Kubernetes at Belle II: Orchestrating the Computing Infrastructure for Particle Physics

Hideki Miyake (KEK)
Dhiraj Kalita (KEK)

KEK-CSCS Joint Meeting, 2025 Oct 29th

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

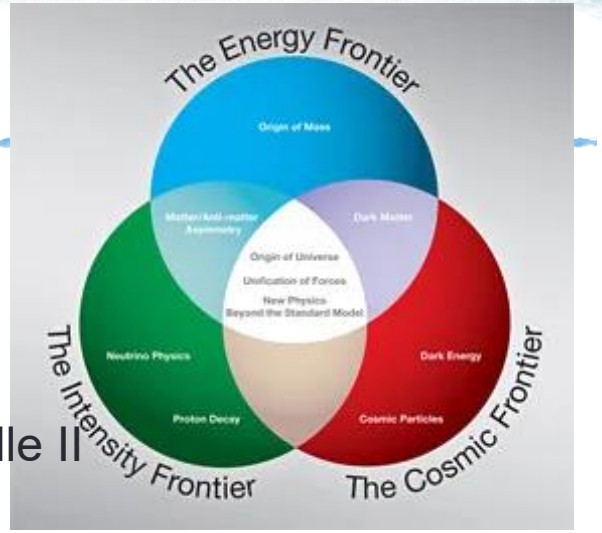

- First data: 2019
- Finished long shut down in 2024
- Start 2025 data taking soon

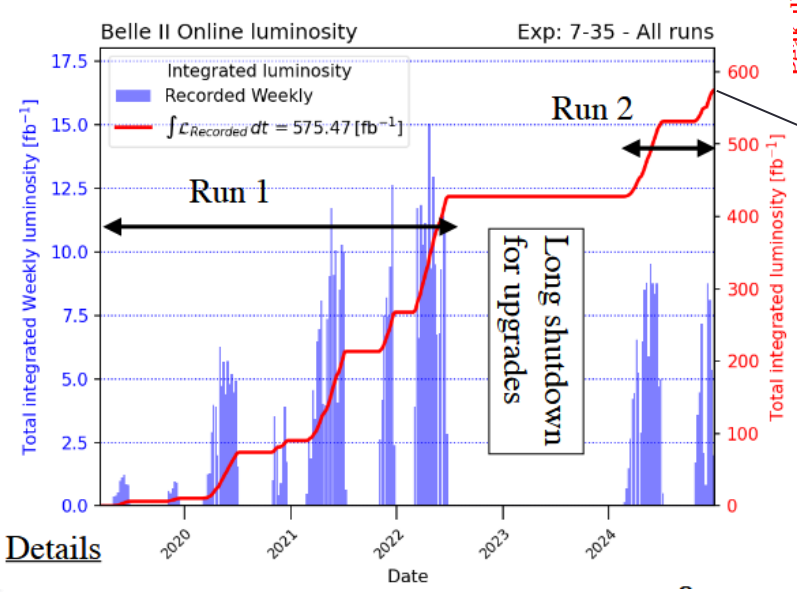

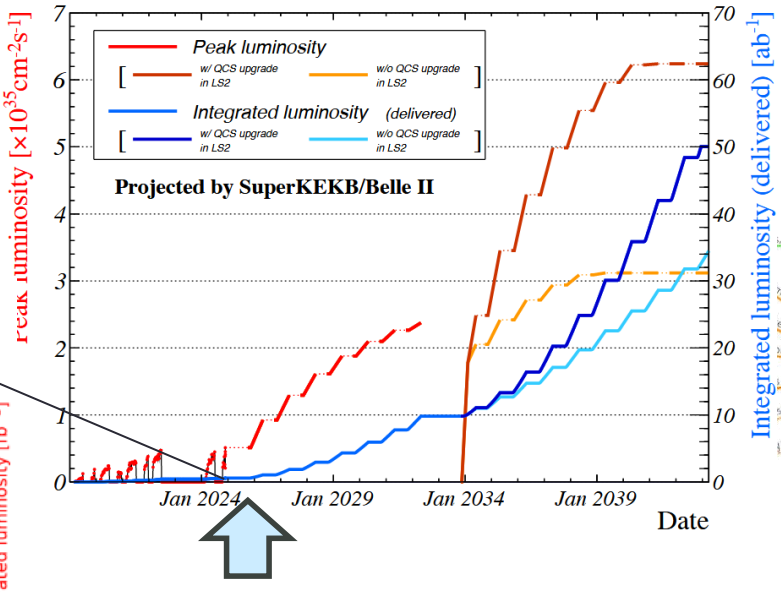

We're here

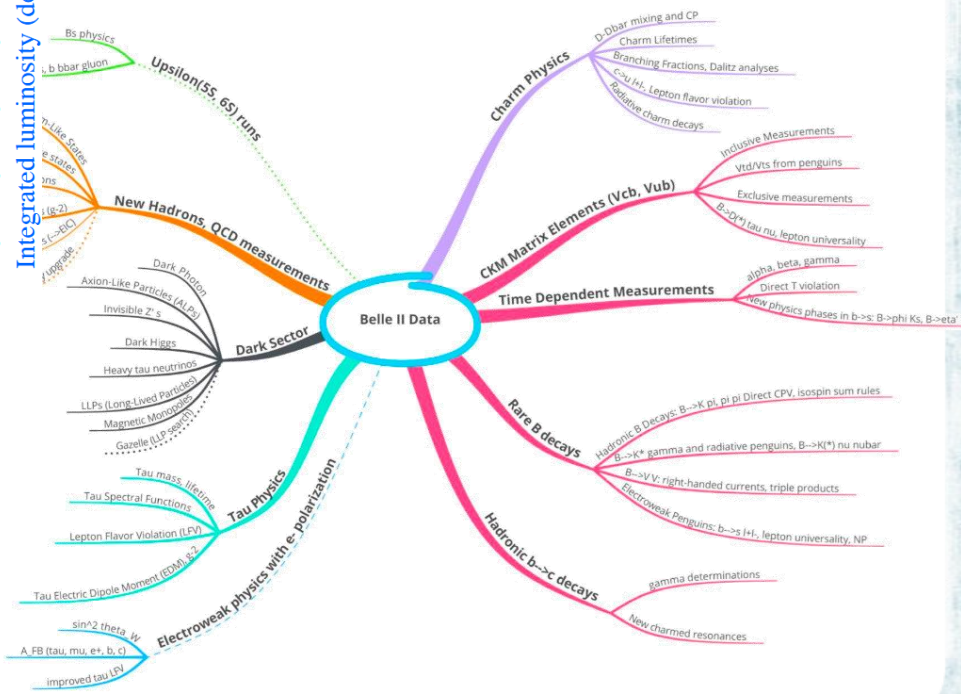

SuperKEKB + Belle II

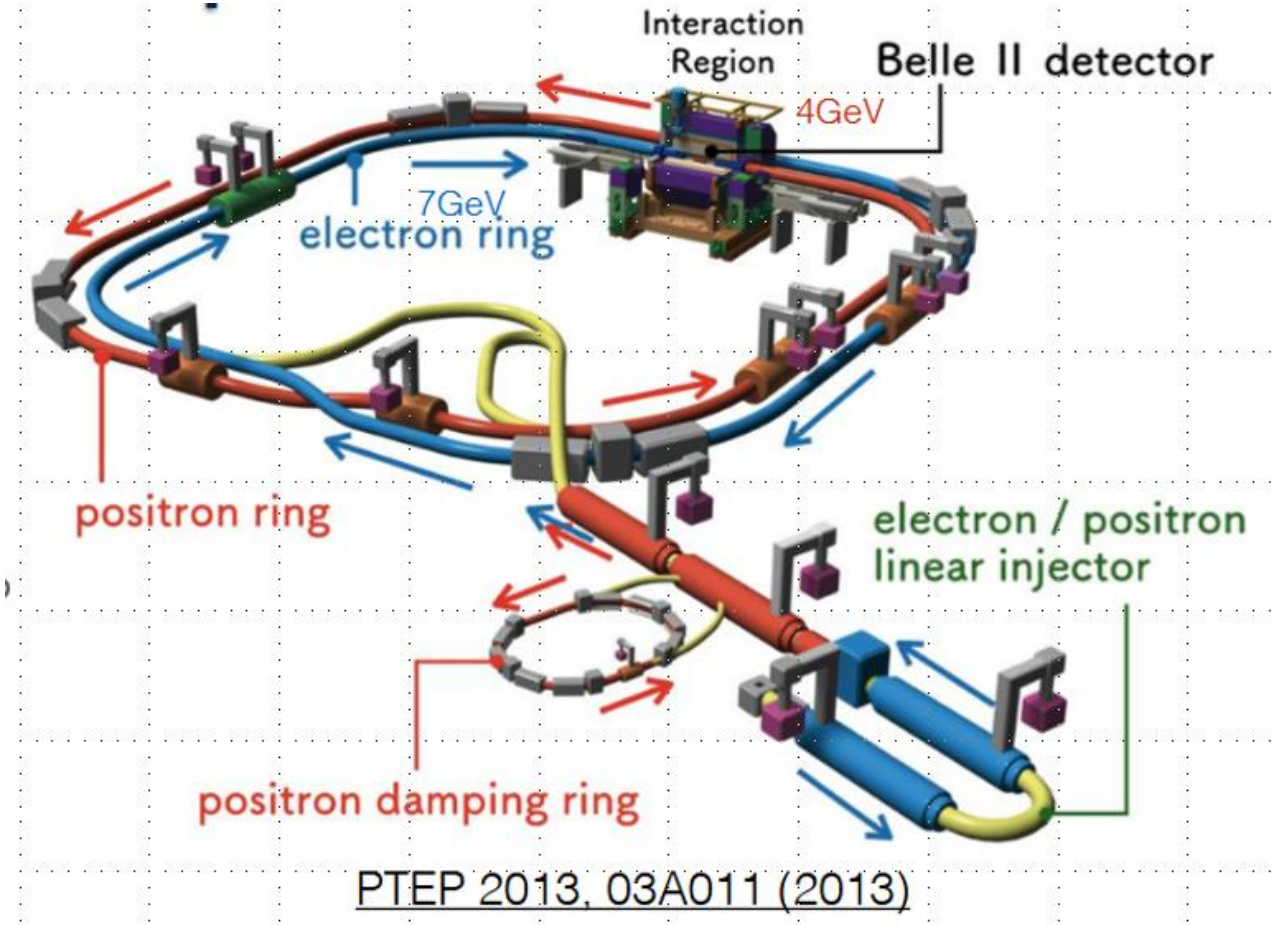

PTEP 2013, 03A011 (2013)

