

# University of Sussex site report

Leonardo Rojas

Experimental Particle Physics group  
School of Mathematical and Physical Sciences

HEPSYSMAN – RAL 2017

# Our team

## **ITS Research team**

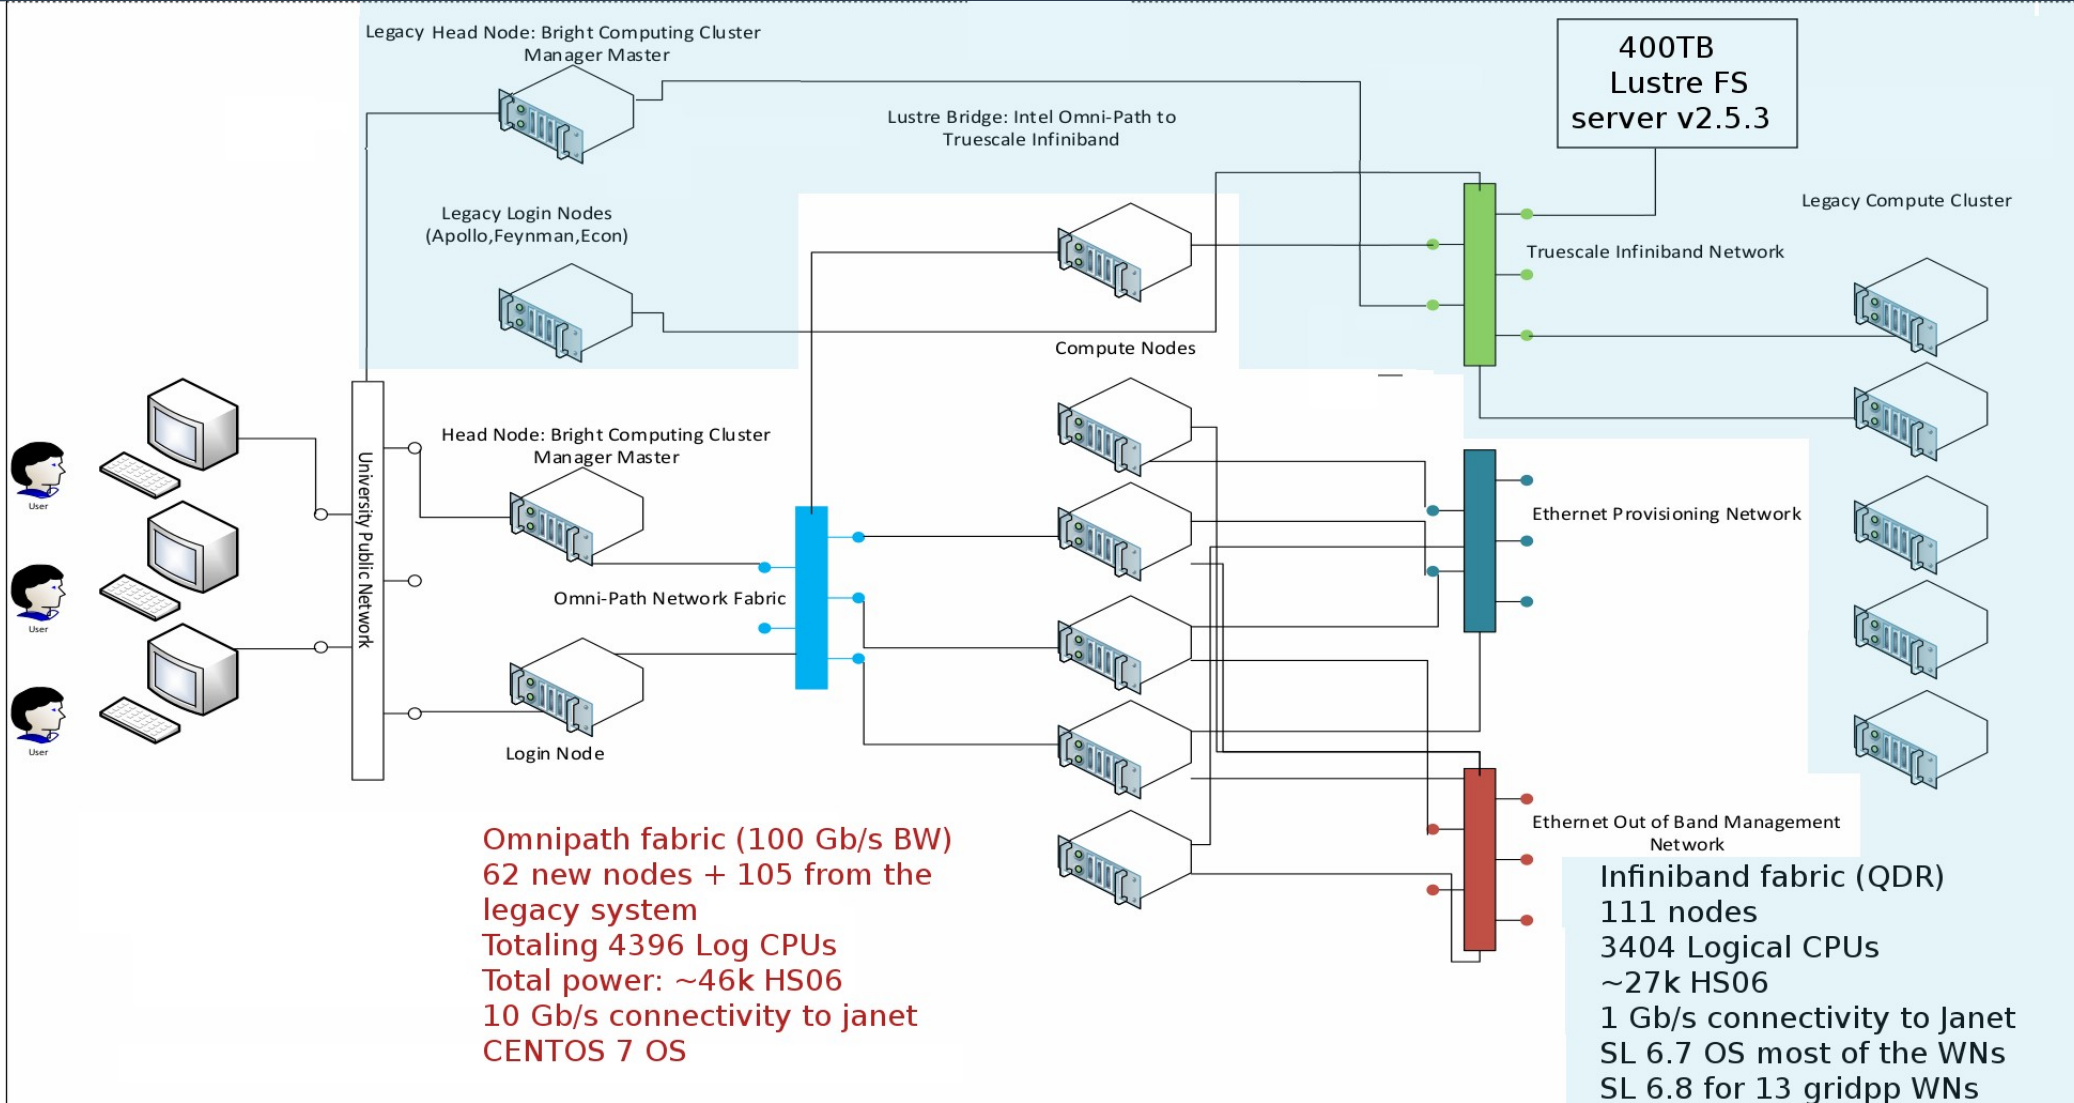
University's HPC development, implementations and maintenance.

