Distributed Computing at Belle II.

Or "What happens when I use gbasf2 for my analysis?"

Michel Hernandez Villanueva
DESY

Belle II Germany Meeting
Sep 19 - 21, 2022

Belle II

Data Taking at Belle II

Let's talk about numbers

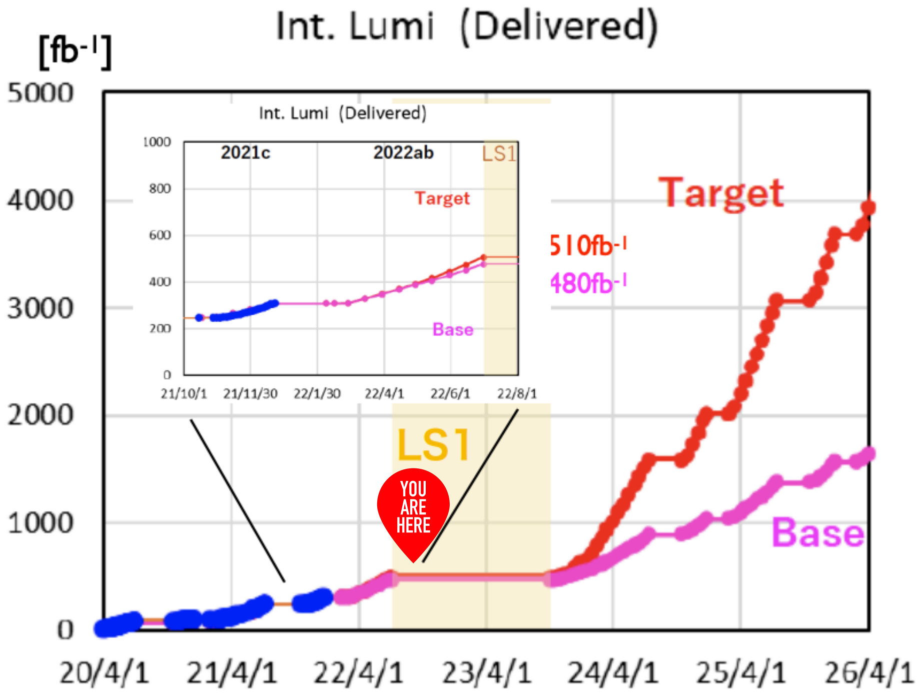

- Integrated luminosity expected by the end of the experiment: 50 ab^{-1}
  (x50 than the previous B factories)

- The estimated size of the dataset collected by the experiment is O(10) PB/year.

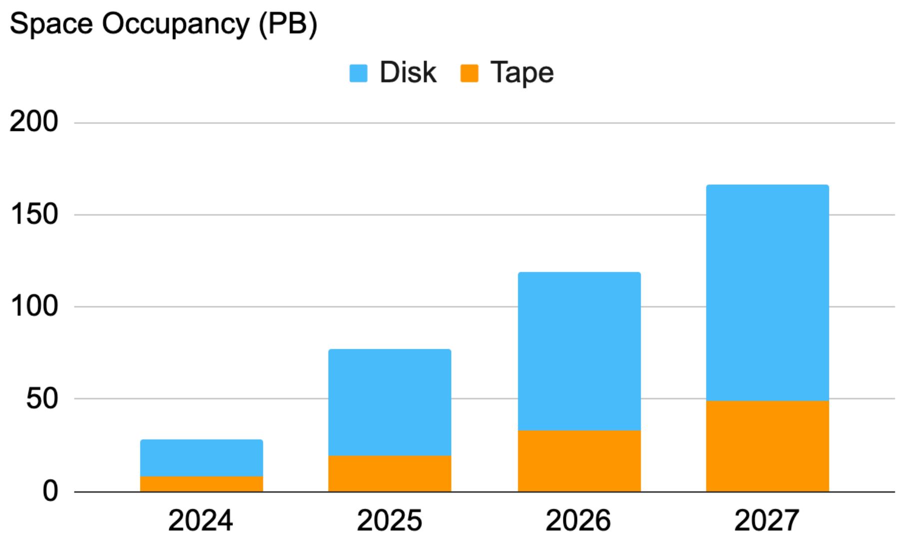

Data Taking at Belle II

Let's talk about numbers

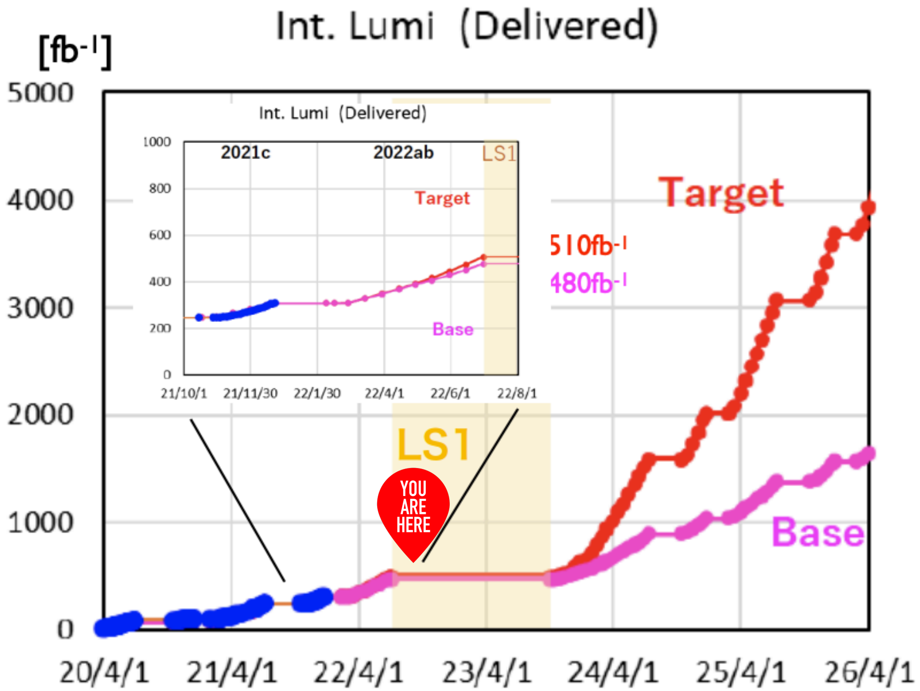

- Integrated luminosity expected by the end of the experiment: 50 ab^{-1}
  (x50 than the previous B factories)

- The estimated size of the dataset collected by the experiment is O(10) PB/year.

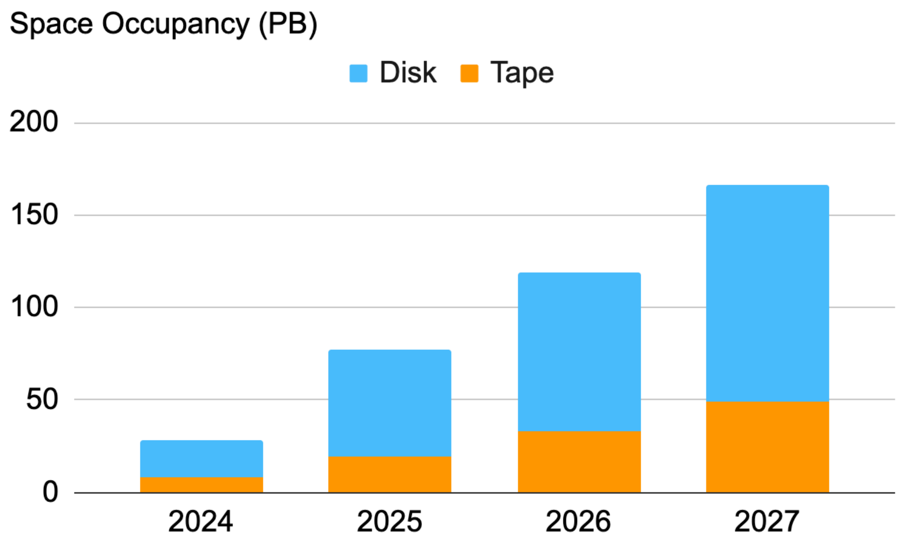

- Not as large when compared to HL-LHC scales, but **corresponds to 10^{12} events,** representing a significant data management challenge.
  - This final Belle II dataset is quite similar to the analysis datasets of the ATLAS and CMS experiments.

Data Taking at Belle II

Let’s talk about numbers

Data is expected to be distributed and analyzed within the whole collaboration.

- This final Belle II dataset is quite similar to the analysis datasets of the ATLAS and CMS experiments.

The Belle II distributed computing system

A.K.A. “The grid”

- A "large virtual system" is composed of many networked loosely computers.
  - Heterogeneous environment.
  - Running Basf2 jobs in several sites is not a trivial task.[1]
- Dedicated data centers keep two copies of the full raw data set.
- Raw data is staged, reprocessed, skimmed and distributed over storage sites.
- Analysers access data sending jobs to the grid and downloading the output.

