

JET Meeting Minutes

May 17, 2011

I. Participants

Guy Almes	TAMU	galmes@tamu.edu
John Baird	DREN	Baird@hpcmo.hpc.mil
Abdella Battou	MAX	abattou@maxgigapop.net
Joe Breen	UofUtah	joe.breen@utah.edu
Ron Broersma	DREN	ron@spawar.navy.mil
Grover Browning	IU	gcbrowni@iu.edu
Joe Burescia	Esnet	joeb@es.net
Shane Cannon	Esnet	scannon@lbl.gov
Rich Carlson	DOE	rcarlson@ascr.doe.gov
Susan Coughlan	ESnet	smc@alcf.anl.gov
Vince Dattoria	DOE/SC	vince.dattoria@science.doe.gov
Sean Donelan	DHS/TIC	Sean.Donelan@dhs.gov
Dale Finkelson	I2	dmf@internet2.edu
Pat Gary	NASA/GSFC	James.P.Gary@nasa.gov
Andy Germain	NASA/GSFC-EOS	andrew.m.germain@nasa.gov
Michael Gill	NIH/NLM	gill@nlm.nih.gov
Hugh LaMaster	NASA/Ames	hugh.lamaster@nasa.gov
Paul Love	NCO	epl@sover.net
Brian Lyles	NSF	jbl2403@earthlink.net
Joe Mambretti	StarLight/MREN.	j-mambretti@northwestern.edu
Kevin McGrattan	Cisco	kmcgratt@cisco.com
Lindon Mercer	NRL	linden@cmf.nrl.navy.mil
Joe Metzger	Esnet	metzger@es.net
Grant Miller	NCO	miller@nitrd.gov
Mark Mutz	NOAA	mark.mutz@noaa.gov
Josh Poltra	UNC/CADA	
Glen Ricar	NLR	GR@nlr.nett
Anne Richeson	Qwest	anne.richeson@qwest.com
Chris Robb	Internet2	chrobb@internet2.edu
Brent Sweeny	GRNOC/Indiana U	sweeny@indiana.edu
Kevin Thompson	NSF	kthomps@nsf.gov
Brian Tierney	ESnet	bltierney@es.net
George Uhl	NASA/Goddard	george.d.uhl@nasa.gov
Alan Verlo	UIC/StarLight	verlo@uic.edu

Action Items

1. Maxine Brown will send the GLIF Map to JET once it has been revised.

Proceedings

This meeting of the JET was chaired by Vince Dattoria of DOE and Kevin Thompson of the NSF.

PerfSONAR Testbed: Joe Metzger

Joe Metzger reported on the status of perfSONAR testing among the Federal agency science networks. They are testing perfSONAR for IPv6 networks this year. The testing group has made progress in enabling v6 and Internet2 is doing interdomain testing with the University of Utah. A Web100 user LAN patch for NDT will be included in the next release. Discussion identified that:

- Andy Germain (NASA) is testing v6 with the University of Utah
- Internet2 nodes are not quite ready yet
- NLR nodes were turned down 5 months ago and do not support testing
- MAX has done v6 testing

Magellan Status Report: Susan Coughlan and Shane Cannon

Magellan is a DOE test bed funded by DOE under the American Recovery and Reinvestment Act (ARRA) to explore cloud computing for science. It is co-located at Argonne Leadership Computing Facility and the National Energy Research Scientific Computing Center (NERSC). Its mission is to determine the appropriate role for private cloud computing for DOE/SC midrange workloads. It deploys a test bed to investigate use of mid-range scientific computing clouds. It is being used to evaluate the effectiveness of cloud computing for a wide spectrum of DOE/SC applications.

