



CAMPUS PANEL: UNIVERSITY OF UTAH RESEARCH/SCIENCE AND PERFORMANCE DMZ NETWORK

JOE BREEN
JOE.BREEN@UTAH.EDU
UNIV. OF UTAH

1/19/2012
Jan 2012 Internet2 Joint Techs Performance Campus Panel

1

UNIVERSITY OF UTAH RESEARCH DRIVERS AND NETWORK ASSETS

Computational Science

- Scientific Computing and Imaging Institute (SCI)
- Institute for Clean and Secure Energy (ISCE)
- Center for the Simulation of Accidental Fires and Explosions (C-Safe)
- Pharmacy modeling (AMBER)
- High Energy Physics
- Computational Chemistry

Network Research

- FLUX/EMULAB – GENI infrastructure

Medical Research

- University Hospital and Clinics

- Huntsman Cancer Institute (HCI)
- Strong genetics research – Mario Capecchi (Nobel Prize)
- Strong genomics research
- Utah Population DB

Local Network Assets

- Utah Education Network
- Metro Optical Network
- New Data Center (building is 75000 sq ft)
- Internet2 and Level 3 PoP conveniently located by airport (within 9 fiber miles)
- Good partnerships with local transportation entities UTA, UDOT
- Consolidated, Redundant 10+Gb/s campus/hospital backbone

1/19/2012

Jan 2012 Internet2 Joint Techs Performance Campus Panel

2

- Metro Optical Network completes in March
- New Data Center completes Jan 26 and then goes into commissioning

WHY? ... PAIN!

- 50Megabit/second transfers from the Texas Advanced Computer Center (10Gig connectivity)
- 12Mbit/sec transfers from Fermi National Labs
- 6.7Mb/s transfers from Oak Ridge National Labs
- Packet drops up to 22% – see it with UDP iperf and video

1/19/2012

Jan 2012 Internet2 Joint Techs Performance Campus Panel



- Started out looking at connections from UEN to the outside world and then moved back into the campus.
- Saw dramatic drop once within the campus border.
- University of Utah backbone is fully redundant with one or more 10Gb/s connecting each distribution node to a redundant core which connects to a redundant WAN which connects to redundant firewalls which connect to redundant Internet Border routers which connect to the Utah Education Network with a 10Gb/s connection apiece.

BEYOND IMMEDIATE PAIN, WHY? **(HINT: \$)**

- **University Mission requirements**
 - Hospital and Clinics (online billing and pharmacy, etc. -- \$\$\$\$)
 - Administrative (payroll, online donations, credit card transactions, etc. -- \$\$\$)
 - Research (access, collaboration, grant deadlines => overhead -- \$\$)
 - Academic (enrollment, classes, -- \$)
- **Diverging Business rule sets**
- **Compliance acronym of the week... PCI, HIPAA, FERPA...**

1/19/2012

Jan 2012 Internet2 Joint Techs Performance Campus Panel

4

- All billing and drug orders, medical records, etc. now handled online. When the network loses connectivity, the hospital has tangible records of \$/min. loss of revenue. People get more than a little grumpy.
- Access to administrative payroll, online billing, online donations, credit card billing, etc. is all online. Less tangible records of lost revenue but still very visible.
- Access to research collaborators, ability to access national labs, ability to move data, ability to submit grants by deadlines, all rely on network stability. Tangible and intangible impacts to research overhead revenue.
- Academics rely on students finding a welcoming online presence. Online classes, online enrollment, online grading, homework submittal, etc. Most of these topics are intangible impacts to the University revenue but still impact it.
- Research == Openness and Collaboration especially with data movement to national labs
- University Hospital and University Administrative business == closed and protected
- Compliance acronym of the week:
 - PCI compliance
 - HIPAA compliance - hospital, clinics
- The “station wagon” effect still rules – faster to send wagon full of DVDs, thumb drives or disks than to use the network.

BEYOND IMMEDIATE PAIN, WHY?

- **Nimbleness to rapidly scale to higher bandwidth connections, i.e. 100Gb/s**
- **Nimbleness to explore early production technologies, i.e. Openflow**
- **Nimbleness to support unique flows**

1/19/2012

Jan 2012 Internet2 Joint Techs Performance Campus Panel

5

- Want to be able to connect to Internet2 at 100Gb/s within the next 1.5-3yrs. Amortization on the firewalls will be approximately 5yrs.
- Ability to prototype gear, i.e. new security gear, new network technology (think OpenFlow), in a pseudo-production environment. Past a development lab scenario but not quite prime-time for the main production network.
- Try to support unique flows, i.e. GENI implementations, that could pose a higher risk than the production environment is comfortable.

HOW? INSTRUMENTATION! COLLABORATION!

