
PSAAP HPC Resources at Lawrence Livermore National Laboratory

Blaise Barney, Tom Spelce, LLNL

University of Texas, Austin
PECOS Center
July 21, 2011

LLNL-PRES-484644

Overview

- Who is the CRT?
- LLNL HPC Platforms
- Accounts & Access
- Software Environment
- File Systems
- Running Jobs
- Queues and DATs
- Support, Documentation & Training
- Alliance Communications
- Usage Stats
- The Future



Who Is The CRT?

- The Computer Resource Team (CRT) is the component of the NNSA PSAAP program that connects Alliance researchers to the High Performance Computing (HPC) resources required to perform their work.
- The CRT is comprised of a representative from each NNSA Lab who is familiar with their lab's computing resources, personnel and policies. The following individuals serve on the CRT:
 - Blaise Barney, LLNL
 - Rob Cunningham, LANL
 - Karen Haskell, SNL
- Our primary purpose is to provide assistance and guidance in all aspects related to the use of HPC resources located at LANL, LLNL, and Sandia.



What Does The CRT Do For You?

- Assist with the establishment and use of computer accounts
- Assist with accessing compute resources
- Provide HPC user documentation
- Provide technical support and referral to in-depth consulting
- Conduct monthly telecons to keep Alliance users up-to-date with account, access, policy, scheduling and technical issues, and to address issues with HPC platform usage
- Interface with other groups within the Labs, such as management, networking, system administration, storage, customer support, etc., to facilitate the effective support of Alliance users
- Track and facilitate the resolution of problems reported to each Labs' customer support "hotline"
- Provide training opportunities
- Collect and distribute monthly machine usage statistics
- Assist with scheduling and supporting special/dedicated runs
- Maintain a balance of machine usage between the Alliances
- Conduct Alliance visits to discuss HPC resources, user issues and to offer technical consultation and/or training
- Showcase Alliance research in the NNSA/ASC research exhibit booth at the annual SC conference



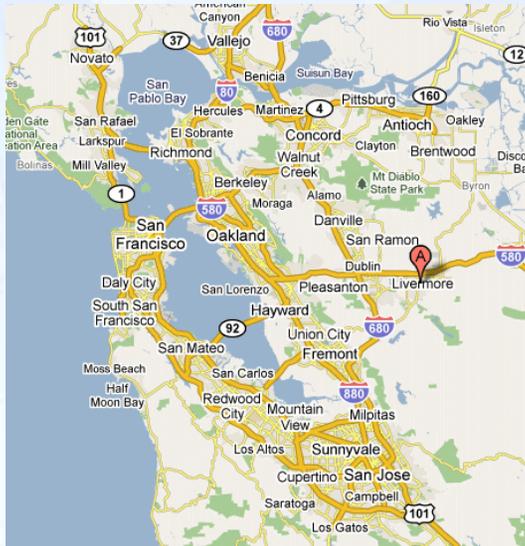
Lawrence Livermore National Laboratory (LLNL)

- Located in Livermore, CA
- Established in 1952
- @ 6,400 employees and contractors
- Managed by Lawrence Livermore National Security, LLC (LLNS)
 - Bechtel National, University of California, Babcock and Wilcox, the Washington Division of URS Corporation, and Battelle
- Budget @ 1.6 billion



LLNL's mission is to ensure the safety and security of the nation through applied science and technology in three key areas:

- **Nuclear Security** - ensure the safety, security and reliability of the U.S. nuclear deterrent and work to prevent nuclear proliferation and nuclear terrorism
- **International and Domestic Security** - develop capabilities to counter terrorism and other emerging 21st-century threats and technologies to enhance U.S. military effectiveness and better protect the soldier
- **Energy and Environmental Security** - advance science to better understand climate change and its impacts and develop technologies supportive of a carbon-free energy future