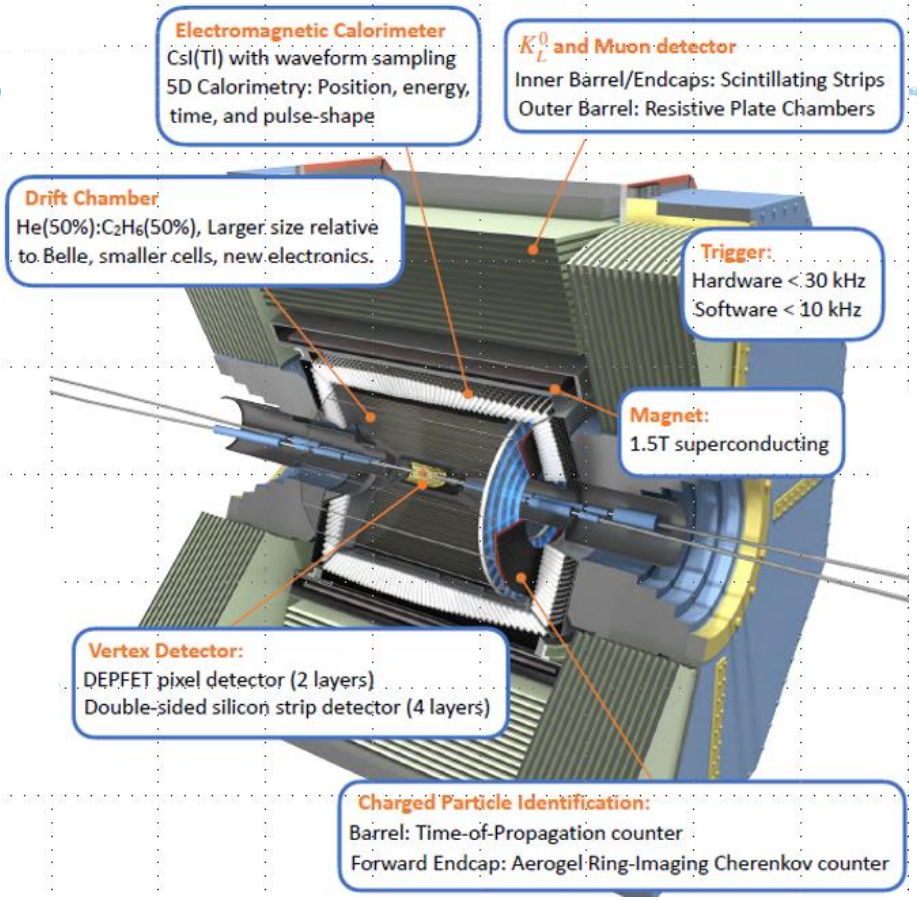

Asymmetric e^+e^- collider at KEK (e^+:4GeV, e^-: 7GeV)
Instant luminosity record of $5.1 \times 10^{34} cm^{-2}s^{-1}$ (x2.4 of KEKB)

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

Belle II collaboration

Belle II collaboration

High-speed network mesh is a key!

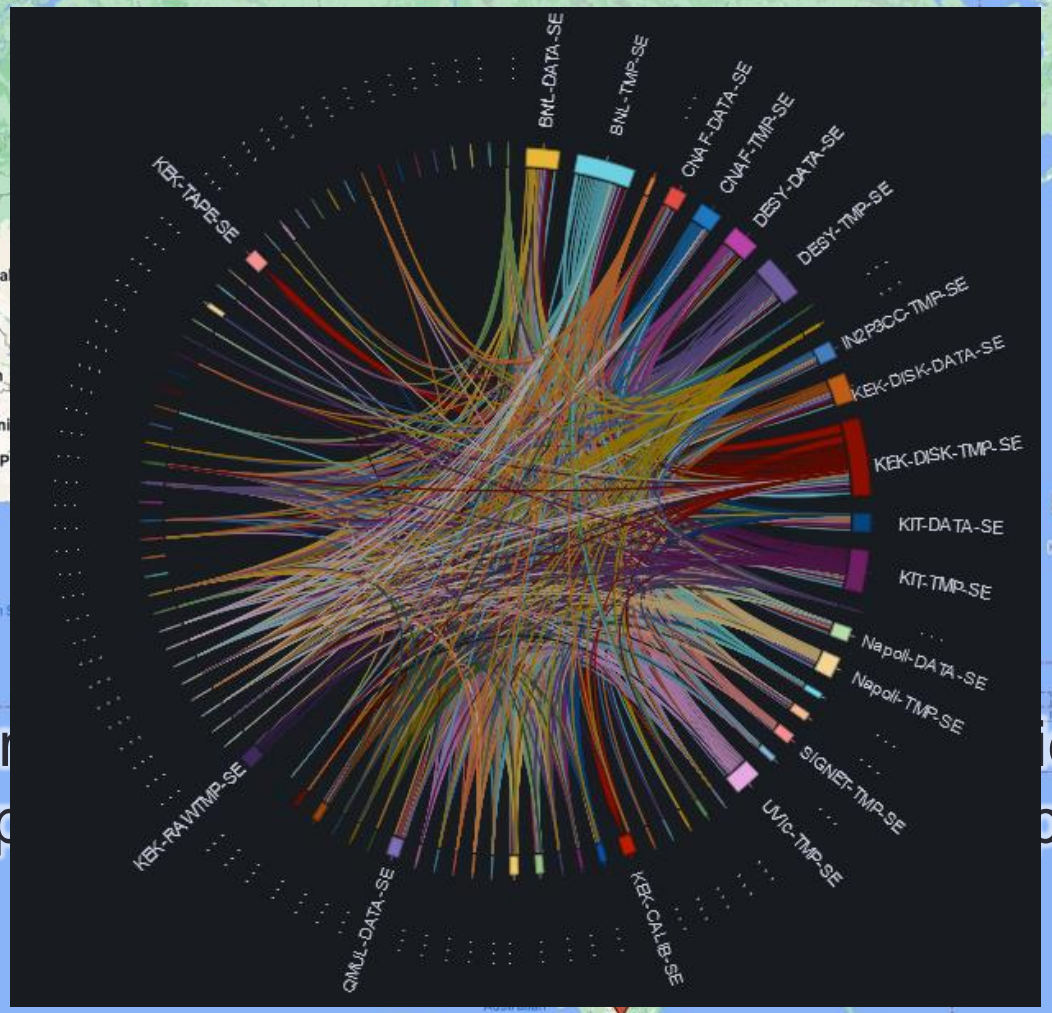

1240 mer ies and regions

Widely sp pe, 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.

