ESnet

ENERGY SCIENCES NETWORK

LHCONE Reachability Measurement

Michael O'Connor moc@es.net

Network Engineering

Energy Sciences Network

Lawrence Berkeley National Lab

LHCOPN-LHCONE Meeting

Taipei Taiwan

March 13, 2016

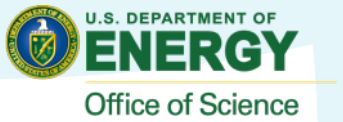

Introduction LHCONE

- The LHC Open Network Environment VRF Overlay Network is intended to be built on shared network infrastructure.
- 100Gbps technologies have reduced bandwidth costs making virtual network overlays cheap, quick and scalable, expect this paradigm to grow.
- The reduced cost, increased speed and the community centric governance models of overlay networks greatly enhance the Globally Distributed Compute Model relied on by international scientific collaborations.
- LHCONE will evolve, connecting is to join in shaping that evolution.
- LHCONE is not like other Internet uplinks, planning and care are essential for a smooth integration.

ESnet

LHCONE Global Expansion

- LHCONE expansion is driving the incorporation of continental transit service to bridge gaps in the existing full mesh BGP architecture.
- On the general Internet, no single provider – no matter how big – connects to all subscribers or to all of the global content.
- To connect it all together (to give all subscribers access to all content), providers must spend money and connect their networks together.
- What options and approaches are available in the research and education networking space to maintain full connectivity among all LHCONE participating institutions?

The challenge is to provide transit services in a zero cost peering framework

ESnet

Brief History of LHCONE Routing Architecture

The early LHCONE implementation was a **full BGP mesh**

- North America
  - CANET, ESnet, Internet2
- Europe
  - GEANT, NORDUnet, CERN

Transit paths were well established along standard provider customer relationships and funding models.

Recently **GEANT** and **Internet2** have begun providing LHCONE transit service.

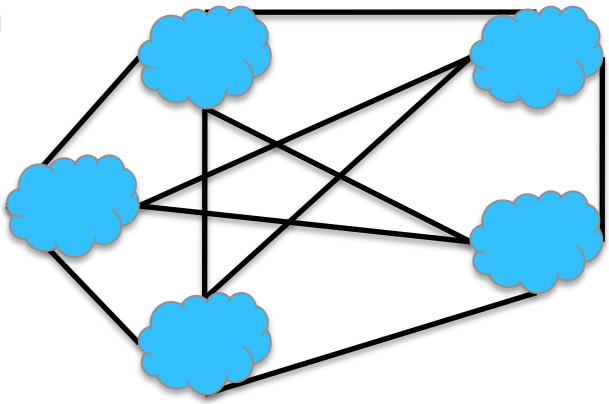

In general, LHCONE NSPs will consider providing transit on a case by case basis

ESnet

Full Mesh LHCONE BGP Routing

Established point to point or multipoint BGP peering

Today the principle LHCONE NSPs form a full BGP Mesh, establishing a fully connected network.

