

ScotGrid Site Report

Mike Kenyon
HEPSysMan - June 2008

Facts 'n' Figures.

The ScotGrid Team

- Glasgow

- Graeme Stewart - Technical Co-ordinator.
- Andrew Elwell - Data Management Expert.
- Mike Kenyon - Glasgow Grid System Manager.

- Edinburgh Compute and Data Facility

- Sam Skipsey.
- Steve Thorn.
- Greig Cowan.

- Durham

- Phil Roffe.
- David Ambrose Pearce.

Hardware

- Glasgow

- 140 Worker Nodes: 2 x Dual Core Opteron 280 CPUs = 560 Job slots.
- 9 Grid servers: 85TB of useable SRM storage (RAID 6 + 1 “hot spare”/server).

- ECDF

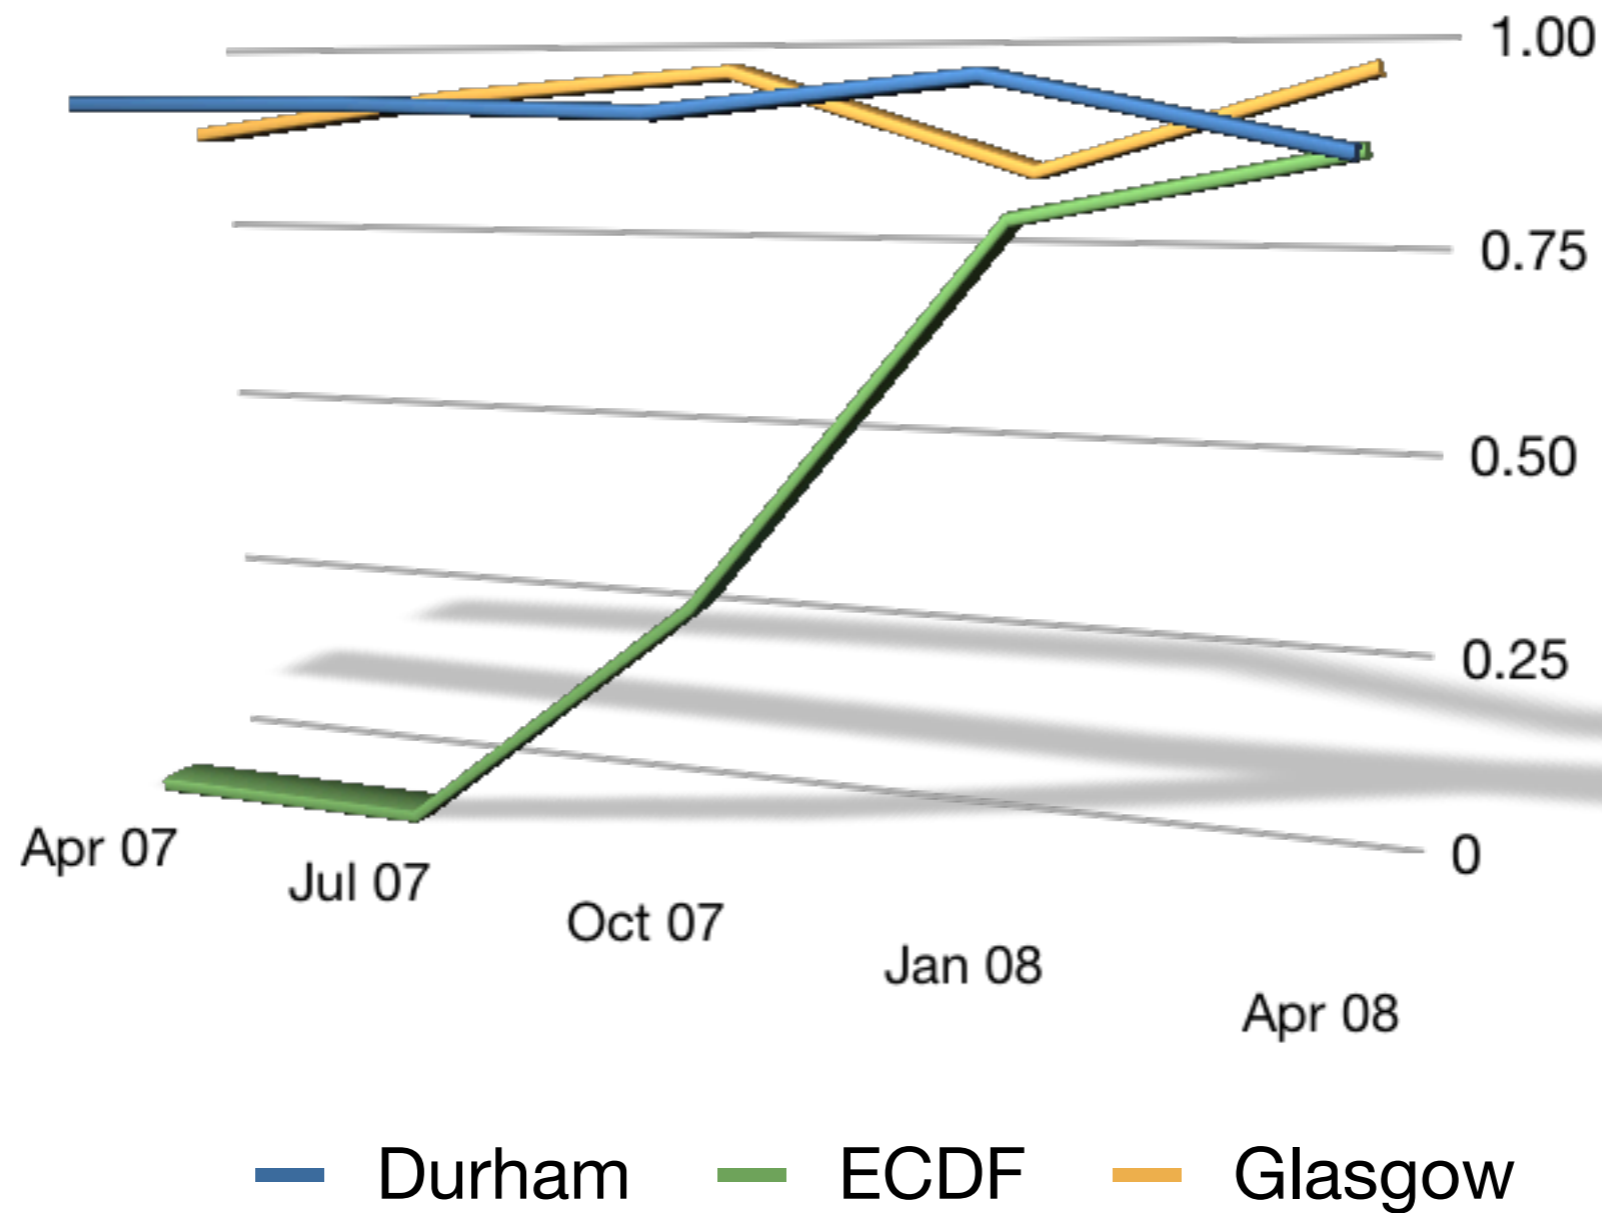
- 230 Worker Nodes: Mix of 2- and 4-CPU 64bit Intels ~1500 Job slots.
- 5TB of useable SRM storage.

