

GridPP

UK Computing for Particle Physics

End User Analysis Support: A Site's Perspective

Mike Kenyon

WLCG Workshop - Imperial College, London


07-07-10




University
of Glasgow

- Context: Where do we fit into the Grand Scheme?
- Supporting our users: how we do this.
 - Site availability monitoring.
 - Documentation & communication.
 - Common support requests.
- “Our” observations/experiences of the VO support mechanisms.
 - GGUS.
 - The Ganga Blacklist.
- General observations.

- Admin team
 - 2 x Technical coordinators.
 - 1 x Systems manager.
 - 2 x GridPP UK Tier-2 data-management system managers (!)
 - 0.5 x User support/grid tool developer.
- Hardware (by end of 2010)
 - 2100 job slots (~20,000 HEPSPEC06).
 - 1.4 PB of storage.
- Local users
 - ~20 cluster users.
- Grid users
 - 400 unique DNs in 2010.
 - /CN=graeme stewart by far the biggest user!





- How do we, as admins proactively support the users?

- Monitoring

- SAM (future? Active MQ)
- Experiment dashboards
- Panglia (PanDA + ganglia)

The Glasgow Dashboards

BDII tests

Status of Last 10 Queries Time of Last Query Overall

BDII 9 8 7 6 5 4 3 2 1 0 Submitted Successful Time Status

LondonGrid 8 8 8 8 8 8 8 8 8 8 08/06/10 14:30 08/06/10 14:30 5:15:35 Good

WMS tests

Status of Last 10 Jobs Time of Last Job Average Time Taken (sec) Overall

Resource Broker Pri 9 8 7 6 5 4 3 2 1 0 Submitted Successful Submit Status Output Total Status

Scot glite WMS 1 4 5 5 5 5 5 5 5 5 5 C C C C 08/06/10 14:30 08/06/10 14:05 4.100 0.100 8.000 2018 Fair

Scot glite WMS 2 5 5 5 5 5 5 5 5 5 5 C C C C 08/06/10 14:30 08/06/10 14:20 3.700 0.500 9.625 2004 Fair

SAM tests

CE	SRIW2	Availability
Hours ago: 9 8 7 6 5 4 3 2 1 0	9 8 7 6 5 4 3 2 1 0	24 Hrs Week Month 6 Mon
UK-SCOTGRID-GLASGOW	████████████████████	100% 100% 97% 98%

ATLAS SAM tests

CE	SRIW2	Availability
Hours ago: 9 8 7 6 5 4 3 2 1 0	9 8 7 6 5 4 3 2 1 0	24 Hrs Week Month 6 Mon
UK-SCOTGRID-GLASGOW	████████████████████	100% 99% 99% 99%

CMS SAM tests

CE	SRIW2	Availability
Hours ago: 9 8 7 6 5 4 3 2 1 0	9 8 7 6 5 4 3 2 1 0	24 Hrs Week Month 6 Mon
UK-SCOTGRID-GLASGOW	████████████████████	100% 100% 97% 98%

LHCb SAM tests

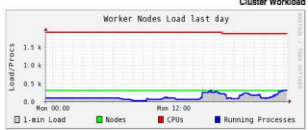
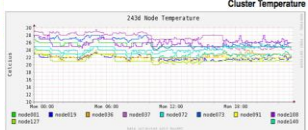



CE	Availability	
Hours ago: 9 8 7 6 5 4 3 2 1 0	24 Hrs Week Month 6 Mon	
UK-SCOTGRID-GLASGOW	████████████████████	100% 100%





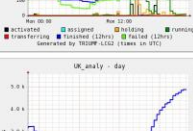
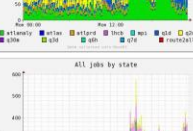
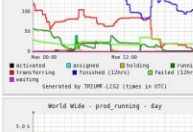
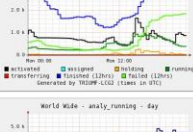

Godashboard Site Status for CMS

Site name	Visible	Jobs	SAM TESTS	Production	Analysis	Site usage	Phylog	Commissioning Links	Produce	Under investigation	Subsites	Supporting CMS
UK-SCOTGRID-GLA	Yes	174	Yes	Yes	Yes	████████████████████	████████████████████	████████████████████	████████████████████	████████████████████	████████████████████	████████████████████

