

End of the Journey Through the Dark Turn on the Light with Wireshark

Matthias Burkhard
IBM



Twitter

: mreeede



SNA Wizards

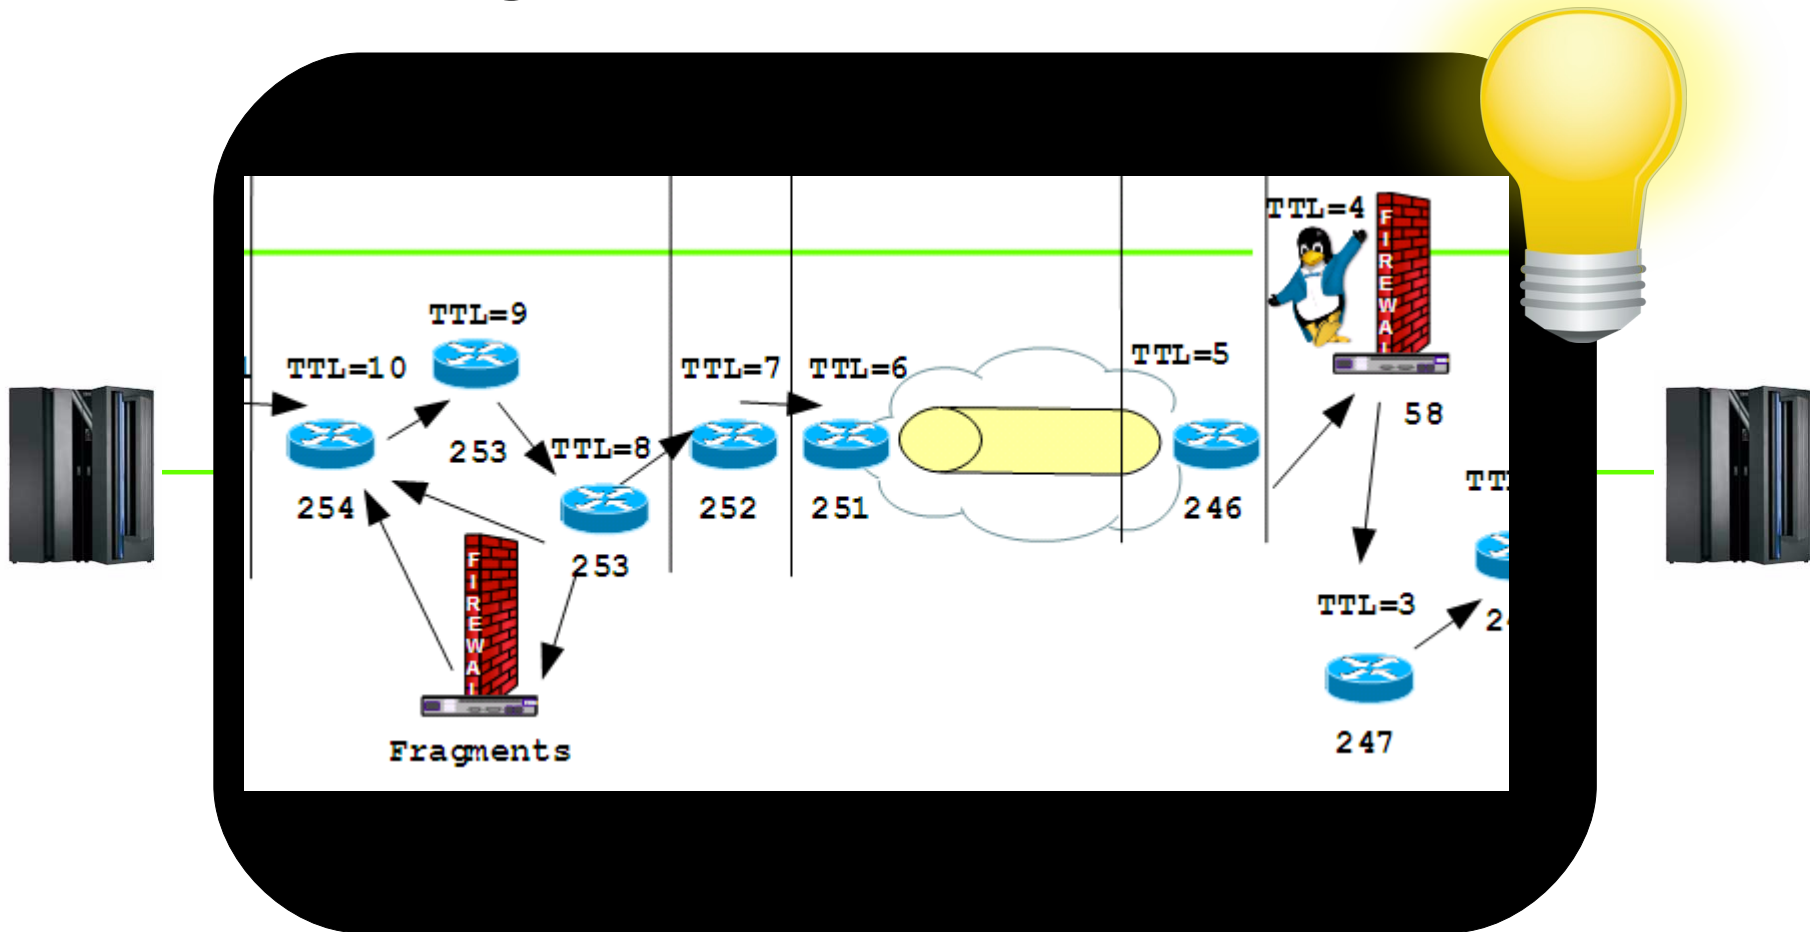
sna.wizards@groups.facebook.com

mburkhar@de.ibm.com

Tuesday, August 9, 2011: 4:30 PM-5:45 PM
Session 9248 – Europe 11

End the journey through the dark

Turn on the light with wireshark



The mother of all IP diagnostics: PING

<http://en.wikipedia.org/wiki/Sonar>

„active sonar is emitting pulses of sounds and listening for echoes. Sonar may be used as a means of acoustic location and of measurement of the echo characteristics of "targets" in the water.“

S9248_1.cap

File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help

Filter: icmp.type==0 or icmp.type==8

| Source | ip.id | ip.df | ip.len | Data | Time | Info |
|-----------------|--------|---------|--------|----------------------|----------------------------|--|
| 205.144.107.201 | 0x8779 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.533325 | Echo (ping) request id=0x6443, seq=1/256, ttl=49 |
| 198.147.171.51 | 0x7ee9 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.533325 | Echo (ping) reply id=0x6443, seq=1/256, ttl=64 |
| 205.144.107.201 | 0x8781 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.659208 | Echo (ping) request id=0x6443, seq=2/512, ttl=49 |
| 198.147.171.51 | 0x7eef | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.778293 | Echo (ping) reply id=0x6443, seq=2/512, ttl=64 |
| 205.144.107.201 | 0x8784 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.778293 | Echo (ping) request id=0x6443, seq=3/768, ttl=49 |
| 198.147.171.51 | 0x7ef0 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.778293 | Echo (ping) reply id=0x6443, seq=3/768, ttl=64 |

Ping comes in with TTL 49

src_ip is 205.144.107.201

Ping reply leaves with TTL 64

src_ip is 198.147.171.51


TTL and Topology I.

S9248_1.cap

File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help

Filter: icmp.type==0 or icmp.type==8 Expression... Clear Apply

| Source | ip.id | ip.df | ip.len | Data | Time | Info |
|-----------------|--------|---------|--------|----------------------|----------------------------|--|
| 205.144.107.201 | 0x8779 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.533325 | Echo (ping) request id=0x6443, seq=1/256, ttl=49 |
| 198.147.171.51 | 0x7ee9 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.533325 | Echo (ping) reply id=0x6443, seq=1/256, ttl=64 |
| 205.144.107.201 | 0x8781 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.659208 | Echo (ping) request id=0x6443, seq=2/512, ttl=49 |
| 198.147.171.51 | 0x7eef | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.778293 | Echo (ping) reply id=0x6443, seq=2/512, ttl=64 |
| 205.144.107.201 | 0x8784 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.778293 | Echo (ping) request id=0x6443, seq=3/768, ttl=49 |
| 198.147.171.51 | 0x7ef0 | Not set | 38 | 20101115113456519193 | 2010-11-15 12:34:56.778293 | Echo (ping) reply id=0x6443, seq=3/768, ttl=64 |



The trace was taken at 198.147.171.51 (TTL 64 = initial TTL of z/OS)
 The incoming PING traveled 15 hops (assuming this host has the same initial TTL)
 The hexadecimal PING payload is date/time – not a common z/OS payload...

Traceroute and TTL

tracerte.pcap

File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help

Filter: ip.addr== 92.123.69.104 or dns Expression... Clear Apply

| Source | ip.id | ip.ttl | ip.df | ip.len | Time | Info |
|--------------|----------------------------|--------|---------|--------|-----------|--|
| 10.0.118.20 | 0xa25c (41564) | 1 | Not set | 92 | 96.509023 | Echo (ping) request id=0x0400, seq=2816/11, ttl=1 |
| 10.0.118.1 | 0x57ed (22509), 0xa25c | 255, 1 | Not set | 56, 92 | 0.001927 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0xa25d (41565) | 1 | Not set | 92 | 0.000234 | Echo (ping) request id=0x0400, seq=3072/12, ttl=1 |
| 10.0.118.1 | 0x57ee (22510), 0xa25d | 255, 1 | Not set | 56, 92 | 0.000689 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0xa25e (41566) | 1 | Not set | 92 | 0.000248 | Echo (ping) request id=0x0400, seq=3328/13, ttl=1 |
| 10.0.118.1 | 0x57ef (22511), 0xa25e | 255, 1 | Not set | 56, 92 | 0.000735 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0xa25f (41567) | 2 | Not set | 92 | 0.989153 | Echo (ping) request id=0x0400, seq=3584/14, ttl=2 |
| 9.155.60.163 | 0x05f7 (1527), 0xa25f | 254, 1 | Not set | 56, 92 | 0.001879 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0xa260 (41568) | 2 | Not set | 92 | 0.000219 | Echo (ping) request id=0x0400, seq=3840/15, ttl=2 |
| 9.155.60.163 | 0x05f8 (1528), 0xa260 | 254, 1 | Not set | 56, 92 | 0.001012 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0xa261 (41569) | 2 | Not set | 92 | 0.000238 | Echo (ping) request id=0x0400, seq=4096/16, ttl=2 |
| 9.155.60.163 | 0x05f9 (1529), 0xa261 | 254, 1 | Not set | 56, 92 | 0.000953 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0xa262 (41570) | 3 | Not set | 92 | 0.995684 | Echo (ping) request id=0x0400, seq=4352/17, ttl=3 |
| 9.155.0.121 | 0x07ad (1965), 0xa262 | 253, 1 | Not set | 56, 92 | 0.001107 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0xa263 (41571) | 3 | Not set | 92 | 0.000251 | Echo (ping) request id=0x0400, seq=4608/18, ttl=3 |
| 9.155.0.121 | 0x07ae (1966), 0xa263 | 253, 1 | Not set | 56, 92 | 0.000980 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0xa264 (41572) | 3 | Not set | 92 | 0.000206 | Echo (ping) request id=0x0400, seq=4864/19, ttl=3 |
| 9.155.0.121 | 0x07af (1967), 0xa264 | 253, 1 | Not set | 56, 92 | 0.001107 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0x0000 (0), 0xa265 (41573) | 252, 0 | Not set | 56, 92 | 0.001140 | Time-to-live exceeded (Time to live exceeded in transit) |
| 9.155.0.25 | 0x0000 (0), 0xa265 (41573) | 252, 0 | Not set | 56, 92 | 0.001140 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0x0000 (0), 0xa266 (41574) | 252, 0 | Not set | 56, 92 | 0.001108 | Time-to-live exceeded (Time to live exceeded in transit) |
| 9.155.0.25 | 0x0000 (0), 0xa266 (41574) | 252, 0 | Not set | 56, 92 | 0.001108 | Time-to-live exceeded (Time to live exceeded in transit) |
| 10.0.118.20 | 0x0000 (0), 0xa267 (41575) | 252, 0 | Not set | 56, 92 | 0.001108 | Time-to-live exceeded (Time to live exceeded in transit) |
| 9.155.0.25 | 0x0000 (0), 0xa267 (41575) | 252, 0 | Not set | 56, 92 | 0.001108 | Time-to-live exceeded (Time to live exceeded in transit) |

Traceroute sets TTL too low to solicit ICMP error messages from the routers in the path

