Status report on KEK CRC

Tomoe Kishimoto

Computing Research Center, KEK

tomoe.kishimoto@kek.jp

KEK Computing Research Center (CRC)

- KEK CRC supports a wide range of research activities at KEK from the computing perspective

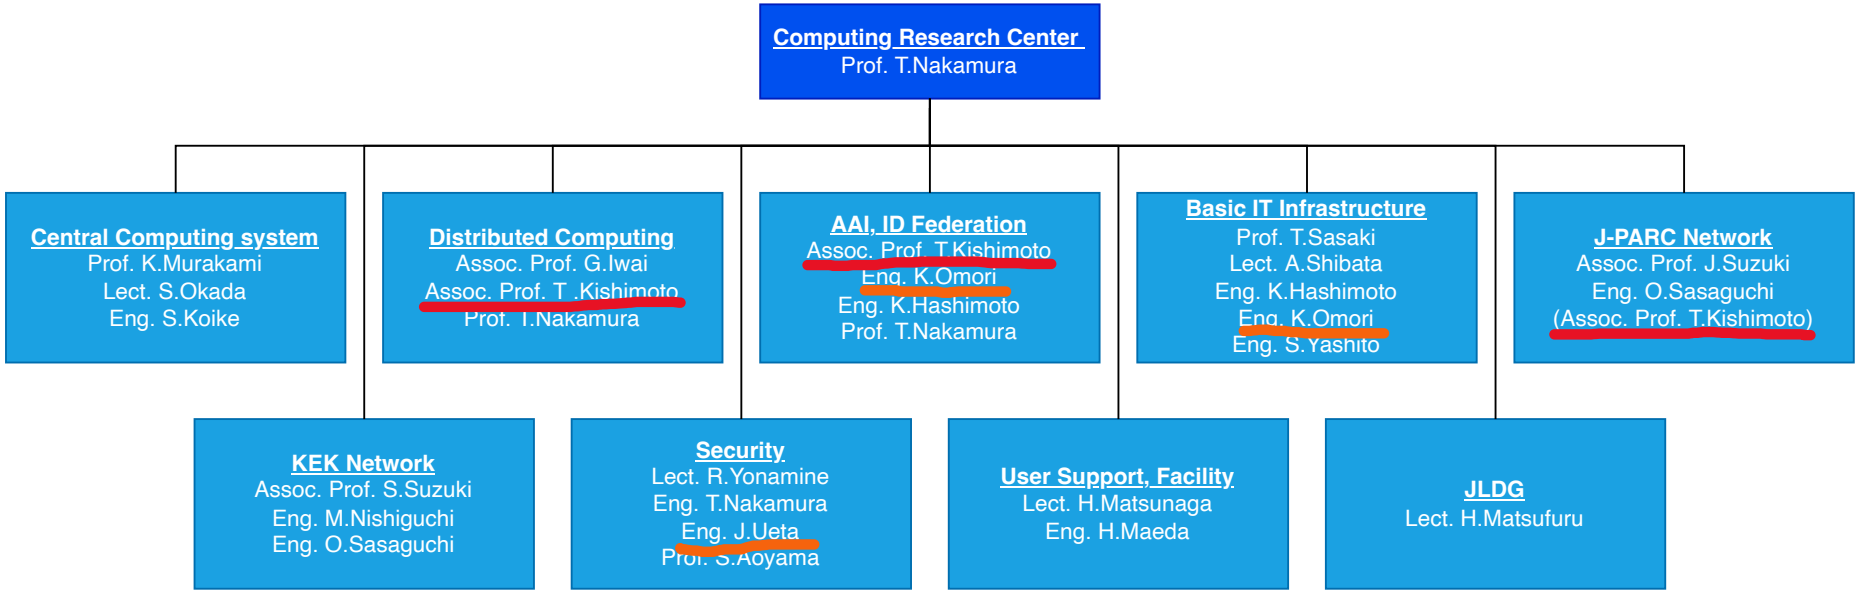

- I have been working on “J-PARC Network”, “Distributed Computing”, and “AAI/ID Federation”
  - This presentation will focus on the activities of J-PARC Network and Distributed Computing
  - AAI/ID Federation and Security will be covered by Konomi and Jo, respectively

What is J-PARC ?

