

All your packets are belong to us

—

Attacking backbone technologies

Daniel Mende & Enno Rey
{dmende, erey}@ernw.de



Who we are



- **Old-school network geeks.**
- **Working as security researchers for Germany based ERNW GmbH.**
- **Fiddling around with devices and protocols makes the majority of our days.**



Agenda

- **Introduction & Dimensions of this talk**
- **BGP**
- **MPLS**
- **Carrier Ethernet**
- **Summary & Outlook**



Dimensions of this talk

- **We want you to reflect on the way \$TECHNOLOGIES work**
 - Some discussion of trust models
 - If you consider this “some esoteric shit”... throw rotten eggs on us ;)
- **We want you to have a mild laughter**
 - That’s why we included that “bingo stuff” (see next slide)
 - But, honestly, quite some time this is not too funny...
- **We want to entertain you**
 - Some demos might help to achieve this (the “Meat!” sections)




Bingo [www.crypto.com/bingo/pr]

[shuffle]
45342

YOU ARE IN VIOLATION OF THE DMCA	WHAT DO YOU HAVE AGAINST US?	OUR SUCCESS SPEAKS FOR ITSELF	YOU'RE JUST AN ACADEMIC	NOBODY WILL EVER TRY TO DO THIS
THAT'S ONLY THERE FOR BACKWARD COMPATIBILITY	YOU'RE PARANOID	WE HAVE CISSP CERTIFIED ENGINEERS	WE EMPLOY TOP SECURITY EXPERTS	NO ONE HAS EVER FOUND ANY PROBLEMS
THIS IS PROBABLY FIXED IN THE NEXT RELEASE	WE READ SCHNEIER'S BOOK	SECURITY PROBLEM EXCUSE BINGO	WHY DO YOU HATE AMERICA?	NOBODY'S PERFECT
IT'S SECURE ENOUGH FOR OUR CUSTOMERS	WE MEET ALL GOVERNMENT STANDARDS	OUR PROACTIVE TECHNOLOGY SOLUTIONS PREVENT THAT	WE THINK IT IS SECURE ENOUGH	YOU'RE BEING IRRESPONSIBLE
YOU'RE ONLY HELPING THE BAD GUYS	WE ALREADY KNEW ABOUT IT	EVERYBODY DOES IT THIS WAY	THE ANTI-VIRUS SOFTWARE DID IT	LA, LA, LA WE'RE NOT LISTENING

2007 JUTTA DEGENER, MATT BLAZE JUTTA@POBOX.COM - PERMALINK



BGP

- **Border Gateway Protocol**
- **Most current version as of RFC 1771 (March 1995)**
- **The glue that keeps the internet together.**
- **Has an interesting trust model.**
- **Was subject of some heavy debate last year.**




BGP - How it works

- **BGP speakers (“peers”) establish relationships with neighboring peers**
 - BGP works over /relies on TCP
 - => no multicasting (=> you can’t easily join a “group of BGP speakers”)
 - No (easy) spoofing
- **Peers announce “Network Layer Reachability Information” (NLRI)**
 - Think: “I know that some network can be reached via some way”
- **NLRIs (+ attributes) serve for *path* building/calculation.**



BGP - Trust Model

- TCP based => mostly configured manually / by script
- => “Intra Operator Trust” [amongst humans]
- Error prone
 - AS7007 Incident
 - YouTube / Pakistan



- Once you're a member of the “old boys club” you might perform all sorts of nasty stuff
 - Pilosov / Kapela 2008



BGP - Security mechanisms

- **MD5 signature, mainly for integrity checking**
 - Uses “generic TCP MD5 Signature Option” (RFC 2385)
- **Certainly that bell in your head just rang... yes: “MD5”**
 - Anybody attended 25C3 recently? ;-)
 - Still, similar attacks would be quite difficult.
 - And “they’re working on it”
 - <http://tools.ietf.org/id/draft-ietf-tcpm-tcp-auth-opt-04.txt>
- **Use of MD5 key secured BGP considered Carrier BCP**
 - Does it really add security value?



Meat!

