


2009 CASTOR F2F

Monitoring at CERN

Miguel Coelho dos Santos

- Overview
- Daemon monitoring
- Probe
- Stager & DLF
- Service Displays
- Dashboard & Cluman
- Summary



- Daemon monitoring and alarms
 - Using lemon to monitor how many daemons should be running and alarm (or have automatic action - restart) when exceptions happen
- Using lemon-sensor-castormon (by Dennis Waldron) to monitor:
 - Process count
 - Process Uptime
 - Total VSZ
 - Total RSS
 - Total FD
 - Thread Count
 - Number of active threads

Useful to detect problems such as memory leaks, file descriptor leaks, etc.

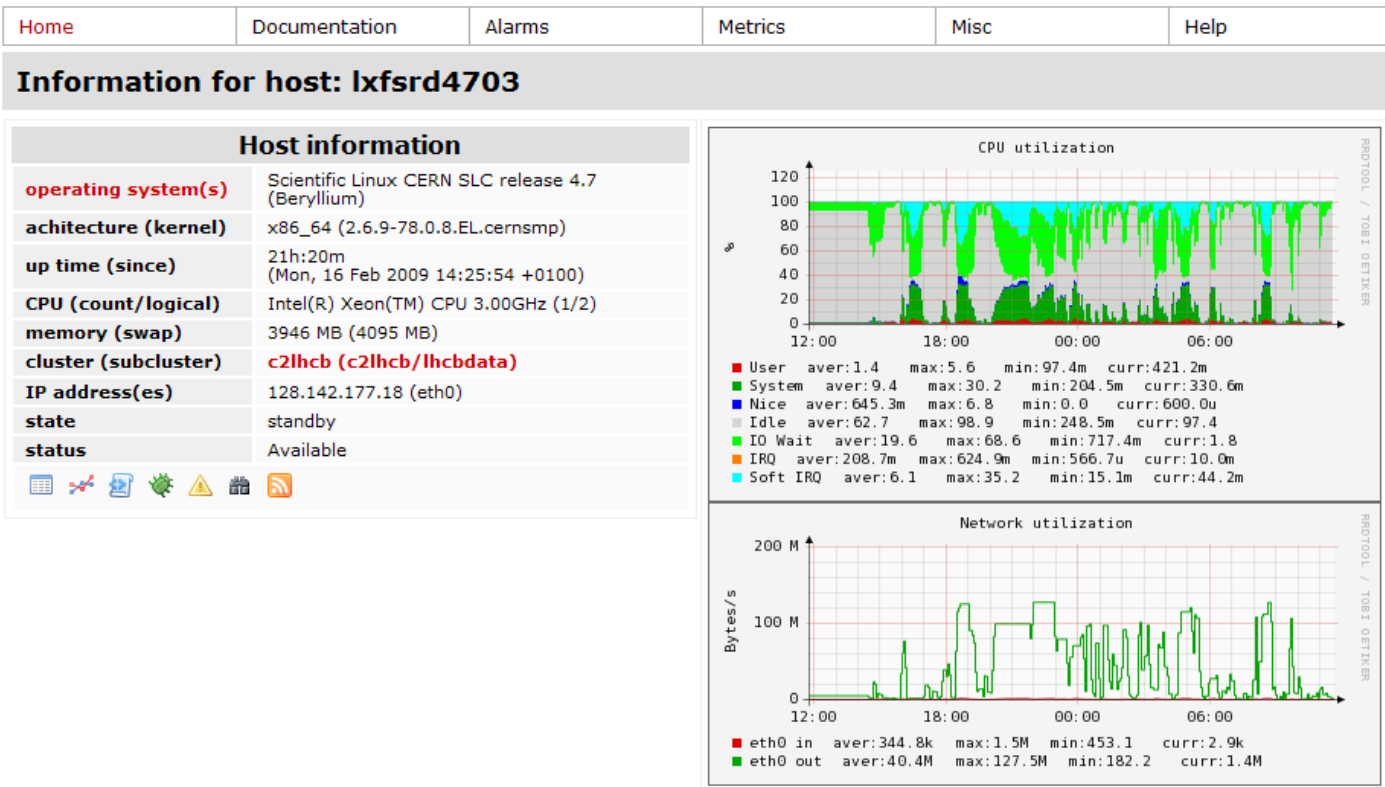
Also useful to identify thread starvation or over dimensioning (each thread usually uses a DB connection so thread counts!)
- Also parses name server and cupv logs for request information

