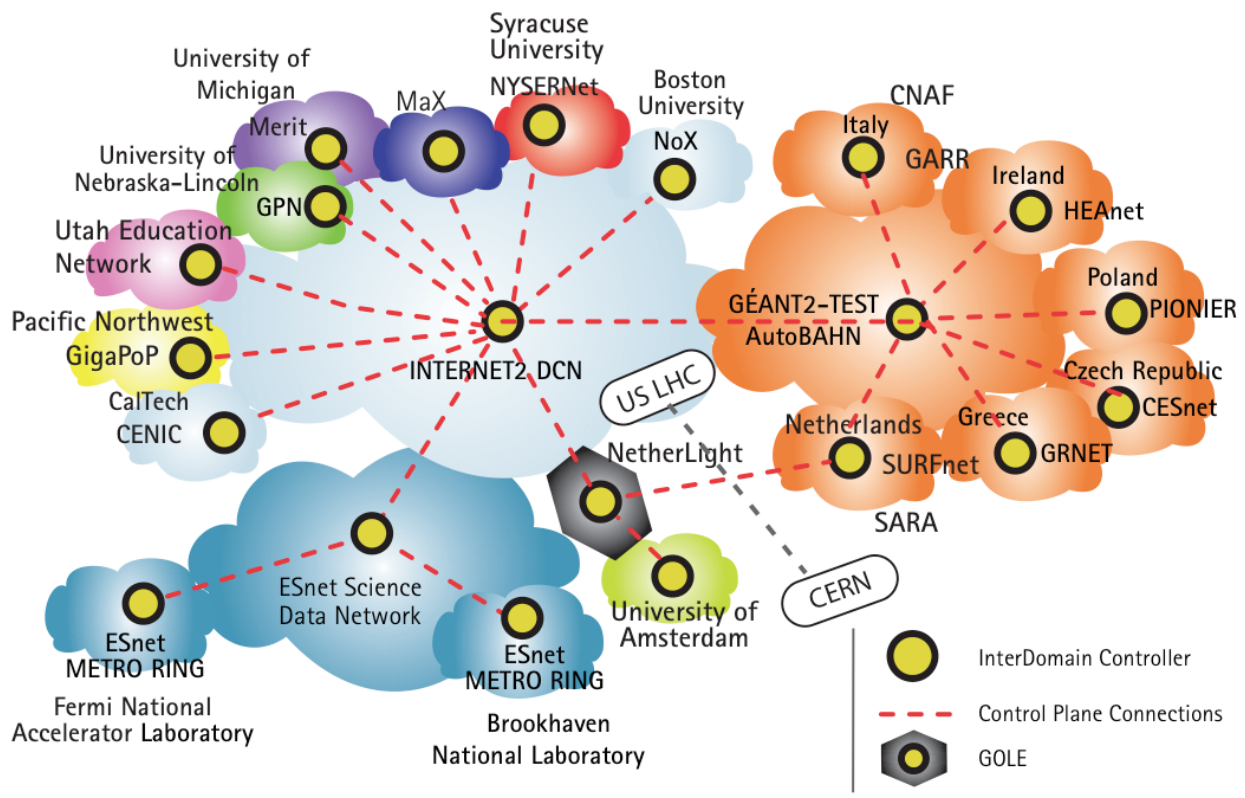


Using the Dynamic Circuit Network (DCN): A Brief Tutorial
 Salon B
 Monday, April 21, 2008
 1:15-2:15 pm

Brian Cashman showed the current global circuit network (what's going on around world) and briefly described various ways in which people are using dynamic circuits.




Brian gave a brief presentation on:

- What is DCN?
- How to connect
- How to request circuits

What is DCN? – The Dynamic Circuit Network (DCN) system creates peer-to-peer circuits, whereby the end-user or application sends requests for connections. The control plane:

- IDC (InterDomain Controller) – accepts requests and coordinates circuit requests with other domains
- DC (Domain Controller) – controls circuits within its own domain.



Version 0.3 of the DCN Software Suite (DCNSS), released on April 18, includes:

- OSCARS IDC (open source project maintained by Internet2 and ESnet)
- DRAGON DC (an NSF-funded, open-source project maintained by MAX, USC ISI East, and George Mason University).

Q: Any interoperability issues between DRAGON and other software packages?

A: No, they're independent of each other.

Q: Does the equipment need to support GMPLS?

A: No.

HOW do I connect? – Physically, you need an Ethernet port at each end. You also need control plane software that allows you to request circuits – usually using the IDC software.

Internet2 members connect via Internet2 Connectors; Universities and Campuses connect to the Connector who links them to the Internet2 IP network. This does not require the Connector to have an IDC. There are two options for this connection:

1. No local IDC
2. Install local IDC

For example, connections to DCN at Boston University and UNL are made without a local IDC – they use Internet2's IDC. For Option 1, you make a request to the nearest IDC (e.g., Internet2's) – note that the initiator of a request isn't necessarily an end-point.