- **ERNW tool “bgp_cli”**
 - Initially research tool for a student writing about trust (Hi Micele!)
 - Can be used to manually inject routes (role of “valid peer” assumed)
 - Can be used to bruteforce MD5 keys
 - In a direct session-based manner
- **ERNW tool “bgp_md5crack”**
 - Written in C => fast!
 - Can work on pcap file...
 - ... or “live” on interface
- **Demos ;-)**



For completeness' sake

- The BGP key used in the campus backbone of a 40K user environment we audited a while ago:

```
(ciscocrack) > ./ciscocrack 070C285F4D06  
Passwd: cisco
```



MPLS

- ***Multiprotocol Label Switching* [RFC 3031 et.al.]**
- **Technology used for forwarding packets, based on *Labels***
Packets may carry multiple labels (for different purposes).
- **Deployed in most carrier backbones.**

- **We are going to cover two subsets of the MPLS technology called “MPLS Layer 3 VPNs” and “MPLS Layer 2 VPNs”**
- **To be found in most \$\$\$ enterpri. for their global networks.**



- **MPLS-based technology [mainly RFC 4364] 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.**



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)

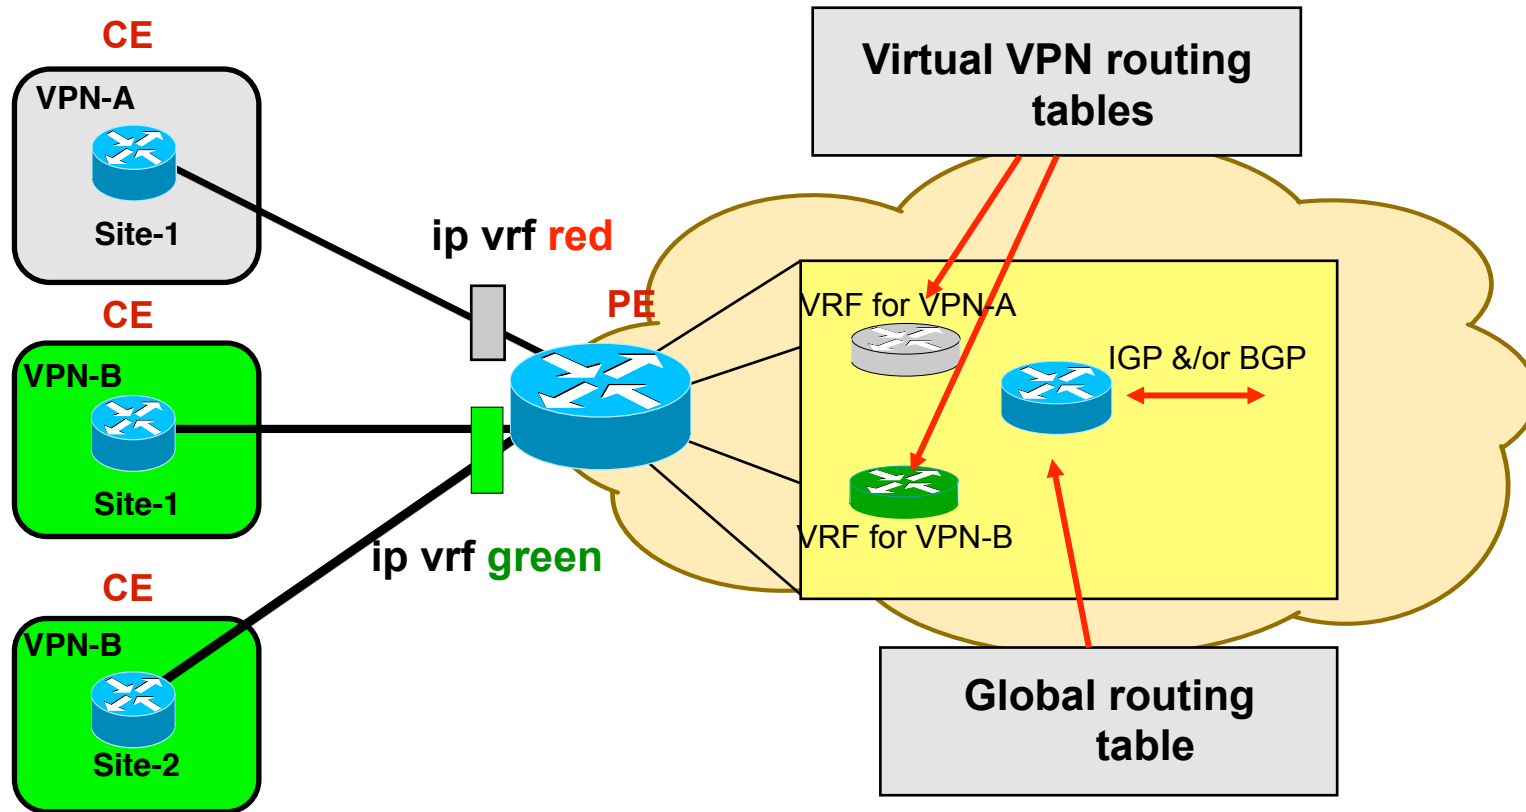
- Router connecting the C network to the PE (may be under control of customer or ISP)



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




MPLS Layer 3 VPNs



MPLS 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, usually 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".