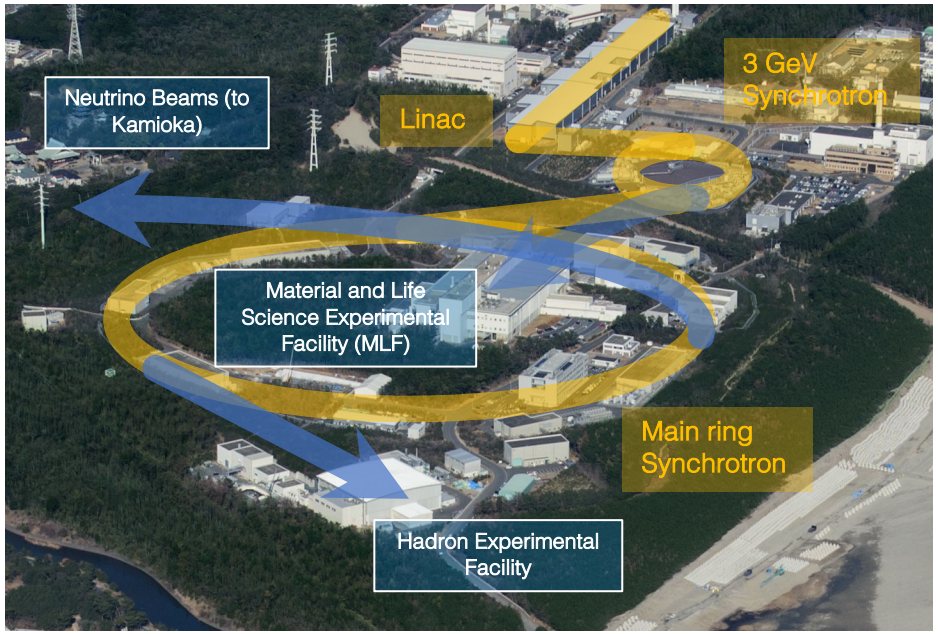

- J-PARC (Japan Proton Accelerator Research Complex)
  - Accelerator-based research facility that provides high-intensity proton beams for a wide range of scientific experiments:
    - Neutrino physics, hadron physics, Material and life sciences, etc
  - Jointly operated by KEK and JAEA (Japan Atomic Energy Agency)
    - The accelerators began operation in 2009
  - Located in the “Tokai” area of Ibaraki Prefecture
    - ~60 km from KEK the Tsukuba campus

大強度陽子加速器施設 J-PARC Japan Proton Accelerator Research Complex

KEK JAEA

加速器だから見える世界。 KEK

Information System Section (ISS)

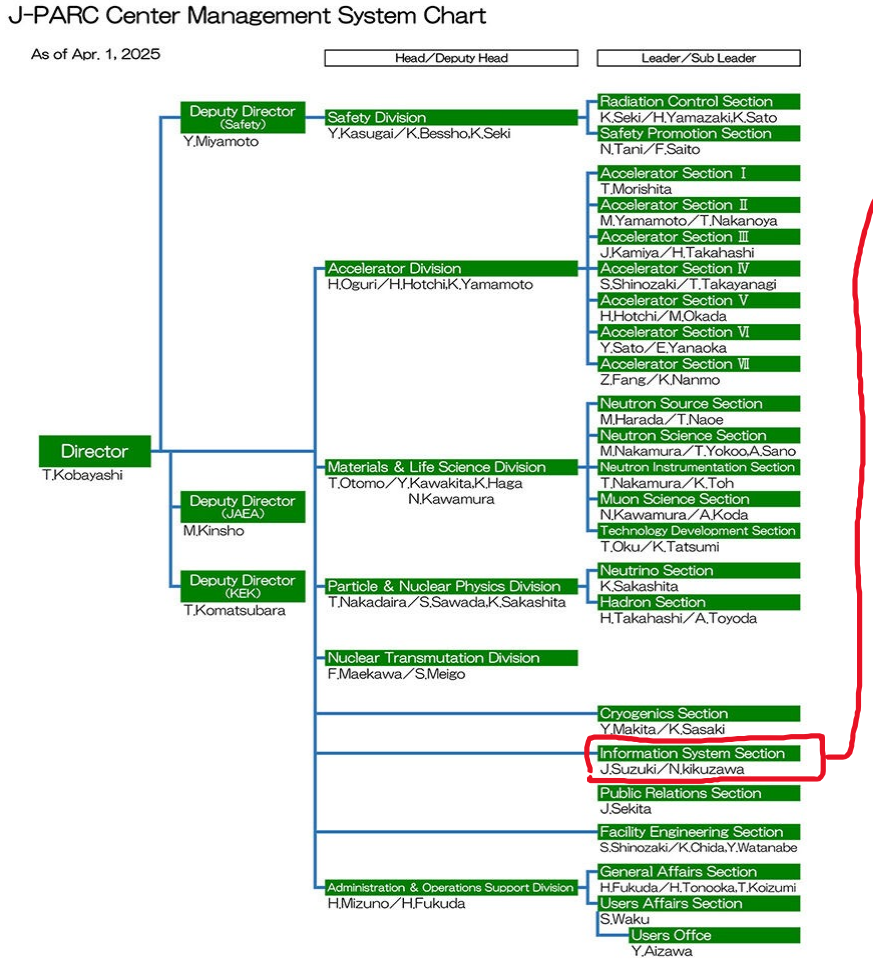

- Role of ISS:
  - Plans, designs, manages, and operates the network and IT infrastructure of J-PARC
    - E-mails, SSL-VPN, Web, User support, etc
  - Including supports to ensure its information security
- KEK Central Computing (KEKCC) system provides large-scale computing and storage resources for J-PARC experiments
  - KEKCC is located at the KEK Tsukuba campus

→ One of the key missions of ISS is to support the reliable transfer of experimental data from J-PARC to KEKCC

Overview of JLAN

- JLAN is the dedicated network infrastructure for J-PARC
  - Designed and operated independently from KEK LAN and JAEA LAN
  - Deployed both J-PARC and KEK Tsukuba campus
  - Edge switches and wireless access points (AP) for each building
- JLAN is connected to Japanese NREN (SINET6)
  - SINET6 provides both Internet access and L2VPN connection between JLAN Tokai and JLAN Tsukuba

→ A high-bandwidth network (20 Gbps) between J-PARC and KEKCC has been established

