LHCC Referee Meeting

September 22, 2015

ALICE Status Report

Predrag Buncic

CERN

ALICE during LS1

- Detector upgrades
  - TPC, TRD readout electronics consolidation
  - TRD full azimuthal coverage
  - +1 PHOS calorimeter module
  - New DCAL calorimeter
- Software consolidation
  - Improved barrel tracking at high p_T
  - Development and testing of code for new detectors
  - Validation of G4
- Re-processing of RAW data from RUN1
  - 2010-2013 p-p and p-A data recalibrated
  - All processing with the ***same*** software
  - General-purpose and special MC completed

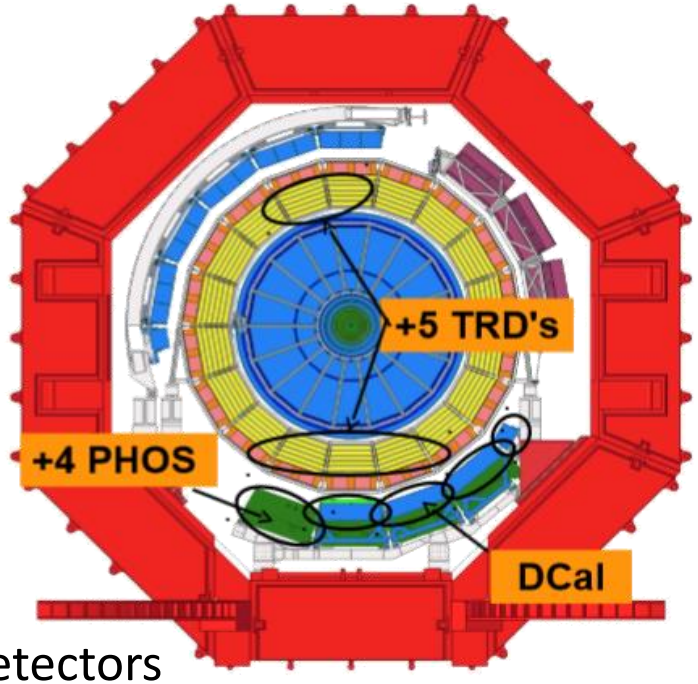

ALICE Grid keeps growing

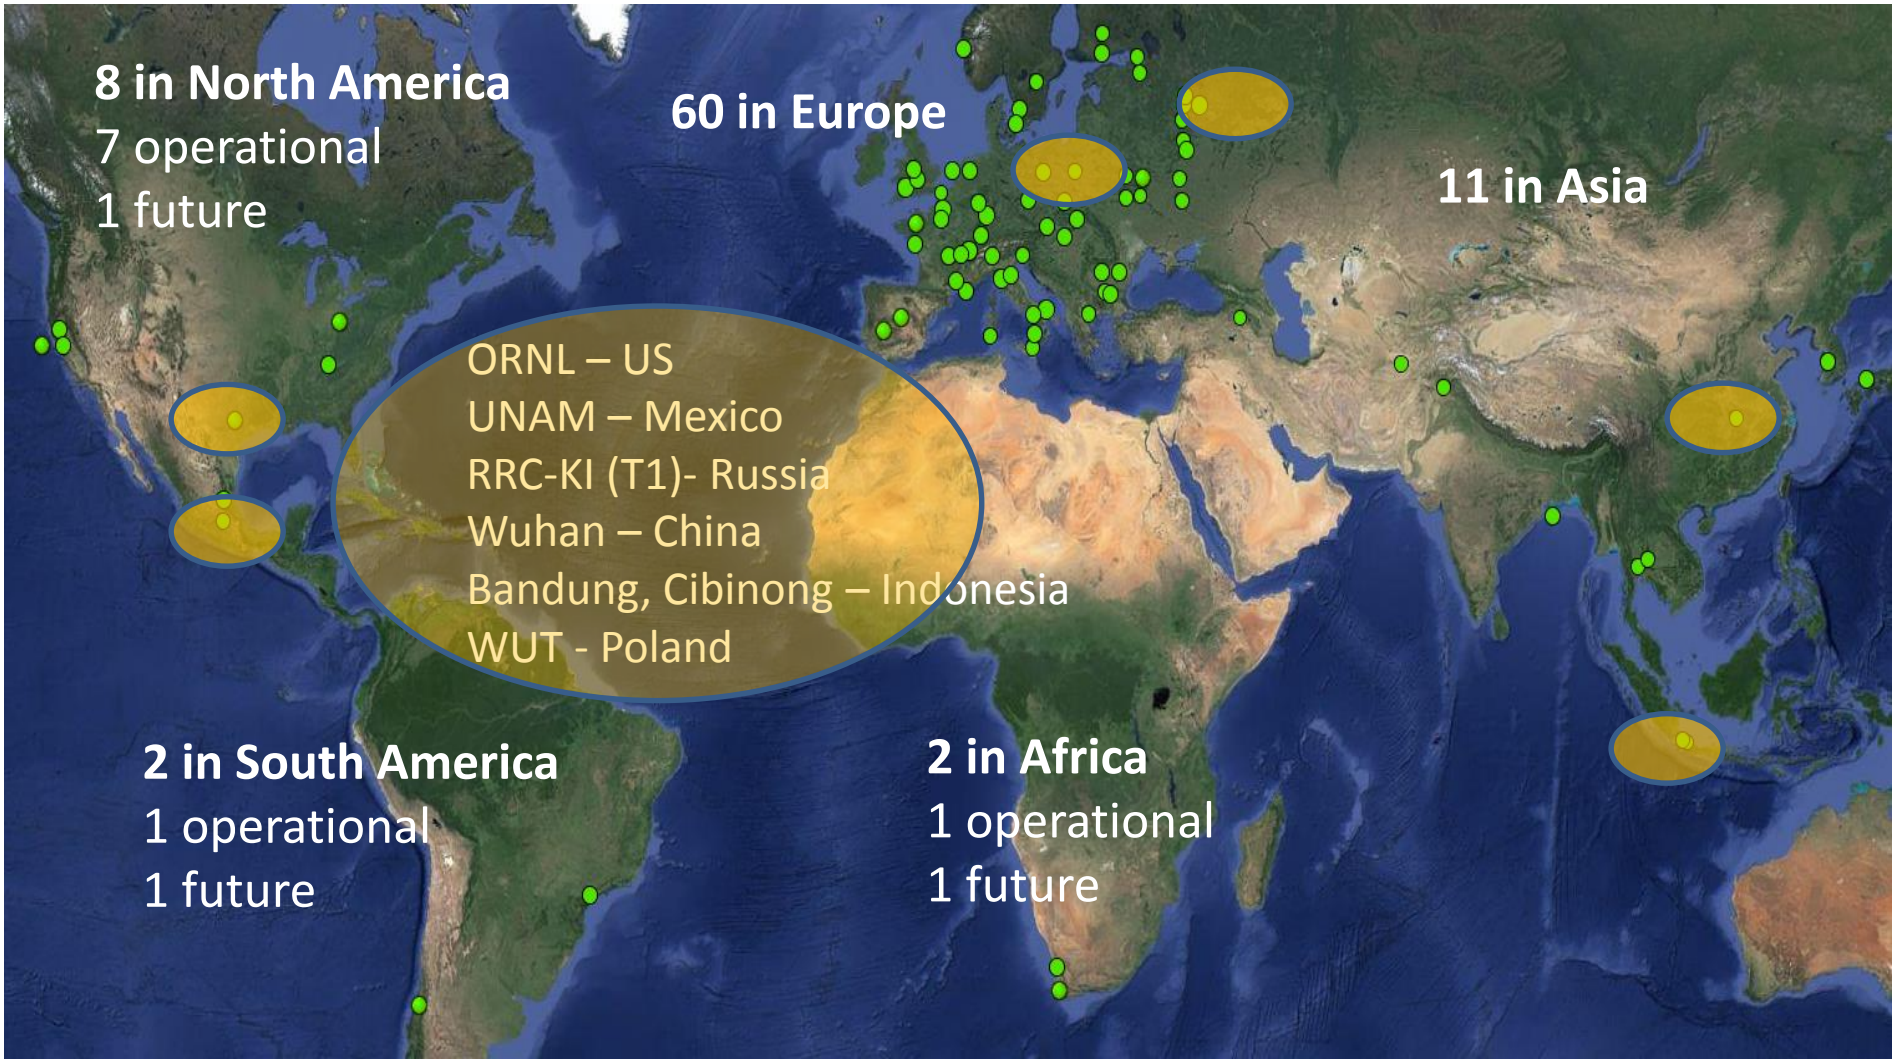

Grid utilization: end of RUN1 - today

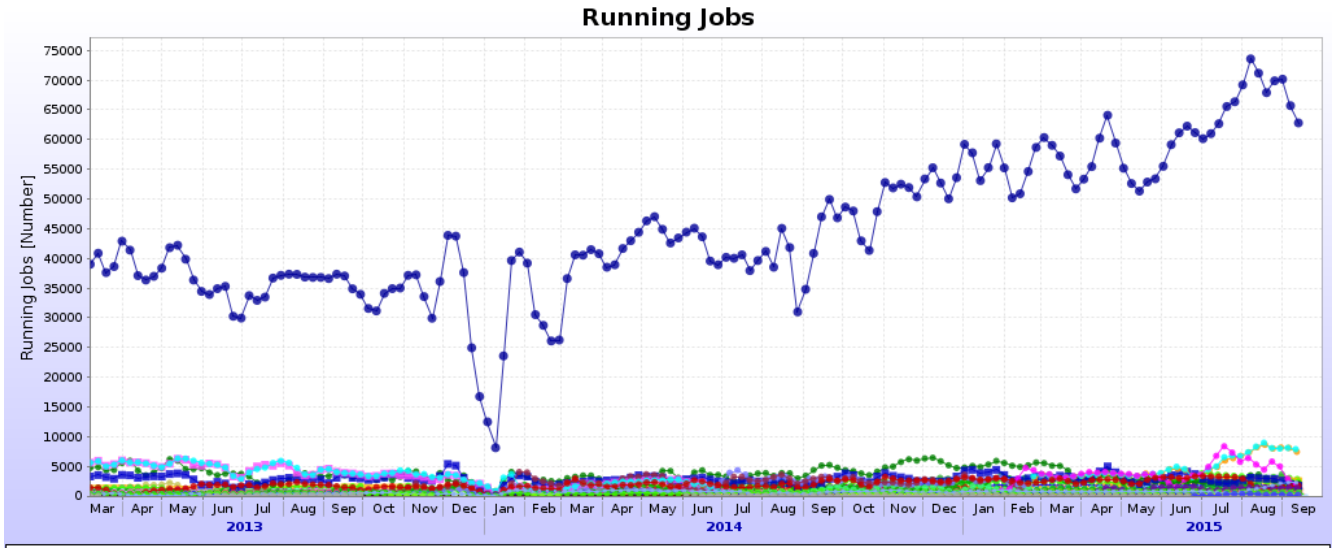

- Reaching new highs
- 96K parallel jobs

- Consistent and very good efficiency in all computing centres

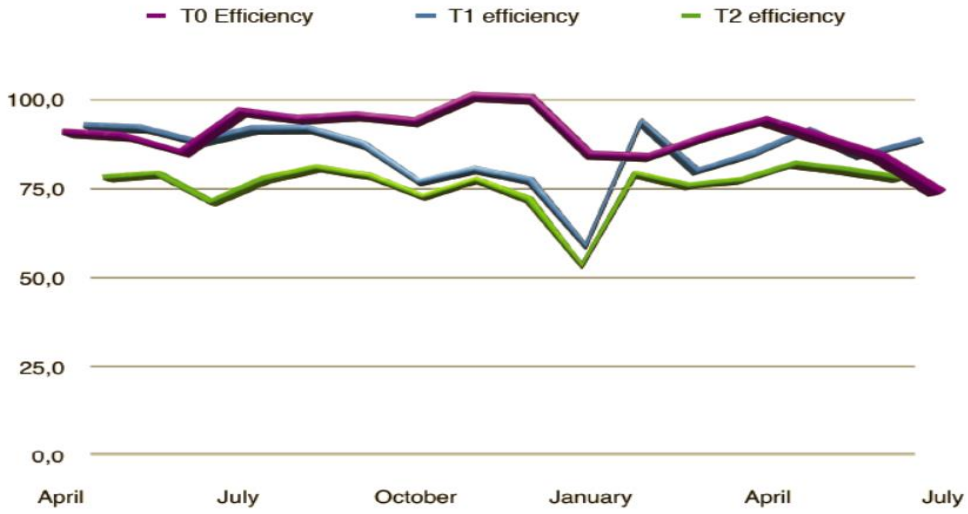

ALICE Resources usage in 2015 (CPU)

