# Chapter 21

# System Operations and Development Unit

## 21.1 Members

Atsuya Uno (Unit Leader)

Hitoshi Murai (Research & Development Scientist)

Keiji Yamamoto (Technical Scientist)

Fumio Inoue (Research & Development Scientist)

Yuichi Tsujita (Senior Technical Scientist)

Mitsuo Iwamoto (Technical Staff)

Katsufumi Sugeta (Technical Staff)

## 21.2 Overview of Research Activities

K computer, a distributed-memory parallel computer comprising 82,944 computing nodes, played a central role in the High Performance Computing Infrastructure (HPCI) initiative granted by the Ministry of Education, Culture, Sports, Science and Technology. The HPCI has achieved the integrated operation of the K computer and other supercomputer centers in Japan. Further, it has enabled seamless access from user machines to a cluster of supercomputers that includes the K computer. Moreover, the HPCI has provided large-scale storage systems that are accessible from all over Japan.

The system operations and development (SOD) unit has conducted research and development on the advanced management and operations of the K computer. While analyzing operational statistics collected during shared use, the SOD has improved the system configuration, including aspects involving job scheduling, the file system, and user environments.

In the fiscal year 2019 (JFY2019), we primarily implemented improvements to the following operational issues:

- Advanced Disactivation of Faulty OSTs in the Global File System
- Active Termination of I/O Racing Jobs Towards High Compute Node Utilization
- Analysis and Utilization of Operation Logs
- Improvement of Information provided to Users
- Power-aware System Operation
- Information Exchanges among HPC Centers

## 21.3 Research Results and Achievements

The K computer was terminated at the middle of August, so the operation period of the K computer in JFY2019 is about 4.5 months. Figure 21.1 shows the resource usage details, and Table 21.1 indicates the major failures of the system down for JFY2019. The number of failures in JFY2019 was only 3 cases, and these suspended time were within 3 hours. The operation of the K computer in JFY2019 was very stable.