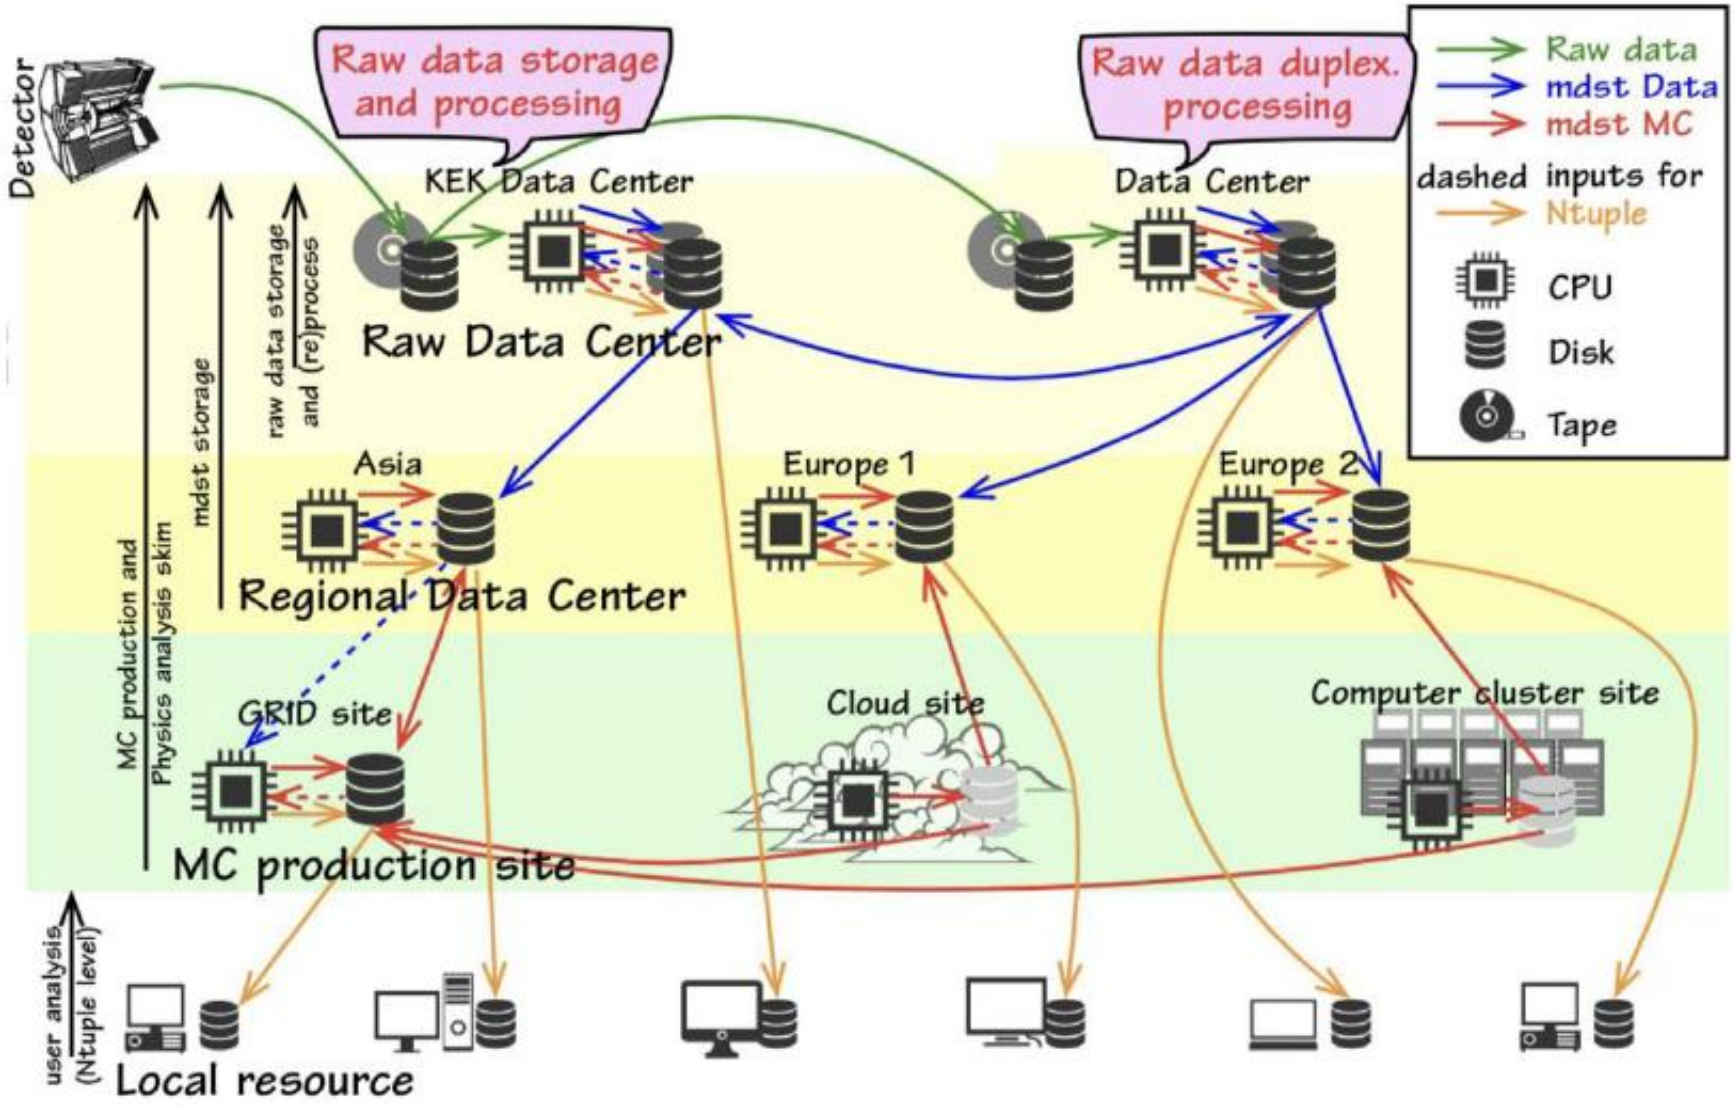

DIRAC

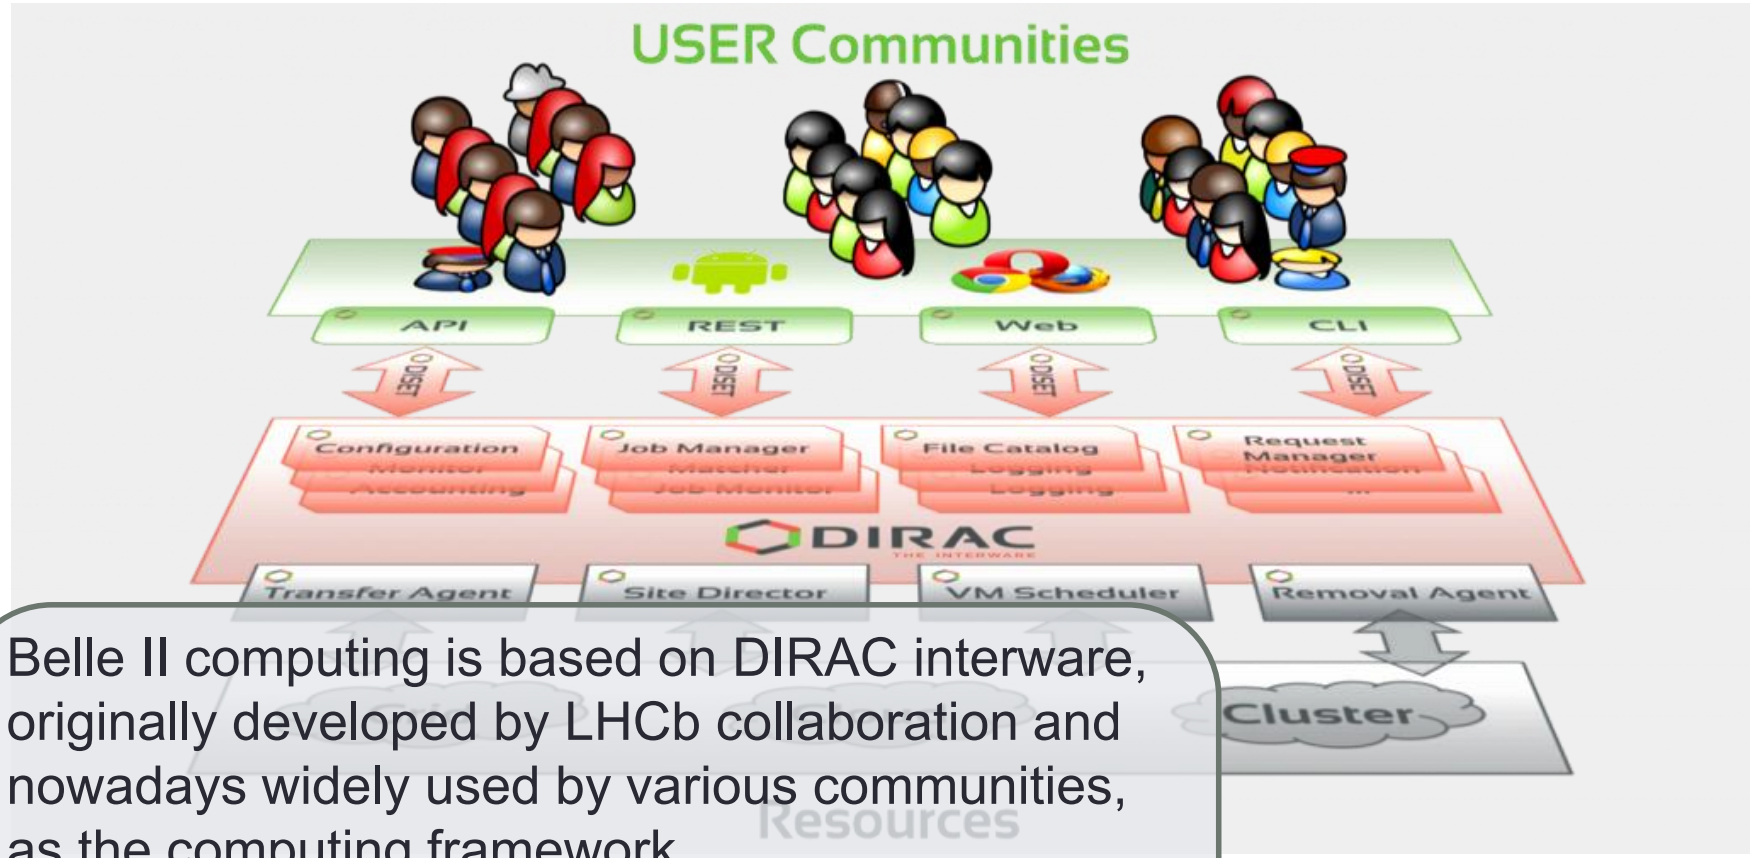

Belle II computing is based on DIRAC interware, originally developed by LHCb collaboration and nowadays widely used by various communities, as the computing framework

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

Similar framework used in other communities

PanDA: ATLAS
JAliEn: ALICE
....

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
      - Provide coherent access to file replicas via LFN (logical file name)
    - As Metadata Catalog
      - We have long used AMGA to store “data of data”
      - Now in final phase to migrate to Rucio

