
The ATLAS Tier-0 Software Suite

Armin NAIRZ
Luc GOOSSENS
(CERN/PH/ATC)

Presented by
Dario BARBERIS
(Genoa Univ./INFN and CERN/PH)

September 5, 2007

CHEP2007, Victoria/Canada

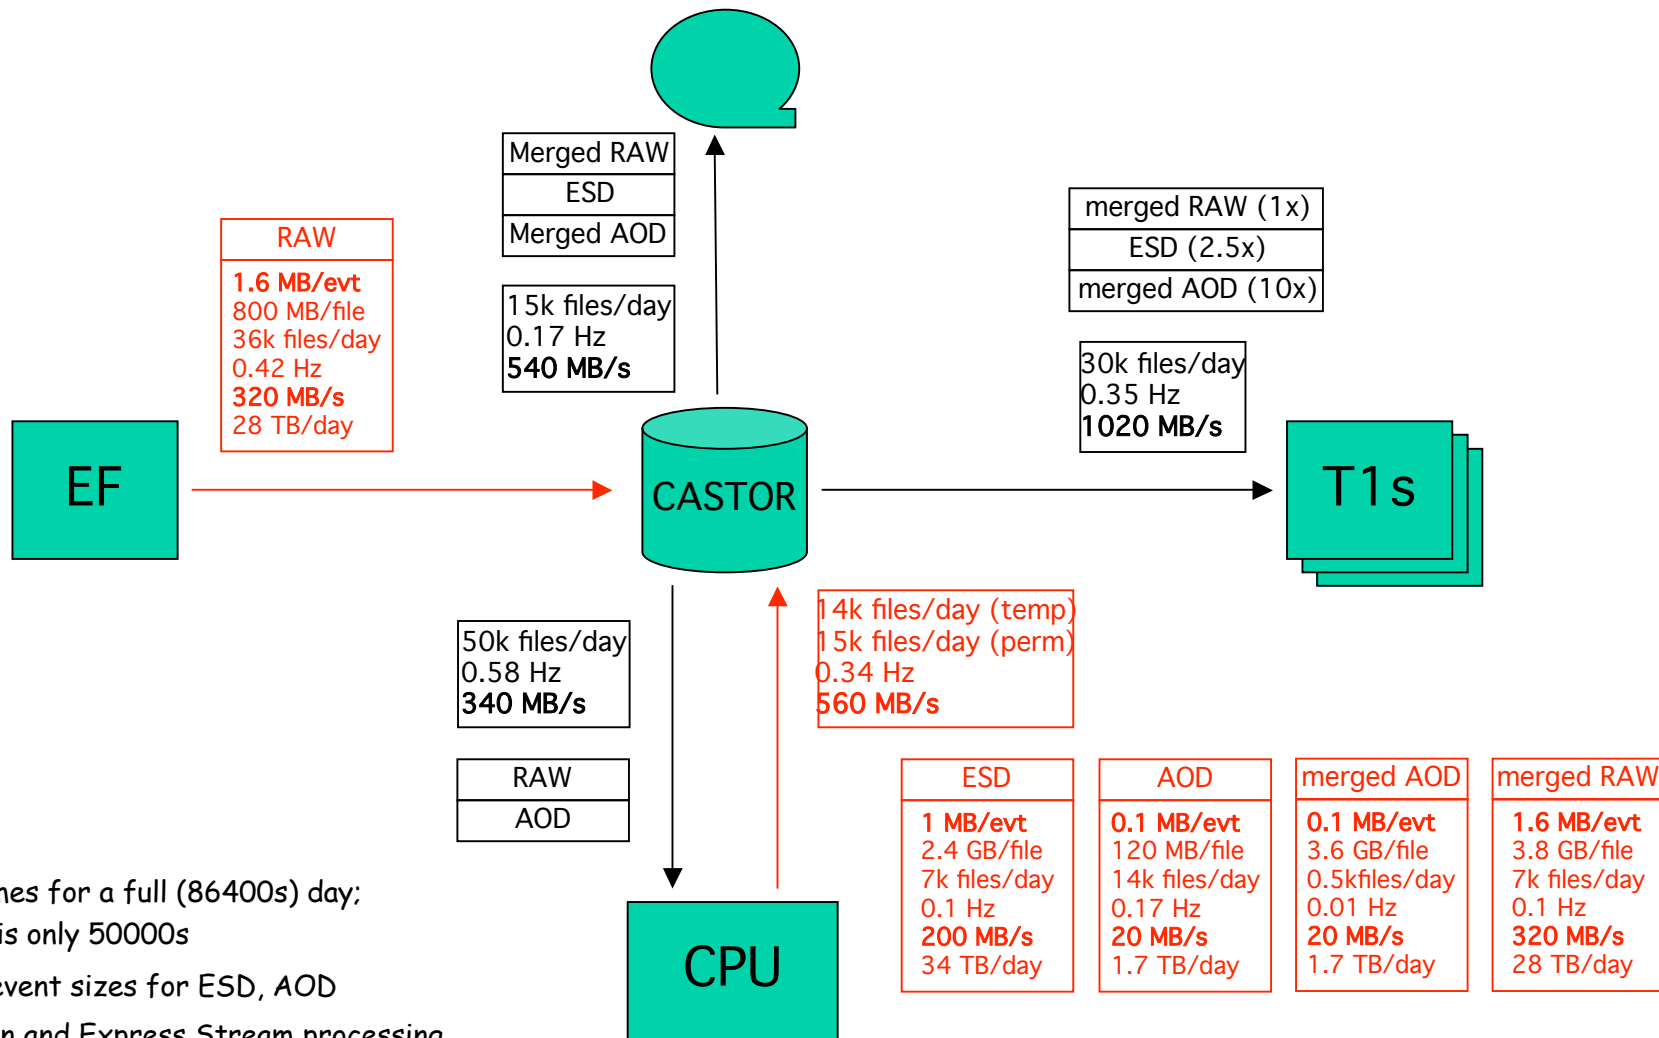


-
- Introduction
 - Tier-0 software suite
 - Design guidelines
 - Shared ProdSys components
 - Tier-0 proper components
 - Tier-0 scaling tests in 2006-2007
 - Summary and outlook

- Tier-0 functions

- Calibration and alignment processing
- First pass ESD/AOD and TAG production (= reconstruction)
- Archiving of first pass ESD/AOD/TAG on tape
- Uploading of TAG files into a TAG database
- Distribution of RAW data and first pass ESD/AOD/TAG to Tier-1s
 - One copy of RAW
 - Two copies of ESD (plus one full copy to BNL)
 - 10 copies of AOD/TAG
- Offline Data Quality Monitoring (oDQM)
- Etc.

-
- Tier-0 nominal scale
 - $O(10K)$ jobs per day
 - $O(10K)$ permanent files per day
 - RAW, ESD and merged AOD
 - $O(100K)$ temporary files
 - Unmerged AOD
 - Disk writing 880 MB/s
 - Disk reading 1900 MB/s
 - Tape writing 540 MB/s
 - Tier-1 export at 1020 MB/s
 - Approximately 3000 recon jobs in parallel for reconstruction

