ADVANCING RESEARCH CREATING SOLUTIONS

February - March, 2010 The Great Plains Network



# GPN

## *Great Plains Network First to Sign with Internet2 to Double Connectivity for Members*

Taking advantage of new options for network connectivity to the Internet2 national research and education network, the Great Plains Network (GPN) has doubled network connectivity on behalf of member universities by adding a second 10 Gigabit connection to Internet2. This connectivity makes it possible for scientists and educators from GPN member universities in South Dakota, Nebraska, Kansas, Oklahoma, Missouri, Arkansas, Minnesota, Wisconsin and Iowa to effectively participate in exciting "big science" projects like the Large Hadron Collider project in Cern, Switzerland, 3D medical applications, advanced climate modeling and interactive educational experiences like those involving GPN membership and National Park Service sites across the US.

GPN was the first to sign an agreement with Internet2 to take advantage of their new 2010 network pricing options, designed to significantly enhance the research and education national infrastructure. GPN was also the very first regional connector to the original Internet2 network in 1998. To understand the significance of the new connection, it has the capacity to download a DVD every 4 seconds. This speed makes it possible to launch a new generation of transformative scientific, educational and medical applications across the GPN region.

## GPN Planning & CI Survey

By Greg Monaco and Bill Mitchell, Great Plains Network

GPN has two planning efforts underway. You can find updates about both the GPN network program planning and GPN cyberinfrastructure (CI) planning at the GPN wiki. For updates on network planning, visit here. For updates on CI planning, visit here.

The CI Advisory Committee is in the final stages of designing a ***CI Needs Assessment Survey***. It will be short and require just a few minutes to fill out. Please check for it in March, and I encourage you to enlist others to complete it.

---

## Network Utilization and Service Summary

By Hank Niederhelm, MOREnet

*GPN connects our university members to the Internet2 "backbone" in Kansas City and that connection is capable of exchanging data at the rate of 10 Gigabits of per second (Gbps). The chart below indicates the utilization of that connection (in Gigabits) during the period from January 1, 2010, to January 31, 2010. The blue line indicates outgoing data and the red line indicates incoming data.*

