

Storage At OIST

Jan Moren, SCDA



OKINAWA INSTITUTE OF SCIENCE AND TECHNOLOGY GRADUATE UNIVERSITY

沖縄科学技術大学院大学

Today:

Our Storage systems



An overview

Research Data Management

Regulations, practical advice

Using the Storage

Best Practices



Storage At OIST

Long-Term Storage

Data, research results

Data Backups and Archive



Off-site storage


Working Storage


Temporary data during
analysis or computing


Transient Storage

Transfer data from instru-
ments or between units













Flash and Work

- Fast
- No backups
- Low latency









Dropbox

- Insecure
- Sharing data with external collabs

Research Data Management

- Data Lifecycle management
- Storage Organization



Both are important




Both are different!

Research Data Management

Handle your data from acquisition to disposal

- What is it, how much is there
- Where is it from, where is it now
- What is it for, who is responsible
- How is it processed, how do you read it
- Who should have access, what is the licence
- Where is the documentation

...




Rules and Regulations

MEXT: "research institutions need to adopt rules requiring researchers to **keep research data for a set period and to disclose the data as needed**"

Science Council of Japan: "Research data used as basis of a research publication need to be **preserved for a period of ten years** after publication."

OIST PRP 4.11: All researchers are asked to **preserve data in a secure, identifiable way** with appropriate backup methods, and disclose these when necessary.

Bucket is the system we provide for this



Rules and Regulations

Information Asset Manager:


The PI. Manages Unit Data

"Accountable for defining the classification level of, authorizing access to, and ensuring an appropriate permission scheme for, the information assets within their remit, in accordance to any and all relevant laws, rules and regulations."

Information Asset Administrator:

The RUA and/or researchers

"The Information Asset Manager may designate an Information Asset Administrator to manage information assets, but accountability remains with the Information Asset Manager."



Research Data: F.A.I.R

Findability

- Discoverable data and metadata, unique identifier

Accessibility

- (Meta)data access by open, standard protocols


Interoperability

- Open, well-defined data formats and cross references

Reusability


- Open license, clear provenance, community standards

<https://www.go-fair.org/fair-principles/>




Data Lifecycle Management

- Specify Requirements
 - Plan for storage, data security
- Ensure Data integrity and reusability
 - Data is tracked, validated, referenced, and credited
- Ensure Compliance and Funding
 - Requirement for grants, for human and animal research




Data Lifecycle Management



Data Management Plan


- What data, how much data?
Storage, bandwidth, cost
- For how long?
No data is stored indefinitely
- Tracking, labeling?
Data annotation, references crucial
- Where will it go?
What to keep, what to dispose of
- Who is responsible?



Manage Unit Storage

As a PI:

- Establish ground rules for the storage
 - The best time is when unit starts.
The next best time is now.
 - Aim to enable data management
- Organize periodic clean-up, archiving
 - Easier to keep clean over time
- Manage access (students, etc)




Manage Unit Storage

As a PI:

- Establish ground rules for the storage
- Organize periodic clean-up, archiving
- Manage access (students, etc)

As a Member:

- Keep your data organized
 - Document your data and processes
- Delete unneeded files
 - Build files, temp files, bad analysis, useless audio/video...



Manage Unit Storage


When you Leave the Unit:

PI:

- Remove the member from the unit storage
- Confirm what data can be deleted, archived and brought out by the member.

Member:


- Clean up garbage, archive old data
- Move data out *before* you leave! You lose all access when account expires.



Bucket Storage

Processes:

- Grant Access
- Get information
- Backup and Archive
- Increase Storage



Bucket Storage


Processes:

- Grant Access
- Get information
- Backup and Archive
- Increase Storage

Unit leaders
are
busy people

→ Assign a member who
can act in their name

Email us for more




Grant access for non-unit users

Who can **grant** access?

- The unit leader
- Any member they assign

<https://highsci.oist.jp>




Grant access for non-unit users

Who can get access?