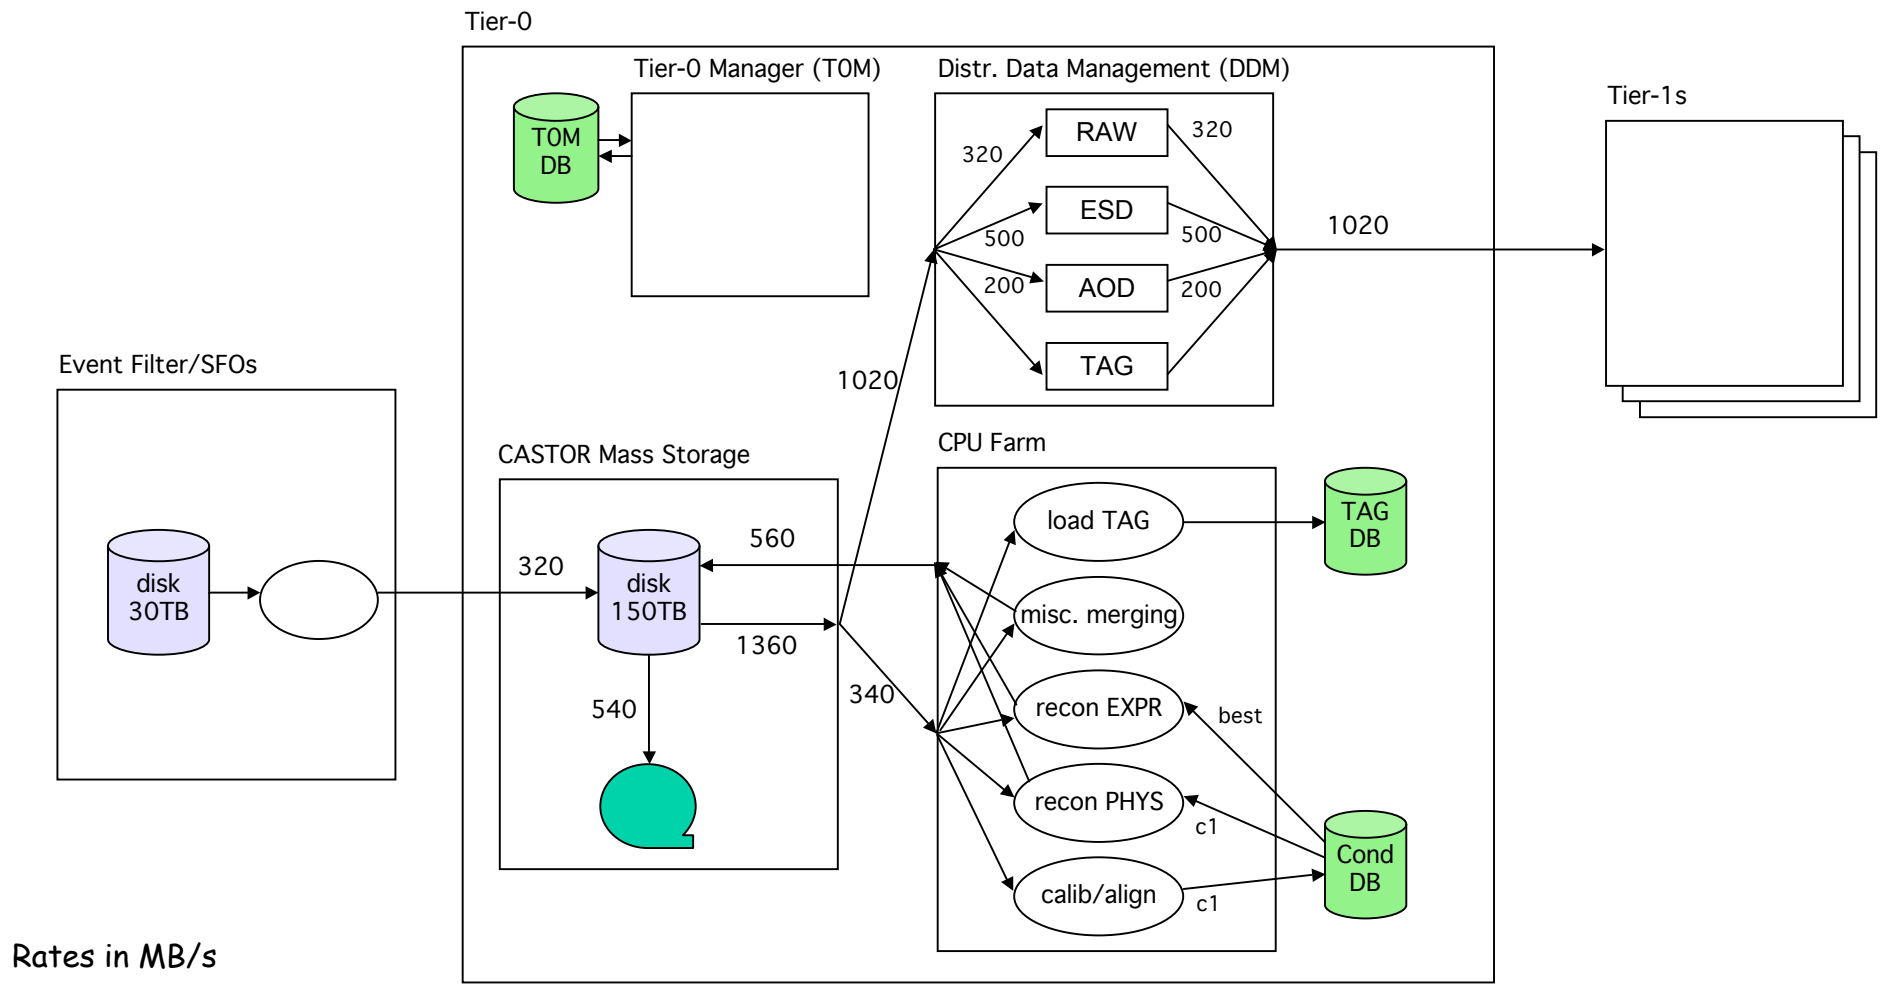


(*) Remarks:

- Data volumes for a full (86400s) day:
"CM day" is only 50000s
- "Target" event sizes for ESD, AOD
- Calibration and Express Stream processing add about 20% to the I/O rates

- Tier-0 architecture

- Event filter (EF) sub-farm output managers (SFOs) write RAW data into CASTOR
- Tier-0 manager (TOM) defines and orchestrates all Tier-0 jobs
 - Data driven
 - Reconstruction jobs, AOD merging jobs, TAG uploading jobs, calibration & alignment jobs, ...
 - Jobs run on LSF batch farm at CERN
- TOM enters Tier-0→Tier-1 export tasks into instance of ATLAS distributed data management system (DDM/DQ2)



- TOM architecture

- Based on ATLAS production system (ProdSys)
 - ProdSys is well-tested and reliable/mature
- ProdSys architecture features a 'facility' neutral supervisor component and 'facility' specific executors
 - Standard supervisor Eowyn is used
 - Custom Tier-0 executor was written interfacing to LSF but also exploiting the peculiarities of the Tier-0 setup
- TOM only needs to define jobs based on the data, ProdSys takes care of running the jobs on LSF

