HPC User Group Meeting

December 3rd, 2021

Jan Moren, SCDA

OIST

OKINAWA INSTITUTE OF SCIENCE AND TECHNOLOGY GRADUATE UNIVERSITY
沖縄科学技術大学院大学

Today:

SCDA
User Survey
Results

System Maintenance

- New User
  Job Limitation
- Horizon VDI is
  going Away

Bioinfo
User Group
Update!

New and Updated Software
New Shared Storage!

SCDA User Survey

n=121

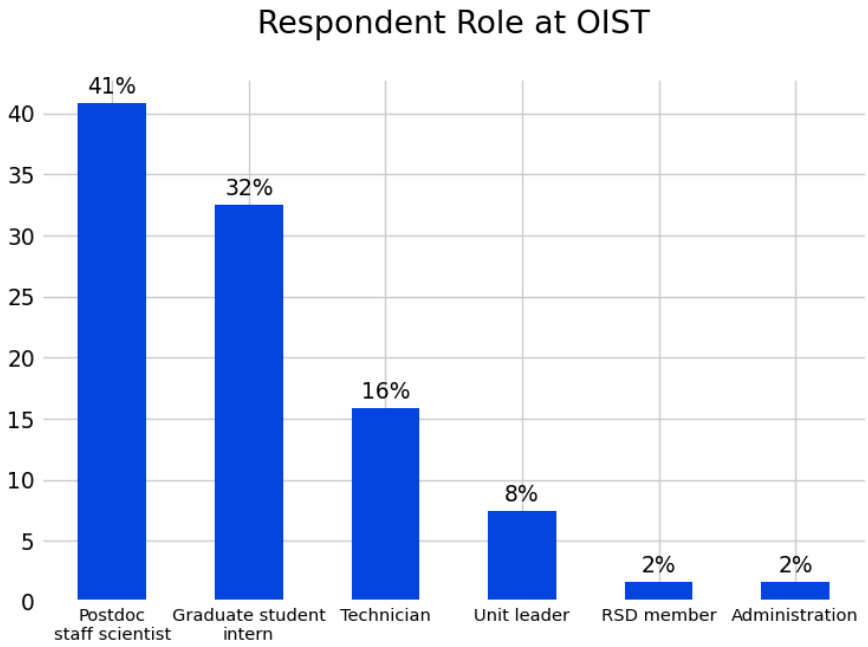

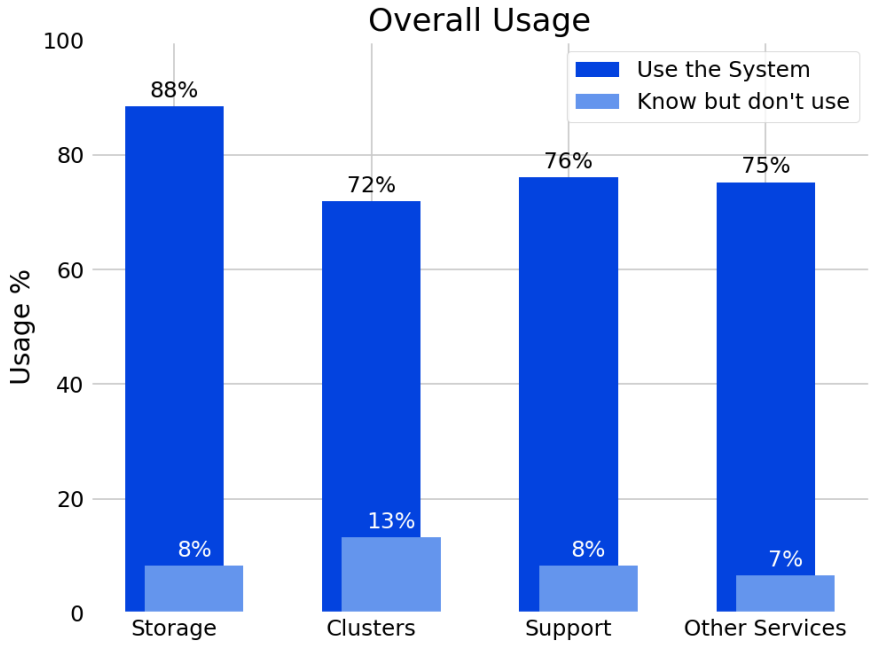

- Researchers overrepresented
  (admin and RSD uses our software)

- Cluster users overrepresented
  (most OIST researchers use storage)