- Durham

- 53 Worker Nodes: Pentium 2.0/2.4/2.6 GHz = 106 Job slots.
- 12TB of useable SRM storage.

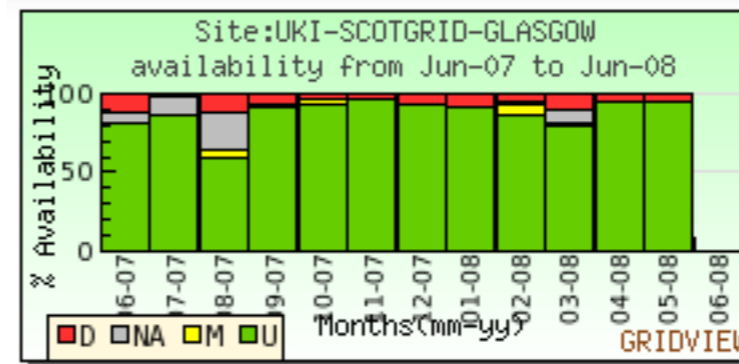
Site Availability

(Steve Lloyd's Tests)

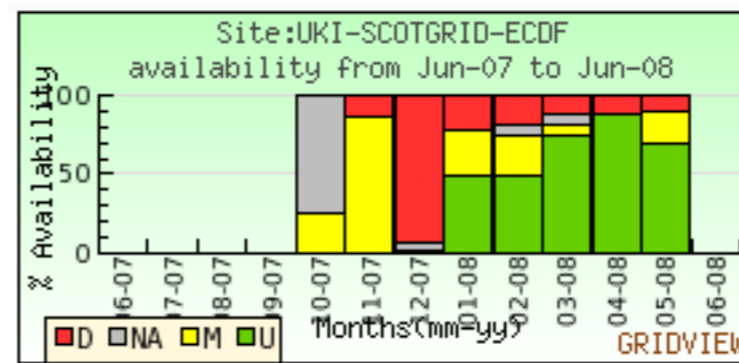