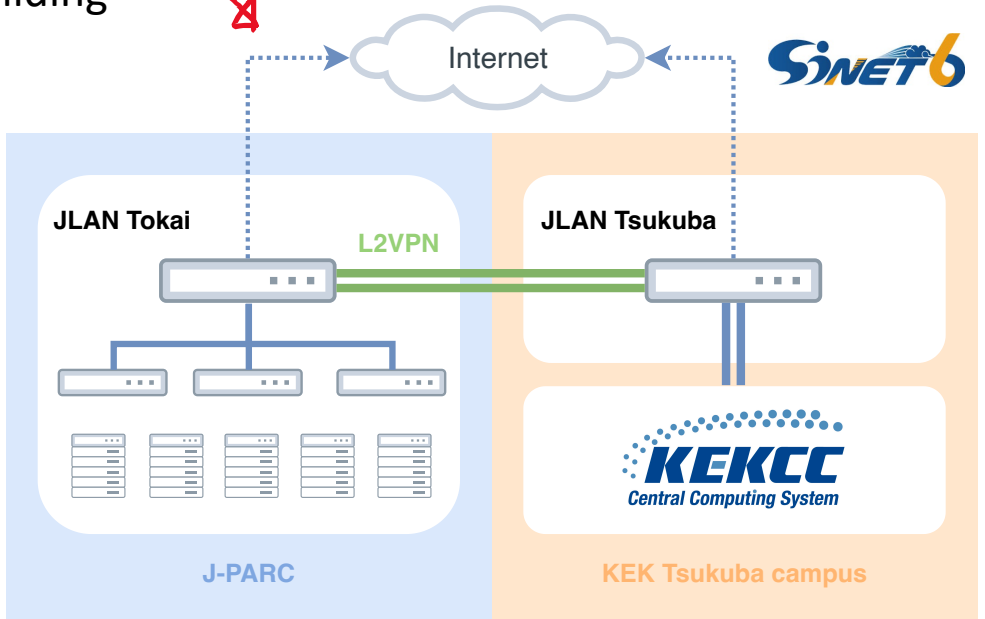

JLAN statistics

- More than 5k hosts, including user PCs, has been registered to JLAN
  - ~2k hosts are consistently connected at the same time

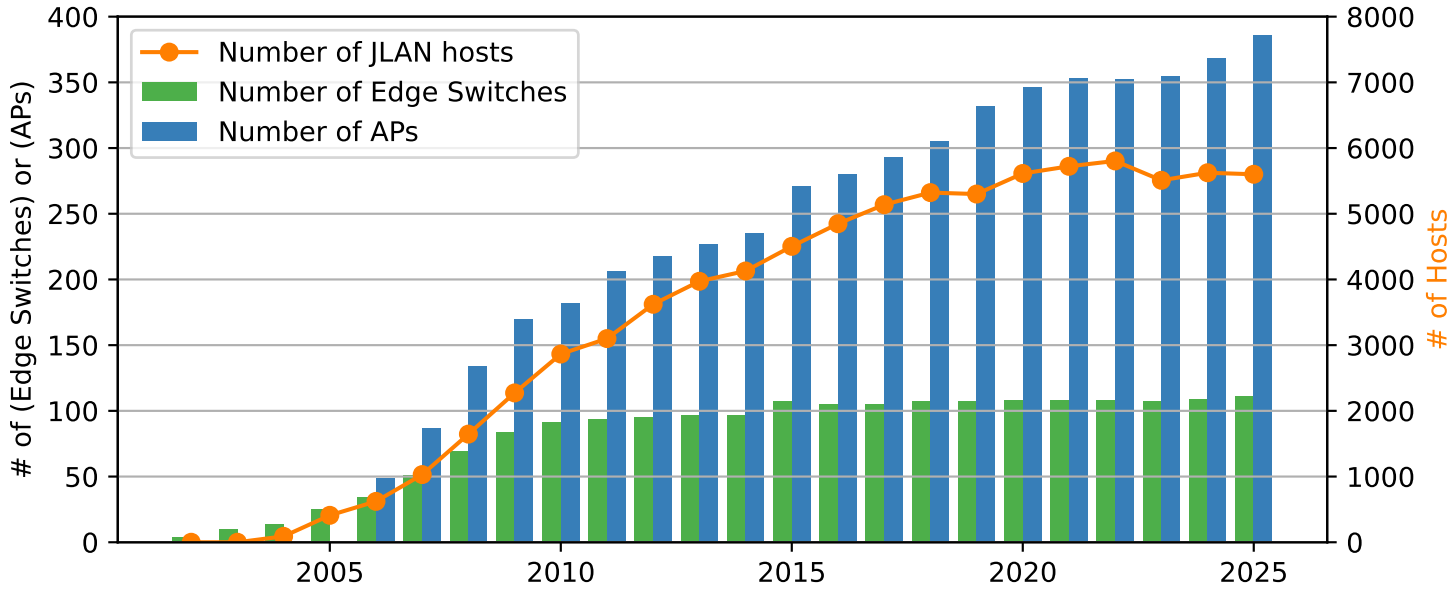

- The number of Edge switches is nearly constant
- The number of APs is still increasing to meet user needs