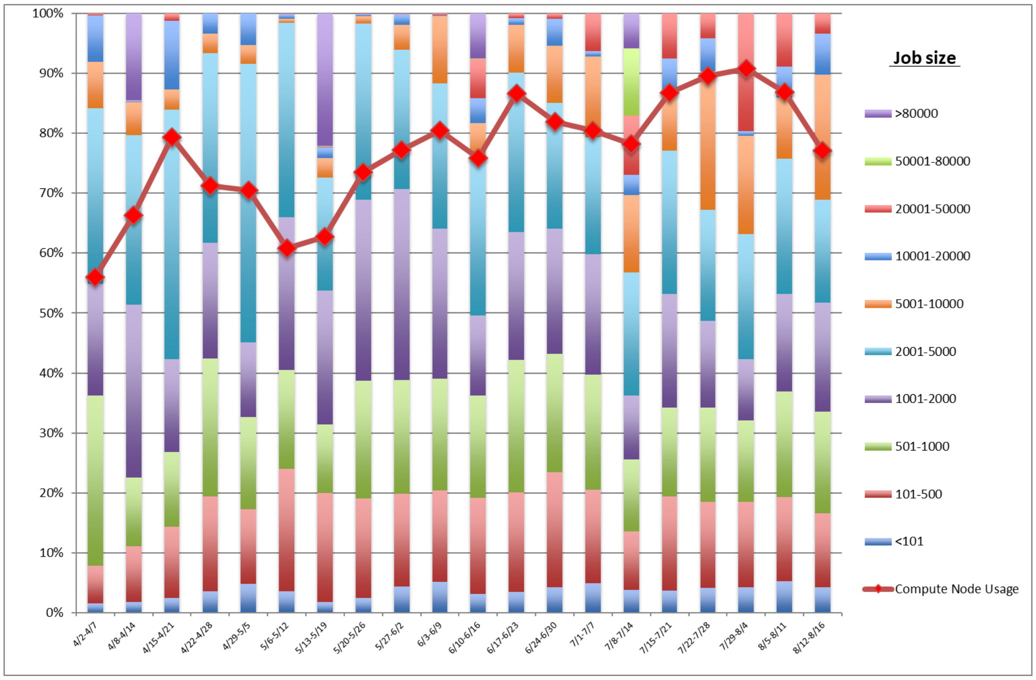


Figure 21.1: Resource usage in JFY2019

As usual, the resource usage is low at the beginning of the fiscal year, and the usage tends to increase towards the end of the term. The resource usage in JFY2019 was the same trend, and we maintained a high utilization rate, approximately 77%, which is nearly equal to that achieved in JFY2018 (77%).

Table 21.1: Major failures in JFY2019

| Date | Details | Maintenance period (H) |
|---|---|---|
| 2019/5/14 | Tofu Network Trouble (SW trouble) | 2.0 |
| 2019/5/26 | Electrical facility trouble (CGS down) | 3.0 |
| 2019/7/2 | halt of pjstat command (SW trouble) | 1.5 |

### 21.3.1 Activities for Stable System Operation

#### 21.3.1.1 Advanced Disactivation of Faulty OSTs in the Global File System

Based on our study about faulty status of HDDs in the Global File System (GFS), we have carried out advanced disactivation of faulty Object Storage Targets (OSTs) that had high failure risks of HDDs inside the same RAID-6 configuration to avoid triple faults of HDDs at the same time. We have not met any severe incident loosing a portion of data due to this kind of faults until the end of the K computer operation.

#### 21.3.1.2 Active Termination of I/O Racing Jobs Towards High Compute Node Utilization

We have deployed a client eviction scheme against the heavy Meta Data Server (MDS) load of the Local File System (LFS) caused by simultaneous I/O accesses at the same directory according to study done in the last fiscal year. I/O operations took very long time in the above situation, and finally such situation led to exceeding elapse time limit and the system operation software could not terminate such job until its I/O operation completion. Terminating affected compute nodes was the promising way, however, it took long time until the nodes were ready, resulting in a loss in compute node utilization. Especially, such negative impact is very big if the number of affected compute nodes is very large. We have focused on client eviction which decline client I/O requests by servers in order to minimize the loss. We have deployed this scheme during the K computer operation as trial, and we have not met any serious problems until the end of the K computer operation.

Similar to the case above, we have considered to cope with the high MDS load due to simultaneous truncate operations from a large number of compute nodes. We demonstrated effectiveness of this approach which issued the client evictions once such truncate operations were observed.

## 21.3.2 Analysis and Utilization of Operation Logs

We collected various operation logs and analyzed them for the stable system operation of the K computer. As one example of this study, we have proposed an I/O characterization analysis framework for further I/O tuning. In this trial, we collected 1) stats of Object Storage Servers (OSSes) of the LFS, 2) I/O sizes at each OST of the LFS, 3) and packet transfer status on Tofu links among I/O nodes.

The analysis framework utilizes those log information with the help of a PostgreSQL database storing job information such as used compute nodes, start and end dates of jobs, and so forth (Figure 21.2, 21.3.) Evaluations of the framework at the K computer provide significant differences in I/O activities at OSTs and packet transfers on the Tofu links between the original MPI-IO implementation and the enhanced MPI-IO implementation named EARTH on K.