- Jeremy Maris
- Alhamdu Bello
- Bernie Broughton

## **MPS Sys Admins**

- Leonardo Rojas  
Grid cluster manager Physics and Maths users HPC admin support and
- A new MPS sys admin, Roberto Scipioni, is actually starting tomorrow.  
Physics and Maths users HPC support, Parallel programming support,  
Storage management for astrophysics.

# Current HPC system and short term development plan



# The Data Center



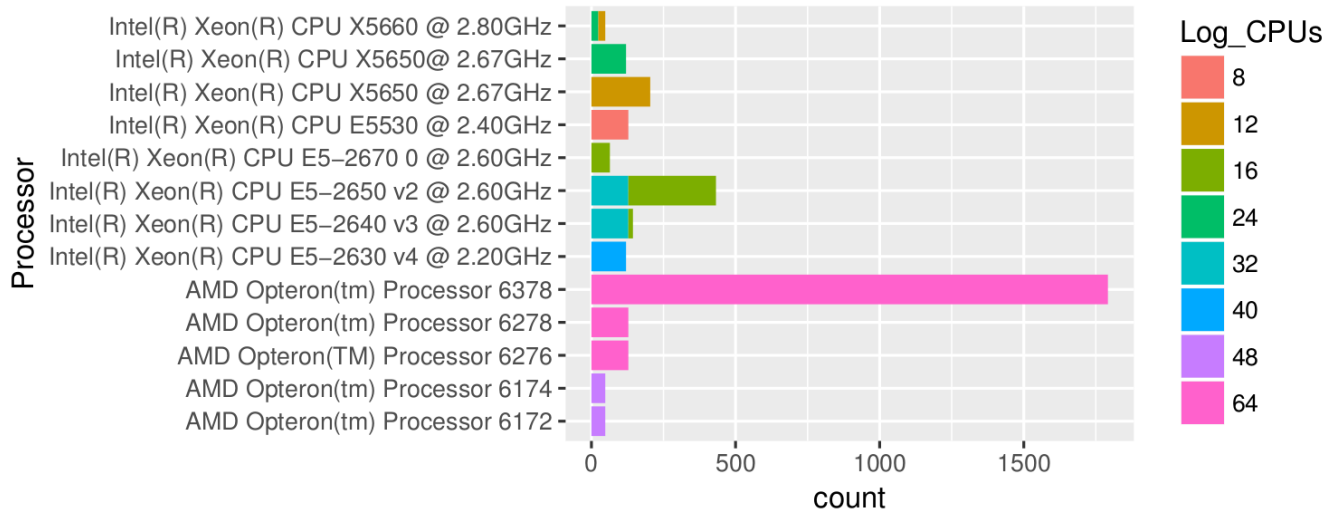
The data center was built in 2011  
The HPC uses 8 Racks with water  
cooled doors



New Omnipath  
fabric rack

# compute power: current legacy HPC system

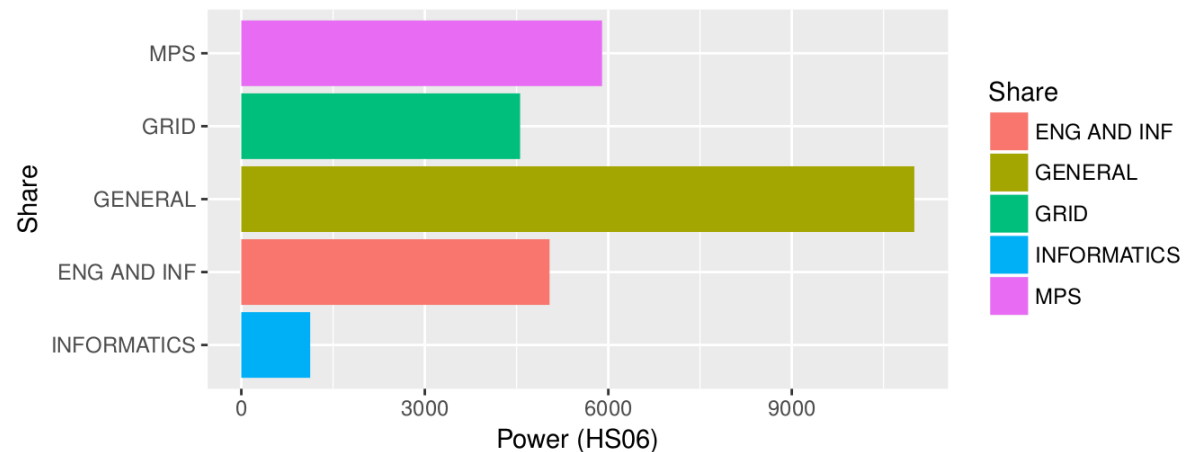
No. Logical CPU's (legacy HPC): 3404



Processor model distribution of our legacy HPC system.

