

the future of research computing and data @ MIT

Dr. James Cuff

Executive Director

20th June, 2023



orcd@mit:~\$ background

University of Oxford: D. Phil. Molecular Biophysics
Neural networks for protein structure prediction

Wellcome Trust Sanger Institute: Group Leader
Built the Informatics Systems Group, (Human Genome)

Broad Institute of MIT & Harvard: Group Leader
Built the Applied Production Systems Group

Harvard: Assistant Dean & Distinguished Engineer
Built the 100,000+ CPU cluster & MGHPCC from scratch

Strategic Scientific Advisor & Independent Consultant
5+ years learning "cheat-codes" in .com

THE RESEARCH MACHINES 380Z COMPUTER SYSTEM



THE RESEARCH MACHINES 380Z A UNIQUE TOOL FOR RESEARCH AND EDUCATION

Microcomputers are extremely good value. The outright purchase price of a 380Z installation with dual mini floppy disk drives, digital I/O and a real-time clock, is about the same as the annual maintenance cost of a typical laboratory minicomputer. It is worth thinking about!

The RESEARCH MACHINES 380Z is an excellent microcomputer for on-line data logging and control. In university departments in general, it is also a very attractive alternative to a central mainframe. Having your own 380Z means an end to fighting the central operating system, immediate feedback of program bugs, no more queuing and a virtually unlimited computing budget. You can program in interactive BASIC or, using our unique Text Editor, run very large programs with a 380Z FORTRAN Compiler. If you already have a mini-computer, you can use your 380Z with a floppy disk system for data capture.

What about Schools and Colleges? You can purchase a 380Z for your Computer Science or Computer Studies department at about the same cost as a terminal. A 380Z has a performance equal to many minicomputers and is ideal for teaching BASIC and Casil. For A Level machine language instruction, the 380Z has the best software front panel of any computer. This enables a teacher to single-step through programs and observe the effects on registers and memory, using a single keystroke.

WHAT OTHER FEATURES SET THE 380Z APART?

The 380Z with its professional keyboard is a robust, hardwearing piece of equipment that will endure continual handling for years. It has an integral VDU interface — you only have to plug a black and white television into the system in order to provide a display

unit — you do not need to buy a separate terminal. The integral VDU interface gives you upper and lower case characters and low resolution graphics. Text and graphics can be mixed *anywhere* on the screen. The 380Z has an integral cassette interface, software and hardware, which uses *named* cassette files for both program and data storage. This means that it is easy to store more than one program per cassette.

Owners of a 380Z microcomputer can upgrade their system to include floppy (standard or mini) disk storage and take full advantage of a unique occurrence in the history of computing — the CP/M™* industry standard disk operating system. The 380Z uses an 8080 family microprocessor — the Z80 — and this has enabled us to use CP/M. This means that the 380Z user has access to a growing body of CP/M based software, supplied from many independent sources.

380Z mini floppy disk systems are available with the drives mounted in the computer case itself, presenting a compact and tidy installation. The FDS-2 standard floppy disk system uses double-sided disk drives, providing 1 Megabyte of on-line storage.

*Trademark, Digital Research.

Versions of BASIC are available with the 380Z which automatically provide controlled cassette data files, allow programs to be loaded from paper tape, mark sense card readers or from a mainframe. A disk BASIC is also available with serial and random access to disk files. Most BASICs are available in erasable ROM which will allow for periodic updating.

If you already have a teletype, the 380Z can use this for hard copy or for paper tape input. Alternatively, you can purchase a low cost 380Z compatible printer for under £300, or choose from a range of higher performance printers.

380Z/32K complete with SINGLE MINI FLOPPY DISK SYSTEM MDS-1
£1787.00

380Z/16K System with Keyboard
£965.00

RESEARCH MACHINES Computer Systems are distributed by RESEARCH MACHINES LTD., P.O. Box 75, Chapel Street, Oxford. Telephone: OXFORD (0865) 49792. Please send for the 380Z Information Leaflet. Prices do not include Carriage or VAT @ 8%.

