Fireside Fridays

Routing and VLANs
Week 6

Thanks to our sponsors!

ACTIVE | COUNTERMEASURES

BLACK HILLS
Information Security
• 2008 •

WILD WEST
HACKIN' FEST

Antisyphon Training

Special Thanks to…

- Hermon 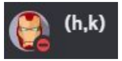
(h,k)
- Emily 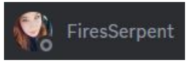
FiresSerpent
- Both gave up many late nights to help with QA and development of this content
- Very much appreciate their efforts!
- Please give them a warm "thanks" the next time you see them online

Lab requirements for this section

- Windows or Linux system
- Labs will be at the command line or terminal

Routing

- Forwards traffic based on Layer 3 info
- Typically destination IP address
- Policy routing may include source IP
- Routing table entries can be static or dynamic
- Static is harder to attack
  - But bogus ICMP redirects can attack both
- Provides better traffic isolation than switches

Hands-on walkthrough - route table

- This walkthrough will vary between OSes
- Slightly different command on each
- Want to view the current routing table
    - Linux/Mac = route -n
    - Windows = route print

Linux routing table

Default route

Network leading to default gateway

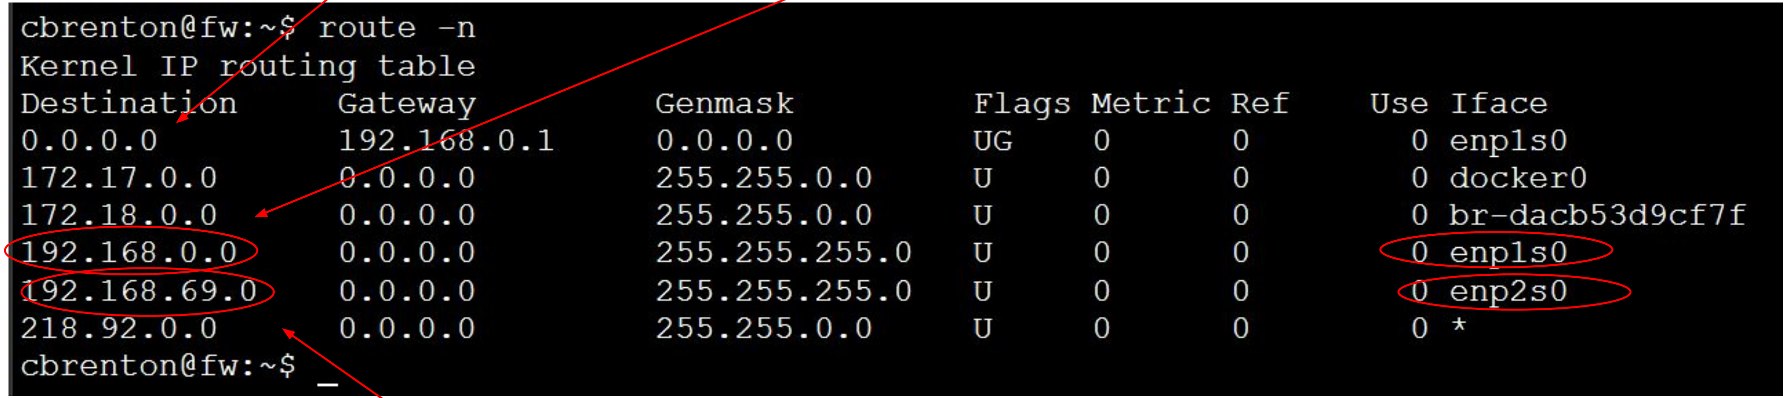

```
cbrenton@fw:~$ route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
0.0.0.0         192.168.0.1     0.0.0.0         UG    0      0        0 enp1s0
172.17.0.0      0.0.0.0         255.255.0.0     U     0      0        0 docker0
172.18.0.0      0.0.0.0         255.255.0.0     U     0      0        0 br-dacb53d9cf7f
192.168.0.0     0.0.0.0         255.255.255.0   U     0      0        0 enp1s0
192.168.69.0    0.0.0.0         255.255.255.0   U     0      0        0 enp2s0
218.92.0.0      0.0.0.0         255.255.0.0     U     0      0        0 *
cbrenton@fw:~$ _
```