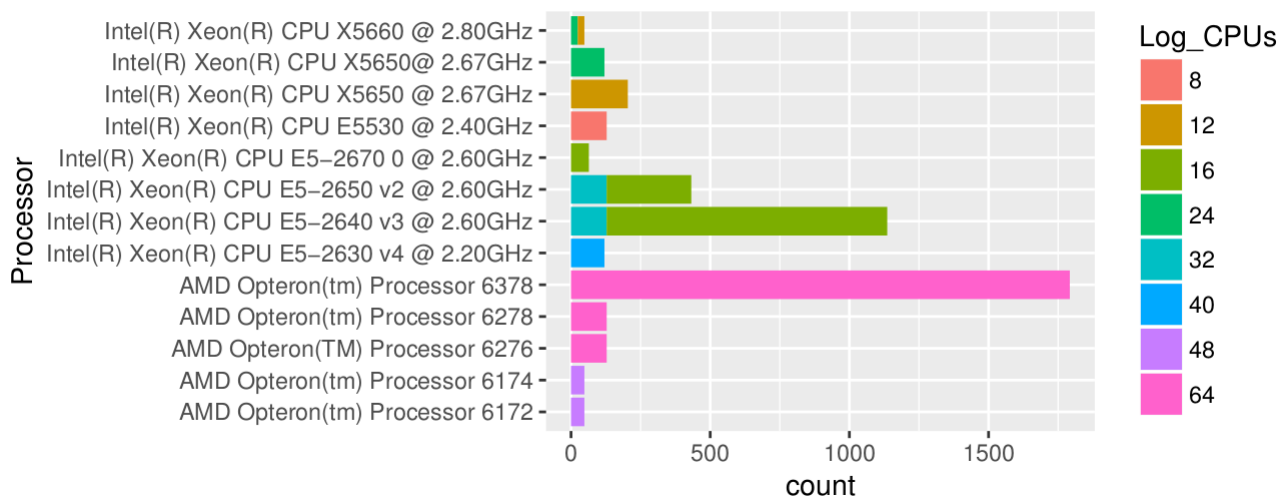
Compute power share amongst our research community

Total power (legacy HPC): 27626



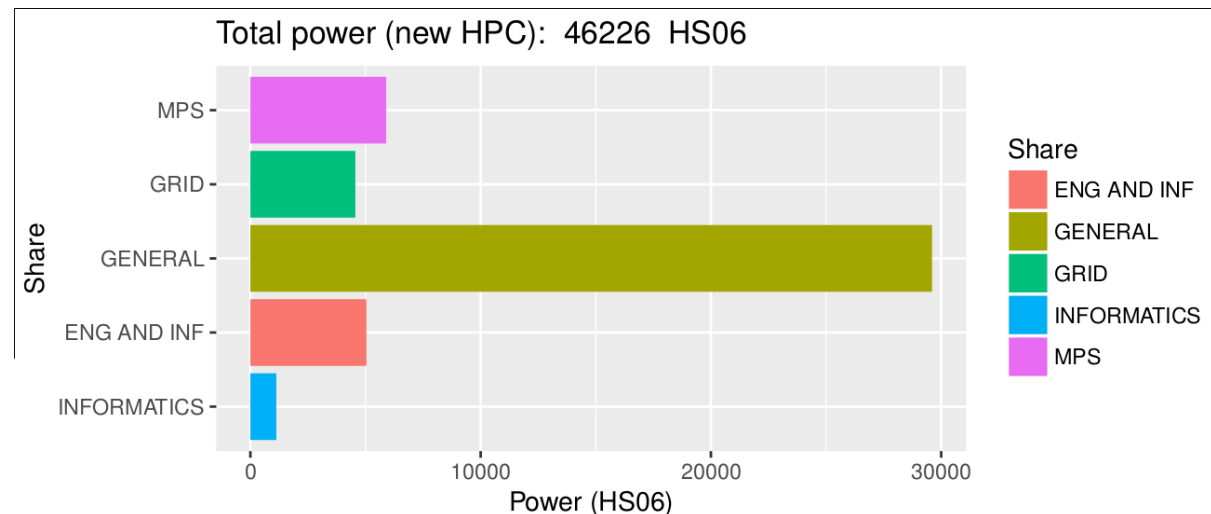
# Compute power: new HPC system

No. Logical CPU's (new HPC): 4396



Processor model distribution for our new consolidated cluster.