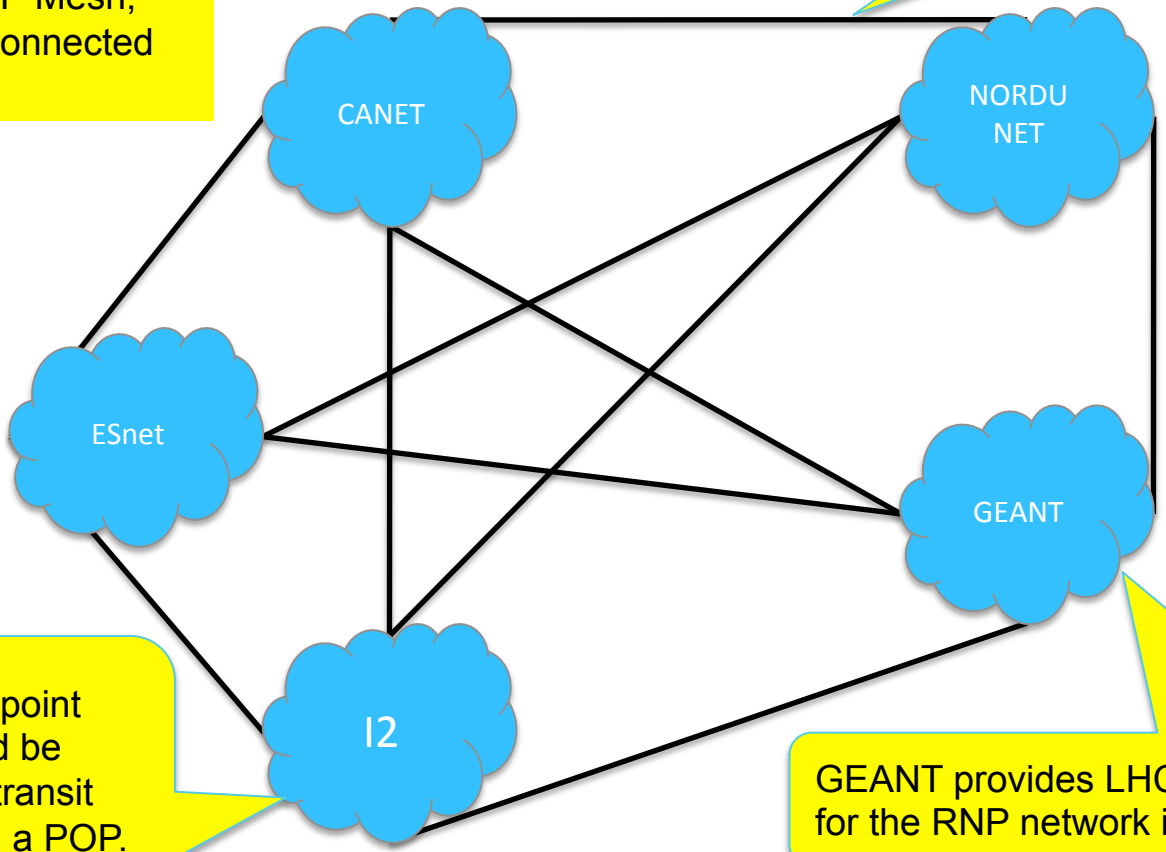

Internet2 hosted multipoint BGP exchanges could be considered a form of transit concentrated within in a POP.

GEANT provides LHCONE transit for the RNP network in Brasil

In general, LHCONE NSPs will consider providing transit on a case by case basis

ESnet

Connecting to Asia

KREONET Hypothetical Example

Example:
If ASGC joins LHCONE by connecting to CANET, ESnet and I2 on the US west coast in Seattle WA