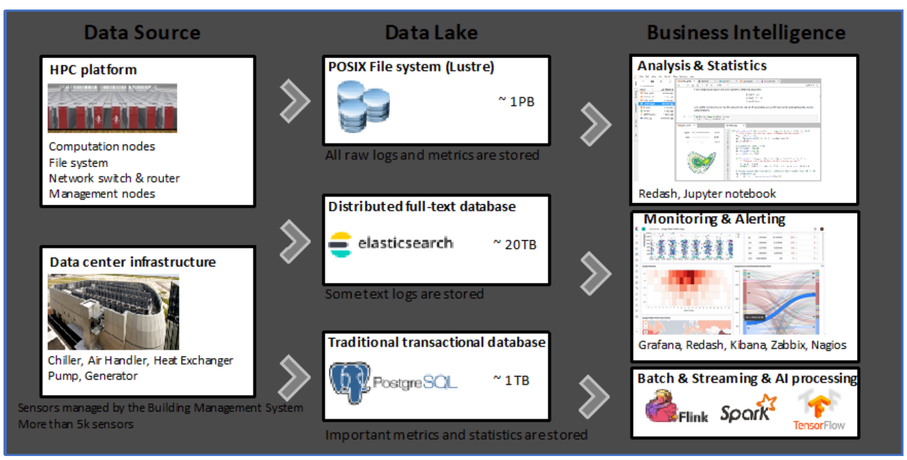


Figure 21.2: Configuration of Data collection framework for Operation log

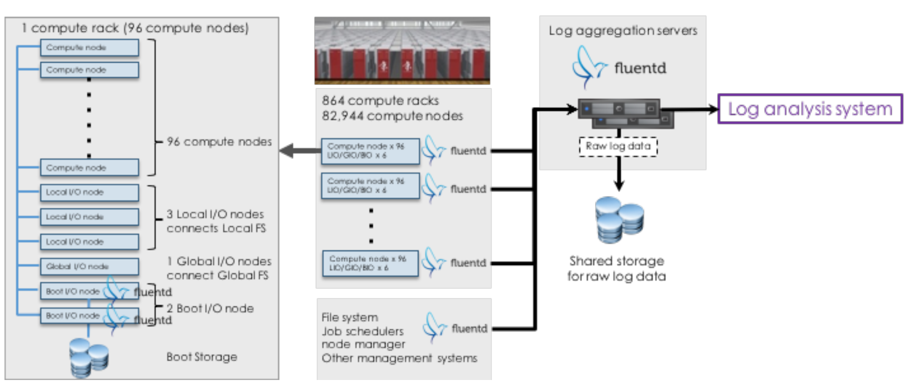


Figure 21.3: Configuration of System log collection

### 21.3.3 Improvement of Information provided to Users

We provided a tool that can display the scheduling status of the user's jobs in the three-dimensional views as the actual scheduling of the K computer. Users could download the scheduling information of the user's jobs from the K user portal, and could see the scheduling status of user's jobs on the local web-browser in the three-dimensional view (Figure 21.4.) We have a plan to provide this tool for the supercomputer Fugaku.

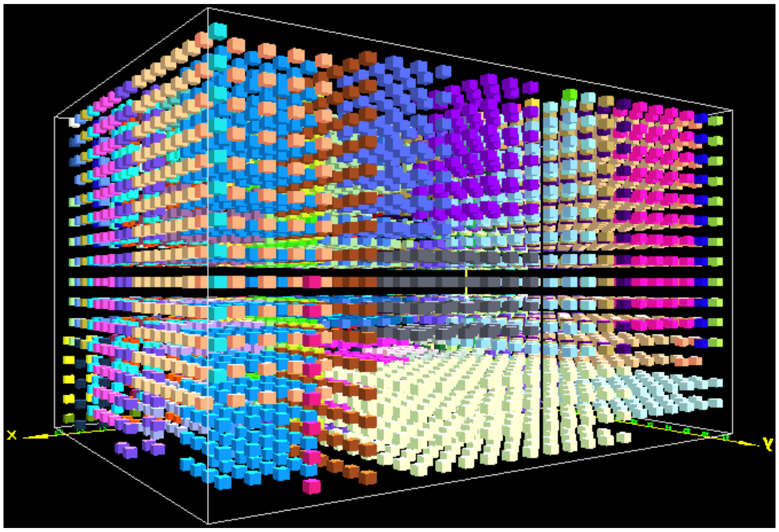


Figure 21.4: Example of 3D mapping of Job scheduling situation

### 21.3.4 Power-aware System Operation

We have studied the job scheduling considering power consumption of each job as excess power prevention. In this method, we estimate the time-series power consumption for each job from the submitted job's parameters and job scripts. In JFY2019, we have evaluated the estimation method using the similarity of job scripts. In this approach, we reclassify the jobs using the similarity of job scripts after classifying the jobs with job's parameters such as the number of compute nodes, declared limit time, and etc. We have found that this method improved estimation accuracy by several percent (Figure 21.5.)

