CDF Data Handling Status: Summer 2003 Plan

Robert D. Kennedy (presenting),
Rick St.Denis, Frank Wuerthwein
for the CDF Data Handling Group

CDF Joint Physics Meeting
15 May 2003

⇒ **DH: Towards a "Services" View, Components Status/Work**

⇒ **CDF dCache - achieving full CAF-scale capabilities (done)**

⇒ **Summer DH Plan - baseline and fall-back, closing note**

⇒ ***Future Talk*: Beyond the Summer - we're planning now...**

DH: Towards a "Services" View

1) From Reaction to Pro-action: DH Goal = data services at specs => Define service levels and how these are monitored.

a) Use simple model to communicate needs, capabilities, resources, actual-use

b) "Golden data sets" - High availability, defined bandwidth maximum, <= 1 shift (8 hours) downtime per month, ... *whatever the implementation.*

c) "Volatile data sets" - DH does its best, but some staging from tape assumed.

d) Implementation details should not be focus. Move towards discussing *service*.

2) Improving/Exploiting User Experience w.r.t. Data Handling

a) cdfdh-announce@fnal.gov - low volume e-mail list for up/down times, etc.

b) DH Users web page - links to simplify "tracking" one's data in delivery.

c) Better central monitoring of CAF job stalling, etc, in the works.

d) Make DH plans accessible to users. Empower users during transitions.

e) Cannot please all the people, all the time... but user feed-back has been crucial.

3) Physics, DH Group Interactions: dialogue to help now/later

a) Physics to DH: Dataset priorities, refinement over time of published needs, ...

b) DH to Physics: Capabilities, statistics on actual use, service level established, ...

c) Not a one-time exchange, but continuous dialogue: DH liaison to Phys Groups, <=5 min DH/Computing status at regular (CDF weekly or Jnt Phys?) meeting

d) After "summer 2003 plans", start together on winter 2004 plans, long-term,....

Components: Status/Work

1) Enstore: Production. Some work underway. Ops issues lately.
   a) Major migration from 9940-A format to 9940-B. 3x capacity and I/O rate.
   b) Tape access contention often comes up, but tolerable (will improve).
   c) Operations: Errors on migrated-recycled tapes (OK now), network reliability

2) Static File Servers: Production. Ops issues few, fixed content.
   a) Limited disk space (25 TB). Access limited to 200 of the 600 total CPUs in CAF
   b) Basically, no new files since Jan/Feb 2003. Does not reflect DFC today.

3) DIM/Kahuna: Legacy. Ops issues at low level. Rootd in use.
   a) Limited disk space (15 TB), due to decrease for static disk to physics groups
   b) Occasional Kahuna restarts - memory consumption, content inconsistencies
   c) Cannot scale to Run 2 data volumes as is, *phase-out is planned* when sensible.

4) dCache: Final Beta. Challenged at full CAF scales. Ready now.
   a) Scales. Content reflects DFC, user requests. Low admin overhead. See next....

5) DFC: Production. Behind-scenes work on schema migration.
6) PNFS: Production. Dec/Jan ops issues hurt dCache, fixed now.

*) Many thanks to CD, DESY, ROOT for work/help/support.

dCache: Achieving CAF scale

1) Dec 2002/Jan 2003: dCache unreliable at full CAF scale.

a) dCache in production, ran acceptably well with O(100) CPU CAF-1

b) Early December, CAF scaled up. O(600) cpus hit ill-prepared dCache.

c) Users: < 50% sections succeeding. DH & task force burned out trying to fix.

d) CAF access to dCache disallowed (*but only CAF*) while problems unraveled.

e) Static file servers hurriedly assembled. Much work to load & make reliable.

2) The Road Back: Hardware, OS, dCache s/w, DHMods, PNFS,...

a) Every component of system had some problem - platform, dCache, client s/w.

b) Example: False RAID failures due to "idle logging" (Western Digital helped).

c) Problems also in "production" static file servers, much work leveraged.

d) dCache was never *off*. Like DIM: only option for some, best option for many.
The patience and perseverence of these beta-testers made this possible!

