東京大学
素粒子物理国際研究センター
International Center for Elementary Particle Physics
The University of Tokyo

ATLAS EXPERIMENT

ICEPP Site Report

22nd Nov. 2022

The 6th Asian Tier Center Forum (ATCF6)

Masahiko Saito, on behalf of the operation team

ICEPP, The University of Tokyo

International Center for Elementary Particle Physics (ICEPP)

ICEPP
The University of Tokyo

Main projects at ICEPP

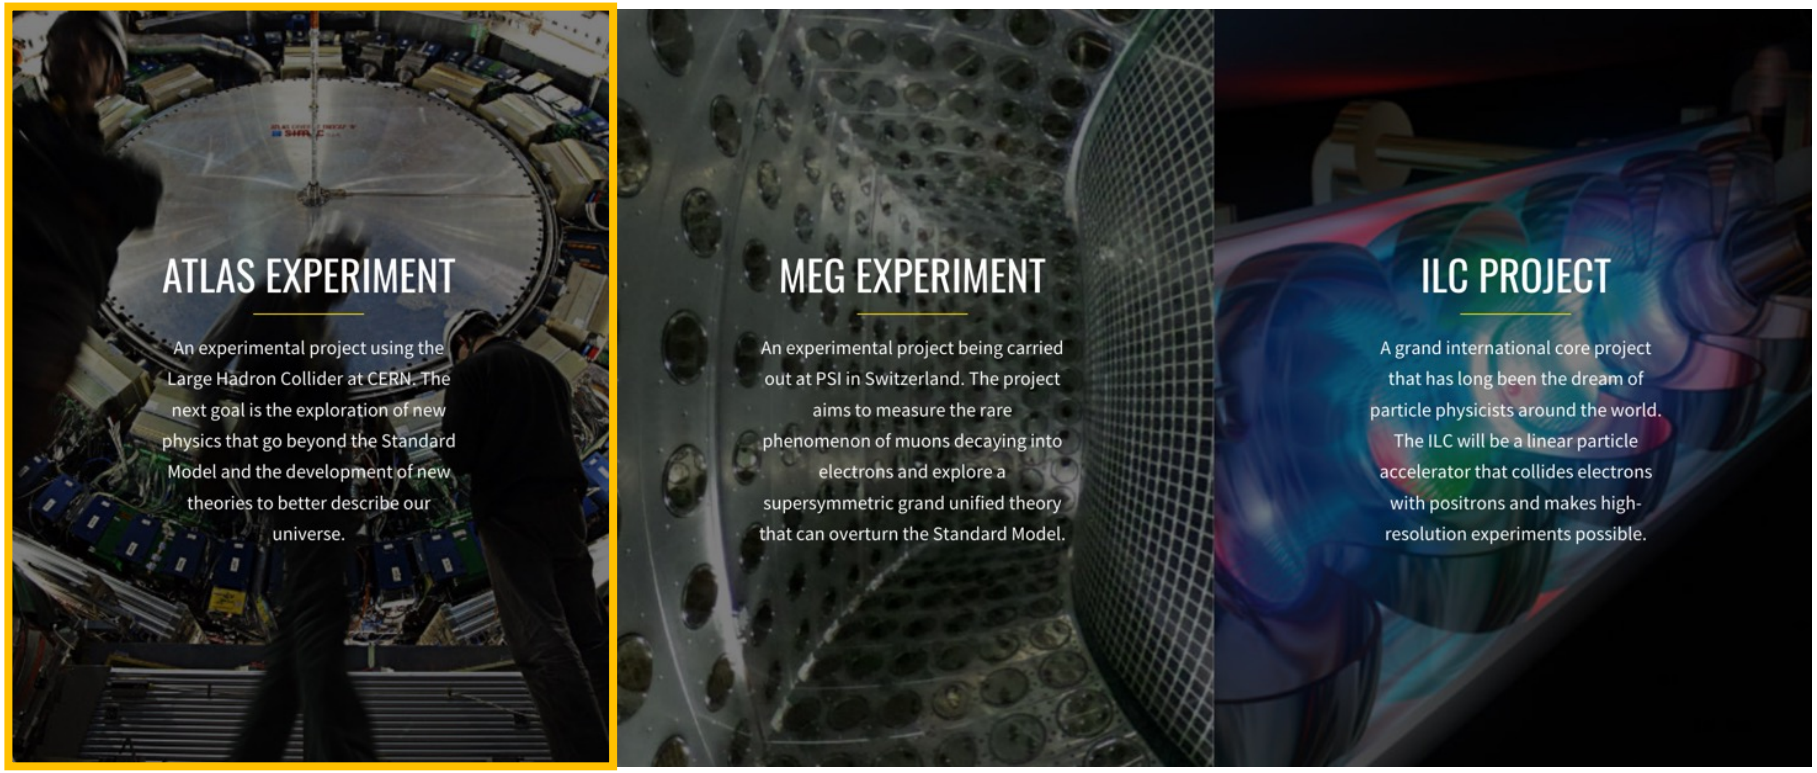

ATLAS-Japan group

- 13 institutes and ~160 members (45 members from ICEPP)
- Contributes to a wide area of the experiment
  - muon triggers, silicon tracker, **Tier2 operation**

➡ ICEPP operates **Tokyo regional analysis center** for ATLAS/ATLAS-Japan

ATLAS Japan

Tokyo regional analysis center

- Support ATLAS VO in WLCG (Tier2) and provide ATLAS-Japan dedicated resources (Tier3)
  - The only WLCG site in ATLAS-Japan

- Hardware is leased and replaced every three years. The current (6th) system started in **January 2022**.

- Tier2 (WLCG) (focus on this presentation)
  - Worker nodes (ARC/HTCondor): ~11,000 cores
  - Storage (DPM): ~15 PB

- Tier3 (ATLAS-Japan)
  - Interactive nodes: ~ 200 cores
  - Worker nodes (HTCondor): ~ 1,800 cores
  - Storage (GPFS): 3 PB
  - GPU resources: V100, T4

Tier2

Provided resources for ATLAS as Tier2

- One of the biggest Tier2 sites
  - CPU: ~5% of all ATLAS resources
  - Disk: ~3% of all ATLAS resources
  - cf. ATLAS-Japan member ratio to author list is ~ 3%