Storage

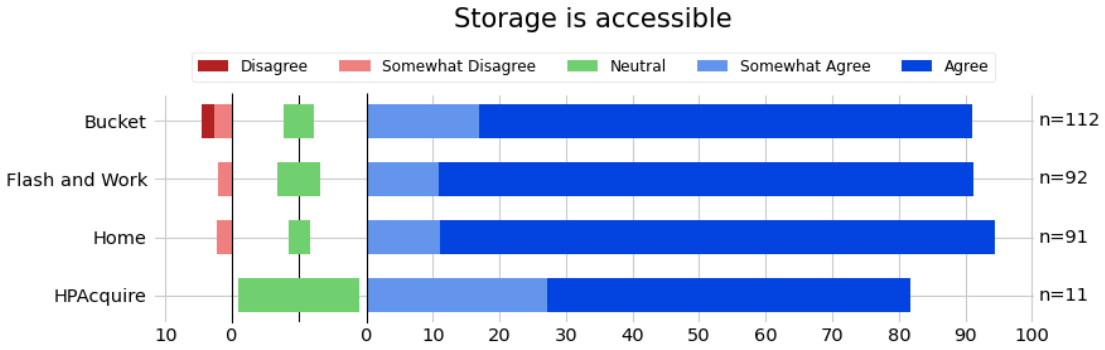

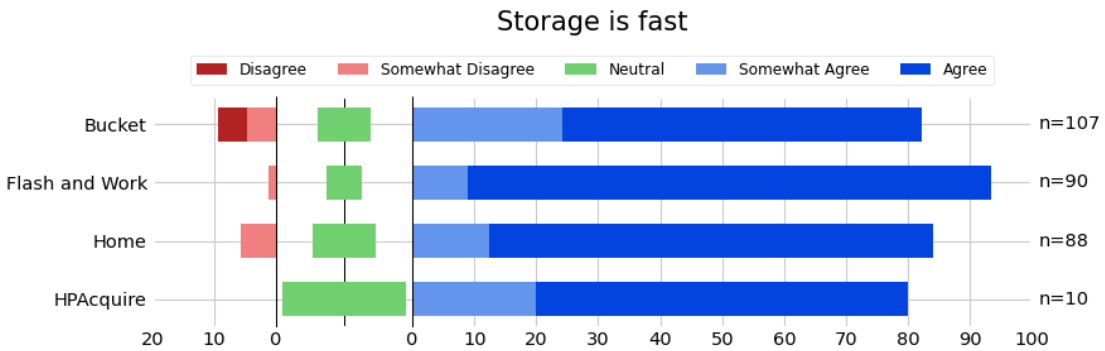

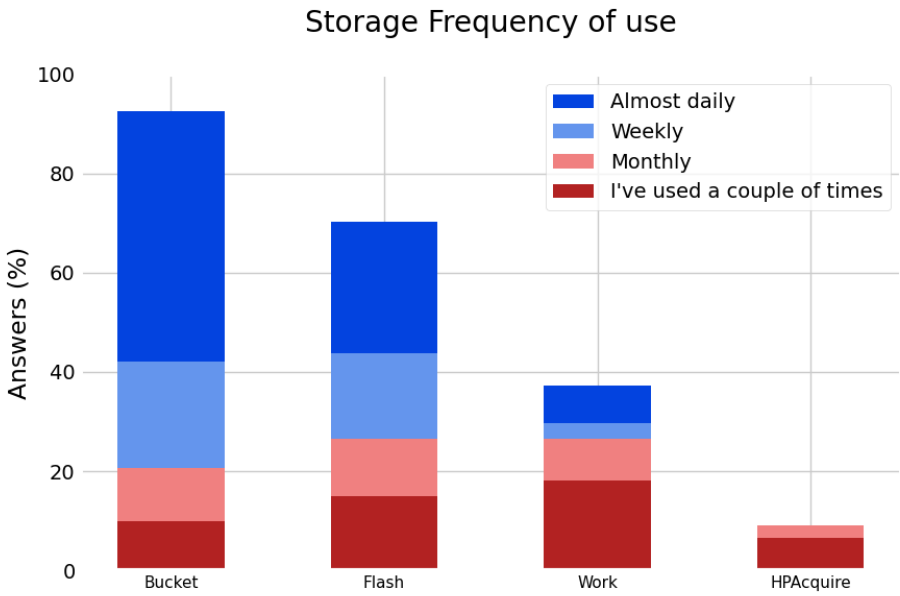

- mostly use Bucket and Flash
- fairly happy with speed and access to storage

OIST

Storage

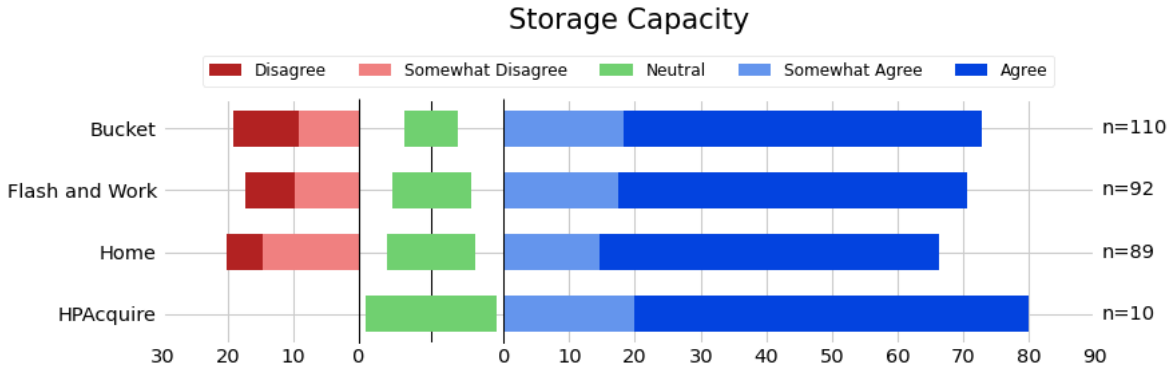

More space on bucket would be very useful.

More space on /bucket and /flash

more space on saion to store results from data without having to write to bucket everytime so that I can load trained neural network models faster for testing.

More storage in Flash for genomic data

More bucket storage would be good. Perhaps groups with dna sequencing work could get an additional amount of space.

More storage capacity

- A question of budget
  - storage costs a lot every year

  → ask OIST to give us the funding
- You can get Bucket allocation increase when necessary
- Flash and Work are fixed by design
  - never store anything here when not running a job!
- We look at other designs for the next cluster

Storage

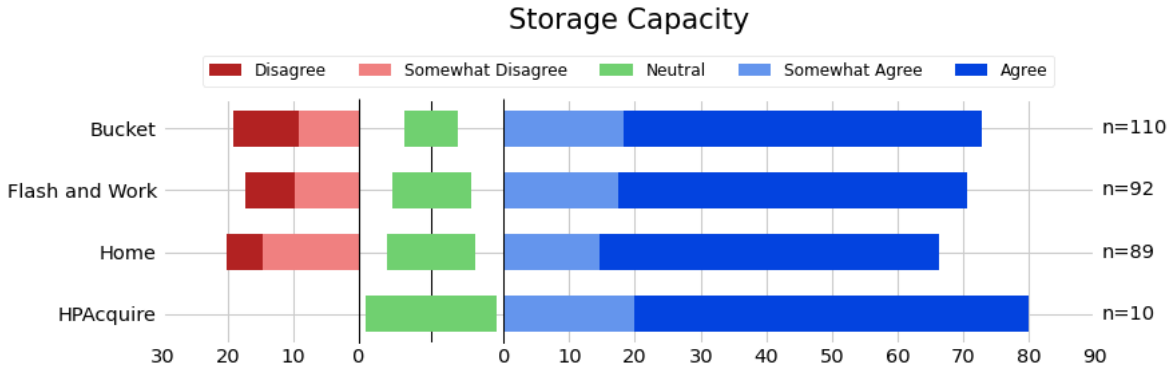

A dedicated method for transferring large files between labs might be useful.

"Big data" storage for sharing online with collaborators or the general public.

Storage with more user access control for saving large quantities of sensitive data (e.g., video recordings).

1) large volume storage option (up to hundreds of GB) that can be accessed by any selected OIST member 2) large volume storage option (up to hundreds of GB) that can be accessed by collaborators outside of OIST

More storage capacity

External data sharing and collaboration

- External access to internal storage is technically and legally difficult
- Dropbox
- Filesender
- Your collaborators' resources

Internal shared, permission-managed storage

- As it happens, we'll bring that up later today ...