- Graduate Students
- Other Unit Collaborators

Who can *not*?

- Non-OIST members



Grant access for non-unit users

Unit members (can't delete)

External user

They can add and remove users

Add and remove users by name


Group Management

example_group



name	admin	delete
Prof Unitleader	☒	
Doc Staffscientist	☐	
Clever Labtech	☒	
Diligent Student	☐	☒

* The account database is updated every 10 min.
* Shown in **Bold** are IT-managed members. Send [request form](#) to delete any of those.

person to add e.g. jan-moren



Storage Information



Storage Information


Storage Inspector

What: Browse unit storage

Who: PI and *assigned*
unit members

Where: Highsc "misc." menu

<https://highsci.oist.jp/monitor/storageinspector.html>




Storage Information


Storage Inspector

Give access on HighSci


The same way as
Storage Access



The screenshot shows a web interface titled "Storage Admins". At the top right are three icons: a question mark, a plus sign, and a minus sign. Below the title is a "Section" dropdown menu and a "refresh" button. A table with two columns, "name" and "delete", contains one row with a greyed-out name and a delete button. At the bottom is a text input field labeled "person to add e.g. jan-moren" and an "Add" button.



Storage Inspector



<https://highsci.oist.jp/monitor/storageinspector.html>



Storage Inspector

Directory/File Tree

Treemap

Usage of Individuals

Show/Hide Columns...

Sort by: Size (descendants)

	Size (descendants)	Size (self)	Duplicate Size (descendants)	Owner	Num of Files (self)	Num of Files (descendants)	Mtime (max)	Mtime (min)	Atime (avg)
▼ on Bucket	268.8T	31.1M	3.7T		21	92254853	2035/10/29 15:32:22	1979/11/30 00:00:00	2022/08/14 11:41:30
▶	74.9T	24.0K	84.0G		3	284027	2024/06/02 10:55:44	2012/04/21 00:42:55	2022/10/10 04:51:06
▶	47.1T	20.0K	0.0		3	185023	2024/05/29 17:36:20	2019/07/03 10:56:22	2024/05/19 05:02:43
▶	23.3T	2.8M	228.0G		8	11887	2023/10/02 11:22:32	2015/03/04 17:14:24	2022/09/02 11:56:20
▶	20.9T	92.0K	682.0G		7	1870676	2023/01/30 22:58:11	2015/12/23 17:59:42	2022/03/14 01:42:07
▶	12.9T	18.0K	0.0		3	15649	2024/07/12 13:43:33	2013/09/06 18:07:58	2024/01/14 14:13:29
▶	11.8T	30.0K	200.0G		3	536371	2024/04/21 18:15:26	1995/10/30 16:55:44	2022/03/22 06:53:48
▶	10.2T	80.7M	141.0G		36	13288716	2023/06/01 00:12:12	2007/01/11 09:16:26	2022/03/03 02:33:06
▶	9.5T	0.0	16.0G		2	19393	2020/07/03 23:10:30	2019/09/04 17:49:22	2022/08/17 10:10:44
▶	9.0T	26.0K	1.5T		2	3871356	2024/06/02 15:44:33	1980/01/01 00:00:00	2022/05/12 14:04:11
▶	8.7T	1000.0M	0.0		7	16358685	2024/06/27 20:22:25	2021/05/18 23:58:26	2024/01/31 12:07:58
(112 more directories/files)									



Mtime: Last modified time

Atime: Last accessed time

Avg: Average time

Min: Oldest file

Max: Newest file

Problematic:

"accessed" can include
non-interactive tools and
processes

Storage Inspector

Directory/File Tree

Treemap

Usage of Individuals

Show/Hide Columns...

