



# The ALICE Computing Status and Readiness

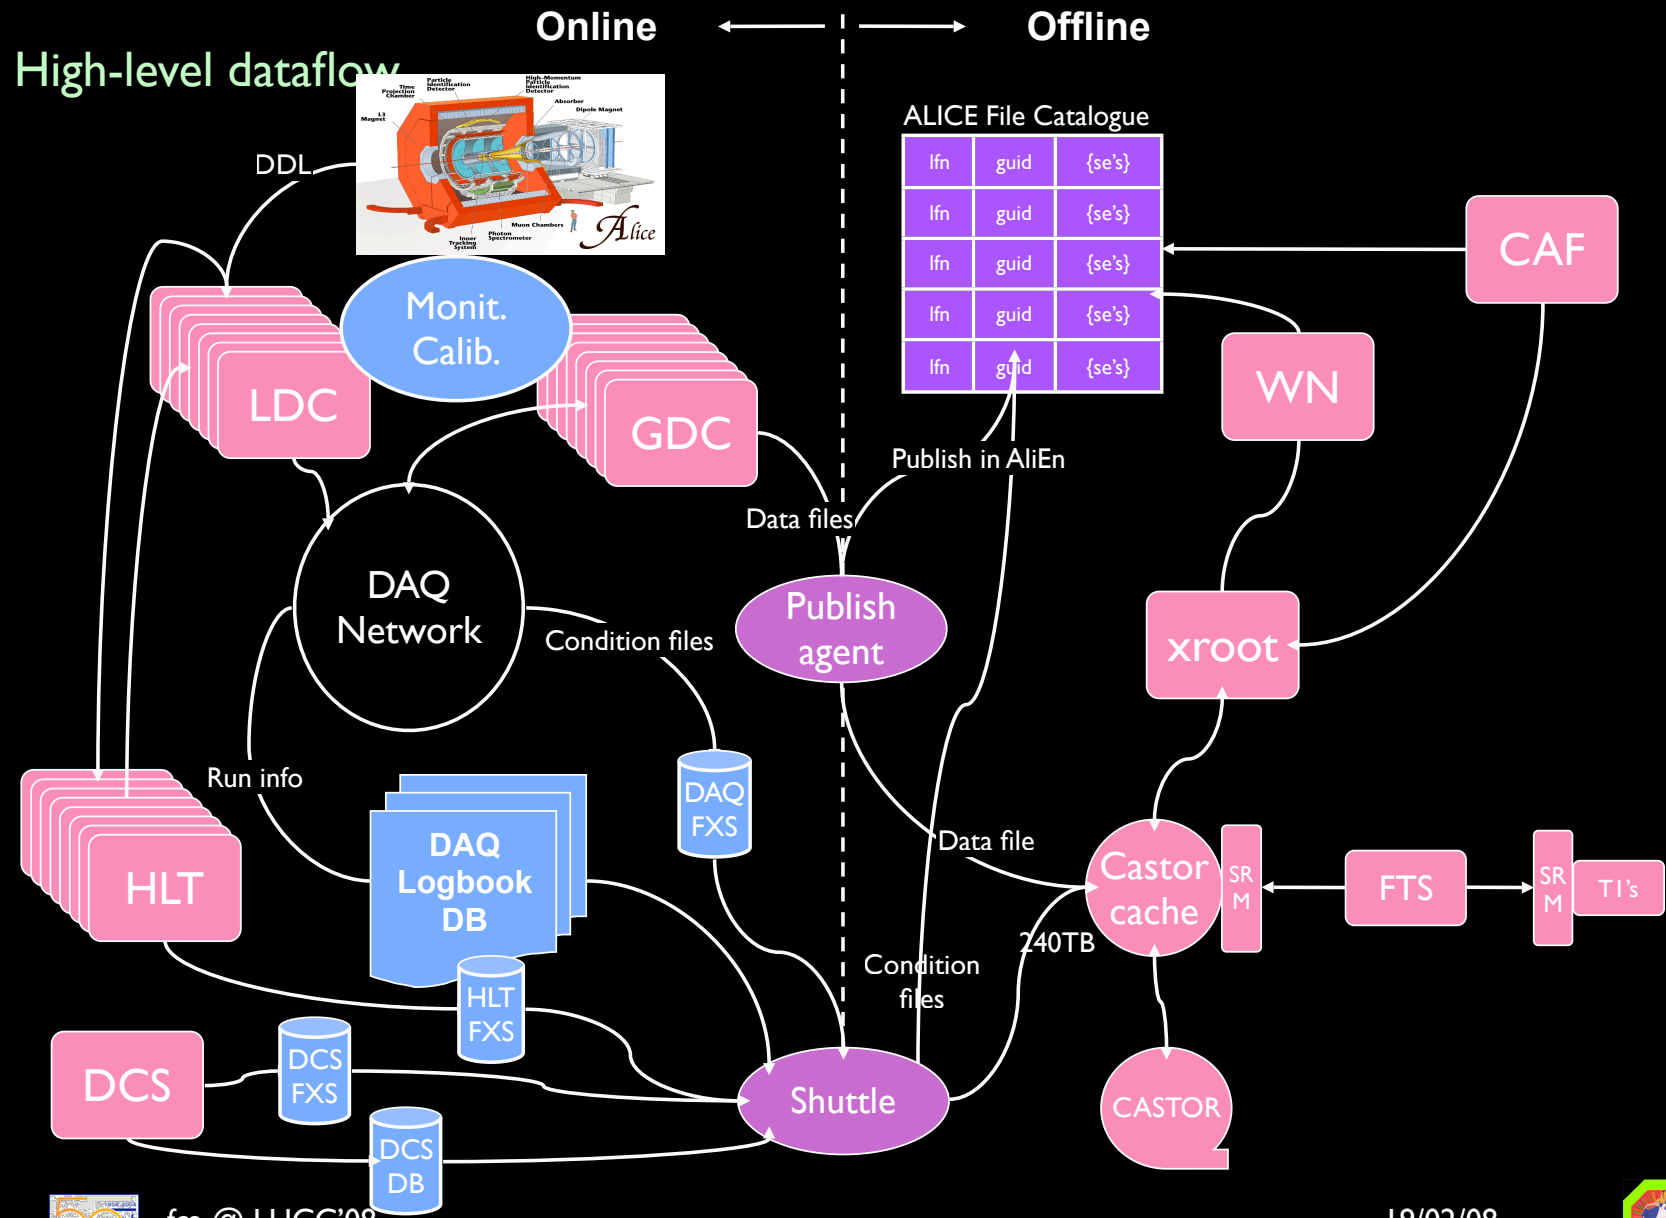
Federico Carminati  
LHCC, February 2008