Cluster

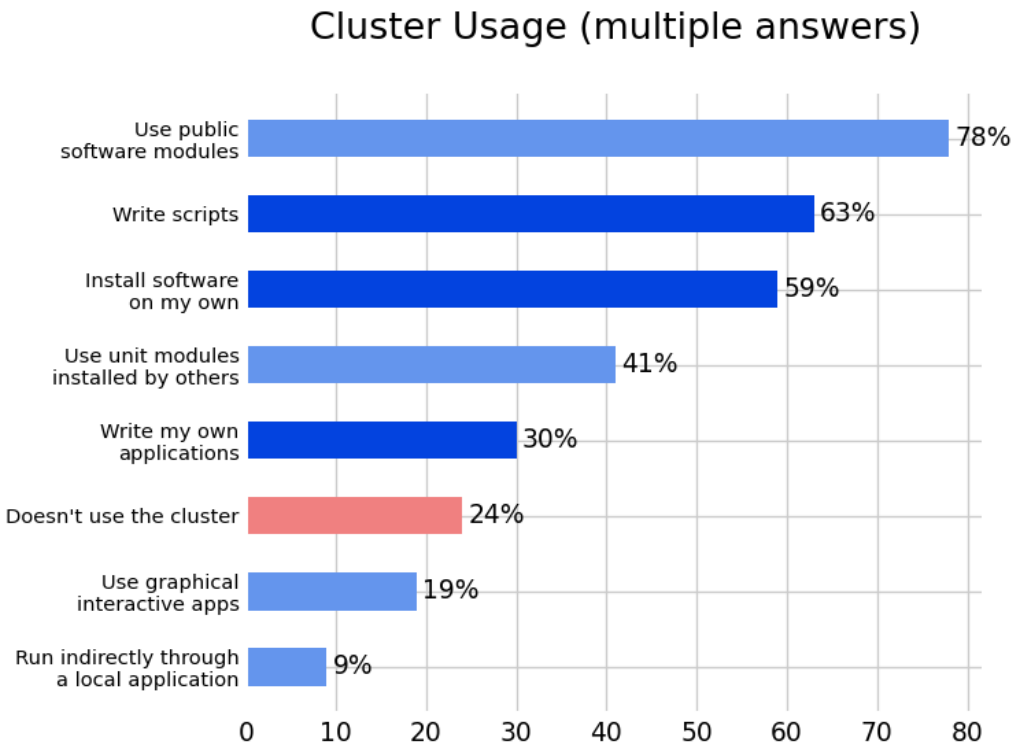

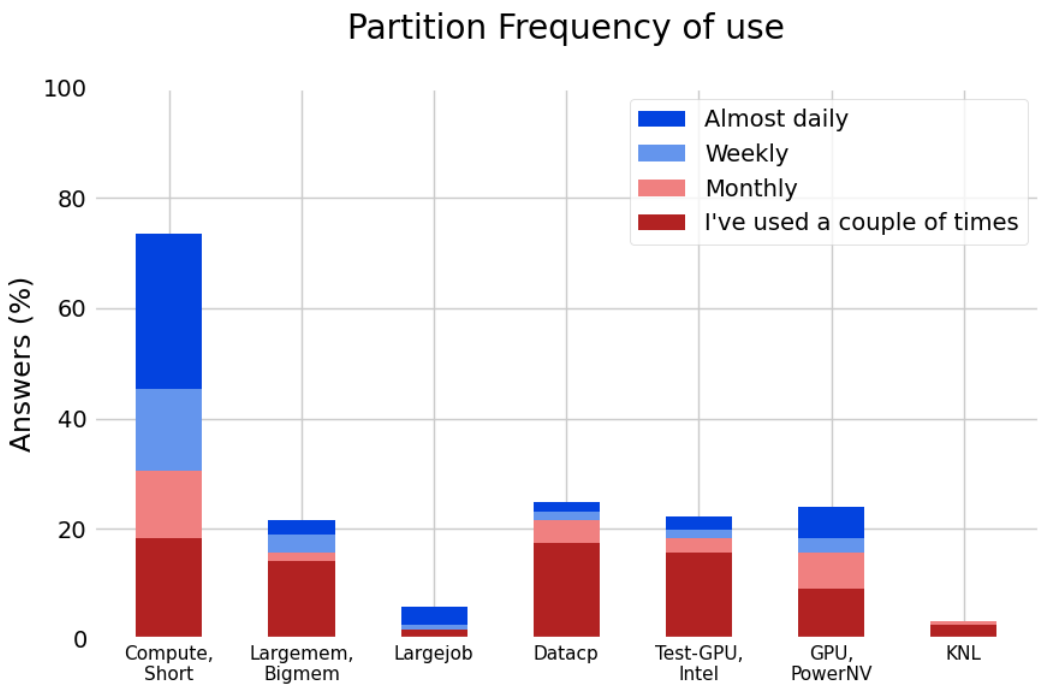

- Most people use compute
- most use prebuilt, preinstalled software

Cluster

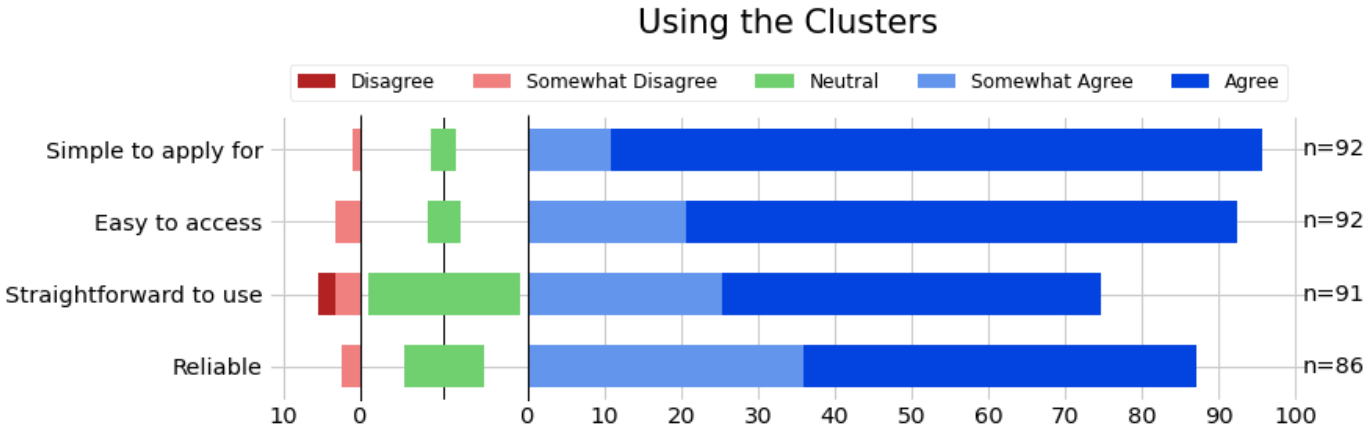

- Mostly happy
- Of course, Deigo is still fairly new...

