LTDA usage

from user and administrative point of view

Marcus Ebert
BABAR

- Hardware/Software
- NFS usage
- Batch system/VM usage
- Performance tests
- Administrative point of view
- User point of view

Hardware overview

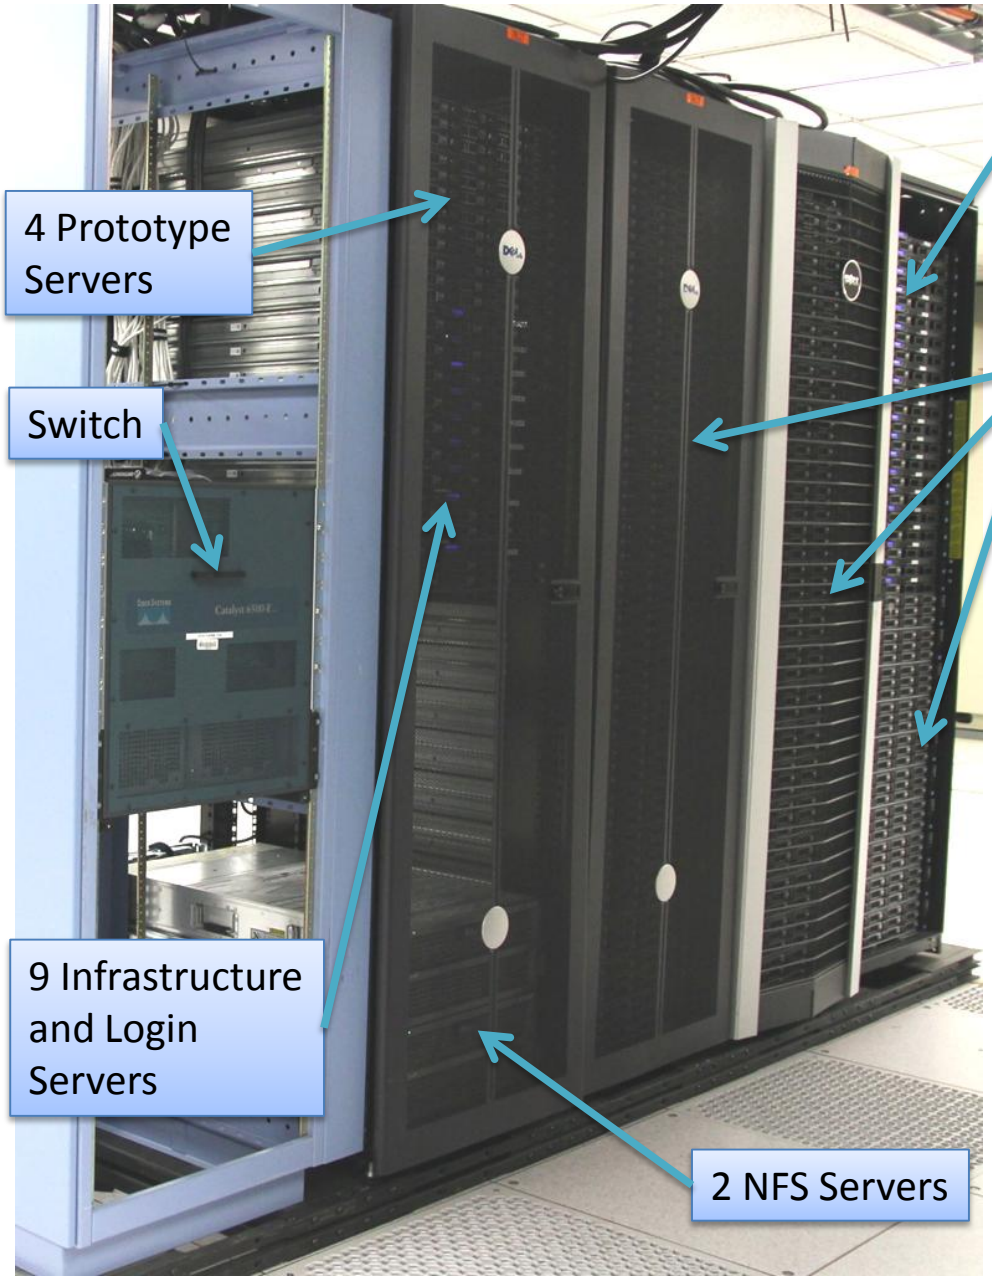

20 Batch Servers (no XROOTD)

50 Batch and XROOTD Servers

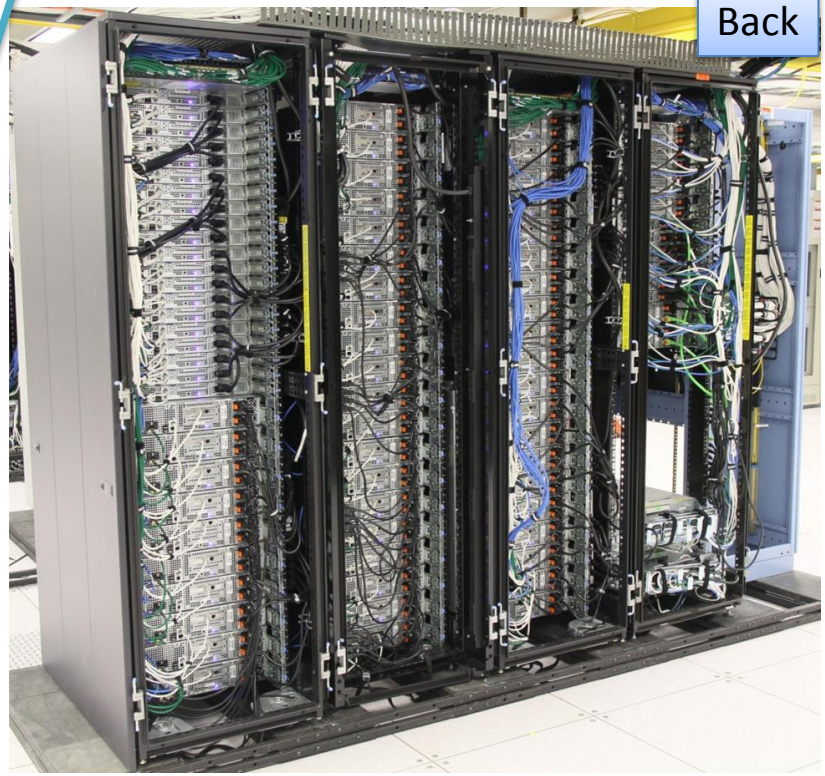

Hardware overview

- infrastructure servers
  (1x cron server, 2x LDAP servers, 2x database servers)

- 3 login machines
  (load balanced pool bbrltda)

