Report from CASTOR external operations F2F meeting held at RAL in February

Barbara Martelli
INFN - CNAF

Outline

- Face 2 Face workshop overview
- Site status reports
- Development status and plans
  - CASTOR
  - Tape
  - Monitoring
- Development agenda
- Open issues
- Load testing
- Conclusions

CASTOR workshop overview

- Held at RAL on February '09 with the following goals:
  - Exchange experience between sites currently using CASTOR about day-by-day operations and issues
  - Plan for 2009 about scaling up the deployments, number of instances, size of disk caches, tape infrastructure, s/w upgrades, migration to SLC5, usage of 32bit architecture
  - Describe and discuss the development plans for 2009 and beyond: planned releases, feature changes, time lines, support for the different software versions
  - Train the newcomers (several new entries in CASTOR staff at various sites)
- ~ 30 participants from CERN, RAL, CNAF and ASGC (via phone conference)

CASTOR Sites Reports

- Staff
  - CERN: CASTOR2 & SRM operations staff decreases from ~3 to ~2 FTEs. Tape team from 4 to 2FTEs
  - RAL: various changes in the staff
  - CNAF: CASTOR2 main expert has left and substituted by a new admin
- Hardware survey

SITE	Tape robots	Tape Drives	Tape Space PB	Disk servers	Disk space PB	Oracle servers
CERN	5	130	37	1050	6	6 per exp
CNAF	2	18 (soon 58)	3.1	45	3 (soon 6)	6
RAL	1 (soon 2)	7 dedicated + 32 shared	5k slots (soon 10k)	225	2.1	8
ASGC	2	8	1.3	53	2.3	5

CASTOR Sites Reports

- Software versions:

SITE	SRM	CASTOR	ORACLE	Tapeservers
CERN	2.7-15	2.1.7-23 and 2.1.8-2	10.2.0.4, nameserver 10.2.0.3, plan to move to 11	2.1.8
RAL	2.7-12	2.1.7	10.2.0.4	2.1.7
CNAF	2.1.7-12 (still SRM v1 for non LHC exp)	2.1.7-17	10.2.0.4	2.1.7
ASGC	2.7-12	2.1.7-19-2	10.2.0.3	2.1.7

- Monitoring in brief:

SITE	
CERN	Lemon, DLF, nagios, c2probe
RAL	Nagios, Ganglia, Cacti, TSBN, checkreplicas
CNAF	Lemon, Nagios
ASGC	Nagios, TLS, Ganglia, RRD base monitoring, C2probe

Databases

SITE	Oracle deployment	Stager deployment
CERN	three 2-nodes RACs per experiment plus a RAC for each service (nameserver, repack)	A stager per experiment
RAL	two 5-nodes RACs plus one single instance for DLF	a stager per experiment
CNAF	one 4-nodes RAC shared by nameserver, stager, dlf. One single instance DB for repack and SRM	A single stager for all experiment
ASGC	one 3-nodes RAC	A single stager for all experiments

- Time slot in the 3D meetings dedicated to CASTOR found very useful
- Working together in order to move all CASTOR databases to CERN 3D Enterprise Manager in order to have a single entry point for CASTOR DB monitoring
- Future development:
  - DB side connection pooling for SRM ?
  - Get rid of the optimiser hints in the code
  - Create SQL profiles via OMS and distribute as part of the Castor software ?

Development (Sebastien Ponce)

Task force from January to June 2008 reviewed the design choices of CASTOR and started 5 projects to address its weak points:

- File access protocols and data latency
  - CASTOR needs to support analysis activity: Small files, many concurrent streams, Mainly disk, with aggregated tape recalls, Low latency for file opening
  - The XROOT protocol and I/O server have been chosen to achieve this with CASTOR specific extensions
- Tape efficiency, repack
  - File aggregation on tape(e.g. Tar/zip files of 10GB)
  - New tape format: less tape marks, no metadata
  - Repack strategies
- Security
  - Goal is to ensure: every user is authenticated (authentication),every action is logged (accounting),Every resource can be protected (authorization),Ensure complete interoperability of Grid and local users
  - Kerberos 5 or GSI authentication for all client connections
- SRM and Database schema
  - Plan to combine the stager and the SRM software, Would allow also to merge the 2 databases
- Monitoring
  - Ease operation
  - Real time, automatic detection of pathological cases