TTL 1: ip.id a25c,a25d,a25e getting response from 10.0.118.1

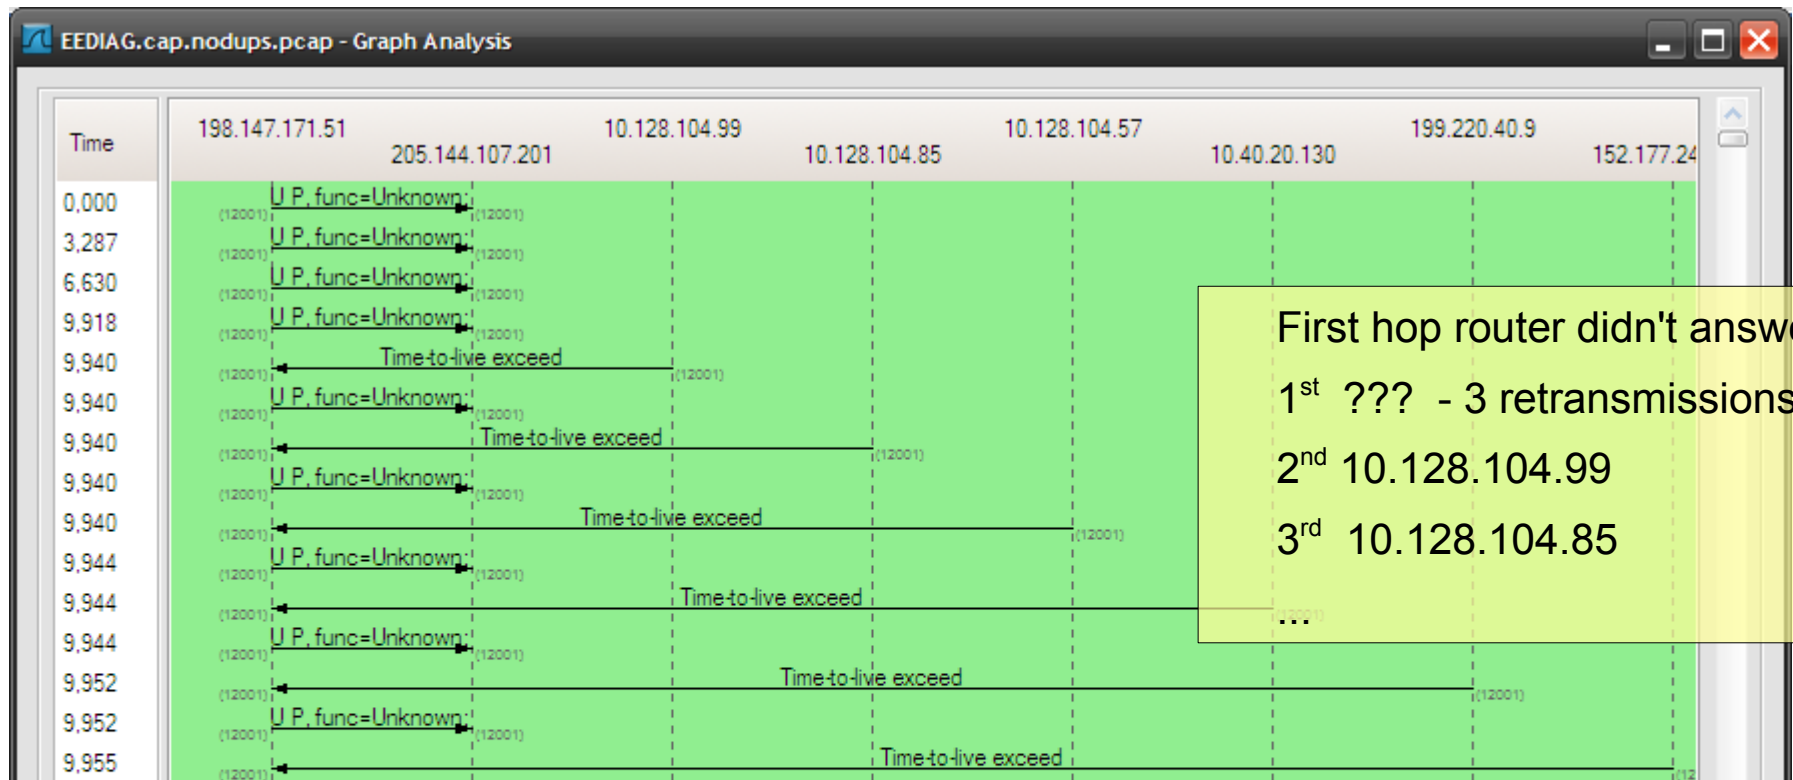
TTL 2 ip.id a25f,a260,a261 getting response from 9.155.60.163

...

Traceroute for HPR/IP EEDIAG TEST=YES

IP Packets are sent to all EE ports with TTL of 1, if no ICMP TTL exceeded response is received the packet is resent with 3.3 seconds interval

If a TTL exceeded message is received, the sender's src_ip and the RTT will be remembered



Traceroute for HPR/IP

EEDIAG TEST=YES

| | | | | | |
|-----|------------------|-----------|------|------------------------------|-----------------|
| 477 | EEDIAG TEST req | 0.000000 | 1 | 0x84c4 (33988) | 198.147.171.51 |
| 482 | EEDIAG TEST req | 3.195164 | 1 | 0x8506 (34054) | 198.147.171.51 |
| 487 | EEDIAG TEST req | 3.388341 | 1 | 0x8586 (34182) | 198.147.171.51 |
| 492 | EEDIAG TEST req | 3.268888 | 2 | 0x8608 (34312) | 198.147.171.51 |
| 497 | EEDIAG RTT reply | 0.044252 | 254, | 1 0x7e13 (32275), 0x8608 | 10.128.104.99 |
| 502 | EEDIAG TEST req | 0.000000 | 3 | 0x860d (34317) | 198.147.171.51 |
| 507 | EEDIAG RTT reply | 0.000000 | 253, | 1 0x4734 (18228), 0x860d | 10.128.104.85 |
| 509 | EEDIAG TEST req | 0.000000 | 4 | 0x8612 (34322) | 198.147.171.51 |
| 517 | EEDIAG RTT reply | 0.008905 | 252, | 1 0x8980 (35200), 0x8612 | 10.40.20.134 |
| 519 | EEDIAG TEST req | 0.000000 | 5 | 0x8617 (34327) | 198.147.171.51 |
| 527 | EEDIAG RTT reply | 0.000000 | 251, | 1 0x0000 (0), 0x8617 (34327) | 199.220.40.9 |
| 532 | EEDIAG TEST req | 0.000000 | 6 | 0x861c (34332) | 198.147.171.51 |
| 538 | EEDIAG RTT reply | 0.059784 | 246, | 1 0xa500 (42240), 0x861c | 152.177.242.210 |
| 539 | EEDIAG TEST req | 0.000000 | 7 | 0x8621 (34337) | 198.147.171.51 |
| 547 | EEDIAG RTT reply | 0.045674 | 58, | 1 0x5e1b (24091), 0x8621 | 205.144.107.201 |
| 550 | EEDIAG TEST req | 0.000000 | 8 | 0x8627 (34343) | 198.147.171.51 |
| 557 | EEDIAG RTT reply | 0.048145 | 247, | 1 0xd78e (55182), 0x8627 | 205.144.107.201 |
| 559 | EEDIAG TEST req | 0.000000 | 9 | 0x862d (34349) | 198.147.171.51 |
| 568 | EEDIAG RTT reply | 0.045130 | 247, | 1 0x77d9 (30681), 0x862d | 205.144.107.201 |
| 569 | EEDIAG TEST req | 0.000000 | 10 | 0x8634 (34356) | 198.147.171.51 |
| 577 | EEDIAG RTT reply | 0.041841 | 246, | 1 0xb6ea (46826), 0x8634 | 205.144.107.201 |
| 580 | EEDIAG TEST req | 0.000000 | 11 | 0x863b (34363) | 198.147.171.51 |
| 587 | EEDIAG RTT reply | 0.050343 | 245, | 1 0xf661 (63073), 0x863b | 205.144.107.201 |
| 592 | EEDIAG TEST req | 0.000000 | 12 | 0x8641 (34369) | 198.147.171.51 |
| 597 | EEDIAG TEST OK | 0.048407 | 49 | 0x66a3 (26275) | 205.144.107.201 |
| 602 | EEDIAG TEST req | 19.757549 | 1 | 0x88c8 (35016) | 198.147.171.51 |

TTL 11 was the last packet getting an ICMP

TTL 12 reached the other end.

The destination is 11 hops away

The TEST OK comes in with a TTL of 49

The initial TTL of the remote host is 60



The traceroute for HPR/IP: EEDIAG TEST=YES

| | | | | | |
|-----|------------------|-----------|-----|------------------------------|-----------------|
| 477 | EEDIAG TEST req | 0.000000 | 1 | 0x84c4 (33988) | 198.147.171.51 |
| 482 | EEDIAG TEST req | 3.195164 | 1 | 0x8506 (34054) | 198.147.171.51 |
| 487 | EEDIAG TEST req | 3.388341 | 1 | 0x8586 (34182) | 198.147.171.51 |
| 492 | EEDIAG TEST req | 3.268888 | 2 | 0x8608 (34312) | 198.147.171.51 |
| 497 | EEDIAG RTT reply | 0.044252 | 254 | 0x7e13 (32275), 0x8608 | 10.128.104.99 |
| 502 | EEDIAG TEST req | 0.000000 | 3 | 0x860d (34317) | 198.147.171.51 |
| 507 | EEDIAG RTT reply | 0.000000 | 253 | 1 0x4734 (18228), 0x860d | 10.128.104.85 |
| 509 | EEDIAG TEST req | 0.000000 | 4 | 0x8612 (34322) | 198.147.171.51 |
| 517 | EEDIAG RTT reply | 0.008905 | 252 | 1 0x8980 (35200), 0x8612 | 10.40.20.134 |
| 519 | EEDIAG TEST req | 0.000000 | 5 | 0x8617 (34327) | 198.147.171.51 |
| 527 | EEDIAG RTT reply | 0.000000 | 251 | 1 0x0000 (0), 0x8617 (34327) | 199.220.40.9 |
| 532 | EEDIAG TEST req | 0.000000 | 6 | 0x861c (34332) | 198.147.171.51 |
| 538 | EEDIAG RTT reply | 0.059784 | 246 | 1 0xa500 (42240), 0x861c | 152.177.242.210 |
| 539 | EEDIAG TEST req | 0.000000 | 7 | 0x8621 (34337) | 198.147.171.51 |
| 547 | EEDIAG RTT reply | 0.045674 | 58 | 1 0x5e1b (24091), 0x8621 | 205.144.107.201 |
| 550 | EEDIAG TEST req | 0.000000 | 8 | 0x8627 (34343) | 198.147.171.51 |
| 557 | EEDIAG RTT reply | 0.048145 | 247 | 1 0xd78e (55182), 0x8627 | 205.144.107.201 |
| 559 | EEDIAG TEST req | 0.000000 | 9 | 0x862d (34349) | 198.147.171.51 |
| 568 | EEDIAG RTT reply | 0.045130 | 247 | 1 0x77d9 (30681), 0x862d | 205.144.107.201 |
| 569 | EEDIAG TEST req | 0.000000 | 10 | 0x8634 (34356) | 198.147.171.51 |
| 577 | EEDIAG RTT reply | 0.041841 | 246 | 1 0xb6ea (46826), 0x8634 | 205.144.107.201 |
| 580 | EEDIAG TEST req | 0.000000 | 11 | 0x863b (34363) | 198.147.171.51 |
| 587 | EEDIAG RTT reply | 0.050343 | 245 | 1 0xf661 (63073), 0x863b | 205.144.107.201 |
| 592 | EEDIAG TEST req | 0.000000 | 12 | 0x8641 (34369) | 198.147.171.51 |
| 597 | EEDIAG TEST OK | 0.048407 | 49 | 0x66a3 (26275) | 205.144.107.201 |
| 602 | EEDIAG TEST req | 19.757549 | 1 | 0x88c8 (35016) | 198.147.171.51 |

The initial TTL of routers is 255

Local Datacenter routers

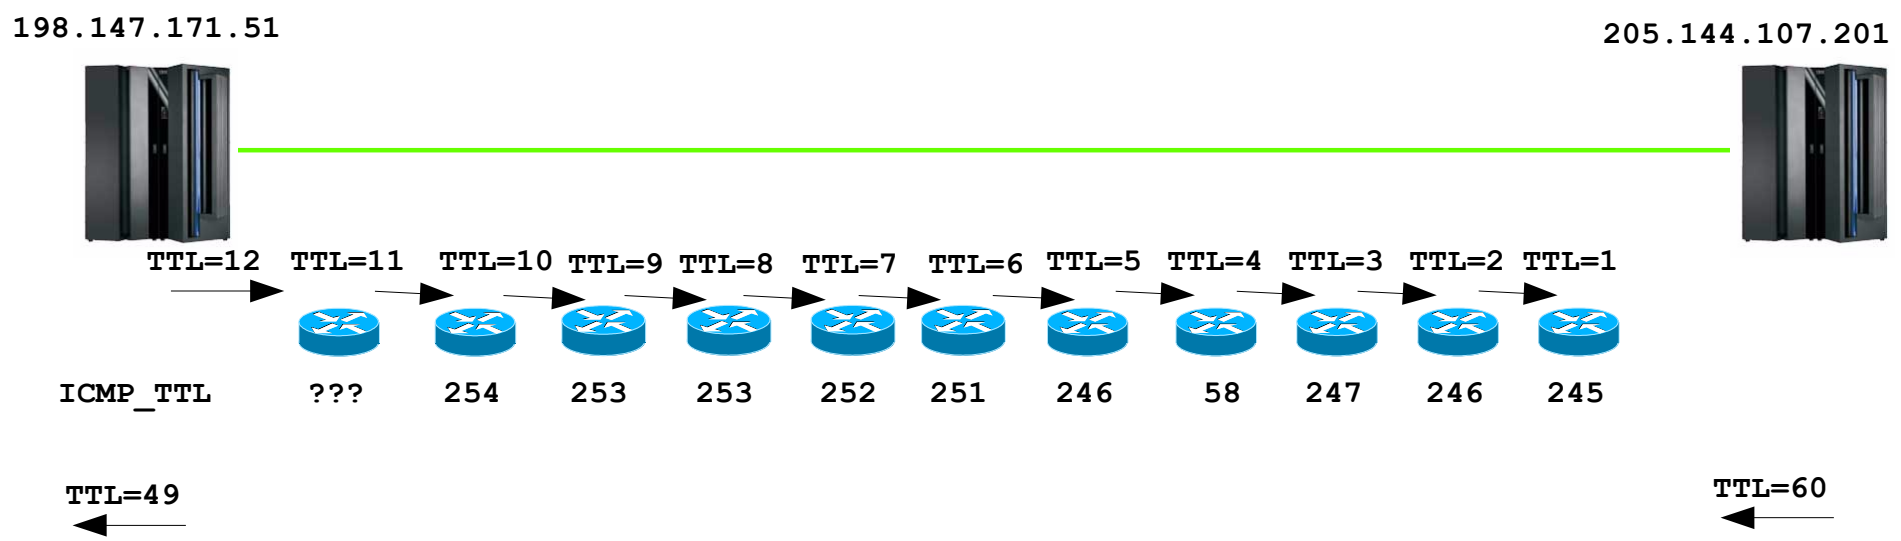
Distant Routers with higher RTT

5 additional hops not seen

Same IP address
Different TTLs, IP.IDs
TTL 58 Not a router
TTL 49 Not a router

TTL and Topology II.