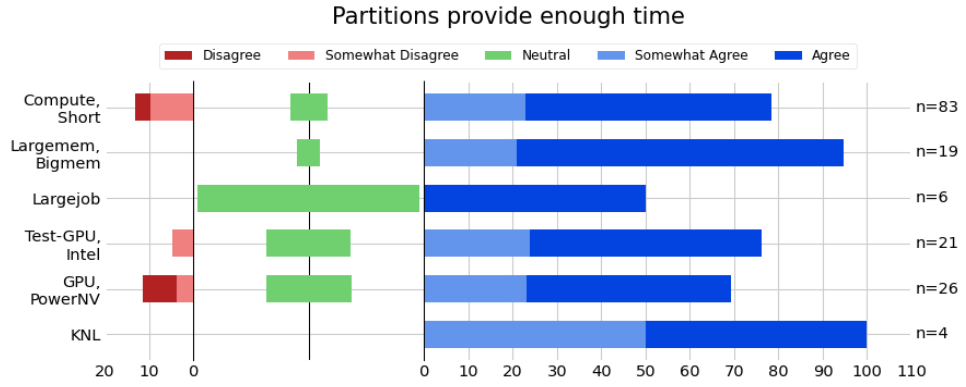

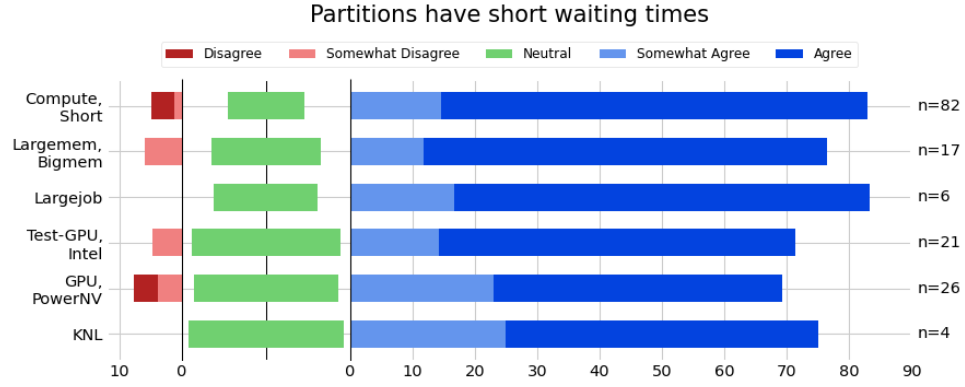

Cluster

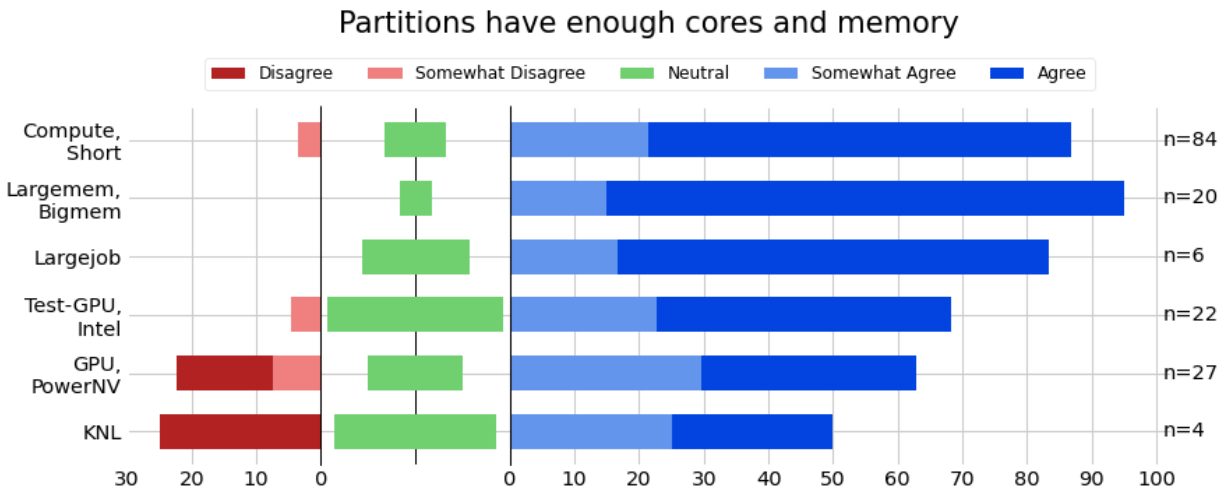

I'd love it if Deigo could have GPU nodes.

For some jobs, 20 days limit is marginal [...]

It would be nice if Intel CPUs are available for the Compute partition.

I would like to have more GPUs on Saion.

Is there any option to use V100 instead of P100?

- No Deigo GPU nodes
  - cluster federation possible in the future
- 20 days is a long time
  - no guarantee nodes stay up that long, we can't wait a month to shut down for maintenance
  - use "largemem"
- Intel CPU on compute might get confusing
  - come to Open Hours – we can figure something out
- You *can* pick a V100 GPU:

```
--gres=gpu:V100:1
```

