



ESnet

ENERGY SCIENCES NETWORK

ESnet's LHCONE Service

March 23rd 2015

Presented by
Jason Zurawski, zurawski@es.net
Science Engagement

Authored by
Joe Metzger metzger@es.net
Network Engineering



U.S. DEPARTMENT OF
ENERGY
Office of Science



What is LHCONE?

LHCONE is a **Global Science Overlay Network** designed for high performance science workflows and dedicated to the HEP community.

- LHCONE provides access to HEP Computing Resources at:
 - LHC Tier 1, and Tier 2 and Tier 3 Centers
 - Belle II Tier 1 Centers (KEK & PNNL) and Tier 2 centers
- Non-science resources are NOT Allowed in LHCONE
 - Desktops, dorms, wireless networks, printers, etc.
- It is designed for performance
 - LHCONE end sites typically connect using the Science DMZ model with security infrastructure designed for science throughput
 - The traffic is clearly segregated so backbone providers can engineer, debug, and track it independently of general traffic.

LHCONE is a **Service** provided by research networks, such as ESnet, Internet2, GEANT, and national R&E networks.

Who is currently participating in LHCONE?

CANET(6509)

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

ESNET(293)

- FNAL(3152)
- BNL(43)
- SLAC(3671)
- PNNL(3428)
- AGLT2(229)*
- MICH-Z(230)*
- UNL(7896)
- UOC(160)
- CALTEC(32361)

I2(11537)

- MIT(3)
- UIUC(38)
- CSUNET(2153)
- VANDERBILT(39590)
- OKLAHOMA(25776)
- INDIANA(19782)
- IUPUI(10680)
- CALTEC(32361)

CERN-LIGHT(20641)

- CERN-WIGNER(61339)
- CERN(513)

GEANT(20965)

- DFN(680)
 - KIT(34878)
 - DESY(1754)ROEDUNET(2614)
- GARR(137)
- ARNES-NET(2107)
- CZECH-ACAD-SCI(2852)
- RoEduNet
- LHCONE-RENATER(2091)
 - IN2P3(789)
 - CEA-SACLAY(777)
- REDIRIS(766)
 - PIC(43115)

NORDUNET(2603)

- NDGF(39590)

SURFSARA(1162)

