

JC, FS, JJ,
ST,GC,Derrek,OvdA,PG,SB,AF,DC,Graeme,DS(part),Catalin(part),MH (part)

09:30

Review of open actions and actions process (30')

Jeremy

- Some actions will be revisited during other items
- How well is the process working & how can it be improved?

JC went through the actions, most seem to be progressing,

PG and FS and AF have sent some info wrt to patching tools.

Editing token for security incident doc with Fraser at the moment.

Dave Britton was talking to Dave Kelsey about the security post which was never appointed, (gridpp funded) science post has been filled, (Roman left).

Security incident prevention still open

AF agreed to add Bham and Cam ganglia entries to the web server.

Follow up on Atlas nos.

Sizing formula to TB support, Roger Jones will inform JC.

Ensure sites are warning site network contacts about data transfers.,probably can be closed all sites should have done this by now.

CMS are they using xrootd /scratch space?Remit of UB so far has been to do allocations for T1, now looking at T2. UB needs to be a better forum for interfacing between Users and DTEAM. They will have a stripped ntuple, how will it be stored?

Currently on home directories, talking about having proof farms.

Questions like this should be added to the wiki issues log.

VO info, needs to be made available, Process needs to be streamlined for addedin new vos.

Dcache gridftp xfer perf ongoing

Nominate site for lfc close

Report to dteam any sites readiness.

Circulate to dteam 90 day ref. needs doing, its in the security policy

Plan for completing weekly CIC reports need to get site admins to do this.

Graeme to distill advice.

Nominate sites for throughput testing.

Security contact info in gocdb, JC to follow up again.

Editing time frame for CIC reports to be extended

Q4 reports

Answers about security update procedure

Dteam mailing list to go to cern ok

Condor support in apel

Ensure SFT has enough info for quarterly reports.

Action Pages are getting very long.

Need to agree to get them closed in a timely manner.

10:00

Status of SRM deployments (15')

Greig, Jens, Graeme

- Has the deployment worked
- Is it sustainable
- What are the major problems that have been encountered
- Which sites are proving "difficult"
- Are we happy with progress
- What difficulties do we expect over the next 6 months?

In general it is working as only 2 remaining sites to go. Sheffield have made some progress but cooling failure and h/w problems have slowed things down. Grieg can test it if they let him know which machine it is. AF could also help if she is given access to the machine. Now only have 0.5 fte. Effort

Brunel fw problems. Will install DPM on head node next week.

Manchester needs certificates, dcache sites should update to 1.6.6.

Not trivial, but there are instructions in the wiki.

Have to down grade for the info system.

New DPM is available but can just wait till 2.7

Need to decommission all classic se's.

Web page shown, storage at sites, Some figures seemed to be incorrect, info systems possible not advertising correctly. Status as permanent, at most sites. SRM definition means that they will not randomly delete files. QM will be marked as Volatile.

Definitions need to displayed on the web page. Possible difference in perception of 'Permanent' by experiments of sys admins.

10:15

SC4 milestones and transfer tests (45')

Jeremy, Graeme and ALL!

- Are any milestones looking unachievable
- What are we learning from the transfer tests
- Have sites started to engage enough
- Is this scheduling process sufficient
- Is the site support sufficient
- Are the coordinators happy with their level of engagement in this area
- What additional tests are required over those being done

Milestones, quite a few sites involved,

All sites working by end of January yes but maybe UCL CC due to equipment move, Sheffield also have problems so are uncertain.

40% of sites using FTS with SRM2.1 by end of Feb. OK providing JRA1 provided support in time, in FTS?

Interoperability tests, not yet started, also need castor to be online, this should happen by end of Jan.

FTS channels all T1-T2 by end of Jan

FTS clients to be configure all looks like its on schedule

Data transfers

Have done 20 % of sites but not all at target rate! Many less than 300Mb/s
(Dec 05)

Could try using http or (even rfio! If desperate) but this hides the real issue. DPM to DPM tests can be done to prove the network can cope, but we need to solve the dcache to dpm problem asap.

Sustained transfers by mid march. Bristol may be a problem as they have limited storage at the moment. UCL CC may also have trouble meeting this dead line.

Sites with 100Mbs bottle necks will clearly not meet the targets.

Sites within a T2 could define one site as the close se rather than itself, if they have good links between them.

Iperf can be used to measure network perf. If transfer test are not meeting target rates.

