Technology Update
Salon B
Wednesday, April 23, 2008
8:45-10 am

Eric Boyd, Internet2, reviewed the agenda for the meeting: a brief overview of the status of technological activities within Internet2 and then turn the meeting over to three community members to discuss how they are using different community-developed tools.

Eric gave a brief overview of Internet2's CyberInfrastructure effort, which is based on the NSF report and requests for support to members. He noted that Internet2 is focusing on building distributed R&E systems – using existing applications, exploiting existing tools and enhancing the inter-domain aspects.

A brief Technology Update included information on middleware infrastructure (focus is on inter-institutional collaboration and scalable authentication/authorized access to remote resources via such tools as Shibboleth, Grouper, Signet and CoManage), performance tools (such as BWCTL for throughput, OWAMP for latency, NDT for common last mile problem diagnosis, and the perfSONAR framework, including the new perfSONAR-PS, a specific software package written in Perl, designed for east of installation and use).

Carla Hunt, MCNC, presented information on the K20 E2Epi, leveraging the I2 tools and infrastructure. She provided an overview of MCNC, the E2Epi, and th eK12 E2Epi pilot. This includes lessons learned from the pilot, revised approach and a proposed architecture that leverages I2 tools.

MCNC owns and operates the NC R&E network; the K12 part of the K20 vision – about a ¼ way through connecting 15 public school systems. MCNC plans to connect 15 of 115 LEAs this year. Key projects include performance; the K12 architecture is locally controlled but parts of it are provided by MCNC. The K12 part needs network performance data and tools to diagnose network performance – MCNC wanted to expand the scope and ability to assess application readiness of each schools network infra structure.

The pilot included a few school systems that were good sample pool – this was NOT a simple arch to implement (one firewall case, another had a non-standard CPE device, and most tech directors in school systems were very busy and didn't have much time to interact with MCNC staff). The new revised approach includes a centralized architecture that polls the LEAs from a central point. This might include using Internet2 tools. The proposed architecture includes a CD image that can be distributed to the LEA (easily installed on any server or computer, simple installation that only requires user to boot from CD, will work on most hardware, and supports remote configuration and management) that include specific tools – for latency, loss, and utilization – and might want to include flow data in future, want to map to perfSONAR, support for federated ID management (with a GUI that is easy to use).

How would it work? – CD installs OS and tools; once installed, automatically creates northbound connection to central server. To get through NDDs, server establishes IPSEC tunnel

to central reporting server. Centralized server would have a GUI interface with a dashboard-like view of the network interface. Centralized server is a VMP concentrator that gives out private IP addresses to endpoints. Next steps are using the Internet2 NPToolkit for endpoint at schools, and leveraging perfSONAR data.

Eric Boyd provided some basic information on the Internet2 DCN provisioning system and then introduced **Dale Finkelson, UNL**, to discuss his experiences in using the DCN with the LHC CMS physicists.

Dale noted that he has maintained flows of 6-8 Gbps on a regular basis. He says his bosses are only concerned about getting their mail and that they can use the network ("it works"). My vision is that the administrators won't think we're making the network 'work' if physicists are using 8+ Gbps of my 10G backbone at any given time. What the DCN allows me to do is to set those flows aside – off the production network – UNL has increased demands for their production networks, so it isn't 'ok' for them to be, at any time, tied up. Dale reported that the police get very nervous if their cameras all go offline because someone is running an 8 Gbps flow so he put a few more switches in place and rerouted the physics data onto the DCN. In a few months, he'll have a few more locations and applications going on line, so he'll need more switches, more fiber, and might have to do some scheduling but the bottom line is that what the technology lets him do is to move these disruptive (but useful) flows onto another infrastructure so that it doesn't disrupt the flows he's expected to have available 24/7.

"I wouldn't care if UNL used 8 of the 10 flows and left the other2 for the remaining states but that doesn't make them happy. This separate infrastructure is a 'win-win' for our campus. In a few months, LHC Atlas is probably going in this way – which will be a 'win-win' for them (they won't have to buy multiple 10G connections)."