Janet Status

RecordStart	RecordEnd	Record Count in Your Database	RecordCount What You Published	Synchronization Status
2010-07-01	2010-07-05	100087	100087	OK (last published 0 days ago: 2010-07-05)
2010-06-01	2010-06-29	548438	548438	OK
2010-05-01	2010-05-31	512027	512021	OK (minor discrepancy in event numbers)

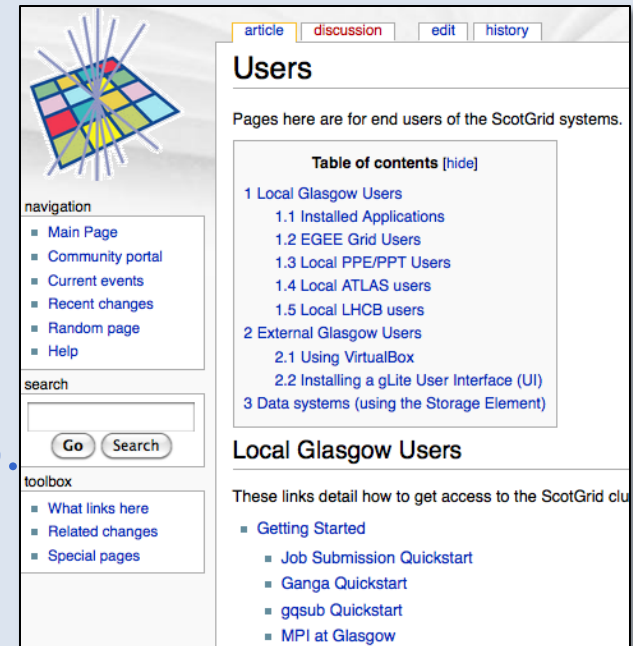
Queue	R	P	C	W	U	Y	Unknown
atlas20	0	0	0	0	1	0	0
atlas21	0	0	0	0	3	0	0
atlas24	0	0	0	0	3	0	0
atlas26	1	0	0	0	0	0	0
atlas27	2	0	0	0	0	0	0
atlas28	1	0	0	0	0	0	0
atlas29	0	0	0	0	1	0	0
atlas30	0	0	0	0	1	0	0
atlas31	0	0	0	0	0	0	0
atlas32	1	0	0	0	0	0	0
atlas33	0	0	0	0	0	0	0
atlas34	0	0	0	0	0	0	0
atlas35	0	0	0	0	0	0	0
atlas36	0	0	0	0	0	0	0
atlas37	201	0	0	0	0	0	0
atlas38	0	0	0	0	0	0	0
atlas39	1	0	0	0	0	0	0
atlas40	1	0	0	0	0	0	0
atlas41	0	0	0	0	0	0	0
atlas42	20	0	0	0	0	0	0
atlas43	1	0	0	0	0	0	0
atlas44	1	0	0	0	0	0	0
atlas45	1	0	0	0	0	0	0
atlas46	1	0	0	0	0	0	0
atlas47	1	0	0	0	0	0	0
atlas48	1	0	0	0	0	0	0
atlas49	1	0	0	0	0	0	0
atlas50	1	0	0	0	0	0	0
atlas51	1	0	0	0	0	0	0
atlas52	1	0	0	0	0	0	0
atlas53	1	0	0	0	0	0	0
atlas54	1	0	0	0	0	0	0
atlas55	1	0	0	0	0	0	0
atlas56	1	0	0	0	0	0	0
atlas57	1	0	0	0	0	0	0
atlas58	1	0	0	0	0	0	0
atlas59	1	0	0	0	0	0	0
atlas60	1	0	0	0	0	0	0
atlas61	1	0	0	0	0	0	0
atlas62	1	0	0	0	0	0	0
atlas63	1	0	0	0	0	0	0
atlas64	1	0	0	0	0	0	0
atlas65	1	0	0	0	0	0	0
atlas66	1	0	0	0	0	0	0
atlas67	1	0	0	0	0	0	0
atlas68	1	0	0	0	0	0	0
atlas69	1	0	0	0	0	0	0
atlas70	1	0	0	0	0	0	0
atlas71	1	0	0	0	0	0	0
atlas72	1	0	0	0	0	0	0
atlas73	1	0	0	0	0	0	0
atlas74	1	0	0	0	0	0	0
atlas75	1	0	0	0	0	0	0
atlas76	1	0	0	0	0	0	0
atlas77	1	0	0	0	0	0	0
atlas78	1	0	0	0	0	0	0
atlas79	1	0	0	0	0	0	0
atlas80	1	0	0	0	0	0	0

