Okinawa Institute of Science and Technology Graduate University
Centralized Shared High Performance Research Data Storage

Contents

1. Background and Scope

OIST seeks proposals for its centralized shared high performance research data storage to provide centralized storage resources to OIST scientific research and education. This research data storage will follow and take over the existing Research data storage. This document gives OIST system requirements to be considered in vendor proposals.

The successful bid will be determined via the proposed core-price value and proposal evaluation, for which eligibility and evaluation criteria are described in the next sections.

Each proposal shall include the total cost for the hardware, delivery and physical installation plan, OS installation, configuration, license costs, staff training, and hardware support.

2. Eligibility Criteria

The vendor must have prior experience in high performance storage deployments in Japan at the scale of this system or larger.

The vendor must have engineers with experience in installing, maintaining and supporting multi-petabyte research data storage systems in Japan. These engineers must be regular employees of the system vendor and cannot be outsourced. Moreover, they will be involved in the design, implementation, operation and support of the delivered research data storage system.

The vendor shall have prior experience of delivering maintenance and support in Japan.

OIST will evaluate the entire proposition using the criteria in the section below.

3. Evaluation Criteria

Refer to the “Score Sheet.xlsx” document.

4. The matters listed in the proposal

The following must be provided as part of the response to the tender submissions (documentation clearness, format and completeness are taken into consideration during the evaluation):

- Evidence for eligibility
    - Relevant experience and demonstrated ability to design, deliver and support HPC system (computing and storage) in Japan.
    - At least three names of engineers having experience installing, maintaining and supporting storage systems of the present scale or larger in Japan.
- A complete set of quotes that include the unit cost for each item in the system
    - A quote must be provided for the whole system.
    - Unit cost shall be at/ offer price (not list price)
    - Maintenance and support costs 6 years (FY2026~FY2031)

- Evidence that the storage system can fulfill the requirements detailed in the specification
- Detailed plan for the data migration and consistency checks of the migrated data
- Detailed faceplate and estimated (at peak performance) maximum power consumption of the system. Total (estimated) power consumption at peak performance should not exceed the power budget limit.
- A basic acceptance testing procedure for the system that includes the stress check
- Evidence of support response time of two or less business days for customer in Japan (from first contact to problem resolution including part replacement lead-time).
- Delivery and acceptance test plan

Detailed information about OIST building facility, data center and existing storage system is available during the public announcement period of the tender (Q&A).

5. Warranty, Maintenance and Support

All systems must be covered by a yearly renewable warranty, and the whole system must have at least 6 years of lifetime from the delivery date.

OIST staff will physically replace the components listed below. The vendor must provide a minimum spare part stock onsite for the items in the following table.

Component	**Minimum number of spare parts per component type** **[Proposed number must be based on annualized failure rate (AFR)]**
HDD	4
SSD	1

All other failed hardware components must be replaced by the vendor onsite within the next two business days.

Technical support in either English or Japanese (whichever is available) must be available by telephone and by email during business hours (business days, excluding weekends Japanese holidays and year-end Dec 29th ~ Jan 3rd, 9:00-17:00 minimum core time, Japan local time).

Maintenance and support for the storage for this new system should be at least equivalent (quality, response time, implementation, etc.) to the maintenance and support provided to actual OIST tiered storage systems (See the attached document “OIST_Storage_Maintenance_Spec_example.docx” for an example maintenance and support specification).

6. High Performance Research Data Storage

The storage system is implemented as a parallel storage. All servers, switches, etc. components of the storage system must have redundant high efficiency power supplies, and all the hardware components must have the latest working firmware and OSes.

The storage system will provide SMB/CIFS and NFS services by connecting to the campus network switches through at least 2 x 100Gb/s LACP per head (refer to section 9-2) connected with SR4 MPO cable. The server components in the storage system should also have the latest Infiniband connectivity support for connection (at least one port per head node) to future HPC systems (HPC-1 and HPC-2).