Second network interface

Linux usually has multiple options

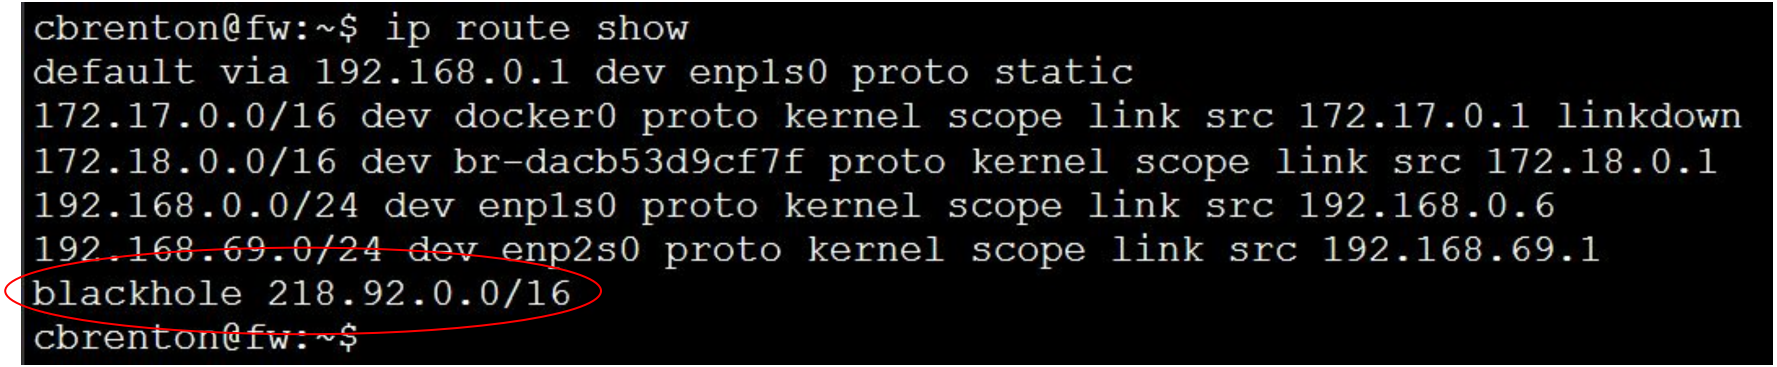

```
cbrenton@fw:~$ ip route show
default via 192.168.0.1 dev enp1s0 proto static
172.17.0.0/16 dev docker0 proto kernel scope link src 172.17.0.1 linkdown
172.18.0.0/16 dev br-dacb53d9cf7f proto kernel scope link src 172.18.0.1
192.168.0.0/24 dev enp1s0 proto kernel scope link src 192.168.0.6
192.168.69.0/24 dev enp2s0 proto kernel scope link src 192.168.69.1
blackhole 218.92.0.0/16
cbrenton@fw:~$
```

Basically the same info

"Blackhole" eats responses to this network

More on blackhole later

What's this?

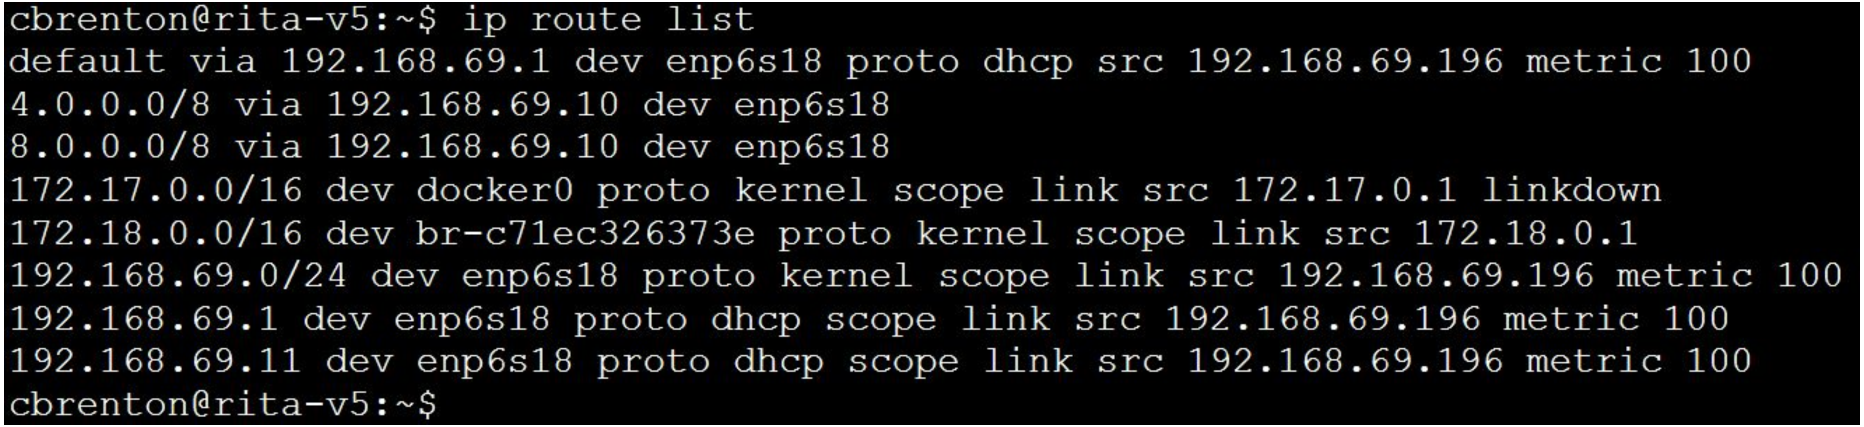

```
cbrenton@rita-v5:~$ ip route list
default via 192.168.69.1 dev enp6s18 proto dhcp src 192.168.69.196 metric 100
4.0.0.0/8 via 192.168.69.10 dev enp6s18
8.0.0.0/8 via 192.168.69.10 dev enp6s18
172.17.0.0/16 dev docker0 proto kernel scope link src 172.17.0.1 linkdown
172.18.0.0/16 dev br-c71ec326373e proto kernel scope link src 172.18.0.1
192.168.69.0/24 dev enp6s18 proto kernel scope link src 192.168.69.196 metric 100
192.168.69.1 dev enp6s18 proto dhcp scope link src 192.168.69.196 metric 100
192.168.69.11 dev enp6s18 proto dhcp scope link src 192.168.69.196 metric 100
cbrenton@rita-v5:~$
```

Anything odd about this output?

