PDSF User Meeting

- PDSF performance (UGE+SLURM)
- announcements
- AOB
- SLURM+CHOS evaluation

aggregated load on PDSF interactive nodes

https://portal-auth.nersc.gov/pdsf-mon/

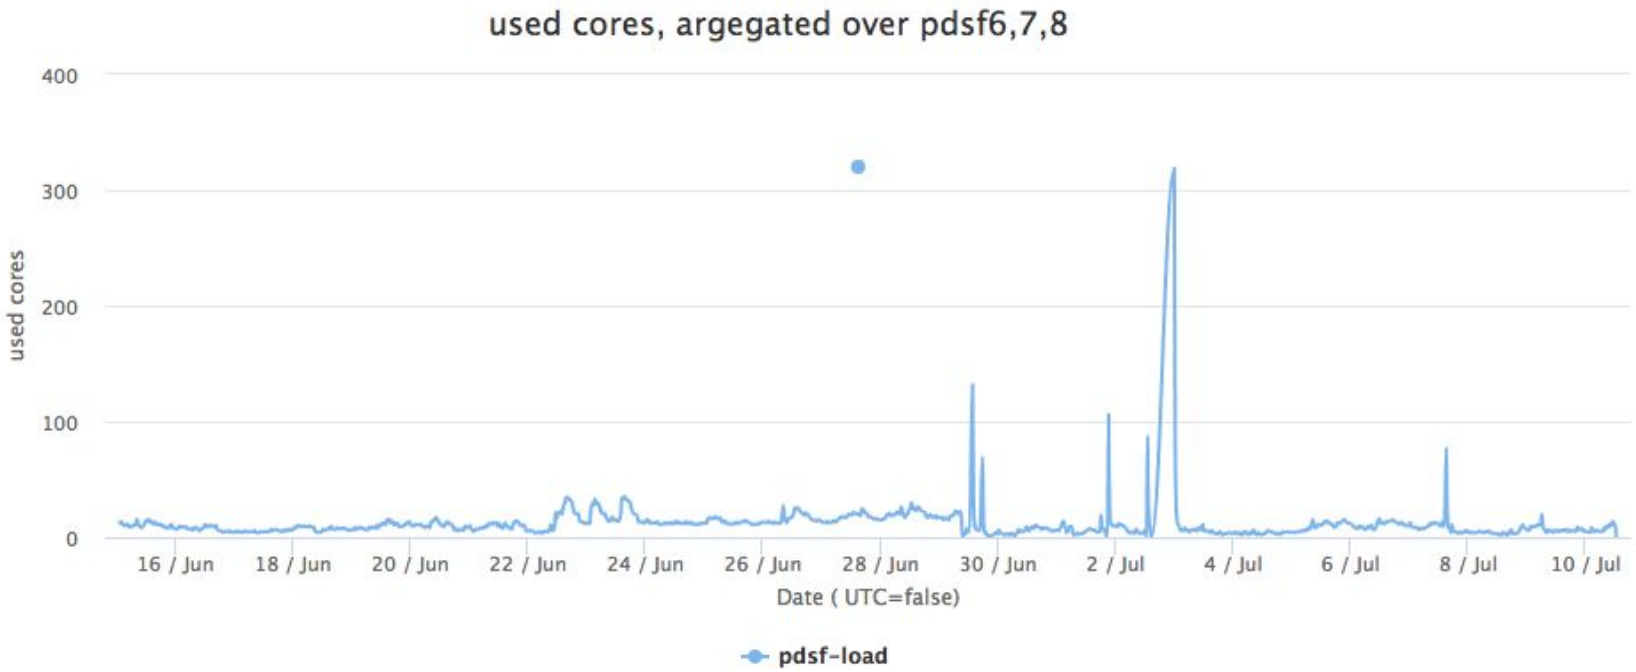

CPU aggregated over month

UGE fully loaded

UGE+CHOS

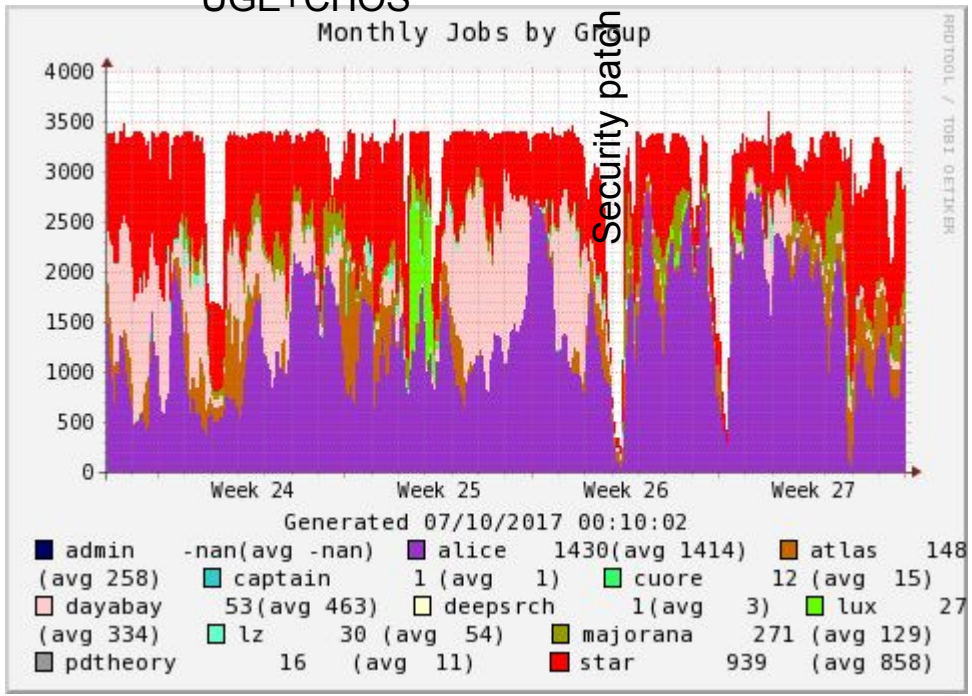

SLURM not used (all load generated by Jan)