6-1. Data storage

Storage should have at least 60 PB of capacity with equivalent RAID6 redundancy, and rack redundancy if the storage uses several racks.

Read and write performance should be at least 200 GB/s and 200 GB/s, respectively. The filesystem must also provide user, group and directory quotas.

The storage must be able to be integrated with OIST existing tape backup/archiving system and workflow (refer to figure 1 below, more information available upon request).

For flexible volume control, the proposed storage system must allow future expansion, either by adding HDDs, enclosures or whole device units.

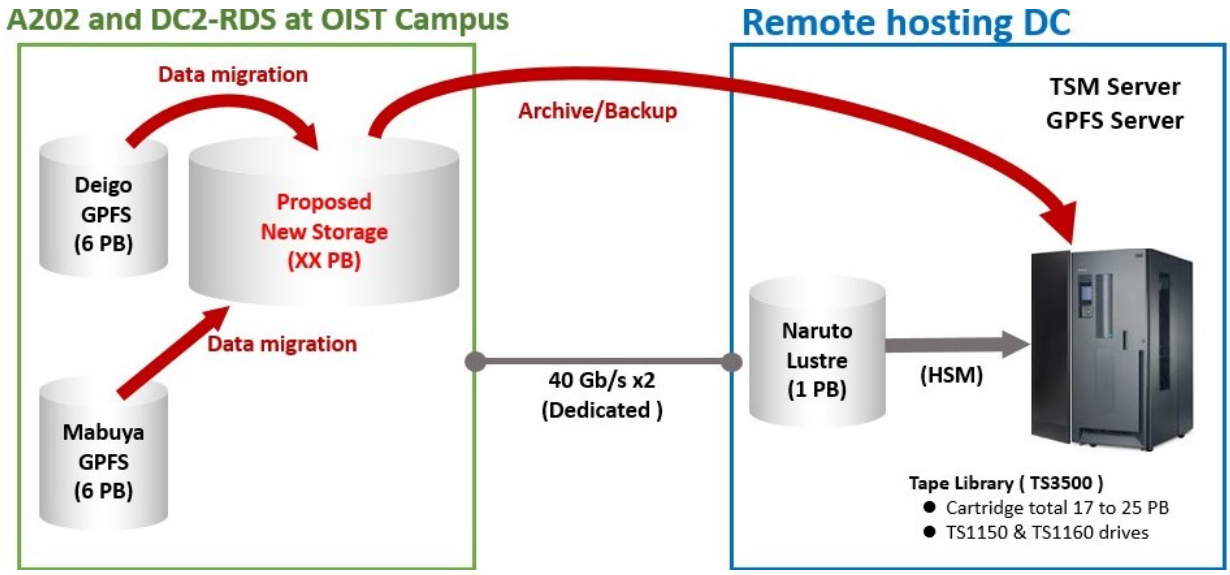

Fig. 1 Storage systems overview

6-2. Data management and monitoring

The storage should have at least one dedicated monitoring server equipped with flash local storage (at least NVMe) in RAID configuration, with monitoring software providing basic monitoring functions for all storage components. Redundant network connection (LACP), 2 x 10GBase-T connection is acceptable.

6-3. Network

OIST will provide the switches for the Campus and Management Ethernet networks, so they should not be included in the delivery.

The vendor must provide all the necessary devices (switches, cables, connectors/transceivers) for the Infiniband network, and for connecting to OIST Campus and Management Ethernet network, except for the transceivers on OIST switches for 100Gb/s and 10Gb/s connections.

6-4. Support for Storage

The vendor will propose the support flow. When having contact from OIST about storage issues, the vendor should fix the issues using remote login, via SSH, except for HW issues when onsite is required.

Manufacturer direct support regarding HW issues will be preferred to reduce the resolution time and therefore the outage time.