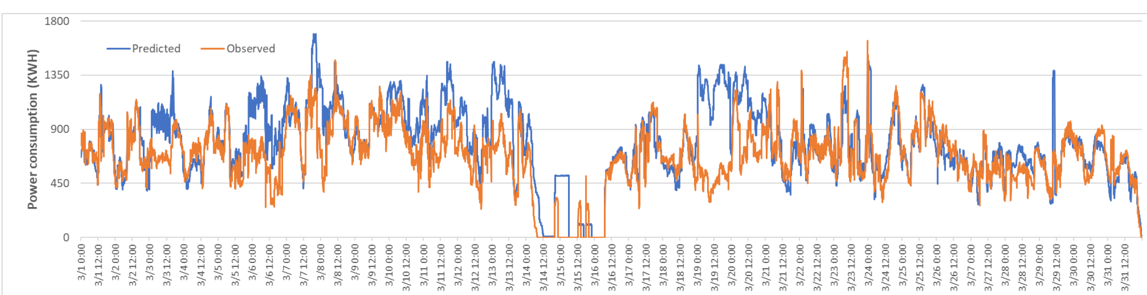


Figure 21.5: Prediction of system power consumption

In addition, we have developed the estimation model using the deep neural network system (Figure 21.6.) This model estimates the average power consumption of jobs, not the time-series power consumption at this time. We make the feature vectror of the job script, user id, number of nodes and etc, and make the model with a one-dimensional convolutional neural network using them as input data. We also estimate the standard deviation in addition to the average power consumption, and give the confidence level of the estimated values. In the future, we plan to schedule jobs considering its power consumption on the supercomputer Fugaku.

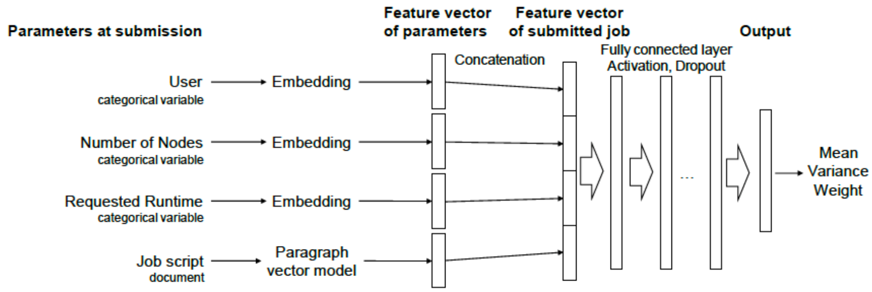


Figure 21.6: Layer structure of power prediction model using Deep Neural Network

### 21.3.5 Information Exchanges among HPC Centers

Recent growth in HPC systems has brings us complexity in HPC system operation. On the other hand, system operation is required to manage a high demand from a variety of HPC users. Each HPC center has been struggling to build own sophisticated operation environment and user support service system. To improve this situation, we started to make some information exchanges about HPC system operation and user support. As the first step of such inter-relationship, we started inter-relationship with the HPC centers at JAMSTEC and TITECH, and we have carried out file system performance evaluation in order to show characteristics of each HPC system in JFY2018. In this fiscal year, we held the workshop for the technical stuffs that are charge in the system operation in the domestic supercomputer centers to share the technical information. We have planned to share the technical information among world's supercomputer centers as the next step.

