

HPC User Group Meeting

March 11 2024

Jan Moren, SCDA



OKINAWA INSTITUTE OF SCIENCE AND TECHNOLOGY GRADUATE UNIVERSITY

沖縄科学技術大学院大学

Today:

Post-Deigo

- What is it?
- What's happening?
- What can we expect?
- When?

Deigo Trouble

Work in Progress

Matlab

What's going on?!

Bioinformatics User Group

The latest updates!

Deigo Trouble

You have probably seen the error:

```
launch failed requeued held
```

“**launch failed**”:

- Slurm reserved nodes for you, and tried to start your job
- Something went wrong — most likely a broken node

“**requeued held**”:

- Slurm put your job back into the queue (“requeue”);
- And marked it “held”: it will not actually start until somebody “*releases*” it again.

Deigo Trouble

“launch failed”:

- Slurm tried to start your job but failed due to a broken node

“requeued held”:

- Your job is requeued and held until released or removed.

What do you do?

- You *could* release the job:

```
$ scontrol release <jobid>
```

but you’ll get the same (broken) nodes again.

- You can’t directly find the node that is broken

→ Tell us the job ID at: ask-scda@oist.jp

- Until the node is fixed, you can *exclude* specific nodes with the “-x” option:

```
$ sbatch -x deigo012345 my_job_script.slurm
```

Deigo Trouble

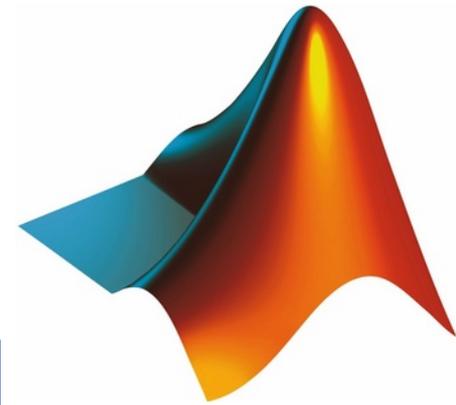
Why does this happen in the first place?

- Sometimes it's a simple problem: the node failed to log you in
- Sometimes it's broken hardware: bad memory, failing CPU...

But, we have a recurring issue where a combination of different jobs can force a node to run out of memory.

- We have some ideas, but not yet a fix
- We may need to upgrade Slurm
- We may have to slightly reduce the maximum amount of memory per node.

Matlab 2023 on Deigo

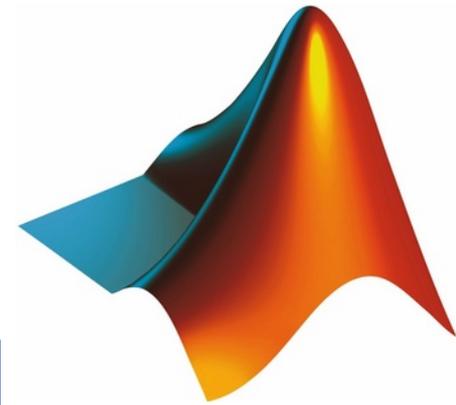


When you start any kind of UI, you get:

```
Inconsistency detected by ld.so: ../elf/dl-tls.c: 517:  
_dl_allocate_tls_init: Assertion `listp != NULL' failed!
```

- Matlab triggers a bug in “glibc” <2.36 (we have 2.28)
- glibc is the front-end for the Linux kernel itself
updating glibc == upgrading the kernel
- Mathworks supports CentOS 8.6 and up.
We have CentOS 8.1.
- Redhat/IBM has dropped all support for CentOS, so
we would have to upgrade to a new, different OS

Matlab 2023 on Deigo

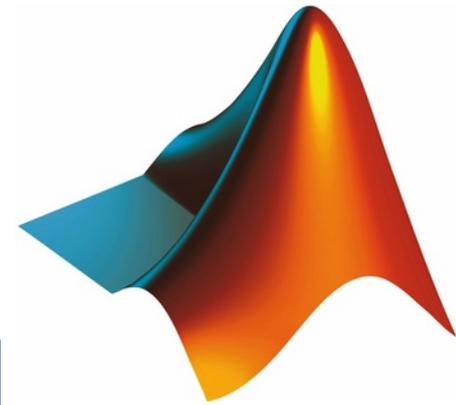


When you start any kind of UI, you get:

```
Inconsistency detected by ld.so: ../elf/dl-tls.c: 517:  
_dl_allocate_tls_init: Assertion `listp != NULL' failed!
```

