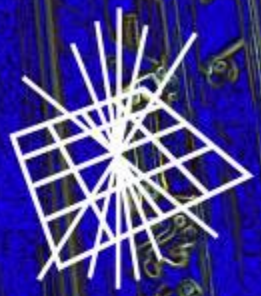

GridPP

UK Computing for Particle Physics

dCache T2 Feedback

Brian Davies

Science & Technology Facilities Council
Rutherford Appleton Laboratory

Outline

- Who Was Polled
- Response
- Current Site Feedback
- Ex Site Feedback
- Conclusion

Who was Polled?

- Tier2s
  - Those who currently run dCache
  - Those who have run dCache in the past
    (Comparison analysis)
- Tier1s
  - Those who run dCache themselves
  - Those who support dCache at associated Tier 2s.
- Regions
  - EGEE and NorduGrid (no OSG response but not sure if they received request)

Response

- Mixed
  - Positive and negative
  - Verbose and succinct
  - Some Unique abut also Repeated views
  - Mostly regarding dCache itself wit only a little on the support structure
    - (Did manner of request skew this ?)
- Not Exhaustive Survey
  - (Is a site more likely to respond negatively if they have issues; rather than a site respond positively if they do not have issue?)
    - Skewed view
  - Change of Scope
    - UK to ROW
    - Mail bounced as spam
    - Not on dCache users list

WARNING WARNING WARNING

- The following views expressed here are those of the sys-administrators polled NOT the talk presenter!!
- Whether the views are correct or incorrect, these comments are the perceived views of the sites.
  - If wrong then need to clarify

Current Site 1 Comments

- ‘happy with it so far, especially since we got rid of pnfs’
- ‘rough edges here and there, but see a clear progression towards more stable and admin-friendlier software.’
  - ‘now SRM seems to not be changing all the time’
- ‘As a somewhat experienced admin I find dCache rather predictable most times’
  - ‘and if not I can usually get good response and support from the dCache developers’

Current Site 1 Comments

- 'I find dCache to be both stable and scalable'
- 'Effort installing might be a bit high'
  - 'as usual with [grid] software, error messages are sometimes quite unclear'
    - But this too has gotten a lot better over the last couple of years,
    - looking forward to the dCache version of a couple of years into the future.

Current Site 2 Comments (1)

- What communication channels do I find useful/difficult?
  - I mostly turn to the user-forum for support; is usually the dCache team that answers.
  - Contacts through V0 with a number of US sites.
  - Own ROC very good. Useful for keeping me abreast of developments/new releases; if not hard core configuration questions.
- How do I find Installing/configuring/updating/upgrading dCache?
  - Better I've given up on YAIM, (not YAIM fault; but the original config on our dCache was so early a lot of stuff is not set up how YAIM expects and it's much easier to do it by hand.

Current Site 2 Comments (2)

- dCache's best feature?
  - Speed and scalability
- What do I wish was better?
  - Permissions, we need ACLs as soon as possible.
- Space Tokens on dCache?
  - It's not exactly trivial but once I've got my head round it I haven't had many problems setting up or updating space tokens.

Current Site 3 Comments

- ‘I don't think we are happy about the situation as it is now’

Current Site 4 Comments (1)

- ‘running without downtimes for very long time’
  - ‘When the system comes up, it doesn't drop by itself’
- ‘Normal operation is easy and smooth’
- ‘Even though we've got a lot of cores dedicated to dCache, we experience timeouts on high loads’
  - ‘pnfs server, which will be changed soon (with move) to chimera.
- ‘Documentation is really poor’
  - ‘"The Book" is obsolete’
  - ‘wiki is not educative’, ’not really trusted to put into production’
  - ‘Many features are hidden’
  - ‘core topics (pool selection mechanisms, p2p transfers, read/write pools, link groups) lack examples’

Current Site 4 Comments (2)

- ‘If I wanted to install a brand-new dCache with chimera; shouldn't that be explained in the book?’
  - ‘Do I still need to mount the pnfs in the doors?’
- ‘There's a new information system. How do I make use of it?’
- ‘Information about new releases has improved a lot.’
- ‘It's good to know what has changed, but I don't upgrade our system so often, and I need to know if I have to change anything (database, filesystem) in my upgrade.’
  - ‘...do I have to read all the release notes of all the intermediate releases...’
- ‘need a tool to verify that database schemas’
  - ‘Isn't the install.sh script able to perform those checks for me?’