Sort by: Size (descendants)

	Size (descendants)	Size (self)	Duplicate Size (descendants)	Owner	Num of Files (self)	Num of Files (descendants)	Mtime (max)	Mtime (min)	Atime (avg)
▼ on Bucket	268.8T	31.1M	3.7T		21	92254853	2035/10/29 15:32:22	1979/11/30 00:00:00	2022/08/14 11:41:30
▶	74.9T	24.0K	84.0G		3	284027	2024/06/02 10:55:44	2012/04/21 00:42:55	2022/10/10 04:51:06
▶	47.1T	20.0K	0.0		3	185023	2024/05/29 17:36:20	2019/07/03 10:56:22	2024/05/19 05:02:43
▶	23.3T	2.8M	228.0G		8	11887	2023/10/02 11:22:32	2015/03/04 17:14:24	2022/09/02 11:56:20
▶	20.9T	92.0K	682.0G		7	1870676	2023/01/30 22:58:11	2015/12/23 17:59:42	2022/03/14 01:42:07
▶	12.9T	18.0K	0.0		3	15649	2024/07/12 13:43:33	2013/09/06 18:07:58	2024/01/14 14:13:29
▶	11.8T	30.0K	200.0G		3	536371	2024/04/21 18:15:26	1995/10/30 16:55:44	2022/03/22 06:53:48
▶	10.2T	80.7M	141.0G		36	13288716	2023/06/01 00:12:12	2007/01/11 09:16:26	2022/03/03 02:33:06
▶	9.5T	0.0	16.0G		2	19393	2020/07/03 23:10:30	2019/09/04 17:49:22	2022/08/17 10:10:44
▶	9.0T	26.0K	1.5T		2	3871356	2024/06/02 15:44:33	1980/01/01 00:00:00	2022/05/12 14:04:11
▶	8.7T	1000.0M	0.0		7	16358685	2024/06/27 20:22:25	2021/05/18 23:58:26	2024/01/31 12:07:58
(112 more directories/files)									



Mtime: Last modified time

Atime: Last accessed time

Avg: Average time

Min: Oldest file

Max: Newest file

self: Only that folder

descendants: Everything below


the folder

Storage Inspector

Area: Size (total)

Color: Age (newest file)

→ Both can be freely changed, same as with File tree



Storage Inspector


Individual Tab

List total usage per unit member

Directory/File Tree Treemap Usage of Individuals				
Sort by: Size ▼				
██████████ on Bucket				
member name	size(byte)		num of files	last scan
██████████	20013236284402	18T	307842	2026-02-25 03:13:57
	18320637836773	17T	74600	2026-02-25 03:10:43
	15951531107913	15T	1105421	2026-02-25 03:13:20
	12585563803477	11T	347062	2026-02-25 03:11:38
	5274900608511	5T	123917	2026-02-25 03:10:56
	1885993718547	2T	66544	2026-02-25 03:11:44
	889388862250	828G	21506	2026-02-25 03:11:05
	854635117288	796G	7126	2026-02-25 03:13:59
	178355897839	166G	677	2026-02-25 03:11:01
	107367774667	100G	2602	2026-02-25 03:10:45
	19108439636	18G	582	2025-09-24 03:39:13
	5365287765	5G	141	2026-02-25 03:13:25
	4466292279	4G	77	2026-02-25 03:11:46
	2741497175	3G	957	2026-02-25 03:11:02
	1395190518	1G	667	2026-02-25 03:11:47
	122475616	117M	14	2026-02-25 03:10:58
	5767968	6M	30	2026-02-25 03:13:23
	0	0	0	2026-02-25 03:13:22
	0	0	0	2026-02-25 03:10:59



Backup Data



Why a new backup system?

The tape backup was designed for ~10PB storage

- At 7PB, backups take ~3 days
- ~15PB, we're out of time and space

Backups are meant for *disaster recovery*

- Only business-critical data needs off-site backups

And human mistakes


For disaster recovery:

- Only critical subset can be restored
- Need good backup organization

For file restoration

- Usually only most recent version
- Need good backup organization

Backup Data →



Backup Folder

- In every units top folder
- 100TB of Local Backup
- Weekly Remote Snapshots
 - 5 snapshots or 6 months
- Just* copy your data over


* terms and conditions apply. Subject to change. Not valid when invalid.