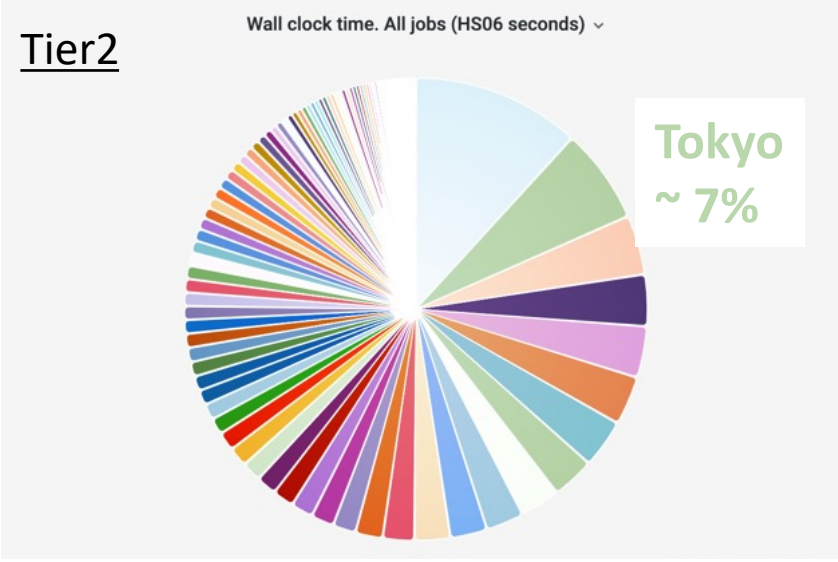

Compute resources

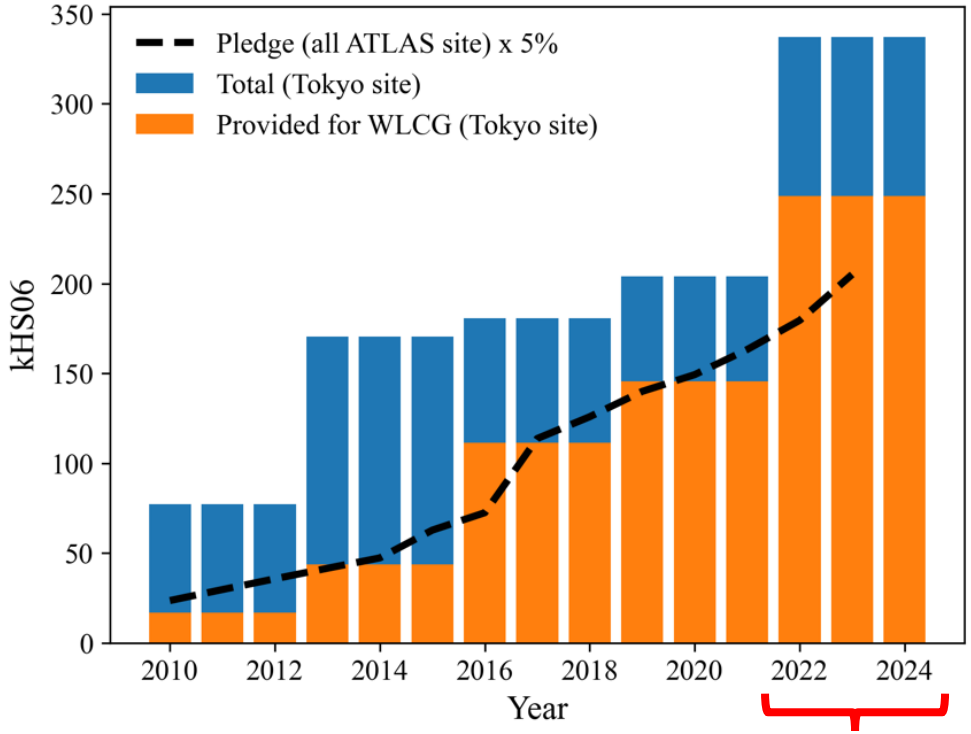

Storage resources

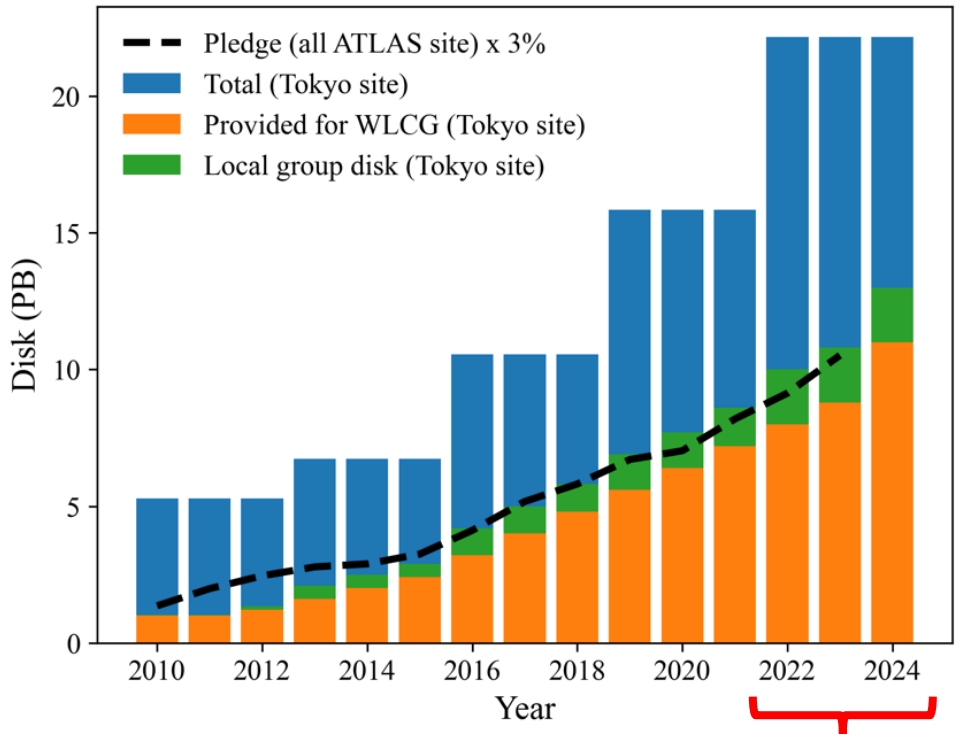

6th system: 2022 - 2024

Availability & Reliability

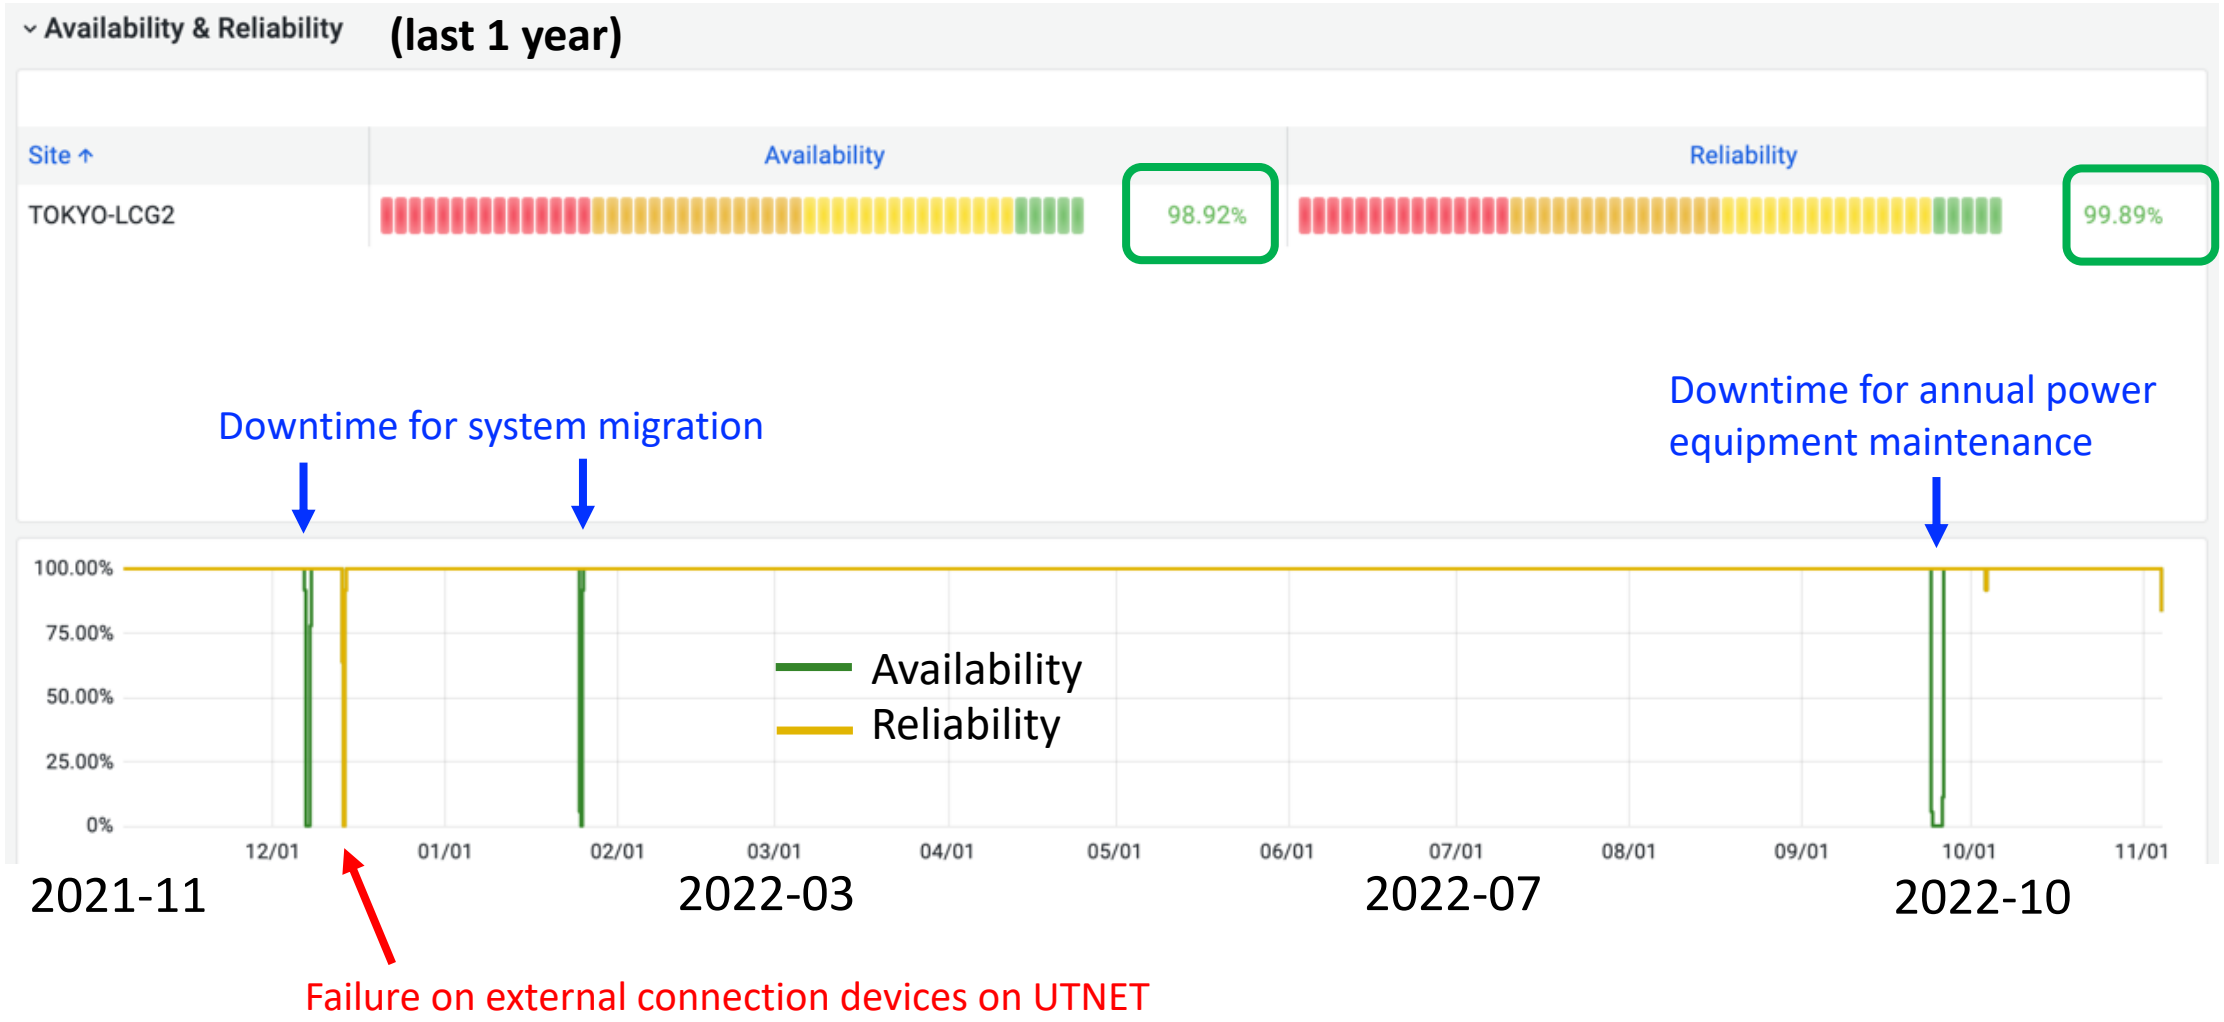

- Operating with high availability (~99%) and reliability (~99.9%)

Computing element (CE)

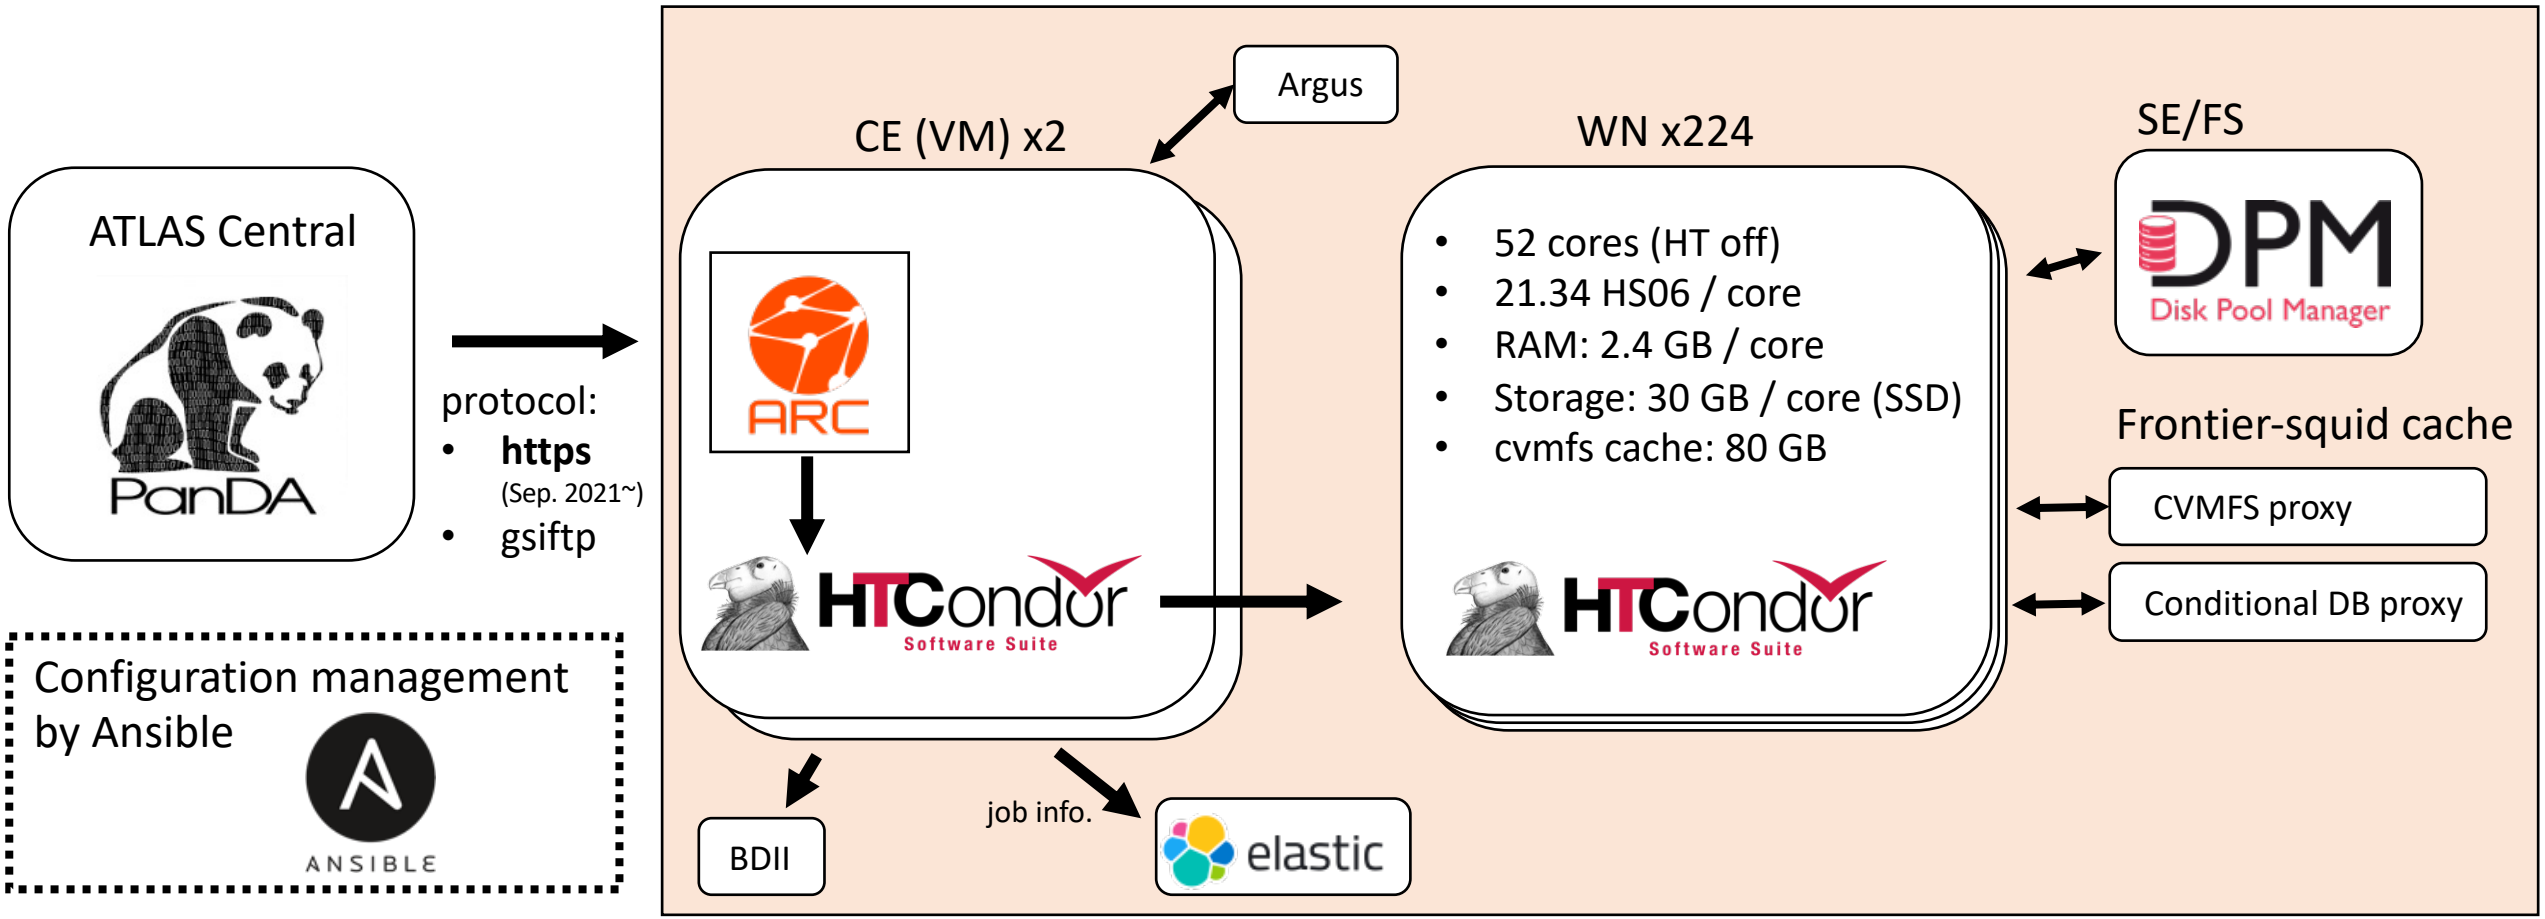

- Grid middleware
  - ARC-CE: Grid front-end
  - HTCondor: Job scheduler
- Two CEs
  - For redundancy
  - Kernel/package update without downtime (rolling update)

Computing element (CE)

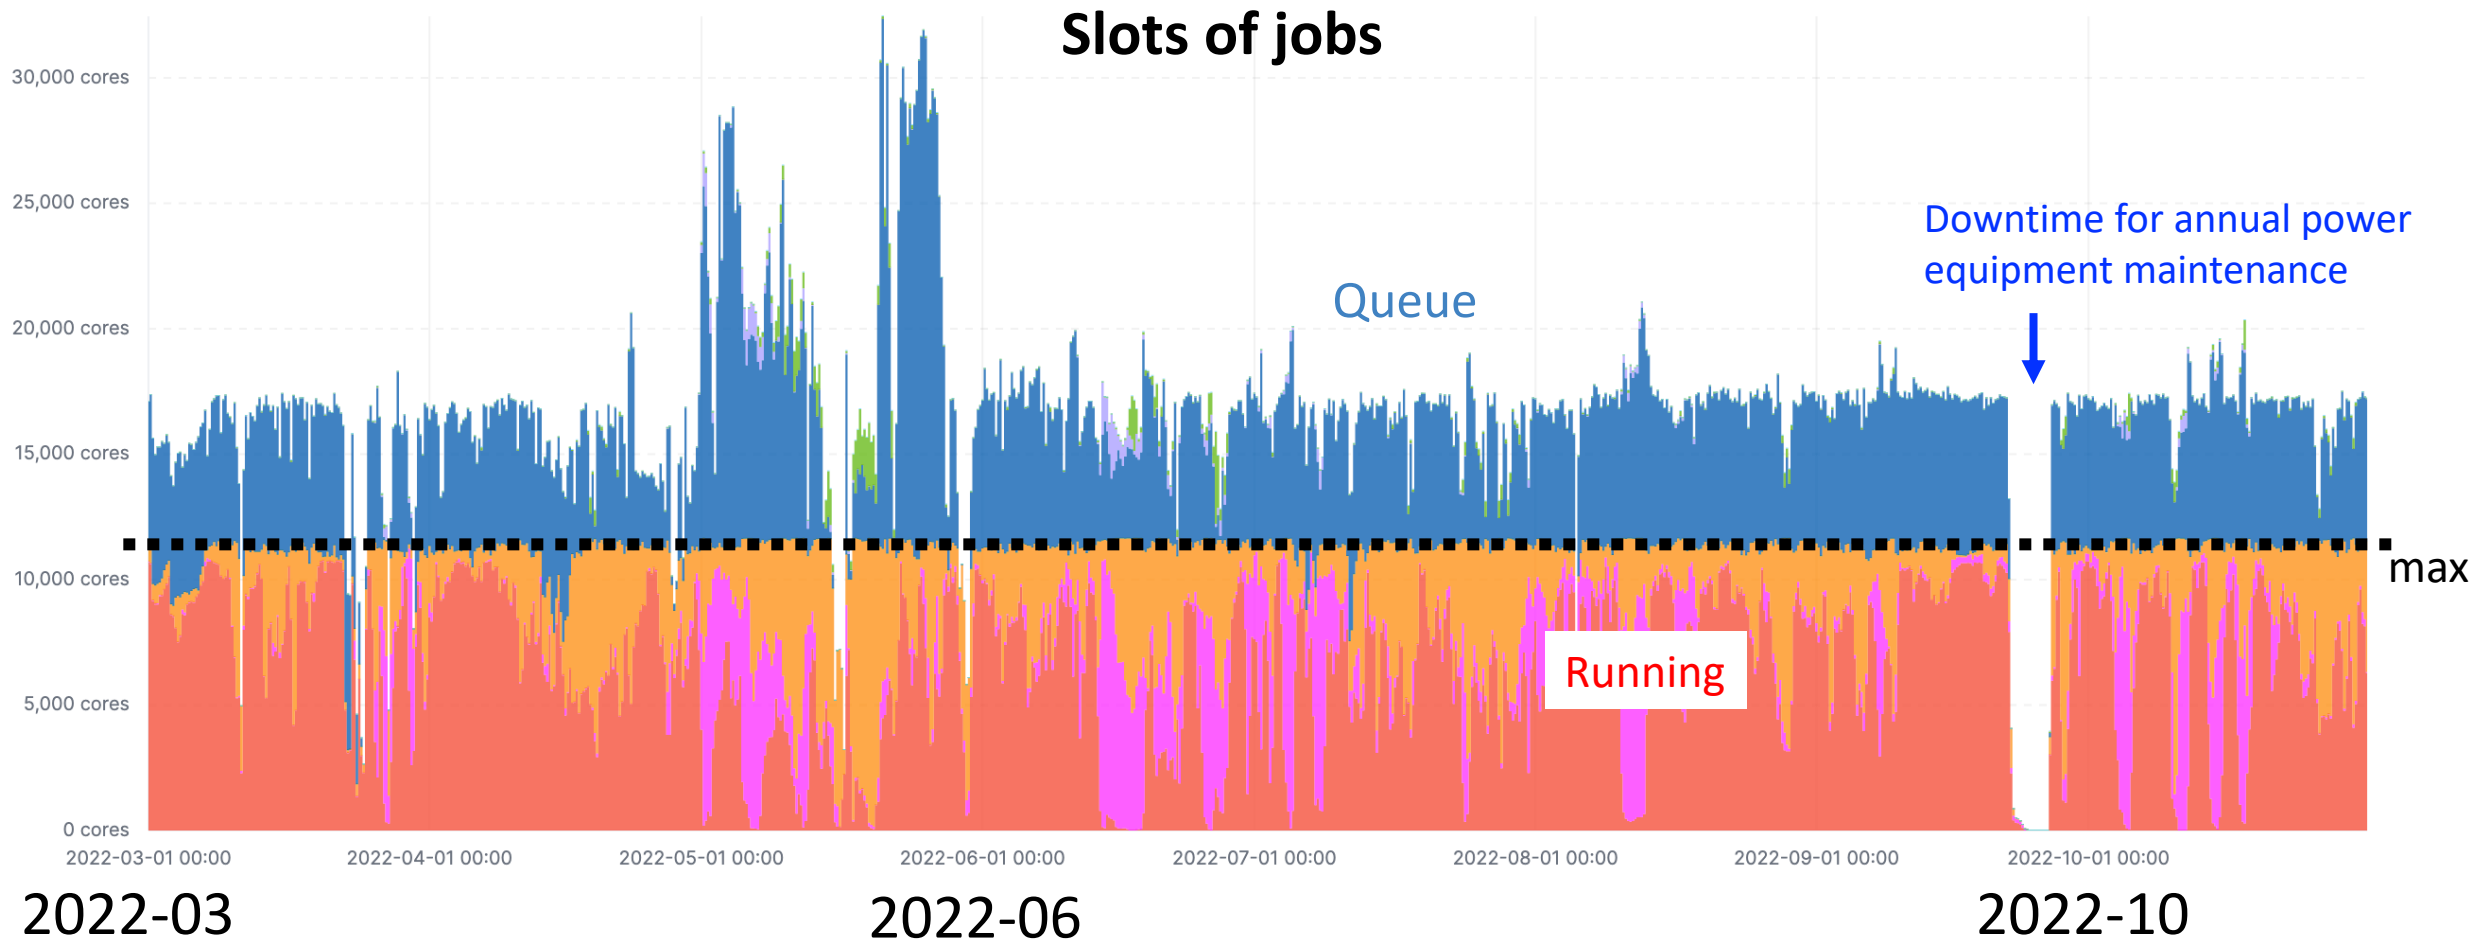

- 11,000 CPU cores are running almost constantly.
- Updated packages without downtime by rolling updates.