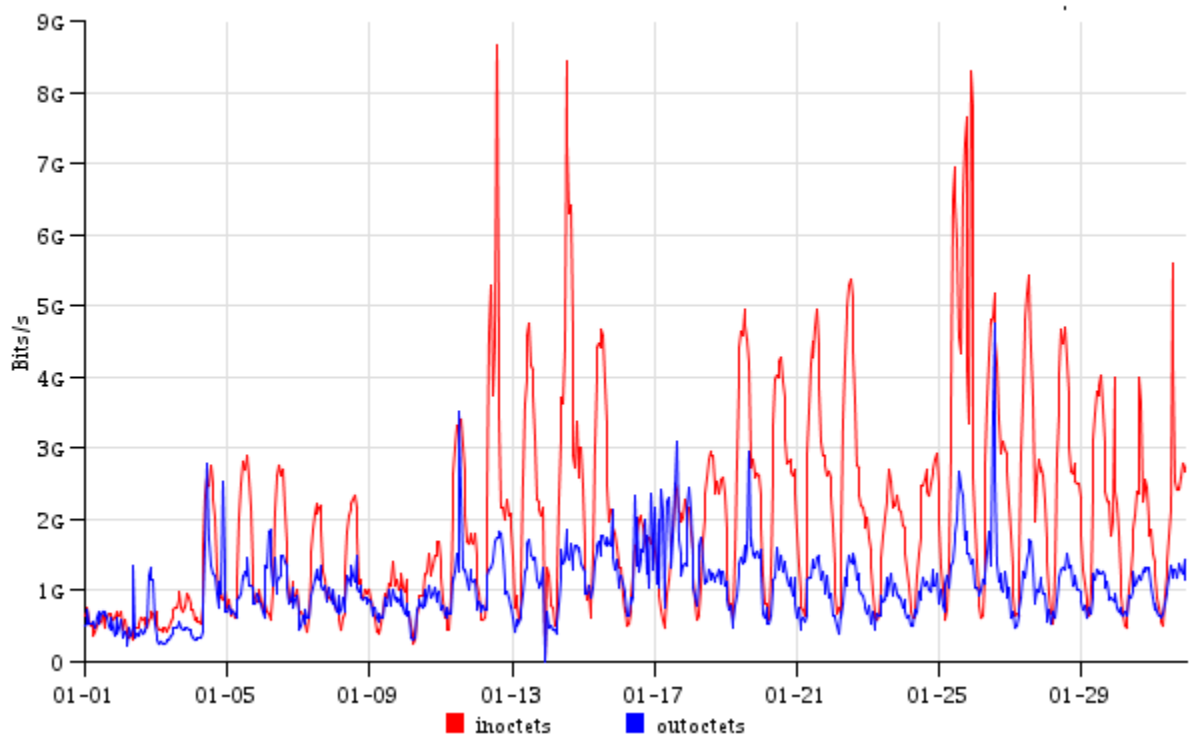


---

## Representative Council Profile: Mark Askren, University of Nebraska, Lincoln

Mark Askren is the first Chief Information Officer for the University of Nebraska-Lincoln. He began on August 10, 2009. His role includes responsibility for the campus-wide leadership of information technology systems and services in support of instruction, research, and outreach. His prior experience includes serving as the Assistant Vice Chancellor for Administrative Computing Services at the University of California Irvine,

where he was also a member of the University of California's Information Technology Leadership Council. Previously, he has served as Assistant Vice President for Application Development and Data Management at the University of Illinois, and as Assistant Dean for Information Technology at the UC San Diego School of Medicine. Mark completed his master's degree in Business Economics and Public Policy at Indiana University, and his bachelor's degree in business from Florida State University. Mark also serves as a member of the GPN Executive Council.

---

# Education

## *GPN K20 Committee Plans Outstanding National Park Service Plan Events for April*

By Kate Adams, Great Plains Network

As a result of the planning efforts mentioned in last month's newsletter, the GPN K20 Committee together with several national parks are planning interactive distance education events for April. These events will be interactive between the parks and classrooms of students and will be streamed live via the Internet. The parks and dates times are:

- Monday, April 19th, 10-11am CST: **Homestead National Monument**, Beatrice, NE
- Tuesday, April 20th, 9-11am CST: **Brown vs. Board of Board of Education**, Topeka, KS
- Wednesday, April 21, 10-11am CST: **Knife River**, Stanton, ND
- Thursday, April 22, 1-2pm CST: **Minuteman Missile**, Wall, SD

Registration is open from March 1 – 30, 2010. ***Registration to be an interactive classroom*** is on a ***first come, first served basis***, and there is still time for classroom teachers to register their classes for the events! There will be multiple classrooms involved for each event.

***Registering to watch via video stream*** (non-interactive), only, ***is still encouraged*** to be sure that the teachers and tech personnel are updated and aware of any special requirements.

A link to a flyer that you can circulate to schools to register is here. You can learn more at http://collaboration.greatplains.net/, or contact kate@greatplains.net.

If you are interested in arranging an event of this nature to demonstrate your research and become a part of this exciting educational process, please contact Greg Monaco, greg@greatplains.net.

---

# Campus Technology

## *New Type of Phishing Attack Threatens K-State Passwords*

By Harvard Townsend, Kansas State University

*Harv Townsend is familiar to all long-time GPNers as an individual who significantly contributed to the overall development of GPN. Harv is the* Chief Information Security Officer *at Kansas State University. This article is condensed and reprinted with his permission.*

Hackers have been VERY successful at tricking K-Staters into giving away their eID password — **in 2009 more than 430 K-Staters replied to phishing e-mails, sending their eID passwords** to criminals who used those to log into K-State's e-mail and send thousands of spam e-mails. The good news is that repeated communications by K-State's IT security team, Help Desk, and IT support staff have slowed the pace of compromised e-mail accounts. **The bad news is the hackers' techniques have evolved accordingly with new, more sophisticated scams that steal your password.**

K-State's mantra for the last two years has been "NEVER provide your password in an e-mail to anyone under any circumstances!" How did the hackers respond? On Jan. 23, they sent a scam e-mail to numerous K-Staters. It doesn't ask you to send your password *in an e-mail.* Instead, it tries to trick you into clicking on a link that goes to a website where they want you to enter your eID and password.

