CASE STUDY: UNIVERSITY OF UTAH RESEARCH/ SCIENCE AND PERFORMANCE DMZ NETWORK

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

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

* Dave Pershing

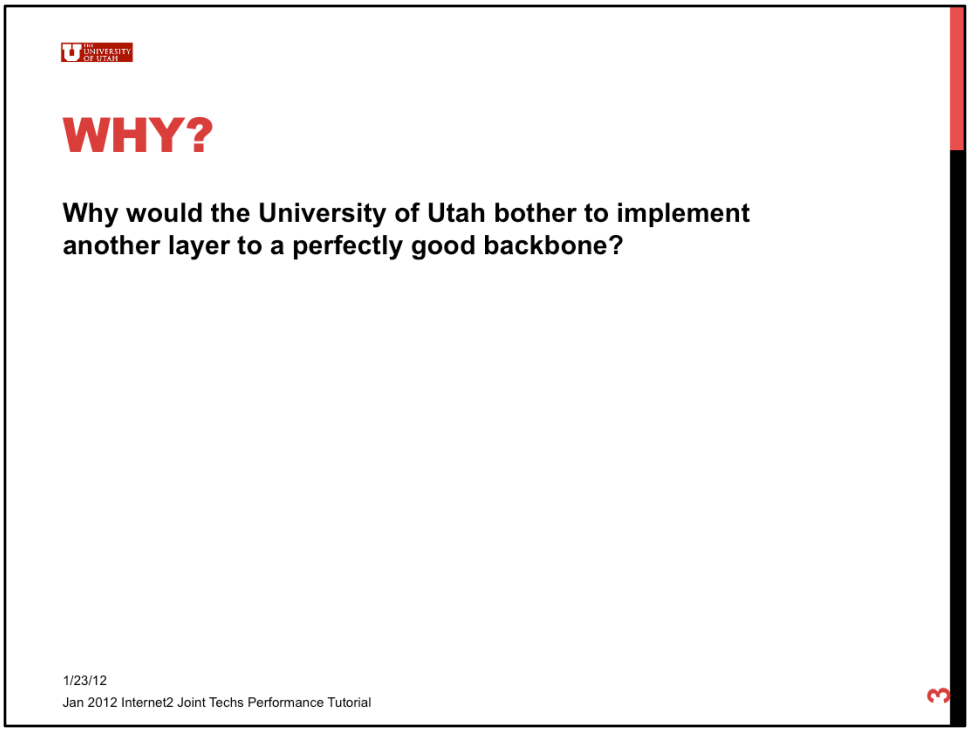

- University of Utah backbone is fully redundant with one or more 10Gb/s connecting each distribution node to a redundant core which connects to a redundant WAN which connects to redundant firewalls which connect to redundant Internet Border routers which connect to the Utah Education Network with a 10Gb/s connection apiece.

WHY? … PAIN!

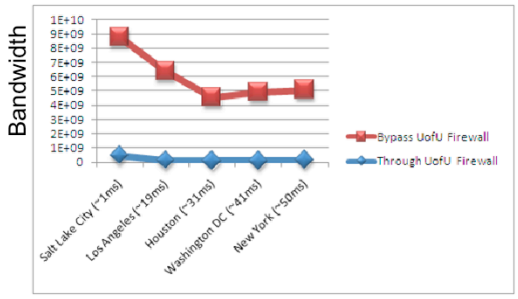

WITH DMZ

WITHOUT DMZ

1/23/12

Jan 2012 Internet2 Joint Techs Performance Tutorial

4

- Starting for a moment with some of the results quickly highlights the pain points…
- Univ of Utah has 2 10Gb/sec links to the Utah Education Network which has 10Gb/s to Internet2
- Red line denotes performance without UofU firewall
- Blue line denotes performance THROUGH UofU firewall

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

- 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.

