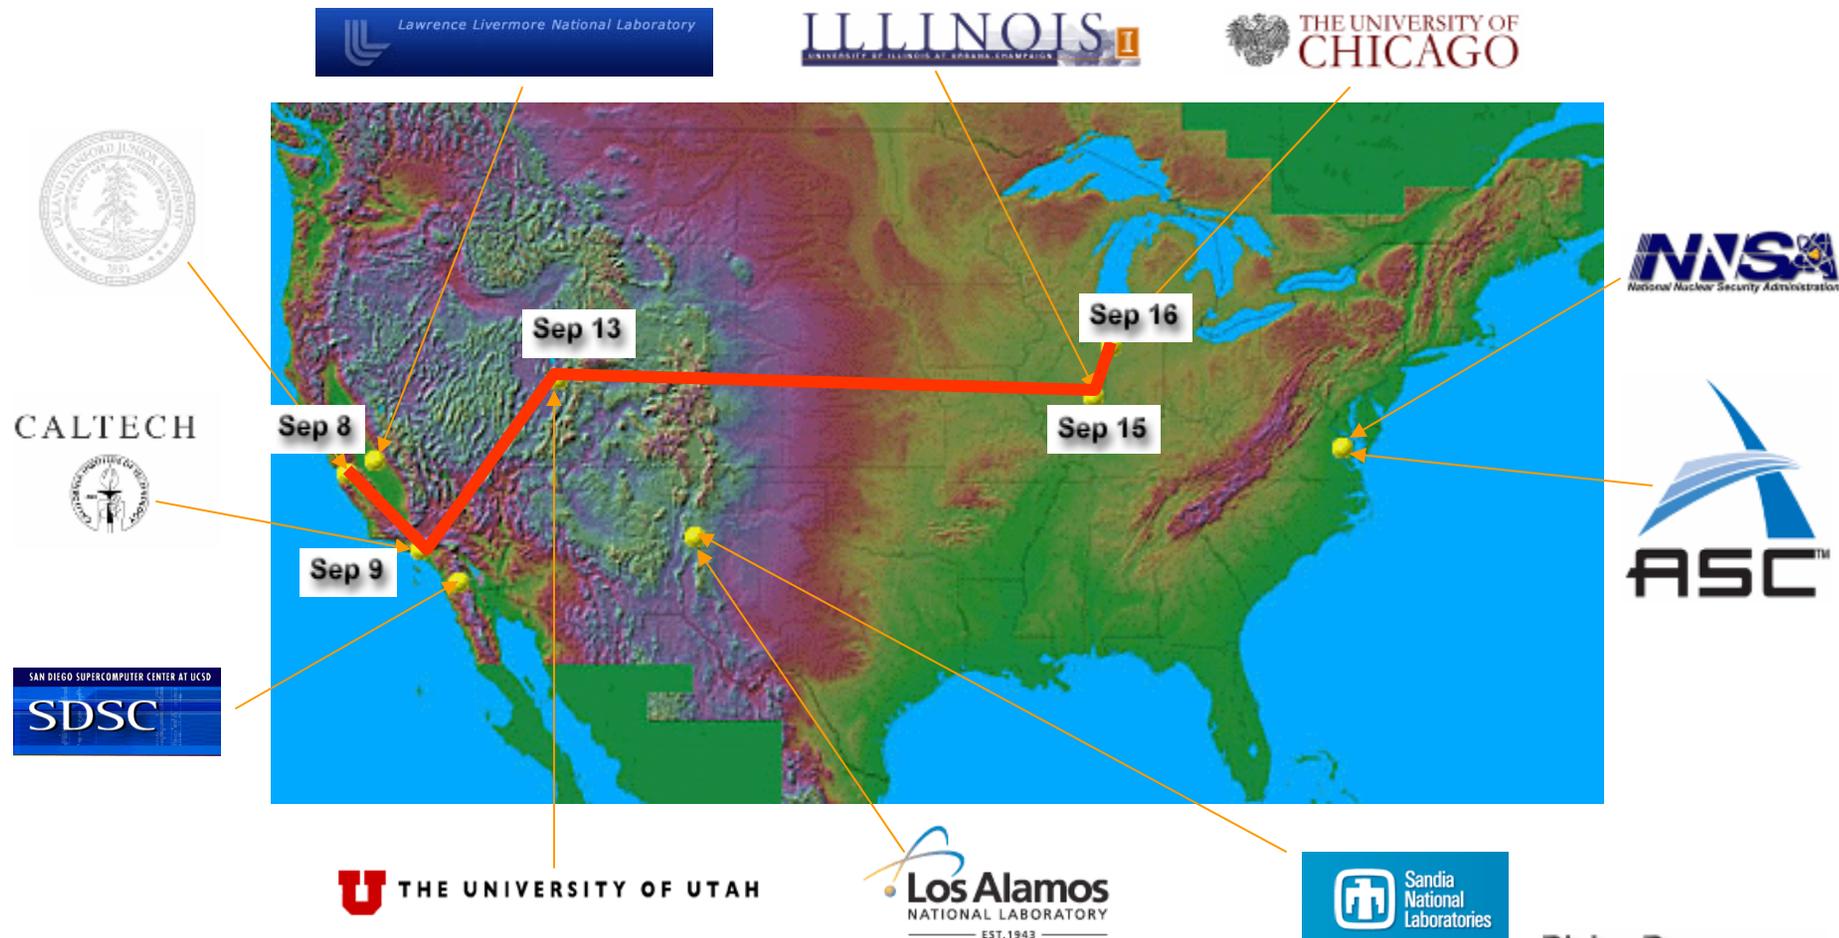


ASC Alliance Visits 2005

LLNL Resources and Environment

Sep 8-16, 2005



Blaise Barney
Dave Dannenberg
Richard Hedges
Jean Shuler

Agenda



- **What's Happened Since Last Year?**
- **Hardware Resources**
 - ALC
 - uP
 - Thunder
 - MCR
 - Parallel File Systems
 - HPSS
- **Usage Statistics**
- **Software Resources**
 - Compilers, Tools, etc
- **Training**
- **Issues We Are Aware Of**
- **Futures**

What's Happened Since Last Year?