# ALICE in February CCRC

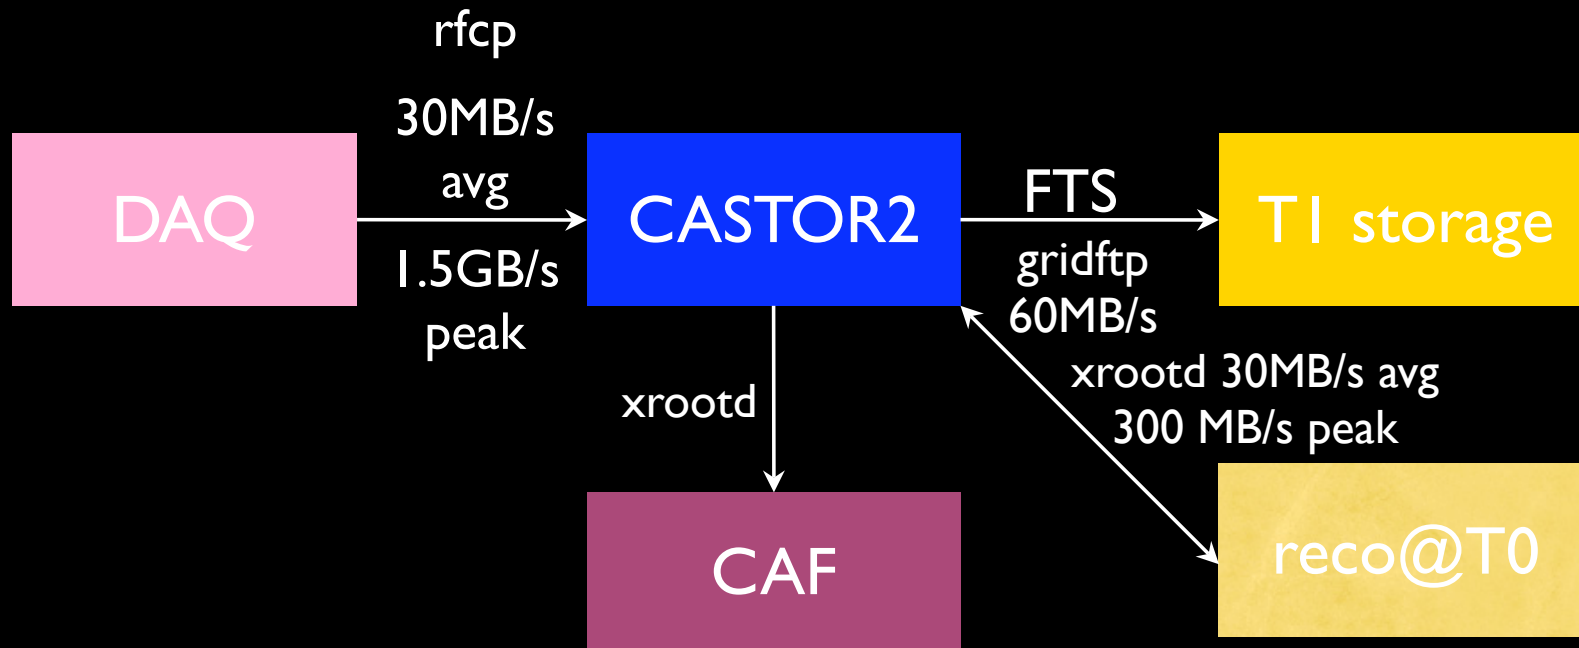
- We are now in the *February* phase of CCRC'08
  - It overlaps with the second ALICE commissioning phase (Feb 15 - Mar 11)
- First ALICE commissioning exercise ran Dec, 10-21 2007
  - Data writing to CASTOR2 and export to T1s (with no target transfer rate) – OK
  - Conditions data gathering and access – OK
  - Reconstruction of RAW at T0 – done in January



