



AFAD 2026

Asian Forum for Accelerators and Detectors

13-17 October 2026 Shanghai, China



KEK Central Computing System (KEKCC):

Operational Updates, Asian Network Shifts, and Readiness for Belle II
Data Challenge 2027

Go Iwaj, T. Kishimoto, K. Murakami, T. Nakamura, and S. Suzuki

High Energy Accelerator Research Organization (KEK)

Computing Research Centre (CRC)



Outline



News and Updates – may be relevant to some of you

- New CAs in production
- Inter-Asia network routing



KEK Central Computing System (KEKCC)

- 2 years of production experience



Readiness for Belle II Data Challenge 2027

- Data transfer test between B2RAWDCs with token-based authentication and authorisation

Outline



News and Updates – may be relevant to some of you

- New CAs in production
- Inter-Asia network routing



KEK Central Computing System (KEKCC)

- 2 years of production experience



Readiness for Belle II Data Challenge 2027


- Data transfer test between B2RAWDCs with token-based authentication and authorisation

20+ years operation


2nd Gen of KEK Grid CA in Production

4.1K user certs


4.9K host certs




- Launched a new CA (KEK Grid CA 2024) in **September** 2025 to continue X.509-based auth. for the next few years
 - Accredited and included in the EGI Trust Anchor v1.137 (Sep 15, 2025)
 - New root certificate has longer validity and is signed with a more secure algorithm, SHA-2
- **IHEP CA 2026** was released in v1.145 (Aug 31, 2026), distributing CRLs with stronger signature algorithms
- **Issue:** Temporarily applying a workaround to accept **SHA-1 signed CRLs** across KEK nodes:
 - `update-crypto-policies --set DEFAULT:SHA1`
- Still, ~100 users in Belle II hold legacy IHEP CA certificates
- Once this user migration is complete, the strict policy will be re-applied




Recap: International Network Route Transatlantic Back Again



- SINET serves as the primary NREN for KEK, providing both domestic and international connectivity
- Siberian 100G route for Euro has been upgraded to the transatlantic route with 100G x4 lines
- Dedicated line for the traffic between Japan and Euro, not shared with traffic for the US
- To minimise the latency:
 - traffic passes through fewer routers on the shortest route close to the Arctic




JGN: Tokyo Hong Kong Line



- Japan Gigabit Network (**AS17934**) has been another LHCONE traffic backbone between **Tokyo and Hong Kong**, ensuring **50+% lower latency** routing than Singapore path, for specific Asian sites, such as: KISTI, ASGC, and IHEP, for a decade
 - JGN is operated by NICT (MIC) but not NII (MEXT), our primary NREN provider
 - Decommissioned in Dec 2025
- Since then, APAN-JP (**AS7660**) has taken over operations
 - e.g.: AS2505(KEK)→AS2907(SINET)→~~AS17934(JGN)~~→**AS7660(APAN-JP)**→AS24489(TEIN)→AS23911(CERNET)x3→AS7497(CSTNET)→AS3460(IHEP)
 - Experienced a brief **LHCONE outage** (a few days) during the transition
 - Quickly **recovered** with help from people at IHEP ☺
- **Acknowledgments:**
 - Special thanks to Tao Cui and Qingbao Hu (IHEP) for their prompt troubleshooting and dedicated support during the outage in Dec 2025
 - Sincere appreciation to JGN for their long-standing provision and maintenance of the dedicated link over the years

Ongoing Issue: Geographically Close, but Farthest Apart in Routing!