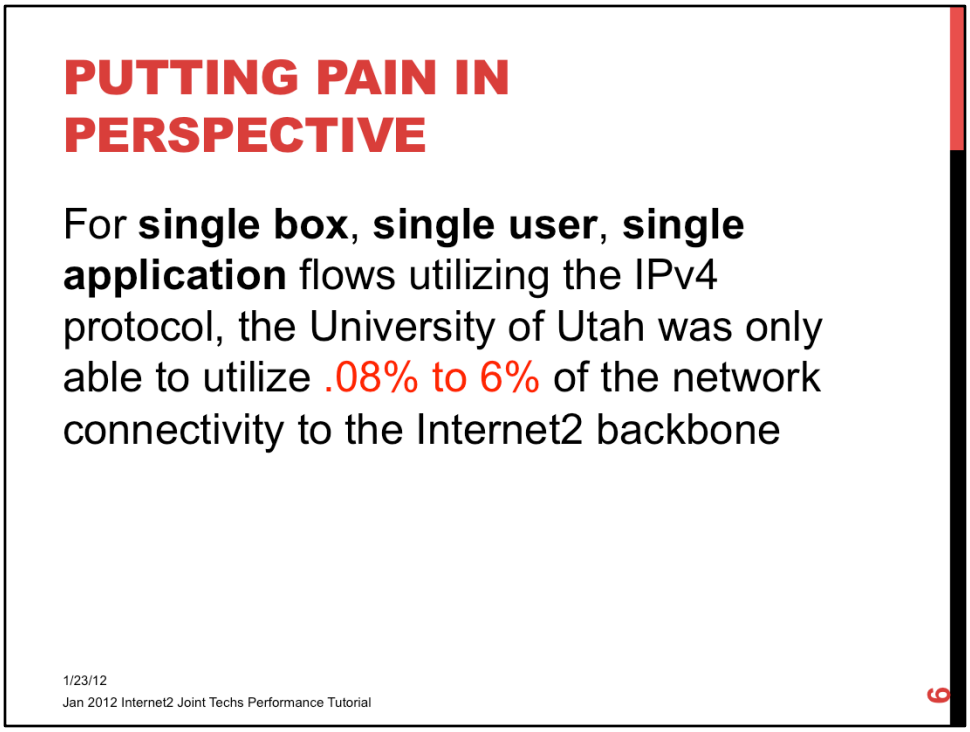

- Used iperf and FDT to test the baseline network and then file transfers.

PUTTING PAIN IN PERSPECTIVE

For multiple box, multiple user, multiple application flows, the Univ. of Utah was hitting ceilings of 20-30% of the available network bandwidth

- Used iperf and FDT to test the baseline network and then file transfers.
- Created multiple parallel flows, both from UEN's perspective and from within the University.

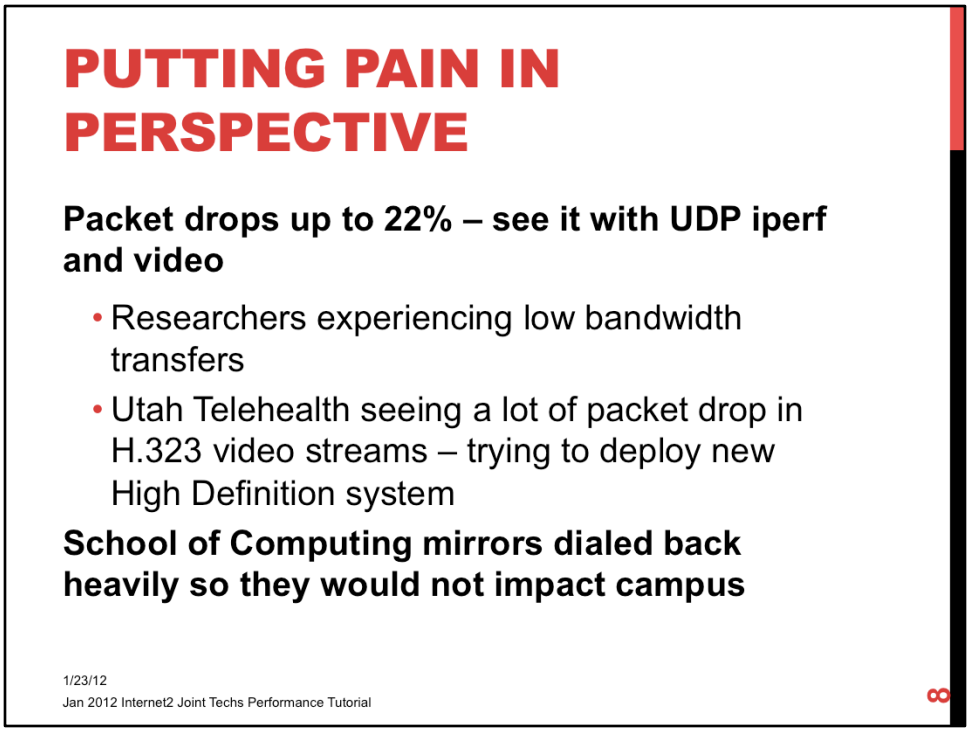

- We didn't start with all of this info at the beginning, we had to dig it up by looking at a lot of aspects of the network.
- Started with pain of large research transfers and kept digging. Utah Telehealth started researching their own issues in parallel.
- Campus saw School of Computing bury the existing firewalls when some of the Linux distros released another distribution. School of Computing wanted 10Gb/s but funding and a bit of concern held campus back from allowing the connectivity.

PUTTING PAIN IN PERSPECTIVE -- $

UofU/BYU/USU/UEN/Montana maintains 2x10Gb/s connection to Internet2 at $525k/yr

The performance issues were preventing the University of Utah from fully realizing the significant investments it is making in the network