The packet with TTL 12 reached the other end.
 The destination is 11 hops away
 The TEST OK comes in with a TTL of 49
 The initial TTL is 60

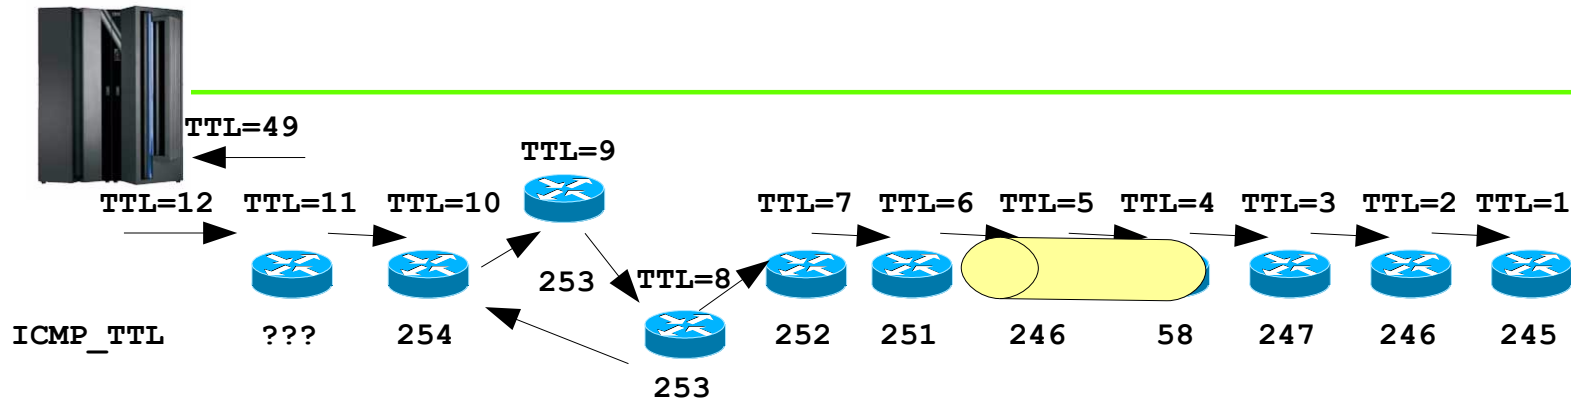


TTL and Topology III.

Looking at the returned TTLs, we can make assumptions as to how the routers are connected.

198.147.171.51

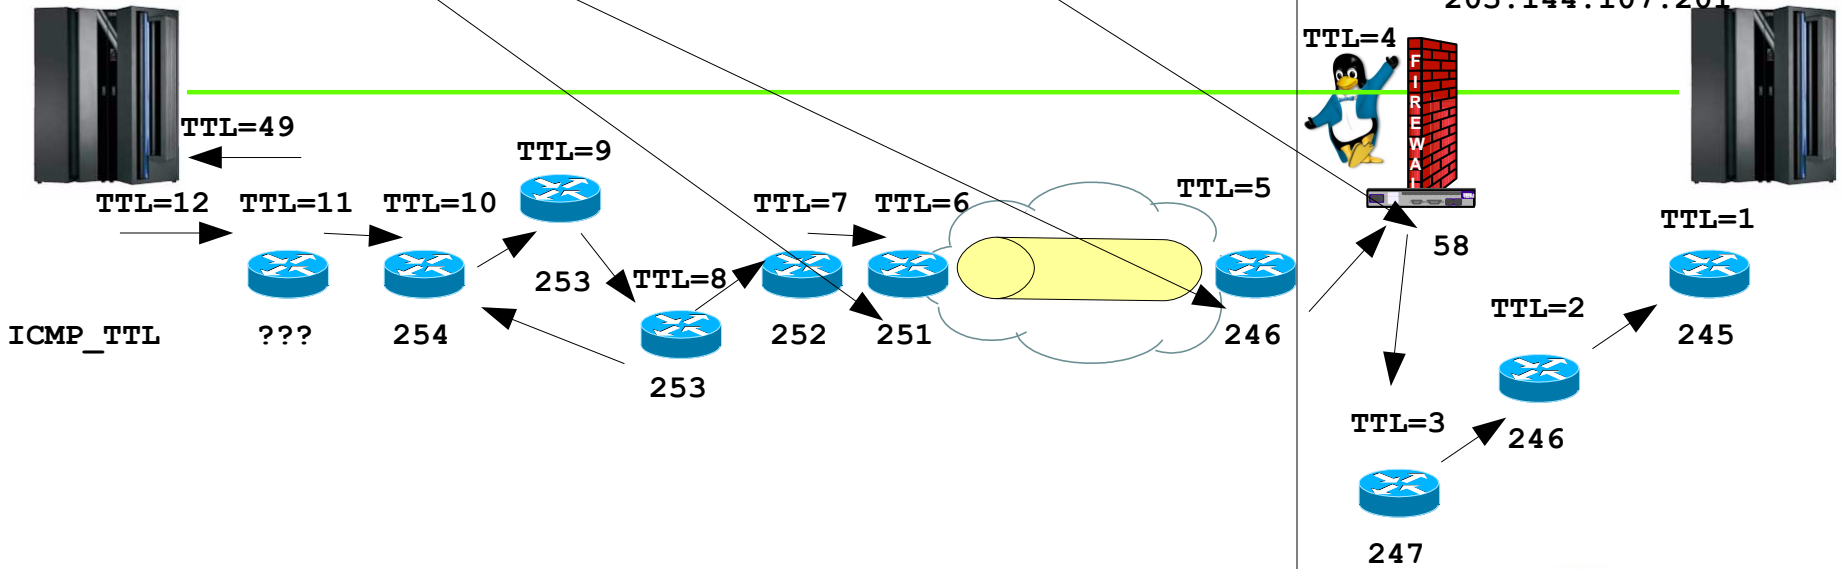
205.144.107.201



TTL and Topology IV.

Routers send with a TTL of 255
 Linux sends with a TTL of 64
 A gap in the TTLs indicates a VPN IPsec tunnel is in the path
 Multiple replies from the same IP address indicate a NAT hiding the original addresses

198.147.171.51



Fragmentation is bad – BAD – BAD



Filter: (ip.flags.mf==1 or ip.frag_offset > 0) && (ip.proto == 17) Expression... Clear Apply

| No. | whazzin . | Time | ip.ttl | ip.id | ip.len | src.addr | dst.addr | dst.port | nhdr |
|-----|---------------|----------|--------|--------|--------|-----------------|----------------|----------|--------------------|
| 445 | Fragmentation | 0.000000 | 48 | 0x6350 | 44 | 205.144.107.201 | 198.147.171.51 | 12003 | d00000000000000000 |
| 446 | Fragmentation | 0.021794 | 48 | 0x6350 | 1494 | 205.144.107.201 | 198.147.171.51 | | |

Internet Protocol Version 4, Src: 205.144.107.201 (205.144.107.201), Dst: 198.147.171.51 (198.147.171.51)

- Version: 4
- Header length: 20 bytes
- Type of service: 0x00 (None)
- Total Length: 44
- Identification: 0x6350 (25424)
- Flags: 0x01 (More Fragments)
- Fragment offset: 0
- Time to live: 48
- Protocol: UDP (17)
- Header checksum: 0x5c50 [validation disabled]
- Source: 205.144.107.201 (205.144.107.201)
- Destination: 198.147.171.51 (198.147.171.51)

User Datagram Protocol, Src Port: 12003 (12003), Dst Port: 12003 (12003)

- Source port: 12003 (12003)
- Destination port: 12003 (12003)
- Length: 1498
- Checksum: 0xd270 [unchecked, not all data available]

Logical-Link Control

Systems Network Architecture

Unassembled Packet: SNAP

```

0000  00 0f a1 00 00 01 00 50 9b 00 00 01 08 00 45 00  ...~...&...
0010  00 2c 63 50 20 00 30 11 5c 50 cd 90 6b c9 c6 93  ...&...&...
0020  ab 33 2e e3 2e e3 05 da d2 70 04 08 03 c2 08 d0  ...3.T.T..K....B.}
0030  00 00 00 00 00 00 00 ff 00 3d                      .....=
  
```

Same ip.id, only first packet has a UDP header! Most Firewalls will drop 2nd packet as it does not match any port filter rule!

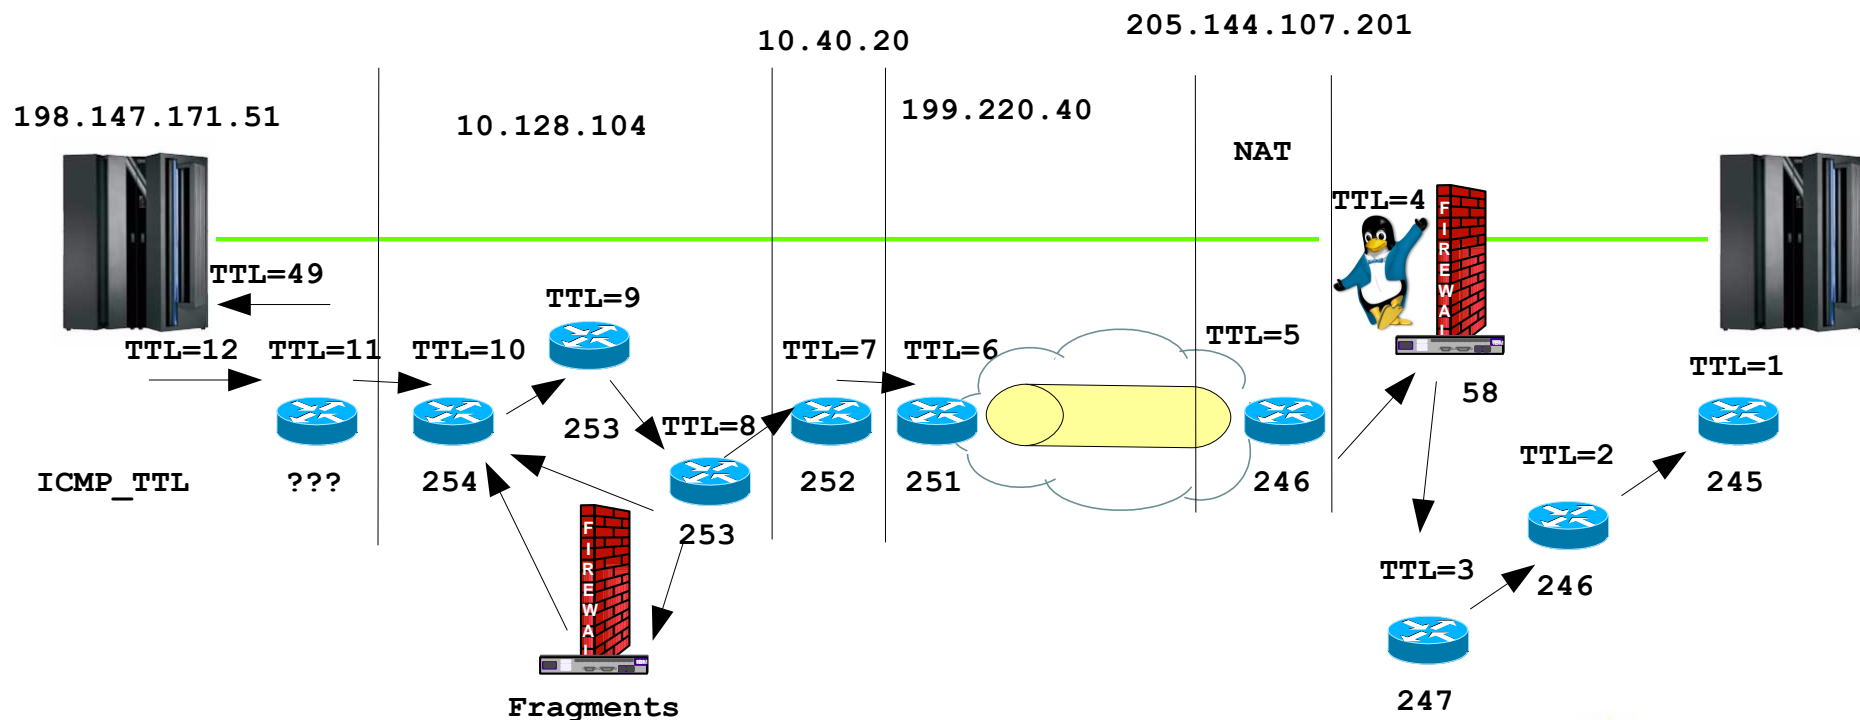
TTL changes to 48!

UDP length is 1498, adding 20 bytes IP header indicates the original packet was 1518 bytes at the sender!