GridView “Overall” Availability

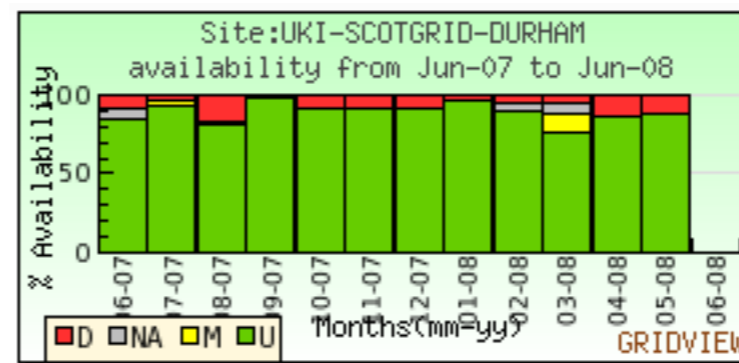
Glasgow



ECDF

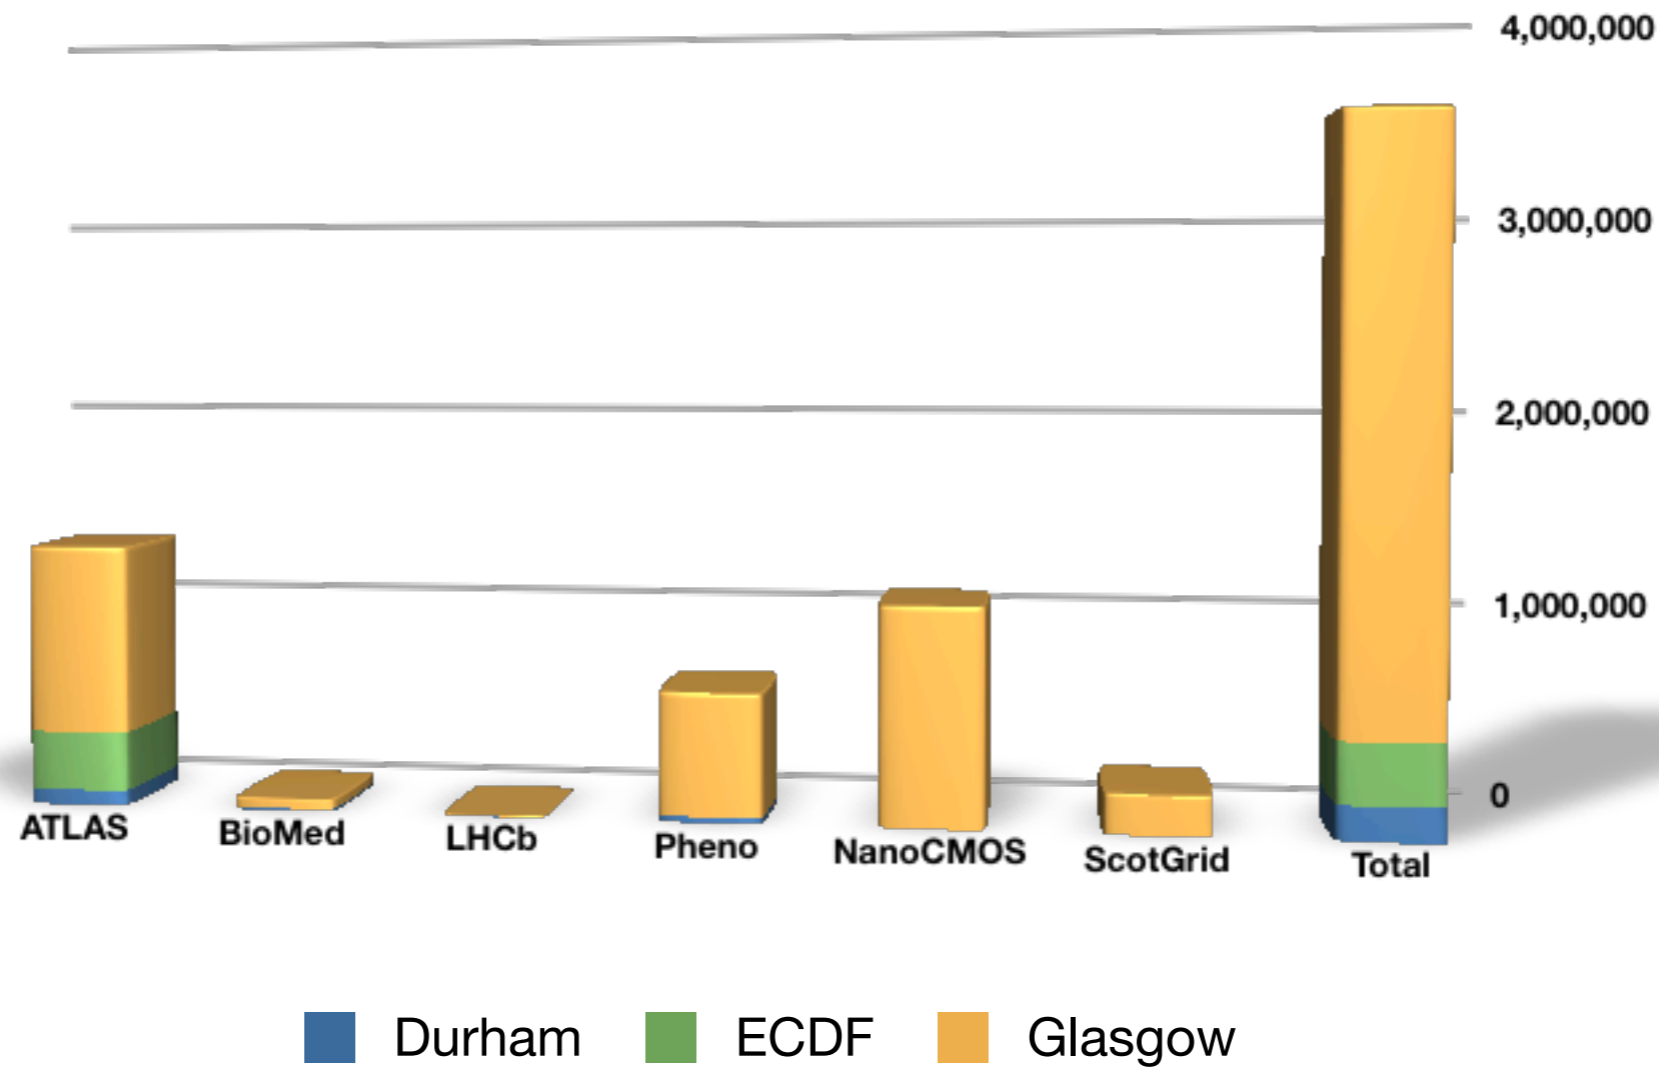


Durham



CPU Used

Normalised CPU time per VO (K SI2000 Hours YTD)



