LHCC Referee Meeting

01/12/2015

ALICE Status Report

Predrag Buncic

CERN

Grid usage

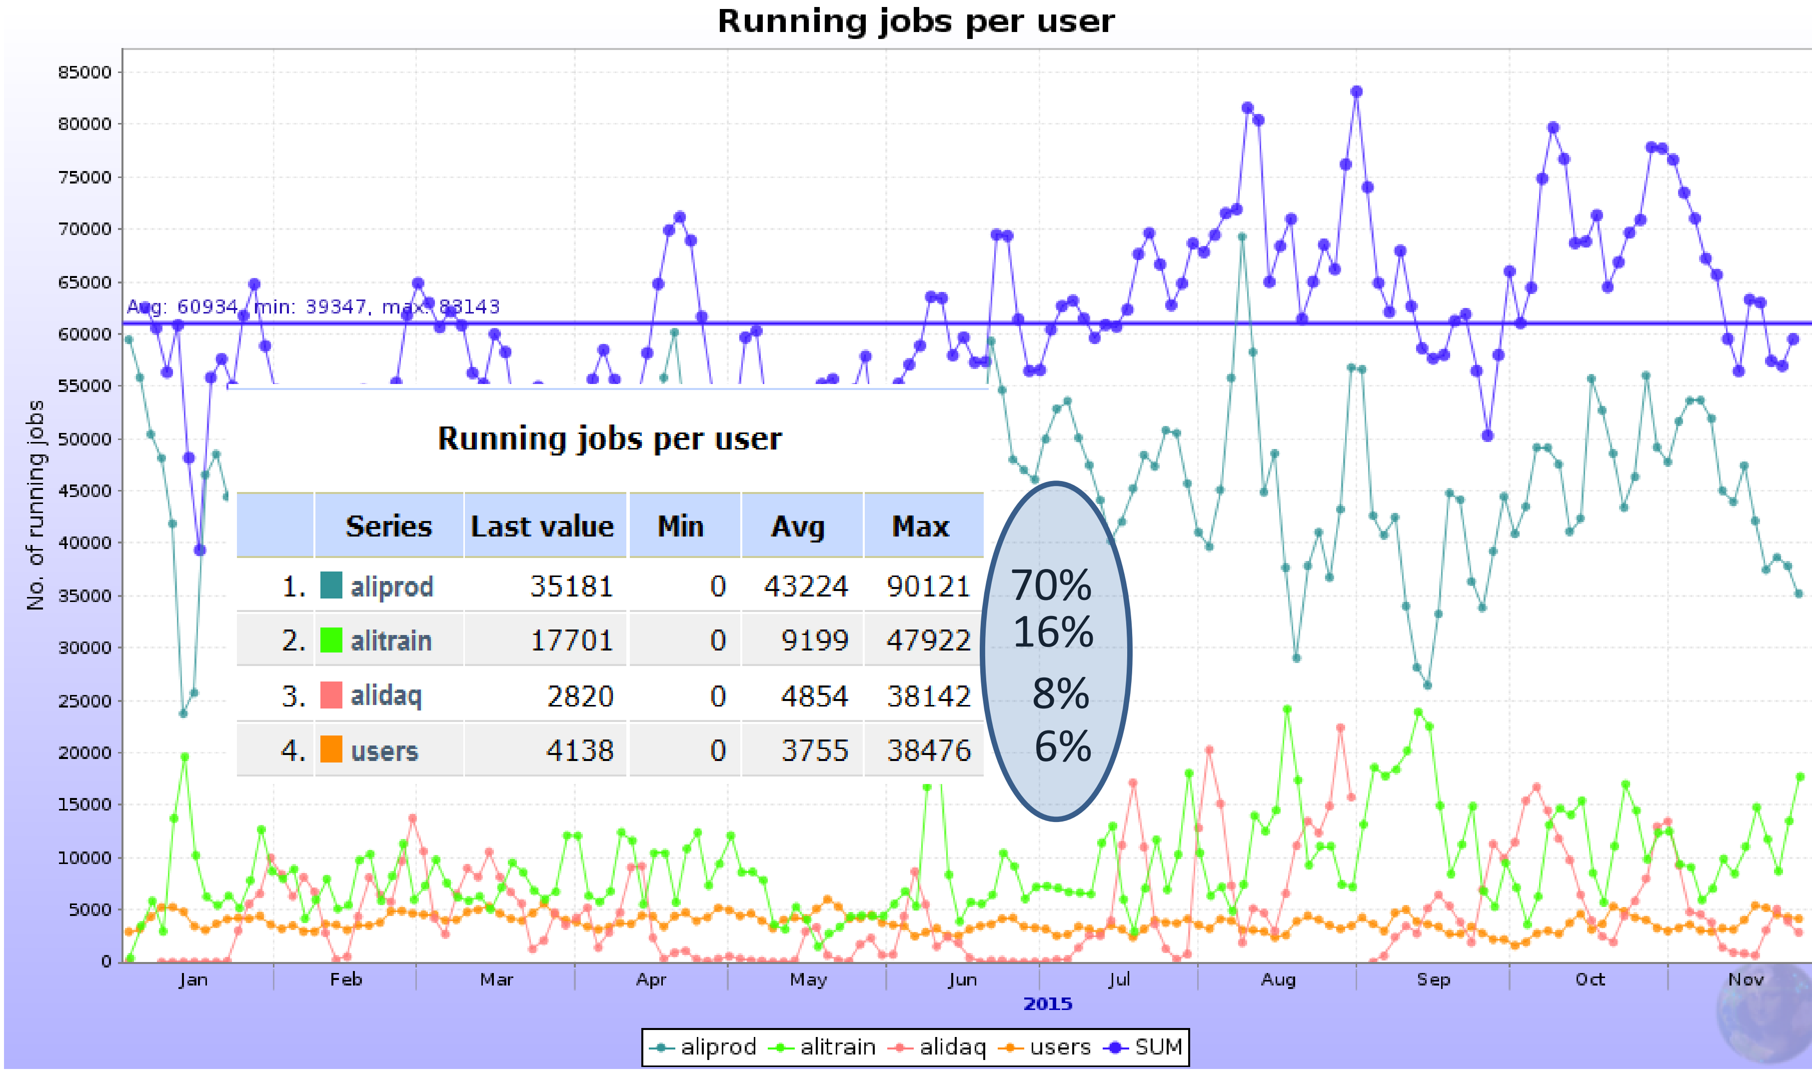

	Series	Last value	Min	Avg	Max
1.	aliprod	35181	0	43224	90121
2.	alitrain	17701	0	9199	47922
3.	alidaq	2820	0	4854	38142
4.	users	4138	0	3755	38476

