








NERSC Job Data



Richard Gerber
Zhengji Zhao
NERSC User Services

June 12, 2014

What Does NERSC Collect?



- ALTD
 - Track library usage both at compile and run time
- Torque Logs
 - Job information, accounting
- ALPS Logs
 - Track applications run time data and options on the Cray systems
- Darshan
 - IO profiling data
- IPM
 - MPI profiling data
- Performance Monitoring
 - Monitoring system performance over the life time of the machines
- LMT
 - Lustre data

Expose job data via the web



- **We try to make as much data available as possible via the web**
 - For users to track usage
 - For users to check resource utilization
 - For users to monitor performance
 - For staff to help debug jobs
 - For summary reports
- **The following are web screen shots**
 - All data collection is transparent to users



JOB LOGS & ANALYTICS

[Queue Status](#)

A listing of jobs running and waiting in the queues (NERSC login required).

[Queue Wait Times](#)

Batch Queue Wait Times

[Completed Batch Jobs](#)

Data from the job scheduler logs, including links to data from the job launcher logs (for Cray) and IPM (NERSC login required).

[Completed Parallel Jobs \(Cray aprun\)](#)

Data from the ALPS job launcher for Cray systems (NERSC login required).

[Usage Reports](#)

Machine Usage Reports

- [Batch Job Statistics](#)
- [Parallel Job Statistics \(Cray aprun\)](#)
- [Hopper Hours Used](#)
- [Edison Hours Used](#)
- [Carver Hours Used](#)
- [Historical Data](#)
- [Hopper Job Size Charts](#)

Displaying 500 jobs that completed between Nov-11-13 00:00 and Nov-18-13 21:59

#	Host	JobID	Job Name	User	Nds	Complete	Wall hrs	Raw Mach Hrs*
0	edison	516239	SIC-n	carnevav	8	11/18/13 21:55	0.905	173.76
1	hopper	6760464	353533-1	yosh	1	11/18/13 21:55	1.010	24.23
2	hopper	6733432	Cu1_In1_Te... [Full Name]	matcomp	2	11/18/13 21:55	168.025	8,065.20
3	hopper	6764387	Li_bcc	angsten	1	11/18/13 21:55	0.043	1.03
4	hopper	6764389	Li_bcc	angsten	1	11/18/13 21:55	0.053	1.28
5	edison	517684	ms12	arindam	8	11/18/13 21:55	0.977	187.63
6	carver	7441883	my_job	aschmah	1	11/18/13 21:55	0.014	0.01
7	carver	7441858	my_job	aschmah	1	11/18/13 21:55	0.013	0.01
8	carver	7441842	my_job	aschmah	1	11/18/13 21:55	0.013	0.01
9	carver	7441851	my_job	aschmah	1	11/18/13 21:55	0.013	0.01
10	carver	7441849	my_job	aschmah	1	11/18/13 21:55	0.013	0.01
11	carver	7441860	my_job	aschmah	1	11/18/13 21:54	0.013	0.01
12	carver	7441839	my_job	aschmah	1	11/18/13 21:54	0.013	0.01
13	carver	7441835	my_job	aschmah	1	11/18/13 21:54	0.013	0.01
14	carver	7441856	my_job	aschmah	1	11/18/13 21:54	0.013	0.01
15	carver	7441854	my_job	aschmah	1	11/18/13 21:54	0.013	0.01
16	carver	7441887	my_job	aschmah	1	11/18/13 21:54	0.013	0.01

IPM summary

Executable	/global/homes/j/jdeslip/codes/BGW_trunk/bin/sigma.cplx.x				
Number of tasks	1,179	Aggregate GFlop/sec	0.0339	Average GFlop/sec/task	0.0000
Average wall secs	3.148e+03	Aggregate memory (GB)	754.2820	Average memory/task (GB)	0.6398
Average MPI secs/task	2.170e+03	MPI time %	68.94	Aggregate MPI calls made	1.878e+08

IPM summary statistics - 1179 tasks

Metric	Sum over all tasks	Average (per task)	Task CV (%)	Task Minimum	Task Maximum
Aggregate Floating Point Operations (Flop x 10**9)	1.067e+02	9.049e-02	6.48	1.152e-02	9.992e-02
GFlop/sec	3.389e-02	2.874e-05	6.48	3.661e-06	3.174e-05
Maximum Memory Usage (GBytes)	7.543e+02	6.398e-01	1.64	5.851e-01	8.286e-01
Time Spent in MPI Routines (sec)	2.559e+06	2.170e+03	3.90	8.829e+02	3.086e+03
Wallclock Time (sec)	3.712e+06	3.148e+03	0.00	3.148e+03	3.148e+03

Memory in units of gigabytes; time in seconds.