- How do we, as admins proactively support the users?
 - Monitoring
 - SAM (future? Active MQ)
 - Experiment dashboards
 - Panglia (PanDA + ganglia)
- } The Glasgow Dashboards
- However
 - Though ATLAS (CMS to some extent) are covered, we have a lack of job-status information for most other VOs.
 - E.g. Undetected LHCb upload problem at some UK sites for months.
 - Sites had no idea until the VO notified us.
 - Diagnosis was hampered by lack of access to failed job logs.
 - Biomed/Camont etc: We'd have no idea if all their jobs were failing on the cluster!

- Questions:
 - What is the post-SAM future? Migrate to WLCG-Nagios.
 - How long will the SAM tests be around?
 - Will WLCG-Nagios replicate SAM tests?
 - Will the UK's Steve Lloyd pages be updated?
 - Could it be extended to non-WLCG VOs?
 - Can admins have single “go-to” page with pointers to all the *critical* monitoring interfaces?

- Most common support requests/complaints:
 - Why won't my job run? It did last week.
 - If the job ran locally, we have more resources to investigate.
 - ATLAS: we can use the PanDA monitor.
 - Otherwise, we direct user to DAST (which we monitor for tips).
 - LHCb: DIRAC job monitoring locked down to the user?
 - Is ATHENA release X.Y.Z installed?
 - ATLAS INFN software pages.
 - Other VOs - software tags.
 - How can I “download” several 100GB of data to run over?
 - Why do you want to do that? (-> DATRI)
 - How do I do <something> with Ganga?
 - Why has feature X changed/disappeared/“broken” in Ganga?

- We maintain a comprehensive, Glasgow-centric and multi-VO WIKI.
 - Aims to be a single hub for users.
 - But I also maintain an extensive list of VO FAQs, HowTo, Tutorials etc.
 - We try to keep contents as current & accurate as possible.
 - Encourage users to contribute.
- Keep users informed of upgrades/changes/service status.
 - There's a balance. E.g. Ganga bugfix releases now installed silently.
(Manage the perception of a “broken” Grid).



The screenshot shows a Wiki page titled "Users" with tabs for "article", "discussion", "edit", and "history". The page content includes:

- Navigation menu: Main Page, Community portal, Current events, Recent changes, Random page, Help.
- Search box with "Go" and "Search" buttons.
- Table of contents (Table of contents [hide]):
 - 1 Local Glasgow Users
 - 1.1 Installed Applications
 - 1.2 EGEE Grid Users
 - 1.3 Local PPE/PPT Users
 - 1.4 Local ATLAS users
 - 1.5 Local LHCb users
 - 2 External Glasgow Users
 - 2.1 Using VirtualBox
 - 2.2 Installing a gLite User Interface (UI)
 - 3 Data systems (using the Storage Element)
- Section: Local Glasgow Users
 - These links detail how to get access to the ScotGrid clu
 - Getting Started
 - Job Submission Quickstart
 - Ganga Quickstart
 - gqsub Quickstart
 - MPI at Glasgow

- Polled the GridPP admin community in advance of today.
 - SITE #1: Only feels “fully connected” with one VO: ATLAS
 - SITE #1: Doesn’t know whether WMS-submitted jobs are failing.
 - SITE #1: Camont/Biomed: how do I check jobs/contact users?
 - SITE #2: “The Thursday ATLAS meetings are informative - and a good way of raising issues. Graeme does a good job of representing the ATLAS position to us.”
 - SITE #2: “...rather than solving problems with the grid middleware, the VOs bodge around them...pilot jobs are a bodge around inadequacies in the monitoring of worker nodes and primarily the CE software.”
 - SITE #3: “Not at all” aware of the communication/support channels within VOs.
 - SITE #3: “Every PhD student that starts complains of difficulties in starting working with the Grid.”