Subject: Account Update.
From: Kansas State University <onlineupdate@k-state.edu>
Date: 1/23/2010 7:12 PM
To: undisclosed-recipients:;

Dear User,

A scheduled maintenance has just been completed on our e-mail system. In order to keep this account active and protected, you will be required to immediately re-login. Kindly click on the "Login to My Account" link stated below:

Login to My Account

We apologize for any inconveniences caused.

Security Department,

E-mail Services.

If you click on the link, you end up on a website *hosted on a server in the Netherlands* that is **an exact replica of K-State's single sign-on (SSO) page:**

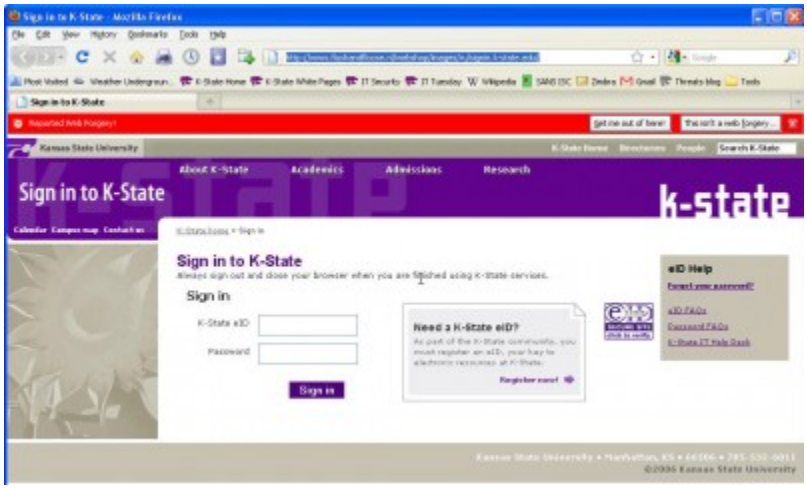


When you enter your eID and password on this fake page, they're sent to the criminals the same as if you provided them in an e-mail. Then your browser is immediately re-directed to K-State's homepage to make you think you successfully logged in.

Fortunately, within a short time of receiving this scam e-mail, Trend Micro OfficeScan's Web Reputation Services blocked access to this malicious replica, as did Firefox's anti-phishing filter, thus limiting the damage done by this particular scam.

This also complicates K-State's prevention message — we can no longer simply say "don't give away your password in an e-mail and you'll be protected." Everyone needs to be on the lookout for new types of scams.

# *Information Technology Disaster Recovery at Missouri University*

By Brien Waage, University of Missouri

*Brien Waage is a Network Systems Analyst Architect at the University of Missouri. He is also a contributor to several GPN initiatives and has served on the GPN Representative Council.*

Disaster recovery for IT systems is very important to MU. The loss of the data network and associated services for an extended time period would cripple the academic and business operations of the university. The MU central IT organization, the Division of Information Technology, or DoIT, has been working for several years to architect and deploy physical infrastructure, network structures, service structures and procedures to recover from events that would interrupt IT services. In the MU IT DR plan, the Telecommunications Building is destroyed through either a natural or manmade disaster. Several years ago the Telecommunications Building was a single point of failure for most MU IT systems and services.

- It housed the data center where the most important applications and services lived.
- It housed all data storage systems except for off-site tape backup.
- It housed the components that provided network connections to the UM System network, Intercampus network, University Hospital network, Internet1, Internet2 and MORENet.
- It housed the network and network security services including routing, firewalling, intrusion prevention, VPN, DNS and DHCP.

Over the last few years MU has steadily reduced the risk of extended service outages associated with either disasters or more mundane incidents like power outages, component failures and system maintenance. Several structures have been put into place to mitigate our risks.

**Geographic Redundancy**

MU now has three data center and network switching node locations instead of just one. The Telecommunications Building site continues to be the primary data center and data network switching node for MU.

