



Fugaku Open OnDemand

Masahiro Nakao (RIKEN R-CCS)

Background

- The issue with HPC clusters is that there is a lot of prerequisite knowledge to use them
 - Command Line Interface (CLI)
 - Generate SSH key pair and register public key
 - Job Scheduler
- Users want to run GUI applications such as remote desktop or JupyterLab on compute nodes, but the procedure is complicated.
 - Requires VNC client and SSH tunneling settings



```
top
Processes: 596 total, 2 running, 594 sleeping, 2418 threads      10:05:40
Load Avg: 2.15, 1.81, 1.63  CPU usage: 4.39% user, 3.8% sys, 92.51% idle
SharedLibs: 643M resident, 120M data, 40M linkedit.
MemRegions: 96880 total, 4270M resident, 425M private, 2994M shared.
PhysMem: 14G used (1893M wired, 364M compressor), 992M unused.
VM: 223T vsize, 4283M framework vsize, 160759(0) swapins, 248728(0) swapouts.
Networks: packets: 10719029/12G in, 2457839/525M out.
Disks: 8336373/93G read, 2912845/60G written.

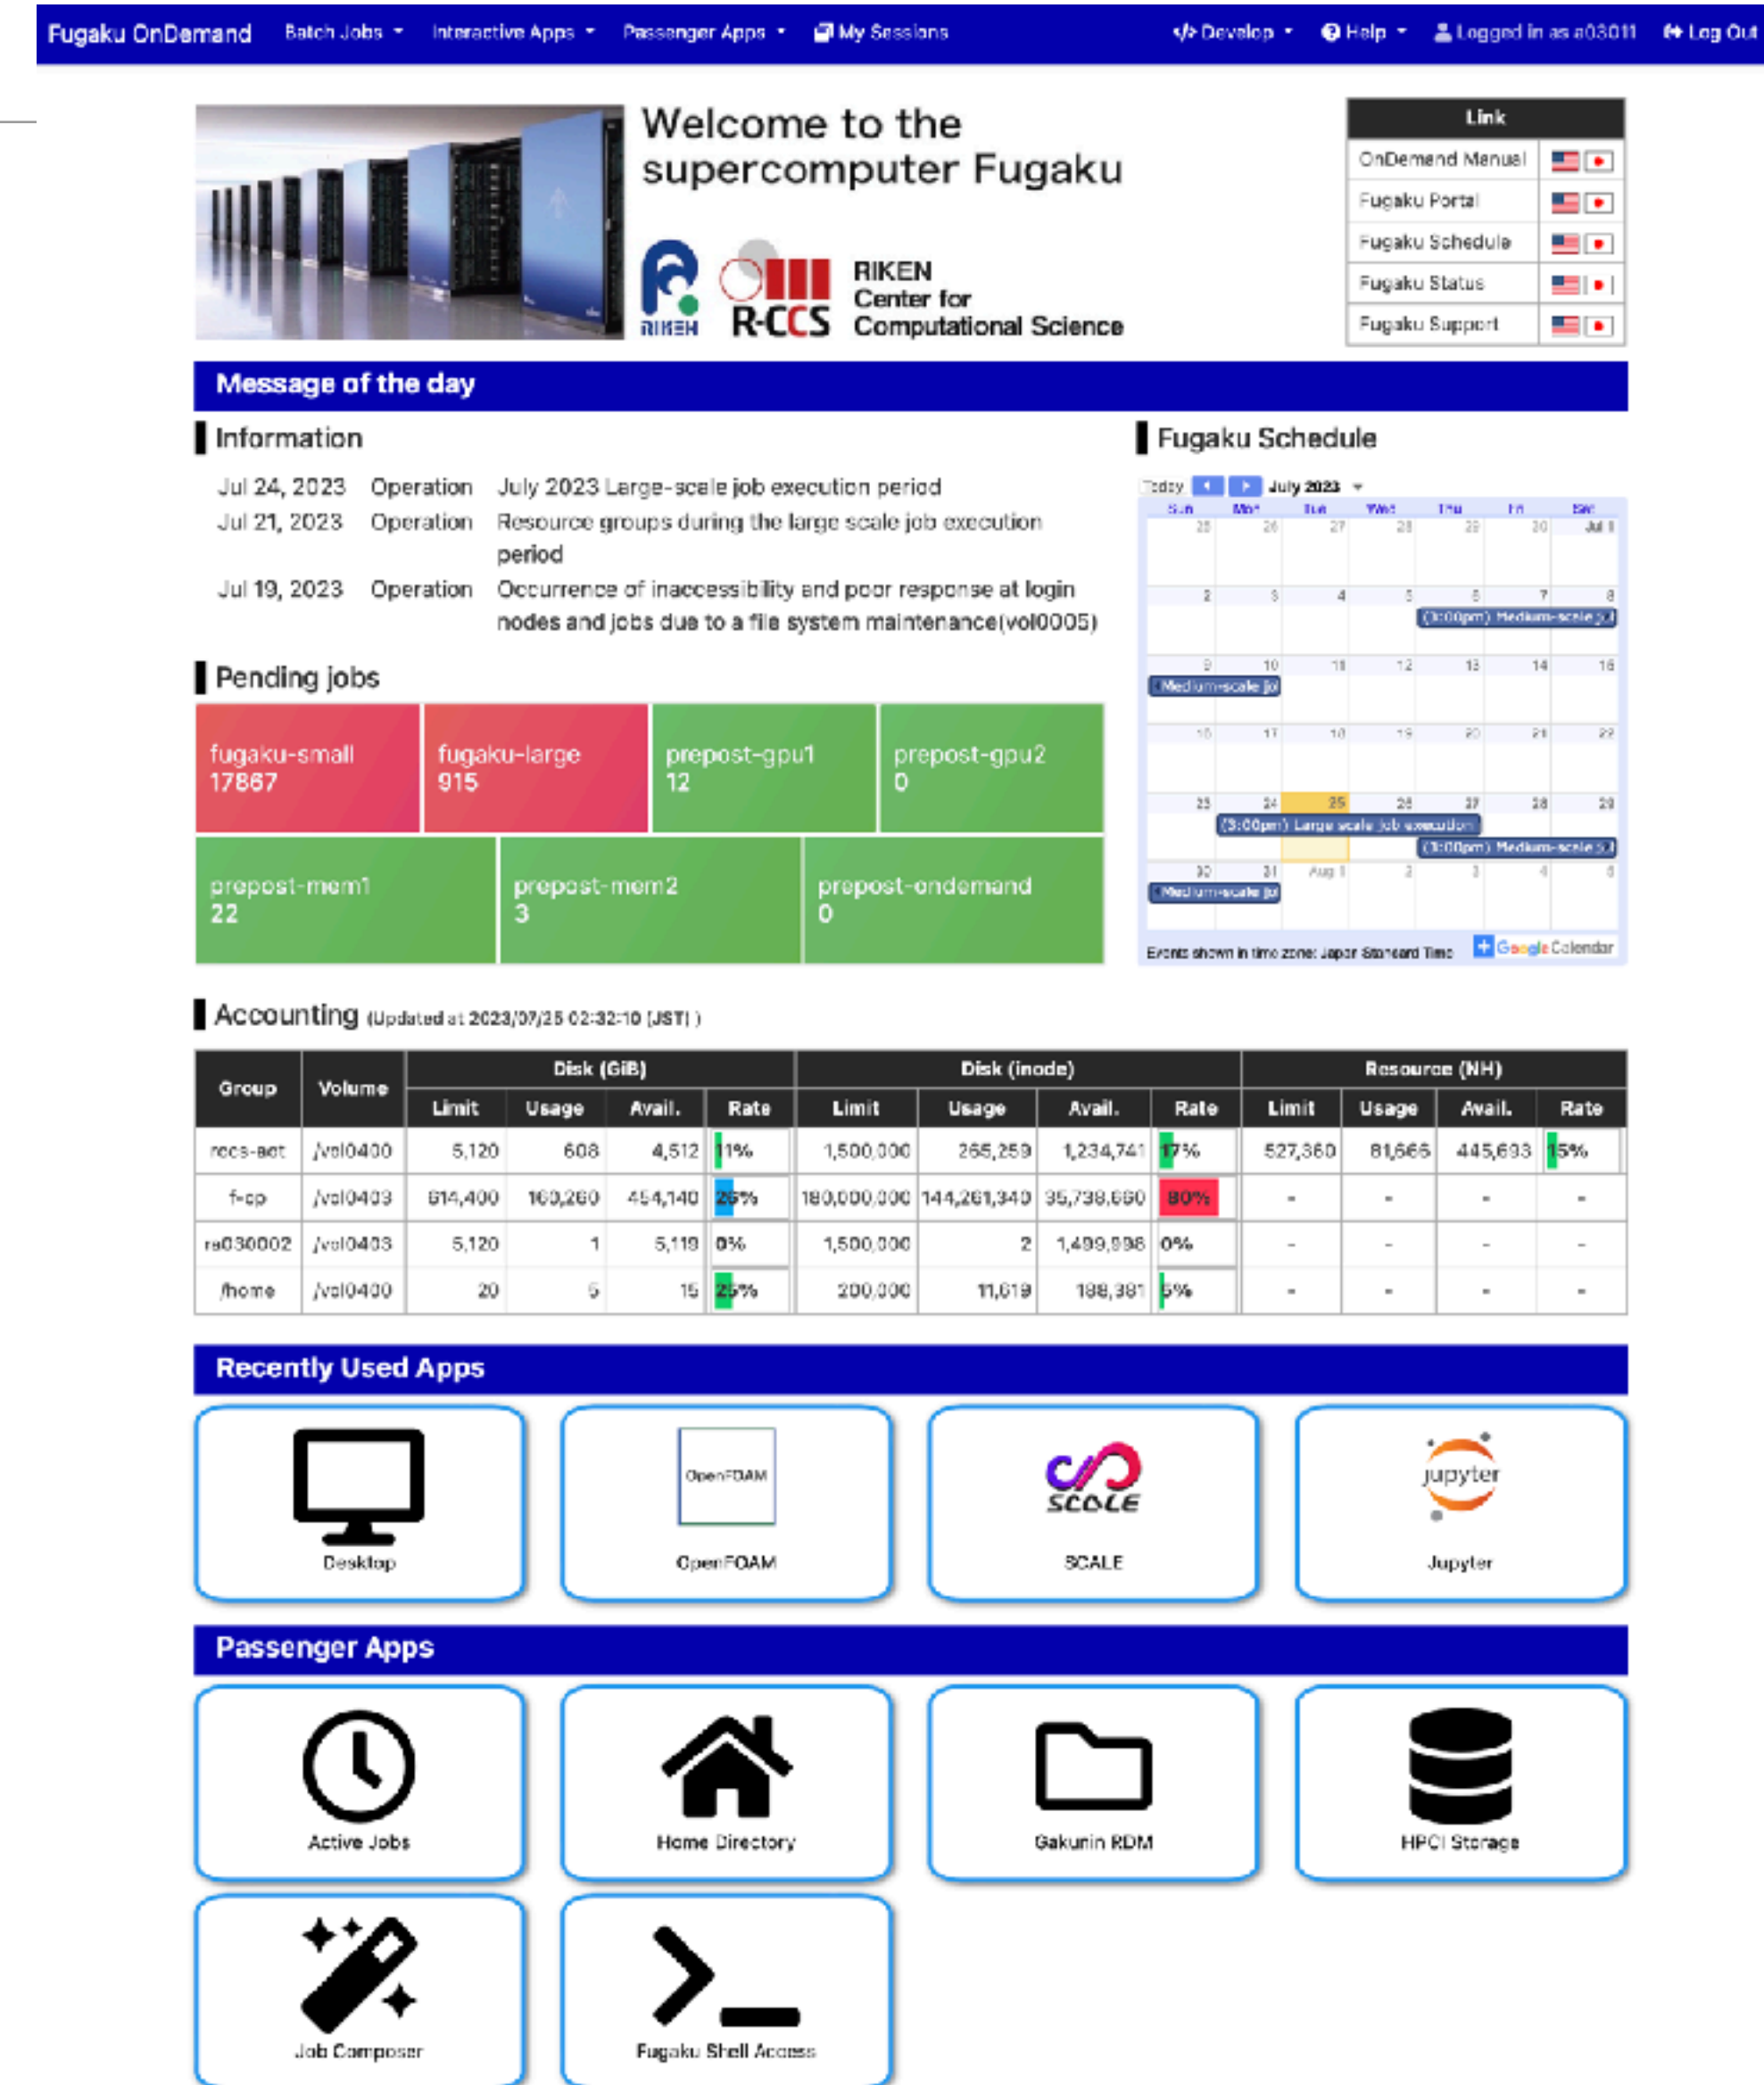
PID    COMMAND      %CPU    TIME    #TH    #WQ    #PORT  MEM    PURG    CMPS    PGRP
177    WindowServer 31.0    03:44:15 26     10     4438+  2247M  223M+  79M    177
0      kernel_task  6.0     01:53:35 484/8  0      0      2002M+ 0B      0B      0
78311  top          5.3     00:00.43 1/1    0      28     5761K  0B      0B      78311
78313  screencaptur 4.1     00:00.33 4       3      71     7618K  752K    0B      1803
1794   iTerm2       3.7     22:59.68 6       3      422-   226M+  13M-   83M     1794
```

