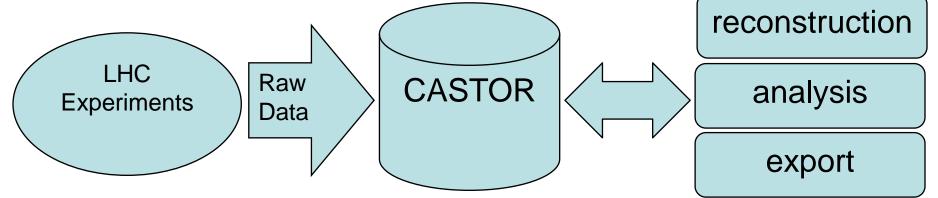# DM
Data Management

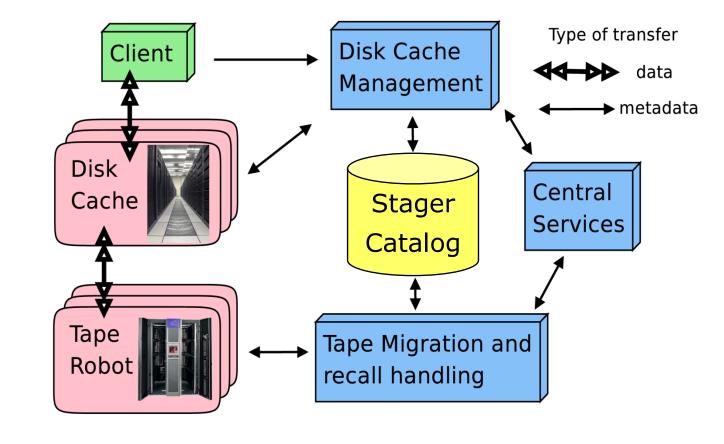# CASTOR end-to-end monitoring system

Witold Pokorski, Theodoros Rekatsinas[1], Dennis Waldron, Dirk Duellmann, Sébastien Ponce, Bartolomeu Rabaçal, Jacek Wojcieszuk

[1]Department of Electrical and Computer Engineering, National Technical University of Athens, Greece

CERN IT Department

The CERN Advanced STORage system (CASTOR) is to be used by the LHC experiments for storing experimental data.

CASTOR architecture is based on the disk cache which serves to transfer data from and to the tapes. The system is able to serve thousands of requests per second, while managing millions of files.
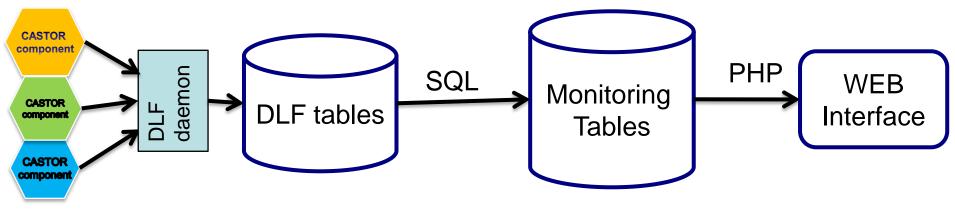
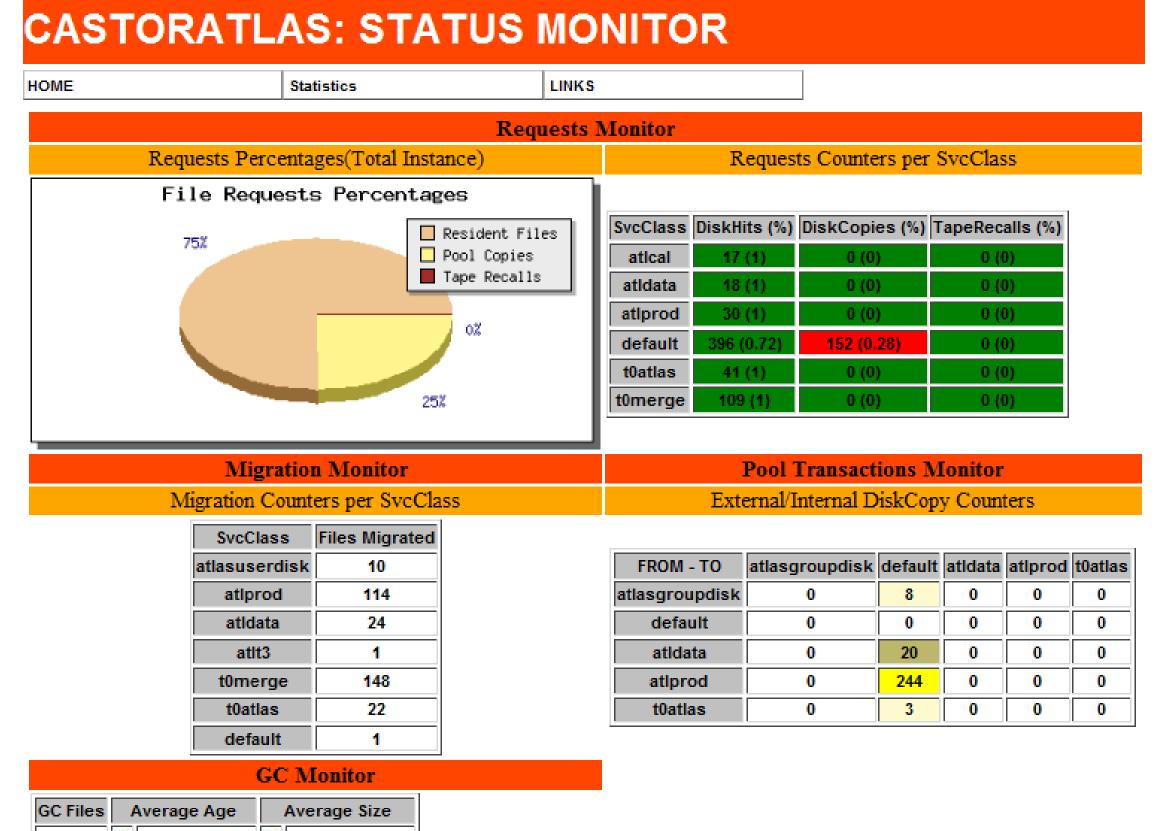Due to the complexity of the system, an end-to-end monitoring system is needed
• for users to check the current status
• for operators, to diagnose any problem
• for developers to observe the effectiveness of different algorithms and policies

