

dCache Briefing - CDF Perspective

Robert D. Kennedy

FNAL Computing Division

for the CDF Data Handling Group

15 May 2003

v1.0

- ⇒ **Data caching at CDF now**
- ⇒ **Open Issues**
- ⇒ **Deployment Plan w/ Fall-back**
- ⇒ **Some Risks**

Now: Disk Caching

CDF Needs: being served by a combination of systems

Production: "static datasets" containing 13 file servers ~ 25 TB

- a) Limited to the datasets loaded on disk once and for all
- b) Manually loaded, difficult to repair if any files are lost (h/w fails, etc)
- c) All datasets are snapshots. No loading of new data as it comes in
- d) Too small for CDF physics needs for summer by at least factor of 2
- e) No load balancing or load "traffic shaping"
- f) As such can only handle worst-case load from 1/3 (200 cpus) of CDF CAF
- g) Must be augmented by some other system...

Final Beta: CDF DCache containing 25 file servers ~ 48 TB

-) See list next page. Can handle 1/2 (300 cpus) of CDF CAF (conservative)

Legacy: CDF DIM/Kahuna containing ~ 15 TB

- a) Can access any CDF dataset
- b) Slow - being served by legacy SGI SMP
- c) Easy to maintain nowadays, but no experts left... we will be phasing out.

dCache Open Issues

1) "Duplicate mover requests" - appears to be resolved... so far

- a) Some client requests to read files had duplicates generated inside of dCache
- b) Duplicates wasted "active read request" queue slots, throttling throughput
- c) Work-around: increase size of queues, but loses some load mgmt benefit
- d) Fix from Patrick Fuhrmann is in the system (Tues - upgraded dCache s/w)
- e) This fix appears to be working (Jon Bakken - Wed aft)
- f) After more experience to be sure, we can reduce queue sizes to control access

2) dCache crashes - few per week dCache service interruptions

- a) Was admin node crashes (kernel panics) with dual-Athlon "worker" admin
- b) Then just software crashes (kernel exceptions) with upgraded kernel
- c) Installed a trusted Dell 2650 server as main admin node - no crashes so far...
- d) **LAST REMAINING ISSUE** for CDF w.r.t. production service
- e) No detailed model for crashes - "bad platform?"... **WAIT AND SEE.**

Deployment and Fall-back

1) Final Pre-production Roll-out - dCache

Begin: now. End: <= 06 June

- a) Already managed 20 TB/day user load (read+staging), stability is key now
- b) Add file servers (total 36+), increase fraction of CAF that can access (2/3+)
- c) Pre-load "static datasets" into dCache. Re-org cache to "pin" these on disk
- d) Invite "break it with load if you can" testing by skeptics

2) Staged Production Deployment - dCache, w/statics for a time

Begin: "stable", <= 06 June End: 01 Sep (end summer)

- a) "golden dataset" service level: will deliver high reliability, one way or another
- b) Stage 1: Retain static file servers "live" through summer (smaller dCache)
- c) Stage 2: Absorb static file servers into dCache, increase golden dataset size
- d) Size of "golden dataset" service is limited by n file servers (static file servers)

3) Fall-back - replicate static file servers, w/dCache to the side

Begin: "unstable", <= 06 June End: 01 Sep (end summer)

- a) Populate copies of golden dataset on static file servers to handle CAF load
- b) Limits size of golden dataset service: CAF load without "shaping", mgmt
- c) Issues regarding staff-intensive maintenance, resilience to h/w failures, etc.

Some Risks

1) Public stability test: failure now could damage perception

- a) **Pressed in schedule: pre-blessings in June for EPS**
- b) **Fall-back is staff-time intensive: must have in place before "physics starts"**
- c) **We have a 20 TB/day day, but need > 10TB/day consistent load to prove.**

2) Service crashes may continue

- a) **The problem -may- be independent of platform**
- b) **We need at least 1-2 weeks of stable operation -at moderate load- to be sure.**
- c) **If continues: basically no time to fix this possibly kernel/java rte problem.**

3) User Experience does not meet "expectations" - other issues

- a) **Minor glitch or tuning issue - taking a conservative (user experience) position**
- b) **Heavy tape demand outside dCache or inside dCache - user access delays**
 - define public tests to use specific datasets known to be in cache
 - define cache sub-structure to preserve this over time
- c) **Heavy demand for files on disk - user access delays**
 - limited # client accesses per pool: leave large for now (max out bandwidth)
- d) **Pilot error - use of older CDF software: education**
- e) **Dismiss it before trying it: education, communication**
- f) **Unsupported access patterns, CAF-specific issues: define judgement criteria!**