http://portal.nersc.gov/project/mpccc/balewski/pdsf3Load/latest/

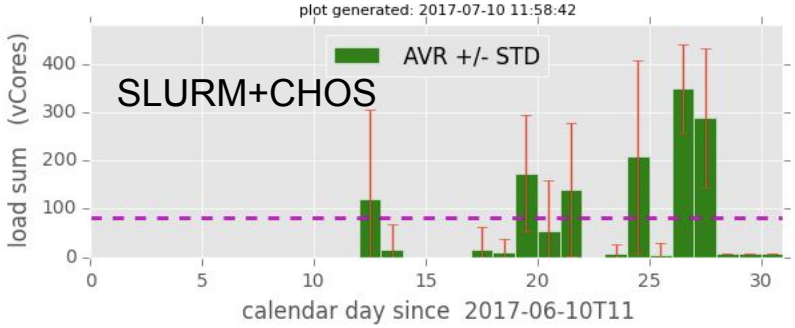

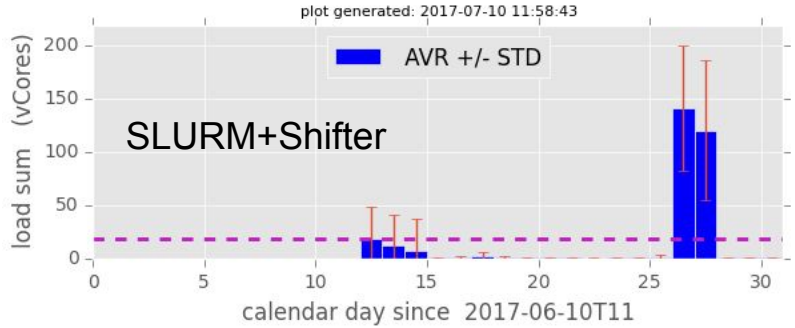

queue(s) load

UGE - looks fine

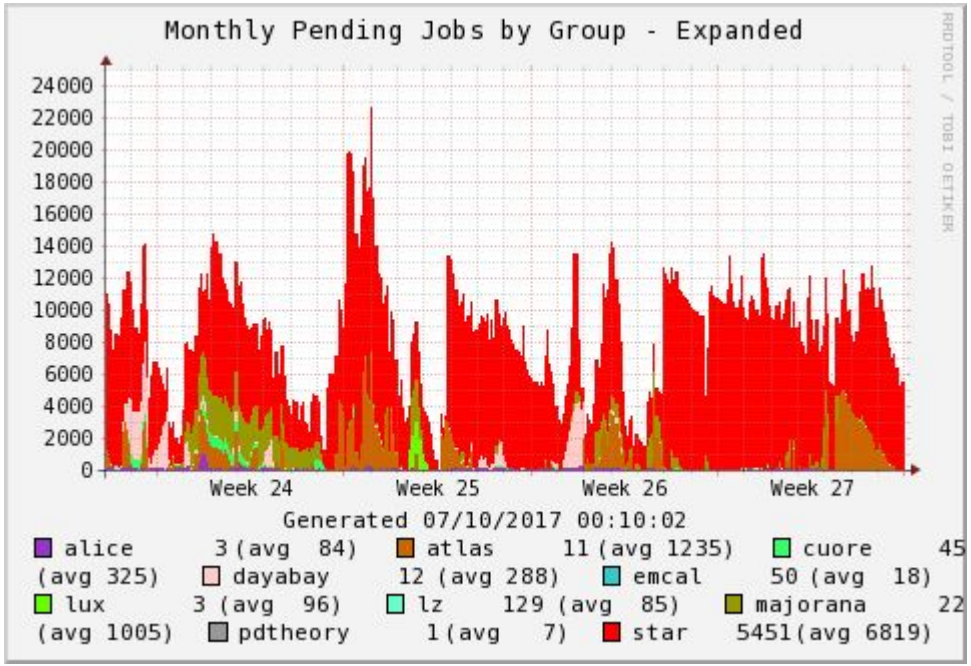

SLURM - no queue limits
(no monitor vs. time)

Partition	Nodes per Job	Physical Cores per Job	Max Walltime per Job
shared-chos	1	1-32	2 days
shared	1	1-32	2 days

Partition	Total vCores [nt1]	Total Nodes	Default Mem per vCore (MB)	Max Mem per vCore (MB)
shared-chos	420	7	1000	2000
shared	176	4	1000	2000

Aggregated utilization of UGE

user \ timeRange	sum last 30_days
0-all-users : owner	251.1 (cpu*year) , 1342121 jobs wallT fract=1.000, CPU/wallT=0.74
alice : project	113.8 (cpu*year) , 99351 jobs wallT fract=0.453, CPU/wallT=0.70
atlas : project	16.6 (cpu*year) , 386661 jobs wallT fract=0.066, CPU/wallT=0.93
dayabay : project	35.2 (cpu*year) , 105573 jobs wallT fract=0.140, CPU/wallT=0.36
dybprod : owner	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU/wallT=0.00
hack : owner	3.3 (cpu*year) , 5687 jobs wallT fract=0.013, CPU/wallT=0.97
lux : project	3.5 (cpu*year) , 4195 jobs wallT fract=0.014, CPU/wallT=0.95
lz : project	2.3 (cpu*year) , 181794 jobs wallT fract=0.009, CPU/wallT=0.92
majorana : project	9.4 (cpu*year) , 141214 jobs wallT fract=0.037, CPU/wallT=0.74
star : project	69.8 (cpu*year) , 415768 jobs wallT fract=0.278, CPU/wallT=0.92
staremb : owner	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU/wallT=0.00

June average
Served 3.1k vCores
Capacity 3.4k vCores
CPU/WallT 0.74

May average
Served 3.1k vCores
Capacity 3.4k vCores
CPU/WallT 0.83

user \ timeRange	sum last 30_days
0-all-users : owner	259.5 (cpu*year) , 1461858 jobs wallT fract=1.000, CPU eff=0.83
alice : project	139.5 (cpu*year) , 135521 jobs wallT fract=0.538, CPU eff=0.81
atlas : project	9.6 (cpu*year) , 266516 jobs wallT fract=0.037, CPU eff=0.99
dayabay : project	33.0 (cpu*year) , 83481 jobs wallT fract=0.127, CPU eff=0.81
dybprod : owner	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00
hack : owner	21.8 (cpu*year) , 17499 jobs wallT fract=0.084, CPU eff=0.99
lux : project	1.7 (cpu*year) , 27385 jobs wallT fract=0.007, CPU eff=0.72
lz : project	3.3 (cpu*year) , 336043 jobs wallT fract=0.013, CPU eff=0.95
majorana : project	4.3 (cpu*year) , 161276 jobs wallT fract=0.017, CPU eff=0.83
star : project	67.7 (cpu*year) , 448828 jobs wallT fract=0.261, CPU eff=0.83
staremb : owner	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00
zmarshal : owner	0.0 (cpu*year) , 29 jobs wallT fract=0.000, CPU eff=0.16