orcd@mit:~# y2k_genomes.pl



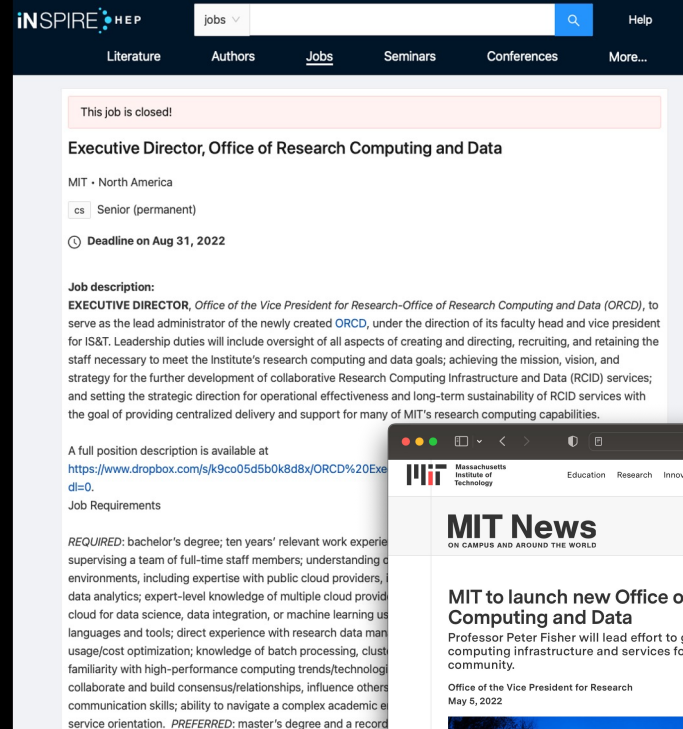
orcd@mit:~\$ jobs

Started 26th Sept 2022

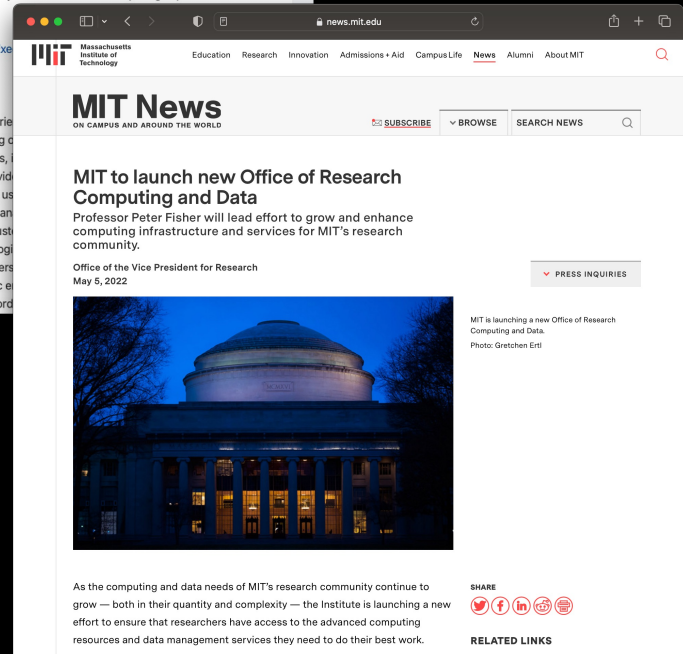
10+ years of RCP set the foundation, epic kudos to Dr. Chris Hill for this!

ORCD now reports to the VPR and Provost

Returning to MIT & MGHPCC has been awesome!



The screenshot shows a job listing on the INSPIRE website. At the top, there are navigation tabs for Literature, Authors, Jobs, Seminars, Conferences, and More... The job title is "Executive Director, Office of Research Computing and Data". It is located at MIT - North America and is a Senior (permanent) position with a deadline of August 31, 2022. The job description states that the Executive Director will lead the newly created ORCD, reporting to the Vice President for Research. The role involves overseeing research computing and data goals, managing staff, and setting strategic direction for research computing infrastructure and data services. A link to a full position description is provided.



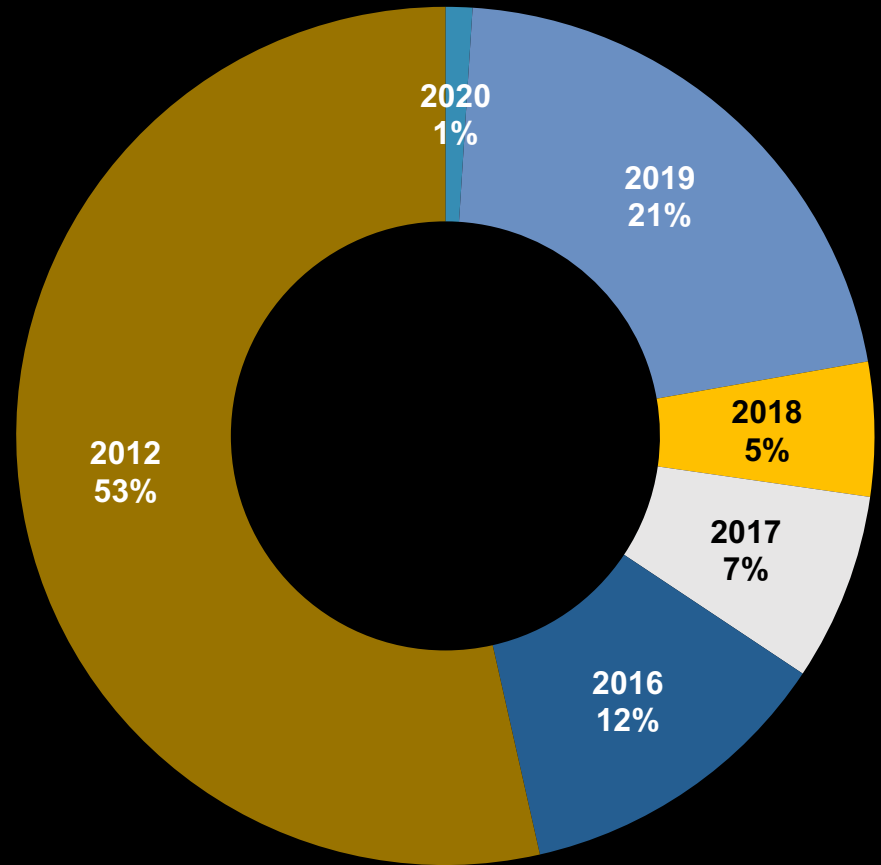
The screenshot shows an MIT News article. The headline is "MIT to launch new Office of Research Computing and Data". The article features a photograph of the MIT dome at night. The text indicates that Professor Peter Fisher will lead the effort to grow and enhance computing infrastructure and services for the MIT research community. The article is dated May 5, 2022. At the bottom, there are social media sharing options and a "RELATED LINKS" section.