- **Instrumentation of the network - perfSONAR**
- **Collaboration! Collaboration! Collaboration!**
 - Internet2, ESnet, Colleagues at labs, UEN NOC, UEN Engineering, UofU NOC, UofU Architecture, UofU Information Security Office, UofU Compliance
 - Troubleshoot, Validate, Document, Design potential solutions, Mitigate risks, Communicate, Create an Acceptable Use and Security Policy

1/19/2012

Jan 2012 Internet2 Joint Techs Performance Campus Panel

6

- Collaboration within the R&E community and leveraging the perfSONAR instrumentation is key to successful troubleshooting.
- Collaboration with campus entities and the regional network were key to localized troubleshooting of the campus and regional networks. The feedback from the various engineers and the multiple sets of eyes helped in faster isolation of issues.
<https://wiki.chpc.utah.edu/display/CyberInfrastructureLab/Network+Performance+Troubleshooting>
- Leveraged info from <http://fasterdata.es.net> and slides from ES Net group
- UofU and UEN Team wrote up collaborative white paper - http://www.chpc.utah.edu/~jbreen/network/performance/2011-05-30_Network_Performance_Issues_at_the_University_of_Utah.pdf
- Work with UofU Information Security Office (ISO) to review thoughts and vulnerabilities
- Work with UofU Architecture to make adjustments to campus backbone directions
- Work with UofU NOC to design and implement campus backbone
- Work with UofU Compliance office to review and validate risk mitigation
- UEN and UofU are collaborating on metro optical network which will mitigate the single 10G link but it exists for now and is a bottleneck. Always important to work with the upstream provider and keep them in the loop regarding activities in which you may be experimenting. Otherwise, your local fast pipe may become an itty, bitty straw above you. PerfSONAR instrumentation helps in identifying some things. Lots of communication helps mitigate them.
- Having a good policy helps with clarification and understanding of all concerned. The policy also helps to give the security team some teeth and protection so they can work closely with the research community.

HOW? COLLABORATION! RESEARCH BUY-IN, CIO BUY-IN

- **Presentations to University Center for High Performance User Council**
- **Presentations to Univ. of Utah Cyberinfrastructure Council**
 - Researchers, CIO/VPIT, Director of IT Operations, Cyberinfrastructure Director, VP Research, Libraries

1/19/2012

Jan 2012 Internet2 Joint Techs Performance Campus Panel

- The HPC community is always looking for ways to improve data flow and get more from cycles. Several of the UofU researchers account for large blocks of hours of the national lab cycles. They are particularly sensitive to moving their data effectively between the sites. Ongoing explanations to the research community let them know we were working on things and kept their involvement and support.
- Presentations to the CI council helped distribute information and highlight issues. The CI council also lent approval and support to the project. Having key people all in the same room disseminated info and addressed questions. Having a lot of data from the perfSONAR instrumentation was also critical.
- CI council includes following representation
 - head of Eccles Medical Library
 - head of university Marriott Library
 - dean of School of Architecture
 - School of Computing
 - Communications
 - Chair of Geography
 - University Information Technology Faculty representative
 - University Information Technology CIO
 - University Information Technology Director of Operations/Assistant CIO for hospital
 - College of Pharmacy
 - Chemical Engineering/ Institute for Clean & Secure Energy
 - Physics
 - Assistant Vice President Information Technology Health Sciences and Biomedical Informatics
 - Huntsman Cancer Institute

- Vice President of Research
- College of Engineering/Electrical Engineering/Assistant Vice President of Research
- University Information Technology Director for Cyberinfrastructure

HOW? COLLABORATION! EDUCATION, IMPLEMENTATION

Educate community regarding tools, i.e. FDT, bbcp, GridFTP, etc.

BGP Null Routing – scripting based on Netflow triggers by UofU security team and NOC

Out of band security – Bro prototype project happening now by UofU security team

Exploration of additional mechanisms for protecting but simultaneously keeping out of the way.

