

Understanding and mitigating the security issues for Ethernet-based networks and services

Enno Rey, erey@ernw.de



Who I am



- “Old-school networker“ with vast carrier background.
- Started working on Layer 2–4 in the early 90s.
- With special focus on security since 1997.
- (Co-) Author of several books, articles and whitepapers and regular speaker on international conferences (incl. Black Hat, Hack In The Box, FutureNet).
- Founder (2001) and CTO of a highly specialized information security consultancy [www.ernw.de] with 12 employees, based in Heidelberg/Germany and Lisbon/PT.



Agenda

- **Definition & Restriction of Term “Carrier Ethernet”**
- **Traditional Attacks in Ethernet Networks**
- **Security Problems in Carrier Ethernet Space**
- **Mitigating Controls**
- **Who is Responsible?**
- **Outlook on the Future**
- **Conclusions**



Definition & Restriction of Term “Carrier Ethernet”

- ***Carrier Ethernet*** basically means that ethernet frames are transported across (at least) one carrier's backbone.
- So ethernet is not (only) used as an *access medium* here, but offered as a *service*.
- This service may be implemented on the basis of different technologies (see below).
- The resulting connection (mostly) shows typical ethernet properties & behaviour.



Properties 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*



Technologies

- **Metro Ethernet**
- **(Mainly) Cisco's Ethernet-over-MPLS (EoMPLS)**
- **VPLS**
- **L2TPv3**

- **From carriers' product space and from customer perspective all this is "Carrier Ethernet".**



No Layer 3 Device as CE

The following discussion assumes that consistent continuous “layer 2 domains“ (*collision domains* in classical ethernet lingo) result from the implementation of *Carrier Ethernet* between customer sites.

=> any layer 3 device (router) as customer edge (CE) “breaks“ this assumption (and “breaks“ *Carrier Ethernet*).



Full vs. Partial Transparency

- Depending on the (carrier's) service/product, potentially the devices used and the configuration of PE and CE the connection may or may not provide full transparency.
- “Full transparency” means, that *all* BPDUs (including e.g. STP, DTP, VTP, GVRP, LACP, 802.1x packets and the like) and *all* Layer2 Headers (incl. VLAN tags, CoS) are transparently transported from one site to another/others across the cloud.
- In contrast “partial transparency” means that some of the BPDUs or header information is filtered/discarded when entering the cloud.



Example of *Carrier Ethernet* in the sense of this talk [copied from *MEF*]

- **Transparent LAN Service (TLS)**
provides

- Intra-company Connectivity
- Full transparency of control protocols (BPDUs)

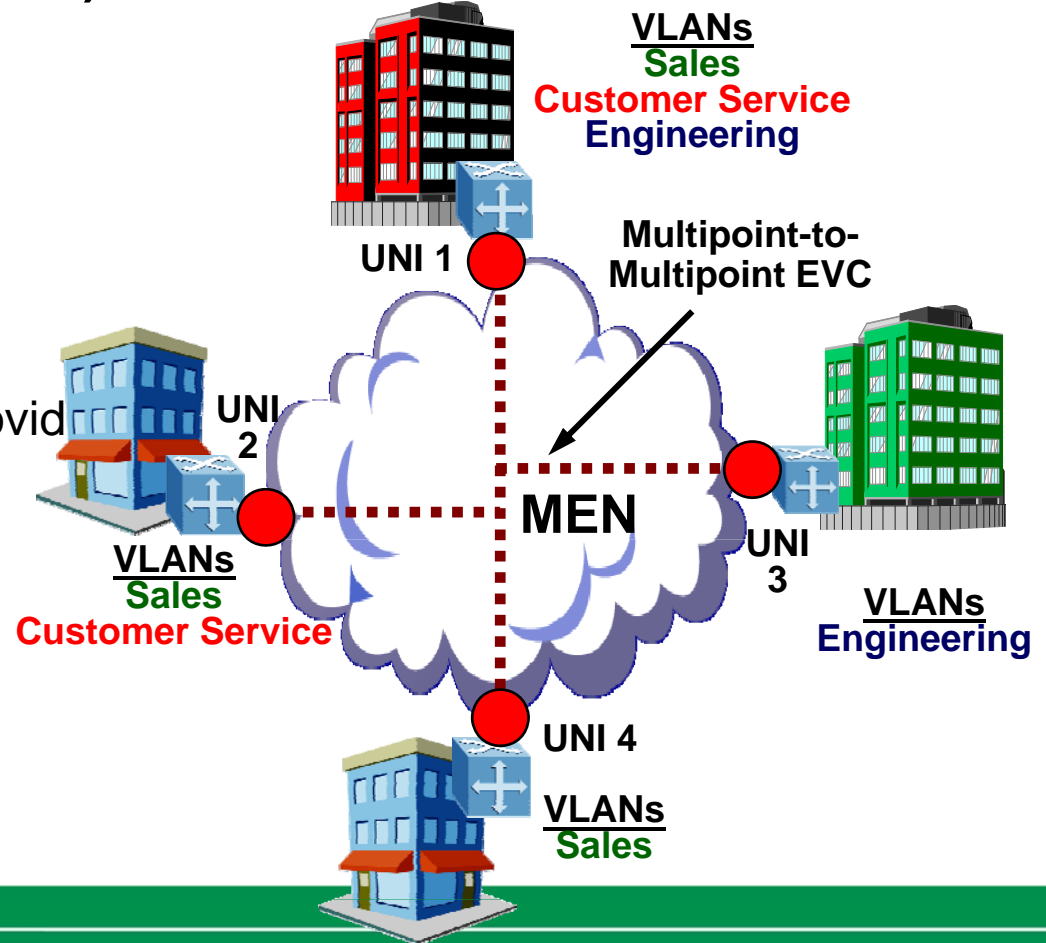
- **New VLANs added**

- without coordination with provider

!!!