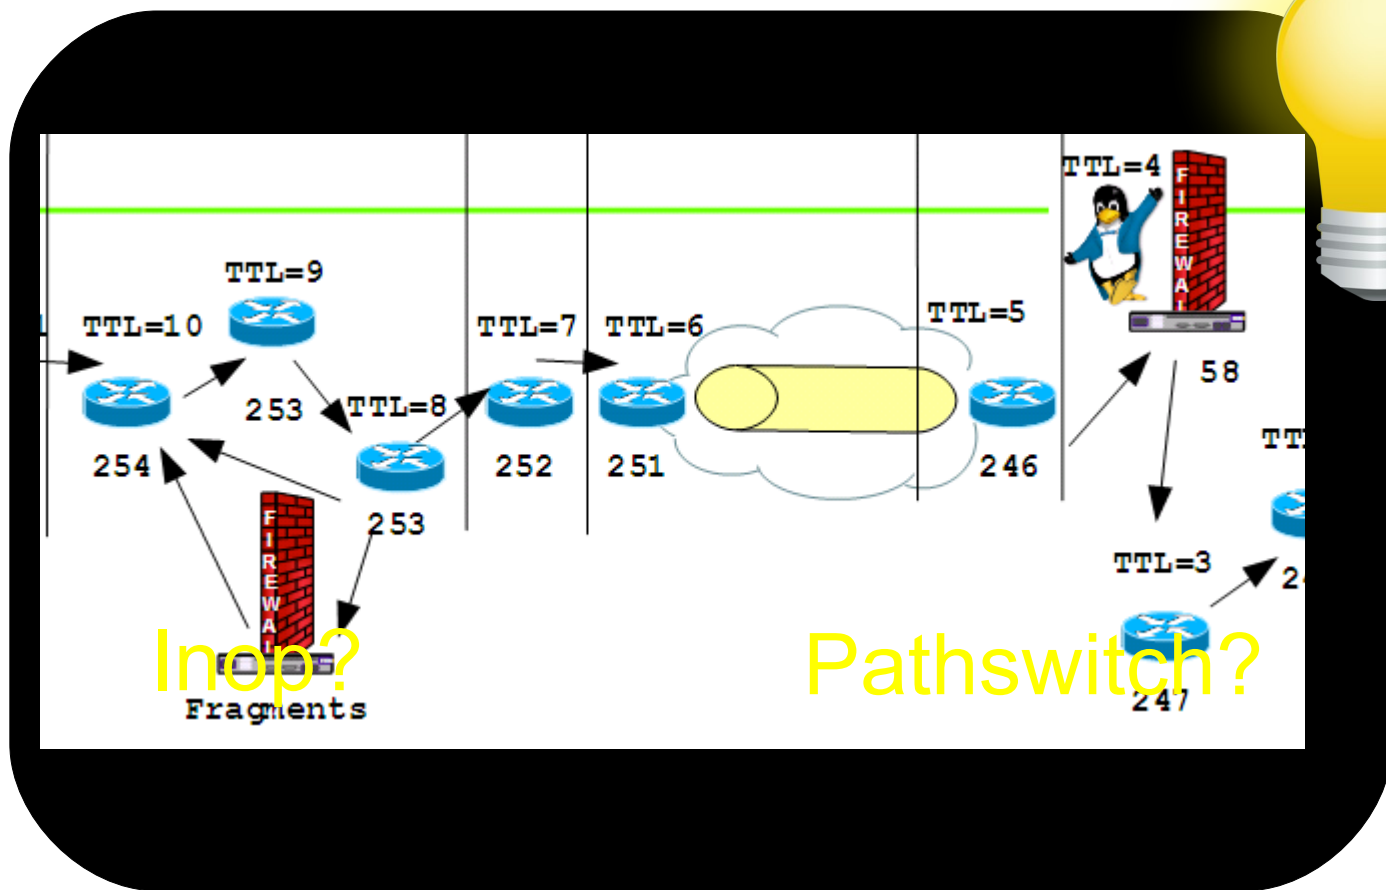


TTL and Topology VI. - Fragmentation

Fragmented IP packets get inspected adding an additional hop to the ip path.



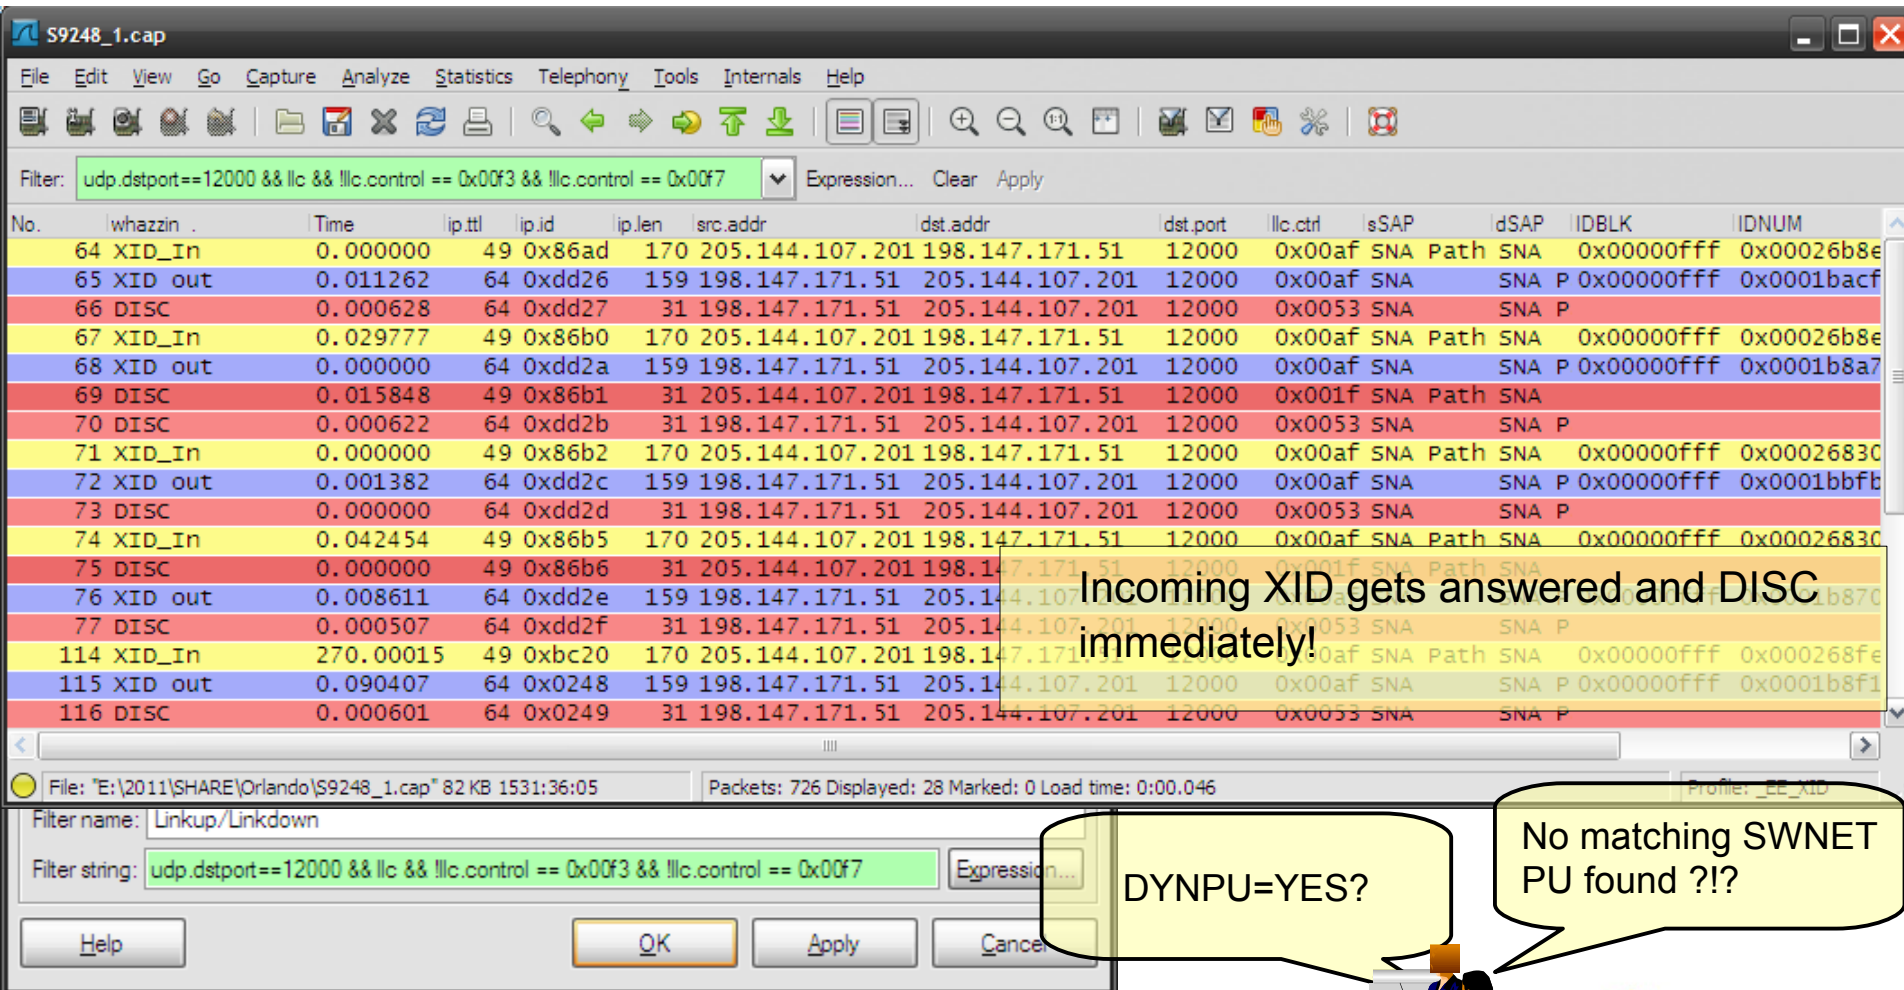
Now we have picture of the environment Time to get started working on the 'problem'



Inop?
Fragments

Pathswitch?

Detecting INOPs with wireshark



Filter: `udp.dstport==12000 && llc && !llc.control == 0x00f3 && !llc.control == 0x00f7`

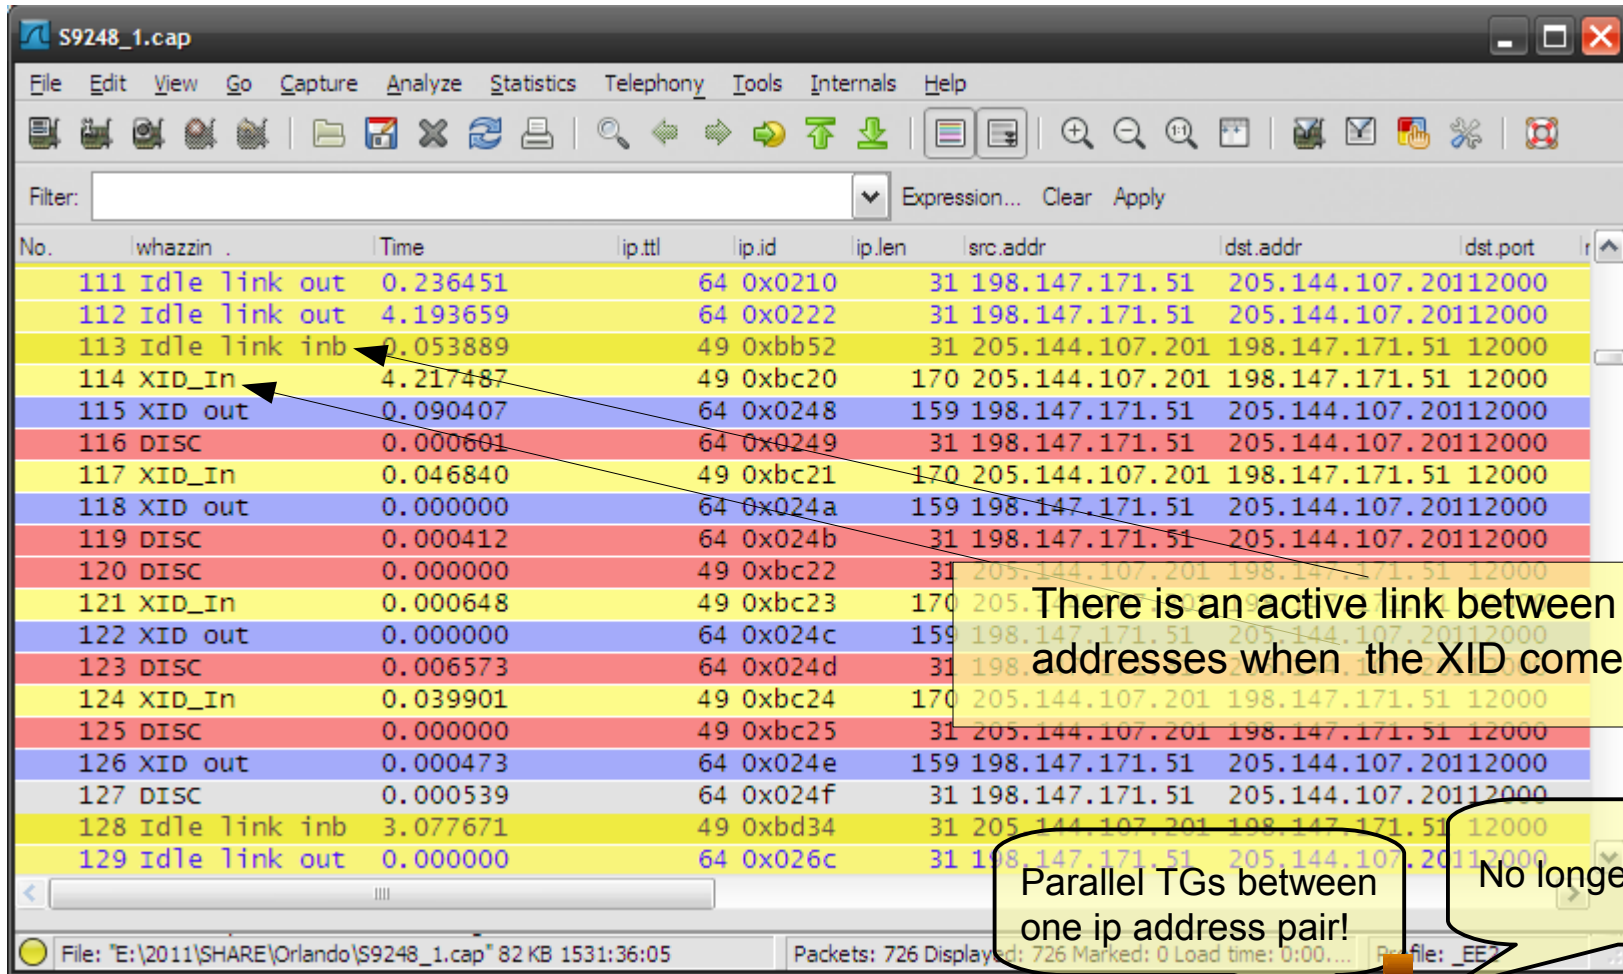
| No. | hwazzin | Time | ip.ttl | ip.id | ip.len | src.addr | dst.addr | dst.port | llc.ctr | sSAP | dSAP | IDBLK | IDNUM |
|-----|---------|-----------|--------|--------|--------|-----------------|-----------------|----------|---------|----------|-------|------------|------------|
| 64 | XID_In | 0.000000 | 49 | 0x86ad | 170 | 205.144.107.201 | 198.147.171.51 | 12000 | 0x00af | SNA Path | SNA | 0x00000fff | 0x00026b8e |
| 65 | XID out | 0.011262 | 64 | 0xdd26 | 159 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x00af | SNA | SNA P | 0x00000fff | 0x0001bacf |
| 66 | DISC | 0.000628 | 64 | 0xdd27 | 31 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x0053 | SNA | SNA P | | |
| 67 | XID_In | 0.029777 | 49 | 0x86b0 | 170 | 205.144.107.201 | 198.147.171.51 | 12000 | 0x00af | SNA Path | SNA | 0x00000fff | 0x00026b8e |
| 68 | XID out | 0.000000 | 64 | 0xdd2a | 159 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x00af | SNA | SNA P | 0x00000fff | 0x0001b8a7 |
| 69 | DISC | 0.015848 | 49 | 0x86b1 | 31 | 205.144.107.201 | 198.147.171.51 | 12000 | 0x001f | SNA Path | SNA | | |
| 70 | DISC | 0.000622 | 64 | 0xdd2b | 31 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x0053 | SNA | SNA P | | |
| 71 | XID_In | 0.000000 | 49 | 0x86b2 | 170 | 205.144.107.201 | 198.147.171.51 | 12000 | 0x00af | SNA Path | SNA | 0x00000fff | 0x00026830 |
| 72 | XID out | 0.001382 | 64 | 0xdd2c | 159 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x00af | SNA | SNA P | 0x00000fff | 0x0001bbfb |
| 73 | DISC | 0.000000 | 64 | 0xdd2d | 31 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x0053 | SNA | SNA P | | |
| 74 | XID_In | 0.042454 | 49 | 0x86b5 | 170 | 205.144.107.201 | 198.147.171.51 | 12000 | 0x00af | SNA Path | SNA | 0x00000fff | 0x00026830 |
| 75 | DISC | 0.000000 | 49 | 0x86b6 | 31 | 205.144.107.201 | 198.147.171.51 | 12000 | 0x001f | SNA Path | SNA | | |
| 76 | XID out | 0.008611 | 64 | 0xdd2e | 159 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x00af | SNA | SNA P | 0x00000fff | 0x0001b870 |
| 77 | DISC | 0.000507 | 64 | 0xdd2f | 31 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x0053 | SNA | SNA P | | |
| 114 | XID_In | 270.00015 | 49 | 0xbc20 | 170 | 205.144.107.201 | 198.147.171.51 | 12000 | 0x00af | SNA Path | SNA | 0x00000fff | 0x000268fe |
| 115 | XID out | 0.090407 | 64 | 0x0248 | 159 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x00af | SNA | SNA P | 0x00000fff | 0x0001b8f1 |
| 116 | DISC | 0.000601 | 64 | 0x0249 | 31 | 198.147.171.51 | 205.144.107.201 | 12000 | 0x0053 | SNA | SNA P | | |