- 54x R510 batch and xrootd storage server
  (3GHz dual 6-core CPUs, 48GB RAM, 24TB local xrootd disk space)

- 20x R410 batch server
  (3GHz dual 6-core CPUs, 48GB RAM, 2TB local disk space)

- 1 test server
  (configured like batch machines without xrootd)

- 2x NFS server
  (32TB NFS storage each, ZPOOL with 4 raidz-2 consist of 11 disks and 2 hot spare disks)

NFS usage

wain061

- used for job output (large root files)
- every AWG has own ZFS for easy administration (13 AWG in total)
- 1TB space for each AWG to start with, some already need more

wain062

- User home directories
  - separate ZFS for each user (620 in total)
  - 1GB space for every home directory
- *BABAR* repository
  - *BABAR* analysis releases
  - external packages and libraries, like ROOT
  - LTDA/batch system related software
  - VM base images
- temporary job output space, shared between all AWGs

Batch/Queue system

- Torque (PBS)/Maui used
- server runs on ltda-cron
- job submission from the login machines
- 3 user queues available
  - default queue batch with 5h walltime can use all job slots
  - long queue with 20h walltime and 100 job slots
  - production queue with 24h walltime and 850 job slots
  - with HT on we have 1668 jobs slots in total

Usage of virtualization

- depending on the used physics release a SL4, SL5, or SL6 base image can be used
  - base image is on NFS
  - only 1 base image for every OS release is needed to start jobs on all batch servers
  - very easy administration
- physics analysis environment is stable in time
- base images are read-only for the batch jobs
  - adds safety against unwanted changes in the base image
  - temporary local image contains all writes (copy-on-write image)
  - temporary image is deleted when a job finishes

Job startup

- user submits job from the login machines
- Maui schedules the job for a specific server
- on the server a prologue script is started
  - creates the temporary local image for writes
  - creates a temporary local image to provide scratch space in the VM
  - reserves an available MAC address from a local database
  - determines job parameters like CPUs, RAM, base image
  - creates a network interface for the VM (tunctl/brctl/ifconfig)
  - starts the VM by executing qemu directly with all above determined parameters
- VM uses DHCP to get the IP based on MAC address
- once the VM is started the user job is transferred to it together with the environment using ssh and is executed

Job end

- once the user job is finished, the ssh connection is closed
- after closing the ssh connection an epilogue script is executed
  - VM is destroyed
  - network interface used for the VM is deleted
  - MAC address is freed in the local database
  - temporary images are deleted
  - temporary job script is deleted
- log file of the user job is copied from local server to the final destination

Performance tests

We tested physics analysis jobs, xrootd performance, CPU intensive jobs, and I/O intensive jobs.

analysis jobs:
- about 1500 jobs used in the same configuration for every test
- used walltime reported by the queue system
- used CPU time reported by the analysis job itself

xrootd performance: on the cluster locally available root files have been read using xrdcp by many batch jobs in parallel

CPU intensive jobs: ROOT benchmarks without graphical or disk output

I/O intensive jobs: dd to /scratch (local) and cp from scratch to NFS

General performance results

- delivering root files using xrootd is possible up to 12GB/s in the cluster
- CPU intensive jobs show no dependency on number of parallel jobs
- I/O intensive jobs as expected slow down with increasing number of parallel running jobs on the same server
- double the number of parallel jobs with HT on doesn't affect CPU intensive jobs;
  for I/O intensive jobs same trend as before

Physics analysis performance results

- gives same physics results on the LTDA as on the central SLAC system
- total CPU time used to finish physics jobs on the LTDA (11 jobs/server) and on the central SLAC system is comparable (-2.7% difference)
- HT on and up to 11 jobs in parallel show no difference for physics jobs
- CPU time to finish for physics jobs: $1 \rightarrow 11$: $+20\%$ $11 \rightarrow 22$: $+30\%$
- time to finish all 1500 jobs, $11 \rightarrow 22$: -14%
- using SL5/SL6 image: about -35% CPU time used with SL6

(comparison with central SLAC batch system was done with the prototype only;
other tests done before extension of LTDA; all with RHEL5.x host OS)

KSM

- Kernel Samepage Merging (KSM) for RAM similar to deduplication for disk space
  - same memory pages are merged together into a single one - among different processes
  - most effective for many identical processes
  - that's what we have: all jobs use same VM image

- switching KSM on has no impact on used time of a job or the job output

- with KSM on used RAM reduced by nearly 50%

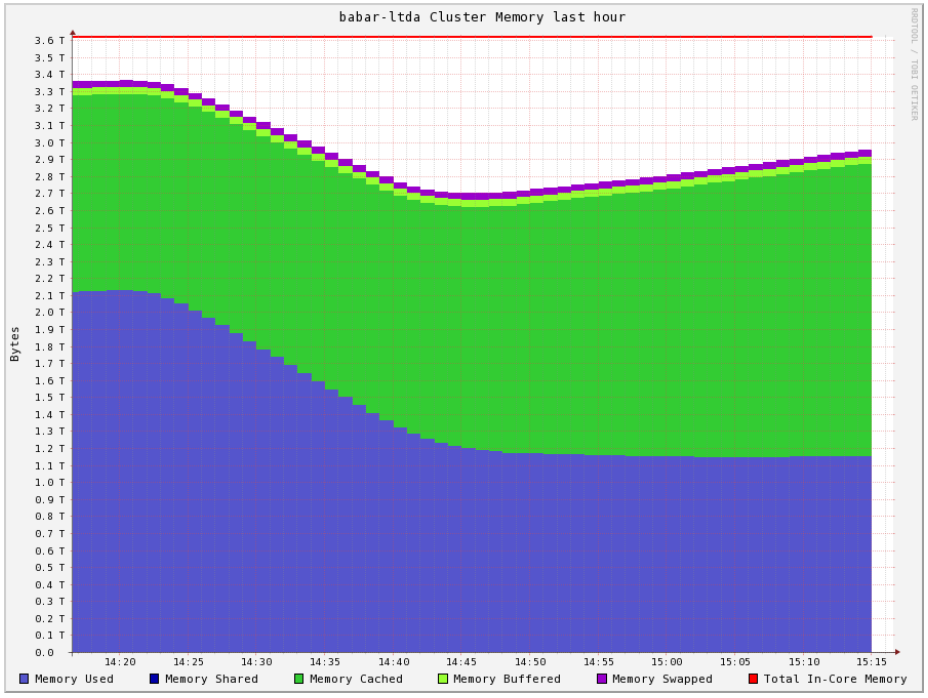

final system:

- use HT on
- use KSM on
- use 22 job slots on machines with xrootd storage
- use 24 job slots on machines without xrootd storage

