

ops team 050711

Present:

PG Chair/Minutes

Stephen Jones

Sam Skipsey

David Crooks

Daniela Bauer

Mathew Doidge

Wahid Bhimji

Rob Harper

Catalin Condurache

Rob Fay

Andrew McNab

Gareth Smith

Raul Lopes

Andrew Washbrook

Elena Korolkova

Chris Walker

Raja Nandakumar

Peter Grandi

Alessandra Forti

Govind Songara

Mingchao Ma

Chris Brew

Santanu Das

Kashif Mohammad (Part)

Brian Davies

11:00 Meetings & updates (20')

- ROD team update

AM: All OK

-- The NGI validation process has been initiated.

- Nagios status

- Tier-1 update

A lot quieter this week than last, in middle of scheduled outage at the moment.

Update to Castor to be able to use t10kc tapes.

Batch update to cvmfs. Batch is down and will come back once storage is fixed.

Networking firmware updated, packet loss on router out of site, being updated, and complete.

FTS is back for external transfers

- Security update
-- including any conclusions from the workshop/HEPSYSMAN meeting
<http://indico.cern.ch/conferenceDisplay.py?confId=145475>

Discussions at HEPSYSMAN, in NGI will have a small security team, to be sorted in a week or two. Will have a mailing list for the group so if Mingchao is away someone else will be available.

CA update hiccup. It is being discussed on tb support. Should not see the warning so quickly. There was a bug introduced in 1.39, this package was rejected and then 1.40 came out, this caused some confusion. Will look to improve the process. This is the last test from the old SAM test, will move to the Nagios test only.

-- T2 issues
Oxford Networking Problems:

Liverpool:
Planning on how to use VMs, to run various services. They have setup a test cluster and submission system.
Reasonably quiet.

Glasgow not much to report, some problems with extra link (10G). Port seems to be having problems. Putting in some QoS rules for that link to protect the service.

Trying to reduce the traffic on the link by trailing cvmfs on a node. Does seem nice and will reduce traffic.
Setting up the Dell switches was pretty straight forward even between Nortel and Dells.

Imperial: CMS hitting the dcache very hard but they are in a busy time trying to transfer a lot of data. Very frustrating, but appear to have reached the limit of the h/w. Updated ce's to Cream ce's but that's not working, and the developers have not answering the emails.

T2k storage group.
Space token ticket request. Not happy that that was the first she heard of it. SS: This request is a result of the storage rroup discussion.
The intention is to make t2k use a space token to limit their use to prevent them using all the rest of the non allocated storage at a site.

IC have dcache so it's difficult /impossible to reliable the existing data as belonging to the token.
She needs a number for the size of the space request.

Lancaster:
Thinking of moving to CVMFS but will wait to follow others. Waiting for glexec re-locatable install.
Fixed a database bug in cream, which affected them as they had moved their sandbox.

Resource usage of atlas jobs, requesting 3GB of ram for atlas production jobs. Not efficient as node that has 2.5GB of ram cannot run a job.
AF: There is a historic dashboard, but for details you have to

check the logs.

Atlas official requirement is 2GB per job, but some do use more.
This is a batch system problem with the shared cluster.

Atlas always relied on the fact that the jobs high memory requirements spike, so jobs should not be killed.

ECDF have a group limit on total allocation per node, and then make the per job limit looser.

QM have a reasonable amount of swap on the node and they over commit wrt memory.

Atlas jobs hardly ever go over 3GB.

SS suggests that Atlas should actually be allocated two CPU's if the need twice the memory requirement.

ECDF as ATLAS cloud offline, not running so many jobs, were running 1500 last week.

Battling with excessive load on cream bupdated process,

Testing new version from developers.

Gap publishing error for Apel.

Broker off for atlas analysis, hoping to get to the bottom of that.

Looking to deploy Argus, and will look for the re-locatable glxec.

QM would like the SGE, fix.

RALPP, lower number of jobs due to T1 downtime. Not much to report.

Manchester:

ARGUS glxec setup, cfengine problem to sort out.

Problems over the weekend with the ce going down. Stayed up in tests but no real jobs running.

Cream ce is unstable, keep restarting pbs, and glite services.

Have to keep monitoring it, and watch the numbers, the problem is not the mysql db, resources getting used.

Have to restart things manually, (tomcat).

Script does not start cleanly, the developers should be ticketed.

May be this is causing the instability in the pbs server. Time waits etc

RHUL CMS issues, not much to report.

Sheffield, Apologies for not being able to give a talk at HEPSYSMAN.

Glxec and argus tests now passing this morning. Is there still a APEL publishing problem?

Yes many sites not published since the start of the month.