Data processing in 2015

- MC – 96 cycles – usual number of productions
  - 1,952,334,979 events
  - p-p, p-Pb, Pb-p, some Pb-Pb, some G4
- RAW – reprocessed Run1 data

Production	Description	Status	Run Range	Runs	Chunks	Size	Chunks		Size		Events
LHC10g_pass4	LHC period LHC10g - Full production pass 4, ALIROOT-5311	Completed	135941 - 136193	10	5,158	13.17 TB	5,080	98%	1.069 TB	8%	18,997,194
LHC10f_pass4	LHC period LHC10f - Full production pass 4, ALIROOT-5311	Completed	133005 - 134304	26	32,502	85.78 TB	32,374	99%	8.696 TB	10%	106,533,766
LHC10e_pass4	LHC period LHC10e - Full production pass 4, ALIROOT-5311	Completed	127712 - 130850	166	108,038	282.4 TB	106,107	98%	30.47 TB	10%	314,214,914
LHC10d_pass4	LHC period LHC10d - Full production pass 4, ALIROOT-5311	Completed	122372 - 126437	107	66,827	174.6 TB	65,566	98%	19.95 TB	11%	245,147,842
LHC10c_pass4	LHC period LHC10c - Full production pass 4, ALIROOT-5311	Completed	118503 - 121040	91	37,843	98.47 TB	37,715	99%	16.16 TB	16%	162,461,274
LHC10b_pass4	LHC period LHC10b - Full production pass 4, ALIROOT-5311	Completed	114751 - 117222	83	10,526	25.63 TB	10,455	99%	2.854 TB	11%	47,475,443
					260,894	**680.1 TB**	**257,297**		**79.2 TB**		**894,830,433**

