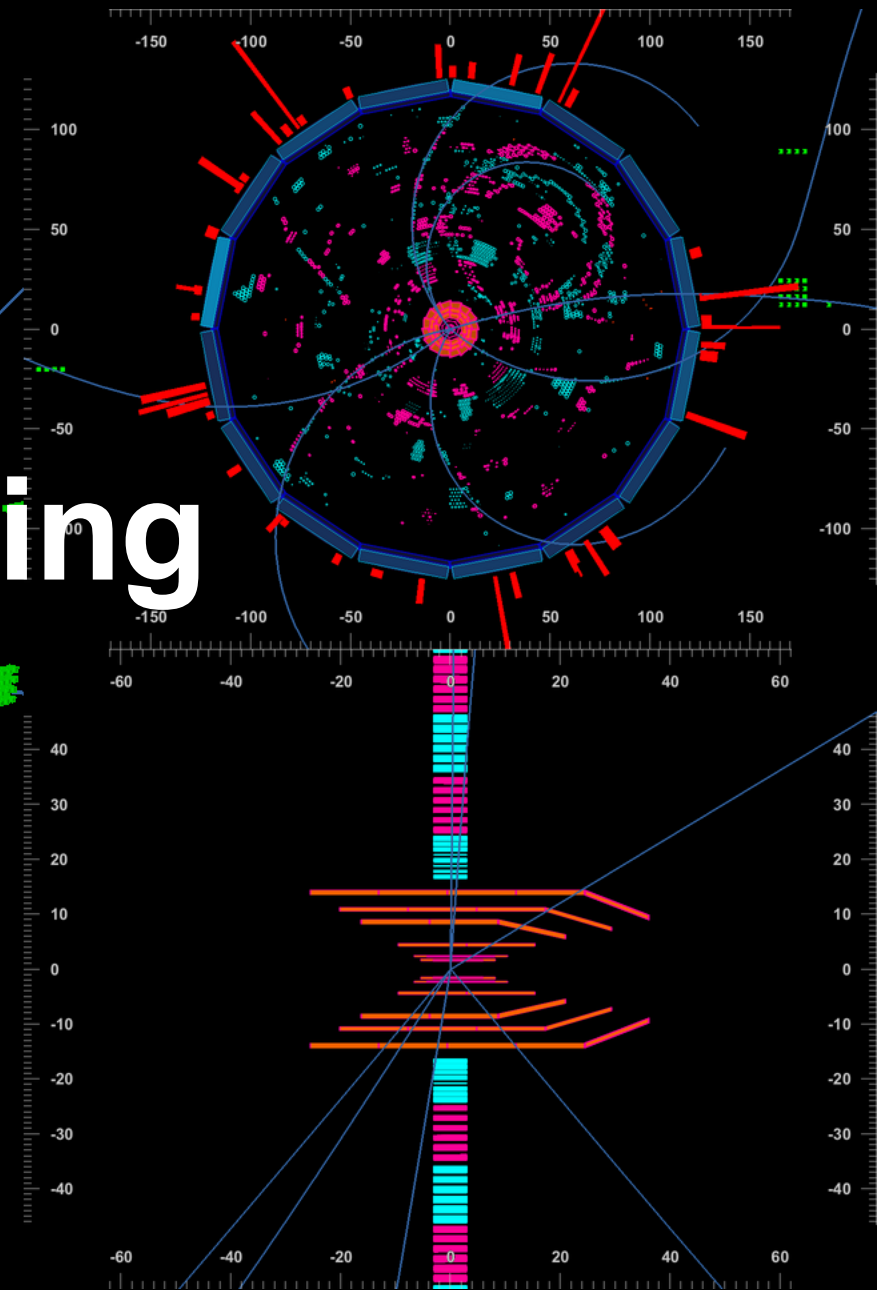




Belle II computing

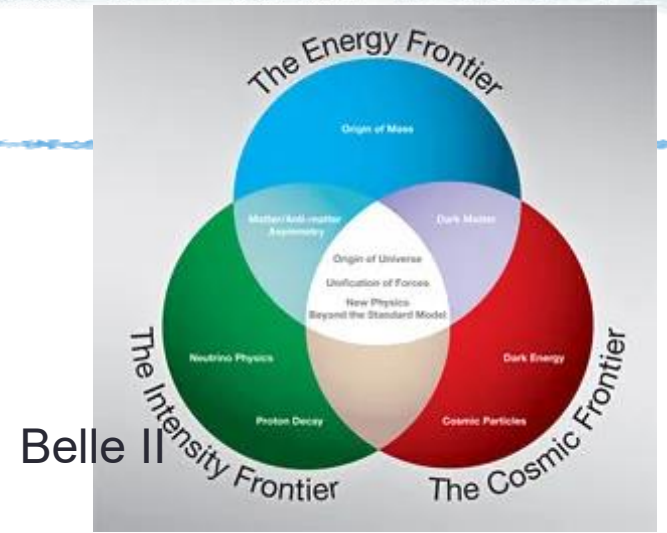
Hideki Miyake (KEK)

Run Type : physics
Run Number: Exp 42 Run 686
Created at: 2026 May. 14, 09:13:07 JST

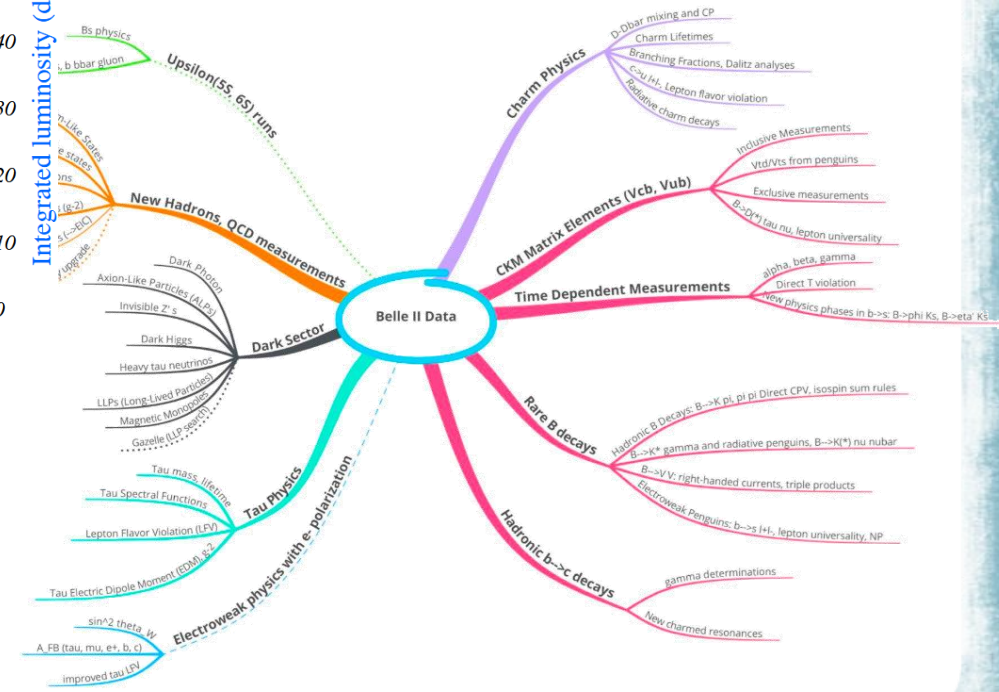
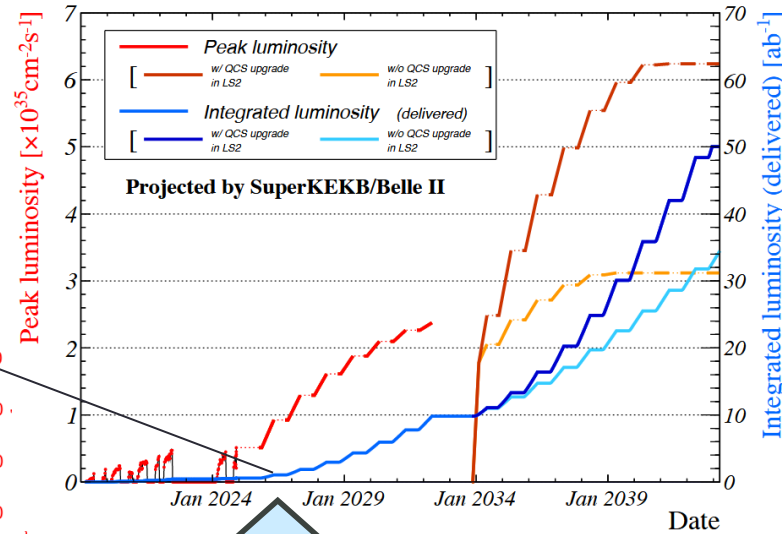
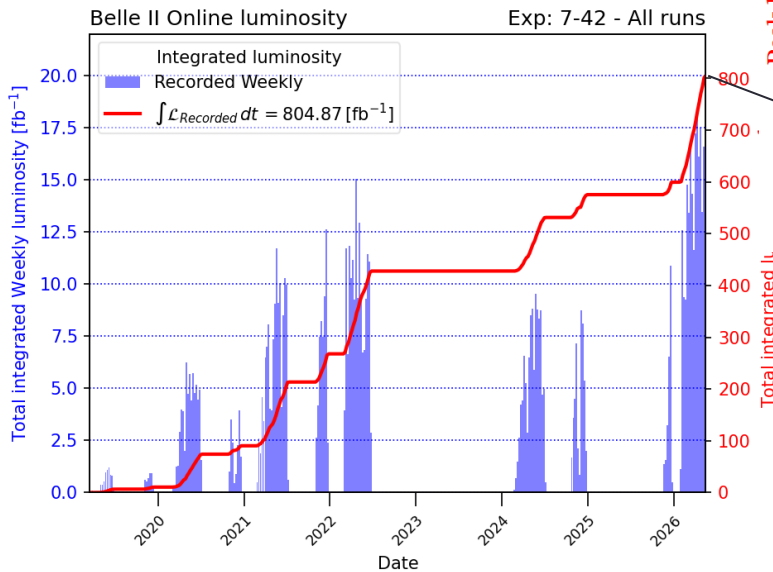