Filter name: Linkup/Linkdown
Filter string: `udp.dstport==12000 && llc && !llc.control == 0x00f3 && !llc.control == 0x00f7`

DYNPU=YES?

No matching SWNET PU found ??!

Active link – why a new XID?



| No. | whazzin | Time | ip.ttl | ip.id | ip.len | src.addr | dst.addr | dst.port |
|-----|---------------|----------|--------|--------|--------|-----------------|-----------------|----------|
| 111 | idle link out | 0.236451 | 64 | 0x0210 | 31 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 112 | idle link out | 4.193659 | 64 | 0x0222 | 31 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 113 | idle link inb | 0.053889 | 49 | 0xbb52 | 31 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 114 | XID_In | 4.217487 | 49 | 0xbc20 | 170 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 115 | XID out | 0.090407 | 64 | 0x0248 | 159 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 116 | DISC | 0.000601 | 64 | 0x0249 | 31 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 117 | XID_In | 0.046840 | 49 | 0xbc21 | 170 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 118 | XID out | 0.000000 | 64 | 0x024a | 159 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 119 | DISC | 0.000412 | 64 | 0x024b | 31 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 120 | DISC | 0.000000 | 49 | 0xbc22 | 31 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 121 | XID_In | 0.000648 | 49 | 0xbc23 | 170 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 122 | XID out | 0.000000 | 64 | 0x024c | 159 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 123 | DISC | 0.006573 | 64 | 0x024d | 31 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 124 | XID_In | 0.039901 | 49 | 0xbc24 | 170 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 125 | DISC | 0.000000 | 49 | 0xbc25 | 31 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 126 | XID out | 0.000473 | 64 | 0x024e | 159 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 127 | DISC | 0.000539 | 64 | 0x024f | 31 | 198.147.171.51 | 205.144.107.201 | 12000 |
| 128 | idle link inb | 3.077671 | 49 | 0xbd34 | 31 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 129 | idle link out | 0.000000 | 64 | 0x026c | 31 | 198.147.171.51 | 205.144.107.201 | 12000 |

There is an active link between the two ip addresses when the XID comes in.

Parallel TGs between one ip address pair!

No longer supported!



Now we have picture of the environment Time to get started working on the 'problem'

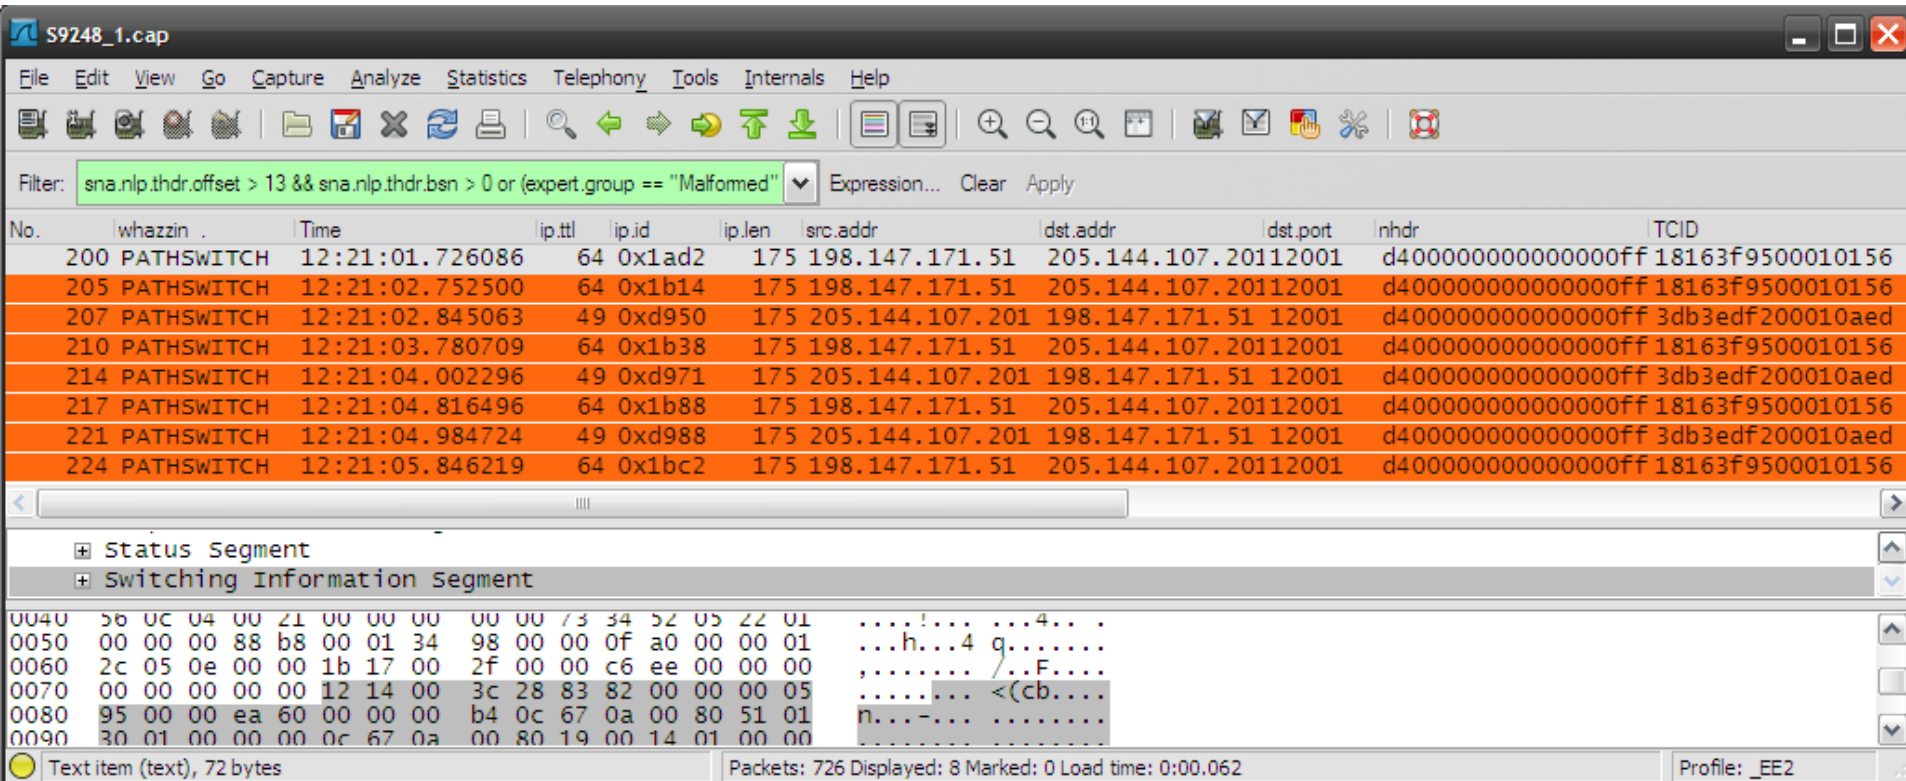
So, what is your problem?



~~In?~~

Pathswitch?

How to find switching pipes



Filter: `sna.nlp.thdr.offset > 13 && sna.nlp.thdr.bsn > 0 or (expert.group == "Malformed")`

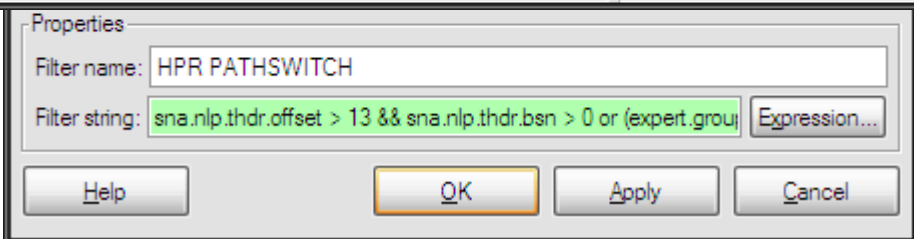
| No. | whazzin . | Time | ip.ttl | ip.id | ip.len | src.addr | dst.addr | dst.port | nhdr | TCID |
|-----|------------|-----------------|--------|--------|--------|-----------------|-----------------|----------|--------------------|------------------|
| 200 | PATHSWITCH | 12:21:01.726086 | 64 | 0x1ad2 | 175 | 198.147.171.51 | 205.144.107.201 | 12001 | d400000000000000ff | 18163f9500010156 |
| 205 | PATHSWITCH | 12:21:02.752500 | 64 | 0x1b14 | 175 | 198.147.171.51 | 205.144.107.201 | 12001 | d400000000000000ff | 18163f9500010156 |
| 207 | PATHSWITCH | 12:21:02.845063 | 49 | 0xd950 | 175 | 205.144.107.201 | 198.147.171.51 | 12001 | d400000000000000ff | 3db3edf200010aed |
| 210 | PATHSWITCH | 12:21:03.780709 | 64 | 0x1b38 | 175 | 198.147.171.51 | 205.144.107.201 | 12001 | d400000000000000ff | 18163f9500010156 |
| 214 | PATHSWITCH | 12:21:04.002296 | 49 | 0xd971 | 175 | 205.144.107.201 | 198.147.171.51 | 12001 | d400000000000000ff | 3db3edf200010aed |
| 217 | PATHSWITCH | 12:21:04.816496 | 64 | 0x1b88 | 175 | 198.147.171.51 | 205.144.107.201 | 12001 | d400000000000000ff | 18163f9500010156 |
| 221 | PATHSWITCH | 12:21:04.984724 | 49 | 0xd988 | 175 | 205.144.107.201 | 198.147.171.51 | 12001 | d400000000000000ff | 3db3edf200010aed |
| 224 | PATHSWITCH | 12:21:05.846219 | 64 | 0x1bc2 | 175 | 198.147.171.51 | 205.144.107.201 | 12001 | d400000000000000ff | 18163f9500010156 |

Expert view: Status Segment, Switching Information segment

```

0040 56 0c 04 00 21 00 00 00 00 00 73 34 52 05 22 01  ....!...  ...4...
0050 00 00 00 88 b8 00 01 34 98 00 00 0f a0 00 00 01  ...h...4 q.....
0060 2c 05 0e 00 00 00 1b 17 00 2f 00 00 c6 ee 00 00 00  ....  /...F....
0070 00 00 00 00 00 12 14 00 3c 28 83 82 00 00 00 05  ....  <(cb...
0080 95 00 00 ea 60 00 00 00 b4 0c 67 0a 00 80 51 01  n...-... ..
0090 30 01 00 00 00 0c 67 0a 00 80 19 00 14 01 00 00
  