Mendel utilization - snapshot

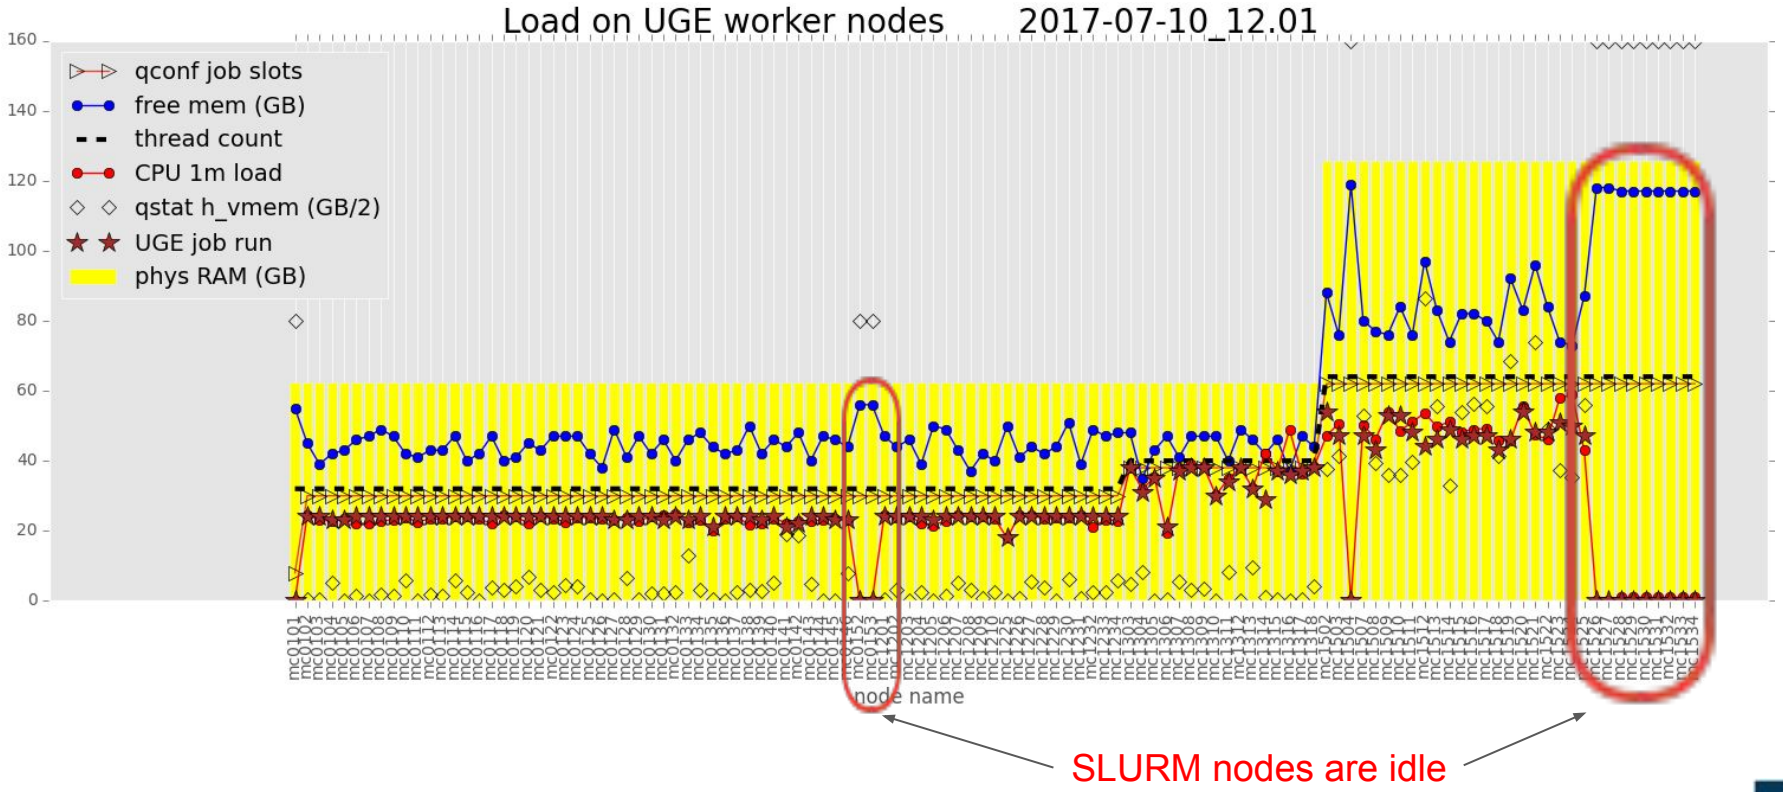

Mendel RAM/task usage - snapshot

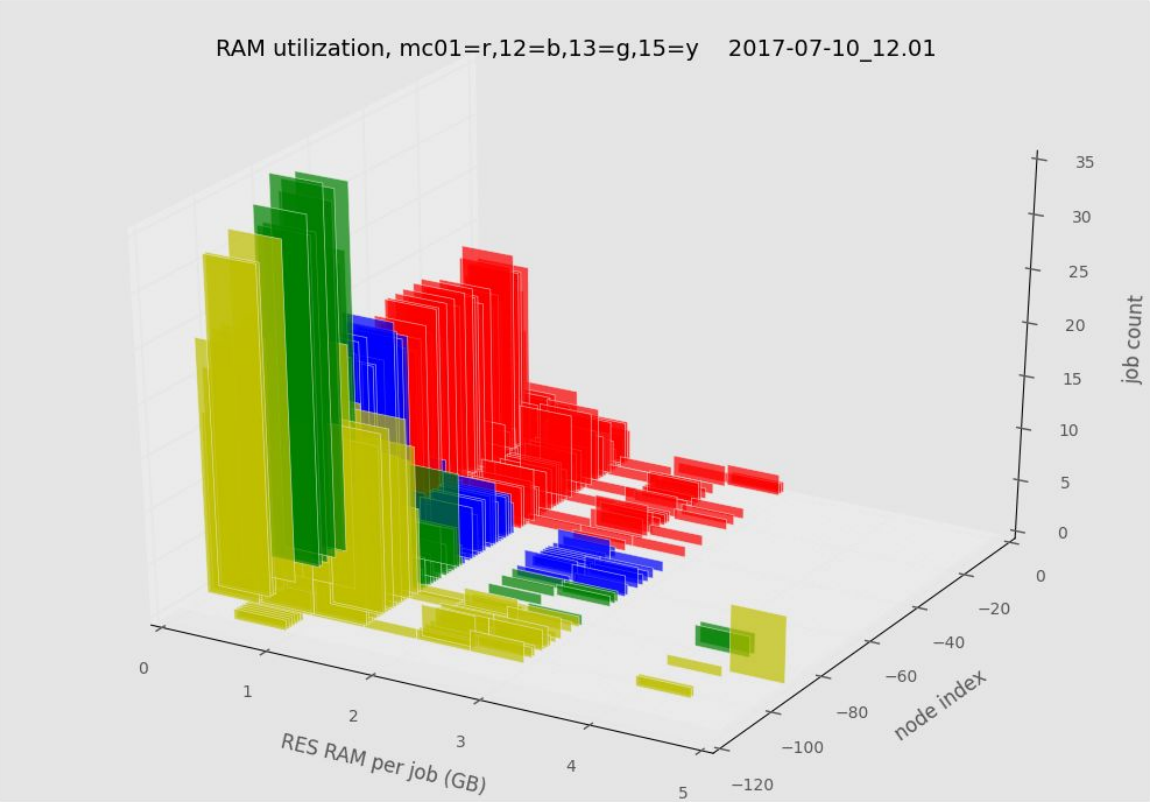

r	qw	Project
1166	1	alice
482	1711	atlas
1	0	cuore
100	10	dayabay
0	3	lux
30	0	lz
133	0	majorana
7	0	pdtheory
1048	7023	star
2967	8748	Totals

XrootD availability - snapshot (1/2)