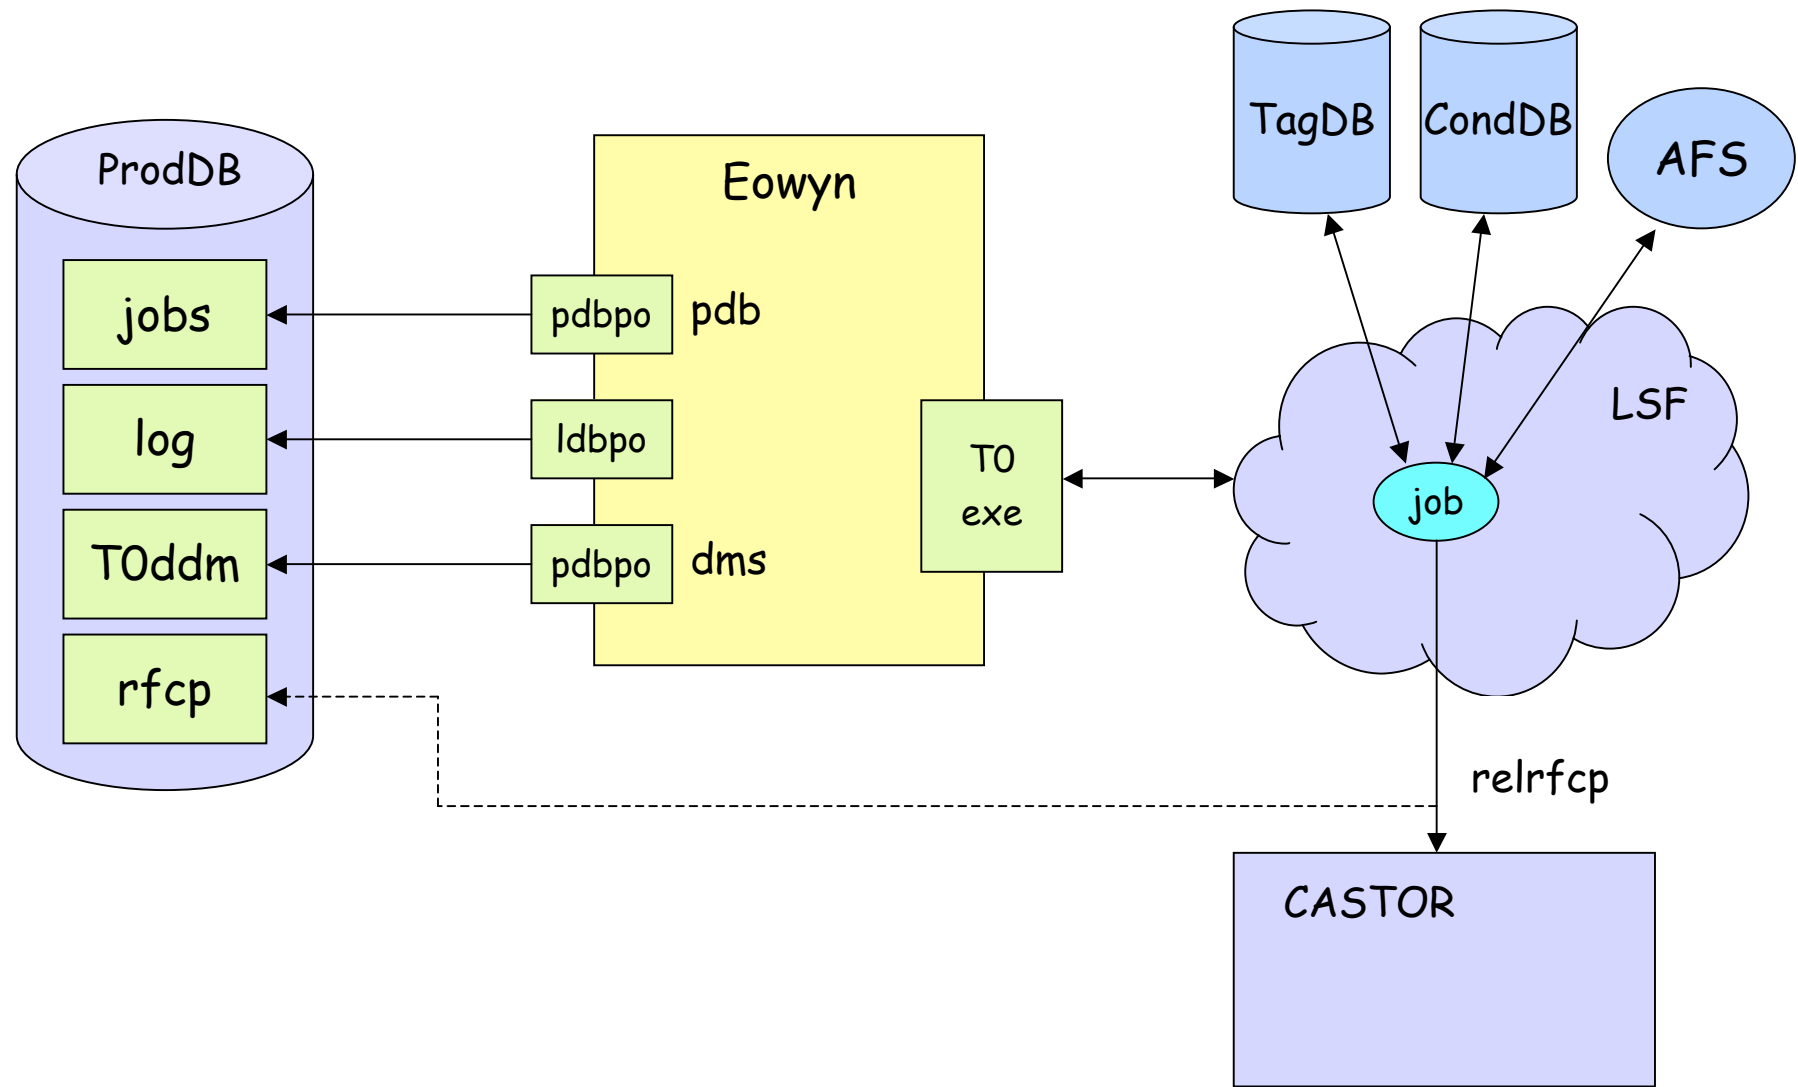
-
- Introduction
 - Tier-0 software suite
 - Design guidelines
 - Shared ProdSys components
 - Tier-0 proper components
 - Tier-0 scaling tests in 2006-2007
 - Summary and outlook

