

First production with the Belle II distributed computing system

T Kuhr for the Belle II Distributed Computing Group

Karlsruhe Institute of Technology, Institut für Experimentelle Kernphysik,
Wolfgang-Gaede-Str. 1, 76131 Karlsruhe, Germany

E-mail: Thomas.Kuhr@kit.edu

Abstract. The next generation B factory experiment Belle II will collect huge data samples which are a challenge for the computing system. To cope with the high data volume and rate, Belle II is setting up a distributed computing system based on existing technologies and infrastructure, plus Belle II specific extensions for workflow abstraction. The system was successfully tested in two production campaigns this year and valuable information for the further development was obtained.

1. Introduction

Belle II [1] is the successor of Belle [2], a B-factory experiment at the KEKB accelerator in Japan, which discovered CP violation in B meson decays together with the BaBar experiment [3] at the PEP-II accelerator in the US. This discovery confirmed the theory of Kobayashi and Maskawa [4] who were awarded the Nobel prize in 2008. Their theory explains the CP violation in the standard model of particle physics, but fails by many orders of magnitude to explain the matter-antimatter asymmetry that we observe in the universe today. The mission of the next generation B-factory experiment Belle II is to search for new particles or forces that could generate the observed asymmetry. As current experiments do not observe significant deviations from the standard model, new physics effects may only be discovered in measurements with very high precision. Therefore Belle II and the upgraded accelerator SuperKEKB are designed for a data rate 40 times higher than Belle.

The upgrades of accelerator, detector, and computing system are actively in progress. The commissioning is planned for early 2015 and the first physics data taking for 2016. SuperKEKB is expected to reach its design luminosity, which is 40 times higher than at KEKB, around 2020 so that a data sample of 50 ab^{-1} can be collected by the end of 2022.

2. Computing model and distributed computing system

A consequence of the drastically increased data rate and volume compared to Belle is an increase in required computing resources by about the same factor. Because the requirements on storage, processing power, and network bandwidth will be at a similar scale as for the LHC experiments and because the collaboration has become more and more international, Belle II has decided to adopt a distributed computing model.

It has a similar hierarchical structure to the Worldwide LHC Computing Grid (WLCG), but is a bit simpler. The raw data is recorded and processed at KEK and copied only to one other site,

PNNL (Pacific Northwest National Laboratory). This provides a backup copy and the possibility to distribute reprocessing tasks over both sites. The output of the raw data processing are mDST files, a format containing all relevant information for physics analyses, which are copied to grid sites around the world. These sites also produce simulated data samples (Monte-Carlo, MC) in mDST format. Physicists process the mDST files on grid site and transfer the output to local resources where the final analysis is performed. The Belle II computing model is illustrated in figure 1.

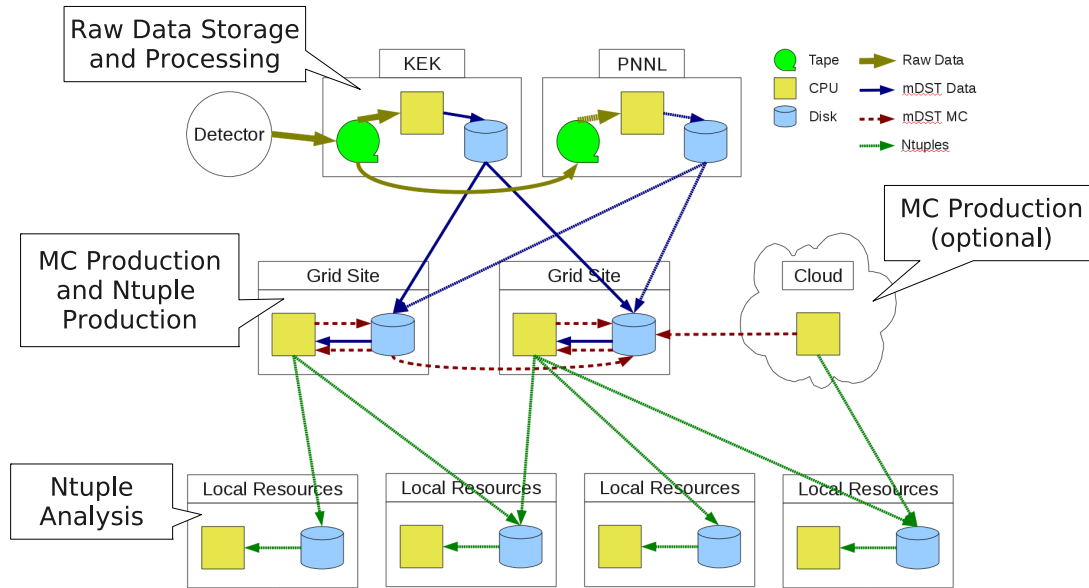


Figure 1. The Belle II computing model.

