Data Science Infrastructure System Operation Procedure

Version 2.0　April 2024

Information Systems Division

Chapter 1　Objective

This document sets the operation procedure ("Procedure", hereafter) for the Data Science Infrastructure system ("System" hereafter) maintained and operated by the RIKEN Information System Division (ISD), based on Article 3 of the Data Science Infrastructure System Usage Agreement ("Agreement" hereafter). In particular, the system configuration and projects that are necessary for users to use the system shall be defined, and the system shall be operated according to the Procedure.

Chapter 2　System Configuration

The system name is HOKUSAI SailingShip (HSS), and the computational resources are as follows. Users are provided with computational resources as tenants on a project basis. In this tenant, a virtual machine (VM) and network environment can be built and used based on a virtualization infrastructure. Tenants can also use it as a storage infrastructure for the RIKEN Research Data Management System (R2DMS) without building VMs.

Section 1　Computational resources (CPU farm)

- 440 node (21,120 core)
  - CPU: Intel Xeon Platinum 8260 2.4 GHz、48 core/node、384 GB/node
    - Hyper-threading is turned on, so there are 2 vCPUs (virtual CPUs) on each core.
    - Two cores per node are used for management purposes, so users can use up to 92 v CPUs (46 cores) per node.
  - Memory: 384 GB/node
  - Local disk (SSD): 1.92 TB/node
  - Network interface: 10GBASE-T x 2 ports
- Cluster configuration
  - The CPU farm is divided into two clusters hssa and hssb with 220 nodes each.

- The hssa cluster is a pool type where users reserve computing resources at all times.
- The hssb cluster is an on-demand type where users use computing resources only when they are available.

Section 2 Shared storage (data farm)

- 30PB (Lustre file system)
  - Used as shared storage for VM.
  - Various protocols are supported.
    - CIFS: Also provided by the RIKEN R2DMS.
    - S3: Also provided by the R2DMS.

Section 3 Virtualization infrastructure

- Mirantis Cloud Platform (MCP)
- A private cloud infrastructure with OpenStack technology at its core.
- Virtualize CPU farms, data farms, and networks to provide VM and virtual network environments.

Chapter 3 Project Users and Management

Section 1 Payment representative

The payment representative, stipulated in Article 3 of the Agreement, is responsible for managing the project as a whole and is required to pay the usage fees and manage the project's users.

Section 2 Project members

Project members, as stipulated in Article 4 of the Agreement, are of four types and have different privileges in the Online Submission System.

- Payment representative: project application and approval, tenant management
- Assistant: same authority as payment representative
- Tenant manager: project application, tenant management
- Sub manager: partial management of tenant (e.g., start/stop VMs)

Section 3 Tenant users

Tenant users, as stipulated in Article 4 of the Agreement, are the users of computing resources in the tenant, such as VMs and shared storage. Account management for VMs and shared storage is basically managed by each project and not by ISD.

Section 4 Security export control

If users are non-residents or otherwise subject to the RIKEN's security export control review, a review form must be drafted and approved in advance by the Research Integrity and Economic Security Division.

Section 5 Ethics regulations

For research that is subject to review by affiliated organizations, such as the conduct of life science/medical research involving human subjects, comply with the ethics regulations, etc., of the affiliated organization. Prior permission for the use of the System must be obtained through an ethics review or other process.

Section 6 Security measures

In accordance with Article 8 of the Agreement, tenants are required to take necessary security measures. Follow the information security-related documents of RIKEN. In particular, VMs that can be accessed from outside the tenant must take more appropriate security measures such as timely updates, and VMs that can be accessed from outside RIKEN must take even stringent security measures.

Chapter 4 Usage Fees

In accordance with Article 5 of the Agreement, the usage fees are set for usage of computational resources and shared storage (Resources hereafter).

Section 1 Point

When using services with the usage fees, users must apply through the Online Submission System. Specifically, users first purchase points and then convert the points into the Resources. The points are valid until the end of the fiscal year and cannot be returned. Also, once the points are converted to the Resources, they cannot be converted back to points.

