

IDC: Inter-Domain Controller API

GLIF Control Plane WG
Jan 19, 2008

Evangelos Chaniotakis (haniotak@es.net)
ESnet Computer Software Engineer

Chin Guok (chin@es.net)
ESnet Network Engineer

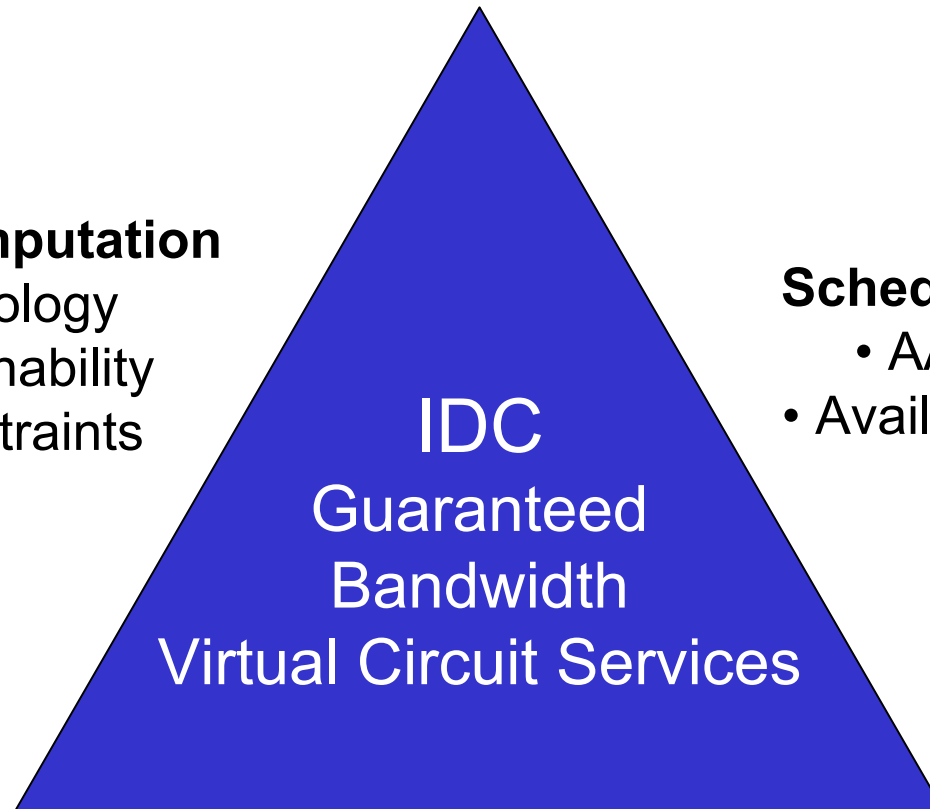
Lawrence Berkeley National Laboratory

Path Computation

- Topology
- Reachability
- Constraints

Scheduling

- AAA
- Availability



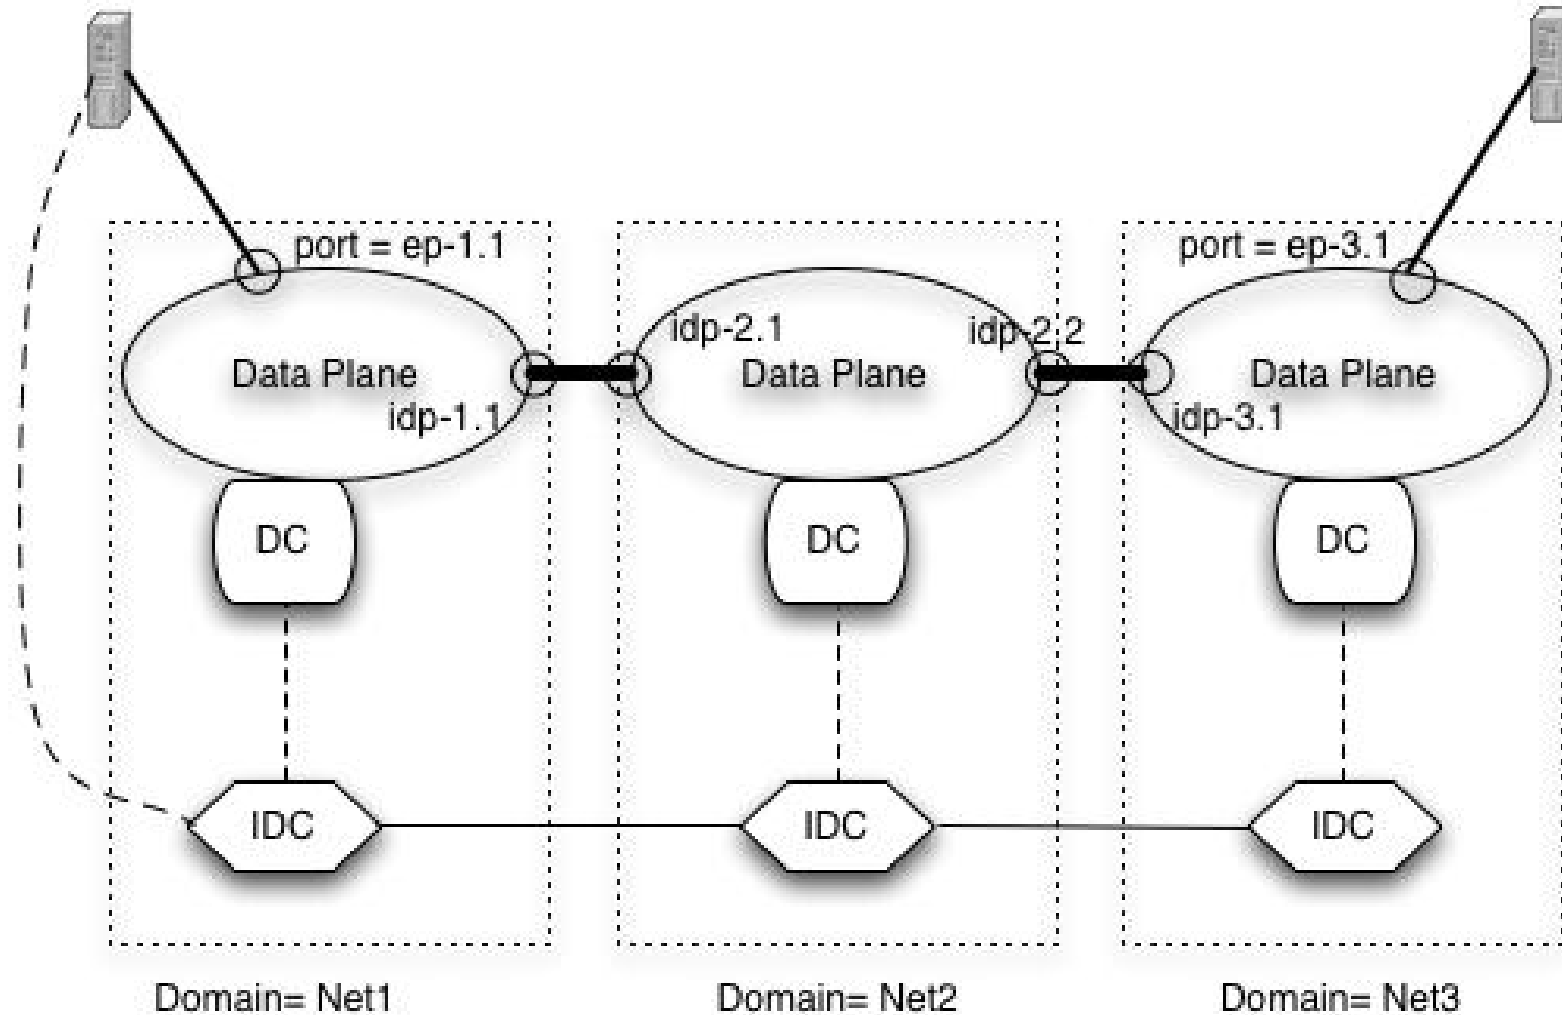
Provisioning

- Signalling
- Security
- Resiliency/Redundancy

➤ IDC Role

host1 -> Server1.schoolA.edu
-> domain=schoolA; node=Server1

host1 -> Server1.schoolB.edu
-> domain=schoolB; node=Server1



- **Design and implementation partners:**
 - DICE (DANTE, Internet2, CANARIE, ESnet)
 - HOPI / DRAGON (Internet2/ISI East)
 - OSCARS / Science Data Network (ESnet)
 - AutoBAHN (GEANT)
 - Network Management Working Group (NMWG)
 - Dynamic Resource Allocation Controller (Nortel)
 - Phosphorus (University of Amsterdam)
 - TeraPaths (BNL)
 - LambdaStation (FNAL)
- **The service has been deployed and is in pre-production**
 - Internet2, DRAGON and ESnet SDN infrastructure managed by IDC instances
 - More than 3000 reservation requests have been processed
 - Large scale interoperability demo at SC07
 - On-going work with TeraPaths, LS, DRAC, UvA interop

- **Core Technology: SOAP / WSDL**

- Ensures wide compatibility and easy adoption
- Message security guaranteed by SSL
- Data types conform to NMWG schema

- **Typical usage scenario**