Run 2 progress

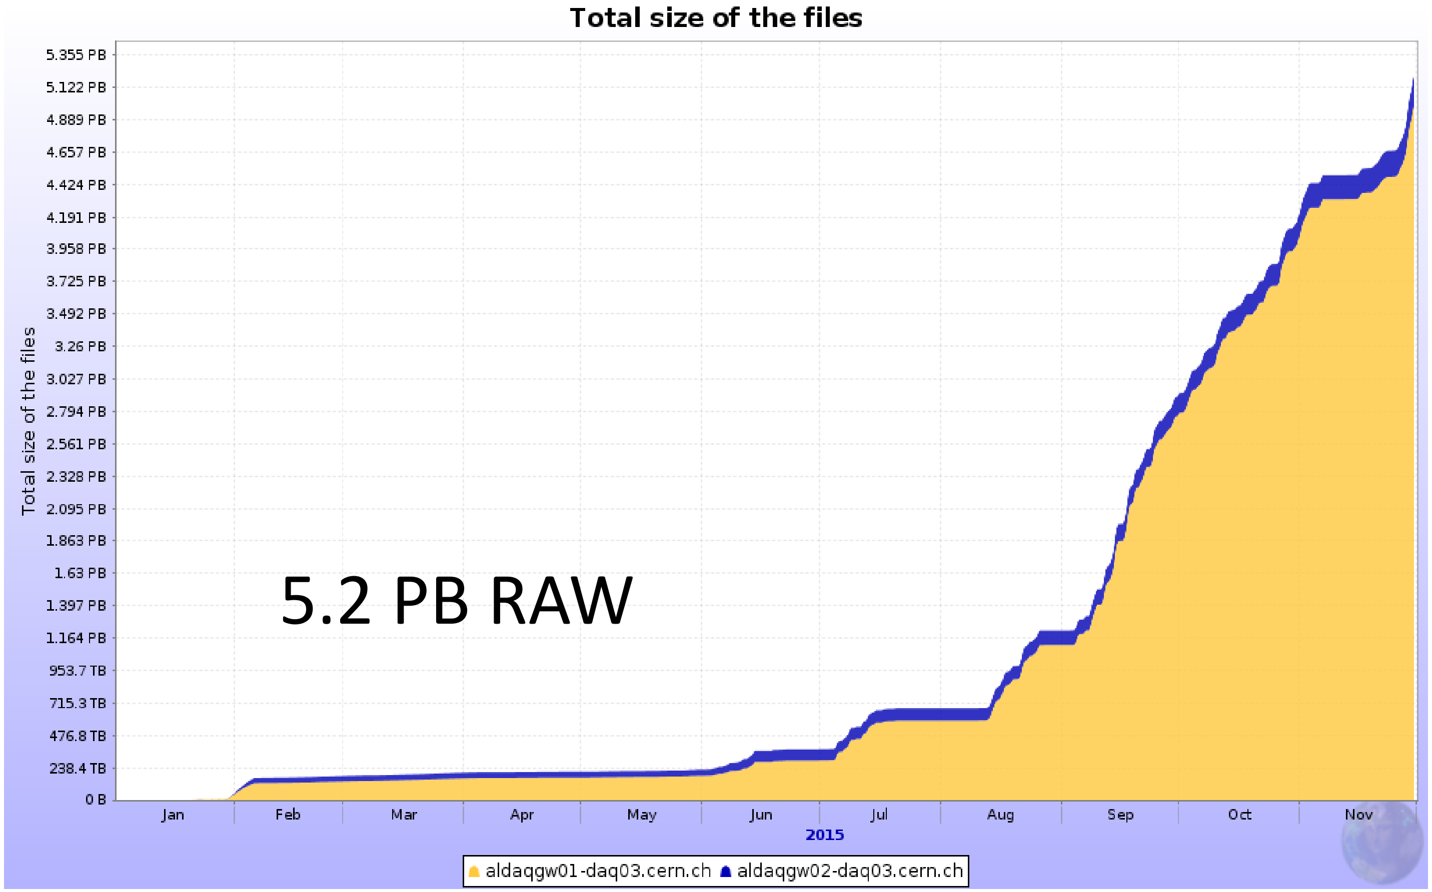

- Registered data volume in 2015 – equal to the sum of 2010-2013

pp $\sqrt{s}$ = 13 TeV

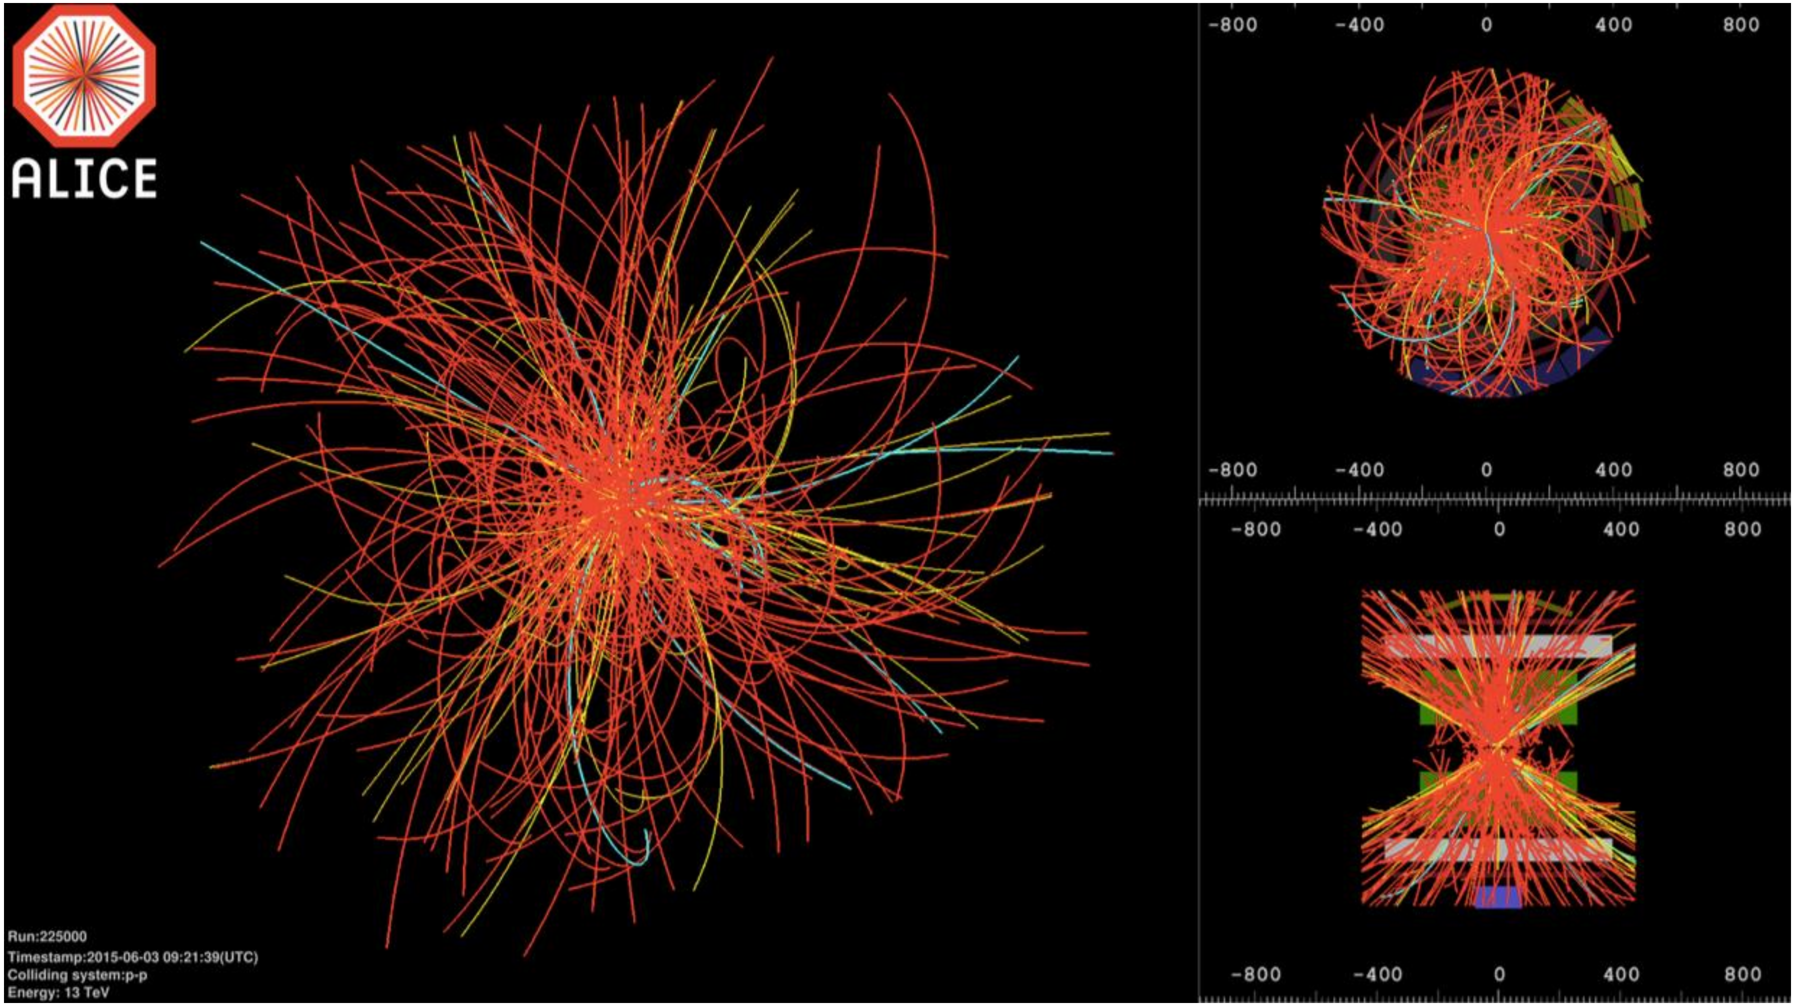

- At highest interaction rates, events with pileup comparable to Pb+Pb collisions
- TPC tracking code was modified to be able to correct for large distortions caused by the space charge build-up

