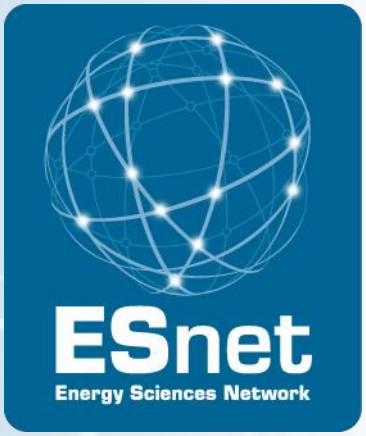

LHCONE Operations Update

Mike O'Connor

ESnet Engineering Group

LHCOPN and LHCONE Joint Meeting

Caltech

Dec. 3, 2013

New LHCONE collaborators

New compute centers

- MERIT MICH-Z(230)
- Vanderbilt University (39590)
- Czech Academy of Science (2852)

Near term

- North East Tier2 – I2 MANLAN Boston University & Harvard
- Kurchatov Institute (59624)– ESnet STARLIGHT, RU Net via Gloriad
  - Limited access to ESnet sites only

LHCONE Collaborating NSPs and Compute Centers

CANET(6509)
- BCNET(271)
- UTORONTO(239)
- UVIC(16462)
- MCGILL(15318)
- TRIUMF(36391)
- UALBERTA(3359)

ESNET(293)
- FNAL(3152)
- BNL(43)
- SLAC(3671)

I2(11537)
- UIUC(38)
- UNL(7896)
- MIT(3)
- AGLT2(229)
- **MICH-Z(230)**
- UOC(160)
- CSUNET(2153)
- ULTRALIGHT(32361)
- **VANDERBILT(39590)**

INDIAN(19782)
- IUPUI(10680)

North America

Europe

CERN-LHC1(20641)
- CERN-WIGNER(61339)
- CERN(513)

DFN(680)
- KIT(34878)
- DESY(1754)

GEANT(20965)
- ROEDUNET(2614)
- ASGARR(137)
- ARNES-NET(2107)
- **CZECH-ACAD-SCI(2852)**

LHC1-RENATER(2091)
- IN2P3(789)
- CEA-SACLAY(777)

NORDUNET(2603)
- NDGF(39590)

Migration to 100GE Substrate

LHCONE is deployed primarily over shared infrastructure.

All participating NSPs have either upgraded, or are in the process of upgrading their core networks to 100GE circuits.

Single 10GE circuits that transport LHCONE along with other traffic are likely to experience periods of saturation.

NSPs should identify their remaining non-aggregated10GE segments and plan to eliminate them, this includes inter domain connections used for BGP peering.

A 100G network substrate will be essential for deploying various kinds of virtualized networks to address the needs of the growing number of distributed scientific collaborations world-wide.

Reporting LHCONE Issues

May 2013, the University of Nebraska Lincoln (UNL) observed an apparent LHCONE issue and in reporting it sent mail to lhconeoperations@cern.ch.

While on the surface sending issues to this list seems reasonable, the LHCONE operations mailer is intended for discussing operations and has no process in place for managing issue resolution.

The method agreed to on the call and at prior face to face meetings was to work through your LHCONE NSP to report problems, however we also agreed that this was not clearly stated.

Issues with LHCONE should be reported to your LHCONE network service provider. This does not include regional providers not providing L3 LHCONE services.

Troubleshooting

LHCONE does not have a centralized NOC. Compute facilities should work through their LHCONE NSP and possibly regional network providers where necessary.

Operations model of the LHCONE benefits from it's close similarity to common routed R&E networking standards and operating procedures.

Information:

- Ticket chain
- Detailed topology
- Measurement infrastructure.

Ticket Chain

Cross domain troubleshooting will generate trouble tickets in each domain.

A series of trouble tickets will be created, generally one per NSP and these numbers will be shared in upstream and downstream communications while troubleshooting the issue. Maintaining a series of tickets, one for each NSP in the troubleshooting path. *LHCONE-operations call 11/18/2013*

Detailed Topology Data

Within each of our respective networks, we rely on topology and measurement data of various kinds to diagnose, locate and address issues using efficient internal communication and managed processes.

Detailed topology contribution from each participating NSP

- Contains VRF edge router and interface information
- Facilitates focused looking-glass queries
- Enables more specific requests that better fit general NOC processes
- Reduce the need for LHCONE specific operating procedures

Issues coordinating cross domain troubleshooting have impacted the effectiveness of our efforts to date.