New compute power will go into the general purpose queues. New kit from research grants will be added to the OP side of the cluster



# Summary

- The new cluster is using the newest versions of Bright Cluster Manager v7.3 and UGE 8.5.
- New compute nodes are Centos 7 machines, old legacy compute nodes needs to be migrated from SL 6.7 to Centos 7, except the gridpp compute nodes. Old Bright cluster master should be decommissioned and all its data should be migrated to the new head.
- MPS and general purpose software stack must be recreated.
- Current 400 TB LustreFS is going to be decommissioned by 2019, current HPC uses server and client v2.5.3. In the mean time, the short term plan for new compute nodes is to use lustre client v 2.8 for compatibility issues (not tested yet).
- By 2019 all the HPC system should be using lustre server and client v2.10, so we can do small files on Meta Data Targets.
- We are going to have 10Gb/s Ethernet connection to Janet from the end points (including grid service nodes)
- We are going to have a ~1000 slots OmniPath island for MPI processing with 100 Gb/s BW, and interoperability with old legacy nodes through an Lnet router.

# Our gridpp site

## Service Overview For Host Group 'site-UKI-SOUTHGRID-SUSX'

Up	Down	Unreachable	Pending	Ok	Warning	Unknown	Critical	Pending
4	0	0	0	28	0	0	0	0
<i>All Problems</i>		<i>All Types</i>		<i>All Problems</i>		<i>All Types</i>		
0		4		0		28		

### Site UKI-SOUTHGRID-SUSX (site-UKI-SOUTHGRID-SUSX)

Host	Status	Services	Actions
grid-argus-02.hpc.susx.ac.uk	UP	1 OK	
grid-bdii-02.hpc.susx.ac.uk	UP	6 OK	
grid-cream-02.hpc.susx.ac.uk	UP	9 OK	
grid-storm-02.hpc.susx.ac.uk	UP	12 OK	



Last Updated: Mon Jun 12 11:09:26 CEST 2017

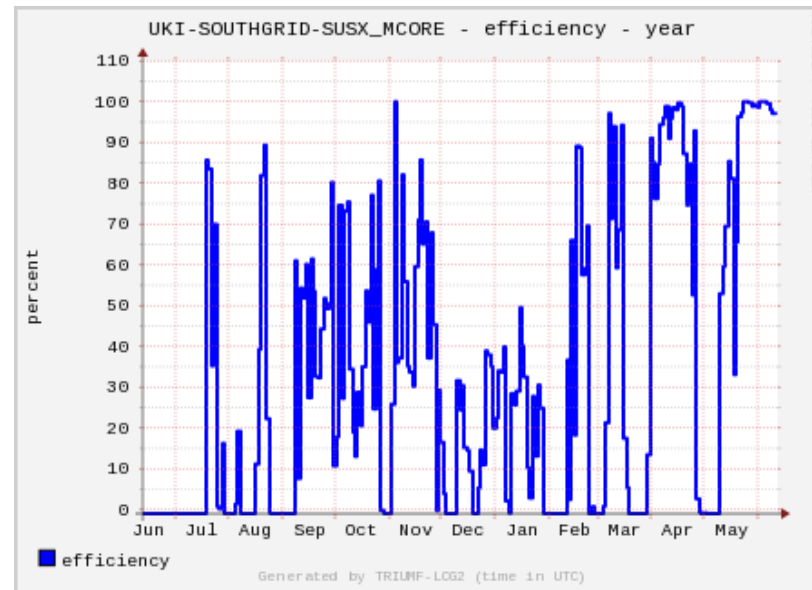
Updated every 90 seconds

Nagios® Core™ 4.3.1 - [www.nagios.org](http://www.nagios.org)