LLNL HPC Platforms - Summary

IBM BG/L and BG/P
AMD Opterons
Intel Xeons

System	Top500 Rank	Program	Manufacture/ Model	OS	Inter-connect	Avg Power Demand (KW)	Nodes	Cores	Memory (GB)	Peak TFLOP/s
Unclassified Network (OCF)										869.5
Sierra	40	M&IC	Dell	Linux	IB QDR	TBD	1,944	23,328	46,656	261.3
uBG/L	12*	ASC+M&IC	IBM BG/L	Linux	IBM	673	40,960	81,920	21,504	229.4
Hera (TLCC)	69	ASC+M&IC	Appro	Linux	IB DDR	492	864	13,824	27,648	127.2
Hyperion	458	ASC	Dell	Linux	IB DDR	248	1,152	9,216	11,520	90.3
Atlas (Peloton)	328	M&IC	Appro	Linux	IB DDR	573	1,152	9,216	18,432	44.2
Ansel	459	M&IC	Dell	Linux	IB QDR	TBD	324	3,888	7,776	43.5
Zeus (Peloton)		M&IC	Appro	Linux	IB DDR	143	288	2,304	6,912	22.1
Aztec		M&IC	Dell	Linux	N/A	TBD	96	1,152	4,608	12.9
uPurple		ASC	IBM SP	AIX	Federation	139	108	864	3,456	6.6
Prism		M&IC	GraphStream	Linux	IB SDR	55	128	256	2,048	1.2
Edge	72	M&IC	Appro	Linux	IB QDR	165	216	2,592	20,736	29.0
Hive		M&IC	Appro	Linux	IB DDR	7	6	96	1,088	0.9
Ebert		M&IC	Appro	Linux	N/A	7	6	96	192	0.9
Classified Network (SCF)										1,512.6
Dawn	16	ASC	IBM BG/P	Linux	IBM	1,260	36,864	147,456	147,456	501.4
BlueGene/L	12*	ASC	IBM BG/L	Linux	IBM	1,077	65,536	131,072	49,152	367.0
Muir	66	ASC	Dell	Linux	IB QDR	TBD	1,296	15,552	31,104	174.2
Juno (TLCC)	48	ASC	Appro	Linux	IB DDR	600	1,152	18,432	36,864	162.2
Graph	91	ASC	Appro	Linux	IB DDR	429	576	13,824	72,960	110.6
Coastal	107	Other	Dell	Linux	IB DDR	399	1,152	9,216	27,648	88.5
Eos (TLCC)		ASC	Appro	Linux	IB DDR	150	288	4,608	9,216	40.6
Minos (Peloton)		ASC	Appro	Linux	IB DDR	308	864	6,912	13,824	33.2
Rhea (Peloton)		ASC	Appro	Linux	IB SDR	157	576	4,608	9,216	22.1
Inca		ASC	Dell	Linux	N/A	TBD	96	1,152	4,608	12.9

System Category	TFLOP/s	% of Total
Unclassified	869.5	37%
Capability	534.9	62%
Capacity	200.3	23%
Serial	14.1	2%
Visualization	29.9	3%
Collaboration	90.3	10%
Classified	1,512.6	63%
Capability	956.8	63%
Capacity	271.0	18%
Serial	0.0	0%
Visualization	284.8	19%

Details: <https://computing.llnl.gov/?set=resources&page=index#hardware1>



LLNL Alliance Compute Resources

Hera

■ Hera Configuration

- AMD quad-core quad-socket
Opteron @2.3 GHz
- 127.2 TFlop system
- 864 nodes; 13,824 cores; 16 cores/node
- 32 GB memory/node
- 4X DDR Infiniband
 - 20 Gbits/sec rating
 - 2 - 2.5 Gbytes/sec non-block pt2pt
MPI in practice
- 64-bit architecture; CHAOS OS
- Lustre parallel I/O file systems (multiple)

■ Usage

- ASC and Institutional resource
- Alliance allocation @20% of the machine
- DAT runs



LLNL Alliance Compute Resources

UBGL

- UBGL = Unclassified BGL
- UBGL Configuration
 - IBM PowerPC 440 @ 700 MHz
 - 229.4 TFlop system
 - 40,960 nodes; 81,920 cores; 2 cores/node
 - 512 MB memory/node
 - Multiple networks
 - Torus for pt2pt MPI (2.1 GB/sec)
 - Tree for MPI collectives (2.1 GB/sec)
 - Barrier
 - 32-bit architecture; lightweight kernel OS
 - 690 TB Lustre parallel I/O file system
- Usage
 - ASC and Institutional resource
 - Alliance allocation @20% of the machine
 - DAT runs



LLNL Alliance Compute Resources

UDawn - **NEW!**

- UDawn = Unclassified Dawn
- UDawn Configuration
 - IBM PowerPC 450 @ 850 MHz - BG/P
 - 13.9 TFlop system
 - 1,024 nodes; 4,096 cores; 4 cores/node
 - 4 GB memory/node
 - Multiple networks
 - Torus for pt2pt MPI (5.1 GB/sec)
 - Tree for MPI collectives (5.1 GB/sec)
 - Barrier
 - 32-bit architecture; lightweight kernel OS
 - 1.2 PB Lustre parallel I/O file system
- Usage
 - ASC and Institutional resource
 - Available to Alliances



Account Requests

- **Begin with Sarape online: sarape.sandia.gov**
 - Most fields are self explanatory
 - Select LLNL and then "hera", "ubgl" and/or "udawn" from pull-down menu
 - Additional hints for completing the form can be found at: asc.llnl.gov/alliances/quickguideLLNLLANL.pdf
- **Coordinated and approved through the PECOS account authorizers Chris Simmons and Bob Moser (anyone else??) Then goes to LLNL.**
- **Account processing for non-US citizens requires additional time and “paperwork” - allow 30-90 days (plan ahead)**
- **Following approval, you will have two "identities" at LLNL**
 - Institutional (lab-wide) Official User Name (OUN) - smith25
 - Livermore Computing machine login name - jjsmith
 - Sometimes they can be the same