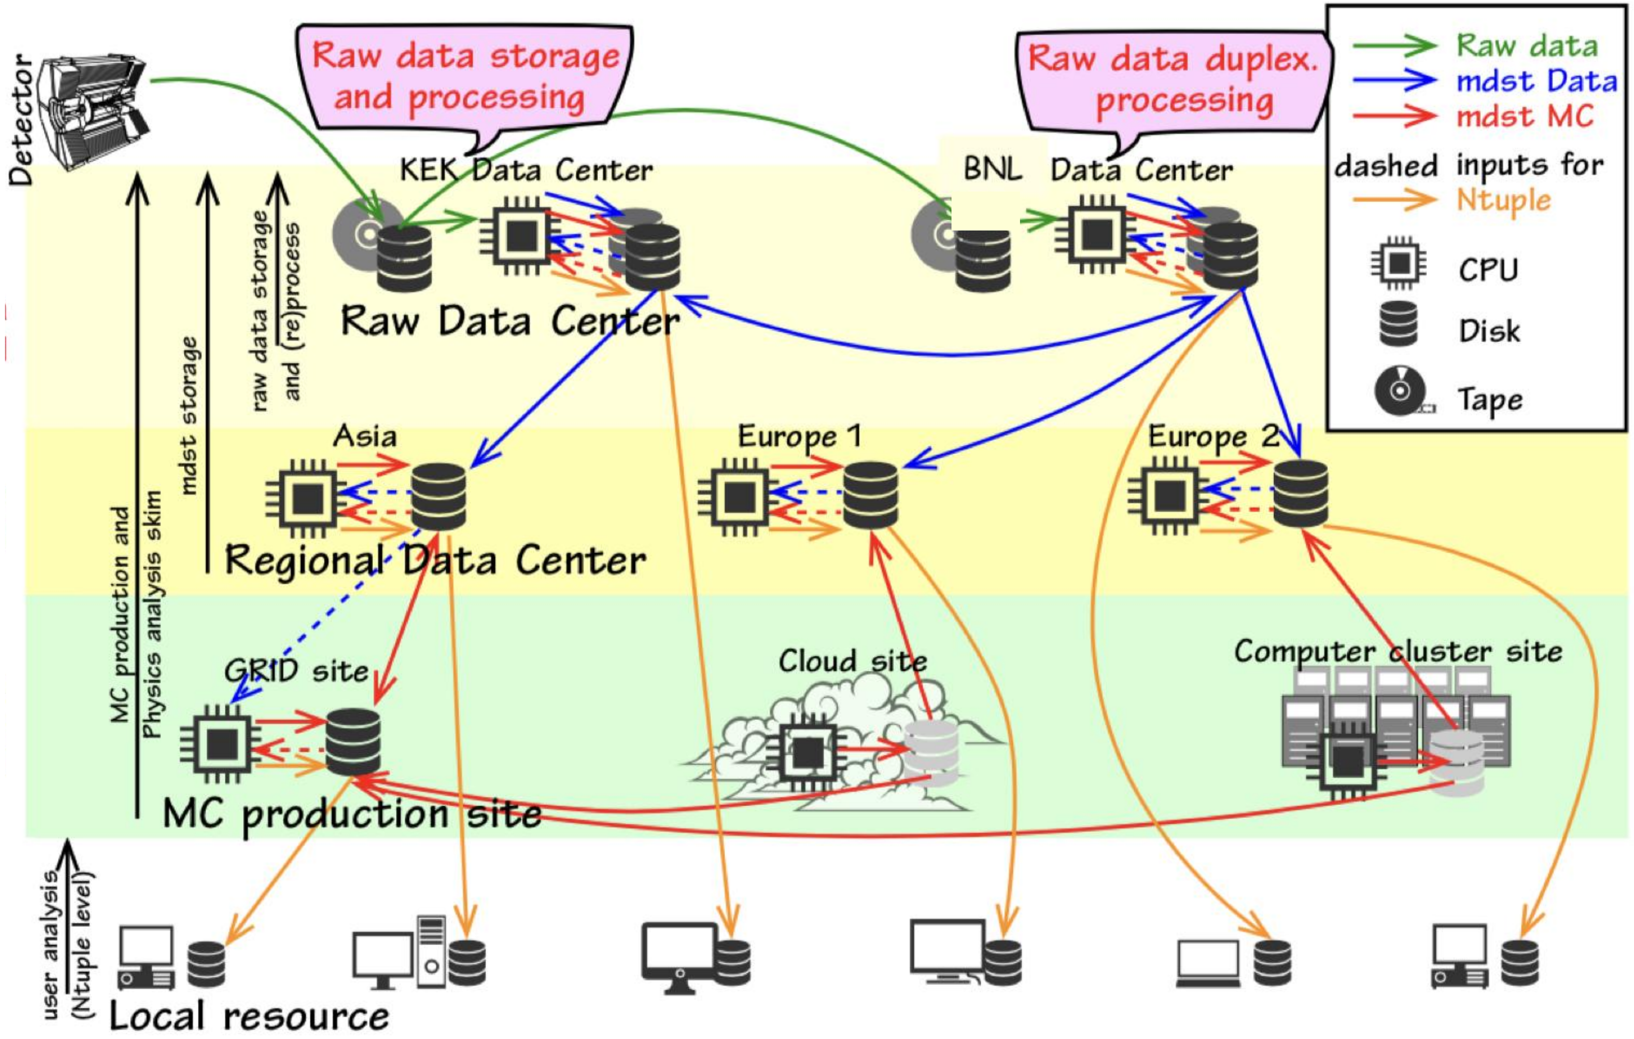

[1] Keep this in mind. At the end I will ask you a favor…

The Belle II Grid

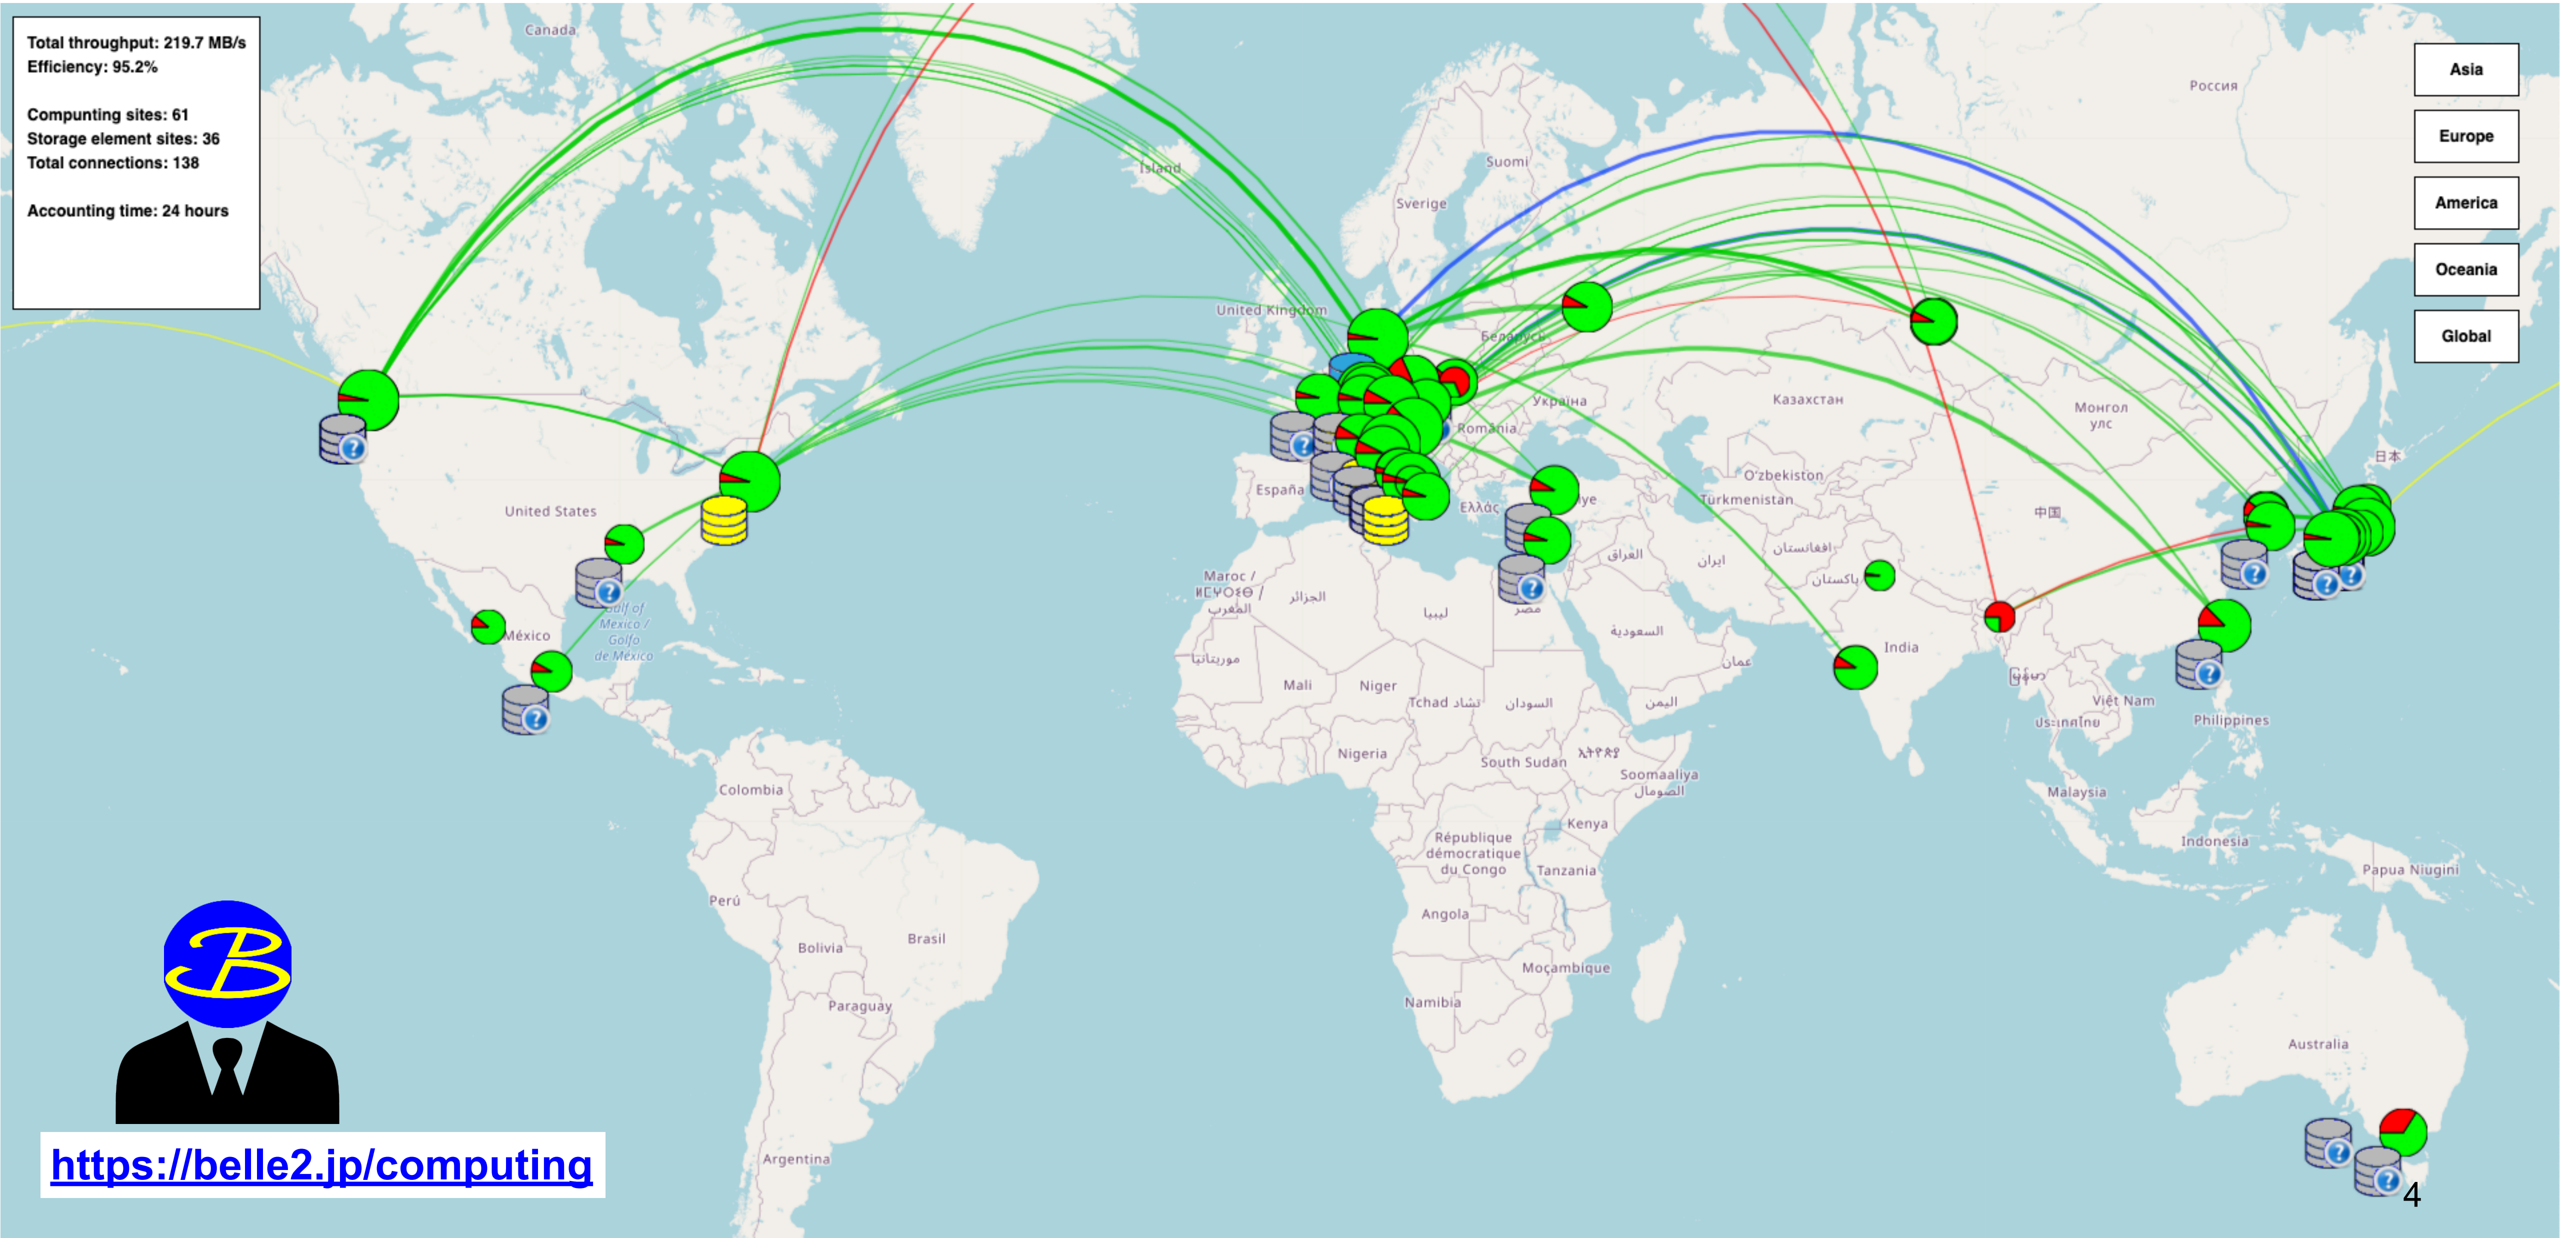

https://belle2.jp/computing

The Belle II Grid

https://belle2.jp/computing

Distributed computing infrastructure at Belle II

Resources available

- **Storage Elements (SEs)**
  - 29 storage sites. 5 Tape systems.
    - 92% of Storage on LHCONE.
    - 8.2 PB reachable via IPv6 over of 13.8 PB.
    - All sites except 3 nominally support HTTP/WebDAV.
- **Sites (CEs)**
  - 55 sites registered in DIRAC.
    Some sites with multiple CEs.
    - 24 Sites Providing Pledged CPUs.
    - 12 Sites Pledged + Opportunistic.
    - 18 Sites Opportunistic Only.
  - Most part of the sites (49) are EL7 based.

Storage	Space (PB)
Disk	13.6
Tape	10.1

CPU	kHS06	Job slots
Pledged CPU	452	31,484
Opportunistic CPU	310	25,377
TOTAL	**762**	**56,861**

Distributed Computing

Architecture Overview

- We adopted DIRAC as the main framework to interact with distributed computing systems.
- **BelleDIRAC** is the extension that handles automated productions, data placement, monitoring, client tools, etc.
- Rucio is a sophisticated, highly scalable data management system.
  - Works as file catalog, takes care of data policies and replication.

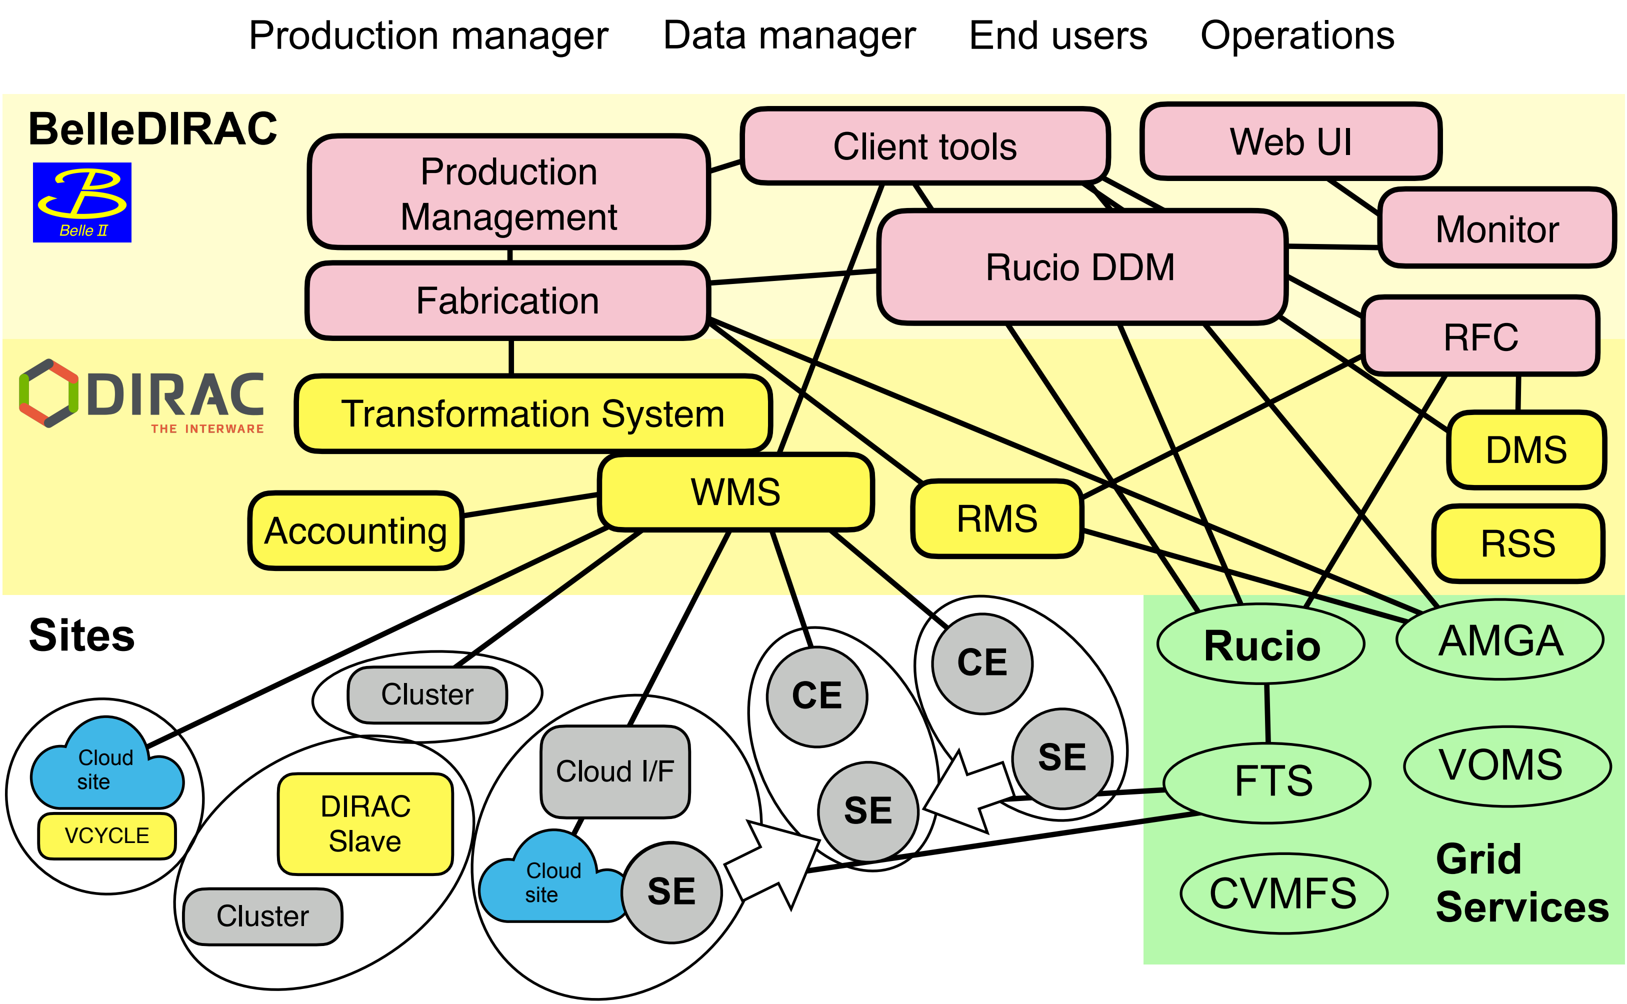

Distributed computing infrastructure at Belle II

Interoperability with DIRAC

- Computing resources in DIRAC are represented by logical Computing Elements with various implementations:
  - Grid: CREAM-CE, ARC-CE, HTCondor-CE
  - Clusters w/o middleware: ssh or DIRAC SiteDirector
  - Cloud: VMDIRAC, CloudSchedule

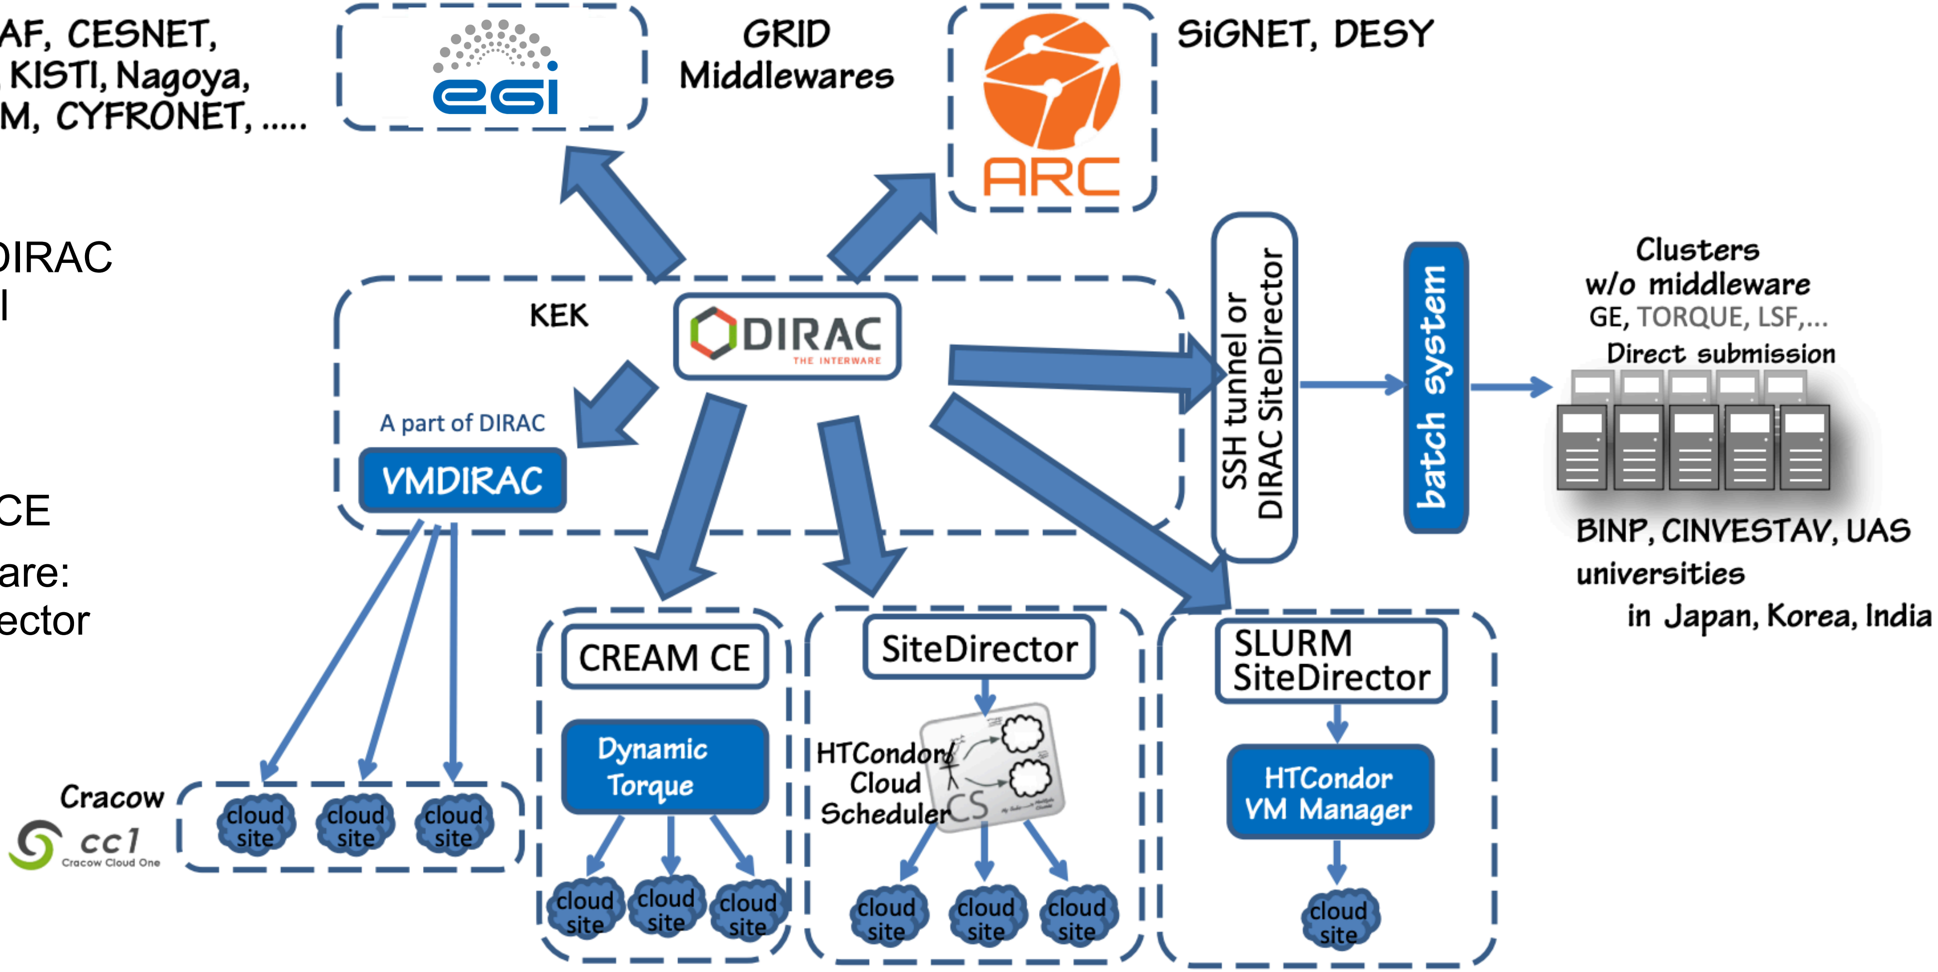

Rucio File Catalog

- Files are stored around the world in the different storage elements.
- A logical file name (LFN) is the unique identifier of a file in the Belle II grid in the form of a unix-like file path:

  `/belle/data_type/some_more_directories/file_name`

- **Rucio File catalog** resolves the LFN, and provides the information of where to find the file(s).
- Then, you only need to provide the LFN(s) of the datasets which are relevant for your analysis.
  - Bonus: integration with Rucio allow the definition of collections (a single reference for a group of datasets of interest).