Data Storage

- Disk pool manager: SRM v2.2 across ScotGrid.
- Glasgow: Migrated to xfs now that x86_64 is supported by DPM.
- ScotGrid now has space tokens.
- Reliability excellent: 0 failures over 12 months for 220 Western Digital 0.5TB

WLCG Nagios Monitoring

<https://twiki.cern.ch/twiki/bin/view/LCG/>

- Prototype Nagios Grid Service monitoring.
- Installed as per TWIKI instructions.
- “Remote & Local” installation.
 - Results from SAM tests & local tests are monitored.
- Biggest problem was getting the MyProxy proxy working.
 - Caused by [bug](#) in the old MyProxy code that runs on myproxy-fts.cern.ch.

Tactical Overview

Tactical Monitoring Overview

Last Updated: Wed Jun 18 00:22:54 BST 2008
 Updated every 90 seconds
 Nagios® 3.0.1 - www.nagios.org
 Logged in as /C=UK/O=eScience/OU=Glasgow/L=CompServ/CN=michael kenyon

Monitoring Performance

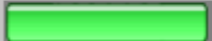

Service Check Execution Time: 0.01 / 60.01 / 0.392 sec
 Service Check Latency: 0.00 / 1.50 / 0.287 sec
 Host Check Execution Time: 0.01 / 3.02 / 0.065 sec
 Host Check Latency: 0.00 / 6.34 / 2.042 sec
 # Active Host / Service Checks: 205 / 1801
 # Passive Host / Service Checks: 0 / 260



Network Outages

0 Outages

Network Health

Host Health: 
 Service Health: 

Hosts

3 Down 0 Unreachable 202 Up 0 Pending

1 Unhandled Problems
 2 Acknowledged

Services

45 Critical 2 Warning 13 Unknown 2001 Ok 0 Pending

16 Unhandled Problems 2 Disabled 7 Unhandled Problems 236 Disabled
 6 Disabled
 12 on Problem Hosts
 1 Acknowledged
 16 Disabled

Tactical Overview

Tactical Monitoring Overview
 Last Updated: Wed Jun 18 00:22:54 BST 2008
 Updated every 90 seconds
 Nagios® 3.0.1 - www.nagios.org
 Logged in as /C=UK/O=eScience/OU=Glasgow/L=CompServ/CN=michael.kenyon

Current Network Status
 Last Updated: Wed Jun 18 00:25:33 BST 2008
 Updated every 90 seconds
 Nagios® 3.0.1 - www.nagios.org
 Logged in as /C=UK/O=eScience/OU=Glasgow/L=CompServ/CN=michael.kenyon

[View History For all hosts](#)
[View Notifications For All Hosts](#)
[View Host Status Detail For All Hosts](#)

Host Status Totals

Up	Down	Unreachable	Pending
202	3	0	0

All Problems | **All Types**

3	205
---	-----

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
2000	2	13	46	0

All Problems | **All Types**

61	2061
----	------

Display Filters:
 Host Status: Pending | Up
 Types:
 Host Properties: Any
 Service Status: Critical
 Types:
 Service: Not In Scheduled Downtime & Has Not Been
 Properties: Acknowledged & Active Checks Enabled

Service Status Details For All Hosts


Host	Service	Status	Last Check	Duration	Attempt	Status Information
ce.glite.ecdf.ed.ac.uk	org.nagios.LocalLogger-PortCheck	CRITICAL	06-18-2008 00:24:11	29d 14h 33m 8s	4/4	(Service Check Timed Out)
ce01.dur.scotgrid.ac.uk	hr.srce.CAdist-Version	CRITICAL	06-17-2008 21:01:52	6d 9h 23m 40s	4/4	Official IGTF version is 1.22. Old CA distribution version found, new version is 8 days old.
	hr.srce.GRAM-Auth	CRITICAL	06-18-2008 00:14:05	1d 7h 41m 27s	4/4	(Service Check Timed Out)
	hr.srce.GRAM-CertLifetime	CRITICAL	06-17-2008 16:59:42	0d 10h 25m 50s	4/4	(Service Check Timed Out)
	hr.srce.GRAM-Command	CRITICAL	06-17-2008 23:44:43	6d 9h 10m 49s	4/4	(Service Check Timed Out)
	org.nagios.BDII-Check	CRITICAL	06-18-2008 00:11:03	1d 7h 59m 29s	4/4	Could not search/find objectclasses in GlueSiteUniqueID=UKI-SCOTGRID-DURHAM,Mds-Vo-Name=UKI-SCOTGRID-DURHAM,O=Grid
	org.nagios.LocalLogger-PortCheck	CRITICAL	06-18-2008 00:21:04	29d 14h 33m 7s	4/4	(Service Check Timed Out)
gallows.dur.scotgrid.ac.uk	hr.srce.CAdist-Version	CRITICAL	06-17-2008 20:00:23	6d 9h 25m 9s	4/4	Official IGTF version is 1.22. Old CA distribution version found, new version is 8 days old.
	hr.srce.GridFTP-Transfer	CRITICAL	06-17-2008 23:50:31	6d 9h 5m 1s	4/4	Upload to remote computer failed failed with error: globus_ftp_client: the server responded with an error500 Command failed. : open/create error: No space left on device



Sites' Headline News.