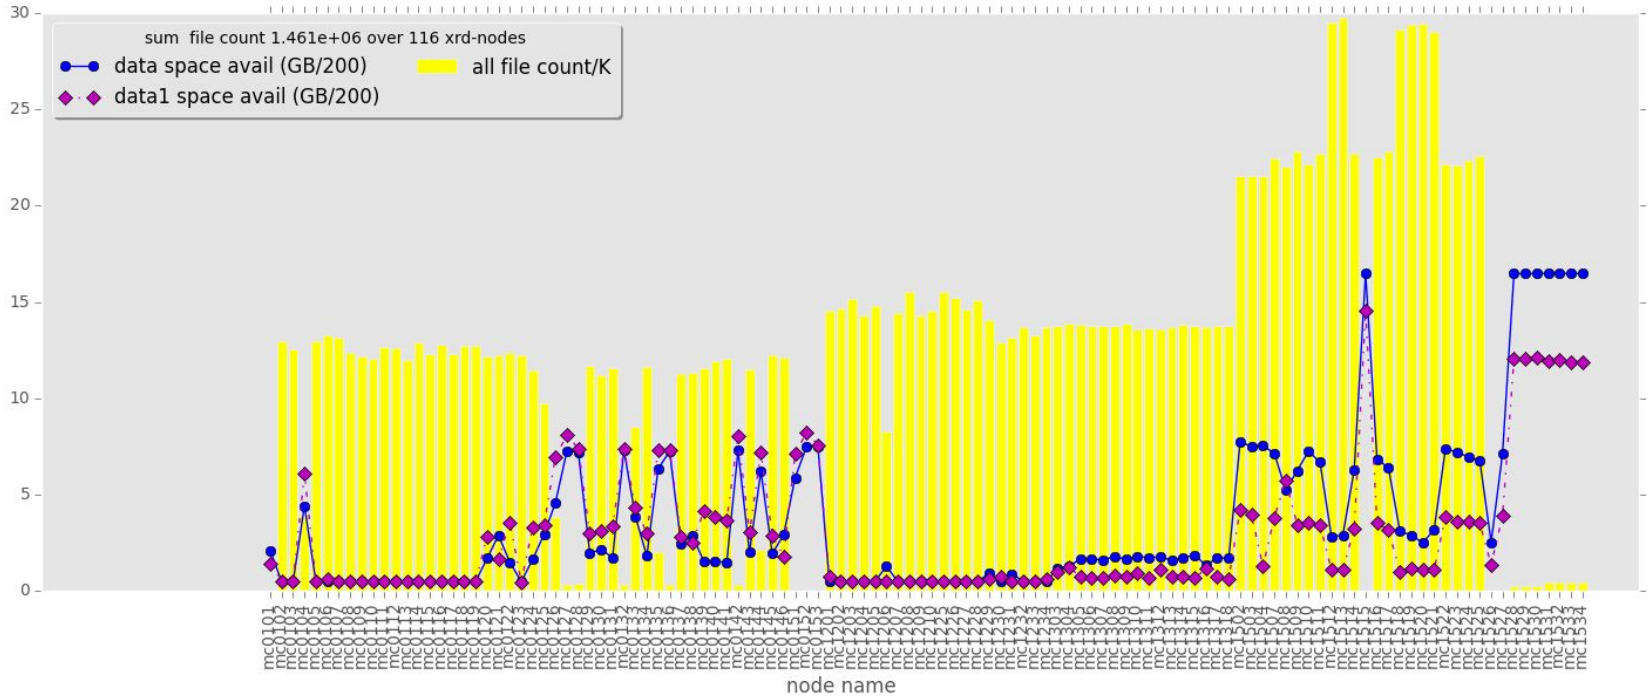

XrootD availability - snapshot (2/2)