Measurement

Measurement information is essential for identifying and locating performance issues. The LHCONE network lacks dedicated measurement infrastructure.

perfSONAR

- Performance baseline bandwidth & latency
- Changes in performance over time

Looking Glasses

SNMP interface metrics

- Load
- Errors
- Discards

Dedicated LHCONE measurement infrastructure implies that it is **only** accessible via LHCONE participating address space.

LHC perfSONAR MPs

Tier	Type	Hostname	IP address
RAL	Latency:	perfsonar-ps02.gridpp.rl.ac.uk	130.246.179.197
	Bandwidth:	perfsonar-ps01.gridpp.rl.ac.uk	130.246.179.196
CC-IN2P3	Latency:	ccperfsonar2-lhcopn.in2p3.fr	193.48.99.78
	Bandwidth:	ccperfsonar-lhcopn.in2p3.fr	193.48.99.79
CERN	Latency:	perfsonar-ps2.cern.ch	128.142.223.237
	Bandwidth:	perfsonar-ps.cern.ch	128.142.223.236
TRIUMF	Latency:	ps-latency.lhcopn-mon.triumf.ca	206.12.9.71
	Bandwidth:	ps-bandwidth.lhcopn-mon.triumf.ca	206.12.9.70
SARA	Latency:	ps.lhcopn-ps.sara.nl	145.100.17.9
	Bandwidth:	ps.lhcopn-ps.sara.nl	145.100.17.9
ASGC	Latency:	lhc-latency.twgrid.org	117.103.105.188
	Bandwidth:	lhc-bandwidth.twgrid.org	117.103.105.187
BNL	Latency:	lhcperfmon.bnl.gov	192.12.15.26
	Bandwidth:	lhcmon.bnl.gov	192.12.15.23
CNAF	Latency:	perfsonar-ps.cnaf.infn.it	131.154.254.11
	Bandwidth:	perfsonar-ow.cnaf.infn.it	131.154.254.12
NDGF	Latency:	perfsonar-ps.ndgf.org	109.105.124.86
	Bandwidth:	perfsonar-ps2.ndgf.org	109.105.124.88
PIC	Latency:	psl01.pic.es	193.109.172.188
	Bandwidth:	psb01.pic.es	193.109.172.187
FNAL	Latency:	psonar2.fnal.gov	131.225.205.141
	Bandwidth:	psonar1.fnal.gov	131.225.205.139
KIT	Latency:	perfsonar2-de-kit.gridka.de	192.108.47.12
	Bandwidth:	perfsonar-de-kit.gridka.de	192.108.47.6

LHCONE Reachable

None of these MPs are located within NRENs.

