

Meeting Minutes from SMM 2011 in Arlington

Performance Working Group Minutes from Face-to-Face Meeting at 2011 Spring Member Meeting, 20-April-2011

AGENDA:

- Internet2 Update -- *Jeff Boote*
 - Indiana University's Data Capacitor Wide Area Network (DC-WAN) – Matt Link, Robert Henschel_
 - Speed Test Tool – [Community Anchor Institution Data Collection and Mapping Working Group](#) - *Randy Stout, Seth Peery*
 - Circuit Monitoring Implementation for OSCARS --*Aaron Brown*
 - Periscope - *Martin Swany and Dan Gunter*
-

Internet2 Update - Jeff Boote

Jeff Boote explained that Carla Hunt and Chris Hawkinson, co-chairs, were unable to attend this WG meeting.

- Internet2 staff members have been involved in recent months with projects relating to NSF grants:
 - IRIS (International Research Instrumentation System)
 - DyGIR (Dynamic Gateway for International Research)
 - DYNES (Development of Dynamic Network Systems)
- Some of the work involves supporting OSCARS (On-Demand Secure Circuits and Advance Reservation System), IDC deployments and perfSONAR deployments in support of those grants
- Largest development effort is getting circuit monitoring integrated with Dynamic Circuit provisioning, to be described by Aaron Brown later in this session
- Ongoing perfSONAR work involves understanding specific needs of IRNC (International Research Network Connections) participants

Indiana University's Data Capacitor Wide Area Network (DC-WAN)

Robert Henschel, Indiana University, presented

Slides are found at: http://www.internet2.edu/presentations/spring11/20110420_henschel_IUdatacapacitor.pdf

- Problem: Handling / moving of large data sets
- Hard to keep track of what is the up-to-date copy
- Often the up-to-date copy of data is stored on someone's laptop
- NSF funded the Data Capacitor project in 2005, and it was implemented in 2006
- IU followed up in 2008 with Data Capacitor WAN, involving nationwide cyberinfrastructure
- Uses Lustre storage, http://wiki.lustre.org/index.php/Main_Page
- Now Data Capacitor (not a researcher's laptop) is the center of the workflow
- Eliminates copying data around
- Research projects served include Gas Giant Planet research, Electron Microscope research, and WIYN One Degree Imager research
- For visualization, IDL software (a commercial application) is used
- Network and file system tuning involves dealing with end point of file system and with the network path in between
- Work with the IU Global Research NOC to understand the network typology
- Tests can reveal firewalls that hold up packets
- Have used Internet2 test servers, have not yet used perfSONAR
- Working with German partners -- Center for Applied Mathematics of Research Center Jülich and the Center for High Performance Computing of the Technische Universität Dresden
 - Using the VampirTrace open source monitoring facility and interactive visualization program, <http://www.vampir.eu/index.html>
 - Use case: 100 Gbit Testbed
 - The data visualization allows drilling down to individual fiber channel ports
 - Can expand the displays, can zoom and look at several milliseconds

Q: Do you look at the visualization when there is problem with the file system performance, or when there is a new wide area link or all the time for quality control?

A: Can capture and visualize data all the time, though we can't link to application in every case. Generally, we run the visualization once when we branch to new location.

Q: For the data capacitor, what network data characteristics are best : low latency? same performance in both directions?

A: Best is when there is symmetric routing, same path back and forth, otherwise it can be necessary to debug on 2 networks. Also, obviously, the lower the latency the better. Performance is important, but we also stress the convenience aspect. Uses TCP for network connection

Q: Is alarming available when there are network performance problems?

A: No, currently it's a diagnostic view. Looking at implementing alarming for a telescope project, where network outages will create serious problems.

Speed Test Tool - Community Anchor Institution Data Collection and Mapping Working Group

James Werle of Internet2 introduced Randy Stout, Kansas Board of Regents, to talk about measuring the broadband capacity of Community Anchors Institutions

Randy Stout, Kansas Board of Regents

- Project underway supports the K20 program goals and supports U.S. UCAN effort
- In K20, there has been collection of connectivity data for SEGPs for several years
- There has been great value in that data
- We need to have more efficient pathways to collect the data and to refresh more dynamically
- With the U.S. UCAN announcement, the K20 Initiative sent a letter to Internet2 CIO Dave Lambert (Sept. 2010) re what they hoped to do to support the work Internet2 will be doing
 - Dave Lambert responded right away, and listed many things the K20 initiative could do to support this, such as:
 - establish baselines data for existing regional networks and the anchors they serve and the Community Anchor Institutions.
 - develop best ideas on how to present the data
 - suggest how to have the data available so members in the network and community can demo applications and the capacity of the networks
- Just before 2010 FMM in Atlanta, a Working Group was formed around this project/initiative
- Working Group activity is tracked on wiki <https://wiki.internet2.edu/confluence/pages/viewpage.action?pageId=2425767>
- Goal is to establish complete, accurate, and consistent baseline data on the current connectivity and broadband capacity of anchor institutions and develop mechanisms for contributing data to similar state and national efforts underway.
 - Team will develop and disseminate an open source speed test tool targeted to CAI sectors around the country to collect actual upload/download speeds and other connectivity metrics.