Windows route command

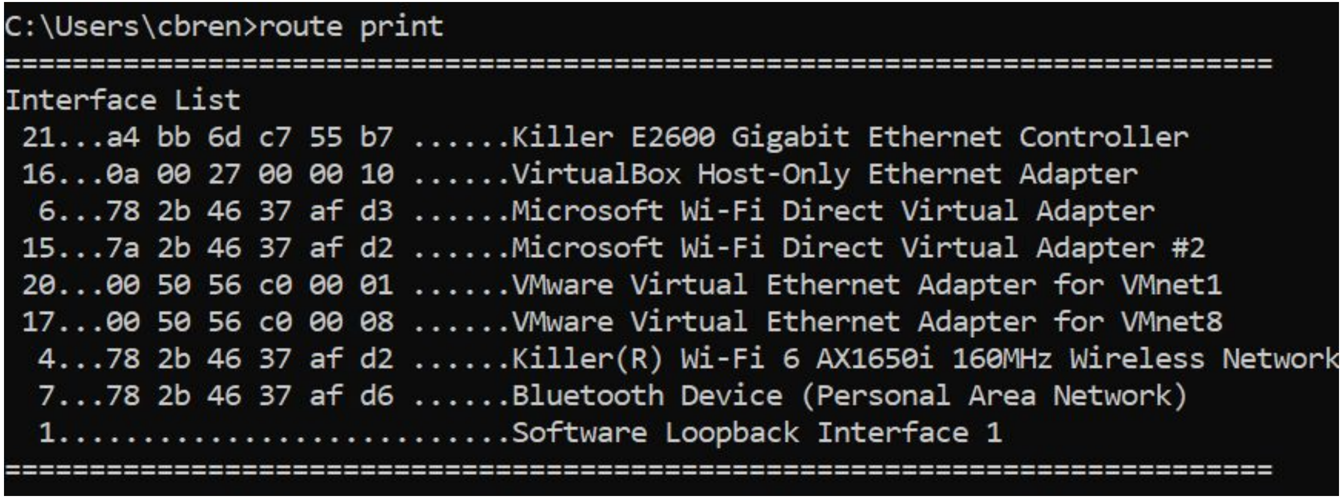

```
C:\Users\cbren>route print
===========================================================================
Interface List
 21...a4 bb 6d c7 55 b7 ......Killer E2600 Gigabit Ethernet Controller
 16...0a 00 27 00 00 10 ......VirtualBox Host-Only Ethernet Adapter
  6...78 2b 46 37 af d3 ......Microsoft Wi-Fi Direct Virtual Adapter
 15...7a 2b 46 37 af d2 ......Microsoft Wi-Fi Direct Virtual Adapter #2
 20...00 50 56 c0 00 01 ......VMware Virtual Ethernet Adapter for VMnet1
 17...00 50 56 c0 00 08 ......VMware Virtual Ethernet Adapter for VMnet8
  4...78 2b 46 37 af d2 ......Killer(R) Wi-Fi 6 AX1650i 160MHz Wireless Network
  7...78 2b 46 37 af d6 ......Bluetooth Device (Personal Area Network)
  1...........................Software Loopback Interface 1
===========================================================================
```

