Summary of the dCache workshop

Andrea Sciabà

WLCG Tier-1 Service Coordination Meeting
31/03/2011
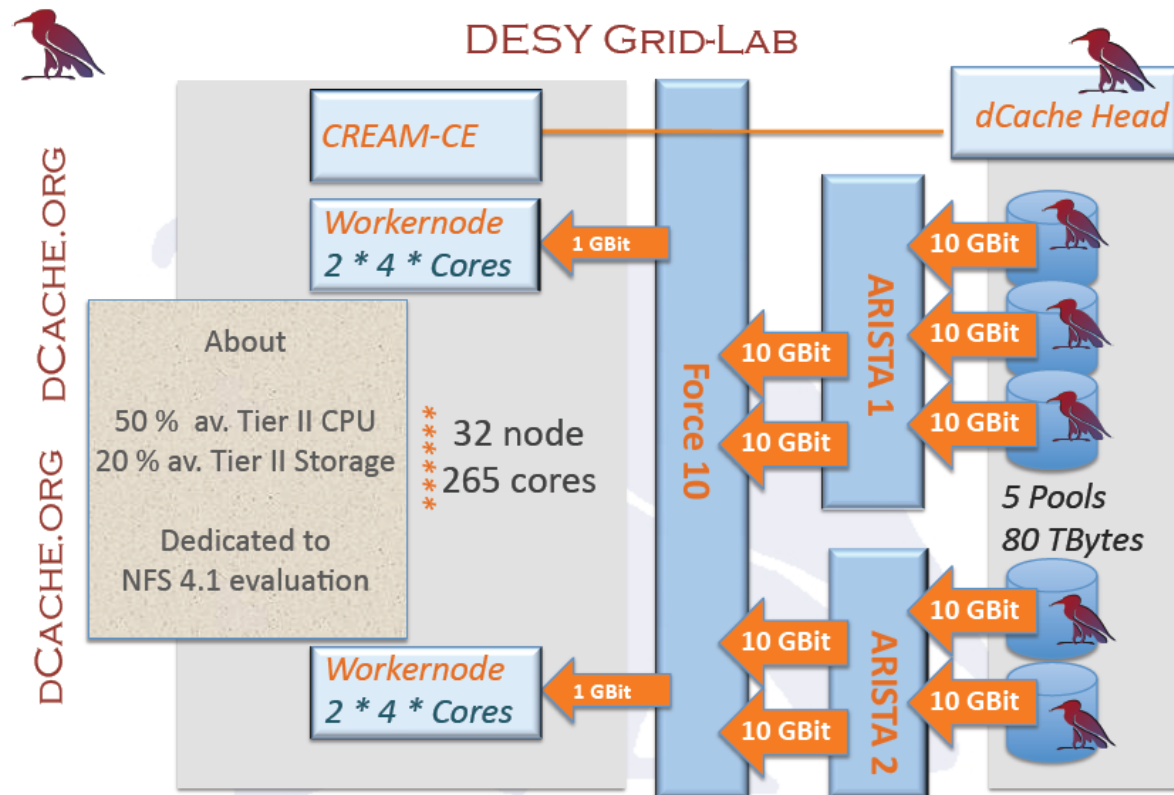
- New configuration

- $1.9.5 \rightarrow 1.9.12$ migration

- DESY Grid Lab

- Protocols

- SRM

- Other news

- Golden releases

- Much more **flexible**

- Much **easier** to understand

- Default values defined in */opt/d-cache/share/defaults* and well **documented**
  - An empty dcache.conf is valid

- Script to **migrate** from old to new configuration

- Script to **check** the configuration for errors without having to start the service