Access

■ One Time Password (OTP) token

- Provided when your account request is approved
- 2-factor authentication: static PIN + OTP passcode
- Used for all LC machines and some web pages
- More information & help: access.llnl.gov/otp



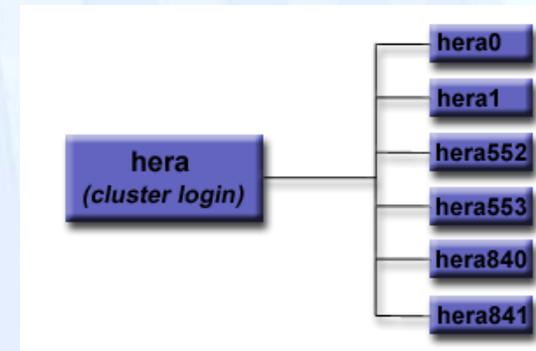
■ Virtual Private Network - Collaborator (VPN-C)

- VPN-C account provided along with your machine account. Required for all remote machine access.
- Install the software on your local machine
- Software, setup instructions and help available at access.llnl.gov/
- Windows, Mac and Linux software available, but no support for Linux.
- Authentication to VPN client is your OUN (smith25) and PIN + OTP passcode
- VPN requirement will be going away soon (discussed later)



Access

- SSH to LLNL machine after VPN is running on your local machine
 - `ssh hera.llnl.gov` `ssh ubgl.llnl.gov`
 - login with your LC machine name and PIN + OTP passcode
- Some internal web pages require similar authentication
- Automatic HPSS archival storage account
 - "ftp storage" from an LLNL machine
- File transfers
 - Use `scp` to/from LLNL machine
 - `ftp` initiated from within LLNL may also work (depends upon your local site permitting outside `ftp` connections)



Software Environment

- LC's Development Environment Group (DEG) – comprised of @ 15 staff responsible for the installation and support of software needed by most code development teams:
 - Compilers - see computing.llnl.gov/code/compilers.html
 - Debuggers
 - MPI libraries (MVAPICH and OpenMPI)
 - Profiling, tracing and performance analysis tools
 - Parallel I/O libraries
 - Misc. utilities, some math libraries
 - See computing.llnl.gov/?set=code&page=software_tools
- dotkit
 - Many of LC's software packages are managed by "dotkit". Same concept as "modules".
 - To see packages: `use -l`
 - To load a package: `use packagename`
 - More info: computing.llnl.gov/?set=jobs&page=dotkit



Software Environment

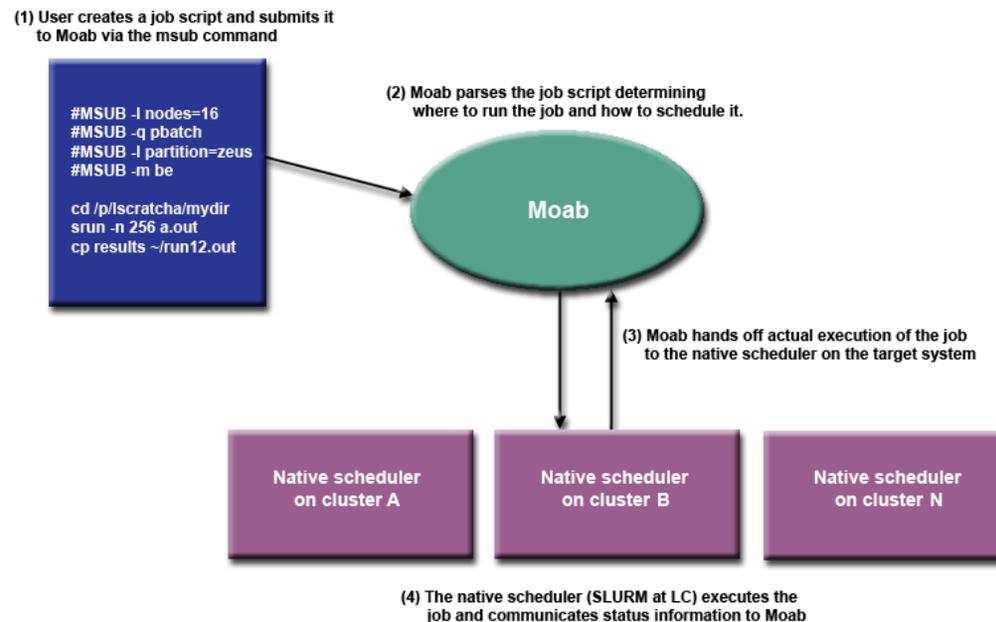
- CHAOS operating system on all Linux clusters
 - Based on Red Hat but differs in the following areas:
 - Modified kernel - to support high performance hardware, the Lustre file system, and other Livermore requirements.
 - New packages - support added for cluster monitoring, system installation, power/console management, parallel job launch, resource management, compilers, etc.
 - Modified packages - a few Red Hat packages are modified to implement timely bug fixes or enable them to work with CHAOS packages.
 - More info: computing.llnl.gov/linux/projects.html#chaos
- /usr/gapps
 - File space provided to code development teams that permits group/world sharing
 - Managed/supported by the individual code teams
- Visualization group – provides support for viz software and hardware, consulting, other viz related services
 - More info: computing.llnl.gov/vis/