7. Data Migration

The data and file structure from the existing storage in Lab1-A202 must be transferred to the new storage in DC2-RDS (confer to Fig. 1). The data migration consists in moving all the data from Mabuya GPFS and Deigo GPFS devices to the new storage device located in DC2 RDS room. The migration should occur with minimum service outage to reduce the impact on OIST research activities. The migration procedure might follow the steps below:

1. Installation of the new storage in A202
2. First copy (rsync) of the data using in parallel multiple compute nodes of the Deigo system (from OIST), over the infiniband (HDR) network. If required HDR ports would be made available on OIST HDR switches, but HDR cable should be prepared by the vendor.
3. Relocation of the new storage from Lab1-A202 to DC2-RDS room
4. Final copy (rsync) of the data that was generated during stage 3, over the Ethernet network

8. Relocation and rebuilding of existing storage system

The GPFS storage devices in racks D1-2 and I3-4 of Lab1-A202 server room shall be relocated to racks A5-A6 of DC2-RDS room. After relocation of the storage devices, the storage will be rebuilt and then configured to act as a standalone GPFS storage device (reusing the existing GPFS licenses).

9. Data centers and rack layout

The layout of Lab1-A202 and the RDS room of DC2 are shown in Fig. 2 and 3, respectively. The RDS room will have a maximum power capacity of approximately 100kW available for the new and the relocated storage systems from a HA-redundant UPS source.

All components of the storage systems installed in the RDS room shall be connected to the UPS source.

9-1. Lab1-A202 data center layout

The location of the existing storage in Lab1-A202 is shown in Fig. 2. The research storage consists of three devices from three different hardware generations located in H1-3, D1-2 and I3-4, respectively, in generation order.