Engaging1 – an example shared ORCD cluster

As of Q4 2022, there now exist over 1,463 machines in the computer cluster called “Engaging-1”

Of those 1,463 machines, 1,097 are currently available for processing. Of those 1,097 available, 797 are between 5 and 11 years old, 300 are less than 5 years old, and 244 remain under some form of warranty and are available for repair.

Of the total machines available, those considered viable for modern computing represent less than 16% of the total available fleet, after taking into consideration damaged, out of warranty, legacy and currently unavailable systems.





CPU	Samsung 32-bit RISC ARM1176JZ(F)-S v1.0 - 620 MHz Underclocked to 412 MHz
GPU	PowerVR MBX Lite 3D GPU
Memory	128 MB eDRAM
Storage	4, 8, or 16 GB flash memory
Battery	3.7 V 1400 mAh Lithium-ion battery
Display	90 mm (3+ ¹ / ₂ in) screen (diagonally) 480x320 pixel resolution at 163 ppi 3:2 aspect ratio 18-bit (262,144 colors) LCD
Sound	Single loudspeaker
Rear camera	2.0 MP with geotagging (not GPS-based)
Connectivity	Quad-band GSM/GPRS/EDGE Wi-Fi (802.11 b/g) Bluetooth 2.0 USB 2.0 / dock connector

June 2007

Did we mention technical debt?



Credit: Heather Sardis of MIT Libraries & MidJourney

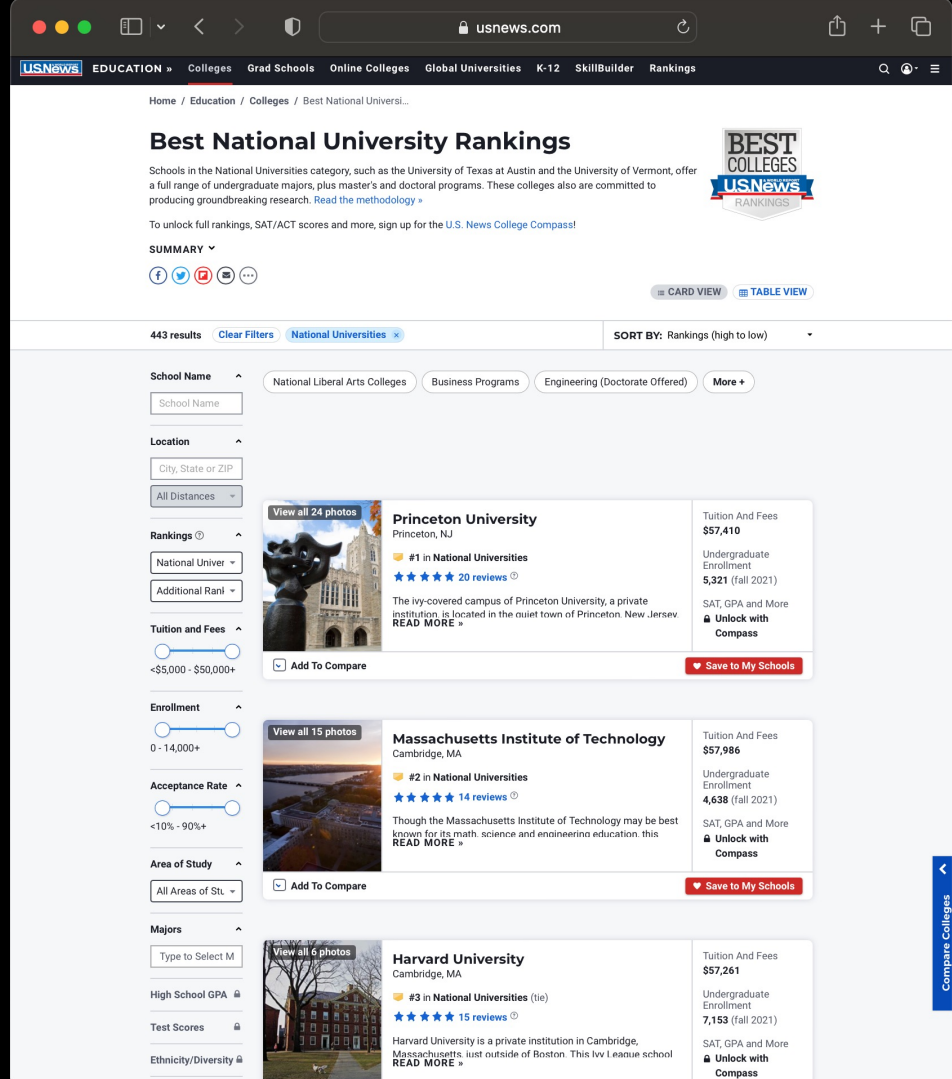
Did we?

*There may, or
may not be
some of these
iPhone 1 vintage
servers still
in production...*



It's clearly not all bad news.

How do we do this?