3) DCache Today: Problems resolved, load tests passed.

a) Professional-grade server as "admin node" fixes service crashes.

b) Configuration re-tuning, quick s/w fix from DESY: rate thottling resolved.

c) 50 TB capacity. 10 TB/day common. 20 TB read stably in one day **by users**.

d) Can be accessed by 50% of CAF now, full CAF soon with more file servers

e) The system is not "finished": optimization, tuning, tapeless datapath, etc.

dCache Bytes Read Per Day

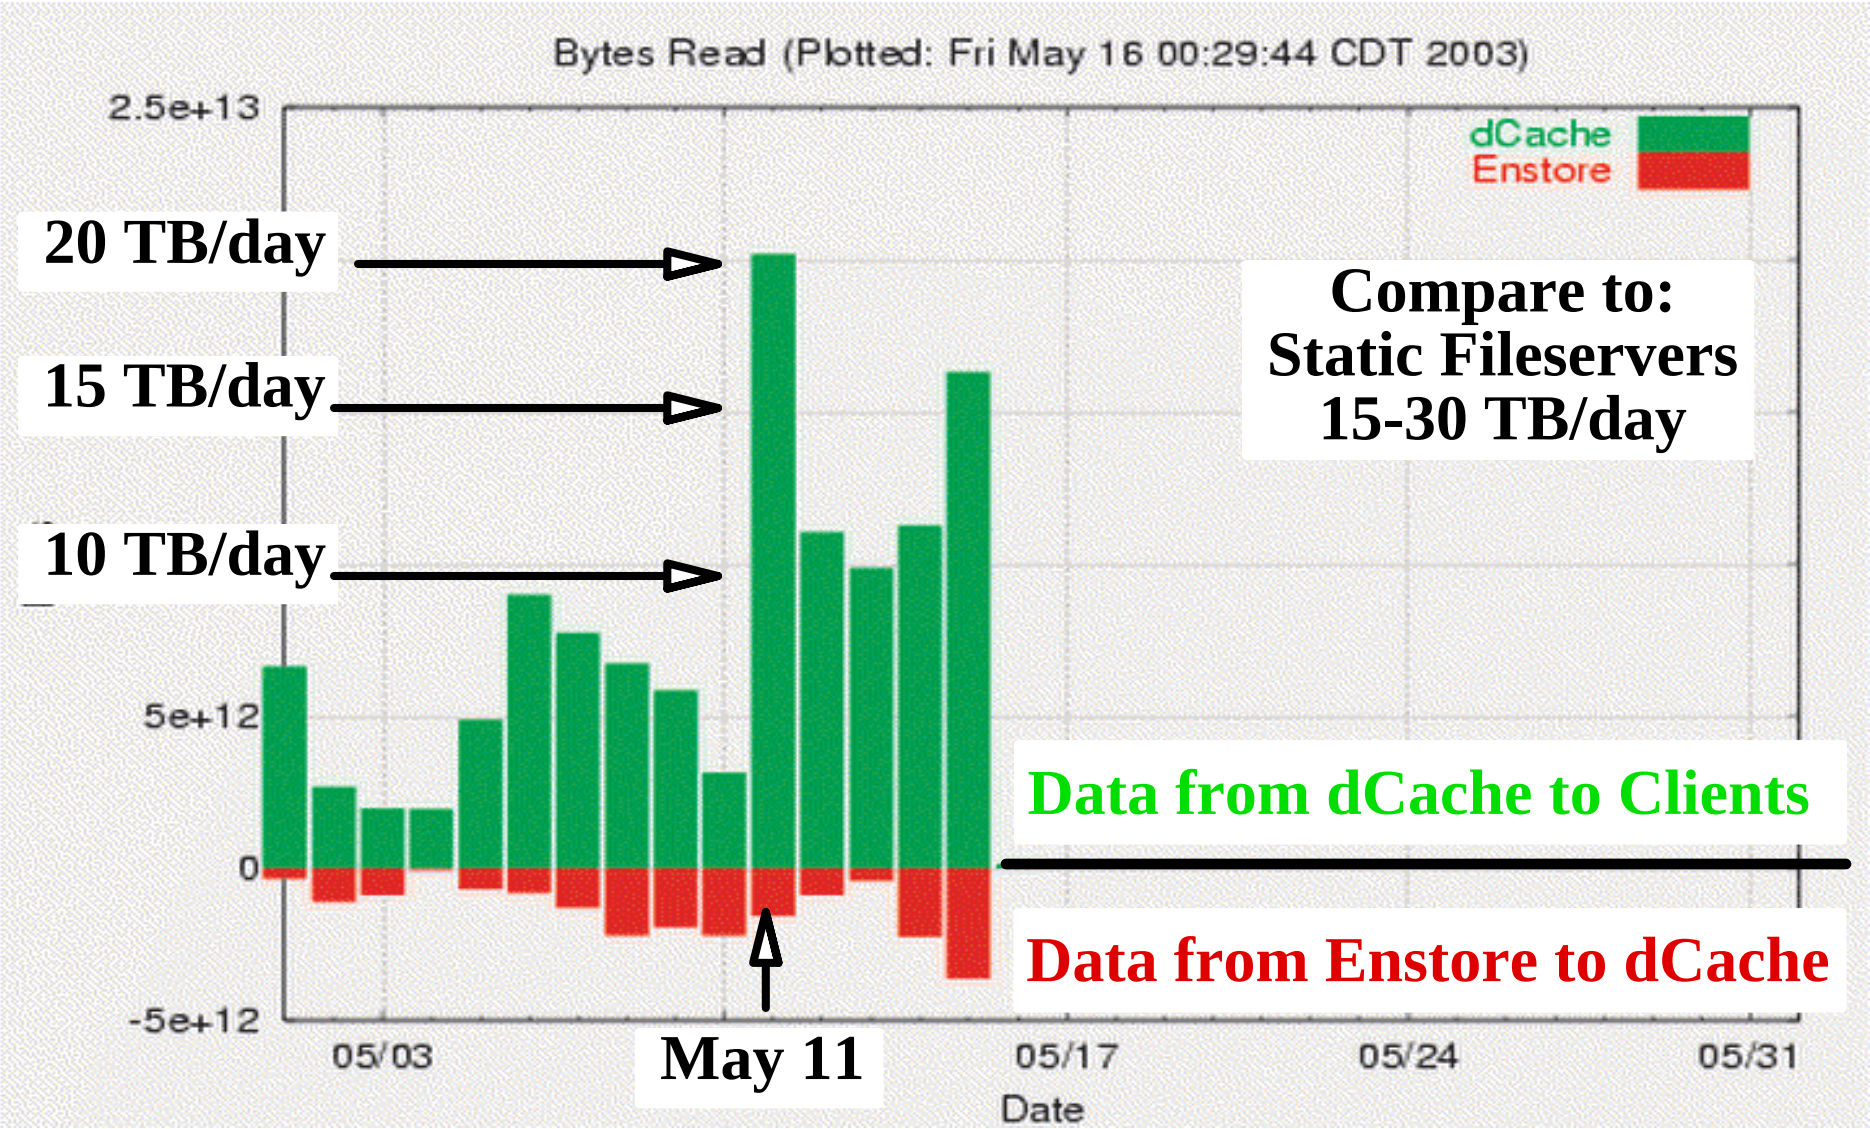

http://cdfdcam.fnal.gov:8090/dcache/outplot?filename=billing-2003.05.daily.brd.png