Belle II experiment



- HEP experiment for the intensity frontier
 - Ongoing rich physics program (not only B, but Tau, dark matter...)
 - Aiming at 50/ab corresponding to ~x50 of Belle ~60PB



- First data: 2019
- Finished long shut down in 2024
- 2026 data taking is ongoing
- Collected 0.8/ab as of today



SuperKEKB + Belle II



Asymmetric e^+e^- collider at KEK (e^+ :4GeV, e^- : 7GeV)
Instant luminosity record of $5.2 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$ (x2.4 of KEKB)

- Composed of 7 major subdetectors
- Hermetic detector
- Good particle ID
- Excellent tracking efficiency and improved vertex resolution

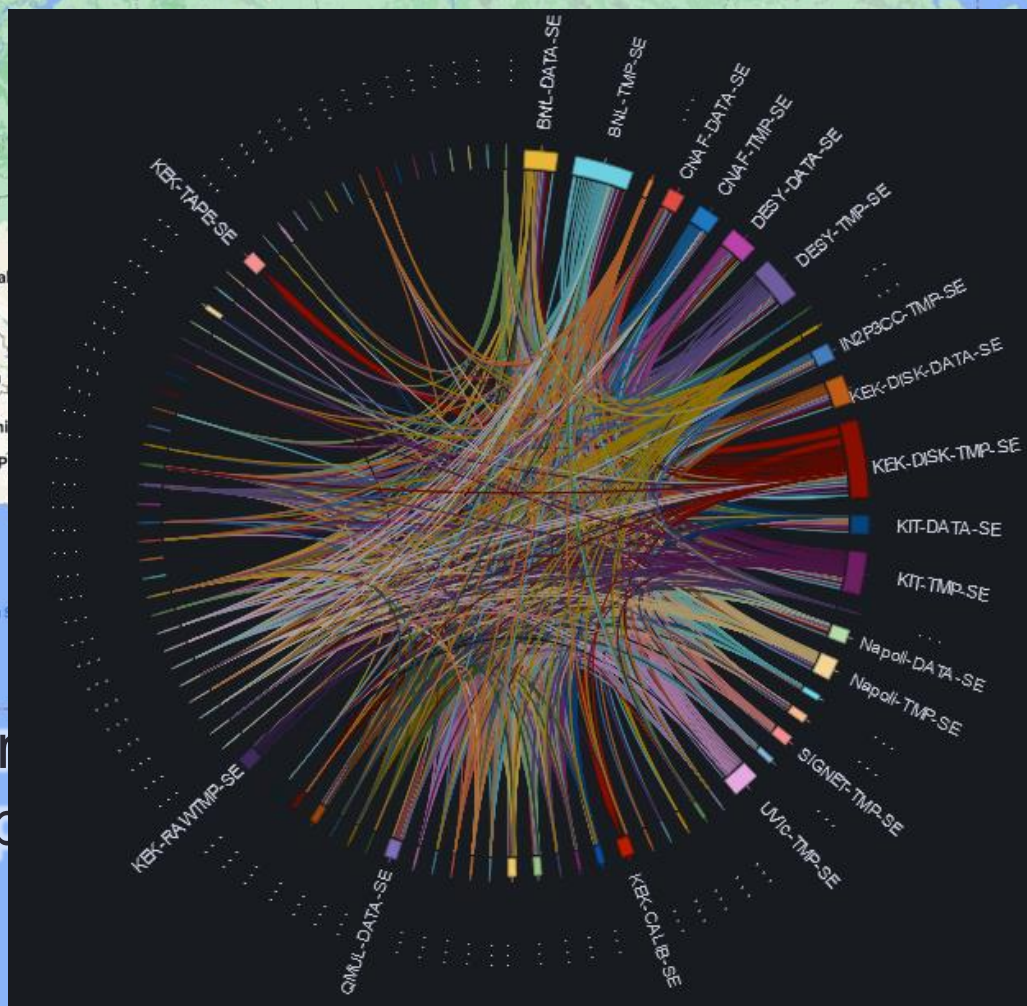
NEW at Mar 19 in this year!

Belle II collaboration

High-speed network mesh is a key!