Software Environment

■ SLURM and Moab Batch System

- SLURM - Native resource manager. One instance runs on each LC cluster. More info: computing.llnl.gov/linux/slurm/documentation.html
- Moab - Workload Manager - top level batch scheduler that "sits on top" of multiple clusters managed by their local SLURM resource manager. More info: computing.llnl.gov/tutorials/moab



File Systems Summary

File System	Description	Backup?	Purge?
Home directory plus special .snapshot directory	NFS mounted from all machines 16GB quota per user. Online backups past 2 days: cd to .snapshot	YES	NO
/tmp	Small (tens of GB) temporary file space local to a single node. Aliased to /usr/tmp and /var/tmp also.	NO	YES
/nfs/tmp2	Shared temporary file space across all platforms 134 TB; User quota in effect (400 GB)	NO	YES
Parallel file systems /p/lscratch{a,b,c,d} /p/gscratchc (ubgl)	Very large parallel file systems specifically for parallel I/O. No quota. Multiple file systems in 700TB - 1.3PB range	NO	YES
HPSS Storage	Virtually unlimited; accessible from all clusters	NO	NO



A Few Words on Parallel I/O

■ Hera

- Use `/p/lscratch{a,b,c,d}`. Maximum aggregate bandwidths:
 - `lscratcha` = 10 GB/sec
 - `lscratchb` = 7 GB/sec
 - `lscratchc` = 20 GB/sec
 - `lscratchd` = 20 GB/sec