Glasgow CE Woes

- In February we suffered a failure of the CE.
- At 17.00: load started to rise to ~100.



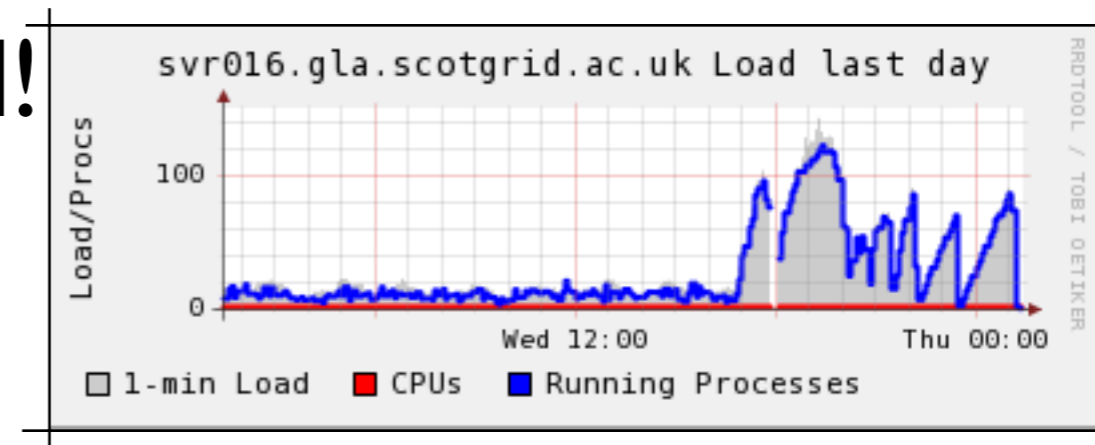
Glasgow CE Woes

- In February we suffered a failure of the CE.
- At 17.00: load started to rise to ~100.
- At 18.00: hard crash.



Glasgow CE Woes

- In February we suffered a failure of the CE.
- At 17.00: load started to rise to ~100.
- At 18.00: hard crash.
- At 22.00: defeat accepted!



“Every cloud...”

- Took the opportunity to bring March’s downtime forward.
- CE & DPM headnode moved to SL4 x8_64.
 - SL4 CE appears to scale much better.
- Batch system upgraded to Torque 2.1.9/
Maui 3.2.6, running on SL4 x86_64.
- Queues reduced to 4, with queue lengths of 30m, 6h, 3d and 7d.

ECDF

- ECDF now online & validated.
- SAM tests were being killed by the batch system.
 - Problems due to SGE “[virtual size limit](#)” on jobs.
- ATLAS software validation kit couldn’t “phone home”.
 - Outbound port 443 was firewalled!

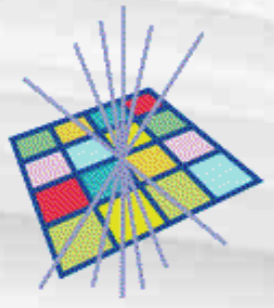
Durham

- March 08: Upgraded CE, SE, UI & WNs to SL4.6.
- [cfengine'd](#): installs packages, config files & runs YAIM.
 - Had previously used “home-brew” installation scripts.
- Conclusion: cfengine is worth the time & effort of installation!

New/Improved Services.

ScotGrid wiki

<http://www.scotgrid.ac.uk/wiki>



navigation

- [Main Page](#)
- [Community portal](#)
- [Current events](#)
- [Recent changes](#)
- [Random page](#)
- [Help](#)

search

toolbox

- [What links here](#)
- [Related changes](#)
- [Special pages](#)

[article](#) [discussion](#) [edit](#) [history](#)

Main Page

Welcome to the ScotGrid Wiki
ScotGrid [↗](#): a member of [EGEE](#) [↗](#), [GridPP](#) [↗](#) and [NGS](#) [↗](#).

- [NorthGrid](#) [↗](#)
- [SouthGrid](#) [↗](#)
- [London Tier2](#) [↗](#)

General Information for Users

These links detail how to get access to the ScotGrid clusters and how to start submitting your first grid jobs.