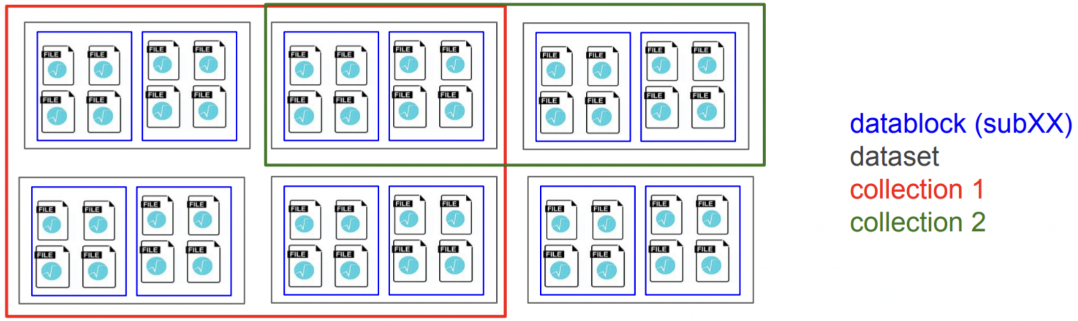

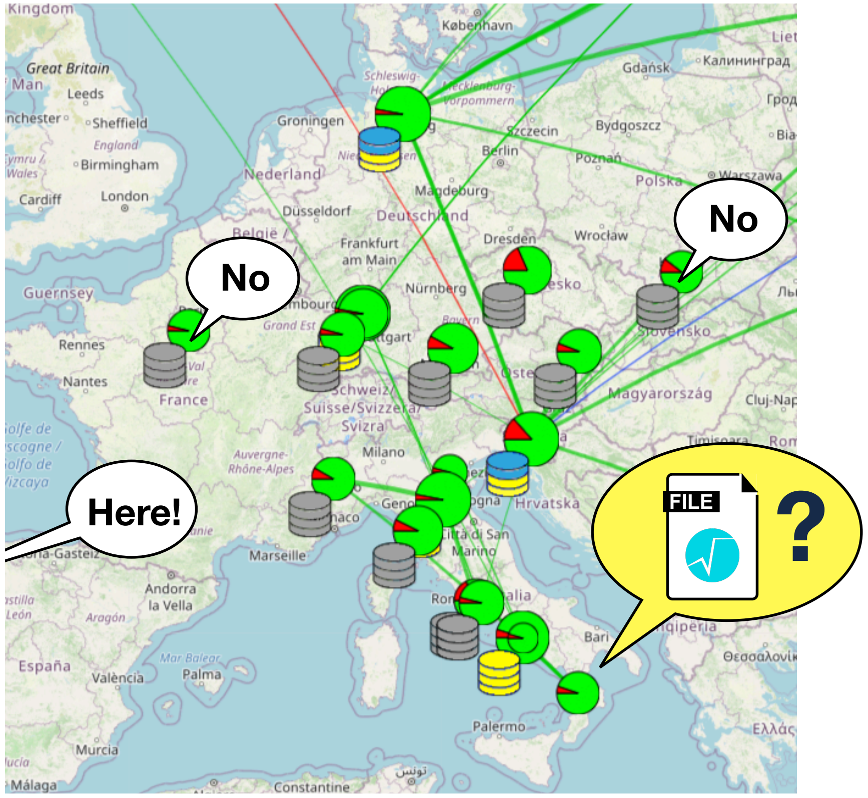

Rucio as Data Management System

- Rucio takes care of interacting with the File Transfers Service (FTS) to move around files, and takes care of the physical deletion of the file replicas.
- Features available:
  - Automatic replication of data according to policy (for MC, data, raw).
  - Automated deletion of intermediate/outdated files.
  - New monitoring tools for transfers and space occupancy.
  - For analysts: replica management, asynchronous deletion.
- We are continuously expanding the functionality.

Data movement between storage sites managed by Rucio

Monitoring System

- A full monitoring system has been developed to maximise the operation efficiency.
- Every certain time, jobs are submitted to each site to check: CPU, memory size, OS, installed rpm packages, cvmfs status, free disk space. It also performs a test execution for Belle II software, upload/downolad of files, etc.
- Similar checks performed for Storage Elements, such as free space and availability.
- Human-readable operation summary provided.

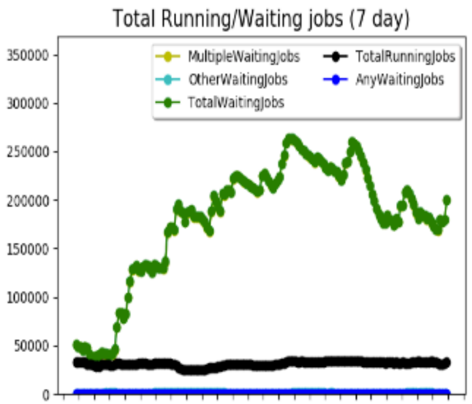

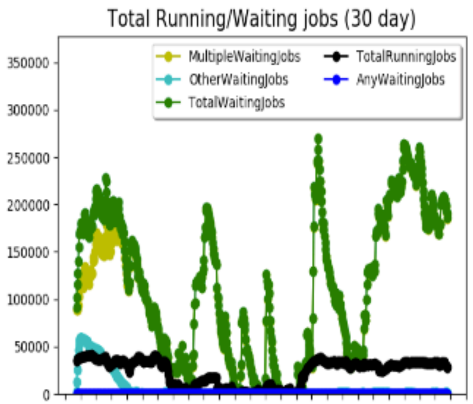

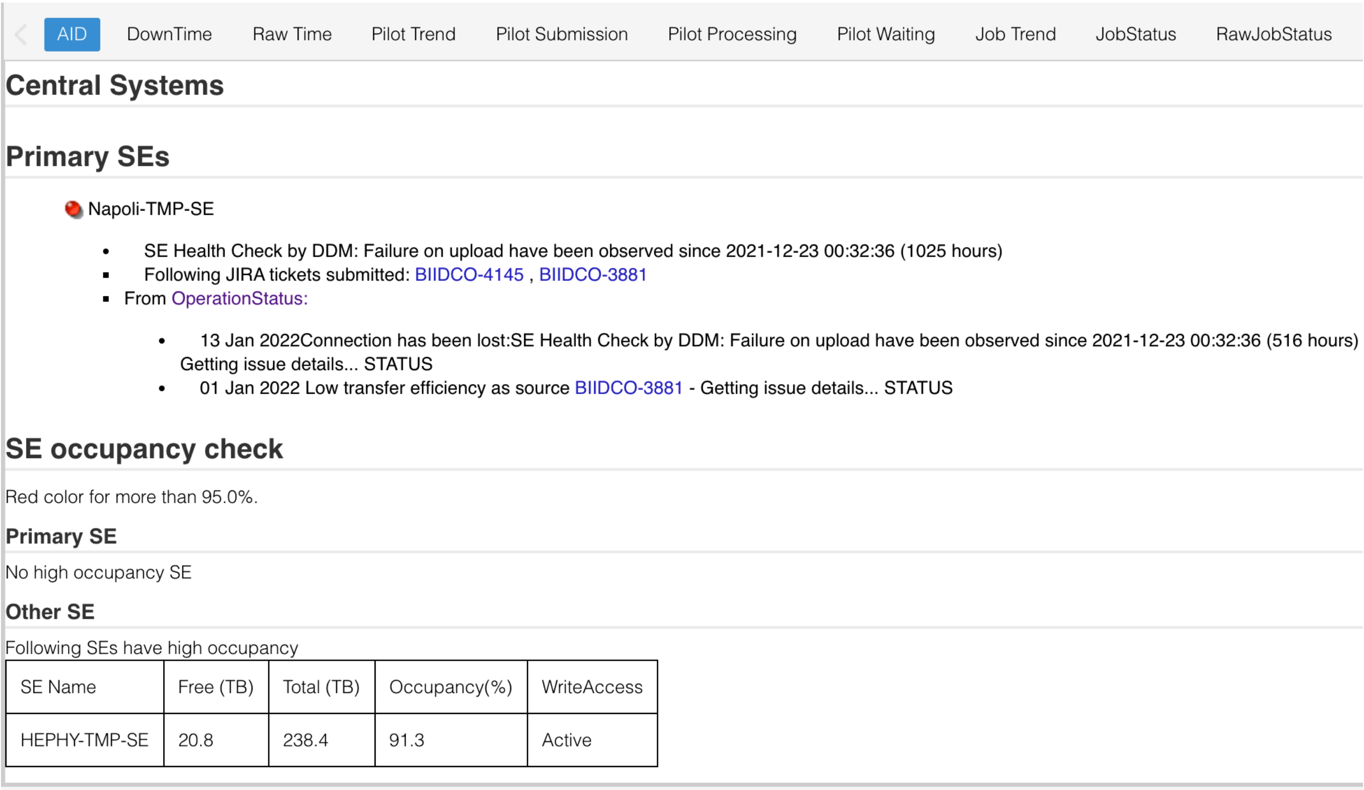
AID DownTime Raw Time Pilot Trend Pilot Submission Pilot Processing Pilot Waiting Job Trend JobStatus RawJobStatus

Central Systems

Primary SEs

Napoli-TMP-SE

- SE Health Check by DDM: Failure on upload have been observed since 2021-12-23 00:32:36 (1025 hours)
- Following JIRA tickets submitted: BIIDCO-4145 , BIIDCO-3881
- From OperationStatus:
  - 13 Jan 2022Connection has been lost:SE Health Check by DDM: Failure on upload have been observed since 2021-12-23 00:32:36 (516 hours) Getting issue details... STATUS
  - 01 Jan 2022 Low transfer efficiency as source BIIDCO-3881 - Getting issue details... STATUS

SE occupancy check

Red color for more than 95.0%.

Primary SE

No high occupancy SE

Other SE

Following SEs have high occupancy

SE Name	Free (TB)	Total (TB)	Occupancy(%)	WriteAccess
HEPHY-TMP-SE	20.8	238.4	91.3	Active

Gbasf2

The distributed analysis client for Belle II

- gbasf2 (grid + basf2) is a command-line tool for users intended to submit grid-based jobs.
- **The same Python steering files used with Basf2, work with gbasf2 on the grid.**
  - Users (analysts) test their code on local resources, then submit the job with same steering file.
- Users perform operations such as **monitor jobs, manage replicas** and **download the output** through a set of command-line tools provided within the gbasf2 environment.

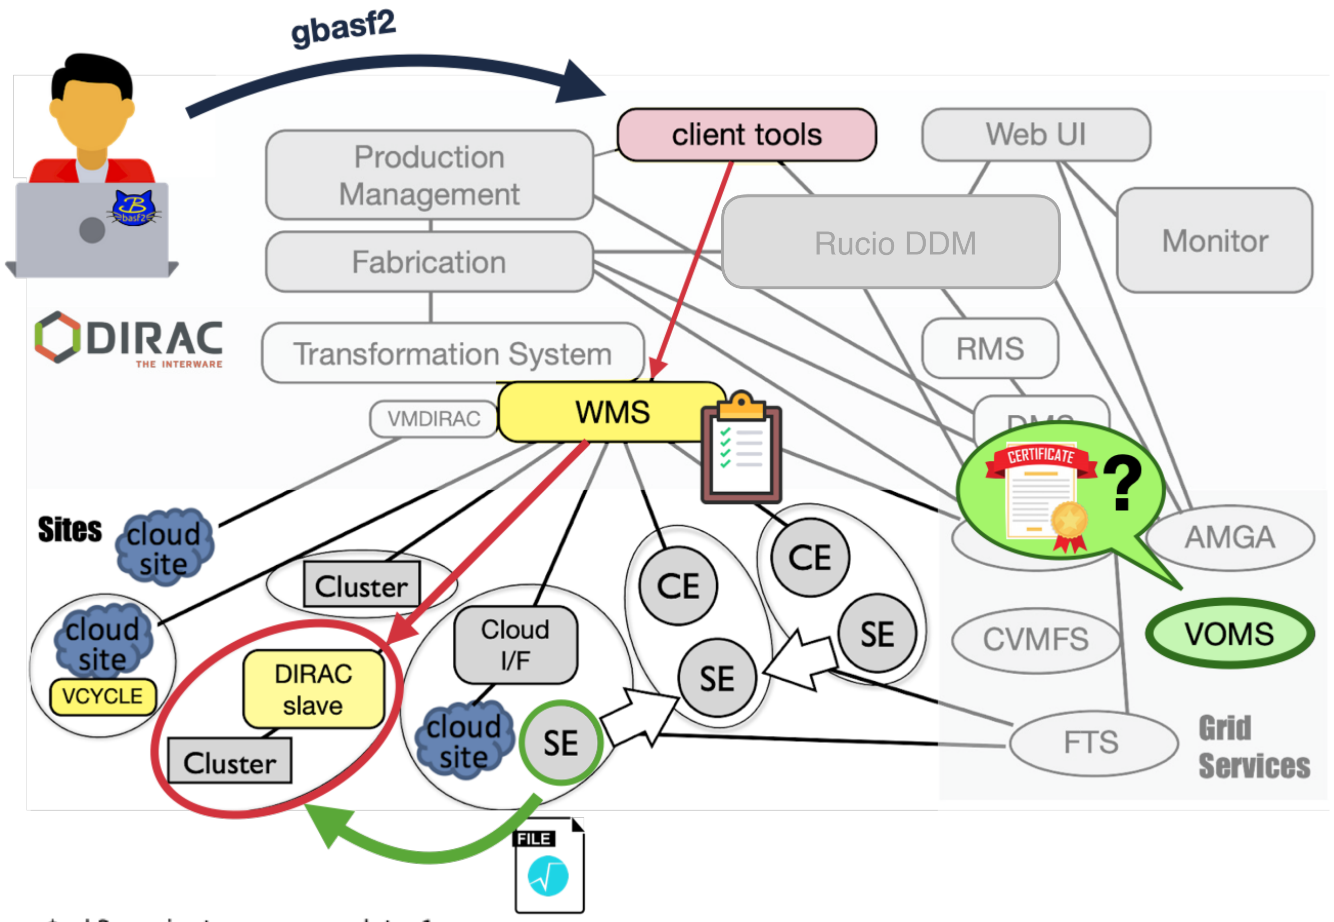


```
~ $ gb2_project_summary --date 1w
      Project               Owner    Status   Done   Fail   Run   Wait   Submission Time(UTC)   Duration
==========================================================================================================
gb2Tutorial_Bd2JpsiKs      michmx   Good     5      0      0     0      2020-07-07 08:41:40    00:18:04
BdJpsiKs_proc11_exp10      michmx   Good     874    0      0     0      2020-07-07 09:29:07    02:24:27
gb2Tutorial_B02JpsiKs      michmx   Good     5      0      0     0      2020-07-07 21:53:12    02:49:34
gb2TutorialProc11Exp10     michmx   Bad      95     779    0     0      2020-07-07 22:32:23    00:34:38
```

Operations on the grid

During 2022

- Size of datasets (w/replicas)
  - **Raw data**
    - 6 PB (4M files)
  - **Moriond 2022 / ICHEP**
    - Proc12: 125 TB (92K files)
    - MC14: 265 TB (716K files)
  - **In progress**
    - Proc13: 192 TB (187K files)
    - MC15: 541 TB (323K files)

- CPU usage
  - MC production: 72%
  - Data processing: 6%
  - Skimming: 1%
  - **User analysis 22%.**

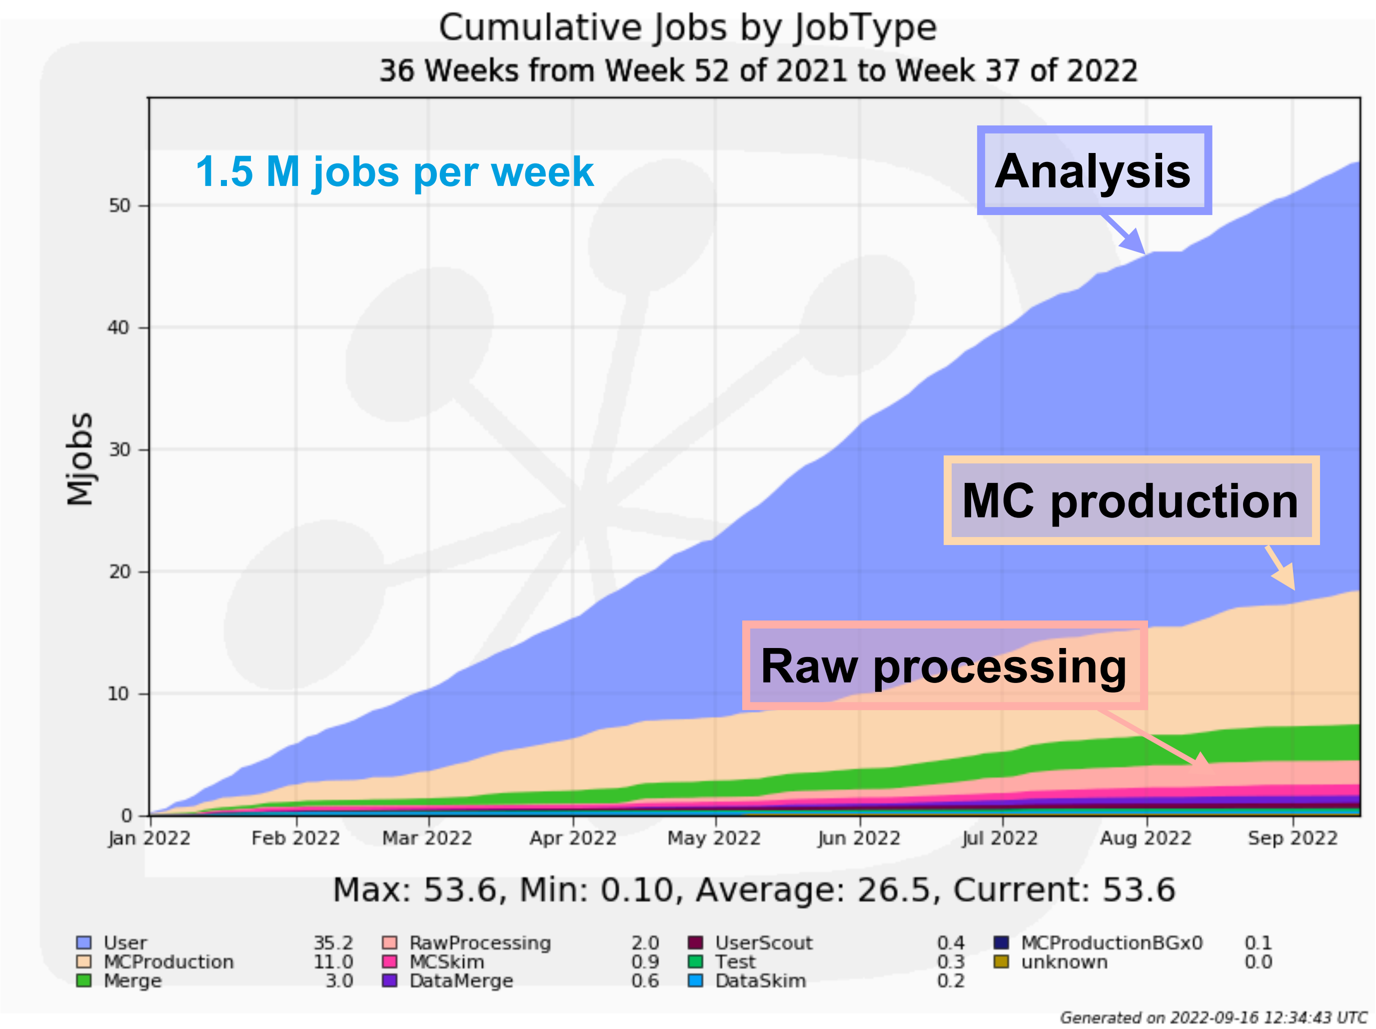

Analysis on the grid

Issues identified

- Analysis with non-skimmed data put a heavy load on the grid services.
  - We will discuss during B2GM how to improve the situation.
- Advanced usage, like training of BDTs, reach limits in sandbox size.
  - Working in a solution for distributing large files.
- Estimated CPU time is not properly set.
  - Default value does not properly represent the complexity of the steering file.
  - We are working in extracting the value from scout jobs.
  - For now, we are advertising to [set it by hand.]