Hardware counter statistics - 1179 tasks

Counter Name	Sum over all tasks	Average (per task)	Task CV (%)	Task Minimum	Task Maximum
PAPI_FP_OPS	3.358604e+14	2.848689e+11	6.48	3.627467e+10	3.145713e+11

MPI time statistics - 1179 tasks

Call	Sum over all tasks	Average (per task)	Task CV (%)	Task Minimum	Task Maximum	% of MPI	% of wall
MPI_Bcast	2.554e+06	2.166e+03	3.70	8.781e+02	2.640e+03	99.809	68.807
MPI_Allreduce	4.808e+03	4.078e+00	752.69	2.588e+00	1.057e+03	0.188	0.130
MPI_Barrier	7.512e+01	6.371e-02	46.29	2.940e-04	1.161e-01	0.003	0.002
MPI_Comm_rank	9.306e-02	7.893e-05	71.32	4.768e-06	2.038e-04	0.000	0.000
MPI_Comm_size	2.849e-04	2.417e-07	180.52	0.000e+00	1.192e-06	0.000	0.000
MPI_Init	0.000e+00	0.000e+00	0.00	0.000e+00	0.000e+00	0.000	0.000
MPI_Finalize	0.000e+00	0.000e+00	0.00	0.000e+00	0.000e+00	0.000	0.000

Task distribution of MPI function calls for JobID 2459214.sdb

Call	Sum	Average	Std. Dev.	CV (%)	Minimum	Maximum	% of MPI	% of wall
MPI_Bcast	2.554e+06	2.166e+03	8.020e+01	3.70	8.781e+02	2.640e+03	99.809	68.807
MPI_Allreduce	4.808e+03	4.078e+00	3.069e+01	752.69	2.588e+00	1.057e+03	0.188	0.130
MPI_Barrier	7.512e+01	6.371e-02	2.949e-02	46.29	2.940e-04	1.161e-01	0.003	0.002
MPI_Comm_rank	9.306e-02	7.893e-05	5.629e-05	71.32	4.768e-06	2.038e-04	0.000	0.000
MPI_Comm_size	2.849e-04	2.417e-07	4.362e-07	180.52	0.000e+00	1.192e-06	0.000	0.000
MPI_Init	0.000e+00	0.000e+00	0.000e+00	0.00	0.000e+00	0.000e+00	0.000	0.000
MPI_Finalize	0.000e+00	0.000e+00	0.000e+00	0.00	0.000e+00	0.000e+00	0.000	0.000

Task distribution of time taken by **MPI_Barrier** - as a percentage of maximum

The MPI rank represented by each cell in the table is the sum of the cell's column and row indices.

Table Row Width:

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
0	0	75	79	57	68	50	46	72	17	42	53	60	94	39	90	83	100	35	32	28	25	15	64	86
24	99	86	79	83	90	93	71	68	60	28	46	32	64	42	57	53	75	50	18	39	35	15	25	21
48	100	90	93	79	83	86	53	75	50	46	68	60	32	42	57	28	64	71	24	15	21	35	38	17
72	94	75	60	83	50	67	71	86	90	42	100	79	57	53	21	46	39	25	28	32	64	15	17	35
96	98	93	90	86	79	83	74	20	38	63	34	15	49	60	31	28	53	42	16	67	46	24	71	56
120	93	41	77	59	45	52	85	82	63	55	70	48	33	27	19	89	98	15	37	30	14	74	66	23
144	99	86	93	83	79	90	67	38	63	24	50	42	46	21	35	71	60	74	56	31	28	53	17	15
168	93	81	70	62	55	51	73	58	66	37	40	22	88	27	30	15	98	85	48	33	19	14	44	77
192	100	86	94	79	90	83	57	71	68	50	46	36	75	22	39	42	26	32	53	29	15	64	60	18
216	58	76	91	94	88	84	51	80	69	48	40	55	62	33	73	44	99	26	65	36	19	22	30	15
240	67	56	60	49	86	41	63	45	38	31	23	52	82	14	27	19	93	34	70	74	14	14	15	78
264	93	60	41	52	67	49	56	45	38	63	27	74	90	34	86	83	99	31	24	79	14	20	71	16
288	100	90	86	94	79	83	53	57	46	68	64	60	72	50	15	39	35	32	28	75	42	21	25	18