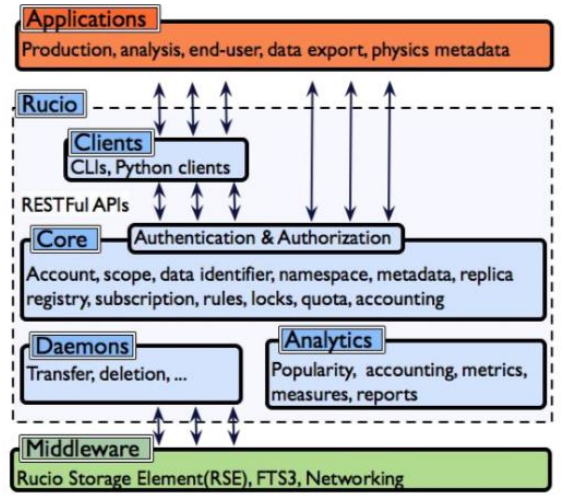

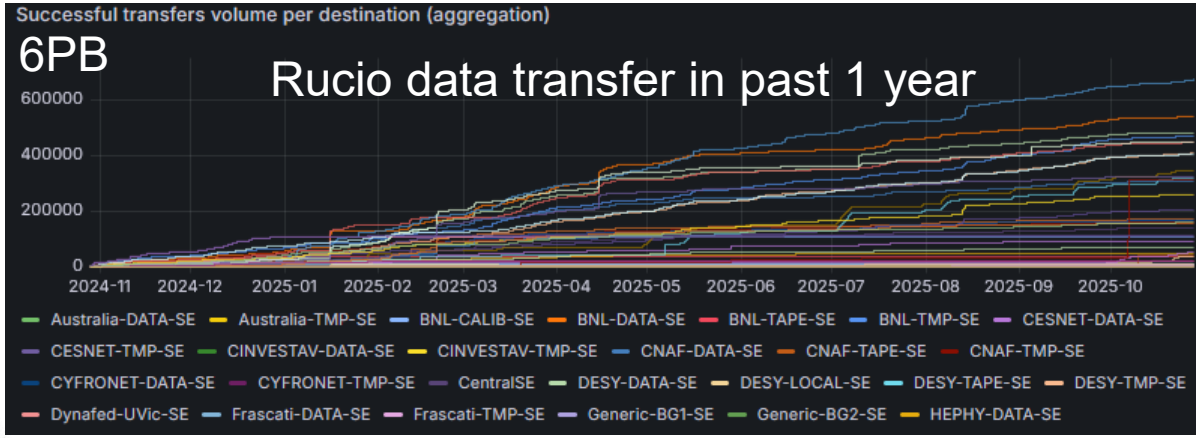

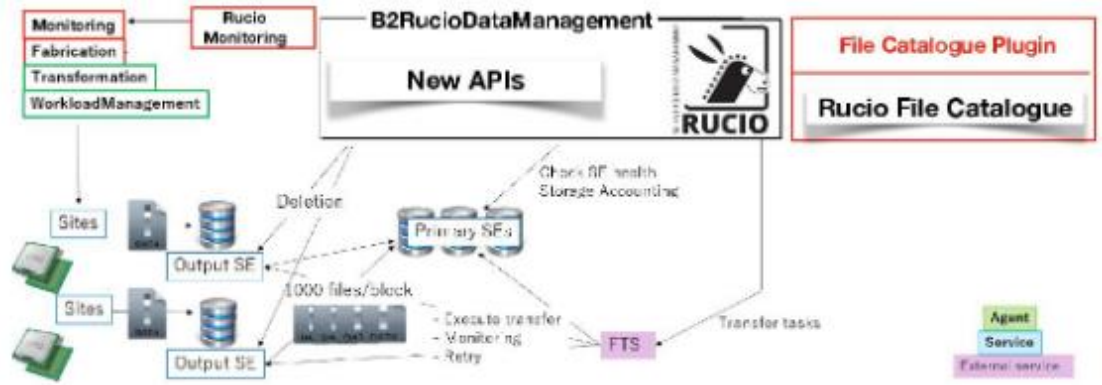

Raw data transfer in Belle II

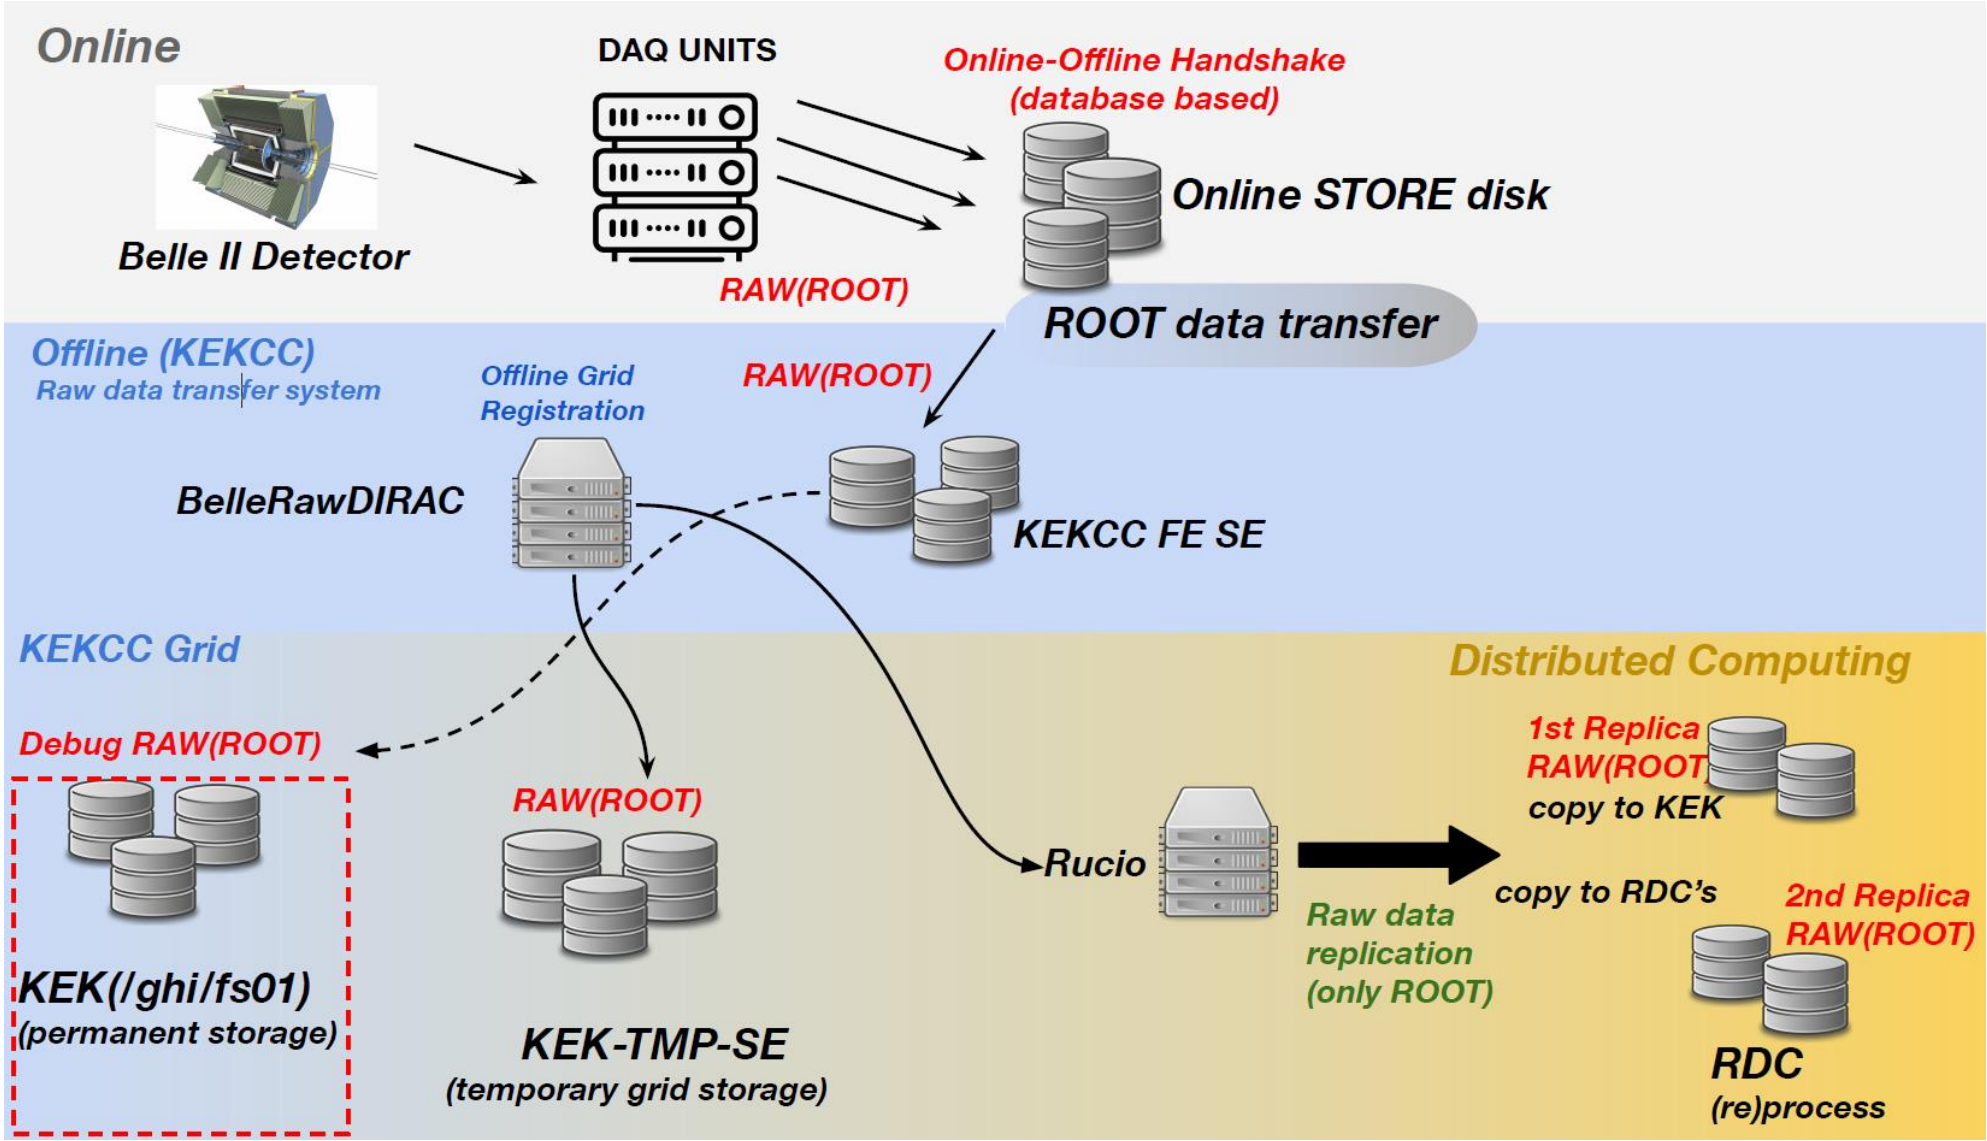

Belle II “core computing” system

Most servers are hosted by KEKCC

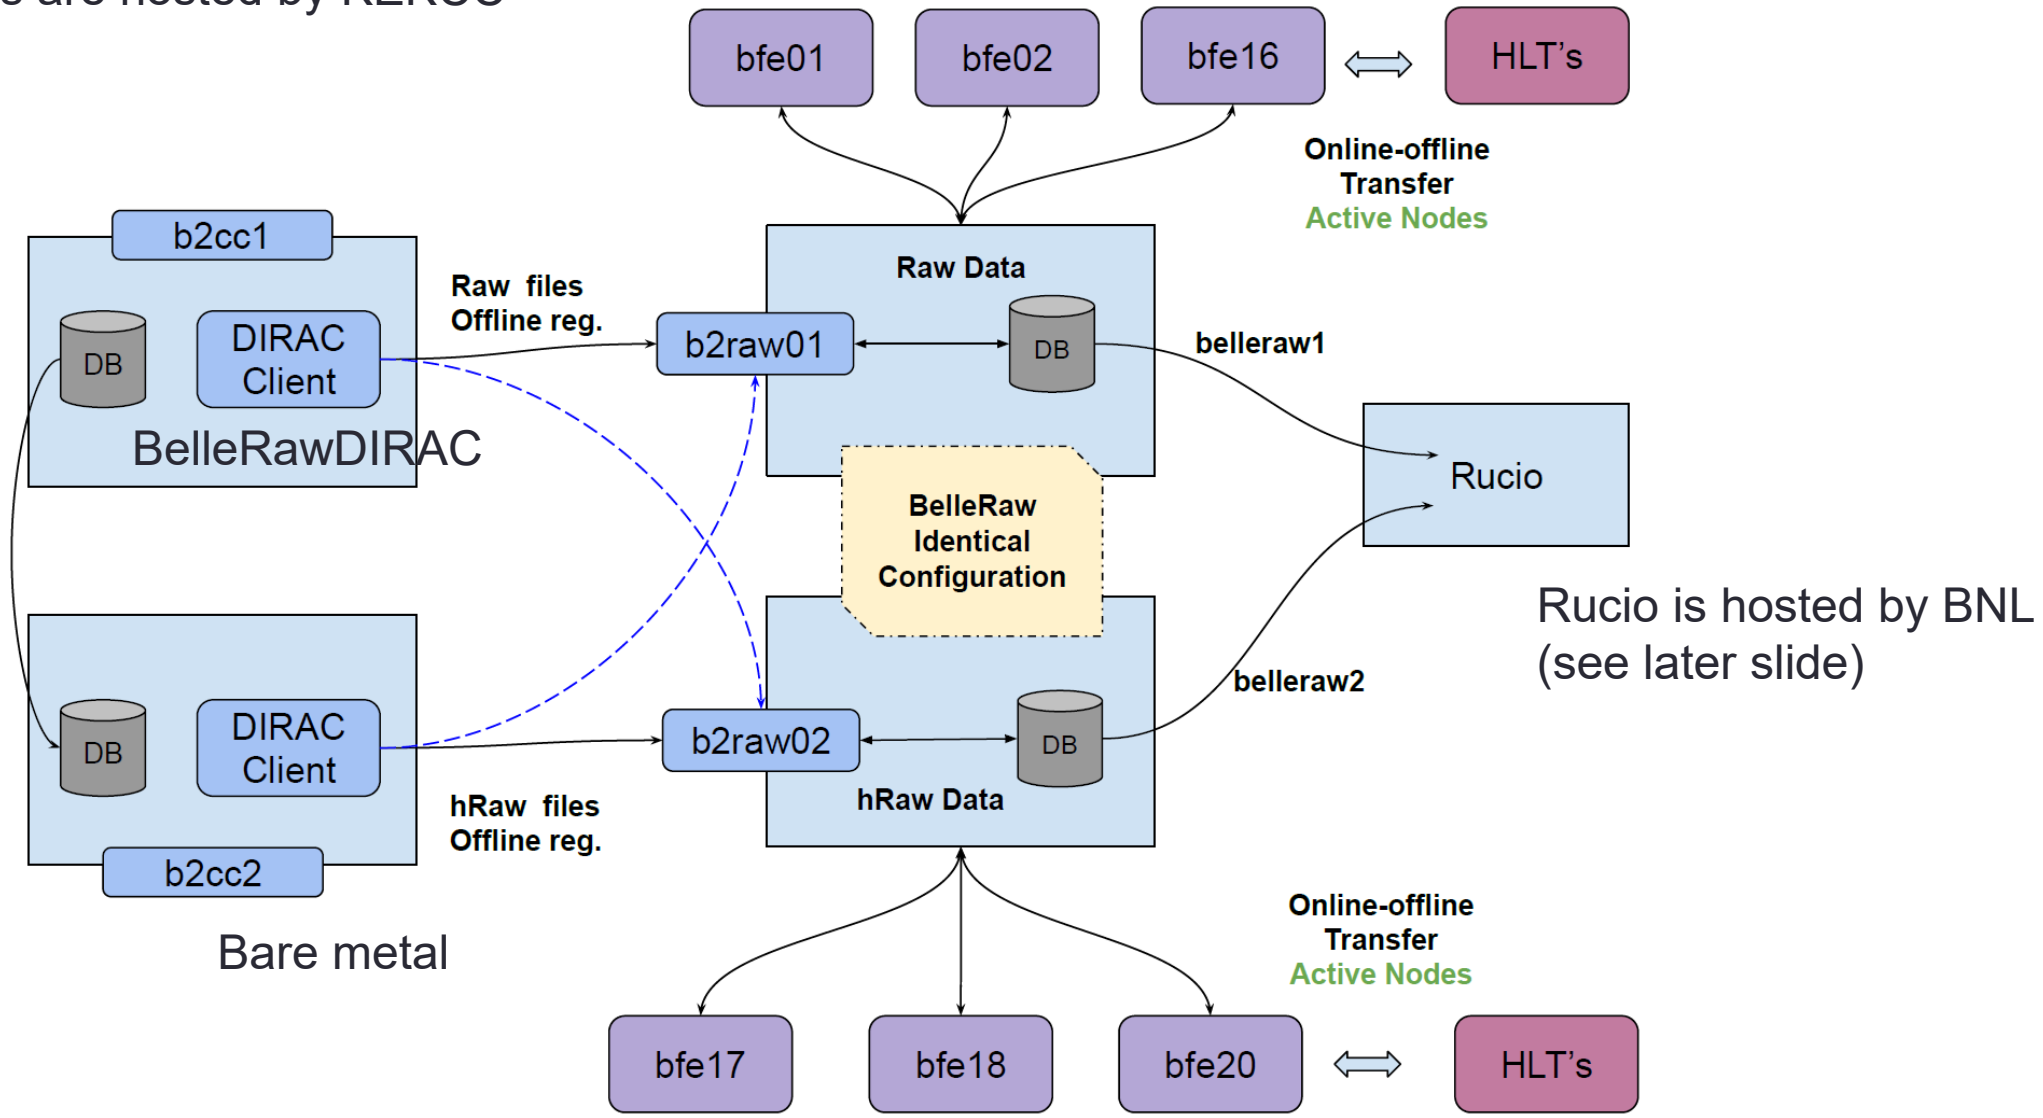

Rucio is hosted by BNL
(see later slide)

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)