- The castor probe looks at the system from the user perspective
- Different calls are timed:
 - `r fio_open_w`, `r fio_write`, `r fio_close_w`, `r fio_open_r`, `r fio_read`,
`r fio_close_r`, `r stager_qry`, `r stager_rm`
- Response times and errors are used to provide a number (availability) that tries to reflect current user experience
- Very helpful to find problems and a good way to feedback system responsiveness information to the users
- Probe writes a simple file with the results.
- CERN uses lemon and SLS do display availability
- Probe is packaged in `castor-hsmttools` RPM.

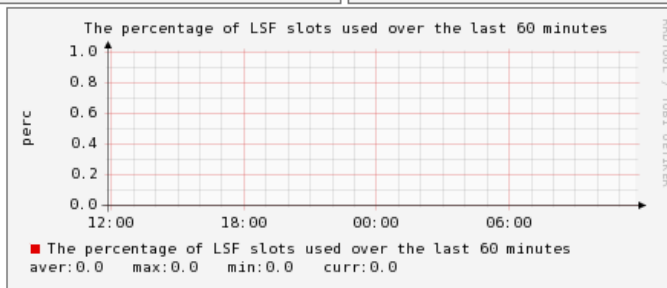
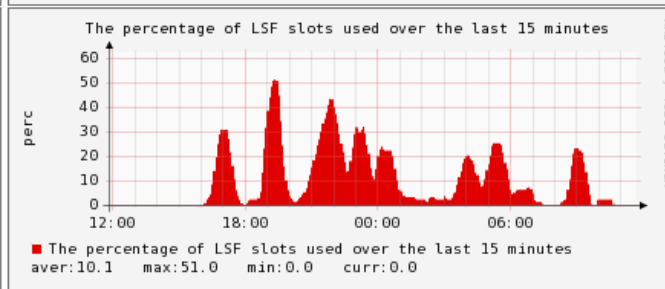

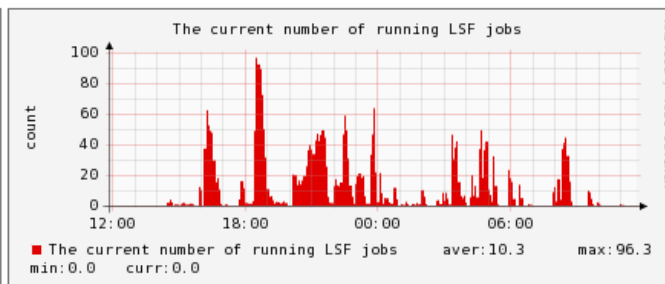
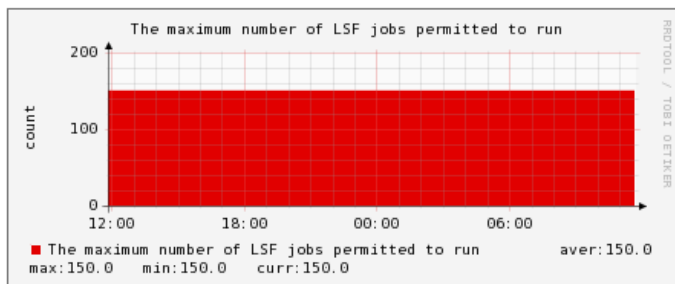
- Stager & DLF DB scheduler jobs gather important service metrics such as: disk cache efficiency, migration queue, recall queue, garbage collection, etc.
- DLF scheduler jobs have been reviewed recently by CASTOR developers, stager jobs should be reviewed next
- DLF scheduler jobs collect a large number of WLCG storage metrics.
- Information is stored in a DB table. A lemon sensor reads this data and feeds it to lemon where it can be read and/or plotted.
- Should be easy to plug nagios in and replace lemon since all the intelligence is server side, the rest is “just plotting”