Magellan began in September 2009. The core system was deployed from January 2010 to February 2010 at Argonne and from December 2009 to January 2010 at NERSC. User access began in March 2010 at Argonne for cloud uses and in April 2010 for cluster use at NERSC and in October 2010 for cloud use at NERSC. Other milestones included:

- Hadoop user access: December 2010 Argonne, May 2010 NERSC
- Joint demonstration between the two sites of MG-RAST: June 2010
- Nimbus was deployed June 2010
- OpenStack was deployed December 2010
- Magellan cloud ends September 2011
- ANI 100G is deployed October 2010
- Magellan ANI ends December 2010

Clouds are used in several different models:

- Infrastructure as a Service (IaaS),
- Platform as a service (PaaS)
- Software as a service (SaaS)

The Magellan distributed test bed employs IBM iDataPlex with 504 nodes at Argonne and 720 nodes at NERSC. Storage includes disk storage, archival storage and two classes of flash storage. The resources were architected for flexibility to support research.

Findings included that DOE HPC cluster software stacks are stable, mature, scalable, with integrated I/O and have a depth of tool availability. The leading open source cloud software stacks were deployed on Magellan and staff evaluations were done on application sets. Evaluation criteria included feature set, stability, infrastructure scalability, usability, manageability and sustainability. Evaluation did not include

performance. OpenStack, Nimbus, and OpenNebula did well and Eucalyptus 2.0 did well except for manageability.

Security frameworks, architectures and monitoring controls were deployed. Cloud systems provide unique challenges beyond traditional HPC clusters. The biggest risks were with the IaaS cloud model.

PaaS models have appeared for processing of large data sets, e.g., Azure and Hadoop. Hadoop for BioInformatics was evaluated using MapReduce. It was found that the new models are useful for addressing data intensive computing. It hides the complexity of fault tolerance and high-level languages can improve productivity.

Performance evaluations identified the importance of high performance networks to tightly coupled applications and commercial offerings optimized for Web applications are poorly suited for even small (64 node) MPI applications.

Evaluation of specific scientific applications (MG-RAST, STAR, and Joint Genome Institute) identified that:

- MG-RAST is well suited to clouds.
- STAR was pleased with real-time capabilities
- Joint Genome Institute: took 1 week for manual provisioning, operation was transparent to JGI users.

Summary findings included:

- IaaS clouds can require significant system administration expertise and can be difficult to debug
- Image creation and management are a challenge
- I/O performance is poor
- Workflow and data management are problematic
- Projects were ultimately successful simplifying future use of cloud computing
- Clouds enable rapid prototyping at larger-scale than desktop
- Clouds support tailored software stacks, different levels of service, and surge computing. They facilitate resource pooling

For the full briefing please see:

<http://www.nitrd.gov/Subcommittee/lsn/jet/material/Magellan-Coghlan-Canon-20110517.pdf>

Trusted Internet Connections (TICs): Sean Donelan of DHS

Sean Donelan gave an update on TIC architecture and requirements. Current TIC architecture version is TIC 2.0. The current TIC implementation responds to agency comments on TIC policies (e.g., what data/information can be exempt from TIC), design concepts, clarified language and intent, and technical clarification and security gaps.

New clarifications include:

- Unrestricted access data has no restrictions on access or usage
- Security default is deny, permit by exception, source egress IP address filtering is implemented, and FedRelay is provided for deaf and hard of hearing individuals
- TIC components are organized into functional blocks with different requirements: FIPS 199 for high and medium impact systems, SCIF facilities governed by DCID 6/9 Physical Security Standards Clarifications include:

- - TICAPs need to meet other Federal/OMB policies, e.g., for DNSSEC, IPv6, HSPD-12, NSTIC, incident reporting
- TICAP routing protocol authentication
- Expand VPN/Remote access alternatives
- E-Mail forgery detection
- DNS query filtering
- Data leak prevention policy documentation
- Network inventory process
- Operational exercise participation
- Vulnerability scanning, continuous monitoring

There are currently 19 TICAPs, largely 2 TICAPs per agency (NASA has 5 by exception). FY11 assessments of the TICAPs are planned based on TIC 1.0 capabilities.

For the full briefing please see:

http://www.nitrd.gov/Subcommittee/lisn/jet/material/TIC_UPDATE_2011-04-13.pdf

JET Roundtable

TransPac: Brent Sweeny

TransPac has no changes

ACE: Brent Sweeny

The ACE RFPs are being evaluated and announcements are imminent. Initially, Washington DC and New York City will connect to Europe.