- SITE #4: dteam meeting provides a regular source of information for all the VOs and the opportunity to ask questions direct to VOs.
- SITE #4: If our site is not running jobs correctly - or banned for some reason - it would be good to get this pushed to the site at the earliest instance.
- SITE #4: ...unclear who to contact in specific cases (e.g software installation/compatibility) so a “who to contact in case of...” contact list.
- SITE #4: “VO-side information (grid submission and monitoring portals) are either not well advertised or are restricted.”

- Not strictly WLCG, but they consume our/your resources.
- Users are not necessarily Unix-aware, certificate-aware, Grid-aware, middleware-aware!
- Cases in point at Glasgow.
 - vo.ssp.ac.uk: 1 local user running MPI Castepp jobs.
 - vo.optics.ac.uk: 1 local user running Lumerical jobs.
- glite-wms-job-submit is too heavy weight for these users, who generally want direct submission to a cluster.
 - gqsub: Wrappers around the glite-wms-* commands.
 - Support for running across the Grid & Myproxy support.
 - Grid-run-castepp: Creates JDL for MPI job and submits to Cluster.
- Remote VOs are harder to contact and service monitoring minimal/non-existent.

- “GGUS tickets to VOs don’t seem to work very well”.
 - 3rd Dec 2009. Site announces via EGEE Broadcast and then GGUS #53843 that an SE was to be decommissioned:
 - “I have attached a list of SURLs of affected files. Please can you replicate the files elsewhere (or even better delete them if no longer required) and let me know you have done so.”
 - Changed concerned VO to: CMS & Assigned to VO support
 - 15th April 2010. TPM asks VO: “Did anyone delete those SURLs?”
 - 6th May 2010. TPM asks VO: “What is the status of this ticket?”
 - 21st May 2010. TPM tells VO: “Please remove those SURLs”.
 - 21st May 2010. Steve Traylen: “Passing to CMS first to see if they care.”
- Biomed (#56679): Banned user spamming site.
- ATLAS (#54912): Not only the VOs that are unresponsive.

The Ganga blacklist.

- We're blacklisted?? Why? How? When?
- First indication often comes from local user who can't submit to site.
- I believe developments are planned, but, let me illustrate what seems to be a common experience for admins...

- Chase the error...

Glasgow on Ganga Blacklist (total of 1 UKI tokens on list) Inbox | X

★ **root** to me [show details](#) 03:30 (18 hours ago) ↩ Reply

<http://gangarobot.cern.ch/blacklist.html> ANALY_GLASGOW

↩ Reply ➔ Forward



- Chase the error...

Glasgow on Ganga Blacklist (total of 1 UKI tokens on list) Inbox | X

★ **root** to me

<http://gangarobot.cern.ch/blacklist>

↩ Reply → Forward

There is no automatic black-listing activated for Panda jobs, i.e. ANALY queues. Panda queues are manually turned on or off. Automatic blacklisting is only activated for glite WMS jobs in Ganga.

Sites currently failing jobs by the GangaRobot for Panda in 15.6.9:

ANALY_DESY-HH
ANALY_PNPI
ANALY_TAIWAN
ANALY_WEIZMANN

Sites with job submission failures by the GangaRobot for Panda in 15.6.9:

ANALY_CSCS brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_CSCS'] are inactive (no pilots)
ANALY_NCG-INGRID-PT brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not found
ANALY_LIP-Lisbon brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not found
ANALY_UAM brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_UAM'] are inactive (no pilots)
ANALY_FREIBURG brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_FREIBURG'] are inactive (no pilots)
ANALY_RAL brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_RAL'] are inactive (no pilots)
ANALY_TW-PTT brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_TW-PTT'] are inactive (no pilots)
ANALY_INFN-MILANO brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not found
ANALY_CNAF brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_CNAF'] are inactive (no pilots)

Sites currently failing the GangaRobot jobs for Panda in 15.6.3.6:

Sites with job submission failures by the GangaRobot for Panda in 15.6.3.6:

Sites currently failing the GangaRobot jobs for Panda in 15.6.7:

Sites with job submission failures by the GangaRobot for Panda in 15.6.7:

Sites currently failing the GangaRobot jobs for Panda in Frontier/Squid test in 15.6.6:

ANALY_DESY-HH
ANALY_FZK
ANALY_FZU
ANALY_GLASGOW
ANALY_LANCS

- Chase the error...