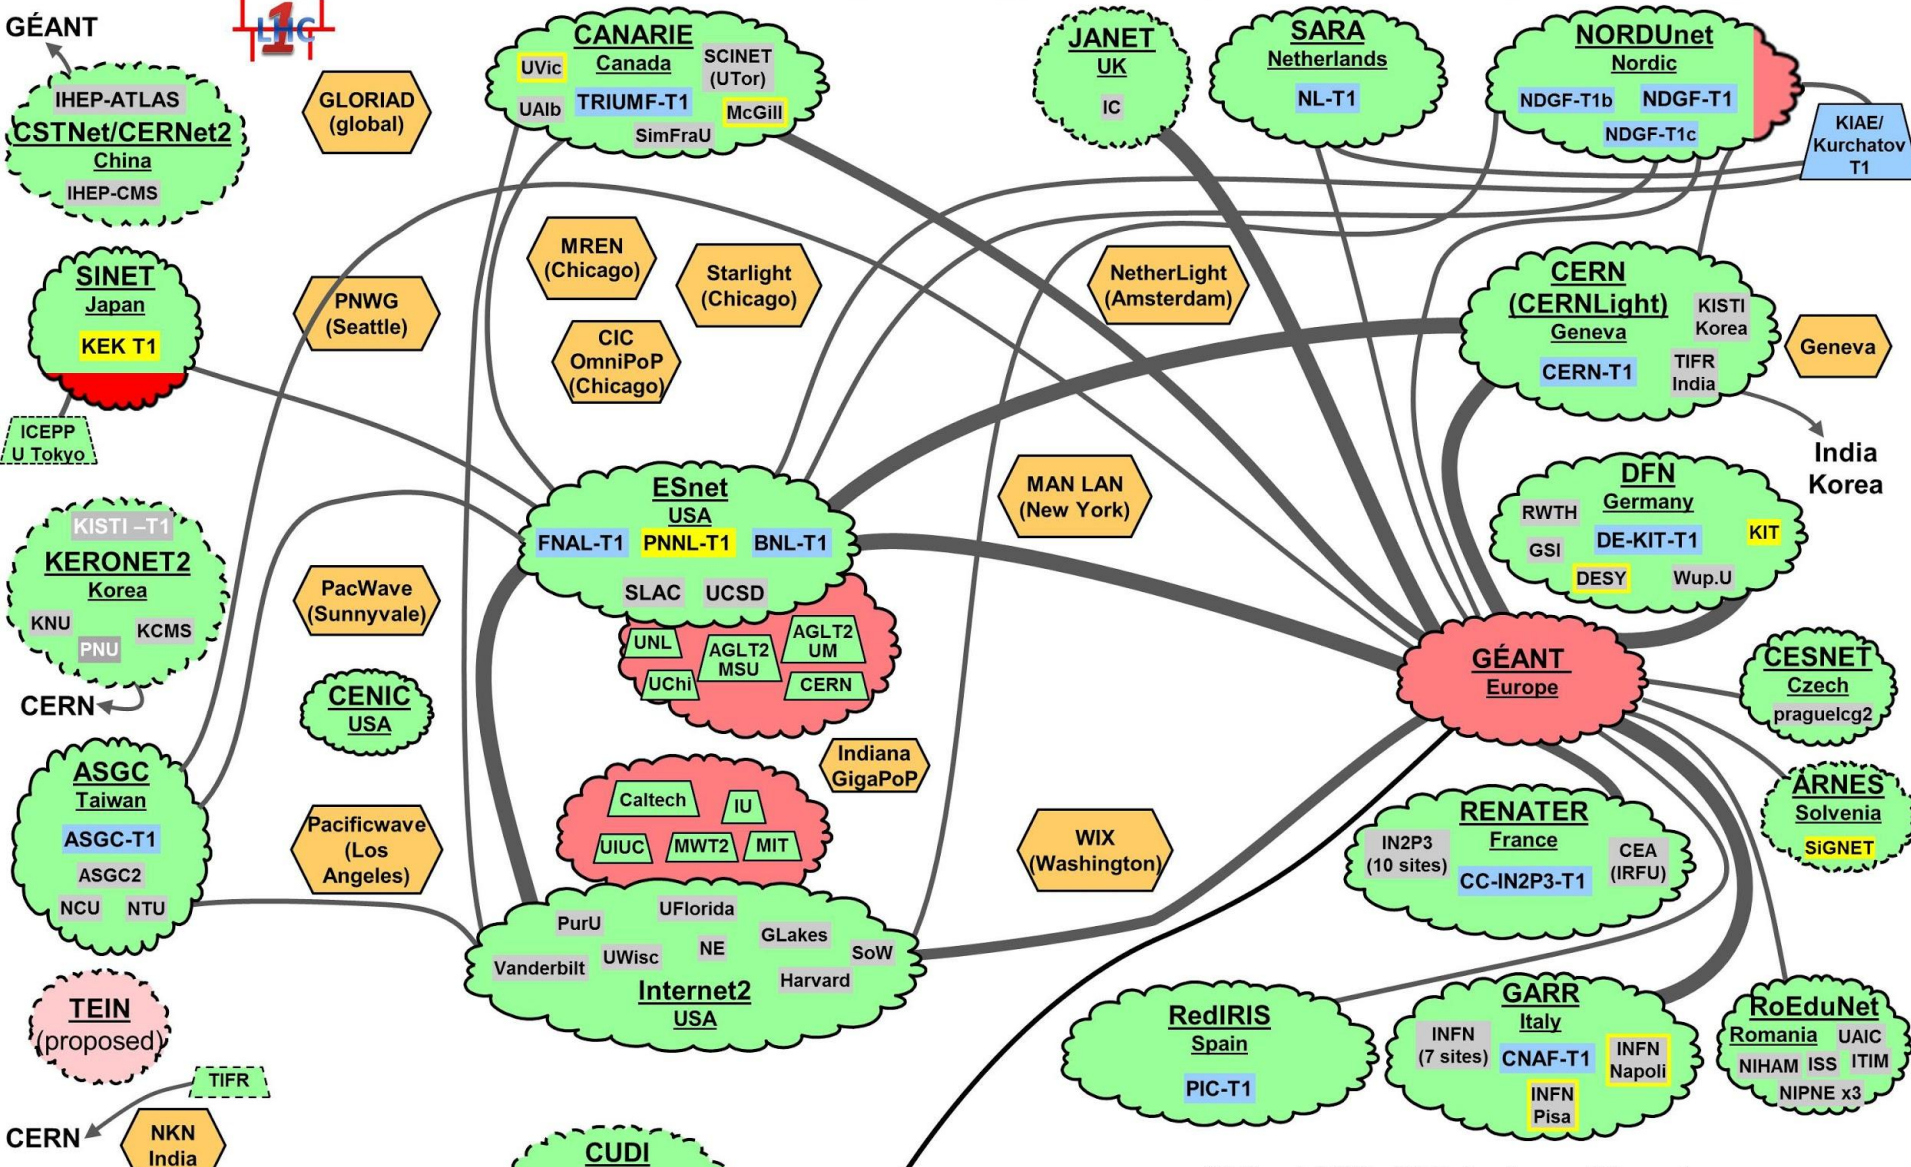
- NIKHEF(1104)