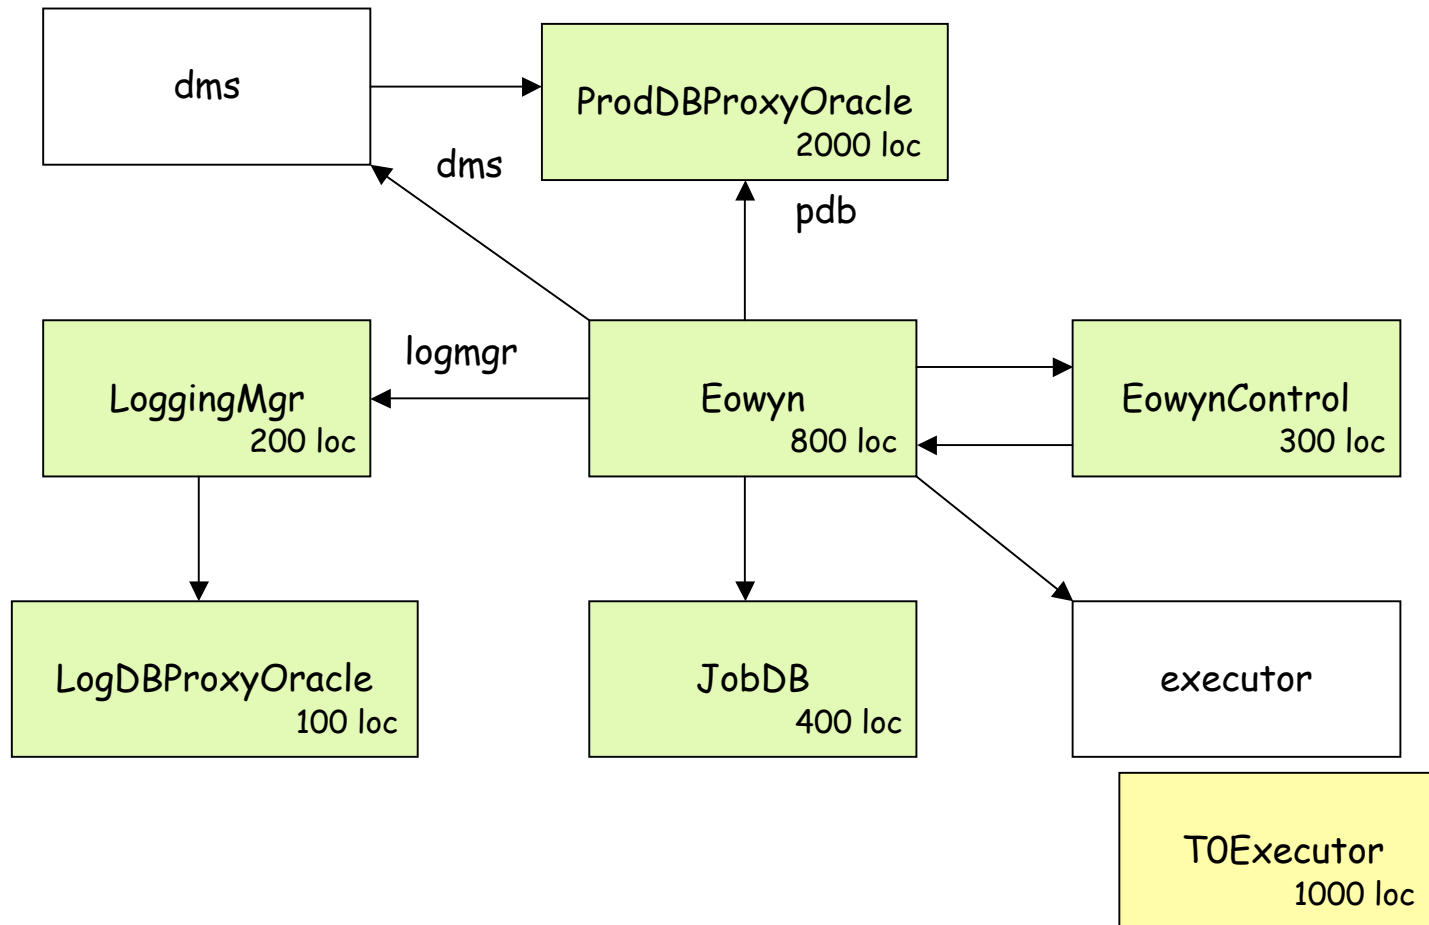
-
- Design guidelines
 - Keep it as simple as possible
 - Minimal external dependencies
 - Services
 - AFS, LSF, CASTOR
 - Code
 - Oracle Instant Client, cx_Oracle
 - Decoupling and modularity
 - Robustness through redundancy
 - Double AFS structure

-
- Introduction
 - Tier-0 software suite
 - Design guidelines
 - Shared ProdSys components
 - Tier-0 proper components
 - Tier-0 scaling tests in 2006-2007
 - Summary and outlook

■ History

- Architecture proposed in Aug 2003 and unchanged since
- Evolutions
 - ProdDB
 - Moved from shared Oracle PDB service to ATLAS RAC
 - Only recently migrated through first non-trivial schema change
 - Supervisor
 - Retired Windmill implementation in favour of Eowyn implementation (mid 2005)
 - Dropped Jabber communication in favour of direct Python API
 - Data management system
 - Migrated from DQ to DQ2
 - Still ongoing
 - Executors
 - Lxor (LCG)
 - Dulcinea (ARC)
 - Capone (OSG) retired in favour of PanDA