Our old system:

Only remote snapshots
(hard to get files back)

Only 2 snapshots and 6
months

Backup Data →




Backup Folder

- In every units top folder
- 100TB of Local Backup
- Weekly Remote Snapshots
 - 5 snapshots or 6 months
- Just* copy your data over

* terms and conditions apply. Subject to change. Not valid when invalid.


As long as the data is in Backup/ there is **always** at least one snapshot!

Backup Data




Backup Folder

- In every units top folder
- 100TB of Local Backup
- Weekly Remote Snapshots
 - 5 snapshots or 6 months
- Just* copy your data over




Backup Data




How to Use:

- Put your data in Backup/
- *Only* archive files:
 - Zip, tar, tgz, ...
 - archive directly into Backup/
- Add a separate description file



Backup Data




Create (or update) a zip archive

```
# go to your bucket storage
$ cd /bucket/UnitU


# create the backup archive
$ zip -r Backup/Jan_neuron.zip neurondata/
```

Where we put our backup

The folders we
want to back up




Backup Data




For a new version:

- overwrite the previous archive
- keep the same name

```
# update the backup archive — same command!  
$ zip -r Backup/Jan_neuron.zip neurondata/
```



Backup Data



Description file

- Same name as archive, ".txt" ending
- Contains:
 - **Name:** User name
 - **Project:** Project name
 - Free-form text description

The purpose is for
future you and your PI to
know what the backup contains!

Archive:

jan_neuron.**zip**

Text file:


jan_neuron.**txt**

Name: jan-moren

Project: neuron spike rebound model

Model code (NEST) in git. Branch
"publish" is version for the paper
...

Backup Summary (CLI)



1. Archive file: Backup/jan_neuron.zip


```
$ zip -r Backup/Jan_neuron.zip neurondata/
```

2. Description file: Backup/jan_neuron.txt


```
$ nano Backup/Jan_neuron.txt
```

3. Update the backup:

```
$ zip -r Backup/Jan_neuron.zip neurondata/
```




Backup Summary



1. **Archive file:** Backup/jan_neuron.zip
2. **Description:** Backup/jan_neuron.txt

Other files than these
will not get a remote
snapshot! Do not put
other data here.



Backup Summary




1. **Archive file:** Backup/jan_neuron.zip
2. **Description:** Backup/jan_neuron.txt

Changed archives are copied
to tape weekly

Up to 5 versions are stored offsite

Up to 6 months old.

Always at least 1 copy





Data Recovery

You've lost a file – what to do?

Your most recent backup is in Backup/
→ get your file back directly

For older versions contact us
and tell us:

- Name and path of the file
- Name of the archive
- The date you want



Increase Bucket Storage

Initial allocation is 50TB per unit


The unit can ask for an increase

👉 Only once you actually need it
Do let us know your needs well in advance!

🙏 Please **clean up** and archive old data

👉 Typically we increase by **50TB** at a time

<https://scda-answers.oist.jp/books/storage/page/request-bucket-storage>



Manage files

Permissions on Bucket are very (too) simple:

owner



entire unit



(everybody)

Default:

You can:

- Read
- Write
- Run as a program

myfile.data	
Owner:	Entire unit:
✓read	✓read
✓write	✗write



Manage files

Doable:

Owner



Entire unit



Simple perms:

- Read
- Write
- Run program

Not Doable:


No only part of unit



No Individual perms



No External access



No complex perms:

- add but not change
- create but not delete
- Read but not copy
- . . .