Open OnDemand

<https://openondemand.org>

- Web portal for HPC clusters
- Developed mainly by Ohio State University
 - Open source software
 - Used by over 250 institutions worldwide
- Operate HPC clusters from your web browser
 - No other software installation required
 - File upload/download, job submission and monitoring, shell access via CLI
 - GUI applications on compute nodes can be launched easily

 Improving user convenience



The screenshot displays the Fugaku OnDemand web portal. At the top, a navigation bar includes links for 'Fugaku OnDemand', 'Batch Jobs', 'Interactive Apps', 'Passenger Apps', and 'My Sessions', along with a 'Develop' button and a user login status. The main header area features a 'Welcome to the supercomputer Fugaku' message, the RIKEN Center for Computational Science logo, and a 'Link' menu with options like 'OnDemand Manual', 'Fugaku Portal', 'Fugaku Schedule', 'Fugaku Status', and 'Fugaku Support'. Below this, a 'Message of the day' section provides information about the July 2023 large-scale job execution period and a login issue. A 'Pending jobs' section shows a grid of job status boxes for various resource groups. To the right, a 'Fugaku Schedule' calendar is visible. The 'Accounting' section, updated at 2023/07/25 02:32:10 (JST), contains a table with disk and resource usage statistics. At the bottom, there are sections for 'Recently Used Apps' (Desktop, OpenFOAM, SCALE, Jupyter) and 'Passenger Apps' (Active Jobs, Home Directory, Gakunin RDM, HPCI Storage, Job Composer, Fugaku Shell Access).

Group	Volume	Disk (GiB)				Disk (inode)				Resource (NH)			
		Limit	Usage	Avail.	Rate	Limit	Usage	Avail.	Rate	Limit	Usage	Avail.	Rate
recs-bot	/vol0400	5,120	608	4,512	11%	1,500,000	255,259	1,234,741	17%	527,360	81,665	445,695	15%
f-cp	/vol0403	614,400	160,260	454,140	26%	180,000,000	144,261,340	35,738,660	80%	-	-	-	-
ra030002	/vol0403	5,120	1	5,119	0%	1,500,000	2	1,499,998	0%	-	-	-	-
/home	/vol0400	20	9	15	45%	200,000	11,619	188,381	5%	-	-	-	-

Dashboard of Fugaku Open OnDemand

- A. External links (Fugaku manual, etc.)
- B. Failure information, operation information, etc.
- C. Number of waiting jobs in each queue using Grafana
- D. Operational calendar using Google Calendar
- E. User disk and budget utilization
- F. Apps that run on recently used compute nodes
 - Interactive application (GUI jobs and Batch Jobs)
 - Other apps can be selected from a navigation bar
- G. Utilities that work with Open OnDemand server
 - Passenger application
 - File upload, job monitoring, etc.

<https://ondemand.fugaku.r-ccs.riken.jp>

The screenshot shows the Fugaku Open OnDemand dashboard. Annotations point to specific features:

- A**: Points to the 'Link' menu in the top right corner, which includes links to the OnDemand Manual, Fugaku Portal, Fugaku Schedule, Fugaku Status, and Fugaku Support.
- B**: Points to the 'Message of the day' section, which contains information about recent operations and a 'Fugaku Schedule' calendar.
- C**: Points to the 'Pending jobs' section, which displays a bar chart showing the number of pending jobs for various resource groups like 'fugaku-small', 'fugaku-large', 'prepost-gpu1', 'prepost-gpu2', 'prepost-mem1', 'prepost-mem2', and 'prepost-ondemand'.
- D**: Points to the 'Fugaku Schedule' calendar, which shows a monthly view of the schedule.
- E**: Points to the 'Accounting' section, which displays a table of disk and resource usage for various groups.
- F**: Points to the 'Recently Used Apps' section, which shows icons for Desktop, OpenFOAM, SCALE, and Jupyter.
- G**: Points to the 'Passenger Apps' section, which shows icons for Active Jobs, Home Directory, Gakurin RDM, HPC Storage, Job Composer, and Fugaku Shell Access.