Storage element (SE)

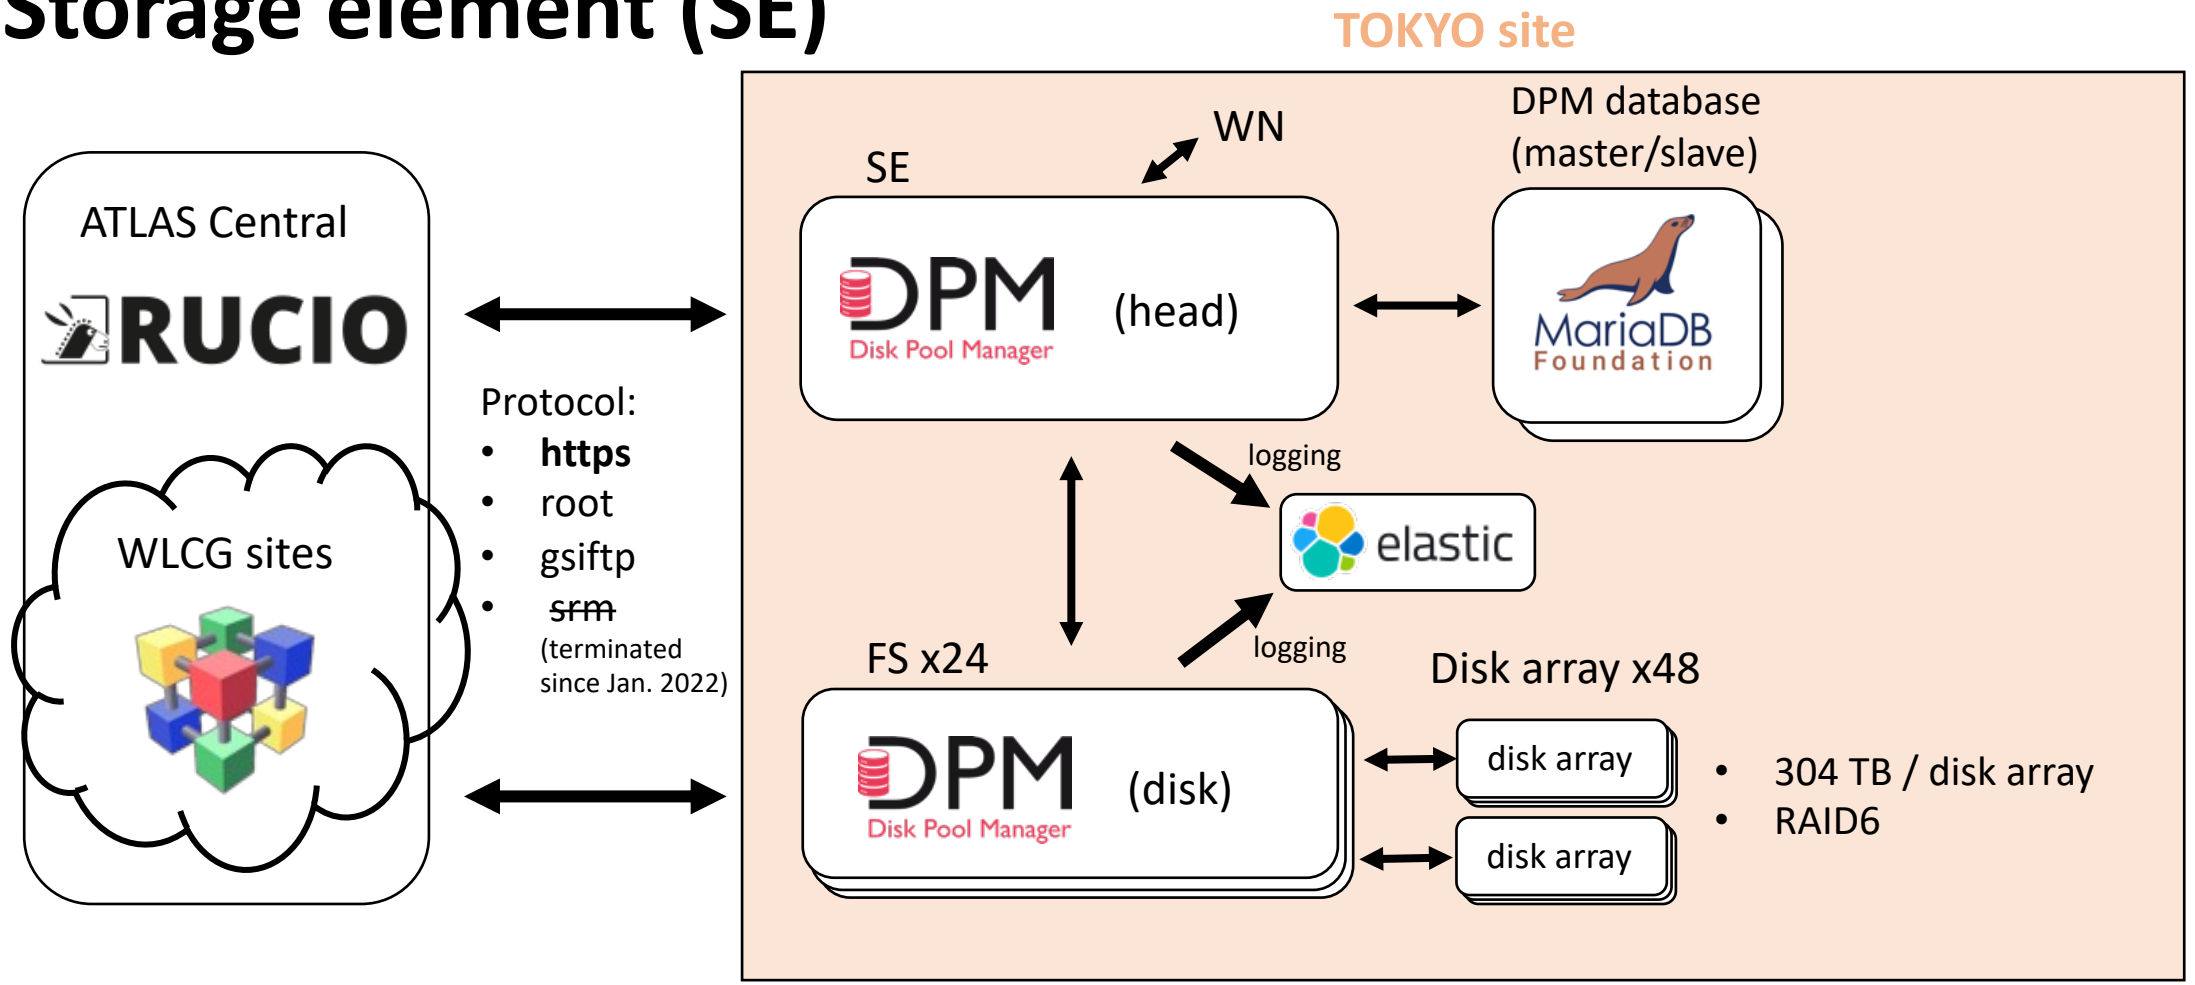

- Grid middleware
  - DPM
    - Plan to move to dCache
- Disk storage
  - 14.6 PB (provided 8 PB)

Storage element (SE)

The 5th system installed

The 6th system installed

Storage volume provided for ATLAS DATADISK

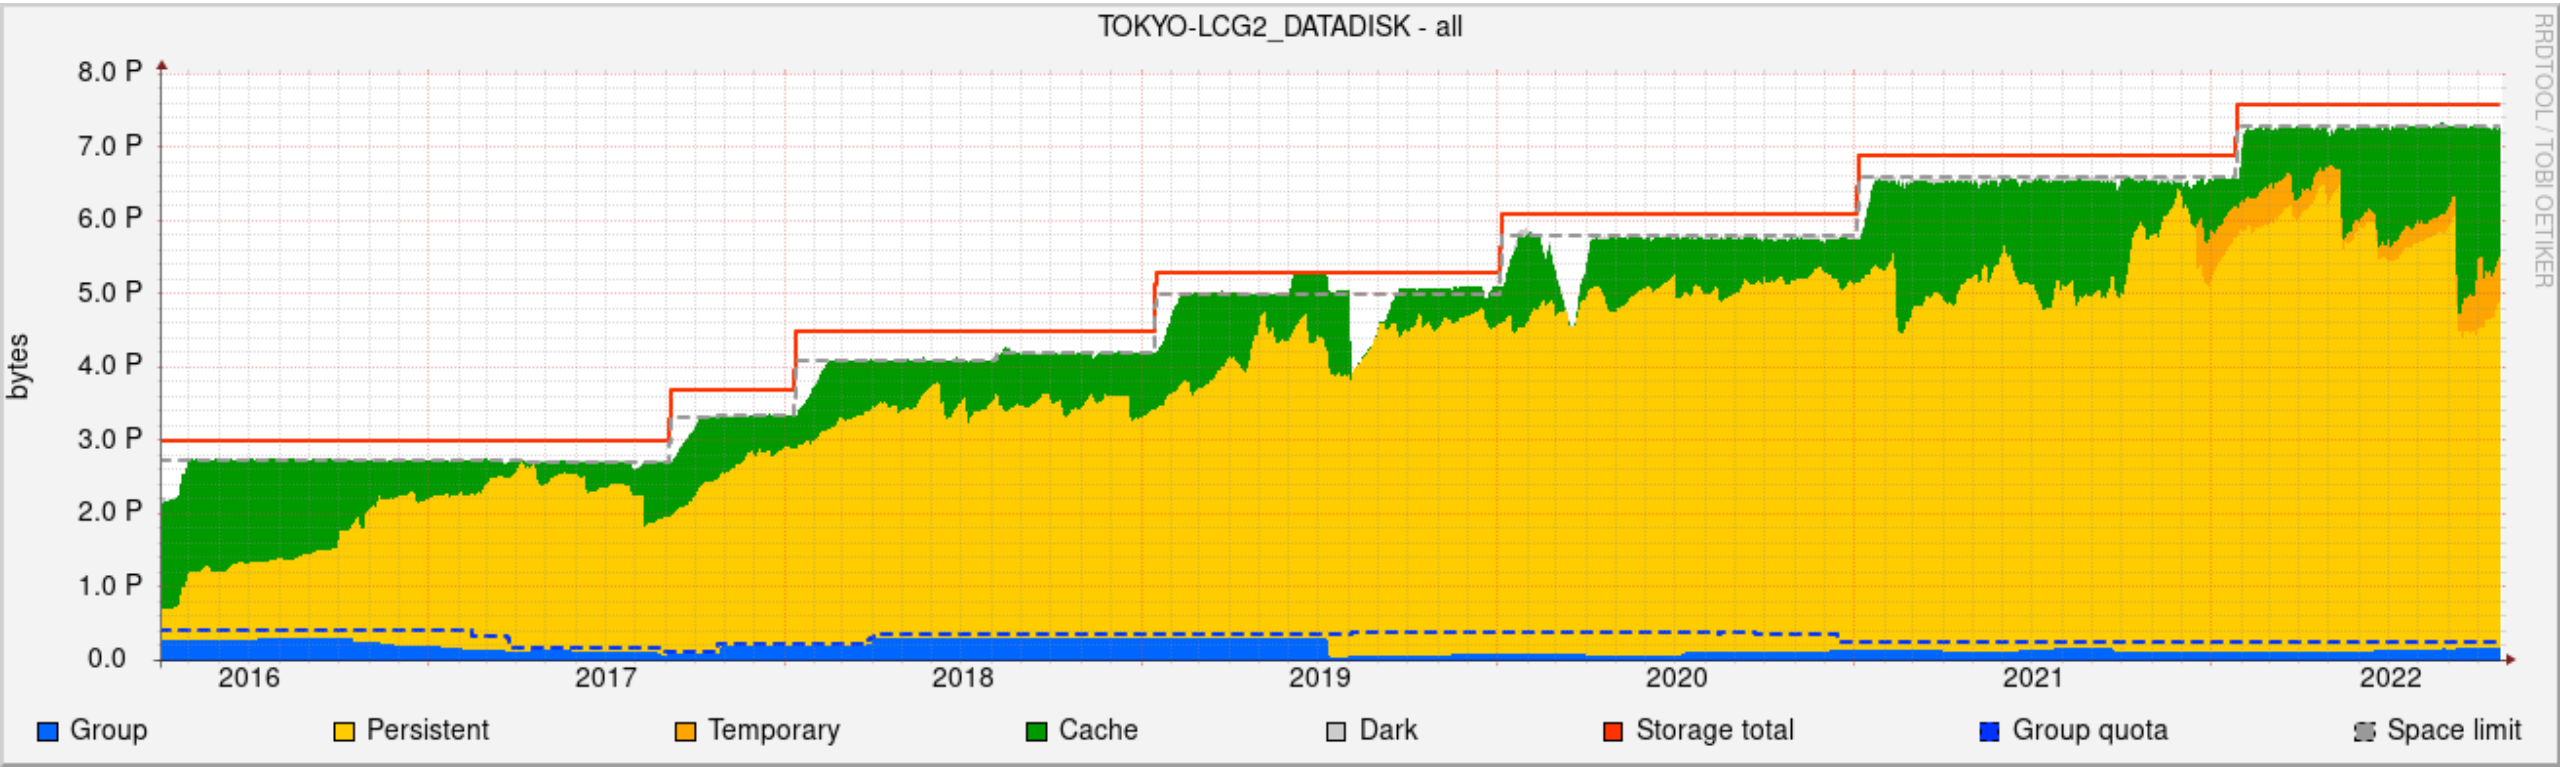

- Increase provided storage volume year by year
- Almost all of the quota are used

DPM EOL

2006 – 2022(?)

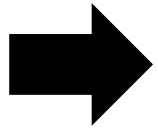

2023(?) -

- DPM EOL: summer 2024
- Plan to migrate to dCache next winter

- We are (probably) the biggest DPM site.
  - 8 PB, 70 M objects

```
> info /dpm/icepp.jp/home
/dpm/icepp.jp/home
------------------
File type:   Folder
Size:        8616297582060998B
Status:      Online
```

~8PB

```
MariaDB [cns_db]> select count(*) from Cns_file_metadata;
+-----------+
| count(*) |
+-----------+
| 69348759 |
+-----------+
```

~70M objects

- Preparation for migration is ongoing
  - Updated DPM to the latest version 1.15.2 (which includes migration script)
  - Applied DPM “DB inconsistency check”
    - Took 6 hours (actual wall time ~ 2h)
    - Found 20k dark objects & 750 lost objects. Most of them were old files (before 2013) or test files. Deleted/cleaned them after confirming deleting the files is OK.
  - Calculated the missing checksum: 275 k files, 16.8 TB. Took 6 hours.