dCache Bytes Read per Hour

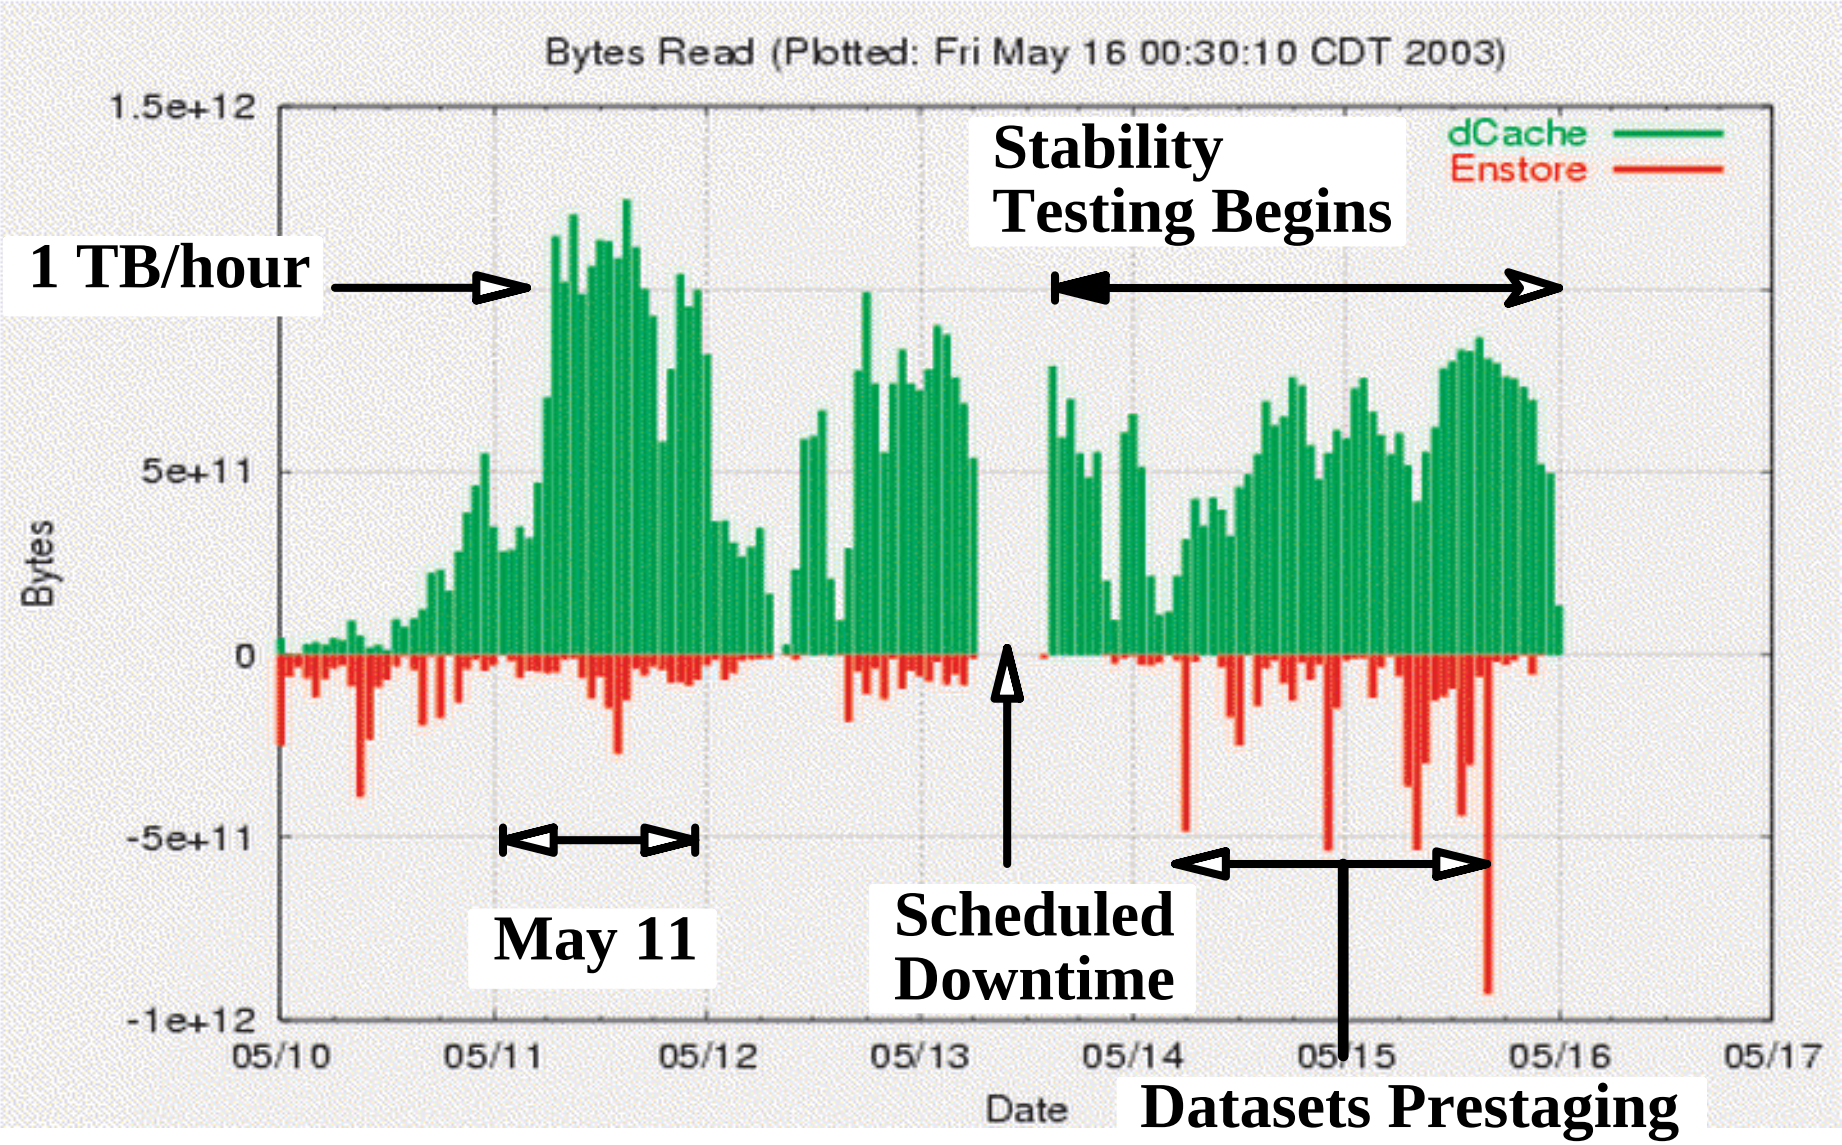

http://cdfdcam.fnal.gov:8090/dcache/outplot?filename=billing.week.brd.png

dCache: Use and Monitoring

1) Link/run current cdfsoft2: 4.9.1hpt3+, 4.10.1+, 4.10int1+, 4.11.0
   a) Must use improved CDF client interface to dCache.
   b) For CAF, must submit to short_dcache, medium_dcache, or long_dcache
   c) Most user reports of dCache errors nowadays are due to (a) and (b) failures.
   d) 4.9.1hpt3 = 4.9.1hpt1 + just the needed dCache and ROOT net I/O changes

2) AC++ TCL: "cache set Dcache" - CDF 5981 "Intro to dCache"

3) Sub-caches have been in use - RawStreamA, RawStreamOther
   a) Prevents cycling through raw data from displacing all other data
   b) Also limits bandwidth to raw data consumers to protect other data access

4) "Golden Pools" - Effective pinning of "golden datasets"
   a) Allow higher read bandwidth since only infrequent write from Enstore.
   b) Cannot be entire dCache. Must have adequate "volatile" space for other needs.