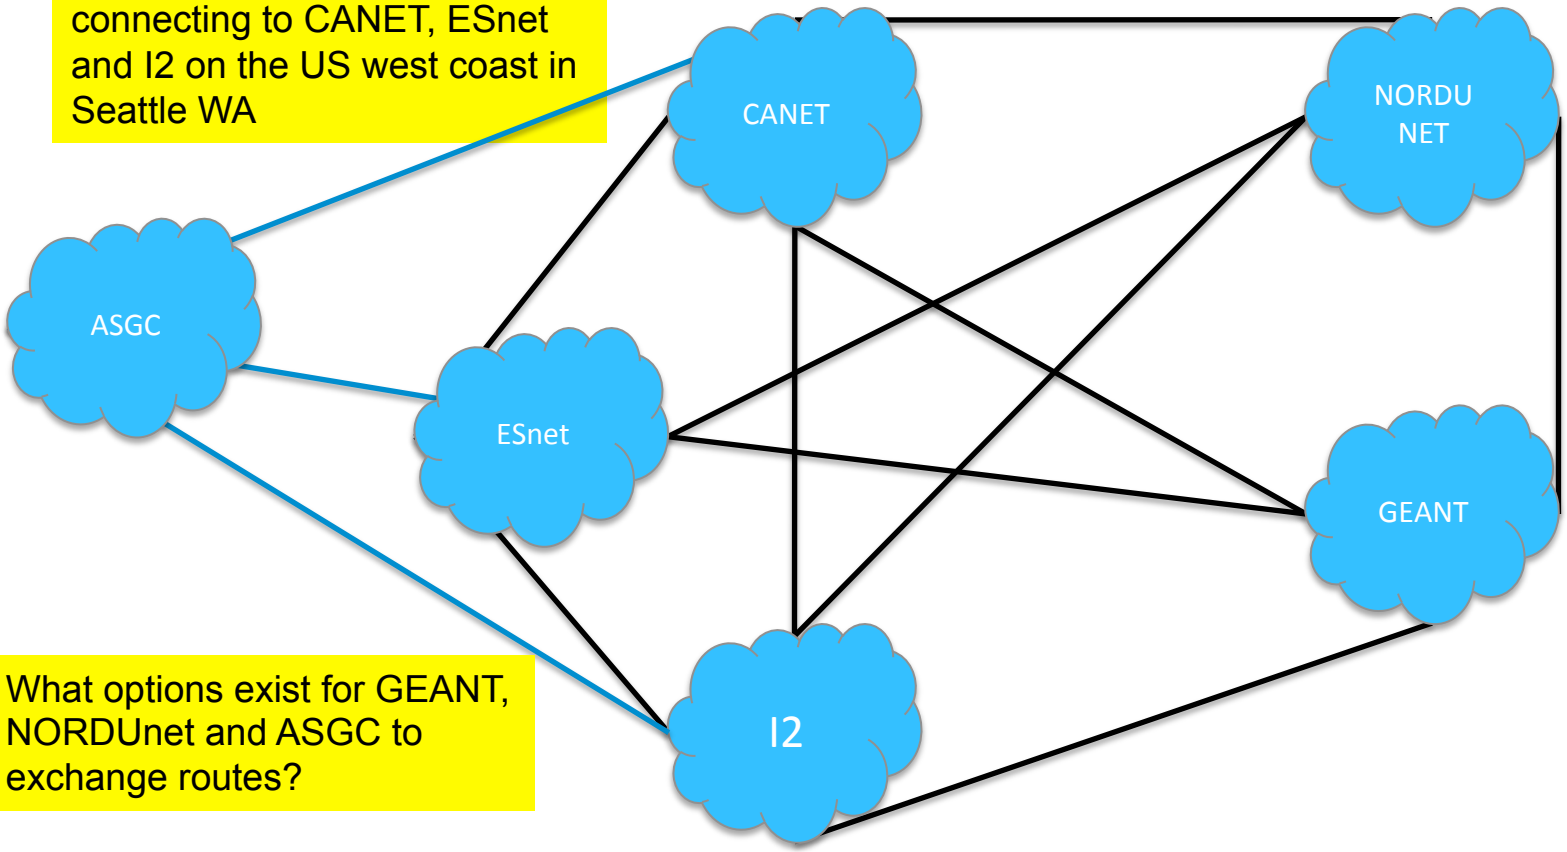

What options exist for GEANT, NORDUnet and ASGC to exchange routes?

ESnet

Connecting to Asia

KREONET Hypothetical Example

Example:
If ASGC joins LHCONE by connecting to CANET, ESnet and I2 on the US west coast in Seattle WA