General Architecture

Our new high-level CASTOR-dedicated monitoring system consists of a Database back-end for event processing and storage and a Web Interface for visualization. The stream of log messages, generated by the different CASTOR components, is used as input. Those raw messages are initially collected by the Distributed Logging Facility (DLF), and then combined into a high-level, compact monitoring data.

• PL/SQL monitoring procedures pre-process incoming DLF messages accumulating and storing information for every type of event
• Monitoring tables correspond to specific event oriented metrics
• Web interface provides aggregated visualization (histograms, pie-charts, etc) of the monitoring data

CASTORATLAS: STATUS MONITOR

HOME	Statistics	LINKS

Requests Monitor

Requests Percentages(Total Instance) — Requests Counters per SvcClass

File Requests Percentages
- 75%
- 6%
- 25%

Legend: Resident Files, Pool Copies, Tape Recalls

SvcClass	DiskHits (%)	DiskCopies (%)	TapeRecalls (%)
atlcal	17 (1)	0 (0)	0 (0)
atldata	18 (1)	0 (0)	0 (0)
atlprod	30 (1)	0 (0)	0 (0)
default	396 (0.72)	152 (0.28)	0 (0)
t0atlas	41 (1)	0 (0)	0 (0)
t0merge	109 (1)	0 (0)	0 (0)

Migration Monitor — **Pool Transactions Monitor**

Migration Counters per SvcClass — External/Internal DiskCopy Counters

SvcClass	Files Migrated
atlasuserdisk	10
atlprod	114
atldata	24
atlt3	1
t0merge	148
t0atlas	22
default	1

FROM - TO	atlasgroupdisk	default	atldata	atlprod	t0atlas
atlasgroupdisk	0	8	0	0	0
default	0	0	0	0	0
atldata	0	20	0	0	0
atlprod	0	244	0	0	0
t0atlas	0	3	0	0	0

GC Monitor

GC Files	Average Age	Average Size
112	625.55 (0.3504)	5668.648 (0.977)

Monitoring DB schema

The Database Schema is composed of:

➢ Monitoring Tables, where each entry contains the full descriptions of the respective CASTOR event

➢ Materialized views, which implement CASTOR status monitors

Monitoring tables (DB) isolate the front-end (web interface) from the back-end (transport for log messages) offering design flexibility and the possibility to:

➢ switch system input from DLF transport to SYSLOG

➢ feed the data into other existing general monitoring systems like LEMON (LHC Era Monitoring) and similar

➢ implement several web interfaces simultaneously
•summary dashboard for experiments
•detailed monitoring for developers and operators

PL/SQL Procedures

PL/SQL procedures run periodically every 5 minutes to extract the information about significant CASTOR events from the incoming data stream. Some of the monitoring information comes directly from individual log messages (e.g. facility errors) while for the monitoring data like
• file read requests
• file migration to tape
• file garbage collection
one needs to combine multiple daemon logs and store them as a single entry.