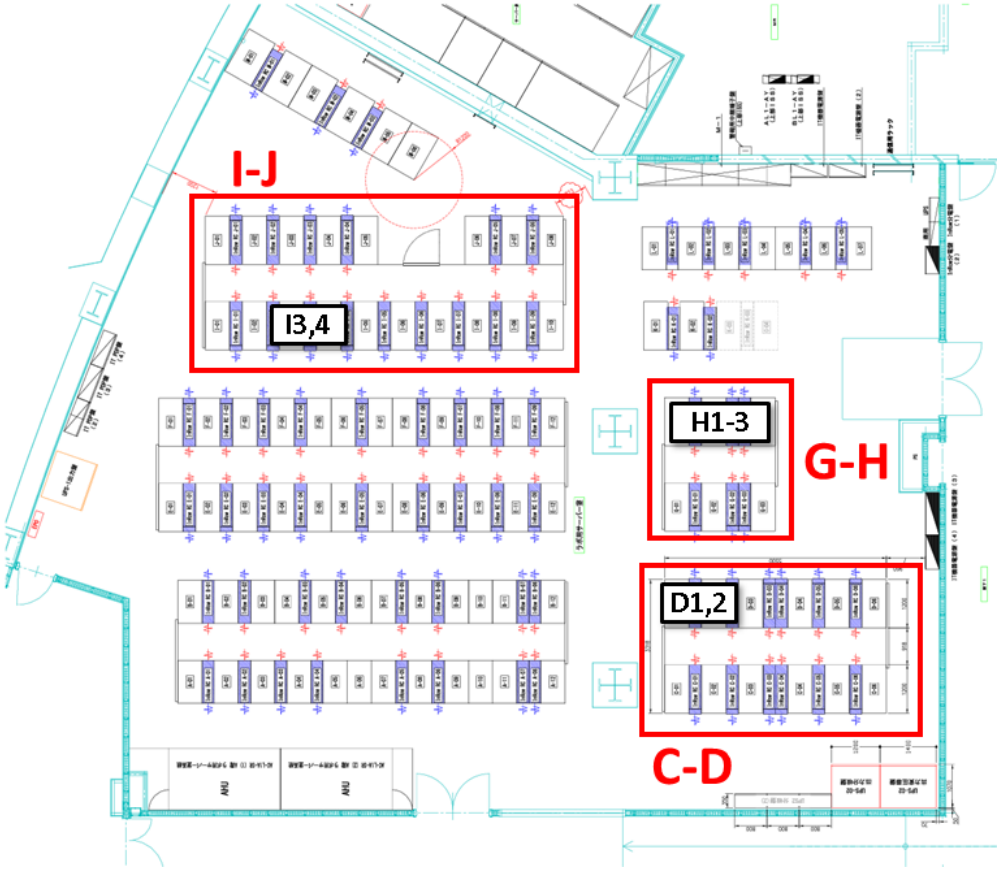

Fig. 2 Top view rack layout of OIST data center Lab1-A202 and location of existing storage system in racks D1-2, I3-4 and H1-3.

OIST will provide the racks and PDU for the I-J, C-D and H-G rows. Electrical power will be available at the PDU socket to operate the system components.

9-1-1. Rack description

Racks in Lab1-A202 are grouped into rows to form HACS (Hot aisle containment system) islands.

- Front width: 600 mm
- Depth: 1200 mm
- Height: around 2000mm (42U usable)
- Cooling capacity: between 22kW and 24kW

9-1-2. PDU description

In each rack, two types of PDUs will be used.

0U vertical mount PDU (4 in H-G, 6 in C-D and I-J):

- Rated input voltage: Single phase, 200V
- Rated input current: 30A

- Nominal power: 6kW
- Output outlets: IEC C19(6), C13(36)

2U horizontal mount PDU (2 only in C-D and I-j):

- Rated input voltage: Single phase, 200V
- Rated input current: 30A
- Nominal power: 6kW
- Output outlets: IEC C19(4), C13(12)

9-2. OIST DC2 and RDS room layout

The RDS room consists of 12 racks. OIST will provide racks, PDU and Ethernet switches in the RDS room. Electrical power will be available at the PDU sockets to operate the system components.