Logged in as /C=UK/O=eScience/OU=Sussex/L=PhysicsAndAstronomy/CN=leonardo rojas






# Sussex gridpp site compute power

node	DELL kit	Server Model <u>PowerEdge</u>	Processor model	Delivery date	Log/ CPUs ( <u>HT</u> on)	HS06/ Node	
101	2 x 	C6145	AMD <u>Opteron</u> Processor 6276	16 Jun 2011	64	515	
102					64	515	
113	2 x 	C6145	AMD <u>Opteron</u> Processor 6378	4 Feb 2014	64	558	Power HS06/ LogCPU
114					64	558	
				Total	256	2146	8.38



Last 12 moths multi-  
core efficiency plot  
<http://panglia.triumf.ca>

# Sussex gridpp site compute power

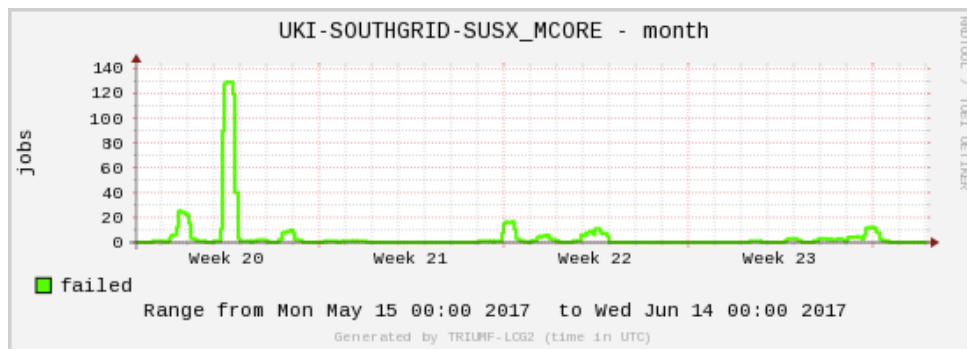
node	DELL kit	Server Model <u>PowerEdge</u>	Processor model	Delivery date	Log/ CPUs (HT on)	HS06/ Node	
101	2 x 	C6145	AMD <u>Opteron</u> Processor 6276	16 Jun 2011	64	515	
102					64	515	
113	2 x 	C6145	AMD <u>Opteron</u> Processor 6378	4 Feb 2014	64	558	
114					64	558	
210	3 x 	C6220 II	Intel E5-2650v2	8 Jan 2015	32	358	
212					32	358	
211					32	358	
216	3 x 	R430	E5-2640 v3	4 Feb 2016	32	372	
215					32	372	
214					32	372	
202	3 x 	R430	E5-2630 v4	19 Feb 2017	40	416	Power HS06/ <u>LogCPU</u>
203					40	416	
201					40	416	
				Total	568	5584	9.83