Data transfers from J-PARC to KEKCC

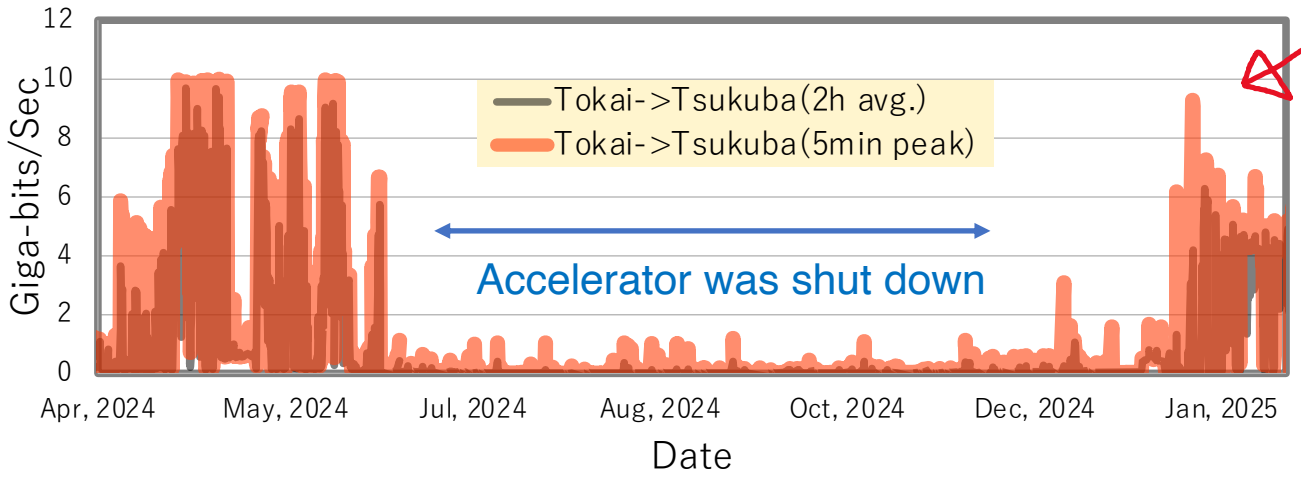

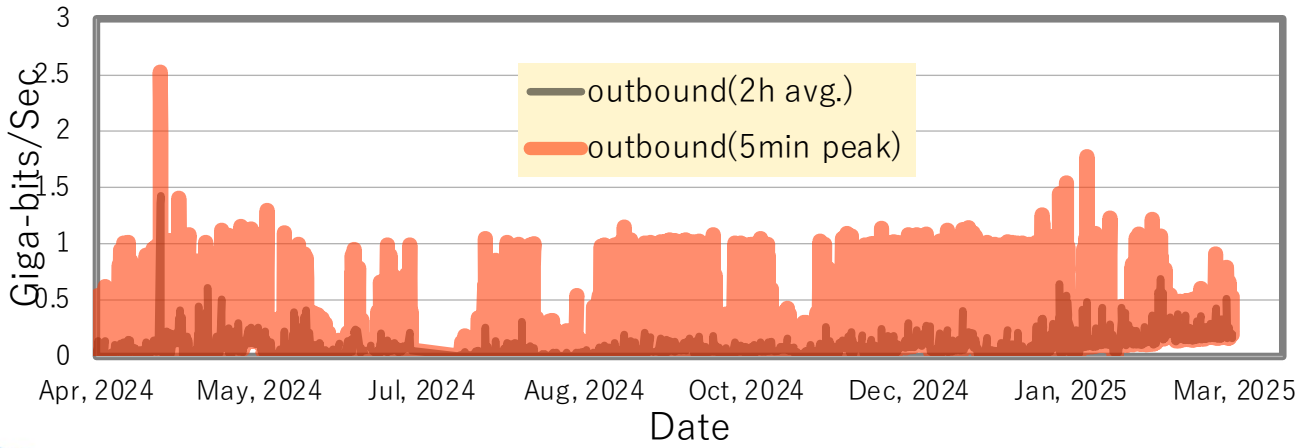

- Large volumes of data are transferred from J-PARC to KEKCC during accelerator operation
  - Edge switches (e.g. Hadron facility) occasionally reach their maximum throughput limit of 10 Gbps
- Data transfers from JLAN to Internet peak around 1Gpbs over a 5-minute interval

加速器だから見える世界。
KEK

MLF experiment (g-2/EDM)

- A project for the precision measurement of g-2 and electric dipole moment (EDM) is in progress at J-PARC
  - They conducted a data transfer test over JLAN

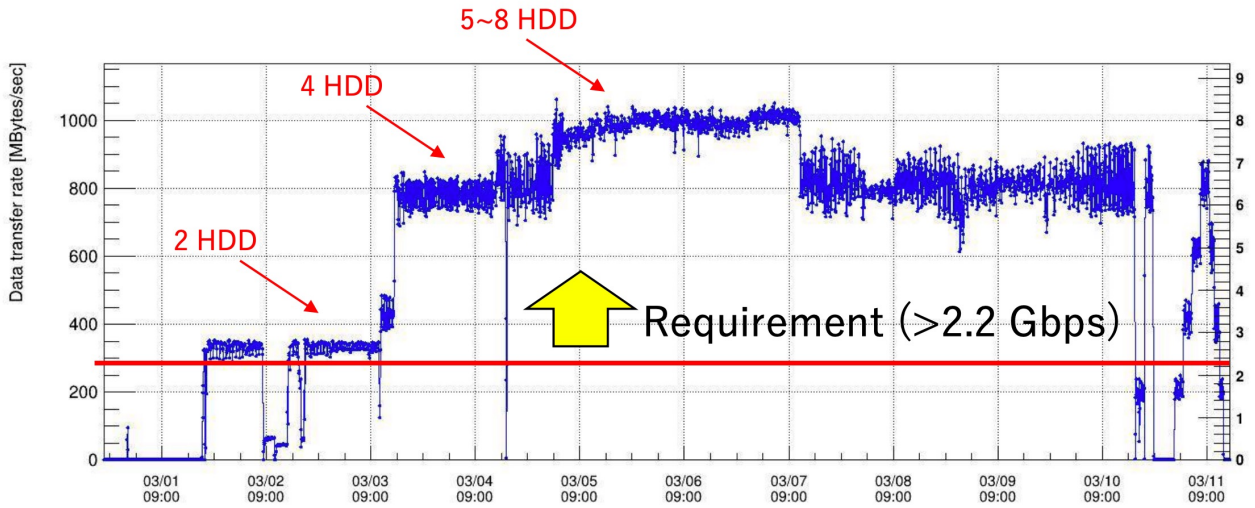

→ The data transfer rate exceeded the experimental requirement!