## 21.4 User support

The K computer executed approximately 160 projects in JFY2019. The number of active daily users was approximately 90. We supported users through the K support desk and provided them with technical information regarding the K computer, including information regarding its system environment, system tools, and software libraries. In addition, we performed user registrations, failures investigation, software installation, etc. We offered our consulting services together with the HPC Usability Development Unit and Application Tuning Development Unit. Figure 21.7 presents the number of issues addressed in JFY2019, showing the number of new issues in JFY2019 to be approximately 85. The number of resolved issues was approximately 94.

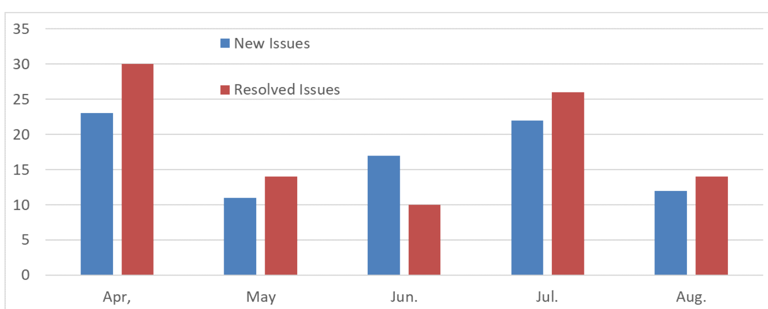


Figure 21.7: Number of issues addressed in JFY2019

## 21.5 Schedule and Future Plan

In this fiscal year, we analyzed the operation of the K computer and examined the operation improvement, and also studied the prediction method of the time series power consumption for each job about the power excess. In addition, we have proposed the framework that utilizes the operation logs and output information that in not obtained from the general profile, and we evaluated the difference in I/O activities at OSTs and packet transfer on the Tofu links between the original MPI-IO implementation and the enhanced MPI-IO implementation.

The operation of the K computer was terminated in this fiscal year. We will work for the stable operation of the supercomputer Fugaku with the experience we've gained through the operation of the K computer.

## 21.6 Publications

### 21.6.1 Conference Papers

[1] Motohiko Matsuda, Hiroya Matsuba, Jorji Nonaka, Keiji Yamamoto, Hiroshi Shibata, Toshiyuki Tsukamoto, Modeling the Existing Cooling System to Learn its Behavior for Post-K Supercomputer at RIKEN R-CCS, In proceedings of Energy Efficient HPC State of the Practice Workshop, 2019

### 21.6.2 Posters

[2] Shigeto Suzuki, Michiko Hiraoka, Takashi Shiraishi, Hiroyuki Fukuda, Takuji Yamamoto, Shuji Matsui, Atsuya Uno, Power prediction with probabilistic topic modeling for HPC, ISC High Performance 2019, 2019
[3] Jorji Nonaka, Keiji Yamamoto, Akiyoshi Kuroda, Toshiyuki Tsukamoto, Kazuki Koiso, Naohisa Sakamoto, A View from the Facility Operations Side on the Water/Air Cooling System of the K Computer, SC19
[4] Jorji Nonaka, Motohiko Matsuda, Hiroya Matsuba, Keiji Yamamoto, Yasumitsu Maejima, Toshiyuki Tsukamoto, CPU Water Cooling Temperature Effects on the Performance and Energy Consumption, ISC19
[5] Shigeto Suzuki, Michiko Hiraoka, Takashi Shiraishi, Enxhi Kreshpa, Takuji Yamamoto, Hiroyuki Fukuda, Shuji Matsui, Masahide Fujisaki, Atsuya Uno, Power Prediction for High-performance Computing, The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC19), 2019

### 21.6.3 Oral Talks

[6] Atsuya Uno, Approaches to the Power Consumption Problem on the K computer, PowerStack Seminar, 2019
[7] Keiji Yamamoto, Operational Data Analytics, Birds of Feather at SC19, 2019

### 21.6.4 Software

[8] Job scheduling 3D viewer
[9] EARTH on K