1240 members
Widely spread

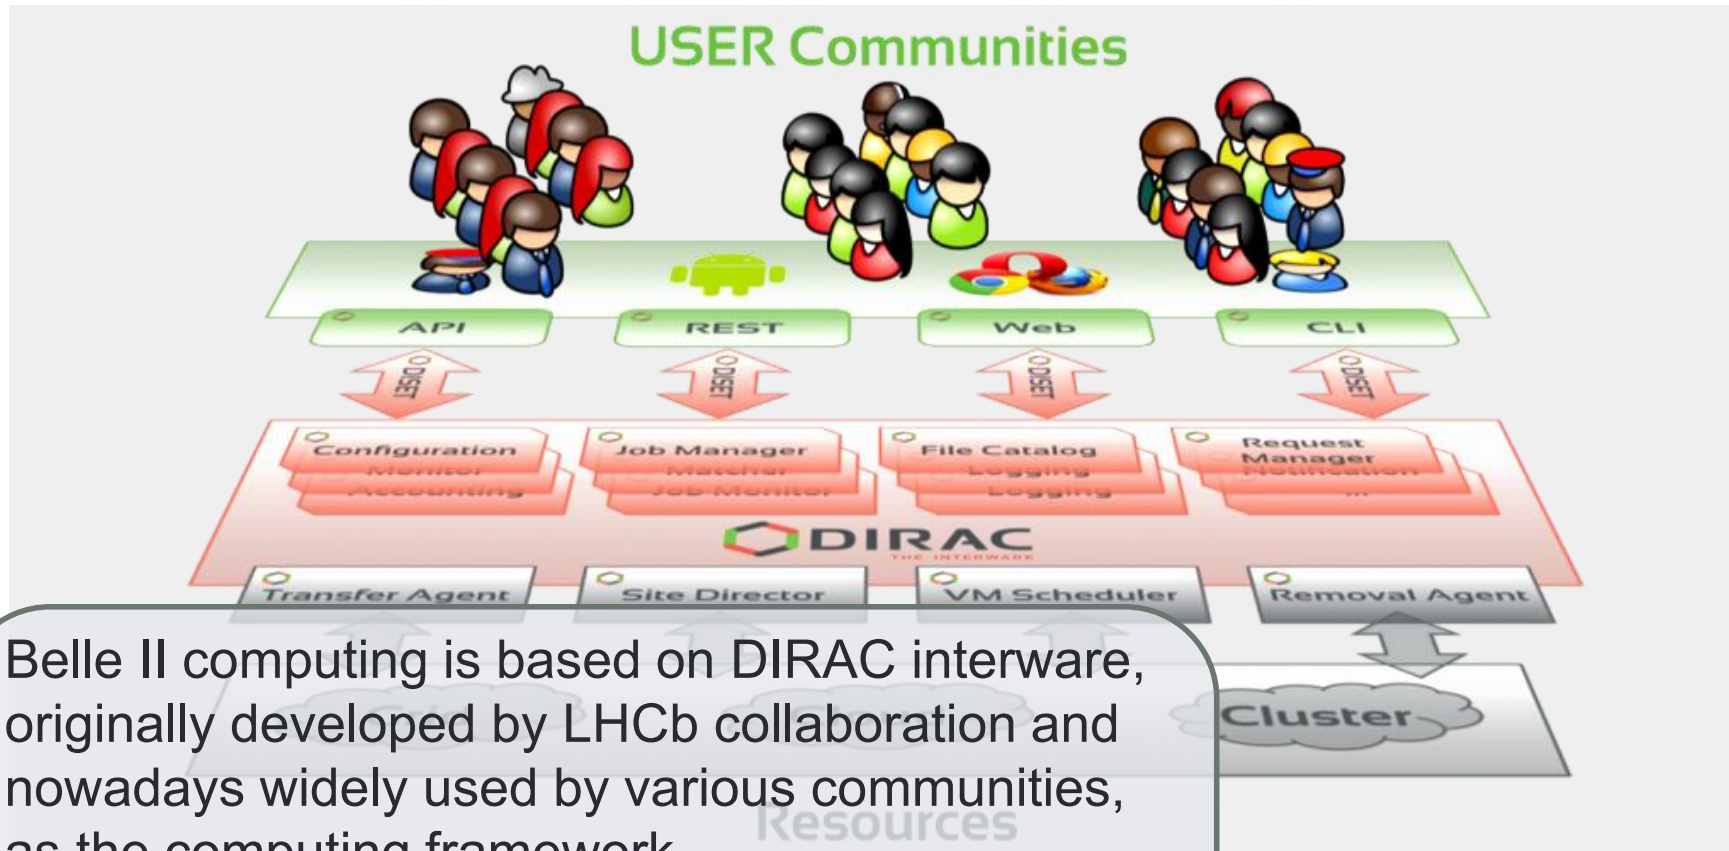


Europe and regions
(Europe, Asia, and NA)

Belle II computing model

- Data is transferred from the offline servers to the KEK data center.
- Raw data centers around the world keep a second copy of the full raw data set.
- Raw data is processed, skimmed and distributed over Raw Data Centers and Regional Data Centers.
- MC production is performed at all sites.
- Analyzers access data and MC sending jobs to the grid and downloading the output to local resources.





Similar framework used in other communities

PanDA: ATLAS
JAlEn: ALICE

....

Belle II computing is based on DIRAC interware, originally developed by LHCb collaboration and nowadays widely used by various communities, as the computing framework

- single interface to handle various resources
- modular design: community can easily extend the functionality

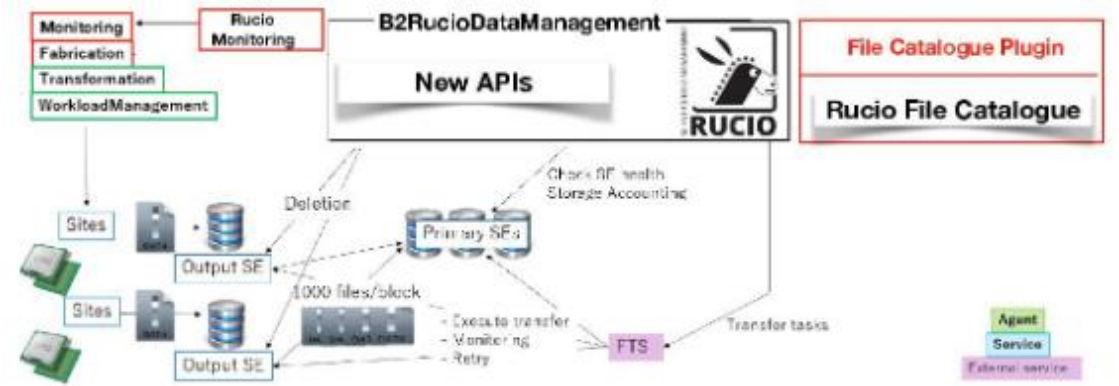
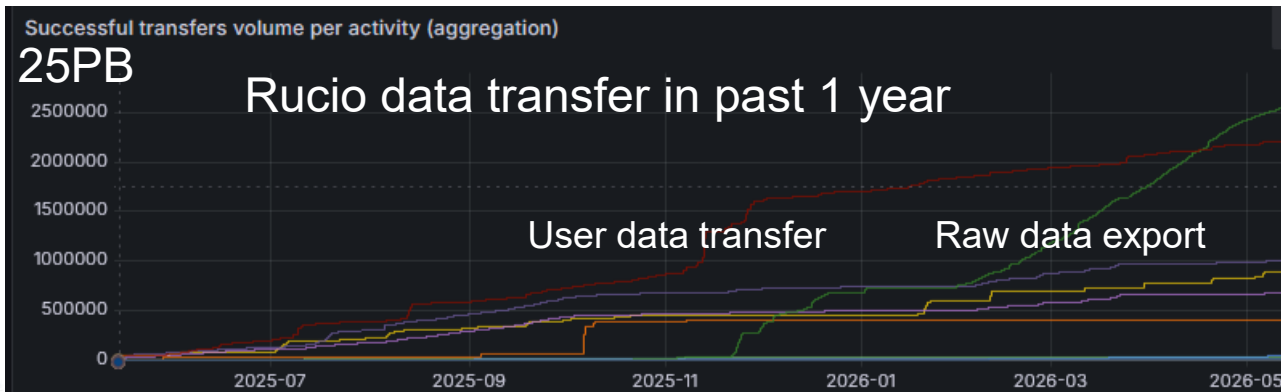
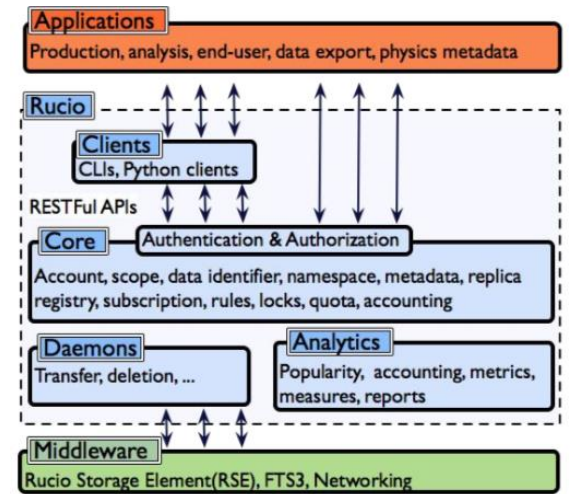
In Belle II, two different DIRAC systems are utilized