# Sussex gridpp site, multi-core queue behaviour

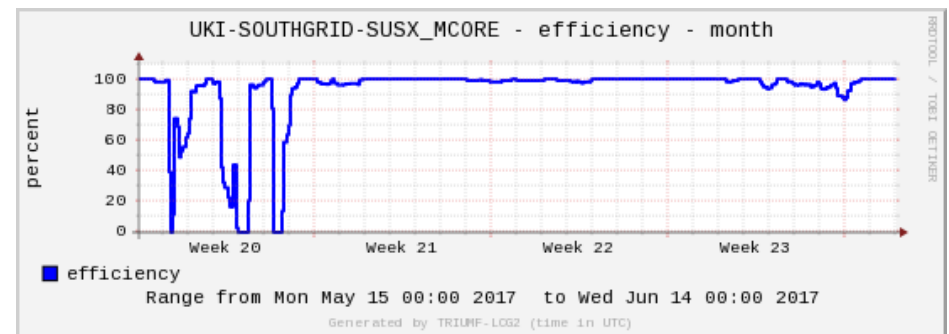


Last month number of multi-core grid jobs running at Sussex

Last month no. of failed jobs



Last month efficiency



# Sussex gridpp site, purge jobs test CE issue

State	Type / Reason	Time	% Total Time	% Known Time
OK	Unscheduled	24d 20h 38m 3s	80.193%	80.193%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	<b>Total</b>	<b>24d 20h 38m 3s</b>	<b>80.193%</b>	<b>80.193%</b>
WARNING	Unscheduled	0d 0h 0m 0s	0.000%	0.000%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	<b>Total</b>	<b>0d 0h 0m 0s</b>	<b>0.000%</b>	<b>0.000%</b>
UNKNOWN	Unscheduled	0d 0h 0m 0s	0.000%	0.000%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	<b>Total</b>	<b>0d 0h 0m 0s</b>	<b>0.000%</b>	<b>0.000%</b>
CRITICAL	Unscheduled	6d 3h 21m 57s	19.807%	19.807%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	<b>Total</b>	<b>6d 3h 21m 57s</b>	<b>19.807%</b>	<b>19.807%</b>
Undetermined	Nagios Not Running	0d 0h 0m 0s	0.000%	
	Insufficient Data	0d 0h 0m 0s	0.000%	
	<b>Total</b>	<b>0d 0h 0m 0s</b>	<b>0.000%</b>	
All	<b>Total</b>	<b>31d 0h 0m 0s</b>	<b>100.000%</b>	<b>100.000%</b>

## Service Availability Report

Last Updated: Mon Jun 12 11:32:34 CEST 2017

## Service 'emi.cream.CREAMCE-JobPurge-ops' On Host 'grid-cream-02.hpc.susx.ac.uk'

05-12-2017 11:24:17 to 06-12-2017 11:24:17

Duration: 31d 0h 0m 0s

## Service State Breakdowns:



## Service Log Entries:

[\[ View full log entries \]](#)

Event Start Time	Event End Time	Event Duration	Event/State Type	Event/State Information
05-08-2017 00:00:00	05-08-2017 00:28:54	0d 0h 28m 54s	SERVICE OK (HARD)	OK: job purged
05-08-2017 00:28:57	05-08-2017 03:29:50	0d 3h 0m 53s	SERVICE OK (HARD)	OK: job purged
05-08-2017 03:29:52	05-08-2017 06:28:57	0d 2h 59m 5s	SERVICE OK (HARD)	OK: job purged
05-08-2017 06:28:59	05-08-2017 09:30:27	0d 3h 1m 28s	SERVICE OK (HARD)	OK: job purged
05-08-2017 09:30:30	05-08-2017 12:30:46	0d 3h 0m 16s	SERVICE OK (HARD)	OK: job purged
05-08-2017 12:53:14	05-08-2017 15:30:14	0d 2h 37m 0s	SERVICE CRITICAL (HARD)	Received timeout while fetching results.

# Sussex gridpp site, SRM status

































Services of Host grid-storm-02.hpc.susx.ac.uk 29 rows /C=UK/O=eScience/OU=Sussex/L=PhysicsAndAstronomy/CN=leonardo rojas (user) 10:00  