Windows first prints a list of known interfaces

Windows routing table

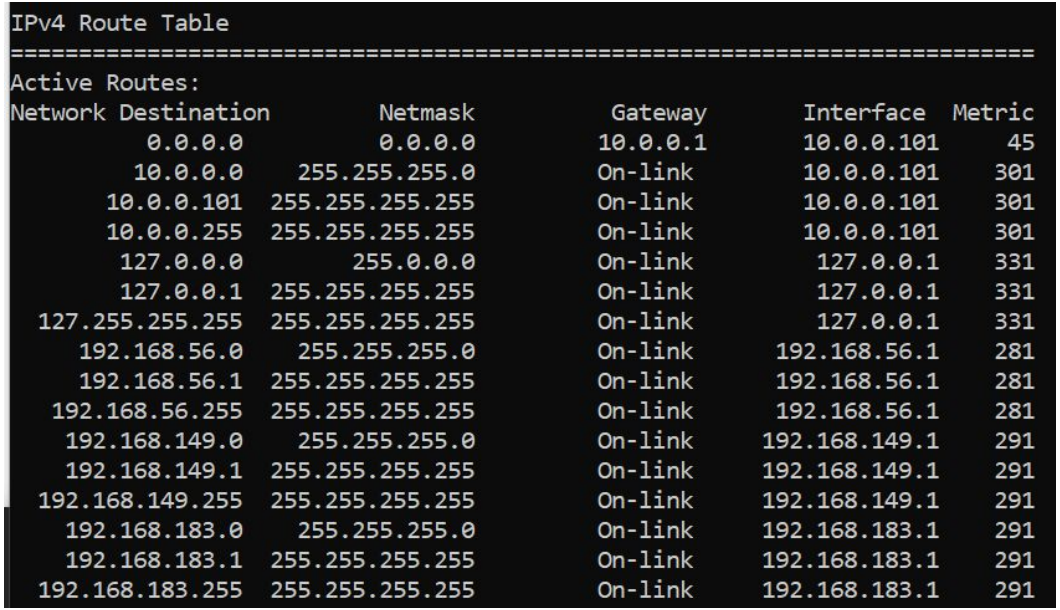

```
IPv4 Route Table
===========================================================================
Active Routes:
Network Destination        Netmask          Gateway       Interface  Metric
          0.0.0.0          0.0.0.0         10.0.0.1      10.0.0.101     45
         10.0.0.0    255.255.255.0         On-link       10.0.0.101    301
       10.0.0.101  255.255.255.255         On-link       10.0.0.101    301
       10.0.0.255  255.255.255.255         On-link       10.0.0.101    301
        127.0.0.0        255.0.0.0         On-link        127.0.0.1    331
        127.0.0.1  255.255.255.255         On-link        127.0.0.1    331
  127.255.255.255  255.255.255.255         On-link        127.0.0.1    331
     192.168.56.0    255.255.255.0         On-link     192.168.56.1    281
     192.168.56.1  255.255.255.255         On-link     192.168.56.1    281
   192.168.56.255  255.255.255.255         On-link     192.168.56.1    281
    192.168.149.0    255.255.255.0         On-link    192.168.149.1    291
    192.168.149.1  255.255.255.255         On-link    192.168.149.1    291
  192.168.149.255  255.255.255.255         On-link    192.168.149.1    291
    192.168.183.0    255.255.255.0         On-link    192.168.183.1    291
    192.168.183.1  255.255.255.255         On-link    192.168.183.1    291
  192.168.183.255  255.255.255.255         On-link    192.168.183.1    291
```

Lower metric number is more preferred route

It then prints known routes

Routing protocols (1 of 2)

- Static
  - Hard set into system
  - Usually learned via DHCP
- Distance vector
  - Build tables based on neighbor announcements
  - Don't know what the whole network looks like
  - Memory efficient but not always accurate
  - Easy to deploy
  - Used by smaller networks

Routing protocols (2 of 2)

- Link state
  - Each router draws a map of the networks they connect to
  - This info is shared with other routers to jigsaw together
  - Full picture uses more resources but provides better recovery
  - Used internally by larger networks
  - Arguably most popular option after static
- Path-vector Routing
  - Focuses on path rather than hop count
  - Useful when storing individual routes wouldn't fit in RAM
  - This is how routing on the Internet works

Popular routing protocols

- Routing Information Protocol (RIP, all versions)
  - Distant vector based
  - One of the older routing protocols
- Open Shortest Path First (OSPF)
  - Link state based
  - Popular internal routing option
- Border Gateway Protocol (BGP)
  - Pre-vector based, runs the Internet