Current Site 4 Comments (3)

- ‘dCache says you're able to leverage high loaded pools with p2p transfers. What they don't say is that you need read-only pools’
  - I have to dedicate extra pools just for that, and those would become bottlenecks again’
- ‘ACLs were promised for the end of last year. Where are they?’
  - other options to implement request?
- ‘…this is a centrally controlled system, starting and stopping dCache services should also be controlled from the gui.’ (including creating the pools)
  - ‘I've tried to automate dCache installation & configuration, and having more than a single entry point makes the whole thing impossible’

Current Site 5 Comments

- 'Moving away from dCache'
  - (Resiliency vs Space Tokens)

Current Site 6 Comments

- general monitoring of the system is poor and not operations-friendly
  - eg. the web interface is polling its info every minute or so
- log files are some times far too verbose for normal sysadmin work
  - they look useful for developers but, our debugging is at system level
- configuration of the system should or could be improved
- Much of the activity has to happen on the command line with tools/options that are little documented
  - we need to migrate data from se05 to se21… how do we do that?
    - the answer is not that straightforward even though the need for it is well-understood and fairly common.
- Tier1s strive to get one on-site developer for dCache maintenance. Tier2/3s are much smaller, also in terms of manpower and need also support. Currently this happens through good personal contacts, but this model is not sustainable for support at a "business level".
  - Tier2s to have a dCache/WLCG expert person for reference?

Ex Site 1 Comments

- "DPM simpler and easier"
  - 1) Configuration monumentally simpler and easier to setup
  - 2) (at the time, at least) better support for space tokens, publishing and amending
  - 3) Better stability and resource usage
  - 4) Better organised and more logical logging
  - 5) File ACLs and more fine grained user mapping
- "dCache tuneable"
  - 1) More configuration options (but see 1 above)
  - 2) nfs-style direct POSIX access (although not without its problems)
  - 3) easily scriptable cli interface
  - 4) Better support for replication (although we hardly need it)

Ex Site 2 Comments (1)

‘An overly tunable ‘House of Cards’’

- ‘prone to breaking after minor or major upgrades’
  - shift to srm2.2
- ‘PNFS was Black Magic’
- ‘HSM storage capability a cause of peril’
  - (accidently enablement of HSM-only option
- ‘dCache/PNFS was resource intensive’
  - Splitting head node a necessity not an option
- ‘Never achieved space tokens and publishing them to work’
- ‘Metadata namespace/physical file mappings easier in DPM’

Ex Site 2 Comments (2)

- '..when something did break for the dCache although the developers did they best; almost every problem resulted in a few days
  of downtime or reduced service as you picked apart the server to see
  what was wrong in the first place.'
- 'It wasn't all bad'
  - The replications were tunable.
  - As were the pool draining mechanisms.
  - Pool file-loading was better balanced.
- 'dCache was to us an overly complex solution and at times in its development cycle it was overly fragile…'.
- '….DPM whilst sturdy, is occasionally short on options.'

Conclusions

- dCache both seen as Success and Failure by Sites
  - 'you can't please everyone'
- Ex sites; although having problems that are solved by other implementations; still see the advantages that dCache had over subsequent chosen SRM implementation.
  - 'Not all Bad'
- Support methods vary in variance and depth.

Will any support mechanism ever please everyone?

Will ANY SRM implementation have all functionality implemented that will please everyone?