In the design of our distributed computing system we rely on existing and well-proven infrastructure and software. By choosing DIRAC [5], which was originally developed by LHCb, as the basis for our distributed computing system we can use resources provided via different middleware interfaces, including WLCG, Open Science Grid (OSG), and cloud resources. The location of output files is stored in the Logical File Catalog (LFC). For each file we record metadata in an AMGA database [6]. The installation of offline software releases is done conveniently via CVMFS (CERN Virtual Machine File System) [7], a system for mounting and caching a file system from a remote server, on most sites. The CVMFS stratum-0 server is kindly provided by CERN. A stratum-1 server is running at GridKa. At those sites without CVMFS support, mainly OSG sites, the software is distributed via installation jobs.

On top of these components Belle II has developed a user interface tailored to the needs of the Belle II collaborators. The interface is called gbasf2 and is designed to provide an easy transition from an analysis done with the offline software framework basf2 [8] to an analysis performed on the grid. A guiding design principle was to provide an abstraction of the workflow from single files and jobs to datasets and collections of jobs, called projects, as shown in figure 2. The idea is that the user selects a set of input files based on metadata criteria and then processes them with the same analysis code. Instead of submitting multiple jobs for multiple files, the user starts just one project for the input dataset and then can refer to the collection of individual jobs via a chosen project name. This project-level management of jobs comprises the submission of jobs, the monitoring of jobs, the cancellation of jobs, and the retrieval of their output sandboxes. The

workflow abstraction is also applied to the output files of a project which can be collectively referred to as an output dataset.

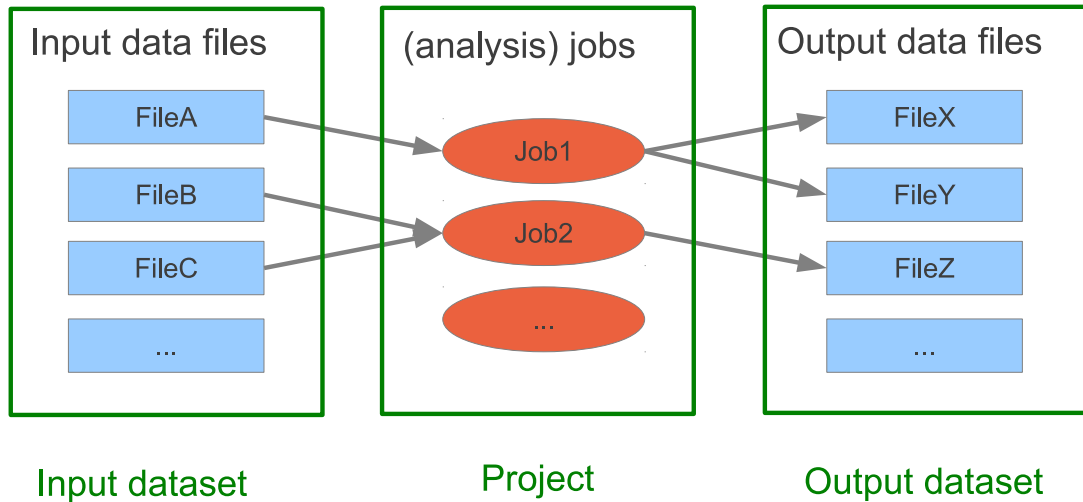


Figure 2. Abstraction of workflow with datasets and projects.

3. Production campaigns and data challenges

To evaluate the computing model and the distributed computing system, Belle II has performed two MC production campaigns so far. The first campaign started on February 8th, 2013. In a first stage $B\bar{B}$ events were generated with EvtGen [9] and the detector response simulated based on a geometry implemented with Geant4 [10]. The output are raw data files. They are uploaded to storage elements and their metadata information is registered in AMGA. In a second stage the raw data files were the input to reconstruction jobs.