Glasgow on Ganga Blacklist (total of 1 UKI tokens on list) Inbox | X

★ **root** to me

<http://gangarobot.cern.ch/blacklist>

↩ Reply → Forward

There is no automatic black-listing activated for Panda jobs, i.e. ANALY queues. Panda queues are manually turned on or off. Automatic blacklisting is only activated for glite WMS jobs in Ganga.

Sites currently failing jobs by the GangaRobot for Panda in 15.6.9:

ANALY_DESY-HH
ANALY_PNPI
ANALY_TAIWAN
ANALY_WEIZMANN

Sites with job submission failures by the GangaRobot for Panda in 15.6.9:

ANALY_CSCS brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_CSCS'] are inactive (no pilots)
ANALY_NCG-INGRID-PT brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not fo
ANALY_LIP-Lisbon brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not fo
ANALY_UAM brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_UAM'] are inactive (no pilots)
FREIBURG brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_FREIBURG'] are inactive (no pilots)
RAL brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_RAL'] are inactive (no pilots)
TW-PTT brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_TW-PTT'] are inactive (no pilots)
INFN-MILANO brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not fo
CNAF brokerage gave wrong PandaSiteID:ERROR : No candidate. ['ANALY_CNAF'] are inactive (no pilots)

currently failing the GangaRobot jobs for Panda in 15.6.3.6:

with job submission failures by the GangaRobot for Panda in 15.6.3.6:

currently failing the GangaRobot jobs for Panda in 15.6.7:

with job submission failures by the GangaRobot for Panda in 15.6.7:

currently failing the GangaRobot jobs for Panda in Frontier/Squid test in 15.6.6:

DESY-HH
FZK
FZU
GLASGOW
LANCS

- [December 2009](#)
- [January 2010](#)
- [February 2010](#)
- [March 2010](#)
- [April 2010](#)
- [May 2010](#)
- [June 2010](#)
- [July 2010](#)

- Chase the error...

Glasgow on Ganga Blacklist (total of 1 UKI tokens on list) Inbox | X

★ **root** to me

<http://gangarobot.cern.ch/blacklist>

↩ Reply ↪ Forward

There is no automatic black-listing activated for Panda jobs, i.e. ANALY queues. Panda queues are manually turned on or off. Automatic blacklisting is only activated for glite WMS jobs in Ganga.

Sites currently failing jobs by the GangaRobot for Panda in 15.6.9:

ANALY_DESY-HH
ANALY_PNPI
ANALY_TAIWAN
ANALY_WEIZMANN

Sites with job submission failures by the GangaRobot for Panda in 15.6.9:

C	0	0	0	C	0	0	0	ANALY_BHAM
C	0	0	0	N	0	0	0	ANALY_CAM
0	0	0	0	0	0	0	0	ANALY_ECDF
F	C	0	0	F	C	0	0	ANALY_GLASGOW
F	C	0	0	F	C	0	0	ANALY_LANCS
F	C	0	0	F	C	0	0	ANALY_LIV
C	C	0	0	C	C	0	0	ANALY_MANC
F	C	0	0	F	C	0	0	ANALY_OX

No candidate. ['ANALY_CSCS'] are inactive (no pilots found): ERROR : No candidate. Release:Atlas-15.6.9 was not found. No candidate. Release:Atlas-15.6.9 was not found. No candidate. ['ANALY_UAM'] are inactive (no pilots found): ERROR : No candidate. ['ANALY_FREIBURG'] are inactive (no pilots found): No candidate. ['ANALY_RAL'] are inactive (no pilots found): No candidate. ['ANALY_TW-PTT'] are inactive (no pilots found): ERROR : No candidate. Release:Atlas-15.6.9 was not found. No candidate. ['ANALY_CNAP'] are inactive (no pilots found).

Panda in 15.6.3.6:

GangaRobot for Panda in 15.6.3.6:

Panda in 15.6.7:

GangaRobot for Panda in 15.6.7:

Panda in Frontier/Squid test in 15.6.6:

- Chase the error...