J-PARC muon
g-2/EDM experiment

Figure and pictures from "Keisanki workshop@ICEPP"

Overview of KEKCC

- KEKCC provides computing resources for various research activities at KEK and J-PARC
  - Current KEKCC started operations in Sep. 2024
  - Linux computing cluster (LSF) + storage system (GPFS/HSM)
  - Grid system (ARC-CE, StoRM, etc)
- CPU: ~12k cores
  - AMD EPYC 9654 (96cx2/nodes)
- Disk storage: 30PB
  - 20PB: GPFS for user groups
  - 10PB: GPFS-HPSS interface (GHI) as HSM cache
- Tape: 120PB as maximum capacity

KEKCC Central Computing System

Monitoring dashboard

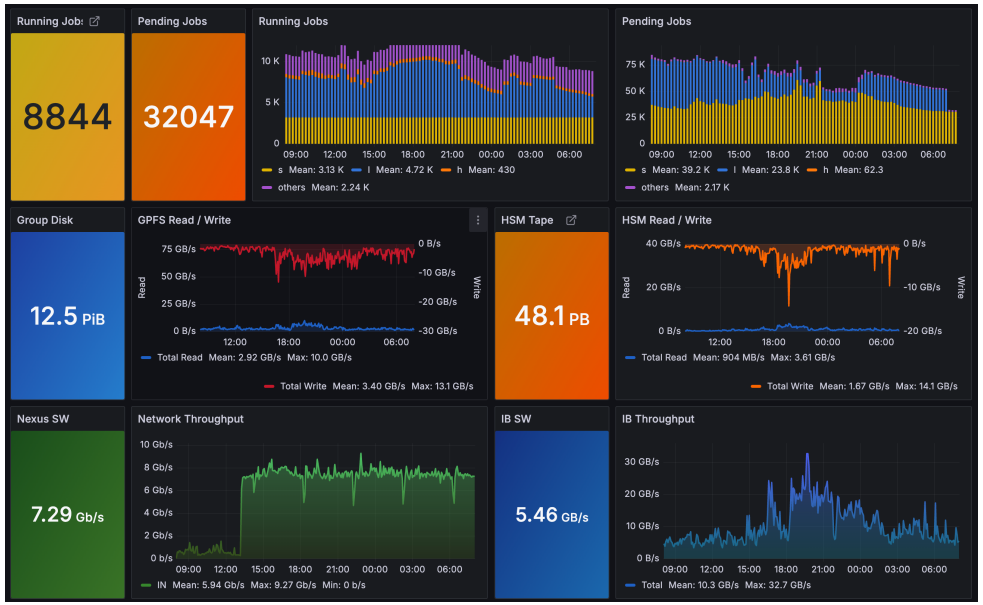

加速器だから見える世界。
KEK

Resources utilization @KEKCC

CPU usage

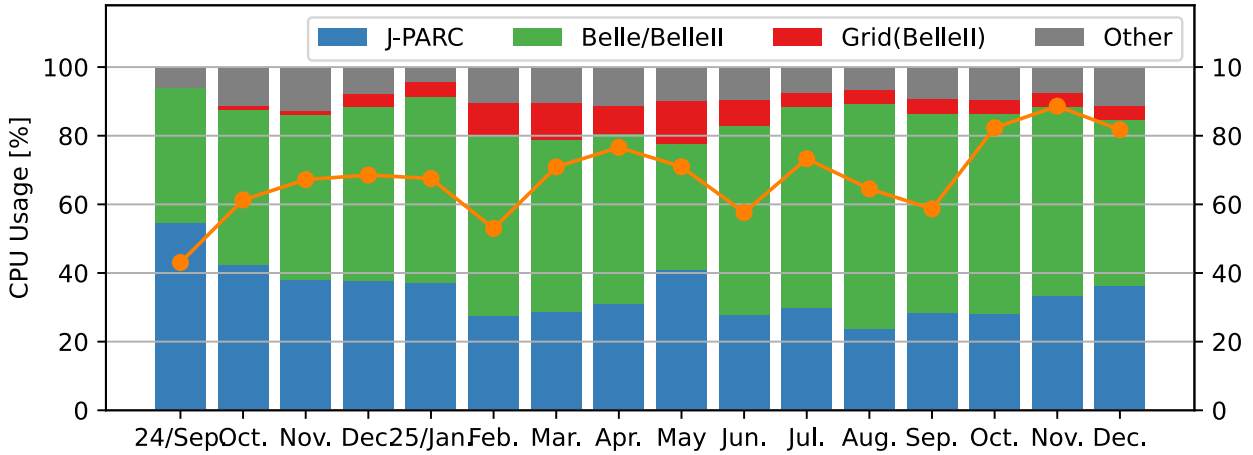

Tape usage

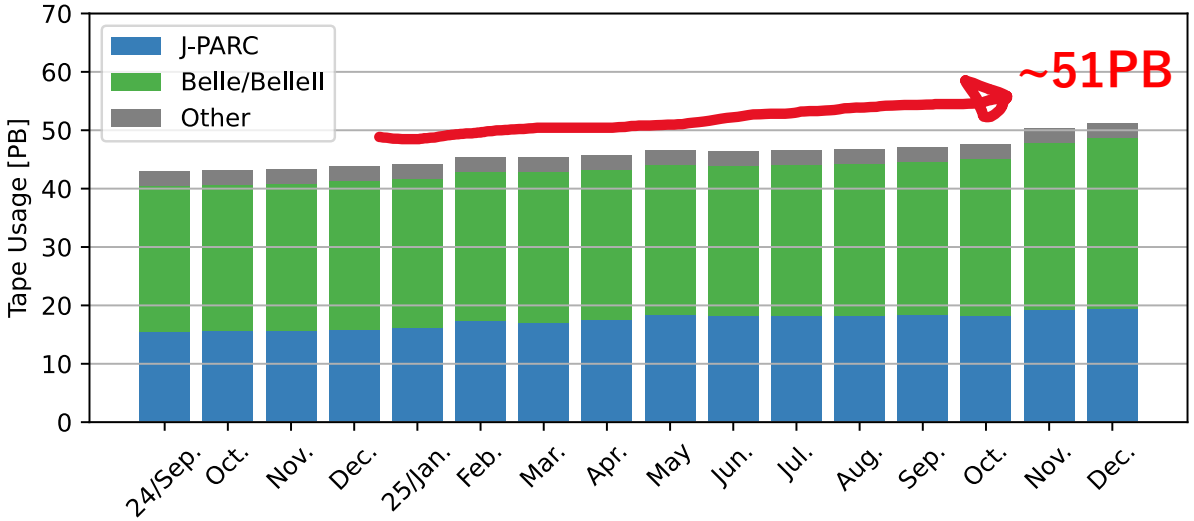

- The CPU resources are used efficiently
  - J-PARC experiments are not yet fully adapted to the Grid
    - Hyper-Kamiokande experiment is scheduled to participate in the next WLCG data challenge
    - Efforts are ongoing to enable J-PARC experiments to utilize the Grid
- The total tape usage has reached 50PB

Status of Grid system

Ongoing migration campaign to RHEL9
As of September 2025

as Belle2 dedicated

Migrated to SciTags-enabled DTNs on RHEL9