- [User Guide](#)

Information for System Administrators


- [Cluster Administration](#)

Site Specific Admin Info

- [Durham](#)
- [Edinburgh](#)
- [Glasgow Logbook I Dev Boxes](#)

Other ScotGrid News and Links

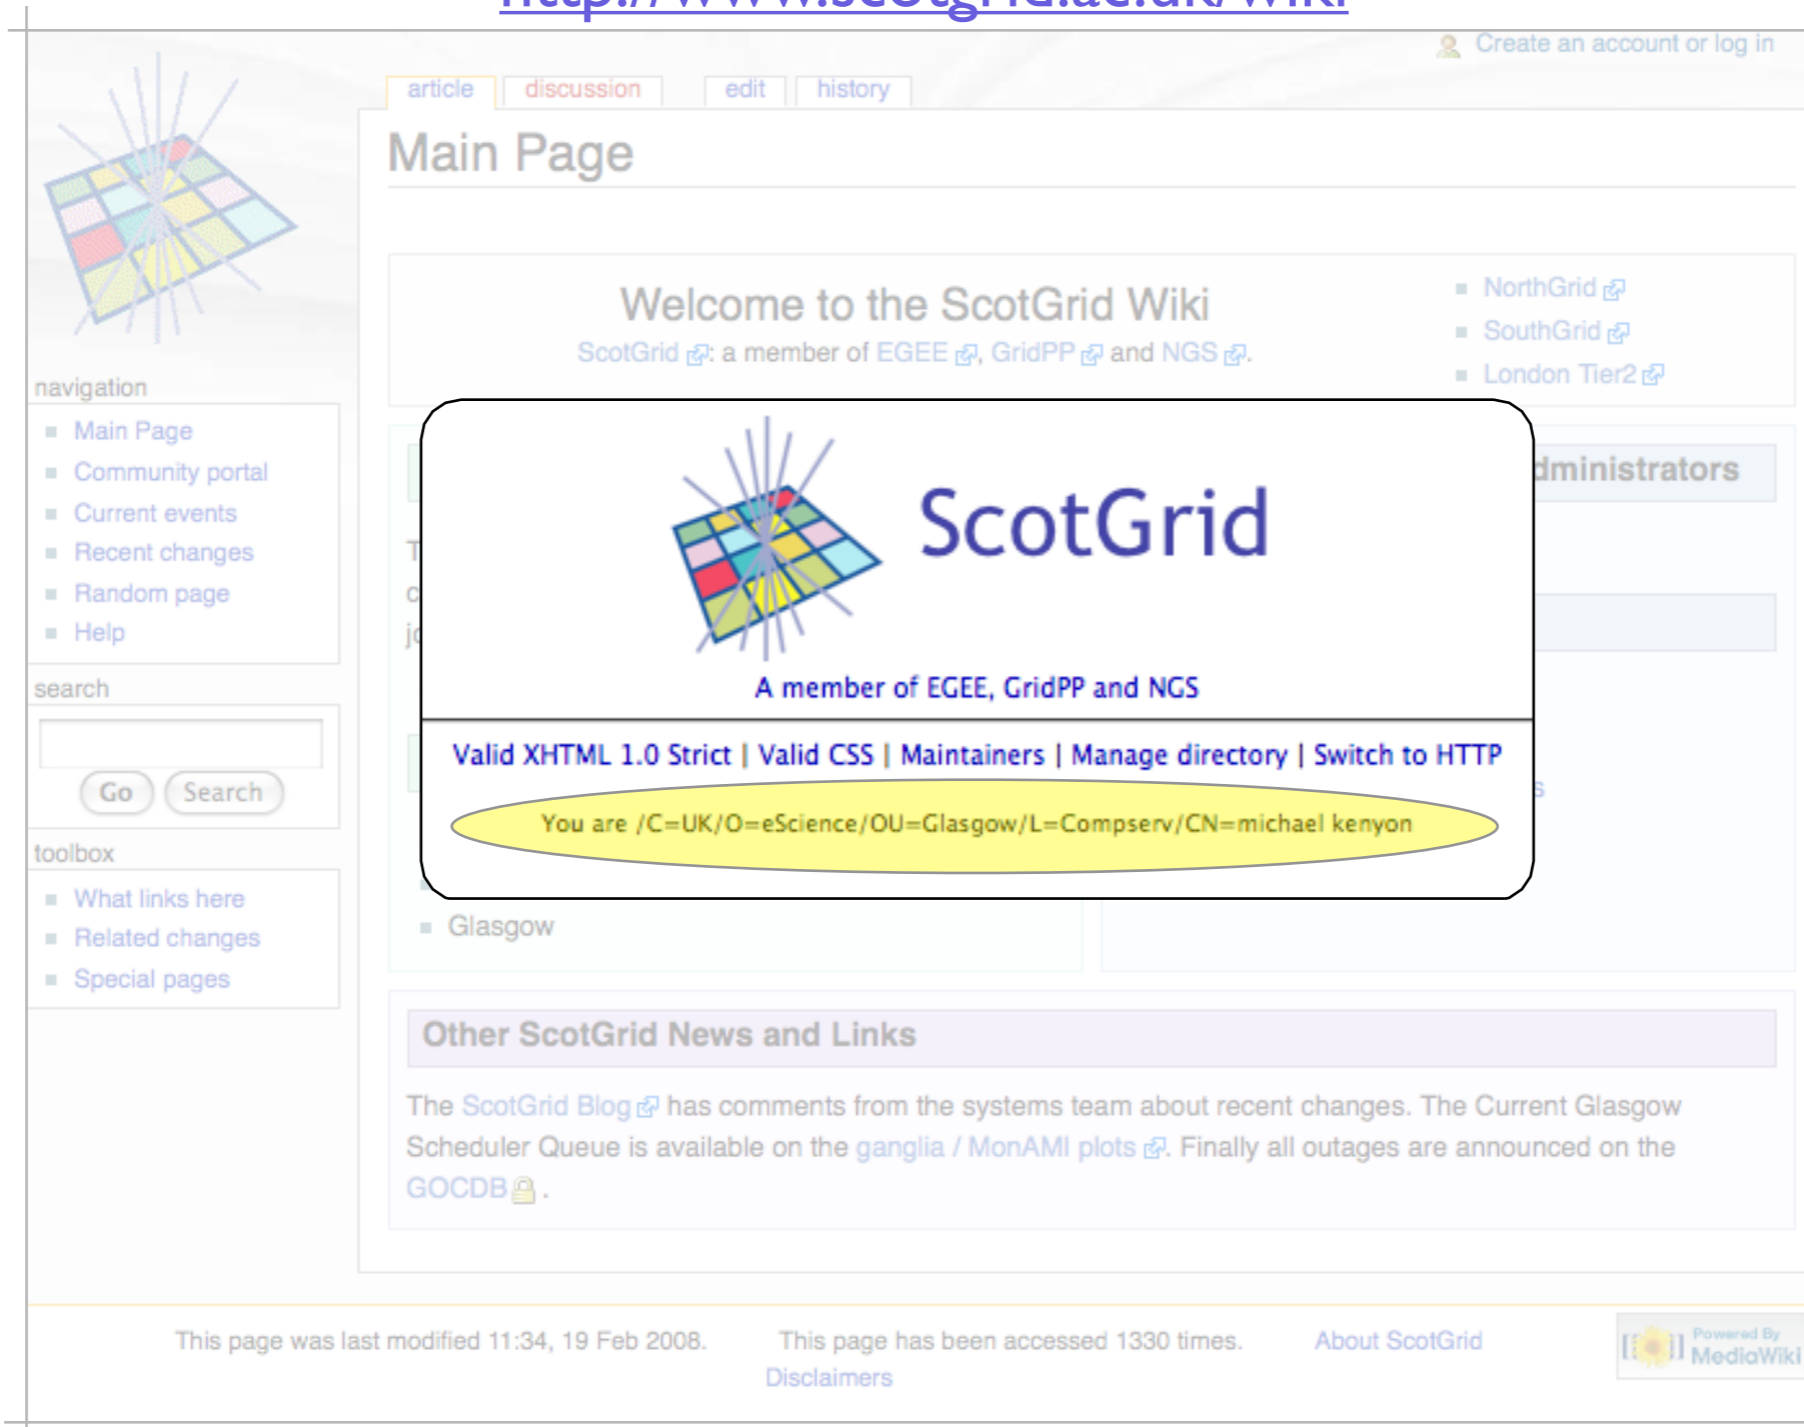
The [ScotGrid Blog](#) [↗](#) has comments from the systems team about recent changes. The Current Glasgow Scheduler Queue is available on the [ganglia / MonAMI plots](#) [↗](#). Finally all outages are announced on the [GOCDB](#) [↗](#).