Administration of running system

- Maui and Torque bring all needed tools for interacting with the batch system/scheduler
- we wrote our own scripts for monitoring
- on-disk and tape backups in place for wain061/wain062
- backup against failure of wain062 in place
- we created a validation system for Red Hat updates

Monitoring

System and usage monitoring are in place.

system monitoring

- client process on the machines which need to be monitored collects information
- same script for all machines
- client data is written to NFS
- server process on ltda-cron analyzes this data and creates a web page together with batch system information and informs about problems

usage monitoring

- daily plots for running jobs, queued jobs, number of users, and production activity
- monthly plots available for queue usage by running jobs
- monthly plots with daily averages available

System monitoring

Server	last updated	uptime	ssh to VM	2. ssh to VM	running kvms	wain061	wain062	/var usage	/tmp usage	/scratch usage	sda	sdb	sdc	sdd	sde	sdf	sdg	sdh	sdi	sdj	sdk	sdl	sdm
bbrltda01	03/18/2013-10:39	6days	0	0	0	OK	OK	13%	2%	1%	OK	OK	OK	OK	n	n	n	n	n	n	n	n	n
bbrltda02	03/18/2013-10:39	6days	0	0	0	OK	OK	15%	2%	1%	OK	OK	OK	OK	n	n	n	n	n	n	n	n	n
bbrltda03	03/18/2013-10:39	6days	0	0	0	OK	OK	14%	2%	1%	OK	OK	OK	OK	n	n	n	n	n	n	n	n	n
ltda-cron	03/18/2013-10:39	146days	0	0	0	OK	OK	60%	2%	4%	OK	OK	n	n	n	n	n	n	n	n	n	n	n
ltda-srv001	03/18/2013-10:39	3days	22	22	22	OK	OK	15%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
ltda-srv002	03/18/2013-10:39	3days	22	22	22	OK	OK	11%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
ltda-srv003	03/18/2013-10:39	3days	22	22	22	OK	OK	15%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
ltda-srv004	03/18/2013-10:39	3days	22	22	22	OK	OK	14%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
ltda-srv005	03/18/2013-10:39	3days	0	0	0	OK	OK	14%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
ltda-srv006	03/18/2013-10:39	3days	0	0	0	OK	OK	10%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
ltda-srv007	03/18/2013-10:39	3days	0	0	0	OK	OK	11%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
ltda-srv008	03/18/2013-10:39	3days	0	0	0	OK	OK	15%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
ltda-srv009	03/18/2013-10:39	3days	0	0	0	OK	OK	11%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	n
ltda-	03/18/2013-	3days	0	0	0	OK	OK	11%	2%	1%	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK

System monitoring

ltda-srv073	03/18/2013-10:49	3days	0	0	0	OK	OK	11%	2%	1%	OK	OK	n	n	n	n	n	n	n	n	n	n	n
ltda-srv074	03/18/2013-10:49	189days	15	15	15	OK	OK	11%	3%	1%	OK	OK	n	n	n	n	n	n	n	n	n	n	n
ltda-srv075	03/18/2013-10:49	5:45	0	0	0	OK	OK	10%	2%	1%	OK	OK	n	n	n	n	n	n	n	n	n	n	n

Queue	Max	Tot	Ena	Str	Que	Run	Hld	Wat	Trn	Ext	T	
testing	0	0	yes	yes	0	0	0	0	0	0	E	
long	100	0	yes	yes	0	0	0	0	0	0	E	
opr	0	12	yes	yes	0	12	0	0	0	0	E	
production	850	202	yes	yes	0	202	0	0	0	0	E	
mpiq	0	0	yes	yes	0	0	0	0	0	0	E	
interactive	0	3	yes	yes	0	3	0	0	0	0	E	
batch	1644	421	yes	yes	0	421	0	0	0	0	E	

interactive jobs on:

---

jobs exceeded walltime today
number of jobs in batch queue: 0
number of jobs in production queue: 0

successful finished jobs today: 1161

System monitoring

```
FairShare Information

Depth: 8 intervals   Interval Length: 12:00:00   Decay Rate: 1.00

FS Policy: [NONE]
System FS Settings:  Target Usage: 0.00    Flags: 0

FSInterval        %     Target       0
FSWeight       ------- -------  1.0000
TotalUsage      100.00 ------- 13401.6

USER
-------------
afilippi          6.06 -------    6.06
arossi           12.62 -------   12.62
tqn              23.28 -------   23.28
fwilson           7.46 -------    7.46
biplabd          13.85 -------   13.85
babaropr          7.72 -------    7.72
oberhof           4.51 -------    4.51
ebert             8.23 -------    8.23
susmita           0.00 -------    0.00
buenger           1.40 -------    1.40
mchrzasz         14.88 -------   14.88

GROUP
-------------
br               84.81 -------   84.81
bf                0.00 -------    0.00
ec               15.19 -------   15.19

CLASS
-------------
batch            71.43 -------   71.43
production       12.62 -------   12.62
testing           0.51 -------    0.51
opr               7.72 -------    7.72
interactive       7.72 -------    7.72
```

```
maui active for    3:14:15:38  stats initialized on Wed Dec 31 16:00:00

Eligible/Idle Jobs:                    0/0         (0.000%)
Active Jobs:                           638
Successful/Completed Jobs:       2173668/2173668   (100.000%)
Avg/Max QTime (Hours):              0.01/52.60
Avg/Max XFactor:                    0.00/9.83

Dedicated/Total ProcHours:     626129.22/13268461.83 (4.719%)

Current Active/Total Procs:          647/1684      (38.420%)

Avg WallClock Accuracy:          15.245%
Avg Job Proc Efficiency:         87.596%
Est/Avg Backlog (Hours):            0.00/0.00
```