- Upgrading the Deigo OS would take several months:
 - Pick replacement, confirm vendor support
 - test with drivers and user level software
 - Plan downtime and plan for failure (how to revert)
 - Do the upgrade

Matlab 2023 on Deigo



When you start any kind of UI, you get:

```
Inconsistency detected by ld.so: ../elf/dl-tls.c: 517:  
_dl_allocate_tls_init: Assertion `listp != NULL' failed!
```

- Upgrading the Deigo OS would take several months:
 - But post-Deigo is coming
→ we're already starting the process
 - We lack people (2 open system engineer positions)

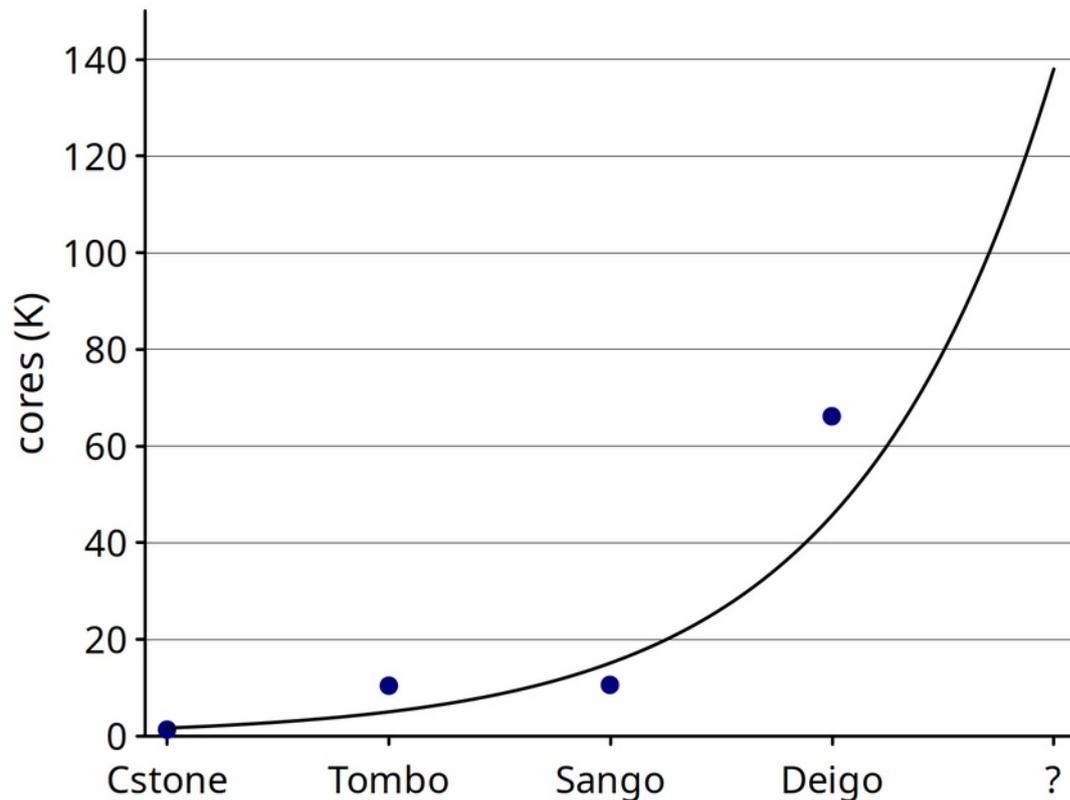
**Is it worth it? *Can* we do it with current resources?
While also planning and implementing post-Deigo?**