Click the icon to launch the application

Interactive Applications on Fugaku Open OnDemand

- GUI Jobs (Operated interactively)

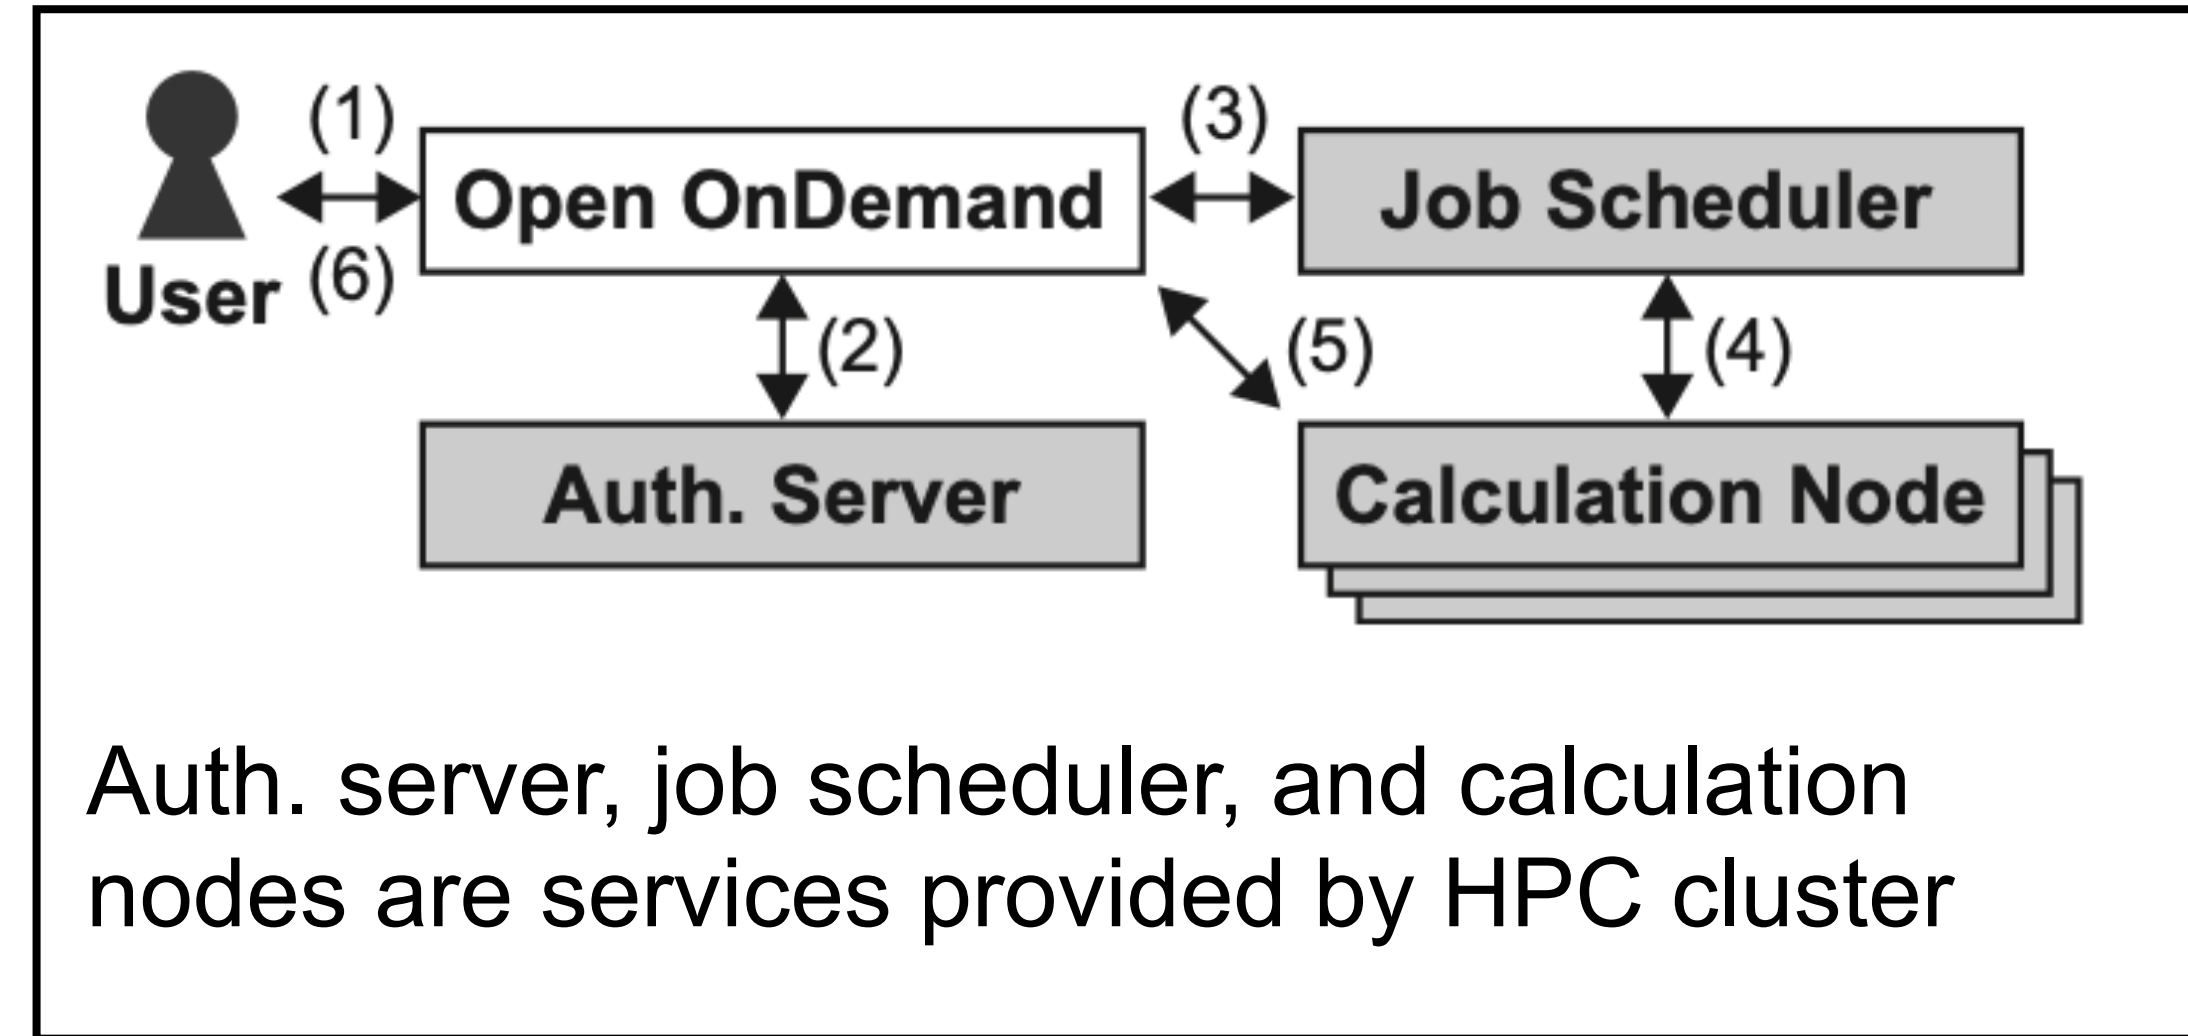
Category	Application
Development	Remote Desktop, JupyterLab, MATLAB, RStudio, VSCode
Profiler	NVIDIA Visual Profiler, NVIDIA Nsight Compute*, NVIDIA Nsight Systems, Vampir
Viewer	AVS/Express, C-Tools, GaussView, ImageJ, OVITO, Paraview, PyMOL, SALMON view, Smokeview, VESTA, VMD, VisIt, XCrySDen
Workflow	WHEEL

- Batch Jobs (Not operated interactively)