System monitoring

user	running jobs	running procs	currently using procs hours	completed jobs	%	proc-hours requested	%	dedicated proc hours	%	fairshare	AvgXF	MaxXF	AvgQT	Efficiency	WCAcc
aperez	0	0	0.00	33	0.00	165.0	0.00	57.8	0.00	-----	0.35	0.00	0.00	39.34	35.01
vsantoro	0	0	0.00	28	0.00	145.0	0.00	44.9	0.00	-----	0.32	0.00	0.03	43.70	31.88
cartaro	0	0	0.00	37	0.00	33725.0	0.03	13.6	0.00	-----	0.02	0.00	0.00	32.26	1.56
martelsl	0	0	0.00	8	0.00	40.0	0.00	0.3	0.00	-----	0.01	0.00	0.00	72.87	0.65
zoso	0	0	0.00	3	0.00	72.0	0.00	0.1	0.00	-----	0.00	0.00	0.00	70.42	0.20
demori	0	0	0.00	1	0.00	5.0	0.00	0.1	0.00	-----	0.02	0.00	0.00	77.45	1.70
griess	0	0	0.00	1	0.00	5.0	0.00	0.1	0.00	-----	0.02	0.00	0.00	74.39	1.70
stracka	0	0	0.00	2	0.00	10.0	0.00	0.1	0.00	-----	0.01	0.00	0.00	68.00	0.75
tomo	0	0	0.00	1702	0.08	34040.0	0.03	9983.1	0.51	-----	0.30	0.00	0.17	99.20	29.33
afilippi	0	0	0.00	3466	0.16	17330.0	0.02	1568.3	0.08	-----	0.10	0.00	0.06	87.07	9.10
beaulieu	0	0	0.00	62263	2.86	351188.0	0.33	44880.6	2.30	-----	0.14	0.00	0.02	88.40	14.23
lueckt	0	0	0.00	153964	7.08	775590.0	0.72	284569.2	14.58	-----	0.37	0.00	0.02	90.22	36.86
dana	0	0	0.00	58110	2.67	342272.0	0.32	42155.7	2.16	-----	0.10	0.00	0.01	91.87	9.47
anulli	0	0	0.00	11	0.00	55.0	0.00	0.2	0.00	-----	0.00	0.00	0.00	79.35	0.42
rsobie	0	0	0.00	1	0.00	5.0	0.00	0.0	0.00	-----	0.02	0.00	0.06	80.54	0.40
narnaud	0	0	0.00	1	0.00	5.0	0.00	0.0	0.00	-----	0.01	0.00	0.02	71.96	0.30
arossi	202	202	3192.61	4089	0.19	97927.0	0.09	34739.5	1.78	-----	0.34	0.00	0.01	98.32	33.79
chcheng	0	0	0.00	28939	1.33	154387.0	0.14	69806.6	3.58	-----	0.49	0.00	0.02	94.47	48.81
rid	0	0	0.00	121	0.01	605.0	0.00	267.2	0.01	-----	0.44	0.00	0.00	97.94	44.17
fransham	0	0	0.00	5	0.00	65.0	0.00	6.8	0.00	-----	0.09	0.00	0.00	21.16	9.26
manoni	0	0	0.00	26	0.00	325.0	0.00	6.3	0.00	-----	0.04	0.00	0.21	83.32	1.99

System monitoring

patrign	0	0	0.00	1	0.00	5.0	0.00	0.0	0.00	-----	0.00	0.00	0.00	82.13	0.30
ebert	19	28	12553.49	102595	4.72	95517494.0	88.73	134820.6	6.91	-----	0.16	0.00	0.03	61.14	15.47
buenger	0	0	0.00	29953	1.38	159710.0	0.15	50034.8	2.56	-----	0.34	0.00	0.03	91.05	33.31
benhaim	0	0	0.00	1	0.00	5.0	0.00	0.0	0.00	-----	0.00	0.00	0.00	100.00	0.30
mchrzasz	0	0	0.00	42430	1.95	236180.0	0.22	107911.4	5.53	-----	0.49	0.00	0.04	95.74	48.01
homer	0	0	0.00	99024	4.56	1144610.0	1.06	136021.2	6.97	-----	0.11	0.00	0.01	93.89	10.96
gapon	0	0	0.00	3	0.00	15.0	0.00	5.3	0.00	-----	0.35	0.00	0.00	36.51	35.00
gcasa	0	0	0.00	49	0.00	9825.0	0.01	2.4	0.00	-----	0.01	0.00	0.00	72.18	0.91

```
diagnosing node table (5120 slots)
Name                  State  Procs     Memory         Disk          Swap      Speed  Opsys   Arch Par   Load Res Classes                        Network                 Features

ltda-srv074         Running   0:24   47771:48251  1766496:1782025  67645:97403   1.00  linux [NONE] DEF   0.05 015 [batch_24:24][production_24:24 [DEFAULT]                [opr][interactive]
ltda-srv001            Busy   0:22   47547:48251  1875301:1877272 118246:139483  1.00  linux [NONE] DEF  21.49 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv002            Busy   0:22   47547:48251  1875527:1877274 122992:139483  1.00  linux [NONE] DEF  21.34 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv003            Busy   0:22   47547:48251  1875639:1877274 125198:139483  1.00  linux [NONE] DEF  21.24 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv004            Busy   0:22   47547:48251  1875650:1877274 125386:139483  1.00  linux [NONE] DEF  20.73 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv005            Idle  22:22   48251:48251  1877079:1877274 137990:139483  1.00  linux [NONE] DEF   1.18 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
WARNING:  node 'ltda-srv005' has been idle for 5:33:51 but load is HIGH.  load:  1.180 (check for runaway processes?)
ltda-srv006            Idle  22:22   48251:48251  1877079:1877274 137975:139483  1.00  linux [NONE] DEF   0.32 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
ltda-srv007            Idle  22:22   48251:48251  1877079:1877274 138000:139483  1.00  linux [NONE] DEF   0.28 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
ltda-srv008            Idle  22:22   48251:48251  1877079:1877274 137971:139483  1.00  linux [NONE] DEF   0.22 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
ltda-srv009            Idle  22:22   48251:48251  1877079:1877274 138125:139483  1.00  linux [NONE] DEF   0.62 001 [batch_22:22][production_22:22 [DEFAULT]              [testing]
WARNING:  node 'ltda-srv009' has been idle for 7:58:59 but load is HIGH.  load:  0.620 (check for runaway processes?)
ltda-srv010            Idle  22:22   48251:48251  1877079:1877274 137837:139483  1.00  linux [NONE] DEF   0.63 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
WARNING:  node 'ltda-srv010' has been idle for 8:41:18 but load is HIGH.  load:  0.630 (check for runaway processes?)
ltda-srv011            Idle  22:22   48251:48251  1877079:1877274 138110:139483  1.00  linux [NONE] DEF   0.75 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
WARNING:  node 'ltda-srv011' has been idle for 8:23:51 but load is HIGH.  load:  0.750 (check for runaway processes?)
ltda-srv012            Idle  22:22   48251:48251  1877079:1877274 137968:139483  1.00  linux [NONE] DEF   0.34 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
ltda-srv013            Idle  22:22   48251:48251  1877079:1877274 137924:139483  1.00  linux [NONE] DEF   0.54 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
WARNING:  node 'ltda-srv013' has been idle for 8:49:22 but load is HIGH.  load:  0.540 (check for runaway processes?)
ltda-srv014            Idle  22:22   48251:48251  1877079:1877274 138024:139483  1.00  linux [NONE] DEF   0.57 001 [batch_22:22][production_22:22 [DEFAULT]              [batch]
WARNING:  node 'ltda-srv014' has been idle for 8:31:35 but load is HIGH.  load:  0.570 (check for runaway processes?)
ltda-srv015            Busy   0:22   47547:48251  1875650:1877274 125069:139483  1.00  linux [NONE] DEF  20.85 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv016            Busy   0:22   47547:48251  1875649:1877274 125070:139483  1.00  linux [NONE] DEF  21.83 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv017            Busy   0:22   47547:48251  1875676:1877274 125129:139483  1.00  linux [NONE] DEF  22.32 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv018            Busy   0:22   47547:48251  1875667:1877274 125247:139483  1.00  linux [NONE] DEF  21.50 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv019            Busy   0:22   47547:48251  1875644:1877274 125121:139483  1.00  linux [NONE] DEF  21.70 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv020            Busy   0:22   47547:48251  1873277:1877274 125205:139483  1.00  linux [NONE] DEF  21.60 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
ltda-srv021         Running   1:22   47579:48251  1875729:1877274 125864:139483  1.00  linux [NONE] DEF  19.99 021 [batch_1:22][production_22:22] [DEFAULT]              [batch]
ltda-srv022            Busy   0:22   47547:48251  1875672:1877274 125014:139483  1.00  linux [NONE] DEF  21.43 022 [batch_0:22][production_22:22] [DEFAULT]              [batch]
```