■ UBGL

- Use `/p/gscratchc`. Maximum aggregate bandwidth = 10 GB/sec

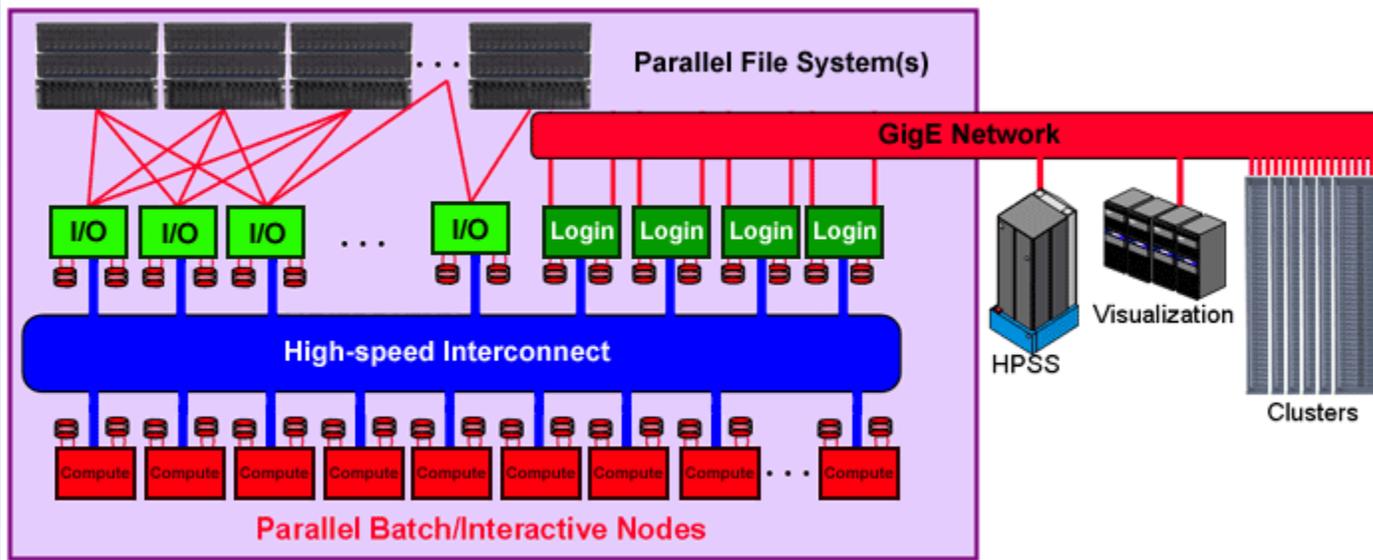
■ Useful information & helpful hints

- `/usr/local/docs/lustre.basics`
- `/usr/local/docs/lustre.striping`
- `/usr/local/docs/FileSystemUse`
- `/usr/local/docs/PurgePolicy.linux`
- <https://computing.llnl.gov/LCdocs/> --> I/O Guide
- Best performance: large (1MB) "chunks" of contiguous I/O
- Worst performance: small "chunks" of I/O to random offsets; writing lots of small files.
- Do NOT use nfs mounted directories (home directory, `/nfs/tmp2`) for parallel I/O



Running Jobs

- Typical LC machine configuration



- Most jobs require submission to the batch (pbatch queue) compute nodes via a Moab job control script. See computing.llnl.gov/tutorials/moab for details.
- Interactive (pdebug queue) jobs can be run without Moab on hera.
`srun -n64 -ppdebug a.out`
- Job limits: `news job.lim.hera` `news job.lim.ubgl`

Running Jobs

■ Login nodes

- Shared by multiple users
- Primarily used for interactive work such as editing files, submitting batch jobs, compiling, running GUIs, etc.
- Interactive use exclusively - login only nodes do not permit any batch jobs.
- DO NOT run production jobs on login nodes! It will impact other users.

■ Debug (pdebug) nodes

- Intended for short running, small, interactive jobs
- Not shared - owned exclusively by the running job
- Can be logged into while your job is running

■ Batch (pbatch) nodes

- Intended for all production work
- Comprise the majority of all nodes on a given cluster
- Not shared - owned exclusively by the running job
- Can be logged into while your job is running - useful for debugging
- Jobs must be submitted via Moab

■ I/O and service nodes - not available to users



Example Moab Job Script

```
#!/bin/csh
##### These lines are for Moab
#MSUB -l nodes=16
#MSUB -l partition=hera
#MSUB -l walltime=2:00:00
#MSUB -l gres=lscratchc
#MSUB -q pbatch
#MSUB -m be
#MSUB -V
#MSUB -o /p/lscratchc/joeuser/par_solve/myjob.out

##### These are run-time shell commands
date
cd /p/lscratchc/joeuser/par_solve

##### Note that srun command is required to launch parallel jobs
srun -n 256 a.out
echo 'Done'
```

See the LC Moab tutorial for details: computing.llnl.gov/tutorials/moab



Summary of Batch System Commands

Command	Description	Source	Scope
<u>msub</u>	Submit a job control script to the batch system. Many options used to specify job requirements and behavior.	Moab	Job
<u>srun</u>	Launch a parallel job from within a Moab job script or interactively	SLURM	Job
<u>mshow</u>	Display running, idle and blocked jobs	Moab	Grid
<u>showq</u>	Display running, idle and blocked jobs	Moab	Grid
<u>mjstat</u>	Display queue summary and running jobs	LC	Cluster
<u>squeue</u>	Display running jobs. Numerous options.	SLURM	Cluster
<u>checkjob</u>	Display detailed information about a single job	Moab	Job
<u>mdiag -j</u>	Display running, idle and blocked jobs	Moab	Grid
<u>sacct -j</u>	Display information about a running job, including multiple job steps	SLURM	Job
<u>sinfo</u>	Display a concise summary of queues and running jobs	SLURM	Cluster
<u>ju</u>	Display a concise summary of queues and running jobs	LC	Cluster
<u>mjobctl -h</u>	Place a queued job on user hold	Moab	Job
<u>mjobctl -u</u>	Release a user held job	Moab	Job
<u>mjobctl -c</u>	Cancel a running or queued job	Moab	Job
<u>canceljob</u>	Cancel a running or queued job	Moab	Job
<u>scancel</u>	Cancel a running or queued job	SLURM	Job
<u>mjobctl -m</u>	Change a job's parameters. Multiple options.	Moab	Job
<u>mshare</u>	Display bank/account allocations, usage statistics, and priorities.	Moab	Grid
<u>mdiag -u</u>	Display a user's bank/account information	Moab	Grid
<u>sreport</u>	Report usage information for a cluster, bank, individual, date range, and more.	SLURM	Grid
<u>mdiag -f</u>	Display fair-share scheduler statistics	Moab	Grid
<u>showstart</u>	Guesstimate when a job will start	Moab	Job
<u>showbf</u>	Show what batch resources are immediately available	Moab	Grid
<u>showstate</u>	Display a text map of jobs and which nodes they are running on	Moab	Grid
<u>smap</u>	Display a text map of jobs and which nodes they are running on	SLURM	Cluster



