

Local Monitoring at SARA

Ron Trompert

SARA

- Monitors nodes for
 - Load
 - Memory usage
 - Network activity
 - Disk usage

- Monitors running jobs


SARA HPC > Matrix Cluster > Joblist

Batch	Jobs	Nodes	Cpus
Capacity		36	72
Running	28	15	28
Queued	7	7	7
Total	35	22	35
Free		21	44
View	7	7	7




Last updated: Tue 12 Sep 2006 16:54:28

Cluster queue usage



Id	S	User	Queue	Name	Req. CPU time	N/P	Queued	Started	Runningtime
156960	Q	esrsgm	short	STDIN	4:00:00 hours	1/1	Tue 12 Sep 2006 16:49:18		
156959	Q	esrsgm	short	STDIN	4:00:00 hours	1/1	Tue 12 Sep 2006 16:49:14		
156958	Q	esrsgm	short	STDIN	4:00:00 hours	1/1	Tue 12 Sep 2006 16:49:13		
156957	Q	esrsgm	short	STDIN	4:00:00 hours	1/1	Tue 12 Sep 2006 16:49:12		
156954	Q	cms011	short	STDIN	4:00:00 hours	1/1	Tue 12 Sep 2006 16:23:47		
156948	Q	atlas023	medium	STDIN	1 day - 9:00:00 hours	1/1	Tue 12 Sep 2006 14:46:17		
156946	Q	atlas023	medium	STDIN	1 day - 9:00:00 hours	1/1	Tue 12 Sep 2006 14:44:17		



■ Network monitoring

- Cern router (ping)
- SARA-r1 router (ping)
- Cern fts machines (ping)

■ Monitoring services

- gridftp doors
 - ▶ checks service listening on port 2811
- srmdoor
 - ▶ checks service listening on port 8443
- gsidcapdoor
 - ▶ checks service listening on port 22128
- Rfiod (cxfs/dmf clients)
 - ▶ checks service listening on port 5001

- Checks the services/network at user-specified intervals
- Sends email when user-specified error conditions occur
- Standard build-in tests
 - Tcp/udp
 - http,snmp,imap, sql databases
- Possibility to sound an alarm
- Possibility to invoke services rather just do a tcp check
- Supports graphing

- It's free 😊
- Easy to configure 😊
- No SL rpms ☹
 - Needs some perl modules and other programs not in SL.

Top:Alles

name Alles

status up

Name	Status
CERN FTS	fts101 fts102 fts103 fts104
Routers	CERN router SARA r1
Storage nodes	ant1 ant2 ant3 ant4 ant5 srm
World Reachability	Root Servers
dm nodes	dm1 dm2

Status: up since Tue 5 Sep 17:24:57 2006

	start	elapsed time	% up	% down	times down
Today	Mon 11 Sep 00:00:00 2006	0:00:00	0.00	0.00	0
Yesterday	Sun 10 Sep 00:00:00 2006	23:59:35	100.0	0.00	0
2 Days Ago	Sat 9 Sep 00:00:00 2006	1d 0:00:00	100.0	0.00	0
This Month	Mon 4 Sep 15:09:33 2006	6d 8:50:02	89.77	10.23	4
Last Month	Tue 1 Aug 00:00:00 2006	31d 10:26:51	89.47	10.53	7

Top:Alles:Storage_nodes:ant1

name ant1

status up

Name Status

[Ping](#) up

[gridftp](#) up

[gsidcap](#) up

[srm](#) up

Status: up since Tue 5 Sep 13:11:15 2006

	start	elapsed time	% up	% down	times down
Today	Mon 11 Sep 00:00:00 2006	0:00:00	0.00	0.00	0
Yesterday	Sun 10 Sep 00:00:00 2006	23:59:35	100.0	0.00	0
2 Days Ago	Sat 9 Sep 00:00:00 2006	1d 0:00:00	100.0	0.00	0
This Month	Mon 4 Sep 15:09:33 2006	6d 8:50:02	99.89	0.11	1
Last Month	Tue 1 Aug 00:00:00 2006	31d 12:26:51	99.84	1.16	1

Top:Alles:Storage_nodes:ant1:TCP_8443_ant1.grid.sara.nl


name TCP
status **up**
hostname ant1.grid.sara.nl
port 8443

Status: up since Tue 5 Sep 13:11:15 2006


	start	elapsed time	% up	% down	times down
Today	Mon 11 Sep 00:00:00 2006	0:00:00	0.00	0.00	0
Yesterday	Sun 10 Sep 00:00:00 2006	23:59:35	100.0	0.00	0
2 Days Ago	Sat 9 Sep 00:00:00 2006	1d 0:00:00	100.0	0.00	0
This Month	Mon 4 Sep 15:09:33 2006	6d 8:50:02	99.90	0.10	1

- Monitoring disk-disk throughput with home-grown python scripts and the dCache billing files.
 - Information is extracted from billing files every hour and put in a postgres db
 - ▶ Was written before dcache logged the billing data in postgres
 - Scripts running in the cron every hour extract info from db and use gnuplot to generate the graphs
 - Requires Gnuplot development release and a patch for histogram support

- Host-based and VO-based graphs
 - Overall and per VO and per host
- Graphs of throughput, number of transfers and amount of data moved



- Monitoring disk-tape throughput with home grown python scripts and DMF's dmdlogs
- Same operation as disk-disk



■ Settings and agent status

- Scripts uses services.xml file to get info about myproxy information and ChannelManagement service
- Run in the cron every 5 minutes
- Scripts renew proxy when needed
- Scripts get info about channel settings with the glite-transfer-channel-list command
- Does not require access to the oracle db. Able to monitor other people's FTS.