CASTOR monitoring web interface
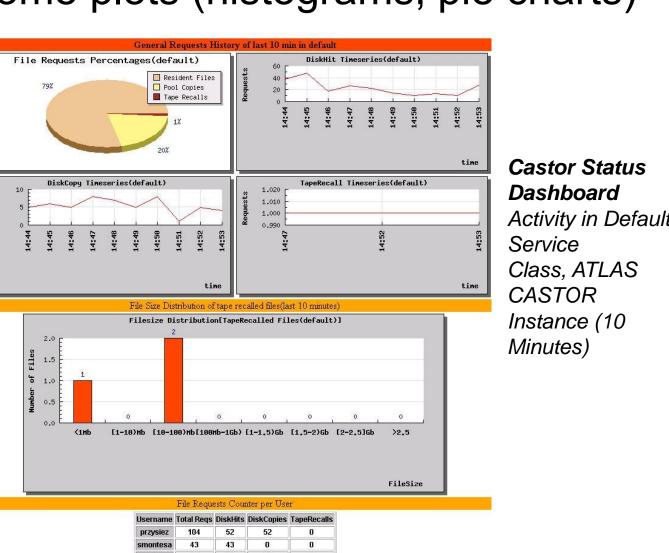
The Web interface consists of two different levels
•Current Status Dashboard
•Statistical Information

➢The Dashboard feature uses tables with number indicators to present system's current status. Some plots (histograms, pie charts) are used to display short-term historical information.
Source: DB Materialized Views

➢More detailed statistical and historical Information are in place to monitor CASTOR activity in time. Simple queries containing group, aggregate or analytic functions implement the desired metrics, which are displayed using JpGraph libraries.
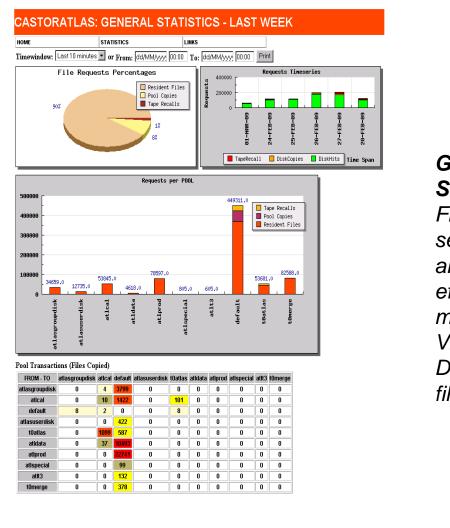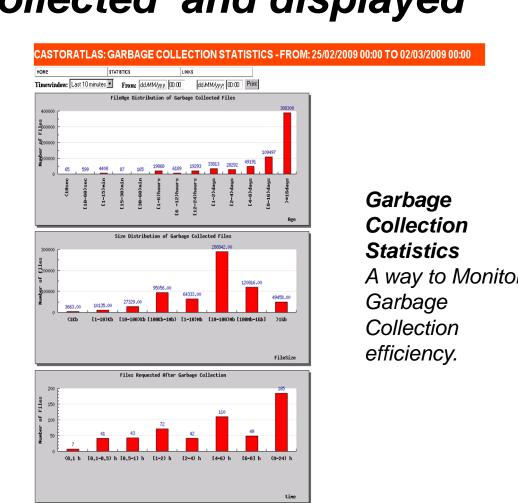In order to serve multiple user requests, retaining DB load to a minimum, JpGraph's built-in caching mechanism is used to cache images for repeated requests.

Castor Status Dashboard
Activity in Default Service Class, ATLAS CASTOR Instance (10 Minutes)

Examples of different metrics collected and displayed
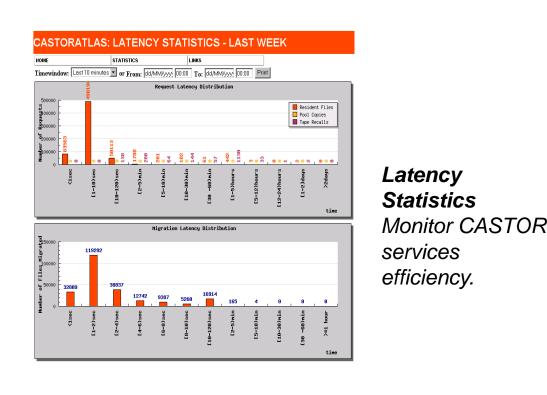
General Activity Statistics
File read request time series and disk cache efficiency metrics are displayed. Visualization of Disk-to-Disk file copies.

Garbage Collection Statistics
A way to Monitor Garbage Collection efficiency.

User Specific Metrics
Monitor user activity in time.

Latency Statistics
Monitor CASTOR services efficiency.

Performance and Stress Tests

To ensure that the performance requirements are met, the new monitoring system has been extensively tested using simulations of real life scenarios.
✓ PL/SQL procedures were tested by processing real CASTOR metadata. Their scalability was tested by replaying past data in higher speed (factor of 10). A C++ program was implemented to reconstruct daemon logs, compress and inject them into DLF.
✓The web interface was tested by simulating load scenarios with WAPT, a load, stress and performance testing tool for web sites.