- Tutorial for the migration to 1.9.12
  - Procedure was well documented and easy, worked flawlessly
- Large sites can update first head nodes and then each pool, as a rolling upgrade
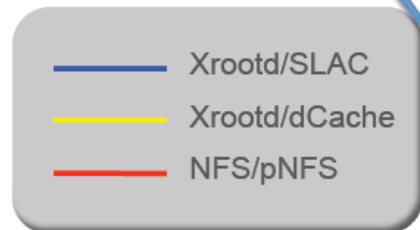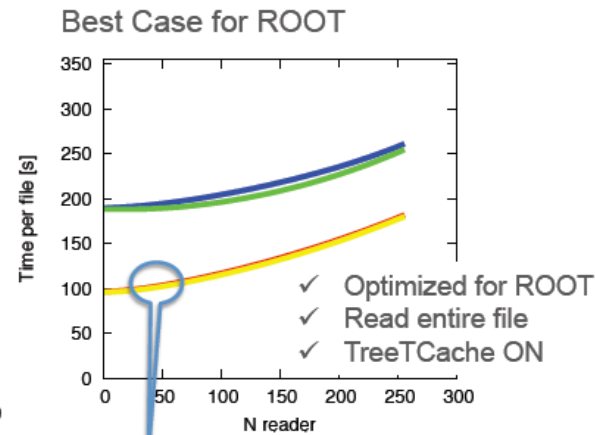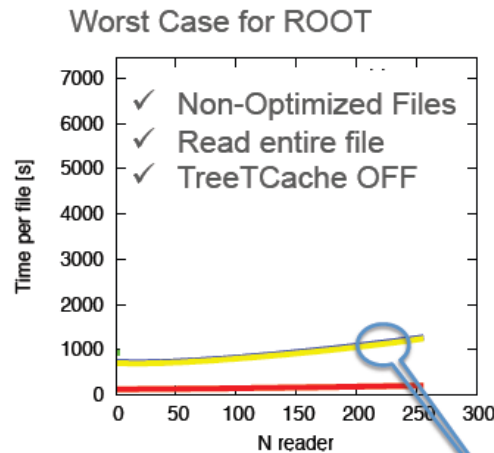
- Used to compare different protocols and different clients
  - ROOT, Hammercloud, CMS analysis

DESY GRID-LAB

DCACHE.ORG DCACHE.ORG DCACHE.ORG

CREAM-CE

Workernode
2 * 4 * Cores

1 GBit

About

50 % av. Tier II CPU
20 % av. Tier II Storage

***** 32 node
***** 265 cores

Dedicated to
NFS 4.1 evaluation

10 GBit
10 GBit

Workernode
2 * 4 * Cores

1 GBit

10 GBit

Force 10

ARISTA 1

10 GBit
10 GBit
10 GBit

ARISTA 2

10 GBit
10 GBit

dCache Head

5 Pools
80 TBytes

NFS 4.1

- pNFS not in SL6, but will be in SL6.2 (backported from kernel 2.6.38)
- DESY Linux pNFS kernel for SL5 already available
- Performance tests show it is (almost) never worse than xroot, often much better

Reading entire file.

Worst Case for ROOT

✓ Non-Optimized Files
✓ Read entire file
✓ TreeTCache OFF

Best Case for ROOT

✓ Optimized for ROOT
✓ Read entire file
✓ TreeTCache ON

For full file read, NFS behaves as good as SLAC/xRoot

If setting is bad for ROOT, SLAC and dCache xroot implementation behave the same. = The longer you are working on a file, the closer both implementations are.

— Xrootd/SLAC
— Xrootd/dCache
— NFS/pNFS

Mar 16, 2011

CERN**IT**
Department

- # Xroot
  - ## Several performance improvements
    - Java native IO for disk operations, block read-ahead, vector reads for ROOT files, …)
  - ## Tested using script from R. Brun on
    - Non-optimized ATLAS files (no basket ordering)
    - Optimized ATLAS files
    - Different TTreeCache sizes
    - Different # of branches
  - ## Results show only slightly slower than SLAC xrootd

- # Global namespace
  - ## dCache integration in B. Bockelman's Global Redirector schema via CMSD protocol

- Smaller memory footprint

- Supports SLL instead of GSI (prototype)