- **Frost**
- **New Machine for Alliance Use**
 - uP
- **Existing Machines now enabled for Alliance Use**
 - MCR
 - Thunder
- **Result: significant net increase in computing resources available to Alliances at LLNL since last year. You spoke (loudly)...and we listened.**
- **ALC problems noted during last visit:**
 - System instability, not being able to make complete runs, Lustre problems
 - Giving different answers than other platforms
 - Not getting full allocation – competition with non-Alliance users



What's Happened Since Last Year?

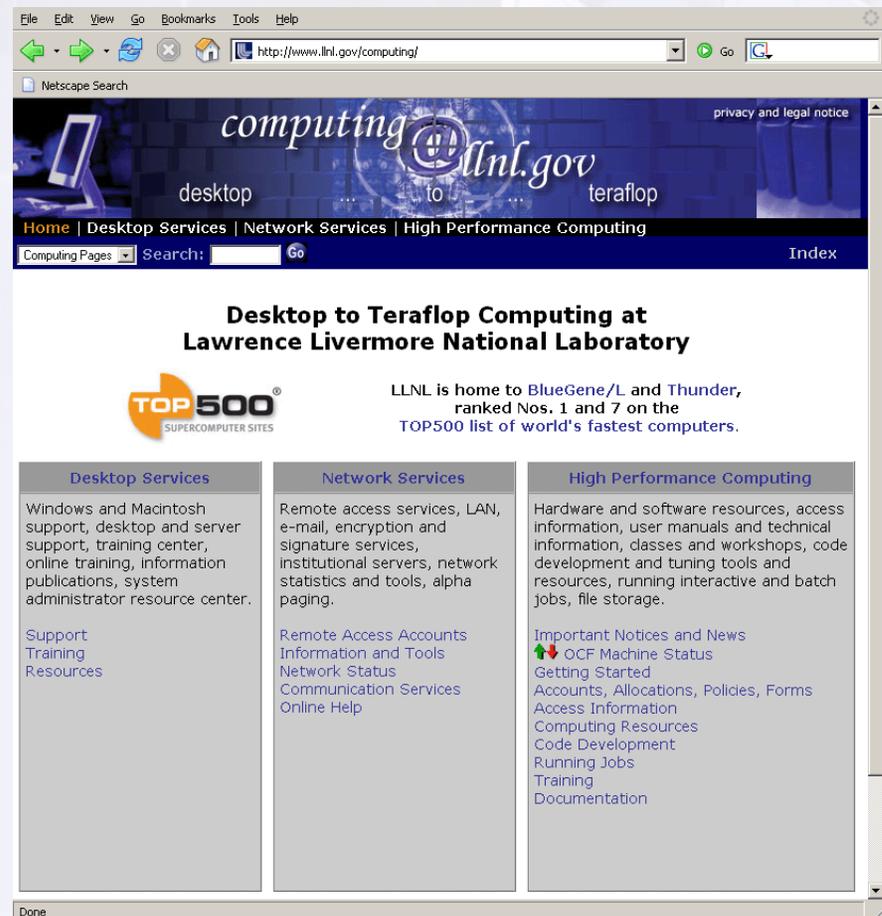


- **Progress on ALC problems:**
 - LC management provided Richard Hedges as POC to work directly with Alliance users and their code issues. After several months, the major issues had been resolved and codes are now running effectively on ALC.
 - LC “fixed” its batch allocation mechanism to insure that the Alliances received their full allocation (50% of machine) and got rid of scheduling problems with non-Alliance users.
 - Lustre and CHAOS have matured, resulting in greater stability over the past year.
 - “Wrong answers” were later discovered to be application related, not due to ALC specifically.
- **DAT times penalized post-DAT usage on LC machines. This was a common complaint among the Alliances.**
 - LC management agreed that DAT time should not penalize post-DAT usage. New banks were setup for DAT times only, so DAT usage no longer penalizes post-DAT runs.
- **Move into Terascale Simulation Facility (Feb):**
 - Bad news: major service outages for all users – including Alliances
 - Good news: move went smoother than expected

What's Happened Since Last Year?



- **Export controlled codes not permitted on ALC**
 - This still holds true for ALC
 - However, uP has no restrictions on running export controlled codes
 - MCR and Thunder (for DATs) can also run export controlled codes
- **Better method of announcing system changes, software upgrades, etc.**
 - Machine status email lists, login banners, news items
 - LC Homepage – especially “Important News and Notices” and “Machine Status” web pages.
 - Technical bulletins, viewgraphs from monthly LC Customer meetings
 - ...but there will always be times when we all get surprised!



What's Happened Since Last Year?



- **BG/L is going to be classified**

- Last year, we were soliciting candidate applications from the Alliances for running on the full 65K node (367 Tflop) unclassified machine.
- Since then, it has been decided that BG/L will ultimately become a classified resource with a very small number of programmatic specific applications permitted to run on it.
- There will be a “window” (approx. 11/01/05 – 2/1/06) when the full machine will be unclassified, and selected Alliance applications *may* be able to be run during this time. Current estimate: 20-35 days total *may* be available for the Alliances.
- After BG/L goes classified, there will be no unclassified BG/L resources available at LLNL for Alliance use.
- **QUESTIONS:** Would this be of interest? Have you run your codes on other BG/L platforms? Would you need much assistance getting your code to run on BG/L if interested? Would a “hands-on” workshop at LLNL be useful? How much of the machine would you want to run on (1/4, 1/2 or all 65K nodes?)
- See accompanying handout for the details...

What's Happened Since Last Year?



- **IPA access removed :**
 - Replaced by VPN-C
 - Some exceptions made for allowing IPA on a case-by-case basis
 - **Question:** any problems currently with VPN-C access? Or IPA, if you're one of the few exceptions?
- **SSH1 support gone (@ 9/7/05) – only SSH2 now supported**
- **Request to install GridFTP, netcdf and/or pnetcdf on ALC**
 - Now installed