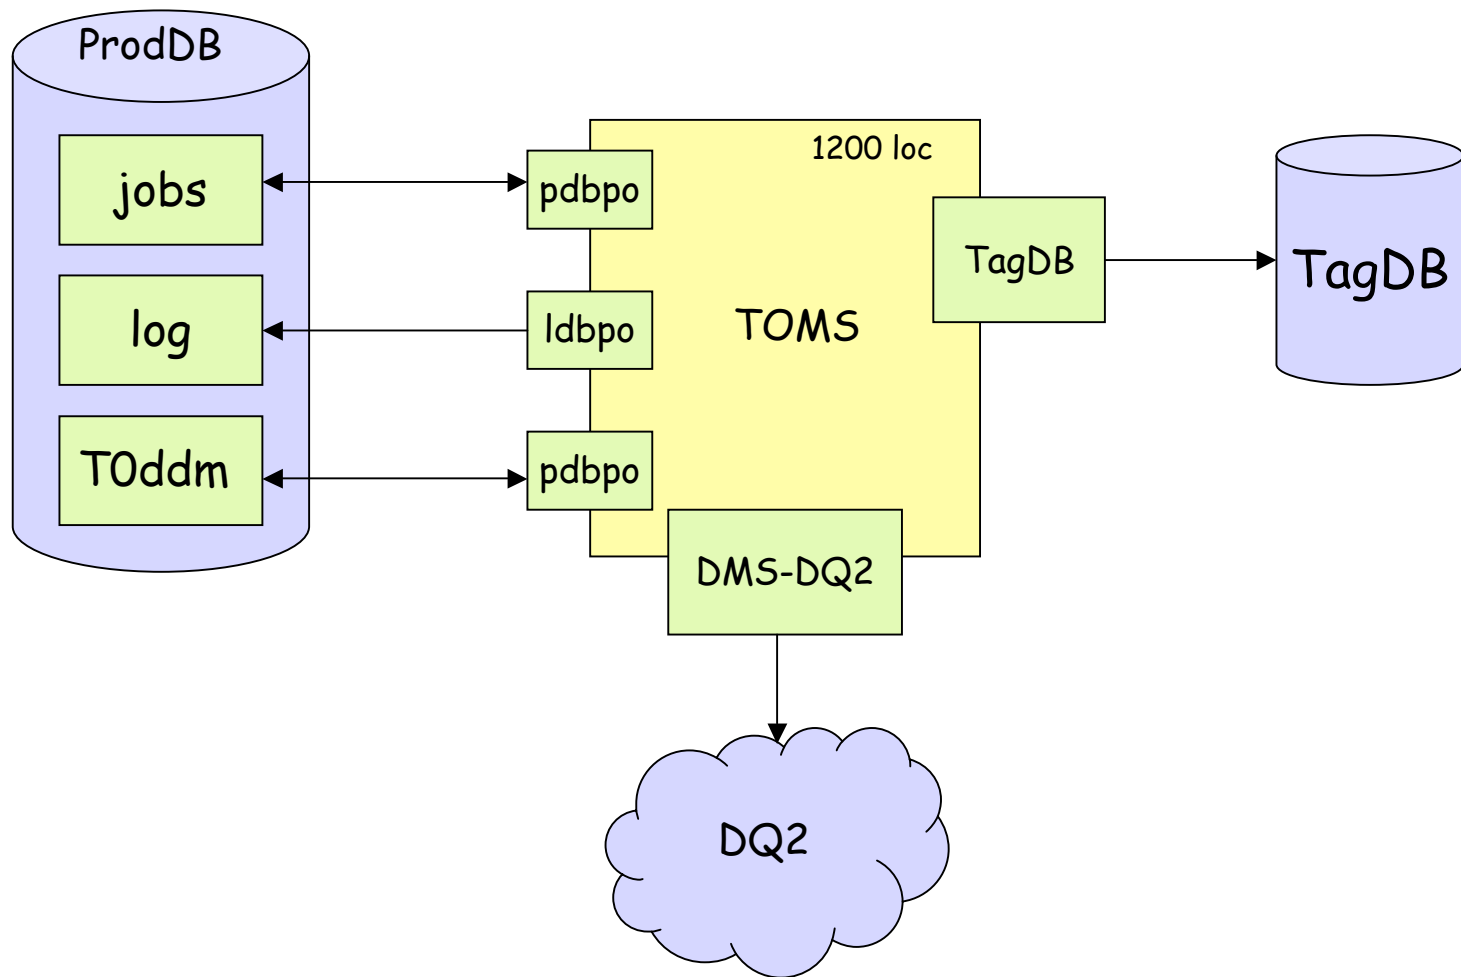




-
- Eowyn
 - Bulk state processor
 - All jobs move through their state diagram
 - Grouping of operations in transverse dimension
 - Coarse interleaving of operations
 - Get status of all active jobs, find jobs to be done, ...
 - Multi-threading does not work very well in Python
 - Easy operations
 - Handful of parameters to set
 - Instructions fit on 3 pages (A4)
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/EowynDocumentation>
 - Deployed world-wide with < 10 instances
 - Mostly run in on/off mode (with default settings)
 - Proven reliability and robustness
 - Processed ~10 million jobs (2 million CPU days)
 - Supervisor process runs stably for $O(1 \text{ week})$
 - No indications yet that it cannot run longer