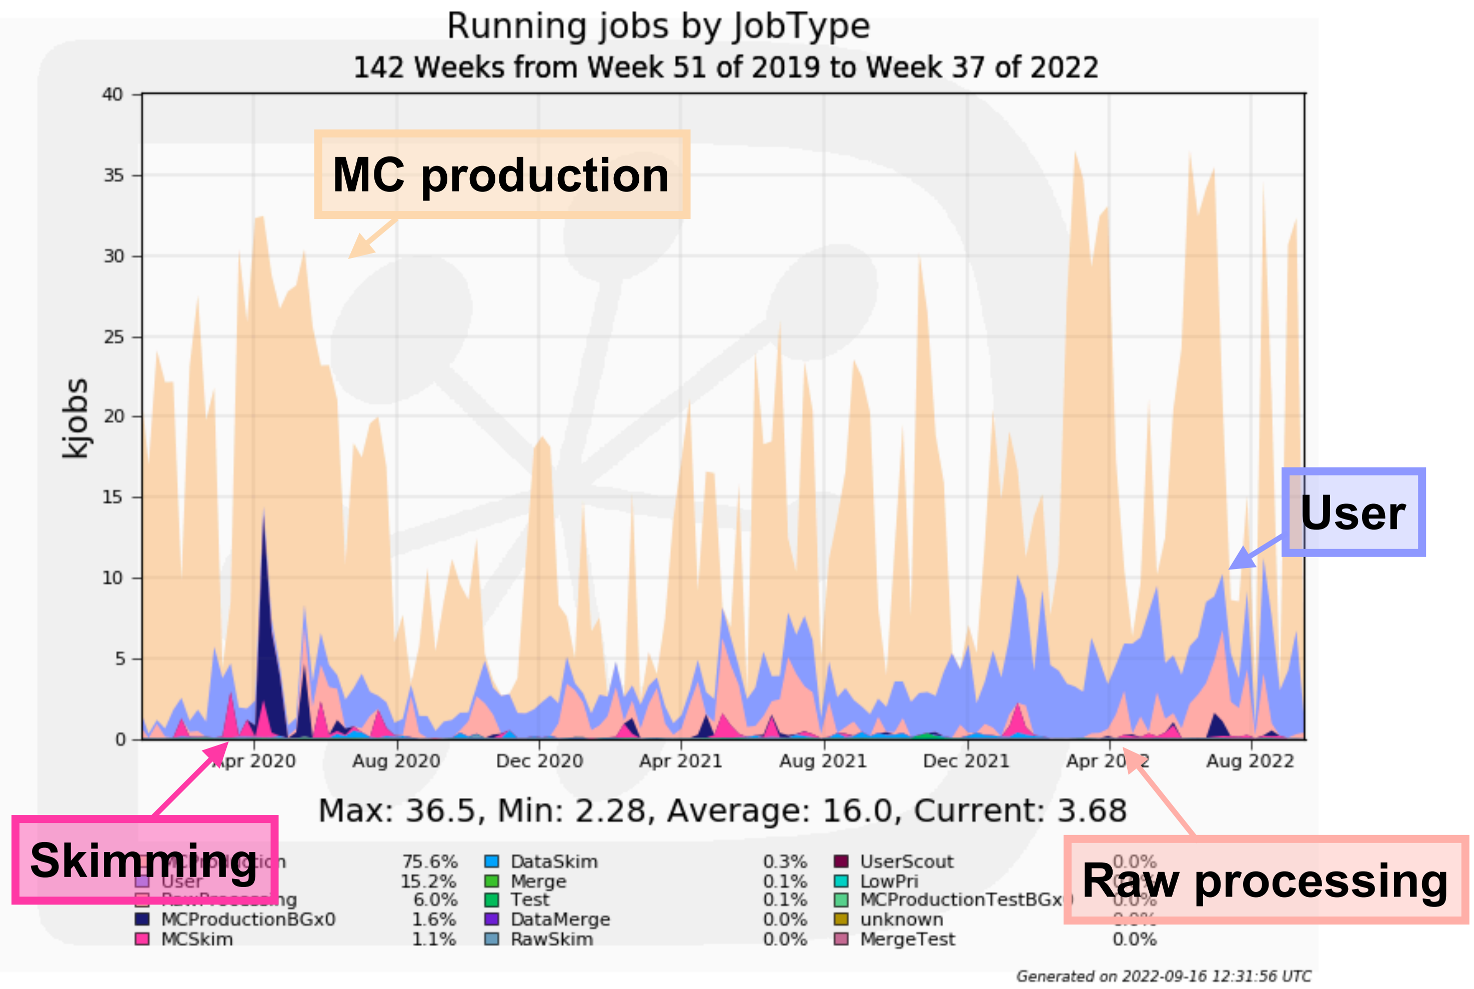

Development

BIIDCD on JIRA

- Significant improvements in documentation. End-users manual built in Sphinx: gbasf2.belle2.org
- Rucio features integrated in our tools/operations.
  - Async operations. Dataset collections integrated in job submission/client tools.

- **Priorities for LS1:**
  - Automatization of raw data staging/unstaging during campaigns.
    - Staging from disk to tape for reprocessing has a significant impact on free space available.
  - BelleDIRAC & gbasf2 APIs
    - Relevant for automatic submission of data/MC production, validation, systematics framework jobs.
  - Token based authentication (replacing x509 certificates).
    - We must ensure that all our grid systems/services are prepared for the switch.
  - Improvements in the Production system
    - Multicore jobs.
    - Removal of run-boundaries per production unit.

We need your help

- Have you ever experience that
  - Your window has gone somewhere?
  - Your keyboard is frozen?
  - Your wifi connection is down?

- Have you seen this before?

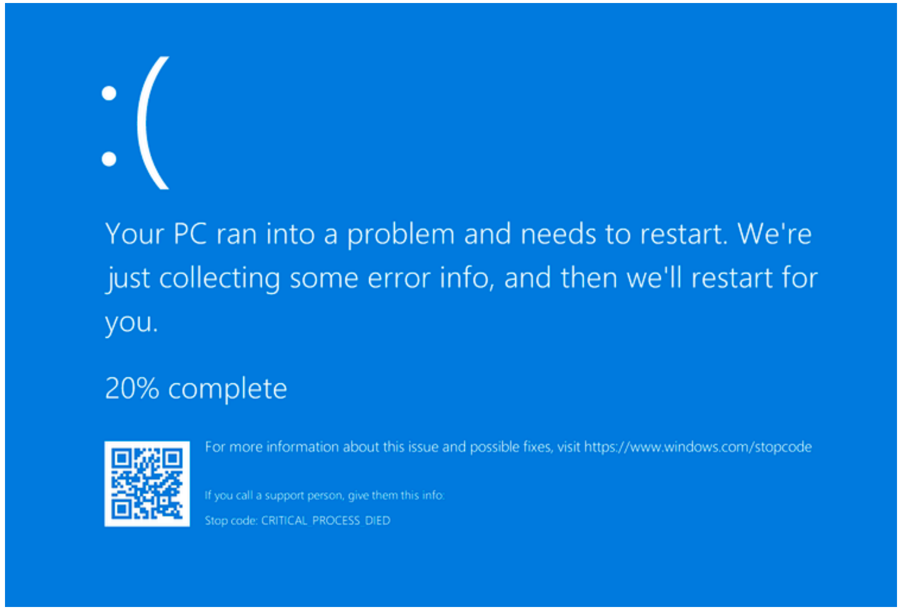

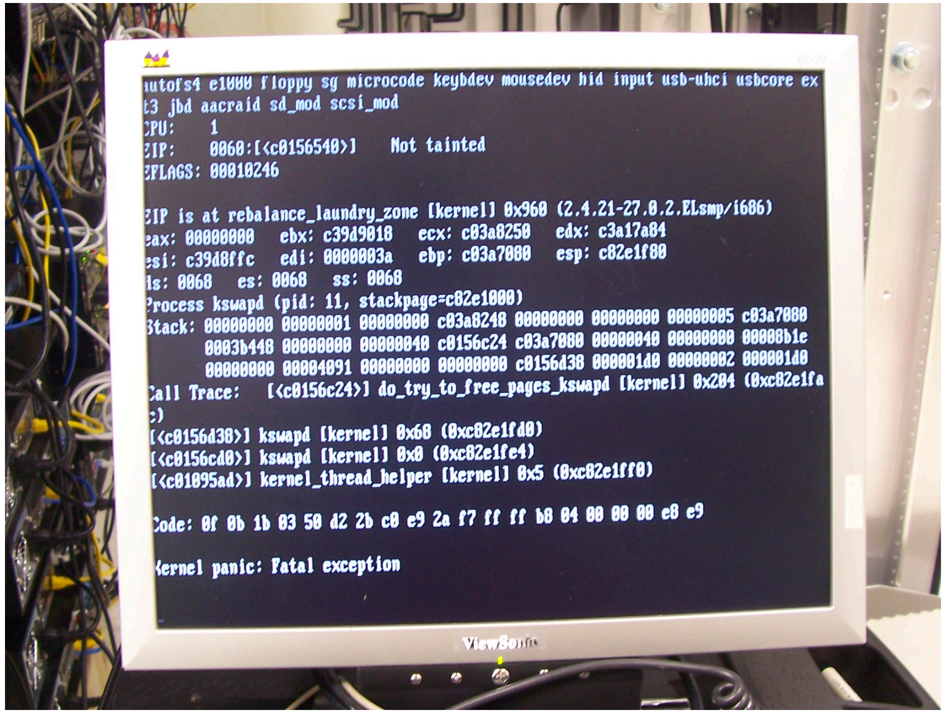

We need your help

- Computers are not so smart. Sometimes, they fail.
  - **"Sometimes" x Huge Resources = "Often"**
  - The computing system need 24 hour x 7 day care.

- Please help us as a Data Production Shifter. You can book at shift.belle2.org
  - A nice manual is already prepared.
-
- If you have some experience as data production shifter, please become an Expert Shifter.
  - The Expert Shifter training course is open.

- You will learn a lot about the computing system, and it is a very important service to the collaboration.

Summary

- Belle II is expected to produce tens of petabytes of real and simulated data per year.
- DIRAC pilots allow to integrate heterogeneous computing resources with a single interface (gbasf2).
- The Integration of the DDM system with Rucio was successfully performed last year. Automatic replication and deletion has been enabled. Multiple protocols supported.
- Computing will be a critical component to deliver results with high-precision in the coming years. Help us:
  - Taking data production shifts.
  - Becoming an Expert Shifter.
  - Looking into the open tasks. If you are interested in taking one contact us (comp-dirac-devel@belle2.org)

Thank you

Backup

Analysis Workflow

A reminder

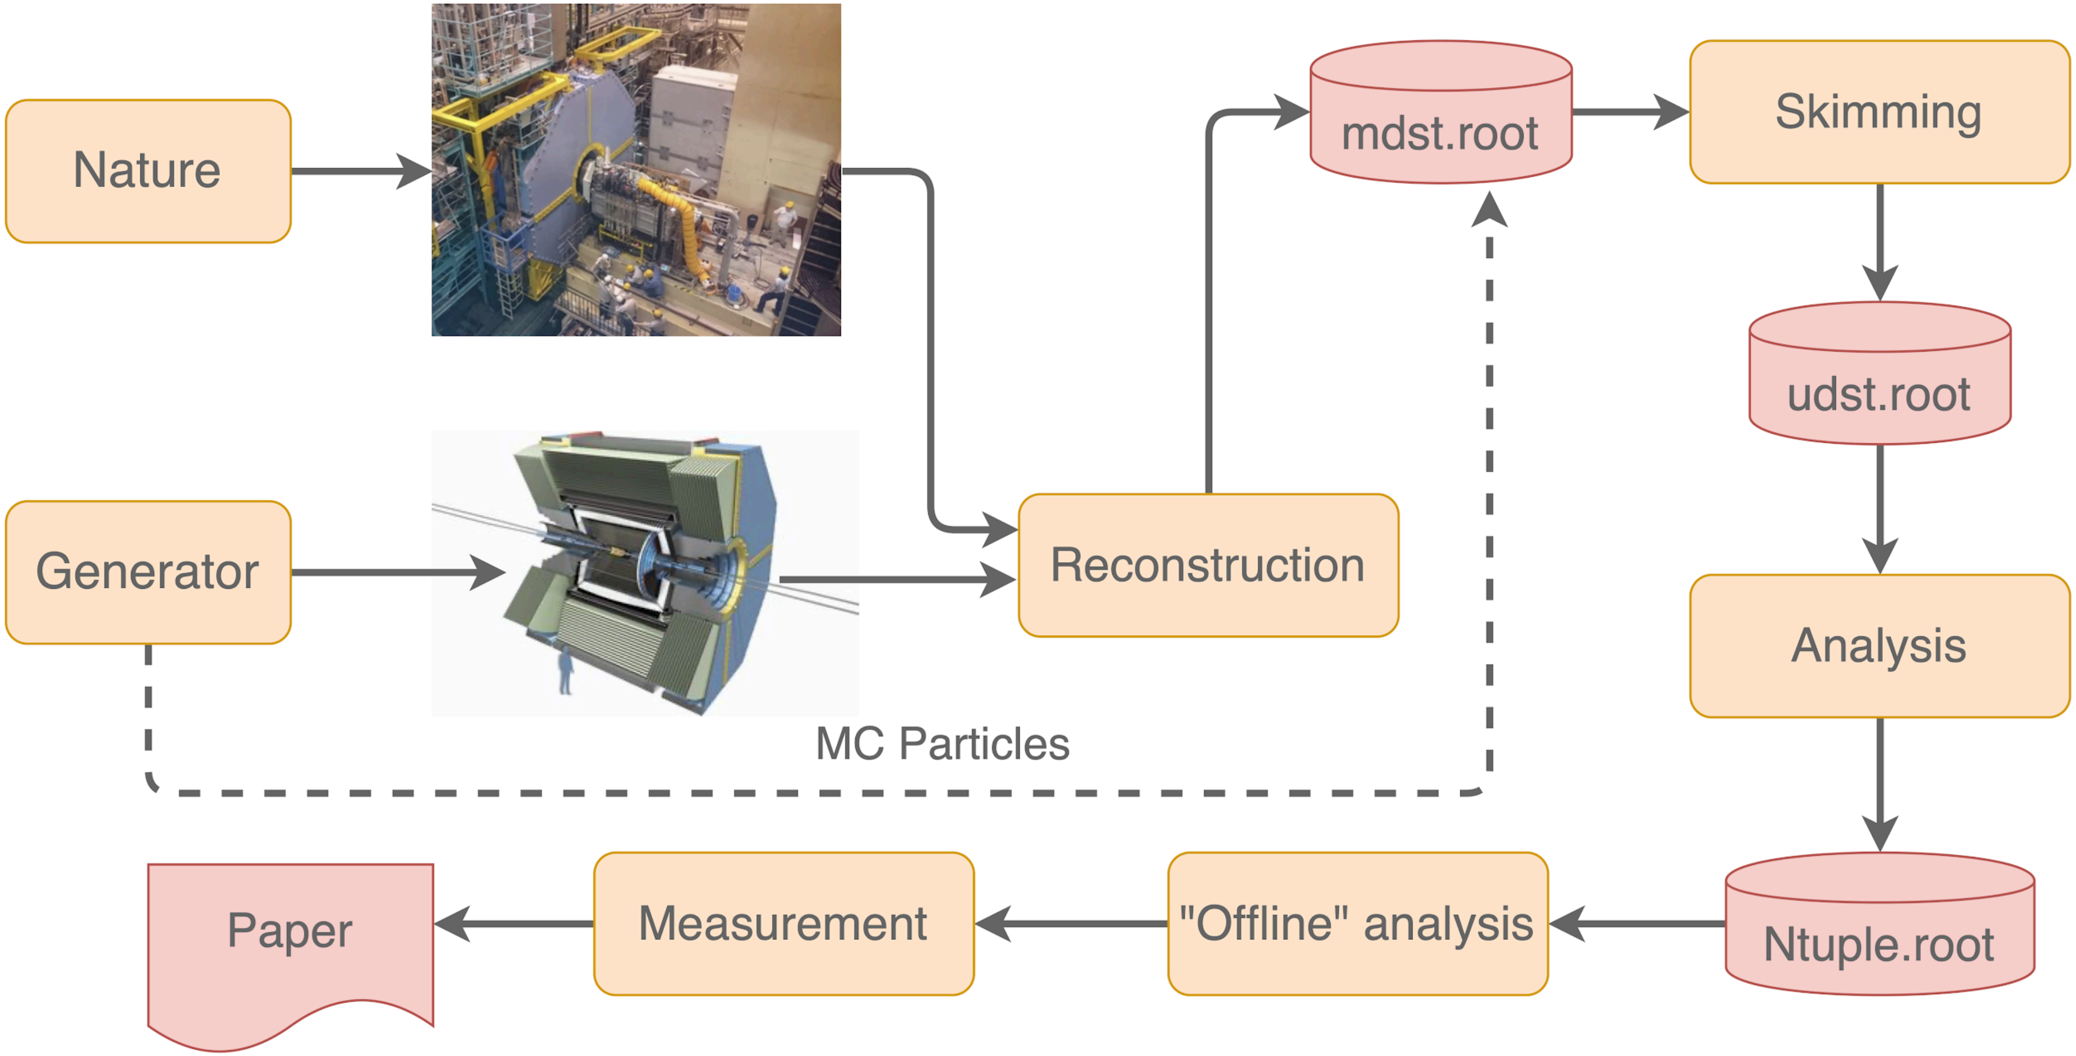

Distributed computing infrastructure at Belle II

Central services

- **Production**
  - 11 DIRAC servers + 4 DB servers + 2 Web servers (KEK)
    - SiteDirectors for SSH sites (Nagoya)
    - SiteDirector for cloud (University of Victoria); Vcycle (Napoli); TARDIS (KIT).
    - ReqProxy (KEK, Nagoya, Napoli, …)
  - Rucio server (BNL)
  - FTS servers (KEK & BNL)
  - CVMFS (KEK) for DIRAC tar-ball distribution.