The screenshot shows the US News website's "Best National University Rankings" page. At the top, there's a navigation bar with "USNews" and "EDUCATION" menus. The main heading is "Best National University Rankings" with a sub-heading explaining that schools in this category offer a full range of undergraduate, master's, and doctoral programs. A "BEST COLLEGES USNews RANKINGS" badge is visible in the top right. Below the heading, there are social media icons and view options for "CARD VIEW" and "TABLE VIEW". A filter bar shows "443 results" for "National Universities" and a "SORT BY: Rankings (high to low)" dropdown. On the left, there are several filter sections: "School Name" (with a search box), "Location" (with a search box and "All Distances" dropdown), "Rankings" (with "National Univer" and "Additional Rani" dropdowns), "Tuition and Fees" (with a range slider from $-\\$5,000$ to $-\\$50,000+$), "Enrollment" (with a range slider from 0 to 14,000+), "Acceptance Rate" (with a range slider from <math><10\%</math> to $90\%+$), "Area of Study" (with "All Areas of Stu" dropdown), and "Majors" (with "Type to Select M" search box). The main content area displays three university cards: Princeton University (#1 in National Universities, 20 reviews, tuition of \$57,410), Massachusetts Institute of Technology (#2 in National Universities, 14 reviews, tuition of \$57,986), and Harvard University (#3 in National Universities, 15 reviews, tuition of \$57,261). Each card includes a photo, a brief description, and a "Save to My Schools" button. A vertical "Compare Colleges" button is on the far right.

orcd@mit:~# sinfo -a

The screenshot shows the MIT SuperCloud website. The header includes the MIT logo and the text 'MIT SuperCloud'. A search bar is visible. The main content area is titled 'Requesting an Account' and features a section for 'Account Request Process'. The process is detailed in a numbered list with sub-points. A sidebar on the left contains navigation links such as 'Home', 'Requesting An Account', 'Getting Started', 'How To Use', 'Jupyter Portal', 'Online Courses', 'Publications And Software', 'Acknowledging The MIT Supercloud', 'Contact', and 'FAQs'. At the bottom of the sidebar, it says 'Login using Touchstone'.

SuperCloud

The screenshot shows the MIT Research Computing Project website. The header includes the MIT logo and the text 'MIT Research Computing Project'. Navigation links include 'What's available', 'Facilities', 'Projects', 'Questions', 'About', and 'Quick Links'. The main content area is titled 'Request Account' and contains a form for requesting an account. The form is divided into two sections: 'About You' and 'Your Account'. The 'About You' section includes fields for 'Name and Affiliation(s)', 'Sponsor Name and Affiliation(s)', 'Email', 'Phone', and 'Group'. The 'Your Account' section includes a field for 'Username' and a 'Description' field.

C3DDB

The screenshot shows the Satori Portal website. The header includes the text 'Satori Portal'. Below the header, there is a message: 'To access the OnDemand portal, please click below to authenticate with an Identity Provider'. A prominent green button labeled 'Authenticate with Globus' is visible. The background of the page features a photograph of a modern building with large glass windows and a curved walkway in the foreground.

Satori

orcd@mit:~# !!

The OpenMind Computing Cluster

A Shared Resource For The MIT Brain Research Community.

Introduction

Getting Started

Tutorials

- An Introduction
- Applications and Containers
- Slurm Job Scheduler and Best Practices
- MATLAB for High Performance Computing

Examples

Contact

Introduction

The OpenMind computing cluster is operated by the Department of Brain and Cognitive Sciences (BCS) and the McGovern Institute. It provides the MIT brain research community with access to state-of-art computing resources. The cluster is housed at the Massachusetts Green High Performance Computing Center (MGHPCC) in Holyoke, MA, with a 10G link to the MIT campus at

OpenMind

Engaging Cluster Documentation

Search...

- 1. Logging into the cluster
 - Logging into Engaging...
 - Logging into FastX
 - Logging in via ssh
- 2. Replacing or Adding an S...
- 3. Slurm
- 4. Software
- 5. Storage
- 6. Virtual Environments
- 7. Best Practices
- 8. Frequently Asked Questions

LOGGING INTO ENGAGING OPEN ONDEMAND

1. Start by going to engaging-ood.mit.edu in your web browser.
2. Select one of the Single Sign On authentication options.
3. Select the institution your Engaging account email is associated with.
 - The email associated with your Engaging account is your MIT Kerberos account.
4. Follow the authentication procedure required by the selected institution and upon success you will be redirected to the Engaging Open OnDemand Portal.

If you have issues with this process, please contact engaging-admin@techsquare.com

Example

Here are the steps illustrated if you are using the "Globus" option and are an MIT user:

Engaging

Submit Portal Account Status SSH Keys Cluster Status User Guide

Account Status

Provisioned Status:	provisioned
Username:	jcuff
User ID:	10375
Home Directory:	/home/submit/jcuff

SubMIT

orcd@mit:~# !!

tig.csail.mit.edu

MIT accessibility info

Search...

SHARED COMPUTING

General purpose shell access

login.csail.mit.edu ([more information](#)) is available for general purpose access. It should be used for computationally non-intensive tasks, such as reading email, editing text files, compilation of small source packages, etc. It should not be used for long running or computationally demanding processes.

Info

Do not run MATLAB on login.csail

OpenStack

TIG maintains an OpenStack based 'cloud' platform for all lab members.