Service	OS	VM/Bare metal	Ethernet	IPv6	HA	UPS
StoRM	RHEL9 (Sep '25	Bare metal	10GE	✓	✓	
VOMS	RHEL7 + ELS	VM	10GE	✓	✓ SIOS LifeKeeper	✓
IAM	RHEL9	Bare metal	10GE	✓	✓	✓
AMGA	RHEL7 + ELS	Bare metal	10GE	✓	✓ SIOS LifeKeeper	✓
Top BDII	RHEL9	VM	10GE	✓	✓	✓
Site BDII	RHEL9	VM	10GE	✓	✓	✓
FTS3	RHEL9	VM	10GE	✓	✓	✓
ARC-CE	RHEL7 + ELS	Bare metal	10GE	✓	✓	
CVMFS Stratum Zero	RHEL9	Bare metal	10GE	✓	✓	
CVMFS Stratum One	RHEL9	Bare metal	10GE	✓	✓	
CVMFS publisher	RHEL9	VM	10GE	✓		
Frontier Squid HTTP Proxy	RHEL9	VM	10GE	✓	✓	✓

Slide from Go Iwai

- The BelleII experiment is the primary user of the Grid system
- Some hosts are still running RHEL7 + ELS…
  - New CE needs to be evaluated: ARC-CE or HTCondor-CE
  - AMGA will be decommissioned
  - VOMS should be decommissioned

→ Complex work related to authentication and authorization is required …

KEK Grid CA

- 20 years operation, 3.8K user certs and 4.5K host certs
- Initially only for domestic users, then extended the service for specific experiment members outside Japan (e.g. BelleII)
- CA root certificate had an expiration date of Nov 2025
  - We aimed to decommission the KEK Grid CA by this date by migrating from X.509/VOMS to JWT-based/IAM authentication
  - However, this was not archived, as was the case with VOMS