Hardware Resources



Alliance Resources

System	Program	Manufacturer & Model	Operating System	Interconnect	Nodes	CPU's	Memory (GB)	Peak GFLOP/s
Unclassified Network (OCF)								
427,859								
BlueGene/L	ASC	IBM	Linux	IBM	65,536	131,072	32,768	367,002
Thunder	M&IC	California Digital	CHADS	Elan4	1,024	4,096	8,192	22,938
MCR	M&IC	Linux Network	CHADS	Elan3	1,152	2,304	4,608	11,059
ALC	ASC	IBM xSeries	CHADS	Elan3	960	1,920	3,840	9,216
uP	ASC	IBM SP	AIX 5.3	Federation	108	364	3,456	5,566
uBG/L	ASC	IBM	Linux	IBM	1,024	2,048	512	5,734
Gauss	ASC	GraphStream	CHADS	IB	256	512	2,048	2,458
Sphere	ASC	Rackable Systems	CHADS	Elan3	96	192	192	1,075
iLX	M&IC	RAND Federal	CHADS	N/A	67	134	268	678
PVC	ASC	Acme Micro	CHADS	Elan3	64	128	128	614
GPS	M&IC	HP GS32/ES45/ES40	Tru64 5.1f	N/A	33	160	356	277
Vertex	ASC	GraphStream	CHADS	IB	16	32	64	128
Snowhart	M&IC	IBM SP	AIX 5.2	Colony	16	128	64	114
Classified Network (SCF)								
131,125								
Purple	ASC	IBM SP	AIX 5.3	Federation	1,536	12,288	49,152	93,389
White	ASC	IBM SP	AIX 5.2	Colony	512	8,192	8,192	12,288
Lilac (xEDTV)	ASC	IBM xSeries	CHADS	Elan3	768	1,536	3,072	9,186
UM (pEDTV)	ASC	IBM p655	AIX 5.2	Federation	128	1,024	2,048	5,144
UV (pEDTV)	ASC	IBM p655	AIX 5.2	Federation	128	1,024	2,048	5,144
Acc	ASC	Rackable Systems	CHADS	N/A	160	320	640	1,792
GVz	ASC	Rackable Systems	CHADS	Elan3	64	128	256	717
Ice	ASC	IBM SP	AIX 5.2	Colony	28	448	448	672
Tempest	ASC	IBM Power5	AIX 5.2	N/A	12	84	480	555
Klein	ASC	GraphStream	CHADS	Elan4	10	20	40	136
SC Cluster	ASC	HP ES45	Tru64 5.1f	N/A	8	32	256	64
Tidalwave	ASC	SGI Onyx2	Linux 6.5.13f	16 IR2 Pipes	1	64	24	38

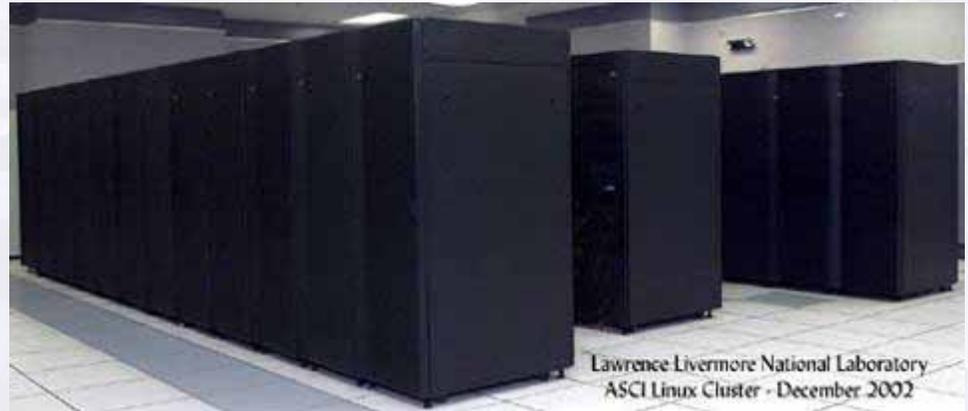
Unclassified		
Linux	420,902	58%
UNIX	6,957	2%
Classified		
Linux	11,831	9%
UNIX	119,294	51%

Unclassified		
Capability	393,674	52%
Capacity	26,841	6%
Serial	1,067	0%
Visualization	4,275	1%
Classified		
Capability	105,677	81%
Capacity	22,146	17%
Serial	2,411	2%
Visualization	891	1%



- **ASC Linux Cluster Details**

- Unclassified component of ASC Purple
- 9.2 TFlop system
- 960 nodes
- Each node has 2 Xeon processors at 2.4 GHz
- 4 GB memory per node
- Quadrics Elan3 switch
- 32-bit architecture running CHAOS operating system
- 72 TB Lustre parallel I/O file system



- **Alliance Resource**

- 1/2 of ALC is devoted to Alliance use
- 1/2 remains in testing and development mode for the Lustre file system

- **Configuration (still being called “experimental”)**

- 454 node batch pool for users. Max nodes TBD, max time is 24 hr
- 8 node debug pool. 30 min time limit weekdays, 2 hr limit weekends
- 462 node lustre testing pool
- DAT runs (including full system) commonly requested by Alliances



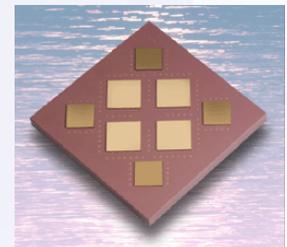
- **uP = unclassified Purple Details**

- A welcomed surprise to the Purple contract!
- 6.6 TFlop system
- 108 nodes
- Each node has 8 IBM Power5 p5-575 processors at 1.9 GHz
- 32 GB memory per node
- Federation switch
- 64-bit architecture running AIX operating system
- 140 TB GPFS parallel I/O file system



- **Alliance Resource**

- Became generally available on 7/19/05. Currently being used by the Alliances.
- Allocation @ 35% of the machine (currently at 20% due to another project)
- More than replaces the capacity lost by Frost's retirement:
200K hrs/mo @ 1.9 GHz vs 300K hrs/mo @ 375 MHz = @3.4 times the cycles





- **Running on uP – a few helpful hints:**

- Very similar to running on Frost:
 - AIX
 - POE
 - LCRM
 - User environment, file systems...
 - IBM compilers and MPI
- Some differences:
 - Must enable large page support (ledit command, -llpdata link option, or LDR_CNTRL=LARGE_PAGE_DATA=Y environment variable)
 - SLURM replaces LoadLeveler as native, underlying scheduler
- See /usr/local/docs/up.basics for some additional, important information and known problems.