- BelleRawDIRAC for “Core Computing”
- BelleDIRAC for Distributed Computing

Rucio



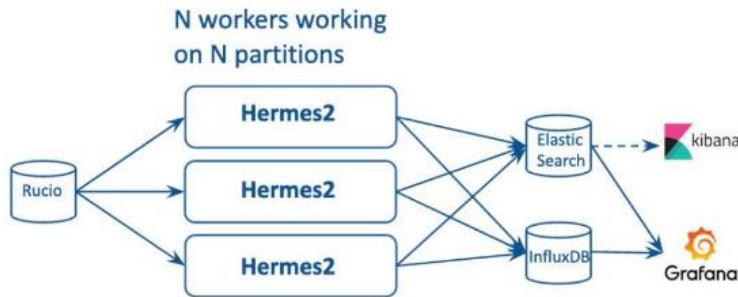
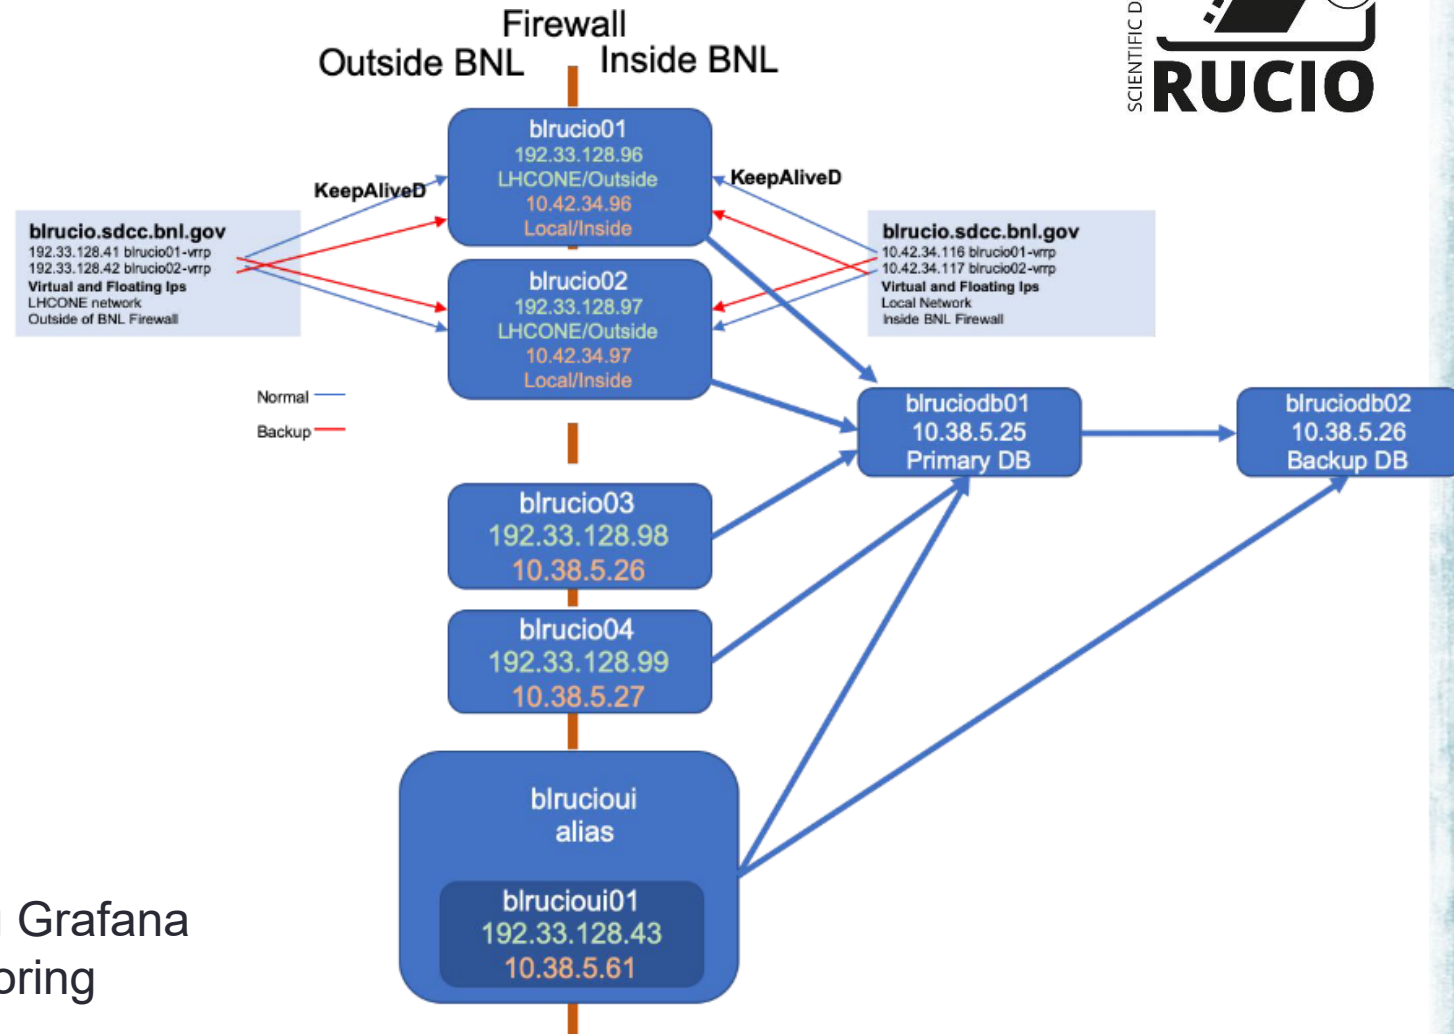
- Rucio is a highly-scalable policy-driven data management system
 - Originally built for ATLAS experiment and adopted by Belle II since 2021
- In Belle II, Rucio is used for several purposes
 - As Distributed Data Management System
 - Transfers between site using policies engines (rules and subscriptions)
 - Monitoring and transfers, deletions, SE occupancy
 - As File Catalog
 - As Metadata Catalog
 - We have long used AMGA to store “data of data”
 - Now in final phase to migrate to Rucio