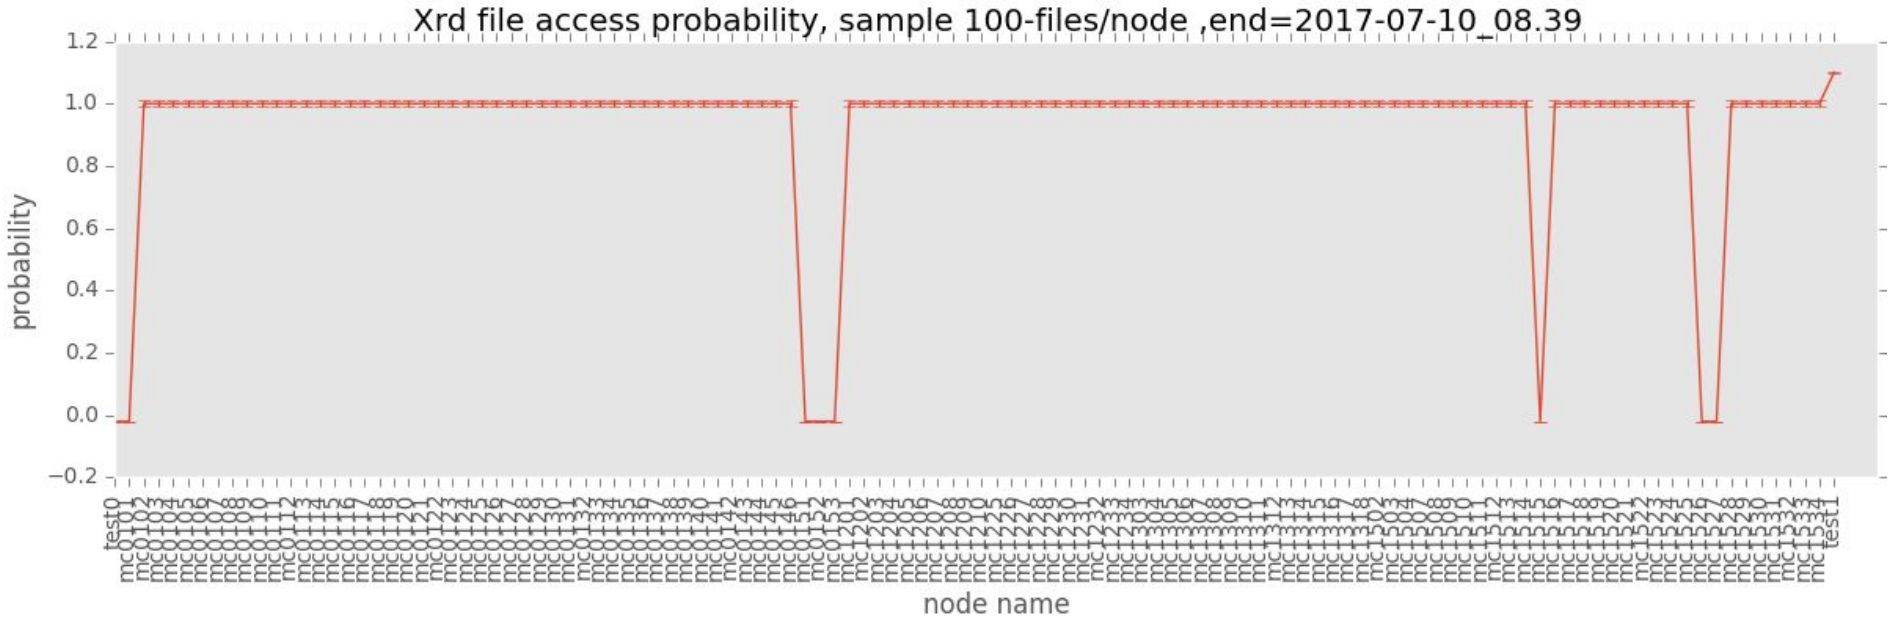

/project(a) utilization - snapshot

http://portal.nersc.gov/project/star/jthaeder/diskUsage/overview/indexExt.html
https://my.nersc.gov/data-mgt.php

```
pdsf6 $  prjquota dayabay
          ---------- Space (GB) ---------     ------------- Inode -------------
Project         Usage     Quota    InDoubt        Usage      Quota     InDoubt
----------  ---------  ---------  ---------    ----------  ----------  ----------
dayabay        755275    870400          2     130338118   150000000         280

pdsf6 $  prjaquota dayabay
          ---------- Space (GB) ---------     ------------- Inode -------------
Project         Usage     Quota    InDoubt        Usage      Quota     InDoubt
----------  ---------  ---------  ---------    ----------  ----------  ----------
dayabay        708577    716800          0       2292229    10000000           0

pdsf6 $  prjquota majorana
          ---------- Space (GB) ---------     ------------- Inode -------------
Project         Usage     Quota    InDoubt        Usage      Quota     InDoubt
----------  ---------  ---------  ---------    ----------  ----------  ----------
majorana        37225     40960          0       2294117     4000000           0
pdsf6 $  prjaquota majorana
          ---------- Space (GB) ---------     ------------- Inode -------------
Project         Usage     Quota    InDoubt        Usage      Quota     InDoubt
----------  ---------  ---------  ---------    ----------  ----------  ----------
majorana        56540     61440          1       4096124    10000000         280
```

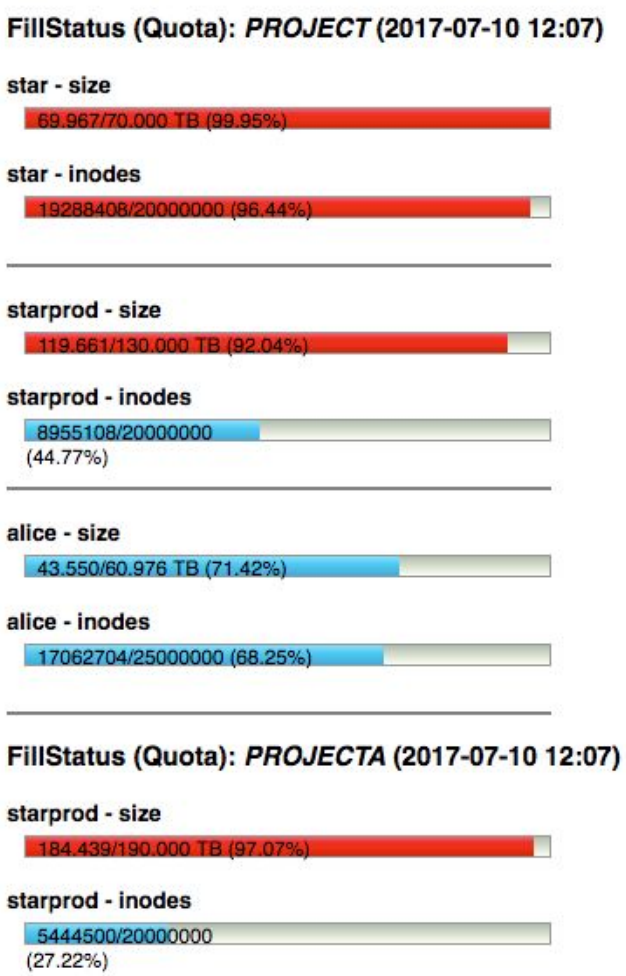

Announcements

Bi-weekly office hours 12:30 -2:30pm
Thursday, July 20, August 3, 59-4016-CR

PDSF user meeting

- Tuesday, August 8, 11am - 12pm, 59-3034-CR

SC Meeting: July 24
Outages :
Cori maintenance: this Wednesday
Edison major upgrade: July 17-31
Outage for Quarterly Maintenance: October 10, 2017

Stress-test of SLURM at PDSF

Please ask your users to test SLURM+CHOS at scale.