→ New CA (KEK Grid CA2024) has been launched…

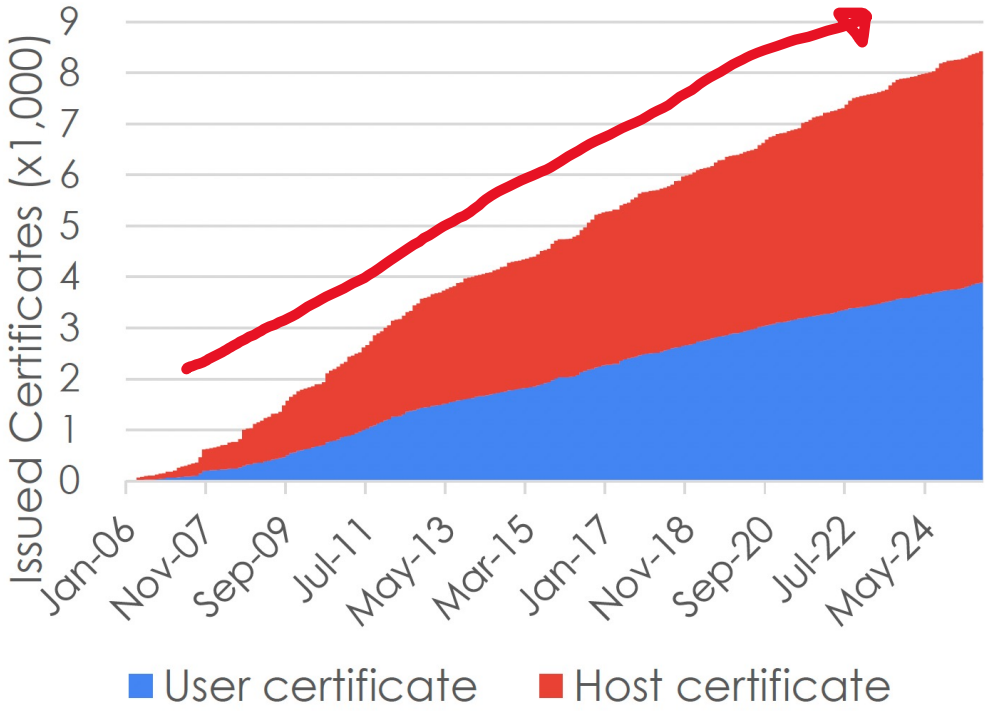

Status of INDIGO IAM deployment

- INDIGO IAM: Authentication and authorization infrastructure implements the VO concept
  - OpenID Connect provider
- IAM instances have been deployed for the BelleII experiment
  - Users are synchronized from VOMS
  - Then, users can obtain tokens from IAM

→ We have verified some basic use cases (e.g. TPC with FTS) using token-based authentication

→ Next steps: VOMS and X.509 should be removed…

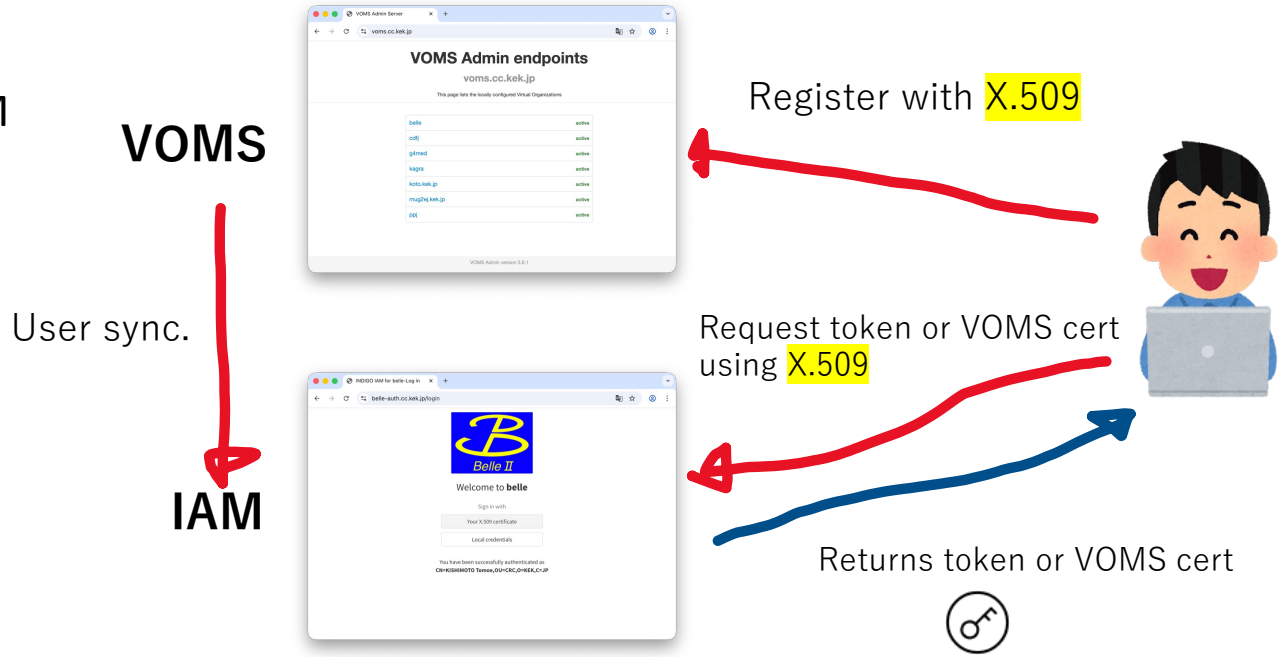

Decommission of VOMS

- As a first step, VOMS will be decommissioned:
  - Users will be registered directly in the IAM
  - Users who have X509 certificate will be allowed to submit the registration request via the web

→ At this phase, we are still dependent on X.509, and there is no particular benefit in terms of user experience…