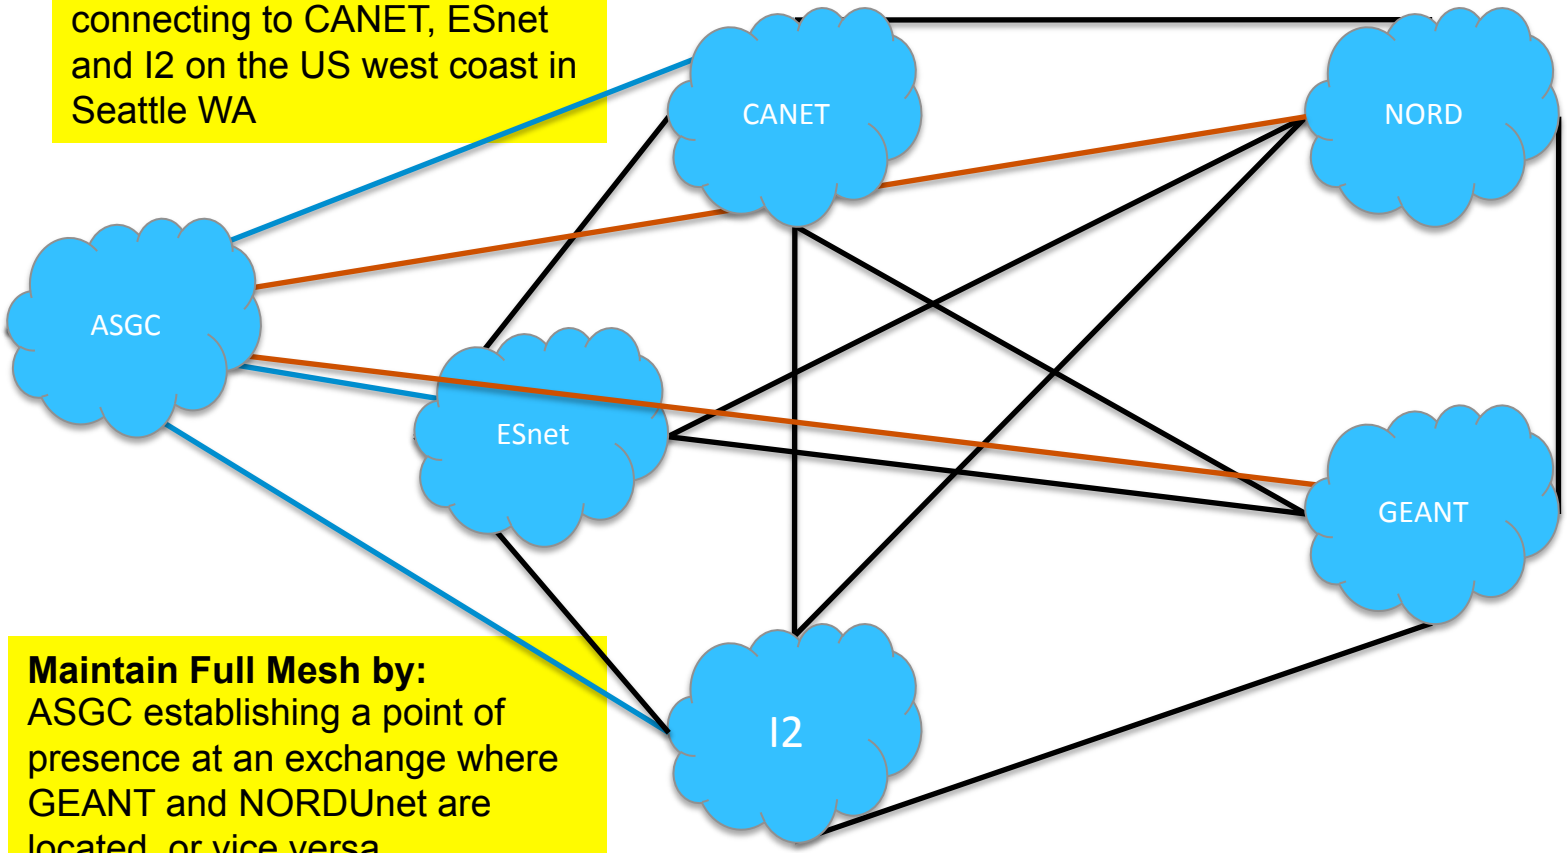

Maintain Full Mesh by:
ASGC establishing a point of presence at an exchange where GEANT and NORDUnet are located, or vice versa.

ESnet

Connecting to Asia

ASGC Hypothetical Example

- Dedicated transcontinental circuits are no longer required.
- Transit service also provides alternate paths to CANET and ESnet.