System monitoring

```
ltda-srv041             Idle  22:22   48251:48251  1877079:1877274 137927:139483  1.00  linux [NONE] DEF   0.28 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
ltda-srv042             Idle  22:22   48251:48251  1877079:1877274 137958:139483  1.00  linux [NONE] DEF   0.94 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv042' has been idle for 8:48:45 but load is HIGH.  load:  0.940 (check for runaway processes?)
ltda-srv043             Idle  22:22   48251:48251  1877079:1877274 137925:139483  1.00  linux [NONE] DEF   0.66 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv043' has been idle for 8:32:00 but load is HIGH.  load:  0.660 (check for runaway processes?)
ltda-srv044             Idle  22:22   48251:48251  1877079:1877274 137967:139483  1.00  linux [NONE] DEF   0.58 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv044' has been idle for 8:15:01 but load is HIGH.  load:  0.580 (check for runaway processes?)
ltda-srv045             Idle  22:22   48251:48251  1877079:1877274 137963:139483  1.00  linux [NONE] DEF   1.12 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv045' has been idle for 8:57:09 but load is HIGH.  load:  1.120 (check for runaway processes?)
ltda-srv046             Idle  22:22   48251:48251  1877079:1877274 137870:139483  1.00  linux [NONE] DEF   1.04 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv046' has been idle for 2:08:39:41 but load is HIGH.  load:  1.040 (check for runaway processes?)
ltda-srv047             Idle  22:22   48251:48251  1877079:1877274 137941:139483  1.00  linux [NONE] DEF   0.55 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv047' has been idle for 2:08:22:32 but load is HIGH.  load:  0.550 (check for runaway processes?)
ltda-srv048             Idle  22:22   48251:48251  1877079:1877274 138065:139483  1.00  linux [NONE] DEF   0.92 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv048' has been idle for 8:05:40 but load is HIGH.  load:  0.920 (check for runaway processes?)
ltda-srv049             Idle  22:22   48251:48251  1877079:1877274 137956:139483  1.00  linux [NONE] DEF   0.67 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv049' has been idle for 8:48:24 but load is HIGH.  load:  0.670 (check for runaway processes?)
ltda-srv050             Idle  22:22   48251:48251  1877079:1877274 137954:139483  1.00  linux [NONE] DEF   1.20 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv050' has been idle for 8:30:50 but load is HIGH.  load:  1.200 (check for runaway processes?)
ltda-srv051             Idle  22:22   48251:48251  1877077:1877272 137911:139483  1.00  linux [NONE] DEF   0.25 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
ltda-srv052             Idle  22:22   48251:48251  1877079:1877274 137995:139483  1.00  linux [NONE] DEF   0.35 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
ltda-srv053             Idle  22:22   48251:48251  1877079:1877274 137887:139483  1.00  linux [NONE] DEF   0.38 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
ltda-srv054             Idle  22:22   48251:48251  1877079:1877274 138037:139483  1.00  linux [NONE] DEF   0.91 000 [batch_22:22][production_22:22 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv054' has been idle for 3:05:53:53 but load is HIGH.  load:  0.910 (check for runaway processes?)
ltda-srv055             Idle  24:24   48251:48251  1781829:1782025  96193:97403   1.00  linux [NONE] DEF   0.05 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv056             Idle  24:24   48251:48251  1781829:1782025  96186:97403   1.00  linux [NONE] DEF   0.00 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv057             Idle  24:24   48251:48251  1781829:1782025  96190:97403   1.00  linux [NONE] DEF   0.06 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv058             Idle  24:24   48251:48251  1781829:1782025  96184:97403   1.00  linux [NONE] DEF   0.02 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv059             Idle  24:24   48251:48251  1781829:1782025  96190:97403   1.00  linux [NONE] DEF   0.09 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv060             Idle  24:24   48251:48251  1781829:1782025  96189:97403   1.00  linux [NONE] DEF   0.00 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv061             Idle  24:24   48251:48251  1781829:1782025  96191:97403   1.00  linux [NONE] DEF   0.00 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv062             Idle  24:24   48251:48251  1781829:1782025  96188:97403   1.00  linux [NONE] DEF   0.32 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv063             Idle  24:24   48251:48251  1781829:1782025  96191:97403   1.00  linux [NONE] DEF   0.26 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv064             Idle  24:24   48251:48251  1781829:1782025  96189:97403   1.00  linux [NONE] DEF   0.10 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv065             Idle  24:24   48251:48251  1781829:1782025  96186:97403   1.00  linux [NONE] DEF   0.10 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv066             Idle  24:24   48251:48251  1781829:1782025  96189:97403   1.00  linux [NONE] DEF   0.00 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv067             Idle  24:24   48251:48251  1781829:1782025  96189:97403   1.00  linux [NONE] DEF   0.00 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv068             Idle  24:24   48251:48251  1781829:1782025  96187:97403   1.00  linux [NONE] DEF   0.05 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv069             Idle  24:24   48251:48251  1781829:1782025  96191:97403   1.00  linux [NONE] DEF   0.71 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
WARNING:  node 'ltda-srv069' has been idle for 1:08:10:27 but load is HIGH.  load:  0.710 (check for runaway processes?)
ltda-srv070             Idle  24:24   48251:48251  1781829:1782025  96187:97403   1.00  linux [NONE] DEF   0.09 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv071             Idle  24:24   48251:48251  1781829:1782025  96195:97403   1.00  linux [NONE] DEF   0.00 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv072             Idle  24:24   48251:48251  1781829:1782025  96191:97403   1.00  linux [NONE] DEF   0.09 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv073             Idle  24:24   48251:48251  1781829:1782025  96189:97403   1.00  linux [NONE] DEF   0.00 000 [batch_24:24][production_24:24 [DEFAULT]                [batch][production]
ltda-srv075             Busy   0:16   23501:24013  185472:185837  69251:73165   1.00  linux [NONE] DEF  14.16 017 [batch_16:16][production_16:16 [DEFAULT]              [testing]
-----                    --- 1025:1684 3573787:3594587 136775202:137199129 8919527:9553307

Total Nodes: 75  (Active: 32  Idle: 43  Down: 0)

---

Diagnosing blocked jobs (policylevel SOFT  partition ALL)

---
```