Seth Peery, Virginia Tech

- Seth Peery has background developing a framework using network performance tools and GIS systems
- Hope is to scale up what was done in Virginia to apply nationwide to community anchor institutions (CAIs)
- Some CAI's are connected to R&E , some route thru Internet2, some are not
- The goal is to get good benchmarking data about the CAIs and the R&Es
- Plan to measure the R&Es separately from the CAIs
- Kicking around idea of forward deploying performance testing devices out into the R&E networks to which CAIs are connecting
- These can serve as endpoints for user initiated tests (NDT for example)
- Where the client is located at the CAI
- Hope is to test the connection between the CAI , their WAN within their network, and out to their R&E
- Then to control for rest of it, have measurement appliances test against one and other
- Idea is instead of having network performance measurement appliances near core, push those out so we test the segments closer to the CAI -- to get benchmarks
 - Conceptual architecture is shown on this diagram: http://www.internet2.edu/presentations/spring11/20110420_peery_caidata_archi.pdf
 - Issues:
 - This approach requires willingness of the stakeholders to install the performance device.
 - Also requires user-initiated active testing , could be a marketing challenge
 - Is there a passive approach that would work?
- Platform used in Virginia was based on NDT
- NDT is less suitable for measuring CAIs (schools and libraries), due to firewalls, etc.
- If we deploy an active performance measurement tool that is user initiated, it needs to just work, shouldn't require a lot of tech support
- There is an FCC Broadband map program showing the access that different communities have. <http://www.broadband.gov/maps/availability.htm>
- Is this kind of comparative data useful in R&E community?
- Bo Lowrey stated that the CAI community - from Jefferson County Public Schools in Louisville , KY -- is interested.
- Need to explain to users what they might be able to accomplish with increased capacity and improved connection to the R&E network
- We want to provide the user/tester something in return from taking the time
- Perhaps translate the broad network performance results into "what can I do with it"

Q: Is there a concern that this approach requires a self-selected group to do the measurements and this creates a bias?

A: Yes, there is potential for systemic abuse, some people could try to put out a certain kind of picture, would test at a certain time of day, etc. We will have to think about how to control human factor

Q: What is the incentive for folks to do this performance measurement testing more than once?

A: The working group needs to discuss this issue of what will be the carrot? Possibly providing people data, so that they can make a statement about their network performance -- relative to average for the area/region -- is some incentive

Circuit Monitoring Implementation for DYNES

Aaron Brown, Internet2, presented

Slides are found at <http://www.internet2.edu/presentations/spring11/20110420-brown-circuitmonitoringforDYNES.pdf>

- Scientific disciplines are coping with growing data sets and time constrained operations
- Cross-domain dynamic circuits, like those provided by Internet2 ION Service, can be used to improve the connectivity
- DYNES is an NSF-funded project to deploy dynamic circuit networking
- DYNES will bring 33 new sites into the dynamic circuit network "cloud" with hundreds of users
- If a user looks at a circuit now, they don't see enough data to fix failures
- Goal is to give end users better measurements, domain by domain, to fix problems
- Need to enable domains to export monitoring data about circuits AND
- Need to enable users to discover domains that make up their circuit, and see the relevant data
- Solution involves perfSONAR Topology Service and perfSONAR Measurement Archive
- Web UI Client is provided

Q: What is the development roadmap?

A: Circuit monitoring will be part of OSCARS Vers 0.6, to be released this summer

Q: Can we use this for international multi-domain work?

A: This will work on domains that run OSCARS. Also, we are working with DICE partners to create compatibility. This will take time, due to typology sharing issues.

Periscope

Dan Gunter, Lawrence Berkeley National Lab, presented

Slides are available at: http://www.internet2.edu/presentations/spring11/20110420-swany_gunter-periscope.pdf

- Martin Swany and Dan Gunter have been collaborating on the Periscope Project
- There is some overlap between this work and the OSCARS / DYNES work that Aaron just presented
- The Periscope project uses XSP (eXtensible Session Protocol)
- XSPd is the daemon that implements the protocol front end
- Uses ESCPS to allocate a circuit and to monitor circuit status
- The router has a filter for policy-based routing
- Info works its way into the Periscope Monitor Agent
- Periscope includes a perfSONAR type service, incorporating some analysis
- Monitoring Agent can visualize