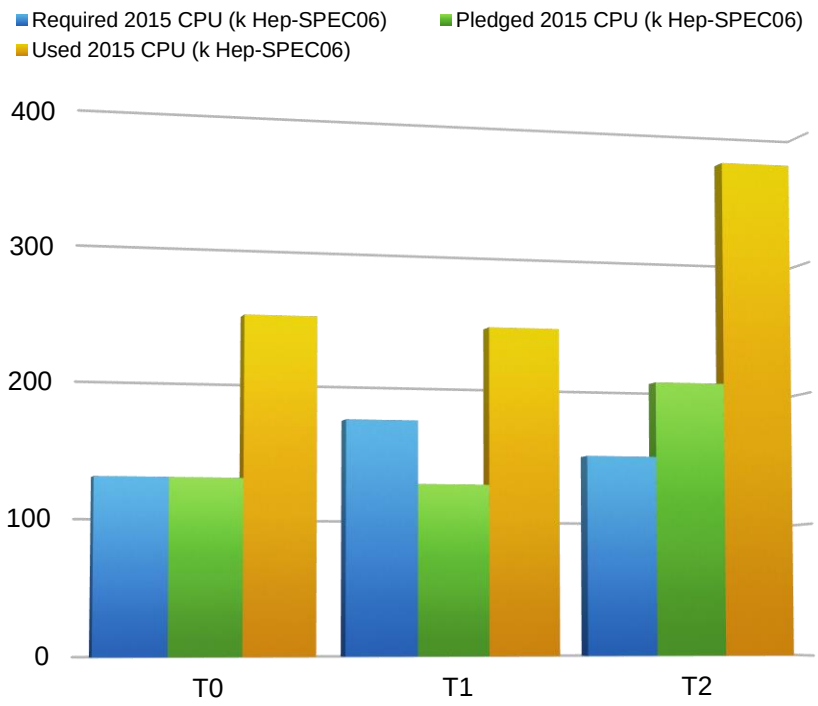

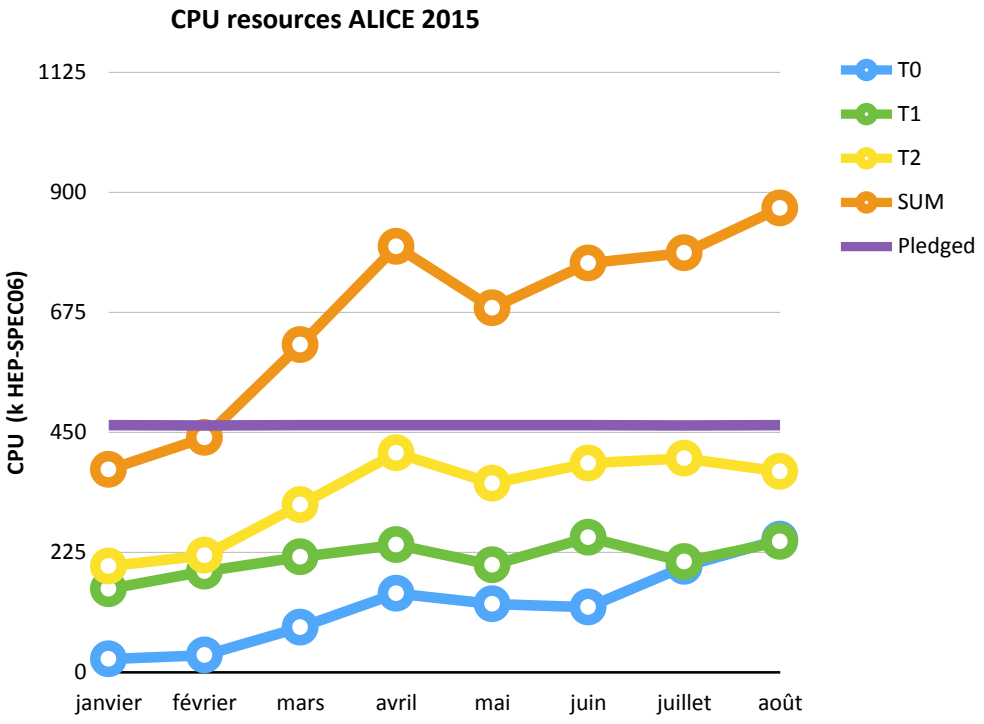

ALICE Resources usage in 2015 (DISK)