Usage monitoring

last modified: Mar 20 07:20 (SLAC time) **successful finished jobs:** 1719 **all finished jobs:** 1721 **jobs exceed walltime**: 0

previous day

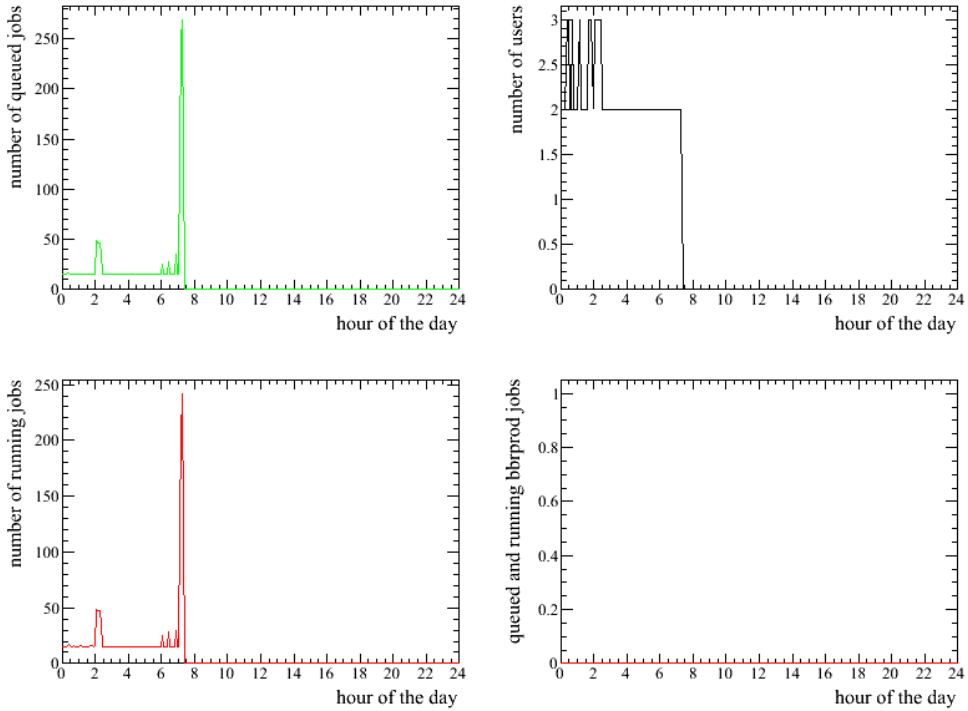

last modified: Feb 5 23:55 (SLAC time) **successful finished jobs:** 7792 **all finished jobs:** 10104 **jobs exceed walltime**: 64