- User performs a SOAP call at their local IDC instance,
- SOAP call is locally processed and forwarded downstream to peer IDC instance as needed,
- Message keeps going downstream until last IDC on path reached,
- SOAP responses start coming back upstream,
- Local IDC returns appropriate SOAP response.

- **Resource scheduling:**

- Create new reservation (immediate / in advance)
- Cancel existing reservation
- Query status
- List all / by criteria

- **Data plane:**

- Path setup and teardown (synchronous)
- Path refresh (end-to-end data plane status check request, implicit teardown on failure)

- **Topology exchange (IDC-to-IDC only):**

- Retrieve local IDC's view of the network topology
- Signal IDC to initiate topology pull

- **Reservations:**

- Creating a reservation is very opaque to the user. Little information on what they can do, and limited feedback.
- User can't change reservation parameters; must cancel or wait to expire then resubmit. Service is interrupted.

- **Data plane setup:**

- Network operations are slow & expensive; devices may not be able to keep up. We often observe up to 3-4 mins of setup time. Requests would occasionally come in during SC07 at a higher frequency rate.

- **Topology exchange:**

- Some domains will only want to expose an abstracted view of their local topology, and only to trusted peers.
- How do we handle topology changes / maintenance / outages?

- **Operations:**

- Difficult to debug reservations across multiple IDC instances

➤ Proposed API changes

- **Reservations:**

- Modify reservation parameters (duration, bandwidth...)
- Detailed reservation history
- Constraints checking

- **Data plane setup:**

- Path setup / teardown: asynchronous, scheduled by user and/or IDC

- **Topology exchange**

- A comprehensive model is being designed by the DICE group (lead by Mauro Campanella and Jerry Sobieski)

➤ Other Design Issues

- **Technologies**

- Considering RSVP for path setup and maintenance rather than SOAP messages; how do we tie this into AAA infrastructure? What if RSVP is rejected at the domain border?
- Considering using SNMP to catch local network status changes and to then notify peer IDCs.
- Topology exchange and interdomain pathfinding: At first look is similar to BGP functionality; could we leverage existing BGP implementations?

- **Policies**

- Once the service goes in full production, issues like link costs and usage authorization will come into play. How will domains handle this information?

- **Schemas**

- NMWG evolved to NML-WG. IDC engineers are participating. New standards may require changes to API.

➤ Contacts

- **Chin Guok, ESnet (chin@es.net)**
- **David Robertson, LBNL (dwrobertson@lbl.gov)**
- **Evangelos Chaniotakis, ESnet (haniotak@es.net)**
- **John Vollbrecht, I2 (jrv@internet2.edu)**
- **Tom Lehman, ISI East (tlehman@east.isi.edu)**
- **Andy Lake, I2 (alake@internet2.edu)**