Queues and DATs

■ Queues

- Each machine has its own unique queue structure
- Specific limits based on job size, duration, and interactive vs. batch
- To view queue information, use the command:

```
news job.lim.machinename
```
- “Fair share” algorithm employed for resource delivery
 - Use more than your allocation and your service decreases
 - Use less, and it increases
 - Available cycles don’t go to waste

■ DATs

- Dedicated Application Time
- Scheduled runs that exceed normal queue limits
- Usually allocated on weekends
- To request: computing.llnl.gov/?set=forms&page=ASC_dat_form
- Coordinated through the Livermore Computing (LC) Hotline



Support, Documentation & Training

■ Livermore Computing Hotline

- Issues concerning use of LC's supercomputers
- Email, phone (925-422-4531) and walk-in during normal business hours: Mon-Fri, 8:00 am–noon, 1:00–4:45 pm Pacific time.
- Accounts/passwords: lc-support@llnl.gov
- Technical questions: lc-hotline@llnl.gov
- Referral to subject experts and vendors
- Problem tracking/database system
- Operations team provides support during off-hours

■ Machine status email lists, technical bulletins

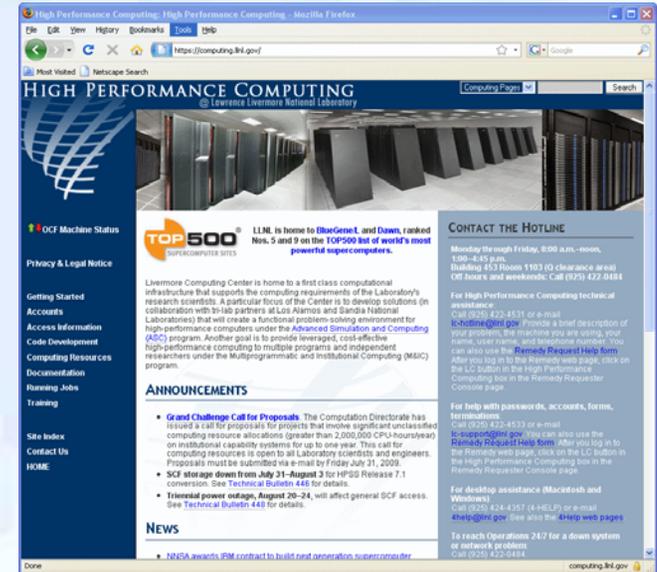
■ 4HELP Institutional Help Desk

- Desktop issues, VPN, network, OTP tokens
- 4help@llnl.gov 925-424-4357 (4-HELP)



Support, Documentation & Training

- **Documentation:** computing.llnl.gov
 - Most of what users need to know is available online via LC web pages
- **Training:** computing.llnl.gov/training
 - Online tutorials
 - Local workshops, seminars
 - Remote workshops
- **Recommended for new (and even old) users:**
 - computing.llnl.gov/tutorials/lc_resources
 - computing.llnl.gov/tutorials/moab
 - computing.llnl.gov/tutorials/linux_clusters



Support, Documentation & Training - **NEW!**

- **mylc.llnl.gov** - user portal
- Available now
- Work underway to develop job (batch) management tools
- Future work will provide tools for interacting with specific user applications

The screenshot displays the LORENZ user portal interface. At the top, it says "LORENZ your gateway into LC resources" and shows the user is logged in as "blaise" on "5/20/2011 10:18 AM". The interface is divided into several panels:

- current machine loads:** A list of machines with their current batch job status and free nodes. For example, "ansel: pbatch@91%, 27 nodes free".
- cluster utilization:** A section titled "Cluster Utilization for the past 14 days (hourly)" showing line graphs for various machines like ansel, atlas, aztec, etc.
- my accounts:** A grid of buttons representing different user accounts such as alator, ansel, atlas, aztec, etc.
- my groups:** A grid of buttons representing different groups like lglddev, blaise, csigweb, dweb, etc.
- my disk usage:** A table showing disk usage for various file systems. A note explains that orange bolded rows mean the quota is between 90% and 97%, and red rows are at 98% or above.

File System	# Files	# Bytes	Quota	pct
/g/g0	13070	7.3 GiB	16.0G	45%
/ns/lmp2	6	0.00 b	2.0T	0%
/p/lscatcha	0.0K	4.0K	n/a	0%
/p/lscatchb	0.0K	4.0K	n/a	0%
/p/lscatchc	0.0K	4.0K	n/a	0%
/p/lscatchd	0.0K	4.0K	n/a	0%
/usr/global	2606	291.6 MiB	n/a	0%
archive	155	14.7Gb	1.0Tb	1%
- my jobs:** A section for viewing job entries, currently showing "No data available in table".
- run a command:** A form to execute a command on a specific host. The host is set to "oslic" and the command is "top-b-n 1".
- my links:** A section with "LC Links" (Home, Resources, Status, Mgmt (IDM) Tool, Computer Coordinators) and "Custom Links" (No user links found).
- my historic CPU usage:** A section for viewing CPU usage history, currently showing "No data available in table".