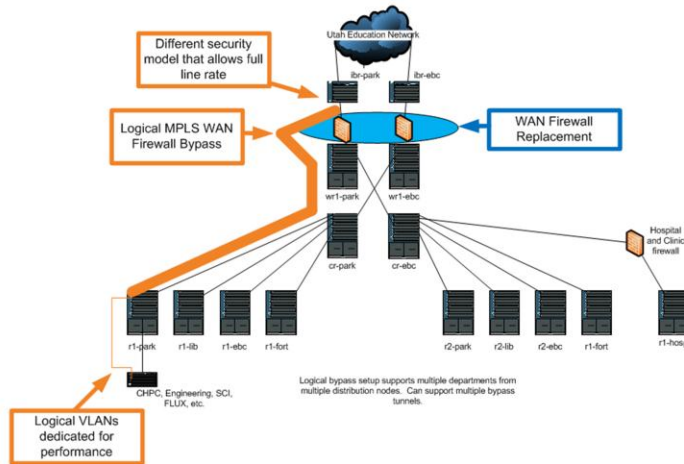
1/19/2012

Jan 2012 Internet2 Joint Techs Performance Campus Panel



- Education regarding tools is a continual process
 - Still heavy use of scp, rsync, etc.
- Implement optimized tools and make easy, i.e. HPN ssh
- Use of parallel rsync streams – somewhat effective
- Security must walk hand-in-hand with the performance. Not even the hard core performance research groups want to hang out in the wind without a little protection.
- The collaborative team has worked and is working to develop out-of-band solutions that scale. The security group already had scripting for the existing firewall based on Netflow information. The security group successfully prototyped expanding this scripting to manipulate Blackhole routing for problematic hosts and networks. This project is now going production.
- The collaborative team is now exploring Bro and setting up a prototype. Bro allows scaling in an out-of-band manner.
- The collaborative team will be exploring an OpenFlow implementation in the upcoming months and see where we can leverage these ideas too.

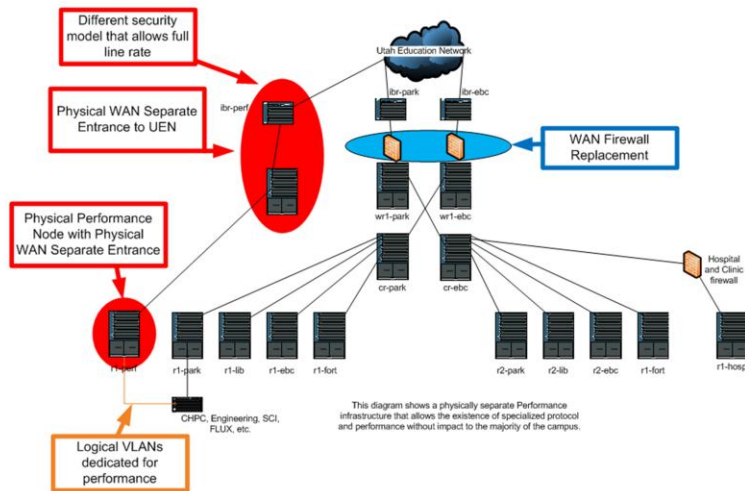
WHAT IS THE RESULT TODAY?



1/19/2012
Jan 2012 Internet2 Joint Techs Performance Campus
Panel

- Partial snapshot of campus backbone with a MPLS tunnel providing a backbone path that goes from a distribution node, through the core to the WAN router, around the firewall and terminates traffic on the Internet Border Router. The traffic ingresses/egresses directly on the IBR and on the distribution router. The end customer provides own routing or routes on the distribution router.

WHAT MIGHT THE FUTURE BE, IF FUNDING ALLOWS?



1/19/2012
Jan 2012 Internet2 Joint Techs Performance Campus
Panel

10

- New physical IBR in order to separate the performance research/science DMZ network traffic from the rest of the U WAN traffic in order to mitigate risk
- Complete separate infrastructure – NOT 5 nines, no dual-homing (under discussion), possibly different network vendor infrastructure.

ISSUES ALONG THE WAY

7yr firewall hardware

Are you testing IPv4 or IPv6 with your active measurement infrastructure?

What really is the path MTU?

Ability to release the bottlenecks at University can potentially flood upstream provider – Make sure you are collaborating tightly!

1/19/2012
Jan 2012 Internet2 Joint Techs Performance Campus
Panel



- Operational graphs did not reflect the packet drops and did not show the limited throughput.
- Graphs of firewall throughput looked like existing firewalls were within expected parameters, though, some anomalies had arisen.
- Dual-stack is nice for servers but problematic for measurement infrastructure. What are you really testing?
 - MPLS overhead causing mismatch in MTU, etc.
 - New firewalls have different MTU max than previous firewalls.
- UEN has temporary single 10Gb/s feeding Level 3 PoP which houses Internet2 connectivity and multiple Commodity Internet connections.
- Waiting on metro optical network to relieve bottleneck.
- Filling research pipes causes commodity to slow down dramatically leading to some concern.
- Resources available
 - Timing with major data center project
 - Timing with other major projects

PANACEA? NOPE, AT LEAST NOT YET.

- **Still working with MTU issues with new firewalls, MPLS tunnels and router settings**
- **Still educating and trying to get researchers to use high performance transfer tools**
- **Research/Science DMZ Network *NOT* a Network Development Sandbox - Need Network development sandbox testbed too**

1/19/2012

Jan 2012 Internet2 Joint Techs Performance Campus Panel

12

- Trying to finish policy
- Trying to obtain funding
- Still seeing changes in the world affect transfers