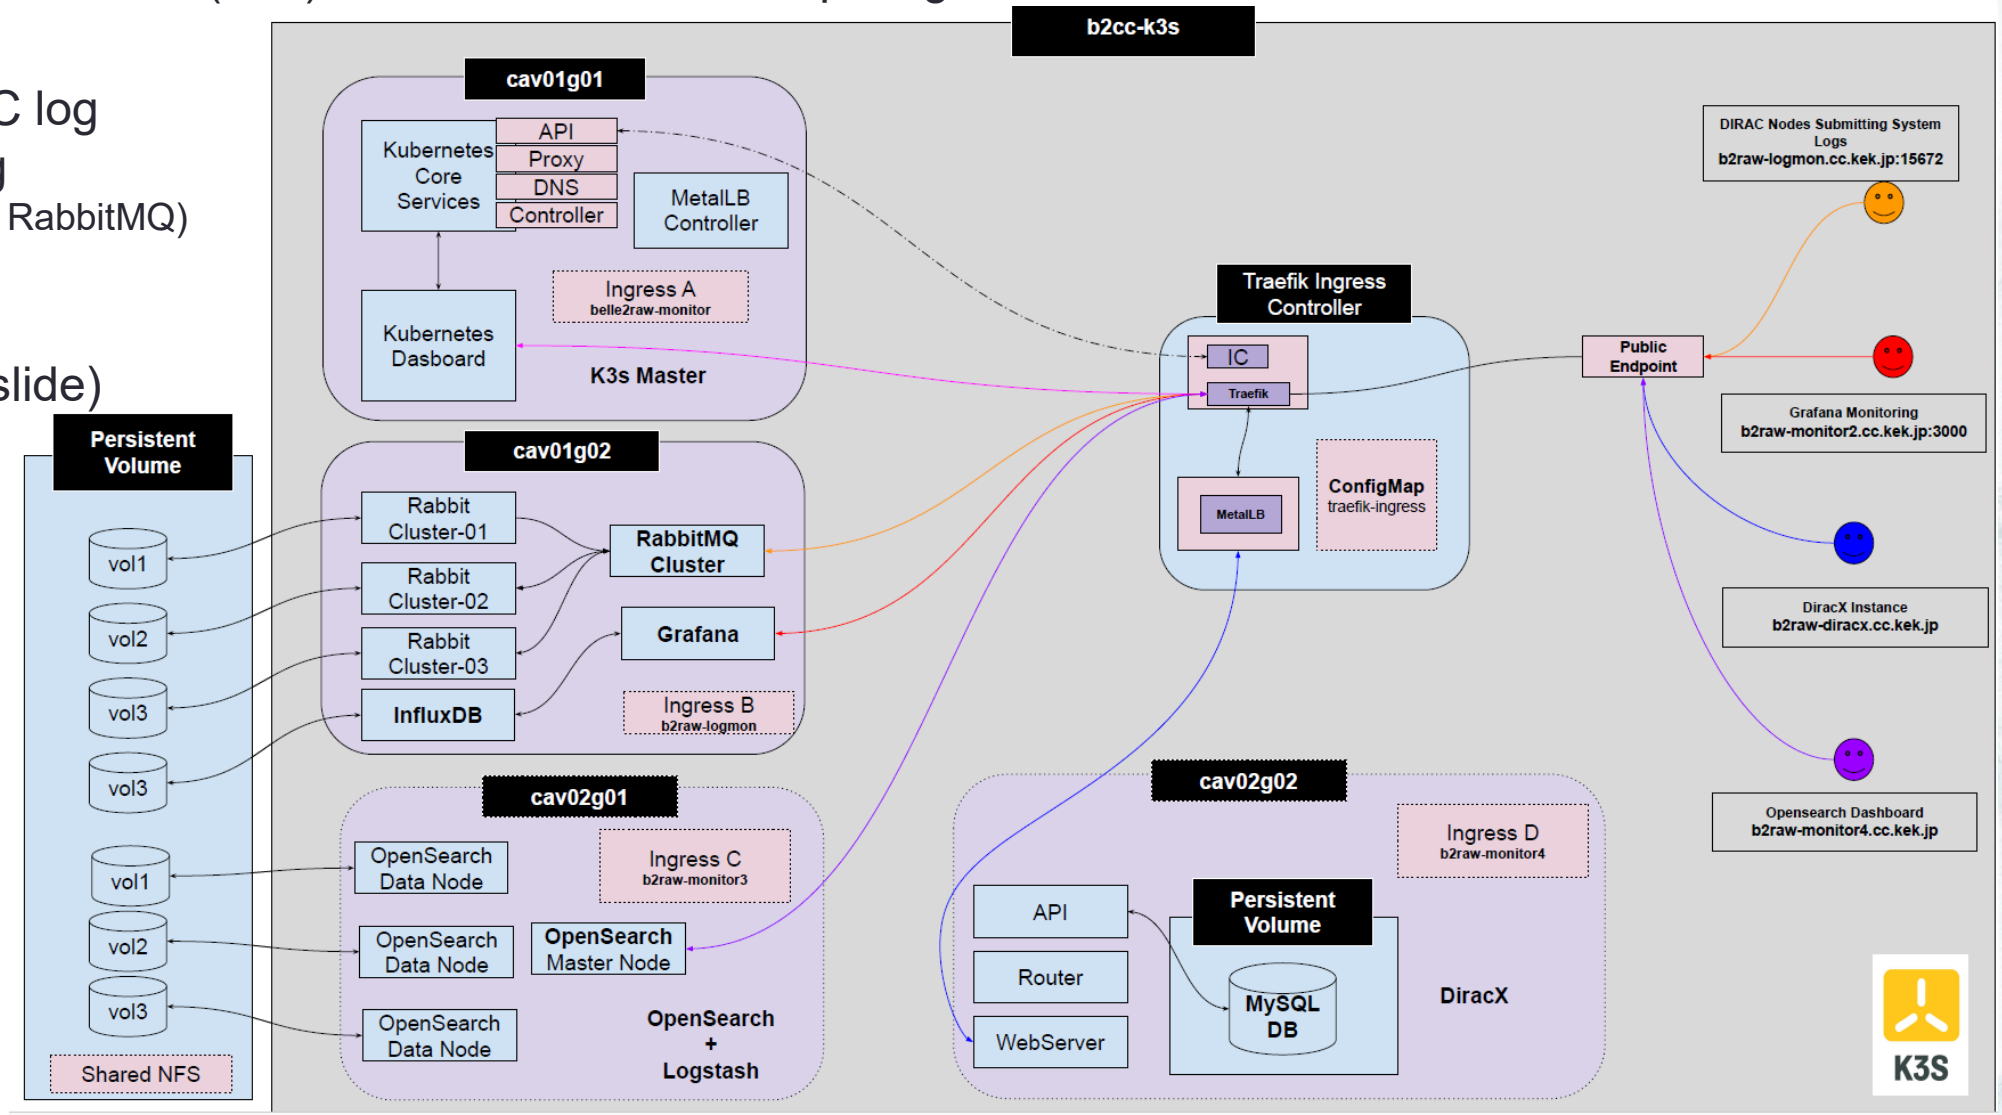

Distributed computing in Belle II

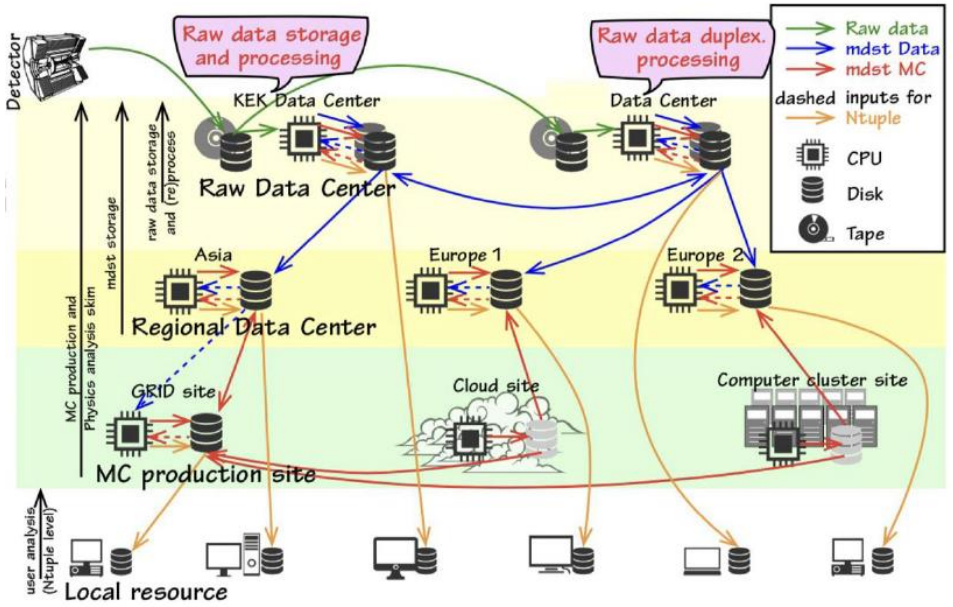

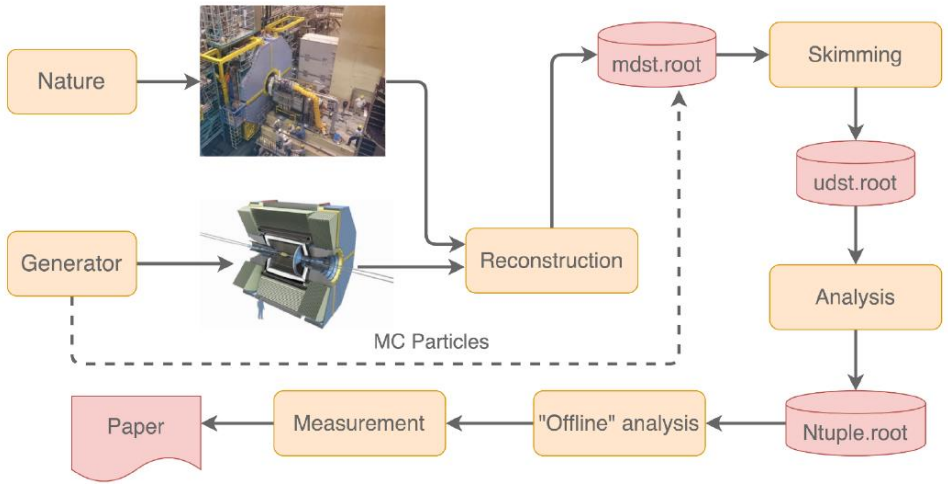

After Raw data is transferred to Raw data centers, apply various processing and filtering (“Skimming”) toward final analysis result

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

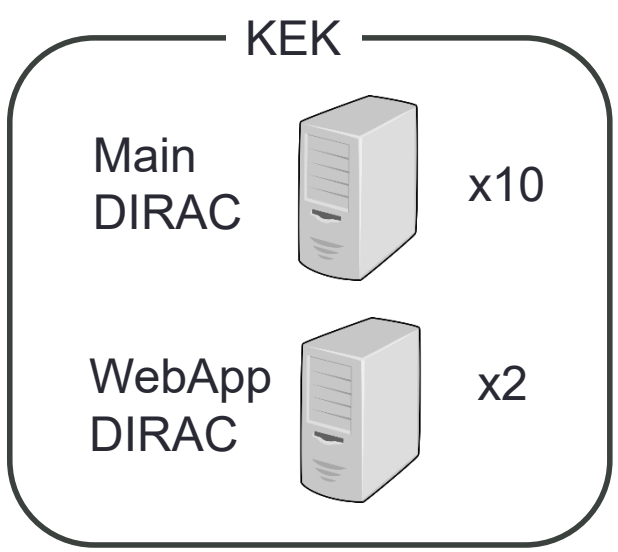

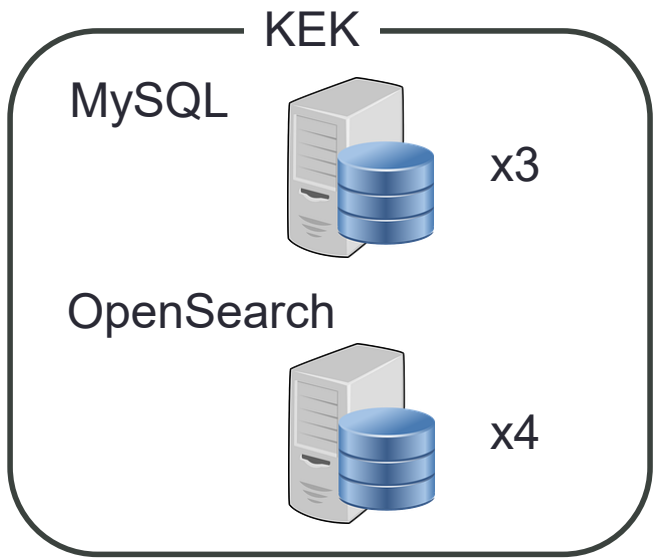

BNL

UVic

NDU

(for SSH)