```

Text item (text), 72 bytes | Packets: 726 Displayed: 8 Marked: 0 Load time: 0:00.062 | Profile: _EE2

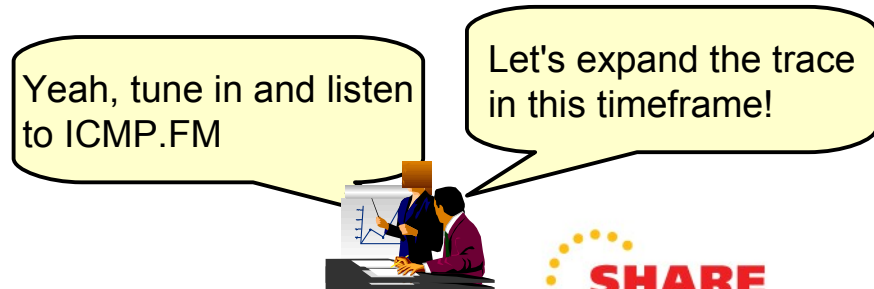


Properties

Filter name: HPR PATHSWITCH

Filter string: `sna.nlp.thdr.offset > 13 && sna.nlp.thdr.bsn > 0 or (expert.group == "Malformed")`

Buttons: Help, OK, Apply, Cancel



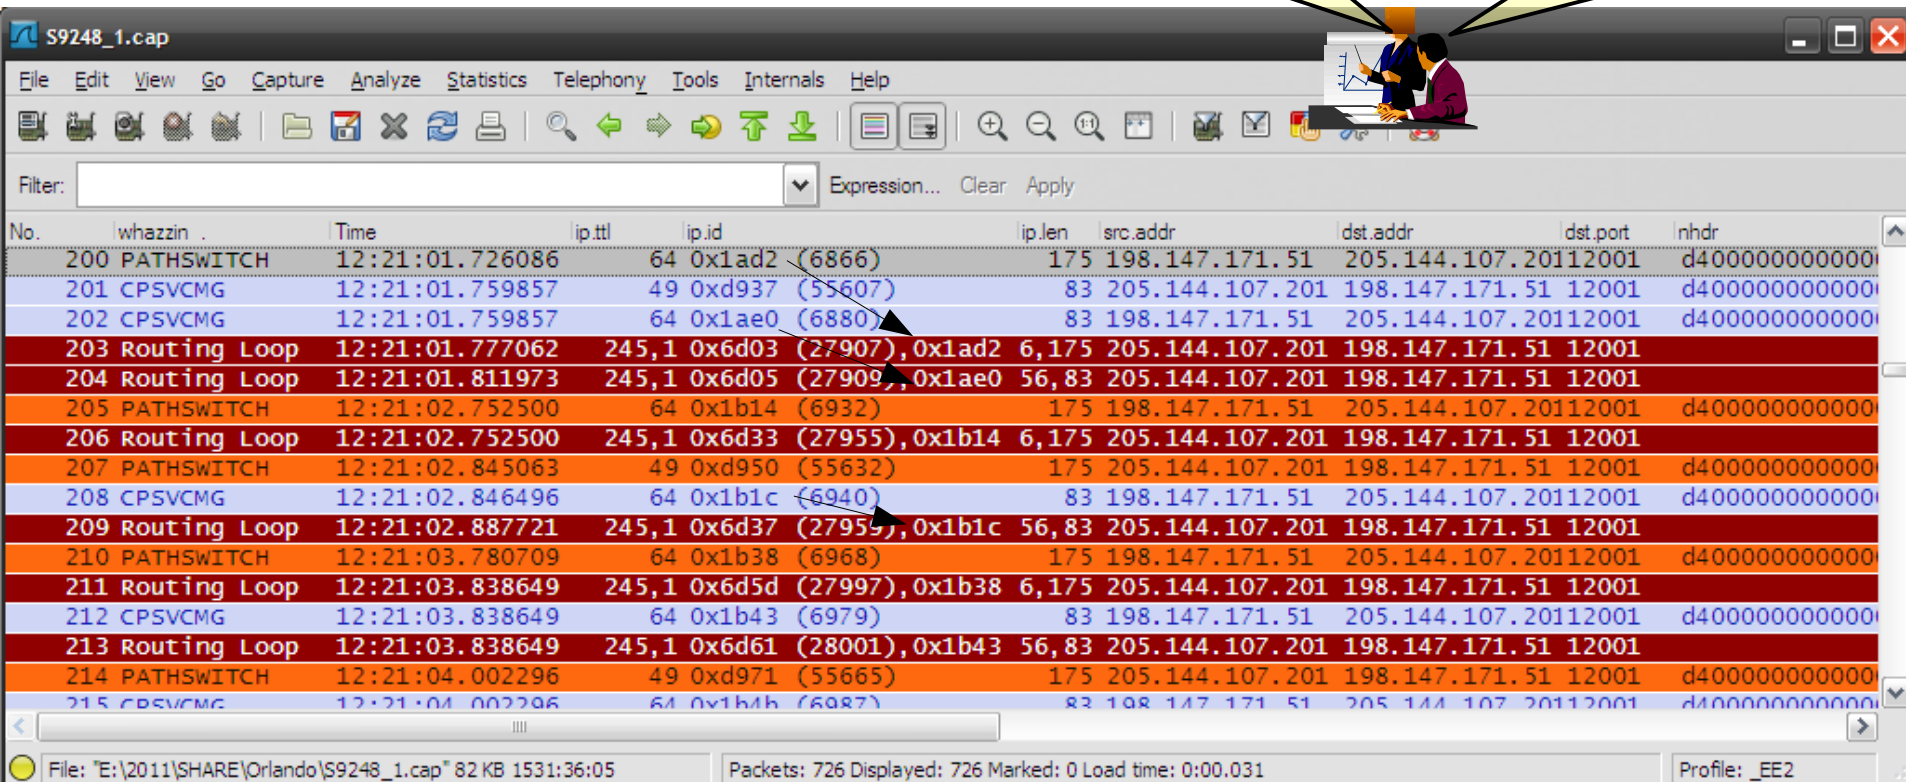
Yeah, tune in and listen to ICMP.FM

Let's expand the trace in this timeframe!

PATHSWITCH due to routing loop

Yes, if they don't make it to the remote RTP a PATHSWITCH is the logical consequence

Our outbound NLPs die in a routing loop!



| No. | whazzin . | Time | ip.ttl | ip.id | ip.len | src.addr | dst.addr | dst.port | nhdr |
|-----|--------------|-----------------|--------|------------------------|--------|-----------------|-----------------|----------|----------------|
| 200 | PATHSWITCH | 12:21:01.726086 | 64 | 0x1ad2 (6866) | 175 | 198.147.171.51 | 205.144.107.201 | 12001 | d4000000000000 |
| 201 | CPSVCMG | 12:21:01.759857 | 49 | 0xd937 (55607) | 83 | 205.144.107.201 | 198.147.171.51 | 12001 | d4000000000000 |
| 202 | CPSVCMG | 12:21:01.759857 | 64 | 0x1ae0 (6880) | 83 | 198.147.171.51 | 205.144.107.201 | 12001 | d4000000000000 |
| 203 | Routing Loop | 12:21:01.777062 | 245,1 | 0x6d03 (27907), 0x1ad2 | 6,175 | 205.144.107.201 | 198.147.171.51 | 12001 | |
| 204 | Routing Loop | 12:21:01.811973 | 245,1 | 0x6d05 (27909), 0x1ae0 | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 | |
| 205 | PATHSWITCH | 12:21:02.752500 | 64 | 0x1b14 (6932) | 175 | 198.147.171.51 | 205.144.107.201 | 12001 | d4000000000000 |
| 206 | Routing Loop | 12:21:02.752500 | 245,1 | 0x6d33 (27955), 0x1b14 | 6,175 | 205.144.107.201 | 198.147.171.51 | 12001 | |
| 207 | PATHSWITCH | 12:21:02.845063 | 49 | 0xd950 (55632) | 175 | 205.144.107.201 | 198.147.171.51 | 12001 | d4000000000000 |
| 208 | CPSVCMG | 12:21:02.846496 | 64 | 0x1b1c (6940) | 83 | 198.147.171.51 | 205.144.107.201 | 12001 | d4000000000000 |
| 209 | Routing Loop | 12:21:02.887721 | 245,1 | 0x6d37 (27957), 0x1b1c | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 | |
| 210 | PATHSWITCH | 12:21:03.780709 | 64 | 0x1b38 (6968) | 175 | 198.147.171.51 | 205.144.107.201 | 12001 | d4000000000000 |
| 211 | Routing Loop | 12:21:03.838649 | 245,1 | 0x6d5d (27997), 0x1b38 | 6,175 | 205.144.107.201 | 198.147.171.51 | 12001 | |
| 212 | CPSVCMG | 12:21:03.838649 | 64 | 0x1b43 (6979) | 83 | 198.147.171.51 | 205.144.107.201 | 12001 | d4000000000000 |
| 213 | Routing Loop | 12:21:03.838649 | 245,1 | 0x6d61 (28001), 0x1b43 | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 | |
| 214 | PATHSWITCH | 12:21:04.002296 | 49 | 0xd971 (55665) | 175 | 205.144.107.201 | 198.147.171.51 | 12001 | d4000000000000 |
| 215 | CPSVCMG | 12:21:04.002296 | 64 | 0x1b4b (6987) | 83 | 198.147.171.51 | 205.144.107.201 | 12001 | d4000000000000 |

Routing Loop: TTL exceeded

Filter: `ip.ttl<10 and udp and !udp.length==56`

| No. | whazzin . | Time | ip.ttl | ip.id | ip.len | src_addr | dst_addr | dst_port |
|-----|--------------|-----------------|--------|------------------------|--------|-----------------|----------------|----------|
| 59 | Routing Loop | 12:13:17.725312 | 246,1 | 0x256d (9581), 0xdba4 | (56,31 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 61 | Routing Loop | 12:13:21.893922 | 246,1 | 0x25fd (9725), 0xdc31 | (56,31 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 107 | Routing Loop | 12:17:36.060766 | 245,1 | 0x64ce (25806), 0x00c6 | 56,31 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 182 | Routing Loop | 12:20:58.315697 | 245,1 | 0x6c5e (27742), 0x1a09 | 56,95 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 184 | Routing Loop | 12:20:58.349276 | 245,1 | 0x6c5f (27743), 0x1a0d | 6,464 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 186 | Routing Loop | 12:20:59.230052 | 245,1 | 0x6c80 (27776), 0x1a46 | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 189 | Routing Loop | 12:20:59.394710 | 245,1 | 0x6c8e (27790), 0x1a69 | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 192 | Routing Loop | 12:20:59.526941 | 245,1 | 0x6c9b (27803), 0x1a6b | 56,31 | 205.144.107.201 | 198.147.171.51 | 12000 |
| 194 | Routing Loop | 12:21:00.055715 | 245,1 | 0x6cb4 (27828), 0x1a79 | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 197 | Routing Loop | 12:21:00.587253 | 245,1 | 0x6ccd (27853), 0x1a9a | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 199 | Routing Loop | 12:21:00.864111 | 245,1 | 0x6ce1 (27873), 0x1ab1 | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 203 | Routing Loop | 12:21:01.777062 | 245,1 | 0x6d03 (27907), 0x1ad2 | 6,175 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 204 | Routing Loop | 12:21:01.811973 | 245,1 | 0x6d05 (27909), 0x1ae0 | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 206 | Routing Loop | 12:21:02.752500 | 245,1 | 0x6d33 (27955), 0x1b14 | 6,175 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 209 | Routing Loop | 12:21:02.887721 | 245,1 | 0x6d37 (27959), 0x1b1c | 56,83 | 205.144.107.201 | 198.147.171.51 | 12001 |
| 211 | Routing Loop | 12:21:03.838649 | 245,1 | 0x6d5d (27997), 0x1b38 | 6,175 | 205.144.107.201 | 198.147.171.51 | 12001 |

The ICMP message comes in with a TTL of 245

Who's that? Let's check our picture!

```

Internet Protocol Version 4, Src: 205.144.107.201 (205.144.107.201), Dst: 198.147.171.51 (198.147.171.51)
Internet Control Message Protocol
  Type: 11 (Time-to-live exceeded)
  Code: 0 (Time to live exceeded in transit)
  checksum: 0x3d6d [correct]
Internet Protocol Version 4, Src: 198.147.171.51 (198.147.171.51), Dst: 205.144.107.201 (205.144.107.201)
User Datagram Protocol, Src Port: 12001 (12001), Dst Port: 12001 (12001)
  
```



PATHSWITCH due to routing loop

Where is it?

So, here is your problem!