DM

XROOT overview

- Client connects to a redirector node
  - This redirector finds out where the file is
  - It handles a cache of recent files for effici
- Client then connects directly to the node holdi

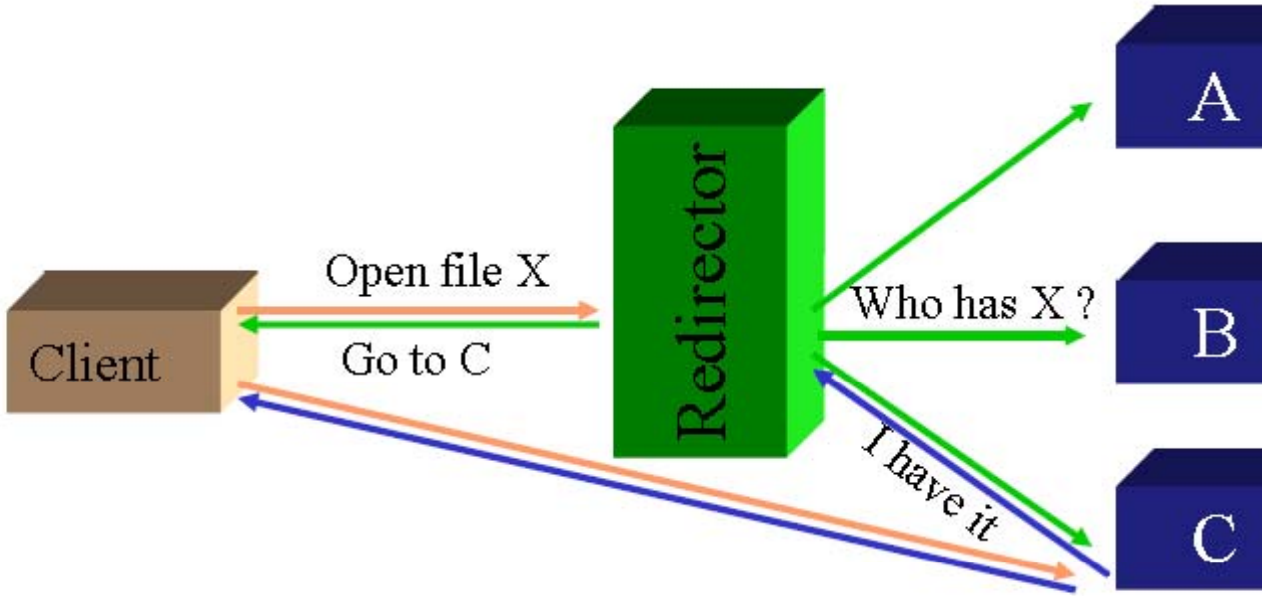

Disk S

DM

XROOT in CASTOR

- Client connects to a redirector node
- The redirector asks CASTOR where the file is
- Client then connects directly to the node holdi
- CASTOR handles tapes in the back