Belle II Rucio

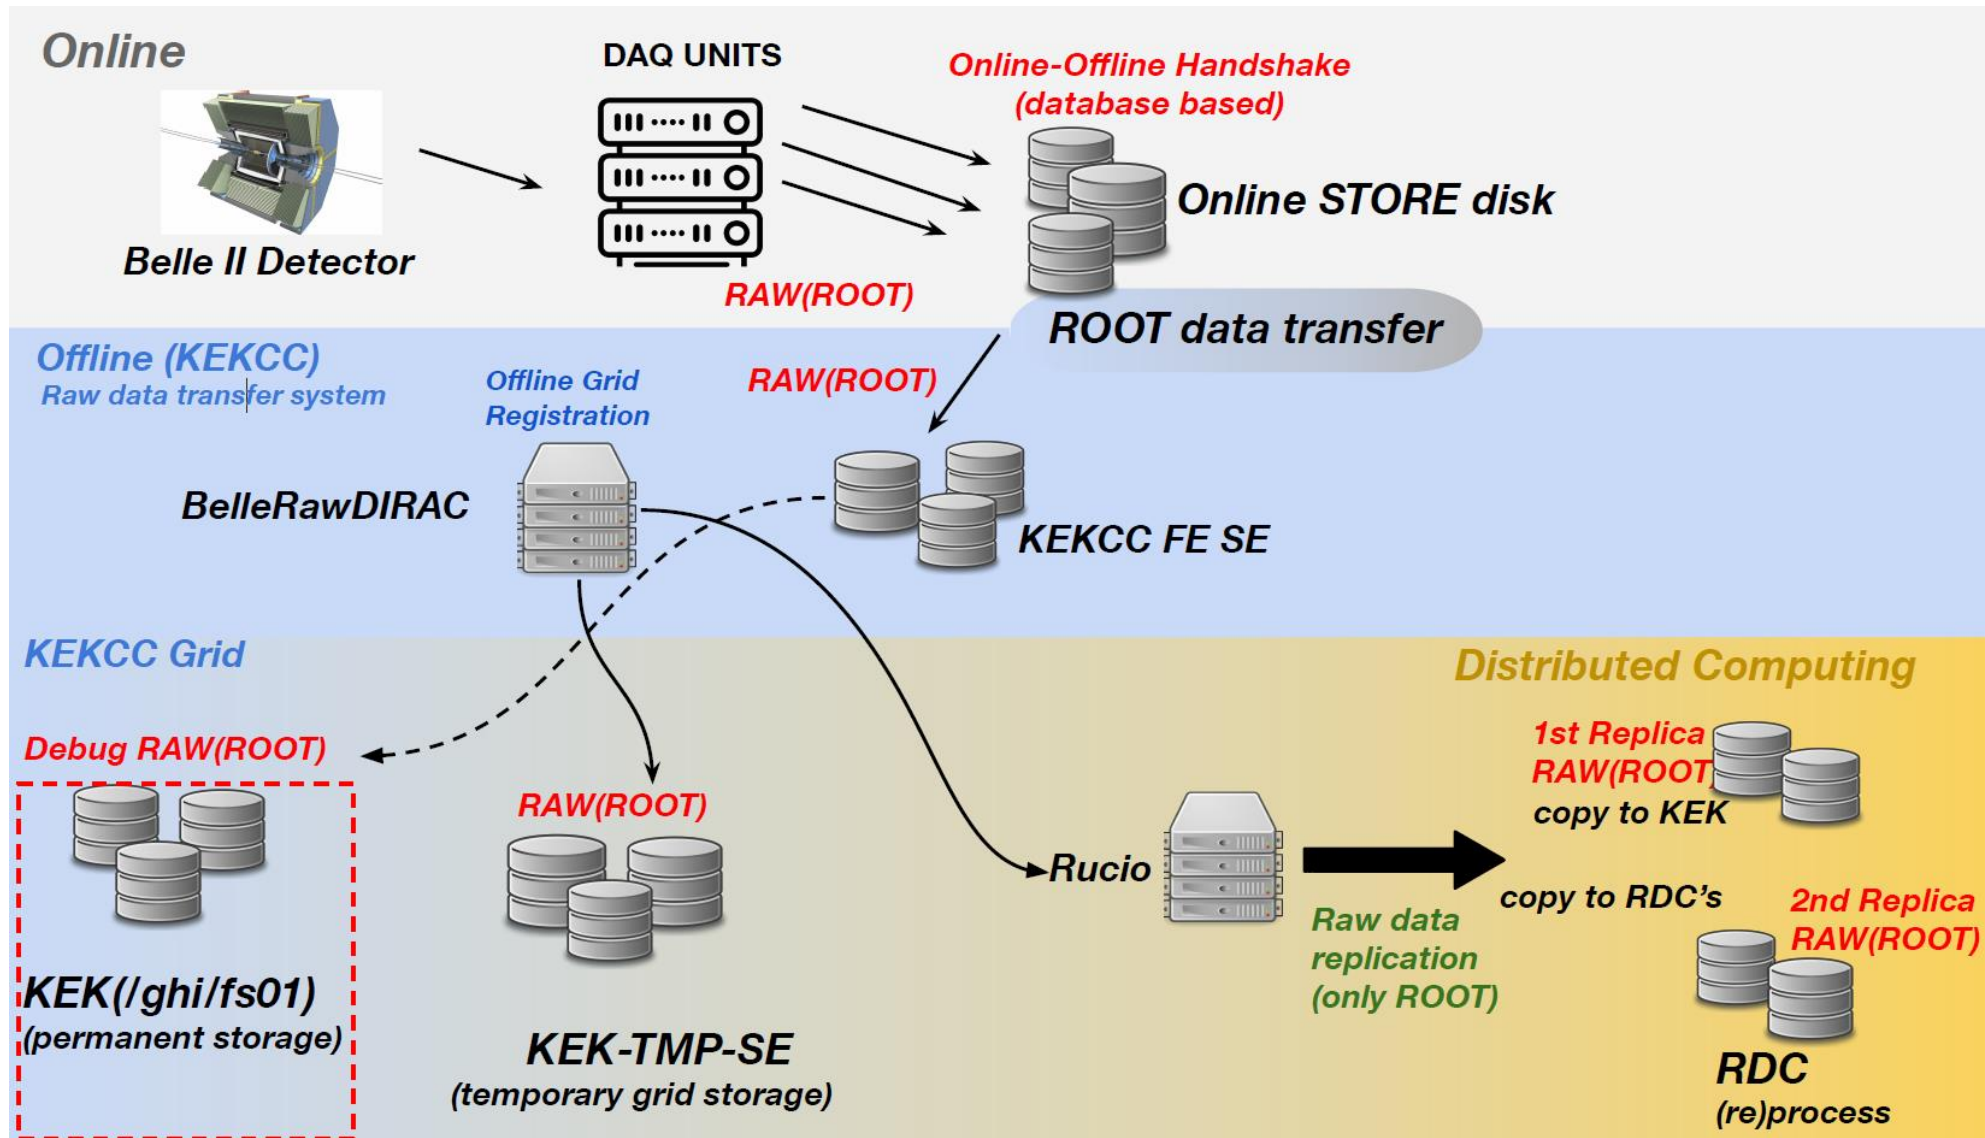


- Belle II Rucio is hosted by BNL
- Initially on RHEV, then moved to OpenShift in this year
 - 4 servers, 2 DB and 1 UI
- In BNL OpenShift, conditions DB (e.g. calibration constant) is also deployed