- Synchronous getTURL
  - Falls back to asynchronous if it takes too much

Other news

- **Migration module**
  - Very powerful, can be used to balance space usage, supports cksum validation, preserves pins…
- **dCache "Book" being completely updated**
- **gPLAZMA II**
  - Plugins, flexible configuration
  - Compatible with gPLAZMA1 configuration
  - Will map DN $\rightarrow$ UID and FQAN $\rightarrow$ GID, GID, …
  - Ready for ARGUS integration (e.g. for DN blackisting)
- **New logging system**
- **GLUE 2.0 supported**
- **Improved Pin Manager**

Golden releases

- 1.9.5: current golden release, supported until the end of the next LHC winter technical break

- 1.9.12: new golden release, from April to at least 2012

- Proposal: sites upgrade during next technical break

Conclusions
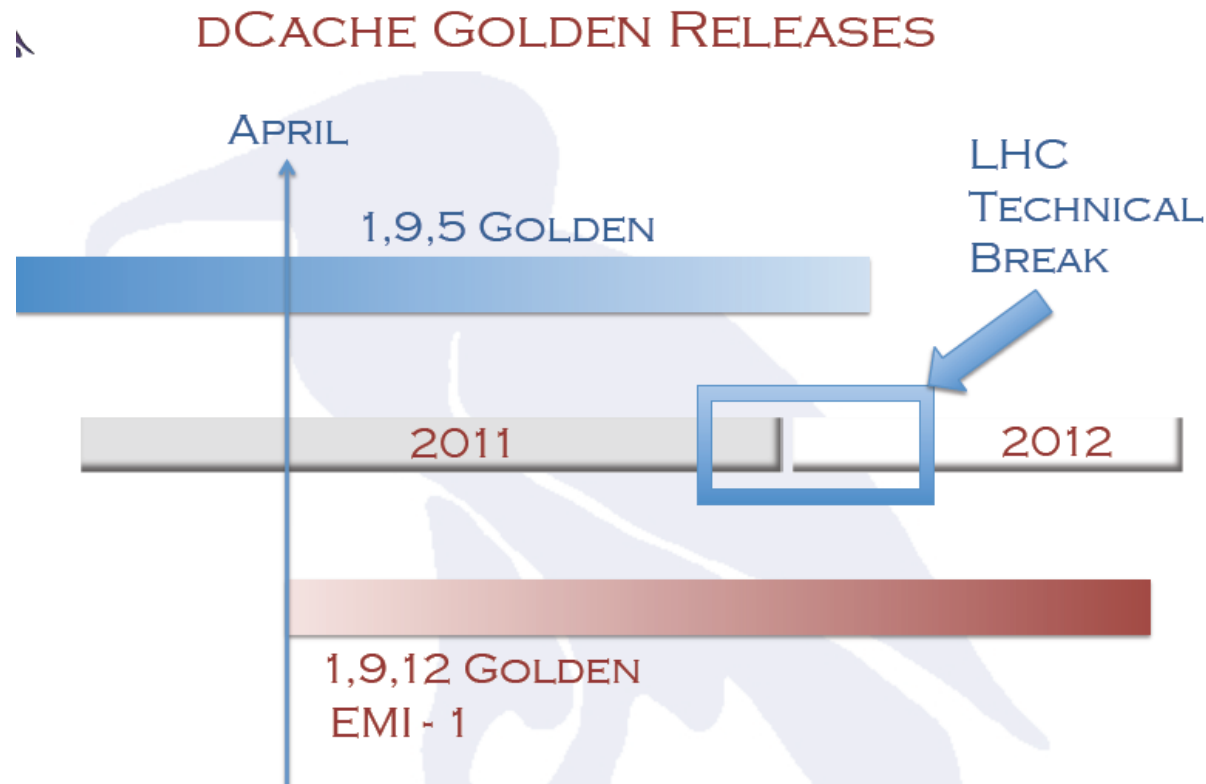
- An overview of what comes with 1.9.12

- All European Tier-1's present

- Several improvements from the administrator's point of view

- Very interesting performance tests

  – Choice of protocol is not really performance-driven

- The proposed goal is to migrate to it before the 2012 data taking