- **Configuration (experimental at this time)**

- 99 node batch pool with 32-node/12 hr job limit
- 2 node debug pool with 2 hr limit
- DATs are available for request

Thunder



- **Thunder Details**

- M&IC Resource – not ASC
- 23 TFlop system
- 1024 nodes
- Each node has 4 Intel Itanium2 processors at 1.4 GHz
- 8 GB memory per node
- Quadrics Elan4 switch
- 64-bit architecture running CHAOS operating system
- 185 TB Lustre parallel I/O file system



- **Alliance Resource**

- Primarily for DAT runs

- **Configuration (usual)**

- 986 node batch pool with max job size of 986 nodes. Max run time 12 hrs on weekdays, 24 hrs on weekends.
- 16 node debug pool with 30 min time limit
- DAT runs commonly requested and available to Alliances via special DAT bank



- **MCR Details**

- M&IC Resource – not ASC
- 11.2 TFlop system
- 1154 nodes
- Each node has 2 Intel Xeon processors at 2.4 GHz
- 4 GB memory per node
- Quadrics Elan3 switch
- 32-bit architecture running CHAOS operating system
- 90 TB and 91 TB Lustre parallel I/O file systems



- **Alliance Resource**

- Primarily for DAT runs

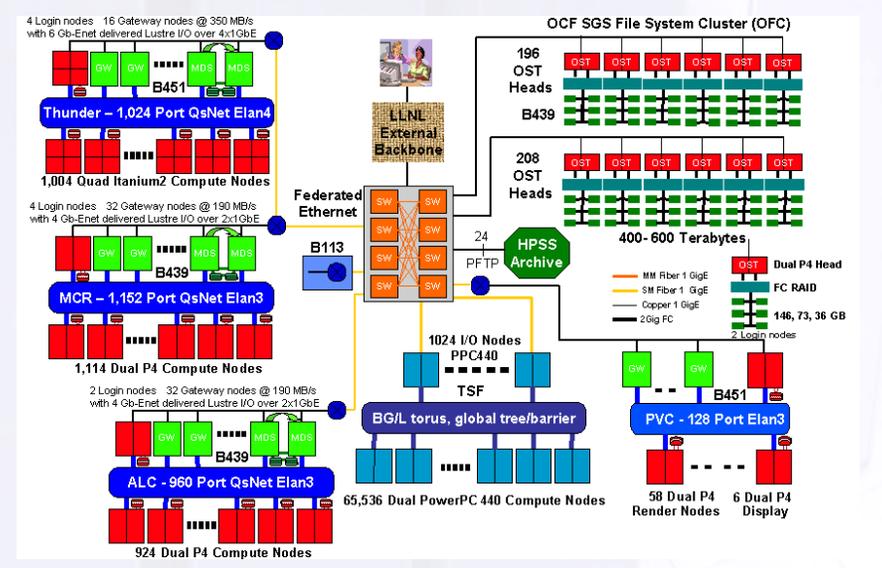
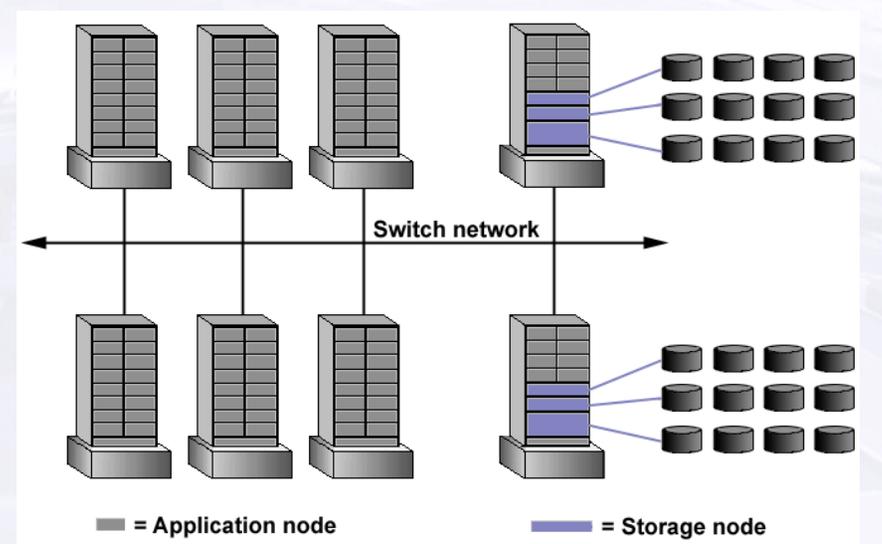
- **Configuration (usual)**

- 1048 node batch pool. Max job size/time varies:
513 nodes/12 hrs on weekdays and 768 nodes/entire weekend
- 64 node debug pool. Max nodes/time varies: 16 nodes/30 min on weekdays, 64 nodes/2 hrs on weekends.
- DAT runs commonly requested and available to Alliances via special DAT bank

Parallel File Systems

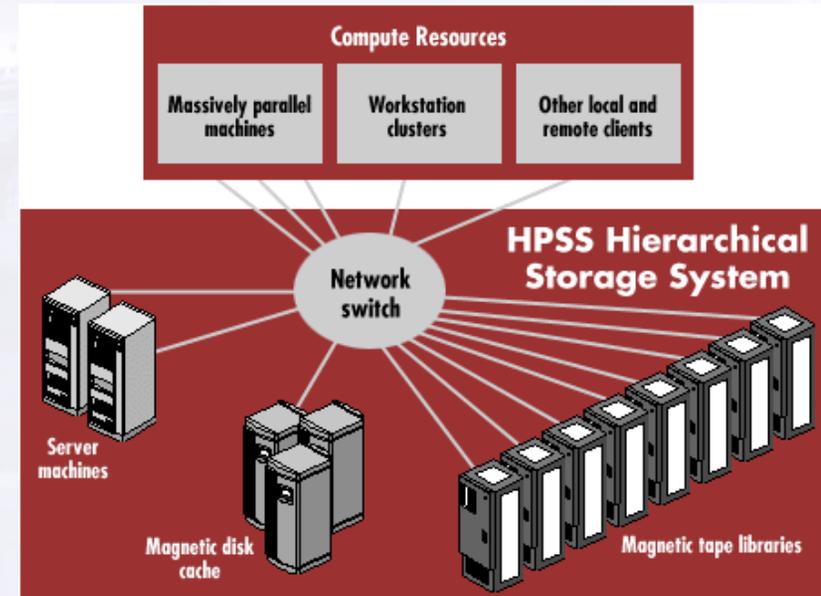


- Only file systems that should be used for parallel I/O!
- Subject to purging!
- GPFS
 - IBM's General Parallel File System
 - All LLNL ASC production systems have their own, multi-terabyte GPFS file system(s)
- Lustre
 - Linux cluster based parallel file system from Cluster File Systems, Inc.
 - Used on most LC Linux production systems. Also used on BG/L systems.
 - LC goal: Lustre based SAN
 - Still being developed and tested
 - See www.lustre.org for more info