Section 2 Pool type computational resources usage

Users will apply for the number of vCPUs to be reserved and reserve them until the end of the fiscal year. The fee for the use of vCPUs is monthly, and no discount will be given for the use of vCPUs from the middle of the month.

Section 3 On-demand type computational resources usage

Users apply for the product of the number of vCPUs and the amount of time they plan to use. When a user reserves a computing resource, it can be used if there is space available. The applied computing resources are valid until the end of the fiscal year.

Section 4 Shared storage usage

Users must apply for the capacity they need in TB units and reserve it by the end of the fiscal year. The usage fee is calculated on a monthly basis, and no discounts will be offered even if you start using the service in the middle of the month.

Section 5 Budget numbers and transfer

Points can only be purchased from the budget with a RIKEN budget number on it, and payment will be made by budget transfer. External funds should be checked with the External

Fund Office to see if they can be paid in advance. When points are purchased, a transfer will be made to the budget number specified at the time of point purchase.

Chapter 5 Project Application and Approval

Project applications and approvals will be made through the Online Submission System. Some applications and approvals to be specified separately will be made via email.

Section 1 Create new project

If users want to use a new project, users can apply to create a new project. When users apply, users will enter the information of the payment representative, the budget number and its name, the content of the use, select the type of use, and, if necessary, register the project members.

Section 2 Approval of new project

After application to create a new project, when the application is approved by the payment representative, the project application is completed. After the approval of the payment representative, points can be purchased, but only a certain amount of computing resources can be used in the unapproved stage.

Section 3 Pre-approval use of new project

Before approval of a new project, you may purchase 10,800 points. After approval, the point purchase limit will be removed.

Section 4 Selecting CPU farm usage

When applying for a new project, users must choose between pool and on-demand type usage, which cannot be changed later. Note that shared storage can be used for both types.

Section 5 Registration of project members and ss accounts

Users must have ss accounts to register project members in the project. ss accounts can be obtained by logging into the Online Submission System and setting a password.

Section 6 Application and approval for project

Changes to project information, changes to project members, point purchases, and resource purchases can be made through the Online Submission System. All requests, except for resource purchases, are effective upon approval by the payment representative or assistant. Resource purchases are reflected immediately without an approval process.

Section 7 Project expiration date and continuation

Projects are valid until the end of the fiscal year. At the end of the fiscal year, the user is asked to confirm whether the project is to be continued or terminated, and if the user chooses to continue the project, the project will be continued in the following fiscal year.

Section 8 Trial use

Trial use is available to new users who are the primary users of a new project. If the user declares that they are a trial use at the time of creating a new project, they may purchase 10,800 points for a certain period of time to try out the system. If the user stops using the system within this period, the usage fee will be waived. If they continue to use the system, the points purchased during the trial period will also be subject to the usage fee.

Chapter 6 VM Usage

VMs can be managed through the OpenStack Management Console.

Section 1 OS image for VM (Instance) booting

We provide the default images published by OpenStack officials for Rocky Linux and Ubuntu as OS images available on the system. We also provide an image with various configurations to use the shared storage. In this case, the shared storage will be mounted and used as a home area (/home).

Section 2 Flavor

The CPU farm is used by choosing from the flavors available with a minimum of 2 vCPU and a maximum of 92 vCPU. In this case, memory and local disk are allocated in proportion to the amount of vCPU (memory: 4 GB/vCPU, local disk 20 GB/vCPU). Therefore, if you need a large amount of memory or local disk, you need to select a flavor with large number of vCPUs.

Chapter 7 Storage Usage

Local disk and shared storage are available for storage.

Section 1 Local storage

Local disk is available as local storage for each VM and the capacity is determined by the flavor. For images that do not use shared storage, the home area (/home) is included in the local storage.

Section 2 Shared storage

Shared storage part of the data farm and is an online storage area that can be used on a project basis. In the image of using shared storage, the home area (/home) becomes the shared storage, and the application area (/APL, read only) of the shared storage is also mounted. The home area is subject to a usage fee, and users can apply for the capacity you want to reserve through the Online Submission System.