BGP drill down

1.1.1.1 doesn't need exact route info to get to 2.2.2.2

Limited number of peer points between providers

Only needs to know to send it to AS3 and let that router figure out exact path

BGP routes based on reachability rather than exact hop count

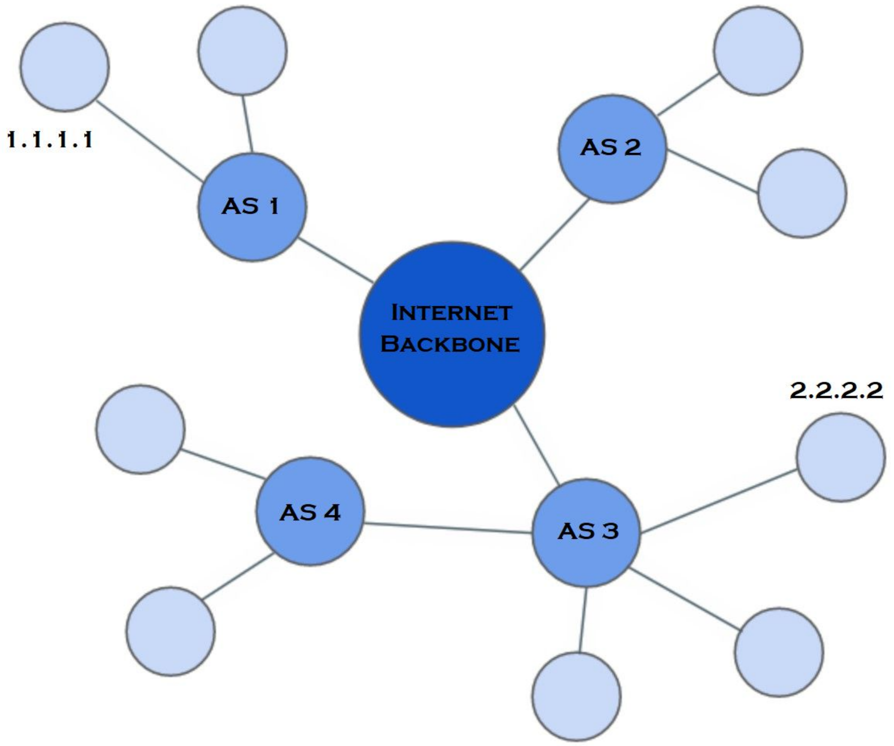

BGP hijacking

- Ass-u-me's all other BGP routers are always truthful
- No way to validate which BGP routers go with which networks
- Hijacker simply advertises:
  - More specific route
  - Shorter path to get there
- This funnels traffic through a network of their choice
- Traffic can be monitored, hijacked or blackholed
- Hard to detect - attacks can last for days
- sBGP can fix this, but ISPs have no interest in deploying it

BGP attack example

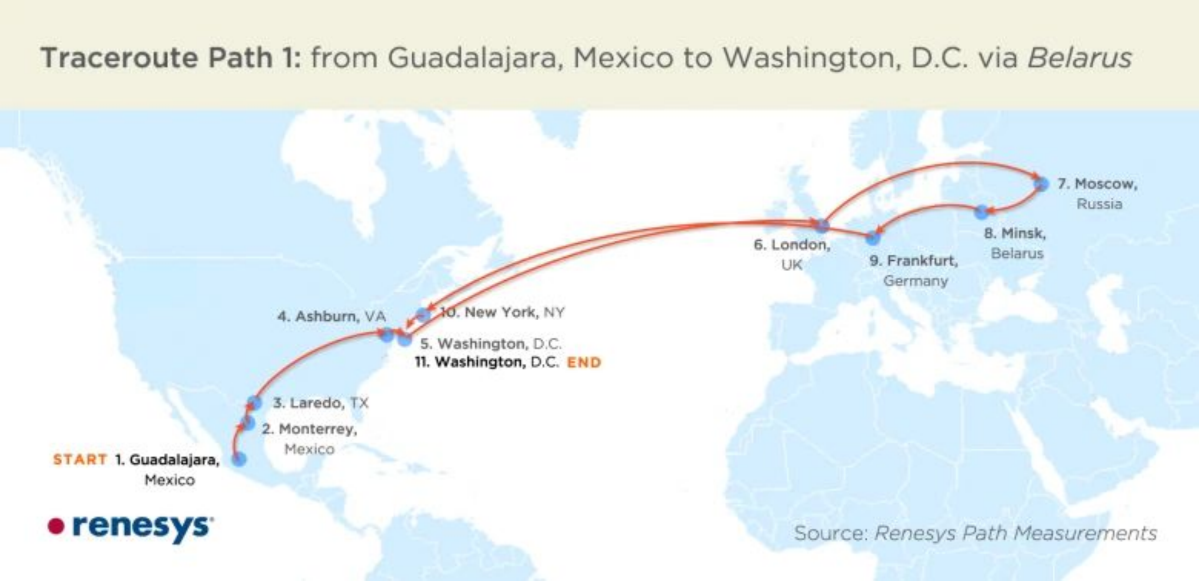

Traceroute Path 1: from Guadalajara, Mexico to Washington, D.C. via *Belarus*

Source: *Renesys Path Measurements*

Aug, 2013. 38 times, traffic from Mexico to US government agencies, diplomatic offices of multiple countries, and credit card transactions were routed through Belarus and Russian networks.

Blackholing networks

- Bans communications to a specified network
- Block networks originating lots of malicious activity
- Also useful when "customers" are geographically defined
- Typically implemented on routers or possibly firewalls
- Created using bogus route entries
  - Packets still get in, replies are not returned
  - Far more processor and memory efficient than firewalling

Blackhole Chinanet-Backbone

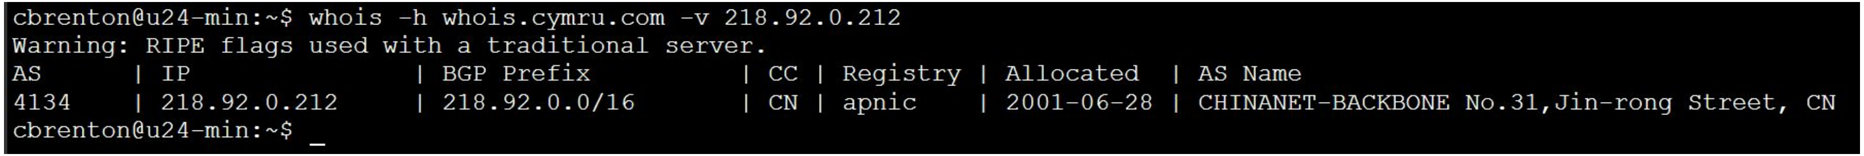

```
cbrenton@u24-min:~$ whois -h whois.cymru.com -v 218.92.0.212
Warning: RIPE flags used with a traditional server.
AS      | IP               | BGP Prefix          | CC | Registry | Allocated  | AS Name
4134    | 218.92.0.212     | 218.92.0.0/16       | CN | apnic    | 2001-06-28 | CHINANET-BACKBONE No.31,Jin-rong Street, CN
cbrenton@u24-min:~$ _
```