- Daemon, Probe, Stager and DLF information is imported into Lemon (the Fabric Monitoring tool).
- This data can then be used by different consumers:
 - Lemon display pages: service and server monitoring, ‘eye’ correlations
 - SLS: a standard way to provide users with basic (most relevant) service information
 - CASTOR Dashboard: a simple page that gives an overview of all the instances. The “Hi CASTOR, I’m an admin, show me how you’re doing this morning” page. Also used to fast prototype and select important plots that should make it to the user via SLS and/or Lemon
 - Cluman: fast plotting, cluster overview, metric correlation and hopefully in the future to perform actions

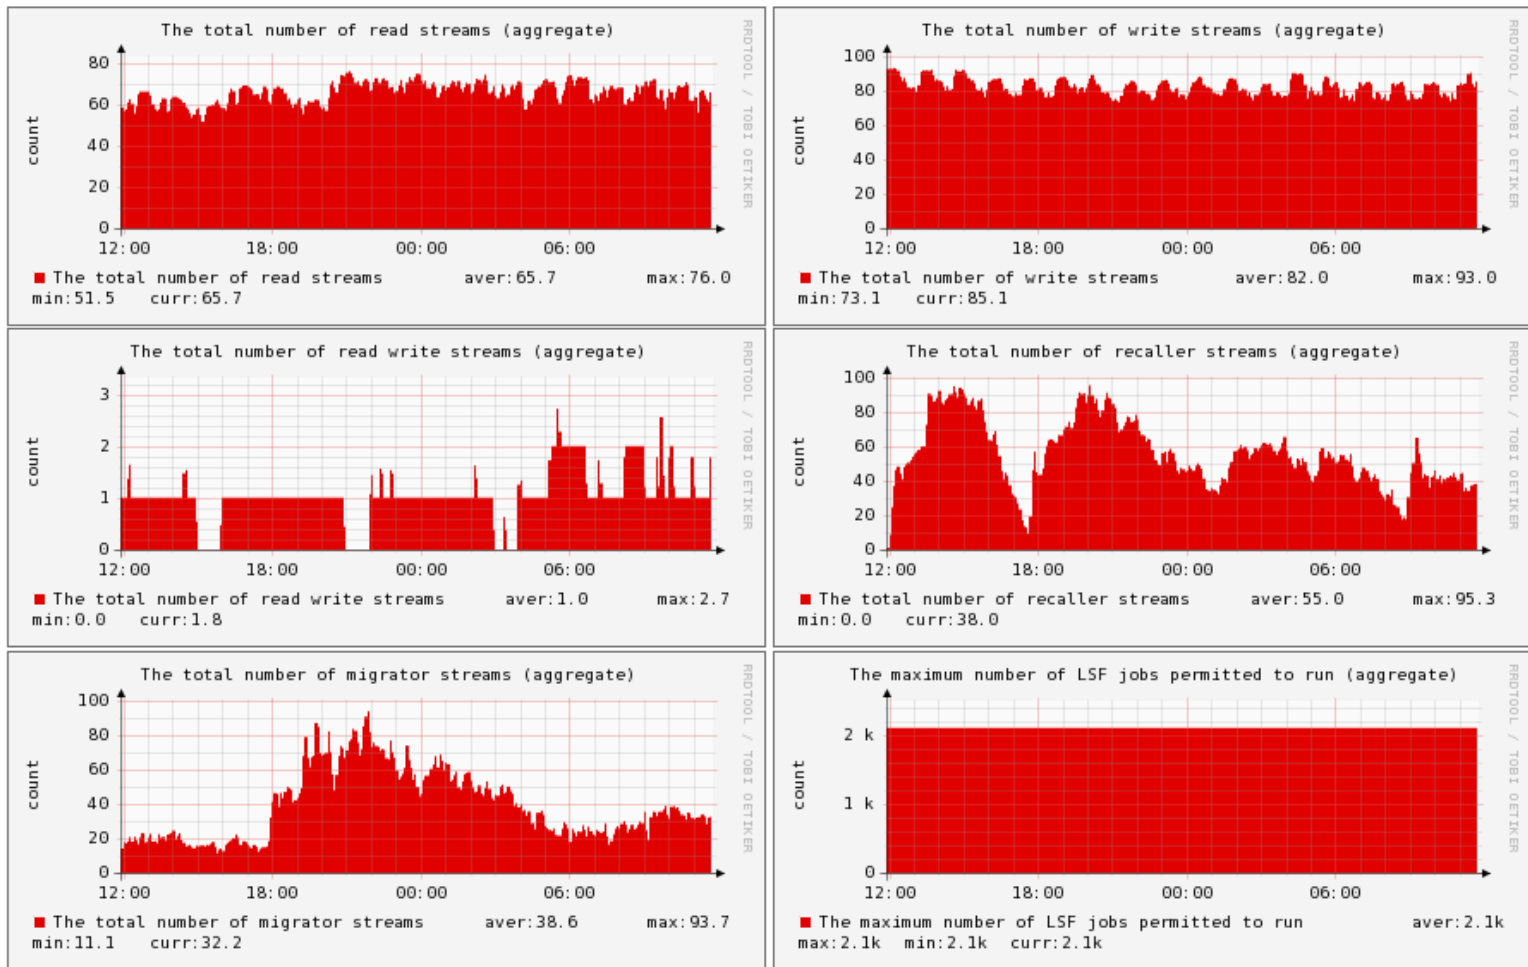
- Typical lemon disk server page




- Slot usage over a diskserver



- Aggregation of several disk server to have a stream usage overview for a cluster (c2repack)