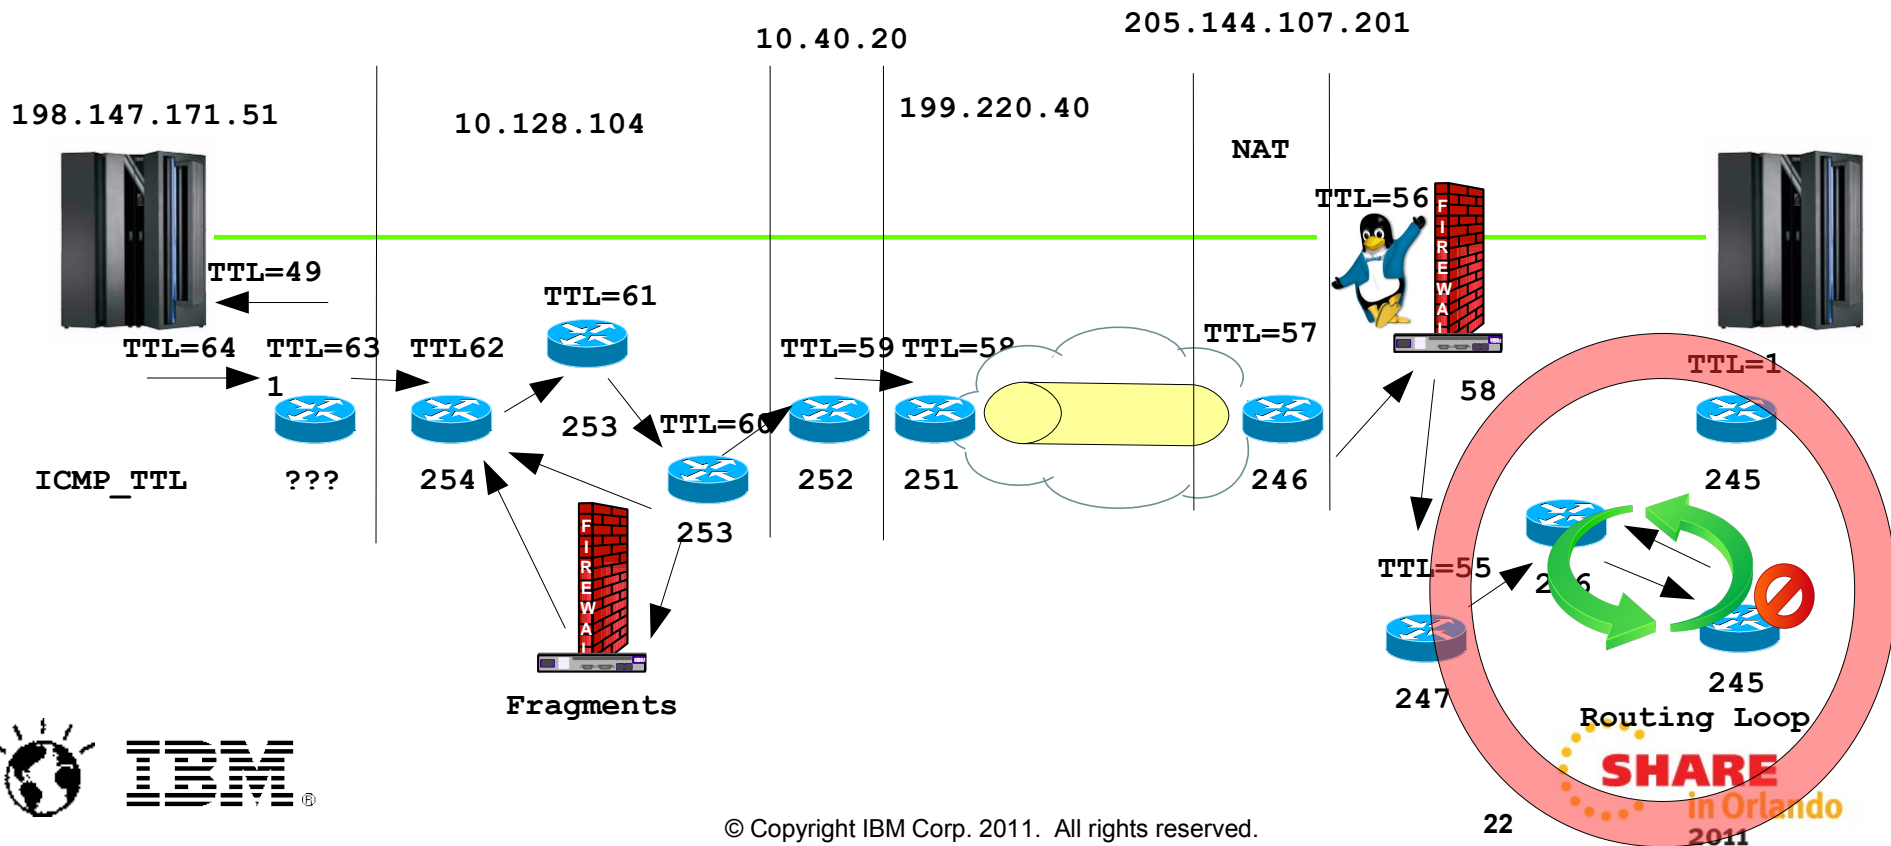
In?


Pathswitch!



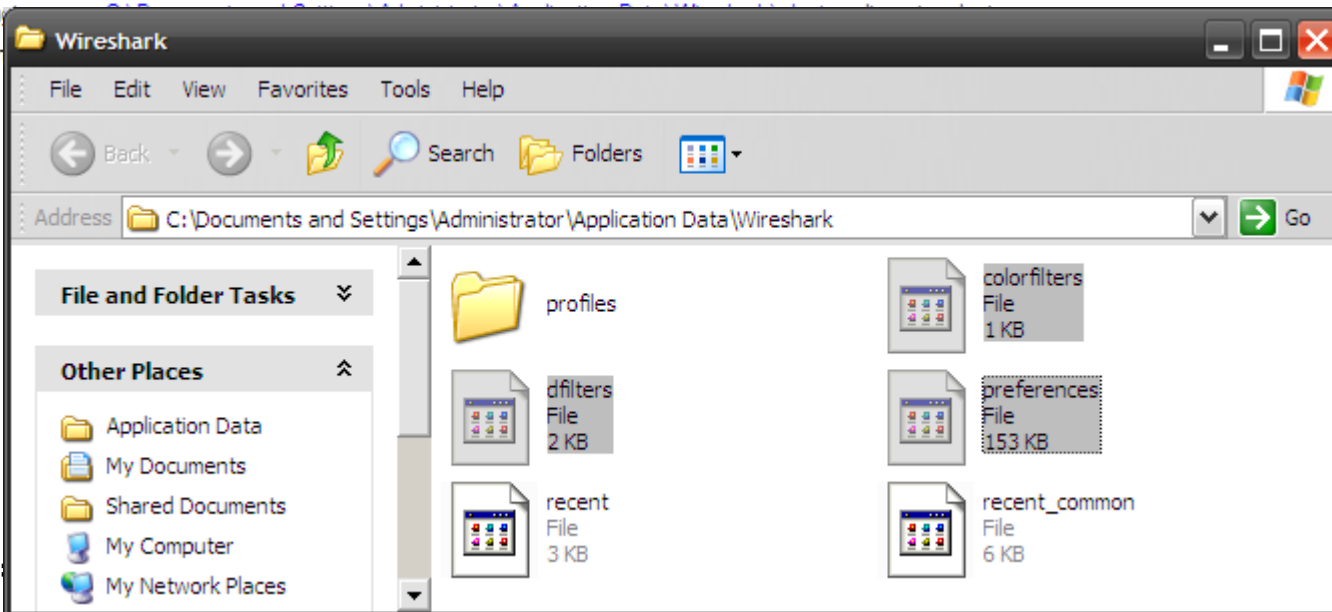
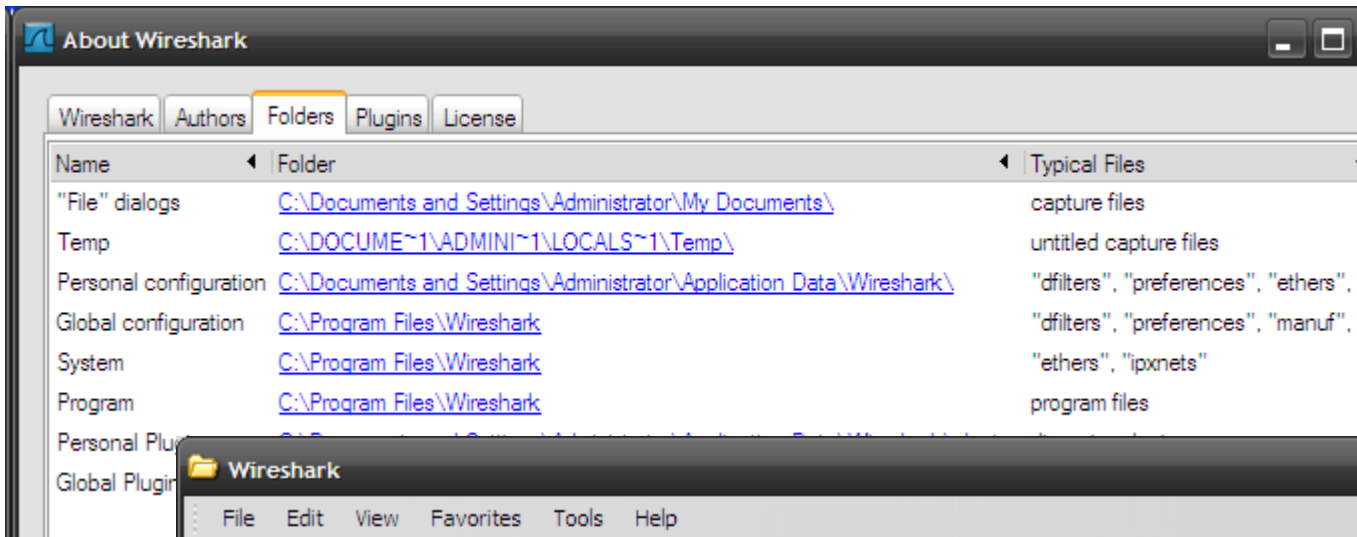
TTL and Topology VII. - Here's the problem