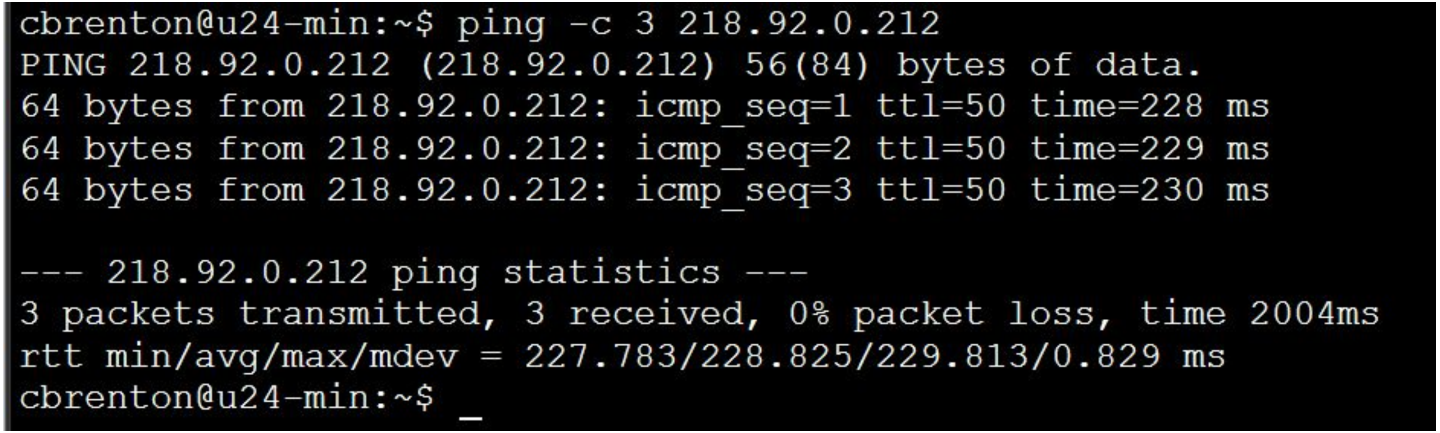

```
cbrenton@u24-min:~$ ping -c 3 218.92.0.212
PING 218.92.0.212 (218.92.0.212) 56(84) bytes of data.
64 bytes from 218.92.0.212: icmp_seq=1 ttl=50 time=228 ms
64 bytes from 218.92.0.212: icmp_seq=2 ttl=50 time=229 ms
64 bytes from 218.92.0.212: icmp_seq=3 ttl=50 time=230 ms

--- 218.92.0.212 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2004ms
rtt min/avg/max/mdev = 227.783/228.825/229.813/0.829 ms
cbrenton@u24-min:~$ _
```

ID malicious host. Verify it's reachable.
We'll go after it's entire network.

Blackhole route on the firewall

Block all traffic to the entire Chinanet backbone

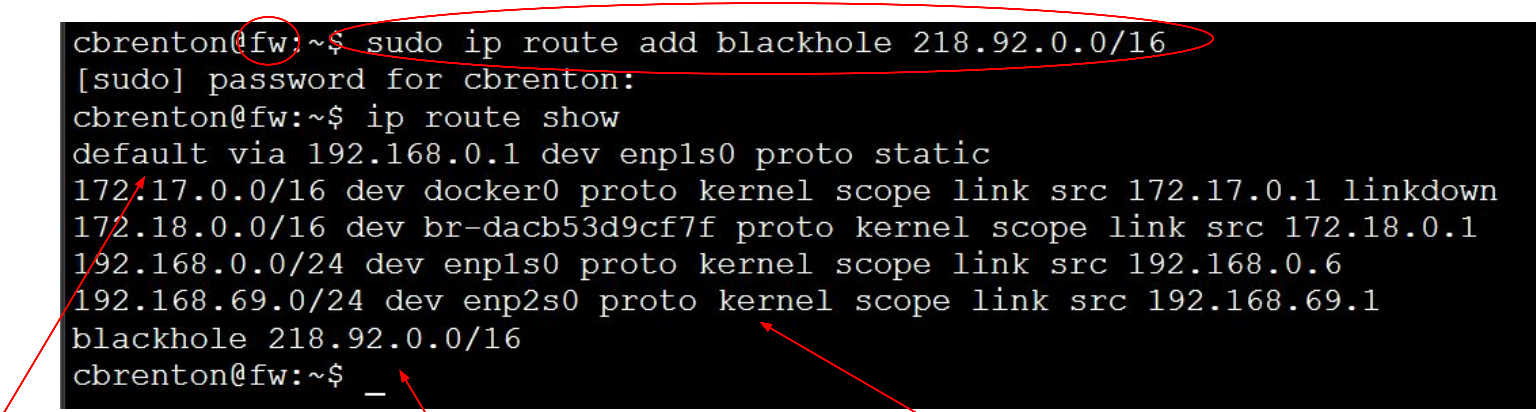

```
cbrenton@fw:~$ sudo ip route add blackhole 218.92.0.0/16
[sudo] password for cbrenton:
cbrenton@fw:~$ ip route show
default via 192.168.0.1 dev enp1s0 proto static
172.17.0.0/16 dev docker0 proto kernel scope link src 172.17.0.1 linkdown
172.18.0.0/16 dev br-dacb53d9cf7f proto kernel scope link src 172.18.0.1
192.168.0.0/24 dev enp1s0 proto kernel scope link src 192.168.0.6
192.168.69.0/24 dev enp2s0 proto kernel scope link src 192.168.69.1
blackhole 218.92.0.0/16
cbrenton@fw:~$ _
```