Network

- External (WAN):
  - Connected to SINET with 40 Gbps (since 2019.10)
    - plan to upgrade to 100 Gbps within a few months.
  - SINET6 international network
    - Tokyo – Amsterdam: 100 Gbps
    - Tokyo – Los Angeles: 100 Gbps x 2
    - Tokyo – Singapore: 100 Gbps

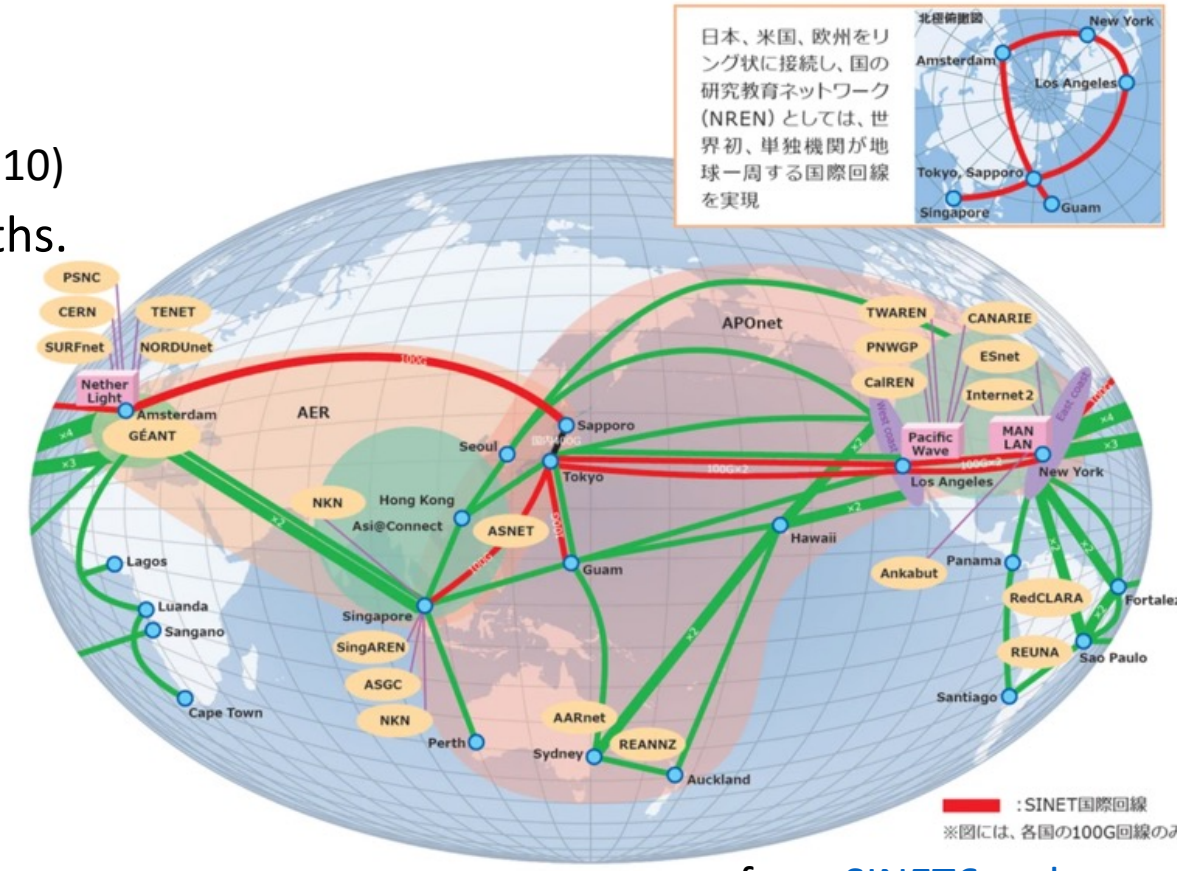

from SINET6 webpage

- Internal (LAN):
  - Core switch ← (25 GbE) → File server
  - Core switch ← (40 GbE) → Edge switch ← (10 GbE x 16) → Worker node x 16

Network (WAN)

DPM file servers ↔ LHCONE/Internet

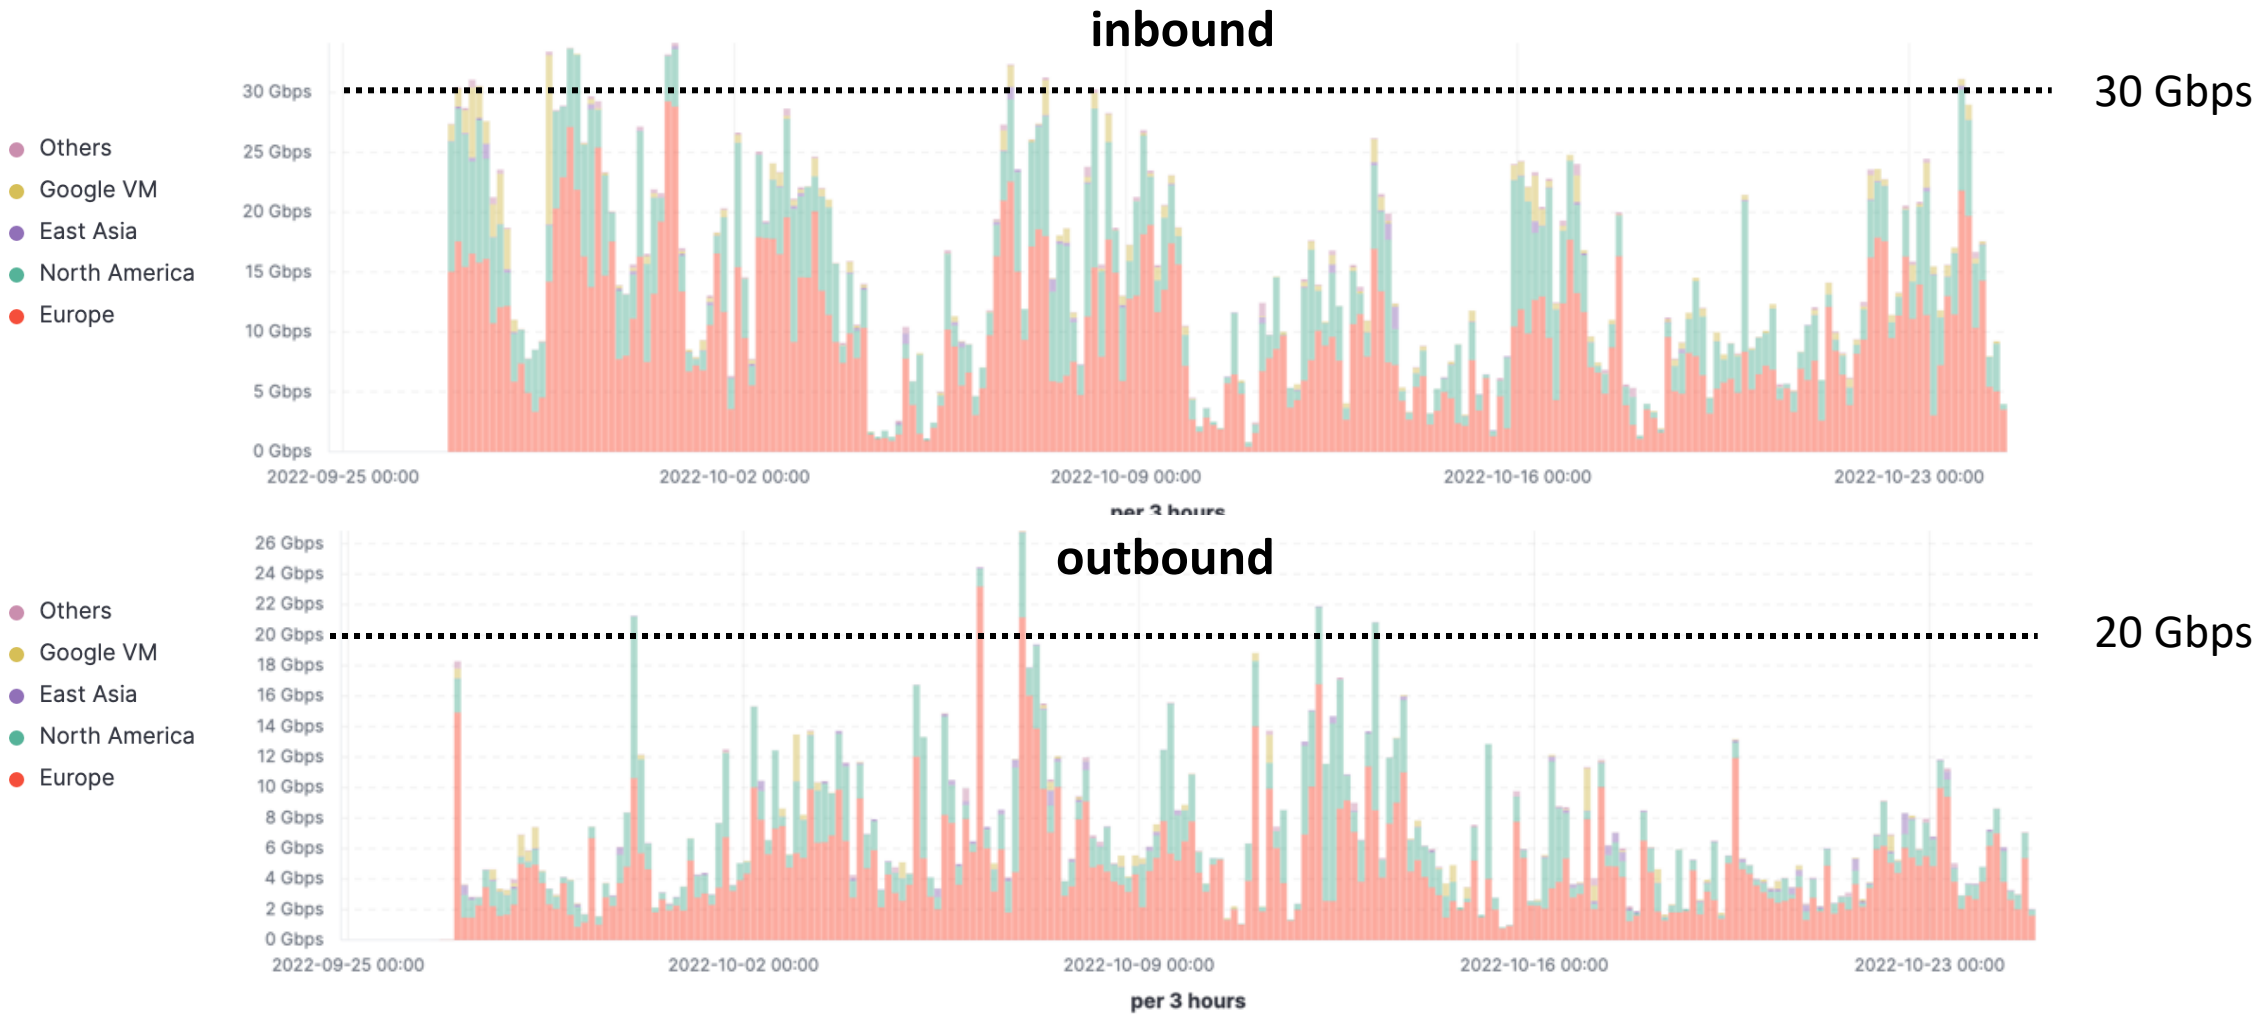

- Data transfer volume: ~ 120 TB / day
- Dominant transfer region is Europe, followed by North America.
- Sometimes transfer rate is rate-limiting at 40 Gbps. Expected to improve by the upgrade to 100 Gbps.

Network (LAN)

DPM file servers ↔ ICEPP host machines

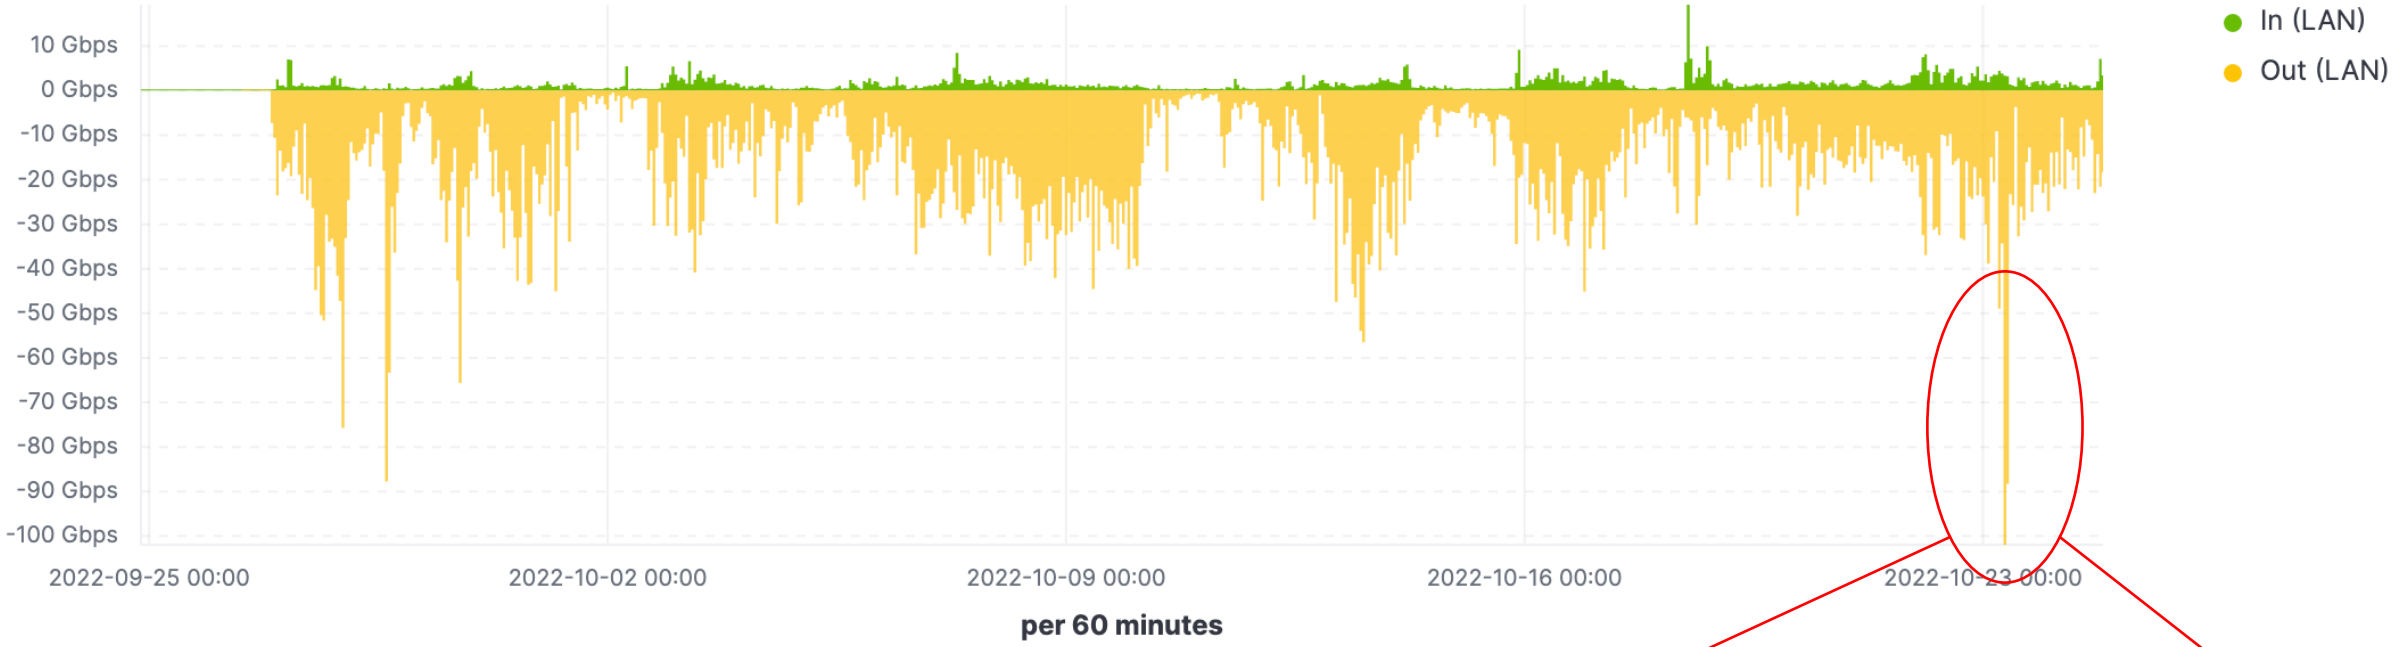

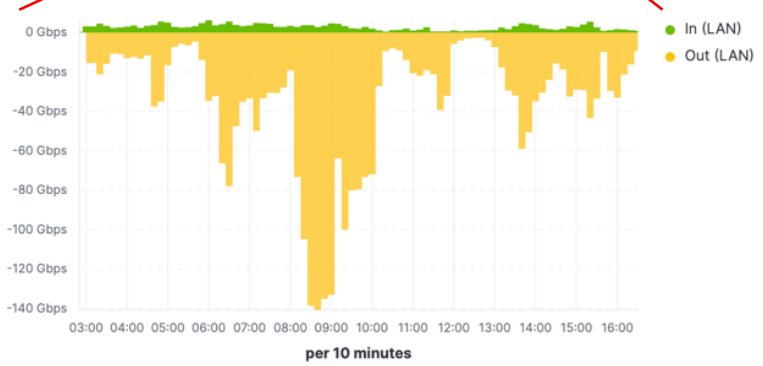

- Most LAN connections are from worker nodes.
- Constantly 10 ~ 30 Gbps, peak: 140 Gbps

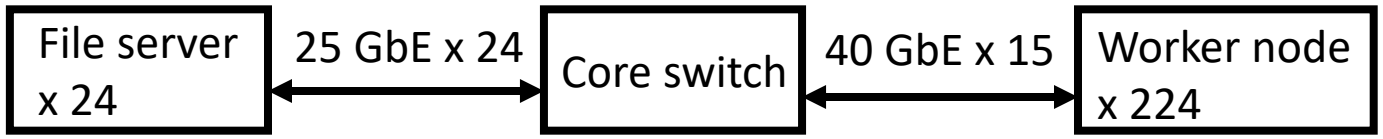

System migration

from the 5th system (2019 - 2021)
to the 6th system (2022 - 2024)

System migration (Dec. 2021 – Feb. 2022)

- Hardware is replaced every three years.
  - need to migrate all head nodes, worker nodes, storage
- To avoid a long downtime, we set a “reduced system” phase (~ 2 months)
  - All services run on the 5th system’s hardware.
    - Reduced worker nodes. Gradually added worker nodes of the new hardware after being ready.
    - Reduced network bandwidth for production file servers (25 Gbps → 10 Gbps)
  - Copied data (DPM/GPFS, ~8.5 PB) from old disks to new disks as a background process.