- **Test servers at BNL**
  - Certification: validation of new BelleDIRAC releases.
  - Migration: test of base DIRAC upgrades.

- **Development**
  - Multiple instances at KEK, BNL, Mississippi, etc.

Data Management with Rucio

- Rucio is an advanced Distributed Data Management System initially developed for the ATLAS experiment.
- Last year, our Distributed Data Management system was successfully integrated with Rucio.
- Rucio takes care of interacting with the File Transfers Service (FTS) to move around files, and takes care of the physical deletion of the file replicas.
- It is able to use multiple protocols to interact with the Storage Element, in particular WebDAV.

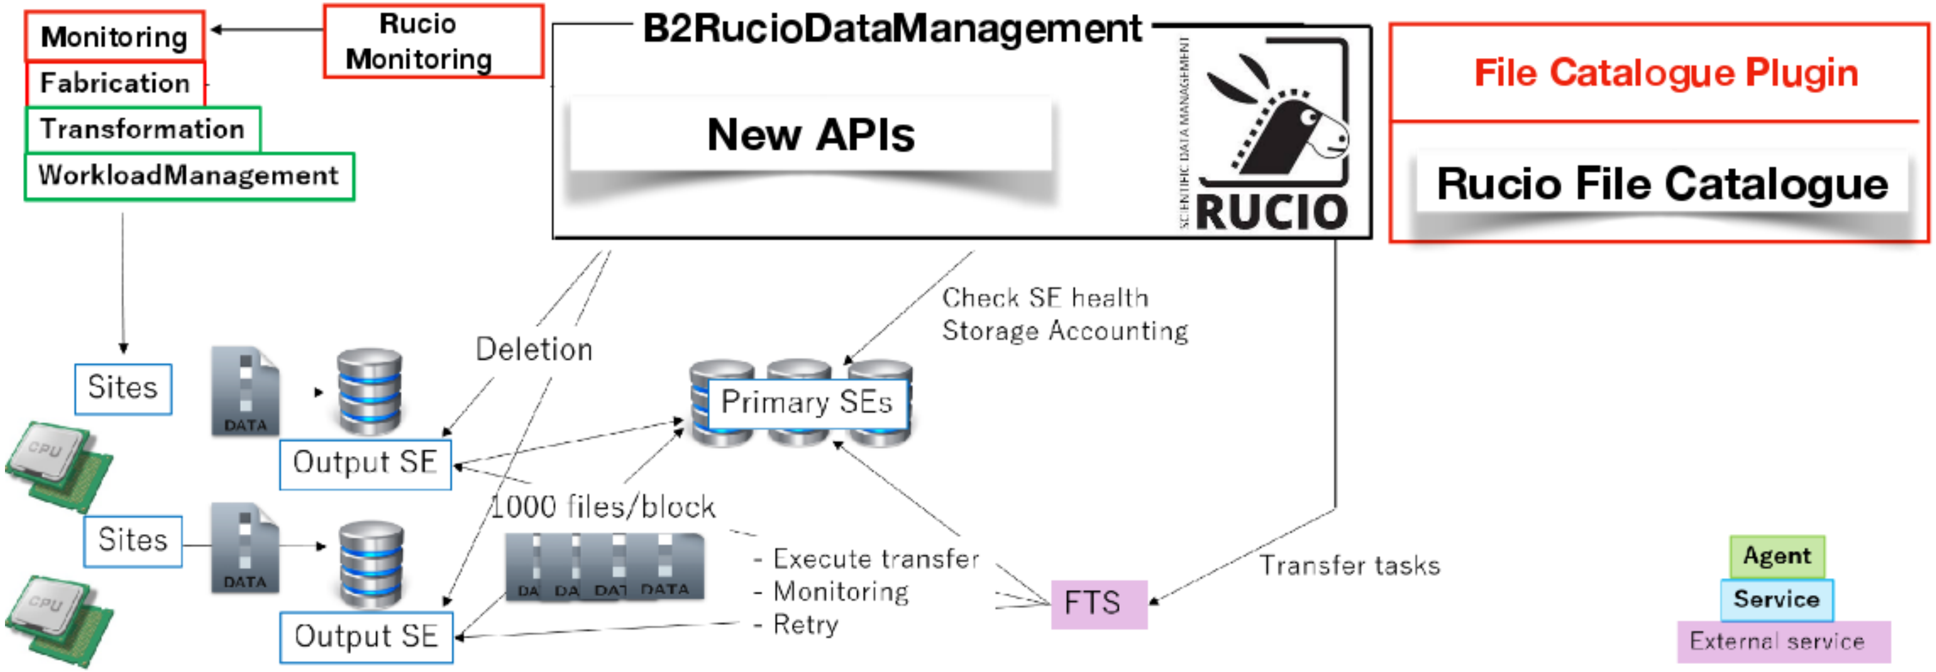

DIRAC WMS and Computing Resources

- DIRAC pilot jobs are submitted to computing resources by the WMS.
- The pilot sets the local DIRAC configuration, and launches an entity for matching jobs.
- Pilots retrieve the jobs from the task queue and steer their execution on the worker nodes, including data uploading.
  - This model allows to integrate heterogeneous computing resources.
- A pilot can run on every computing resource, e.g.: on CREAM Computing elements, on DIRAC Computing elements, on Virtual Machines in the form of contextualization script.

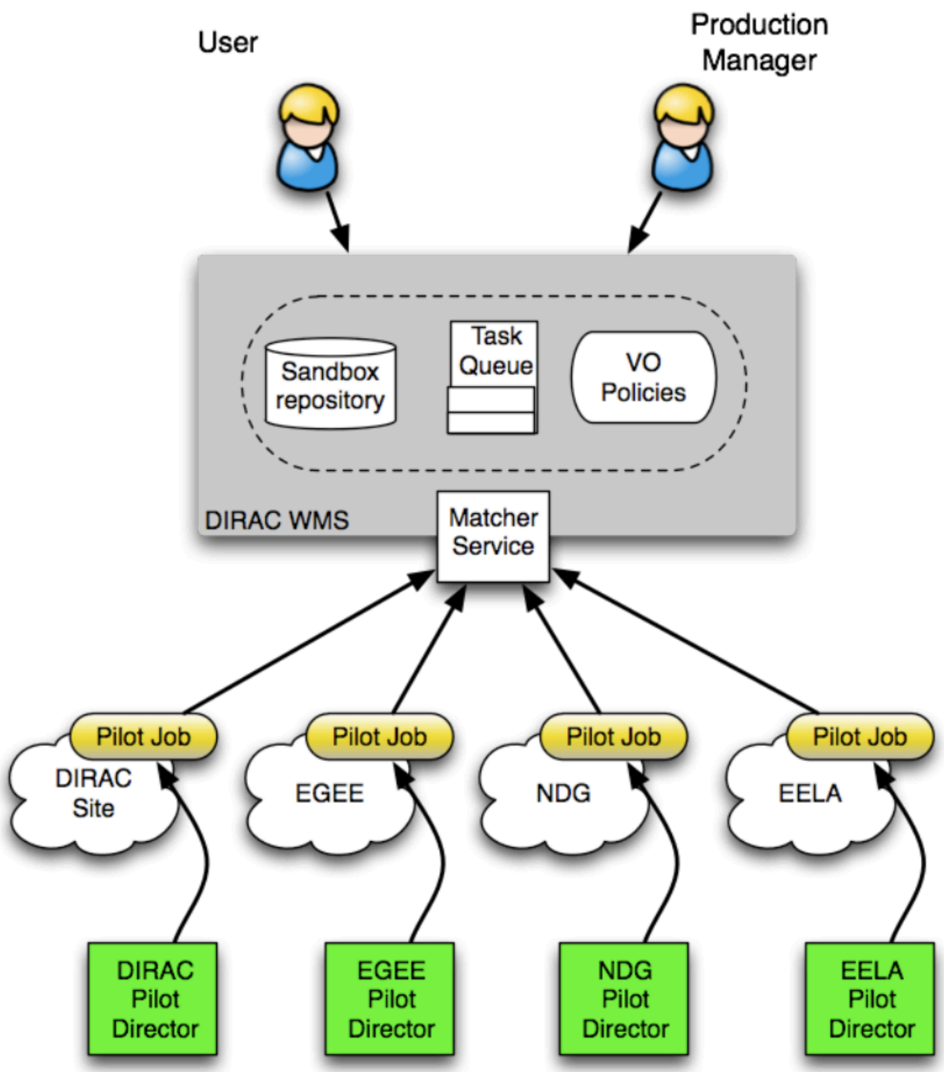

Data Formats

In general, Belle II output is stored in ROOT files containing subsets of objects.

- **RAW**: raw data containing detector information.
  - ~70 kB/event
  - Raw data set during 2019-2021 operation: 5 PB
- **cDST**: calibration Data Summary Table
  - ~120 kB/event
  - Contains objects needed for calibration. Locally produced.
- **mDST**: mini Data Summary Table
  - ~15 kB/event
  - Strictly controlled subset of objects necessary for analysis.
- **uDST**: user data summary table
  - ~20 kB/event
  - mDST objects + analysis objects (ParticleLists).
  - Produced from skims and intended for physics results.

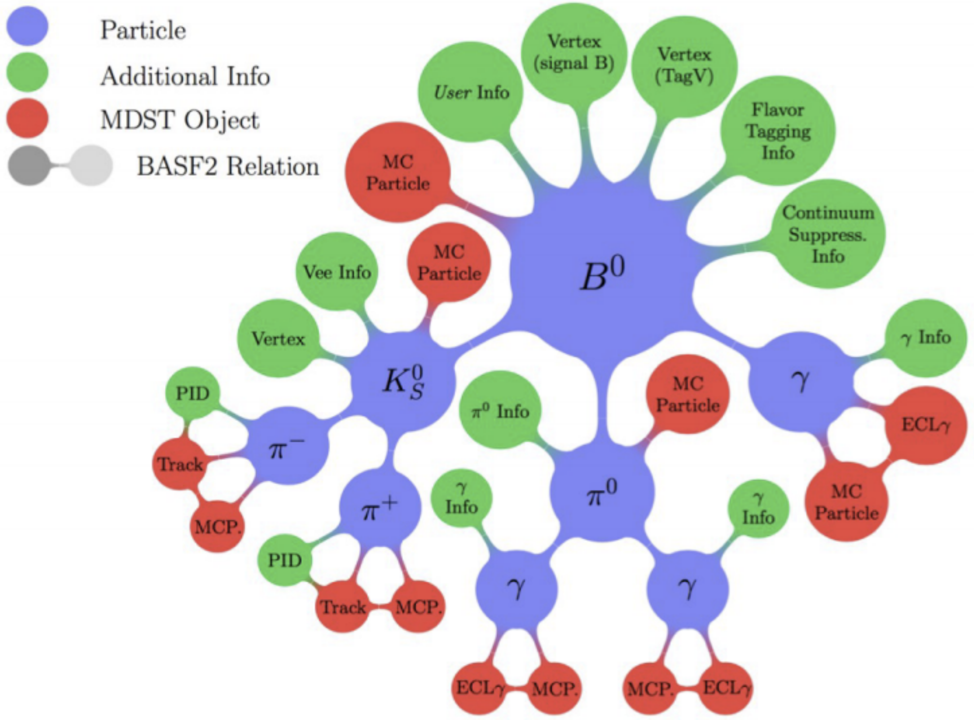

Data Processing

Scheme

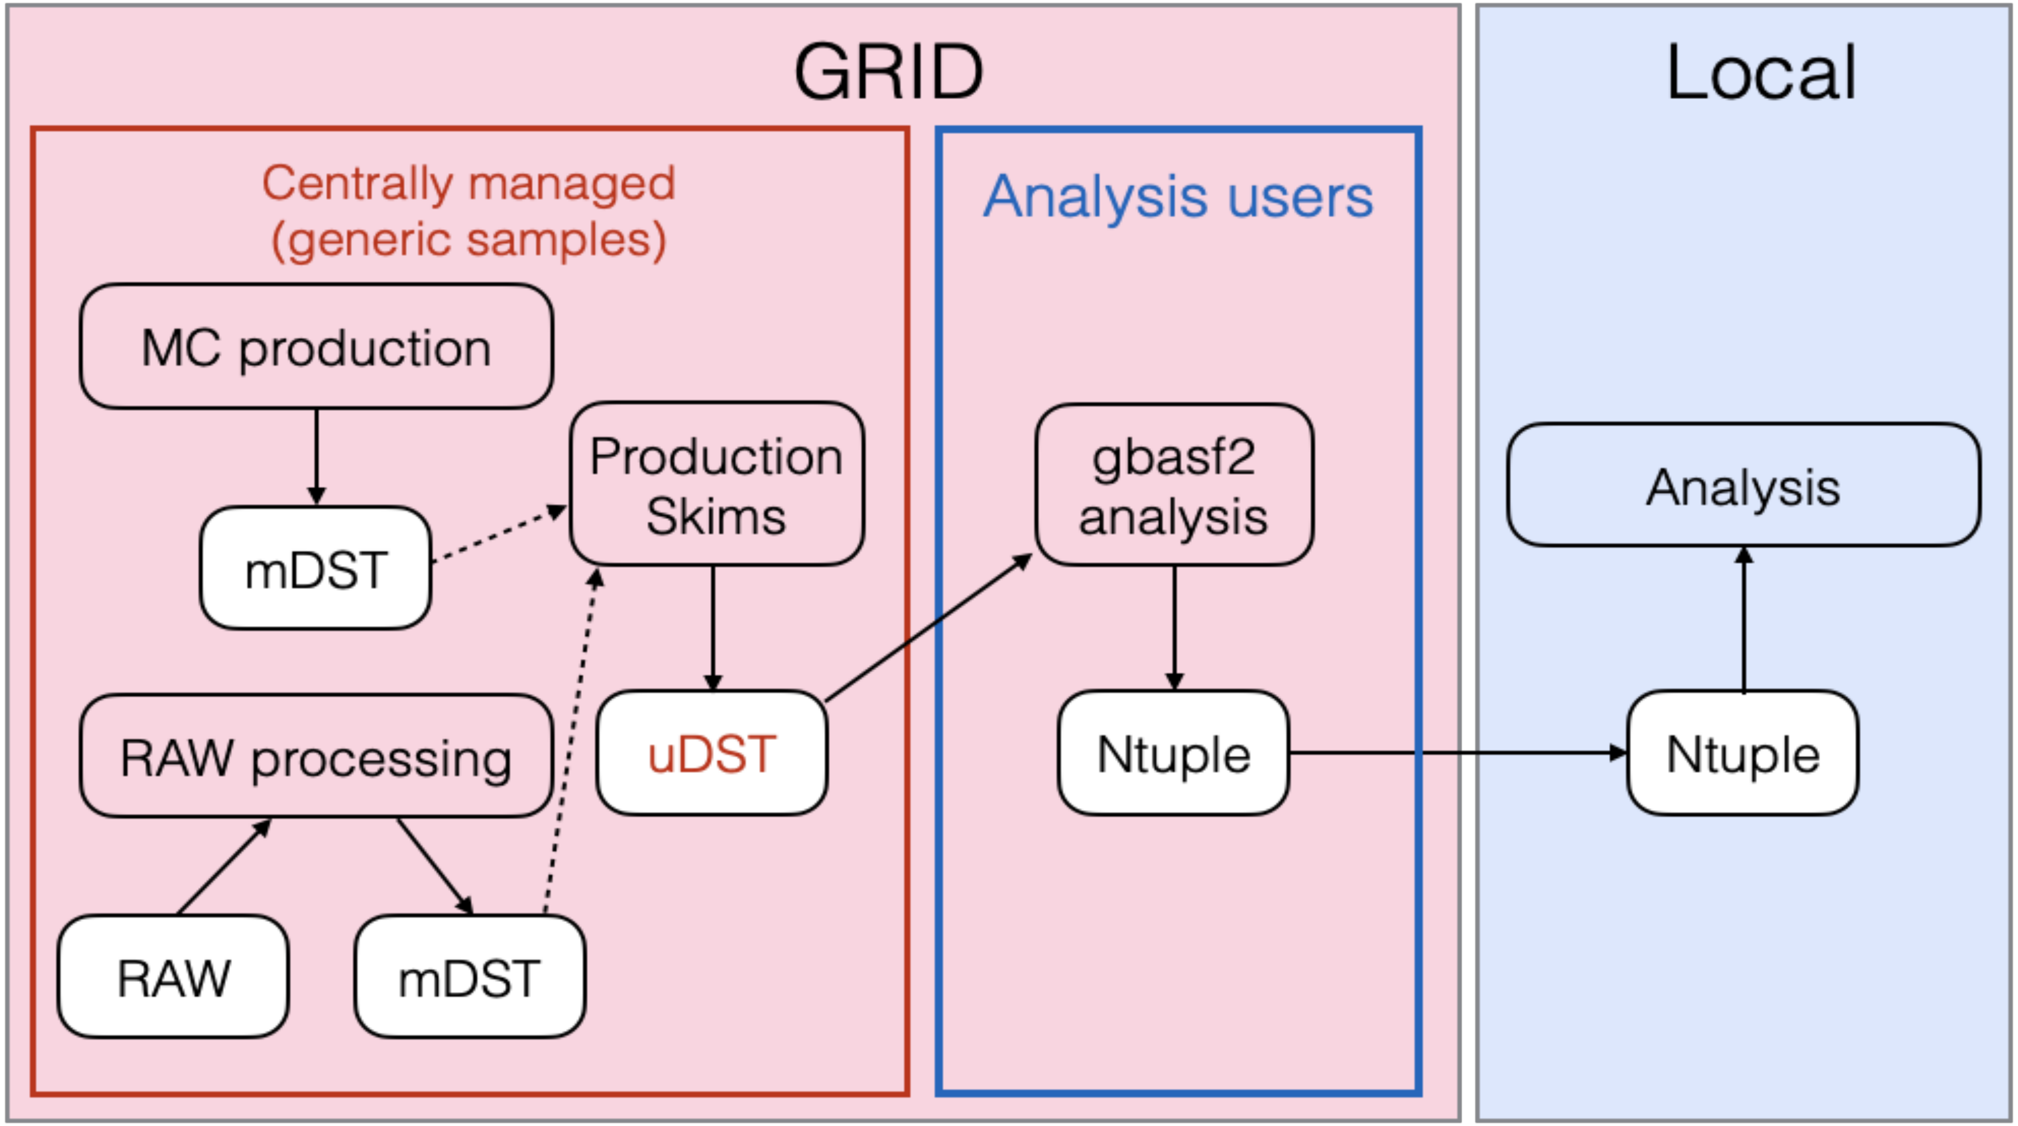

Distributed Computing

Gbasf2

- Gbasf2 uses client tools to submit jobs to the grid.
- The Virtual Organization Management System (VOMS) verifies if the user have the correct permissions.