TLS makes the MEN
look like a LAN

Transparent LAN Service



Scenarios and their Security Impact

[from an ERNW customer project]



Technology	CE = Router	CE = Switch, Router behind CE managed by \$CUST	CE = Switch, no Router behind CE
Metro	does not apply	configure router as "secure perimeter device"/NACS element (like CE_template)	Alarm!
Point-to-Point Ethernet, like EoMPLS	if working as L3 device, same as MPLS L3 VPN	configure router as "secure perimeter device"/NACS element (like CE_template)	try contractual controls, take care of L2 sec
Multi-Point Ethernet (e.g. VPLS)	if working as L3 device, same as MPLS L3 VPN	configure router as "secure perimeter device"/NACS element (like CE_template)	take even more care of L2 sec. Here all the stuff from checklist comes into play.



Security Goals

- **RFC 3871, sect. 1.4:**

A secure network is one in which:

- o The network keeps passing legitimate customer traffic (availability).
- o Traffic goes where it is supposed to go, and only where it is supposed to go (availability, confidentiality).
- o The network elements remain manageable (availability).
- o Only authorized users can manage network elements (authorization).
- o There is a record of all security related events (accountability).
- o The network operator has the necessary tools to detect and respond to illegitimate traffic.



- **Eavesdropping & Traffic Redirection**
=> Breach of Confidentiality/Integrity
- **Denial-of-Service (DoS)**
=> Loss of availability
- **Eavesdropping & DoS “on a large scale“**
- **Attacks against infrastructure/devices/physical sec**



Eavesdropping in Ethernet Networks



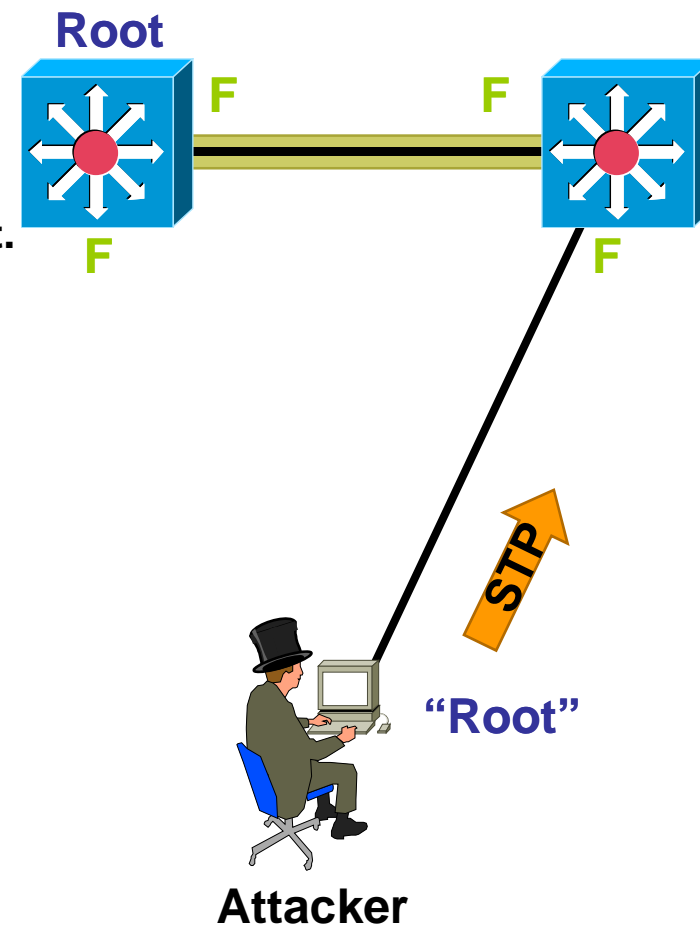
- ***Shared Medium***
- **No built-in encryption**
- **Well known (malicious) redirection options by**
 - ARP spoofing/poisoning
 - Attacks on multicast based routing protocols



Example of Denial-of-Service in Ethernet networks: Attacks on Spanning Tree (STP)



- Attacker sends BPDUs to enforce recalculation
- Impact mostly DoS
- If *Rapid Spanning Tree* is used much less impact.