		Before 07 Dec. 2021	Reduced system phase	After 25 Jan. 2022 (Tier2) 7 Feb. 2022 (Tier3)
Head nodes	5th	**prod @ main rack**	**prod @** tmp rack	-
	6th	setting @ outside	setting @ main rack	**prod @ main rack**
Worker nodes	5th	**prod @ main rack**	**prod @ tmp rack** (reduced)	-
	6th	setting @ outside	**prod @ main rack** (added)	**prod @ main rack**
File servers	5th	**prod @ main rack**	**prod @ main rack**	-
	6h	setting @ outside	setting @ main rack	**prod @ main rack**

~ 1 month

System migration (Dec. 2021 – Feb. 2022)

1st Downtime 07 Dec. 2021, 14 hours

- Moved head nodes and a part of compute nodes to a temporary rack.
- Carried out almost the 5th system's servers except for storage servers and carried in the 6th system's servers.

2nd Downtime 25 Jan. 2022, 28 hours (Tier2)
6-7 Feb. 2022, 24 hours (Tier3)

- Migrated all head nodes and storage servers to the 6th hardware.
- Service configuration are deployed via Ansible.

		Before 07 Dec. 2021	Reduced system phase	After 25 Jan. 2022 (Tier2) 7 Feb. 2022 (Tier3)
Head nodes	5th	**prod @ main rack**	**prod @** tmp rack	-
	6th	setting @ outside	setting @ main rack	**prod @ main rack**
Worker nodes	5th	**prod @ main rack**	**prod @ tmp rack** (reduced)	-
	6th	setting @ outside	**prod @ main rack** (added)	**prod @ main rack**
File servers	5th	**prod @ main rack**	**prod @ main rack**	-
	6h	setting @ outside	setting @ main rack	**prod @ main rack**

~ 1 month

System migration (Dec. 21 – Feb. 22)

5th system (- Nov. 2021)

Computer room

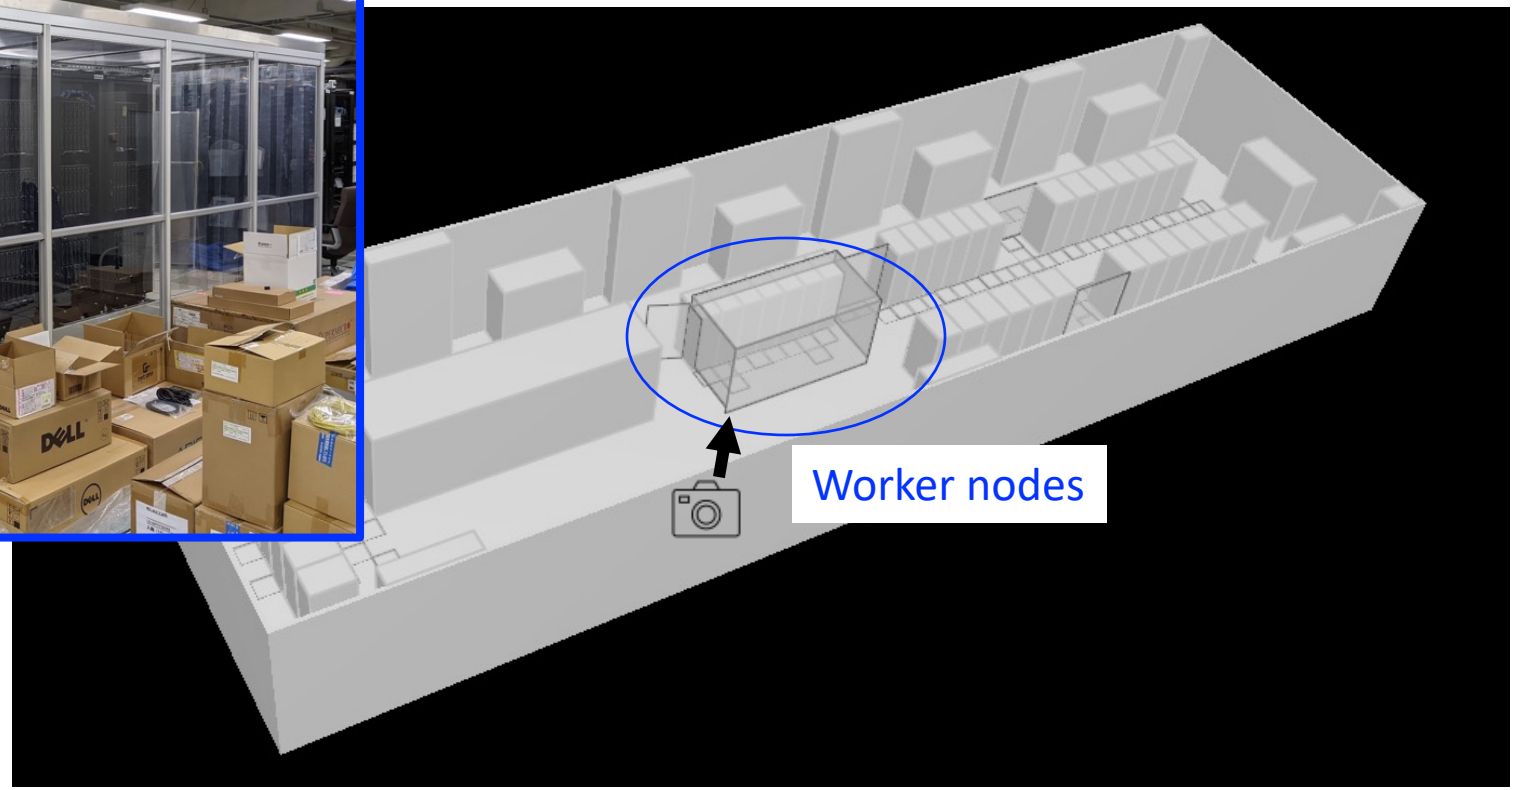

System migration (Dec. 21 – Feb. 22)

1st downtime

5th system (- Nov. 2021)

6th system (Feb. 2022 -)

Reduced system phase

During the reduced system phase, we operated with fewer servers

- 2 / 15 for Tier2
- 1 / 2 for Tier3

Running CPU cores during the reduced system phase (Tier2)

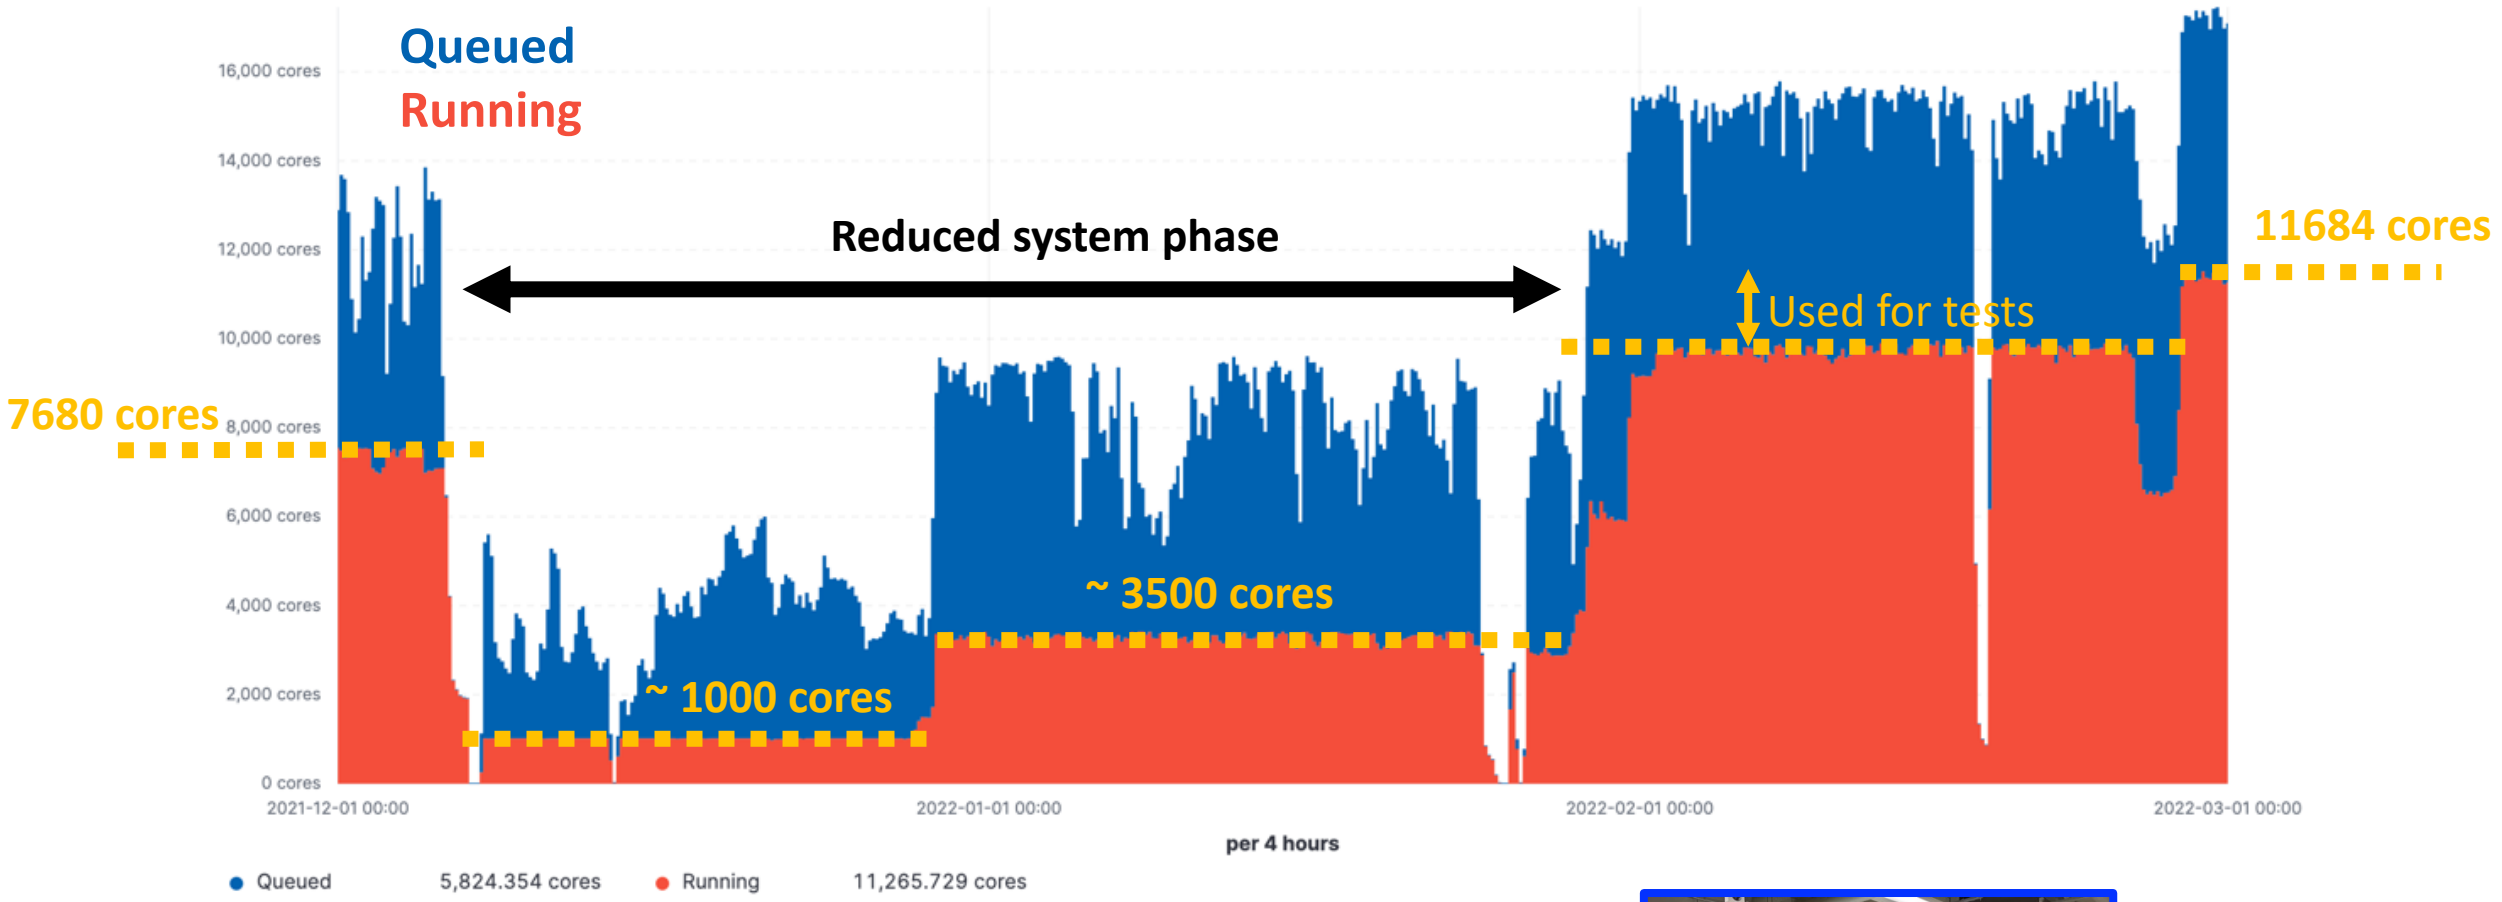

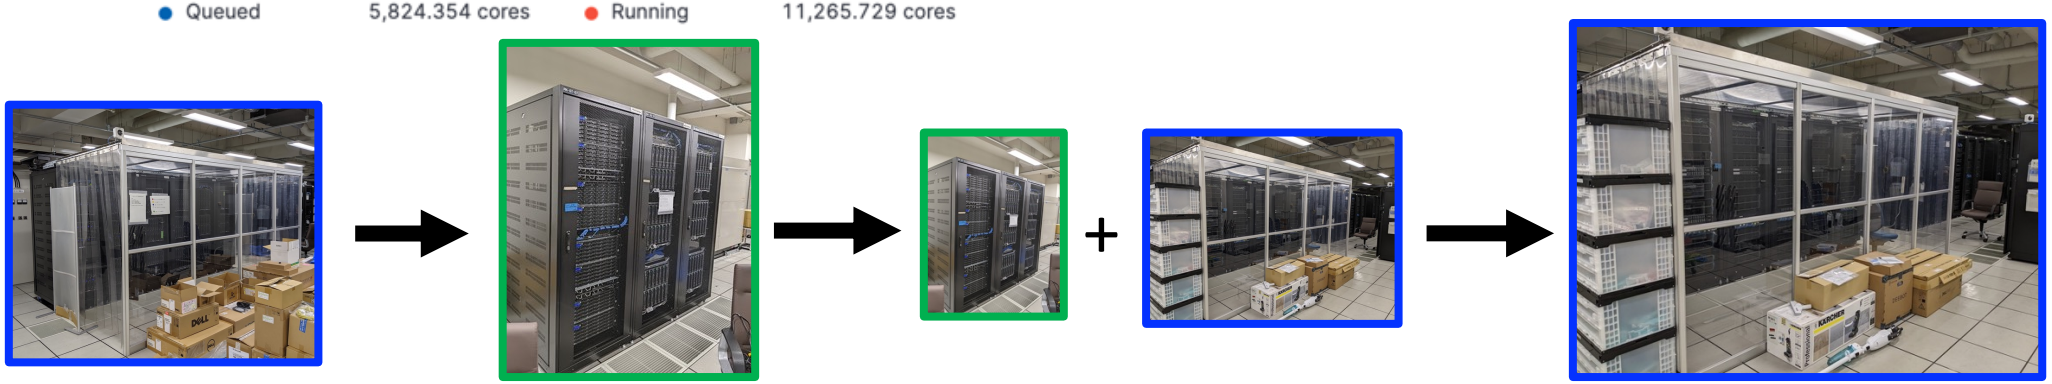

System migration (Dec. 21 – Feb. 22)