Up to the end of the first MC production campaign 240k jobs were executed which consumed 40 kHEPspec06 days of CPU resources. 60M events were simulated and reconstructed. The total output size was 190 TB. The job failure rate was at a level of about 20%, mainly caused by errors during the metadata registration process, failures of input data retrieval in the reconstruction stage, and crashes of the offline software.

With the experience of the first MC production campaign at hand, the distributed computing and offline software were developed further to address the observed issues. Another assessment of the system was then done in a second MC production campaign from July 23rd to September 8th, 2013. In this campaign the event generation, the detector simulation, and the reconstruction were executed within one job. As the output was stored in mDST format, the ratio of output size to CPU usage is much lower than in the first campaign. A further difference with respect to the first campaign is that background data is mixed to the simulated signal events. This requires the import of background data files at the beginning of each job. The type of jobs executed in the second MC production campaign is similar to the type of jobs that will produce MC data samples for Belle II physics analysis later during the data taking phase.

In total 630k jobs consumed 700 kHEPspec06 days and produced 560M events, corresponding to an output size of 8.5 TB. Compared to the first campaign the used CPU resources could thus be increased by more than an order of magnitude as illustrated in figure 3. At the same time the job failure rate could be reduced considerably. While in the beginning it was at a level of 10% it reached about 1% at the end. Inefficiencies were caused by issues or downtimes of individual

sites, the expiration of proxys, overloaded servers, and human errors like the submission of jobs with wrong parameters. It should be noted that no crash of the offline software was observed.

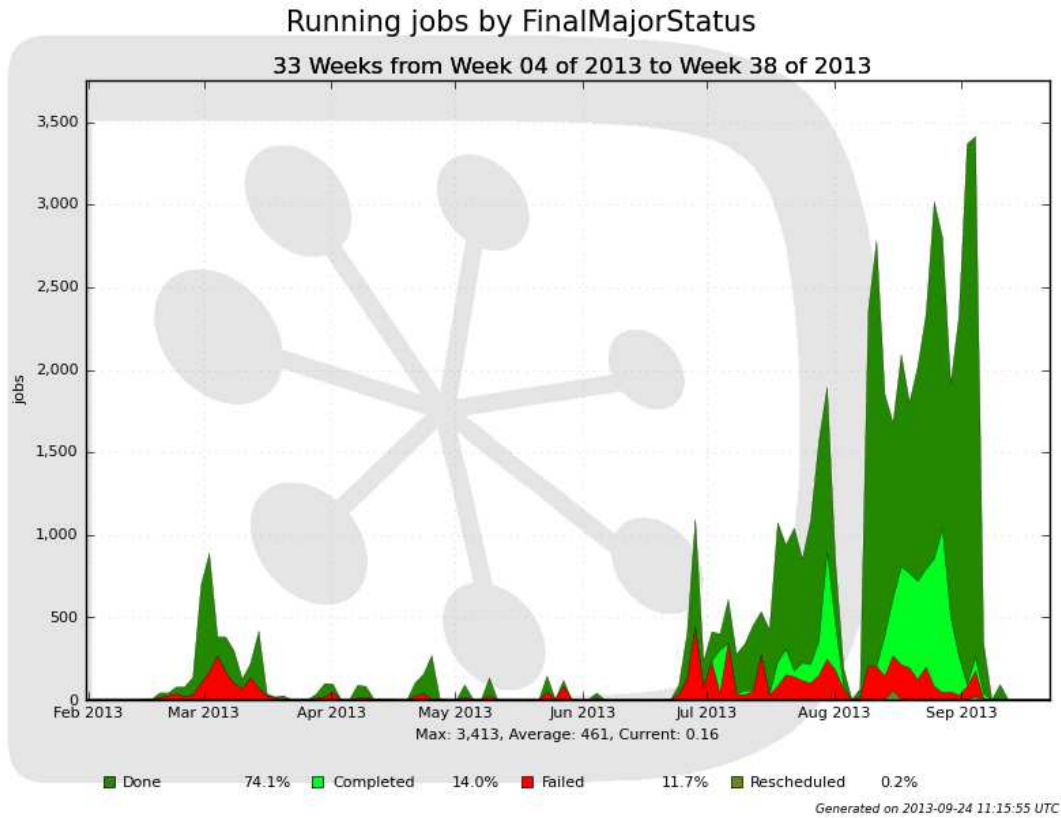


Figure 3. The number of running jobs for the first MC production campaign in March and the second campaign in July to September. Successfully completed jobs are shown in green and failed jobs in red. Light green indicates that the failover mechanism for the upload of the output was triggered.