LFC

Southgrid and Scotgrid are fine, North grid behind, LT2 half way there.

VO Boxes

Assumptions are being made that say Birmingham will support Alice, (Is this true).

As may expts only require 1 VO box per T2 as apposed to at each site.

ATLAS had trouble with their VO model and are now considering installing a cluster at CERN, LFC nodes were primarily used by Atlas so it may be that we will not need them in the long run. So this situation is in a state of flux.

ACTION: Graeme Check up on whether ATLAS and other expts will really need LFC boxes at T2's given that ATLAS are moving VO functionality back to CERN.

Awaiting tools for testing of LFC

Depending upon expt responses to questionnaires,

VOBs for 1 site in each T2 by mid jan?

DC talked about cms. requirements CMS does not have a VO box as such but does require some software, (which probably should be in a vo box)

T1 have installed ATLAS, lhcb, and alice.

Learning interesting things about disks. Disk tuning is probably site dependant. Some general suggestions can be made but may be not applicable at every site. Sites have different raid 5 solutions. Storage group are discussing this sort of thing. Join the email list/meetings....

DC It is generally a good thing to do, only be doing the test can you have confidence that it works. The tests are quite stressful for the SRMs involved. 50 active streams, if it stays up its good. We need to do continuous tests to give a realistic idea of a real use scenario.

One thing they have noticed is that tcp does a very good job. JJ its is an experimental science, getting these things to work together.

Next week there will be no t1 tests? May be just intra t2 tests.

Redoing SC3 tests next week. After that week can continue t1-t2 tests.

T2C's are happy so far that it is working, and they have sufficient notice. Often network people are enthusiastic to see what throughput can be achieved.

Any struggling sites?
Additional test. Iperf details going to be on the wiki.

11:00

(email) BREAK (10')

11:10

Experiment disk allocations at Tier-2s (50')

Discussion

- What is the request & future role of UB
- What was suggested at the PMB/DB meetings
- What are our concerns
- What additional requirements information is needed
- How do we implement the requests
- How do we start/improve the experiment interaction to change perceptions about Tier-2 disk reliability
- Can we write a summary of our understanding (of terms)
- What else can WE do to get better experiment collaboration

JC: Trying to get expts for engaged with t2s There is a shortage of disk at t1 but an excess at t2s so try to balance it out. UB want to allocate t2 disk, eg give ATLAS x TB at these centres.

Eg 10TB at Northgrid how would we arrange for this to be met.

FS It depends on the sites setup, which SRM they have. Olivier: for DPM we would need to allocate a particular pool.

Problem making changes to fast. T1 has much more effort than T2. So making changes more often than 6 months or even a year would be hard. If you have just bought new disk it is easy, but if you have to ask a vo to remove disk usage and re allocate to some one else much more effort and time required.

We have seen when decommissioning classic se's how difficult it is to get people to remove data from a server.

Initial suggestion would be ed lhcb, ICL cms, Lancaster ATLAS. SB Crucial point is the definition of permanent storage. Expts are worried about sys admins deleting data. T2Cs generally feel this is unlikely to happen. JC we think there fears are unfounded. Permanent means that you will not deliberately delete data. But it does not say anything about reliability.

Much discussion....., DC says cms are happy to use T2 disk space, Dave Newbold (head of UB) said at the meeting yesterday that the experiments don't trust t2 disk.

Roger Jones is telling us that they want to try allocating disk space at T2s.

We could implement something, there needs to be statements of what the h/w is.

AF does not agree, G Storage is the most un grid like part of the grid.

In London their could be a problem because they use poolfs at QM.

Engage students and users. DC does ATLAS blacklist sites?

Talked about how we get software installed at a site. Atlas has a web page .

Other groups have to email the software manager.

T2Cs should possible email the expt software installers and coordinate installation at all there sites.

What else can we do to get better collaboration. AF Start talking to people we know.

DC there should be a joint meeting between dteam and experiment people. But who...

12:00

TPM - the UKI contribution (30')

David Spence

- What does it mean to be a Ticket Process Manager
- How will the deployment team contribute
- When do we need to support
- What training is going to be provided

DS gave a talk:

Changes to UKI ROC help desk in approx 1 week. 23rd jan

New automatic interface between ggus and footprints server.

New name from ggus to UKIROC email ukiroc@grid-support.ac.uk

Should be just one email now, should no longer need to login to ggus.

