

DCN Case Studies  
Salon H (netcast)  
Wednesday, April 23, 2008  
1:15-2:30 pm

John Vollbrecht introduced the speakers and indicated that they were going to speak about their experiences in using different applications with dynamic circuits.

**Bill Owens, NYSERNet** – “One application that we have a researcher using DCN with is the LIGO project; we certainly think we have other projects, such as the LambdaStation and TeraPaths projects.” Bill reports that they have been polling folks in NY state to see who needs to use dynamic circuits. NYSERNet decided to go with a dynamic configuration – in NYC, where we connect to the DCN as one point and Syracuse (where Bill is located) as the other. Syracuse and Rochester are both interested in large transport issues (video and other) as well as a researcher in NYC who is doing high-speed data transfer research (nothing to push across the network but wanting to understand how data travels).

Bill spent a lot of time getting the software deployed – the software was free and came with better support than you get with most commercial software – and, while he ran into some difficulties, these were overcome and he is quite happy with the direction NYSERNet is going.

**Phil Demar (Fermi)** – Fermi has been using static circuits for a while but are now moving into the dynamically-allocated circuits. They have a functional dynamic circuit system deployed at Fermi, University of Nebraska-Lincoln (UNL), and Caltech. The interface is to ESnet and then to Internet2’s DCN for requesting dynamic circuits. Using Nebraska as a test case – Caltech is nearly identical – there is a routed path between Fermi and UNL; the “dynamic” capability between those locations is really between the points via Internet2 and ESnet.

Since Fermi created LambdaStation, they are, of course, using it – it does local modifications as well as setting up/tearing down the circuit. You have to provide it with the location details for end points. Characteristics include selective forwarding on a per flow basis, graceful cutover and fallback on path changes – physicist can start his transfer, the circuit will switch the data over transparently as needed – on-demand from applications (i.e., the AA is built in).

“When we started, we could make the LambdaStation call for circuits, but it was a bit high-maintenance. We’ve been experimenting with the idea that the network should be ‘aware’ instead of having the applications be aware. We adapted the network to recognize the need, call LambdaStation to request/setup a circuit.” Phil provided a demo of the sequence of events when LambdaStation calls a circuit.

They can run circuits concurrently and monitor them at the same time – this was a bit difficult but they’ve managed to graph this along with the IP traffic. He reported that, as more and more folks want to use the software, Fermi will have to adopt this for other users – multiple platforms, enhance interaction with OSCARS/ESnet and DCN/Internet2, develop management configuration tools, convergence with related Brookhaven Terapaths software and provide better documentation.

**Artur Barczyk, Caltech** – Working for US LHCnet, he is connecting US Tier1 locations with CERN. He reports that they have four PoPs – CERN, Starlight, ManLan, and Netherlight that have a 80 Gbps by 2010 – a requirement of LHC effort. Artur showed the data grid hierarchy. “We’re not just streaming data from CERN to the other Tier1s – T1s are responsible to collating and reprocessing the data.” T0 > T1 (reprocess)>T2 >T3>T4s. Very dynamic environment – there is a great need for DCNs because the data can saturate the IP connections very quickly. T1 > T2 allows the data to be redistributed for analysis – this is hierarchical in that they redirect the resources to the highest-level requestor at any given time.

**All** of the data transfers will be multi-domain transfers – the simplest will cross only three domains. Interent2’s DCNSS was tested for a demo here at SMM transferring data from Caltech to CERN – while this demo did work, it isn’t a *realistic case* since the traffic would flow in the other direction, with a T1 in the middle.

**Inder Monga, Nortel** – Inder talked about the history of dynamic applications – “the first application we heard about was in 2001, where nightly data backups began using lambdas.” Then, DRAC-driven hybrid network appeared. (more data; currently unavailable)

**Martin Swamy, University of Delaware** –Phoebus and how this works with DCN was the subject of Martin’s presentation. The important issue is that, once the data goes through the Phoebus Gateway at the edge of the DCN, it is not tied to TCP and often is using rate-based transport protocols. It goes back to TCP when it leaves the Gateway at the end of the DCN link and travels to the end-user location.

Martin gave an overview of the Session Layer approach – an end-to-end composition of the segment-specific transports and signaling. Martin talks about the simple Phoebus authentication system. He talked about ways different apps could use Phoebus – this is enabled by the Phoebus library. You could change all your socket calls () or, use the Linux “LD-PRELOAD” to automatically do a function overwrite to change the socket calls – w/o you having to do anything! This doesn’t require you to recompile to begin using. The important other way doesn’t require you to change anything – users shouldn’t have to know that anything is happening; they’re just getting better performance.

The Phoebus XIO driver is taken from the GlobusXIO driver (from the Globs toolkit) greatly simplifies using the Phoebus gateways. This allows 3<sup>rd</sup> party setup. He also gave an update on the use with GridFTP and how they are available through all but the Windows users.

#### **Q&A:**

Q: what kind of users are going to use it first? What kind of policy is going to be needed? What is necessary for general use? What doesn’t work now and needs to be fixed?

A: (Bill) We need patient users – there are glitches – use case includes anyone who needs to send massive data files (vs. shipping hard drives around!) – think there will be a lot of cross-over into the arts, who are interested in HD video streams, high-bandwidth, multi-channel experiments.

A: (Phil) Answering through LHC-colored glasses – anyone who is sending large blocks of data – big science or anyone who needs guaranteed bandwidth (high or low). Policy – there are some

concerns – right now, the technology is way ahead of everything else but no one is thinking out the policies that'll need to be implemented. Actually, all non-technical issues as well as the long-term tech perspective, we're not thinking that we're just trying to get it up and running and working. Phil also pointed out that DCN was working in 'early adopter' mode, with lots of activity, bugs being found and fixed in new applications, number of "believers" growing, etc.

A: (Artur) Also very LHC-biased, so computing centers and high-bandwidth applications. Even within LHCnet, Atlas is performing online calibration. That, for this, users need protected, but lower bandwidth, paths.

A: (Inder) Ultimately, we're talking about a scarce resource, so you need to have a policy and tools to manage those resources. A high blocking rate would be detrimental to using

A: (Martin) Rather than negotiating the details, you want to be able to say what you need to transfer by when and allow the system to schedule.

John noted that there needs to be some way to protect the user from having a DDOS – IM agreed and said there needs to be a way for them to pay, for you to be sure they are using what they ask for, system of checks and balances.

John asked Artur about charging or proving 'use' or hierarchy of use. In packet networks you? Artur noted that there needs to be any policy that allows the system to 'rate' users to be sure to block abuses.

Q: Scheduling tools – could those be built in?

A: Yes, they are looking to incorporate them.

John asked how people get access – Phil and Bill were talking about on-ramps. What other options are folks thinking of. Bill felt that, because they felt there was a need to extend the access to the DCN in NYSERNet, there are more options for their campuses – it seems as though there is more of a need for static switches. Phil felt that the 'Q' in QoS stands for 'quantity'.