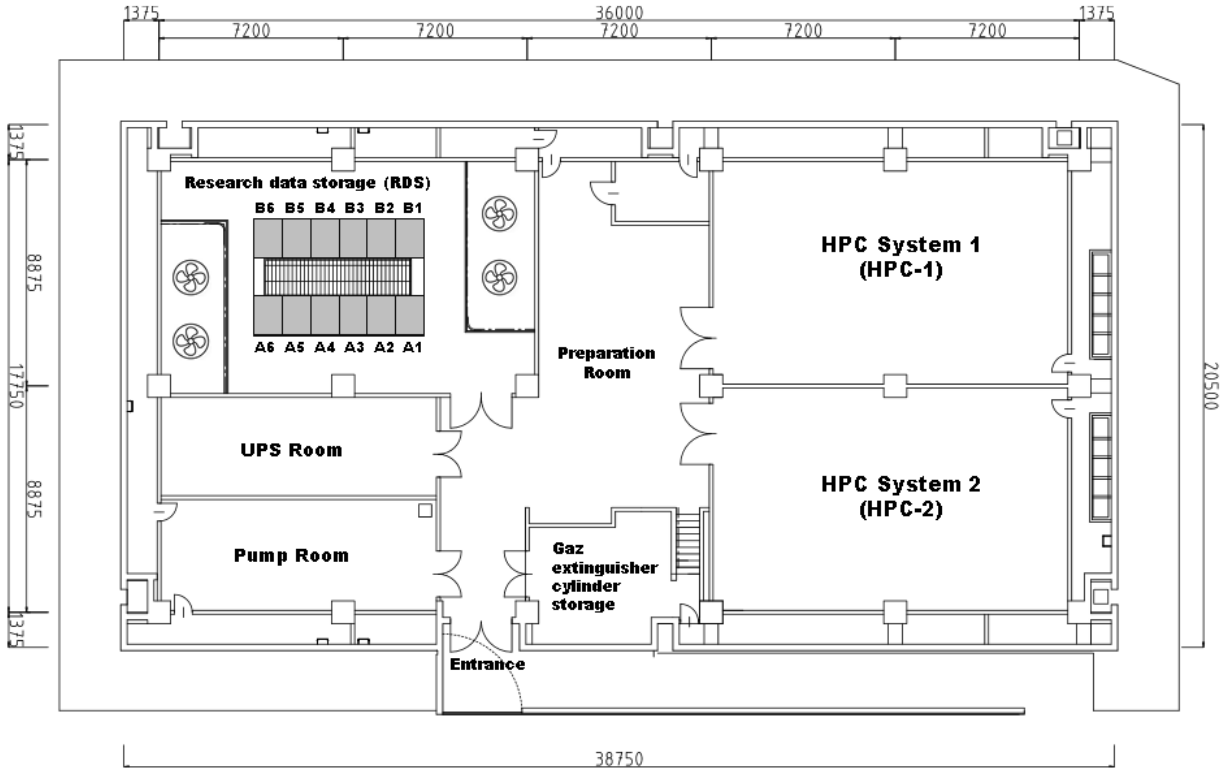

Fig. 3 Top view of OIST 2nd data center rack layout. The RDS room will host the new storage system.

The assignment of racks to be used by the proposed system is as follows: racks B1, B2, B4, B5 and B6 are to be used for the new storage, rack B3 is for the Infiniband network switches and necessary servers for the new storage (part of the rack B3 will be used for OIST switches, so devices placement in the rack should be coordinated), and, racks A5 and A6 will be used for the relocated storage system (OIST campus and management network switches will also be present in the top of those racks).

9-2-1. Rack description

The RDS room consists of 2 rows of 6 racks configured to form a CACS (cold aisle containment system).

- Front width: 800 mm
- Depth: 1200 mm
- Height: around 2000mm (42U usable)
- Width of cold aisle: 1800mm (600x600 tiles)
- Maximum cooling capacity: 30kW

9-2-2. PDU description

- 4 PDU per rack
- Rated input voltage: Single phase, 240V
- Rated input current: 30A
- Nominal power: 17.3kW per PDU
- Output outlets: C13/C15/C19/C21(24), C13/C15(24)

9-2-3. Network switches description

- Service network (private) 48-port 1Gb/s switch (per rack)
- Management network (private) 48-port 1Gb/s switch (per rack)
- Campus network (uplink to OIST core) 25 Gb/s switch (per rack)
- RDS spine switches 100 Gb/s

10. Physical Installation and acceptance

A spreadsheet will be used to track any issues and changes happening during the physical installation and acceptance, and a daily report will be provided to OIST using this spreadsheet.

10-1. Physical installation

The vendor should bring all additional tools, power meters, parts, etc. required for the installation and acceptance. Dedicated clean clothing must be used inside OIST data center and the vendor personnels must wear the required protection for the work. All the servers and components will be unpacked outside the data centers, and no card boxes, plastic bags, bubble wraps, polystyrene foams, paper sheets, etc. shall be introduced inside the data center.

A kick-off meeting will be held prior to the installation to discuss final technical aspects of the delivery, such as the final layout of the system components (storage, servers, switches), etc.

10-2. Acceptance

The acceptance will consist of the verification of the entire requirement and the proposed items agreed during the proposal evaluation. Moreover, the following checks must be cleared for the acceptance of the deliverable.

- The vendor will perform storage performance evaluation (the head servers can be used for the evaluation)

As part of the acceptance, the following system details must be provided to OIST:

- Documentation on the installed storage system (layout, configuration, etc.).
- MAC addresses of all NIC devices together with their node/server names (before the delivery, OIST will provide the name of the system that will be used for the nodes/servers naming)
- Operation manual for the storage that include emergency safe shutdown procedure

The vendor will provide a report for each verification. In the event of failure to clear a verification, the vendor should provide remedy at no cost.

10-3. Delivery deadline and place

The system must be delivered and accepted by February 27, 2026.
The system delivery place is OIST 2nd data center, at the address below:

Okinawa Institute of Science and Technology Graduate School
1919-1, Tancha, Onna-son, Kumigami-gun, Okinawa 904-0412, Japan