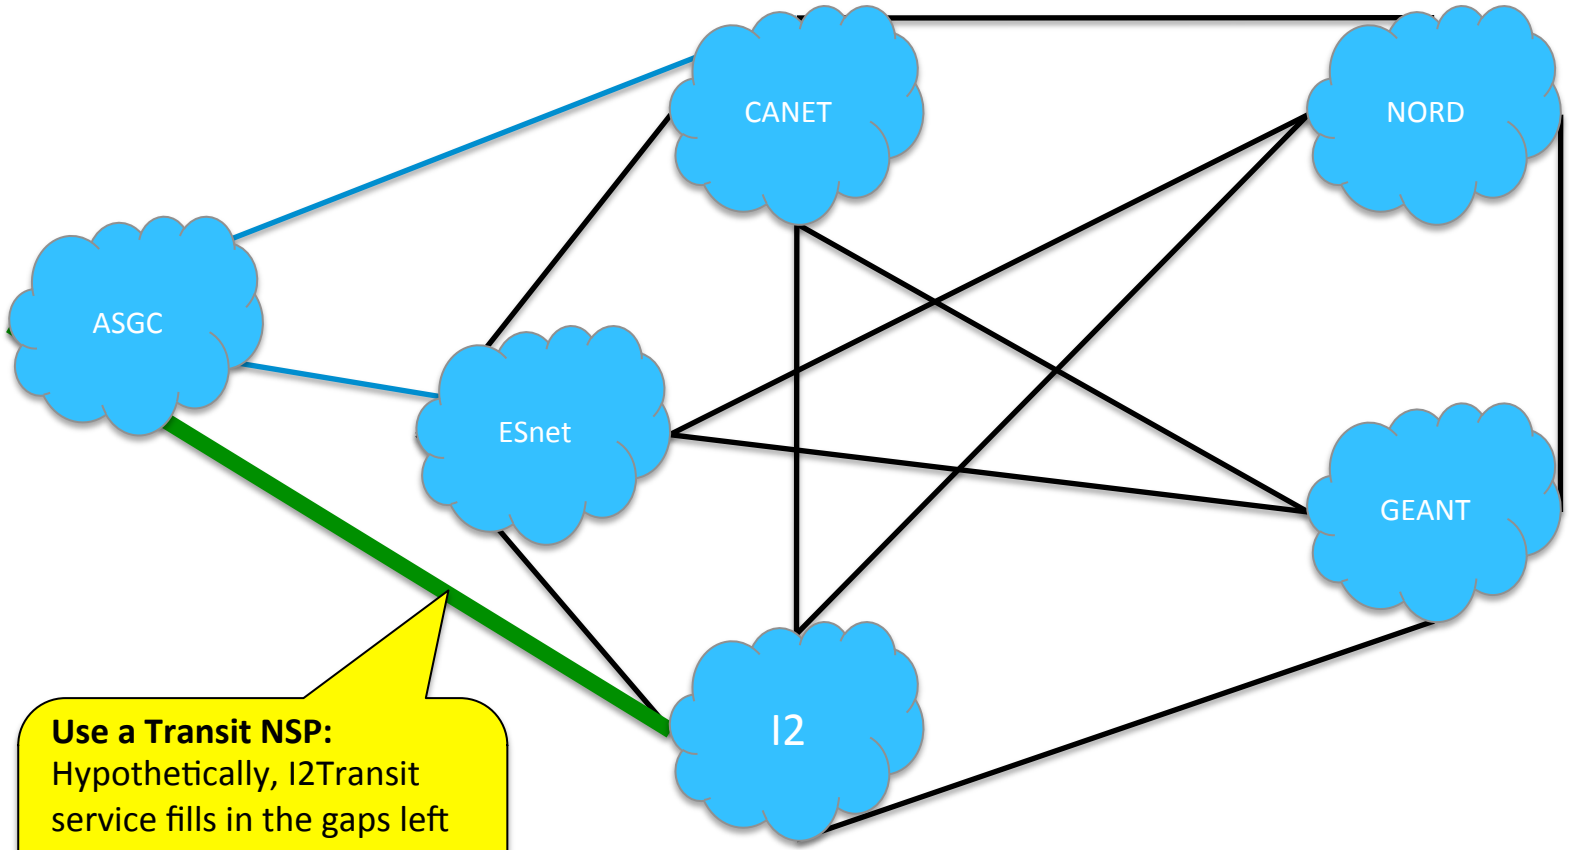

Use a Transit NSP: Hypothetically, I2Transit service fills in the gaps left in the partial mesh by exporting the full routing table to ASGC.

ESnet

Connecting to Asia

How will this scale?