- UofU has 10Gb/s+ backbone
- UofU has two 10Gig connections to UEN

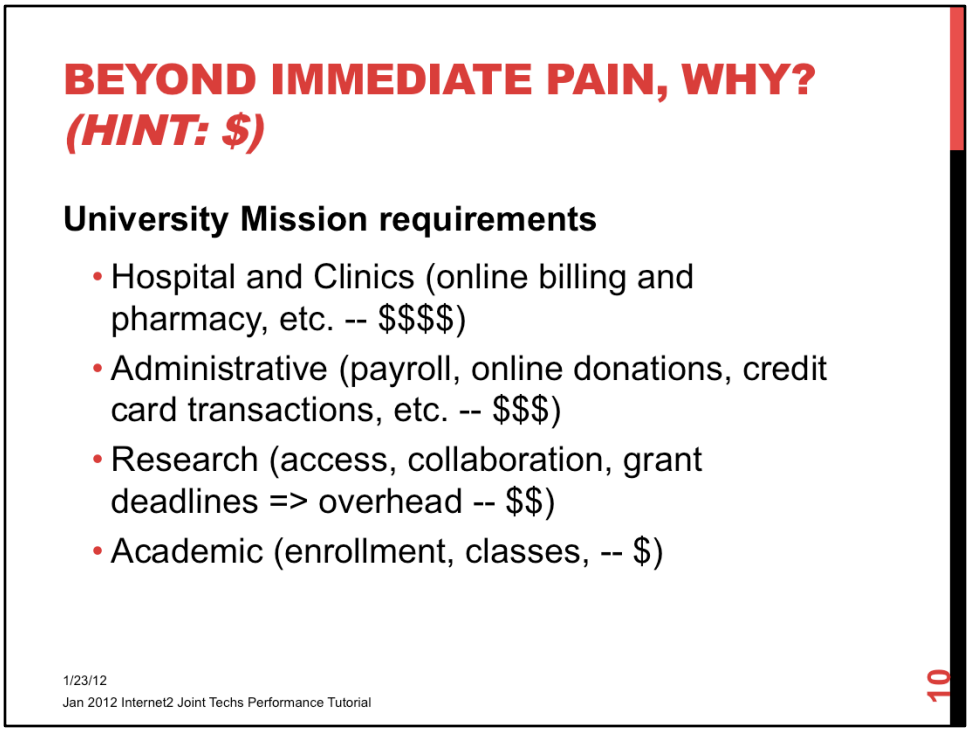

- All billing and drug orders, medical records, etc. now handled online. When the network loses connectivity, the hospital has tangible records of $/min. loss of revenue. People get more than a little grumpy.
- Access to administrative payroll, online billing, online donations, credit card billing, etc. 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.

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

Diverging Business rule sets

- Research == Openness and Collaboration especially with data movement to national labs
- University Hospital and University Administrative business == closed and protected
  - PCI compliance -
  - HIPAA compliance - hospital, clinics
  - Compliance acronym of the week…

1/23/12
Jan 2012 Internet2 Joint Techs Performance Tutorial
11

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? *(HINT: $)*

- **Operations:**
  - Longer amortization of redundant WAN equipment and redundant WAN firewalls → $

- 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.

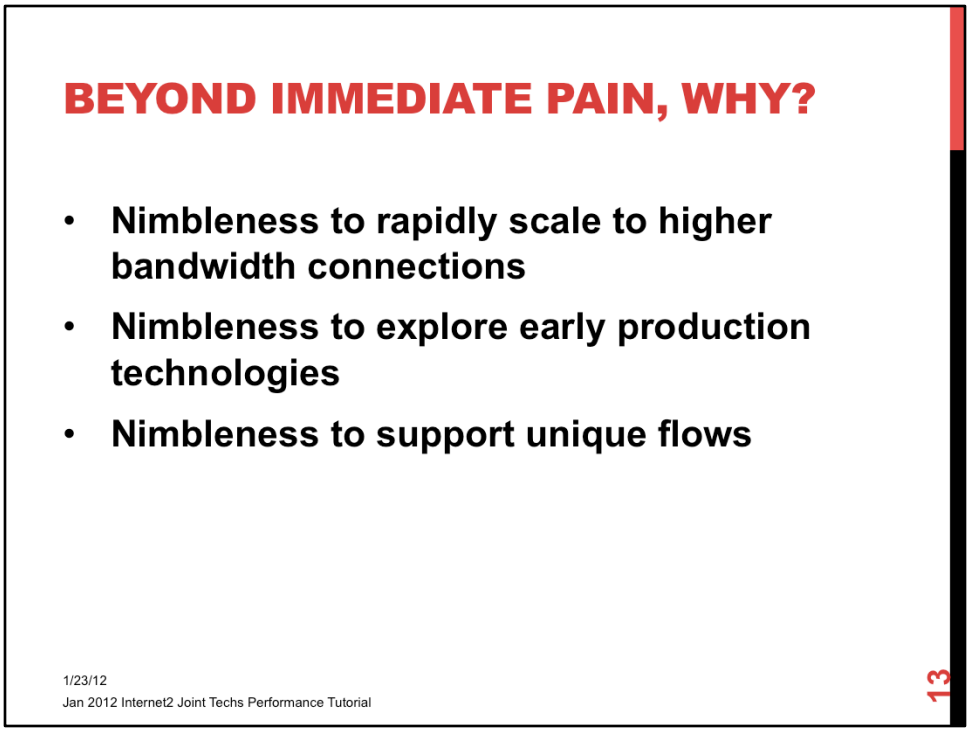

- Want to be able to connect to Internet2 at 100Gb/s within the next 1.5-3yrs. Amortization on the firewalls will be approximately 5yrs.
- Ability to prototype gear, i.e. new security gear, new network technology (think OpenFlow), in a pseudo-production environment. Past a development lab scenario but 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.

NOW THAT WE HAVE ESTABLISHED THE WHY, HOW?
FIRST, INSTRUMENT THE NETWORK.