# High-level dataflow



# Data paths



# Storage requirements for February

Tier I	Disk (TB)	MSS (TB)
CCIN2P3 (15%)	2	9
CNAF (15%)	2	9
GridKA (45%)	5.5	27
NDGF (15%)	2	9
RAL (5%)	0.6	3
TI-NL (5%)	0.6	3
<b>Total</b>	<b>13</b>	<b>60</b>

- This corresponds to 50% of the total transfer data volume and is also 50% of the standard p+p data taking rate
- Volumes and rates compatible with 50% accelerator duty cycle assumed for CCRC'08



# Storage requirements

- Three storage classes at TIs
  - TID0 for RAW
  - T0DI for ESDs hot copy
  - TID0R for complementary data (keeping 60MB/sec constant rate out of T0)
- The TID0R is recyclable at the site's discretion
- The classes are implemented as separate directories in all types of storage
  - Allows for optimisation of access to tapes, especially important for RAW
  - Already implemented at CERN, CCIN2P3, GridKA, SARA, NDGF
  - Under way at CNAF



# Activity table


Activity	Week1	Week2	Week3	Week4	Week5
Data Taking			✓	✓	✓
Data replication to T1s	✓	✓	✓	✓	✓
Raw -> ESD, P1@T0 & P2@T1s			✓	✓	✓
MC production	✓	✓	✓	✓	✓

- Week 1, 2 – detector integration with control systems
- **Week 5 (new: 03/03 to 09/03)** – global cosmics trigger run with all detectors (high statistics, calibration/alignment)
- Replication and reconstruction plans are adjusted with respect to the priorities of the ALICE commissioning exercise
  - **In particular with the quick feedback needed for detector calibration**



# CASTOR2+xrootd

- CASTOR2+xrootd fully in production at CERN since beg Feb
- Minor issues observed with prepare request on disk servers
- The release was extensively tested before deployment



## Many thanks to

CASTOR2 development and deployment teams

xrootd development team

Rosa Garcia IT/DM  
Andreas Peters IT/DM  
Fabrizio Furano IT/GS





# Site readiness

- Storage (dCache - FTS/xrootd) set up at CCIN2P3, GridKA, NDGF
  - The necessary storage capacity is installed
- T-NL (dCache) – storage deployed, VO-box requires reinstallation with SLC4/gLite 3.1 (in progress)
- CNAF (CASTOR2) – CASTOR2+xrootd installed, testing phase
- RAL (CASTOR2) – discussions ongoing
- 15% of the ALICE T2 sites have deployed DPM or dCache as disk storage
  - Already in use for analysis tasks, ESDs replication in progress



# ALICE T1 VOBox SLA

- GridKA, CNAF - **approved**
- RAL - advanced draft (**we can approve it if RAL is OK**)
- NDGF - **advanced draft**
- T1-NL - **nothing yet, waiting for the centre**
- CCIN2P3 - **nothing yet, waiting for the centre**
- SLAs reflect operational experience with T1s in past years
  - VO-box support is good, responsibilities of the centres and ALICE are well-defined
  - No VO-box malfunction (VO-box dead), issues are on software/service levels
  - VO-box setup is rather simple – services can be restored quickly