Test on 400 KHz interaction rate

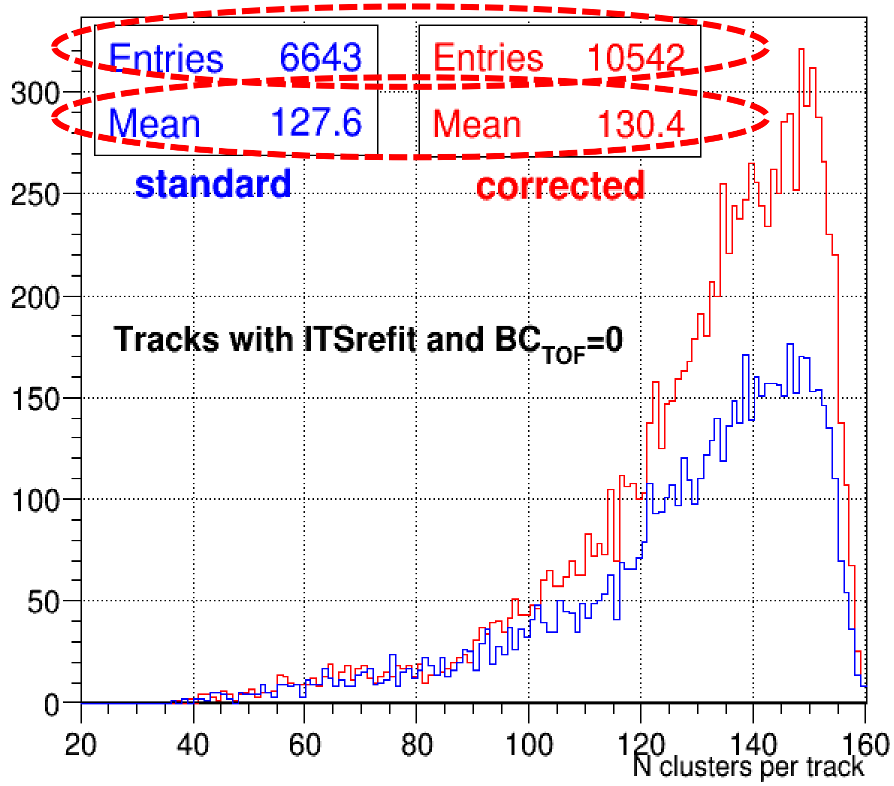

- New calibration step required to take these (time dependent) distortions
- Comparing standard reconstruction with reconstruction using new correction:
  - ITS Matching rate for triggered BC increases by ~60%
  - N TPC clusters per track also increases

Pb-Pb at $\sqrt{s_{NN}}$ = 5 TeV

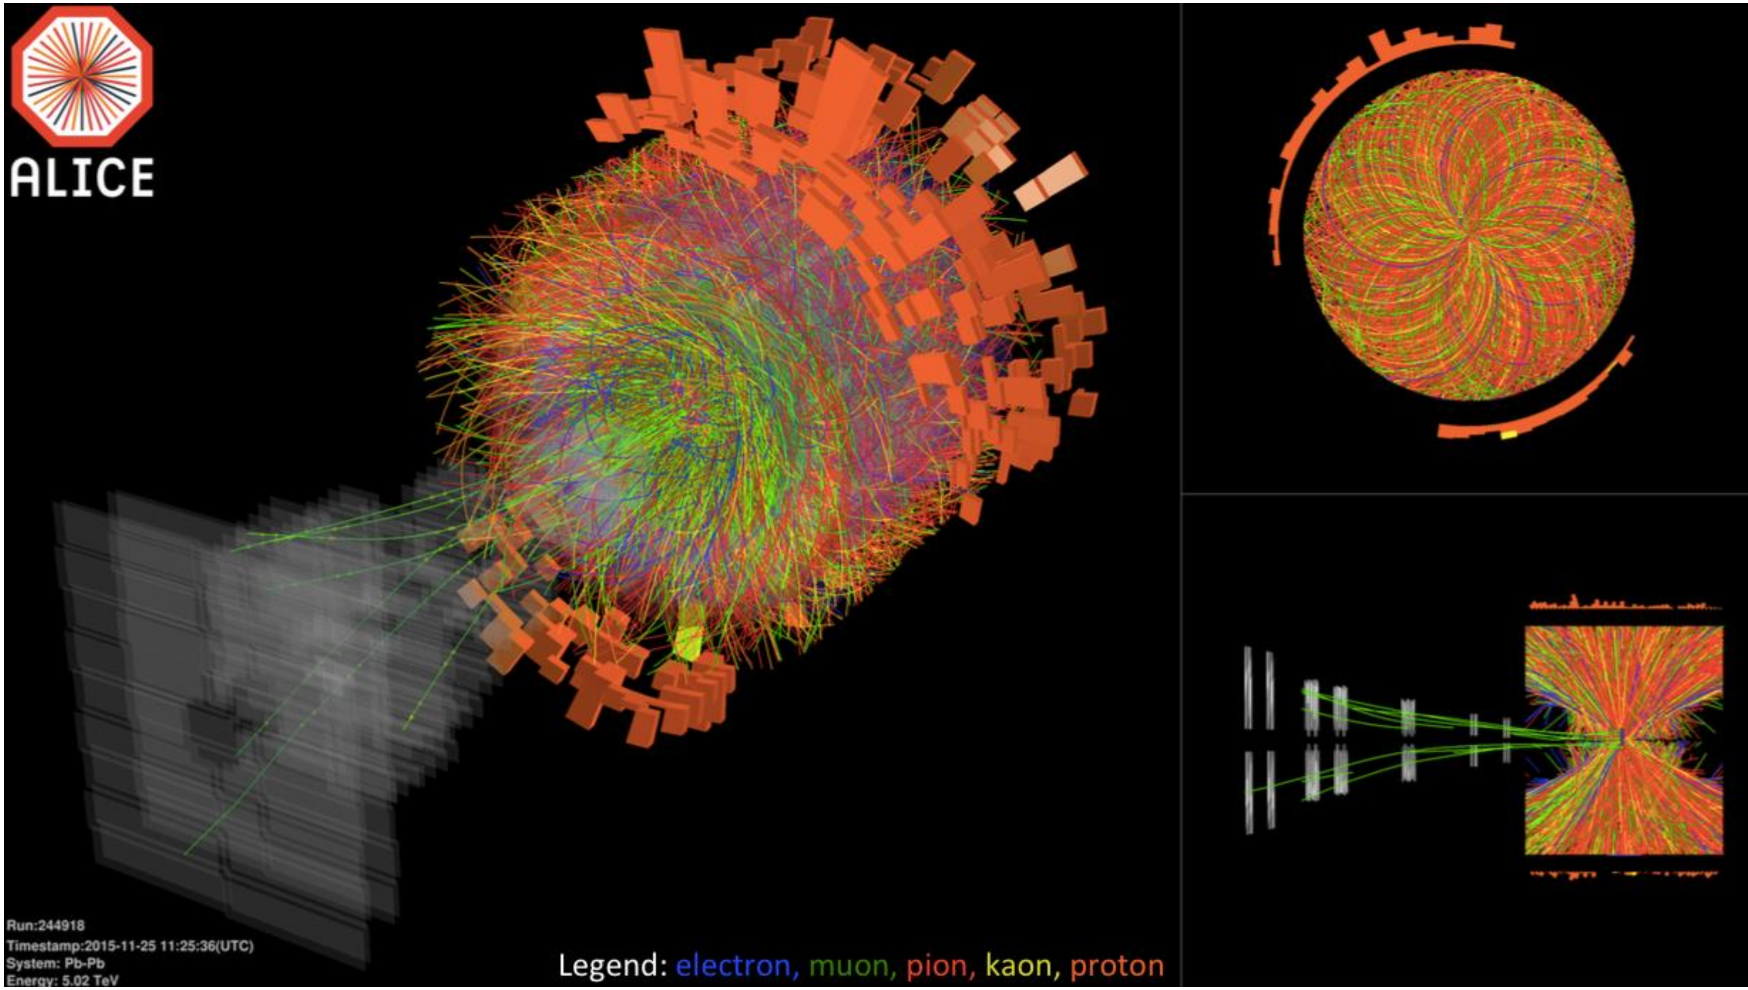

- Initial track seeding in the presence of distortions required to open the cuts in track search leading to increased memory usage
  => Request for temporary large memory queues on T1 sites