This page was last modified 11:34, 19 Feb 2008. This page has been accessed 1330 times. [About ScotGrid](#)  Powered By MediaWiki

[Disclaimers](#)

ScotGrid wiki

<http://www.scotgrid.ac.uk/wiki>



The screenshot shows the main page of the ScotGrid wiki. At the top right, there is a link to "Create an account or log in". Below this are tabs for "article", "discussion", "edit", and "history". The main heading is "Main Page".

On the left side, there is a navigation menu with the following items:

- Main Page
- Community portal
- Current events
- Recent changes
- Random page
- Help

Below the navigation menu is a search box with "Go" and "Search" buttons. Underneath is a toolbox with the following items:

- What links here
- Related changes
- Special pages

The main content area features a "Welcome to the ScotGrid Wiki" message, stating that ScotGrid is a member of EGEE, GridPP, and NGS. To the right of this message are links to "NorthGrid", "SouthGrid", and "London Tier2".

Below the welcome message is a large box containing the ScotGrid logo and the text "ScotGrid" and "A member of EGEE, GridPP and NGS". Below this box are links for "Valid XHTML 1.0 Strict", "Valid CSS", "Maintainers", "Manage directory", and "Switch to HTTP". A yellow oval highlights the current user's path: "You are /C=UK/O=eScience/OU=Glasgow/L=Comperv/CN=michael kenyon".

Below the yellow oval is a list of "Administrators" with a link to "Glasgow".

At the bottom of the main content area is a section titled "Other ScotGrid News and Links" with the following text: "The ScotGrid Blog has comments from the systems team about recent changes. The Current Glasgow Scheduler Queue is available on the ganglia / MonAMI plots. Finally all outages are announced on the GOCDB."