Q: By statically configuring, you're saying you've got a preset lambda; is this a 'mother may I?' situation?

A: This would be a VLAN but, yes, in a static configuration the connectivity is preset.

For Option 2, install a local IDC (using for example the IDC that is part of Internet2's DCN SS). The example is that network1 installs its own DC and IDC so the end-user makes a request to the

local (network1's) IDC. Each DC has control over its own domain – if one domain doesn't want others to connect, the request stops there.

Q: What if the network in the middle (network2) didn't have an IDC?

A: Then end network is connecting IDC to IDC – middle network connection would have to be a static connection.

Q: Would Internet2 entertain a request for a connection directly to an individual user?

A: No. We'd contact your Connector to try to make it happen.

Requesting a Circuit

- Who? (Clients, requests, and interfaces)?
- What ?
- Where?

Who makes a request?

CLIENTs

- 1) User-initiated – using the OSCARS web page or simple command-line tools
- 2) Program-initiated – using the Web Services API. Examples of this are: Phoebus, LambdaStation and TeraPaths. Each program transparently requests circuits upon data transfer initiation. Alternatively custom applications can be written using this API. .

REQUESTS (What is needed to make a request? Includes When? and Where?)

Need a minimum set of information (source, destination, start and end time, bandwidth and description) and additional fields (VLAN range, path, etc.).

INTERFACES (How does the requestor make the request?)

Web user interface (WBUI) is currently a java serv-let interface used by the OSCARS web page; it is not intended for use by other applications – this is a 'proof of concept'. Another option is the Web Service API (an xml-based API intended for use by applications, such as Phoebus, LambdaStation and TeraPaths).

WHAT method do I use to request a circuit?

Brian provided a demo of the WBUI On-demand Secure Circuits and Advance Reservation System web pages and the command-line tool for creating a reservation. Brian provided instructions for folks to try out the online WBUI form and request circuits.

In quick summary, to use the Internet2 DCN, users need a physical connection as well as control plane software (campuses need to go through their Connectors). DCN software is available to install and run – and users don't have to run a local IDC to use DCN. An API exists to help modify applications to use DCN.

WHERE can you learn more? Brian provided pointers to:

- DCN SS – <https://wiki.internet2.edu/confluence/display/DCNSS/Home>
- Java Client API – <https://wiki.internet2.edu/confluence/display/CPD/OSCARS+Client+Java+API>
- Test IDC Guide – <https://wiki.internet2.edu/confluence/display/DCNSS/Internet2%27s+Test+IDC>
- Obtaining a Test Certificate – <https://wiki.internet2.edu/confluence/display/CPD/How+to+Request+an+IDC+User+Certificate>

Q: You mentioned 3 applications that use this API – what are they?

A: Phoebus, Terapaths and Lambda Station.

C: Phil Demar (Fermi) is going to present more details on LHC Applications at the Tools Tutorial session.

Q: Are there any costs?

A: For Connectors to connect to DCN there are no costs at present; whether Connectors charge their users is their decision.

Q: What does a Connector need to enable clients to use this?

A: Simply a connection to Internet2's DCN (Ethernet connection), best would be an Ethernet switch. At its simplest, scientist > campus > regional could be static switches; when they hit Internet2's DCN, there could be dynamic circuits across the backbone.

Q: How do you get data to Fermi?

A: There's a demo today – static to the Internet2 network; dynamic across Internet2 and ESnet; and static or dynamic out to Fermi (depending on their situation).

Q: How do I know that I'm not getting a connection left over from some lab that they shouldn't be connected to? How do you make sure both sides have said 'yes' to this connection? What is the authorization method?

A: When the request comes in, each side has to say 'yes'. Basic trust model in place now – when request comes in, the network affected has to determine if that user is allowed to have that bandwidth. End-points have to determine this before they say 'yes'.

A: Security aspects are lagging behind the technical ones – LambdaStation hasn't needed to add policy and security issues because of our trust model.

Q: What are the methods for scheduling and policy?

A: It is a schedulable resource – if someone comes in subsequently, they're out of luck. There's no system, yet, to see if the 2nd person should have preemption rights, higher policy, etc. This relies upon a centralized domain for scheduling. For multi-domain scheduling, you need to have all the domains working with each other.

Q: Where can people get involved to affect policy and direction of this effort?

A: The NTAC has proposed a Working Group (which is meeting for the first time tomorrow) and there are several other openings – talk to council members, come to sessions this week, come to the NTAC meeting and give input.

Q: Is there any work being done on data plane feedback – i.e., if this connection hasn't had any traffic in 15 minutes, hang it up?

A: Yes, there are folks working on that – that gets into the perfSONAR monitoring we're looking at but right now it is not as fluid as automated monitoring. Also, there's a policy issue for the community – do we say this is a scarce resource so if you're not using it, we take it away? Or "you pay your \$, you get to use or not use it as you wish."