Belle II report

Cédric Serfon (BNL)
on behalf of the Belle II computing team

September 26, 2025

@BrookhavenLab

Introduction

- SuperKEKB
  - e+e− collider with energies 4 GeV and 7 GeV operating around Y 4S resonance.
  - Achieved world-record peak (December 2024)
  - luminosity of L = 4.7 × $10^{34} cm^{-2} s^{-1}$

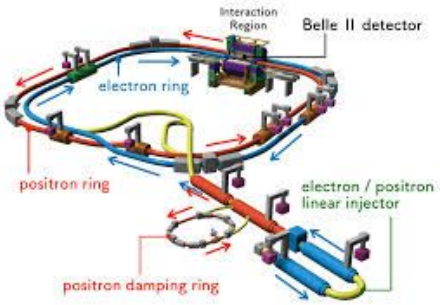

- Belle II
  - A general purpose particle detector with almost full solid angle coverage
  - Tracking, PID and photon reconstruction capabilities

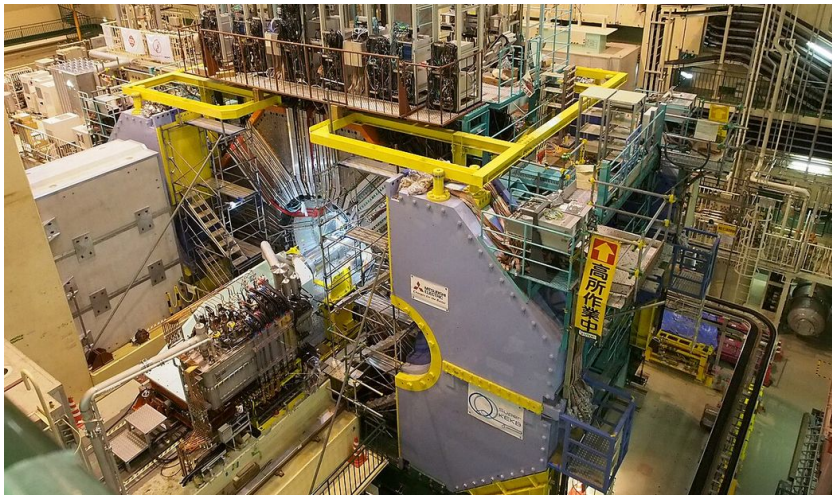

Brookhaven National Laboratory

Belle II collaboration

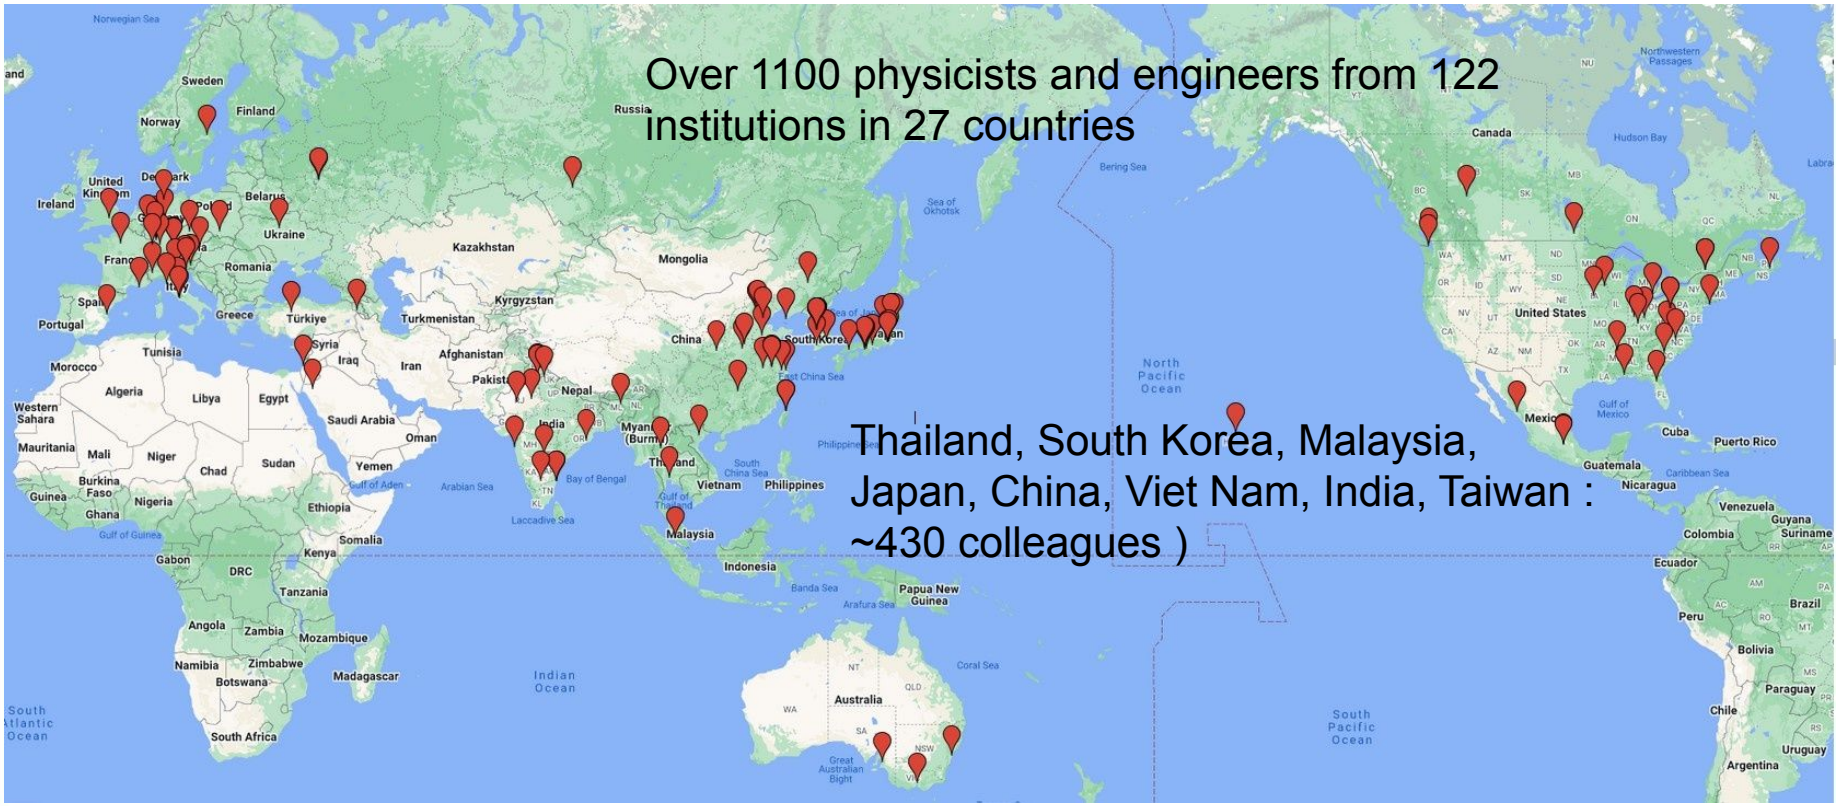

Brookhaven National Laboratory

Belle II status

- Belle II started its run 2 in February 2024 after a long shutdown (July 2022 - Jan. 2024) when many parts of the detector and accelerator were upgraded