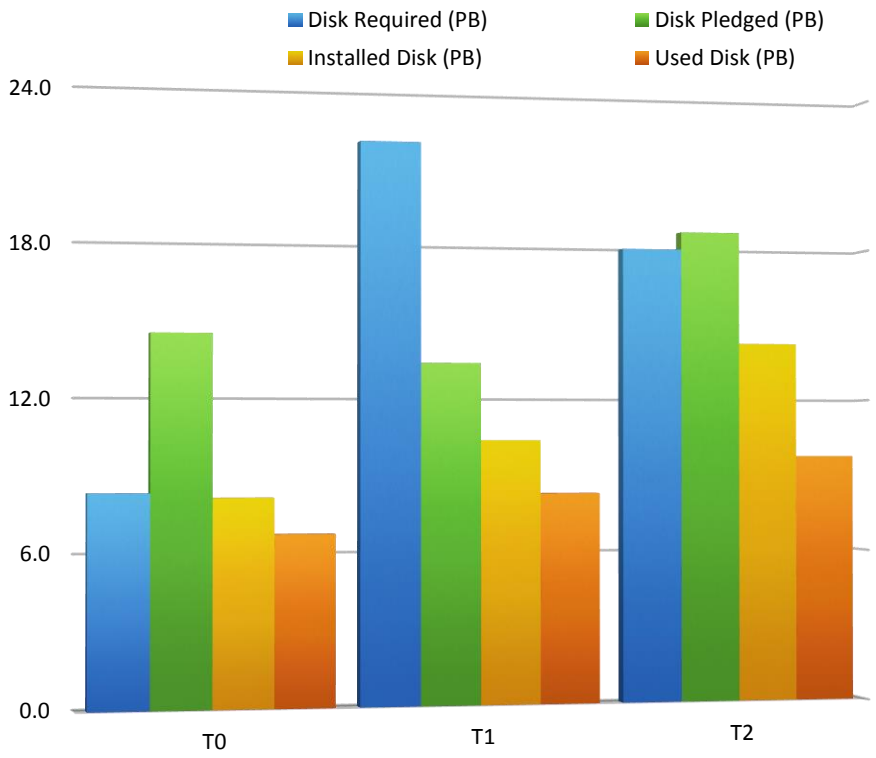

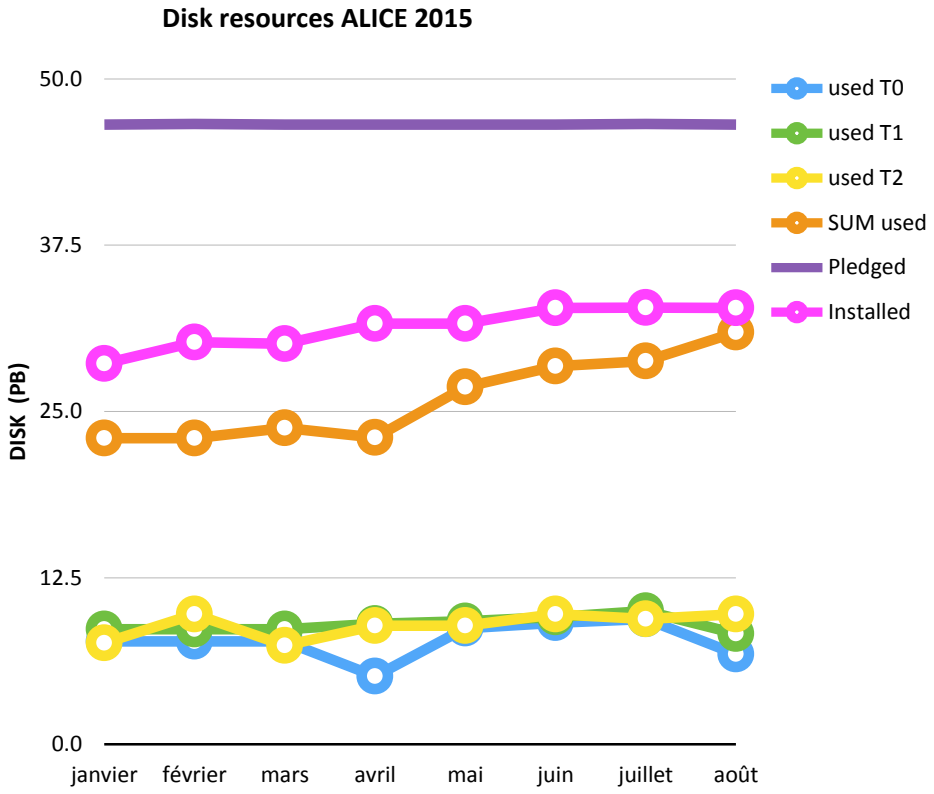

Data popularity

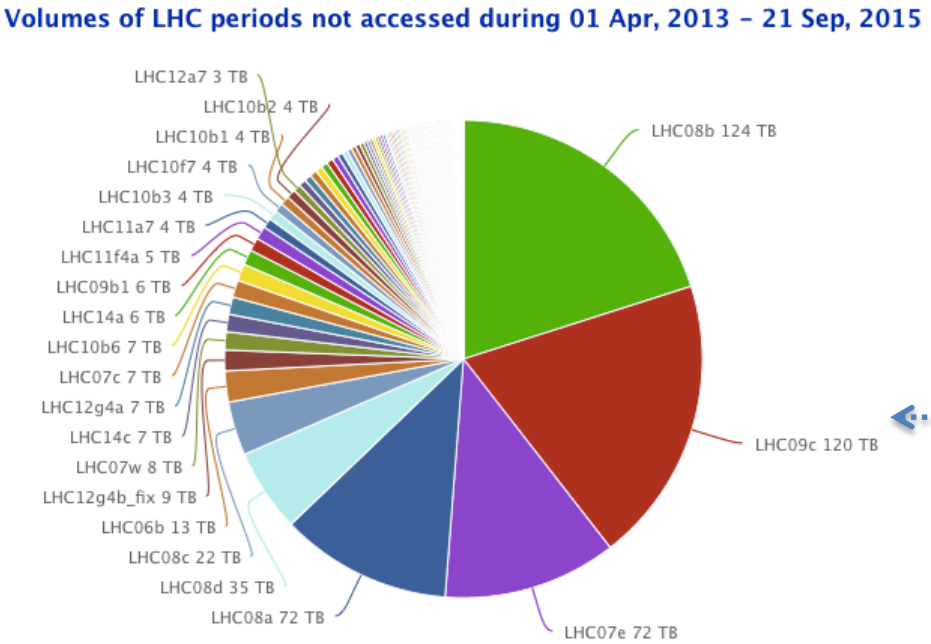

- Regular storage cleanup
- less than 5% of unpopular data on disk

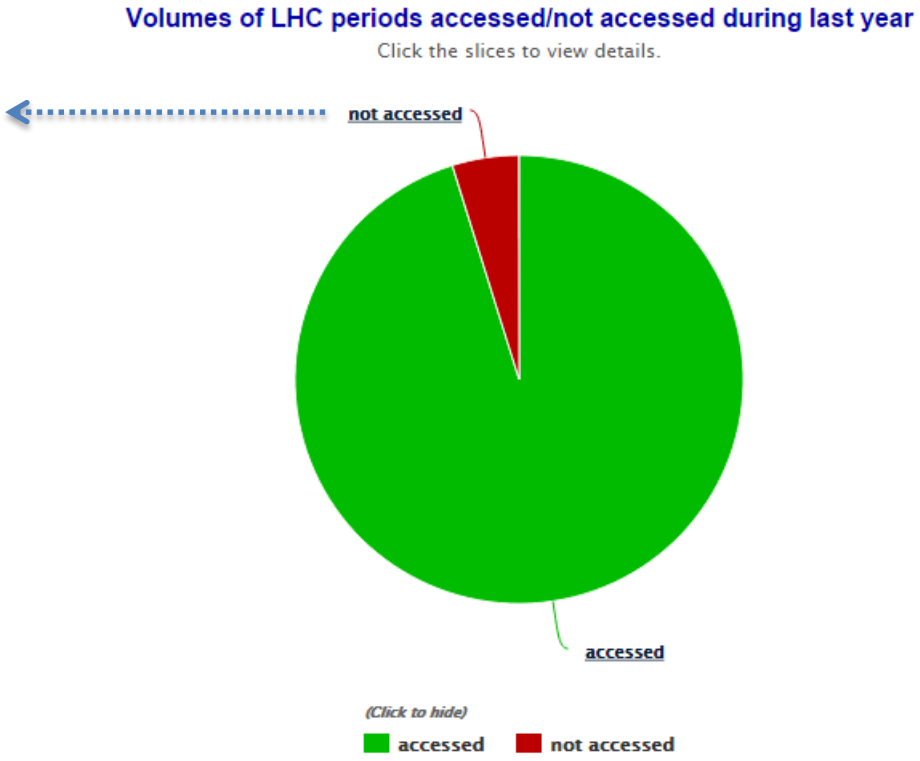

Run2 progress

- Steady data taking
  - Fast calibration and reconstruction
  - Quasi-online data processing
    - Without backlog