Post-Deigo

Deigo (and Tombo) were big leaps forward

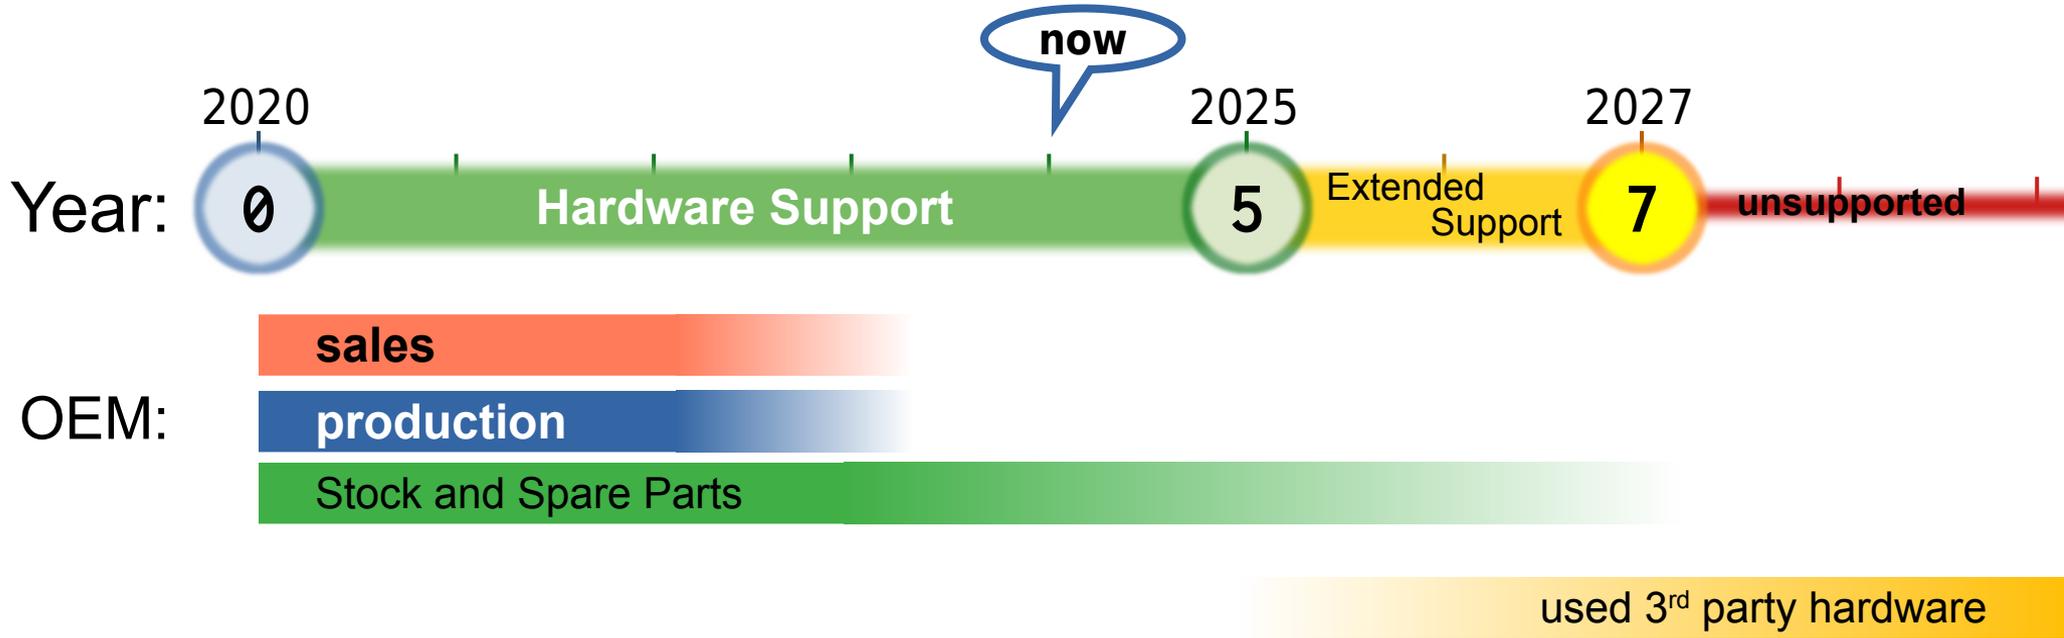
- Tombo and Sango had ~same number of cores
- Deigo is larger than we hoped and expected