- Integrated into the OCF and SCF gigabit ethernet networks
- High-speed, optimized file transfers between LC machines and HPSS storage.
- **Some metrics**
 - OCF: 2.2 PB capacity; @1.03 PB used
 - SCF: 3.3 PB capacity @1.16 PB used
 - 250 MB/s aggregate writes
 - 150 MB/s on a per file basis



OCF System Specifications							
Software	Used Capacity	Archive Capacity	Disk Cache	#STK Robots	#9840 Drives	#9940A Drives	#9940B Drives
HPSS 5.1	1.03 petabytes	2.2 petabytes	45 terabytes	3	35	25	40

SCF System Specifications							
Software	Used Capacity	Archive Capacity	Disk Cache	#STK Robots	#9840 Drives	#9940A Drives	#9940B Drives
HPSS 5.1	1.16 petabytes	3.3 petabytes	50 terabytes	5	40	30	55

Usage Statistics



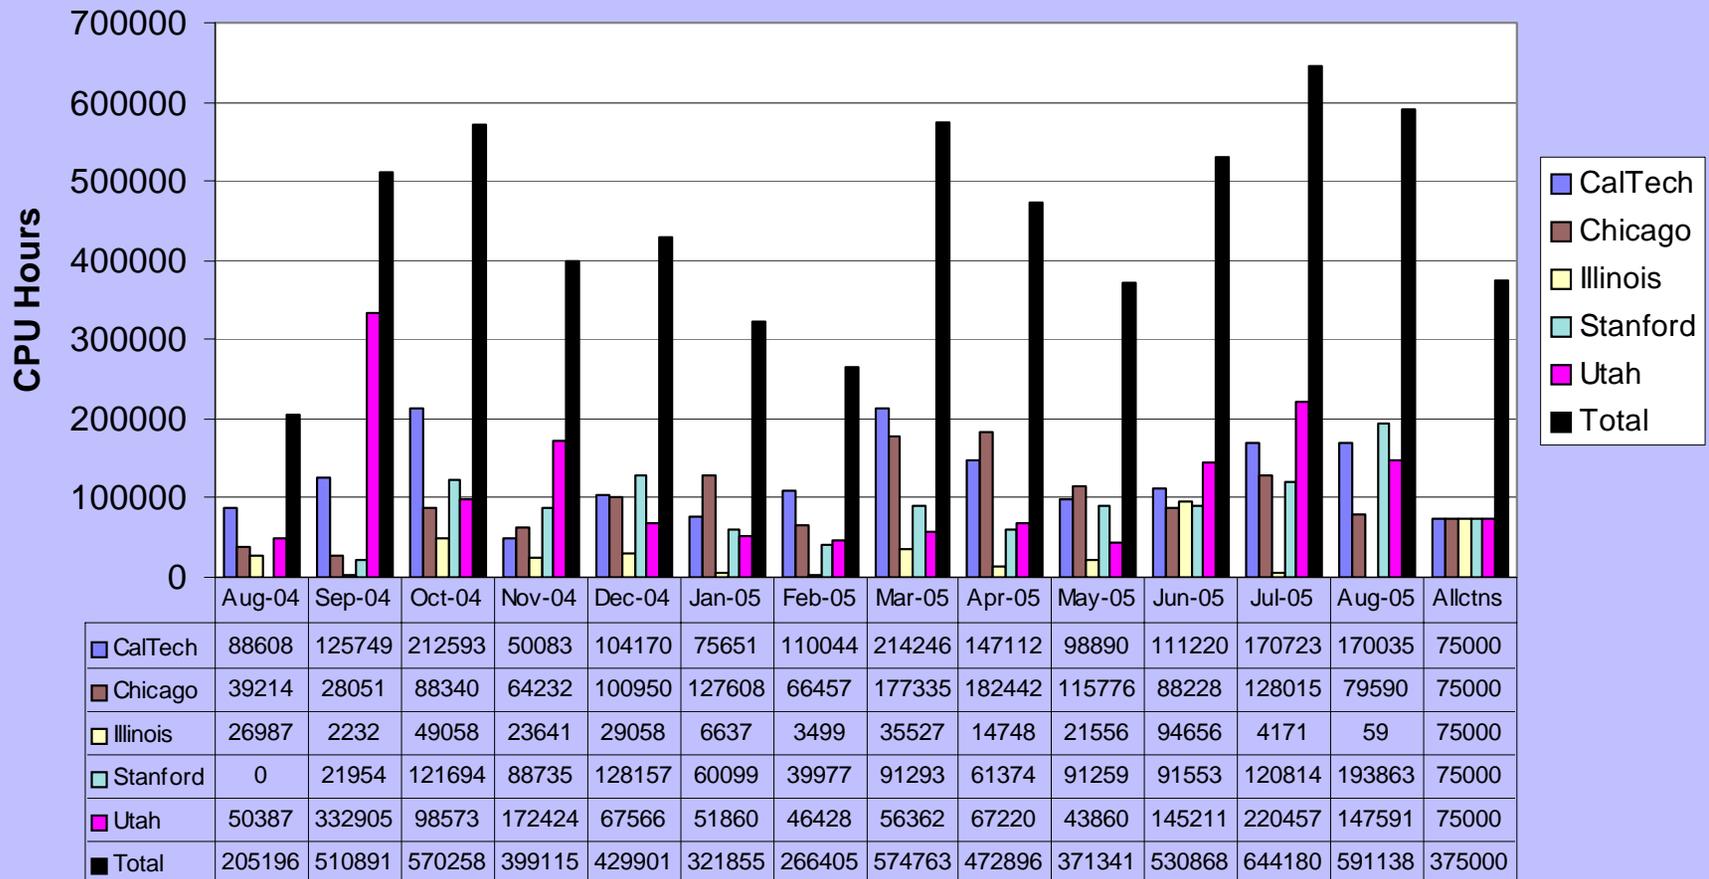
- **Frost – retired on 6/30/05. Heavy usage up until that point.**
- **ALC – each Alliance is allocated 75000 cpu hours/month. Between 8/04 and 8/05, the Alliances used this much of their allocation:**