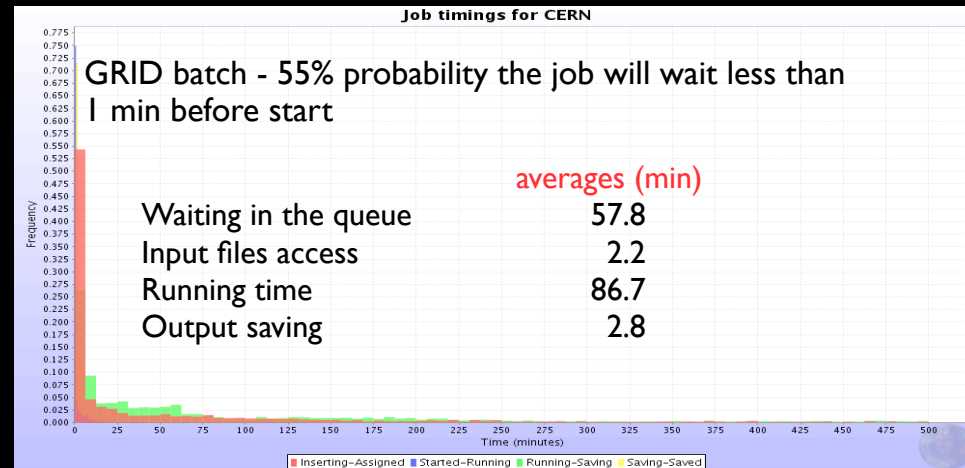
# Application software

- Strict weekly release schedule for AliRoot
  - Coordinated with detector groups
  - Essential for integration and validation of detector algorithms in DAQ/HLT systems (shared code)
- Simultaneous deployment on the Grid for RAW production and ESD analysis



# Analysis

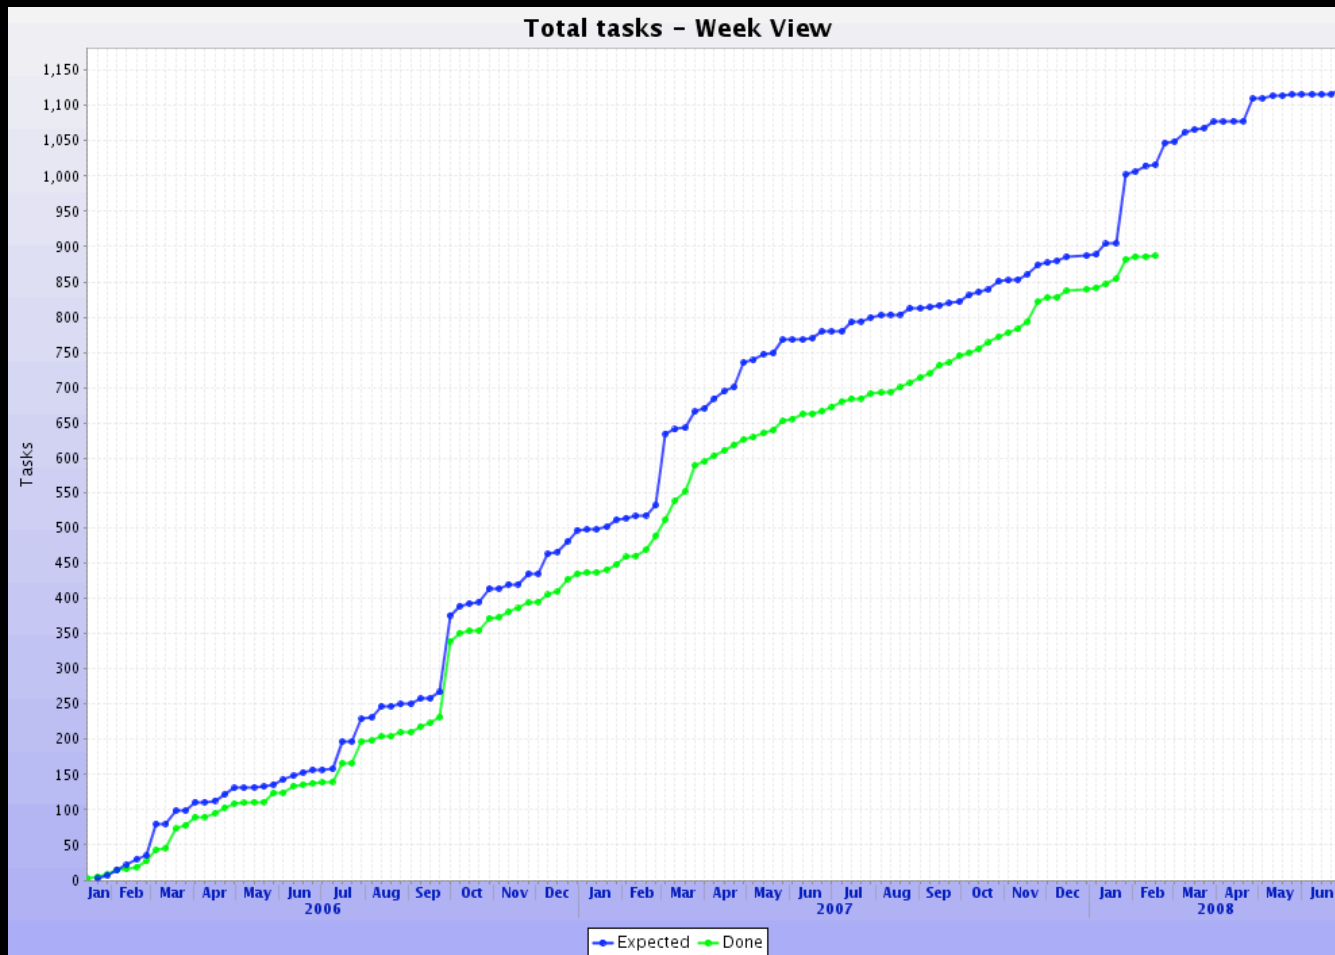
- Quick re-processing of RAW data (calibration)
- Analysis is carried out immediately
  - Batch on Grid - optimisation of user jobs priorities in central queue



