

Transcript of Workshop on dCache Reliability

Indico link: <http://indico.cern.ch/conferenceDisplay.py?confId=45966>

Author: Jos van Wezel, Xavier Mol

Date: January 14 and 15, 2009

Location: FZK Karlsruhe

Participants: Jon Bakken (FNAL), Gerd Behrmann (DESY), Gerard Bernabeu Altayó (PIC), Giacinto Donvito (CMS), Patrick Fuhrmann (DESY), Silke Halstenberg (FZK/KIT), Christopher Jung (FZK/KIT) Simon Liu (TRIUMF), Paco Martinez (PIC), Xavier Mol (FZK), Doris Ressmann (FZK/KIT), Trompert Ron (SARA), Pedro Salgado (BNL), Jonathan Schaeffer (IN2P3), Killian Schwarz (15/1 only)(ALICE), Lionel Schwarz (IN2P3), Reda Tafirout (TRIUMF), Erik Mattias Wadenstein (NDGF), Jos v. Wezel (FZK/KIT), Onno Zweers (SARA)
Flavia Donno, Harry Renshall and Jamie Shiers (connected remotely for a presentation by Jamie) (CERN/WLCG)

Topic: Current reliability of storage/dCache

Some sites, outspoken SARA/NIKHEF, still experience a lacking communication/cooperation between experiments and the storage experts. U

Topic: How to secure the future of dCache

Although dCache is there now, it is still a good idea to keep an eye open for new developments and have these evaluated. E.g. FNAL (Jon) does this to make sure it's not locked on a single solution. Everybody agrees this should be done but time and resources are limited.

If dCache development would be significantly reduced or stopped one could run the system for some time awaiting a solution. At least a T2 or straight-forward dCache installation can continue operation for a long time. The problem of the T1s is that their installations are not in steady state, issues are more complex and assistance from experts is regularly needed.

It is concluded that most T1 support issues end at dcache.org. The current expertise at the T1 is at a very high level but new issues and requirements crop up for which also the developers have no canned solution. The iterative process to tune, update or fix issues and requires in dept knowledge. NDGF has countered with a full FTE involved in direct development and part of the dcache.org team. This 'support' solution has benefited NDGF and other T1s significantly.

The reliance on dcache.org must be reduced, either by direct development or support involvement. Documentation is still lacking in many aspects but is felt severe for problems T1 sites are confronted with. Everybody is willing to contribute documentation, and already did so, but a framework is missing.

Jon says that site administrators could widen their scope if they would take the source code and compare it with their problem at hand. Not all problems require a detailed code review but

many times errors or settings can be checked by glancing at the source code. Of course one needs to know how to find it.

Jonathan and most others are basically against looking at the code because this is a step too far. The docs should explain the protocol and behaviour of the (sub)system.

Conclusion 1: It was agreed that it helps to have someone on site that is familiar with the layout of the source code and could assist in first aid situations. This does not mean every site should join in development and send in new code. In fact finding the proper location in the code may be the most difficult step. (See Conclusion 2 below)

Conclusion 2: dCache.org will organize several master classes (format and frequency to be discussed) which target at accessing and understanding the layout of the dCache source code with focus on self help.

The weekly phone conferences with developers are a mixed blessing because problems are discussed but not looked at directly. E.g. throughput issues come and go and are of course gone when the conference starts. For emergencies and reduced support cycles it helps if developers have a login at the site.

Topic enhancements/features

dCache is a relatively young project with not many users and only a few sites. Support for T1 installations is lots of work. At the same time dcache.org tries to incorporate as many wishes for enhancements and features as possible. Even those that are only helpful at a specific T1 site. The question from dcache.org to the T1 admin is to investigate the possibility to return the favours from developers, when requesting site specific features, with either money or work,. Both mean improved code and documentation.

Note: some improvements come automatically. Because dCache is written in java, it profits from improvements in java (e.g. library updates).

Conclusion 3: all sites not already significantly contributing will relay the question to their management in order to allow more time for expertise building and not just operations. T1 site that are directly contributing at the moment are FNAL, NDGF, FZK

Topic Self-solving operational problems

Many times it is difficult to tell why dCache ran into problems either because dCache lacks crucial logging information or the logging is not specific or the relevant message is drowned in other messages. Reply from dcache.org: there is no documentation for specific error messages, but everyone can access the dCache.org wiki and write one. Exchange of setup files could also help. E.g. Jon copies the setup files to a public readable place.

Problem tickets signal issues that colleagues have with their system. It is possible to get access to the ticket system of dcache.org.

Conclusion 4:

dCache.org will provide a clean wiki page where T1 sites can start entering information in the form of Setup files, How-to's, FAQ's etc. to improve knowledge consolidation and exchange.

Conclusion 5:

dCache is an open source project with small user community and therefore some drawbacks must be accepted. T1 admins should be able to understand the source code at entry level as documentation and not always rely on external help.

Topic: Versioning/updates.

Upgrades that actually involve new code and extensive functionality changes are steered through dcache.org in communication with the T1 sites. The weekly phone-con is a meeting point to synchronise. NDGF is already publishing and running updated versions before these are made public on the distribution page of dcache.org. FNAL never upgrades, unless it has to. “never change a working system”.

Upgrade assistance can be expected from the upcoming migration module.

Site sessions

NDGF: Mattias Wadenstein shows two scripts he uses for monitoring dCache data transfer and I/O operations using the billing files.