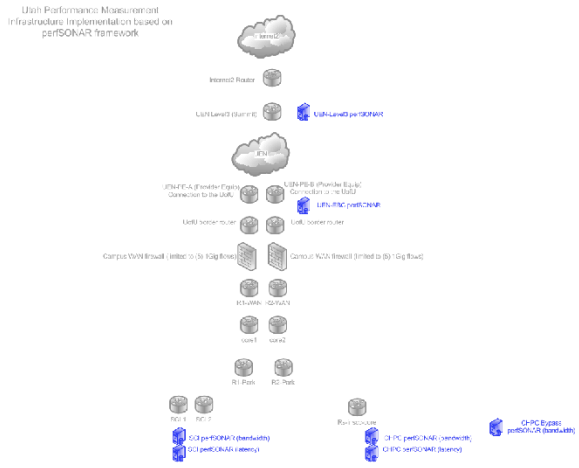

1/23/12

Jan 2012 Internet2 Joint Techs Performance Tutorial

14

- Deployed perfsonar nodes in UEN
  - Immediately outside campus
  - Immediately before Internet2
- Deployed in campus space
  - Within CHPC
  - Within SCI
  - On bypass network

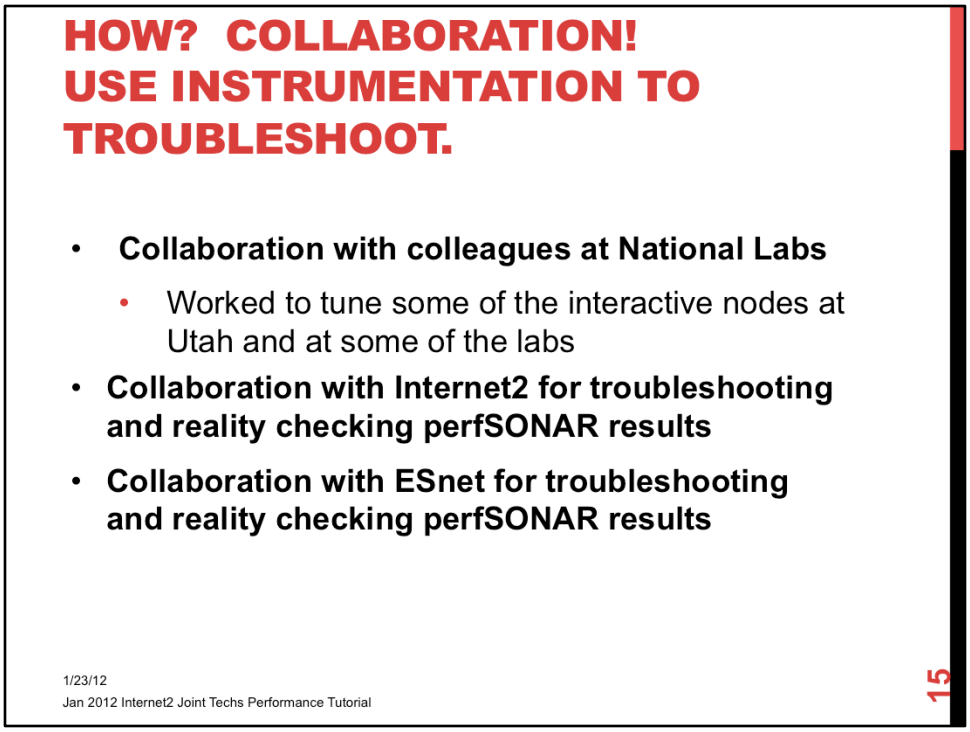

- Collaboration within the R&E community and leveraging the perfSONAR instrumentation is key to successful troubleshooting.