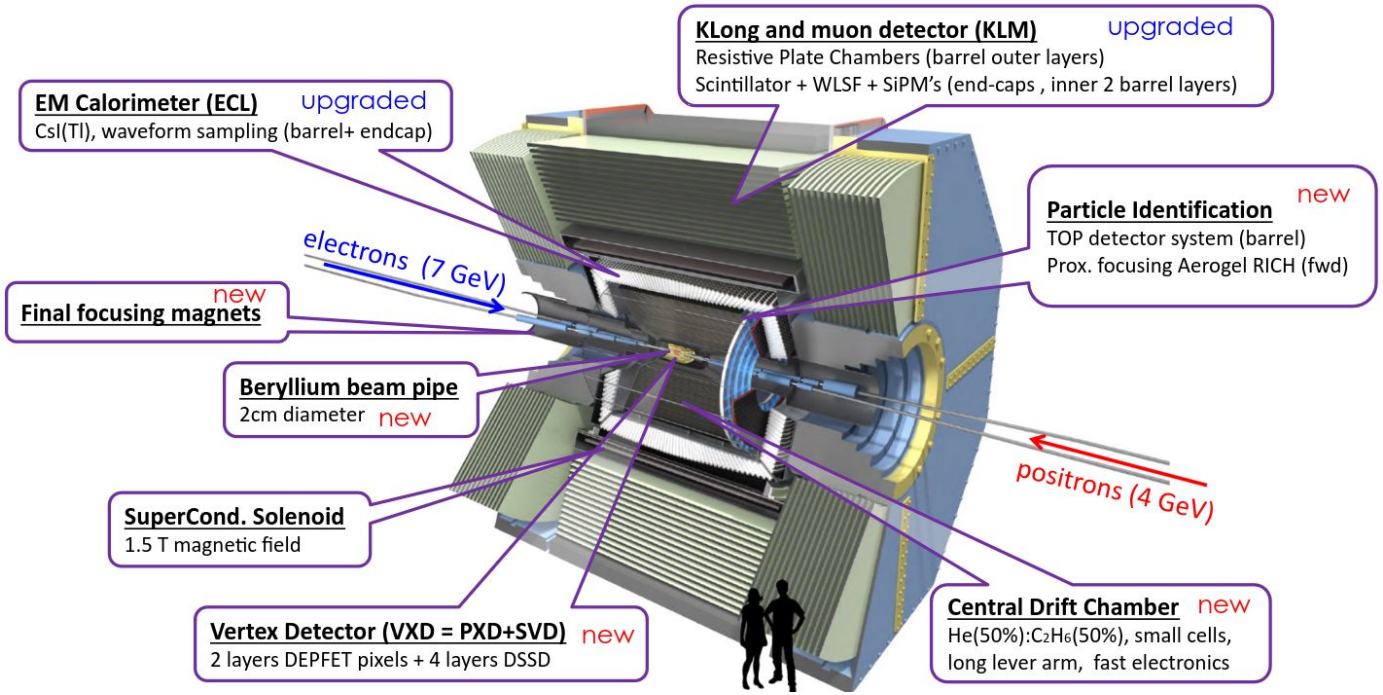

Brookhaven National Laboratory

Luminosity since the start of run 2

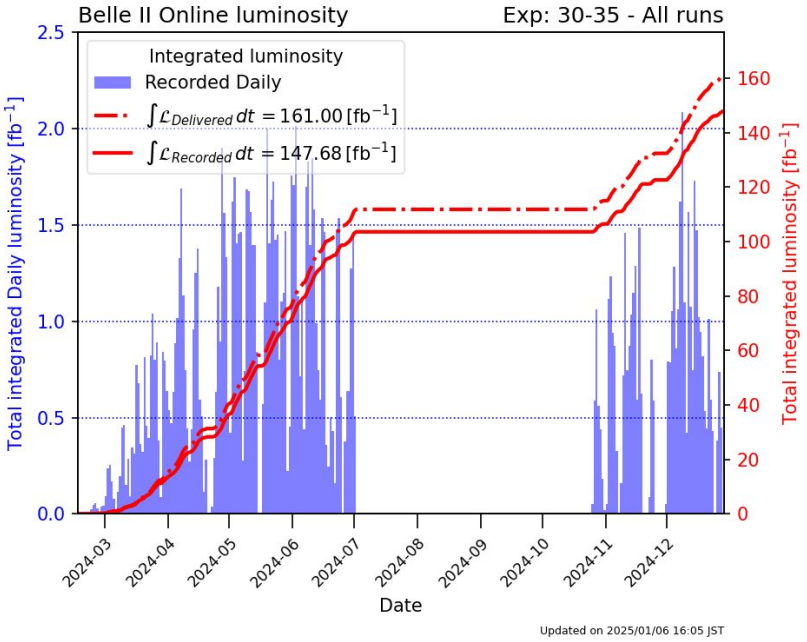

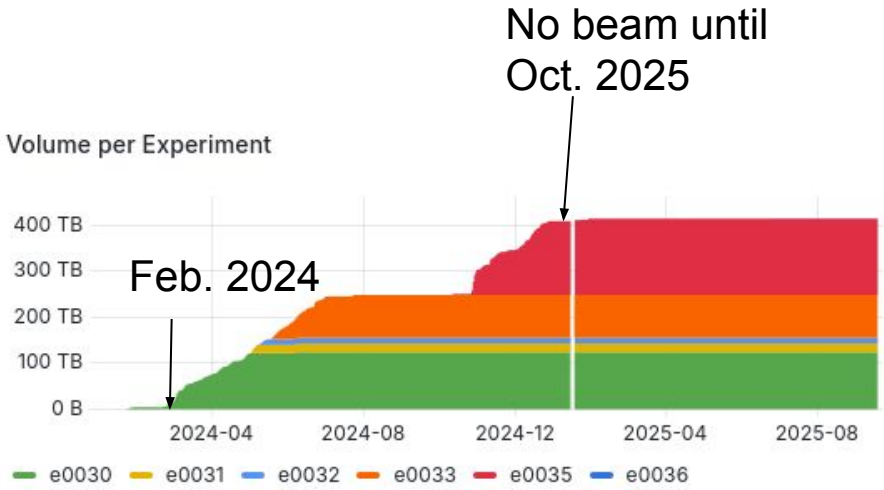

- Run 2 started in Feb. 2024. No beam since Dec. 2024. Data taking will resume next month
- About 150 fb^{-1} (~400 TB) accumulated since the start of run 2

Brookhaven National Laboratory

Belle II computing

- Belle II uses a distributed computing infrastructure (grid) based on set of standard tools developed and used by the WLCG collaborations
  - DIRAC : Workload Management