Glasgow on Ganga Blacklist (total of 1 UKI tokens on list) Inbox | X

★ **root** to me

<http://gangarobot.cern.ch/blacklist>

There is no automatic black-listing activated for Panda jobs, i.e. ANALY queues. Panda queues are manually turned on or off. Automatic blacklisting is only activated for glite WMS jobs in Ganga.

Sites currently failing jobs by the GangaRobot for Panda in 15.6.9:

ANALY_DESY-HH										inactive (no pilots for last 3 hours).
ANALY_DBNP										inactive (no pilots for last 3 hours).
ANALY_GLASGOW	NONAME							1407	failed	buildJob , runAthena
ANALY_LANCS	NONAME							1358	failed	buildJob , runAthena
ANALY_LIV	NONAME							1356	failed	buildJob , runAthena
C	0	0	0	N	0	0	0			ANALY_CAM
0	0	0	0	0	0	0	0			ANALY_ECDF
F	C	0	0	F	C	0	0			ANALY_GLASGOW
F	C	0	0	F	C	0	0			ANALY_LANCS
F	C	0	0	F	C	0	0			ANALY_LIV
C	C	0	0	C	C	0	0			ANALY_MANC
F	C	0	0	F	C	0	0			ANALY_OX

... No candidate. ['ANALY_TW-FIT'] are inactive (no pilots for last 3 hours).
 ERROR : No candidate. Release:Atlas-15.6.9 was not found.
 No candidate. ['ANALY_CNAP'] are inactive (no pilots for last 3 hours).

Panda in 15.6.3.6:
 bot for Panda in 15.6.3.6:
 Panda in 15.6.7:
 bot for Panda in 15.6.7:
 Panda in Frontier/Squid test in 15.6.6:

- Chase the error...

Glasgow

★ roo

PandaID, Owner, V

PandaID,

100

Johann

ANALY_GLAS

ANALY_LANC

ANALY LIV

```
-----
Core dump from CoreDumpSvc on node063.beowulf.cluster

Caught signal 11(Segmentation fault). Details:
  errno = 0, code = 1 (address not mapped to object)
  pid   = 4, uid = 4157497954
  value = (-137068064, 0xf7d481e0)
  vmem  = 570.152 MB
  rss   = 443.023 MB
  total-ram = 7983.95 MB
  free-ram  = 6150.47 MB
  buffer-ram= 221.836 MB
  total-swap= 2047.3 MB
  free-swap = 2019.13 MB
  addr    = 0x4

Event counter: 0
Current algorithm: <NONE>
Algorithm stack: <EMPTY>
-----
```

manually turned on or off

inactive (no pilots
as-15.6.9 was not
15.6.9 was not fou
ctive (no pilots f
'') are inactive (r
ctive (no pilots f
re inactive (no pi
-15.6.9 was not fo
nactive (no pilots

10

DE

gl

0

F C 0 0

F C 0 0

F C 0 0

C C 0 0

F C 0 0

ANALY LIV

ANALY_MANC

ANALY_OX

anda in Frontier/Squid test in 15.6.6:

- As a site, we'd like to see...
 - Notification of when we're blacklisted.
 - "Issue" tracking (eLog/Savannah/GGUS).
 - Site admin only has to look at the ticket to understand why they're listed.
 - Quicker removal from blacklist. I believe the current value is 24 hours?
 - A Twiki page explaining what is/is not important (PanDA queues aren't blacklisted by Ganga).
 - Escalation procedures.
 - I usually mail Dan/Johannes. Not scalable.

- The larger VOs have the resources to lead the way and help the admins help the smaller VOs.
 - The Ganga Hypernews forum is not open to non-LHC (ATLAS) VO members.
 - We have a couple of local communities who regularly used Ganga (Mice & NanoCMOS).
 - Larger VOs could produce “generic” advice applicable to smaller VOs.
 - It would be nice to tell users of the Standard Way of doing X.
- ATLAS/LHCb: Analysis jobs on demand to test site?
- LHCb
 - Dirac monitoring pages restricted?
 - We can't trace your users jobs (compare with PanDA).
 - Can admins have VO membership for testing submission & helping your users, please?



GridPP

UK Computing for Particle Physics

Reserve Slides