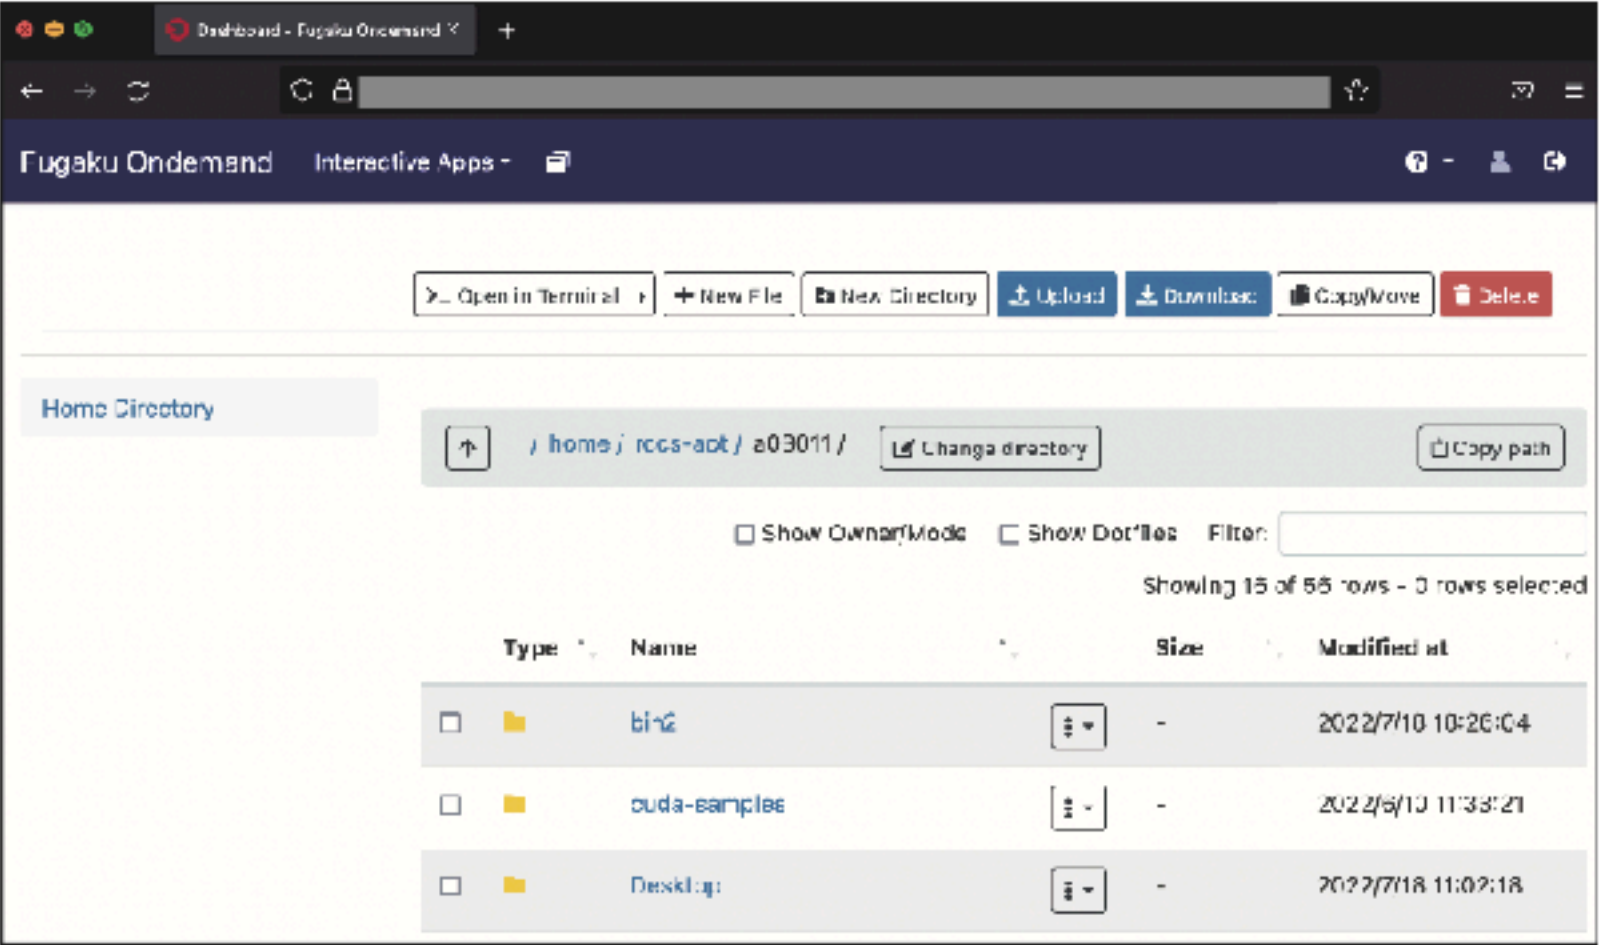
Category	Application
Climate	SCALE
Computer Aided Engineering	FDS, FrontFlow (blue/X), FrontISTR, OpenFOAM (Foundation/OpenCFD)
Condensed Matter Physics	ALAMODE, AkaiKKR, HΦ, mVMC, OpenMX, PHASE/0, Quantum Espresso, SALMON
Molecular Dynamics	GENESIS, GROMACS, LAMMPS, MODYLAS
Quantum Chemistry	ABINIT-MP, Gaussian, NTChem, SMASH
Quantum Simulation	braket

Operation Flow of Interactive Application

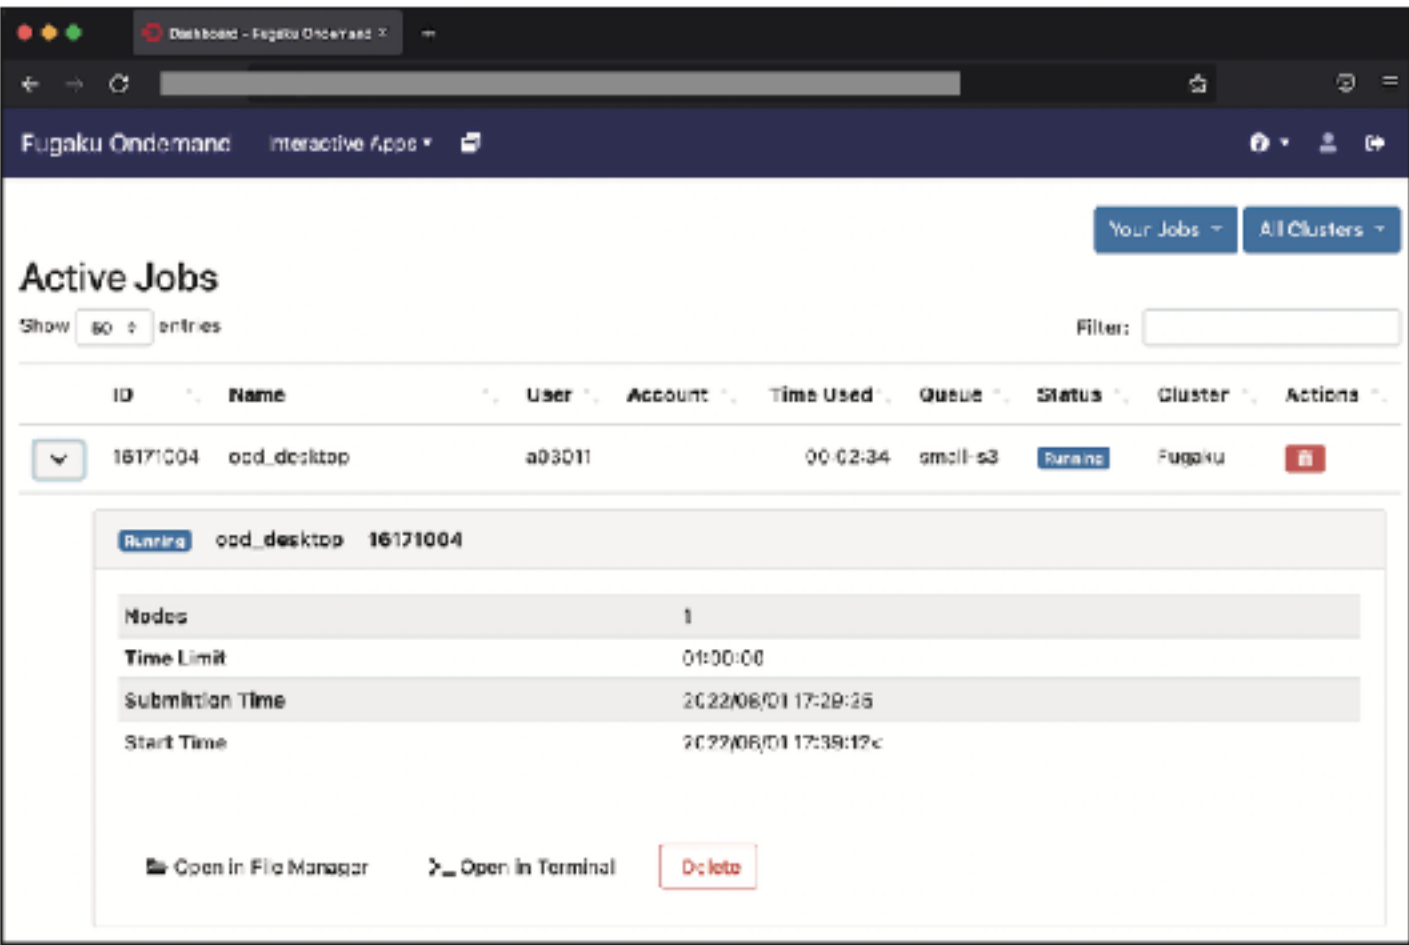
1. Log in to the Open OnDemand server using your web browser
2. User authentication for login
3. When issuing an execution command, the job is submitted to the calculation node
4. Wait until job is executed
5. When a job is executed, information such as the IP address of the calculation node is sent to the Open OnDemand server and a reverse proxy is set (for batch jobs, only the job is executed on the calculation node and ends here).
6. Connect to the calculation node inside the HPC system from a web browser using the reverse proxy URL