Eavesdropping & DoS “on a large scale”



- **Term here used to designate all the methods that enable an attacker to perform eavesdropping and/or DoS in large parts of a network (e.g. in/on several VLANs simultaneously).**
- **These attacks mostly work on proprietary protocols in *Cisco* space (DTP, VTP).**



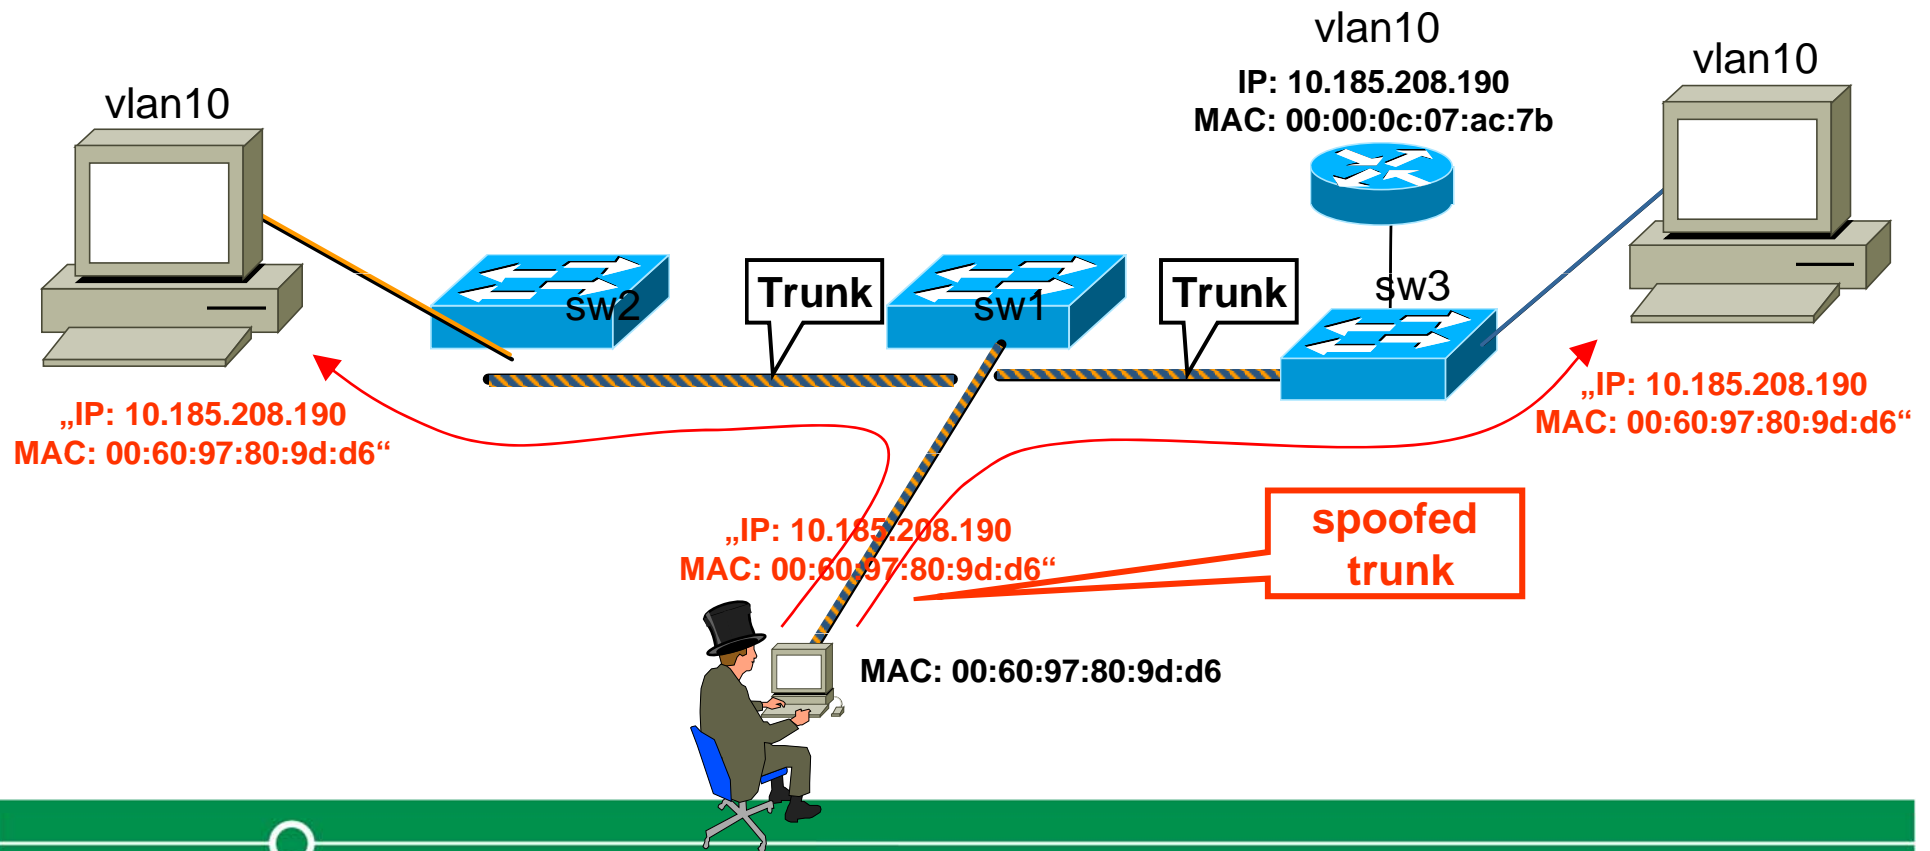
Attack against DTP

```
erey@mobile:~  
----- yersinia 0.5.3 by Slay & tomac - DTP mode ----- [17:43:54]  
Neighbor-ID  Status      Domain      Iface  Last seen  
000BFDB648AF  03 (ACCESS/DESIRABLE)  ernw    eth0    08 May 17:40:46  
048D226BAE78  03 (ACCESS/DESIRABLE)  ernw    eth0    08 May 17:40:55  
000BFDB648AF  83 (TRUNK/DESIRABLE)  ernw    eth0    08 May 17:43:26  
048D226BAE78  83 (TRUNK/DESIRABLE)  ernw    eth0    08 May 17:43:29  
----- Attack Panel -----  
No  DoS  Description  
0   0    sending DTP packet  
1   1    enabling trunking  
-----  
Total Packets: 0  
Those strange: 0  
----- DTP Fields -----  
Source MAC 04:00:00:00:00:00  
Version 01  Domain  
Status 03  Type A5  Neighbor-ID 048D226BAE78  
-----  
Select attack to launch ('q' to quit)  
-----  
Spoofing [X]
```