→ Post-Deigo unlikely to be as big a jump



Post-Deigo

All Good Things Must Come To An End

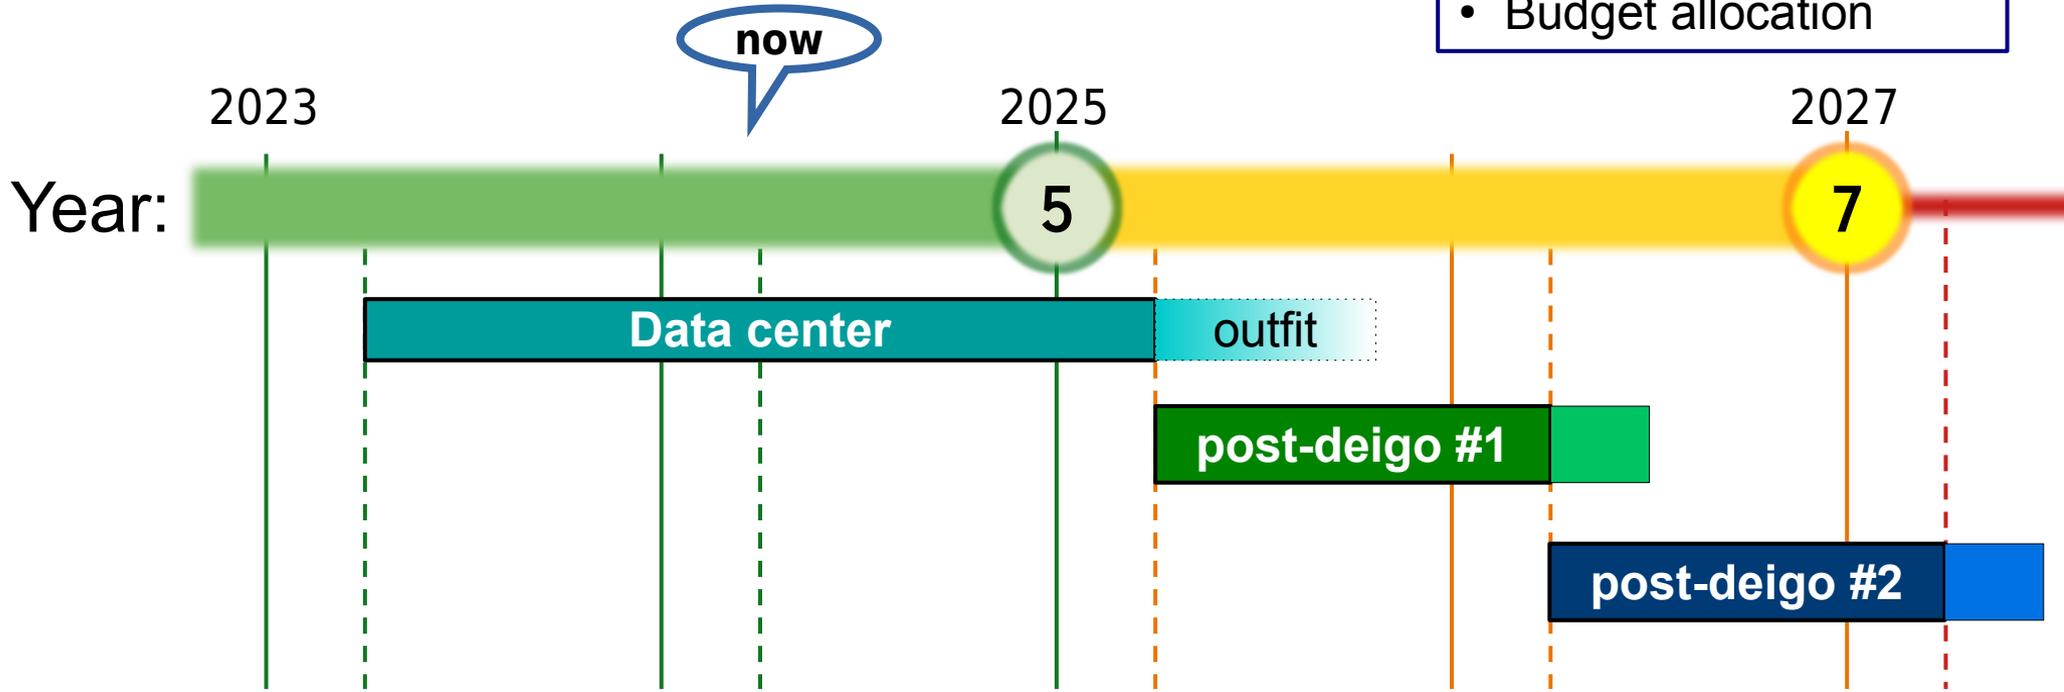


Post-Deigo

Rough timeline(s)

#1 or #2?
That depends on:

- Datacenter completion
- Budget allocation



Data center

Why new data center?