No special access
for the unit leader




User Data Access

- A former member has left data on Bucket (or Flash or Work)
 - Nobody can move it or delete it
-
- Email us at ask-scda@oist.jp (CC your PI)
 - Tell us **who** you want to own the data

We change the
ownership




They deal with
the data



Flash And Work

Short-term storage while running a job



Flash And Work

deigo

/flash




500TB scratch
10TB/Unit

Allocated:
1200 TB

saion

/work



250TB scratch
10TB/Unit


Allocated:
900 TB



Do you have a
job running?

NO?

you should have



0 bytes

on Flash and Work



HPAcquire


These Systems are a *Problem*:

Instruments




Don't get updates
No antivirus

Publicly
Accessible



Security risk

Shared accounts
Non-SSO accounts




No accountability



HPAcquire

Isolates the OIST storage
from insecure systems



- Per-user storage
- Limited retention time (21 days)
- Manual movement
- Possible to copy data back




Comspace

Remote destop filesystem with
fine-grained access controls

- ✓ 5TB per unit
- ✓ Managed from Windows
(Windows-style permissions)
- ✓ Any number of workspaces
- ✓ Fully controlled by the PI
- ✗ No direct access from the
clusters or other storage
- ✗ No management from Mac
or Linux
- ✗ Can't add external people



Share Data



Share Data

`/datashare`

On login nodes:

`/datashare`

- Copy data over
→ Is readable by all
- Recipients copy data
- Autodeleted

Useful for:

- Share data with user from another unit
- Distribute materials for a course
- Give us code or data for debugging



Share Data

Filesender


Go to: <https://filesender2.oist.jp/>

- Upload data (up to **500G**)
- Enter email address
- They get a download link
- Active for **14 days**
- **Can also let others send data to you**



Bucket File Management

- Many small files are *bad*.
 - Take long time to copy, move, find
 - Remote folders become useless
- Instead, use 'tar' or 'zip':
 - A single file/archive is easy to manage



Administrative files

Bucket: *No, No, NO!!*

Not secure enough! Too easy to give access to unit members by mistake!

Units have exposed confidential information – **individual salaries, evaluations, medical leave of absence** – to everyone in the unit.

- Personnel files
 - Salaries, evaluations, contracts etc.
- Applications
 - Grants, ethics, ...
- Committee documents
- Budget, purchases, travel ...



Administrative files

Bucket: **No, No, NO!!**

Dropbox: **Absolutely Not**

Not secure enough! We see all documents. Easy to share with wrong person by mistake.

Only licensed for sharing external research data.

- Personnel files
 - Salaries, evaluations, contracts etc.
- Applications
 - Grants, ethics, ...
- Committee documents
- Budget, purchases, travel ...



Administrative files

Bucket: ***NO!!***

Dropbox:
Absolutely Not

ComSpace: **Maybe**

Good security and permissions control. Need a Windows computer to manage.

If you already use Comspace it's a solid alternative.



Administrative files

Bucket: ***NO!!***

ComSpace: **Maybe**

Dropbox:
Absolutely Not

SharePoint: **Yes!**

This is what it's for. Widely accessible, fully supported for all OIST document processing.



Administrative files

Bucket: ***NO!!***

ComSpace: **Maybe**

Dropbox:
Absolutely Not

SharePoint: **Yes!**

This is sensitive data

Use the appropriate place
to store and process it



PC Backups and Personal Files

Not on Our Systems

Bucket etc. is for *research data*.

Also: Your backups and personal data includes software, operating system files, media and so on that are copyrighted and *not allowed* on our storage.




Unlicensed Software Unlicensed Movies, Music etc.

We Will Report You

**You Will Lose Access
to All Our Systems**

Pirated content on our storage would
become a *major* legal problem for OIST.

You will take the blame.



When You Leave

- Data in unit storage belongs *to the unit*
You+PI decide what you can bring with you
- Your home becomes owned **by OIST**, not you.
- Getting access **after you left**
Complicated – and is often **not granted**
- We **can't transfer data** to hard-drives or similar




Figure this out *before you leave*