Attacker will now be able to participate in all VLANs.



ARP spoofing across VLAN boundaries



Attacks against VTP

```
erey@mobile:~  
----- yersinia 0.5.3 by Slay & tomac - VTP mode ----- [17:44:44]  
Code      Domain      MD5      Iface Last seen  
04 (JOIN)  ernw        eth0     08 May 17:40:56  
04 (JOIN)  ernw        eth0     08 May 17:40:56  
04 (JOIN)  ernw        eth0     08 May 17:41:22  
01 (SUMMARY) e  
04 (JOIN)  e  
04 (JOIN)  e  
----- Attack Panel -----  
No  DoS  Description      8 May 17:41:01  
0   0    sending VTP packet      8 May 17:41:26  
1   X    deleting all VTP vlans  8 May 17:41:26  
2   X    deleting one vlan  
3   0    adding one vlan  
-----  
Total Packets  
Those strange  
VTP Fields  
Source MAC 00: _____ Select attack to launch ('q' to quit) _____ 00  
Version 00 Code 00 Domain  
MD5 00000000000000000000000000000000 Updater 000.000.000.000  
Revision 00000000 Timestamp Start value 00000000  
Followers 000 Sequence 000 VLAN
```

Do I really have to discuss the impact of this?

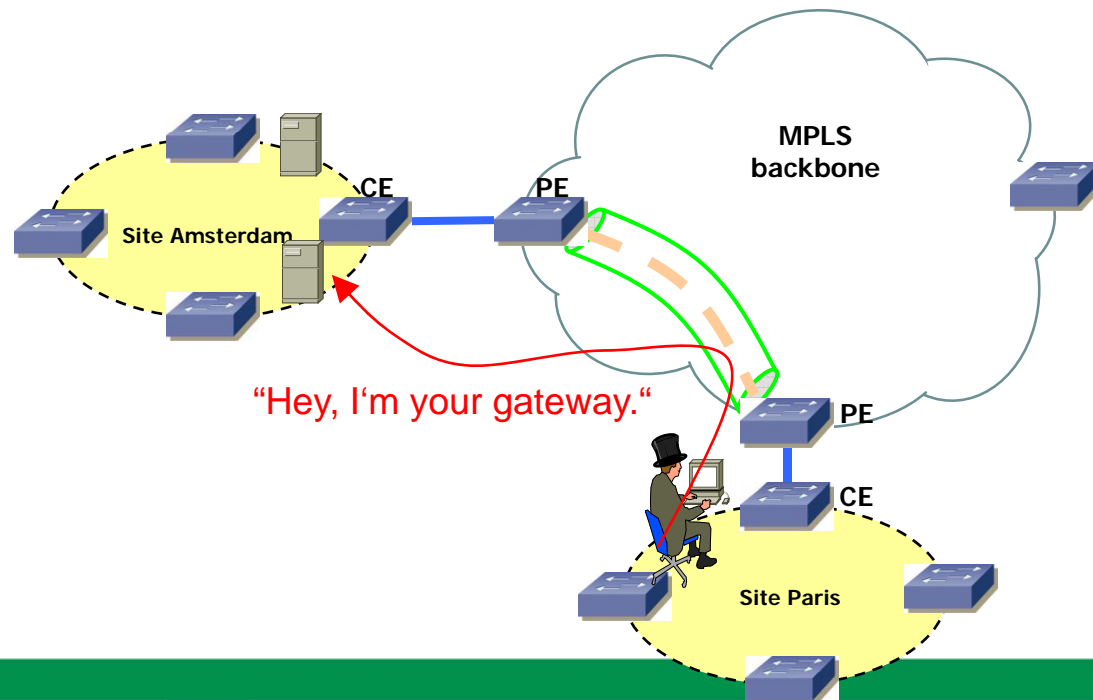


- **Nothing new so far (and still 20 min of my talk left ;-)) ...**
- **So what about security aspects of *Carrier Ethernet*?**
- **Carrier Ethernet might provide larger surface for traditional attacks.**
- **New problems from merger of LAN & WAN might arise.**
- **New problems due to technologies and associated device capabilities/features.**



Traditional Ethernet Attacks “over the cloud”

- Depend highly on the level of transparency a “VPLS cloud” provides.
- Given full transparency (as in *Juniper*-based testbed we used)...
- ... you can perform any traditional 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...