Memory consumption

Tests on the worst LHC11h data chunk

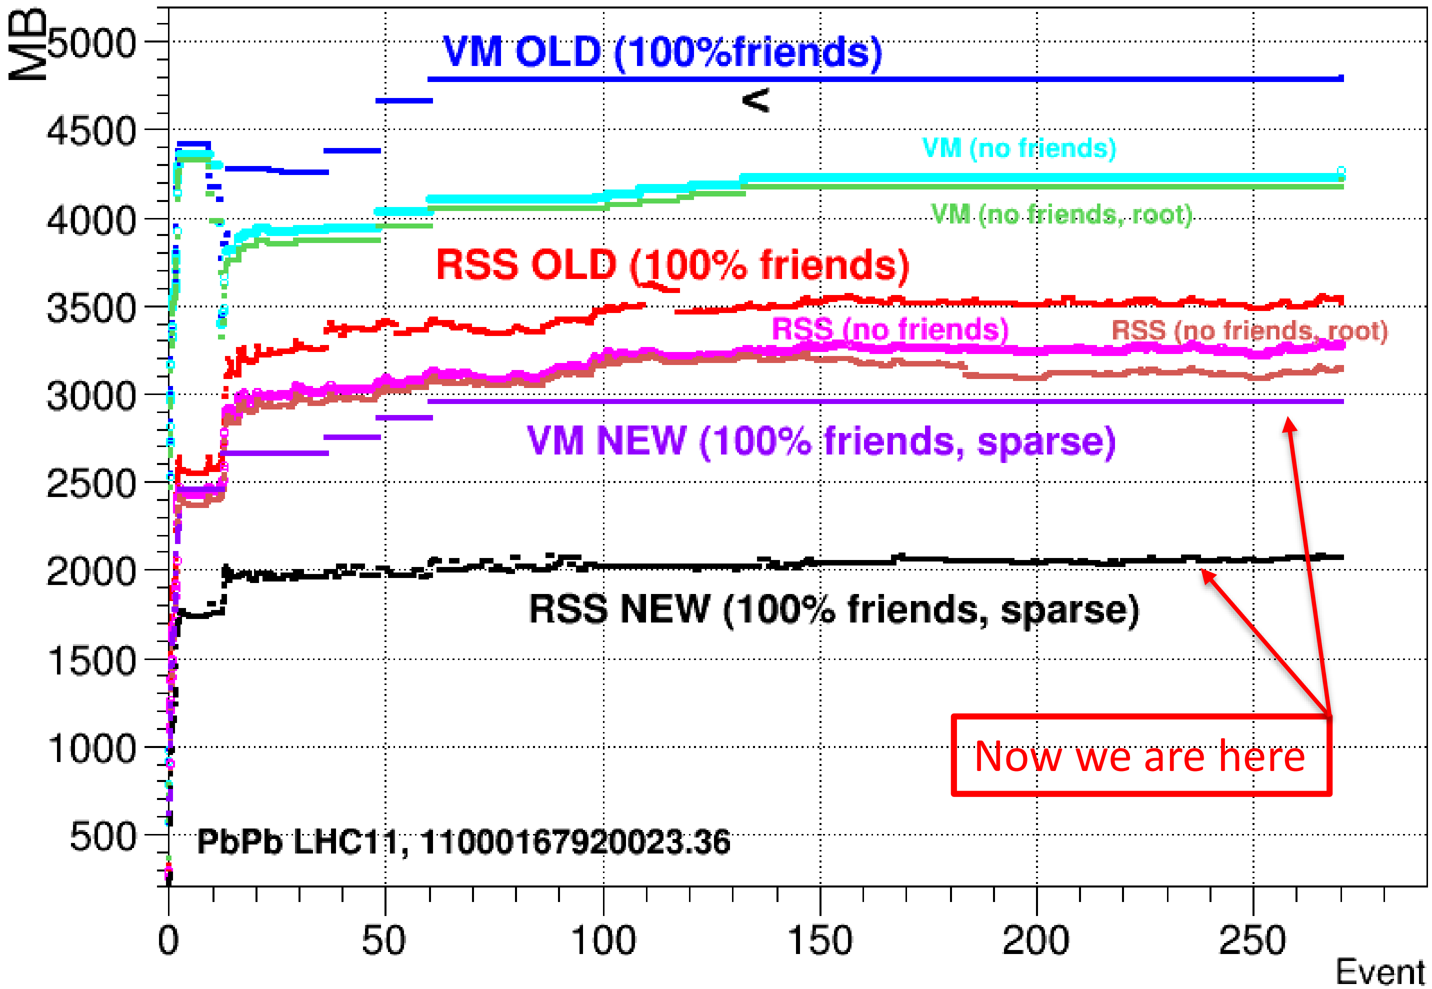

- Last results obtained with modified code
  - We remain under 2 GB/job limit for reconstruction