previous day next day

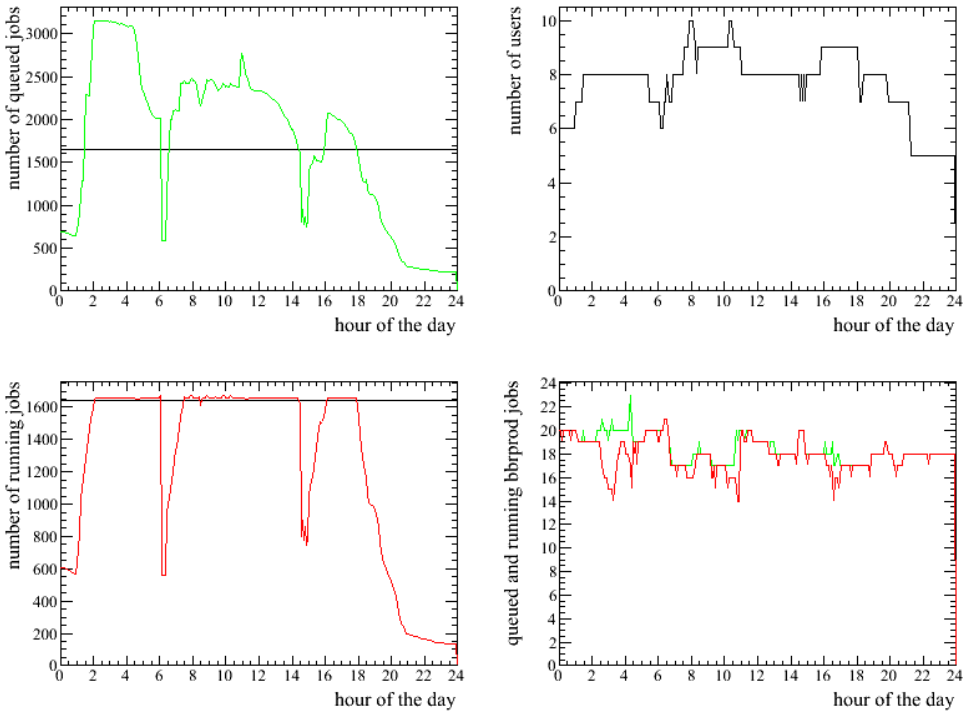

Usage monitoring

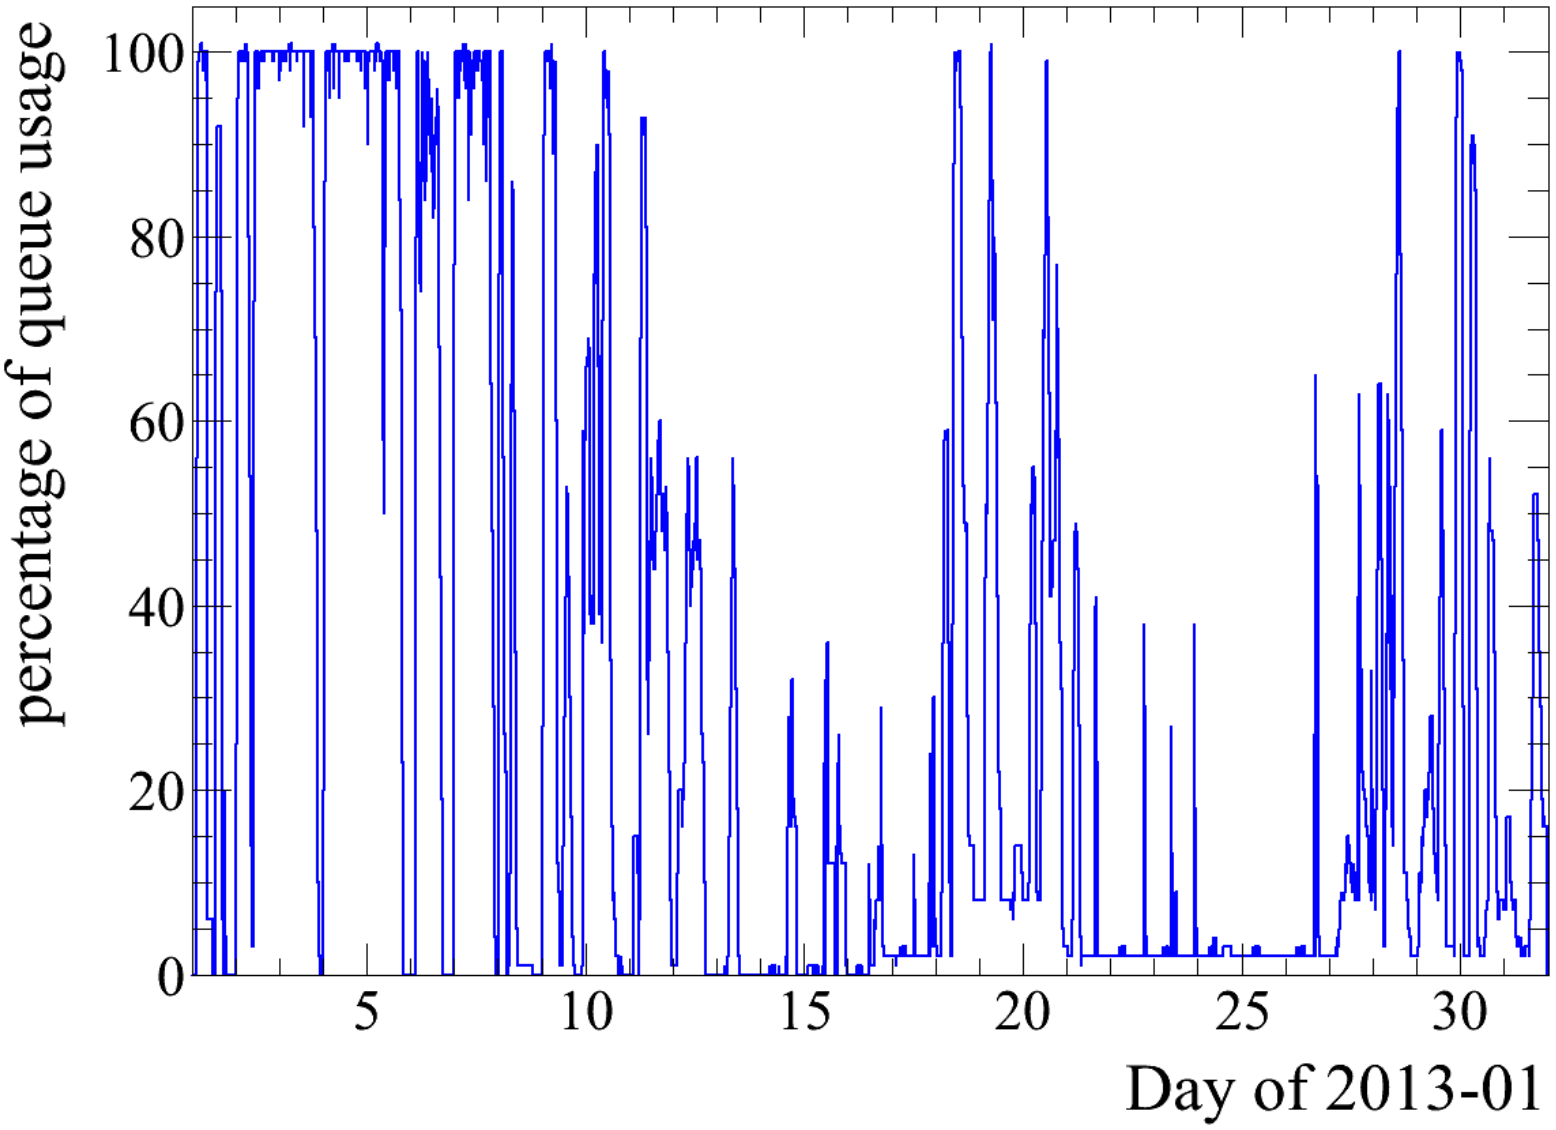

Usage monitoring

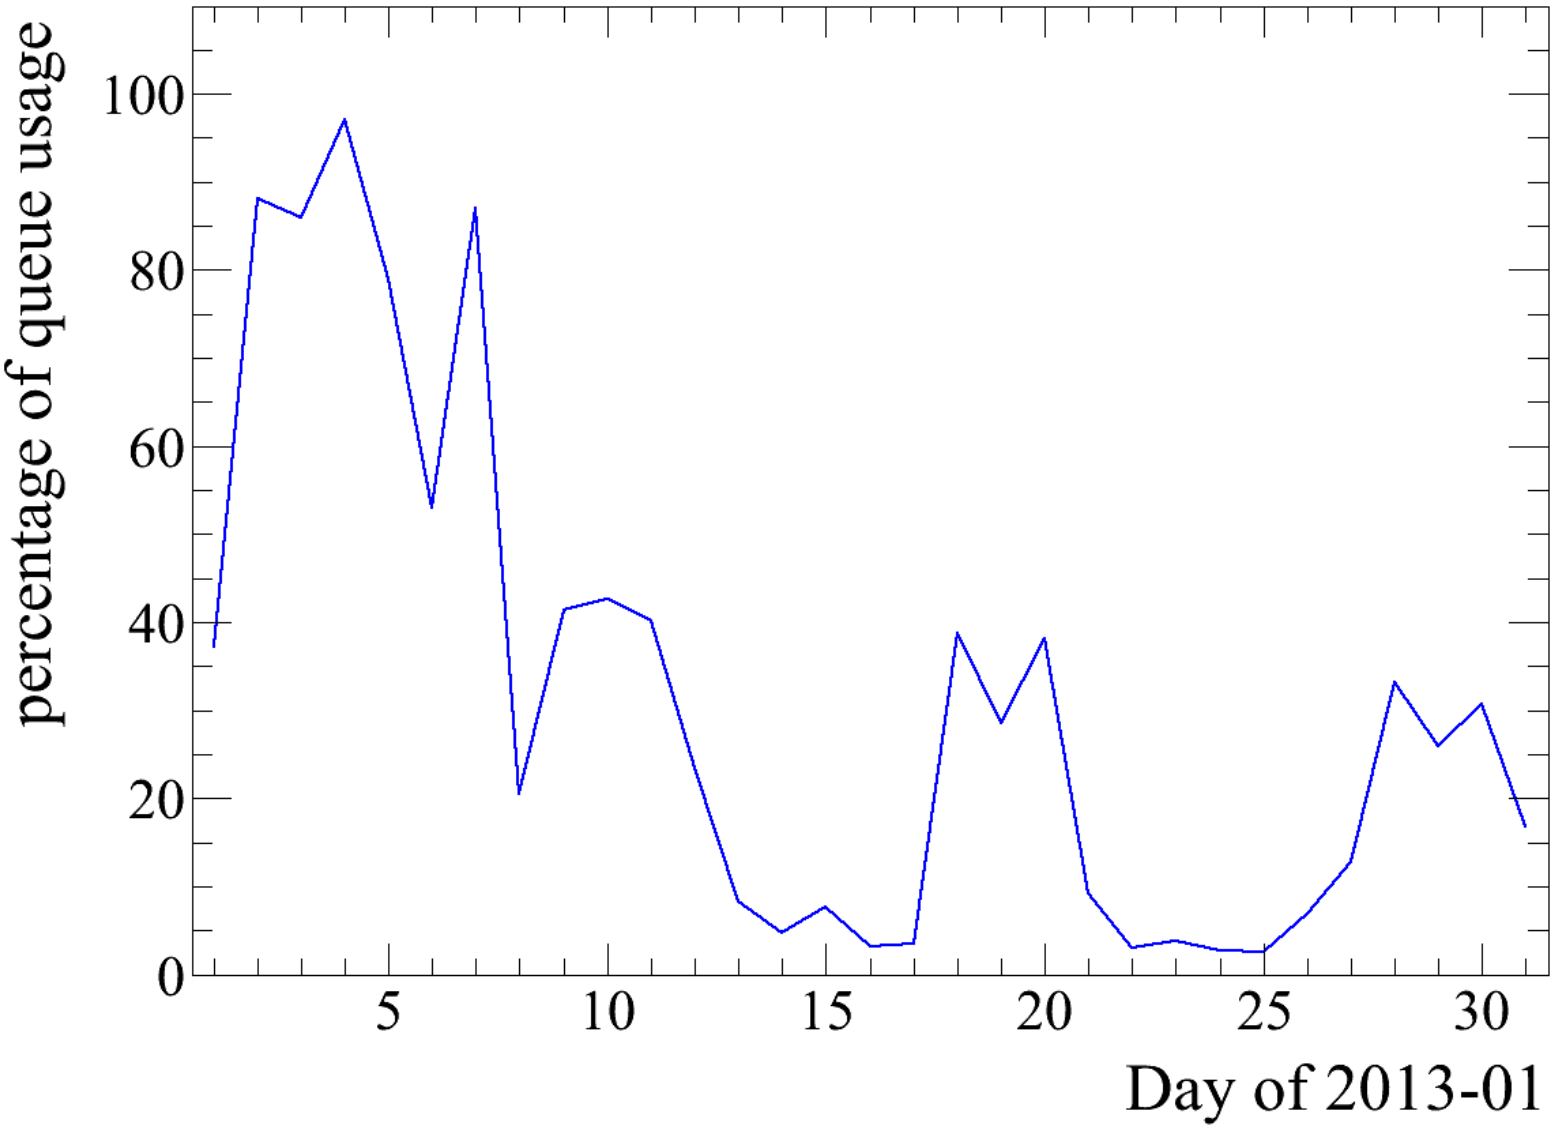

Backup system

- daily tape backups of wain061 and wain062 without any root files

- on-disk home backup using ZFS snapshots
  - frequent snapshot every 15min, overwritten every hour
  - frequent snapshot taken at the full hour becomes hourly snapshot, overwritten every day
  - hourly snapshot taken at midnight becomes daily snapshot, overwritten every month

- on-disk backup of /BFROOT where our releases, packages, and VM images are hosted
  - same as for home backups, but without 15-min snapshots
    (same backup script can be used for /home and /BFROOT backups)

- on-disk backup of the AWG space
  - due to content of large and often changing files, only one daily ZFS snapshot

- users can easy access these snapshots
  - their own home-snapshots are located at: $HOME/.zfs/snapshot/
  - AWG backups for example under /awg/breco/.zfs/snapshot/

Backup against failure of wain062

Problem: wain062 hosts home directories, VM images, code repository,...
If wain062 has a hardware failure resulting in an outage (motherboard, CPU, network interface,...) nobody can work on the LTDA.

Solution:
- daily snapshot of everything under /home and /BFROOT is send to wain061
- wain061 keeps only one copy of /home and /BFROOT, renewed once a day
- in case of failure, wain061 can serve /home and /BFROOT
- all machines use autofs which eliminates any static NFS mounts

Validation system

Problem: There have been outages due to the system not working after updates from Red Hat have been applied.

Solution: Develop a validation system to make sure everything is working with newest updates.

Validation system:

- test server is setup like batch machines and gets automated updates
- reboot after updates have been installed and then run many test jobs (VM, NFS access)
- if everything is fine after many cycles of processing, new updates are written to an approved list