KEK DIRAC servers are on VM
MySQL and OS run on physical machine
No Kubernetes yet

DIRAC servers are maintained by Belle II
KEKCC maintains other essential computing resources hosted by KEK
BNL hosts FTS, Rucio, and relevant services

BNL
- Rucio
- FTS

KEKCC
- VOMS
- IAM
- FTS
- AMGA
- CVMFS

Monitoring
- Grafana+VictoriaMetrics at KEK
- Grafana+InfluxDB+Elasticsearch at BNL

Belle II grid job status

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

Operation with DIRAC is successful, but…

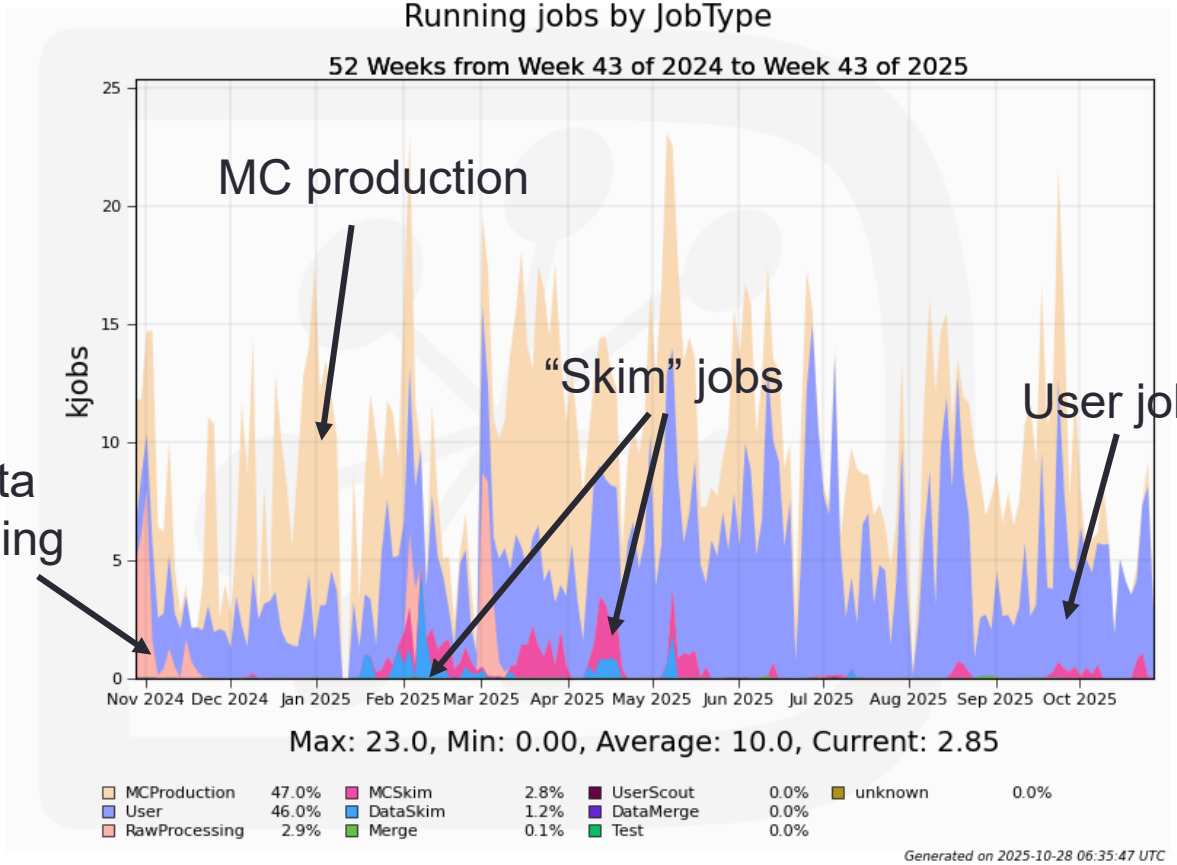

Belle II Rucio

- 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

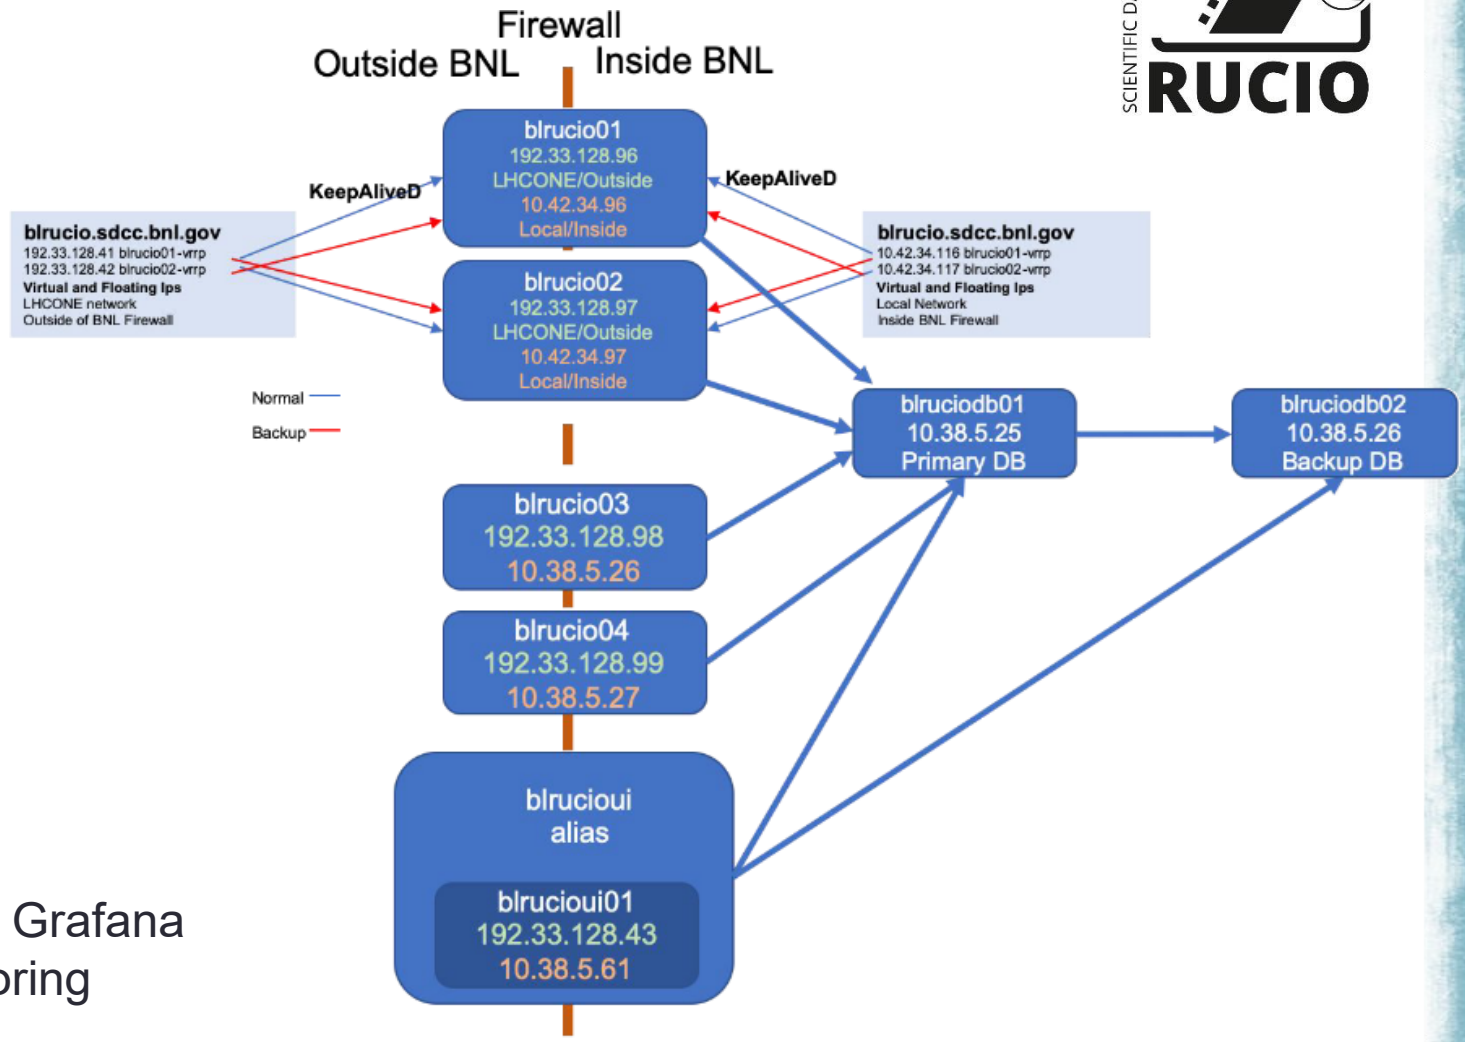

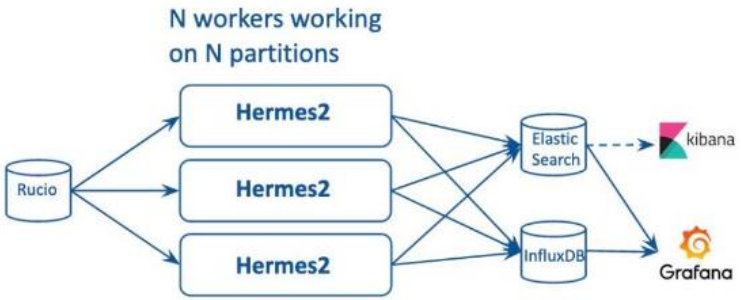

Belle II monitoring stack

Providing Grafana for monitoring

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

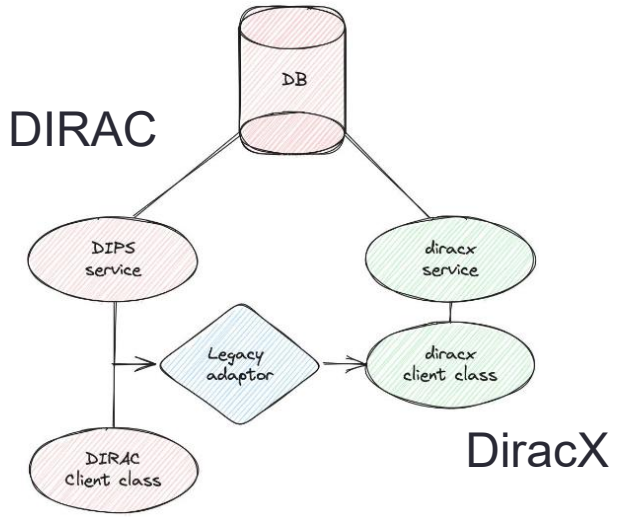

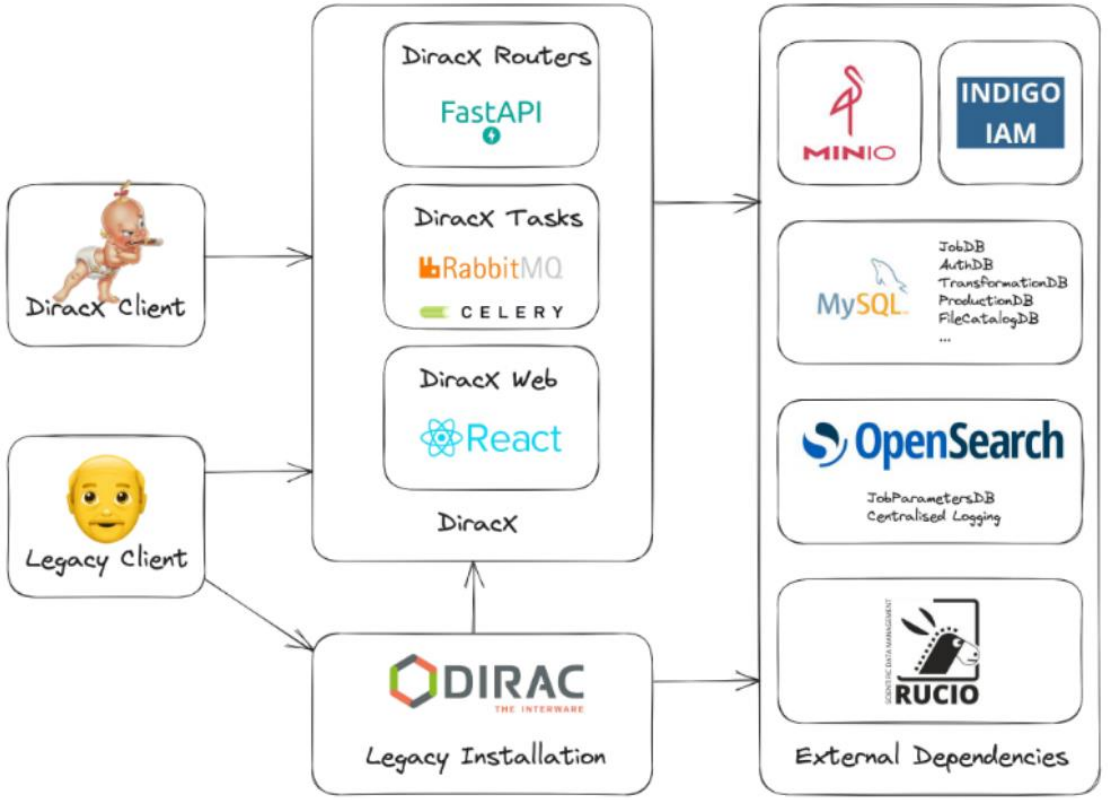

Current status

- LHCb already deployed DiracX in production (v0.1.0) with DIRAC v9 in this year
- Belle II core computing (BelleRawDIRAC) also deployed DIRAC v9 (bare metal) with DiracX on test K3s
- Belle II distributed computing (BelleDIRAC) is still running DIRAC v8 and no DiracX
- We plan to deploy “production K3s”