QM

Big news new colleague Dan Trainer.

Lustre meeting on 14th July please sign up asap.

CREAM ce problems seem to fix themselves over the weekend.

Network intervention this morning to enable gigabit in and out at same time.

Se under strain.

1.7 storm advanced copy to try hopefully.

Govind, cluster under used. not very busy this week.

Santanu

Nothing much to report. Trying to fund the Argus server.
Planning to upgrade DPM this month, 50TB new disk space in rack
but not yet on line.
S15 64bit 1.8.0 DPM is fine.

Brian looking into new golden release of dcache.

-- General notes.

As reminded by Chris, please could sites update
<http://www.gridpp.ac.uk/wiki/HEPSPEC06>.

Useful to run on all nodes as can spot configuration issues such
as not max perf but min power. PG reported one of his WN's lost a
CPU which resulted in the HS06 result being half normal. Power
cycle fixed it , but it is a worry that a CPU could effectively
disappear. May we should be running HS06 tests regularly to spot
such problems.

Peter Grandi: Has anyone used the new node type cluster.
NO

11:20 **Experiment problems/issues (20')**

Review of weekly issues by experiment/VO

- LHCb

Nothing to big to report, jobs running much better at RAL over
last week or so. Question wrt to Manchester to install cvmfs.

- CMS

- ATLAS

3 tickets for storage or networking problems, QMUL was in down
time for router upgrade this morning. May have to put them in down
time as there may still be a problem.

Oxford Networking problems,

UCL analysis queue back on line, prod offline due to lack of s/w
releases. Storage problems reappeared.

T1 in down time for several hours for castor and LFC upgrade.

Cloud taken offline.

Request at ADC to enable CERN in squid configurations. To allow
generators job to access these servers. Graeme tested last week,
Lancaster firewall problem was spotted. There is now a test you
can run to test your setup, AF emailed yesterday.

Deadline mid July but it should be easy unless you have to deal
with University FW people.

Number of sites going in and out on broker off, ECDF for example, UCL, Durham, IC.

- Other

Tickets from t2k,

- Experiment blacklisted sites

- Experiment known events affecting job slot requirements

- Site performance/accounting issues

- Metrics review

11:40 Free discussion time (10')

- Opportunity for sites to follow up on any specific problems being experienced

11:50 AOB (01')

Chat Window:

[11:24:47] Peter Grandi as to ATLAS job limits, I summarized in a mailing list post recently, and it is 2GiB for physical and 4GiB for virtual.

[11:25:26] Peter Grandi my advice is to limit to a bit less than 2GiB and a bit less of 4GiB to allow full resource occupancy

[11:25:37] Mohammad kashif Hi Daniela, can you check CA version at your WN's

[11:26:28] Peter Grandi again: it is 2GiB of physical and 4GiB of virtual, which includes swap.

[11:28:36] Brian Davies joined

[11:28:37] Peter Grandi But alessandra but you *cannot know in advance* that there will be 23x1GiB unless the 23 jobs know it in advance

[11:28:49] Peter Grandi But alessandra but you *cannot know in advance* that there will be 23x1GiB unless the 23 jobs know it in advance and declare 1GiB

[11:28:53] Alessandra Forti no it is all based on stats

[11:29:15] Alessandra Forti I don't have a limit and I really have few memory errors problems

[11:29:57] Andrew Washbrook hi Sam, yeah we apply a slack limit of -l h_vmem=6000M

[11:33:34] Alessandra Forti off for 1 minute

[11:34:55] Alessandra Forti back

[11:37:11] Matthew Doidge Same here

[11:41:02] Mohammad kashif left

[11:46:01] Elena Korolkova .thanks for your comments

[11:49:46] Peter Grandi I don't have microphone audio

[11:49:50] Peter Grandi]so typing in here.

[11:50:20] Peter Grandi OK, some new from Durham: CREAM CE and new BDII installed, and I no longer work at Durham.

[11:51:22] Peter Grandi The CREAM CE and gLite 3.2/SL5 CE seem to be working fine.

[11:51:29] Peter Grandi The CREAM CE and gLite 3.2/SL5 BDII seem to be working fine.

[11:59:48] Rob Fay left

[12:01:04] raul lopes left

[12:06:23] Wahid Bhimji left

[12:10:41] Elena Korolkova RAL is unblacklisted

[12:10:57] Elena Korolkova because DT is over according to AGIS

--

Peter Gronbech Senior Systems Manager and Tel No. : 01865 273389
GridPP Project Manager Fax No. : 01865
273418

Department of Particle Physics,
University of Oxford,

Keble Road, Oxford OX1 3RH, UK E-mail : p.gronbech1@physics.ox.ac.uk