An important contribution to the increased job efficiency is the output storage failover mechanism provided by DIRAC. If a file cannot be uploaded to the designated storage element (SE) it is first stored on the local SE and later transferred to its final destination. One reason for failing uploads was a too restricted number of pool accounts on the KEK SE that could be fixed during the campaign. The proxy expiration issue was solved by increasing its lifetime to 168 hours. The high load on the DIRAC server was reduced by decreasing the frequency of communications between the jobs and the server and by distributing DIRAC services over more nodes. Issues with the load and memory consumption of the AMGA server are under investigation.

Several sites around the world participated in the second MC production campaign. These are KEK, KMI [Nagoya] (both Japan), KISTI (South Korea), GridKa, DESY (both Germany), CESNET (Czech Republic), CYFRONET (Poland), SIGNET (Slovenia), UA-ISMA (Ukraine), PNNL, and Virginia (both US). In addition to these sites at Belle II collaborating institutions a sizable share of the resources was kindly provided by the two OSG sites Nebraska and FNAL. A summary of the contributed CPU resources is shown in figure 4.

Like the support of the sites the support of shifters was another essential ingredient for the success of the production campaign. More than twenty volunteers kept production running by

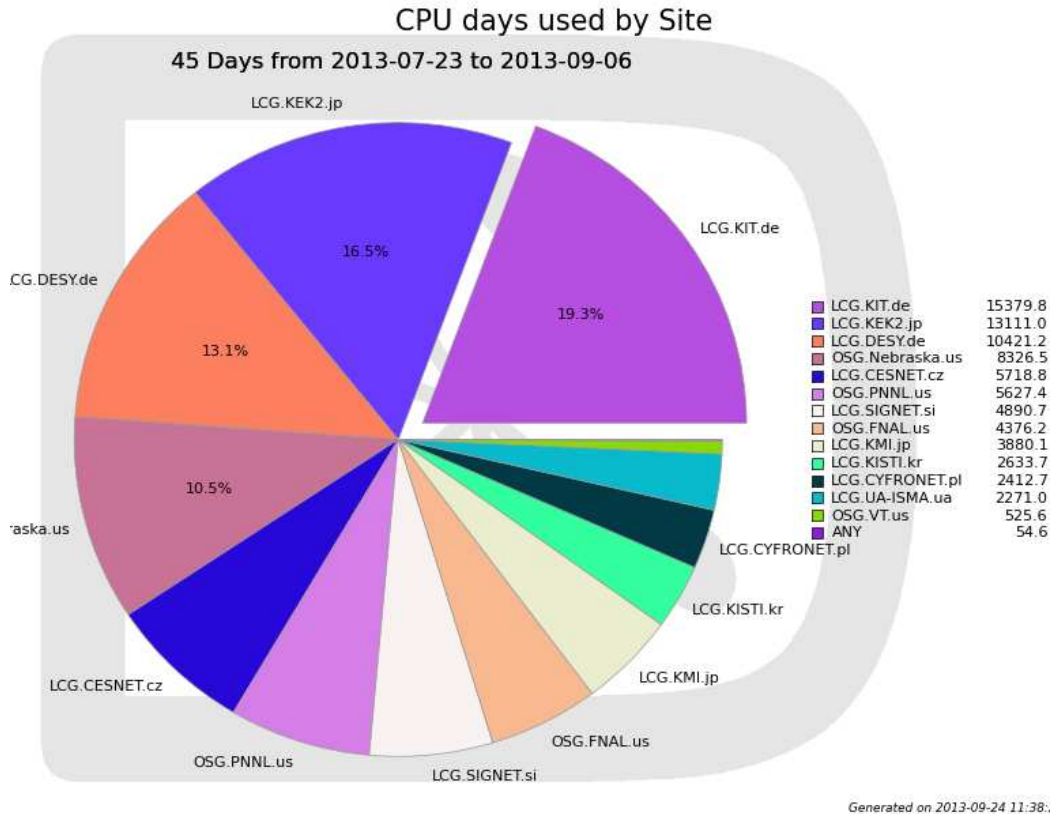


Figure 4. The contributions of sites to the second Belle II MC production campaign.

submitting jobs, monitoring them, and resubmitting failed jobs. Shifts were assigned in 8 hour blocks so that continuous coverage could be achieved by day shifts in Asia, Europe, and America. Log book information was recorded in a twiki. For communication, in particular at the hand over between shifts, it turned out that a video/audio conference with chat using SeeVogh/EVO provided a very efficient and productive solution.

