only insider attacks possible**

4. **Operate Securely**

   → **Avoid insider attacks**