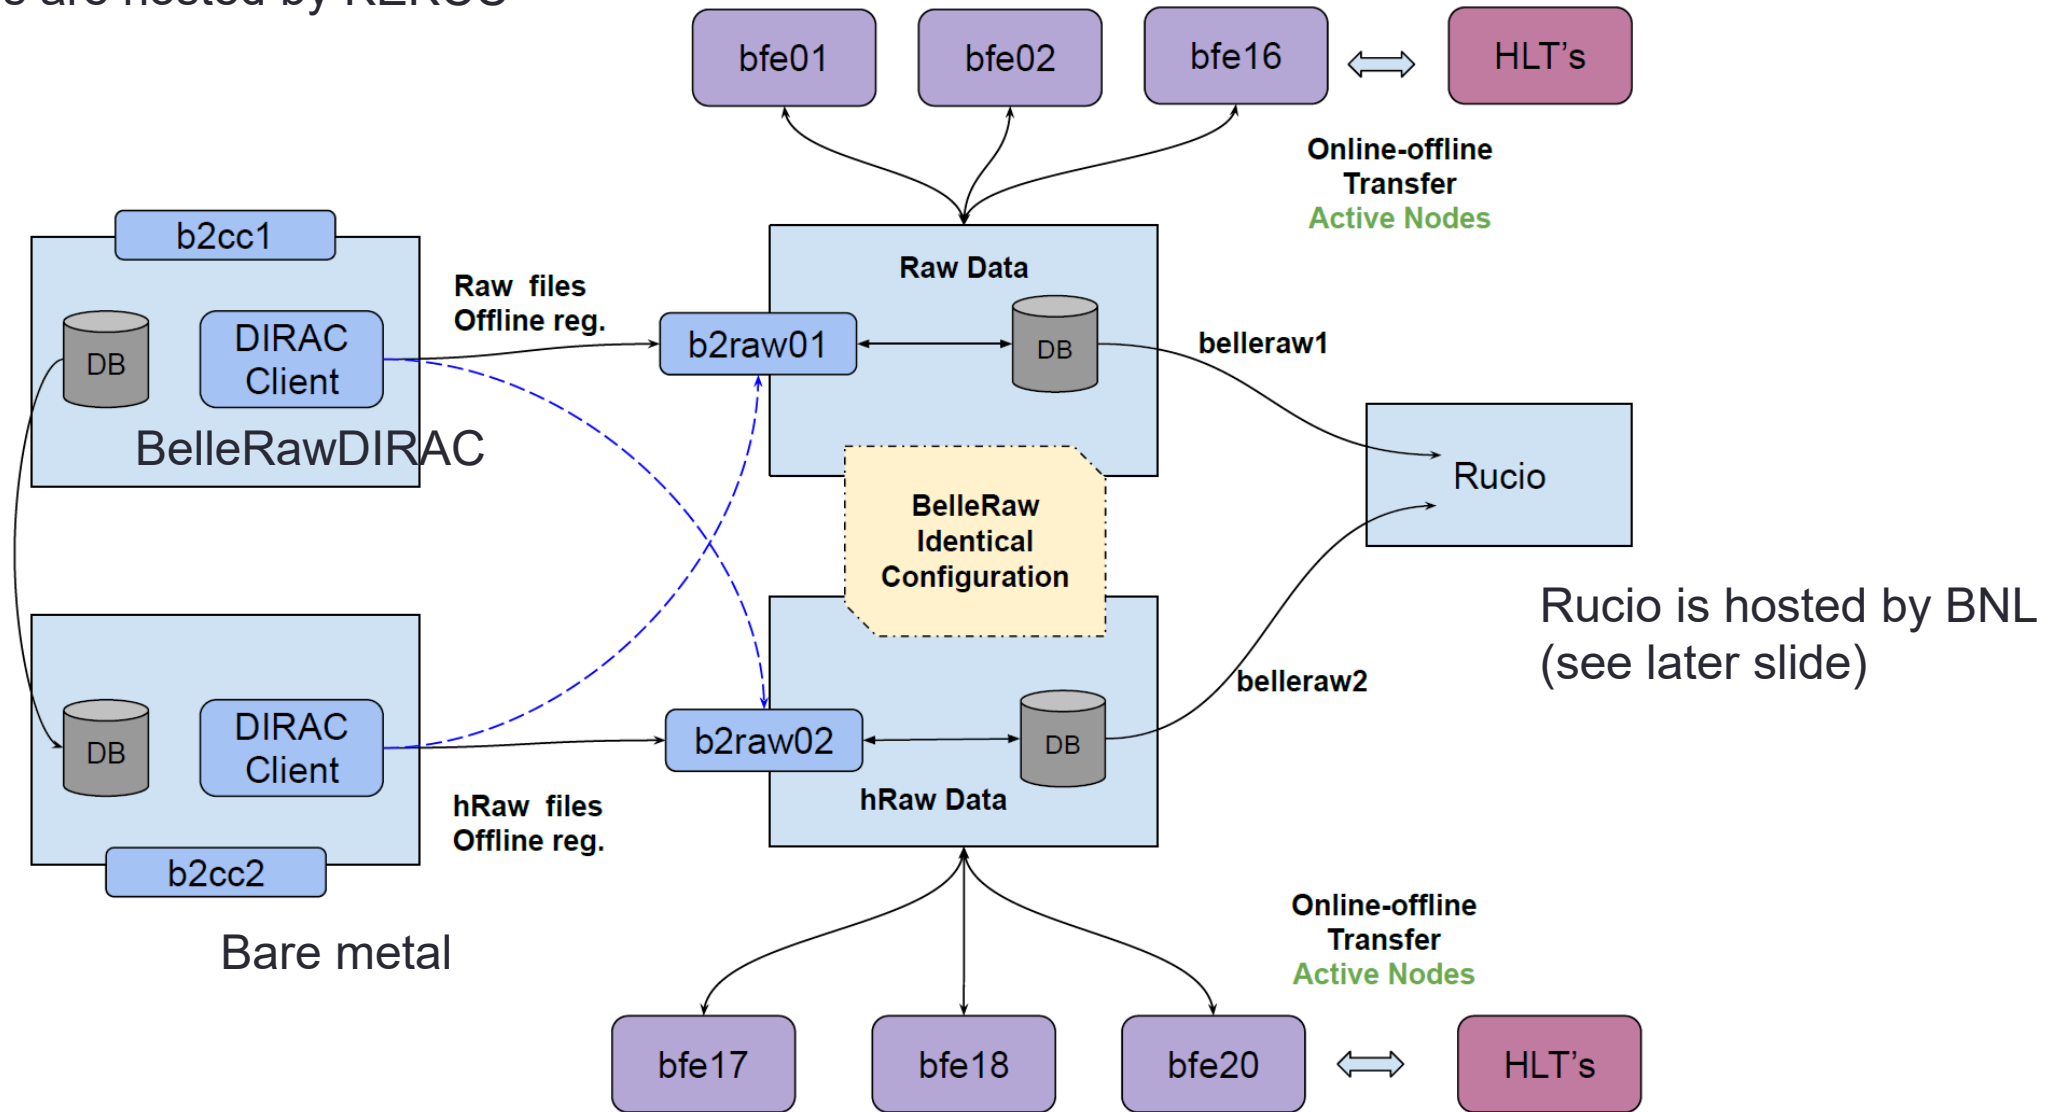
Providing Grafana for monitoring

Raw data transfer in Belle II



Belle II “core computing” system

Most servers are hosted by KEKCC




Kubernetes for core computing

We started to use Kubernetes (K3s) to host some core computing functionalities

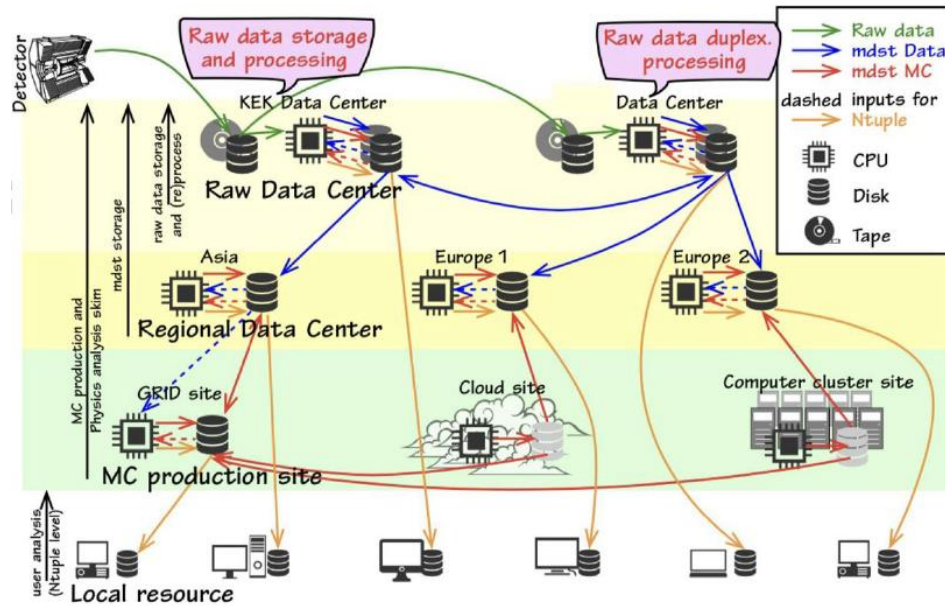
- Centralized DIRAC log
- Centralized syslog (OpenSearch, LogStash, RabbitMQ)
- Grafana
- DiracX (see later slide)

Infra for evaluation

4 VMs in 2 nodes
for now
(see later slide)



Distributed computing in Belle II



After Raw data is transferred to Raw data centers, apply various processing and filtering (“Skimming”) toward final analysis result

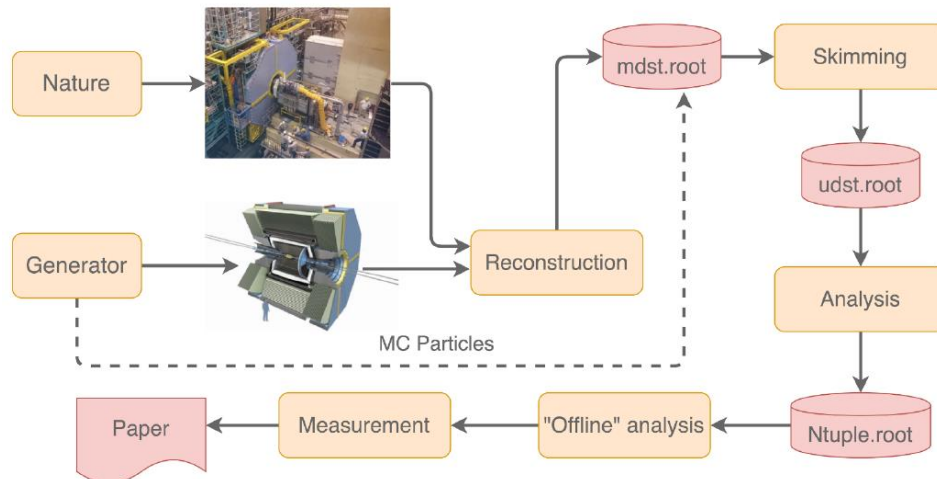


Raw data (or MCProduction)