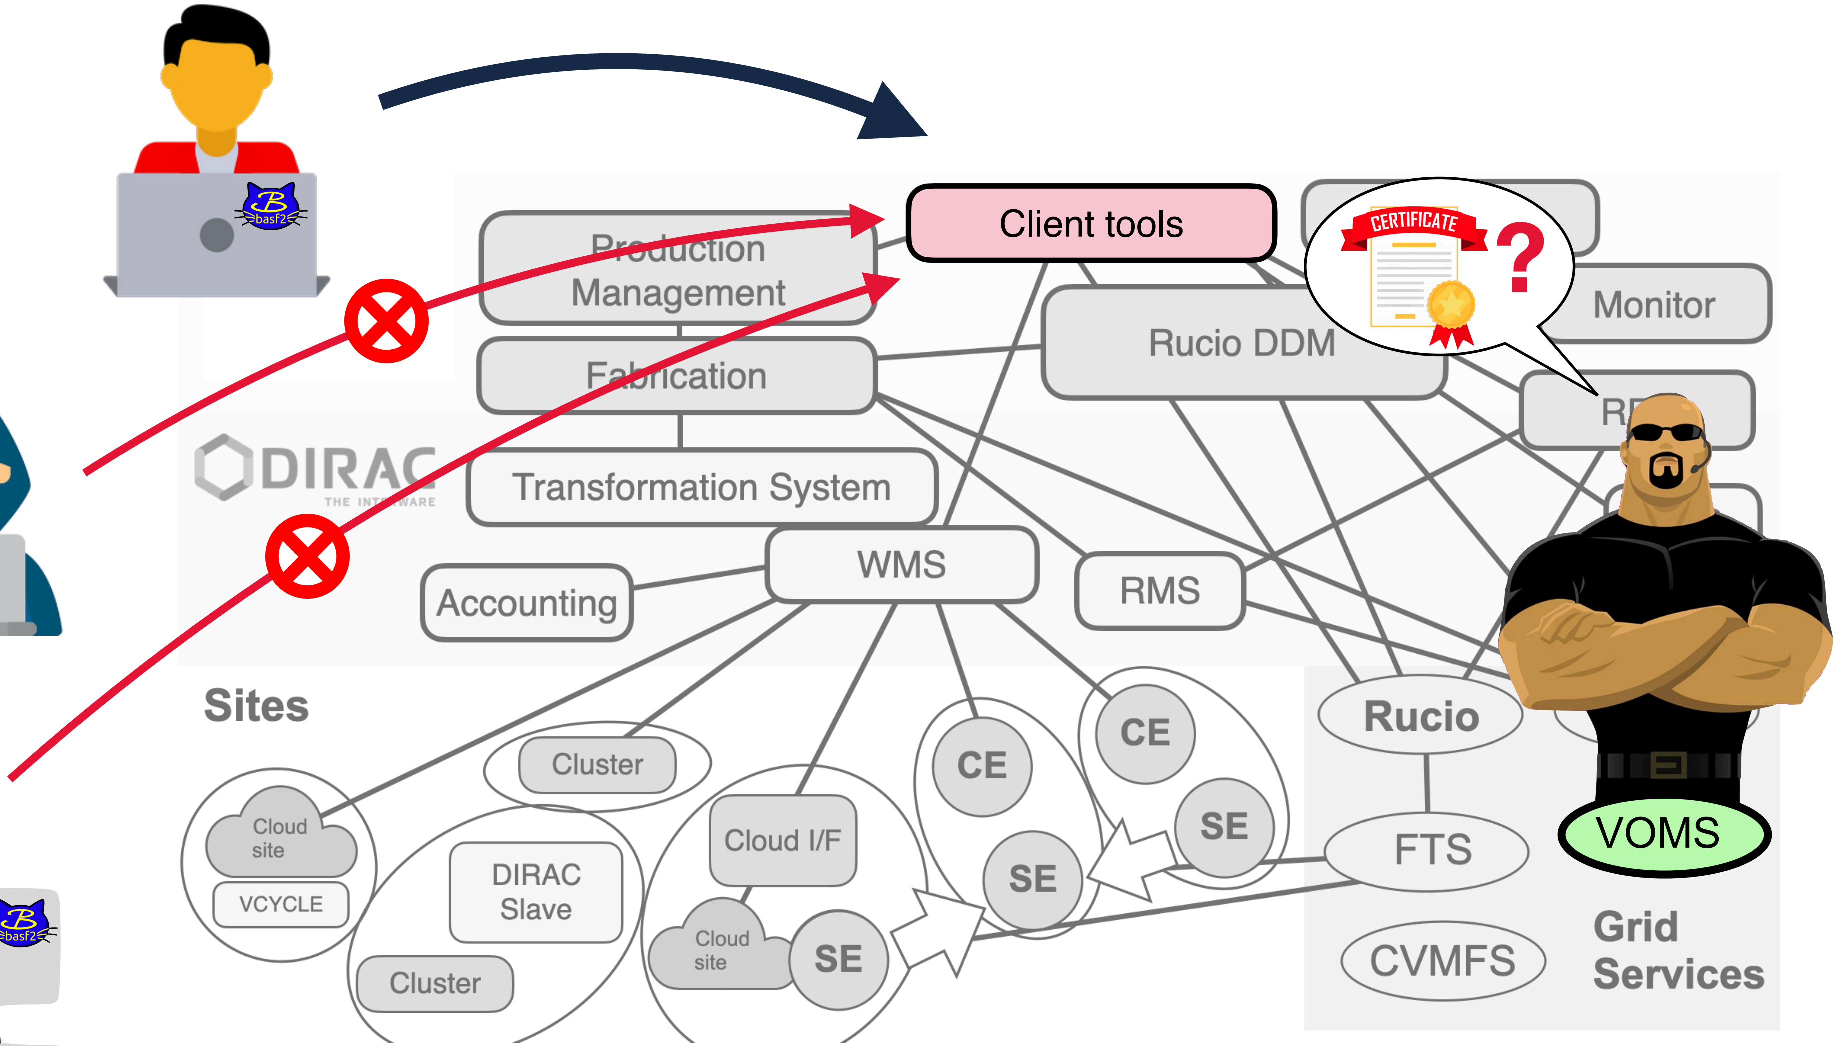

Distributed Computing

Gbasf2

- The jobs are submitted and queued into the DIRAC Workload Management System (WMS).
- Input files are distributed in Storage Elements (SE).
- **In gbasf2, input files are specified using the Logical File Name.**
- The Rucio File Catalog indicates the physical location of the file.

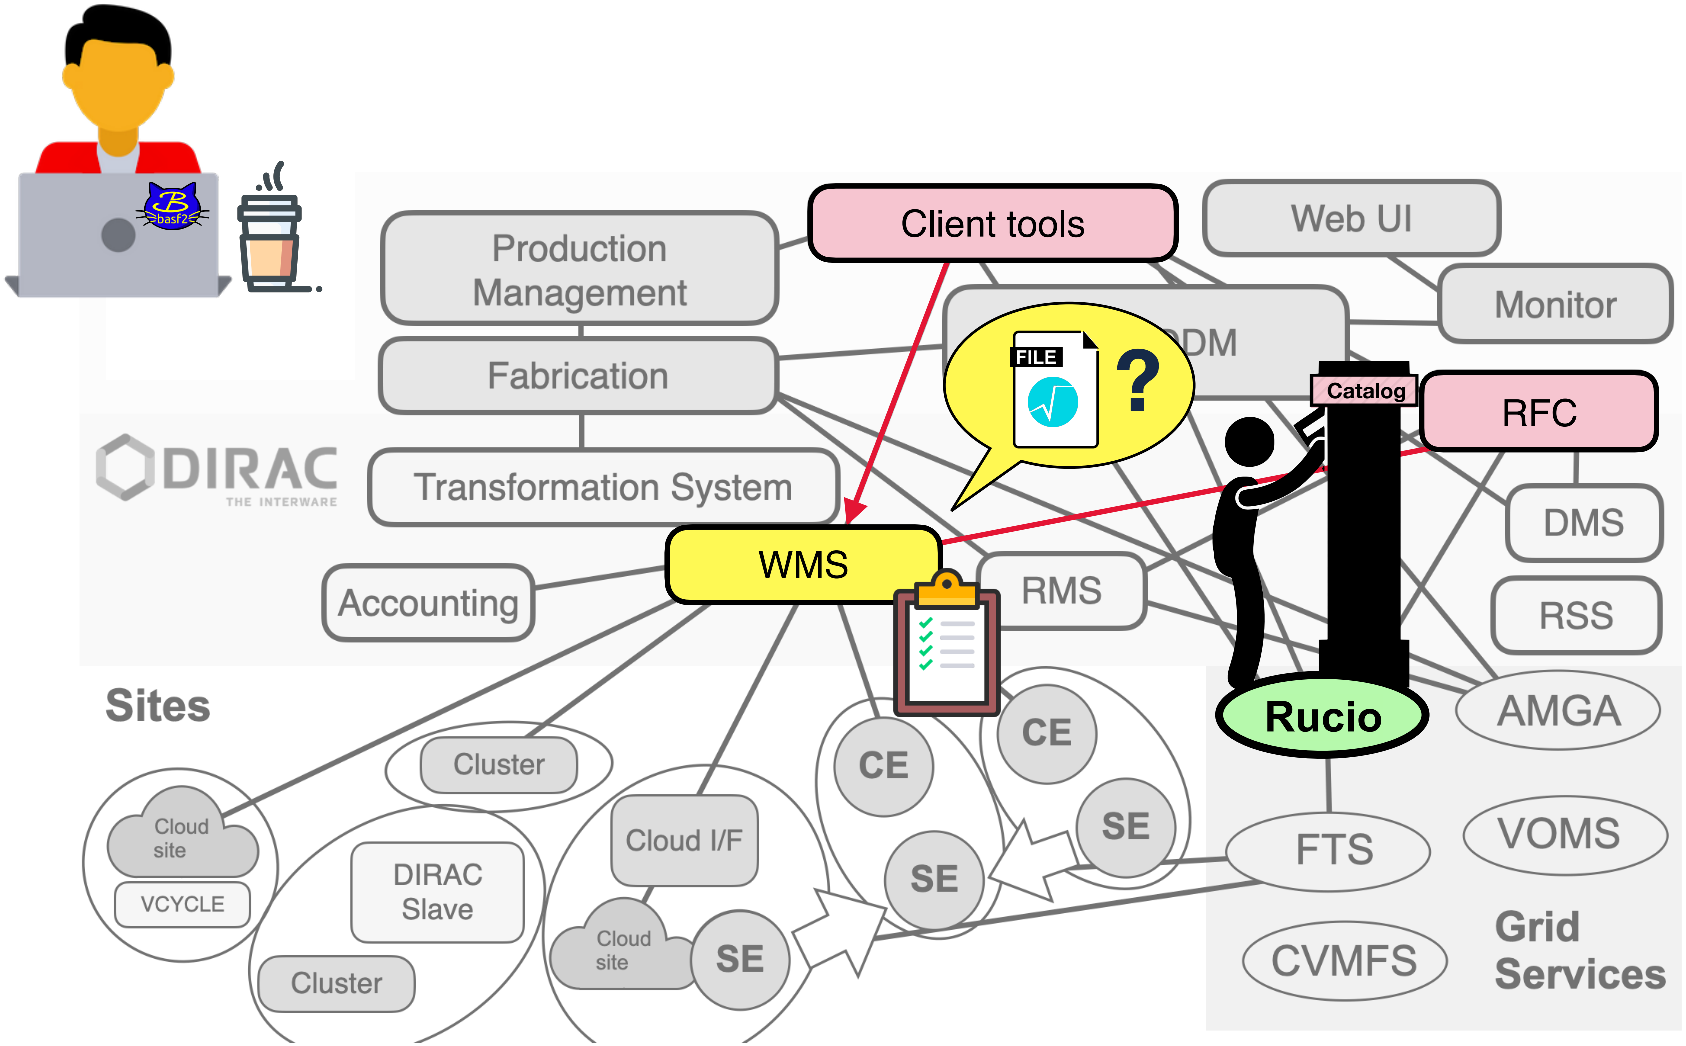

Distributed Computing

Gbasf2

- DIRAC **WMS** will take jobs from the queue and will choose a site to run them, depending of the location of the input files.
- Once the site is chosen, **the pilot** running on the site pulls the job and executes it.

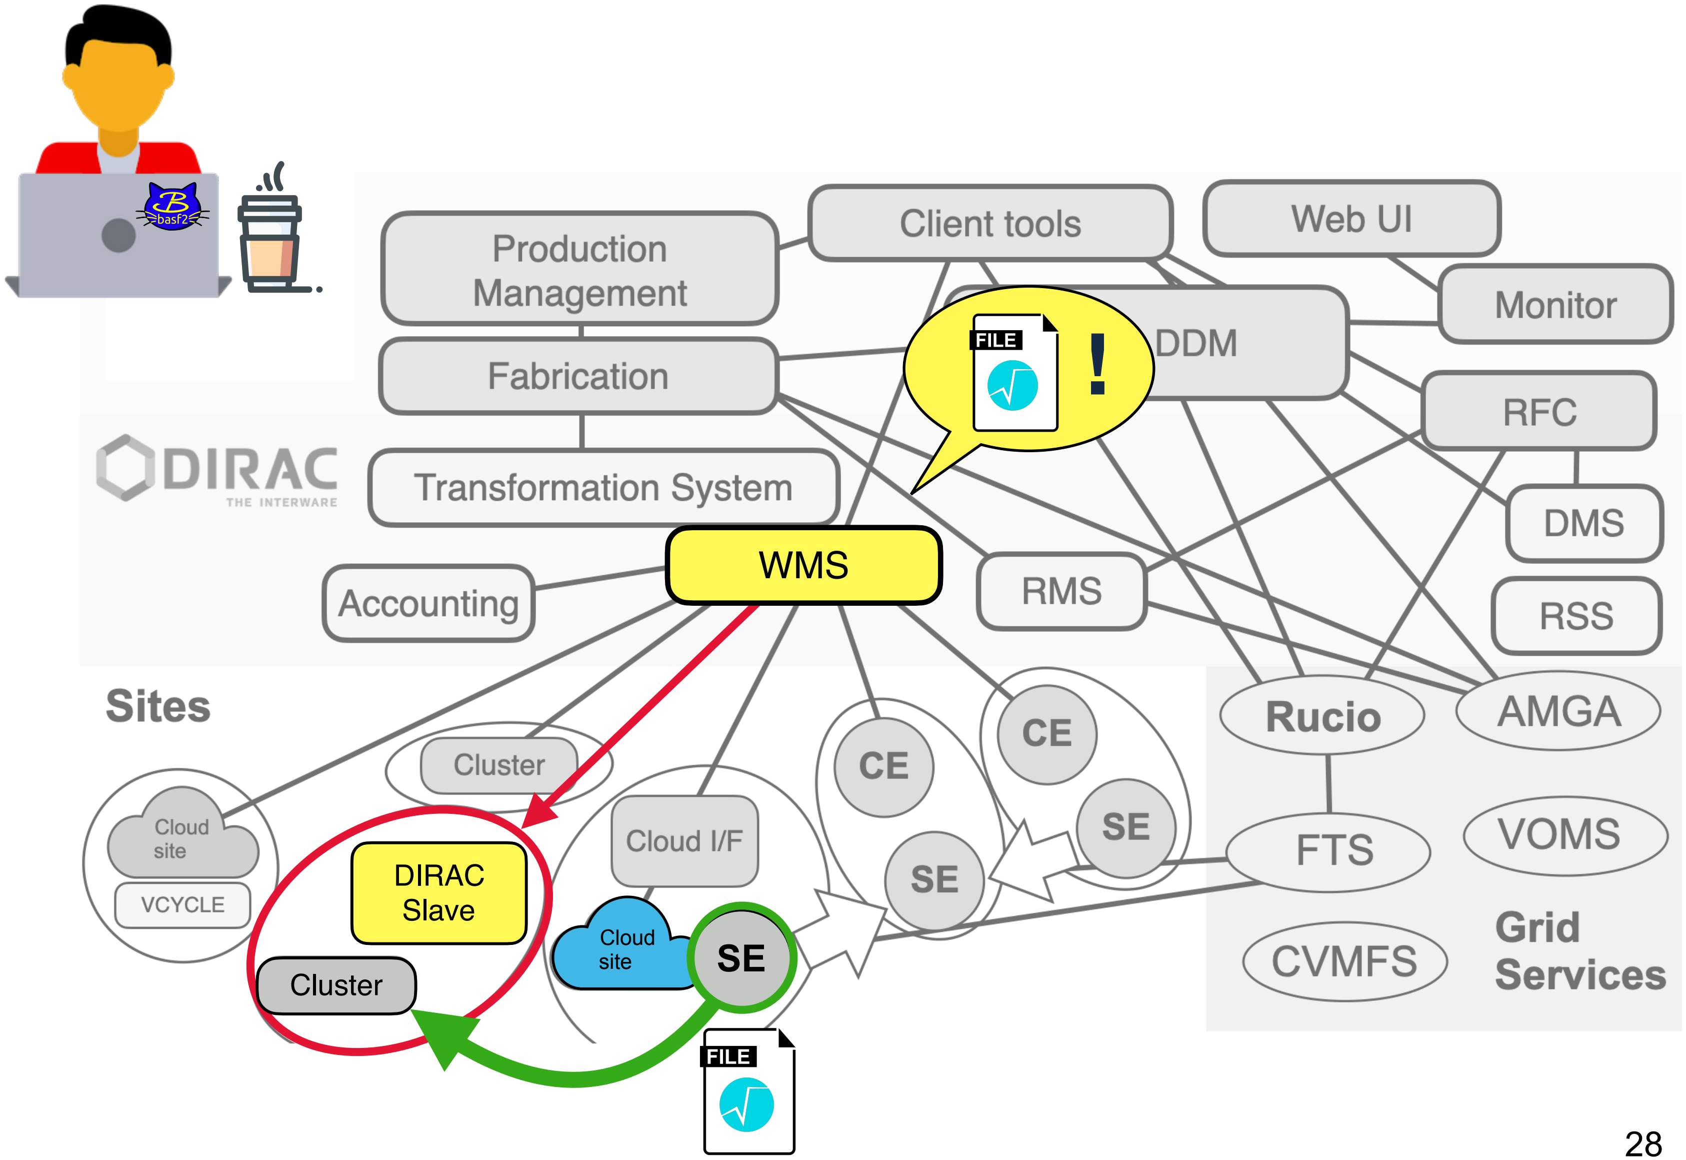

Distributed Computing

Gbasf2

- Once the job finishes, it sends the output file to a SE.
- You will be able to download the files once all the jobs in your project finish successfully.