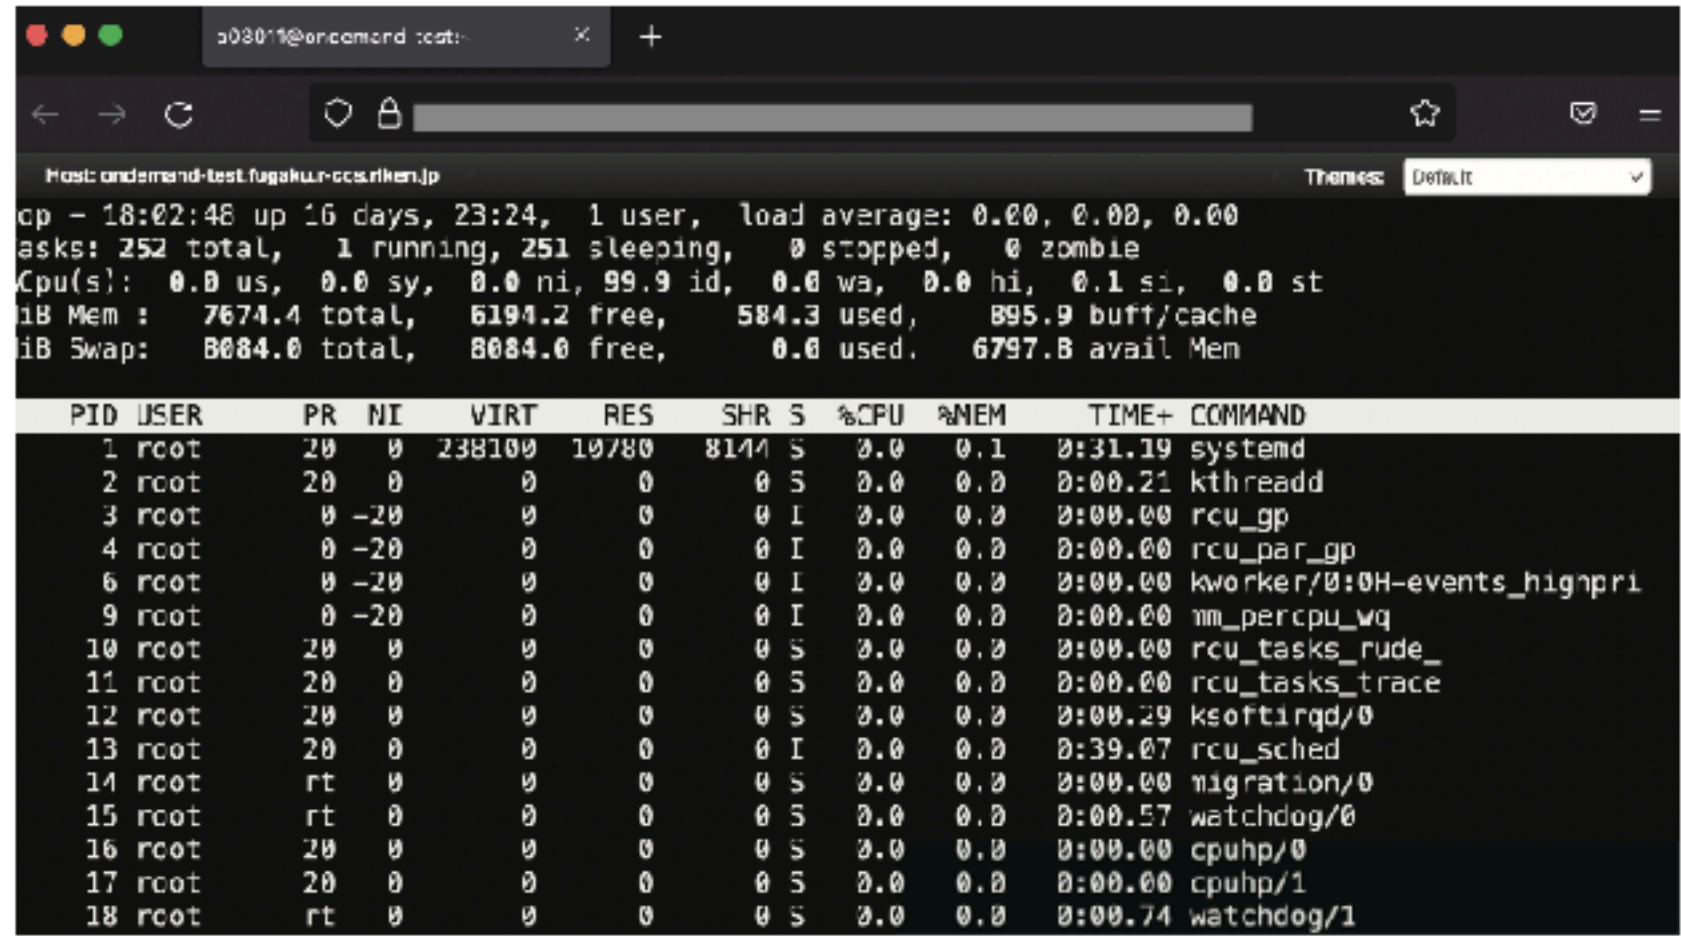
Passenger Applications on Fugaku Open OnDemand



Home Directory



Active Jobs




Shell

- Home Directory : Files can be uploaded, downloaded, and edited. By combining with rclone, it is also possible to transfer data with external storage such as Amazon S3.
- Active Jobs : Job monitoring
- Shell : Web-based Terminal

You can also develop new applications by using frameworks of Open OnDemand.

How to login Open OnDemand on Fugaku

- Please log into Fugaku Open OnDemand using your web browser



<https://ondemand.fugaku.r-ccs.riken.jp>

Fugaku OnDemand Batch Jobs Interactive Apps Passenger Apps My Sessions Develop Help Logged in as a03011 Log Out

Welcome to the supercomputer Fugaku

RIKEN R-CCS RIKEN Center for Computational Science

Link

- OnDemand Manual
- Fugaku Portal
- Fugaku Schedule
- Fugaku Status
- Fugaku Support

Message of the day

Information

- Jul 24, 2023 Operation July 2023 Large-scale job execution period
- Jul 21, 2023 Operation Resource groups during the large scale job execution period
- Jul 19, 2023 Operation Occurrence of inaccessibility and poor response at login nodes and jobs due to a file system maintenance(vol0005)

Pending jobs

Job Name	Count
fugaku-small	17887
fugaku-large	915
prepost-gpu1	12
prepost-gpu2	0
prepost-mem1	22
prepost-mem2	3
prepost-ondemand	0

Fugaku Schedule

Accounting (Updated at 2023/07/25 02:32:10 (JST))

Group	Volume	Disk (GiB)				Disk (inode)				Resource (NH)			
		Limit	Usage	Avail.	Rate	Limit	Usage	Avail.	Rate	Limit	Usage	Avail.	Rate
reco-bat	/vol0400	5,120	608	4,512	11%	1,500,000	265,253	1,234,747	17%	527,360	81,565	445,893	15%
f-cp	/vol0403	614,400	160,260	454,140	26%	180,000,000	144,261,343	35,738,656	80%	-	-	-	-
ra030002	/vol0403	5,120	1	5,119	0%	1,500,000	2	1,499,998	0%	-	-	-	-
rhome	/vol0400	20	0	19	0%	200,000	11,019	188,981	5%	-	-	-	-

Recently Used Apps

- Desktop
- OpenOnDemand
- SCALE
- Jupyter

Passenger Apps

- Active Jobs
- Home Directory
- Gakurin RDM
- HPC Storage
- Job Composer
- Fugaku Shell Access

How to launch Remote Desktop (for Meeting)

Click

Fugaku Schedule

Group	Volume	Usage	Limit	Avail.	Rate
rocs-aet	jvol0400	5,120	618	4,504	12%
f-op	jvol0403	723,184	184,772	538,412	35%
re030002	jvol0403	5,120	1	5,119	0%
/home	jvol0400	20	5	15	6%

Recently Used Apps

- WHEEL
- Desktop
- Desktop (for Meeting)
- Jupyter

Passenger Apps

- Active Jobs
- Home Directory
- Gakunin RDM
- HPC Shared Storage

Desktop (for Meeting)

This app will launch an **Xfce** desktop.