CPU performance

Tests on the same LHC11h data chunk

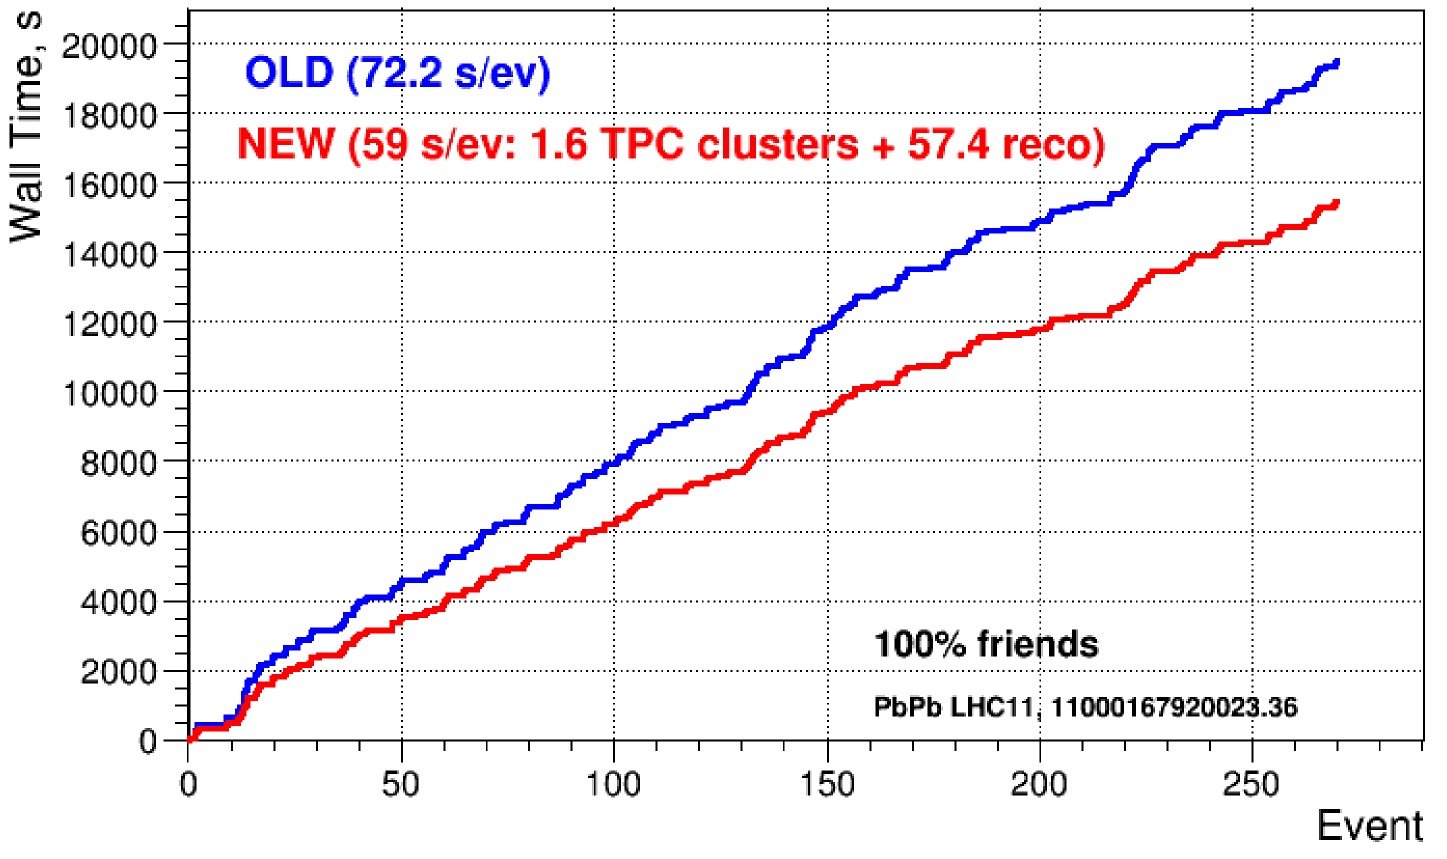

- Gain of ~20% as by-product of memory optimization and optimization of TPC code
- Unfortunately this gains were quickly lost due to increased complexity of events in Run 2

Status of Pb-Pb processing

- Data taking
  - Smooth operation
  - 40% of the data is replicated to T1s

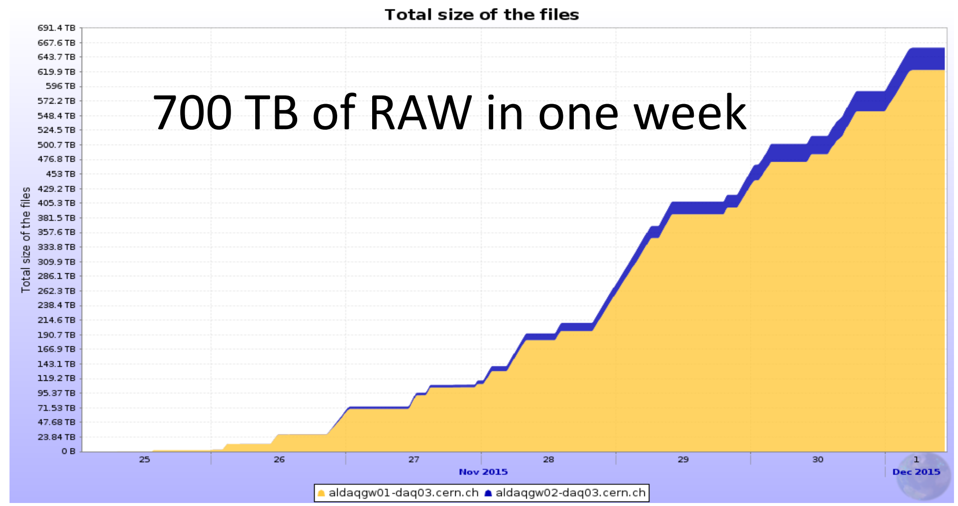

- Processing
  - Fast cycle for Muon analysis and calorimeter calibration
  - Full offline calibration cycle for global processing
  - Special selection of runs for first physics

Description	Status	Run Range	Runs	Chunks	Size	Chunks		Size		Events
LHC period LHC15o - Muon+Calorimeters reconstruction pass 1	Running	244917 - 245497	40	422,960	394.8 TB	419,132	99%	5.524 TB	1%	94,707,948
LHC period LHC15o - Full production pass 1	Running	244918 - 245145	2	18,641	17.43 TB	18,591	99%	6.481 TB	37%	4,948,932
LHC period LHC15o - CPass1 (reconstruction) for pass 1	Running	244917 - 245496	30	265,690	248 TB	208,255	78%	906.5 GB	0%	1,658,381
LHC period LHC15o - CPass0 (reconstruction) for pass 1	Running	244917 - 245497	39	405,135	379 TB	353,225	87%	1.252 TB	0%	12,592,106

Follow-up to O2 TDR

- 6-Oct: Meeting convened by the DRC with the 4 experiments to understand the global requirements
- 16-Oct: Meeting of ALICE (with IT and the WLCG