A second data center and data network switching node, the North Node, has been added on campus and acts as the "high availability" site. This site is about one mile from the Telecommunications Building. There is diverse fiber between the Telecommunications and North nodes. Redundant connections to the UM System network, Intercampus network, University Hospital network, Internet1, Internet2 and MORENet are located at this site. Redundant network services including DNS and DHCP are located at this site. Redundant routing, load-balancing, firewalling, Intrusion Prevention and VPN services are located at this site. The data network and services are architected such that if the Telecommunications Building is destroyed the remaining network segments and services will automatically continue to function out of the North Node.

MU has a reciprocal agreement with the Missouri University of Science and Technology that provides data center services for MU on the MUST campus. This data center acts as the "Disaster Recovery" site. The MUST campus is more than one hundred miles away from the MU campus providing significant geographic segmentation. The MU and MUST sites are connected via a 10Gbs optical network. This network has full diversity, redundancy and optical layer protection. Data storage systems are replicated to this site. Development systems for services are located at this site. Services like PeopleSoft and E-Mail can be brought up at this site if necessary. Today this is a manual process. We are working towards automatic fail-over for some systems and services.

**High Availability**

There are multiple layers of electrical power sources. Major services are connected to both city and campus power sources. UPSs and generators provide power backup.

Major services are deployed in a high availability architecture where possible. Multiple NICs connect to different Ethernet switches. There is load-balancing between multiple servers connected to different Ethernet switches. Servers in a load-balanced team are located in different data centers.

The data network is architected for high availability. Data center Ethernet switches have redundant uplinks to core layer-3 TCP/IP switches located in different data centers. Load-balancers are located in different

nodes. Every network has redundant firewalls located in different nodes. Every network has redundant routers located in different nodes.

**Lessons Learned**

1. Disaster recovery, business continuity and high availability are very expensive.
   - Building and maintaining even a small second data center and switching node costs a lot up front and in ongoing expenses. UPS, Generator, electrical systems, HVAC, fiber cabling, copper cabling all add up.
   - Redundant network routers, switches, firewalls, IPS systems add significant cost up front and in ongoing maintenance, licensing and support costs.
   - Additional servers, software and licenses add significant cost up front and in ongoing maintenance and support costs.
2. Disaster recovery, business continuity and high availability require significant staff time.
   - Someone needs to be responsible for documenting systems and procedures, organizing and coordinating system tests, setting priorities etc.
   - DoIT has a staff of two dedicated to DR.
   - Service and system administrators are required to spend significant time designing and supporting more complex systems, documenting and testing.
3. High availability is complex.
   - High availability features are generally available on data network and network security systems.
     - These features can be complex to configure and don't necessarily work as advertised.
     - Believe nothing and test everything.
   - High availability features for applications are not necessarily available, can be complex to configure and don't necessarily work as advertised.
4. Security is more difficult. Maintaining network firewall, host firewall, ACL and IPS rules when there are more devices and more sites involved is challenging.
5. Long distance data replication is not the same as short distance data replication. We found that in many cases the TCP settings on the servers or storage systems doing the replication had to be tweaked to work effectively over long distances at high speeds. In some cases we saw a tenfold increase in performance by tuning TCP.
6. Have an integrated plan.
   - It is difficult to provide a comprehensive DR plan for a fragmented entity like a University.
     - Servers and storage distributed all over campus and remote sites.
   - MU worked to put into place certain policies and services to better enable DR and to reduce cost.
     - Move servers and services to the data centers.
       - DoIT provides data center space at no or low cost for departments.
       - Departments benefit from power systems like multiple providers, UPS and generators.
       - More secure environment. Access provided 24 x 365.
     - Eliminate direct access storage and move to shared storage.
       - More efficient use of storage.
       - Enables better data replication and backup methods.
     - Virtualize as many services as possible.
       - More efficient use of hardware.
       - Eases the process of making services available at other sites.
7. Testing your plans is painful but necessary.
   - Too many times after spending a lot of money and effort the DR and HA facilities fail when they are needed. The UPS or generator fails, application or network component code was updated and the features no longer work or a systems configuration drifted over time.
   - Test everything using the defined procedures at least yearly or after significant system changes.
   - This can actually create system outages but at least it's at a time of your choosing.