Elapsed Time (1 - 3 hours)

1

Number of nodes (1 - 4)

1

Number of procs (1 - 192)

1

Launch

* The Desktop (for Meeting) session data for this session can be accessed under the **data root directory**

Click

Enter the required information and submit the job.

Desktop (515) 1 node | 1 core | Running

Host: >_globus.cloud.r-ccs.riken.jp

Created at: 2023-10-11 10:15:25 JST


Time Remaining: 59 minutes

Click

Launch Desktop

View Only (Share-able Link)

Wait until job starts



Remote desktop can be operated from a browser

Active Jobs : Monitor jobs

Passenger Apps

Active Jobs

Click

Fugaku Shell Access

GakuNin RDM

HPCI Shared Storage

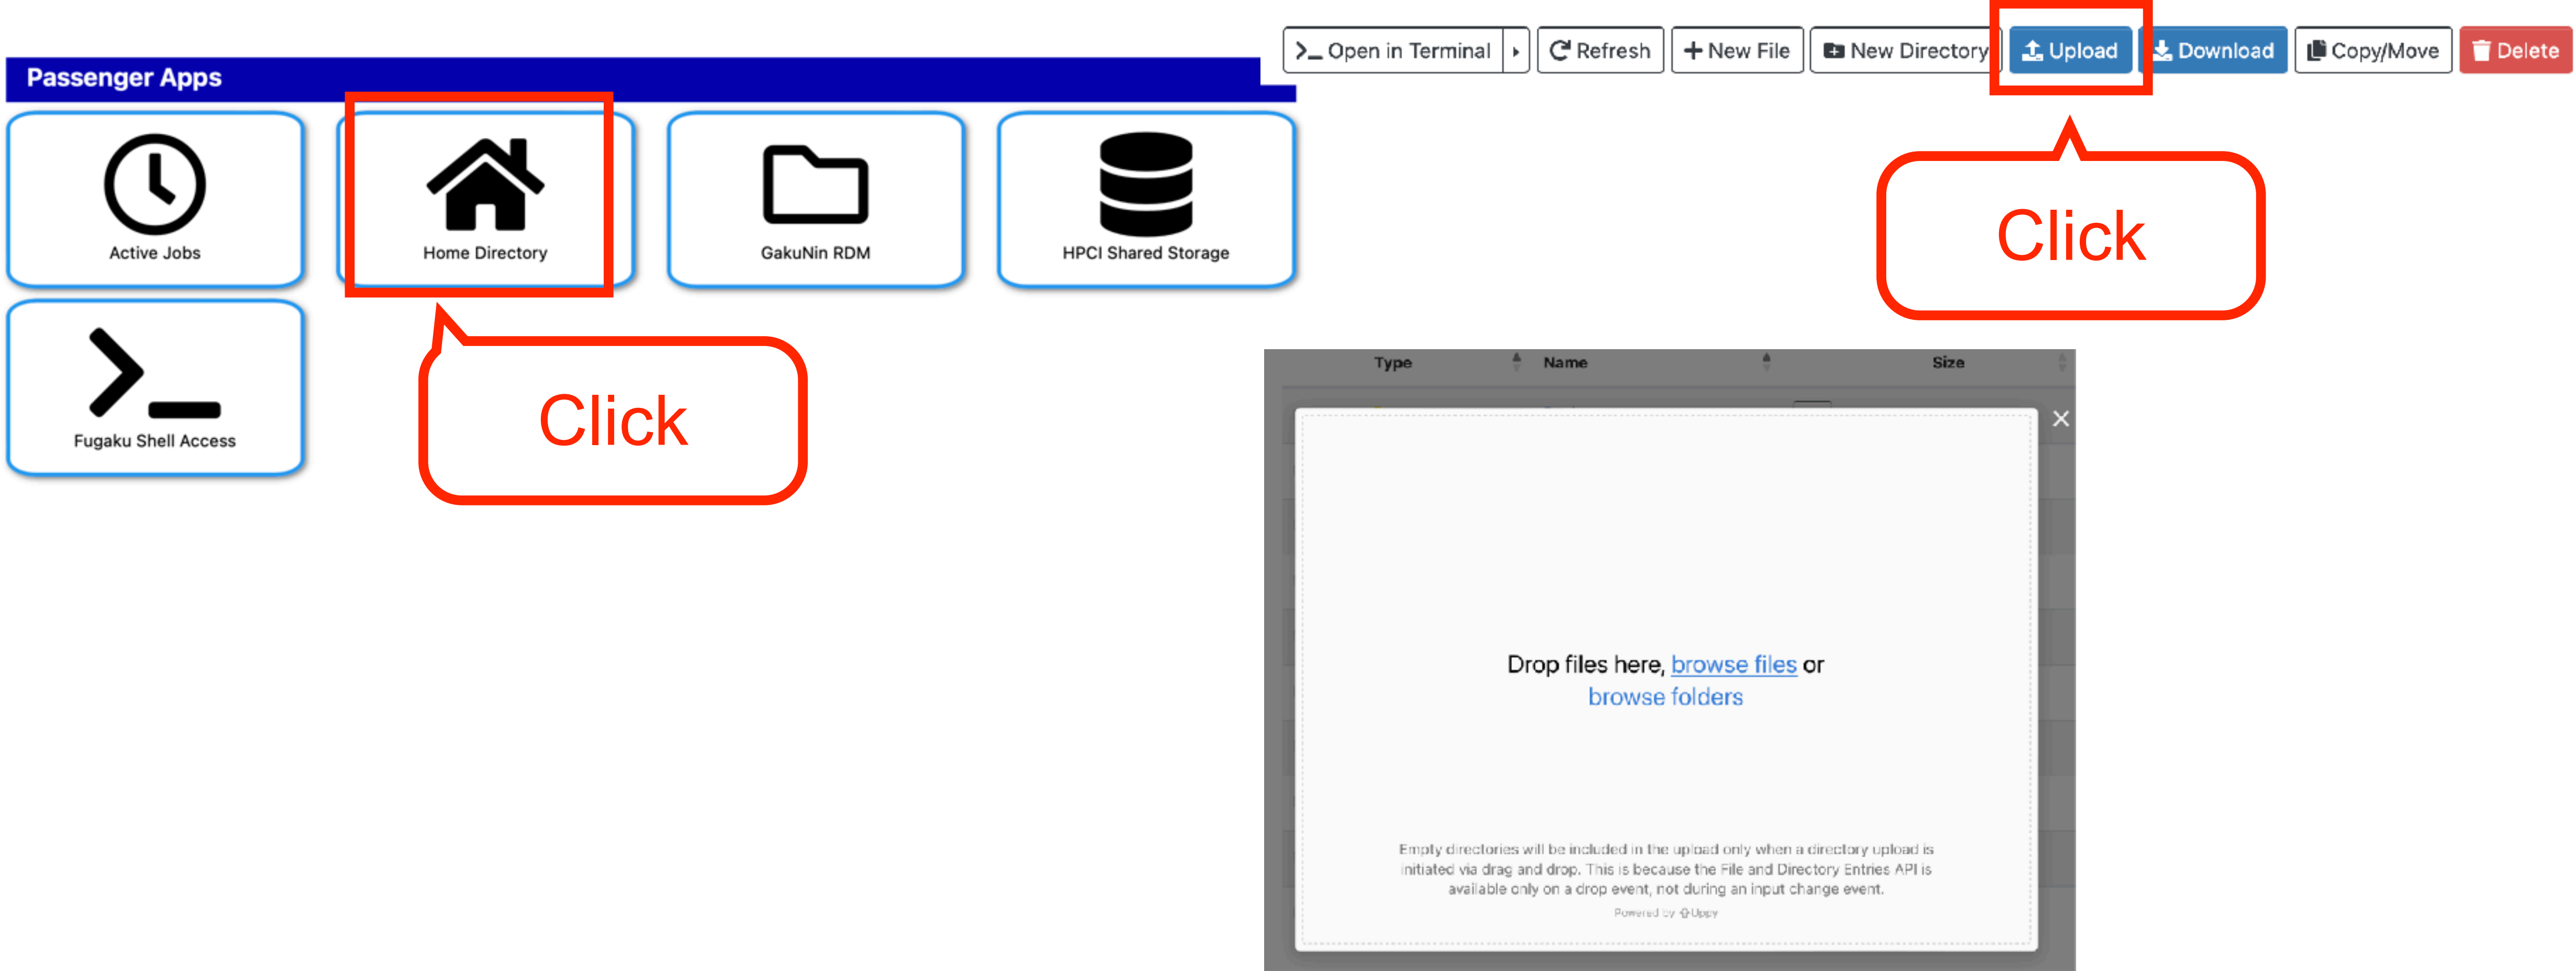
Click

510 sys/dashboard/sys/bc_desktop mnakao 00:00:21 debug Running Tenpozan

Click

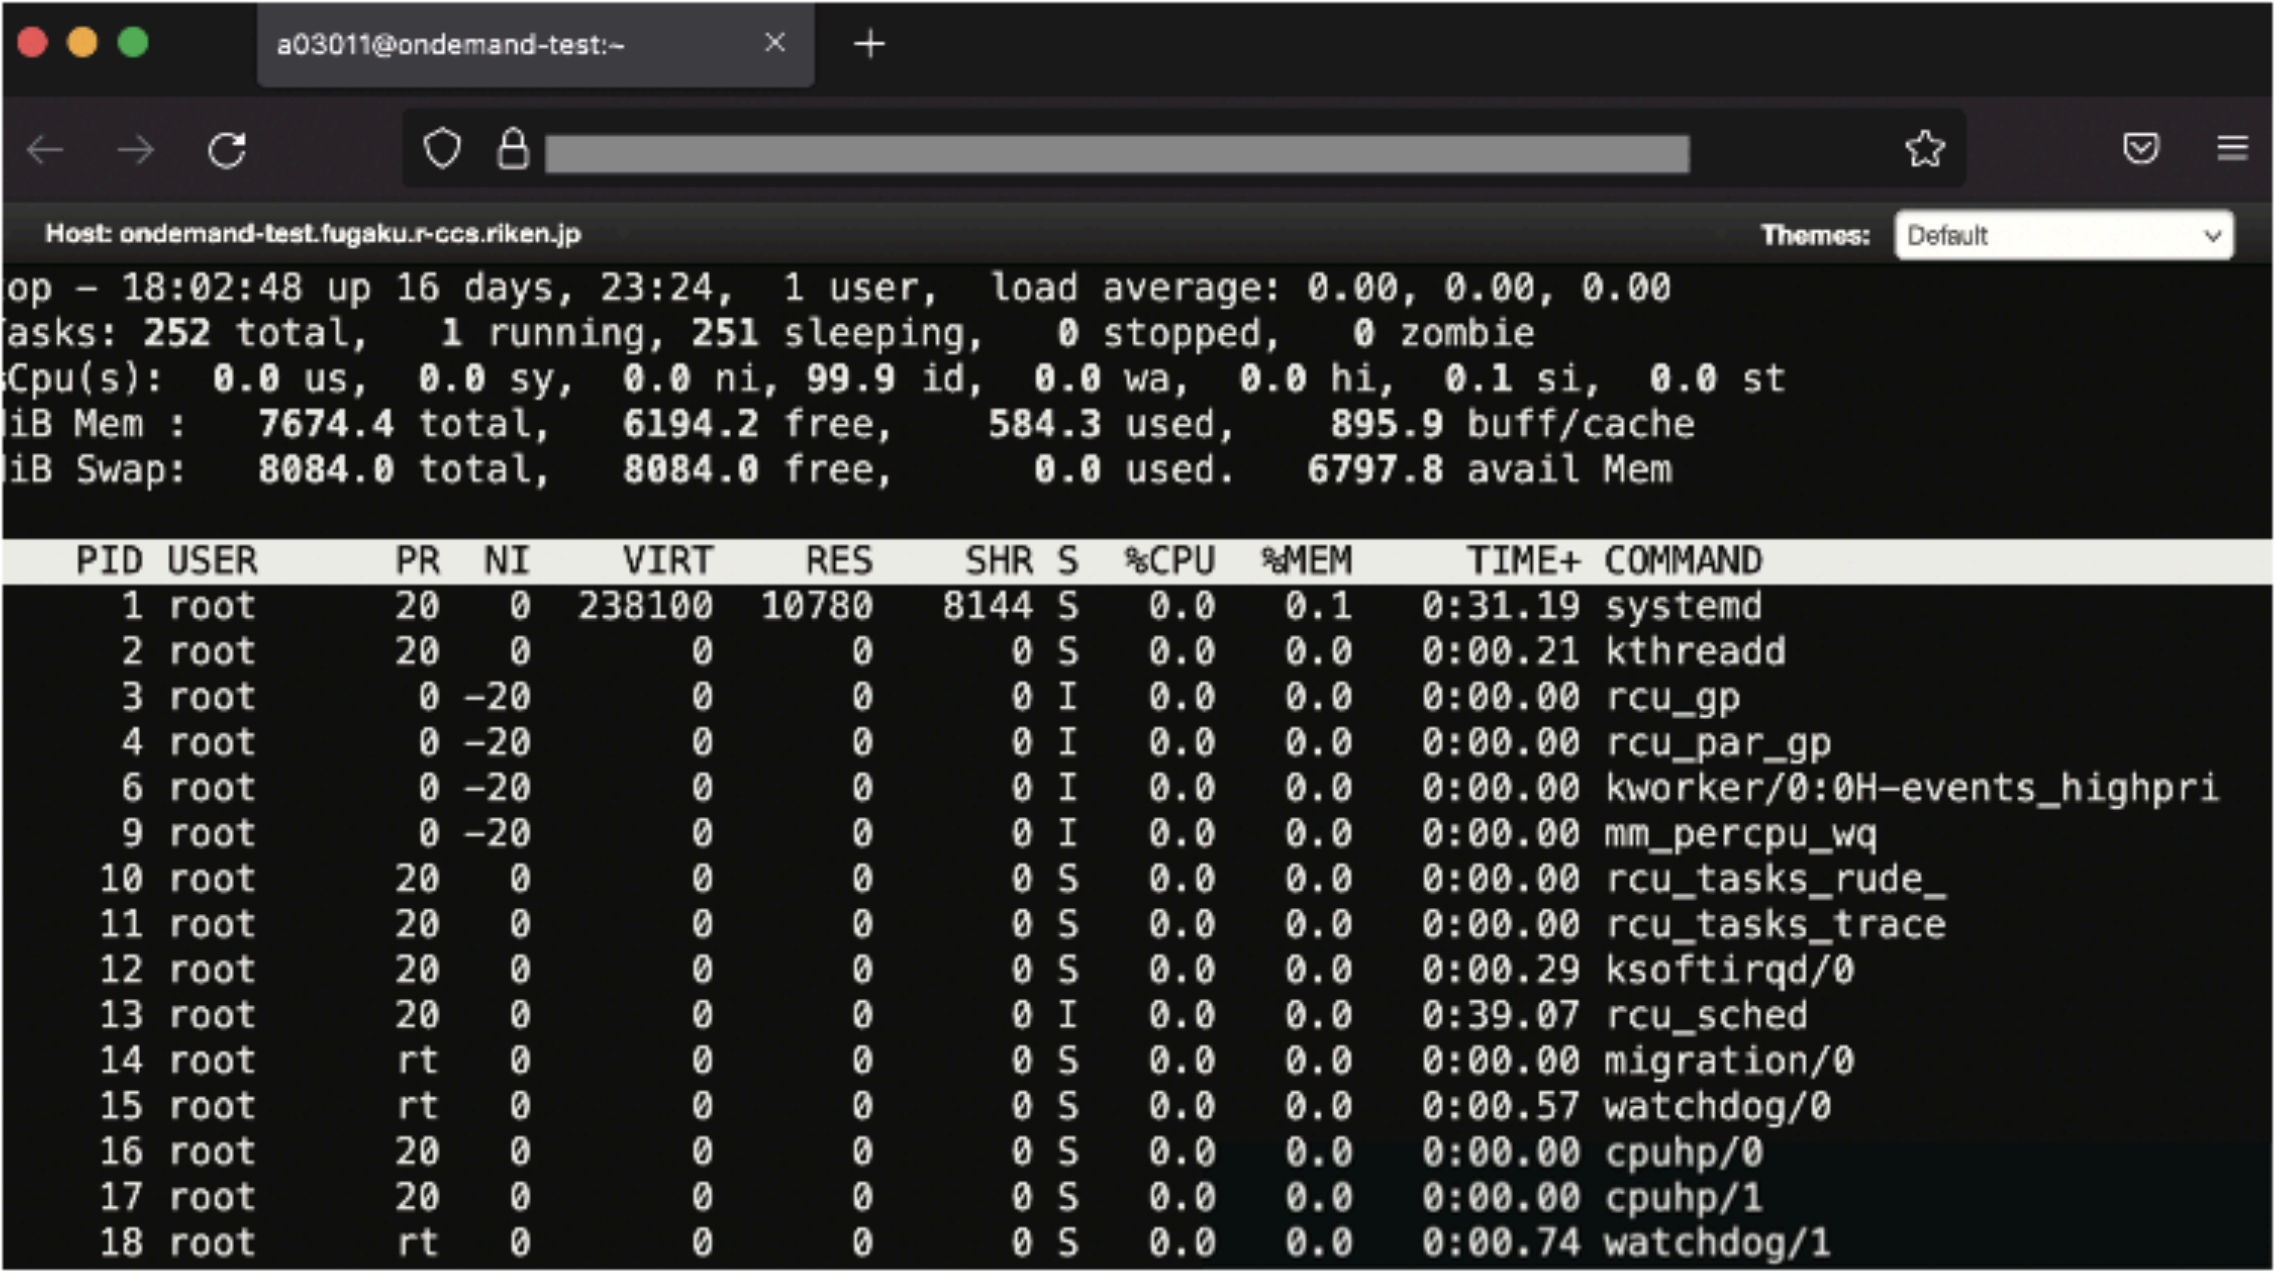
- Remote desktop job may be visible
- Detailed information of the job is shown by clicking the left icon
- You can stop the job by clicking the trash can icon on the right