  - Rucio : Distributed Data Management
  - FTS : File transfers
  - VOMS : VO management system, planned to be replaced by IAM
  - CVMFS : Software distribution
  - AMGA : Was used as metadata service until few months ago, now replaced by Rucio
- Computing resources and storages spread over more than 30 sites

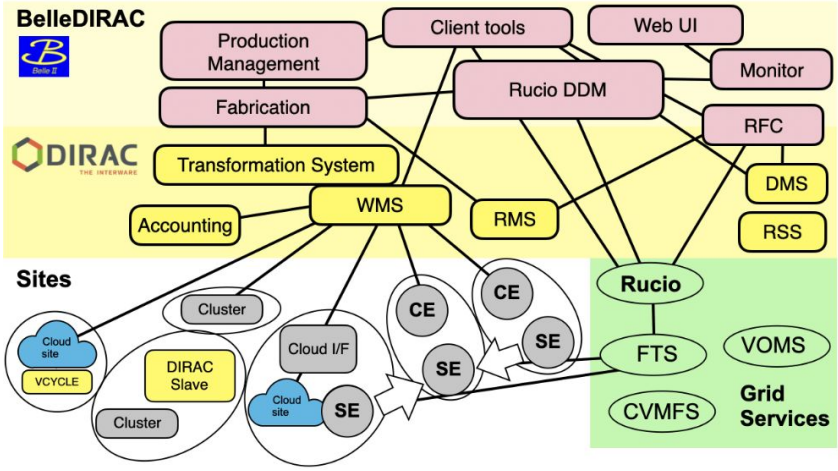

Belle II computing infrastructure

Storage situation

- Storage ressources split between disk (~20 PB) and TAPE (~10 PB)
- TAPE split between KEK and 6 RAW Data Centres in US, CA, DE, FR, IT