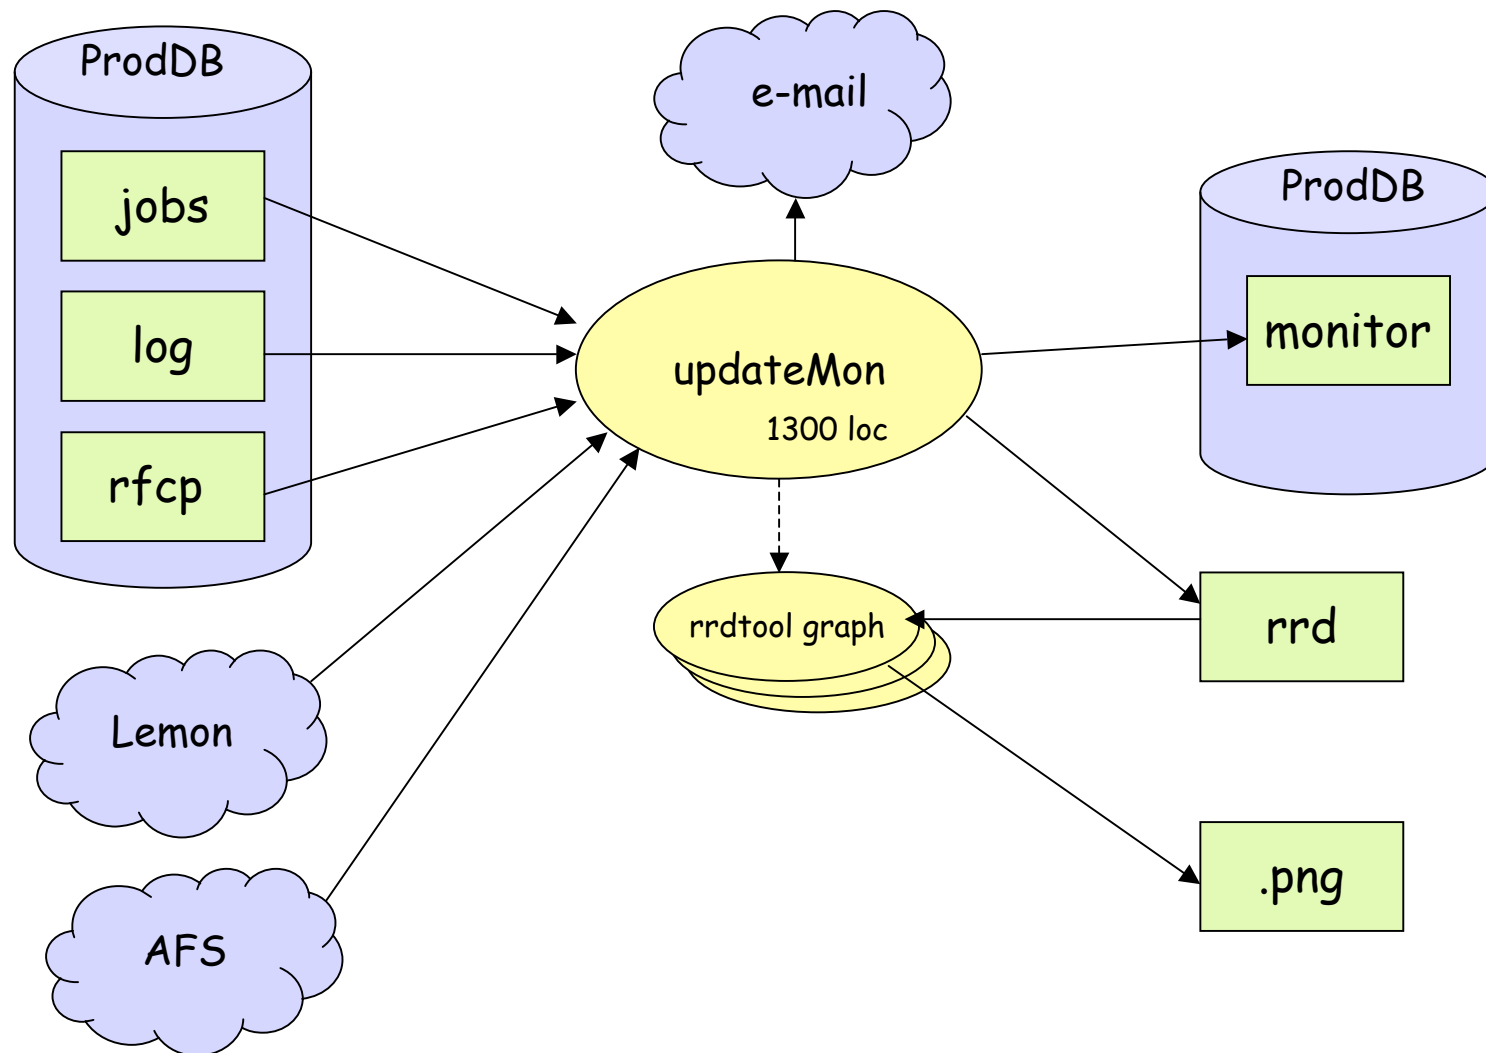
-
- Introduction
 - Tier-0 software suite
 - Design guidelines
 - Shared ProdSys components
 - Tier-0 proper components
 - Tier-0 scaling tests in 2006-2007
 - Summary and outlook