ASGC

KREONET

SINET

LHCONE: A global infrastructure for the High Energy Physics (LHC and Belle II) data management



22 March 2015 – WEJohnston, wej@es.net

	LHCONE VRF domain		LHC Tier 1/2/3 ALTA and CMS Belle II Tier 1/2	} yellow outline indicates LHC+Belle II site
	LHCONE VRF aggregator network		LHC ALICE or LHCb	
	Regional R&E communication nexus or link/VLAN provider		Sites that manage their own LHCONE routing	
	Communication links: 1, 10, 20/30/40, and 100Gb/s		See http://lhcone.net for details.	



LHCONE P2P

LHCONE P2P is an experiment with coordinating the scheduling of compute, storage and networking

- Kickoff: Sept 2014
- Goal: demonstrate an implementation of the LHCONE Point2Point Experiment with a number of LHC sites, based on the Automated GOLE* infrastructure
 - Activity 1: Connecting LHC sites and AutoGOLE
 - Activity 2: Middleware integration
- Current Status (Feb 2014)
 - SURFsara, DE-KIT, Caltech Connected
 - Brookhaven and Fermilab pending.
 - SURFsara – NetherLight – GÉANT – DFN – DE-KIT – Requesting bandwidth through NSIv2

Credit: Gerben van Malenstein (SURFNET)