PIC: Gerard Bernabeu shows slides of their dCache setup and presents tools for graphical monitoring of pool availability. This leads to a discussion of memory problems for SRM and poolmanager. Conclusion: the pool manager normally does not run out of memory, although it queues a lot of requests for longer time. (the pool manager keeps all data in memory.)

FZK: Artem Trunov shows the redundancy of the setup. E.g. the pool failover which allows pools to be served on peer nodes using the shared disks in each group of pool hosts.

Conclusion 6: Information flow: Sites need quicker access to relevant info (updates, problems that may be seen at other sites)

Conclusion 7: Pool is down: If a pool is down, it should report an error, when dCache wants a file from it, instead of waiting and retrying. The behaviour could be different for different protocols. Preliminary answer from Patrick: This change is significant and something to be discussed in depth also with the experiments.

Conclusion 8: More tools for command administrative tasks are needed

Answer Patrick: Remember that there are Jython templates available for everyone to write own tools Existing tools regretfully are not distributed to the world.

Operations: PostgreSQL and Databases (Moderated by Paco)

For dCache:

- PostgreSQL database for PNFS must be on the same host as the dCache domain although it might work for other services.
- Databases for PostgreSQL always on separate disks

Jon: FNAL changes PNFS+Companion to Berkley with performance improvement up to 4. Little layer in dCache has to be adjusted/rewritten.

Paco: There can be a warm standby for PostgreSQL (failover will lose up to 1 minute actions) using the write ahead logs of PostgreSQL. If more precision is needed, than only third party tools (e.g. slony + heartbeat) provide this. Recovery and fail-over can't be automated except with third party tools.

Operations: HSM and tape systems (Moderated by Ron)

NDGF presentation (Mattias)

TSM needs long lists of requests for files to be efficient. But dCache recalls files on singular basis.

It is possible to collect lists for recalls, but it's difficult to maintain these lists to fit to the size of linked read pools. TSM can't manage more files in one session than a read pool can contain on disk.

Another limit is the number of outstanding get-scripts with respect to memory. Each get needs a certain amount of memory space and a lot of gets consume all memory.

TRIUMF presentation (Simon)

Tapeguy uses a knowledge base which presumes disciplined experiment users and agreements on pnfs paths.

Sustainability

Ron: Maybe a workshop giving insight in the structure of dCache code might be of help for understanding dCache.

Patrick: When dCache has a current problem, developers would be happy to receive message of it and get access to the system in trouble. This gives very valuable information. dCache.org is in need of help with stress tests.

Patrick: Sites are responsible for understanding use patterns and limit actions of experiments, if they are malicious. Sites are especially responsible for more alertness in the management board meetings, where requirements of sites and experiments towards dCache are discussed.

All: Are there other denial of service attacks beside calls like srmls?

dCache.org decides on core pieces of the code, explains it to the sites, which in turn converts this explanations into documentation.

WLCG/Experiments

Presentation Giacinto Donvito on behalf of Flavia Donno: Summary of experiments open issues with storage management and dCache

- Pinned files can be removed (from disk) which causes problems. That is, files are removed from fileserver but not the dCache management systems.
- Problems with VOMS proxy in gPlazma, caching time of proxies.

HSM interface redesign needed for (at least) three sites: NDGF, TRIUMF, IN2P3. Happy for now are SARA and FZK.

Jamie Shiers, WLCG, presentation

- Things look good most of the time
- Problems known to Jamie are not specific to dCache/T1 sites but they are definitely affected.
- Reprocessing must be possible but till today only ATLAS tested it.
- If experiments don't start, sites should define requirements and ask the experiments to participate in testing them.

Lionel's suggestion to parse Flavia's/Giacinto's presentation and extract questions from it, that either sites or developers answer on for the next meeting. (d.o = dcache.org)

- Is it possible to refuse connections to a loaded server instead of waiting for the connection to proceed? d.o
- What services currently crash under load? T1
- What can we do against DOS attacks? d.o
- Are all tuning/configuration parameters of dCache documented in relationship to the behaviour of the whole system? d.o
- How to limit interference of different VOs? T1
- Can activities be prioritized? d.o
- When is pre-staging successful? exps. (TB/hr or files/hr?) number of files over time, number of TB over time. exps.
- How can you prove that you can write to tape at the required rate for all VOs at the same time? T1
- Do you have enough monitors/plots for exps. to verify? T1
- Is the same pre-staging rate expected at all times (because this depends on data writing)? (depends on the site) exps.
- What do you expect to happen when throughput limit is reached? T1/WLCG
- Can the name server perform better?
- Can we setup T1 directions for site use (T1-MAXINT)?

Follow-up

Reaction on the question: Should we organize such T1-ws on regular basis or some more times and was it useful:

- Jon prefers not to travel but maybe via video conference
- Quite important from Giacinto
- Patrick would come a second time
- Every six months? (PIC)
- Master classes sooner than six months.
- Meetings not speaking only about dCache

- Next time we need to limit slides even more. Slides of nice topics about sites' work are not able to focus the participants on problems. (Jos: moderation of this kind of discussions takes a lot of preparation for the organisers)

Final remarks

Jos: We 'missed several important topics and questions' that are listed on the agenda page. It would be good to look into these during a next meeting. He also apologizes for not having proper remote participation equipment because of a coincidence. Thanks all participants for their input and time. Hope to see everybody soon.

END