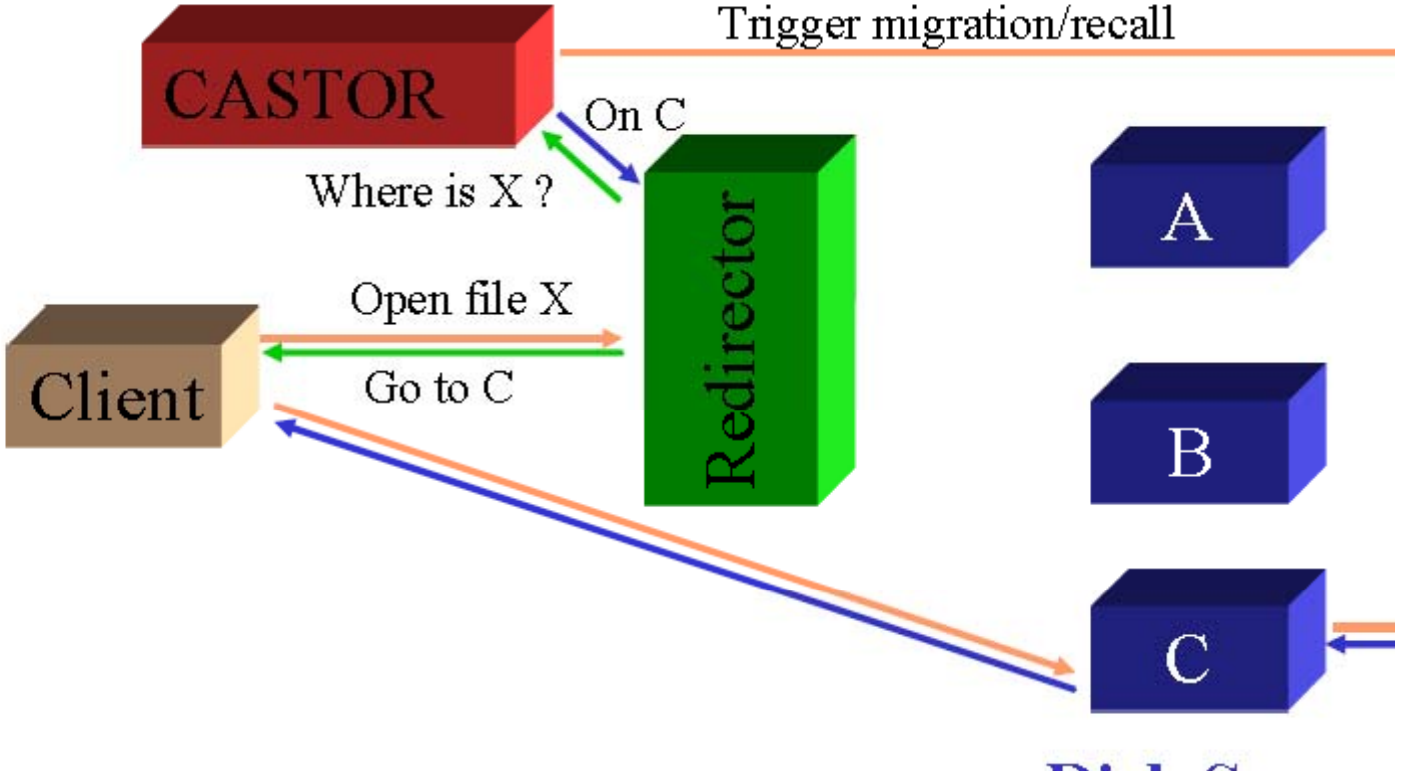

Tape Development (German Cancio's presentation)

- Recall/migration policies :write more data per tape mount:
  - Hold back requests based on the amount of data and elapsed time
- Writing small files on tape is slow, due to tape format: ANSI AUL format
  - 3 tape marks per file
  - 9 seconds per data file independent of its size
- New tape format reduce the metadata overhead:

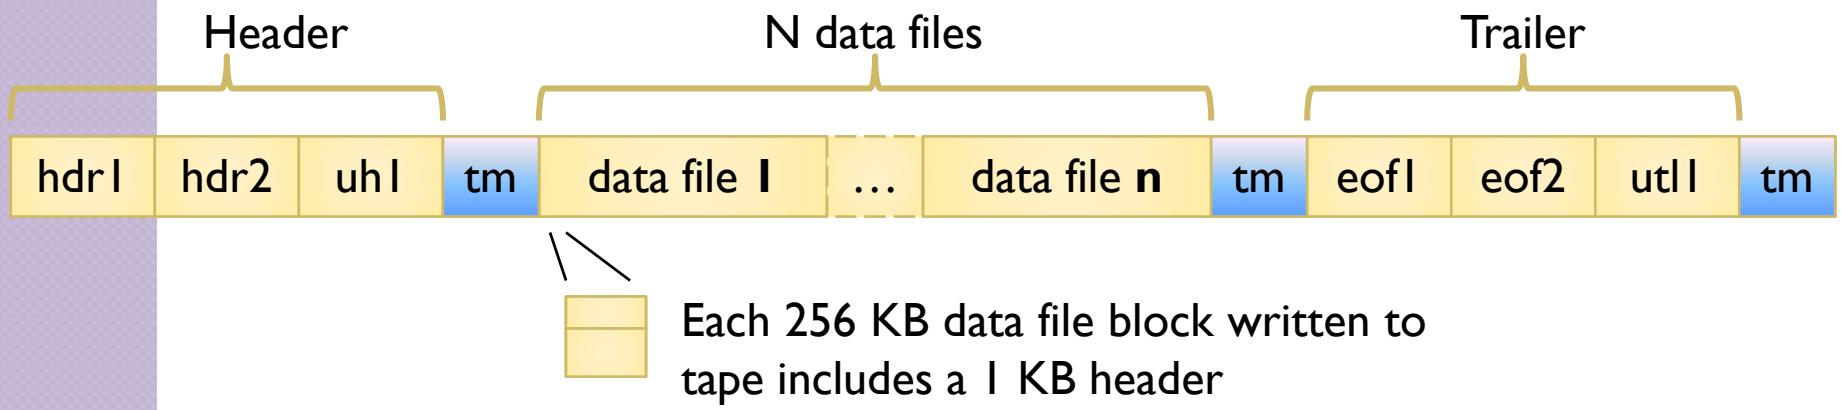

- Multi-file block format within the ANSI AUL format:
  - Header per block for "self description"
  - 3 tape marks per n files

Monitoring Development

- Monitoring CASTOR today requires to collect different information from many different tools
- Both CERN and RAL are putting a big effort on enhancing and consolidating existing monitoring tools
- RAL is starting from CMS monitoring code to incorporate information scattered at RAL but also through the 4 sites in a single view (see Brian Devies's presentation)
- New monitoring system developed at CERN is part of release 2.1.8-6
- Strategy: use the existing logging system (DLF) as the starting point, eventually improve and extend the existing log messages to contain any missing information (see Dennis's presentations)