2 30s
 Host history
 Availability


grid-storm-02.hpc.susx.ac.uk




























State	Service	Icons	Status detail	Age	Checked	Perf-O-Meter
OK	org.atlas.DDM-srm-Del-/atlas/Role=production	 	Del run successfully on all tokens!	2017-05-16 12:39:54	19 min	
OK	org.atlas.DDM-srm-DelDATA-/atlas/Role=production	 	OK: ATLASDATADISK: File deletion successful, 0:00:00.685334 s	2017-05-16 12:39:54	19 min	
OK	org.atlas.DDM-srm-DelGROUP-/atlas/Role=production	 	Token not in VOInfoDictionary, dummy test successful!	2016-06-14 12:00:04	19 min	
OK	org.atlas.DDM-srm-DelLOCALGROUP-/atlas/Role=production	 	OK: ATLASLOCALGROUPDISK: File deletion successful, 0:00:00.084133 s	2017-05-16 12:39:54	19 min	
OK	org.atlas.DDM-srm-DelSCRATCH-/atlas/Role=production	 	Token not in VOInfoDictionary, dummy test successful!	2016-06-14 12:00:04	19 min	
OK	org.atlas.DDM-srm-Get-/atlas/Role=production	 	Get run successfully on all tokens!	2017-05-16 12:27:09	70 sec	
OK	org.atlas.DDM-srm-GetDATA-/atlas/Role=production	 	OK: ATLASDATADISK: File successfully fetched	2017-05-16 12:27:09	58 sec	
OK	org.atlas.DDM-srm-GetGROUP-/atlas/Role=production	 	Token not in VOInfoDictionary, dummy test successful!	2016-06-14 12:09:51	58 sec	
OK	org.atlas.DDM-srm-GetLOCALGROUP-/atlas/Role=production	 	OK: ATLASLOCALGROUPDISK: File successfully fetched	2017-05-16 12:27:09	58 sec	
OK	org.atlas.DDM-srm-GetSCRATCH-/atlas/Role=production	 	Token not in VOInfoDictionary, dummy test successful!	2016-06-14 12:09:51	58 sec	
OK	org.atlas.DDM-srm-LsDir-/atlas/Role=production	 	LsDir run successfully on all tokens!	2017-05-16 12:27:09	58 sec	
OK	org.atlas.DDM-srm-Put-/atlas/Role=production	 	Put run successfully on all tokens!	2017-05-16 12:34:42	25 min	
OK	org.atlas.DDM-srm-PutDATA-/atlas/Role=production	 	OK: ATLASDATADISK: File copy successful, 0:00:01.318824 s	2017-06-05 18:04:57	24 min	
OK	org.atlas.DDM-srm-PutGROUP-/atlas/Role=production	 	Token not in VOInfoDictionary, dummy test successful!	2016-06-14 11:50:06	24 min	
OK	org.atlas.DDM-srm-PutLOCALGROUP-/atlas/Role=production	 	OK: ATLASLOCALGROUPDISK: File copy successful, 0:00:02.339138 s	2017-05-16 12:34:42	24 min	
OK	org.atlas.DDM-srm-PutSCRATCH-/atlas/Role=production	 	Token not in VOInfoDictionary, dummy test successful!	2016-06-14 11:50:06	24 min	





# Sussex gridpp site, Webdav frontend status

Services of Host `grid-storm-02.hpc.susx.ac.uk` 29 rows /C=UK/O=eScience/OU=Sussex/L=PhysicsAndAstronomy/CN=leonardo rojas (user) 10:00  

    2 30s  Host history  Availability ...

grid-storm-02.hpc.susx.ac.uk

OK	<a href="#">webdav</a> .HTTP-All-/atlas/Role=production	  	OK: the HTTP interface is usable	2017-05-24 19:23:42	39 min	
OK	<a href="#">webdav</a> .HTTP-DIR_GET-/atlas/Role=production	 	OK: GET on dir	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-DIR_HEAD-/atlas/Role=production	 	OK: HEAD on dir	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_DELETE-/atlas/Role=production	 	OK: DELETE on file	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_DELETE_ON_NON_EXISTENT-/atlas/Role=production	 	OK: DELETE on a file that does not exist, should return 404	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_GET-/atlas/Role=production	 	OK: GET file	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_HEAD-/atlas/Role=production	 	OK: HEAD on file	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_HEAD_ON_NON_EXISTENT-/atlas/Role=production	 	OK: HEAD on a path that does not exist, should return 404	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_MOVE-/atlas/Role=production	 	OK: MOVE file to another path	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_OPTIONS-/atlas/Role=production	 	OK: OPTIONS on file	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_PROPFIND-/atlas/Role=production	 	OK: PROPFIND on file	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-FILE_PUT-/atlas/Role=production	 	OK: PUT a file	2017-05-24 19:24:20	38 min	
OK	<a href="#">webdav</a> .HTTP-TLS_CIPHERS-/atlas/Role=production	 	OK: Verifies the server provides a reasonable set of TLS ciphers	2017-05-24 19:24:20	38 min	

refresh: 30 secs

# Sussex gridpp site, to do list

- All our service nodes are VMs, running on KVM environments.
- Grid-cream-02, grid-argus-02, grid-bdii-02 and grid-apel-02 are running on a old host which need to be decommissioned, so they need to be migrated to a new host server.
- Once the two cluster are consolidated, grid service nodes will have 10Gb/s connectivity to Janet.
- Publishing is not working right now. Cream was updated to advertise the new kit, but rebus system has not taken this into account, first it showed old values, now is showing 0 on every field.
- I need to separate a slot exclusively for ops jobs to see if purge tests stop flapping

**Thanks!**