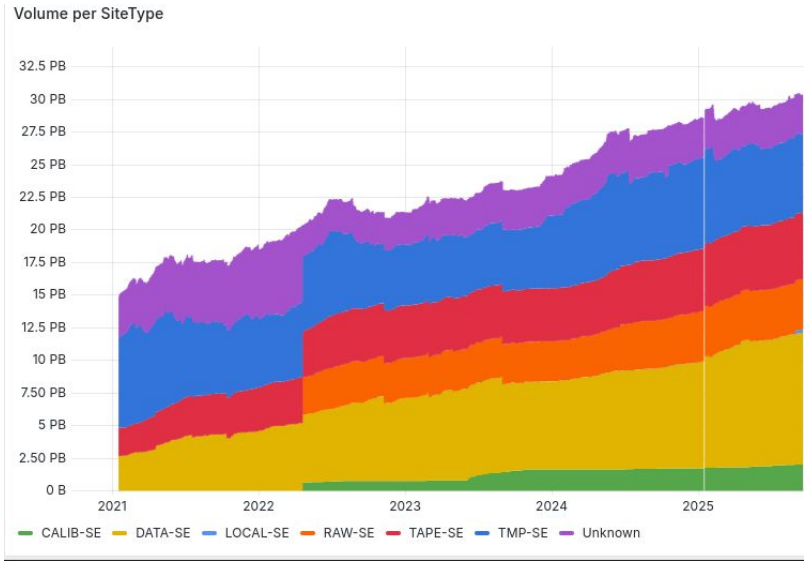

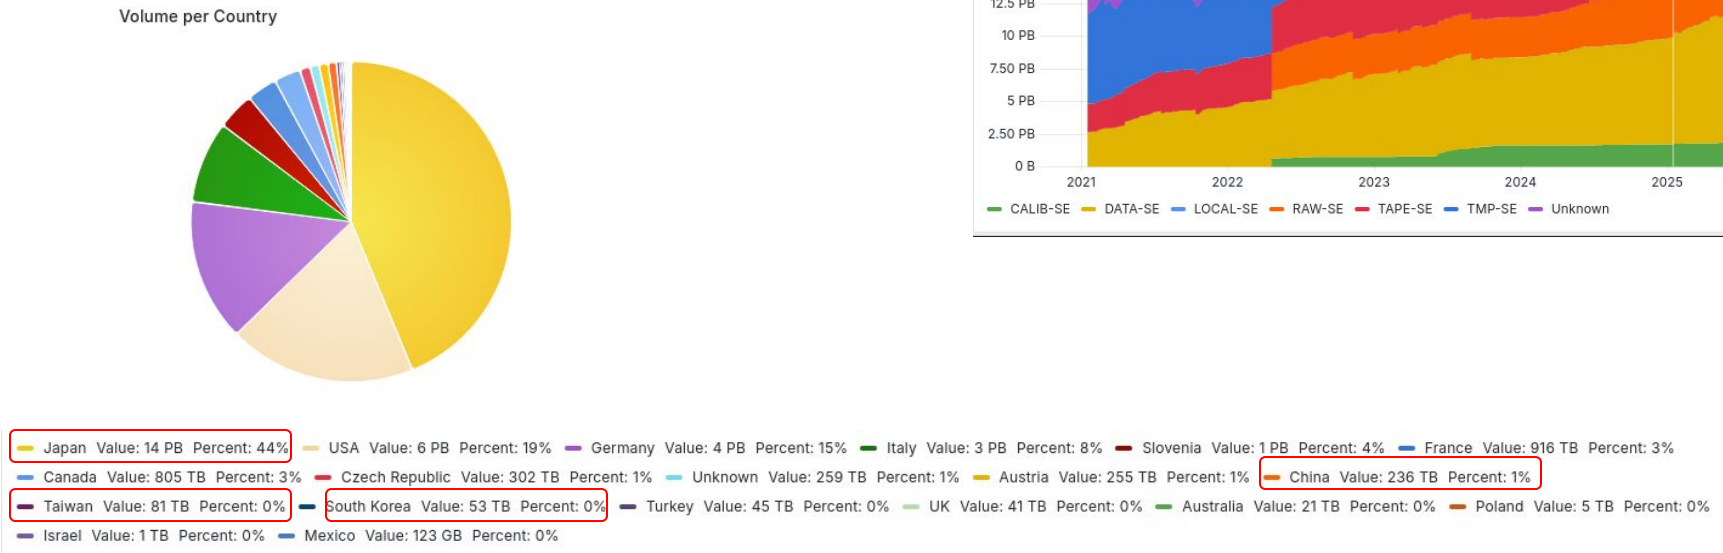

Brookhaven National Laboratory

Transfer overview

- Transfer reaching 40 PB since switch to Rucio in Jan 2021. The whole service runs smoothly
- In February 2024, Belle II took part of the DC24 challenge organized by WLCG and demonstrated the current infrastructure can handle 5 times the expected throughput at the end of the decade

