DCN Around the World
Salon B
Tuesday, April 22, 2008
1:15-2:30 pm

John Vollbrecht opened the meeting by talking about the world-wide effort to interconnect DCN efforts. He noted that there are handouts and poster sessions about this topic throughout the meeting. He showed the DCN Weathermap. He identified all the connections and said it was Interent2's intention to leave these connections up.

John introduced the three speakers who would provide presentations on their recent experiences:

- Afrodite Sevasti – GRNET
- Joe Metzger – ESNET
- Tom Lehman – USC ISI East

Afrodite Sevasti, GRnet, presented what is going on with GEANT2 (of which GRNET is a member) – the AutoBAHN network is intended to be a research activity for engineering, automating and streamlining the inter-domain setup of guaranteed capacity Gbps) connections. AutoBAHN is an overlay of GEANT2.

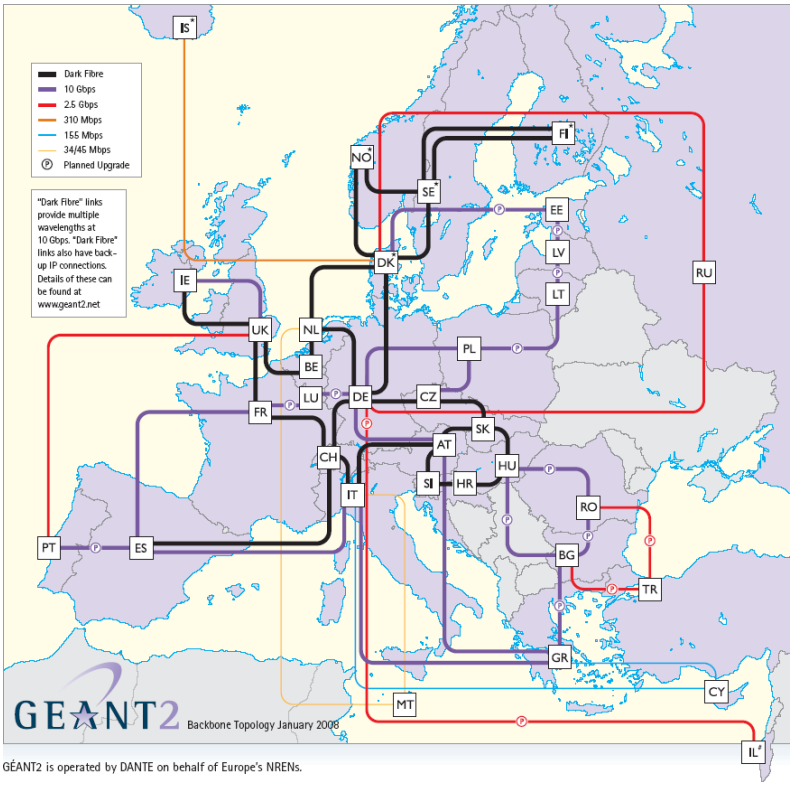

End-to-end paths around the world – in the EU, there are over 7 different administrative domains, with different data plane technologies, systems, and protocols in place. How could they make this interoperable? The approach was to create an environment and in which control and provisioning has to be distributed; business-layer related interactions include AA policies and advance reservations, etc. In addition, privacy and control of intra-domain resources must be safeguarded and separate from each other. Afrodite gave an overview of the AutoBAHN system, identifying how they met the system requirements.

Afrodite spoke about the current deployments (6 networks) and about their attempts to make this more widely available. AutoBAHN is the work of a subset of the R&E networks in the EU – a dozen networks and Dante as a coordinator of the effort.

Joe Metzger, ESnet, spoke about the OSCARS project – he gave a brief overview of the effort and a description of how it is being used. This included a description of the mechanisms by which OSCARS provides DC services. He showed a map of the ESnet4 IP and SDN (dynamic) networks and reported that they have about ½ the new SDN routers deployed. They expect to be bringing up the remaining routers to allow dynamic circuits everywhere in ESnet during the next year.

They are working closely with DICE to ensure interoperability, being involved in the IDC development and deployment, as well as continuing to collaborate with standards bodies like OGF and GLIF.

Tom Lehman, USC ISI East, gave an overview of the I2 DCN, focusing on the software suite, DRAGON control plane, and OSCARS IDC.

DRAGON is a GMPLS control plane that contacts the CoreDirector to setup circuits on the fly. Key elements include the VLSR and NARB. To allow others to interoperate, what do you have to do? There are a set of .wsdl requirements (see the web service definitions at: https://wiki.internet2.edu/confluence/display/CPD/OSCARS+Web+Service+Definition).