OIST

Services and Support

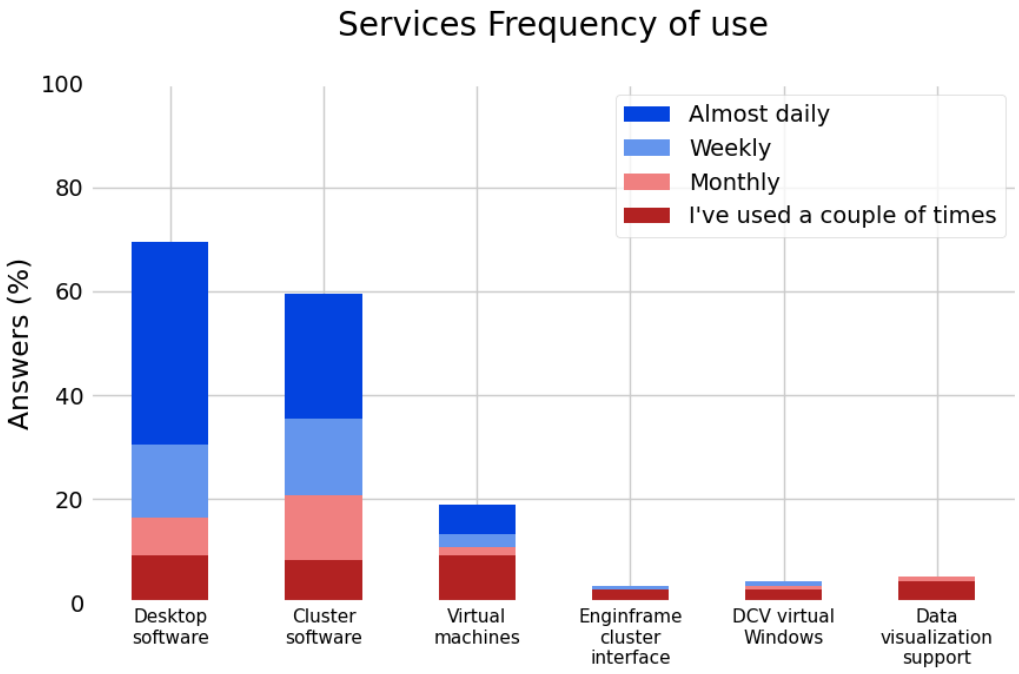

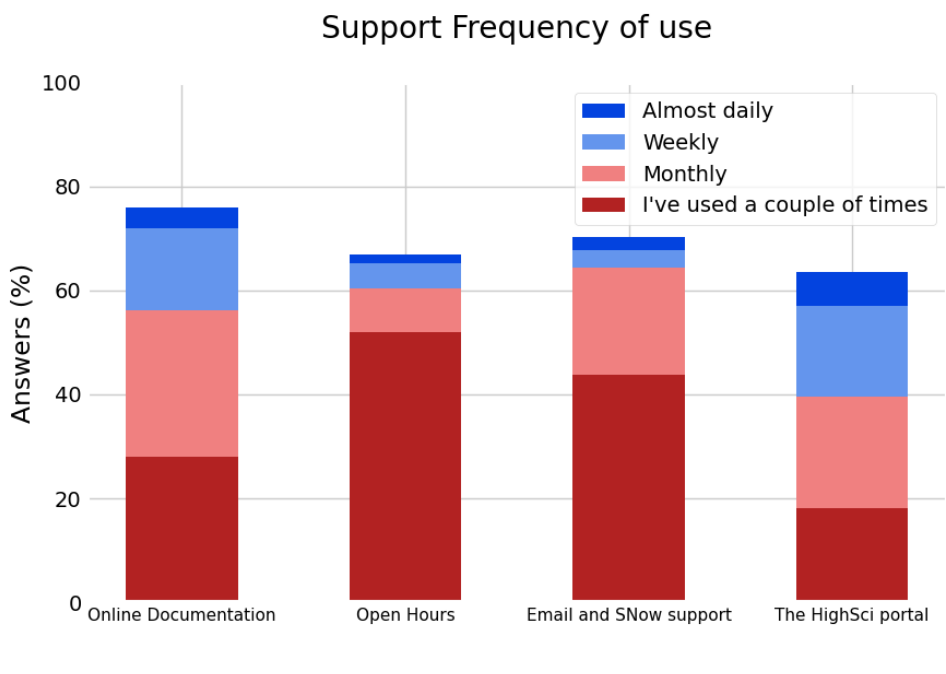

SCDA Services

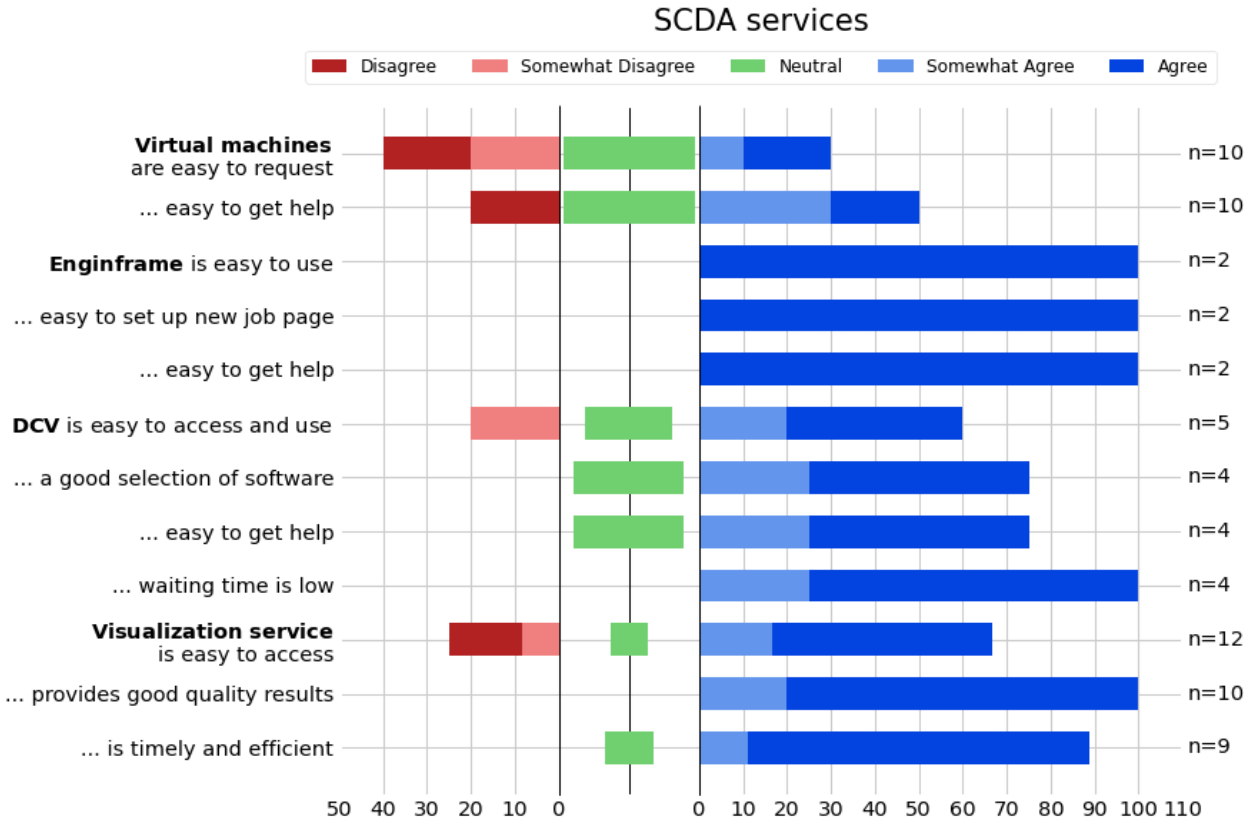

I wish it was possible to get an OIST license for software our Unit uses a lot [...] Available software [...] is unusable and so our Unit has to spend our own budget on the tools we need.

- OIST buys software used by multiple units. Single unit users must buy for themselves

Softwares needs a regular updates

- please let us know if you need updates

More visibility of services such as Data visualization support

I didn't know most of these services existed. Please give us more details about these services and maybe include some in the training and workshops

SCDA Services

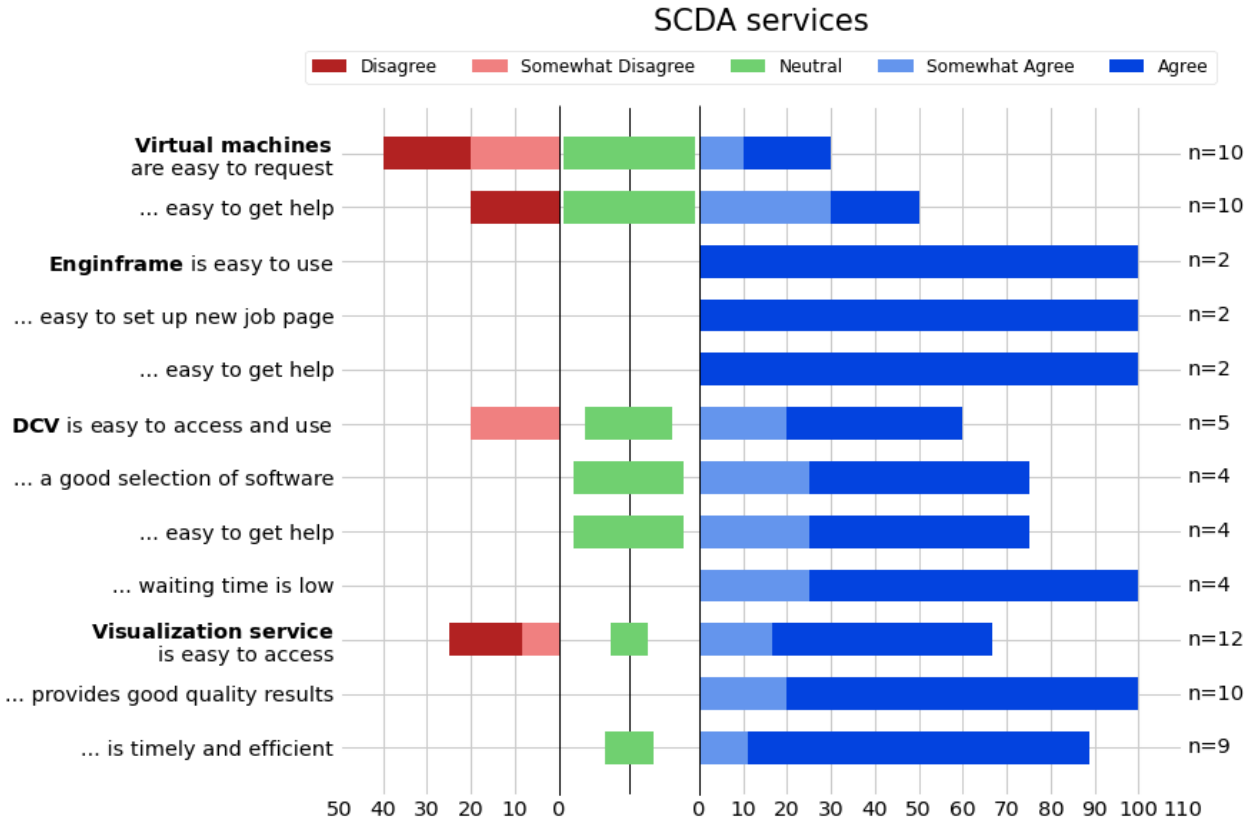

- Virtual machines
  - ongoing set of projects
  - will be improved
- Enginframe
  - run cluster software through web interface, Linux desktop on cluster
- DCV
  - virtual windows desktop for specific software
- Visualization
  - Get help on publication and presentation visualizations.