Closing tickets in footprints should automatically close in ggus.

Can escalate to a new responsible unit.

Some problems:

History limited, ensure new comments tell whole story

Attachments are not copied (footprints lack of functionality)

GGUS only has one solution where as footprints is more of a conversation.

TPM talk:

Ticket Process Management

Team of 2-3 people from a roc and cern.

PG not keen . JC feels it is not a big job. This would be your job for that week....

AF not keen.

JC it's a dteam activity, SF thinks it will be ok.

12:30

LUNCH (1h00')

13:30

Building an effective testzone (and the 2.7.0 release) (45')

Jeremy/All

- First, how are we going to work next week for 2.7.0 testing!?
- What are the testzone objectives
- How many machines were purchased & delivered (& still available!)
- How are they allocated locally (what is each machine going to do)
- What do we need in place to make the testzone work?

Expected Monday. Good plan for components, Graeme DPM LFC, FTS clients Grieg, dcache, server and clients AF VOMS at Manchester.

BDII /RB IC, RGMA RALPPD SB and ST can help, Many sites will install a mini site (eg one per T2).

Give feedback within 3 days, and then further feedback over the week.

Probably will do fresh installs, although an upgrade would be more realistic. But time constraints mean we will not do this this time.

Feed back will go back to a list to be provided by cern.

Scotgrid got systems for catalogue testing, in addition (Total 9)

Most sites could also take out a worker node or two if necessary.

Use the wiki to describe what's going on.

PPS Status, Bham set up a site, IC are setting one up. Yves experience is described on the wiki.

Glasgow are part of the PPS, they run the FTS, in2p3 did an upgrade based on Glasgow's experience. Cnaf and nikef provide voms service for pps.

Voms is the same software on glite as 2.6.0 so no benefit from testing this.

SB PPS is centrally managed, there are regular meetings, not like production system where we have the uki structure. People should join the mailing list, and announce their intentions. Yves found it easy to install so did not ask for any help.

Glite 1.5 may well stay on the PPS even though it has no long term future.

Glite 3 will be the next major release but that is basically lcg 2.7.0 plus some workload management parts from glite. PDG asked where pressure to join was coming from. Pressure to join came from Ian Bird (egee) but JC will find out if it is still useful to encourage sites to join.

NGS Scotgrid investigating what will be required to join. Core sites have to run a specific software core. VDT. (part of LCG) so it should be easy to allow NGS jobs to run on LCG clusters but more difficult to run LCG on NGS clusters. Oxford has started talking to OXGrid but not yet to NGS.

14:15

Areas the Tier-1 can help Tier-2s (30)

Steve/Matt/Andrew

- What expertise is available
- What has been done that might be useful outside of the Tier-1 (e.g. efficiency & benchmark tools)
- What processes are being followed for channels and transfers
- Are there any concerns about current requests
- What is the VOMS support model
- Any other grid services?

Catalin and Matt Hodges supporting VO box 's.

ST Need to put info on web sites so others can learn. Ganglia info for example. PG asked about XFS T1 had looked at it but then decided not to use it as Red Hat removed support from the kernel. ST finds it difficult to persuade people to put results on the web, because they don't want the hassle of answering questions (How did you get that results...)

There has been criticism that the T1 has done useful things but they are not publicised anywhere formally or informally. Graeme very happy with T1 setting up FTS channels very quickly. No pressing requests.

Ganglia for example.

There is some new info in the wiki on the tier 1 section.

Catalin and Matt have been attending some experiments meetings, a list of the meetings T1 is represented at should be put up on the wiki.

Large expts such as CMS and atlas get involved with the t1 as they have special requirements. The smaller expts just get on and work!

ACTION: JC follow up with Andrew about T1 representation at experiments meetings

No problems setting up channels so far. If the load gets higher would need scripting. If there is an FTS upgrade would that cause problems, would have to advertise to users. Upgrade to 1.4 should not break anything! So would take it down for a few hours.

VOMS support model, Manchester will setup things for new VO's . RAL had done some work but Manchester is doing a good job.

Pheno grid central FC is at sara. Any new VOs who want FC can be supported at T1.

SB Question about RB, there is one at RAL, New one at IC had problems but the testzone one is OK so they will swap them over.