List of aprun commands executed in this job


Number of *aprun* commands: 25

APID	Command	Nodes	Tasks	Threads per Task	Tasks per Node	Max Task Mem (MB)	Run Time (secs)	Start	Complete	Scratch I/O Activity*	Network Topology
1842838	gribmeanp.x	3	64	1	24	235	5	Nov-17-13 22:46:54	Nov-17-13 22:46:59	N/A	Display
	Command line:	/opt/cray/alps/default/bin/aprun -n 64 -N 24 /project/projectdirs/incitell/whitaker/edison/newstuff/bin/gribmeanp.x /scratch1/scratchdirs/whitaker/gfsenkf_t126_1999iau_biascor9//2000022618/ 240 121 2000022618 p fg 06									
	Node list:	1982,3625,6026									
1842840	gribmeanp.x	3	64	1	24	235	5	Nov-17-13 22:46:54	Nov-17-13 22:46:59	N/A	Display
	Command line:	/opt/cray/alps/default/bin/aprun -n 64 -N 24 /project/projectdirs/incitell/whitaker/edison/newstuff/bin/gribmeanp.x /scratch1/scratchdirs/whitaker/gfsenkf_t126_1999iau_biascor9//2000022612/ 240 121 2000022612 p anl									
	Node list:	1627,3469,3991									
1842837	gribmeanp.x	3	64	1	24	389	6	Nov-17-13 22:46:54	Nov-17-13 22:47:00	N/A	Display
	Command line:	/opt/cray/alps/default/bin/aprun -n 64 -N 24 /project/projectdirs/incitell/whitaker/edison/newstuff/bin/gribmeanp.x /scratch1/scratchdirs/whitaker/gfsenkf_t126_1999iau_biascor9//2000022618/ 384 190 2000022618 sflx fg 00									
	Node list:	5459,5467,5483									
1842836	gribmeanp.x	3	64	1	24	389	6	Nov-17-13 22:46:54	Nov-17-13 22:47:00	N/A	Display
	Command line:	/opt/cray/alps/default/bin/aprun -n 64 -N 24 /project/projectdirs/incitell/whitaker/edison/newstuff/bin/gribmeanp.x /scratch1/scratchdirs/whitaker/gfsenkf_t126_1999iau_biascor9//2000022618/ 384 190 2000022618 sflx fg 03									
	Node list:	5422,5433,5448									


IO Summary from Darshan

Start	End	Wallclock (secs)	MB Read	MB Written	Estimated I/O Rate (MB/sec)	Estimated Percent Time Spent in I/O
04-17 19:35:04	04-17 20:13:25	2,301	0.0	215.2	25.48	0.37%

Number of Reads Per Size Range



Number of Writes Per Size Range



The libraries used by this executable

The executable was compiled by user *whitaker* on *2013-08-28*.


