

Unexpected traffic in LHCONE ?

Bruno Hoefft / KIT
Kars Ohrenberg / DESY

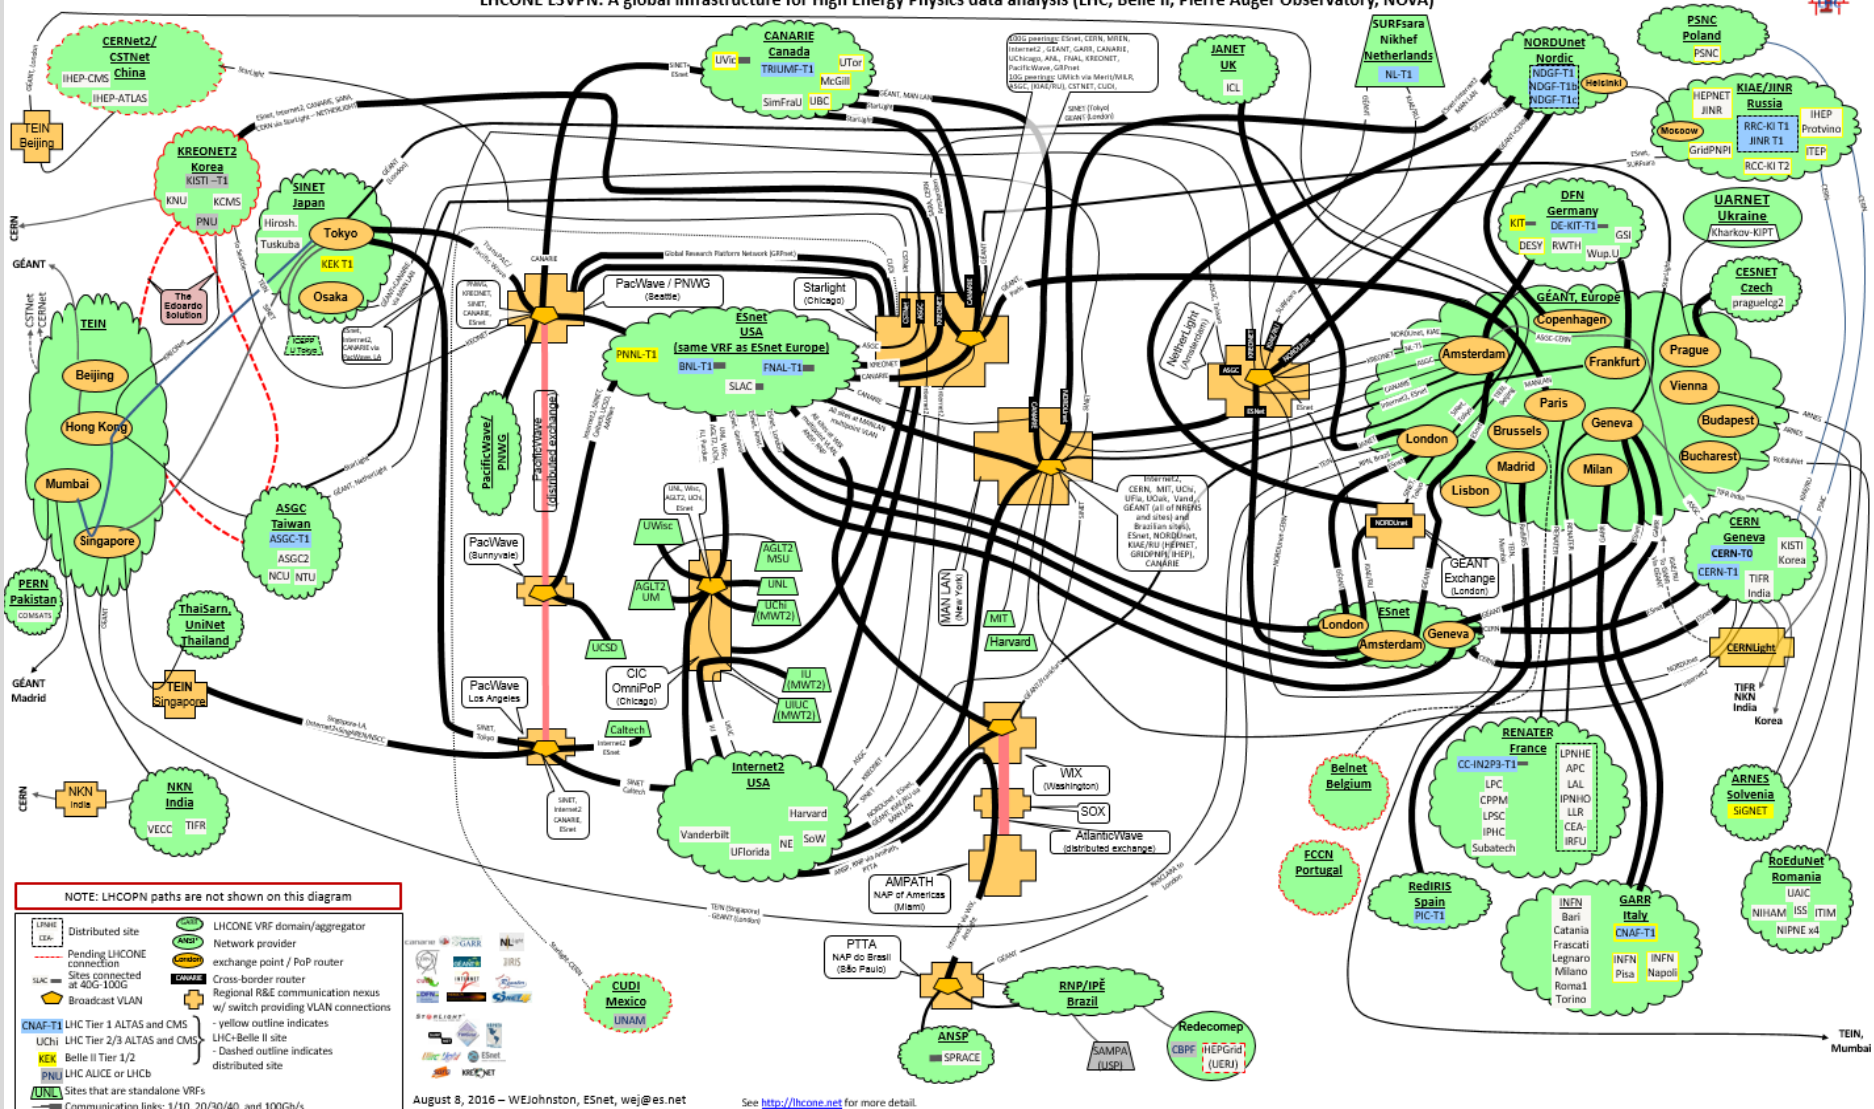
STEINBUCH CENTRE FOR COMPUTING - SCC



LHCONE



LHCONE L3VPN: A global infrastructure for High Energy Physics data analysis (LHC, Belle II, Pierre Auger Observatory, NOvA)



LHCONE policy



- VPN
- Data exchange only between connected sites
- Cidr exchange via BGP
- Follow the aup
 - Participating collaborations
 - Scope of sites/nodes
 - Roles and Responsibilities

LHCONE miss-routed

- At DESY / KIT / ASGC
 - Data from non LHCONE received
 - Multiple captures over 4 to 24 hours
 - KIT → 135 ipv4 (/24 ?) subnets + 16 ipv6 (/64 ?) subnets
 - incl. private addresses (10.x.x.x, 172.16.x.x, 192.168.x.x)
 - DESY → 203 ipv4 (/24 ?) subnets + 26 ipv6 (/64 ?) subnets
 - incl. private addresses
 - Limited overlapping only at DESY and KIT (approx. 15%)
 - ASGC → reported miss-routed data

- Ports:
 - icmp
 - pcsync-https (8443), domain (53), 20000-25000
 - eyetv (2170)
 - ...