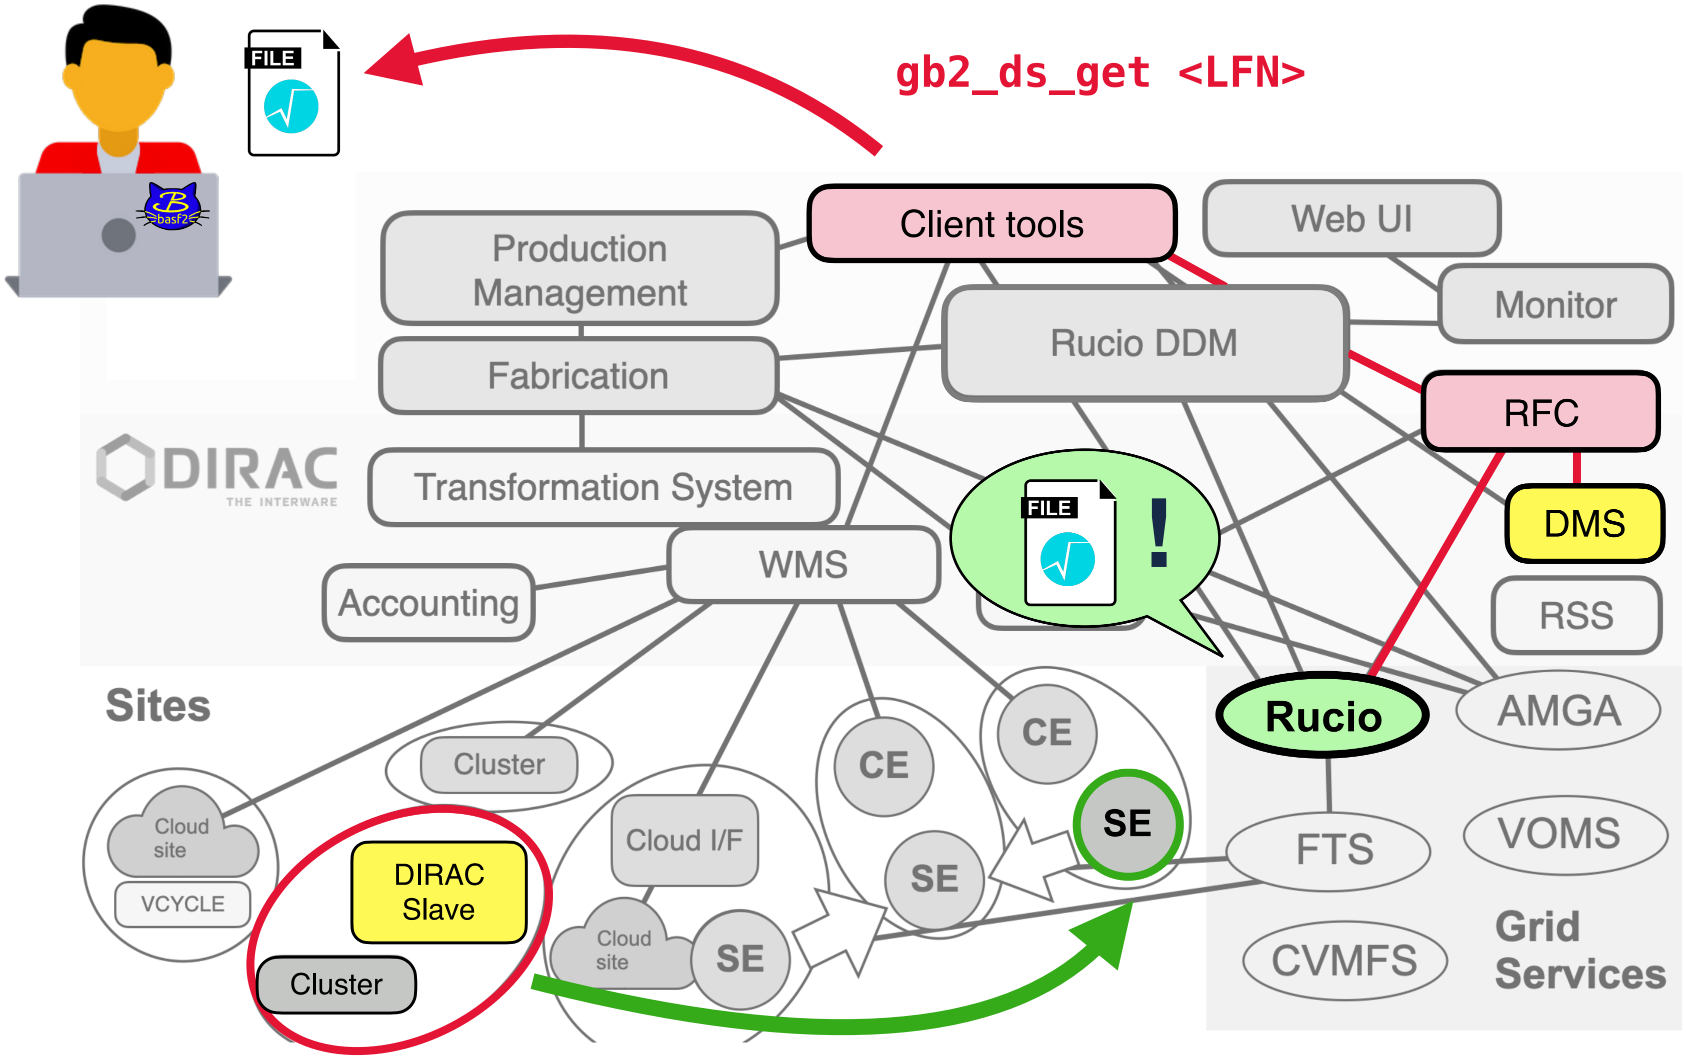

Distributed Computing

Gbasf2

- … and perform the analysis offline.

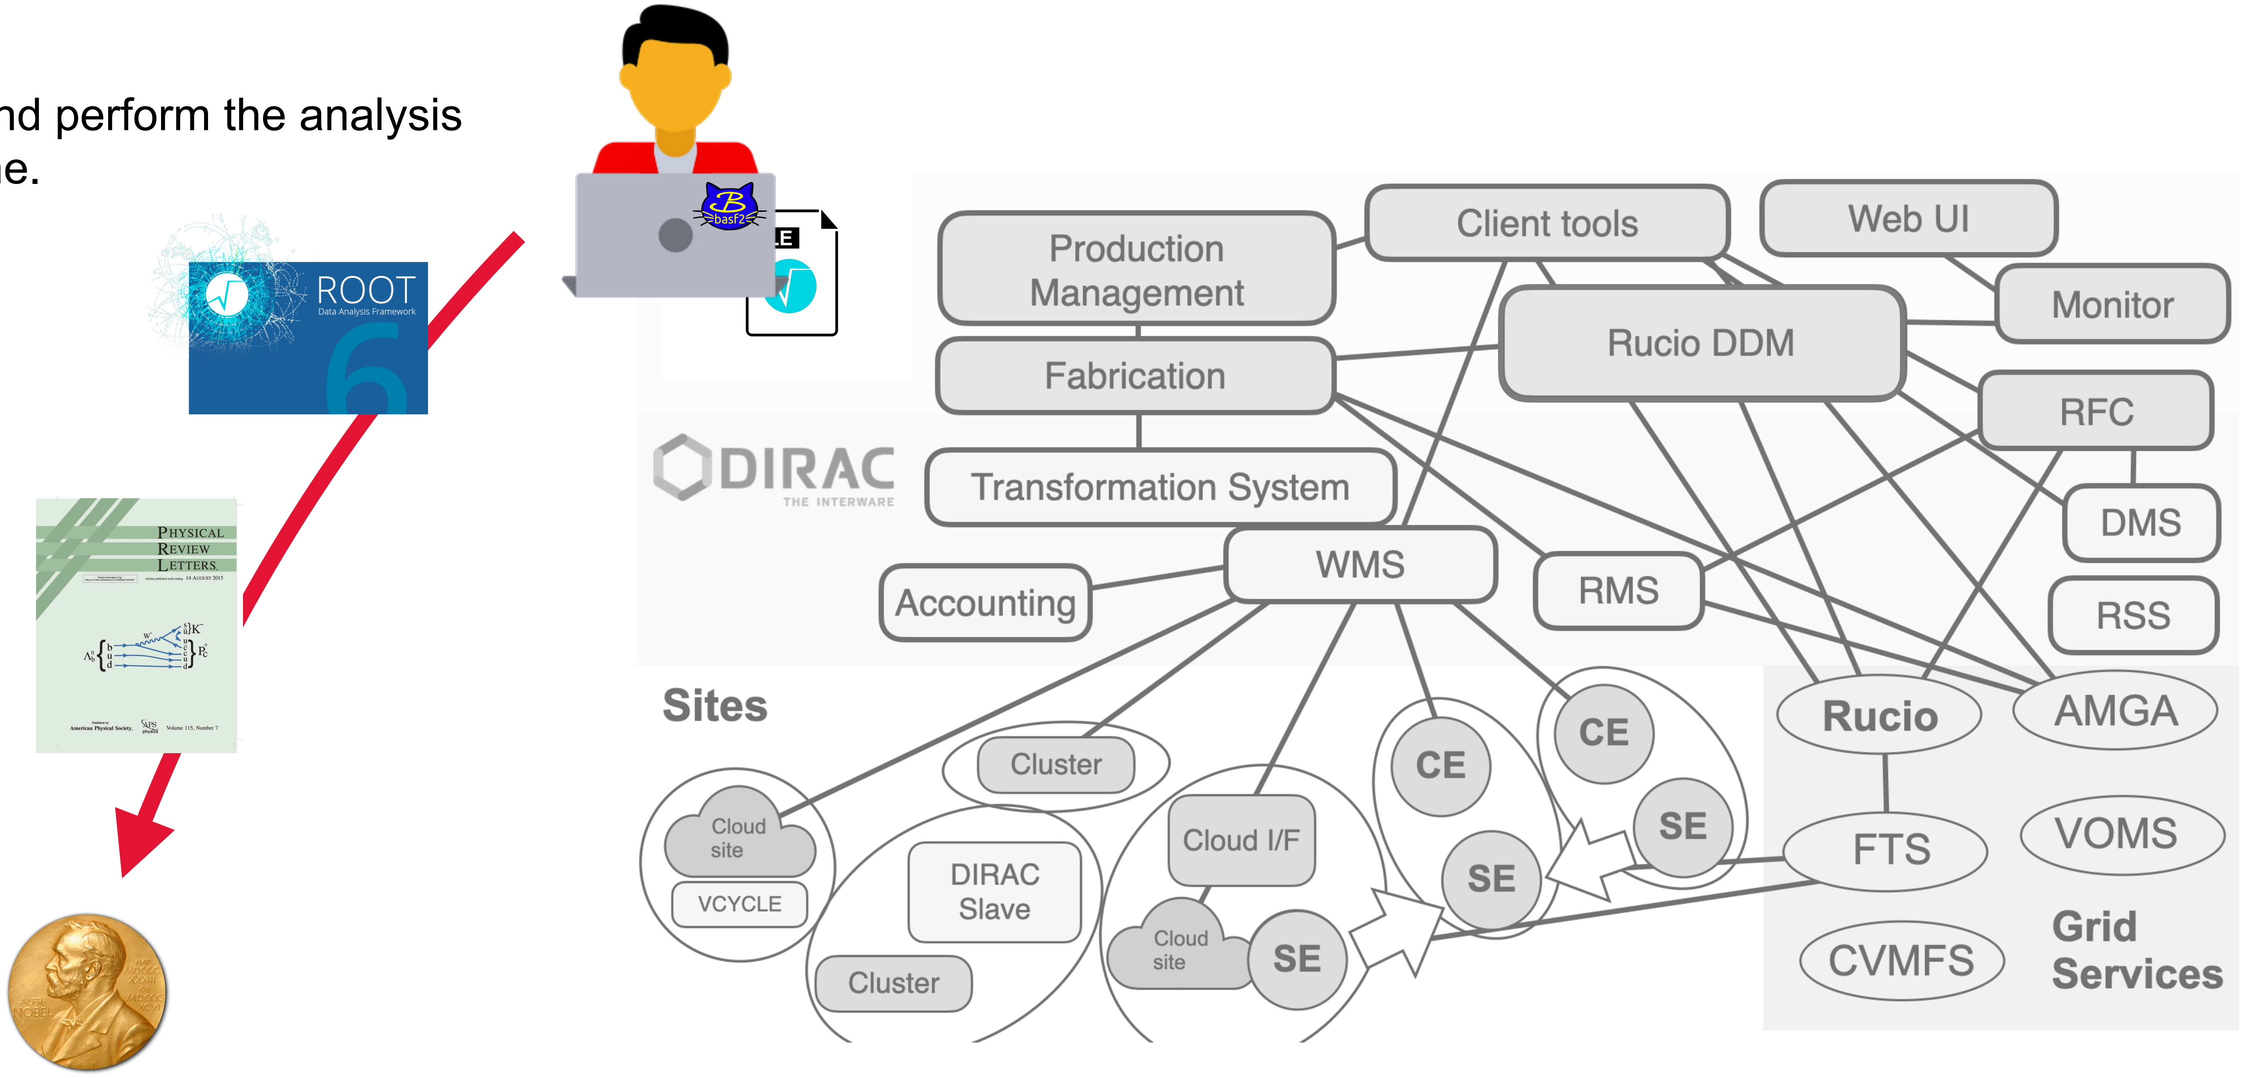

Raw Data Workflow

Online - Offline - Grid

- BelleRawDIRAC is another extension of DIRAC, partially independent from the production activities (**BelleDIRAC**).
- It takes care of the **upload**, **registration,** and monitoring the **replication** for raw and HLT skim data files.
- Once the replication to both sites is finished, the integrity of the file is verified in the raw data centers.