OpenStack is a free and open software suite for providing Infrastructure as a Service (IaaS) both for private clouds like we're implementing and large public clouds. We are currently using "Nova" compute services, "Keystone" identity service, "Horizon" web dashboard, "Cinder" volume service, "Glance" image management, "Neutron" networking service, and "Heat" orchestration service. We are currently using the "Mitaka" release (aka 2016.1 released May 2016) 16.04 LTS "Xenial" with KVM as the virtualization layer. We also have an OpenStack "Swift" and AWS S3 compatible object store based on [Ceph] with 168TB available storage (505TB raw less 3x replication)

As of June 2018 we have 40 physical nodes, with 1720 cores and 11T RAM.

Detailed information on account self-signup and usage is available on the

- > Get Started and Best Practices
- > Accounts and Authentication
- > Email and Communicating
- > Operating Systems
- > Data Storage
- > Servers, Workstations, and Hardware
- ✓ **Shared Computing**
 - > OpenStack
 - > SLURM Compute Cluster
- > Network and Wireless
- > Web Services
- > Print, Copy, and Scan
- Rooms, Video Conference, & Event
- > Calendar
- > Building Operations
- Communications and News

CSAIL - TIG

bateslab.mit.edu

MIT

Massachusetts Institute of Technology

Bates Research and Engineering Center

Contact Us

Search

High Performance Research Computing Facility

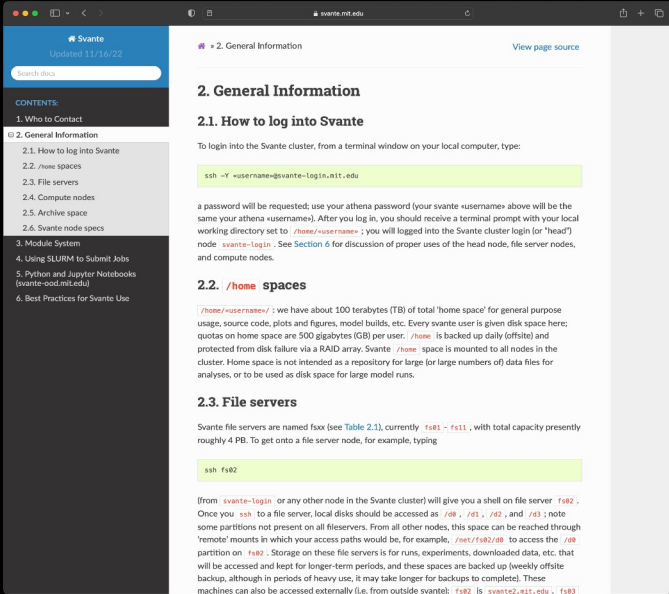
The HPRCF at the Bates Laboratory consists of 71 water-cooled racks, each of which can supply up to 12 kW of power and cooling, and a high-speed 10 Gb/s network link to campus. The construction, racks and network link represent an investment by MIT of about \$7.0M.

The CERN CMS experiment is assigned 32 of the racks for a Tier 2 computing site; currently 24 are in use (one rack houses computers bought with non-CMS funds), with more computers being purchased on a planned schedule over the next several years.

- About
- Research & Engineering
- Photo Galleries
- High Performance Computer Facility**
- Community Outreach
- Projects
- News
- Publications
- Physics Links

Bates

orcd@mit:~# !!



The screenshot shows the Svante website with a navigation menu on the left and a main content area. The main content area is titled "2. General Information" and contains sub-sections: "2.1. How to log into Svante", "2.2. /home spaces", and "2.3. File servers".

2. General Information

2.1. How to log into Svante

To log into the Svante cluster, from a terminal window on your local computer, type:

```
ssh -Y -i username@svante-login.mit.edu
```

a password will be requested; use your athena password (your svante `username` above will be the same your athena `username`). After you log in, you should receive a terminal prompt with your local working directory set to `/home/username`; you will be logged into the Svante cluster (login for "head" node: `svante-login` - See Section 6 for discussion of proper uses of the head node, file server nodes, and compute nodes).

2.2. /home spaces

`/home/username`: we have about 100 terabytes (TB) of total 'home space' for general purpose usage, source code, plots and figures, model builds, etc. Every svante user is given disk space here: quotas on home space are 500 gigabytes (GB) per user. `/home` is backed up daily (offsite) and protected from disk failure via a RAID array. Svante `/home` space is mounted to all nodes in the cluster. Home space is not intended as a repository for large (or large numbers of) data files for analyses, or to be used as disk space for large model runs.

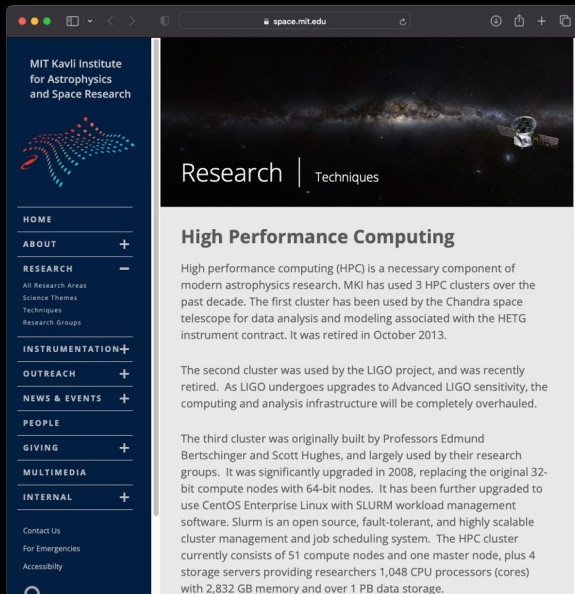
2.3. File servers

Svante file servers are named fox (see Table 2.1), currently `fx81 - fx11`, with total capacity presently roughly 4 PB. To get onto a file server node, for example, typing

```
ssh fx82
```