MPLS VPNs, Trust Model

- **Trusted Core** is assumed.
- **No attacks from outside the core possible.**
- **No additional security controls available**
 - “Trust my blue eyes!”
 - Oh yes, there is MD5 protected LDP... please, would anybody mind explaining us the underlying threat model?
- **Source of grim debates between \$Corp_Global_NW_Team and \$Corp_Info_Sec.**

NOBODY
WILL EVER
TRY TO
DO THIS

YOU'RE
PARANOID




Meat!

- **ERNW Tool “mpls_redirect”**
 - Assumes attacker has access to traffic path (in core).
 - Command line tool
 - Modifies “VPN labels” of packets
 - => Redirects traffic from one customer to another “customer”
[yes, you clever guys, that’s what the name came from...]


- **Demo**




(Bi-directional) Modification of VPN Labels




PING Beer to Beer




PING Beer to Spliff



Some magic [mushrooms?] comes into play ;-)



PING Beer to Spliff with some magic



What does this mean?

- **Attacker can get into VPNs.**
 - Attacker can set up fake “central authorization portal” and re-direct an enterprise’s traffic to it.
 - Same for DNS
 - Same for LDAP
 - Same for ...
- **Use your imagination ;-)**
- **Still, we can only re-label existing traffic. Wouldn’t it be nice to ...**



more meat!

(`"meat!: no such file or directory" ;-)`)



■ ERNW Tool “mpls_tun”

- Assumes attacker has access to traffic path (in core).
- Creates a virtual interface that is “part of a given MPLS VPN”.
- So far only tested with Linux.
- Now attacker has “VPN enabled” network stack.
 - Use all your favorite attack tools “into” some VPN, against various sites.

■ Demo



Mitigating controls

- **“Trust your carrier”**

- This was not a joke ;-) ... if you do, that’s ok. We’re ok, too.
- Contractual controls might kick in.



- **“Authenticate everything”.**

- Breaks approach of “trusted networks”

- **Implement “borders of trust” (e.g. L3 devices) that encrypt/decrypt all inbound traffic on a site level.**