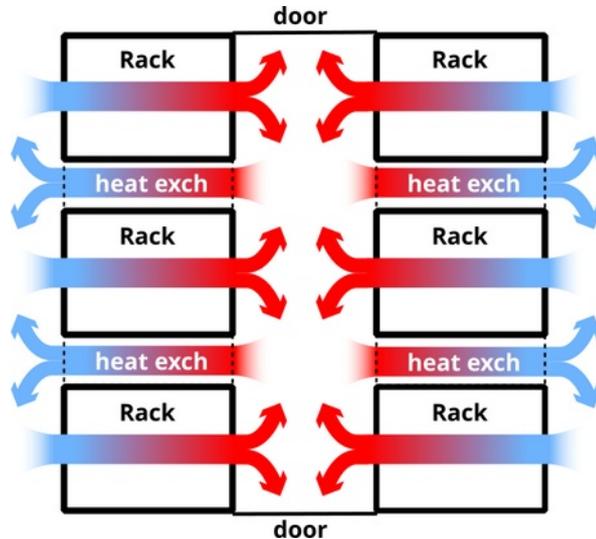
- **More Power**
 - Current site (Lab 1) is restricted
- **Better facilities** and access
 - An actual loading dock
 - Separate rooms for compute and storage
- **Water cooling**



Cooling Systems

Air cooling

Draw cold air through the nodes, cool the hot air with heat exchanger.

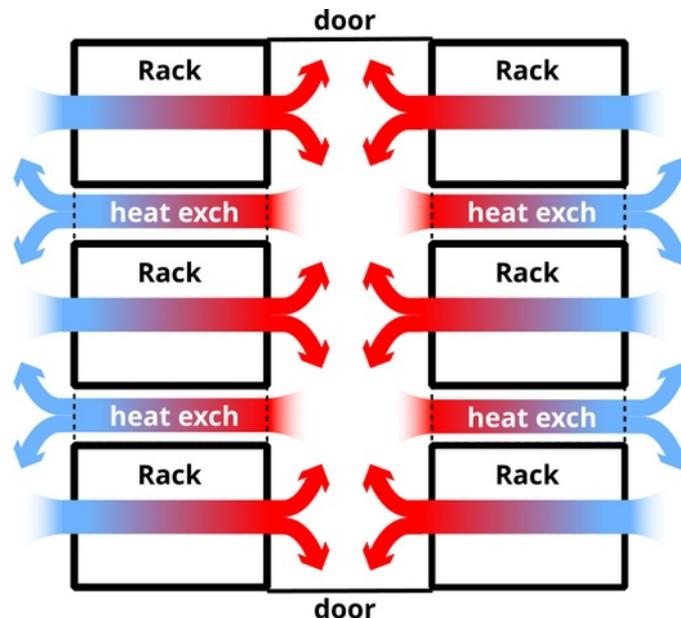


Cooling Systems

Air cooling

Draw cold air through the nodes, cool the hot air with heat exchanger.

- What we currently use
- Takes ~50% extra power
- Easy maintenance
- Completely standard
- (relatively) inexpensive



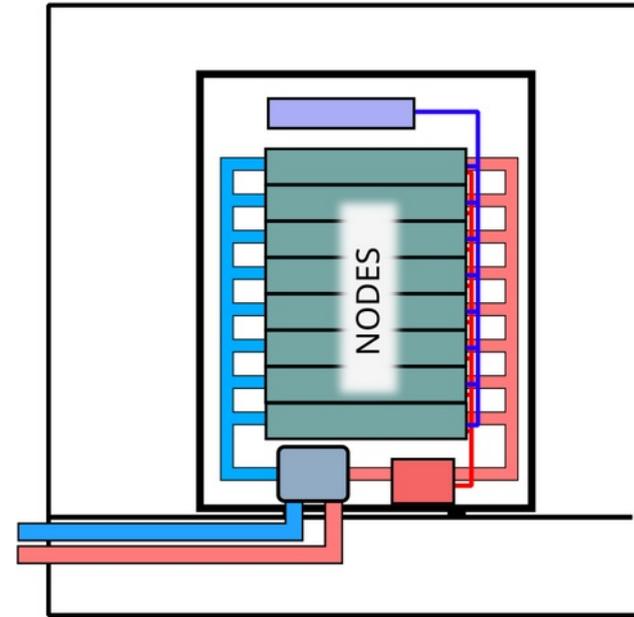
Not feasible for next generation systems

Cooling Systems

Direct liquid cooling

Pumped liquid cools each component, heat exchanger inside every rack.

- Only ~10-15% power
→ 40+% less total power!
- Costs more, lack of standards
- More difficult maintenance



Most likely for post-Deigo

Cooling Systems

Immersion cooling

Immerse entire system in liquid, “aquarium” pump cools and filters.