- CalTech	155%
- Chicago	124%
- Illinois	32%
- Stanford	139%
- Utah	109%
- **MCR and Thunder: now available for Alliance DAT times**
- **uP: has only been GA (generally available) since 7/19/05. DAT time is available, but none requested to date.**
- **DAT time – only recently began keeping track of time separately using the “new” DAT banks scheme.**

Usage Statistics



**Alliance ALC Usage Data
Aug 2004 - Aug 2005**

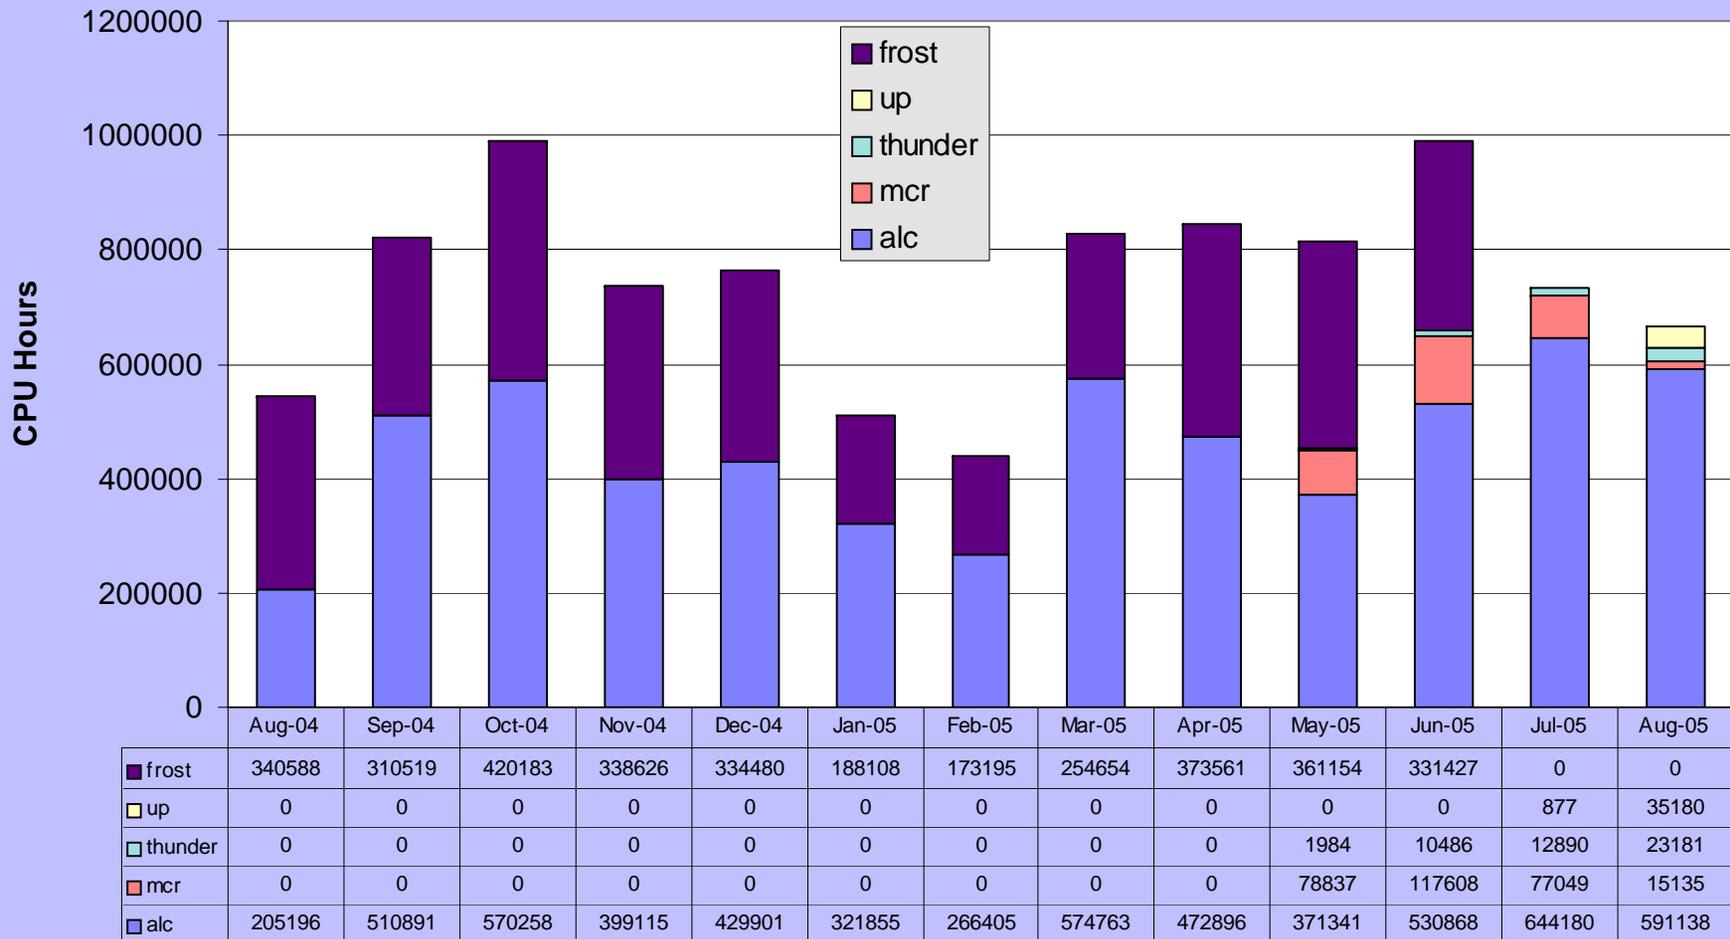


* chart includes DAT times

Usage Statistics



Total Alliance Resource Usage
Aug 2004 - Aug 2005



Frost cpu-hr = 5.4 Tflop/hr
Thunder cpu-hr = 20.2 Tflop/hr
ALC cpu-hr = 17.3 Tflop/hr