Alliance Communications

- **Monthly telecons and email list (asap-crt@lanl.gov)**
 - Toll-free number: 866-914-3976 code: 187522#
 - Participation open to all Alliance team members, LLNL, LANL and Sandia
 - Forum for discussion/questions on user topics such as accounts, access, technical issues, machine schedules, etc.
 - Minutes and usage stats are distributed via our email list to all Alliances, ASC HQ and various staff & managers within the Labs
 - Let us know if you want anyone else at your Center added to our list - initially it includes only your POC and PI

- **LC Hotline**
 - User support as described previously

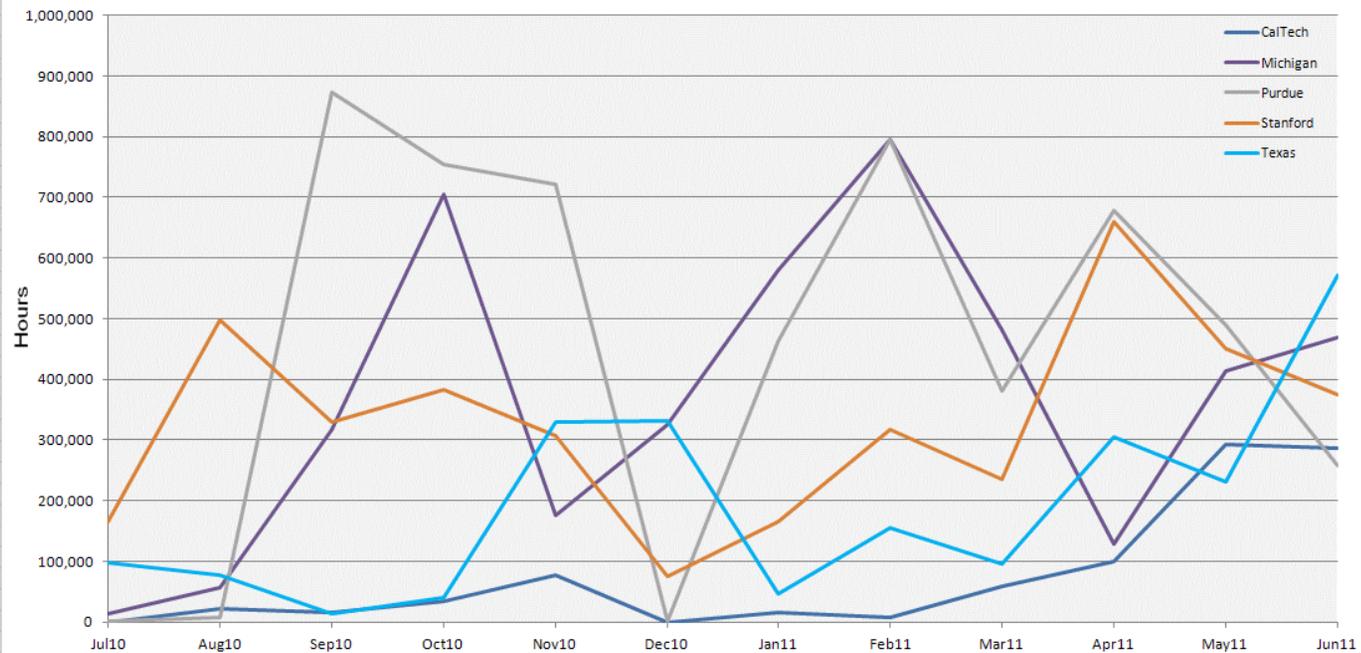
- **Your LLNL POC**
 - Blaise Barney blaiseb@llnl.gov 925-422-2578



Usage Stats

- Sent monthly with the telecon minutes
- User-level detail by Alliance per machine; Ad hoc reporting - ask your CRT
- Texas is #4 in Alliance usage on Hera.
- Past 12 months: 3 active Hera users, 0 active UBGL users