Tom showed how hybrid networks control planes interoperate:

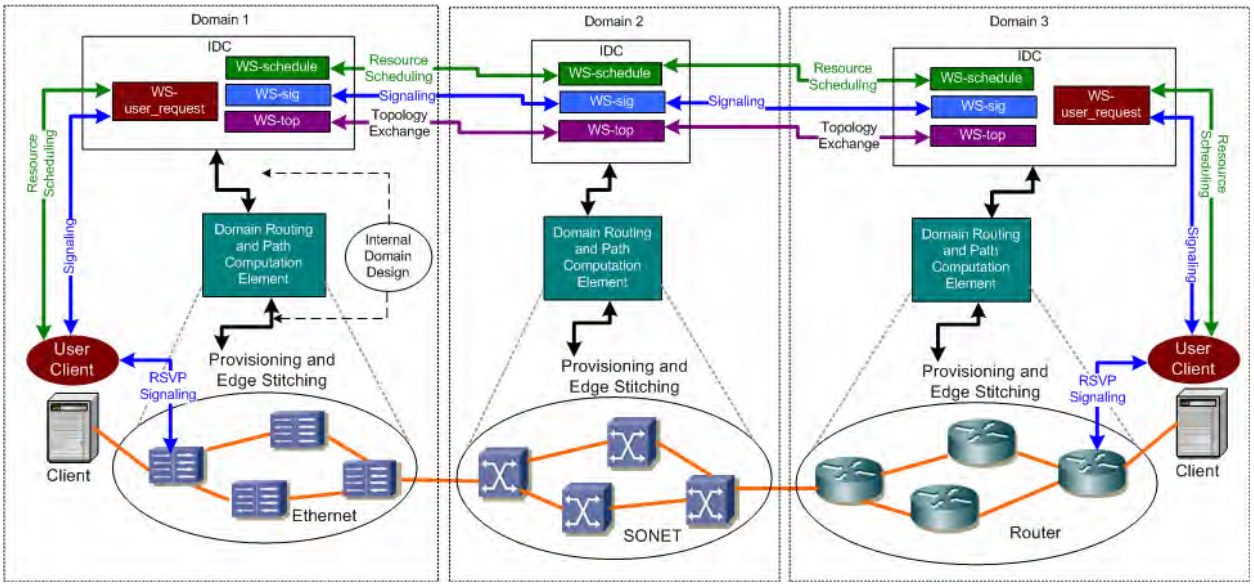

Tom noted that this is all very new, working in Layer 2. Trust models are required. He noted that the web services request location had seen over 300 hits/requests in the past few days.

Q: if this is a resource that is expected to get more and more use, how do we scale the scheduling?
A: local domains are scheduled but inter-domain is distributed. Initiating domain will suggest 5-10 possibilities but it will continue to be a distributed request system.

Q: Inter-domain information, is there a common understanding about what the abstraction should be like?
A: Advertise our topology using the same schema – how often do you do updates/ are the incremental? We're still discussing the topology protocol but at least there is a similar schema so that's a good base point.

Q: How do you determine how much information is provided?
A: each domain has the right to show as much or little information as they choose.

Q: if you are in 1 domain, can you look to see what the path would be – could you specify which path you'd prefer?
A: the hooks are in the web service request so, if folks knew about it, they could do so – at this point, we pick a path we think would work but you could specify a path.

Q: By putting servers for testing at all points in the network you can see whether your path is clean. Some of the BW is too expensive to give away for general use – it seems you need to determine if the requestor can do the transfer.
A: That calls for the infrastructure to monitor the requests and make decisions. This has been discussed but no action taken.

Q: Please say something about how to standardize beyond DICE – or do we need to?
A: (Joe Metzger reported that) the perfSONAR Working Group has decided to move the protocols we've been working on into the OGF; the DCNs are using protocols that could also benefit from being discussed in the OGF community.
A: (Afrodite Sevasti reported that) this is mainly applications data plane layer, for folks sitting at the NOCs – from that point-of-view, it would be nice to address some network engineering issues, perhaps at IETF.
A: (Tom Lehman reported that) we need more acceptance in OGF and GLIF – there is a lot of work going on in a variety of communities, via working groups, that are not parallel efforts but are concurrent developments. We can take a lot into the OGF/GLIF but we're going to see some other ideas, as well.

Q: if you're going to move these into production, you're going to have to charge this back to the end user – with multiple networks around the world, how are we going to standardize the pricing of these services.
A: (Afrodite Sevasti stated that) we are network engineers so we're primarily interested in the development of the tools/services but, as demand rises, someone else is going to need to figure this out.