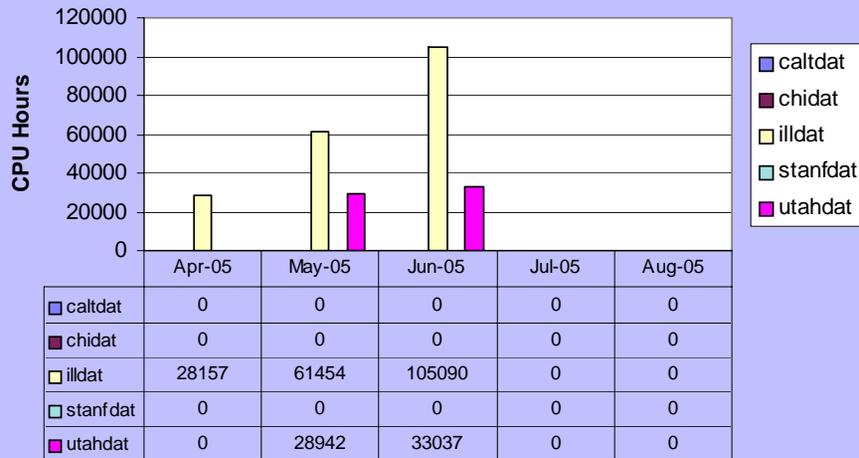
uP cpu-hr = 27.4 Tflop/hr
MCR cpu-hr = 17.3 Tflop/hr

* chart includes DAT times

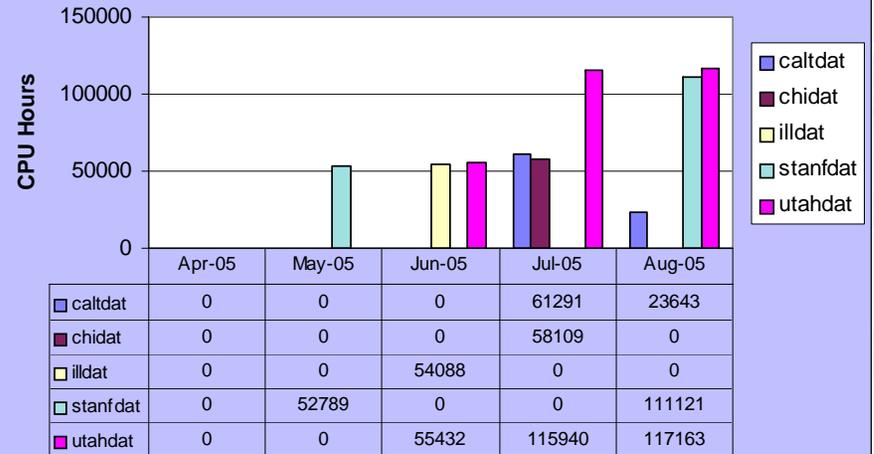
Usage Statistics



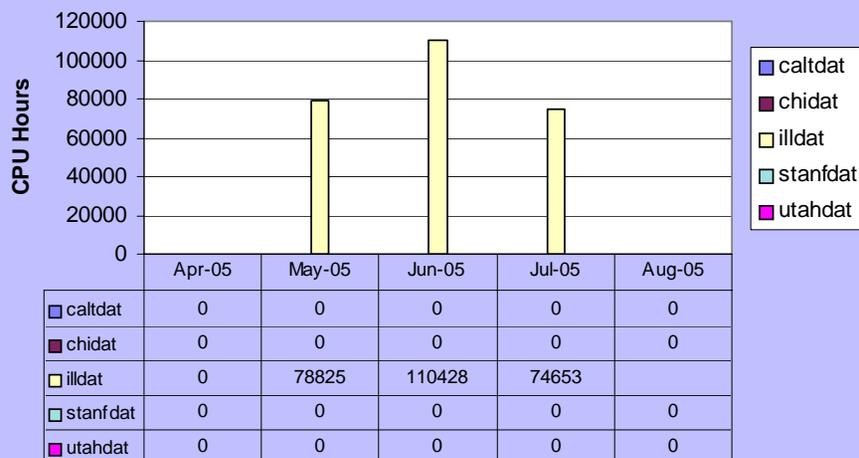
Frost DAT Usage



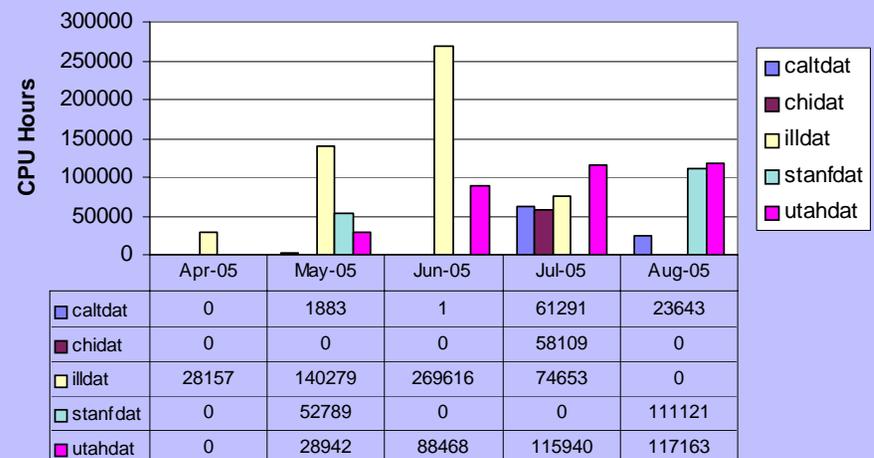
ALC DAT Usage



MCR DAT Usage



All Alliance DAT Usage



Software Resources



- **uP**

- Currently running AIX 5.2
- Parallel Environment 4.2
- Upgrade to AIX 5.3 planned (Sep-Oct)

IBM AIX 5L
UNIX OPERATING SYSTEM

AIX 5L™

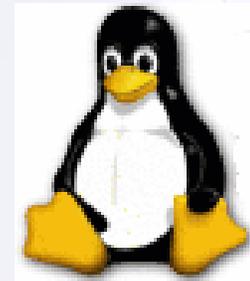
- **ALC, Thunder, MCR**

- Currently running CHAOS 2
- LC's developmental Linux cluster OS
- Based upon RedHat Linux
- Plans for CHAOS 3 upgrade (Oct-Dec)



- **LCRM / SLURM batch/schedulers**

- LCRM upgrade to version 6.13 (tech bulletin #389)
- SLURM: Simple Linux Utility for Resource Management
 - Collaboration between Livermore and Linux NetworX
 - Used on all LC Linux clusters and ASC Purple Power5 systems (like uP)
 - Native scheduler that interfaces to LCRM



Software Resources



- **Compilers**

- uP: usual IBM C/C++ (version 6.0) and Fortran (version 8.0) compilers. Gnu compilers (version 3.3) also available, as are newer IBM compilers.
- ALC, Thunder, MCR: usual Intel C/C++ (version 8.1) and Fortran (version 8.1). Version 9.0 Intel compiler and GNU compilers (version 3.3.3) available on all three Linux platforms also. Portland Group compilers (multiple versions up to 6.x) available on MCR and ALC.