(from `svante-login` or any other node in the Svante cluster) will give you a shell on file server `fx82`. Once you `ssh` to a file server, local disks should be accessed as `/s8`, `/s1`, `/s2`, and `/s3`; note some partitions not present on all file servers. From all other nodes, this space can be reached through 'remote' mounts in which your access paths would be, for example, `/mnt/fx82/s8`. To access the `/s8` partition on `fx82`, Storage on these file servers is for runs, experiments, downloaded data, etc. that will be accessed and kept for longer-term periods, and these spaces are backed up (weekly) offsite backup, although in periods of heavy use, it may take longer for backups to complete. These machines can also be accessed externally (i.e. from outside quantum): `fx87` is `svante.s11.mit.edu`; `fx88`

SVANTE



The screenshot shows the website for the MIT Kavli Institute for Astrophysics and Space Research. It features a navigation menu on the left and a main content area with a header image of a galaxy and a satellite. The main content area is titled "High Performance Computing" and contains text about the institute's HPC clusters.

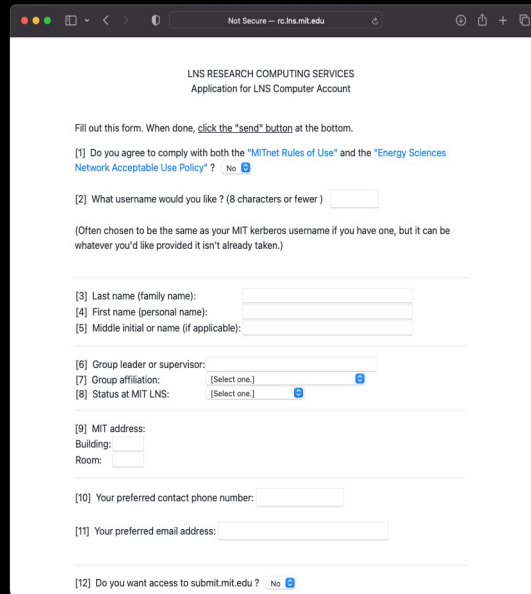
High Performance Computing

High performance computing (HPC) is a necessary component of modern astrophysics research. MKI has used 3 HPC clusters over the past decade. The first cluster has been used by the Chandra space telescope for data analysis and modeling associated with the HETG instrument contract. It was retired in October 2013.

The second cluster was used by the LIGO project, and was recently retired. As LIGO undergoes upgrades to Advanced LIGO sensitivity, the computing and analysis infrastructure will be completely overhauled.

The third cluster was originally built by Professors Edmund Bertschinger and Scott Hughes, and largely used by their research groups. It was significantly upgraded in 2008, replacing the original 32-bit compute nodes with 64-bit nodes. It has been further upgraded to use CentOS Enterprise Linux with SLURM workload management software. Slurm is an open source, fault-tolerant, and highly scalable cluster management and job scheduling system. The HPC cluster currently consists of 51 compute nodes and one master node, plus 4 storage servers providing researchers 1,048 CPU processors (cores) with 2,832 GB memory and over 1 PB data storage.

KAVLI



The screenshot shows a web form titled "LNS RESEARCH COMPUTING SERVICES Application for LNS Computer Account". The form contains several sections with input fields and checkboxes.

Fill out this form. When done, [click the "send" button](#) at the bottom.

[1] Do you agree to comply with both the "MITnet Rules of Use" and the "Energy Sciences Network Acceptable Use Policy"? No Yes

[2] What username would you like? (8 characters or fewer)

(Often chosen to be the same as your MIT kerberos username if you have one, but it can be whatever you'd like provided it isn't already taken.)

[3] Last name (family name):

[4] First name (personal name):

[5] Middle initial or name (if applicable):

[6] Group leader or supervisor:

[7] Group affiliation:

[8] Status at MIT LNS:

[9] MIT address:
Building:
Room:

[10] Your preferred contact phone number:

[11] Your preferred email address:

[12] Do you want access to submit.mit.edu? No Yes

LNS

orcd@mit:~\$ CTRL-C

11 clusters, each independently staffed, funded & operated – they do share one thing though... a TLD.

```
# whois.arin.net

NetRange:      18.0.0.0 - 18.31.255.255
CIDR:          18.0.0.0/11
NetName:       MIT
NetHandle:     NET-18-0-0-1
Parent:        NET18 (NET-18-0-0-0)
NetType:       Direct Allocation
OriginAS:      AS3
Organization:  Massachusetts Institute of Technology (MIT-2)
RegDate:       1994-01-01
Updated:       2021-12-14
Ref:           https://rdap.arin.net/registry/ip/18.0.0.0

OrgName:       Massachusetts Institute of Technology
OrgId:         MIT-2
Address:       Room W92-167
Address:       77 Massachusetts Avenue
City:          Cambridge
StateProv:     MA
PostalCode:    02139-4307
Country:       US
RegDate:
Updated:       2017-01-28
Ref:           https://rdap.arin.net/registry/entity/MIT-2
```

orcd@mit:~\$ quota -v

Engaging-1 today *could* be divided up equitably as:

10,000 TB \approx 1TB dedicated per MIT Researcher

45,000 CPU \approx 4 CPU dedicated per MIT Researcher

Does not include, SuperCloud (32,000 cores, 450GPU), Satori, CSAIL, Kavli, OpenMind, Svante, Submit, LNS & Bates (70+ racks)

And... 100's more systems in Closets, Clouds and Chaos!