A routing problem at the remote end is causing our NLPs to be discarded



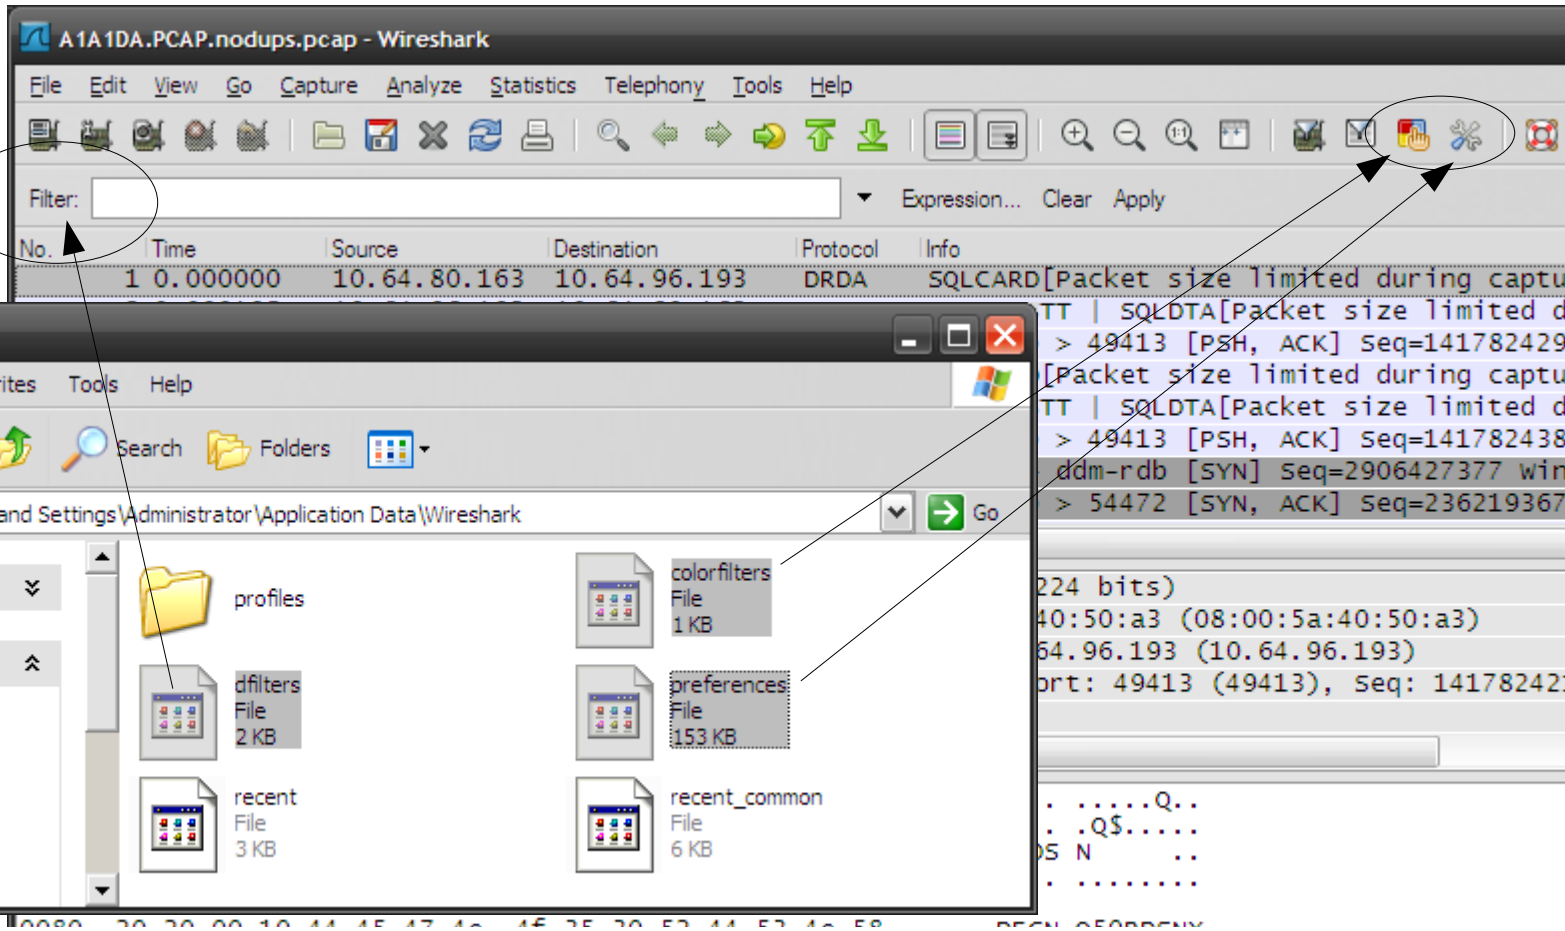
Wireshark

Personal Configuration Files - Profiles



Wireshark

Personal Configuration Files - Profiles



Questions

- IP Wizards on Facebook



IP Wizards

ip.wizards@groups.facebook.com

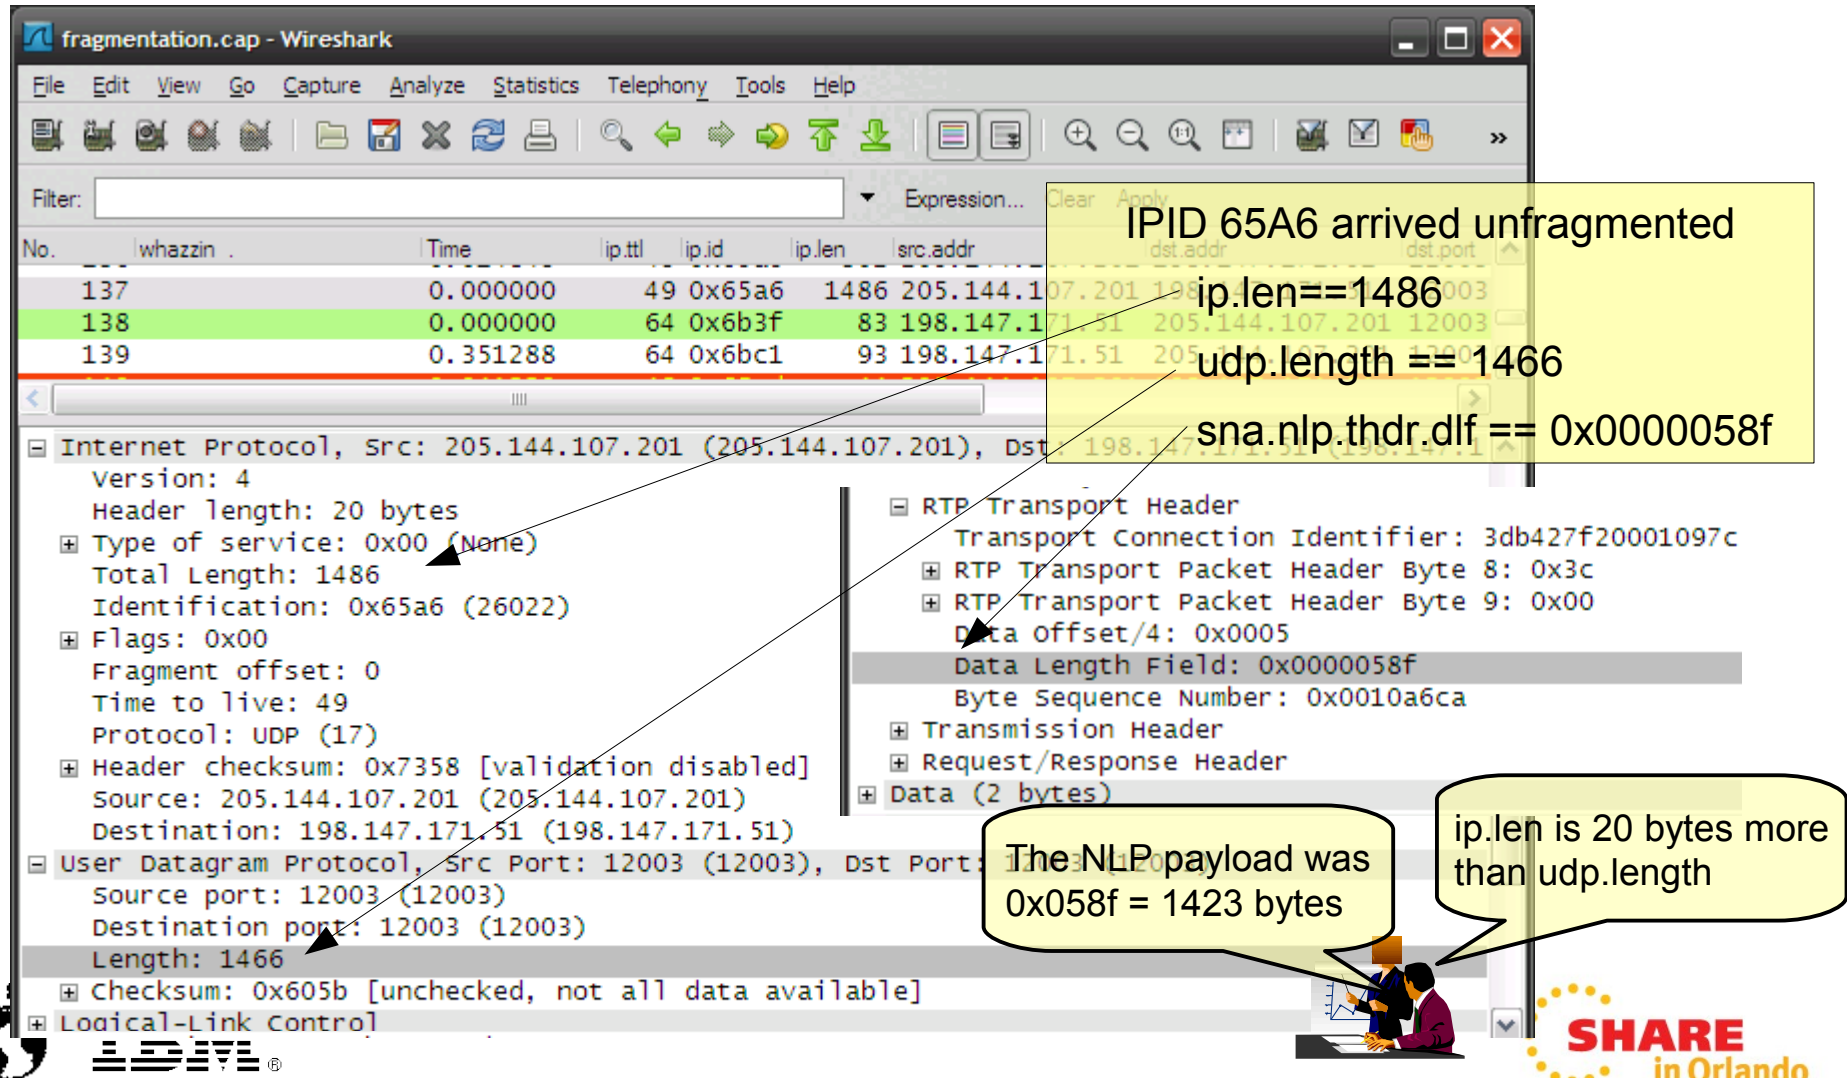
- Wireshark Bootcamp 2011
 - Germany: <http://tinyurl.com/ZOWIE0DE>
 - Canada : <http://tinyurl.com/ZOWIE0CE>

Appendix

- IP Fragmentation

Fragmentation: Why ? – Part I.

An unfragmented packet arrives



The image shows a Wireshark capture window titled "fragmentation.cap - Wireshark". The packet list pane shows three packets. Packet 138 is highlighted in green and has the following details:

| No. | whazzin . | Time | ip.ttl | ip.id | ip.len | src.addr |
|-----|-----------|----------|--------|--------|--------|-----------------|
| 137 | | 0.000000 | 49 | 0x65a6 | 1486 | 205.144.107.201 |
| 138 | | 0.000000 | 64 | 0x6b3f | 83 | 198.147.171.51 |
| 139 | | 0.351288 | 64 | 0x6bc1 | 93 | 198.147.171.51 |

The packet details pane for packet 138 shows the following structure:

- Internet Protocol, Src: 205.144.107.201 (205.144.107.201), Dst: 198.147.171.51 (198.147.171.51)
 - Version: 4
 - Header length: 20 bytes
 - Type of service: 0x00 (None)
 - Total Length: 1486
 - Identification: 0x65a6 (26022)
 - Flags: 0x00
 - Fragment offset: 0
 - Time to live: 49
 - Protocol: UDP (17)
 - Header checksum: 0x7358 [validation disabled]
 - Source: 205.144.107.201 (205.144.107.201)
 - Destination: 198.147.171.51 (198.147.171.51)
- User Datagram Protocol, Src Port: 12003 (12003), Dst Port: 12003 (12003)
 - Source port: 12003 (12003)
 - Destination port: 12003 (12003)
 - Length: 1466
 - Checksum: 0x605b [unchecked, not all data available]
 - Logical-Link Control
- RTP Transport Header
 - Transport Connection Identifier: 3db427f20001097c
 - RTP Transport Packet Header Byte 8: 0x3c
 - RTP Transport Packet Header Byte 9: 0x00
 - Data Offset/4: 0x0005
 - Data Length Field: 0x0000058f
 - Byte Sequence Number: 0x0010a6ca
- Transmission Header
- Request/Response Header
- Data (2 bytes)

Annotations and callouts:

- A yellow box highlights the packet list entry for packet 138 with the text: "IPID 65A6 arrived unfragmented".
- Arrows point from this box to the "ip.len == 1486" and "snr.nlp.thdr.dlf == 0x0000058f" fields in the details pane.
- Another yellow box highlights the "Data Length Field: 0x0000058f" field in the RTP Transport Header details pane with the text: "The NLP payload was 0x058f = 1423 bytes".
- A third yellow box highlights the "Length: 1466" field in the User Datagram Protocol details pane with the text: "ip.len is 20 bytes more than udp.length".

Fragmentation: Why ? – Part II.

What was the original size of the packet?

| No. | whazzin . | Time | ip.ttl | ip.id | ip.len | src.addr | dst.addr | dst.port |
|-----|-----------|----------|--------|--------|--------|-----------------|----------------|----------|
| 140 | | 0.000000 | 48 | 0x65ad | 44 | 205.144.107.201 | 198.147.171.51 | 12003 |

- Type of service: 0x00 (None)
- Total Length: 44
- Identification: 0x65ad (26029)
- Flags: 0x01 (More Fragments)
- Fragment offset: 0
- Time to live: 48
- Protocol: UDP (17)
- Header checksum: 0x59f3 [validation disabled]
- Source: 205.144.107.201 (205.144.107.201)
- Destination: 198.147.171.51 (198.147.171.51)
- User Datagram Protocol, Src Port: 12003 (12003), Dst Port: 12003 (12003)
- Source port: 12003 (12003)
- Destination port: 12003 (12003)
- Length: 1498
- Checksum: 0x128d [unchecked, not all data available]
- Logical-Link Control
- Systems Network Architecture
- Network Layer Packet Header
- Network Layer Packet Header Byte 0: 0xc2
- Network Layer Packet Header Byte 1: 0x08
- Automatic Network Routing Entry: d000000000000000ff
- Reserved

| | | | | | | | | | | | | | | | | | |
|------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-------------|--------------------|
| 0000 | 00 | 0f | a1 | 00 | 04 | 00 | 50 | 9b | 00 | 00 | 04 | 08 | 00 | 45 | 00 |&..... | |
| 0010 | 00 | 2c | 65 | ad | 20 | 00 | 30 | 11 | 59 | f3 | cd | 90 | 6b | c9 | c6 | 93 | ... [... .3., IF] |
| 0020 | ab | 33 | 2e | e3 | 2e | e3 | 05 | da | 12 | 8d | 04 | 08 | 03 | c2 | 08 | d0 | .3.T.T... ..B.] |
| 0030 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | ff | 00 | 3d | | | | | | | = |

IP Header 1st fragment

- ip.len==44
- ip.id == 0x65ad
- ip.flags.mf == 1

UDP Header

- udp.length == 1498

ANR Header

- sna.nlp.nhdr.anr == d0:00:00:00:00:00:00:00:ff

RTP THDR:

- sna.nlp.thdr.tcid == 3d:?:?:?:?:?:?:?:?:?:?

ip,len at the sender must have been 1498+20 bytes!



Fragmentation: Why ? – Part III.

What was the original DLF of the NLP?



```

+ Internet Protocol, Src: 205.144.107.201 (205.144.107.201), Dst: 198.147.
- User Datagram Protocol, Src Port: 12003 (12003), Dst Port: 12003 (12003)
  Source port: 12003 (12003)
  Destination port: 12003 (12003)
  Length: 1498
  + Checksum: 0x128d [unchecked, not all data available]
+ Logical-Link Control
- Systems Network Architecture
  + Network Layer Packet Header
  + [Unreassembled Packet: SNA]
  
```

```

0000 00 0f a1 00 00 04 00 50 9b 00 00 04 08 00 45 00
0010 00 2c 65 ad 20 00 30 11 59 f3 cd 90 6b c9 c6 93
0020 ab 33 2e e3 2e e3 05 da 12 8d 04 08 03 c2 08 d0
0030 00 00 00 00 00 00 00 ff 00 3d
  
```

```

+ Type of service: 0x00 (None)
  Total Length: 1494
  Identification: 0x65ad (26029)
+ Flags: 0x00
  Fragment offset: 24
  Time to live: 48
  Protocol: UDP (17)
+ Header checksum: 0x7446 [validation disabled]
  Source: 205.144.107.201 (205.144.107.201)
  Destination: 198.147.171.51 (198.147.171.51)
  
```

+ Data (60 bytes)

```

0000 00 0f a1 00 00 04 00 50 9b 00 00 04 08 00 45 00
0010 05 d6 65 ad 00 03 30 11 74 46 cd 90 6b c9 c6 93
0020 ab 33 b4 27 f2 00 01 09 7c 3c 04 00 0d 00 00 05
0030 8f 00 10 ac 5a 03 22 c5 58 00 06 db 80 00 00 00
0040 00 05 0e 00 00 01 00 02 00 00 01 a7 00 00 00
0050 00 00 00 00 00 5c 00 03 02 00 00 00 00 34
  
```

IP Header: 2nd fragment

ip.len == 1494

ip.id == 0x65ad

RTP THDR:

sna.nlp.thdr.teid == 3d:b4:27:f2:00:01:09:7c

sna.nlp.thdr.offset == 0x000d

sna.nlp.thdr.dlf == 0x0000058f

sna.nlp.thdr.bsn == 0x0010ac5a

Optional Segments

sna.nlp.thdr.optional.type == 0x22

sna.nlp.thdr.optional.type == 0x0e

NLP payload was 0x058f
Same as before!
But additional segments!