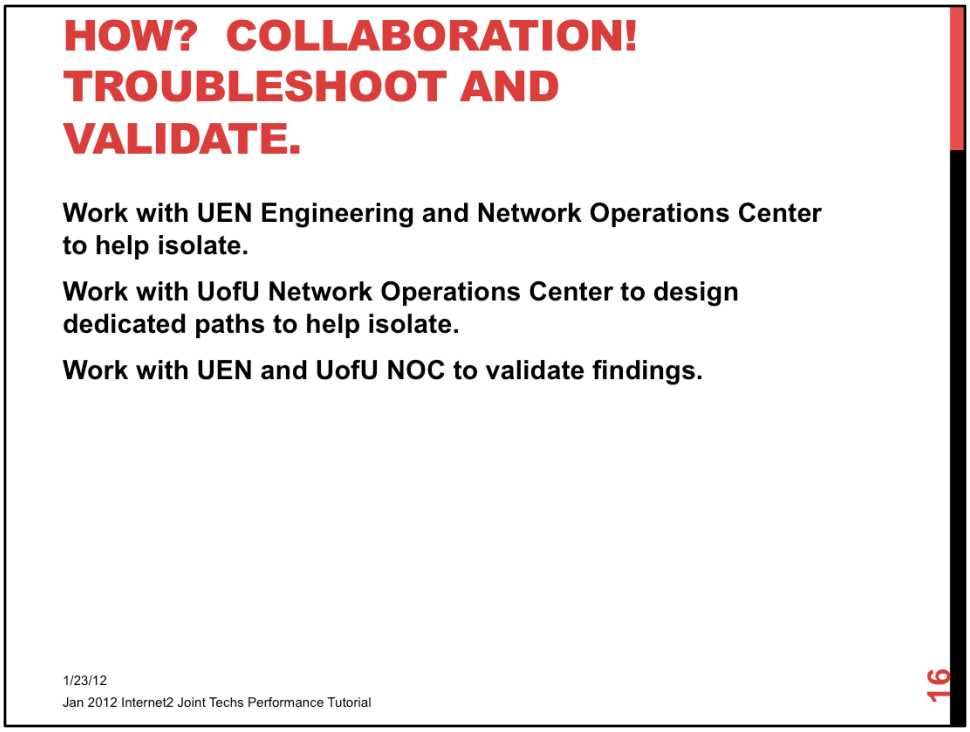

- Collaboration with campus entities and the regional network were key to localized troubleshooting of the campus and regional networks. The feedback from the various engineers and the multiple sets of eyes helped in faster isolation of issues.

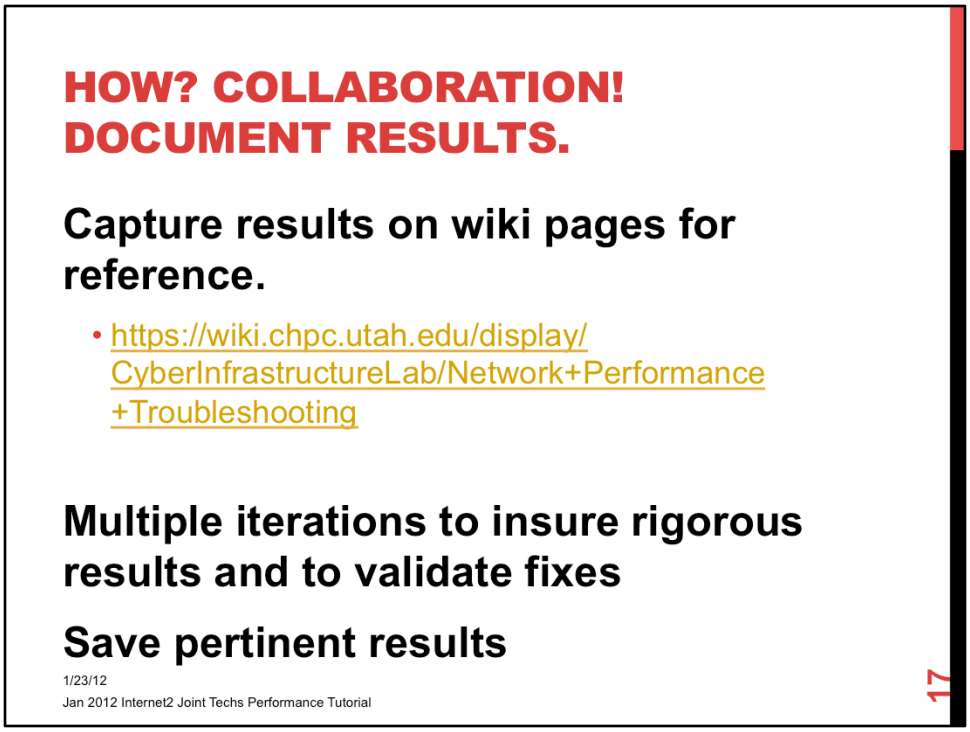

- Documenting notes on wiki really helped in putting together results that we could look back on and see improvement. Also helped when we saw things go worse. For example, we found out the firewalls were affecting IPv6 packets worse than IPv4 quite by accident. We did not realize that Internet2 had fixed some DNS records and our tests were utilizing DNS names instead of IP addresses. The traffic started using IPv6 instead of IPv4 because we had a full IPv6 path. Traffic took a dive.

HOW? COLLABORATION! DOCUMENT RESULTS.

- Leveraged info from http://fasterdata.es.net and slides from ESnet 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