Home Directory : Upload Files



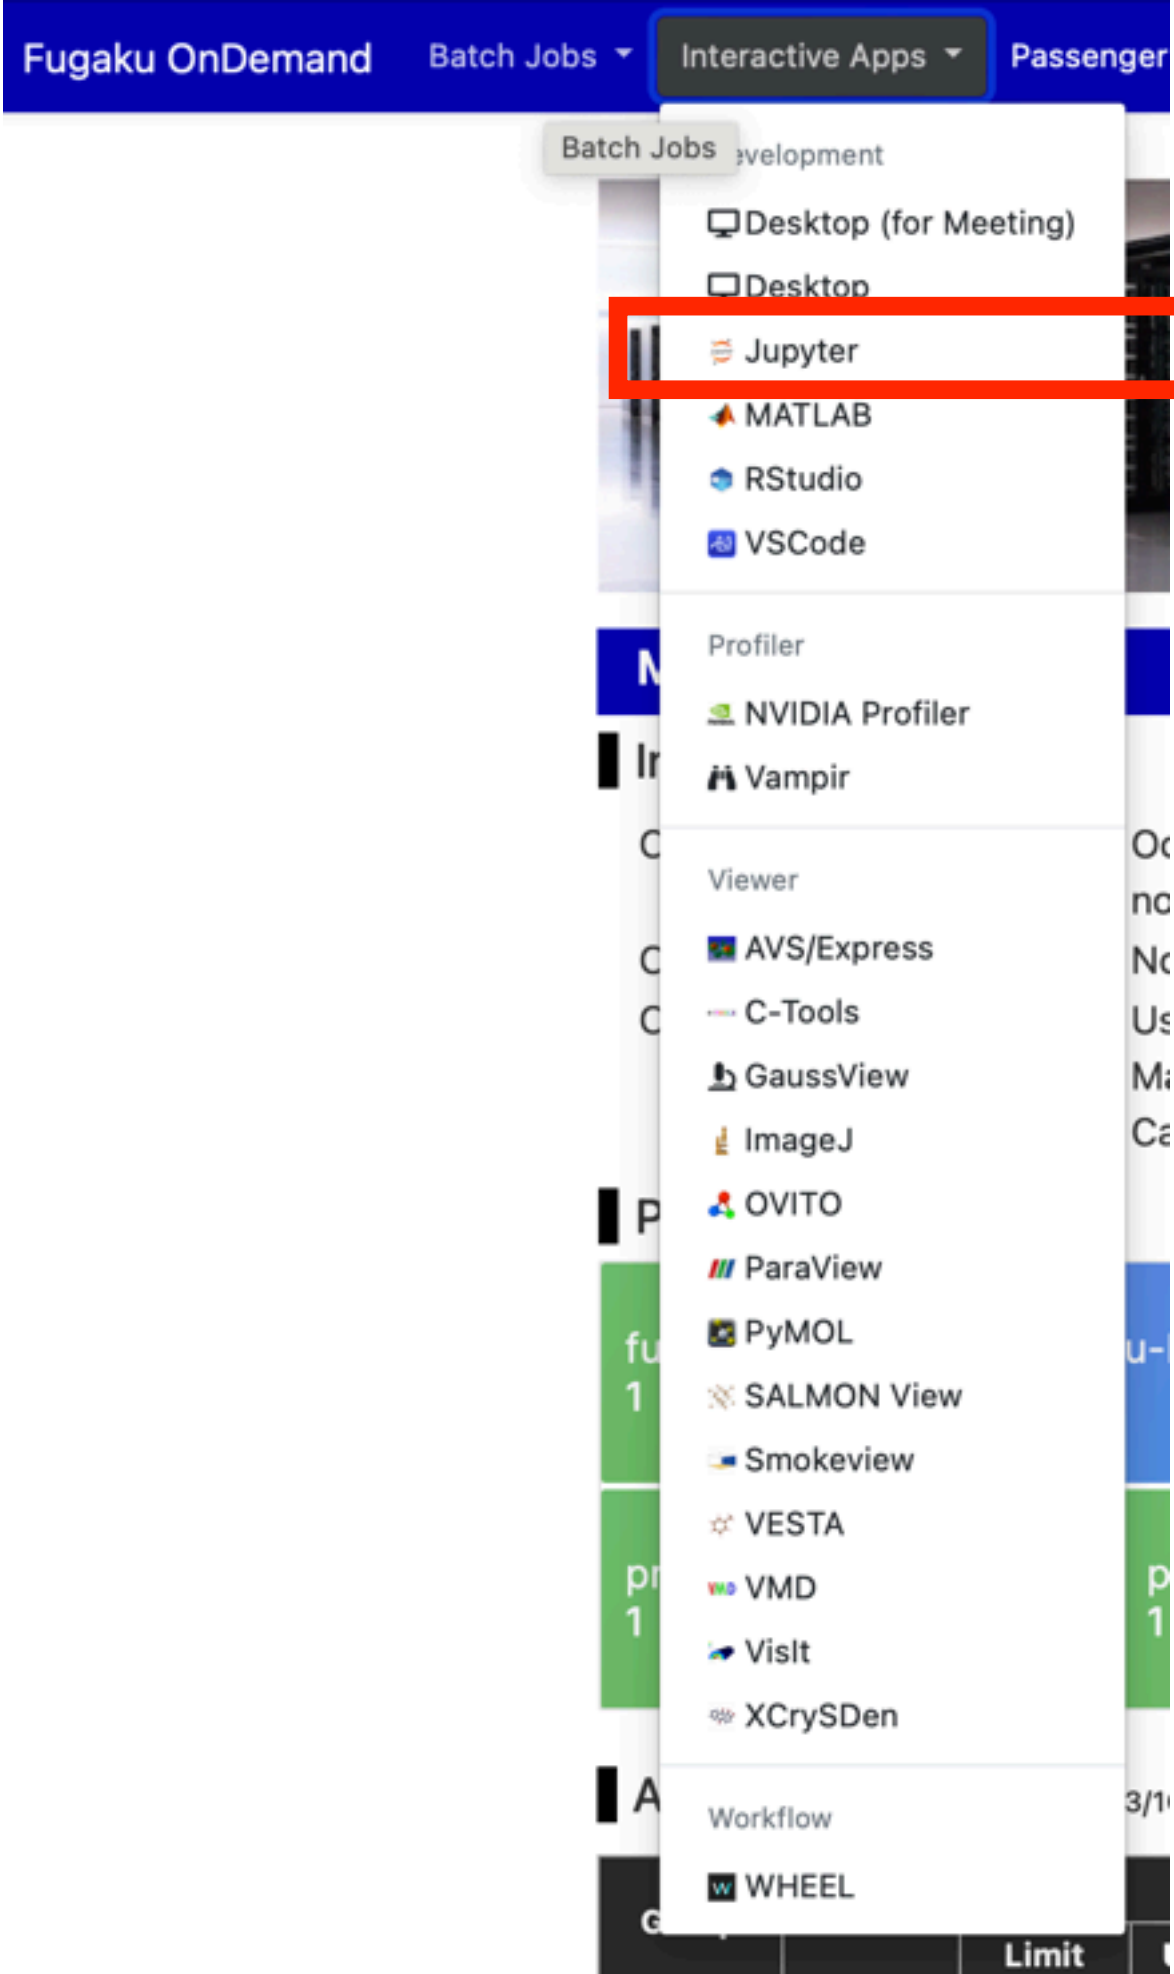
You can upload files by Drag and Drop.
Let's upload an appropriate file.

Job Composer : Web-based Terminal



Let's execute shell commands

Other applications



Jupyter
Jupyter is a web-based interactive computing platform.

Queue
prepost-gpu1

Elapsed time (1 - 3 hours)
1

Number of CPU cores (1 - 72)
1

Required memory (10 - 186 GB)
10

Number of GPUs (0 - 2)
0

Working directory
/home/rccs-aot/a03011

Select Path

Email (You will receive an email when it starts)

Launch

* The Jupyter session data for this session can be accessed under the [data root directory](#).

Please select prepost-*** queue