5) dCache Monitoring - http://cdfdca.fnal.gov - More extensive...
   a) Much content, CDF-centric interface page to "popular plots" in the works.
   b) Now in dcache: hbhd08 = 11 TB (heavy use), hbhd09 = 3 TB
      Large datasets already in use in dCache, and must be dealt in planning.
      Decide priorities as *interest expands to 08 and 09*, not just shifts from 08 to 09.

Summer DH Plan

1) Baseline: dCache (50TB->90TB) + existing static file servers

a) Lesson learned: Maintain "hot back-up" system. Smooths user transition.

b) After stability tests pass, add file servers, re-org data on disks (golden pools).

c) Production status declared: Expect ~75 TB, access by 75%->100% of CAF.
Static file servers remain as is (no space to add), access by 33% of CAF.

d) Later: Assimilate more file servers as available -> 90 TB (100% CAF load)

2) Fall-back: Replicate statics x 3 to allow full CAF access

a) Can only support 25-30 TB of different datasets with full CAF access this way.

b) Severely downscaled dCache used to serve low priority datasets "as-is".

c) Huge effort to *set-up* and *maintain*. May require help from Phys Grps....

d) Any 1 file server is a *single-point-failure* for system. Mostly just lucky so far....

e) Needs smaller event sizes to really work, *required* to expand any further.

3) Decision Point: Is dCache reliable enough for full CAF load?

a) Test: uninterrupted dCache service for at least a week, starting late 13 May.

b) Lure users to increase load: "winter golden" data sets into dCache (see URL)

c) Define criteria carefully: hiccups happen everywhere, service crashes cannot.

*)Whatever happens, Summer Solution in place by *06 June, 2003.*

http://www-cdf.fnal.gov/upgrades/computing/dh/cdfdh.html - under CurrentNews - Golden Datasets

Summer DCache Service Levels

1) dCache Overall Service Level: DH monitors and reports....

a) Capacity: 50 TB now, 75 TB soon, 90 TB by 06 June.

b) Downtime - planned/unplanned: <= 1 shifts (8 hours) per month planned.
Scheduled downtimes will be first Thursday of month, 24 hours notice.
<< 1 shift per month unplanned. Will be announced and tracked.

b) CAF access: scales with # file servers: 50 TB = 50%, 75 TB -> 100% of CAF.

c) Bandwidth sustained: roughly scales: 50 TB = 20 TB/day, 90 TB = 50 TB/day

d) Bandwidth peak: roughly scales: 50 TB = 400 MB/sec, 90 TB = 800 MB/sec

*) CDF clients consume data from 0.1 MB/sec to 5+ MB/sec. Difficult to model.

2) dCache Golden Service: Basically, alway in cache.

a) We start with winter golden dataset, alter per Physics Group input.

b) Capacity: undefined now, 25 TB soon, will consider <= 50% of cache space.
Capacity cannot exceed fall-back plan limits, until CDF "burns that bridge".

c) Bandwidth: Since no writes, reads limited by file server stable max I/O.

3) dCache Volatile Service: Best effort. Raw data = special case.

a) By definition: all non-golden datasets get volatile service.

b) Raw data in separate 10 pools (2/3 TB each) across 10 file servers (experience)

c) Must throttle read bandwidth a bit to allow stages from Enstore at full speed.

Summary, Closing Note

1) Summer Plan Baseline: dCache + existing static file servers

2) Final Pre-production dCache Test Now: Stability >= 1 week, under representative user read load, some staging from Enstore.

3) Fall-back Plan Defined: Replicated static fileservers

4) Whichever - in place and stable by 06 June for the summer

5) Meanwhile: improve/exploit communication, info exchange

To keep scaling: Commodity technology *.and.* event size reduction

a) DH system is under much lower load than will be coming in 3,6,12 months.

b) Users will get used to current service: easy access to 14 TB dataset on disk without the "bureaucracy" of resource prioritization in Physics Groups.

c) Technology can only take us so far due to the cost of conservative testing. Must test for weeks at new scales to prove the demanded reliability.

d) Event size reduction requires hard choices and hard work. Need to plan and test now to avoid a brick wall in the future.