- **Debuggers, correctness tools**

- Assure
- TotalView
- décor
- Umpire
- Insure++
- ZeroFault

- **Profiling, tracing, and performance analysis**

- TAU
- PAPI
- Dimemas
- PE Benchmarker
- Vtune
- HPM
- prof
- Paraver
- gprof
- MPX
- mpiP
- Xprofler
- IBM HPC Toolkit
- Vampir/Guideview
- Jumpshot

- **What are we missing? Send a request to the hotline.**
- **Any problems with using any of the above? Let us know.**

Training



- **Regular introductory workshops at LLNL**

- Parallel programming
- Linux & Compaq clusters
- POE
- Pthreads
- TotalView
- LC resources and environment
- IBM hardware/software
- MPI
- OpenMP
- ...

- **Other workshops**

- Performance analysis tools and topics for the IBM SP
- MPI performance topics
- Vampir/GuideView, Paraver, Dimemas
- Advanced MPI
- Advanced TotalView
- Python, Linux topics
- Intel compilers, Vtune

- **Tri-lab and Alliance workshops**

- Combined training for multiple ASC platforms held at any Tri-lab or Alliance location
- Customized workshops delivered at Alliance's location



Training



- **Collaborative HPC training via the Access Grid and informal HPC Training Consortium**

Berkeley Nat'l Lab

Sandia Nat'l Lab

Los Alamos Nat'l Lab

Ohio Supercomputer Ctr.

Texas Adv. Comp. Ctr.

Argonne Nat'l Lab

Livermore Nat'l Lab

Princeton Plasma Physics Lab

NCSA

SDSC

U. Hawaii

U. Michigan

Purdue U.

General Atomics

Maui HPC



HPC-training-consortium@purdue.edu



Issues We Are Aware Of



- **Fortran I/O to Lustre can be extremely slow (not buffered).**
 - Use `-assume buffered_io` compile flag
- **Lustre I/O failures when some disks become full.**
 - Has impacted DAT runs on ALC. Consultants now have a tool to check file system capacity prior to DAT runs.
- **File system failures still occurring**
- **Data transfer rates to LLNL to/from Alliances and other labs reported as “OK” last year. Chicago has recently reported good transfer rates using GridFTP. Any news to report this year?**
- **Frustration with the batch scheduling system**
 - It's impossible to predict when my job will run!
 - My small job won't run even though there are lots of spare nodes!
 - Running in “standby” mode gives no guarantee job won't get killed
 - The truth: batch scheduler is complex and dynamic. What is true at this moment may not be true two minutes from now – see next slide.

Issues We Are Aware Of



One example output from the pstat command – may provide some helpful predictive information, but it can change at any time depending upon a number of factors.

```
alc1{blaise}45: pstat -m alc -s PRIORITY -o jid,name,nodes,priority,timeleft,user,status,maxcputime
  JID NAME                NODES PRIORITY  TIMELEFT  USER      STATUS      MAXCPUTIME
38922 batchcom2           128   0.751    23:21    jhnieder  RUN         24:00
35550 methanelm_291d     450   0.736     N/A     nguyen   *WCPU      24:00
39273 turbineFull        152   0.701     N/A     schluter *WCPU      24:00
36111 crosswind           100   0.647     N/A     spinti   *WCPU      24:00
38924 batchcom3           64   0.621    13:21    jhnieder  RUN         14:00
38919 batchcom             64   0.621     3:18    jhnieder  RUN          8:00
39024 run_zvch              96   0.614     N/A     shunn1   *WCPU      24:00
38894 afl_res              75   0.571    23:11    calder   RUN        23:50
38686 runpF3d_alc         32   0.569     N/A     dorr     *WCPU      12:00
36141 runpF3d_alc         32   0.569     7:18    dorr     RUN        12:00
38834 methane_small       23   0.491     N/A     spinti   *WCPU      24:00
38795 sl2.a4                32   0.470     N/A     matheoul *WCPU      12:00
39275 dlrFlameA             32   0.470     6:28    mihme    RUN        11:00
39277 dlrFlameB           32   0.469     6:58    mihme    RUN        11:30
38844 sl2.a3              16   0.438    11:21    matheoul RUN        12:00
38800 jet35                  4   0.433     7:18    smartin  RUN        12:00
38796 sl2.a2                4   0.414     1:18    matheoul RUN          6:00
38804 jet39                 4   N/A       N/A     smartin  *DEPEND    12:00
39237 jet31                 4   N/A       N/A     smartin  *DEPEND    12:00
38817 runpF3d_alc         32   N/A       N/A     dorr     *DEPEND    12:00
38594 sl2.a5              128   0.000     N/A     matheoul *HELDu     12:00
38920 sl2.a2                4   N/A       N/A     matheoul *DEPEND     6:00
38805 jet30                 4   N/A       N/A     smartin  *DEPEND    12:00
38801 jet36                 4   N/A       N/A     smartin  *DEPEND    12:00
```

Suggestions? Change ALC job limits policies?

The Future



- **No major new machines for Alliance use on the near horizon.**
- **Tri-lab, multi-platform training? Or, individual Red Storm and ASC Purple (uP) training sessions? Where? When?**
- **Interest in BG/L work during 11/1/05 – 2/1/06 window?**
- **SC|05**

SC|05 - ASC Exhibit Booth



SC|05 Gateway to Discovery

Washington State Convention and Trade Center Seattle, WA
November 12-18 2005

CONTACT

About | Interactive Schedule | Programs | Registration | Exhibits | Initiatives & Challenges | News & Press | Hotel & Travel

Register Now
Interactive Schedule



- Seattle, Nov 12 – 18. Exhibits 15th – 17th
- Call for participation!!
<http://www.lanl.gov/conferences/sc05/index.html>