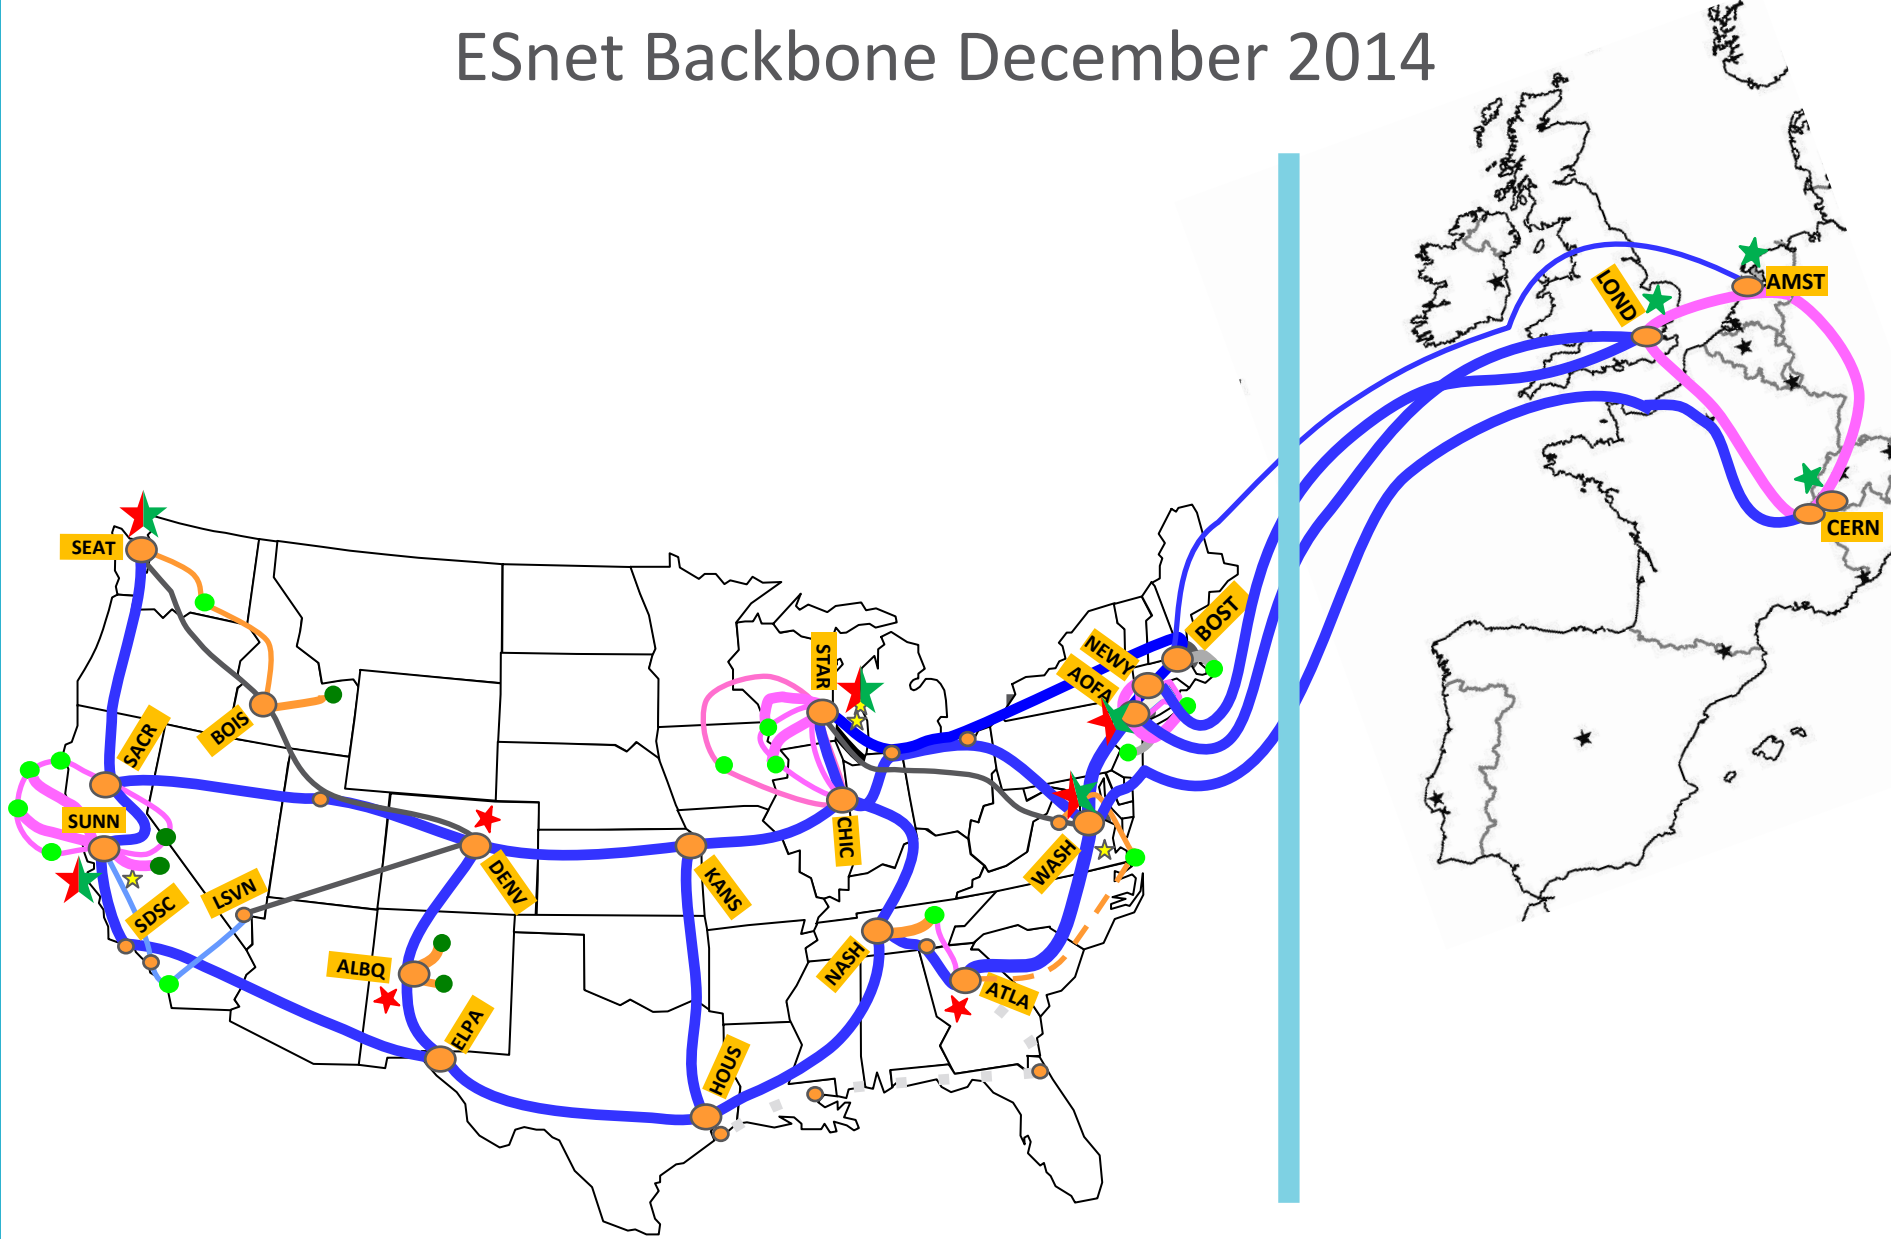
*GOLE = GLIF Open Lightpath Exchanges

ESnet's LHCONE Service

- Provides connectivity to the global LHCONE Overlay Network.
 - From LHC resources at US Universities to LHC centers in the US and abroad.
 - Between US Universities LHC centers.
- Is implemented as a BGP/MPLS VPN/VPRN on top of ESnet's International backbone.
- Is managed & controlled by representatives of the US Atlas and US CMS Experiments.
- Is funded by DOE as part of the ESnet mission.
 - There is no charge to university participants.
- Is managed like other ESnet network services
 - Same infrastructure
 - Same support model
 - Same performance expectations

ESnet is DOE's science network. ESnet is able to provide LHCONE services to US Universities because the LHCONE overlay network's use is constrained to a science mission that is approved and supported by DOE.

ESnet Backbone December 2014



SUNN ESnet PoP/hub locations

100 ESnet 100G routers

★ R&E network peering locations – US (red) and international (green)

★ commercial peering points

Geography is only representational

100 Gb/s

40 Gb/s


Express / metro / regional



Overview of ESnet LHCONE Service Turnup Process

1. The **Experiment Coordinators** provide ESnet with a Technical Contact for a LHC Tier2 or Tier3 center.
Atlas Experiment Site Coordinator are: **Michael Ernst of BNL and Rob Gardner of U Chicago**
CMS Experiment Site Coordinator are: **James Letts of UCSD and Kevin Lannon of Notre Dame**
2. ESnet sends out a template and schedules a 1 hour conference call with the Experiment Coordinator, the Technical Contact(s), and appropriate engineers representing the University, and regionals as necessary.
3. The Technical Contact fills out the template saying they agree with the LHCONE AUP, ESnet AUP, and provides the technical details necessary for implementing the service.
4. The service is tested and turned up.

ESnet is working on approximately 5 universities in parallel. When we finish one, we will move to the next university on the prioritized list provided by the US LHC Experiment management.