Packets are literally enjoying a round-the-world trip 😊

Global Detour via CERN

- Traffic is routed through Europe (GÉANT) instead of the APAN-JP shortcut
- BGP prefers the shorter GÉANT path over the APAN-JP path:
 - Prefers a shorter path length like AS2505 → AS2907 → AS20641 → AS3460 rather than AS2505 → AS2907 → AS7660 → AS24489 → AS23911x3 → AS7497 → AS3460
- Asking NOC (SINET) to adjust Local Preference to restore the Hong Kong route

Trans-Pacific Journey

- Traffic crosses the Pacific Ocean to the US (ESnet) before heading back to Asia

On the Right path through Hong Kong

- Traffic is routed via LHCON through Hong Kong
- Hoping the upcoming Kitakyushu-Korea link with much shorter latency will give us a significant improvement

Outline



News and Updates – may be relevant to some of you

- New CAs in production
- Inter-Asia network routing



KEK Central Computing System (KEKCC)

- 2 years of production experience




Readiness for Belle II Data Challenge 2027

- Data transfer test between B2RAWDCs with token-based authentication and authorisation

KEKCC: Just Begun the 3rd Year of the 4-year Contract

A largest-scale computer system




- Linux Cluster + Storage System (GPFS/HSM)
 - Accommodate multiple experiments in a single system
- In production since September 2024
- CPU: 12,096 cores
 - AMD EPYC 9654, 2.4 GHz (Base clock, up to 3.7 GHz)
 - 192 cores/node (2 CPU/node, 96 cores/CPU)
 - 63 calculation nodes
 - 382K HS23 (31.56 HS23/core)
 - 6K HS23/node @3.5 GHz Measured w/o SMT
 - LINPACK: R_{max} : 0.52 / R_{peak} : 0.45 (PFlop/s)
- Memory: 56 TB
 - 4 GB/job
- Disk: 30 PB
 - 20 PB: GPFS for experimental groups
 - 5.4 PB for Belle II, and 2.4 PB for Belle
 - 10 PB: GPFS-HPSS-Interface (GHI) as an HSM cache
- Tape: 120 PB as maximum capacity

Grid instances are running in KEKCC

Site Scale Evolution




Resource History (Last 4-Gen)




Runtime in the Entire System



~20% contribution to the overall Belle II computing resource




86% more CVEs in 2026 vs. 2025



Frequent countermeasures against the massive amounts of vulnerabilities

86% more CVEs in 2026 vs. 2025

CVE






* Belle2 RAW data centres: BNL, UVic, CNAF, DESY, KIT, and CCIN2P3



6.3 PB of Belle2 RAW data for 0.9 ab⁻¹ of total int. luminosity

Int. L: 0.9 ab⁻¹




Belle II Data Challenge 2027 co-locating with WLCG DC27

The goal of Belle II DC27 is to make sure:

- KEK can distribute **post LS2 (2035)** Belle II RAW data to B2RAWDCs, while the shared network backbone carries 50% of the HL-LHC (Run 4) traffic load
 - Average: **2.26 Gbps**
 - Peak: **19.35 Gbps**
- Transfers use the **IPv6** protocol and **token-based authentication and authorisation** as much as possible

We are here as of today

Oct 15, 2026



Outline



News and Updates – may be relevant to some of you

- New CAs in production
- Inter-Asia network routing



KEK Central Computing System (KEKCC)

- 2 years of production experience




Readiness for Belle II Data Challenge 2027

- Data transfer test between B2RAWDCs with token-based authentication and authorisation


Overview

- Job Size: Each job attempts to transfer 100 individual files (800 MB each)
- Concurrency: 6 active jobs are sustained per direction simultaneously, i.e., 36 jobs for outbound and another 36 jobs for inbound are running during the test
- Duration: Continuous execution for 8 hours
 - Started at 15:00 on Sat, May 30th, 2026
 - Ended at 17:00 on Sun, May 31st, 2026
- The main purpose of the test is to confirm **NO significant performance difference** among the three different authentication and authorisation (authNZ) methods below:
 1. 8 hours for X.509 authNZ with a VOMS proxy certificate
 2. 8 hours for JWT authNZ (in **managed-mode**) with the default token lifetime, i.e., 1 hour
 - Which means FTS attempts to obtain the refresh token every hour on behalf of token's owner
 3. 8 hours for JWT authNZ (in **unmanaged-mode**) with a long-lifetime token (24 hours)
 - Need to extend the access token timeout in advance: 86400 seconds