Mon Sep 11 01:01:19 CEST 2006

Channel	Nominal throughput	Number of VO shares	streams	State	Number of files	Bandwidth	vo shares			
							atlas	dteam	lhcb	alice
ASCC-SARA	0	4	5	Active	5	0	25	25	25	25
BNL-SARA	0	2	5	Active	5	0	25	25		
CNAF-SARA	0	4	5	Active	5	0	25	25	25	25
GRIDKA-SARA	0	4	5	Active	5	0	25	25	25	25
IN2P3-SARA	0	4	5	Active	5	0	25	25	25	25
ITEP-SARA	0	2	6	Active	5	0	25	25		
JINR-SARA	0	2	6	Active	5	0	25	25		
NIKHEF-SARA	0	4	5	Active	5	0	25	25	25	25
PIC-SARA	0	4	5	Active	5	0	25	25	25	25
PRAGUE-SARA	0	4	5	Active	5	0	25	25	25	25
RAL-SARA	0	4	5	Active	5	0	25	25	25	25
RRC-SARA	0	2	6	Active	5	0	25	25		
SARA-Itep	0	2	6	Active	5	0	25	25		
SARA-JINR	0	2	6	Active	5	0	25	25		
SARA-NIKHEF	0	4	5	Active	5	0	25	25	25	25
SARA-PRAGUE	0	4	5	Active	4	0	25	25	25	25
SARA-RRC	0	2	6	Active	5	0	25	25		
SARA-STND	0	2	6	Active	5	0	25	25		

■ Transfers and jobs

- Scripts retrieve information from oracle db of transfers and jobs of the past 24 hours
 - ▶ At the time of writing the glite-transfer-* I/F did not support getting the necessary information about jobs and transfers.
- Displays statistics.
- Displays information about individual jobs and transfers
 - ▶ Status
 - ▶ Source
 - ▶ Destination
 - ▶ Reason of the failure, if applicable

Statistics

The statistics of the transfers performed since 2006-09-09 23:03:04 +00:00

Mon Sep 11 01:03:04 CEST 2006

t_transfer.transfer_state	number
Completed	8
Failed	18

The statistics of the transfers performed since 2006-09-09 23:03:04 +00:00 grouped by Channel

Mon Sep 11 01:03:04 CEST 2006

channel_name	Total	Completed	Failed
IN2P3-SARA	4	4	
RAL-SARA	6	4	2
TRIUMF-SARA	16		16

The statistics of the transfers submitted since 2006-09-09 23:03:04 +00:00 grouped by VO

Mon Sep 11 01:03:05 CEST 2006

vo_name	Total	Completed	Failed
dteam	26	8	18

Transfers

FTS Transfer Status

The status of the transfers submitted during the last 24 hours.


Thu Sep 14 16:03:03 CEST 2006

job_id	logical_name	job_state	source_surl
6662692e-43f7-11db-88c3-8bd46d8ee2fc		Active	srm://srm.grid.sara.nl:8443/srm/managerv1?SFN=/pnfs/grid.sara.nl/data/atlas/jiri/jiri.1GB.1.1
6662692e-43f7-11db-88c3-8bd46d8ee2fc		Active	srm://srm.grid.sara.nl:8443/srm/managerv1?SFN=/pnfs/grid.sara.nl/data/atlas/jiri/jiri.1GB.1.1
0aa89908-43e3-11db-88c3-8bd46d8ee2fc		Failed	srm://dcache.gridpp.rl.ac.uk:8443/pnfs/gridpp.rl.ac.uk/data/lhcb/sc3/test/roberto/zz_zz.f
0aa89908-43e3-11db-88c3-8bd46d8ee2fc		Failed	srm://dcache.gridpp.rl.ac.uk:8443/pnfs/gridpp.rl.ac.uk/data/lhcb/sc3/test/roberto/zz_zz.f
0aa89908-43e3-11db-88c3-8bd46d8ee2fc		Failed	srm://dcache.gridpp.rl.ac.uk:8443/pnfs/gridpp.rl.ac.uk/data/lhcb/sc3/test/roberto/zz_zz.f

Transfers

dest_turl	transfer_time	reason_class	reason
gsiftp://lcg60.sinp.msu.ru:2811//pnfs/sinp.msu.ru/data/atlas/jiri.1GB.1.20060914.1	14-SEP-06 02.02.01.000000 PM +00:00		
gsiftp://lcg60.sinp.msu.ru:2811//pnfs/sinp.msu.ru/data/atlas/jiri.1GB.1.20060914.1	14-SEP-06 01.46.21.000000 PM +00:00	NETWORK_ERROR	Transfer failed due to possible network problem - timed out.
	14-SEP-06 12.59.07.000000 PM +00:00	SOURCE_SRM_ERROR	Failed on SRM get: SRM getRequestStatus timed out on get
	14-SEP-06 12.59.07.000000 PM +00:00	SOURCE_SRM_ERROR	Failed on SRM get: SRM getRequestStatus timed out on get
	14-SEP-06 12.59.07.000000 PM +00:00	SOURCE_SRM_ERROR	Failed on SRM get: SRM getRequestStatus timed out on get
	14-SEP-06		

Future plan

-  Combine job and transfer output so you have information about individual transfers, VOs and channels in one view

■ <http://winnetou.matrix.sara.nl/monitoring/datatransfer>