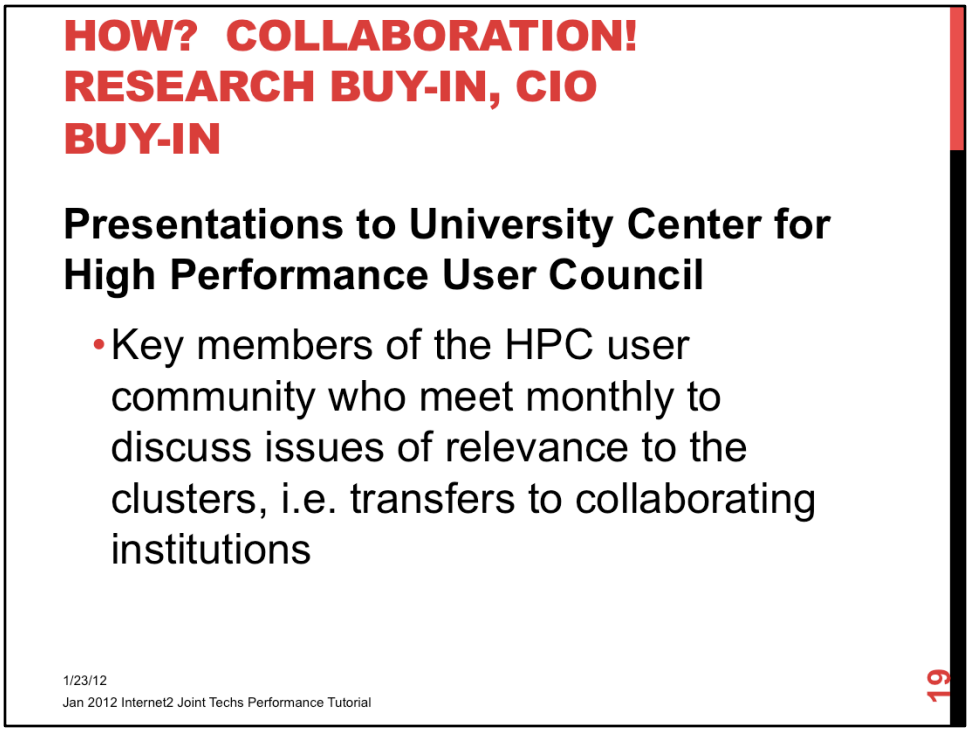

* The HPC community is always looking for ways to improve data flow and get more from cycles. Several of the UofU researchers account for significant use of the national lab cycles. They were particularly sensitive to moving their data effectively.

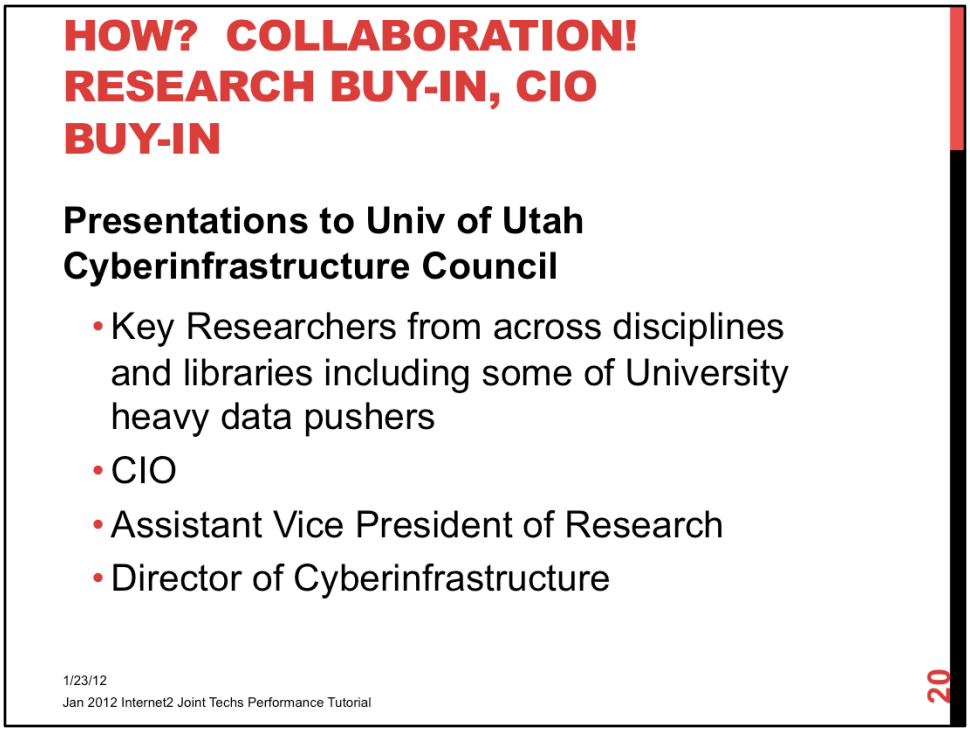

- CI council includes following representation
  - head of Eccles Medical Library
  - head of university Marriott Library
  - dean of School of Architecture
  - School of Computing
  - Communications
  - Chair of Geography
  - University Information Technology Faculty representative
  - University Information 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! DESIGN POTENTIAL SOLUTIONS. MITIGATE RISKS. COMMUNICATE!

Now that the buy-in exists, how do we start putting the pieces together?

- Collaborate with team to identify solutions
- Collaborate with team to identify and mitigate risks
- Communicate!

HOW? COLLABORATION! DESIGN POTENTIAL SOLUTIONS. MITIGATE RISKS. COMMUNICATE!

- 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? EDUCATE.

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

- 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

HOW? COLLABORATION! IMPLEMENT POTENTIAL SOLUTIONS AND PROTOTYPE ADDITIONAL TOOLS.

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.