Route to Internet

Blackhole route added

Internal network

Dead route walking…

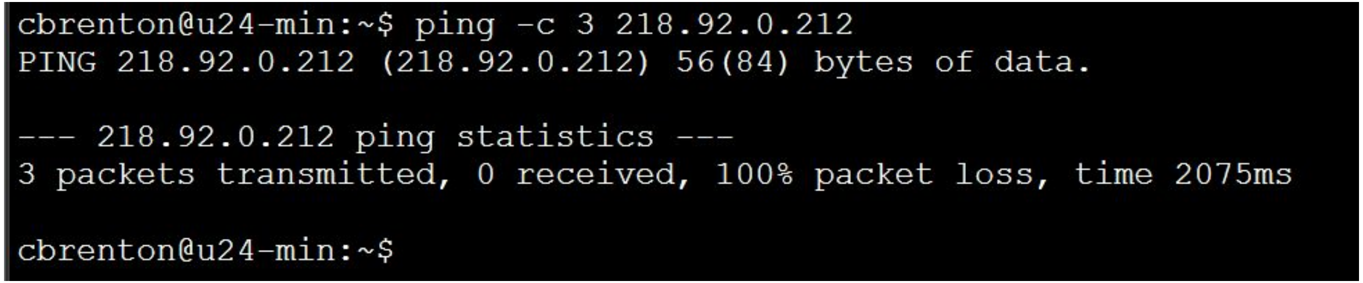

```
cbrenton@u24-min:~$ ping -c 3 218.92.0.212
PING 218.92.0.212 (218.92.0.212) 56(84) bytes of data.

--- 218.92.0.212 ping statistics ---
3 packets transmitted, 0 received, 100% packet loss, time 2075ms

cbrenton@u24-min:~$
```

Packets make it to the firewall

Routing on firewall sends them to /dev/null

/dev/null is device that makes all data disappear

Packets to this network never make it to the Internet

VLANs

- "Feels" like routing, but it's not
- Segregates traffic flow via software
- Permits you to create multiple virtual networks on top of a single physical topology
- Implemented by:
    - Per port, software setting in the switch
    - VLAN tagging
        - 16 byte field added to Ethernet header
        - 12 of those bytes are the VLAN identifier
    - Combo of the above

Example VLAN tag

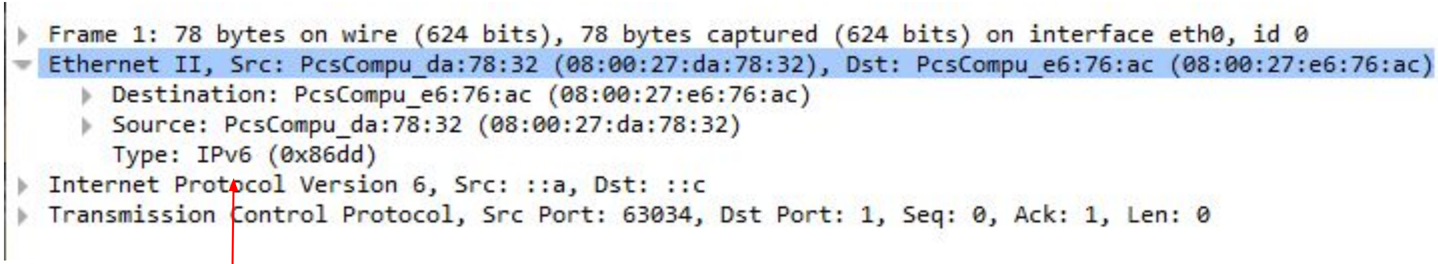

```
▸ Frame 1: 78 bytes on wire (624 bits), 78 bytes captured (624 bits) on interface eth0, id 0
▾ Ethernet II, Src: PcsCompu_da:78:32 (08:00:27:da:78:32), Dst: PcsCompu_e6:76:ac (08:00:27:e6:76:ac)
    ▸ Destination: PcsCompu_e6:76:ac (08:00:27:e6:76:ac)
    ▸ Source: PcsCompu_da:78:32 (08:00:27:da:78:32)
      Type: IPv6 (0x86dd)
▸ Internet Protocol Version 6, Src: ::a, Dst: ::c
▸ Transmission Control Protocol, Src Port: 63034, Dst Port: 1, Seq: 0, Ack: 1, Len: 0
```