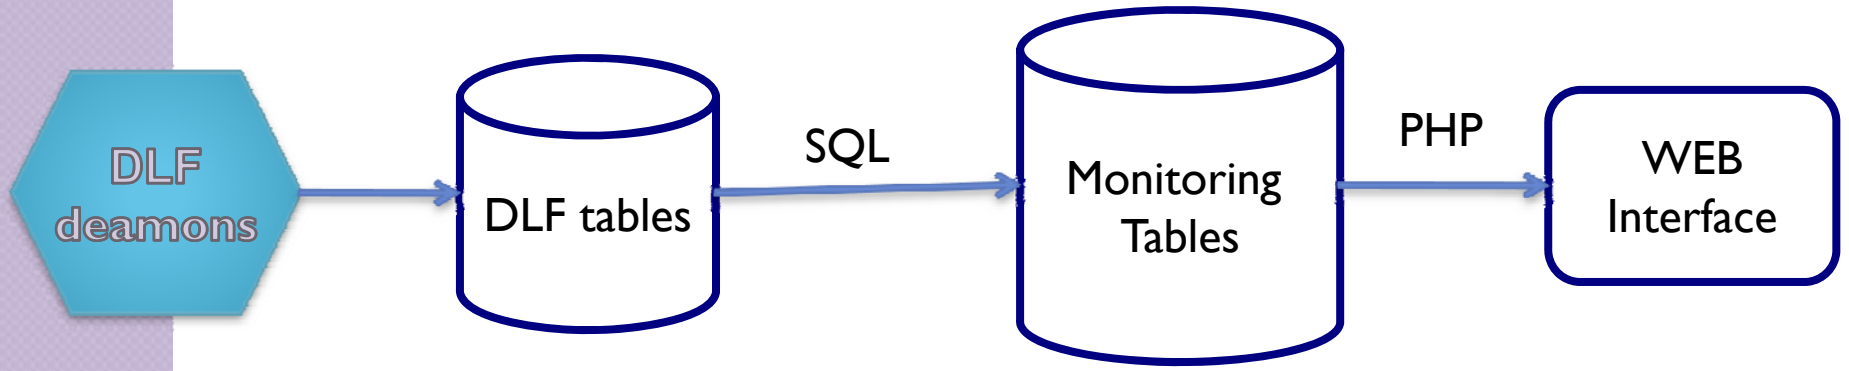

CASTOR Agenda: 2.1.7

- Is in maintenance mode
  - Stable, deployed in production at all sites
  - No new features included for quite long
  - Bug fix releases for major issues
- Will be phased out in the next few months at CERN
  - Repack and analysis setups already running 2.1.8
  - Plans to upgrade the other production instances in March/April
- Would not be supported anymore after the release of 2.1.9 according to current rules
  - i.e. sometime in Spring

CASTOR Agenda: 2.1.8

- Some of the new features:
  - Support for replication on close
  - Basic user space accounting
  - High level monitoring interface
  - Ordering of requests is now respected
  - Support for OFFLINE libraries
- Now stabilized: For ~2 months on CERN's repack and analysis setups
- Proposal of the CERN's operation team to build a 2.1.8-6 with important fixes/features backported and to switch to maintenance mode
- 2.1.8-6 should be available by end of February
- Deployed 2-4 weeks later on CERN production instances

CASTOR Agenda: 2.1.9

- Current development version (head of CVS) will include
  - Improved nameserver
  - Further xroot integration (write case)
  - Revisited build infrastructure
  - Ability to build only client/tape part
- Timelines not yet very precise
  - Spring/Summer 2009
- Deployment before LHC startup unlikely for T0 with current LHC scheduler

Load tests (Dennis)

- Certification setup at CERN based on virtual machine and aimed at certifying CASTOR components functionality
- Many test scripts exist (/afs/cern.ch/project/castor/stresstest)
  - Heavily customized to the CERN environment.
  - Each test should run for 24-48 hours.
- Requires expert knowledge
- Not all elements of CASTOR are tested
- Tests are customized for each CASTOR version
- RAL is thinking of creating a test-bed in order to try out the new CASTOR release in a T1 like environment

Some of the Main Open Issues

- Hotspots: server goes into high IO wait and delivers (almost) no data
  - Can only recover by killing off all RUNning requests (i.e. set server state to DISABLED)
  - Observed correlations to specific RAID configurations -> reducing RAID units shows higher latencies and severely limits I/O performance (see Ignacio talk)
- Very big ID inserted sometimes in the id2type table (Aka BigIDs problem):

```
SQL> select count(*) from id2type where id>10E18;
COUNT(*)
----------
65691
```

  but no "no type found for ID". The daemon is keep trying to process the request in *status=0* which don't have Id2Type.
  - Solution via hand written receipt from Shaun De Witt, but a more stable solution is needed
- Occasional unresponsiveness from JobManager for 2-3 minutes:
  - delay with jobs reaching the job manager from the stager
  - delay with jobs reaching LSF
- Oracle unique constraint violations in RH
- Possible crosstalk between atlas and lhcb stagers (Oracle Bug)

Conclusions

- CASTOR workshop useful to exchange experience between sites in administering and facing the new problems which arise
- One of main topics of discussion is upgrades of CASTOR software releases: non CERN (2.1.8 upgrade or not?)
- Significant effort is being put on development of monitoring integration tools, tape efficiency and core s/w