


Layer-transcending OAM requirements

draft-nordmark-nvo3-transcending-traceroute-00

Erik Nordmark

nordmark@arista.com

Three Main Points



1. Traceroute shows the path even though the user can't control the path in any way
2. Overlay networks might hide underlay path
 - Useful policy in some/many cases
 - Separate policy and mechanism?
3. Easy to rectify - details specified for VXLAN

Bonus: Detect/report PMTU misconfiguration

Revealing the path



```
bash$ traceroute www.dn.se
traceroute to a1910.g1.akamai.net (63.150.12.17), 64 hops max, 52 byte packets
 1  cs2-wifi-epool-vl1070.aristanetworks.com (172.22.227.3)  1.717 ms  6.007 ms  2.755 ms
 2  us-ca-scl-paf001-vl3316.aristanetworks.com (172.22.199.45)  2.118 ms  2.269 ms  2.339 ms
 3  us-ca01-01-sw7124-01-vl550.aristanetworks.com (162.210.130.1)  2.260 ms  3.126 ms  2.697
ms
 4  10ge8-4.core3.fmt2.he.net (216.218.196.189)  4.334 ms  5.611 ms  3.930 ms
 5  10ge10-1.core1.sjc2.he.net (184.105.222.14)  16.673 ms  5.709 ms  12.821 ms
 6  sjo-b21-link.telia.net (213.248.67.105)  4.875 ms  7.110 ms  5.097 ms
 7  qwest-ic-300327-sjo-b21.c.telia.net (62.115.12.94)  4.848 ms  6.488 ms  8.788 ms
 8  * * *
 9  * * *
```

- Useful information for a trouble ticket
- Policy? Could filter or not send ICMPs

Overlay providing L2 service



- Host traceroute and ping show nothing
- Host ARP request may or may not time out
- Need access to ingress NVE to
 - Inspect tables - Mac address to NVE address? Port, vlan to vni id mapping?
 - Run ping/traceroute to destination NVE
- Without NVE access trouble ticket is empty
 - Difficult to troubleshoot temporary conditions

ICMP errors

- Traceroute relies on ICMP
 - ttl exceeded
- Other ICMP errors in underlay can be relayed just like ttl exceeded
- ICMP packet too big if fragmentation issues
 - Adjust ICMP PTB MTU field by size of encaps
 - Send resulting ICMP error to host
 - Helps detect PMTU mismatch overlay to underlay



Next Steps

- A requirement for an mechanism to be able to see underlay errors?
 - Policy whether reflected back to host in overlay
- A requirement to detect PMTU mismatch between overlay and underlay?
 - Ditto policy
- Useful to require a “Uniform ttl” bit in NVO3 header?