Regular fame

VLAN tagging

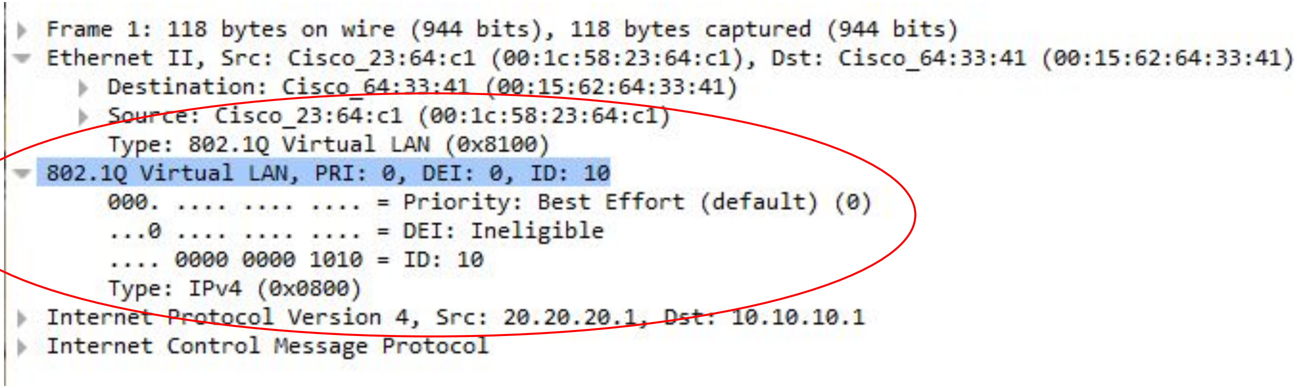

```
▸ Frame 1: 118 bytes on wire (944 bits), 118 bytes captured (944 bits)
▾ Ethernet II, Src: Cisco_23:64:c1 (00:1c:58:23:64:c1), Dst: Cisco_64:33:41 (00:15:62:64:33:41)
    ▸ Destination: Cisco_64:33:41 (00:15:62:64:33:41)
    ▸ Source: Cisco_23:64:c1 (00:1c:58:23:64:c1)
      Type: 802.1Q Virtual LAN (0x8100)
▾ 802.1Q Virtual LAN, PRI: 0, DEI: 0, ID: 10
      000. .... .... .... = Priority: Best Effort (default) (0)
      ...0 .... .... .... = DEI: Ineligible
      .... 0000 0000 1010 = ID: 10
      Type: IPv4 (0x0800)
▸ Internet Protocol Version 4, Src: 20.20.20.1, Dst: 10.10.10.1
▸ Internet Control Message Protocol
```

VLAN weaknesses

- Per port is pretty solid
- VLAN tagging vulnerable to MITM attacks
    - Backbone access can still sees everything
    - Tag ID can be modified/changed, no authentication
    - Requires local access to exploit
- VPN technologies encrypt, but don't prevent tag spoofing
    - HTTPS, IPSec, etc.
    - These work at layer 3 and above (No auth for Ethernet fields)
    - Will not recognize malicious tag modifications

Multilayer switches

- Sometimes incorrectly called layer 3 switching
- Really it's just a router, but faster than traditional
- Trade off is less routing functionality
  - Fine for LAN, usually insufficient for dynamic WAN
- Hardware based while classic routers are software based
  - Harder to patch when vulnerabilities are identified
  - Adds a layer of security PITA

Next week on Fireside Fridays!!!

- We are taking a break next week
- I'll be in Las Vegas at Right of Boom
- Next class will be on the 28th
- We will be discussing
  - IP addressing
  - Table conversions
  - A bit of data obfuscation

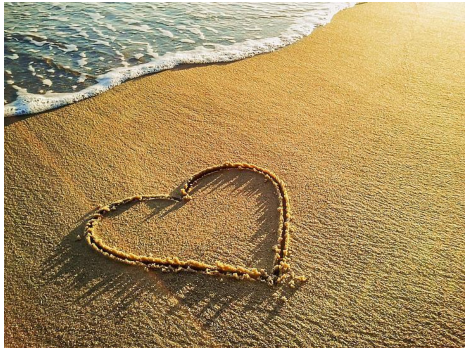

- Thanks for spending your Valentine's Day with us!

Wrap up

- Thank you for attending!
- Certs & video usually go out in 24 hours
- If you have any lingering questions, the Discord channel will remain active
  - Also a good chance to socialize with others in the class
  - Have other tips and tricks? Please share with others!
  - Posting screenshots can be helpful :-)