- Protocol information and filesystem information (not deployed in production yet, preprod cluster)



- Standard user oriented service displays (i.e. data from lemon display in SLS for user convenience)

CASTO T1TRANSFER service class of CASTORCMS instance





CASTO Service information

availability: ██████████

percentage: 100%

full name: **T1TRANSFER service class of CASTORCMS instance**
short name: T1TRANSFER Service Class
group: IT-FIO

email: **Castor.Support@cern.ch**
web site: <http://cern.ch/castor>

service managers: Miguel Marques Coelho Dos Santos 
Ignacio Reguero 
Jan Van Eldik 
Olof Barring 

availability: ██████████

percentage: 100%

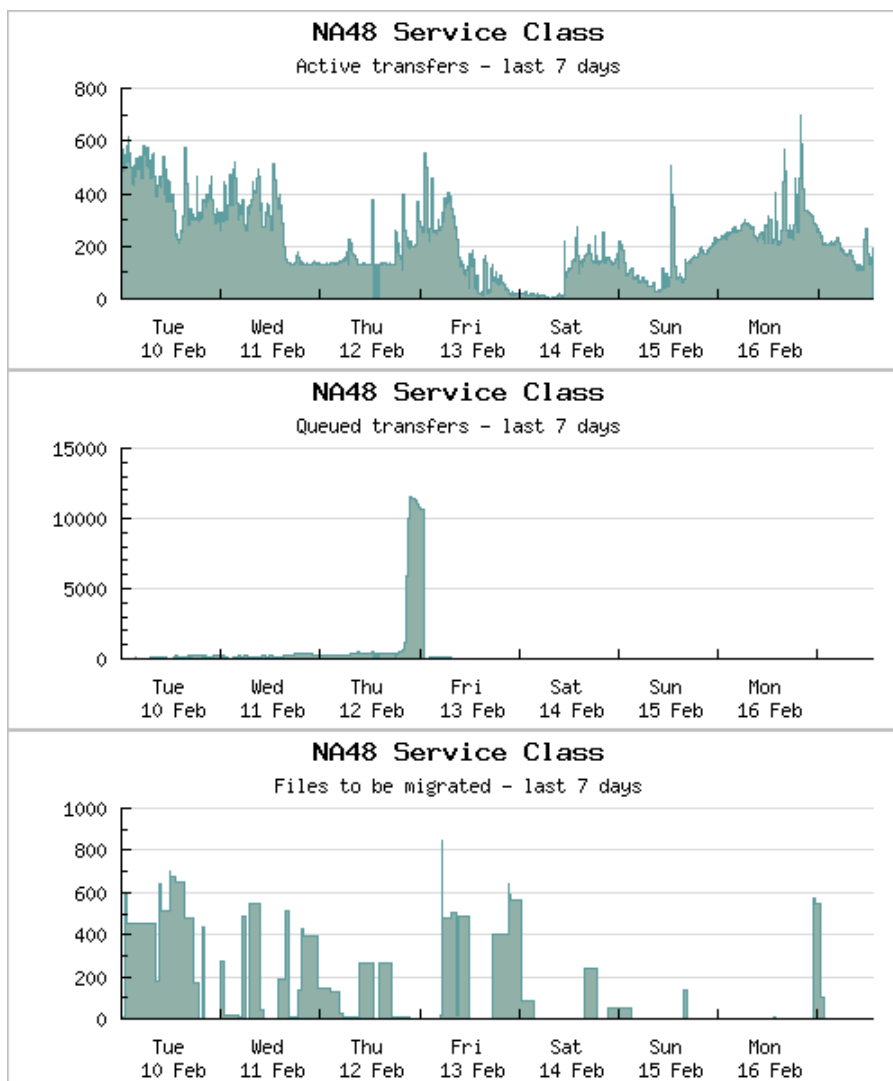
availability info: Transfers running/pending: 28/0
Files to be migrated/recalled 4 /0
Free/Total Space 25.78/170.22 (15.1 %)
Disk0Tape1