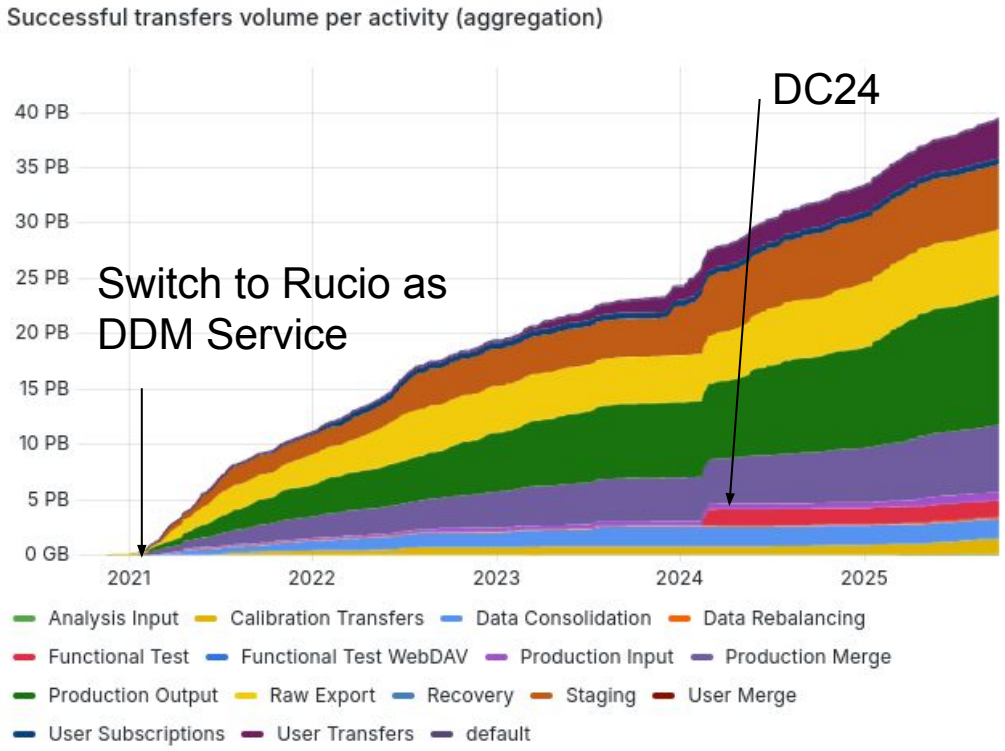

Computing resource situation

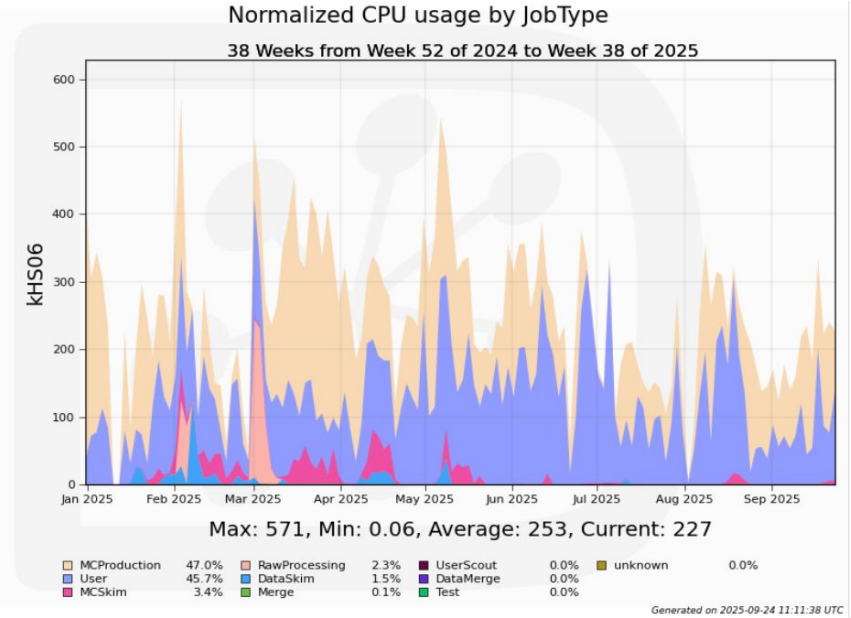

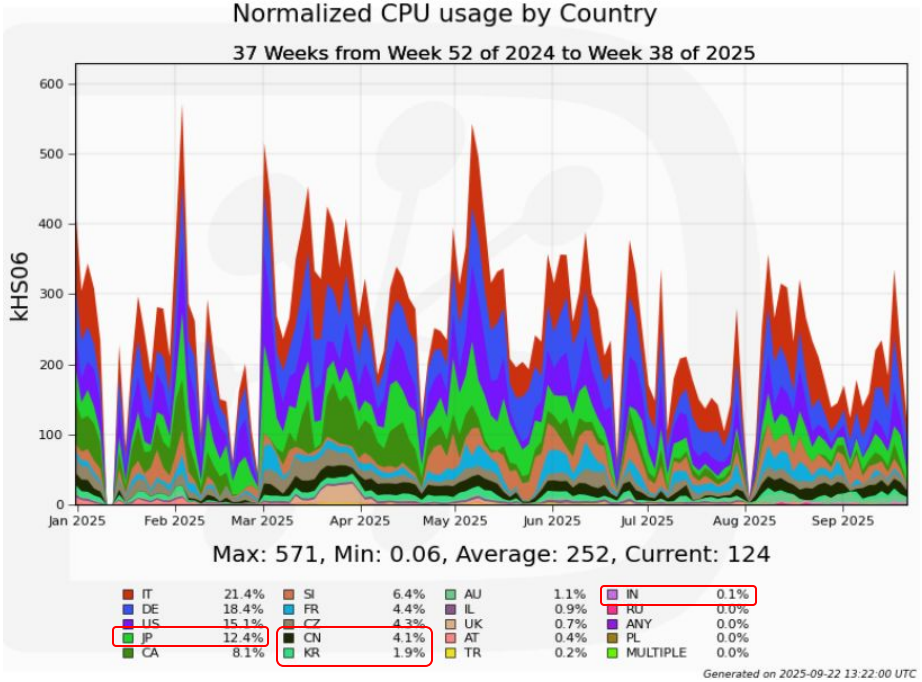

- During the last year, CPU usage dominated by production and user analysis. Reprocessing now conducted every 2 years represents a small fraction of our jobs
- Contribution of Asian sites represents ~20% of the CPU. New sites are being configured in particular in India

Brookhaven National Laboratory

SRM and GridFTP decommissioning

- SRM (Storage Resource Manager) used to be the standard tool for the interaction with Storage Element until mid-/late 2010s
  - SRM had staging capabilities to recall data from TAPE
    - Now replaced by TAPE REST API, enabled on all Belle II Tape endpoints (except KEK)
  - SRM also provided space reporting