Support Resources

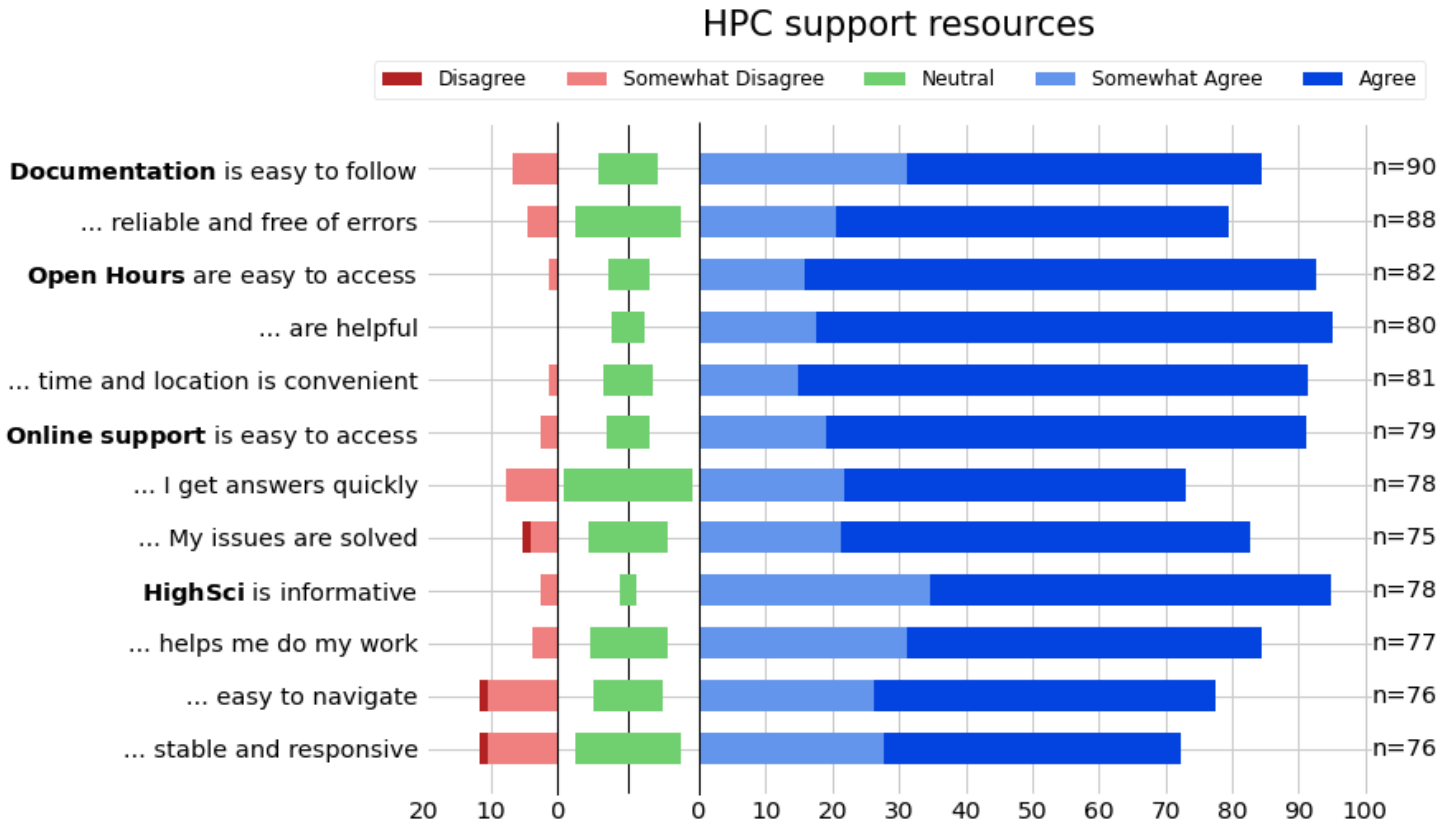

I would like a good search engine for the online documentation.

A lot of the documentation is out of date. Also, I have problems loading some of the figures on the documentation pages.

HighSci doesn't seem accessible from some web browsers, such as Brave web browser.

I wish HighSci would show me how much of my unit's bucket I am occupying (and maybe other details about bucket).

Please let us know about issues!

Training

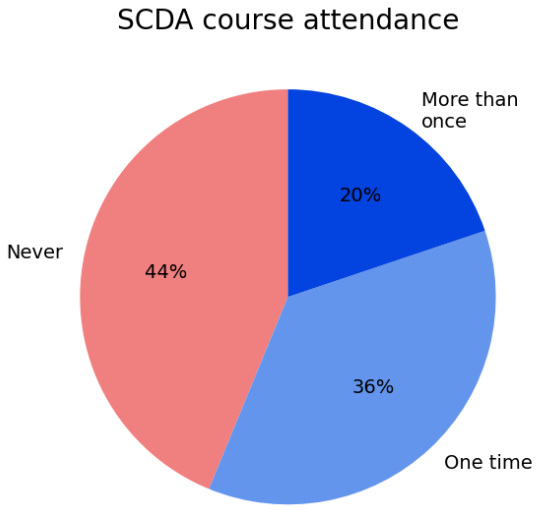

- Mostly happy about the training
- time to create new courses is limited...

Wider range of programming and data analysis courses accessible on a more regular basis.

Field-specific courses?

A course / training on practicing running jobs in parallel several times in Python and Julia, [...]

It would be great if these trainings could be early morning (8-10), or early evening (5-7).