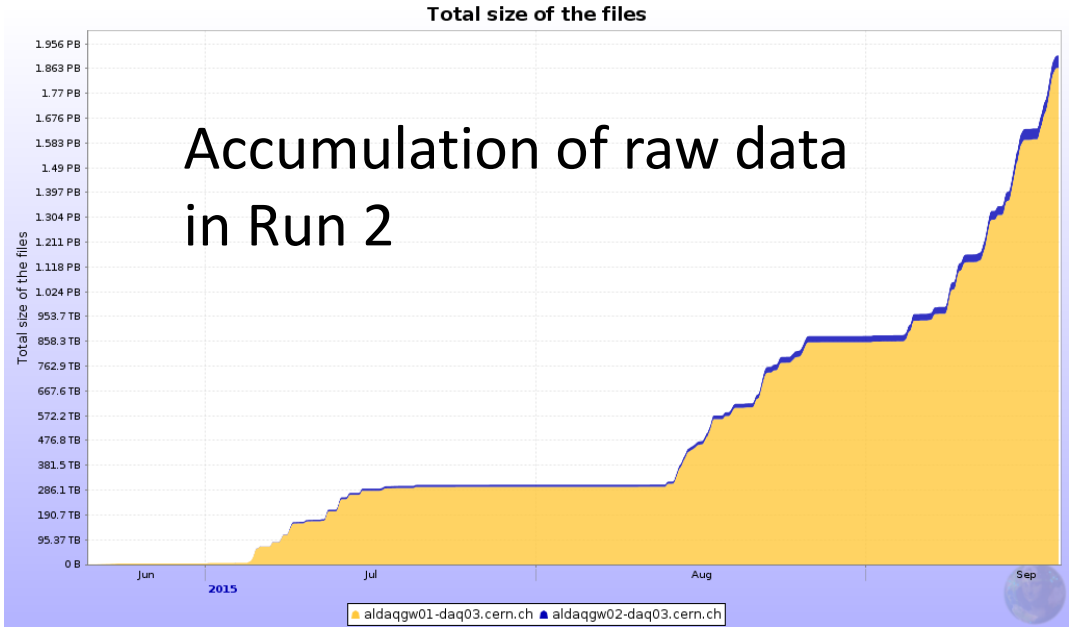

Production	Description	Status	Run Range	Runs	Chunks	Size	Chunks		Size		Events
LHC15i_pass1	LHC period LHC15i - Full production pass 1	Running	235196 - 235886	25	101,081	164.8 TB	83,245	82%	8.665 TB	6%	62,007,426
LHC15h_pass1	LHC period LHC15h - Full production pass 1	Completed	232914 - 234050	68	327,386	544.9 TB	293,861	89%	38.96 TB	7%	213,863,587
LHC15g_pass1	LHC period LHC15g - Full production pass 1	Completed	228855 - 230292	31	26,567	37.65 TB	26,171	98%	6.125 TB	16%	20,766,687
LHC15f_pass1	LHC period LHC15f - Full production pass 1	Completed	224895 - 226532	45	18,857	21.9 TB	16,542	87%	12.1 TB	62%	84,564,615
LHC15e_pass1	LHC period LHC15e - Full production pass 1	Completed	223270 - 224772	59	15,648	9.16 TB	11,595	74%	1.685 TB	24%	73,262,707
LHC15d_pass1	LHC period LHC15d - Full production pass 1	Completed	220139 - 222966	100	6,148	5.513 TB	5,234	85%	656.5 GB	13%	29,817,237
					495,687	783.9 TB	436,648		68.17 TB		484,282,259

Resource requirements

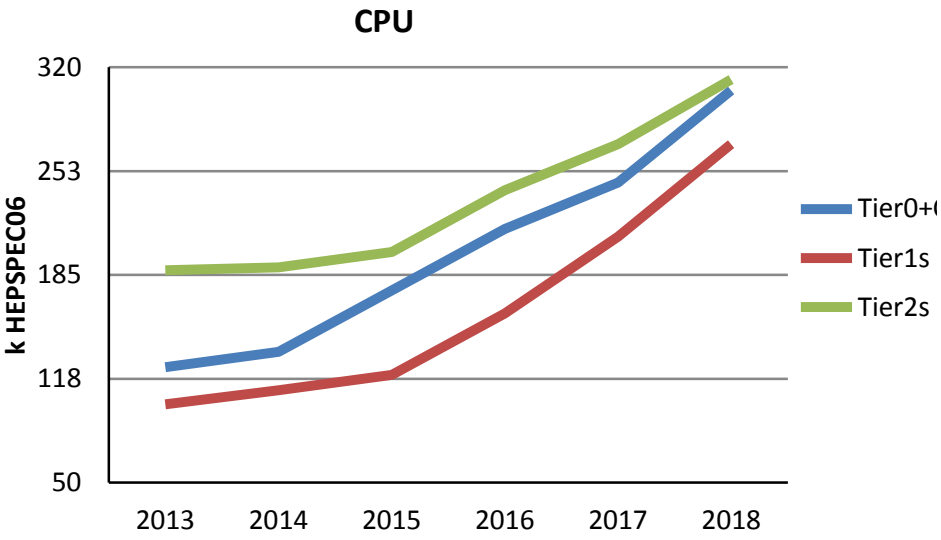

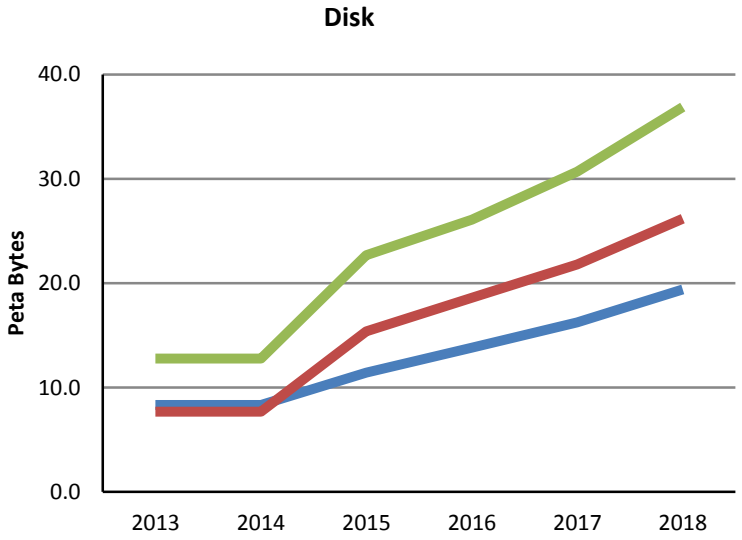

CPU (kHEPSPEC06)

	Tier0	CAF	Tier1s	Tier2s
2015	130	45.0	120	200
2016	170	45.0	160	240
2017	200	45.0	210	270
2018	260	45.0	270	312

Disk (PB)

	Tier0[1)]	CAF	Tier1s[2)]	Tier2s
2015	11.1	0.34	15.4	22.7
2016	13.4	0.44	18.6	26.1
2017	15.7	0.54	21.8	30.7
2018	18.8	0.64	26,2	36.9

[1)] Excluding the 3.0 PB of disk buffer in front of the taping system
[2)] Excluding the 2.35 PB of disk buffer in front of the taping system