status: **available**

Additional service information (more)

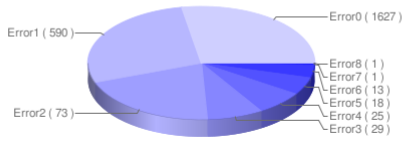
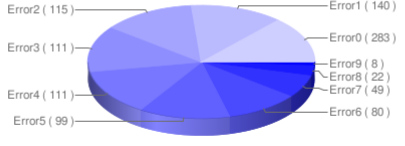
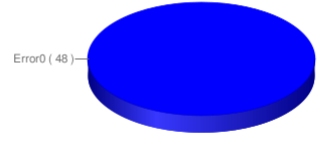
Active transfers:	28
Queued transfers:	0
Files to be migrated:	4
Files to be recalled:	0
Total Space TB:	170.22
Free Space TB:	25.78
Percentage Free Space:	15.1

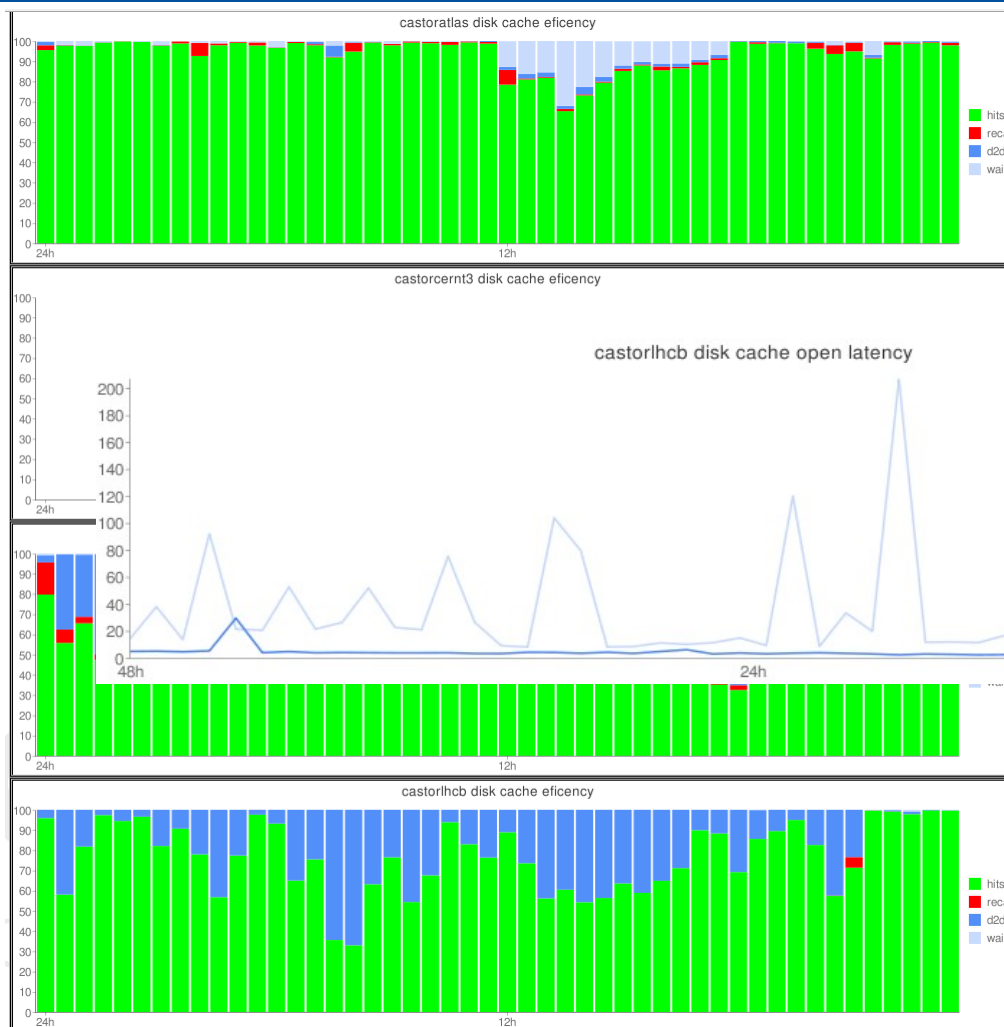
Storage Class: Disk0Tape1




- For debugging and fast-plotting (using data from lemon 'on demand', no data is stored)

CASTOR Dashboard

	availability	Files		Jobs		DiskCopy		
		Pend+ToBeMig/Recalled		Running/Pending		Staged/WV_D2D/WV_Recall		
c2alice								
allimdc		0+0		0	0	174137	0	0
default		2+9		10	0	1865341	0	0
								
		Error (count)		Message				
		Error0 (1627)		JobManager: Job-terminated-due-to-scheduling-error				
		Error1 (590)		Rechandler: Error-in-the-executing-the-policy-script				
		Error2 (73)		SRMServer: SOAP-Error				
		Error3 (29)		SRMServer: Request-Failed				
		Error4 (25)		job: error				
		Error5 (18)		SRMServer: Framework-message				
		Error6 (13)		SRMServer: Stager-returned				
		Error7 (1)		SRMDaemon: Framework-message				
		Error8 (1)		SRMServer: Admin-request-to-release-space-received				
c2atlas								
atl3		0+0	0	0	0	34315	0	0
atldata		0+0	0	0	0	572275	0	0
atlprod		5+0	0	5	0	2832191	0	0
t0merge		0+0	0	7	2	261106	0	0
atlasgroupdisk		0+0	0	27	0	511794	0	0
atlasuserdisk		0+0	0	7	0	501271	0	0
t0atlas		281+5	0	3	0	211150	0	0
atlspecial		0+547	0	0	0	2058	0	0
default		60+3		268	0	397573	0	13
atcal		0+0	0	52	0	678765	18	0
								
		Error (count)		Message				
		Error0 (283)		SRMServer: Request-Failed				
		Error1 (140)		SRMServer: Stager-returned				
		Error2 (115)		SRMServer: Framework-message				