What is MIT's total? Over 1,000,000 CPU cores? More?

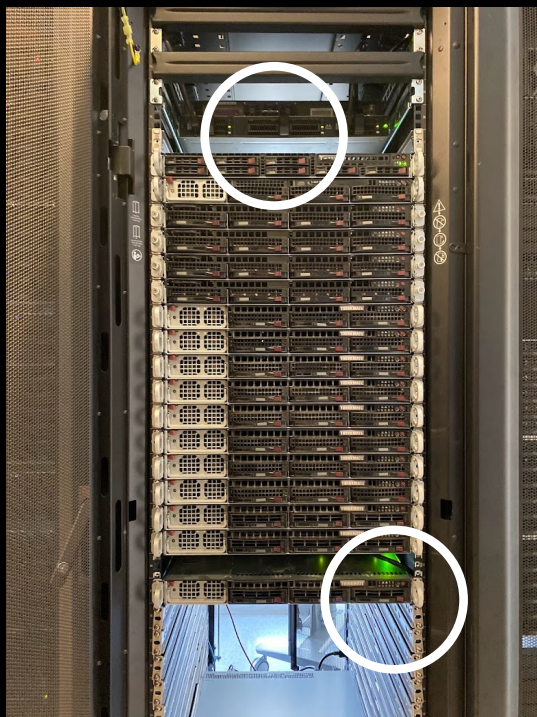
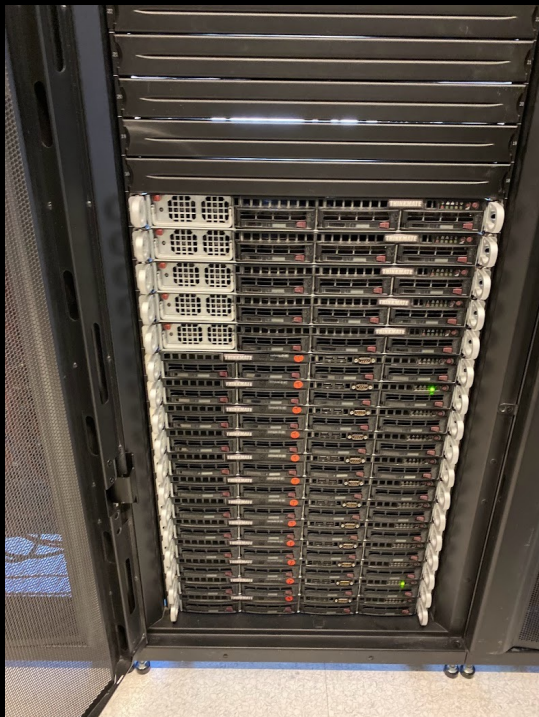
```
orcd@mit:~# find ./mit -name hpc_cluster\*
```

The screenshot shows the MIT Office of Research Computing and Data (ORCD) website. The header includes the ORCD logo and navigation links: About ORCD, Work with ORCD, Collaborations, Resources, and News and Events. A search bar is present in the top right. The main content area features a large image of a modern building with a glass facade. On the left, there are two columns of links: 'Resources' (MIT Information Systems and Technology (IST), MIT Libraries Data Services, Massachusetts Green High Performance Computing Center (MGHPCC)) and 'Related Links' (Work with ORCD, ORCD Resources, News and Events, Jobs at ORCD). A 'Read the latest news from ORCD' link is also visible. The footer contains contact information for the Massachusetts Institute of Technology Office of Research Computing and Data, including a telephone number and email address.

The screenshot shows the MIT Undergraduate Research Opportunities Program (UROP) website. The header features the MIT logo and the UROP logo, along with the text 'Undergraduate Research Opportunities Program'. The main content area has a dark background with a large purple banner at the bottom. The banner text reads: 'Attention Students! Summer ✨ UROP Direct Funding (funds from the UROP Office) Application Deadline ⚡ Tuesday, 04.18.2023'. Above the banner, there is a section titled 'Immerse yourself in real-world faculty research' with a 'CONTACT US' button and a circular graphic containing the text 'We're more than just a URL'. Below the banner, there are two buttons: 'Find Projects' and 'Apply Online'. The footer contains the URL: <https://urop.mit.edu/deadlines/summer-urop-direct-funding-deadline/>.

The screenshot shows the MIT Climate & Sustainability Consortium (MCS) website. The header includes the MCS logo and navigation links: OUR WORK, MEMBERS, GET INVOLVED, PEOPLE, NEWS & EVENTS, and ABOUT US. A search bar and a 'LOGIN' button are in the top right. The main content area features a large image of the Earth from space. The text reads: 'A NEW KIND OF ACADEMIA-INDUSTRY COLLABORATION'. Below this, there is a paragraph of text: 'Working together to vastly accelerate the implementation of large-scale, real-world solutions, across sectors, to help meet global climate and sustainability challenges. Helping to lay the groundwork for one critical aspect of MIT's continued and intensified commitment to climate helping large companies usher in, adapt to, and prosper in a decarbonized world.' A 'READ MORE ABOUT OUR WORK' button is located below the paragraph. The footer contains a 'V' icon.


```
orcd@mit:~# go ./urops.go
```



Wendy Zhang & Yichen Gao

```
orcd@mit:~# future --predict
```

```
ORCD_8ball: Future Looks Cloudy
```

```
Try again? (y/n)
```

```
orcd@mit:~$ ssh root@openai.com /usr/bin/irssi
```

“The future of research computing is expected to be shaped by several factors, including advances in hardware and software technology, changes in research practices, and new data-driven research methods.