- Only a vage suggestion : AMS-IX



Reverse check of KIT/DESY

- most source addresses (no private addr.) „reachable“ (ping)
- Packets not routed to LHCONE
 - traceroute from DESY/KIT not via LHCONE



Suggestions: how to avoid

Avoiding (only some ideas):

- Sites configure routing AUP accordingly
- Reverse path forwarding (RPF) check at sites/NREN
- At Regional internet registries hosted LHCONE prefix table
 - NREN controlled and maintained table
- ...

Questions Suggestions Discussion

UNI-Mainz to ASGC

```
[root@f01-120-126-e ~]# traceroute 117.103.103.71 -l
traceroute to 117.103.103.71 (117.103.103.71), 30 hops max, 60 byte packets
 1 192.108.68.2 (192.108.68.2) 0.613 ms 0.775 ms 0.979 ms
 2 * * *
 3 * * *
 4 cr-fra2-hundredgige0-9-0-1-704.x-win.dfn.de (188.1.153.65) 3.042 ms 3.116 ms 3.272 ms
 5 geant-lhcone-gw.mx1.fra.de.geant.net (62.40.126.184) 2.917 ms 2.948 ms 2.953 ms
 6 62.40.126.229 (62.40.126.229) 20.562 ms 20.735 ms 20.709 ms
 7 asgc-lhcone-gw.ams.nl.geant.net (62.40.126.230) 20.092 ms 20.128 ms 20.173 ms
 8 e-1-3.r0.chi.asgc.net (117.103.111.202) 121.144 ms 121.170 ms 121.184 ms
 9 e-1-2.r2.tpe.asgc.net (117.103.111.206) 303.109 ms 303.001 ms 302.998 ms
10 * * *
11 f-arc02.grid.sinica.edu.tw (117.103.103.71) 303.004 ms 303.001 ms 302.997 ms
```

```
[root@mgftp1 ~]# traceroute -l 117.103.103.71
traceroute to 117.103.103.71 (117.103.103.71), 30 hops max, 60 byte packets
 1 g1341-ms3d-1.zdv.Uni-Mainz.DE (134.93.168.125) 1.086 ms 1.743 ms 2.401 ms
 2 134.93.255.177 (134.93.255.177) 0.341 ms 0.376 ms 0.449 ms
 3 g1341-0033-1.zdv.Uni-Mainz.DE (134.93.254.5) 0.720 ms 0.780 ms 0.839 ms
 4 g-uni-mz-1.rlp-net.net (217.198.241.33) 3.293 ms 3.560 ms 3.714 ms
 5 g-uni-klinik-mz-1.rlp-net.net (217.198.241.6) 1.693 ms 1.926 ms 2.121 ms
 6 g-interxion-1.rlp-net.net (217.198.240.101) 0.985 ms 0.924 ms 0.939 ms
 7 te-0-0-6-561.core1.fra1.ix.f.man-da.net (82.195.78.5) 1.014 ms 0.938 ms 1.136 ms
 8 te-0-2-0.peer1.sara.ams.man-da.net (82.195.67.90) 6.305 ms 6.214 ms 6.222 ms
 9 * * *
10 e-1-3.r0.chi.asgc.net (117.103.111.202) 103.192 ms 103.175 ms 103.161 ms
11 e-1-2.r2.tpe.asgc.net (117.103.111.206) 289.224 ms 289.224 ms 289.219 ms
12 v3967.n7k.tpe.asgc.net (117.103.111.233) 286.798 ms 285.475 ms 285.456 ms
13 f-arc02.grid.sinica.edu.tw (117.103.103.71) 289.221 ms 289.210 ms 289.190 ms
```