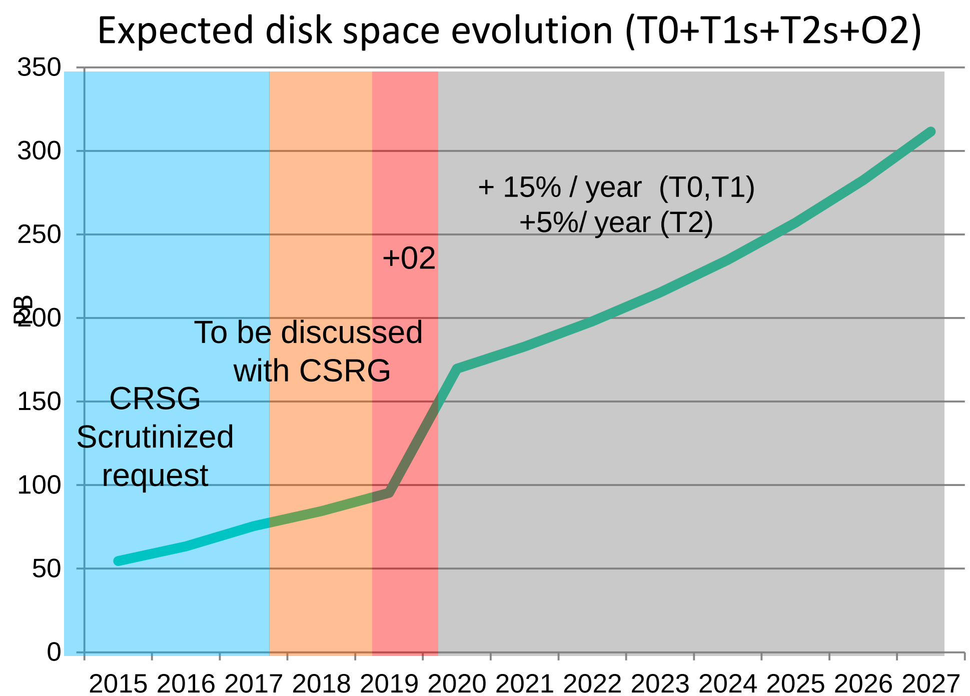

- Conclusions
  – Clarified projected evolution of resources at CERN
  – O^2 architecture compatible with the WLCG.
  – Bandwidth need from P2 to the Tier 0 will need to be increased.
    - precise costing is being established.
  – Bandwidth needs to the Grid will also increase and the requests to the Tier 1's will be done in the near future.

First O2 milestones

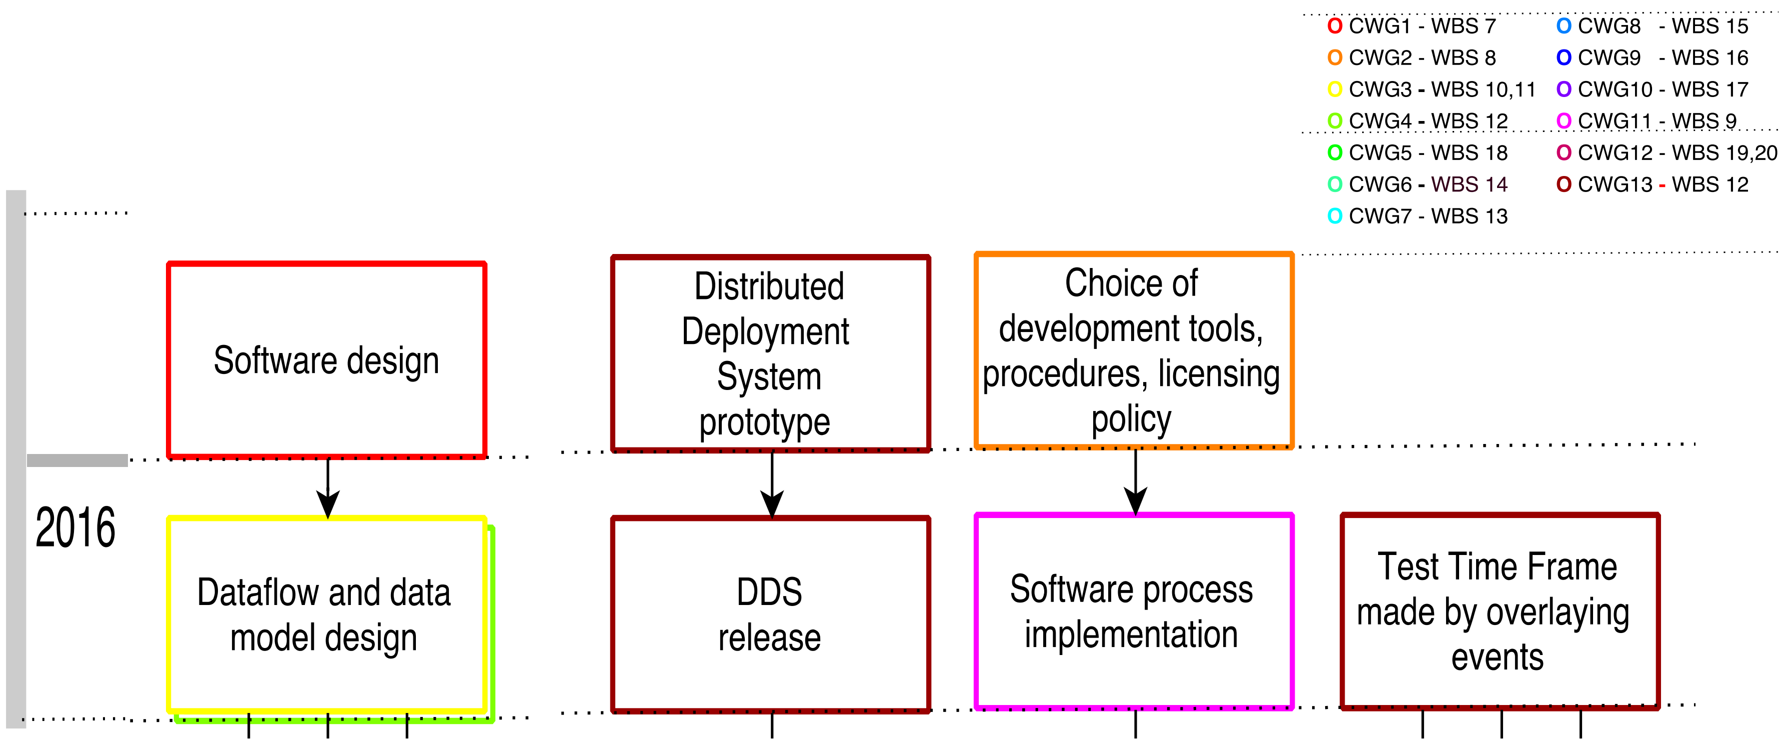

Dynamical Deployment System (DSS)