New problems from merger of LAN & WAN



- **As the “natural boundary“ that WAN transport constitutes for most layer 2 infrastructure protocols disappears, interesting new scenarios – with previously unknown security implications – emerge.**
- **Let’s have a look at some examples (and always remember: I assume consistent collision domains here).**



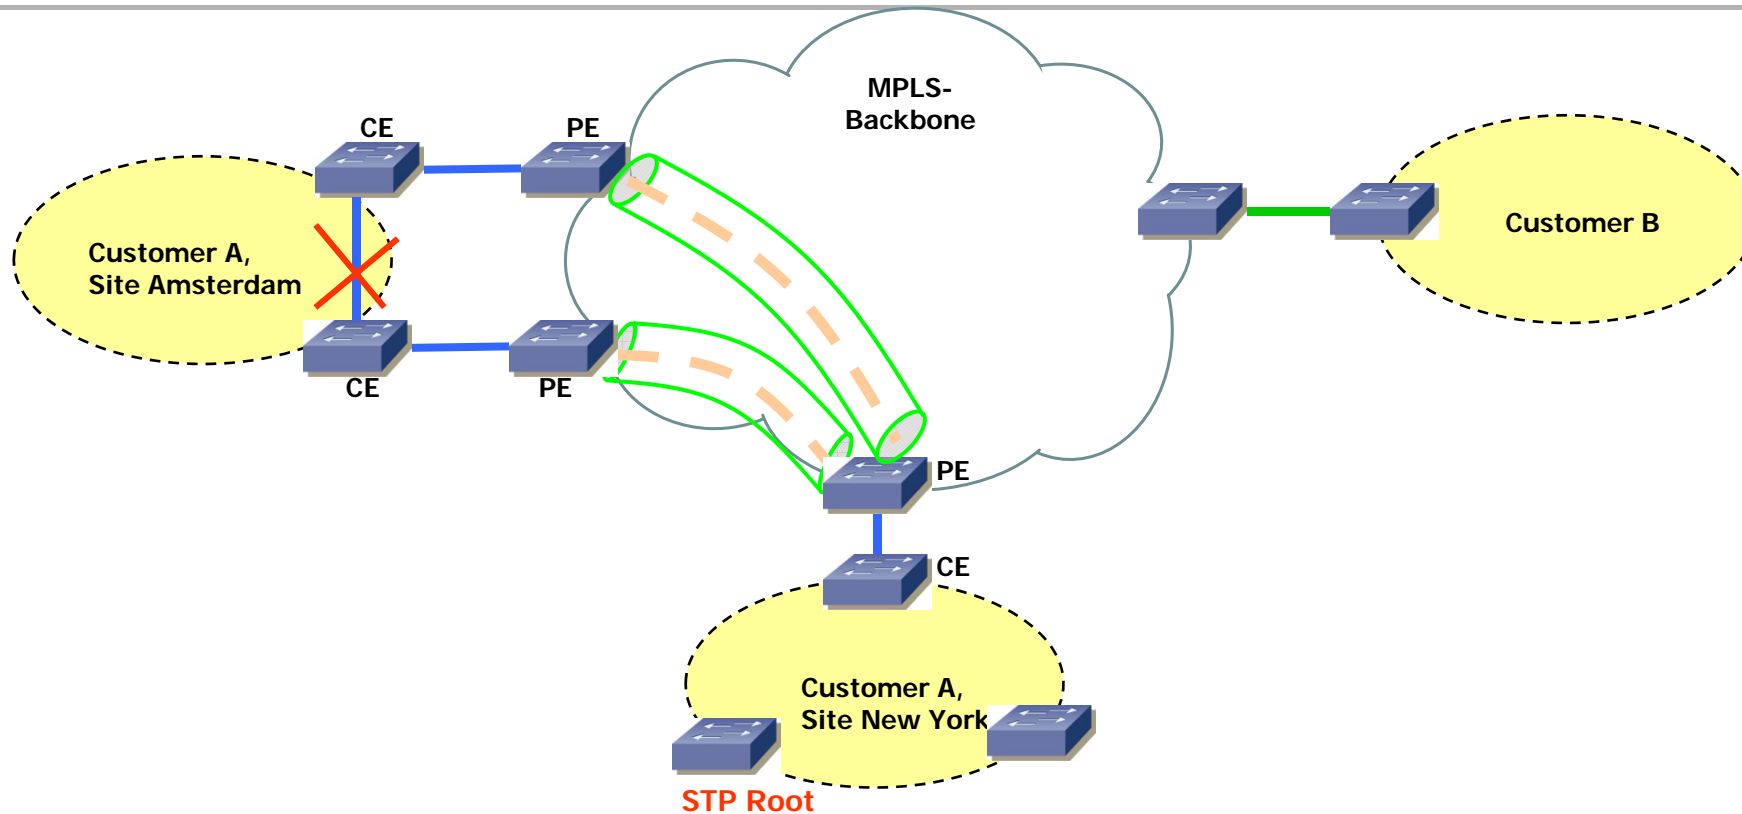
STP Root Election



Btw: RFC 4762 sect. 4.4 explicitly mentions this case.