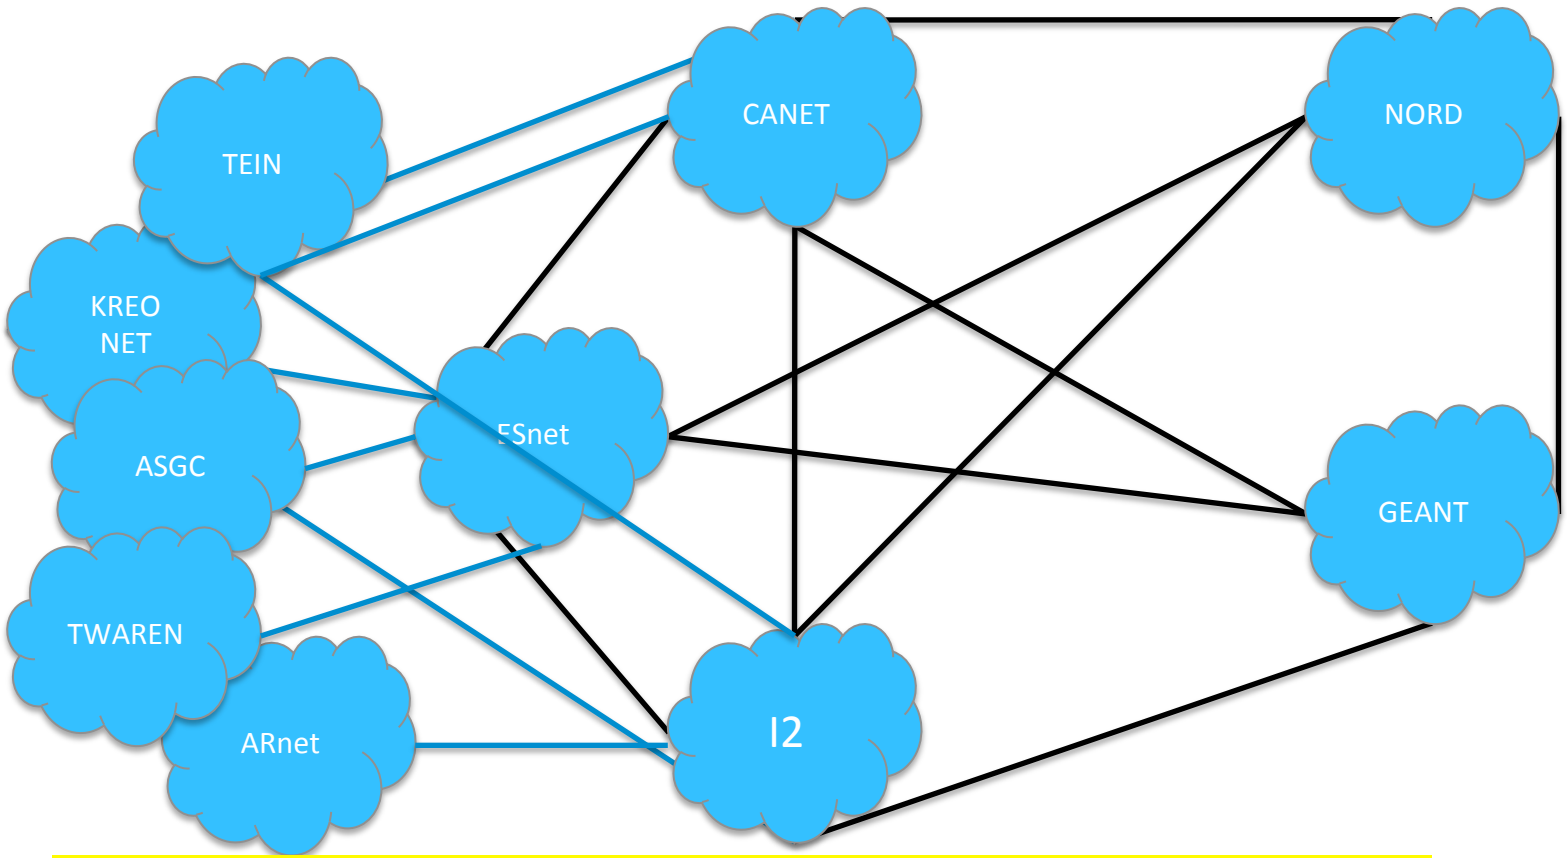

Clearly an Asian regional transit service similar to that provided by GEANT in Europe would simplify LHCONE expansion into Asia.