PIs plan discussed At SC Meeting June 9, 2017

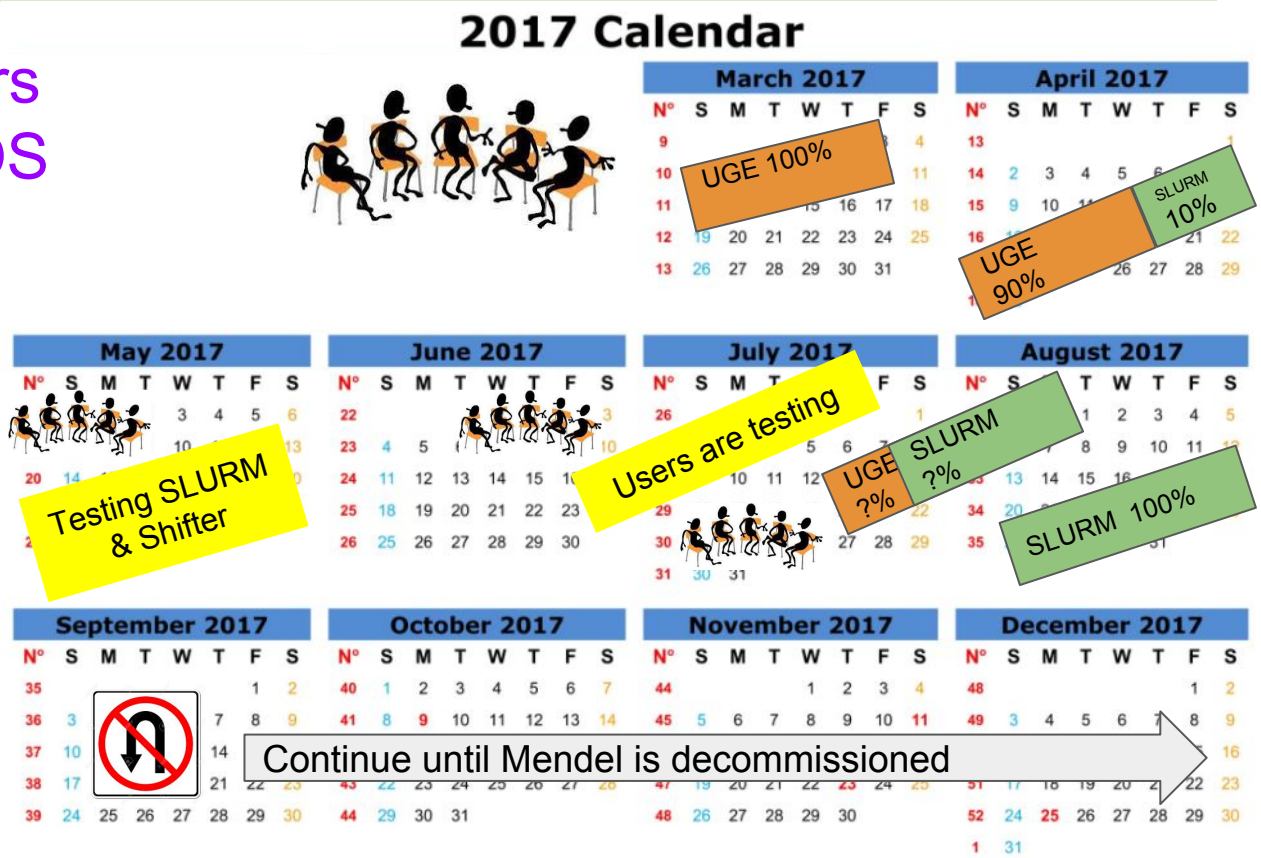

SLURM+CHOS: 60 root4star/node works

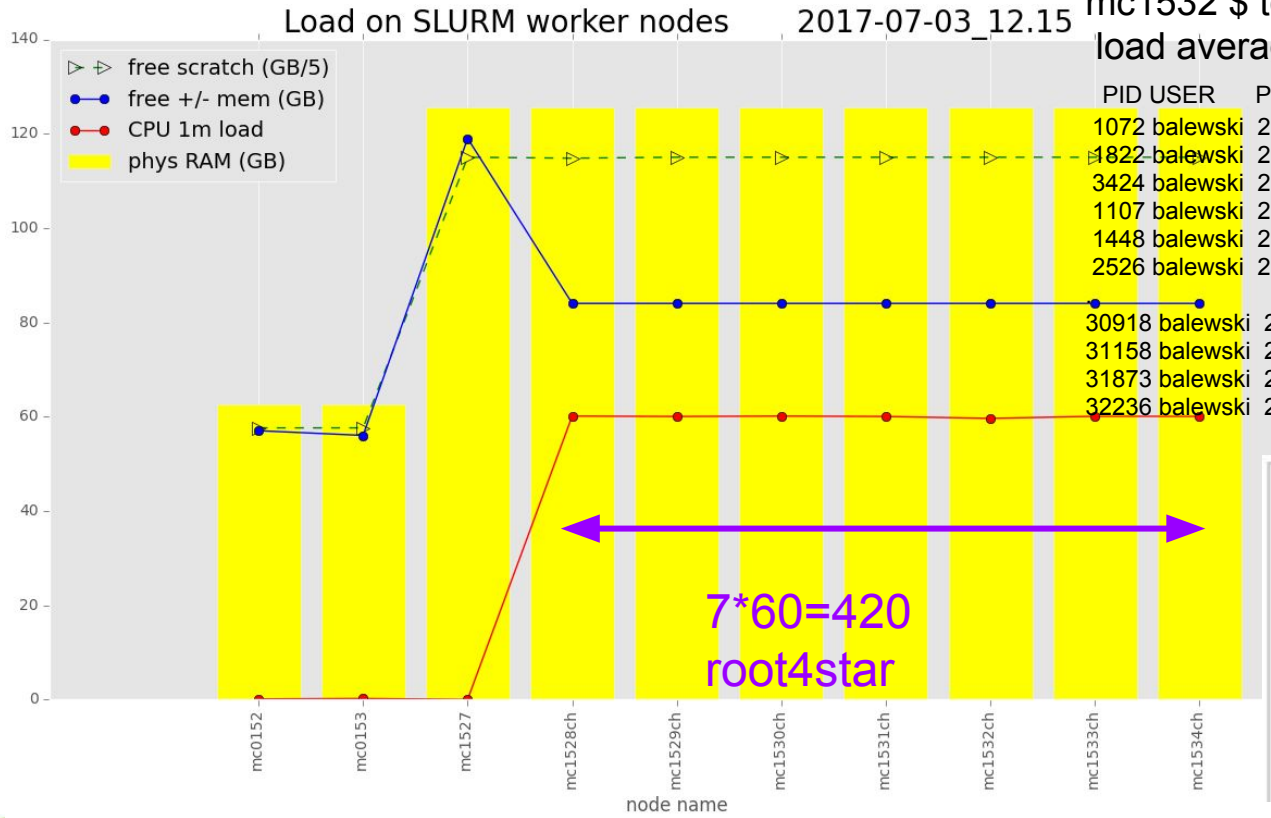

mc1532 $ top ibn1
load average: 60.07, 59.52, 49.50

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMA
1072	balewski	20	0	937m	643m	71m	R	100.0	0.5	23:07.02	root4star
1822	balewski	20	0	943m	649m	71m	R	100.0	0.5	22:19.00	root4star
3424	balewski	20	0	972m	677m	71m	R	100.0	0.5	22:55.48	root4star
1107	balewski	20	0	942m	653m	70m	R	99.7	0.5	40:38.25	root4star
1448	balewski	20	0	958m	670m	71m	R	99.7	0.5	23:10.59	root4star
2526	balewski	20	0	937m	642m	71m	R	99.7	0.5	23:02.94	root4star
				<snip>							
30918	balewski	20	0	939m	645m	69m	R	94.5	0.5	30:26.97	root4star
31158	balewski	20	0	960m	669m	71m	R	94.5	0.5	23:00.85	root4star
31873	balewski	20	0	951m	660m	70m	R	94.5	0.5	39:59.76	root4star
32236	balewski	20	0	987m	690m	71m	R	94.5	0.5	23:51.42	root4star

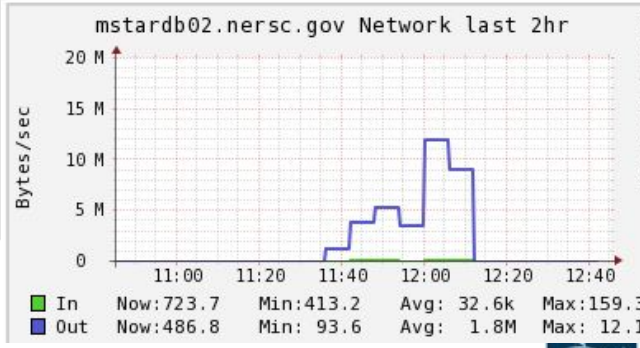

STAR : optimized DB use at cold start

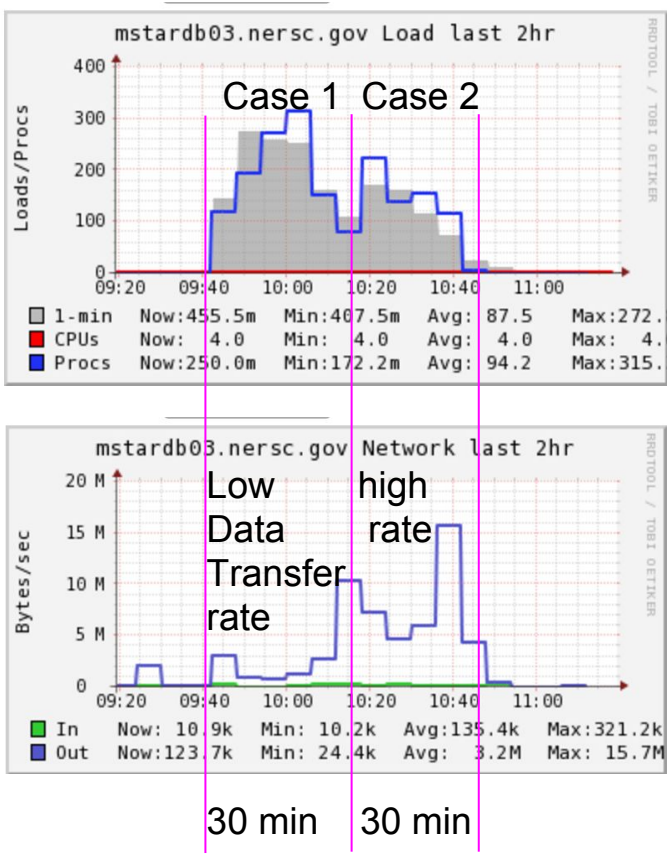

Cold start of 300 1-core tasks
sbatch --array=1-300 myJob.slr

Case 1:

sleep $(($RANDOM % **60**))
root4star (bhla….)