5th system (- Nov. 2021)

Computer room

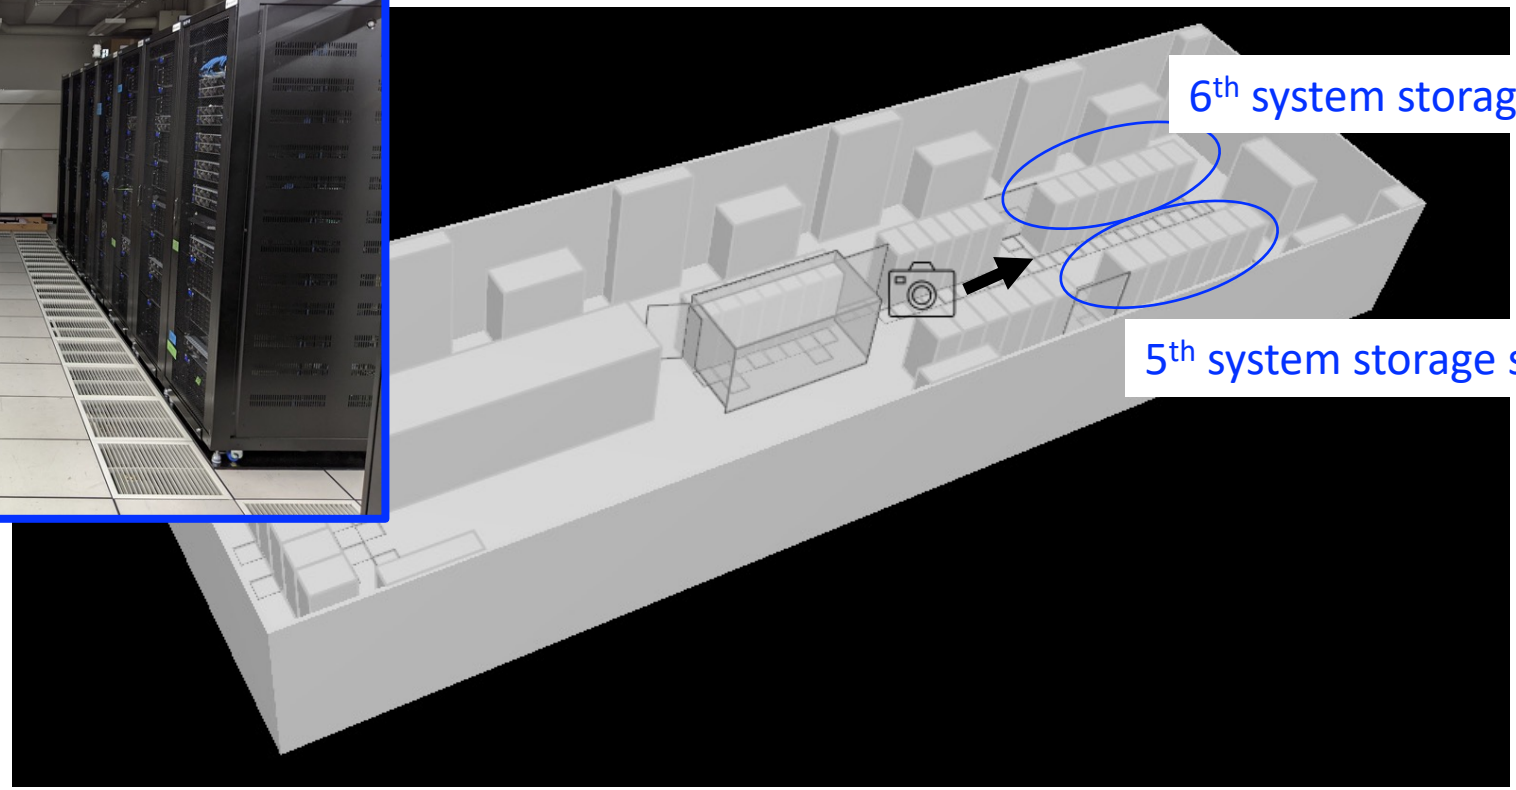

System migration (Dec. 21 – Feb. 22)

5th system (- Nov. 2021)

Computer room

Reduced system phase

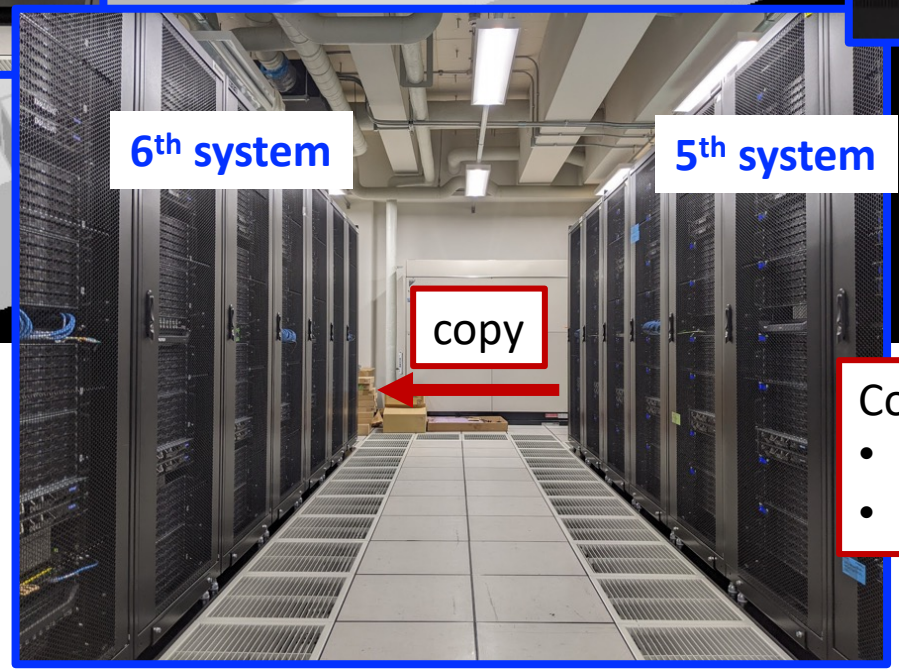

6th system (Feb. 2022 -)

Copied data during production operations
- Tier2: ~ 5 days to be completed (7.3 PB)
- Tier3: ~ 8 days to be completed (1.2 PB)

System migration (Dec. 21 – Feb. 22)

All services run on the 6th system's hardware.

WAN

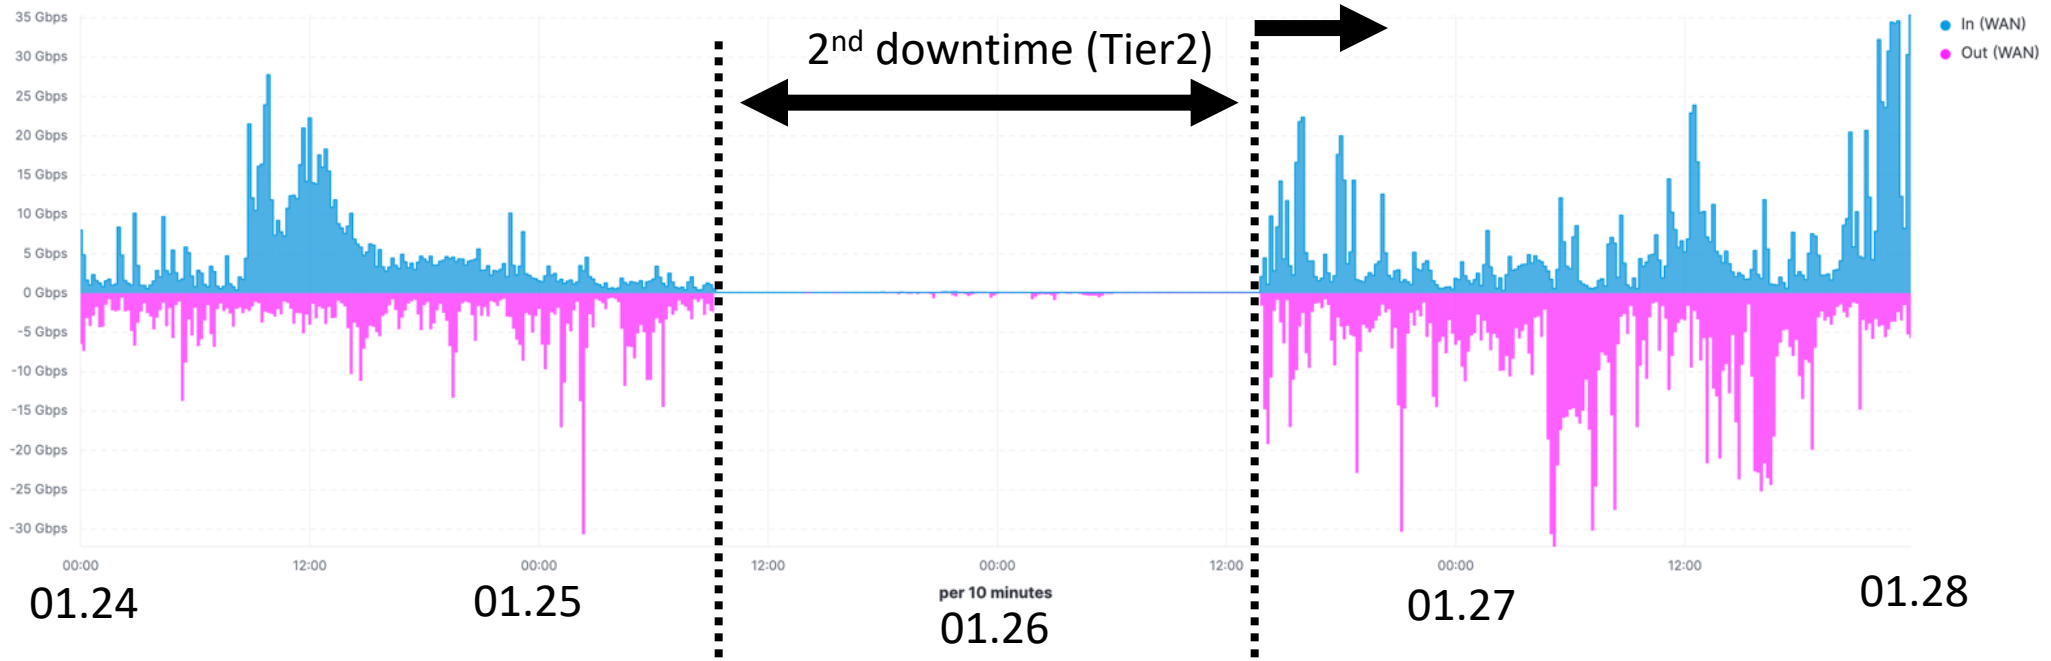

LAN

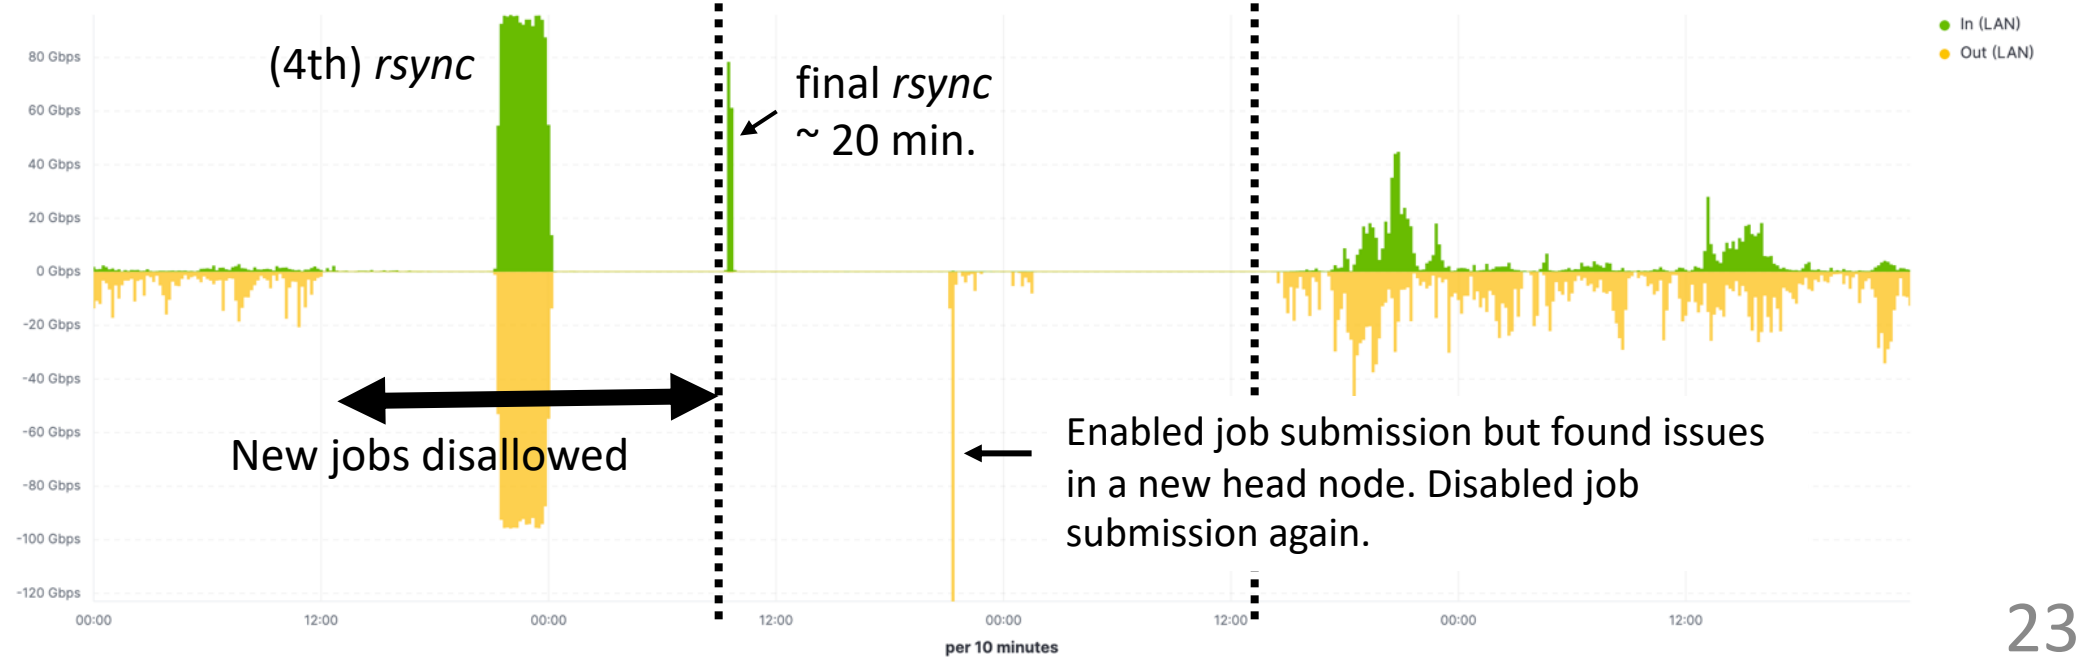

HPC resources

HPC

- HPC resource is one of the main contributors to ATLAS computing resources.

