

## CCJ operations in 2018

S. Yokkaichi,<sup>\*1</sup> H. En'yo,<sup>\*1</sup> T. Ichihara,<sup>\*1</sup> W. Nakai,<sup>\*1</sup> and Y. Watanabe<sup>\*1</sup>

### Overview

The RIKEN Computing Center in Japan (CCJ)<sup>1)</sup> commenced operations in June 2000 as the largest off-site computing center for the PHENIX<sup>2)</sup> experiment being conducted at RHIC. Since then, CCJ has been providing numerous services as a regional computing center in Asia. We have transferred several hundred TBs of raw data files and nDST<sup>a)</sup> files from the RHIC Computing Facility (RCF)<sup>3)</sup> to CCJ.

Many analysis and simulation projects are being conducted at CCJ, which are listed on the web page <http://ccjsun.riken.go.jp/ccj/proposals/>. As of December 2018, CCJ has contributed to 43 published papers and 43 doctoral theses.

### Computing hardware and software

The network configuration and the computing hardware (nodes) and software (OS, batch queuing systems, database engine, etc.) are almost same as described in the previous APR.<sup>1)</sup> We have two login servers, one main server (users' home directory, NIS, DNS and NTP), and two disk servers whose disk sizes are 13 and 26 TB. The main server has an external SAS RAID (21 TB) for the home and work regions of users, and system usage. Also, the server has a RAID with built-in disks (13 TB) which can be used temporarily by users and the system.

We operate 26 computing nodes, of which 16 nodes were purchased in Mar. 2009 and 10 nodes were purchased in Mar. 2011. Thus, in total, 368 (= 8 × 16 nodes + 24 × 10 nodes) jobs can be processed simultaneously by these computing nodes using a batch queuing system LSF 9.1.3.<sup>4)</sup> Table 1 lists the number of malfunctioning SATA or SAS disks in the HP servers, namely, computing nodes and NFS/AFS servers.

One database (postgreSQL<sup>5)</sup>) server and one AFS<sup>6)</sup> server are operated in order to share the PHENIX computing environment. It should be noted that only the SL5<sup>7)</sup> environment is shared by the computing nodes,

Table 1. Number of malfunctioning HDDs in HP servers during 2011–2018.

| Type(Size)  | total | 2018 | '17 | '16 | '15 | '14 | '13 | '12 | '11 |
|-------------|-------|------|-----|-----|-----|-----|-----|-----|-----|
| SATA(1 TB)  | 192   | 16   | 18  | 8   | 14  | 11  | 16  | 20  | 9   |
| SATA(2 TB)  | 120   | 2    | 10  | 2   | 10  | 0   | 2   | 5   | 4   |
| SATA(4 TB)  | 10    | 0    | -   | -   | -   | -   | -   | -   | -   |
| SAS(146 GB) | 38    | 3    | 1   | 5   | 3   | 2   | 0   | 1   | 1   |
| SAS(300 GB) | 26    | 0    | 1   | 0   | 1   | 1   | 0   | 0   | 1   |

<sup>\*1</sup> RIKEN Nishina Center

<sup>a)</sup> term for a type of summary data files in PHENIX

which have approximately 0.9 TB of library files. We have two data-transfer servers, on which the grid environment<sup>8)</sup> is installed for data transfer to/from RCF. Two new data-transfer servers will be in operation in early 2019 and replace two servers being currently used. Data transfer of the order of 100 TB from J-PARC and BNL will be performed in the future. In addition, we operate a dedicated server for the RHICf group<sup>9)</sup> and two servers for the J-PARC E16 group,<sup>10)</sup> in order to keep their dedicated compilation and library environments along with some data.

A 10-KVA UPS (supplied by NTT Facilities) was discarded without battery exchange owing to its high cost of exchange. The power wiring was relocated in Nov. 2018. Subsequently, a total power in our UPS system is 30 KVA and a typical average load factor is 51% without jobs and 59% with maximum number of jobs in the calculation nodes.

### Joint operation with ACCC/HOKUSAI

CCJ and the RIKEN Integrated Cluster of Clusters (RICC) have been jointly operated since July 2009. In April 2015, a new system named “HOKUSAI Greatwave” was launched by RIKEN ACCC<sup>11)</sup> and the joint operation with CCJ continued, including a new hierarchical archive system in which approximately 900 TB of CCJ data are stored. Subsequently, the “HOKUSAI BigWaterFall” IA cluster, which has 840 nodes/33600 CPU cores, was launched in 2017 by ACCC. Then a dedicated usage of legacy RICC 10 nodes by CCJ was ended and the direct connection between CCJ and HOKUSAI with the two 10G Ethernet links was also ended. CCJ has not started to use the cluster because it does not support the NFS to share the computing environment of PHENIX unlike RICC. Instead, we are in preparation to use one of the “container technologies” (such as “Docker”<sup>12)</sup>) to share the environment.

### References

- 1) S. Yokkaichi *et al.*, RIKEN Accel. Prog. Rep. **51**, 188 (2018).
- 2) <http://www.phenix.bnl.gov/>.
- 3) <https://www.racf.bnl.gov/>.
- 4) [https://www.ibm.com/support/knowledgecenter/en/SSETD4\\_9.1.3/lfs\\_welcome.html](https://www.ibm.com/support/knowledgecenter/en/SSETD4_9.1.3/lfs_welcome.html).
- 5) <http://www.postgresql.org/>.
- 6) <http://www.openafs.org/>.
- 7) <http://www.scientificlinux.org/>.
- 8) <http://www.globus.org/toolkit/docs/latest-stable/gridftp/>.
- 9) Y. Itow *et al.*, arXiv:1409.4860 (Proposal).
- 10) S. Yokkaichi, RIKEN Accel. Prog. Rep., in this report.
- 11) <http://accr.riken.jp/>
- 12) <https://www.docker.com/>