ESnet perfSONAR Planning

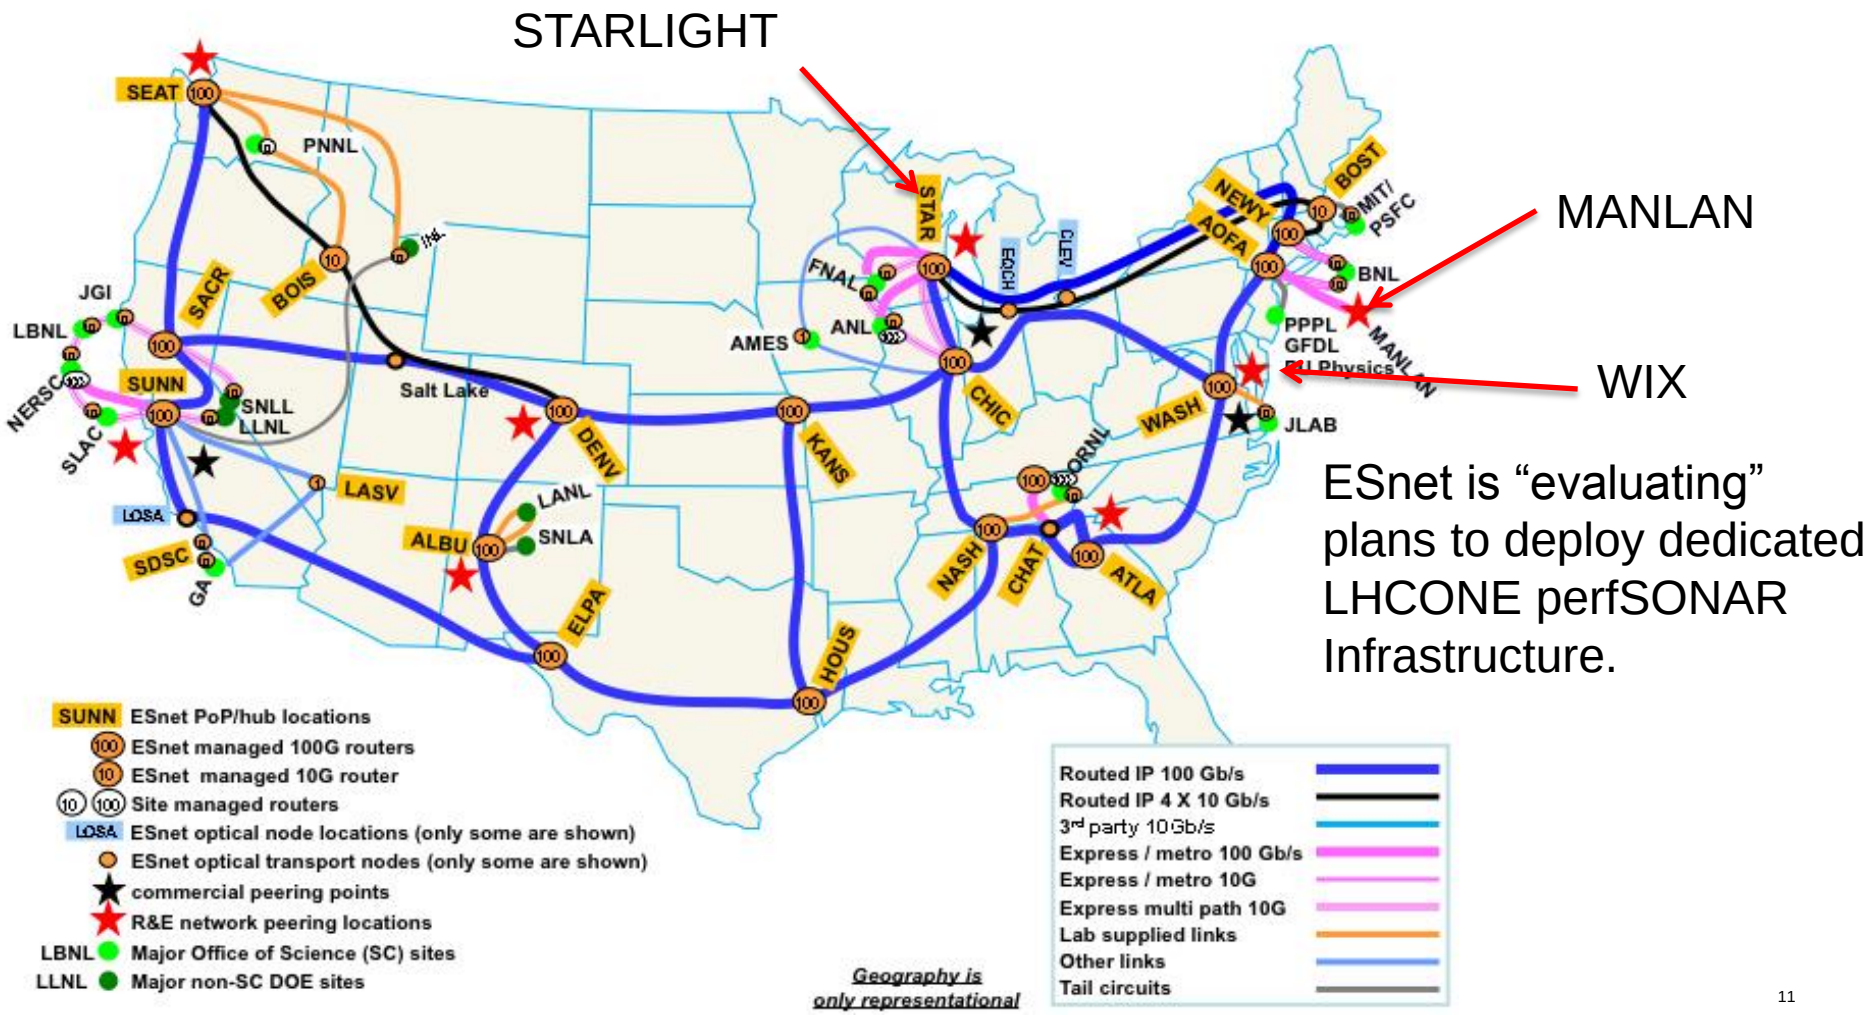

Questions?

Michael O'Connor

ESnet Network Engineer

moc@es.net

631 344-7410