mDST Reconstructed event info

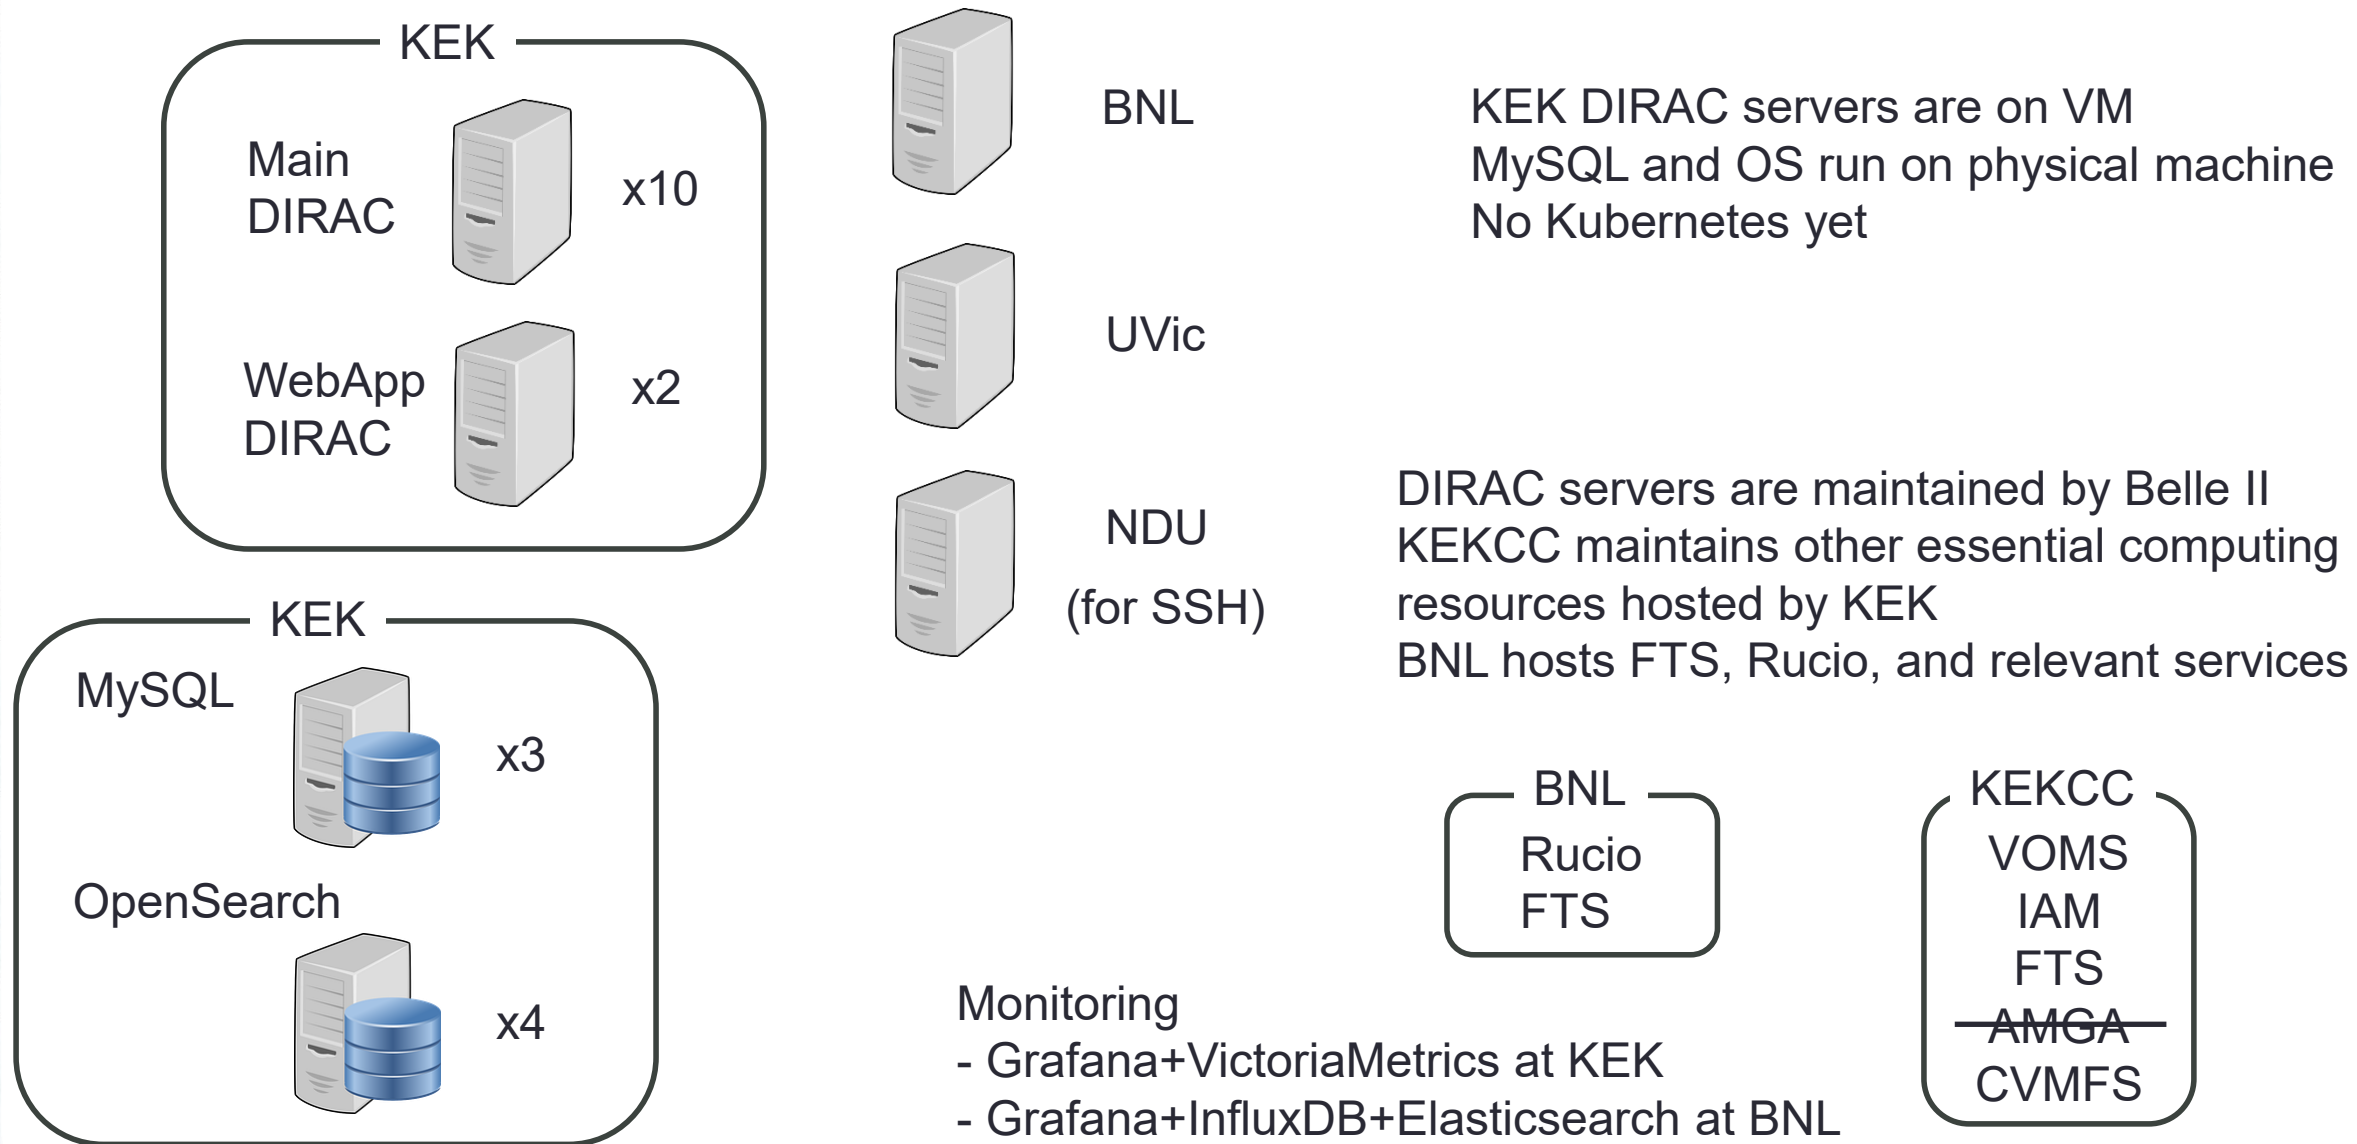
uDST mDST+ particle list (high level objects)

nTuple Analysis results (only values)



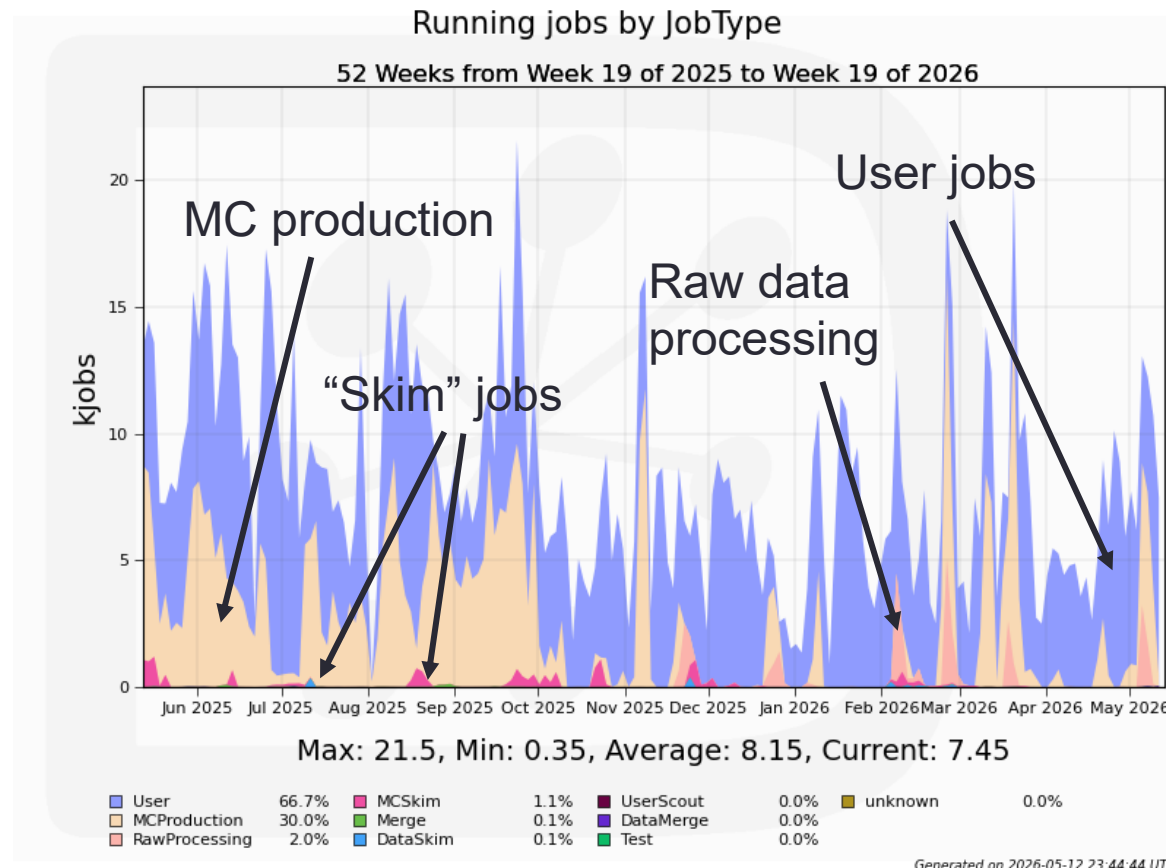
Typically the last step is done by individual (or group) while the others are processed by mass production system centrally