DiracX on test K3s

BelleRawDIRAC and BelleRawDiracX can access each other

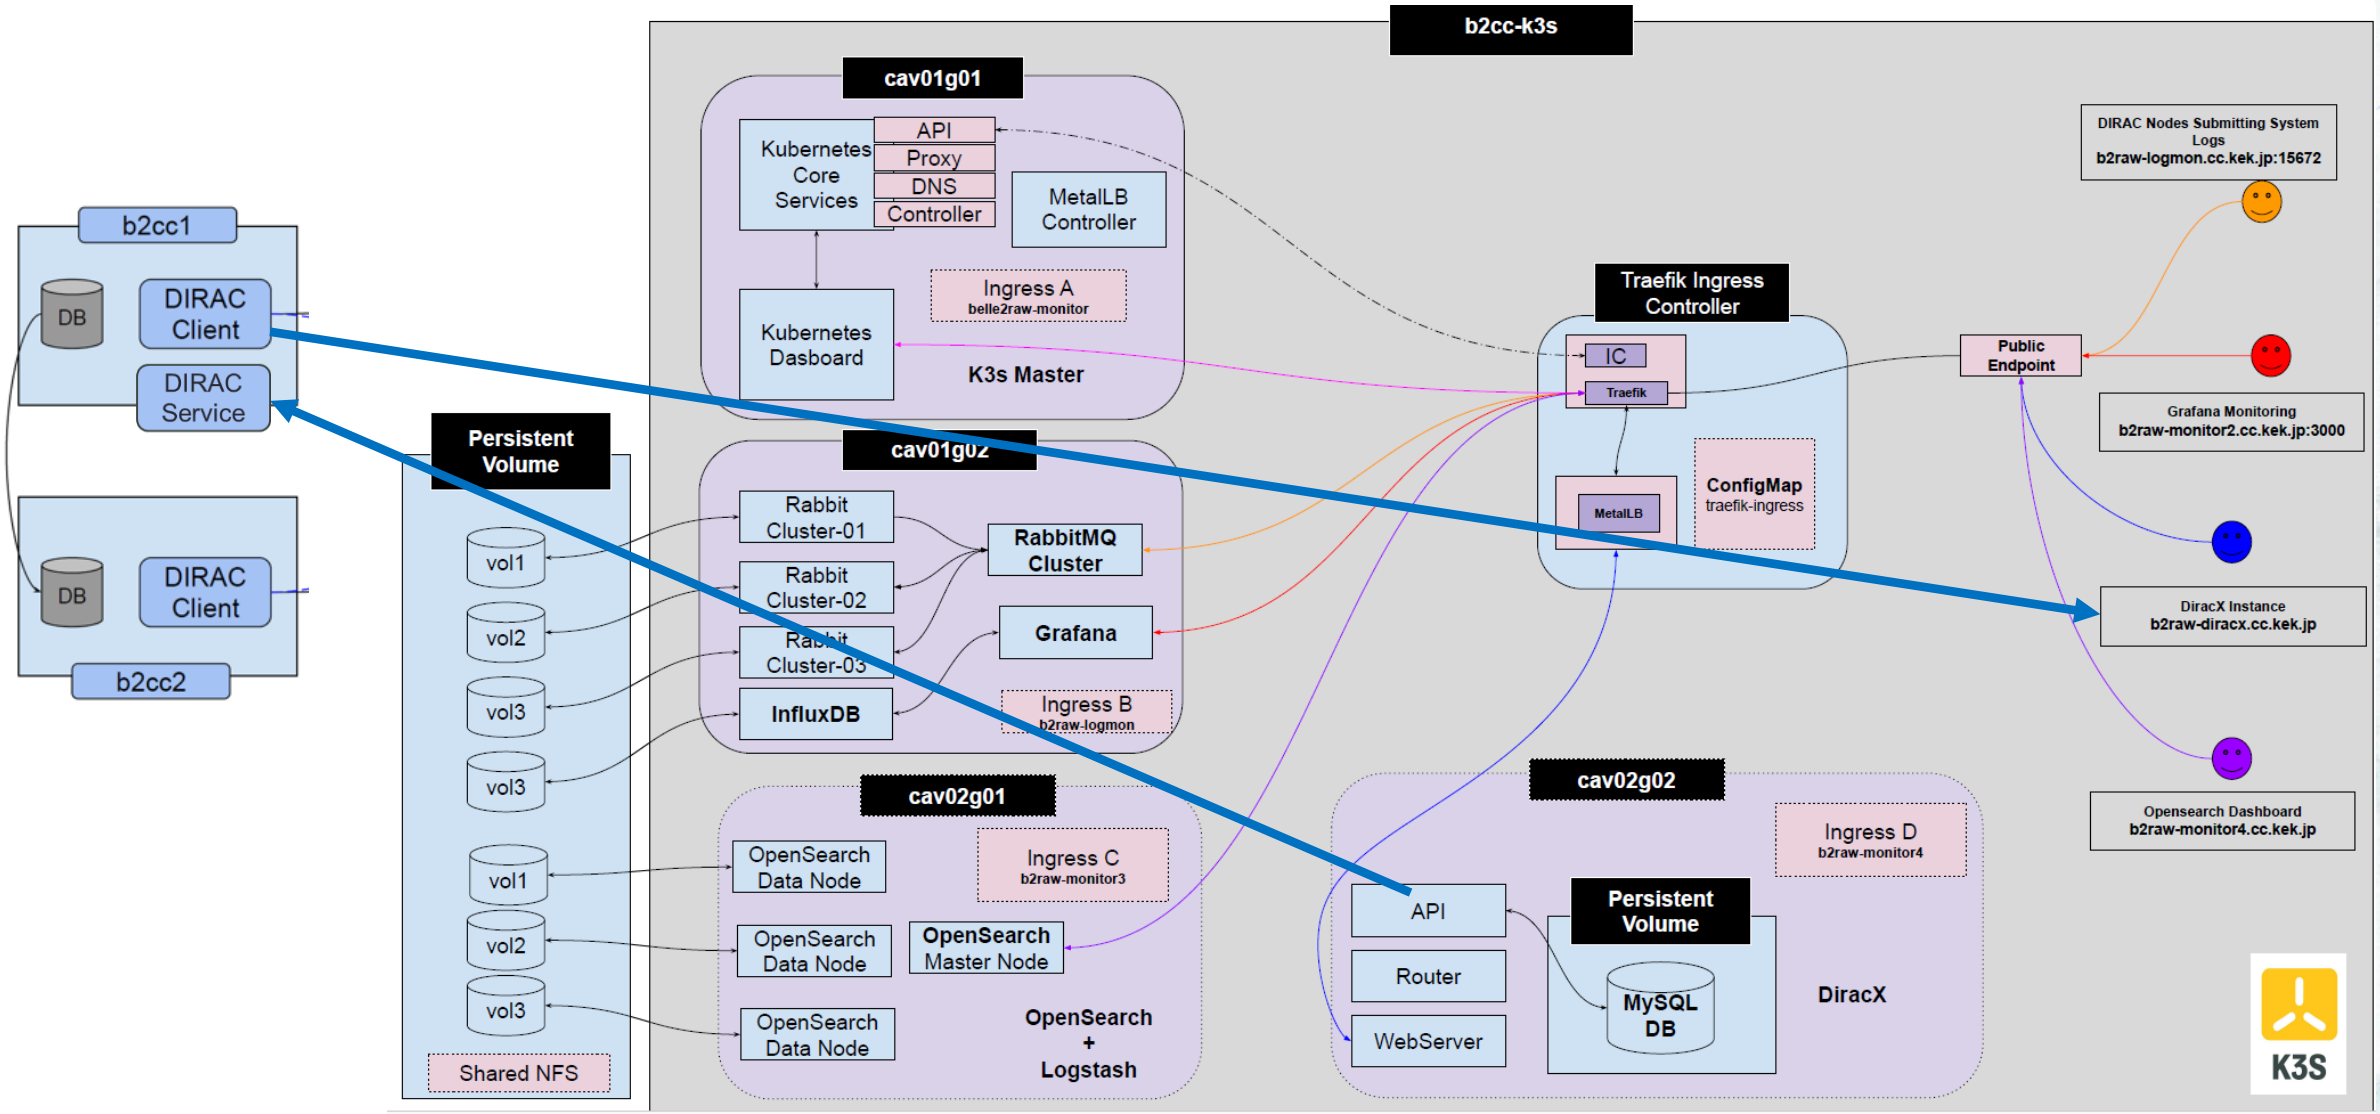

Production K3s: plan

- We share K3s infrastructure between BelleDIRAC and BelleRawDIRAC

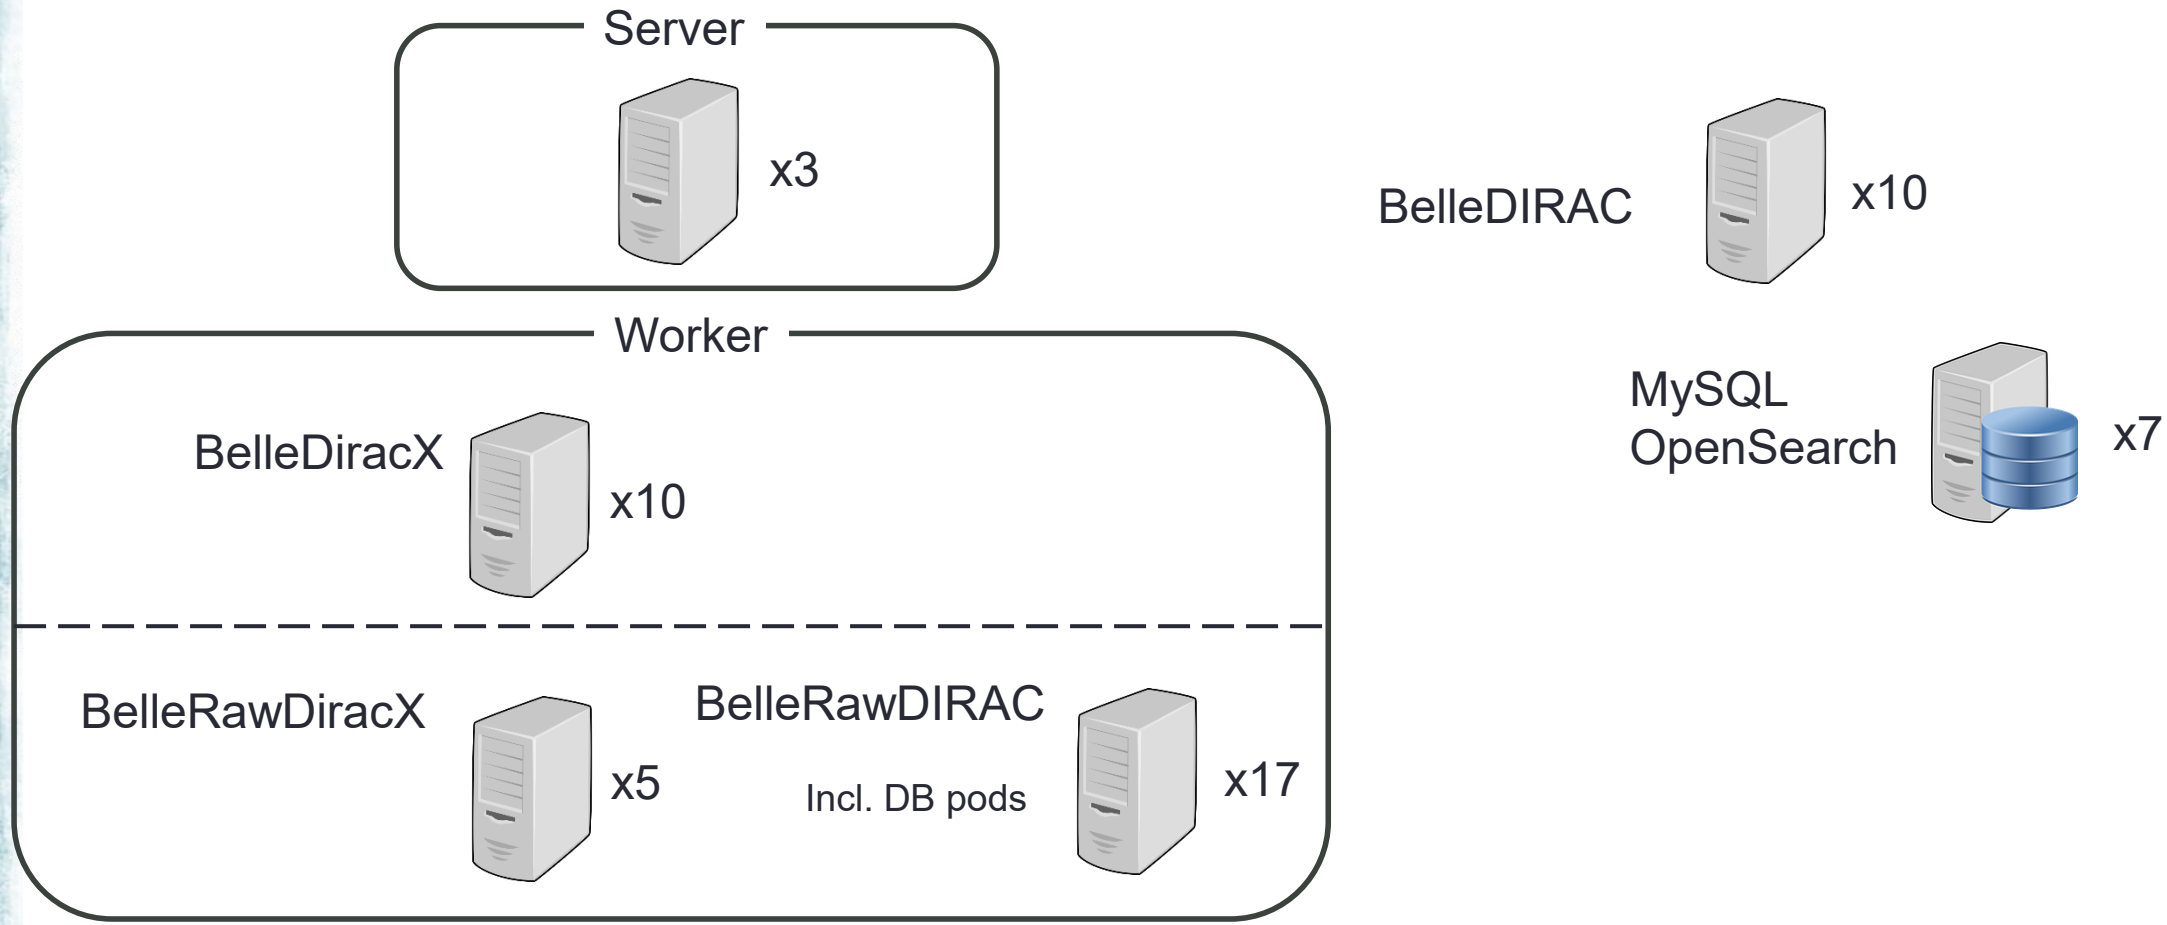

BelleDiracX and BelleRawDiracX are deployed in different namespace

Some issues for production K3s

- Since we don’t have much experience for Kubernetes, we have spent time to decide some technology choices
- Although we are finalizing our production K3s design, still we are not confident if our design is correct
- We would appreciate it if you could give us any advices

- Our issues
  - Load balancer
  - Stateless operation (→persistent storage)

K3s: load balancer

- Choice of load balancer depends on each environment
  - In our case, load balancer is a part of experiment-specific infrastructure
- We initially prefer MetalLB (or Kube-VIP) L2 mode
  - Require no additional hardware and special network configuration

- MetalLB (no HA on control plane)
  - Configuration is rather simple

Concern in L2 mode: Is our Worker tolerable with massive data access?

Now under evaluation

- Kube-VIP (allow HA on control plane)
  - Configuration is complicated
  - Newer project: potential bug, less documentation…