Section 3 Non-shared storage areas of the data farm

The data farm also stores images and snapshots. This area is separate from the capacity reserved for shared storage, is not subject to usage fees, and is not required an application for use.

Section 4 Managing storage

The data in the shared storage does not guarantee the preservation of data after use unless the user indicates an intention to continue using it.

When there is a need to transfer the data belonging to the user (for reconfiguring the System, for instance), ISD will make a request to users and their project representatives to delete, compress or consolidate their data. However, if the user and the project representative do not reply to repeated requests by email, ISD will delete and reprocess the data on its own accord.

Section 5 Use with CIFS

Shared storage is also available via CIFS GW as a storage infrastructure for R2DMS. When a project is created, a CIFS area is automatically created to store data for CIFS and a CIFS account is issued. This CIFS area is independent of other methods of use, and the data is stored in a separate location and cannot be accessed by other methods. The combined capacity of the other method usages will be available within the reserved shared storage capacity.

Section 6 Use with S3

Shared storage is also available via S3 GW. An S3 area to store data for S3 will be created for users who wish to use it, and an access key pair will be issued. This S3 area is independent of other methods of use, and the data is stored in a separate location and cannot be accessed by other methods. The combined capacity of the other method usages will be available within the reserved shared storage capacity.

Chapter 8 Network Usage

There are four types of networks available within a tenant. There are two internal networks and two Floating IPs (FIPs) for access from outside the tenant.

Section 1 Project network

It is necessary to set the project IP address to the VM when communicating within the tenant. It is also necessary to be able to access outside of RIKEN and set up FIP. If you need additional IP addresses, you can add them by sending an email.

Section 2 Storage network

This is a network for accessing shared storage. If you want to use shared storage, you need to set a storage IP address. If you need additional IP addresses, you can add them by sending an email.

Section 3 RIKEN private network

A network that can be accessed from inside RIKEN. You need to set up a RIKEN private FIP address (PFIP) to access the VMs from inside RIKEN. If you need more FIPs, you can add more quotas if you have a good reason to do so.

Section 4 RIKEN global network

This network can be accessed from outside of RIKEN. If you want to access the VM from outside of RIKEN, you need to set up a RIKEN global FIP address (GFIP). If you want to use this network, you can add an additional FIP address if you have a good reason to do so.
However, to access a VM from outside of RIKEN, users must first set up a GFIP to the VM, and then apply for and pass a security audit in the same way as normal applications for a RIKEN global IP address.

Chapter 9 Service Level

Section 1 Availability

We aim to provide at least 97% of the annual service hours.
In the event of a planned outage in the Wako area, which is usually held on a consecutive holiday in October, the system will be shutdown/activated on the weekdays before and after the planned outage. In addition, scheduled maintenance may be performed for a few days before or after a scheduled outage or at the end of the fiscal year. In the event of a planned outage, users will be notified one month in advance. Sudden power outages, momentary voltage drops, system problems, etc. may cause the system to stop suddenly. In this case, we will contact the user immediately and explain the cause and the recovery process after the problem is solved.

Section 2 Performance of VM

We strive to provide the performance that VMs are inherently capable of.

Section 3 Capacity and data preservation

The storage device is located in the RIKEN Wako area, and the data is stored with redundancy (equivalent to RAID 6), and all the controllers and communication paths of the storage devices are redundant. The storage is also connected to an uninterruptible power supply (UPS) in case of power-related problems; however, since there is no emergency generator connected, the storage device is shut down in the event of a power failure to ensure data reliability. However, the integrity and availability of data at any time cannot be guaranteed.

Section 4 System operation period and transition

The SS is scheduled to be in operation from June 2020 to the end of May 2026. It is assumed that the SS will be continued to the next system, but this has not yet been determined. We

will give as much consideration as possible for the transition to another system at the end of the operation.

Section 5 Scope of support

Basically, each project is responsible for managing its own internal tenants. However, we provide technical support for managing project members and starting up VMs (instances) using images provided by us, as well as technical support for using ISVs in application area using images provided by us.

Chapter 10 Others

In addition to the Procedure, other materials necessary for the use and operation of the system will be provided separately by the ISD Director.