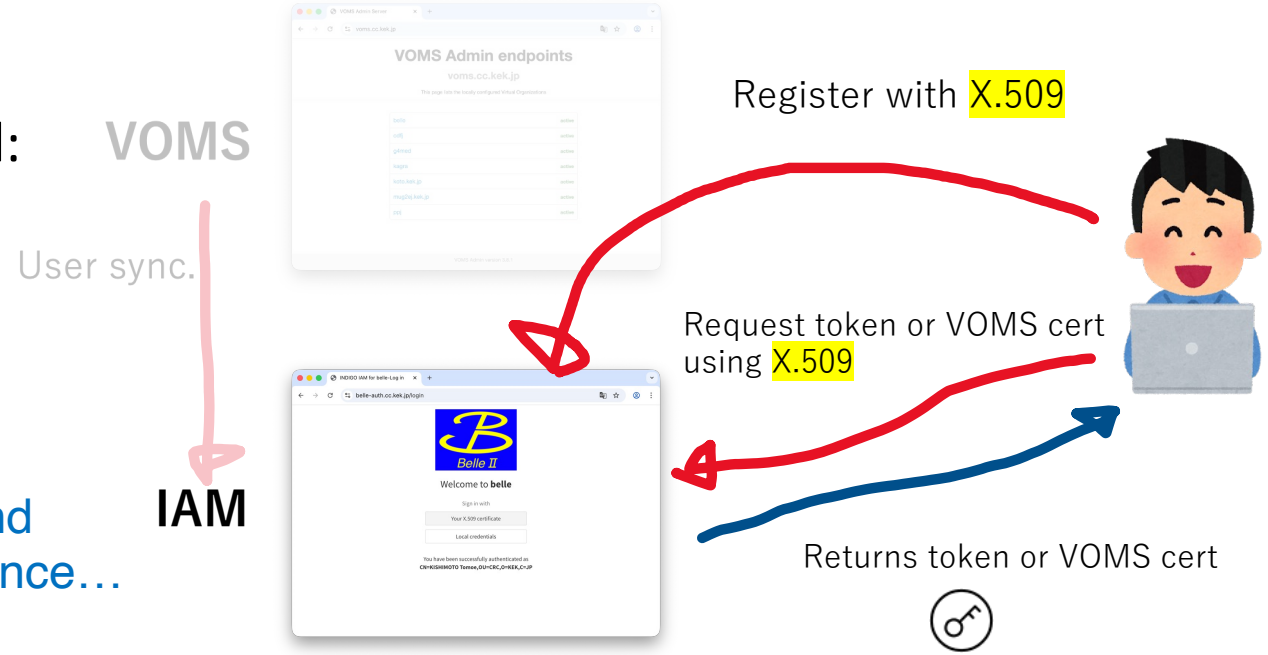

- In the current system, identity verification is performed when obtaining an X.509 certificate from the CA
  - E.g.) KEK CA requires an in-person meeting and a photo ID

→ A new procedure will be required if we remove CA and X.509

ID federation

- In a straightforward approach, KEK could perform the identity verification for all users:
  - However, this would impose a significant workload… (This should be done in the future.)
- ID federation is one possible solution
  - Users obtain an ID from their home institute or university, and the identity verification should be performed
    - Criteria for the identity verification need to be defined, we assume **that Identity Assurance Level 2 (IAL2) is equivalent to the Grid criteria**
  - eduGain in Europe and GakuNin in JAPAN

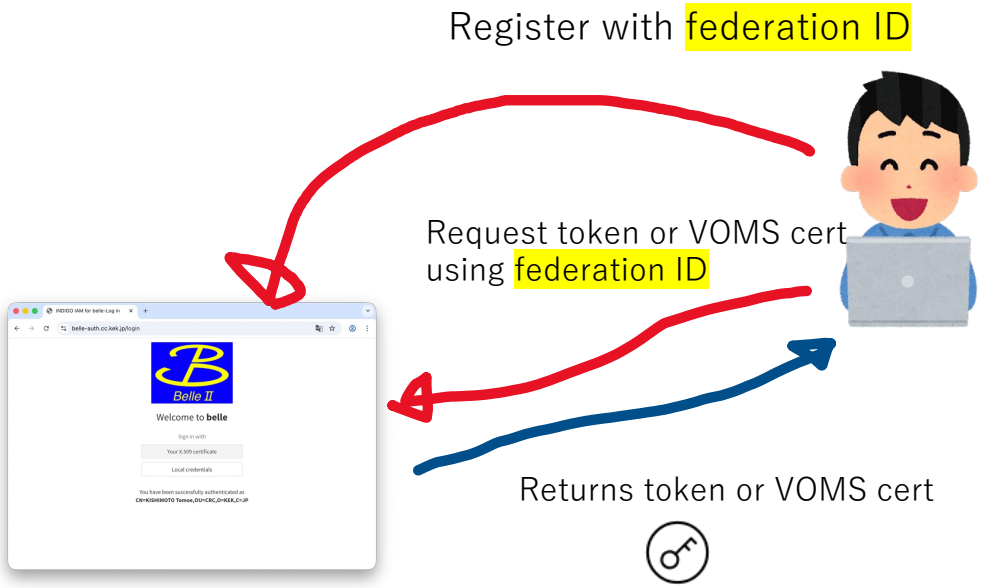

→ As a first step, we deployed an IdP at KEK and joined the GakuNin federation last year
(Konomi will report the details tomorrow)

Summary

- J-PARC Network (JLAN) activities:
  - Stable operation of the high-bandwidth network (20Gbps) connecting J-PARC and KEK Central Computing
- Grid System and Authentication:
  - Migrations are ongoing to RHEL9; VOMS and AMGA should be decommissioned
  - New KEK Grid CA was launched to bridge the gap in the transition to token-based authentication
  - Future: Aiming for ID Federation (GakuNin, eduGain) to streamline identity verification

→ Since CC-IN2P3 and KEK CRC share many challenges, I hope to have meaningful discussion with you through FJPPN.