DREN: Ron Broersma

The RFP responses for outsourcing the next generation DREN are being reviewed. DREN is holding a one-day meeting on IPv6 on August 16 at the DREN conference. DREN lost connectivity to Vicksburg due to a fiber cut this morning.

ESnet: Joe Metzger

ESnet is moving forward with the implementation of the ANI test bed. It is currently facing paperwork challenges.

Internet2: Dale Finkelson

SOX moved a 10 Gbps connector from Atlanta to Chicago. They retain an additional 10 Gbps connector in Atlanta.

New Network:

New network is implementing its initial installations this week along the East coast. They are placing new equipment in the PoPs to implement the Internet2 UCAN cooperation upgrades.

NIH: Mike Gill

Nothing new to report

NOAA: Mark Mutz

The Nwave buildout is proceeding. All 5 supercomputer sites have been connected. The Seattle Core-installation takes place next month. They are investigating connectivity to Hawaii.

NREN: Hugh LaMaster

NREN is working to participate in World IPv6 Day. NISN, Ames and Marshall are cooperating in this participation. August 16-17 there will be an IP Summit held at Ames Research Center.

NLR: Grover Browning:

Nothing new to report

ATDnet: Lindon Mercer

BOSnet was turned down and the equipment was redeployed to ATDnet sites. ATDnet has a 100 Gbps connection to DISA in operation.

Exchange Points

StarLight: Alan Verlo

StarLight is continuing work with GENI. An LHC1 community link to Chicago is being implemented. StarLight is supporting a demonstration between San Diego and the Czech Republic. A virtual overlay network is being deployed. Maxine Brown is upgrading the GLIF Map.

AI: Maxine Brown will send the GLIF Map to JET once it has been revised.

MANLAN

MANLAN is working with LHC1 to provide access through MANLAN.

MAX: Abdulla Battou

Max is working on 100 Gbps connectivity

C-Wave: Kevin McGrattan

C-wave is trying to implement IPv6 to participate in IPv6 World Day

USIGNITE: Glen Ricart

USIGNITE is working with OSTP to implement a future network. It is targeted for early this fall. It will provide advanced applications over advanced networking.

IGNITE is developing a test bed to provide:

- Advanced applications for homes and small businesses at 100 Mbps+, plus storage plus networking. It will utilize GENI resources including networking, storage, racks, perfSONAR and OpenFlow. It will build alliances between communities and campuses.

- Advanced test beds taking advantage of networks at the edges. Starting with GENI campuses and cities to build infrastructure. Initial cities are (Chattanooga, Lafayette, Salt Lake City, Cleveland, Navy yard (Philadelphia) and DCnet).
- Public-private partnerships at 200 campuses with 200 cities.

The NLR, GENI, and NewNet will provide networking for IGNITE. Currently IGNITE is looking for Phase 2 universities and cities. NSF will be providing funding for competitions to deliver advanced applications over advanced networking. July 7 Educause will hold a workshop for CIOs from universities to enlarge university participation in IGNITE. They are seeking to implement 100 Mbps+ at Layer 2 using OpenFlow, wireless, and open standards. They seek to create virtual network slices incorporating networks, campuses, and city resources. Uncompressed video is a focus for early demonstrations of capabilities.

Meetings of interest:

July 10-13	ESCC/Internet2 Techs Workshop, Fairbanks, AK
July 13-14	ESCC, Fairbanks, AK
July 26-28	GENI Engineering Conference 11, Chicago, IL
August 15-17	NASA IT Summit, San Francisco, CA
August 15-18	DREN Networking and Security Conference, Denver, CO
August 22-26	APAN, New Delhi, India
October 12-14	ARIN, Philadelphia, PA
November 12-18	SC11, Seattle, WA
March 2012	US-India workshop on network cooperation
January-February 2013	Techs in Paradise, Honolulu, HI

Next JET Meetings

June 21, 11:00-2:00, NSF
 July 11, 7:45-10:15PM, University of Alaska, Fairbanks. This meeting is being held in conjunction with the summer ESCC/Internet2 Techs Workshop