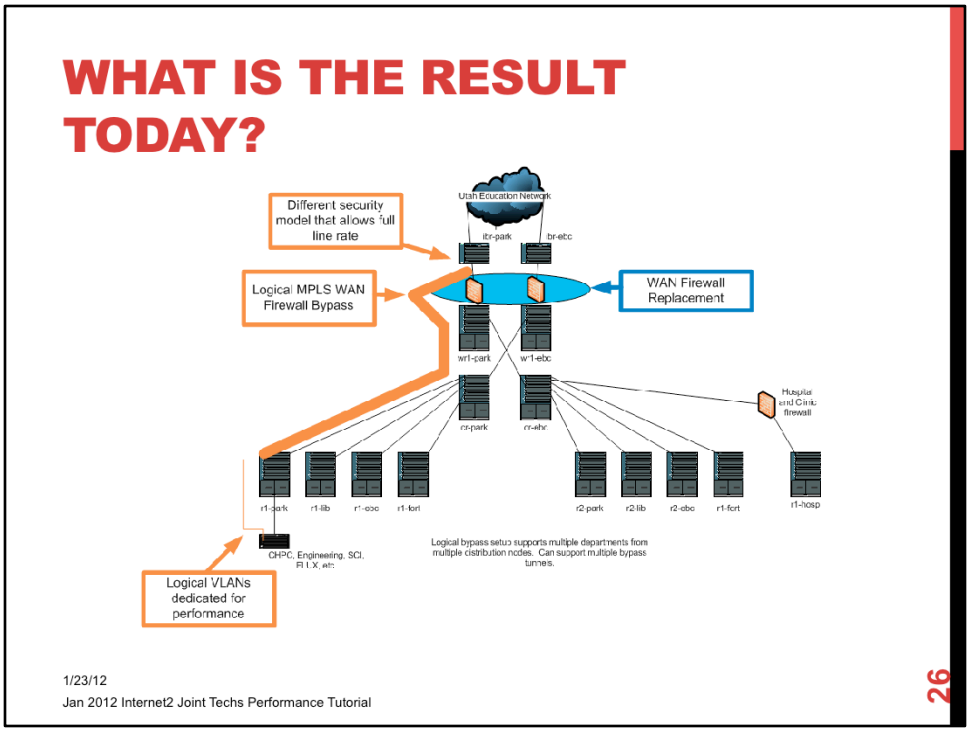

- Partial snapshot of campus backbone with a MPLS tunnel providing a backbone path that goes from a distribution node, through the core to the WAN router, around the firewall and terminates traffic on the Internet Border Router. The traffic ingresses/egresses directly on the IBR and on the distribution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

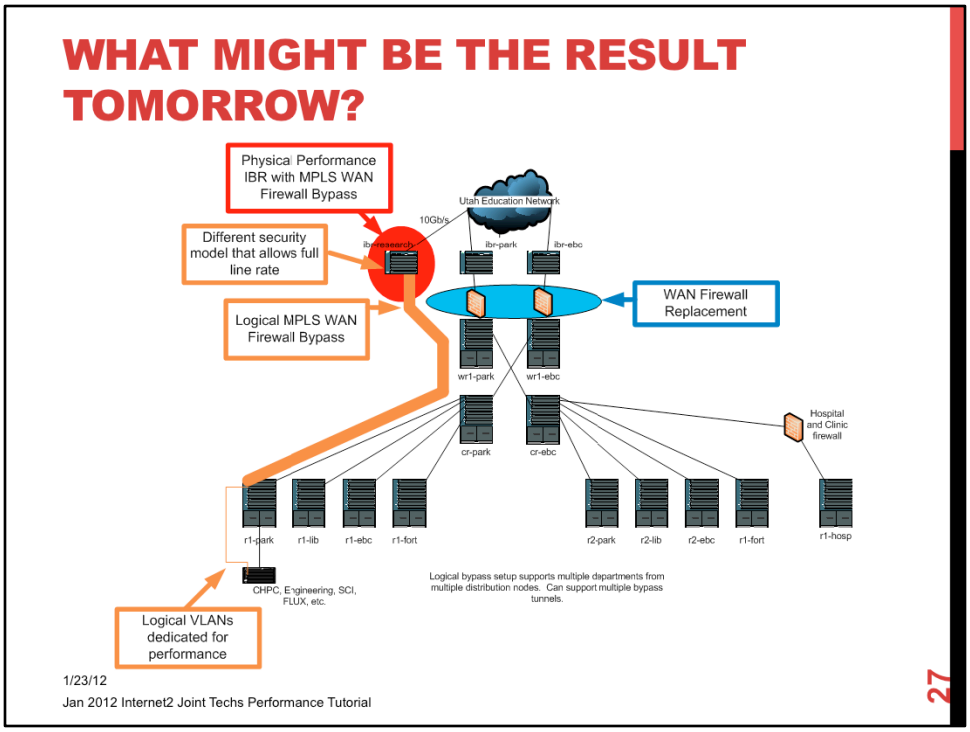

- New physical IBR in order to separate the performance research/science DMZ network traffic from the rest of the U WAN traffic in order to mitigate risk. At first, the idea was to implement a performance distribution node first, but, the WAN is the higher risk, i.e. filling pipe or different security rule gone awry.