    - Replaced now by SRR (Storage Resource Reporting) based on json and enable on all Belle II sites
  - All the sites supporting Belle II are now free to decommission SRM (already done for the majority of the sites)
- GridFTP was the default protocol behind SRM during the same period. Now replaced by WedDAV and/or xrootd

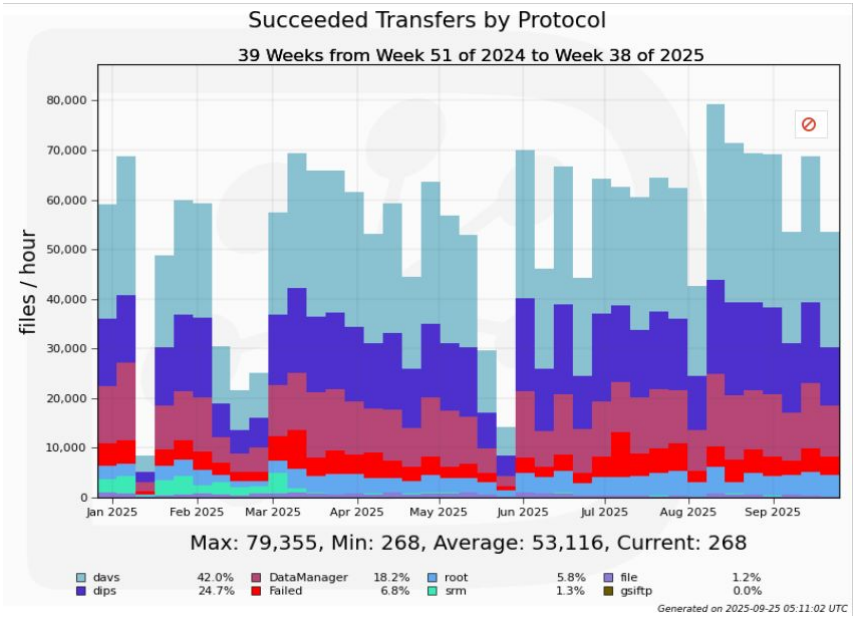

OS migration

- EL7 end of life occurred in July 2024
- Since that date :
  - All our services transition to Alma9 or EL8/9
  - Most of the sites transitioned to EL8/9 or Alma9. Remaining still operating EL7 will migrate in the coming months
- Since one year we use Apptainer, based on the job requirement to launch a container when needed. If user's requirement match OS available on site, do not use container