Most services will be provided via existing connections to Gigapops, Exchanges, or Internet2. ESnet is not building physical infrastructure into Universities.*  **ESnet**

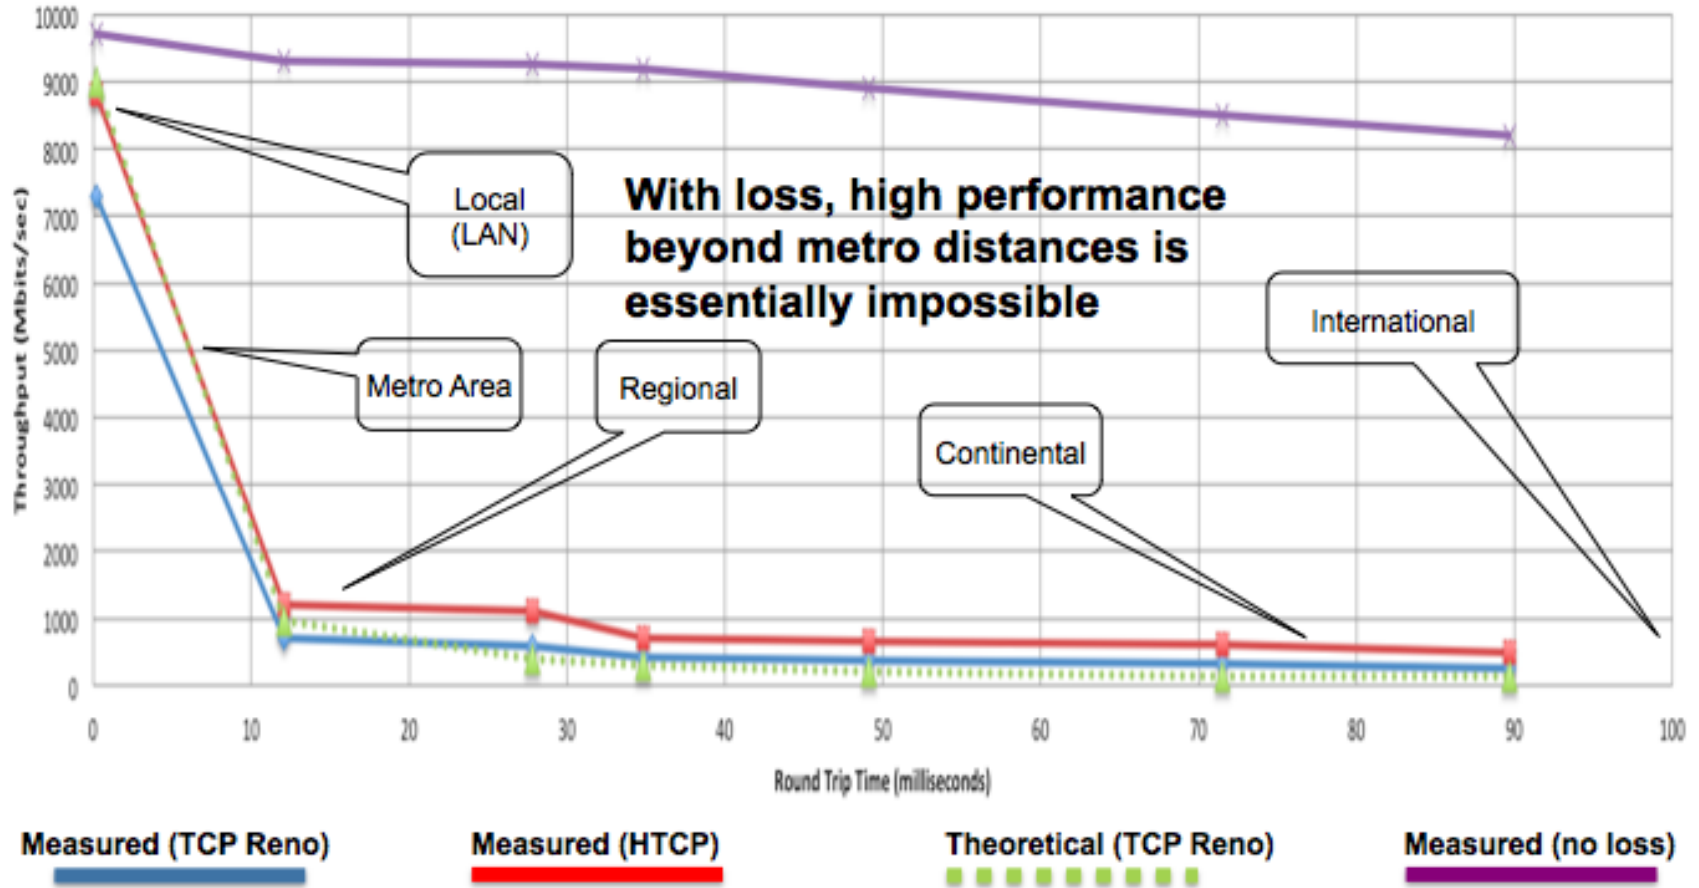
Note on Testing

- ESnet engineers will actively work with the Universities when turning up new services to ensure that the services work correctly.
- If your infrastructure is "Known Good", then this is a very light-weight process.
- If your infrastructure is unknown or suspect, then we can help you characterize it, tune routers and switches, and localize problems.
- ESnet engineers have considerable experience debugging wide area network problems, and have *powerful tools, including 100G test sets*, to facilitate problem identification and isolation.

Please keep in mind that there is a surprising number of hidden performance problems out in the Net that don't impact local traffic, but kill long distance high bandwidth flows. Such as devices supporting speed transitions (40G to 10G, or 100G to 40G) with tiny buffers. See <http://fasterdata.es.net/>

Note on Testing

Throughput vs. Increasing Latency with .0046% Packet Loss



Reference Material

LHCONE AUP

- <https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneAup>

ESnet LHCONE Service Description Document

- <http://www.es.net/assets/ESnetLHCONEServiceDescription.pdf>

LHCONE Web Site

- <http://lhcone.net>

CERN LHCONE WIKI

- <https://twiki.cern.ch/twiki/bin/view/LHCONE/WebHome>

Next LHCONE Face to Face Meeting

- June 1 & 2 at LBNL (Berkeley, CA)
- <https://indico.cern.ch/event/376098/>

Additional Material

ESnet LHCONE Implementation Template (1)

ESnet Ticket #	ESNET-2015xxx-xx
----------------	------------------