The *linkline* was:

```
gribmeanp.x /usr/lib/./lib64/crt1.o /usr/lib/./lib64/crti.o /opt/gcc/4.8.0/snos/lib/gcc/x86_64-suse-  
linux/4.8.0/crtbeginT.o /opt/intel/composer_xe_2013.4.183/compiler/lib/intel64/for_main.o  
/scratch1/scratchdirs/whitaker/ifortkTy55V.o  
/project/projectdirs/incite11/whitaker/edison/newstuff/nwprod/lib/libw3lib-2.0_4.a  
/project/projectdirs/incite11/whitaker/edison/newstuff/nwprod/lib/libbacio_4.a /opt/cray/mpt/6.0.0/gni/mpich2-  
intel/130/lib/libfmpich.a /usr/common/usg/darshan/2.2.7/lib/libdarshan-mpi-io.a /usr/lib/./lib64/libz.a  
/opt/cray/atp/1.6.3/lib/libAtpSigHCommData.a /opt/cray/atp/1.6.3/lib/libAtpSigHandler.a  
/opt/cray/mpt/6.0.0/gni/mpich2-intel/130/lib/libmpich_intel.a /opt/cray/mpt/6.0.0/gni/mpich2-intel/130/lib/libmpl.a  
/opt/cray/xpmem/0.1-2.0500.41356.1.11.ari/lib64/libxpmem.a /opt/cray/ugni/5.0-  
1.0500.0.3.306.ari/lib64/libugni.a /opt/cray/pmi/4.0.1-1.0000.9421.73.3.ari/lib64/libpmi.a /opt/cray/alps/5.0.3-  
2.0500.8213.1.1.ari/lib64/libalpslli.a /opt/cray/alps/5.0.3-2.0500.8213.1.1.ari/lib64/libalpsutil.a  
/opt/cray/udreg/2.3.2-1.0500.6756.2.10.ari/lib64/libudreg.a /usr/lib/./lib64/libpthread.a  
/opt/intel/composer_xe_2013.4.183/compiler/lib/intel64/libimf.a /usr/common/usg/darshan/2.2.7/lib/libdarshan-  
posix.a /usr/lib/./lib64/libpthread.a /opt/intel/composer_xe_2013.4.183/compiler/lib/intel64/libifcore.a  
/opt/intel/composer_xe_2013.4.183/compiler/lib/intel64/libimf.a  
/opt/intel/composer_xe_2013.4.183/compiler/lib/intel64/libirc.a /usr/lib/./lib64/libpthread.a /usr/lib/./lib64/libdl.a  
/usr/lib/./lib64/libc.a /opt/gcc/4.8.0/snos/lib/gcc/x86_64-suse-linux/4.8.0/libgcc.a  
/opt/gcc/4.8.0/snos/lib/gcc/x86_64-suse-linux/4.8.0/libgcc_eh.a /usr/lib/./lib64/libc.a  
/opt/gcc/4.8.0/snos/lib/gcc/x86_64-suse-linux/4.8.0/crtend.o /usr/lib/./lib64/crtn.o
```

ALTD is enabled on all major computing platforms at NERSC



Library usage on Hopper
(June 21, 2012 - Jan 17, 2013)



Applications of ALTD



- Understanding current library usage and plan for future software need
- Providing usage statistics to developers and vendors
- Restoring the program environment where user applications were built
- Assisting with debugging system issues

An ALTD tool to restore the build environment for an application:

```
aryal@edison12:~> linkinfo.sh /global/homes/a/aryal/bin/gvasp5.3.2
User      : zz217
Linked on : 2013-01-03
Executable Name: vasp
Libraries Used :
//usr/lib64/libhugetlbfs.a
../vasp.5.lib/libdmy.a
/opt/cray/atp/1.6.0/lib//libAtpSigHCommData.a
/opt/cray/atp/1.6.0/lib//libAtpSigHandler.a
/opt/cray/libsci/12.0.00/cray/81/sandybridge/lib/libsci_cray_mp.a
/opt/fftw/3.3.0.1/x86_64/lib/libfftw3.a
/opt/cray/mpt/5.6.0/gni/mpich2-cray/74/lib/libmpich_cray.a
/opt/cray/mpt/5.6.0/gni/mpich2-cray/74/lib/libmpl.a
/opt/cray/xpmmem/0.1-2.0500.36799.3.6.ari/lib64/libxpmmem.a
/opt/cray/pmi/4.0.0-1.0000.9282.69.4.ari/lib64/libpmi.a
/opt/cray/ugni/4.0-1.0500.5836.7.58.ari/lib64/libugni.a
/opt/cray/udreg/2.3.2-1.0500.5931.3.1.ari/lib64/libudreg.a
/opt/cray/alps/5.0.1-2.0500.7663.1.1.ari/lib64/libalpslli.a
/opt/cray/alps/5.0.1-2.0500.7663.1.1.ari/lib64/libalpsutil.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libpgas-dmapp.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libu.a
/opt/cray/dmapp/4.0.1-1.0500.5932.6.5.ari/lib64/libdmapp.a
/opt/cray/pmi/4.0.0-1.0000.9282.69.4.ari/lib64/libpmi.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libfi.a
/opt/gcc/4.4.4/snos/lib64/libstdc++.a
/opt/gcc/4.4.4/snos/lib/gcc/x86_64-suse-linux/4.4.4/libgcc_eh.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libf.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libcraymath.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libcraymp.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libu.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libcsup.a
//usr/lib64/librt.a
/opt/cray/cce/8.1.2/craylibs/x86-64/libtcmalloc_minimal.a
//usr/lib64/libpthread.a
//usr/lib64/libc.a
/opt/gcc/4.4.4/snos/lib/gcc/x86_64-suse-linux/4.4.4/libgcc_eh.a
//usr/lib64/libm.a
/opt/gcc/4.4.4/snos/lib/gcc/x86_64-suse-linux/4.4.4/libgcc.a
```