- ICEPP Tier2 site continues R&D for utilizing the HPC resources of the computer center for the University of Tokyo.
  1. Oakbridge-CX
     - Xeon Platinum 8280, 1,368 nodes
     - 6.61 PFLOPS

  → Test operation using production jobs

  2. Wisteria/BDEC01
     - A64FX, 7680 nodes
     - 33.1 PFLOPS

  → Performance study of A64FX

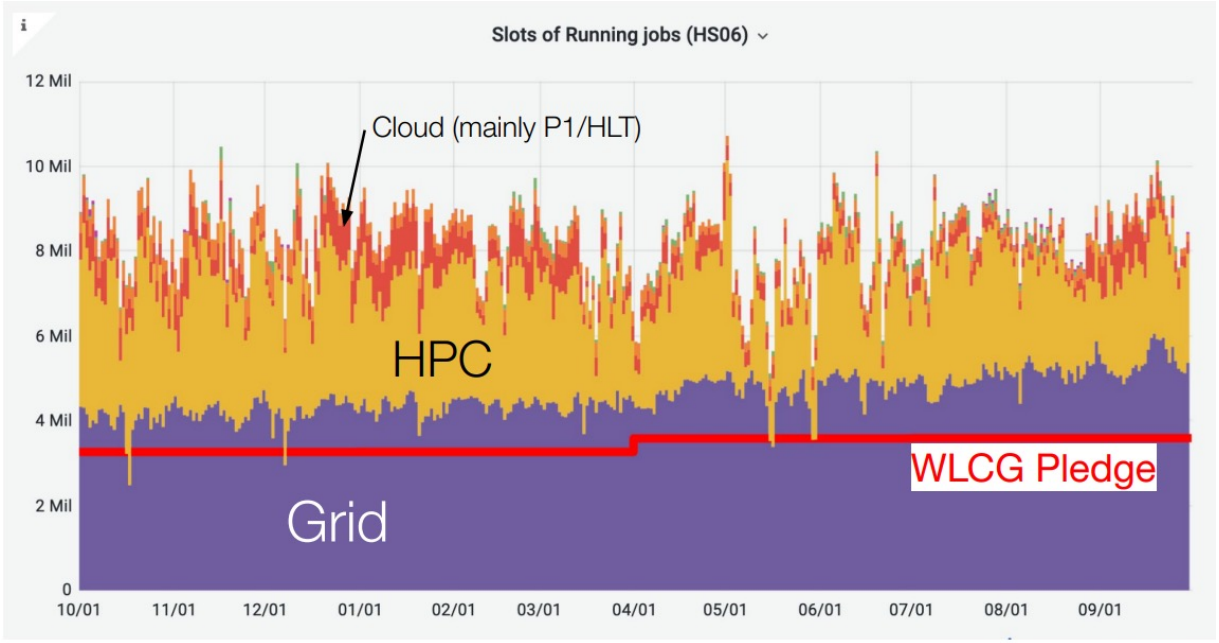

from “ATLAS @ WLCG Workshop 2022”

Fujitsu A64FX
- Armv8.2-A SVE
- 48 cores + assistant core
- used in 富嶽(Fugaku), which is the biggest HPC in Japan.

HPC (Oakbridge-CX)

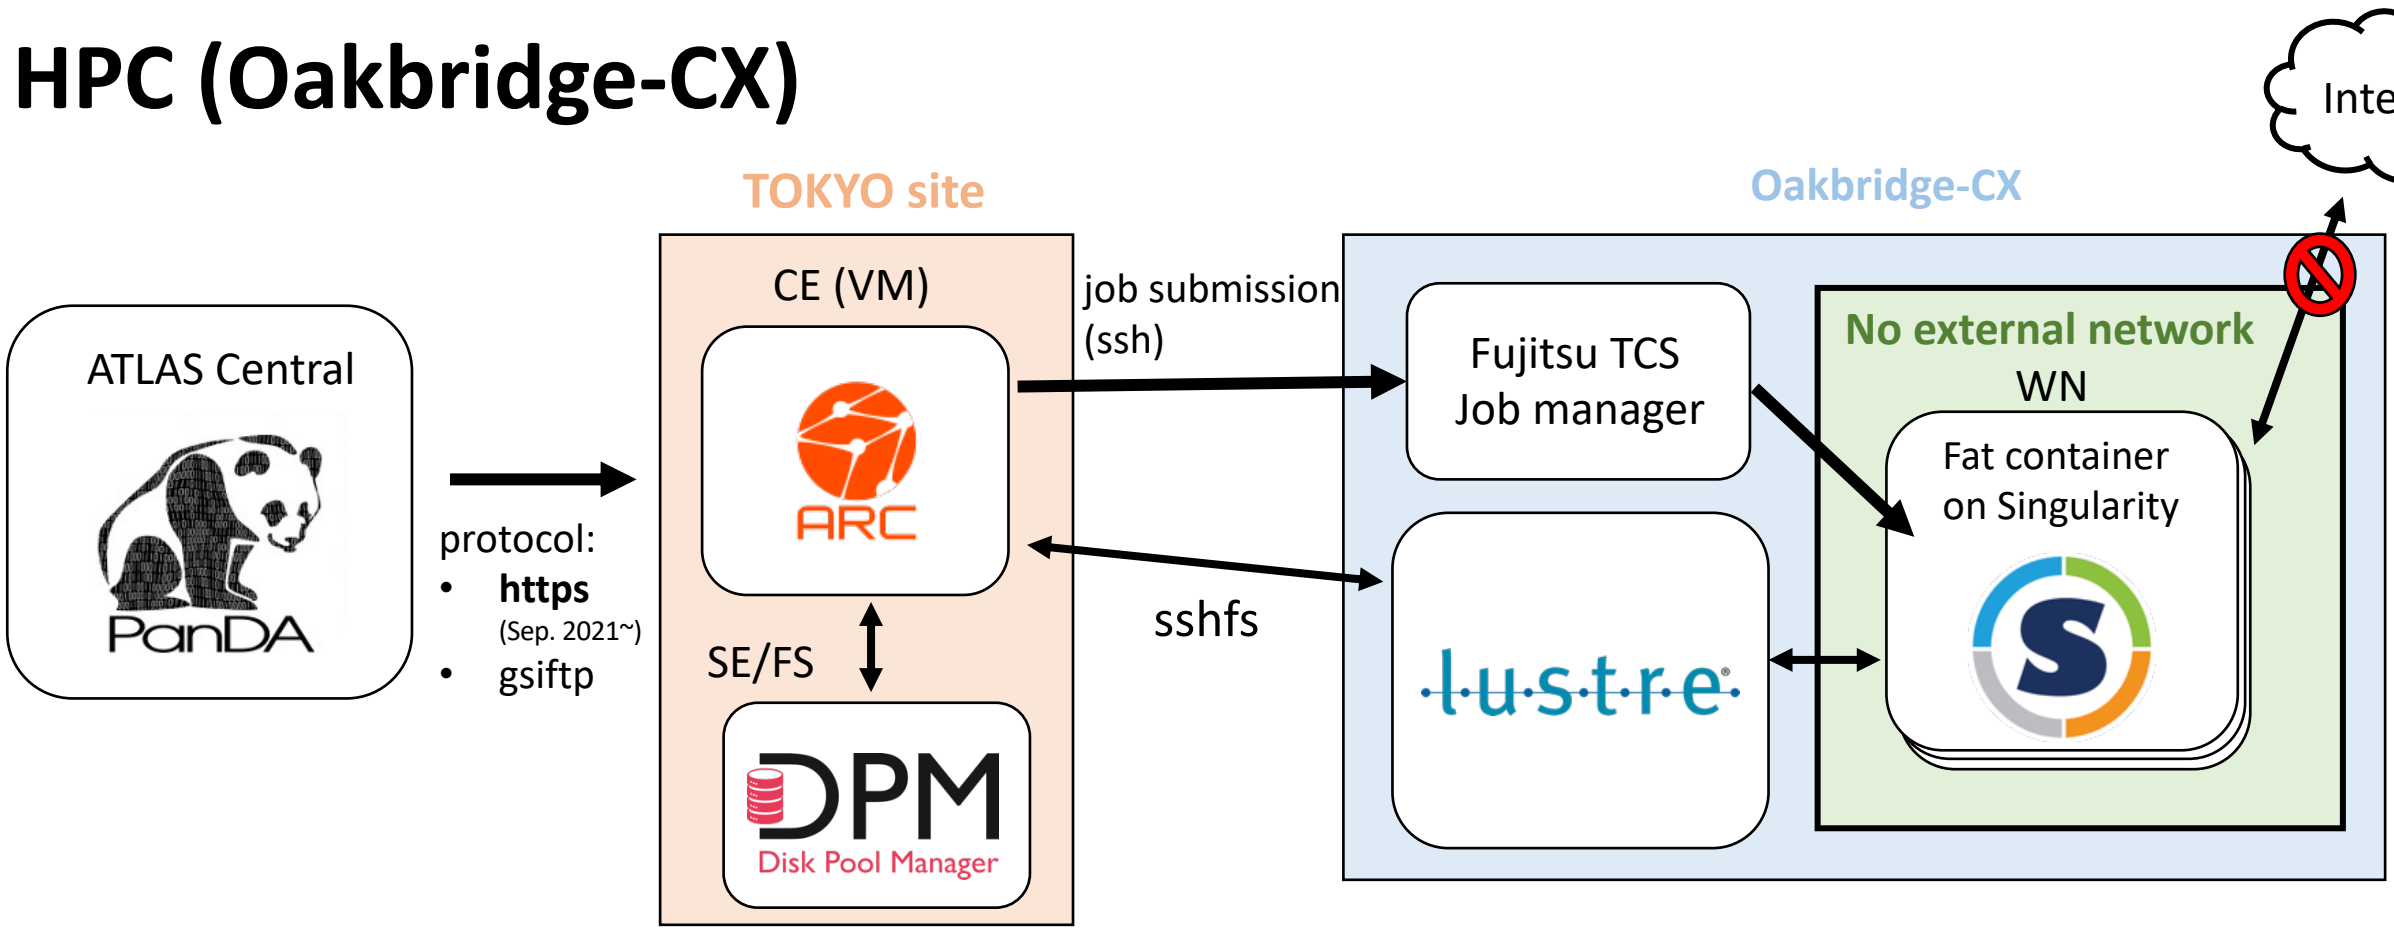

- Only internal connection is allowed
  - no CVMFS, conditional DB, SE access
- Use a fat singularity container image
  - contains all the necessary files
  - processes only simulation jobs

- Input/output files are transferred by ARC.
- ~3% of the CPU resources of the Tokyo Tier2 site

Running jobs in closed network environment

parrot_run in cctools + *cvmfs_preload*
- Pros: All athena(ATLAS software) releases on cvmfs are available
- Cons: Non negligible overhead

Apr 2022.

Fat container on *Singularity*
- Pros: Fast execution
- Cons: Only a few releases available

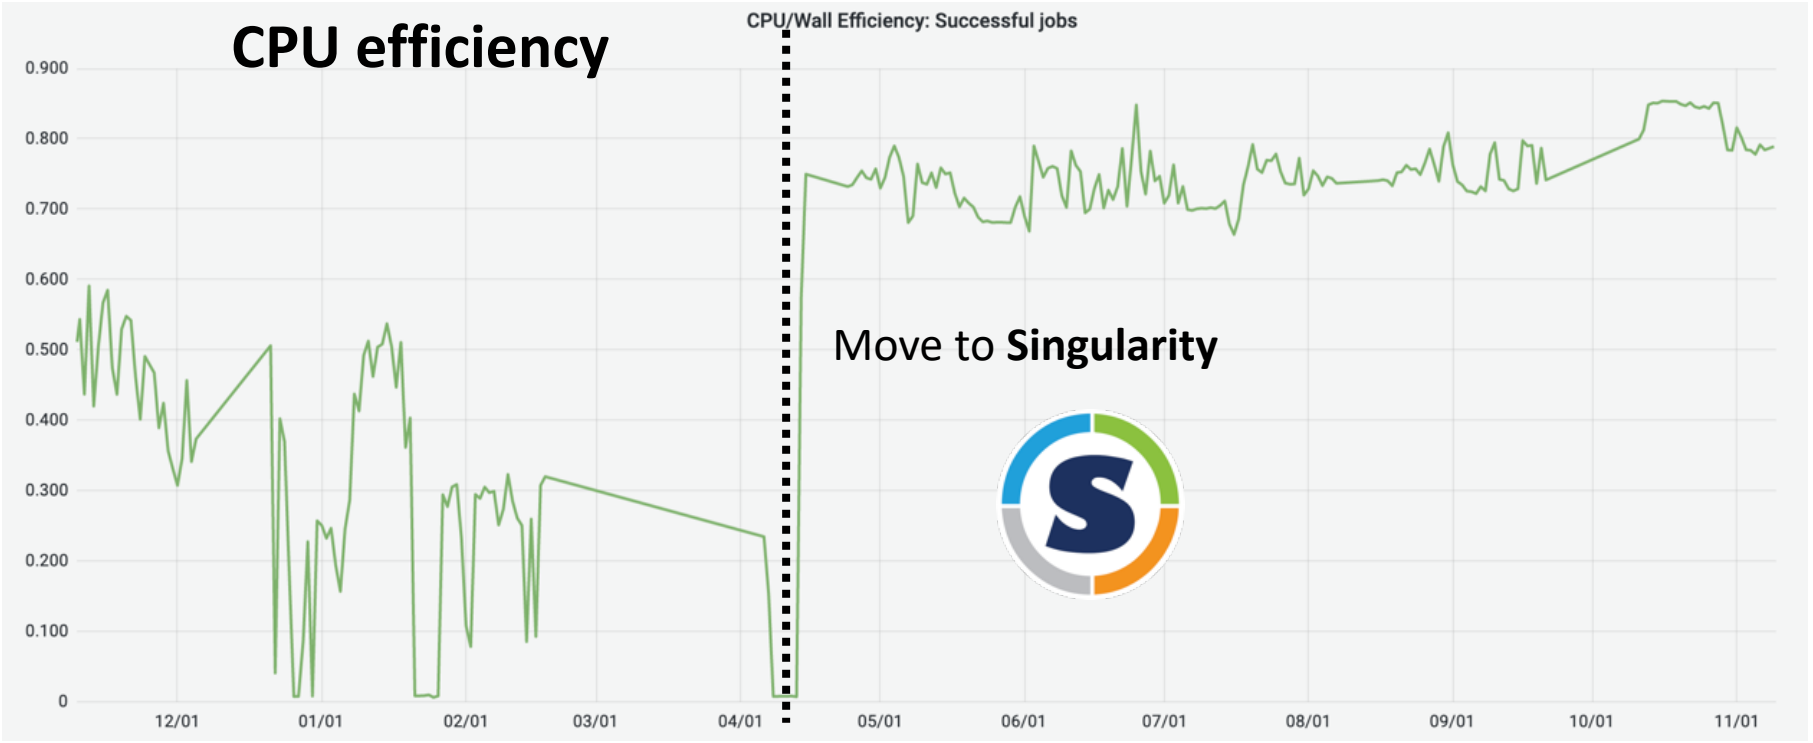

HPC (A64FX)

Testbed: PRIMEHPC FX700 (A64FX x 2)

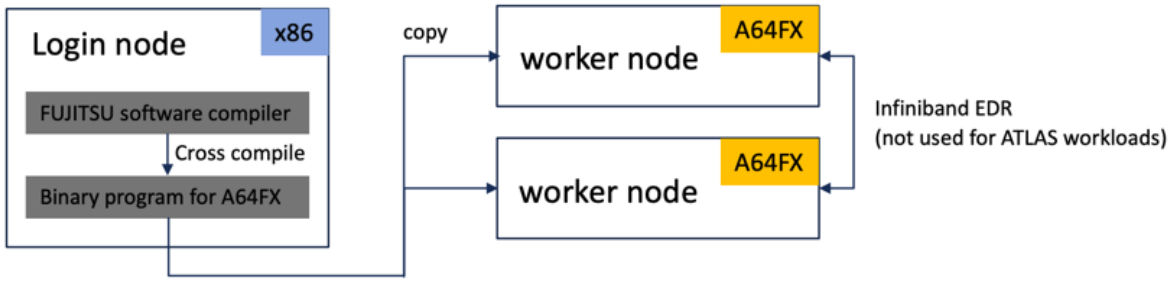

Throughput:

SysBench (prime number calc.)

