ESnet, TRIUMF/CANARIE, CERN/Caltech, GEANT, LHC meeting
CERN, July 7, 2005
Discussion on lhc-us-net@es.net

In order to facilitate specific planning and resource allocation for the CERN – North America connectivity for the LHC experiments, we need to obtain/develop enough information to make a project plan that includes drivers (service challenges, etc.), milestones, dependencies, cost, etc.

Tentative agenda

Morning 1

- Agenda bashing
- Establish as complete view as possible of the timing, capacity requirements, and site involvement for various international, LHC data flows to and from North America
    - Jamie Shiers will present Service Challenges 3 and 4
        - What are the requirements and the timetable?
    - CMS CDR needs
        - Don Petravick
    - Atlas needs
        - who?
- Identify the details of the connection locations and timing for the CERN-Starlight-MANLAN links
    - David Foster, Harvey Newman
- Connection issues for the North American tier 1 sites
    - FNAL
    - BNL
    - TRIUMPH

Morning 2

- Identify peering arrangements needed for various circuits, both existing and needed new ones
    - define the characteristics the need to be determined for "volunteer" circuits that might be available to LHC
        - capacity, access locations, policy, readiness, etc., on a circuit-by-circuit basis
    - how to manage this information? (see note 1)
- Identify the collection of specific links that might be used for backup paths
    - potential role of NSF/IRNC circuits
        - Tom DeFante (Translight), Maxine Brown
    - potential role of GLORAID Greg Cole
    - potential role of Atlantic Wave (Jerry Sobieski or Julio Ibarra)
        - Jerry Sobieski, invited, but late

Afternoon 1

- Engage the GLIF community
  - Presentation of LHC needs and strategy at IGrid2005
    - A panel is being organized with the help of Maxine Brown maxine@uic.edu
    - What should the LHC message be?
- North American continental networking issues
  - BNL-TRIUMPH connectivity requirements
    - René Hatem (?)
  - Internet2/Abilene
  - T1-T2 connectivity
    - Joe Metzger of ESnet is organizing a T1-T2 connectivity BOF at the I2-ESnet Joint Techs meeting in Vancouver
- Topics for follow-on discussion
  - Initially at I2/ESnet Joint Techs meeting in Vancouver, July 19

Afternoon 2

- Evolution beyond 2007
  - What will happen between 2007 (startup) and 2010 (full luminosity)?
    - How will requirements and usage evolve?
      - Harvey Newman and Jamie Shiers
  - Must take into account a 2-3 yr. lead time for major network enhancements
    - planning – funding proposal – engineering – construction – integration with production operation
- Operations and services issues
  - Traffic management paradigm(s)
    - Operational modes; e.g. options for directing flows in case of "overflow" or policy-based special handling of streams
      - ESnet, GEANT, Abilene (Bill Johnston, Roberto Sabatino, Rick Summerhill)
    - How to make real progress?
  - Monitoring and problem reporting
    - new monitoring facilities
      - who ?
    - How to make real progress?

- A Workshop Proposal

Based on some comments by Harvey Newman a discussion then ensued on the interaction of computing, middleware, and networking.

The essential observation is that virtual circuits (in all of their variations for network management, traffic engineering, QoS, etc.) are going to become a managed tool for scientists in the way that computing, storage, and middleware are now. This means that the network (in the

virtual circuits manifestation) will become a much more integrated part of the infrastructure of scientific experiments, and will be directly involved in the operations and management of the experiment. As such, this aspect of networks will become subject to experiment policy and resource management.

This has significant implications for both the scientific community and the networking community, neither of which have any experience with user management of network capabilities.

Further, this sort of network capability entails controlling and monitoring components that are literally scattered all over the world. This immediately raises the same sorts of issues that face Grid based workflow systems that use widely scattered computing clusters. Among other things, the issue of the value and limits of automated management vs. the value and limits of human management pervades all of this.

Bill Johnston suggested that this would be an appropriate topic for a joint DOE-NSF workshop, focused on HEP experiment infrastructure, but looking at the general problem.

The Workshop outcome should be a description of the current state and issues of the roles and interactions of computing, storage, Grid middleware, and especially network elements as managed resources for scientific experiments (physical or virtual). The goal is a published report that can serve as a reference and as motivation for resource allocation for infrastructure and for computer science and networking R&D.

Workshop sponsors might involve Craig Tull (DOE/HEP), Mary Anne Scott (DOE/Computer Science), Jim Whitmore (NFS/HEP) and Kevin Thompson (NSF/Computer Science). Also OSTP? Who?

