LHCONE P2P service meeting notes, SC13 (Nov. 19, 2013)

wej@es.net

Attendees: Tangui Coulouarn, NORDUnet, Stefano Zani, CNAF, David Foster, CERN, Harvey Newman, Caltech, Artur Barczyk, Caltech, Richard Hughes-Jones, DANTE, Joe Mambretti, Northwestern/STARLight, Bruno Hoeft, KIT, Jerry Sobieski, NORDUnet, Ronald van der Pol, SURFNet, Gerben van Malenstein, SURFnet

Version of December 3, 2013 reflecting discussion during the Nov. 25 LHCONE call

In order to make progress, it was agreed that a number of things must be done ASAP.

It was agreed that involving some end-sites in the experiment from the beginning is essential.

It was noted that the LHCONE community is not the first science community to make use of inter-domain NSI circuits: The Radio Astronomy community, at least in Europe, is using a pre-production NSI service involving GEANT, NORDUNet, and SRUFNet.

The essential steps were seen to be:

1. Revise the "LHCONE Point-to-Point Connection Service" document to reflect a simple set of service parameters that are i) common to the LHCONE community, and ii) supported by the current NSI implementations that will be involved in the initial experiment. For example, 1 GE, 10GE, and 100GE VLANS.
    i. This does not necessarily mean that every domain had to support all three services, though at least two adjacent ones would have to.
    ii. The current document authors appear to include Jerry Sobeiski and Dale Finkelson
2. Establish ground rules for the initial experiment..
    a) The participants must:
        i. Name a lead engineer who is involved with, and can work on, the domain controller and the hardware that it controls
        ii. Name a project lead (possible the same as the lead engineer) who will
            a. Work with the other domain project leads to develop a project plan for testing a pre-production service
            b. Develop delivery dates for the elements of the project plan
            c. Deliver progress reports to the LHCONE Arch. List
            (The goal here is not to generate a lot of documents, but to document and push forward the experiment. For example the progress reports might be in the form of email messages to an archived list that all LHCONE participants can gain access to.)
            d. Develop a standard debug/diagnostic suite that can be used by all participants
            e. In the case of the sites, represent the user community
        iii. Commit to the experiment being a pre-production service that is a step toward establishing the service as a production service in their domain.
    b) Develop success criteria - What are the factors for success
        i. A "test harness" that reflects both the NSI standard and user expectations (e.g. guaranteed bandwidth) must be developed and successfully applied in each domain*
        ii. scalability and stress tests need to be agreed upon and metrics defined
        iii. meet minimum circuit setup time

a. 2 min max (based on experience with ISDN circuit setup time and when humans start to conclude that failure is occurring)
b. 20 sec. suggested as the target

iv. Inter-domain circuit setup is complex and failure mechanisms are opaque. Step by step feedback from each domain during the setup process should be provided for progress sanity checks and for debugging

c) Identify a member of the LHCONE community to be the P2P Service experiment project organizer that works with the domain project leaders to ensure that progress is made.

i. The project organizer will produce a periodic overall summary of progress for the community.

* The NSI standard is not yet complete (e.g. topology schema and AAA) are missing) however, these issues were finessed in the GLIF experiments and a similar approach should probably be taken for the LHCONE P2P experiment. The LHCONE community should not "run ahead" of the OSG NSI GW and try and define missing parts of the standard in any formal way. Rather, minimalist work-arounds should be found until the NSI WG completes the standard.