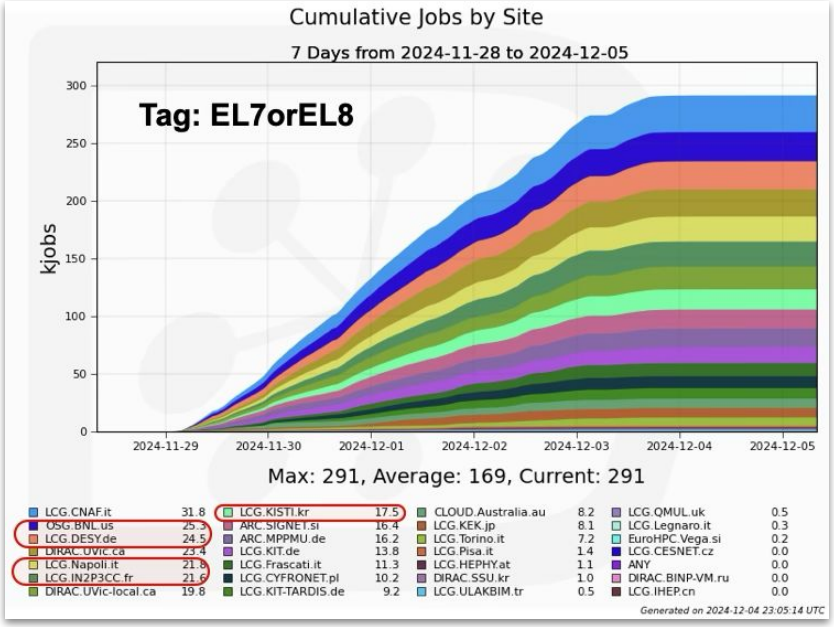

Tokens - KEK CA

- WLCG has been working on a transition from X509 + VOMS towards JSON Web Tokens (JWT) since July 2017
- Token adoption is a slow process. Situation in Belle II :
  - Job submission using token enabled on our certification machine. Enabling submission in production is just being discussed this week
  - Third Party Copy (using FTS+Rucio) and deletion using tokens already enabled on a few sites. Ticketing campaign will be conducted soon to configure tokens on all Storage Elements supported by Belle II
  - No progress on upload and download workflows (pending on situation in WLCG and Rucio)
- Until full support of all the workflows by token, Belle II will still use certificates
  - Until 2 days ago, KEKCC CA still used SHA1 signature. Replaced by new CA. Will need to regenerate the host certificates for all our services

Network: Main Pillars

SINET6 (Science Information Network) links

https://www.sinet.ad.jp/en/news_en/sinet6-upgraded-japan-amsterdam-lines-to-400gbps-april-1-2024

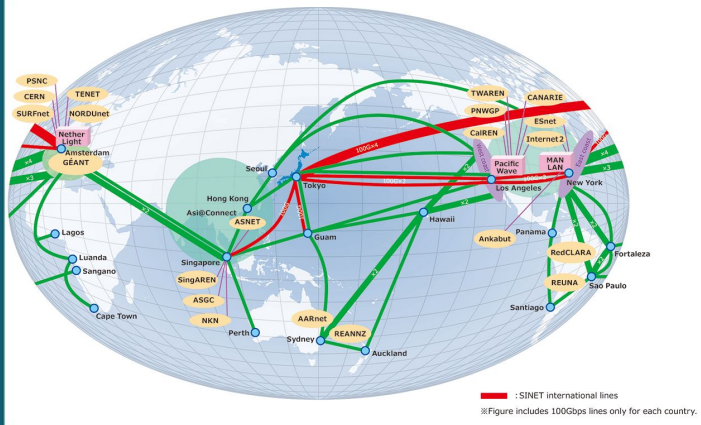

LHCOPN Optical infrastructure that can be used without jeopardizing resources

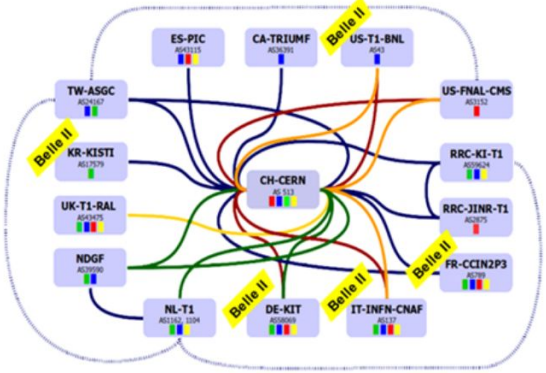

LHCONE L3 VPN

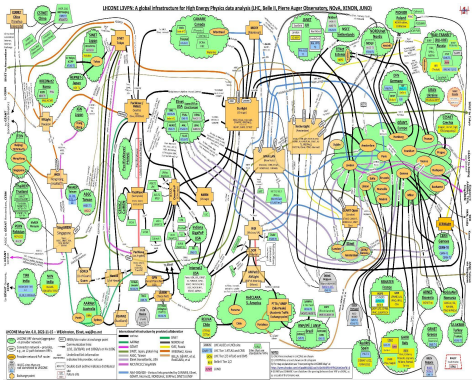

MultiONE

Belle II site Connected to LHCONE

https://twiki.cern.ch/twiki/bin/view/LHCONE/MultiOneBGPcommunities

The activity requires a work site side and NREN side.

Currently, 13 of the 24 Belle II sites connected to LHCONE are seen to publishing BGP communities.

# Belle II Sites	RC Site	BGP Communities
1	IN2P3-CC	ALICE, ATLAS, BelleII, CMS, DUNE, ILC, JUNO, LHCb, NOvA, Pierre, XENON
1	INFN-T1	ALICE, ATLAS, BelleII, CMS, DUNE, JUNO, LHCb, Pierre, XENON
2	FZK-LCG2	ALICE, ATLAS, BelleII, CMS, LHCb
1	praguelcg2	ALICE, ATLAS, DUNE, NOvA, BelleII, Pierre
1	DESY-HH	ATLAS, BelleII, CMS
1	INFN-NAPOLI-ATLAS	ATLAS, BelleII, CMS, LHCb, Pierre, perfSONAR
1	BNL-ATLAS	ATLAS, BelleII, DUNE
2	ARNES	ATLAS, BelleII, LHCONE, Pierre, perfSONAR
1	IN2P3-IRES	perfSONAR, BelleII, ILC
2	CA-UVic-Cloud	ATLAS, BelleII,NOvA

N.B. Network Prefix are kept from CRIC. To doublecheck that belle II resources are configured on the proper subnet.