- all batch machines get only updates which are in this list
- all batch machines get also rebooted after updates have been applied
- if validation fails: notification by email

User point of view

- ssh bbrltda to enter the LTDA, just like ssh yakut
- home directory is not the AFS one, but on the login machines there is access to AFS
- user can checkout an analysis release and needed packages from CVS, and edit their code
- ROOT can be used interactively to analyze ntuples
- to compile users have to submit a job and request the needed OS
- users can submit their analysis jobs to the queue using the PBS commands or using LSF like commands
  - scripts have been written to translate LSF commands and options into PBS commands
  - old job submission scripts have not to be rewritten

  $\Rightarrow$ Working on the LTDA is not much different than to work on any other TierA site.

User point of view

Limitations:

- interactive compile of code on the login machines not possible
- VMs have read-only access to home directories
- VMs have no AFS access and can not access something outside LTDA

How to deal with that:

- users can run interactive VMs to get a shell and compile their code
- for each supported OS a "persistent" VM is available
  - 4CPUs, 8GB RAM for each of currently 3 VMs
  - **usage:** ssh sl4, ssh sl5, or ssh sl6 from the login machines
- analysis code and job output should go to AWG space, but not to $HOME
- AFS access on the login machines available
- job output can be copied outside of the LTDA on the login machines too

Summary

- Torque/Maui have been proved to be a good choice
- using single ZFSs for home directories and AWG together with autofs for easy administration
- using qemu virtualization is stable, easy and fast
- benefits for HT on compensate possible speed issues of single jobs
- improvement of memory usage with KSM
- system and usage monitoring are in place
- very good backup system in place
- emergency plan for complete failure of wain062 in place
- validation system to test system updates in place