# Science

## *Open Science Grid Enables Real-Time Data Analysis for ATLAS*

By John Hover, Brookhaven National Laboratory

*Note: The ATLAS project is one of several large hadron collider (LHC) projects that rely on real-time streaming and processing of data across multiple sites. The University of Oklahoma is one of 5 ATLAS Tier 2 sites. The University of Nebraska at Lincoln is a Tier 2 site for a similar project (CMS). John Hover is the group leader for the Grid Middleware and Services Group within the RHIC/ATLAS Computing Facility (RACF) at Brookhaven National Laboratory (BNL). BNL is the US Tier 1 site for the ATLAS experiment at the LHC in Geneva, Switzerland.*

ATLAS (**A** **T**oroidal **L**hc **A**pparatu**S**) is one of the four particle physics experiments located at the Large Hadron Collider at CERN. With the first sustained operation of the LHC in December, ATLAS has finally taken substantial amounts of real collision data—ultimately at record-setting energy (2.36 TeV).

Subsequently, in late December, 2009, and early January, 2010, U.S. ATLAS made heavy use of its sites on the **Open Science Grid** (OSG) to perform reconstruction and initial analysis of this data. During these processing runs, the ATLAS Tier 1 Center at Brookhaven National Lab, along with 5 Tier 2 sites (and a growing number of Tier 3s) typically ran 10-13,000 jobs at a time, with a maximum throughput of almost 50,000 jobs over a 24-hour period.

Thanks to the OSG, this work was accomplished quickly and with very high efficiencies (~99%). ATLAS plans to perform another reprocessing run (which means re-analyzing the December data with refined calibration information) before the LHC restarts in February.

---

# Meetings

## *Internet2*

### Spring, 2010, Member Meeting

Emerging trends in cyberinfrastructure development and new federal stimulus funding opportunities will share center stage in Arlington, Virginia, for the Internet2 Annual Spring Member Meeting, April 26 - 28, 2010. Learn more at the GPN Wiki or the Internet2 web site.

---

## *GPN Annual Meeting: GPN2010*

This year Memorial Day falls on May 31, so hold the following Wednesday, Thursday and Friday, June 2, 3, and 4, 2010, for the GPN Annual Meeting in Kansas City! Once again we will start with a reception on the evening of Wednesday, June 2, and conclude on Friday, June 4. Contributions will be welcome in the form of panels, posters, individual presentations, BoFs and workshops in networking, collaboration and cyberinfrastructure. If you are interested in serving on this year's program committee, please contact Greg Monaco: greg@greatplains.net.

---

## *MCBIOS Conference Graduate Student Award Winners Sponsored by GPN*

By Greg Monaco

GPN looks for opportunities to advance the research and education agenda of our members by partnering with regional science and education organizations. GPN sponsored Graduate Presentation Awards for the 2010 MCBIOS (MidSouth Computational Biology and BioInformatics Society) Conference. The conference was held on February 19 and 20 at Arkansas State University in Jonesboro, Arkansas. There were three awards in each of three categories: Oral Presentation; Poster Presentation, Computational Merit; and Poster Presentation: Biological Merit. You can see the list of winners, their topics and abstracts, here.

---

## To Contribute to the Newsletter

To contribute a story or to announce an upcoming meeting or event, send your contribution to Bill Mitchell (bill@greatplains.net), Greg Monaco (greg@greatplains.net) or Kate Adams (kate@greatplains.net).

---

## About the Great Plains Network

The Great Plains Network develops and maintains a high-performance network that meets the needs of the membership's research community; is a resource for cyberinfrastructure to the membership; and supports multi-institutional, multi-disciplinary research and education initiatives that require advanced cyberinfrastructure. GPN members includes 24 research institutions in nine states and also connects most other higher education institutions, school districts and public libraries in those states to Internet2.

**Subscription Information**: You can subscribe and unsubscribe to GPN News at http://www.greatplains.net/cgi-bin/majorcool.