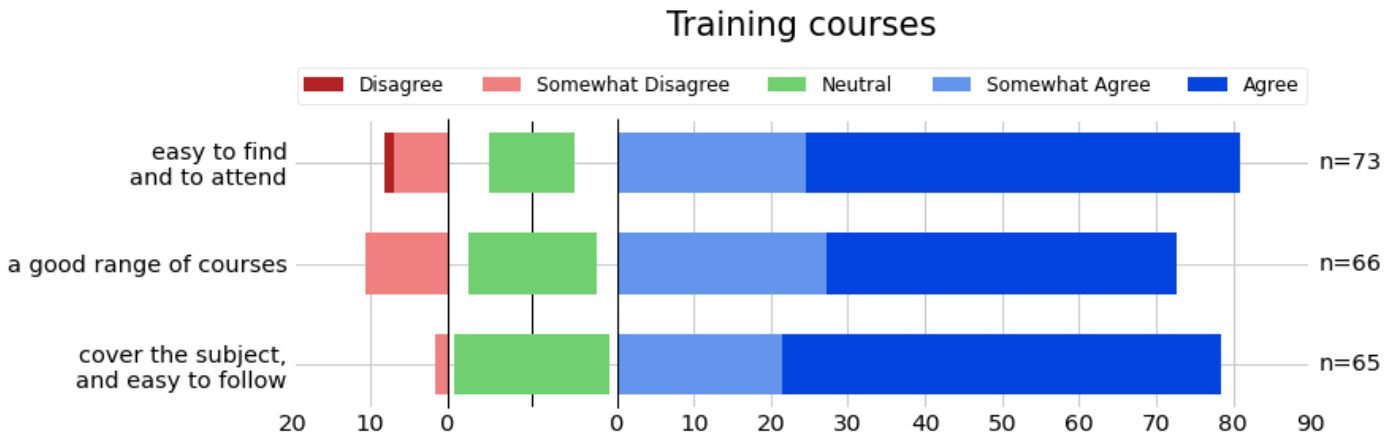

SCDA

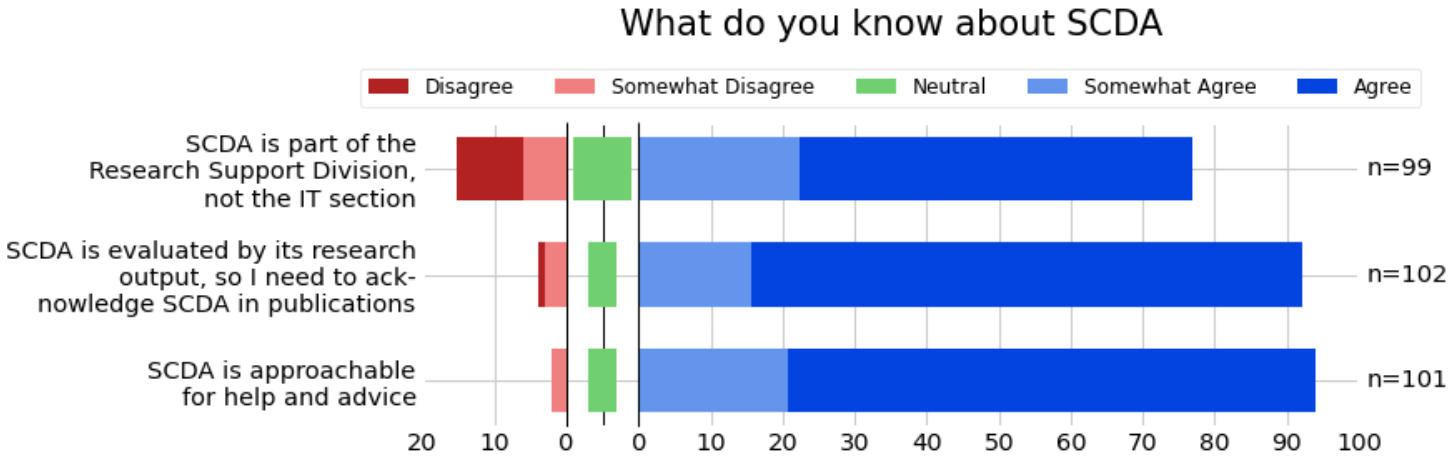

It's great but it would be helpful to have a wider range of support for our daily technical issues, such as

1) support for programming experiments,
2) support for engineering issues (Arduino),
3) support for specific data analysis (physiological data) and
4) support for statistical analysis.

I think the community would benefit from an increased number of staff employed for data visualization.

SCDA is highly valuable for our research and it should probably expand in size (number of staff) as OIST is growing.

- Anything is possible — but needs a budget
- OIST responds to faculty demands
  → Ask your PI to ask OIST to give it priority.

SCDA

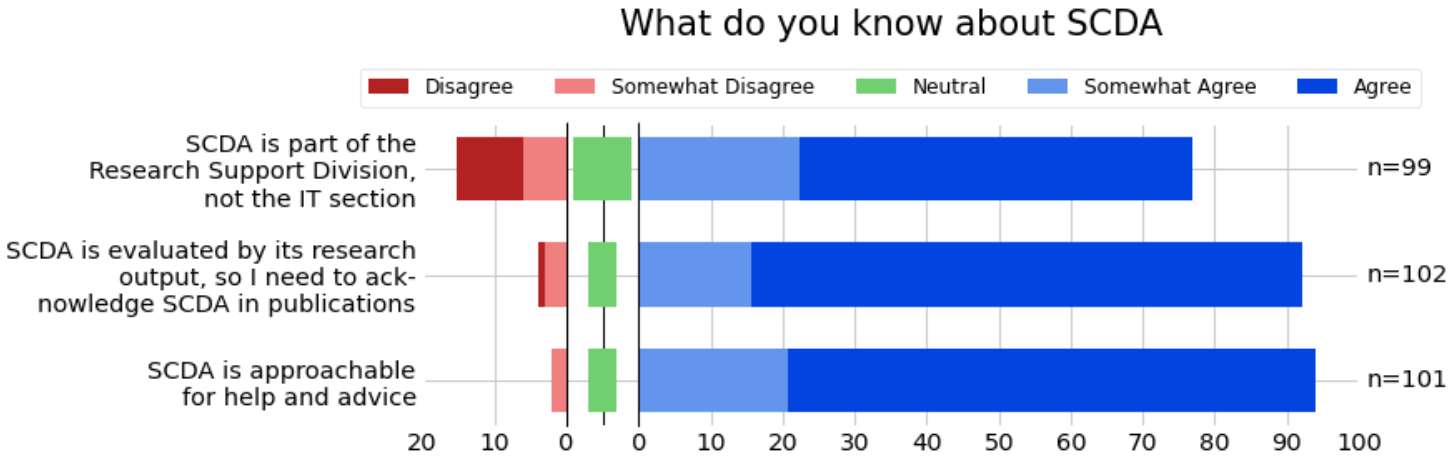

I want to thank you in the acknowledgment section of my paper, but my recent paper got rejected! So, yeah, I want to thank you in the next submission.

Nice knowledgeable people who are always ready to help

Excellent Work overall. Very useful section with sometimes onerous rules that can make it hard to do new things.

- Lots of positive feedback too!
- A respondent that never used any SCDA services:

  Don't know much about you, but your purpose looks great!