Beyond Run2 – the O^2 project

Technical Design Report for the Upgrade of the Online-Offline Computing System

- Submitted to LHCC
- Framework and code development is well under way

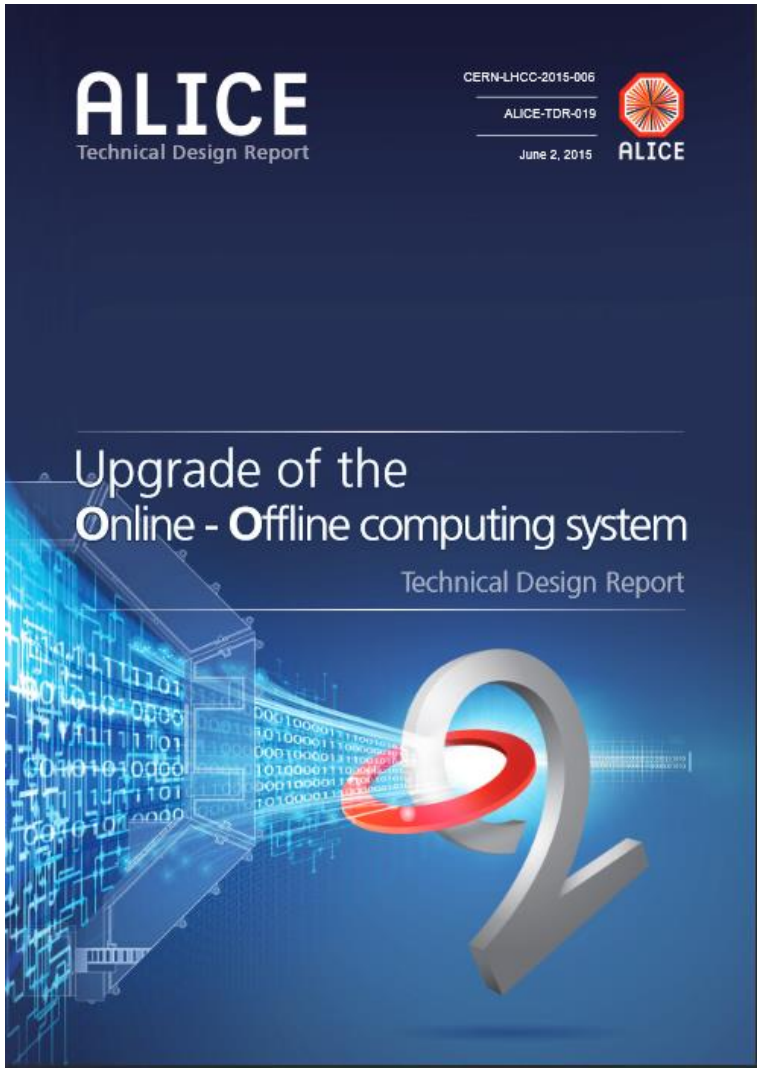

Rebalancing disk/CPU ratio on T2s

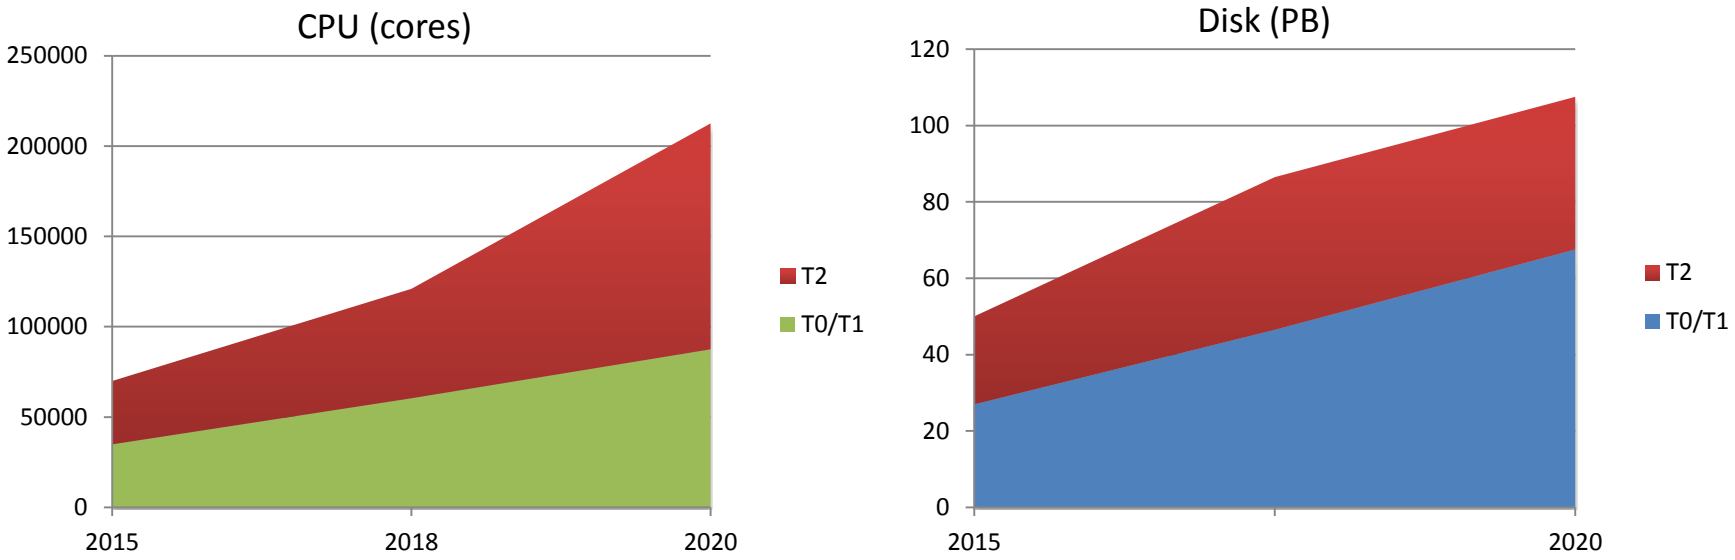

- Expecting Grid resources (CPU, storage) to grow at 20% per year rate
  - Large number of disk will be used by Run 1 and Run 2 data
- Since T2s will be used almost exclusively for simulation jobs (no input) and resulting AODs will be exported to T1s/AFs, we expect to significantly lower needs for storage on T2s

Combined long term disk space requirement (T0+T1s+T2s+O2)