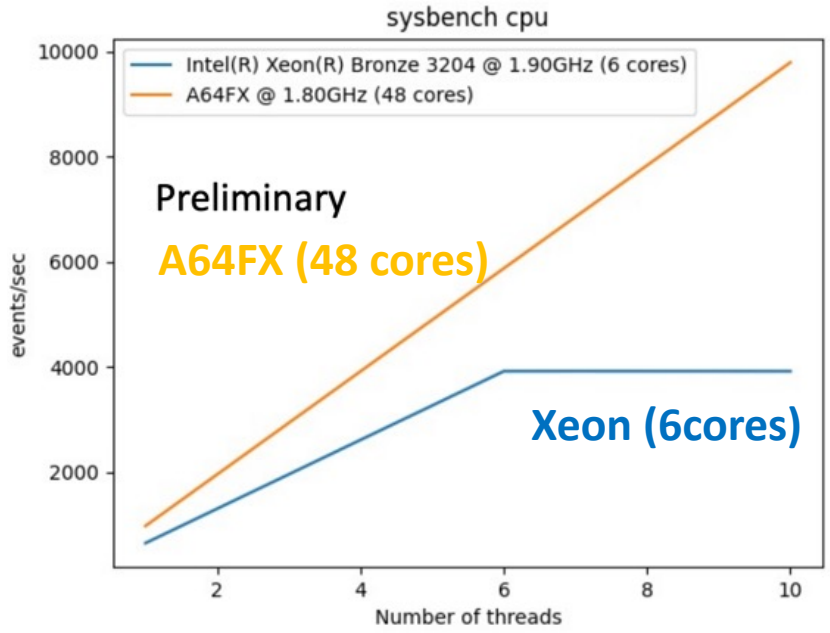

Geant4

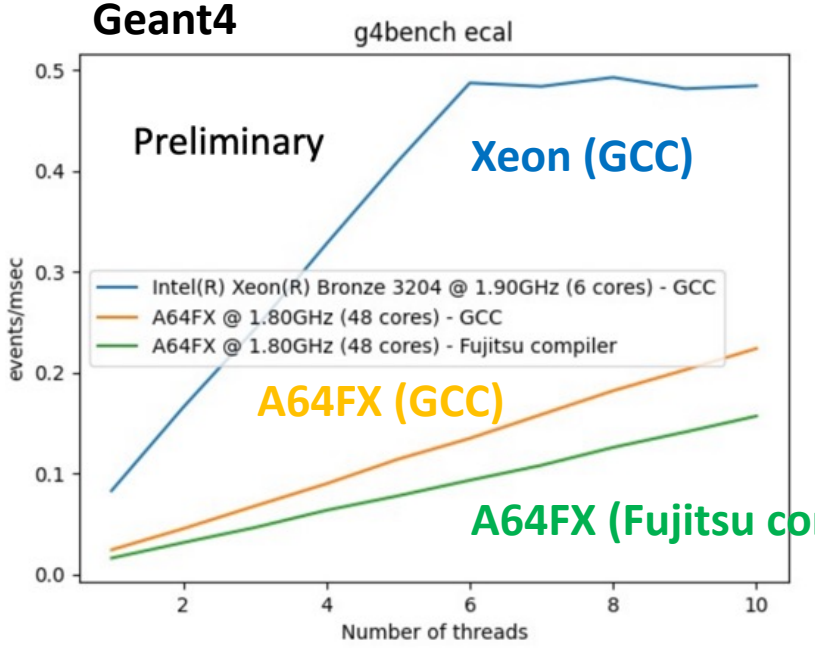

➡ Need to code optimization for Geant4 jobs with A64FX

Summary

- ICEPP regional analysis center is operating stably.
- Contributes to ~5% CPU and ~3% Disk of ATLAS sites
- Hardware replacement was completed successfully.
  - The next replacement is Dec. 2024.
- Near term upgrade plan
  - Storage middleware: DPM → dCache
  - External network: 40 Gbps → 100 Gbps

Backup

Network: IPv4/IPv6 dual stack

IP version to/from file servers

Inbound (WAN)

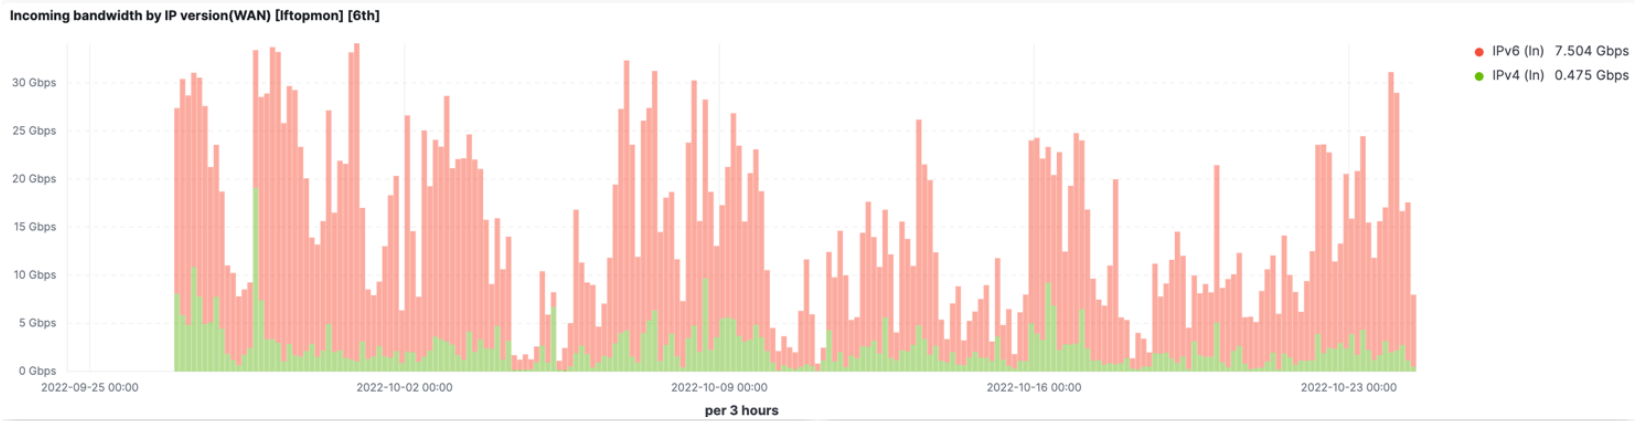

Outbound (WAN)

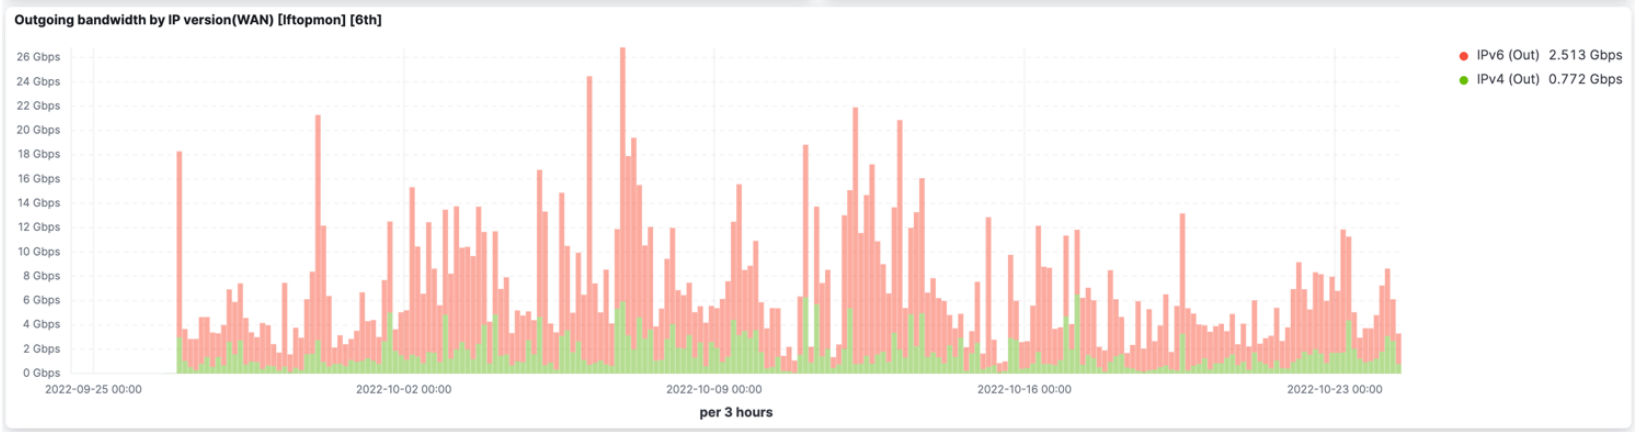

- Added IPv6 to file servers in 2019
- Now, IPv6 is a dominant protocol in data transfer for Tier2
- IPv4 is still used. (ICEPP worker nodes still use only IPv4.)

The 5th system vs the 6th system

		Total	For Tier2
CPU	5th system	336 nodes, 10752 cores (16 cores / CPU) Intel Xeon Gold 6130 2.10 GHz (Skylake) 204 kHS06 1.2 TB HDD x2 / node	240 nodes, 7680 cores 18.97 HS06 / core 3.0 GB RAM / core
	6th system	304 nodes, 15808 cores (26 cores / CPU) Intel Xeon Gold 5320 2.2 GHz (Icelake) 337 kHS06 1.92 TB SSD / node	224 nodes, 11648 cores 21.34 HS06 / core 2.5 GB RAM / core
Disk storage	5th system	72 disk arrays, RAID6 15,840 TB (10TB / HDD)	48 disk arrays, RAID6 10,560 TB (10TB / HDD)
	6th system	72 disk arrays, RAID6 22,176 TB (14 TB / HDD)	48 disk arrays, RAID6 14,784 TB (14 TB / HDD)

Network (LAN)

File servers ↔ File servers

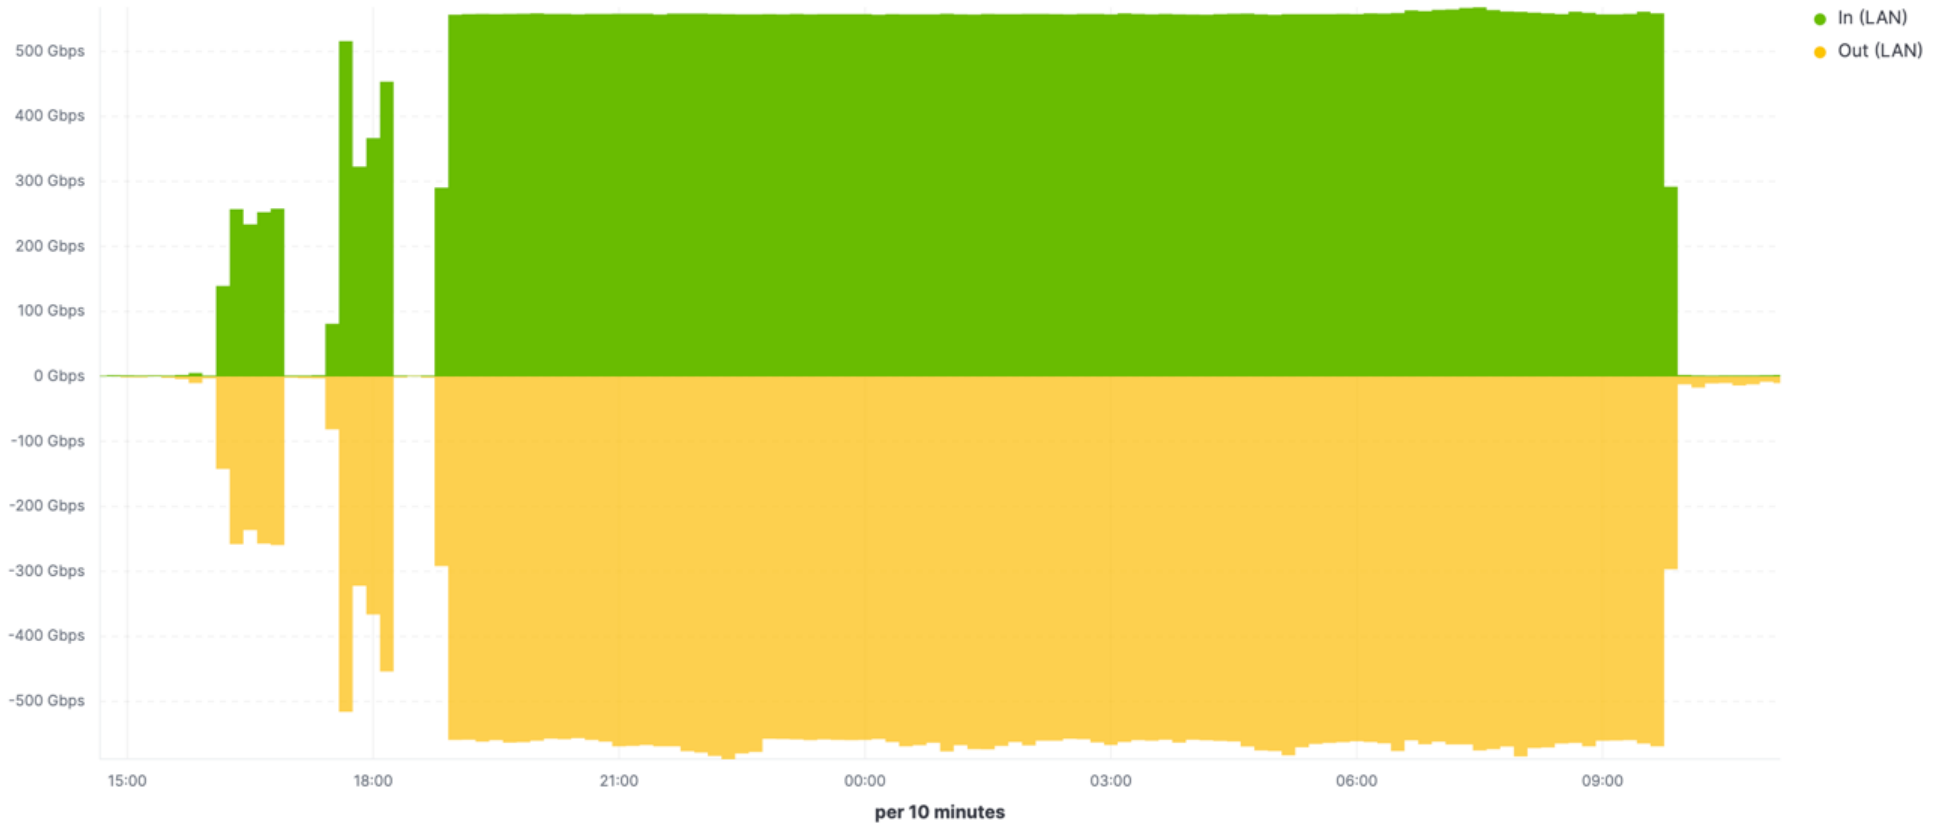

- Performed network stress test between file servers during system migration phase
- Observed: ~560 Gbps. (Ideal bandwidth: 25 Gbps x 24 (file servers) = 600 Gbps)

Disk IO performance

One file server

One worker node

Throughput

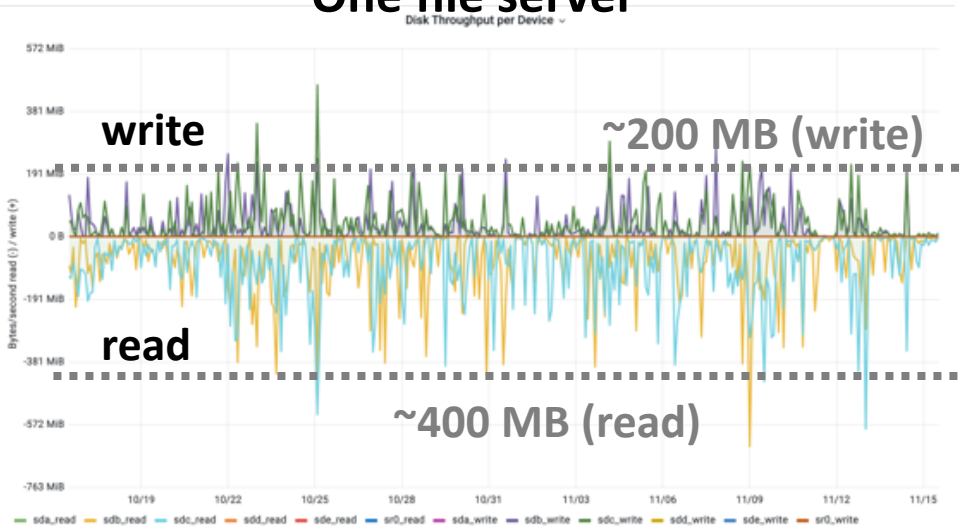

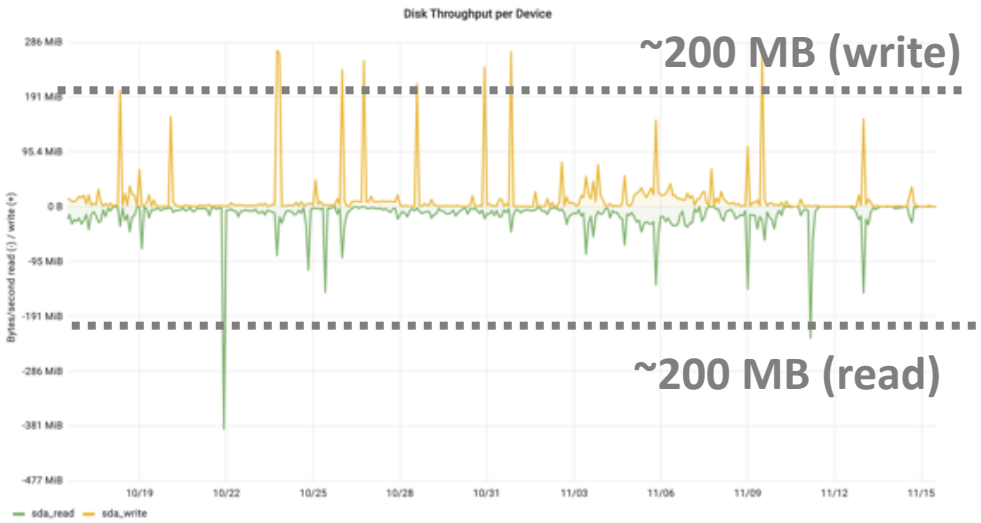

%utilization
(=IOPS x service time)