Belle II distributed computing system




Belle II grid job status


- Intermittent peaks are seen but generally consistent
- Comparable CPU usage by MCProduction and User analysis




Operation



Data Production Shift
- Daily shift



Operation Expert
- Weekly shift



DC experts

Data Production						
11 Mon	12 Tue	13 Wed	14 Thu	15 Fri	16 Sat	17 Sun
Chehra Madan	Chehra Madan	Chehra Madan	Chehra Madan	Melhong Liu	Melhong Liu	Melhong Liu

Data Production Expert						
11 Mon	12 Tue	13 Wed	14 Thu	15 Fri	16 Sat	17 Sun
Jake Bennett	Jake Bennett	Jake Bennett	Hiroaki Ono	Hiroaki Ono	Hiroaki Ono	Hiroaki Ono

- Monitor DC system and activities (e.g. DIRAC, Rucio...)
 - Record shift log
 - Report issues to operation expert
 - Update current situation in Operation Status Page
-
- Support DP shift
 - Issue/manage tickets (GGUS, our issue tracking system...)
 - Escalate to DC experts if necessary

Automatic Issue detector

[All](#)
[DownTime](#)
[Raw Time](#)
[Pilot Trend](#)
[Pilot Submission](#)
[Pilot Processing](#)
[Pilot Waiting](#)
[Job Trend](#)
[JobStatus](#)
[RankJobStatus](#)
[Userlist](#)

Central Systems

SE occupancy check

Red color for more than 95.0%

Primary SE
No high occupancy SE

Other SE
No high occupancy SE

Sites

[EuroHPC Vega AI](#)
 "Pilot Submission Failure" has been observed since 2023-08-14 15:01 UTC (for 24069 hours) [\(details\)](#).
[LCC LANSBIM IV](#)
 "Pilot Submission Failure" has been observed since 2024-10-17 09:20 UTC (for 13755 hours) [\(details\)](#).

Productions

Raw data processing
No problematic input file found

Transformations

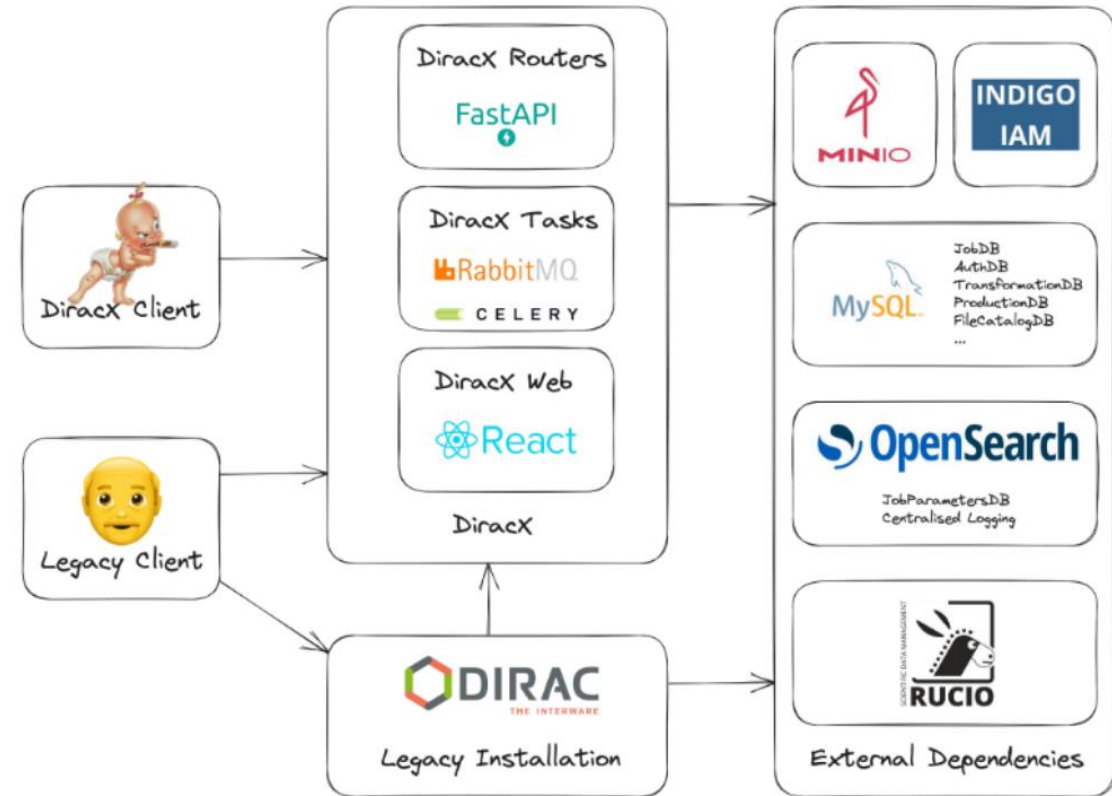
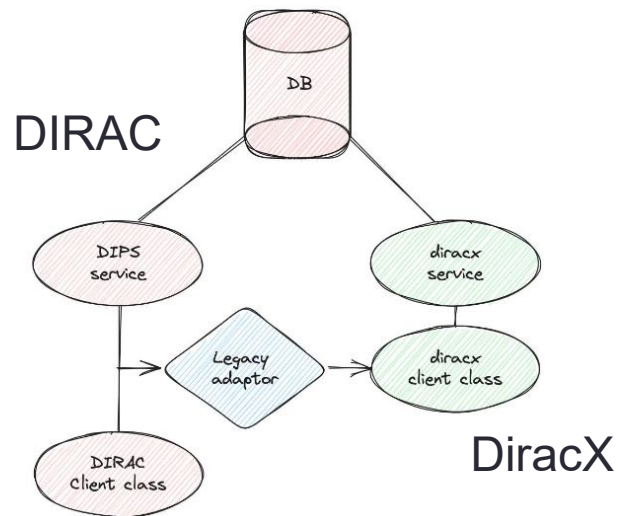
ProductionID	TransformationID	JobType	Progress	Job Done/Failed/Wait/NoSub (Total/Done)	FailRate	Duration	CreatedDate	Description
54896	6548273	MCProduction	Prog=0.00%	0/17/225/15712 (Total/Done)	100.00%	0 days	2026-05-13 12:00:27	MCdThompR9@ (G4Lab, bucket45, 45 v2)
54897	6548279	MCProduction	Prog=0.02%	2/5/210/7861 (Total/Done)	90.36%	0 days	2026-05-13 12:01:06	MCdThompR9@ (G4Lab, bucket45, BeamEnergy 43 v2)
54898	6548285	MCProduction	Prog=0.07%	3/32/128/3281 (Total/Done)	43.13%	0 days	2026-05-13 12:01:06	MCdThompR9@ (G4Lab, bucket45, 45 v2)

Free space accounting (Monitors in Grafana)



DiracX: The neXt DIRAC incarnation

- Current DIRAC has been developed more than 20 years
 - Many outdated technologies and house made tools/protocols
- DiracX
 - Completely different design
 - Standards based
 - A cloud native app
 - Providing transparent migration for users



Current challenges

- Token based authentication
 - IAM server
 - Belle II uses KEKCC instance (belle-auth.cc.kek.jp)
 - Job submission
 - Test is ongoing at validation environment
 - Plan to start it for production
 - File management
 - Sites getting configured to use token for file transfer/deletion
 - Plan to join Data Challenge 2026/2027
- DiracX migration
 - DIRAC and DiracX run simultaneously, and gradually moving to only DiracX
 - DiracX is deployed in Kubernetes (e.g. K3s)
 - BelleRAWDIRAC: already having K3s evaluation cluster
 - BelleDIRAC: not yet
 - Main blocker: extension migration to DIRAC v9 (essential for DiracX)
 - Will start validation at BNL

Summary

- Belle II is collecting collision data steadily and our computing group is processing those data
- Belle II computing group consists of two sub groups, “core computing” and distributed computing and both are centrally managed by DIRAC and Rucio
- Current status is good, and we are adapting new technologies, token based authentication and DiracX

backup