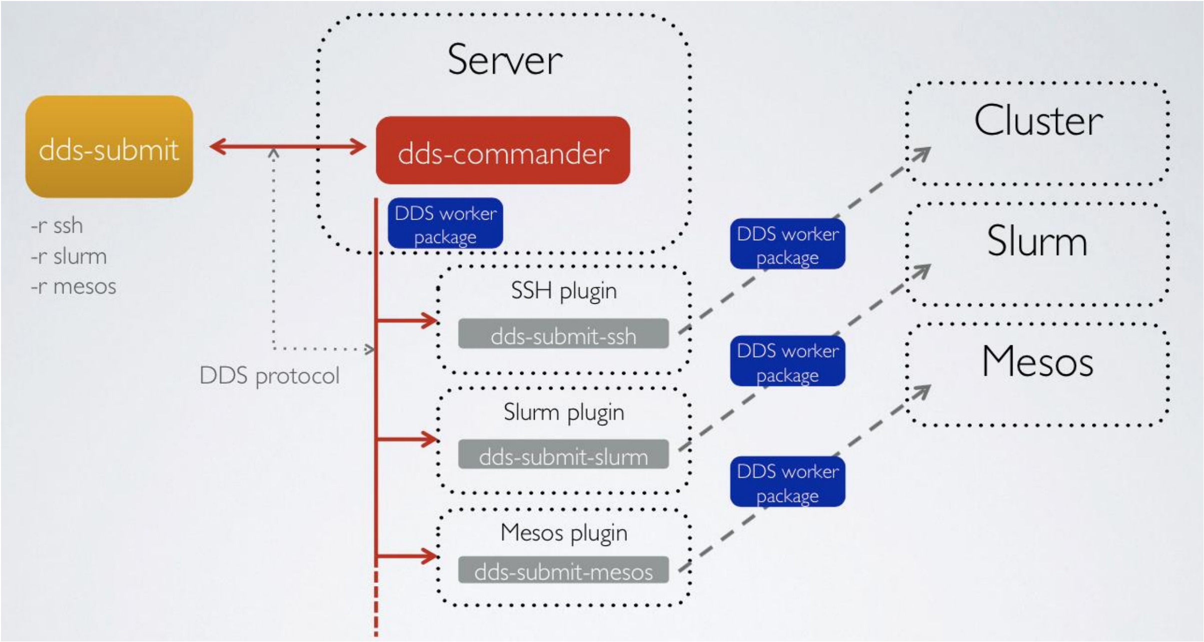

- Current stable release - DDS v1.0 (2015-11-20, http://dds.gsi.de/download.html)
- Home site: http://dds.gsi.de
- User's Manual: http://dds.gsi.de/documentation.html
- DDS now provides pluggable mechanism for external schedulers

Software Lifecycle

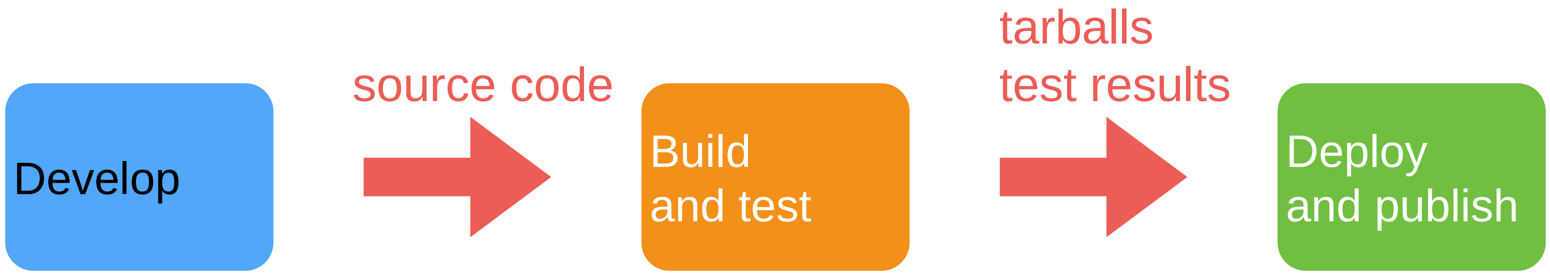

Git-based Integration builds CVMFS, RPM, Ubuntu

- 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: In progress
  - Automatic test and flagging of Pull Requests for O²: In progress
- Basic infrastructure is ready way ahead of time: as of Sep 9, old system

Software Tools and Procedures

ALICE

O² Report
Computing Working Group 11
Software lifecycle

Evaluation of Git hosting platforms

version 1.0

Dario Berzano[a], Giulio Eulisse[a], Barthélémy von Haller[b], Peter Hristov[a]
[a] *CERN, Geneva, Switzerland*

Nov 10, 2015

Abstract

The ALICE O² project will be adopting an *agile* software lifecycle model like all modern software projects. This model is characterized by moving from the proposed source code to a release, and integrating possible patches very quickly: software releases are supposed to happen frequently and quickly as opposed to legacy models favoring long-living releases.

Software lifecycle from code integration to build deployments is composed of several steps, each one of them being error prone and potentially difficult to accomplish. For this reason, we plan to reduce user interaction to the minimum: most operations will be performed automatically, and human interventions are trivialized.

We have chosen Git[1] as version control system. Git itself is very powerful and as such it requires an appropriate workflow to be implemented on top for using it proficiently. The GitHub hosting platform[2] has popularized the *pull requests* model and a web interface standard, also adopted by other platforms, notably GitLab[3]. The web interface is the place where most of the user interaction with the workflow occurs.

In this report we compare GitHub and GitLab in relation to our agile model: finally, we illustrate our final recommendation for the O² project.

ALICE O² C++ Coding Guidelines

Vasco Barroso
Alina GrigoraȘ
Ivana Hřivnáčová
Matthias Kretz
Adriana Telesca
Barthélémy von Haller

This document is based on the work of
B. Weinberger, C. Silverstein,
G. Eitzmann, M. Mentovai
and T.Landray
at http://google-styleguide.googlecode.com,
C++ Google Style guide, Revision 3.274
under the CC-By 3.0 License

Each style point has a summary for which additional information is available by toggling the accompanying arrow button that looks this way: ▶ . You may toggle all summaries with the big arrow button:

▶ Toggle all summaries

▶ Toggle all extra details

Table of Contents

Common license agreed: GPL v3 for ALICE O2 and LGPL v3 for ALFA

Summary

- Data processing in 2015 follows usual pattern
  - All requests have been fulfilled
  - Number of tasks in the pipeline is manageable
  - MC is, as usual, the main resources user (70%), followed by user tasks (22%) and RAW (8%)
- Resources and infrastructure is able to cope with the production load
  - Computing resources are stable and growing
- Unexpectedly large distortions in TPC at high interaction rate required an extra calibration step and additional software development
  - Now under control including memory per job
- Work on O2 has started following the plan presented in the TDR