This isn't to say it is easy – yes, Dale has to lay more fiber, add more VLANs, get authentication and stuff. But there's no way he could have afforded to run fiber all the way to Chicago to support this one physics project so it has been *all win* for UNL.

Eric gave an overview of Phoebus, a tool developed by University of Delaware researcher Martin Swany, and introduced **Bill Owens, NYSERNet**, who has been using Phoebus to improve connections across Internet2 to Syracuse University.

Bill reports that NYSERNet wanted to bring DCN to many of their clients – many of whom cannot afford either a connection to the DCN or high-performance networks on a regular basis. NYSERNet uses Phoebus to increase performance for each user; it is not 'easy to use,' just yet, but it is providing all of their clients with greater performance. Bill gave a brief description of how Phoebus works. The Phoebus Gateways talk to each other – the one in the NYSERNet offices talked to the one at the edge of the DCN node. The key to Phoebus improving performance is turning 1 long path into several shorter ones – keeping the IP connection on the ends to as short a distance as possible. Pair of short TCP connections – edge to laptop on each end – and a 'mystery' connection from Gateway to Gateway that doesn't obey the usual TCP rules – and doesn't need to, so there is greatly increased performance over that portion of the

path. And, due to the need for short TCP connections from laptop to edge Gateway, those paths can be fine-tuned to ensure best performance.

NYSERNet is intending to add its own IDC into the mix; this won't affect the Phoebus connections but will allow them to request circuits.

Phoebus does NOT need the DCN – it worked very well before hand – because it allows you to traverse between 2 Phoebus servers that are well tuned (edge to well-tuned server to well-tuned server to edge speeds up the process). Once the Phoebus Gateway is setup, it understands the paths to which it has access, handles authentication, etc. The only limitation Bill sees is that there is only one entrance to the DCN via a Phoebus Gateway; changing this will take some work but it is on the plate.

Q: ? (quiet and not repeated)
A: There are 2 different ways to do that – there's a shim library that works transparently with the Gateway and, recently, an XIO library that can be loaded into GridFTP.

Q: How much more complicated would this be if there was another network cloud in the picture.
A: once the IDCs have been configured to know what paths it can traverse, they can request circuits.

Q: what sort of traffic can go through these paths? Compressed video?
A: (Bill) the only thing is TCP – Phoebus was designed to improve TCP performance over the long haul.
A: (Dale) if you want to send video, you want DCN connection. If we use the IP connection, it goes through 8 routers; with DCN, it goes through 2 – so the throughput is much better. If latency is an issue, you can certainly create lower-latency streams via DCN. The other people really interested on UNL campus are the fine arts folks – they want to be editing video and are asking for 10G connections to each building.
A: (Joe) you are not getting latency but some jitter; not tested with DCN.

Eric commented that monitoring and debugging the Layer3 network is hard enough – we'd just gotten that straight – and now we have a whole new arena for performance testing!

Q (for CH): How big a load to you see the hand-holding in NC do you see this going to be?
A: We've just started working with this community – they don't have lots of time and we're going to have to see what we can do with as little of their time as possible.

Q: (Brent) ? (question too quiet, not repeated)
A: (Dale) I don't want to be involved in every researcher's desire to get some flow going – so if this really works out and is extensible – and in enough places and has value for a broad a community, the plan would be to make it has dynamic and user-controlled as possible. A number of people are starting to see that this is the direction (service offering) and not related to researched.
C: you can route around firewalls?
D: Yes, you can get around the network-disruptive devices

Q (for EB): You mentioned some things that have commercial interest, and those that aren't 'going away' like perfSONAR. What do you see is of 'commercial interest'?
A: Most of the 'commercial interest' stuff is in middleware and Ken Klingenstein would be the one that would decide if he would let that go out to private development. All these technologies could, eventually, be swallowed by the commercial world, but we have to decide if it make sense to out-source it or run it ourselves.

Meeting ended at 9:55 am