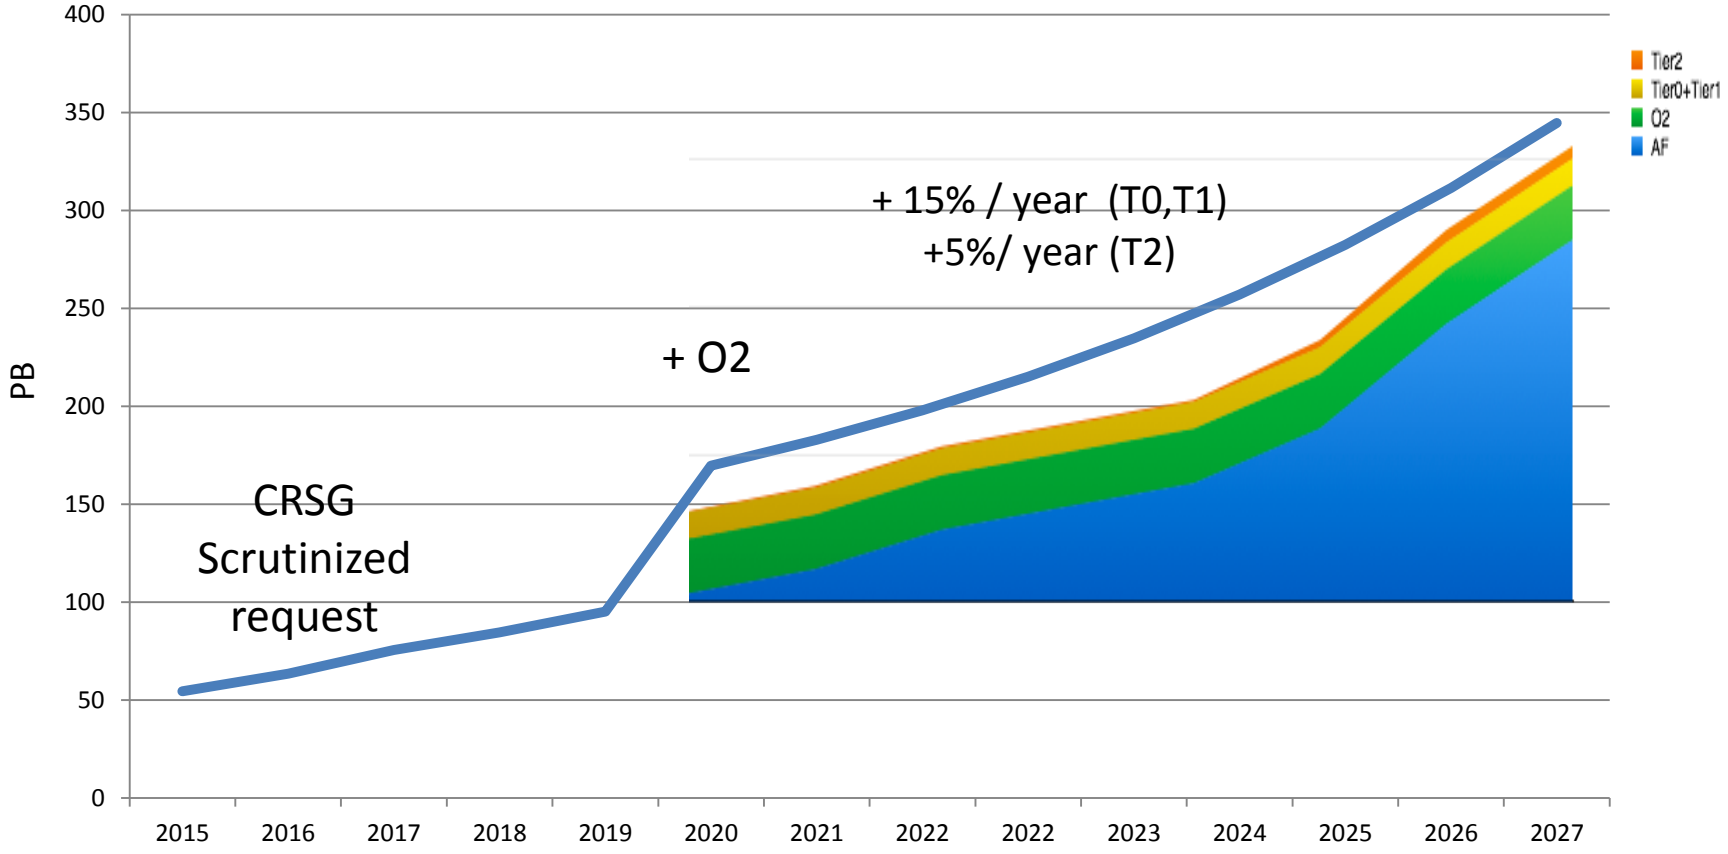

First O2 milestones

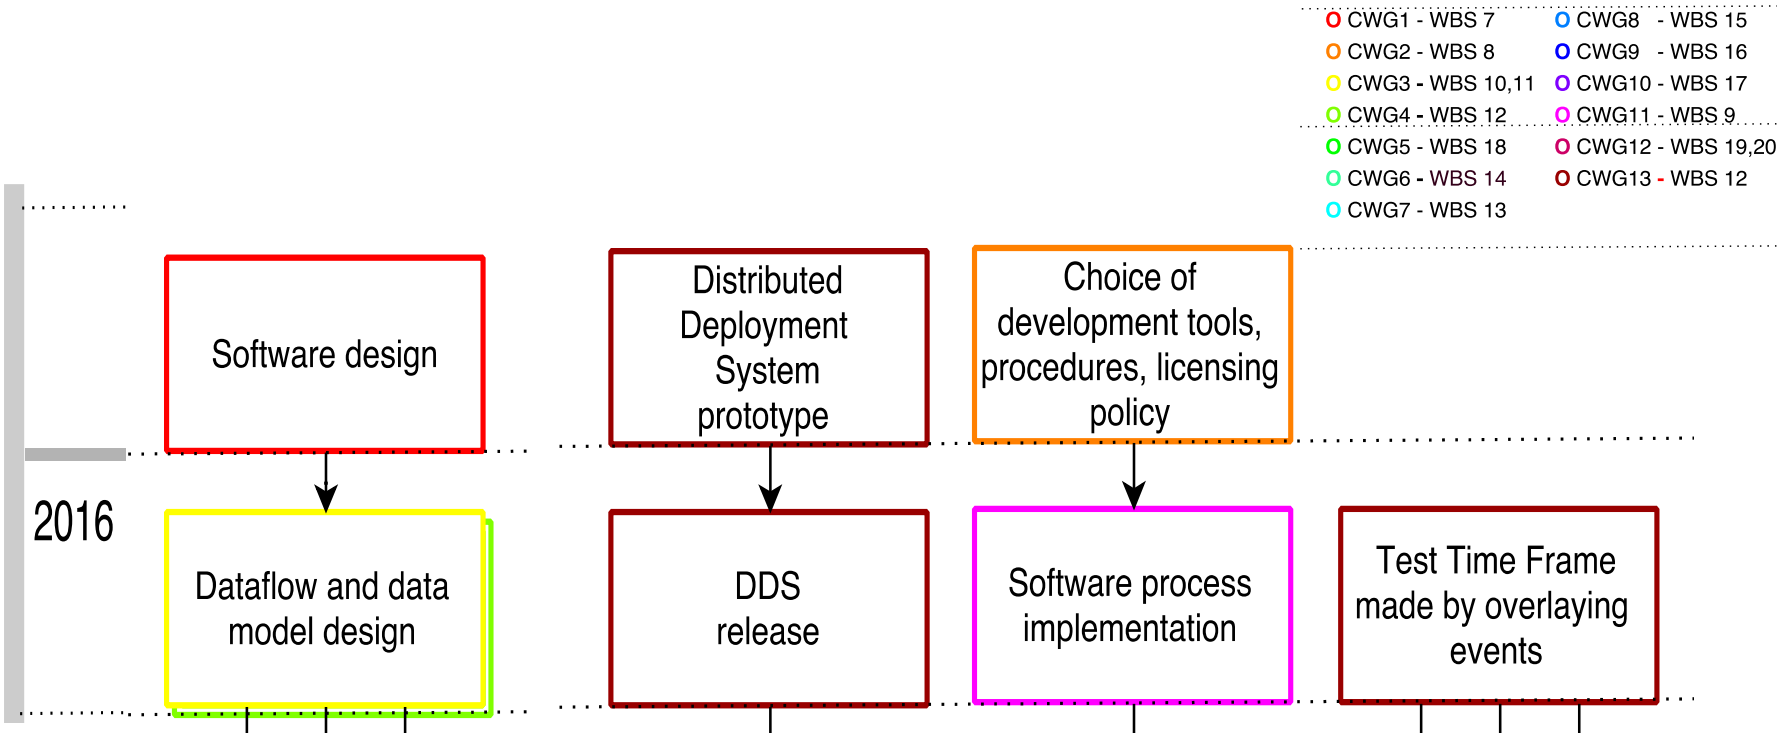

Dynamical Deployment System (DSS)

- Connecting the FairMQ devices/tasks requires
  knowledge of connection parameters that need to be propagated
- startup time - the time which took DDS to distribute user tasks, propagate properties and enter into RUN state

Devices (user tasks)	startup time*	propagated key-value properties
2721 (1360 FLP + 1360 EPN + 1 Sampler)	17 sec	~ $6x10^6$
5441 (2720 FLP + 2720 EPN + 1 Sampler)	58 sec	~ $23x10^6$
10081 (5040 FLP + 5040 EPN + 1 Sampler)	207 sec	~ $77x10^6$

- DDS Topology editor
  - An interactive web tool that allows: creation, modification and visualization of a DDS topology
  - Aims to provide the DDS user with an intuitive platform for creating the desired topology.