Maintenance

Deigo Maintenance:

When: Mid-late January

Duration: ~3-4 days

Major tasks:

- OS update (to CentOS 8.5)
- IB switches firmware update
- Also various deferred maintenance tasks

Saion Maintenance:

When: After Deigo

Duration: ~3-4 days

Major tasks:

- OS update (to CentOS 7.9)
- module system update to same as Deigo
- Fix X11-related issues on the nodes

Job Limitations

The Slurm scheduler is heavily loaded

- more jobs → more work for Slurm
- **Number of users** are increasing
- **Jobs per user** are increasing
- More users × more jobs per user
  = *lots* of jobs
  = too much for Slurm to handle

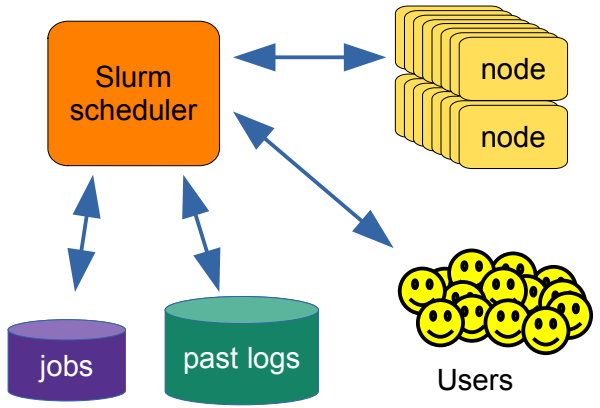

- schedule jobs
- manage partitions
- update logs
- respond to users
- track hardware

…

Job Limitations

Our existing rule is "don't submit thousands of jobs at once"

- Some users have submitted tens of thousands or more – **often by accident**

So, from now on we add a hard limit:

```
submissions <= (core allocation)+16
```

- enough to use all cores + a few extra jobs
- Set on all new users, and will gradually roll out to everybody

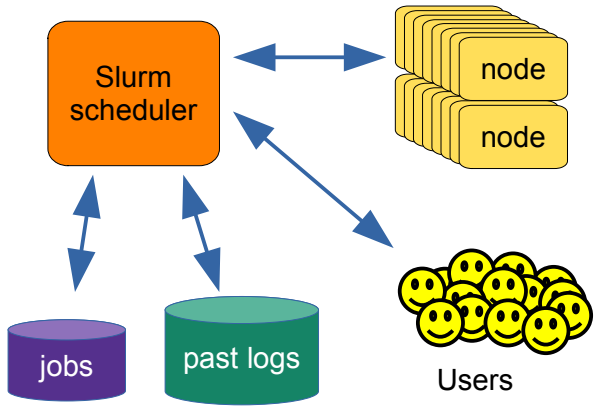

- schedule jobs
- manage partitions
- update logs
- respond to users
- track hardware

…

Job Limitations

A new hard limit:

```
submissions <= (core allocation)+16
```

- For array jobs, use "%" to limit submissions:

```
--array=0-9999%1000
```

- The extra 16 lets you use chain jobs and the like while still make use of all cores

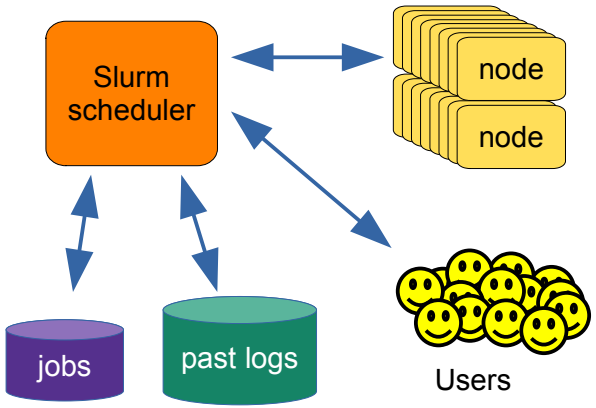

- schedule jobs
- manage partitions
- update logs
- respond to users
- track hardware

…

Horizon VDI

Virtual Windows desktops
with accellerated GPUs

- Set up by IT as a POC (Proof of concept) in 2019
- Meant to be replaced by a production system, but it never happened due to changes within IT.
- IT dropped the service and returned the hardware back to us last month.
- We have no budget, service plan or infrastructure, and the licenses expire at the end of the year

→ Horizon VDI will go away in January

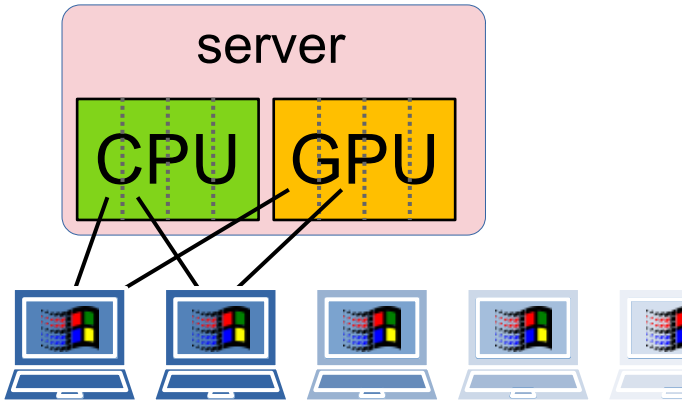

Horizon VDI

Virtual Windows desktops
with accellerated GPUs

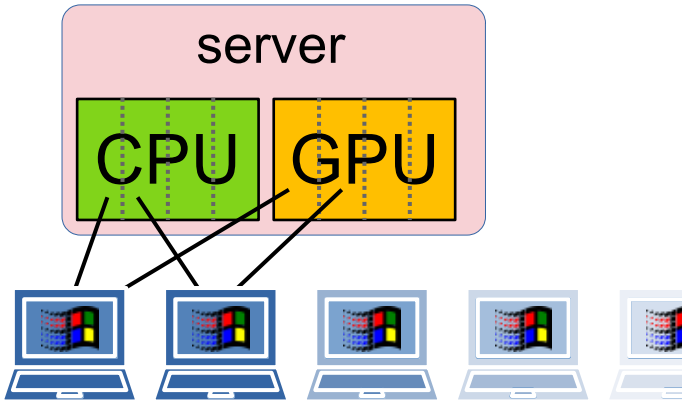

- Horizon VDI will go away in January
  - System will be up for a while, but there is no maintenance or any way to fix things.
- We will expand DCV with more machines – if you used VDI for GPU-accellerated Windows software you can move to DCV
- If you used VDI just to access the OIST internal network, please use the OIST VPN

Bioinformatics User Group!

New Software

- **MestreNova** — spectroscopy software suite
- **Altium Designer** — PCB design software
  - User support by the Electronics Team of the RSD engineering section:

    electronics@oist.jp

- **Gaussian** is updated from version 09 to 16
  - Fixes a lot of issues on the cluster
  - adds NBO v.7 as a module
- **Aivia** — Image 3d-reconstruction
- **AlphaFold** — Protein folding structure prediction
  - Managed by the Bioinformatics User Group
- **MaterialsStudio** 2021 has issues with Deigo and we're working with them to resolve the issues.
- Meanwhile you can now run smaller jobs on Saion. See:

https://groups.oist.jp/scs/software-notes#materialsstudio

New Storage System:

comspace

It's for:

- Sharing data between units at OIST
- Restricted data shares for a subgroup of lab members

It features:

- Any number of distinct shares
- Fine-grained permissions
- Managed by the unit, under control of the unit leader
- 5TB per unit
- GUI-based — mount the shares as remote drives

It's not for:

- Not for external access (that will need a separate solution)
- Not set up to move data on the command line (you might use "`smbclient`" to do it)

It's available:

- Already in testing by a few users
- Gradually rolls out to all units over time

Next training Session:

Introduction to HPC and Scientific Computing

Date: Tuesday, January 25th
Thursday, January 27th

Time: 13:00 — 15:00

Place: C700 (this very room)

A 2+2 hour hands-on introduction to HPC and to using our clusters.

No experience needed! Everybody is welcome.