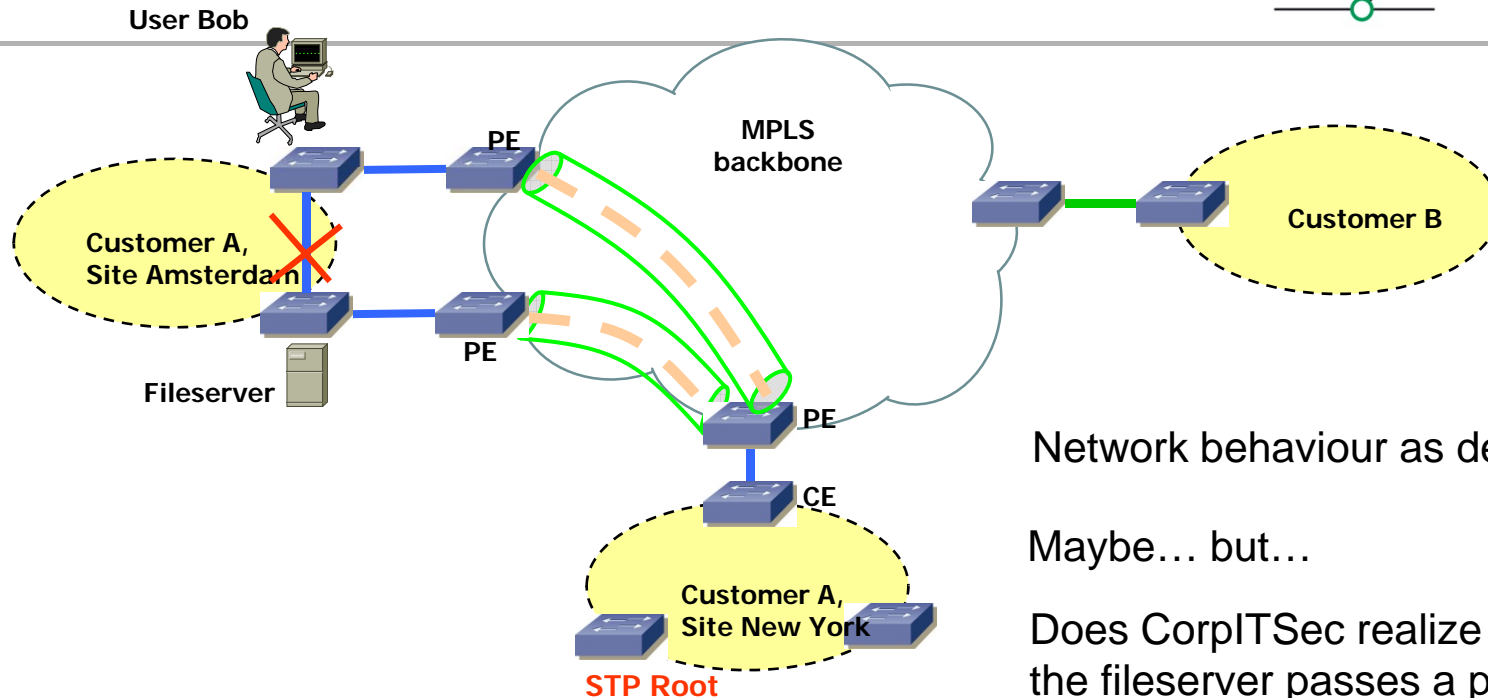
Some customers may want redundant connections...



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?

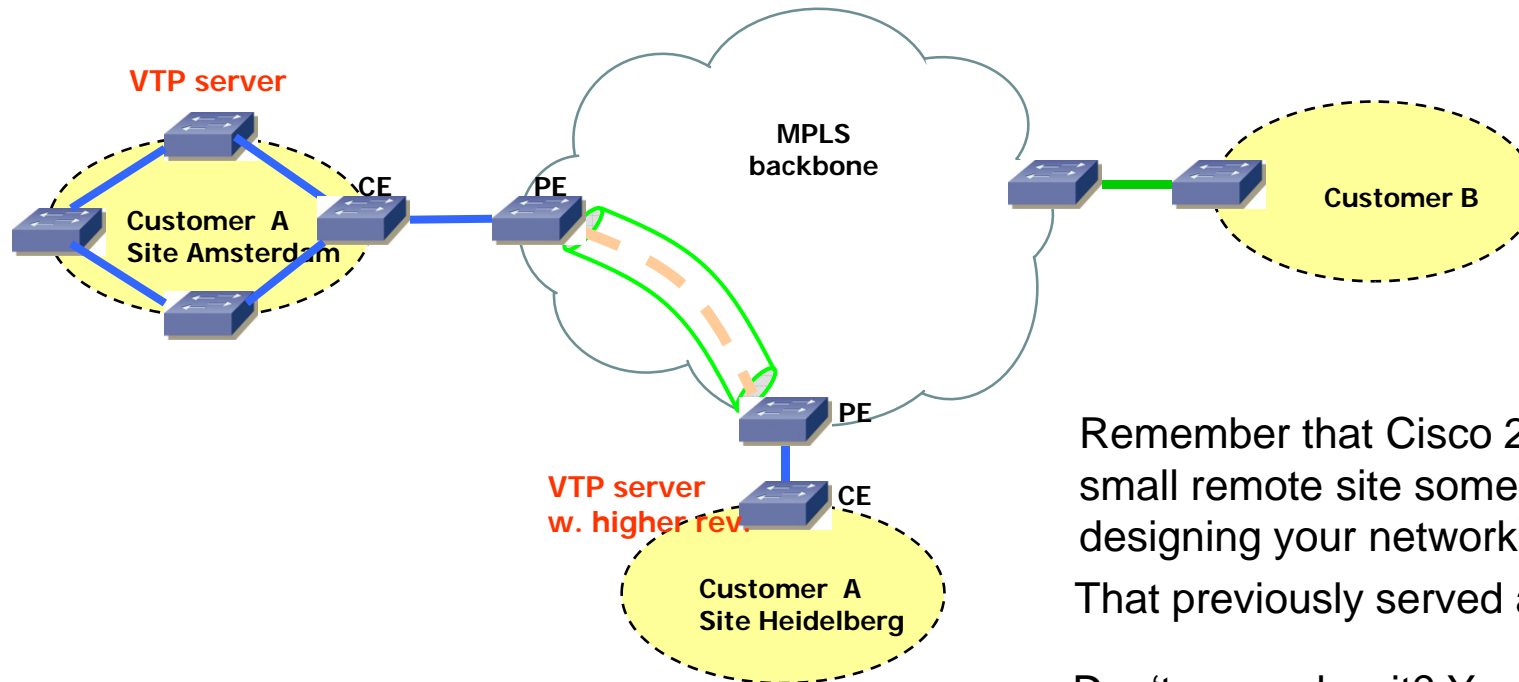
In another country...