ESnet

Asian Transit Agreement
TEIN, ASGC, KREOnet

TEIN, ASGC & KREOnet have agreed to implement LHCONE VRFs and to interconnect them in Hong-Kong where they will transit LHCONE globally.

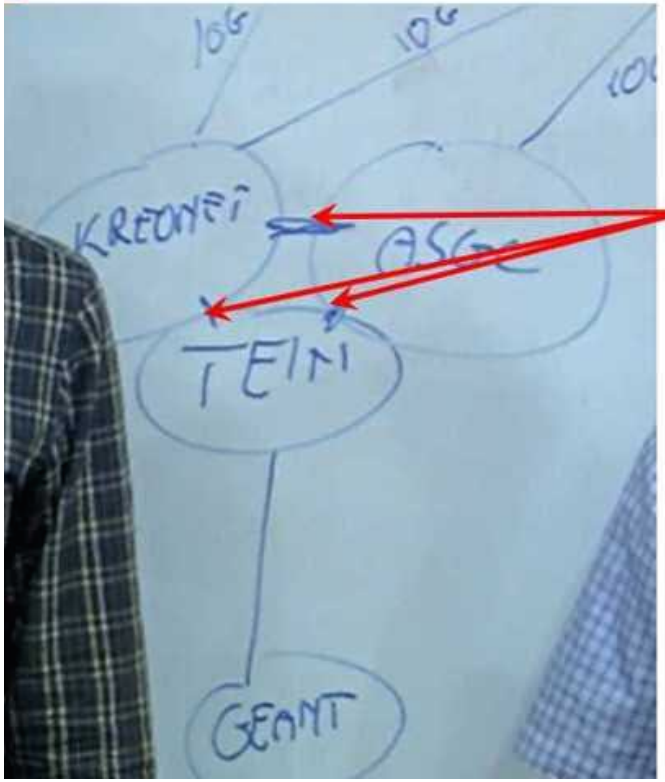

In the context of LHCONE VRF rules and inter-VRF connections

Brazil via GEANT

When RNP (Brazil) peered with the GEANT LHCONE VRF the LHCONE path from ESnet to RNP Crossed the Atlantic twice.

Path:

New York, London, Sao Paulo

ESnet and RNP peered directly in WIX to improve path efficiency by ~60ms

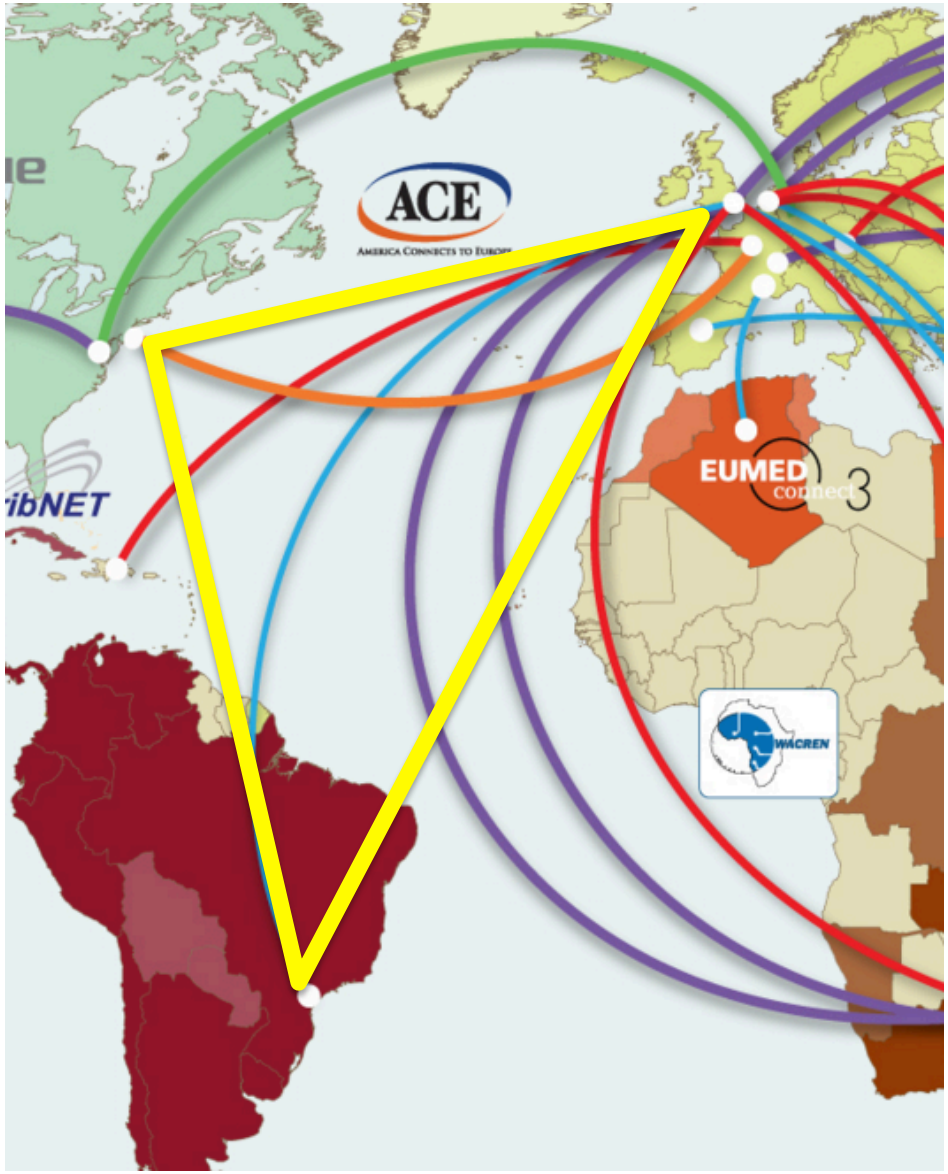

LHCONE paths should be compared with the General Internet Path

ESnet

The Need for Measurement

Goal – To ensure that every LHCONE site is reachable by every other LHCONE site across the LHCONE network.

- Global expansion of the LHCONE network has strained the full mesh architecture between NSPs.
- To maintain full reachability, LHCONE now requires some NSPs to provide transit to other NSPs.
- These transit arrangements are being arranged on a case by case basis, there may be gaps.
- As a high performance network, routing path efficiency will need to be maintained between NSPs.

ESnet

Maintaining LHCONE Reachability Through Measurement

TWiki > LHCONE Web > WebHome (2015-10-24, MichaelOConnorExternal)

twiki.cern.ch/twiki/bin/view/LHCONE/WebHome

LHCONE

Large Hadron Collider Open Network Environment

Monitoring

- perfSONAR
- MaDDash LHCONE
- MaDDash WLCG
- LHCONE BGP Filtering State
- ESnet LHCONE prefixes (JSON)
- ESnet LHCONE IPv6 prefixes (JSON)

- A new report containing LHCONE BGP Ipv4 and Ipv6 route prefixes.
- JSON format facilitates automated route table comparisons.

Proposal:

1. Define an LHCONE route table reporting format.
2. Principle NSPs will share LHCONE route tables.
3. Table comparison will yield an LHCONE reachability metric, identifying isolated prefixes.

ESnet

Summary

- LHCONE success is driving expansion and that is a good thing.
- To manage growth, the community needs to develop:
  - Strategies for maintaining LHCONE reachability.
  - Policies for mitigating the exploitation risk of zero cost transit services.
- What are the transit policies of the principle LHCONE NSPs?
- Are existing transit policies sufficient to support global expansion?
- What technical approaches and tools are required to maintain a fully connected LHCONE network?

ESnet

ESnet

ENERGY SCIENCES NETWORK

Thank You

Michael O'Connor moc@es.net

Network Engineering

Energy Sciences Network

Lawrence Berkeley National Lab

LHCOPN-LHCONE Meeting

Taipei Taiwan

March 13, 2016