Once we discussed BGP mode solution… Is this difficult option?

- Mixed (Kube-VIP for Server and MetalLB for Worker)
  - Should be careful with misconfiguration

K3s: issues against stateless operation

- Some different approaches by BelleDiracX and BelleRawDiracX

- BelleDiracX
  - Run DB outside of K3s cluster
  - Need to access external storage (“sandbox”) (see next slide)

- BelleRawDiracX
  - Run DB inside K3s with persistent volume (volume is shared among Worker by NFS)
  - Need to access GPFS

Is LongHorn good option? (for non-expert operation)

K3s: sandbox

- Yet another persistent volume for BelleDIRAC
- Store input and output files of payload job and pilot
- Allocated volume size: 5TB
- Plan to operate MinIO (on VM, not in K3s)

Recently I'm getting worried
if MinIO is good choice…

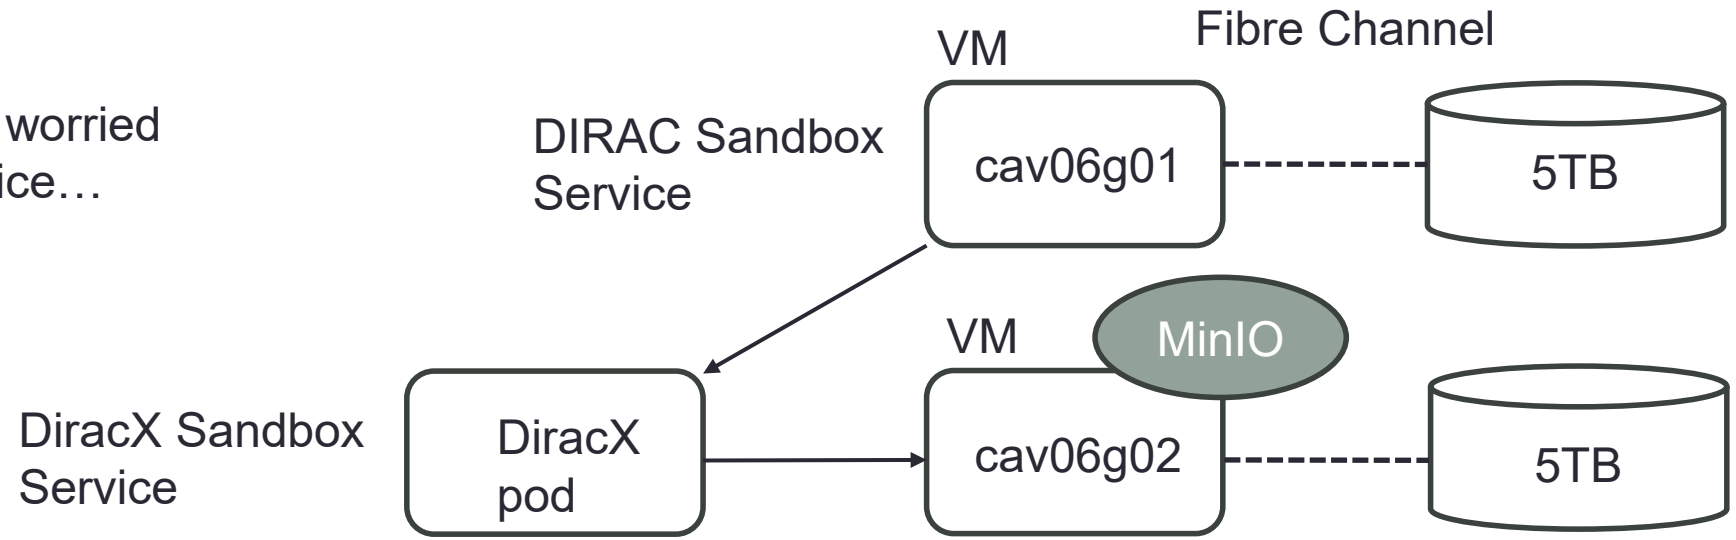

Plan, perspective

- Deploy production K3s in coming months
- BelleRawDIRAC (core computing)
  - Move DIRAC v9 and BelleRawDiracX to production K3s
- BelleDIRAC (distributed computing)
  - Deploy BelleDiracX (isolated) in production K3s
  - Deploy DIRAC v9 for BelleDIRAC (outside of K3s) and BelleDiracX on production K3s
- DIRAC functionalities are being replaced by DiracX
  - Ultimately, everything in DIRAC will be transferred to DiracX on Kubernetes

backup

Planned VMs to use