14:45

The latest on VO-Boxes (15')

Catalin/Steve

- Where does the Tier-1 stand in respect of each LHC experiment
- What are we currently expecting of Tier-2 sites
- Nature of the LCG "working group"

There are 3 vo boxes at t1, alice stated work in Nov but had no disk allocation, needed 1tb, some space is now allocated (2tb) so will start work. Atlas was asking for another box, question about who is going to run the apache web server. They wanted a host cert, ST said no as that would effectively give them root. They gave up.

LHCB are the only group to complete the questionnaire, fully. LHCB vo box was installed 2 days ago. Sgm people contacted.

Cms installed monalisa demon, squid web cache, part of 3d project. Don't require a vo box as such, but ask for these services to be installed.

LHCB require a service certificate, UK CA will issue a lhcb vo box server cert.

Not yet asking for T2 support.

Agenda not yet announced for the Working Group. Will look at how to phase them out, and integrate the services in to LCG middleware.

ALICE is associated with Bham, but its not really clear how it will happen.

15:00

BREAK (05')

15:05

Documentation (25')

Stephen

- What help is needed from the deployment team
- What is the best way to use the wiki and how do we manage it!
- Keeping web-pages up to date
- Feedback on material so far
- Ideas for "care & maintenance guide"

SB not completely clear what his remit is. Mainly being looking over it.

JC explained that management was still concerned that pages are not complete.

SB most pages are reasonable but have not been updated recently. JC admitted that some pages were out of date.

Looked at the old map.jpg in the root of the deployment area. Each person should have a look at the area that they were responsible for.

ACTION: Each Dteam member should have a look at their area.

On grid site you can check the update dates on pages.

On the wiki you have similar tools. Special pages, oldest pages. For example.

GOSC moved to mambo, (A content management system), but JJ says its not easy.

SB does not see the advantages of the wiki, ST finds it easier. G said deployment should go into wiki, user docs in web. SB would prefer that others did not edit documents that he is responsible for, but should ask him to make the changes.

Feedback would be good, particularly from users. But so far nothing, no traffic on users mailing list, was announced at gridpp 15 but again no/ not many users there.

Care and maintenance guide, SB would offer to edit it, but material needs to come from experienced managers. Responding to CIC tickets, dealing with upgrades etc.,

ACTION: SB provide outline of Care and Maintenance Doc/page.

15:30

ROC processes (30)

Jeremy

- Actions required for VO summary information
- How do we deploy and support VOs generally
- GridPP VO. AUP and joining mechanism

- Getting better input from sites for reports
- Actions in the area of security
- What else!?

JC needs to define some things.

What actions are required for VO summary info.

AF has many comments. Some surprise that yaim or a new web based tool would be trying to set the fairshare % in the maui.conf file as each site may have a very complex setup. Its probably best left to sites to sort out.

Looked at the CIC portal VO management, some fields missing from the so called id cards. Gridpp page should possibly have a list of recommended VO's. including zeus and hone etc.

How do we support VOs generally. VOMS certificates are distributed by rpms so new vos have to be added to the rpm.(LCG release?)

What is the gridpp vo for, is it test or a catch all. AF Original idea was to have all VOs be a subset of gridpp. SB thinks it a bad idea, AF also thinks its not so good but did create it, but does not know what its for, may be just testing.

New expts could use it before setting up there own VO.

JC said DK had mentioned the AUP. You need an administrator.

How do we ensure sites support all VOs. FS Its annoying that sometimes you enable a vo and it never gets used. Its more effort at Glasgow as they do not just run the yaim scripts there.

CIC portal may already list which sites support which VOs.

Why sites find it difficult to fill in the reports. Time window is short. Closes at 1000 European time, which is 9 am

16:00

Meeting schedule and external meetings (10')

Jeremy

- Dteam meetings
- UKI ROC meetings
- Perhaps plan a PPS meeting
- Who is going to CHEP
- June workshop on Tier-2 SC work

Dteams on Thursday 11am vrvs

UKI ROC meetings... next feb

JC will try to find out about PPS

CHEP. Who is going AF, Graeme and Grieg. Mona and DC.

June Workshop on SC at Cern 12-14th June.

Actions

ACTION: Graeme Check up on whether ATLAS and other expts will really need LFC boxes at T2's given that ATLAS are moving VO functionality back to CERN.

ACTION: JC follow up with Andrew about T1 representation at experiments meetings

ACTION: Each Dteam member should have a look at their area.

ACTION: SB provide outline of Care and Maintenance Doc/page.