-
- TOMS (Tier-0 Management System)
 - Architecture similar to Eowyn
 - Driven by RAW datasets arriving from EF, defines all subsequent tasks/jobs and datasets in the Tier-0 dataflow
 - Bulk state processor
 - All tasks/datasets move through their state diagram
 - Grouping of operations in transverse dimension
 - Coarse interleaving of operations
 - In infinite cycle
 - DQ2 registration is run in 'quarantine'
 - Sometimes 'hangs'
 - TOMS also (optionally) drives JERRY (the EF emulator)



-
- Tier-0 monitoring
 - Based on acrontab tasks
 - cron job collects information from various sources ...
 - Lemon (CERN IT), ProdDB, AFS
 - ... computes sums, averages, etc. and stores result both in Oracle table and RRD archive (every 5 mins)
 - ... calls RRD tools to create graphs in web-readable AFS directory
 - Monitoring can be done from everywhere
 - In addition an alarm can be associated with each monitored value
 - Notification by e-mail and/or SMS

 - Other acrontab jobs
 - Clean-up and archiving jobs
 - Auto-restart of Tier-0 processes

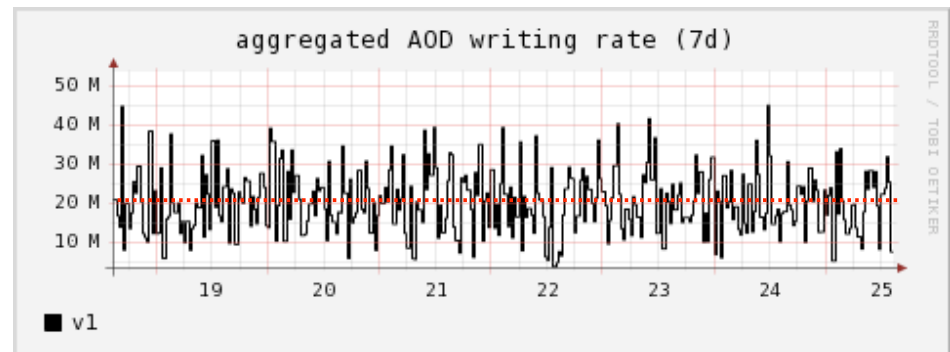




-
- Introduction
 - Tier-0 software suite
 - Design guidelines
 - Shared ProdSys components
 - Tier-0 proper components
 - **Tier-0 tests in 2006-2007**
 - Summary and outlook

-
- Tier-0 test activities started in 2005
 - Test series since then:
 - Nov/Dec 2005
 - Jan 2006; Jun 2006; Sep/Oct 2006
 - Feb 2007; since May 2007 (in continuous "test mode")
 - Tier-0 internal "nominal" transfer rates reached already in Jan 2006
 - Sep/Oct 2006 test: Tier-0 internal achievements
 - Detailed workflow for physics reconstruction and calib/align processing in place
 - Extensive monitoring (more than 90 variables)
 - Live monitoring page: <http://atlas.web.cern.ch/Atlas/tier0/monitoring/>
 - One week of stable running at 140% "nominal" rate (Oct 9-15)
 - 1.5 PB of data moved, 400k jobs executed in total
 - More than 2.6M rfcp's executed successfully

-
- Feb 2007 test: "CASTOR crisis"
 - Common stager for whole of ATLAS overloaded, led to complete breakdown
 - CASTOR/IT Task Force was put in place, to address problems quickly
 - Separate Tier-0 stager was set up in May 2007
 - Test-bed for newly developed CASTOR software
 - Tier-0 tests dedicated to provide fast feed-back to CASTOR developers
 - After successful test period, reverted to common ATLAS stager (Jun 2007)
 - Since then the Tier-0 has been running in continuous "test mode"
 - Providing a "base load" on the stager, sustained over periods of weeks
 - CASTOR performance excellent (better than ever before !)
 - The Tier-0 has already participated in two data-taking exercises
 - "Milestone" cosmics runs M3 (Jun 2007), M4 (end of Aug 2007)
 - Running reconstruction and offline data quality monitoring applications

- Aggregated rates for selected CASTOR transfers
 - Examples from test run in week of June 18-25, 2007
 - RAW (reading and writing), ESD (writing), merged AOD (writing)



----- "Nominal" rates

-
- Introduction
 - Tier-0 software suite
 - Design guidelines
 - Shared ProdSys components
 - Tier-0 proper components
 - Tier-0 scaling tests in 2006-2007
 - **Summary and outlook**

-
- TOM architecture proven adequate to meet its goals
 - Tier-0 software proven reliable, powerful and robust
 - Both ProdSys and Tier-0 specific components
 - Well tested in many exercises since 2005
 - Can cope with expected data rates
 - Can run successfully in real data-taking exercises (M3, M4)
 - Further developments (selection)
 - Improved automated running ("auto-pilot")
 - Shifters' interface
 - Future exercises (short and medium term)
 - "Final Dress Rehearsal" (Autumn 2007, Spring 2008)
 - Full chain tests EF → user analysis, using simulated data
 - Cosmics Runs M5, M6, ...
 - In 2-3 months' intervals

- News of this week: last day of the M4 cosmic ray data-taking run, Monday 3 September
 - Raw data distribution in real time from online → Tier-0 → all 10 Tier-1s
 - First time that the full chain works with all Tier-1s at the same time, and at a reasonable rate!