- Only ~5% power
→ By far most efficient
- Really heavy, lower density
- Maintenance a major problem
- Still experimental



Unlikely for this generation

Post-Deigo Specs

– still very uncertain

Nodes:

If FY 2025: 144-192 cores
~384-768 GB

If FY 2026: 192-256 cores
~512-1024 GB

CPU: X86-64 (AMD or Intel)

System:

Likely ~650 nodes
(yes, same as now)

PCI, memory type,
other partitions,
new storage etc:

???

Applications

We have cancelled the license subscriptions for:

- Fimmwave (waveguide design)
- OMNISIM (photonics)
- MaterialsStudio (simulation)

We buy software when:

- multiple units expect to use it;
- Common licences save money

We cancel when:

- No usage for 1+ year
- Units tell us they don't use it
- Keeping it is an ongoing cost

**Ps: Please ignore “licence is about to expire” messages.
We know, and we will renew in time.**

Dropbox and Your Personal Data

OIST Dropbox accounts are only for sharing research data

- **No personal files**

→ Get a free Dropbox account of your own!

- **No OIST personnel or administrative files**

→ Use Sharepoint — it's explicitly approved for that

No pirated stuff
Seriously, just don't.

Dropbox and Your Personal Data

OIST Dropbox accounts are only for sharing research data

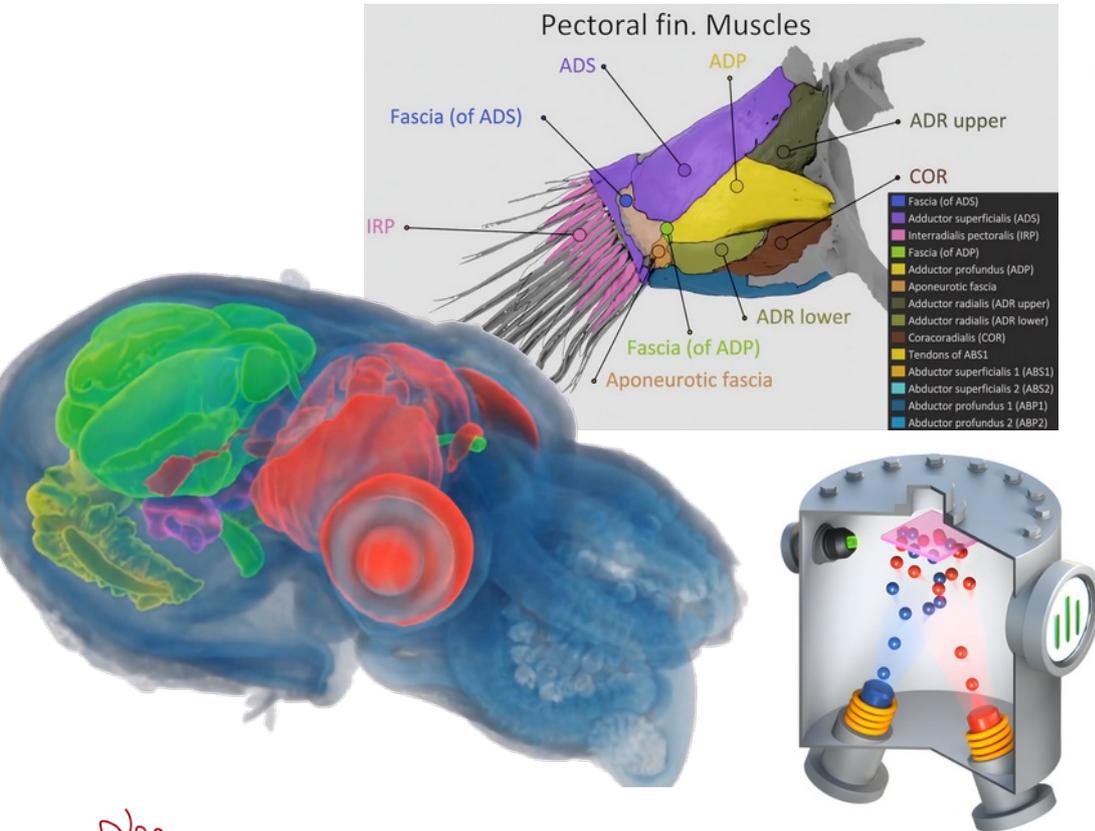
Comspace is well-suited for any restricted data:

- 5TB/unit
- Easy access from desktop
- Share across units
- Granular control over access rights

Partial text in surrounding yellow boxes:

- No → C
- No → U
- at

Do you want *your* papers to look *this* good?