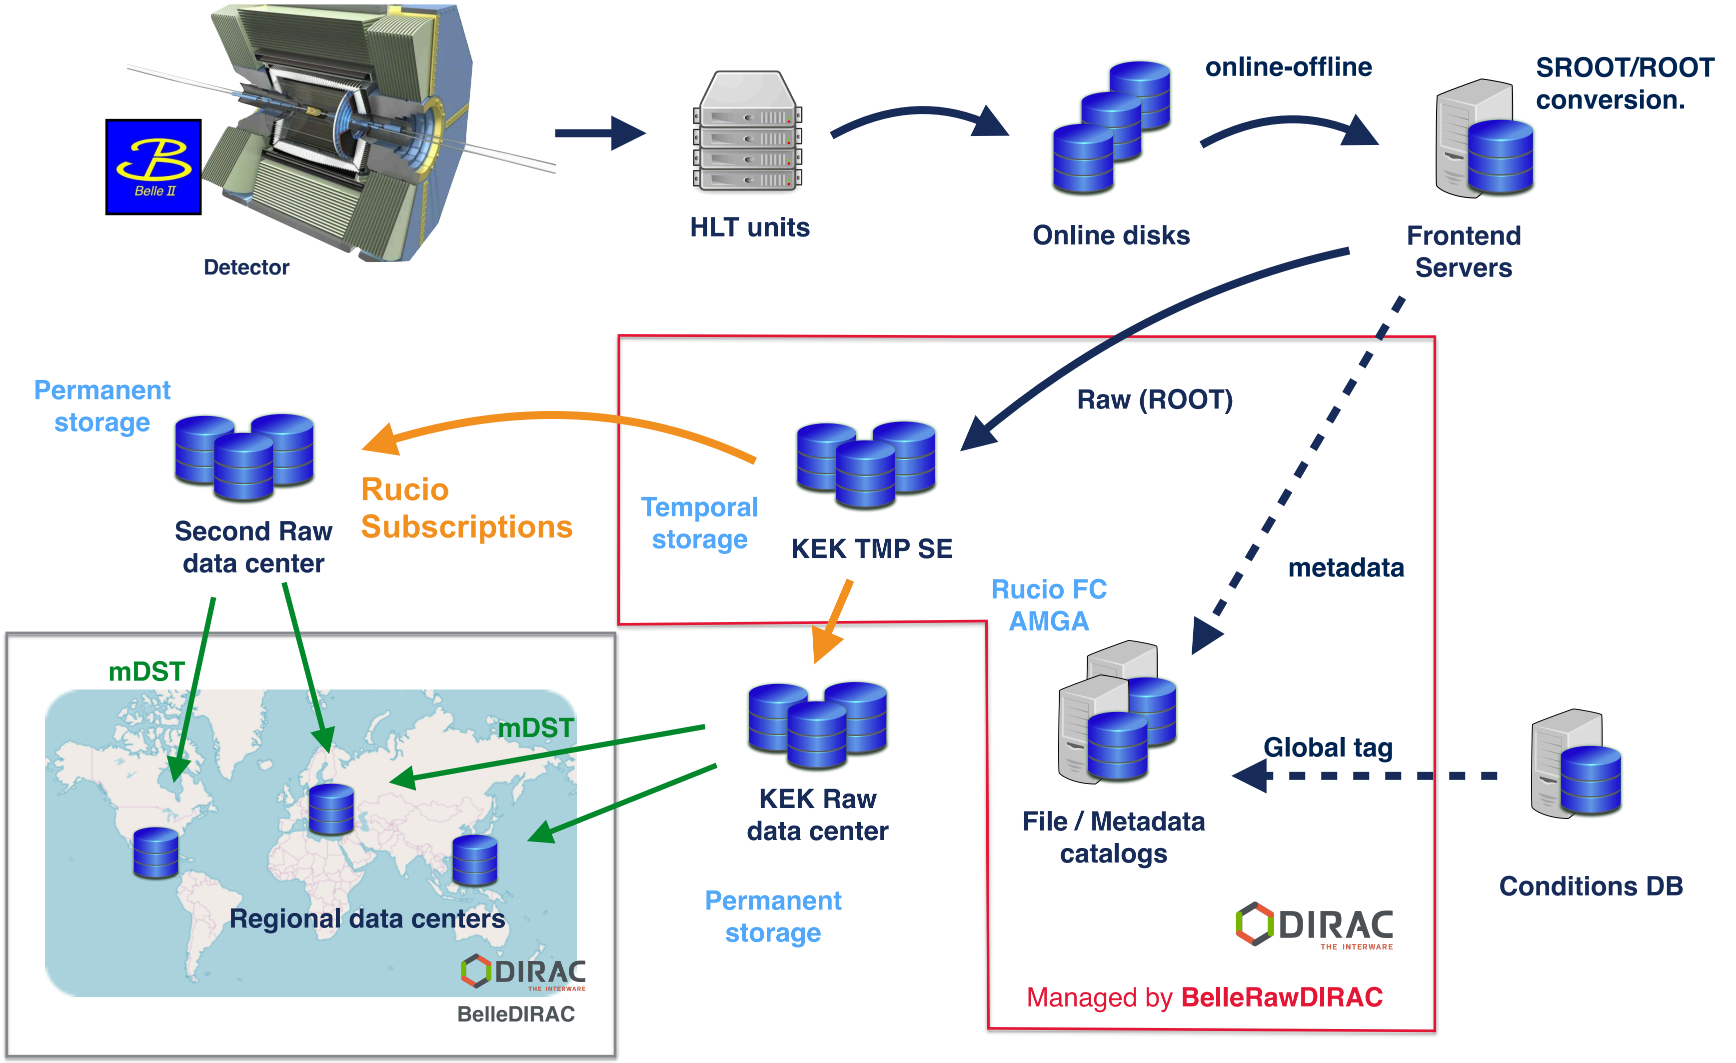

The Systematics Framework

In progress and in the future

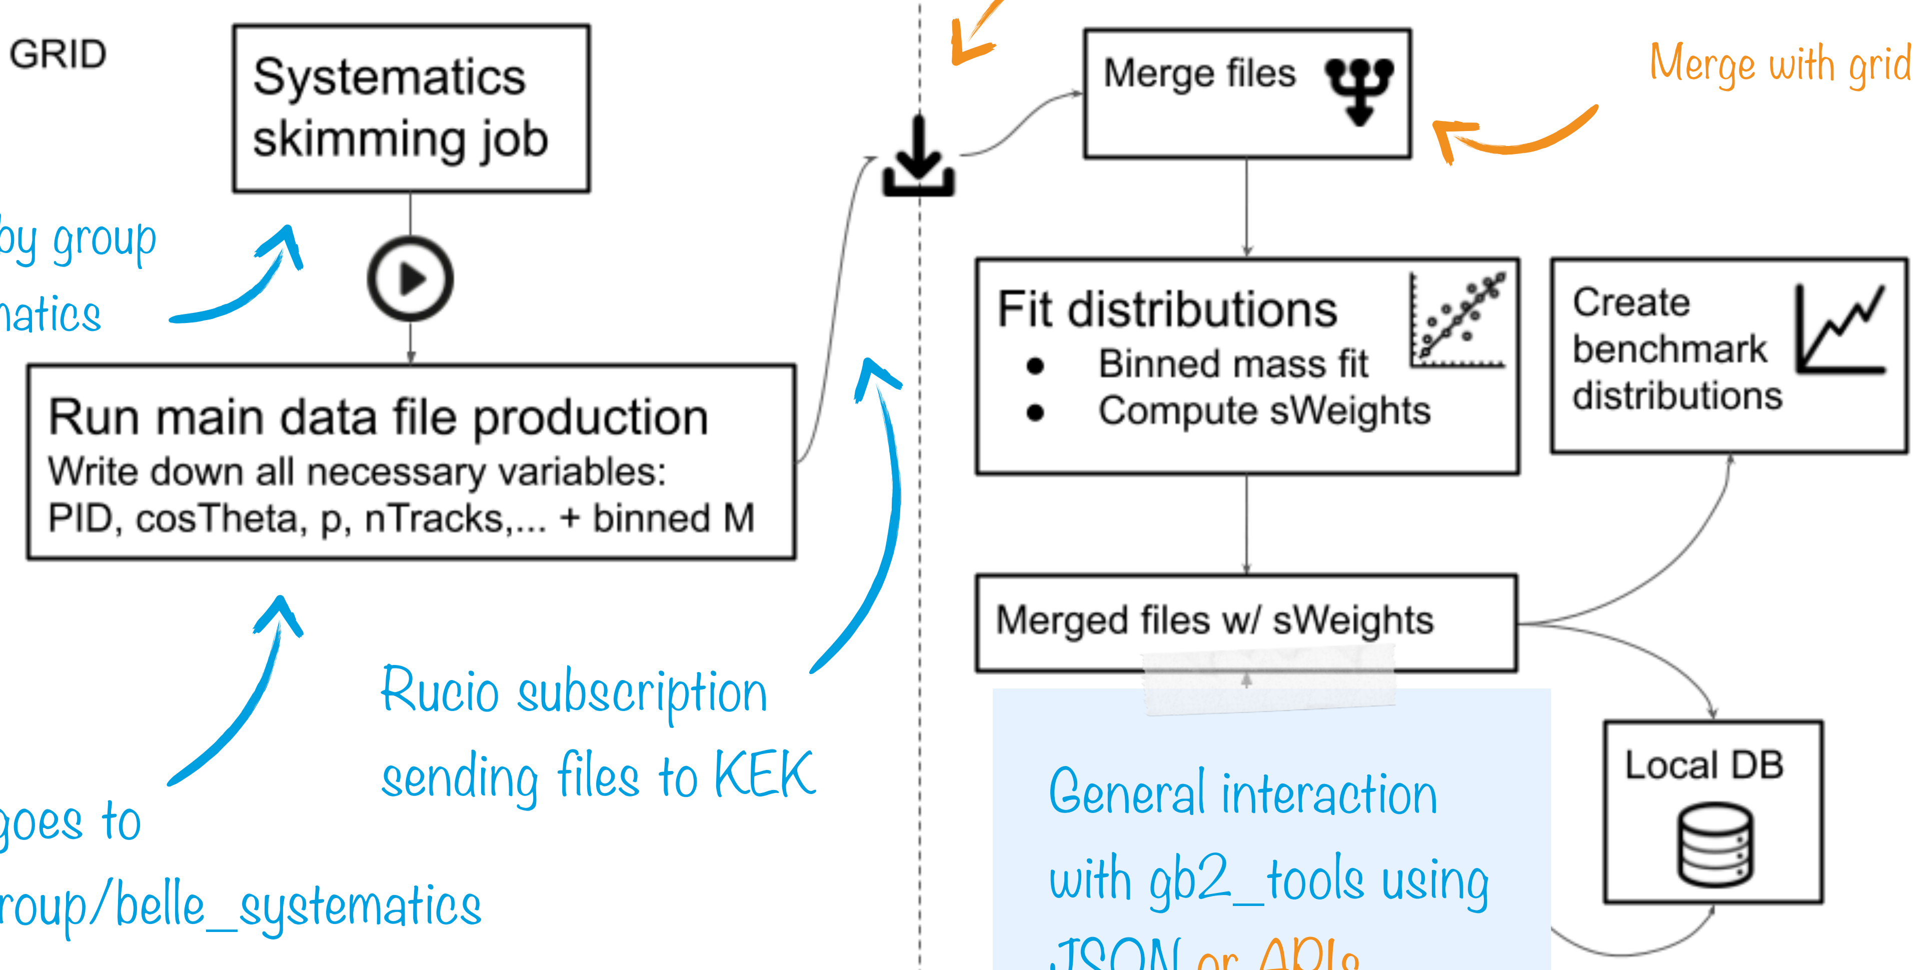

Streaming instead of download.

Merge with grid jobs.

Submission by group belle_systematics

Rucio subscription sending files to KEK

Output goes to /belle/group/belle_systematics

General interaction with gb2_tools using JSON or APIs

https://syscorrfw.readthedocs.io/en/latest/index.html

CPU time estimation

Or why my jobs stay in Waiting for long.

User jobs waiting. Average last month: 14K (last B2GM: 173K)

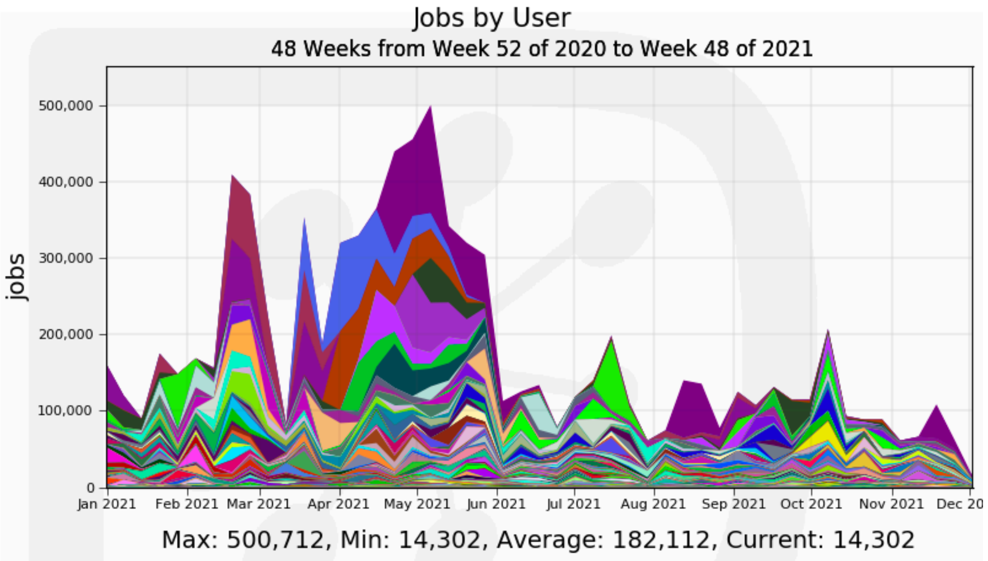

- Currently, default expected CPUtime is determined by parameters in the DIRAC configuration.
- In most cases, this is highly overestimated.
  In consequence, a very large CPUtime is estimated, leaving just a few sites/queues to execute the jobs.
- Two proposed solutions:
  - 1) Provide a better estimation of the default CPUtime based on the WMS accounting.
    - ➡ Work in progress.
  - 2) Determine the expected CPUtime from Scout jobs.
    - ➡ Will require an additional layer between the client and WMS.
- In the meantime, we are advertising to **estimate the CPUtime before the submission** (link to the FAQ).

Scout Jobs

Preventing failed jobs from users

- If the main project has a large number of jobs, a part of them are copied as a group of scout jobs.
    - Main submission proceed only if scout jobs finish without errors.
    - Otherwise, user is notified.
- Successfully implemented in production.

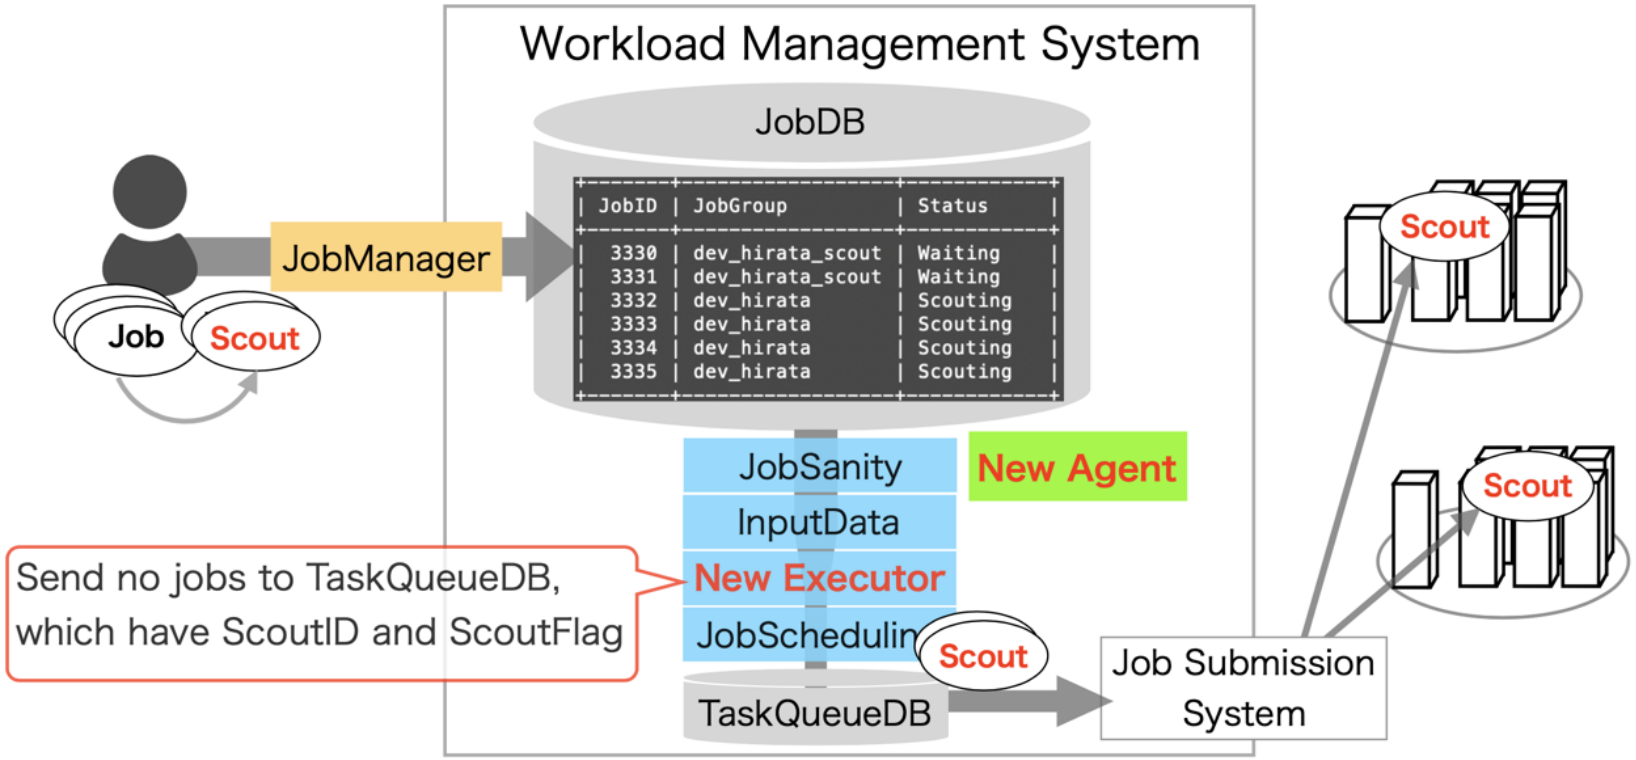

Users manual

gbasf2 from zero to hero

- gbasf2.belle2.org
- Took long since major changes in the client structure were required (BIIDCD-941).
- Maintained inside the BelleDIRAC repository, below `BelleDIRAC/gbasf2/doc`.
- Release-dependency. Easy to maintain.
- Changes in gbasf2 need to be properly updated in the documentation.
  - Enforcing in PRs.
- Plans to extend documentation for data production tools.

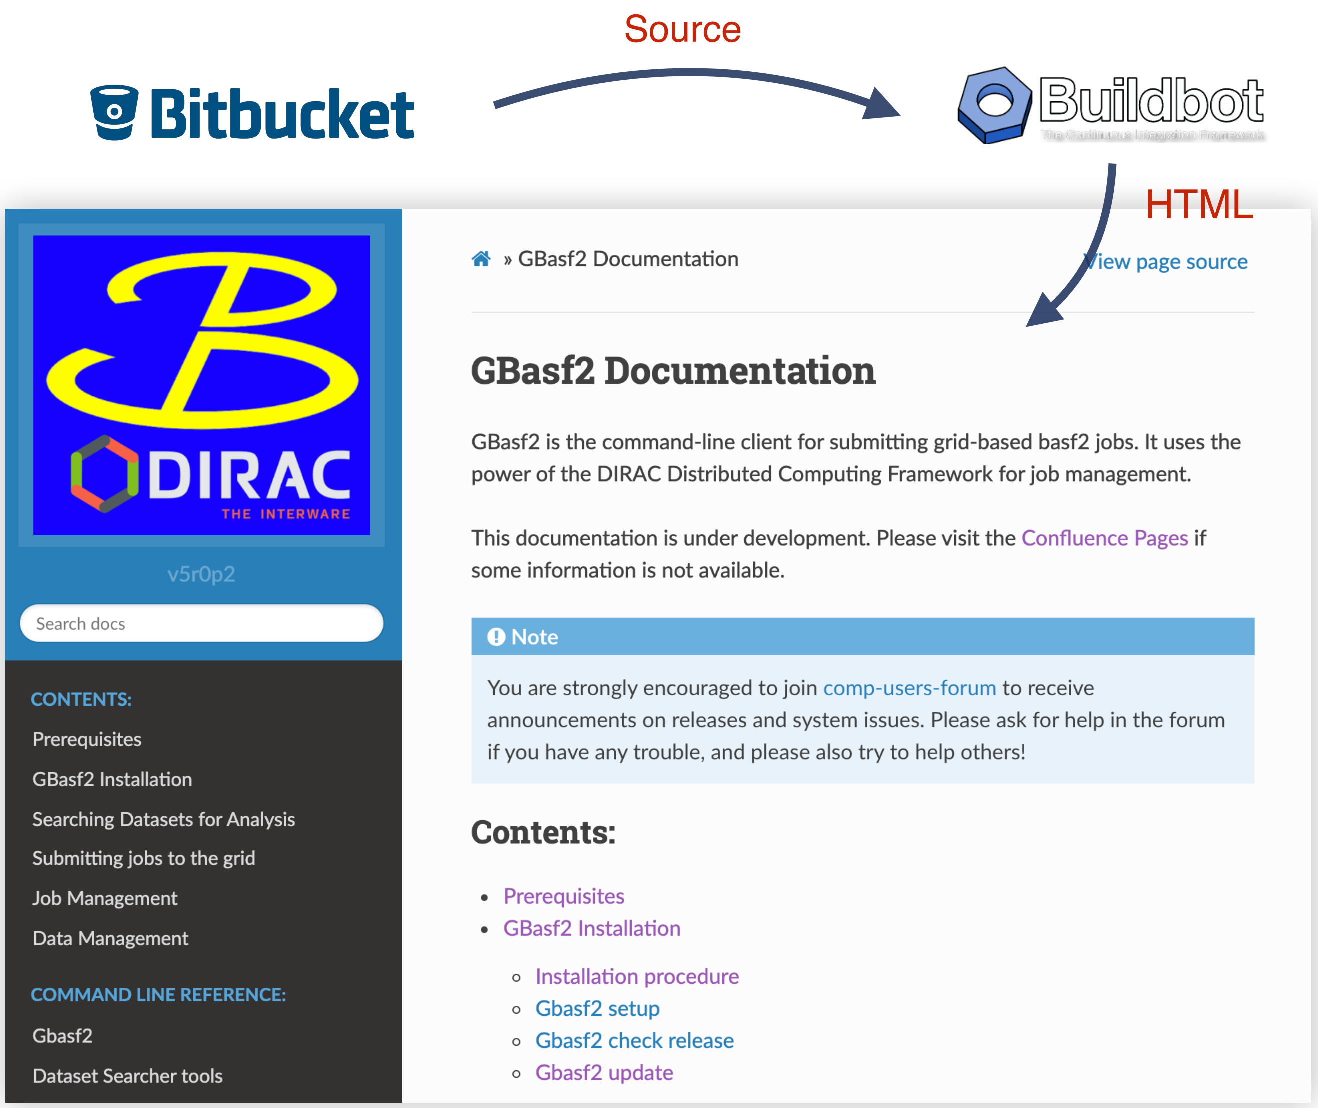

Contact

DESY. Deutsches
Elektronen-Synchrotron

www.desy.de

Michel Hernandez Villanueva
michel.hernandez.villanueva@desy.de
Orcid: 0000-0002-6322-5587