HERA	LLNL Usage by ASC Alliances												Cumulative
	Jul10	Aug10	Sep10	Oct10	Nov10	Dec10	Jan11	Feb11	Mar11	Apr11	May11	Jun11	
CalTech	0	22,563	14,780	33,361	76,580	0	14,979	7,573	57,703	99,218	292,185	286,547	905,489
Michigan	13,225	57,203	317,128	705,842	176,538	325,548	579,243	795,553	481,173	128,632	414,343	468,899	4,463,327
Purdue	1,443	7,744	874,076	754,441	720,787	639	462,759	795,189	380,463	677,556	489,902	258,193	5,423,192
Stanford	165,149	497,275	329,328	381,998	307,206	75,560	165,042	317,426	234,366	659,357	450,383	373,925	3,957,015
Texas	98,309	77,619	12,488	41,096	328,935	331,230	45,442	155,636	94,865	303,953	230,448	571,785	2,291,806
Total	278,126	662,404	1,547,800	1,916,738	1,610,046	732,977	1,267,465	2,071,377	1,248,570	1,868,716	1,877,261	1,959,349	17,040,829



NOTE: As of this chart has reformatted as usage for the Centers only. allocations an



The Future

■ TLCC2 - Tri-Laboratory Linux Capacity Cluster 2

- Next generation Tri-lab procurement. \$39M-\$89M contract has been awarded to Appro. 3-6 Petaflops total between the three Labs.
- Press release: <http://appro.com/press/view.asp?Num=207>
- Will include multiple unclassified Tri-lab systems with Alliance allocations
- Architecture:
 - 1 SU = 154/2464 compute nodes/cores; 49.7 Tflops
 - Dual 8-core Intel Sandy Bridge-EP Xeon processors
 - 32 GB DDR3-1600MHz RAM (8x4 GB DIMMS)
 - Qlogic QDR Infiniband
 - Intel EPSD Motherboard optimized for Appro's GreenBlade 5000 design
 - Intel Patsburg PCH chipset
 - GPU option: NVIDIA Tesla M2090 GPUs
- Availability: varies by cluster - starting later in 2011 thru 2012



The Future

■ Sequoia procurement

- 20 Pflop classified system; IBM BG/Q architecture
- 98,304 compute nodes; 16 cores/node; 1.6 million cores; 1.6 Pbytes of memory
- Available in 2012
- Expected to include a 2 Pflop unclassified system that will have an Alliance allocation
- Would like to say more, but there's this thing call an NDA...



The Future

■ Livermore Valley Open Campus (LVOC)

- NNSA supported collaboration between LLNL and Sandia Livermore to create an open, unclassified research and development space.
- Motivation for the LVOC stems from current and future national security challenges that require increased coupling with the private sector to understand threats and deploy solutions in areas such as energy and environmental security, economic security, cyber security, high performance computing and non-proliferation.
- Strong collaborations between industry, academia and the Labs
- Opportunities exist for university faculty and staff to get involved



The Future

■ Major unclassified network upgrade and changes!

- LLNL's network has been upgraded from 1 Gb/sec to 10 Gb/sec
- Being driven largely by the LVOC (previous slide)
- Open, unclassified collaboration network
- Most of LC's unclassified compute clusters will reside here (incl. hera & ubgl)
- VPN will no longer be required - direct SSH from the Internet permitted
- Foreign nationals (including DOE sensitive country FNs) will also be permitted along with a more streamlined account request process for FNs.
- Coming soon - full deployment expected by Oct 4, 2011. Some portions are already operational.
- **Benefits to Alliances**
 - File transfer rates should improve significantly. Further improvement expected via tuning network parameters.
 - Performance of GUIs (debuggers, performance tools) should improve
 - Non-US citizens permitted
 - Easier access - no VPN
- Who wants to help test/tune file transfers after upgrade is complete?



The Further Future: Exascale

■ PSAAP II

- Next NNSA/ASC Alliance program (2013). Three main focus areas
- Development and demonstration of technologies and methodologies to support effective Exascale computing in the context of science/engineering applications;
- “Predictive Science” based on verification and validation and uncertainty quantification (V&V/UQ) for large-scale simulations; and,
- Discipline-focused research needed both to further predictive science and enable Exascale computing.

■ DOE Lab (+ vendor) Collaborations

- Several DOE Labs are actively collaborating towards future Exascale machines from hardware, system software and application viewpoints:
 - LLNL / ANL / LBL / PNNL / (IBM)
 - LANL / Sandia / ORNL / (Cray)
- There are two primary “swim lanes” being pursued:
 1. Large number of very small power efficient cores
 2. Heavier weight cores tightly coupled to some type of accelerator



The Further Future: Exascale

■ DOE / Tri-lab Exascale Workshops

- Representatives from LLNL, LANL, Sandia, LBL, ORNL and several other institutions, have held two Exascale workshops to date.
- Working groups formed to identify and discuss challenges, strategies, technologies, co-design opportunities/dependencies, etc.

• Programming Models	• Systems Software
• Solvers, Algorithms, Libraries	• Hardware Architectures
• I/O, Networking, Storage	• Tools
• Visualization, Data Analysis	• Applications

■ Exascale activities internal to LLNL

- Workshops, seminars, interviews with code development teams...





Questions?