Case 2:

sleep $(($RANDOM % **600**))
root4star (bhla….)

SLURM+CHOS

SLURM+CHOS : 60 atlas jobs per node

mc1530 $ top ibn1 NOTE the CPU time spread despite common start

top - 20:51:04 up 1 day, 5:28, 1 user, load average: 60.04, 51.24, 34.47

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
10886	shapiro	20	0	1549m	930m	5788	R	100.0	0.7	10:29.10	runLBLttHAnalys
11934	shapiro	20	0	1549m	930m	5788	R	100.0	0.7	10:41.61	runLBLttHAnalys
12082	shapiro	20	0	1549m	930m	5788	R	100.0	0.7	10:30.19	runLBLttHAnalys
13307	shapiro	20	0	1549m	931m	7004	R	100.0	0.7	4:07.18	runLBLttHAnalys
13384	shapiro	20	0	1549m	930m	5788	R	100.0	0.7	10:13.53	runLBLttHAnalys
16841	shapiro	20	0	1549m	930m	5788	R	100.0	0.7	10:09.00	runLBLttHAnalys
16886	shapiro	20	0	1549m	930m	5788	R	100.0	0.7	10:10.67	runLBLttHAnalys
17192	shapiro	20	0	1549m	930m	5788	R	100.0	0.7	10:05.95	runLBLttHAnalys
20043	shapiro	20	0	1549m	990m	65m	R	100.0	0.8	1:51.90	runLBLttHAnalys
20049	shapiro	20	0	1549m	990m	65m	R	100.0	0.8	1:59.63	runLBLttHAnalys
22992	shapiro	20	0	1549m	974m	49m	R	100.0	0.8	2:15.51	runLBLttHAnalys
29130	shapiro	20	0	1549m	932m	8268	R	100.0	0.7	3:10.89	runLBLttHAnalys
30492	shapiro	20	0	1549m	930m	5788	R	100.0	0.7	10:35.90	runLBLttHAnalys
30698	shapiro	20	0	1549m	983m	58m	R	100.0	0.8	2:05.52	runLBLttHAnalys
32417	shapiro	20	0	1549m	983m	58m	R	100.0	0.8	2:10.79	runLBLttHAnalys
5696	shapiro	20	0	1549m	1.0g	100m	R	100.0	0.8	1:58.20	runLBLttHAnalys

user multi-thread jobs can coexist

Atlas (6-thread Madgraph)
[kkrizka@pdsf6]$sbatch --array=1-42 oneMG.slr

```
#SBATCH --mem 10G  --ntasks=6
export WORKDIR=/global/project/projectdirs/atlas/kkrizka/janTmp/job${SLURM_JOBID}
mkdir -p $WORKDIR
cd ${WORKDIR}
sleep $(($RANDOM % 300 ))
cp /global/project/projectdirs/atlas/kkrizka/PROC_xia.tgz .
tar -zxf PROC_xia.tgz
CHOS=sl64 chos  $CODE_DIR/launch.sh
```

STAR (1-thread root analysis)
sbatch --array 1-500 starOne.slr

```
#SBATCH --ntasks=1
sleep $(($RANDOM % 300 ))
CHOS=sl64 chos  $CODE_DIR/r4sTask_bfc.csh
```

SLURM+CHOS

```
mc1532 $ top ibn1
top - 11:41:27 up 5 days, 20:18,  1 user,  load average: 37.98, 20.46, 12.54
  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAN
18448 balewski  20   0  720m 483m  75m R 100.0  0.4   1:14.76 root4star
22983 kkrizka   20   0  110m  67m 2008 R 100.0  0.1   0:02.78 madevent
23117 kkrizka   20   0  110m  67m 2008 R 100.0  0.1   0:01.53 madevent
23216 kkrizka   20   0  110m  67m 2004 R 100.0  0.1   0:00.57 madevent
 8759 kkrizka   20   0 9314m 121m 3852 R 99.4  0.1   1:55.00 python
13978 balewski  20   0  361m 159m  69m R 99.4  0.1   1:08.84 root4star
14127 balewski  20   0  926m 628m  81m R 99.4  0.5   2:49.92 root4star
15647 balewski  20   0  360m 158m  69m R 99.4  0.1   1:04.82 root4star
22977 kkrizka   20   0  110m  67m 2008 R 99.4  0.1   0:02.81 madevent
23138 kkrizka   20   0  110m  67m 2004 R 99.4  0.1   0:01.38 madevent
23159 kkrizka   20   0  110m  67m 2004 R 99.4  0.1   0:01.16 madevent
23169 kkrizka   20   0  110m  67m 2004 R 99.4  0.1   0:01.11 madevent
23176 kkrizka   20   0  110m  67m 2004 R 99.4  0.1   0:00.97 madevent
23184 kkrizka   20   0  110m  67m 1984 R 99.4  0.1   0:00.85 madevent
23190 kkrizka   20   0  110m  67m 2004 R 99.4  0.1   0:00.82 madevent
23144 kkrizka   20   0  110m  67m 2004 R 97.6  0.1   0:01.32 madevent
23148 kkrizka   20   0  110m  67m 2008 R 97.6  0.1   0:01.30 madevent
23165 kkrizka   20   0  110m  67m 1984 R 97.6  0.1   0:01.14 madevent
23199 kkrizka   20   0  110m  67m 1984 R 97.6  0.1   0:00.71 madevent
23219 kkrizka   20   0  110m  67m 1984 R 97.6  0.1   0:00.55 madevent
23107 kkrizka   20   0  110m  67m 2008 R 95.8  0.1   0:01.56 madevent
23207 kkrizka   20   0  110m  67m 1984 R 81.6  0.1   0:00.48 madevent
23288 kkrizka   20   0  110m  67m 2004 R 79.8  0.1   0:00.45 madevent
23094 kkrizka   20   0  110m  67m 2004 R 76.3  0.1   0:01.12 madevent
23129 kkrizka   20   0  110m  67m 2004 R 67.4  0.1   0:00.84 madevent
23308 kkrizka   20   0  110m  67m 1984 R 63.9  0.1   0:00.36 madevent
23310 kkrizka   20   0  110m  67m 1984 R 62.1  0.1   0:00.35 madevent
23290 kkrizka   20   0  110m  67m 1920 R 53.2  0.1   0:00.30 madevent
23293 kkrizka   20   0  110m  67m 1940 R 46.1  0.1   0:00.26 madevent
```