- **Again, our main message is: It’s all about risk [mgmt].**




Definition of *Carrier Ethernet*

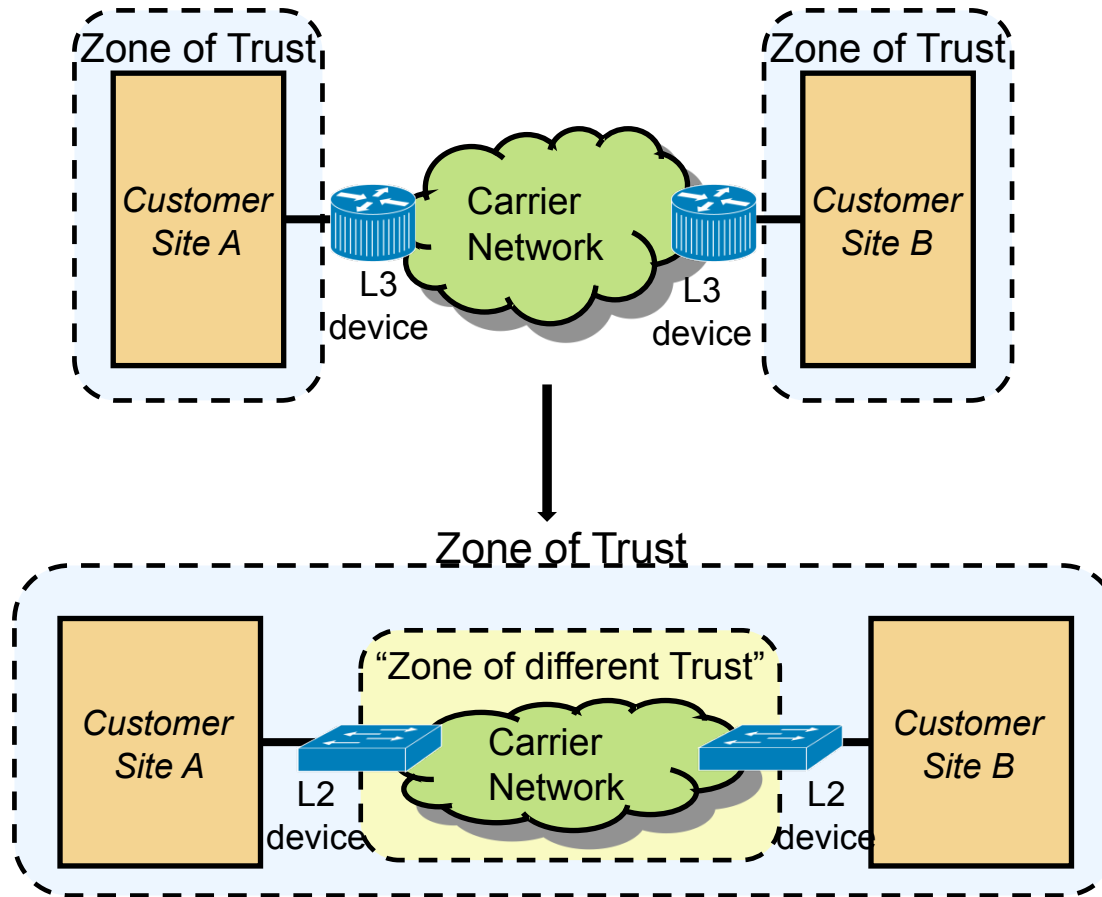
- ***Carrier Ethernet*** basically means that ethernet frames are transported across (at least) one carrier's backbone.
- So ethernet is not (only) used as an *access medium* here, but offered as a *service*.
- **Technologies**
 - Metro Ethernet
 - EoMPLS / VPLS
 - L2TPv3



Example: Ethernet over MPLS



Change of (ethernet) trust model



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.



■ Existing threats have new scope

- Ethernet based attacks may be performed “over the cloud”
 - E.g. attacker in site Brussels might *arp-spoof* (=read) traffic from site Amsterdam.
- Misconfigurations will have larger impact
 - What about that old C2980 with a high VTP rev.-number, accidentally re-plugged in?

■ New threats may show up

- Existing ethernet protocol space not designed for worldwide networks.
 - Spanning Tree dates from 1980s.
 - Again: their whole trust model is built around a concept of “local networks”.
- Segmentation capabilities of technologies involved may not be sufficient for some security needs.



Traditional Ethernet Attacks “over the cloud”

- Depend highly on the level of transparency a “VPLS cloud” provides.
- Given full transparency (as in *Cisco*-based testbed we used)...
- ... you can perform any traditional layer 2 attack over the cloud.
- We tested this successfully with *yersinia*.
- From an attacker’s perspective this is pretty cool: sitting in Brussels and arp-spoofing some boxes located in Amsterdam...



It might not only happen “over the cloud”...

- **... but also “from the cloud” ;-)**
 - `mpls_tun` would do it (see above).
- **For completeness’ sake, one more...**
- **ERNW tool “ldp_cli”**
 - Take part in LDP discovery.
 - Take part in subsequent LDP sessions.
 - Propagate LDP information at your will, e.g. L2 VPN signaling.




Wrap-up on Carrier Ethernet

- **Interesting approach**
(“as networkers” we pretty much like it).
- **Gaining ground commercially.**
 - E.g. for SAN replication.
- **Changes whole trust model of Ethernet**
 - Might have large security implications.



Save the best for last



Some fun with MP-BGP...



Summary & Outlook

- **There are some backbone technologies with a “debatable” trust model.**
 - And “debatable” resulting security controls / control capabilities.
- **Our talk’s intent was to made you aware of that. It’s just that simple ;-)**
- **Oh, btw:**

`www.ernw.de/download/bh09_all_your_packets_tools.tar.bz2`
`0c67d956787d20b2b6d3d265c2acc030eb783c8b3e58ca65200664f6f8293fc3`



There's never enough time...

THANK YOU...



...for yours!



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/rfc/rfc3439.txt>**