where *Carnivore/DCS 1000* stuff 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)?

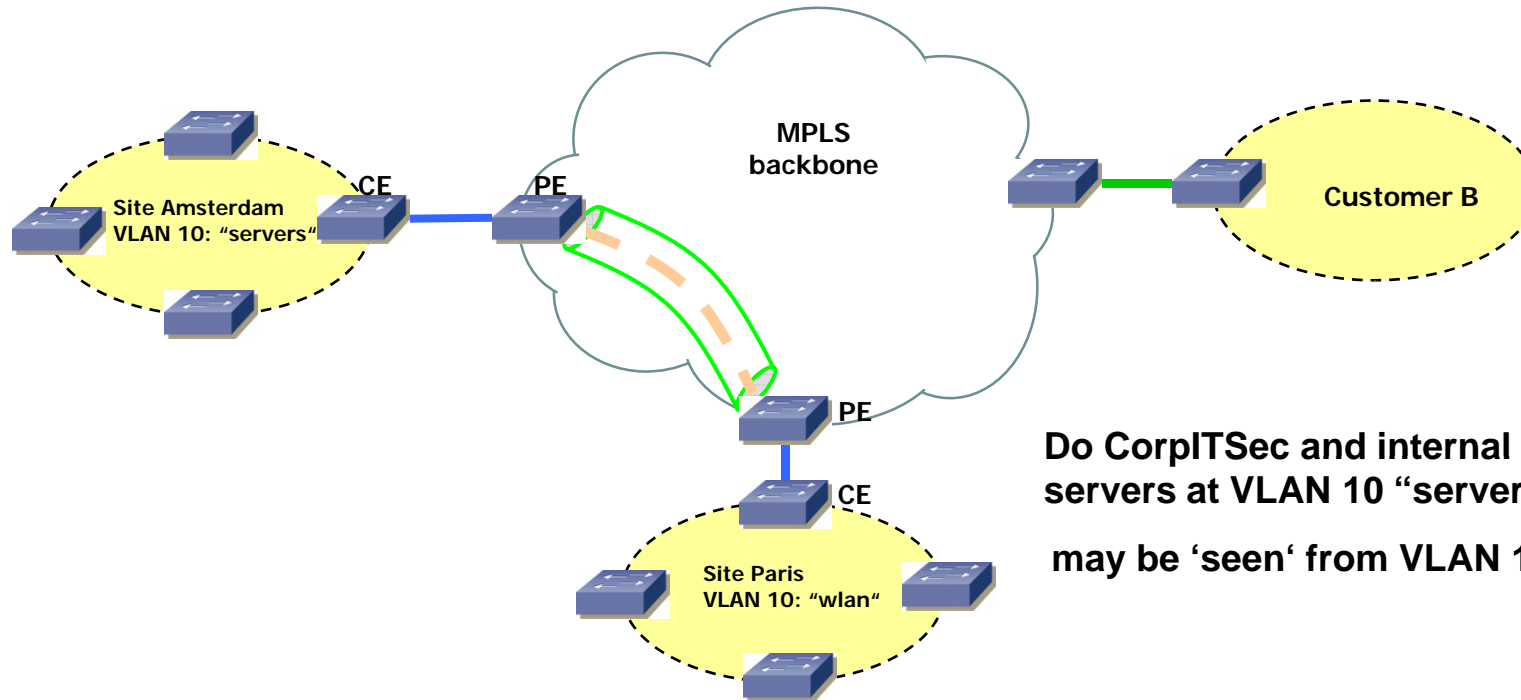
That previously served as *VTP server*...

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).



New problems due to technologies and device features

- **This stuff adds new features (and subsequent [code] complexity) to devices.**
- **Some vendors may not really understand the implications (or implement immature code in first releases).**
- **Again, an example might be helpful...**



Attacks against VPLS-performing devices



- Depend highly on the functions they perform.
- The (quite often used) concept “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].

```
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
```

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

Note:

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



Just a note on PBT here

- Adds even more complexity to devices, especially to (header) parsers.
=> We're anxiously waiting to use our own Layer 2 fuzzing toolkit in a PBT testbed ;-)
- No major security impact (in neither direction) from our perspective.