Notes:

Network Description Files for GLIF

-------- Original Message --------
Subject: Re: [GLIF tech] Network Description Files for GLIF
Date: Fri, 10 Jun 2005 02:34:38 -0700
From: John Graham <johng@nosc.ja.net>
To: Steven S. Wallace <ssw@anml.iu.edu>
CC: <tech@glif.is>, John.Dyer@terena.nl

```
Hi Steve - This is a clever idea; and I was thinking how readily
extensible it would be to allow for a richer data set and specifically
whether different networks and countries might need to express different
metadata(?) and how this could accommodated...

I have been attending the TERENA Networking Conference all this week,
and as is always the case at these events, the trick is to find the gems
that hide in the mass of parallel sessions. I was fortunate enough to
happen on a talk by Victorian Giralt (University of Malago, Spain) who
was describing how they have implemented a customizable staff directory
using LDAP classes and attributes and OpenLDAP Access Control Lists.
This allows a staff member to maintain a single set of directory date
and to construct rules that determine who can view which columns. So I
might want to publicise my cellphone address within my organisation, but
hide it for external queriers. I'm no expert in X.500 stuff, but I
wonder whether your concept might not be suited to the flexibilty
offered by OpenLDAP?
```

John Dyer, the Chief Technical Officer, was chairing the session and agrees there is some mileage in this. I've copied John on this message in case he isn't part of the GLIF Tech list.

Victoriano's slides can be found here:

http://www.terena.nl/conferences/tnc2005/programme/presentations/show.php?pres_id=89

I'll be attending iGrid 2005 and possibly also the Internet2 Autumn member meeting in September. In the meantime, hopefully this message will continue the discussions.

Cheers

John

Steven S. Wallace wrote:
> Here's something I'm working to address similar needs for the Quilt. I
> think this could scale internationally.
>
> Here's the idea:
>
> all.network-map.net. TXT "us-uk-nl-ca"
> ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
> FQDN TXT record showing what country codes contain network map data
>
>
>
> all.us.network-map.net. TXT "IN-OH-FL"
> ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
> FQDN TXT record showing what states (in the US) have network map data
>
>
>
> all.oh.us.network-map.net. TXT "TFN"
> ^^
> FQDN TXT record showing the name of network(s) in Ohio
>
>
> 00001.TFN.oh.us.network-map.net TXT "Link=Columbus, OH-Cincinnati, OH"
> 00001.TFN.oh.us.network-map.net TXT "Status=Existing"
>
> 00002.TFN.oh.us.network-map.net TXT "Link=Columbus, OH-Chicago, IL"
> 00002.TFN.oh.us.network-map.net TXT "Status=Existing"
>
> ^^
> These are records that specify the network segment endpoints and the
> segment's status. An application would fetch sequential records until
> it hit the last one. In this example OARnet (they run the TFN) would
> be delegated authority for TFN.oh.us.network-map.net.
>
>
>
> Each RON would simply maintain their own delegation (zone file). My
> lab would make an application available that walked the DNS data and
> constructed a map. Since there may be multiple TXT records for a given

> FQDN, we could a small set of attributes for each link.

>
>
> What do you all think?
>
>
> Steven
>
> On Jun 2, 2005, at 12:37 PM, Jeroen van der Ham wrote:
>
>> Hello everyone,
>>
>> In the last two months we started working on a way to make a
description
>> of our network and resources at Netherlight and Lighthouse,
>> Amsterdam. We
>> had several goals in mind for making such a description:
>> - For ourselves, to have a machine readable overview of our network,
>> possibly including its current (automatically extracted)
configuration.
>> - For possible users, so that they can see what we resources we have,
>> possibly what is available at the moment. The description also
contains
>> the list of possible entry points to our network. Ultimately, these
>> entry
>> points should point to the network description of the relevant
provider.
>> - Problem detection, if you have a complete overview of the
>> configuration
>> of your network, written in a machine-readable format, configuration
>> problems can be detected more easily. If every network owner were to
>> publish this in a predefined format, then this can also be used to
>> detect
>> inter-domain problems, by pointing at the configuration file of the
>> connecting network in the relevant places.
>> - Publication in GLIF, to make lambda path setup between participants
>> easier or even completely automatic.
>>
>> Yesterday Cees de Laat pointed us at the discussion currently going in
>> the GLIF community regarding the visualization of the networks.
>> We think that if everyone in GLIF (automatically) publishes a
>> description
>> of their network in a predefined machine readable way, this
information
>> can be gathered by applications more easily. This will make it easier
to
>> solve problems like the ones Kees Neggers raised in a recent mail
>> (creating a repository, generating graphics, handling policies).
>>
>> For example the excellent graphics tool Greg Cole can be used as is,
but
>> the database is not filled with information entered through a web
form,
>> but instead automatically by crawlers reading the network
descriptions.
>> Such crawlers can also be used to periodically collect the network
>> information and put this in a single repository to provide an
overview.
>> On the other hand with a machine-readable format, it is also possible
to
>> parse the relevant information on demand.

>>
>> We are currently making a draft version of a schema for describing
>> networks, including information about the capabilities, references to
>> the
>> polices and services. We will post this to the mailinglist soon.
>>
>> Bert Andree, Freek Dijkstra, Paola Grosso, Bas van Oudenaarde and
Jeroen
>> van der Ham.
>>
>>
>

--
Dr. John S. Graham
UKLight Technical Manager

University of London Computer Centre
20 Guilford Street
London
WC1N 1DZ