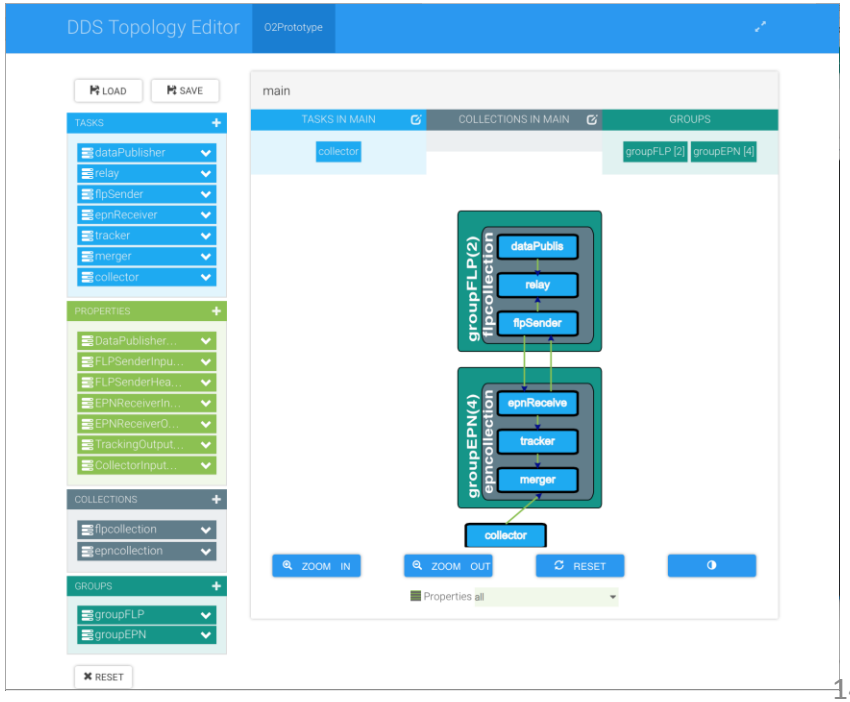

Parameter management

based on Riak key-value store

- high availability
- scalability
- fault tolerance
- configurable

Two storage back-ends were tested:

- Bitcask
  - best latency
  - nodes out-of-memory
- LevelDb
  - similar performance
  - compressed storage

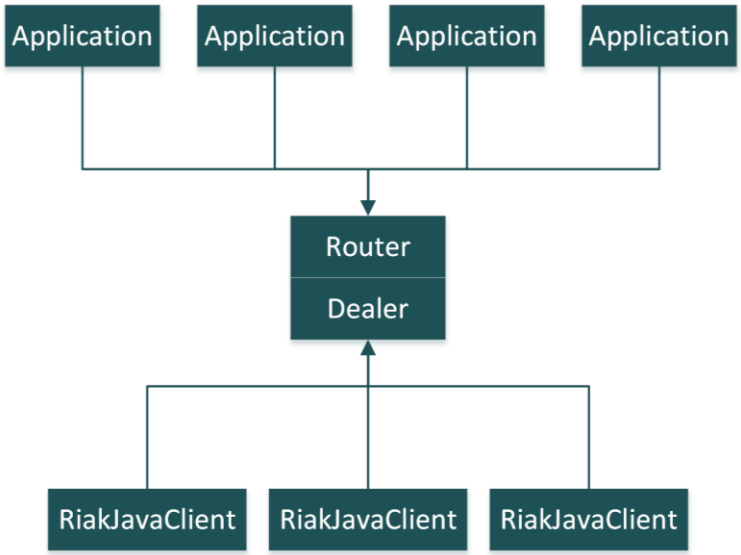

protobuf
Protocol Buffers

Java

ØMQ

Message-Queue based concept make it possible to use directly the native Java client of Riak

Software Lifecycle

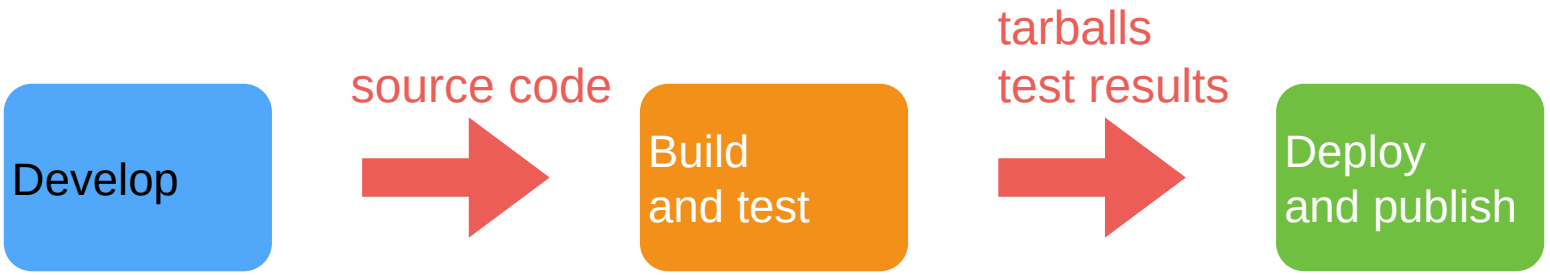

Git-based | Integration builds | CVMFS, RPM, Ubuntu

- Agile lifecycle: from code to published builds quickly, with automated tests
- Philosophy: automatize as much as possible, trivialize human intervention
- Objective for O²: refreshing aging infrastructure and paradigm
  - Setup the infrastructure for running builds and tests: done
  - Basic tools to build+publish current software and O²: done
  - Visualization interface for the tests: Nov 2015
  - Automatic test and flagging of Pull Requests for O²: Nov 2015
- Basic infrastructure is ready way ahead of time: as of Sep 9, old system

Software Lifecycle


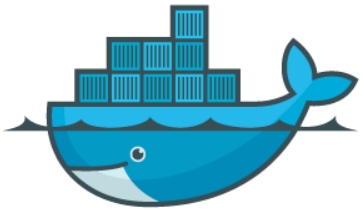
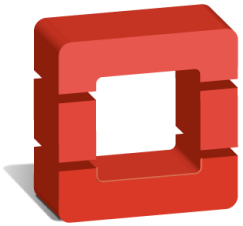


- Quick development of the new infrastructure possible thanks to use of common OpenSource tools supported CERN IT
  - Source code on GitHub and CERN GitLab
  - Builds are scheduled by Jenkins
  - Full infrastructure running on CERN OpenStack virtual machines
  - We can build on multiple platforms seamlessly by using Docker containers
  - Containers are deployed by Apache Mesos
  - Software distribution by CVMFS
  - Visualization of performance test trends via Kibana

Summary

- The main ALICE goals in LS1 in terms of computing and software were achieved
  - Software quality improved and the entire Run1 dataset was re-processed with same software version
- The reprocessing was helped by an effective use of computing resources
  - Often going beyond the pledges for CPU
- Computing resource requests for Run 2 remain unchanged
  - Must keep up in order to match start of Run3
- O2 will absorb the impact on resources at the start of Run3
- Work on O2 has started following the plan presented in TDR
  - First milestones on the horizon

Backup

O2 Milestones