- The “basic building blocks“ (or at least some of them) work here...
- Isolation / Segmentation
- Restriction / Filtering
- Access Control / Authentication
- Encryption
- Hardening of Infrastructure Services
- Secure Management
- Logging / Monitoring



Isolation / Segmentation

- Obviously network segmentation (with VLANs/IP subnets) is a good way to address network security problems.
- Would break our definition of *Carrier Ethernet* though.
=> not covered here.
- However “isolation“ of potentially harmful packets/frames (by using filtering techniques) might be a good idea (see below).



Encryption

- **Best (and mostly) only method if you don't trust the transport path.**
- **Manageability is key.**
- **Think about key mgmt processes *before* deploying.**
- **Performance might be a factor, too.**
=> Undertake risk analysis if highest key length needed.



Some notes on MacSec (802.1ae)



- I recently went through this (not too readable 142-page standard...
- Honestly: who needs this?
- I strongly doubt we will ever see it (deployed).
- Who needs crypto on L2? We have working crypto on L3 (IPsec) and L4 (SSL).
- As for Carrier Ethernet Networks... keep in mind: This will (more precisely: would ;-)) add another layer of complexity to the “Carrier Ethernet Devices“...



- **Control who is able to access / take part in the network.**
- **Preventative mechanism
=> as such a good thing**
- **But might be administrative nightmare**

- **MAC address based (*Port Security*)**
- **Certificate based**
- **802.1x**
- **NAC/NAP (?)**



- **NAC/NAP solutions address mainly two problems:**
 - Missing patches
 - Outdated AV signatures
- **In a perfect world both problems would not exist due to**
 - Good patch management
 - Good signature distribution
- **In an even better world both processes would not be necessary as a result of**
 - Software that doesn't need to be patched every some days
 - Users behaving sensibly



Some Remarks on NAC/NAP here

So...

- **NAC/NAP address deficiencies in processes ... which address deficiencies in some “info processing entities“.**
- **Kind of *the medicine for the symptom of the medicine for the symptom.***
- **Think about that when thinking about spending money on this stuff...**
- **\$\$\$ spent for NAC/NAP might be spent much better elsewhere.**
- **See also: http://www.ernw.de/content/e7/e181/e566/download569/bh07-europe_nacattack_03_ger.pdf**



Logging / Monitoring

- You do this extensively for *compliance reasons*, don't you? ;-)
- If you can't really prevent/control stuff you should at least be able to detect or track.



- Ask yourselves: would you notice “strange layer 2 behaviour“ going on on your (or the customer's) links.



Aside from that general procedures



Examples of dedicated technical controls for *Carrier Ethernet Services*:

- **Disable reception/sending of STP BPDUs**
- **Restrict *allowed VLANs* on PE-CE links**
- ***Storm control* mechanisms**
- **L2 based IPS**
- **Block device discovery, link aggregation, VLAN mgmt prots.**
- **In short: do all the stuff you do in IXPs... but only if this does not break things expected from customer...**
- **See also “Carrier Ethernet Security Guide“, to be found on www.ernw.de from approx. *2007 Oct 15*.**



Filtering – Potentially interesting stuff

- **Spanning Tree / STP**
- **VTP**
- **DTP**
- **Proprietary xDP (e.g. CDP)**
- **LLDP**
- **GVRP**
- **PAgP**
- **LACP**
- **802.1x**
- **WinVista L2 stuff (LLMNR, LLTD etc.)**



Who is responsible?

- **Nice technical stuff... but who is going to take care of that?**
- **Who has the necessary knowledge?**

- **Carrier?**
- **Customer?**

- **In short: might have to be negotiated on a by-case basis.**
- **Even worse: might differ from implementation to implementation (vendors/devices/configuration).**



Carrier Perspective

- Typical *mere conduit* stance (“we only transport”).
- May not have too deep L2 knowledge so far...



Customer Perspective

- **Why don't they protect us?**
["I mean: it's their service, isn't it?"]
- **Why does the same product (*Carrier Ethernet*) behave differently in (seemingly) equal implementation cases?**




What Carriers should do

- **Understand implications**
- **Offer advice**
- **Integrate filtering policies into products**
- **Flexible configuration?**



Outlook on the Future

- **Technical product space still evolving/changing (remember T-MPLS vs. PBT debate at *FutureNet?*).**
- **New layer 2 protocols from *Windows Vista* space.**



New Protocols (mainly in WinVista)

- **Whole new bunch of protocols emerged recently.**
- **Ever heard of LLMNR or PNRP?**
- **Still very little understood, especially as for sec aspects.**
- **Some enabled by default in Win Vista.**
- **Most probably we will see “juicy stuff“ here...**



Conclusions

- ***Carrier Ethernet* (in the sense of this talk) has some interesting (and still not too well understood) security implications.**
- **The mitigate the resulting risks a thorough understanding of the techniques involved is needed on both the carrier side and the customer side.**
- **The whole space is still moving, with constant change from different angles.**



Questions?



Thanks for your attention!



Final Wisdom

Whatever you do... always remember the following two:

- **Ross Callon in *RFC 1925*:**

“Some things in networking can never be fully understood by someone who neither builds commercial networking equipment nor runs an operational network.”

=> If really interested in this stuff get your hands on some devices ;-)

- ***Simplicity Principle* from <http://tools.ietf.org/html/draft-ymbk-arch-guidelines-05>**