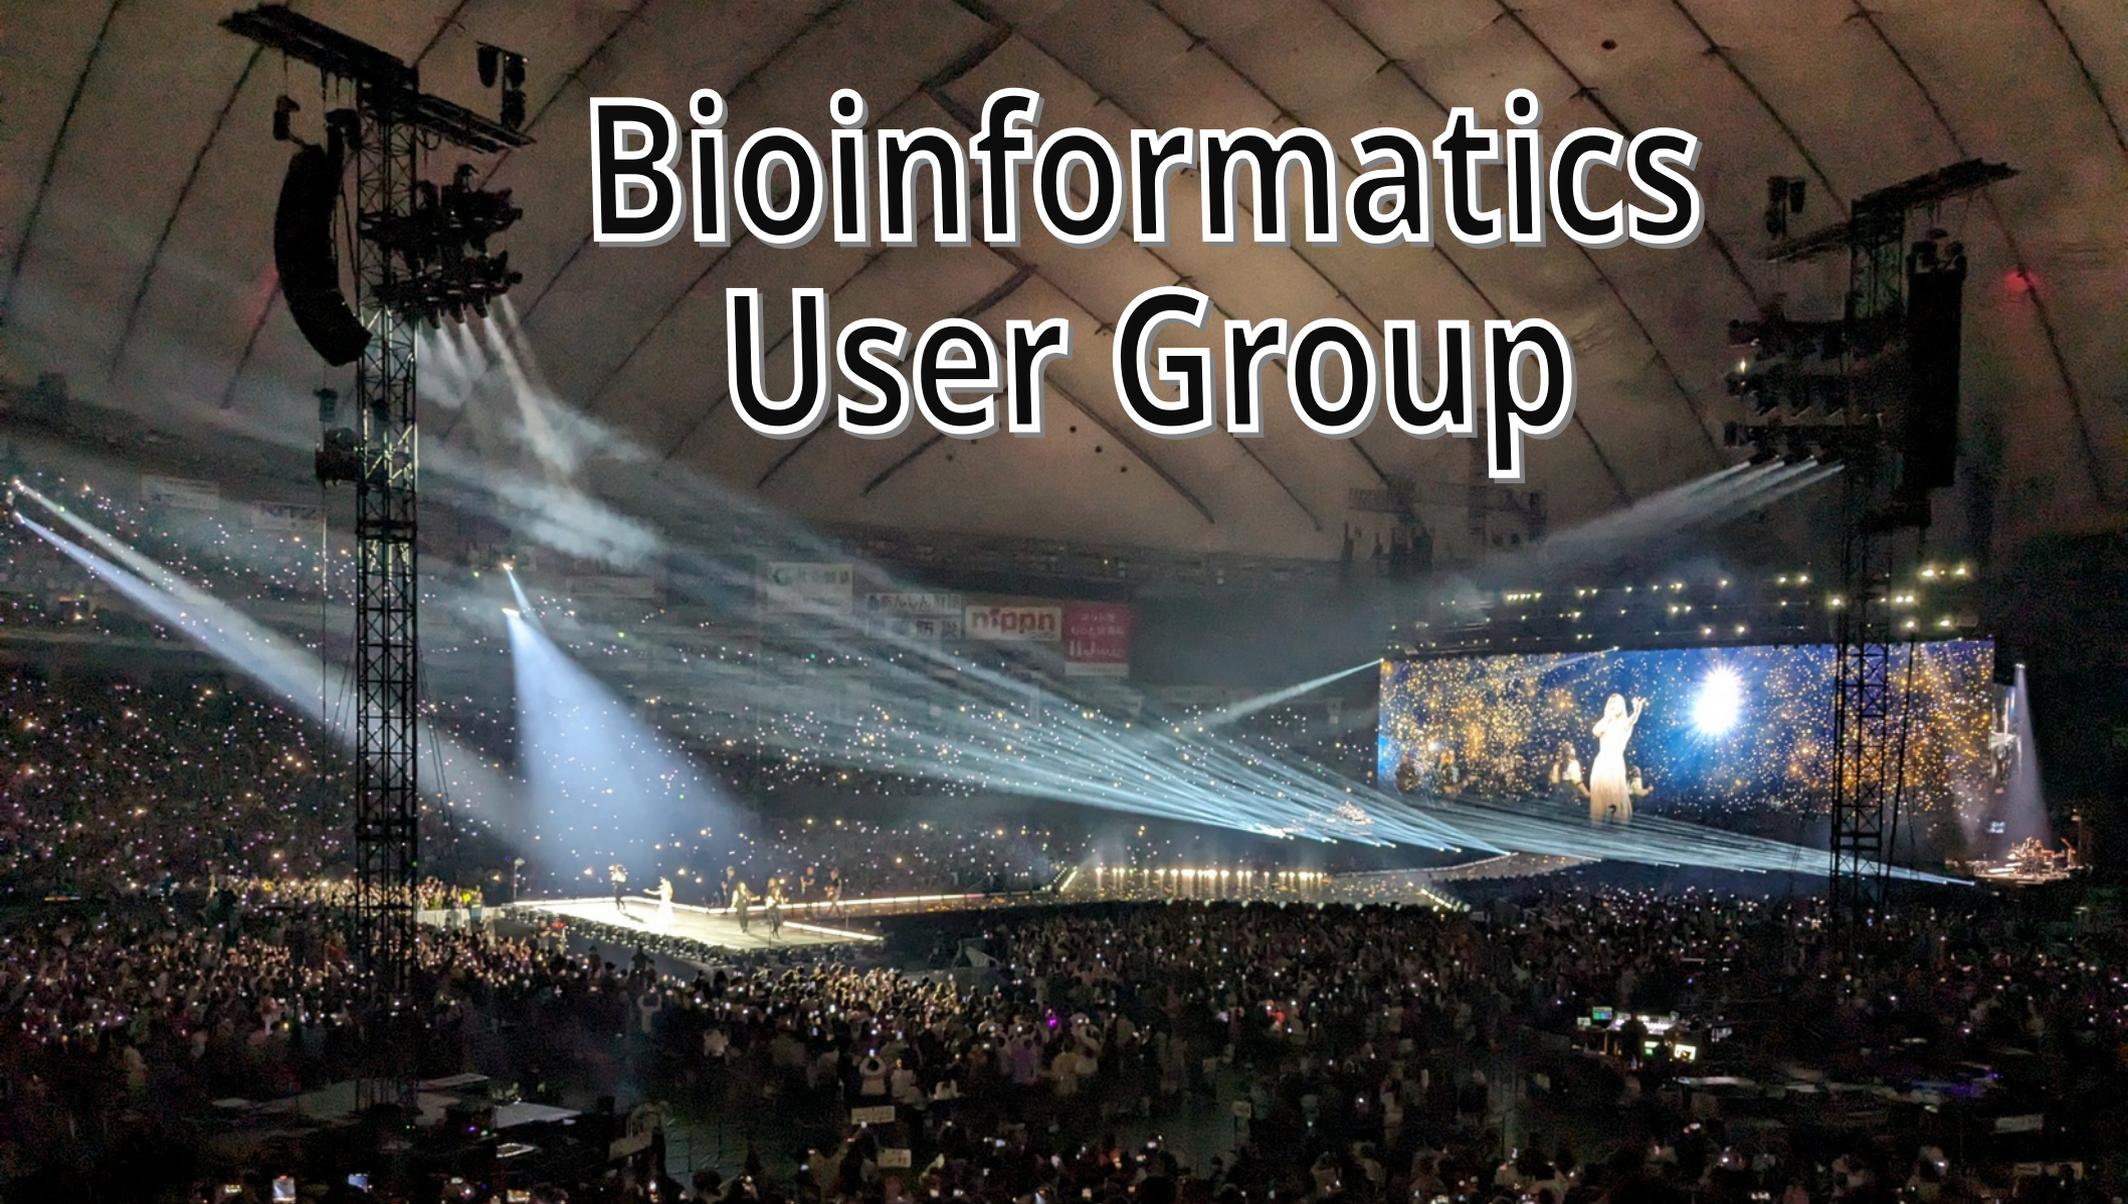


Our visualization service offers:

- Visualization pipelines
- Blender plug-ins for interactive visualization and analysis
- Research project illustrations and animations
- Posters, magazine covers etc.

* CG animations and the like can take several months. Be early!

Bioinformatics User Group



Your Questions!

By the way: If you're

- using Visual Studio Code;
- logging in remotely to Deigo:

please make sure you're not leaving
“node” processes on the login nodes!