Role	Name	Date Completed/Approved
Experiment Site Coordinator	(One of: M. Ernst, R. Gardner, J. Letts, or K. Lannon)	
University Technical Contact		
ESnet Engineer		

Site Name		AS Number:	
Experiment	Atlas or CMS	Target Install Date:	
University Technical Contact	<i>Name</i>	<i>Email</i>	<i>Phone</i>
University NOC Contact	<i>Name</i>	<i>Email</i>	<i>Phone</i>
University Security Contact	<i>Name</i>	<i>Email</i>	<i>Phone</i>

ESnet LHCONE Implementation Template (2)

Do you have an existing LHCONE connection?	
Will your site implement a separate WAN connection for LHCONE? As part of a science DMZ for instance.	
Which ESnet location/demarc will you be connecting to?	<i>(Select from list)</i>
Please describe the path between your University and the ESnet demarc in terms of: Bandwidth, shared/dedicated resource, location and your VLAN ID & MTU.	
List of prefixes you will announce to ESnet's LHCONE Service?	
What proportion (by host address) of the prefixes in your ASN will you announce to LHCONE?	
What types of systems are contained in the prefixes you will announce, other than LHC compute & storage systems?	
What architectural model or protocol policy techniques will be in place to ensure routing symmetry to the LHCONE in accordance with the LHCONE Site Provisioning Guidelines.	
Will your site's perfSONAR infrastructure be included in the prefixes listed above.	
Prefix to be used for BGP Peering with ESnet AS 293	(ESnet assigned.)
<i>How will you implement BCP38, RPF or ACLs?</i>	
What email address should be subscribed to the LHCONE Operations mailing list? (For global LHCONE announcements)	
What email address should be subscribed to status@es.net ? (For ESnet maintenance events)	
What testing process is required?	