3. Proposed initial experiment
   a) Network domain to network domain
   b) Network domain to site
   c) Site to site
      i. L2 circuit terminates on a host
      ii. L2 circuit is a pseudowire connection between two routers at different sites and the sites route subnets to each other

4. It was recognized that there are multiple models that sites might use to participate
   a) Domain controllers may or may not be installed on an internal network that has access to the resources of the LHC community
      iii. The US universities in the NSF CCNIE program "have committed to doing this" (Joe Mambretti)
      iv. NSF Dynes sites that have the physical infrastructure in place to do this (e.g. a ScienceDMZ with domain controller switch)
   b) The NSI connection might terminate at the site boundary and the resulting VLAN connected to an internal mechanism
      i. For example both of the US Tier 1 centers use this model for the LHCOPN traffic. OPN circuits are delivered to the sites as OSCARS Virtual Circuits and those circuits terminate at the site boundary, being presented to the site as a protected VLAN on a port of an ESnet router located at the site.

5. A site-oriented workshop will be organized
   i. Joe Mambretti has agreed to work with the US NSF CCNIE program to organize such a workshop

6. An initial set of sites and networks was proposed. This set is not intended to be an exclusive list, though initial participants should essentially be ready to participate now (have an operating domain controller with an NSI 2.0 interface) or in the very near future. (Re: sites, see point 4, above.)

Table 1. Initial set of sites and networks ("experiment connections" must be participants only)

Site Domains	First hop domain	Domain model	Domain controller	Project engineer	Project lead (User rep.)	Experiment site connections
CERN	NetherLight	NSI?	?	?	?	U. Mich (?)
CERN	StarLight					
SARA	SURFNet	?	?	?		KIT
KIT	GEANT (DFN provides a static circuit KIT to GEANT)	?	?	Bruno Hoeft	Bruno Hoeft	SARA NDGF FNAL
Caltech	StarLight	NSI	?	?	Artur Barczyk (?)	CERN (?)
FermiLab (FNAL) (CMS Tier 1) (participation likely)	ESnet	Custom/static (circuits are used as P2P connections between dedicated routers)	-	Phil DeMar	Phil DeMar	KIT Caltech CERN
Brookhaven (BNL) (Atlas Tier 1)	ESnet	Custom/static	-	?	Michael Ernst	U. Mich CERN
U. Mich	StarLight	NSI?	Dynes	Roy Hock (?)	Shawn McKee (?)	BNL CERN (?)
NDGF	NORDUNet	?	?	?	?	KIT

Network Domains	Domain connections	Domain model	Domain controller	Project engineer	Project lead	Site connections
NetherLight	SURFNet NorduNet GEANT	NSI 2.0	BoD			SARA (?) CERN
GEANT	SURFNet NorduNet ESnet	NSI 2.0	AutoBHAN + NSI Bridge (Q1 2014)			KIT (via DFN)
NORDUNet	SURFNet GEANT NetherLight NDGF	NSI 2.0	OpenNSA	?	?	NDGF
USLHCNet	CERN ESnet	?	?	Artur Barczyk	Harvey Newman	CERN
ESnet	GEANT StarLight	NSI 2.0	OSCARS +NSI Bridge	Chin Guok (?)	Inder Monga (?)	FNAL BNL
Internet2	ESnet GEANT	?	?	?	Dale Finkelson	U.Mich.
StarLight	ESnet GEANT SURFNet	?	?	Joe Manbretti	Joe Manbretti	U.Mich.
MANLAN	ESnet NORDUNet GEANT USLHCNet	?	?	?	Dale Finkelson	

7. Other issues:
    i. How to involve the NRENs (Europe) and RONs (US)?
    ii. The US Dynes sites have the physical infrastructure in place to support the service (ScienceDMZ with resources and with NSI-capable switch) but do not have a working NSI domain controller.
    iii.
8. Attendees at SC13 BOF:

tangui.coulouarn@dtu.dk
stefano.zani@cnaf.infn.it
David.Foster@cern.ch
newman@hep.caltech.edu
Artur.Barczyk@cern.ch
Richard.Hughes-Jones@dante.net
j-mambretti@northwestern.edu
bruno.hoeft@kit.edu
jerry@nordu.net
Ronald.vanderPol@SURFnet.nl
gerben.vanmalenstein@surfnet.nl
wej@es.net