Results: No Significant Performance Difference

WLCG Site Network Input/Output



Oct 15, 2026

WG6 - AFAD2026

Periodic throughput degradation every hour during the test with JWT managed-mode

WLCG Site Network Input/Output




Massive **TokenRefresh** requests are triggered simultaneously **every 60 minutes**, i.e. the default lifetime for the access token


All queued transfer processes fall in **PENDING** for **~5 mins** while obtaining the refresh token

Temporary throughput degradation during the test with JWT unmanaged-mode

WLCG Site Network Input/Output



No significant transfer inefficiency during the test overall



Summary



Inter-Asia network routing

- ☺ Our neighbours are routed quite far at this moment - **Let's improve the situation together!**



KEKCC Status & Operations

- ☺ Successfully entering the 3rd year of the 4-year contract with stable production
- ☺ High utilisation maintained (~80% CPU) with 55 PB of stored data (as of Sep 2026)



Readiness for Belle II Data Challenge 2027

- ☺ Confirmed performance and efficiency parity between X.509 and JWT-based AuthNZ
- ☺ Achieved ~80 Gbps throughput (at peak) between KEK and B2RDCs
- ☺ Almost ready for the Belle II Data Challenge in 2027!



Future Roadmap for Authentication and Authorisation

- ☺ Continued operation of the CA (KEK Grid CA 2024) and VOMS for the time being, until switching the primary user registration point to IAM before X.509/VOMS phaseout in 2028
- ☺ Fortunately or unfortunately, the recent AI-driven massive vulnerability mining trend brought an MFA wave to KEK - hoping that makes KEK more secure!

Many thanks for your Attention!



高エネルギー加速器研究機構


A Long Journey: Migrating X.509/VOMS to JWT/IAM

2025	Q1	
	Q2	
	Q3	Dual operation with VOMS and <u>IAM VOMS Attribute Authority (AA)</u> /cvmfs/grid.cern.ch/etc/grid-security/vomses/belle-belle-auth.cc.kek.jp /cvmfs/grid.cern.ch/etc/grid-security/vomses/belle-voms.cc.kek.jp
	Q4	JWT-supported StoRM
2026	Q1	FTS v3.14 for unmanaged token support
	Q2	JWT-supported ARC-CE
	Q3	
	Q4	
2027	Q1	WLCG Data Challenge 2027
	Q2	
	Q3	
	Q4	
2028	Q1	Completion of the X.509 / VOMS phaseout (user certificate)
	Q2	
	Q3	
	Q4	

Steadily moving toward a pure token-based environment

Action Items toward X.509/VOMS phaseout:

- Switch the primary user registration point from VOMS to IAM (Must be a lot of work behind the scenes)
 - MFA, Non-F2F ID verification, etc
- VOMS service decommissioning by Q1 2028



Ongoing migration campaign to RHEL9

As of Sep 2026

 as Belle2 dedicated

Service	OS	VM/Bare metal	Ethernet	IPv6	HA	UPS
 StoRM	RHEL9	Bare metal	10GE	✓	✓	
VOMS	RHEL7 + ELS	VM	10GE	✓	✓ 	✓
 IAM	RHEL9	Bare metal	10GE	✓	✓	✓
 AMGA (Decommissioned)	RHEL7 + ELS	Bare metal	10GE	✓	✓ 	✓
Top BDII	RHEL9	WIP VM	10GE	✓	✓	✓
Site BDII	RHEL9	WIP VM	10GE	✓	✓	✓
FTS3	RHEL9	VM	10GE	✓	✓	✓
ARC-CE	RHEL9	WIP 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	✓	✓	✓

Decommissioning process in progress by Nov 2026

Deprecated ARC6 instances in the decommissioning process, ended in Nov 2026