In the Belle II computing model we rely on high-bandwidth network connections between sites, in particular between KEK and PNNL for the transfer of raw data. But also for the replication of MC datasets network connections between all sites are needed. To assess the current status of the network resources available to Belle II, we exercised transfers in a data challenge in May 2013. The transfers were controlled by a FTS2 service running at GridKa. One of the conclusions from the data challenge was that the bandwidth for transfers from Japan to Europe is significantly lower than between Japan and the US.

The monitoring of network, computing, and storage resources is essential for the set up of a production system. DIRAC already provides a wealth of information and plots that proved very useful during the MC production campaigns. The DIRAC and AMGA servers are monitored by ganglia [11] and nagios [12]. PerfSONAR [13] servers at several sites were set up to monitor the network performance. To aggregate the various pieces of monitoring information we are developing a HappyFace instance [14]. An additional value of HappyFace is that it provides a rating of the collected information and can thus be a very powerful tool for shifters.

4. Summary

The computing resources required to process and analyze the data of the next-generation B factory experiment Belle II will be of similar scale as the resources requirements of the LHC experiments. Therefore Belle II has adopted a distributed computing model which enables all collaborators to contribute. The distributed computing system is based on existing technologies and extended by Belle II specific features. An essential concept is the workflow abstraction from individual files and jobs to datasets and projects.

This concept and underlying infrastructure were successfully tested in two MC production campaigns this year. This success was only possible by the CPU, storage, and network resources kindly provided to us by sites and NRENs. The campaigns showed the issues and bottlenecks of our system and we managed to address several of them already during the production campaign so that a final job failure rate of about 1% was reached.

The next MC production campaign is planned for the end of this year. This time we expect further sites to contribute and to even include cloud resources. Another important use case of the system is the user analysis on the grid which may show completely different issues than a centrally managed MC production.

Acknowledgments

We are grateful for the support and the provision of computing resources by CESNET, Cyfronet, DESY, GridKa, KEK-CRC, KISTI, KMI, OSG, PNNL, SiGNET, UA-ISMA, and VT. We acknowledge the service provided by ESnet, GEANT, and NII. We thank the DIRAC and AMGA teams for their assistance and CERN for the operation of a CVMFS server for Belle II.

References

- [1] Abe T *et al.* arXiv:1011.0352 [physics.ins-det]
- [2] Abashian A *et al.* 2002 *Nucl. Instrum. Meth. A* **479** 117
- [3] Aubert B *et al.* 2002 *Nucl. Instrum. Meth. A* **479** 1
- [4] Kobayashi M and T. Maskawa T 1973 *Prog. Theor. Phys.* **49** 652
- [5] Casajus A *et al.* [LHCb DIRAC Collaboration] 2010 *J. Phys. Conf. Ser.* **219** 062049; Tsaregorodtsev A *et al.* 2010 *J. Phys. Conf. Ser.* **219** 062029
- [6] Ahn S *et al.* 2010 *Journal of the Korean Physical Society* **57** issue 4 715
- [7] Blomer J, Buncic P, Charalampidis I, Harutyunyan A, Larsen D and Meusel R 2012 *J. Phys. Conf. Ser.* **396** 052013
- [8] Moll A 2011 *J. Phys. Conf. Ser.* **331** 032024
- [9] Lange D J 2001 *Nucl. Instrum. Meth. A* **462** 152
- [10] Agostinelli S *et al.* [GEANT4 Collaboration] 2003 *Nucl. Instrum. Meth. A* **506** 250
- [11] Massie M L, Chun B N and Cueller D E 2004 *Parallel Computing*, **30** 7 817
- [12] <http://www.nagios.org/>.
- [13] Boote J W *et al.* 2005 *Computational Methods in Science and Technology* **11(2)** 91
- [14] Buge V, Mauch V, Quast G, Scheurer A and Trunov A 2010 *J. Phys. Conf. Ser.* **219** 062057