- On CAF@CERN (PROOF) - disk and CPU quotas introduced
- Second interactive analysis cluster deployed at GSI: GSIAF (PROOF) with Grid-aware SE
  - Portions of RAW and ESDs specifically replicated to GSI



# Planning



# Resource overview

Parameter	Now	CTDR	Ratio
pp RAW	2.44MB	0.2MB	22
Pb RAW	35MB	13.8MB	2.5
ESD pp	0.04MB	0.04MB	1.0
ESD Pb	6.3MB	3.0MB	2.1
AOD pp	5kB	16kB	0.3
AOD Pb	1.3MB	0.34MB	3.8
Reco pp	6.8s	6.5s	1.0
Reco Pb	800s	810s	1.0

Pledged by external sites versus required (new LHC schedule) all													
		2007		2008		2009		2010		2011		2012	
		T1	T2	T1	T2	T1	T2	T1	T2	T1	T2	T1	T2
CPU	Requirement (MSI2K)	3.2	4.6	10.1	12.5	19.9	14.3	23.5	25.0	30.5	32.5	39.7	42.2
	Missing %	-45%	-5%	-31%	-37%	-47%	-20%	-39%	-40%	-53%	-54%	-64%	-65%
Disk	Requirement (PB)	1.3	0.7	4.0	1.7	6.8	4.0	12.0	4.3	16.6	5.6	22.4	7.3
	Missing %	-43%	4%	-32%	5%	-37%	-22%	-50%	-3%	-64%	-25%	-73%	-43%
MS	Requirement (PB)	1.4	-	5.8	-	12.4	-	19.7	-	27.0	-	33.7	-
	Missing %	-23%	-	-32%	-	-36%	-	-40%	-	-57%	-	-65%	-



# Staffing situation

- Most of the people in the CORE Offline team are PJAS and Fellows
  - Some of the key PJAS will have to leave end 2008
  - Fellows are now deviating from computing to physics
- The staffing situation of core Offline may soon become critical



# Summary

- ALICE CCRC'08 - focus on data management
- Storage is deployed gradually
  - Very good progress with CASTOR2+xrootd, dCache+xrootd
  - Several T2 sites have deployed DPM+xrootd storage – important for the analysis phase
- The December tests were successful
  - Essentially all building blocks for CCRC'08 were tested
  - Present exercise will add more resources and test the stability
- Site resources should be sufficient to carry out all planned activities
- CCRC'08 is providing a much needed focus for the experiment activities on the Grid
  - In addition to this, ALICE is exercising the complete data taking/calibration and reconstruction system with live data from the experiment