		Error3 (111)		rtcpd: Tape-unavailable				
		Error4 (111)		rtcpd: failed-system-call				
		Error5 (99)		job: system-call-error				
		Error6 (80)		job: error				
		Error7 (49)		SRMServer: SOAP-Error				
		Error8 (22)		migrator: failed-system-call				
		Error9 (8)		SRMServer: Invalid-Request				
c2cern3								
cmscft2		0+0	0	0	0	7	0	0
default		1+0	0	0	0	179	0	0
atlascernuserdisk		0+0	0			6888	0	0
								
		Error (count)		Message				
		Error0 (48)		TapeErrorHandler: failed-system-call				





Project **c2atlas**

Refresh Reload

Project **c2atlas**

Fetches 236 entities. Took 14.251 seconds.

[clear](#)

Map View Tree View

Color cells by:

- state
- CPU Load
- IO wait
- Slot usage
- rpio network timeouts
- Migrator
- Recaller Streams
- CPU temperature
- Number of users
- Uptime in seconds

Display flags:

- Console access warning
- Old/wrong kernel
- High CPU load
- Not in production

CluMan - the cluster management tool - developed by CERN IT/FIO ([more](#))

last released: 10:50, Thursday 04 Dec 2008

- Like other services, CASTOR has a lot of different sources for information
- Lemon sensors are used to collect this data
- Historic data is stored in Lemon
- Different displays for different people
 - SLS is CERN's 'gateway' for users to get service data. Probe information and a few important metrics are key.
 - Lemon display pages show server and service information for power users and admins
 - Dashboard gives a one page overview of all the instances and allows for prototyping of plots (before they are in lemon and SLS)
 - Cluman is a cluster admin tool that allows for cluster plotting of lemon metrics. Easy to configure, useful for correlation and debugging.