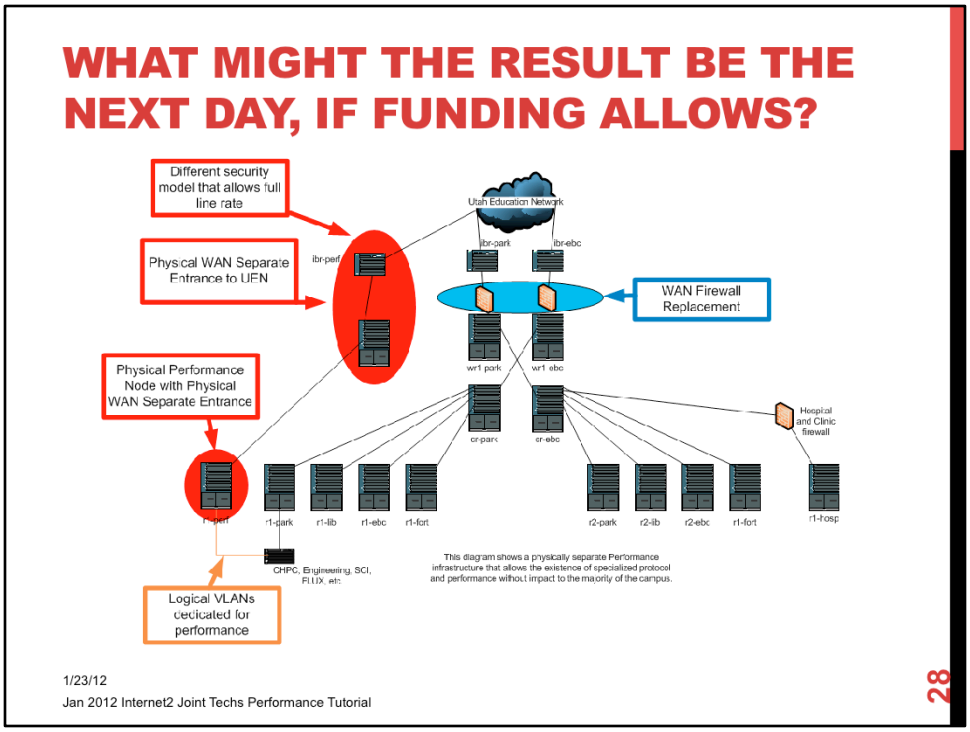

- New physical IBR in order to separate the performance research/science DMZ network traffic from the rest of the U WAN traffic in order to mitigate risk
- Complete separate infrastructure – NOT 5 nines, no dual-homing (under discussion), possibly different network vendor infrastructure.

ISSUES ALONG THE WAY

- **7yr firewall hardware**
  - 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.

ISSUES ALONG THE WAY

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

- Dual-stack is nice for servers but problematic for measurement infrastructure. What are you really testing?

Graphs not showing? What really is the path MTU?

- MPLS overhead causing mismatch in MTU, etc.
- New firewalls have different MTU max than previous firewalls.

ISSUES ALONG THE WAY

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

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

NEED DEVELOPMENT NETWORK TOO

- **Research/Science DMZ Network *NOT* a Network Development Sandbox**
  - Need pseudo-production focused on performance and unique flows
- **Need development sandbox testbed too**
  - Need to play with technologies such as OpenFlow in a network sandbox and then roll to the Research DMZ

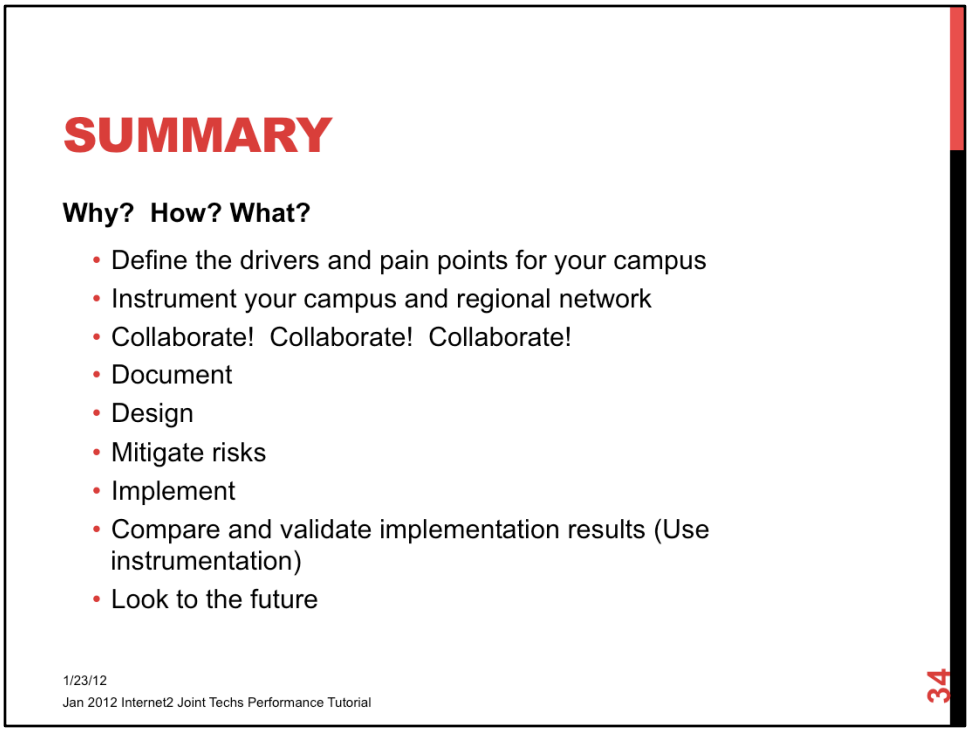

- Make a list of issues that are affecting your campus
- Instrument your campus and regional network with perfSONAR
- Collaborate with your research community, your security group, your NOC, your Compliance group, your IT leadership, your regional NOC, your national backbone provider (I2/ESnet/etc.), your colleagues at peer institutions, …