SLURM : (bad) use patterns

```
$sbatch --array=1-500 myJob.slr

$cat Task.sh

# goto local scratch - good
cd $SLURM_TMP

# very bad:
cp -rp ~/codeDir .

# better:
cp /project/xyz/aa.tgz .
tar -zxf aa.tgz

# this will help to scatter at cold-start
sleep $(($RANDOM % 30))

# never write to /tmp
```

```
PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
4253 kkrizka 20 0 111m 844 756 R 38.6 0.0 0:02.86 cp
3059 kkrizka 20 0 111m 1056 756 D 0.0 0.0 0:02.69 cp
3492 kkrizka 20 0 111m 1056 756 D 0.0 0.0 0:02.78 cp
3524 kkrizka 20 0 111m 1056 756 D 0.0 0.0 0:02.14 cp
3573 kkrizka 20 0 111m 1056 756 D 0.0 0.0 0:02.76 cp
3716 kkrizka 20 0 111m 1052 756 D 0.0 0.0 0:02.18 cp
3731 kkrizka 20 0 111m 1056 756 D 0.0 0.0 0:02.26 cp
3814 kkrizka 20 0 111m 1048 756 D 0.0 0.0 0:02.84 cp
```

```
17882 shapiro   20   0  6480  636  448 R  0.0  0.0   0:00.53 gzip
17883 shapiro   20   0  117m 1244  944 R  0.0  0.0   0:00.11 tar
17892 shapiro   20   0  117m 1244  944 R  0.0  0.0   0:00.11 tar
17896 shapiro   20   0  117m 1244  944 R  0.0  0.0   0:00.12 tar
17897 shapiro   20   0  6480  640  452 R  0.0  0.0   0:00.30 gzip
17928 shapiro   20   0  116m 5000  740 D  0.0  0.0   0:00.97 cp
19329 cvmfs     20   0 62852 5208 1108 D  0.0  0.0   2:54.68 cvmfs2
20811 shapiro   20   0  117m 1244  944 D  0.0  0.0   0:00.64 tar
26015 shapiro   20   0  116m 5004  740 D  0.0  0.0   0:00.93 cp
28277 shapiro   20   0  117m 1444  944 R  0.0  0.0   0:00.94 tar
30492 shapiro   20   0  922m 391m  29m D  0.0  0.3   0:09.45 runLBLttHAnalys
```

SLURM - observed errors related to stuck IO

- - -
FAILED to unmount old image in this namespace, cannot continue.
FAILED to setup image.

FAILED to mount image /chos/common/nsg/images/66b12b5e69ca115ad8b792.squashfs (squashfs) on /var/udiLoopMount
FAILED to loop mount image: /chos/common/nsg/images/66b12b5e20f15ad8b792.squashfs

- - -
File "/usr/lib64/python2.6/threading.py", line 474, in start _start_new_thread(self.__bootstrap, ())
thread.error: **can't start new thread**

- - -
The compilation fails with the following output message:
make: ../../Source/make_opts:23: **fork: Resource temporarily unavailable**

- - -
cp: writing `/scratch/tmp.20588.0/job20588/89/unweighted_events.lhe.gz': **Cannot allocate memory**

- - -
Site pre-mount hook failed. Exiting.
FAILED to properly setup site modifications

- - - The culprint is ‘D”
PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
11284 xxxxx 20 0 111m 1052 752 **D** 0.0 0.0 0:06.00 cp

SLURM RAM/core ratio is 2 GB/task

```
#!/bin/bash
#SBATCH -t 55:00   --ntasks=1 --account atlas
#SBATCH --mem 3008m      ⇐ # GB/task
#SBATCH -J atlas-chos -p shared-chos
```

pdsf8 $ sbatch -a 1-480 atlasOne.slr

```
205  29490_95          R   None        shapiro  atlas-chos
206  29490_96          R   None        shapiro  atlas-chos
207  29490_97          R   None        shapiro  atlas-chos
208  29490_99          R   None        shapiro  atlas-chos
209  29490_9           R   None        shapiro  atlas-chos
210  29490_[223-480]   PD  Resources   shapiro  atlas-chos
```

Ticket INC0103391:
SLURM on PDSF is not accounting for available swap space

pdsf8 $ scontrol show job 28547_300
TRES=cpu=1,mem=3008,node=1
MinCPUsNode=2 MinMemoryNode=3008M
MinTmpDiskNode=0

Math:
120 GB RAM/60 job-slots=2GB/task
7 nodes * 30 task=210

SLURM - solved issues by CSG

Issue	Priority 1=high	Imparies	Experiment	status	Description
Local SLURM_TMP	1	SLURM	ALICE, ATLAS, STAR	Implemented Jun 7, works	SLURM jobs running on PDSF should have local $SCRACH space reserved for each job and free after a job completes regardless if job was successful or crashed
Scalable CVMFS	3	Shifter	ALICE, ATLAS, LZ?	Impl. Jun 7, works	The limit on number of concurrent CVMFS connections should be 64/node, we have few/node, related to loopback limit

Technical issues resolved:

- OK - Madgraph jobs run at scale on slurm+shifter make system erratic, ~half of jobs can't start, ticket: INC0103239 ⇐ see INC0103537
- Ok INC0103389, 2 slurm 32-core nodes do not accept any jobs if '-n2', mc1531 (chos partition) , mc1528 (shifter partition)
- OK- INC0103453: zoo of slurm+shifter job crashes on PDSF ⇐ see INC0103537
- OK - INC0103455: node mc1526+mc1504 can't handle 60 cp operations from /global/home to $SLURM_TMP ⇐ see INC0103537
- OK - INC0103537 : SLURM+CHOS, exactly the same cp in D issue fro slurm+chos partition <= too much IO, reduce task/node and pread
- OK - INC0103567 : SLURM+CHOS, 90/210 jobs crash jobs trashed due to: slurm_script: fork: retry: Resource temporarily unavailable<= too much IO, reduce task/node and pread
- OK -INC0103754 : SLURM+CHOS, worker nodes in drained state, after heavy use <= OOM cause by malloc(10GB)

SLURM - outstanding high priority issues

Issue	Priority 1=high	Imparies	Experiment	status	Description
ganglia	2	SLURM	all	Known cause	Jobs running under SLURM are not reported on ganglia PDSF monitor page, ticket INC0103269
Ignored swap space	2	SLURM	all	discussed	SLURM is not seeing swap space on nodes, either partition (chos/shifter) INC0103391
OSG2	2	SLURM+grid jobs	ALICE, some STAR	Under eval.	deployment of 2nd OSG talking to SLURM on node mpdsfgrid02, will look like ...01. Jeff & James
AliEn VOBox2	2	SLURM	ALICE	In progress	grid jobs submission, monitoring. Need 2nd system for SLURM
Bad tcsh	3	Shifter	STAR	Under eval.	SLURM + Shifter is not working on PDSF for users who have the default NIM shell set to tcsh, ticket INC0101446
Smooth image selection	3	shifter	all	Under eval.	A user logs into PDSF and their shell is run inside the preselected Shifter image without any action on their part.