IPv6

- Almost all central services are IPv6 compliant and configured in dual stack.
- Setup of IPv6 on Condition DB ongoing in those days.

From site report

- 26 Storages endpoint configured in dual stack as of 34
- 23 Sites reported that WNs are configured in dual stack as of 57

On september 2nd IPv4 prefix has been removed from LHCOPN links at BNL and FNAL, and rerouted over LHCONE.
No issue has been observed from belle II.

Multicore Jobs

- Evaluating multi-core job processing
- Currently in certification, when using a resource Tag "MultiCore"
  - Jobs are matched to multicore queues at the raw data centers
  - Tests in progress

```
{
  "ProductionName": "MCTest_charged_BGx0",
  "ProductionGroup": "charged",
  …
  "RunEvents": {"1003": {"0": 1000}},
  "MaxJobEvents": 200,
  "Stream": 0,
  "ResourceTag": "MultiCore,EL7orEL8"
}
```

Site	Configured	MC jobs	Data reprocessing
KIT.de	OK	OK	OK
UVic.ca	OK	OK	OK
IN2P3CC.fr	OK	OK	OK
BNL.us	OK	OK	To-do
DESY.de	OK	OK	To-do
CNAF.it	OK	OK	To-do

KEKCC

- KEKCC is our “tier 0” storing one full copy of all RAW data
- It is entirely replaced every 4-5 years. Last renewal last year, (previous time was in 2020)
  - All services and WNs hardware upgraded
  - OS migration
  - More disk, tape (20%) and CPU power (+40%) than previously
- Transition was challenging since it impacted all our main services. After one year, everything is stable
- Latest change is the deployment of the new KEK CA (has just been done 2 days ago)
- More details in Iwai-san’s slides

Plan for the coming months

- Data taking will resume soon. Expect to accumulate luminosity
- Regarding computing :
  - Focus on enabling tokens on most of the sites
  - A big development effort is foreseen to move to the next DIRAC generation (aka DiracX)
  - Will evaluate the possibility to optimize data distribution based on dataset popularity
  - More and more sites/countries are interested in setting up analysis facilities for the last step of user analysis. Will provide a list of recommendation for them based on what is done on some already existing facilities

Longer term plans

- Ultimate goal: reach 50 ab^{-1} (50 time Belle data) by operating at the design luminosity of 6.10^{35} cm^{-2} s^{-1}
- Run 2 until 2032, followed by Long Shutdown 2 (LS2)
- LS2 upgrade plan being addressed in a detector Technical Design Report (TDR) expected in 2027

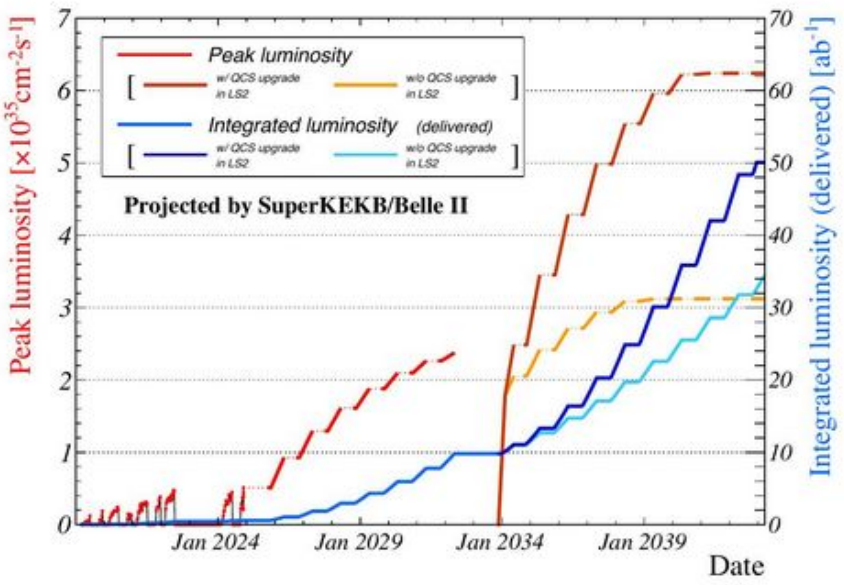

Conclusion

- Belle II experiment will restart data taking soon
- Computing working smoothly even though a lot of changes happened :
  - During the last months (e.g. OS change, KEKCC upgrade)
  - More changes will continue in the future (transition to tokens, move to, DiracX, etc.)
- Experiment plans evolves with LS2 now planned to happen in 2032
  - These changes will impact our computing resources plans
  - But the underlying infrastructure based on our core services like DIRAC and Rucio should stay
- Thanks to my colleagues for providing me with inputs, materials and valuable feedback : I. Ueda, M. Hernandez Villanueva, S. Pardi