The footer contains the following information:

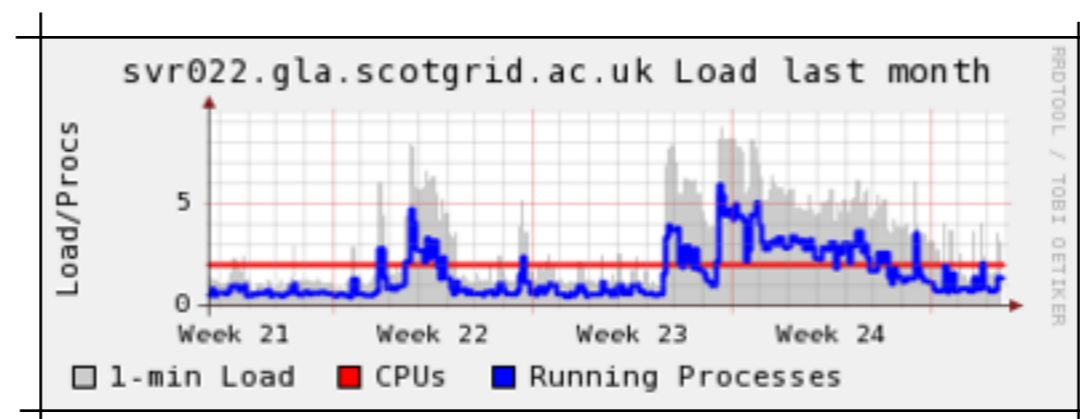
- This page was last modified 11:34, 19 Feb 2008.
- This page has been accessed 1330 times.
- About ScotGrid
- Disclaimers
- Powered By MediaWiki

Glasgow gLite WMSLB

- Combined gLite 3.0 WMS/LB on same (SL3) machine.
- [Installation](#) was relatively pain-free.
 - gLite user must exist.
 - gLite-WMSLB meta-package does not install the gLite-yaim-lb package.

Glasgow gLite WMSLB

- Combined gLite 3.0 WMS/LB on same (SL3) machine.
- [Installation](#) was relatively pain-free.
- However, have observed very high loads.



Glasgow gLite WMSLB

- High load killed WMS in early June.
- Pheno user:
 - Submitted ~20,000 jobs.
 - Without proxy renewal.
 - Couldn't kill their jobs!
- Reverted back to RB.
- Considering an additional WMS.

Our poor WMS was killed from the end of last week when a pheno user submitted about 20k jobs into it. Worse, they hadn't used a proxy renewal service, so their VOMS extension on the proxy expired and so the jobs they had submitted suffered shallow failures, prompting further resubmission attempts and further load.

We have contacted the user, but they were having a great deal of trouble even cancelling the jobs. We've left things for 2 days now, but the situation is really not improving. I can't see much hope for the machine in its current state - we'll probably have to blow it away and start again next week.

For the moment we've asked our users to revert to using ye olde RB, which we were just about to switch off but hadn't actually decommissioned.

There's a serious question now about what level of WMS service we want to provide. It's a complex service and rather difficult to debug when it goes wrong. Certainly an upgrade to the SL4 version should be done, but do we want to have 2 WMS hosts and possibly even a separate LB service?

POSTED BY GRAEME STEWART AT 11:11 AM 0 COMMENTS
LABELS: UKI-SCOTGRID-GLASGOW, WMS

Labels: UKI-SCOTGRID-GLASGOW, WMS
Posted by Graeme Stewart at 11:11 AM 0 comments

Labels: UKI-SCOTGRID-GLASGOW, WMS

Labels: UKI-SCOTGRID-GLASGOW, WMS
Posted by Graeme Stewart at 11:11 AM 0 comments

ScotGrid VOMS Server

- VOMS server [installed](#) at Glasgow.
- Enables us to quickly respond to local users' needs.
- vo.scotgrid.ac.uk, vo.nanocmos.ac.uk,
vo.panda.gsi.de



The screenshot shows a web interface for the VOMS server. It features a header with a VOMS icon and the title "List of VOs on this server". Below the header, there is a list of VOs with their names and base URLs for the web interface.

VO name	Base URL for the web interface
vo.scotgrid.ac.uk	https://svr029.qla.scotgrid.ac.uk:8443/voms/vo.scotgrid.ac.uk
vo.nanocmos.ac.uk	https://svr029.qla.scotgrid.ac.uk:8443/voms/vo.nanocmos.ac.uk
vo.panda.gsi.de	https://svr029.qla.scotgrid.ac.uk:8443/voms/vo.panda.gsi.de

ScotGrid-Glasgow

Tier 2.5



- Interface local (PPE) resources with Grid.
- Local-level access for data, job preparation.
- ScotGrid's DPM ATLAS datasets available via RFIO.
- gLite 3.1 UI available on all Linux desktops.
- Removes need for ScotGrid username (e.g., gla046 \Rightarrow mkenyon).

And Finally...

Vendor Hardware Support

- ClusterVision
 - “...little more than replacing faulty/suspect components.”
 - “...no real support in the sense of problem solving.”
 - “...didn’t even supply PC Doctor or similar.”

Procurement '08

- Tendering for ~ £400,000 of CPU & disk.
- Aiming to bring ScotGrid up to:
 - ~500TB of storage.
 - ~2,000 KSpecInt2000 of CPU.
- Targeting Q4 2008 for switch on.
- Machine-room work to start next week.

Glasgow Machine Room II