- 1) Increased use of cloud computing
- 2) Emphasis on open science
- 3) Expansion of machine learning and AI
- 4) Increased use of big data analytics
- 5) Continued development of quantum computing”

chat.openai.com

New chat

Future of research co

what is the future of research computing?

The future of research computing is expected to be shaped by several factors, including advances in hardware and software technology, changes in research practices, and new data-driven research methods. Here are a few potential trends that could shape the future of research computing:

1. Increased use of cloud computing: Cloud computing offers researchers flexible and scalable access to computing resources, enabling them to perform complex computations and data analysis without having to invest in expensive hardware. As cloud computing becomes more accessible and affordable, it is expected that more researchers will adopt cloud-based solutions for their computational needs.
2. Emphasis on open science: Open science refers to the movement to make research data and publications more transparent and accessible to the broader scientific community. This trend is likely to continue, with researchers using tools and platforms that facilitate data sharing, collaboration, and reproducibility.
3. Expansion of machine learning and AI: Machine learning and artificial intelligence (AI) are becoming increasingly important in research computing, enabling researchers to analyze and interpret large datasets and make predictions based on data patterns. As these technologies continue to advance, they are expected to play an even greater role in scientific research.
4. Increased use of big data analytics: The growth of big data has led to an increase in demand for analytics tools that can process and analyze large datasets quickly and efficiently. Researchers will likely continue to rely on these tools to gain insights from their data and make discoveries that would be difficult or impossible to uncover through traditional methods.
5. Continued development of quantum computing: Quantum computing is an emerging technology that has the potential to revolutionize research computing by enabling researchers to solve complex problems that are beyond the capabilities of traditional computers. While quantum computing is still in its early stages of development, it is expected to become an increasingly important part of research computing in the coming years.

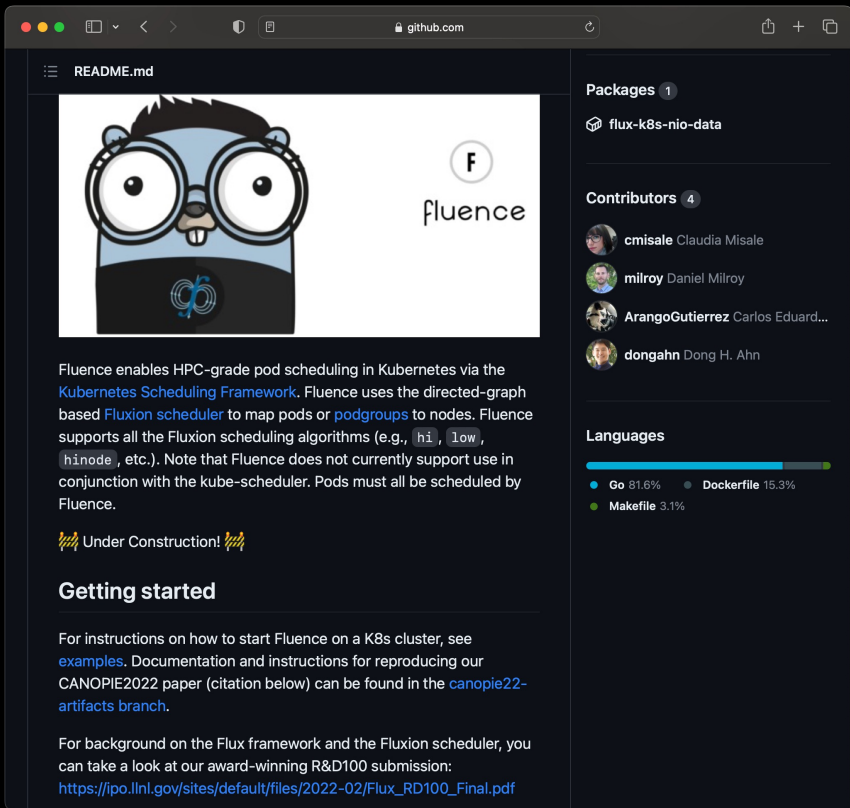
Send a message...

Upgrade to Plus NEW

jcuff@mit.edu

ChatGPT, Mar 23 Version. Free Research Preview. ChatGPT may produce inaccurate information about people, places, or facts.


```
orcd@mit:~# kool_kids_k8.py --hpc --future
```



The screenshot shows the GitHub repository page for 'flux-k8s-nio-data'. The main content is the README.md file, which features a cartoon character with glasses and a blue shirt. The character is wearing a blue 'F' logo. The text in the README describes the Flux Operator and its capabilities. On the right side, there are sections for 'Packages' (1 package: flux-k8s-nio-data), 'Contributors' (4 contributors: cmisale, milroy, ArangoGutierrez, dongahn), and 'Languages' (Go 81.6%, Dockerfile 15.3%, Makefile 3.1%).

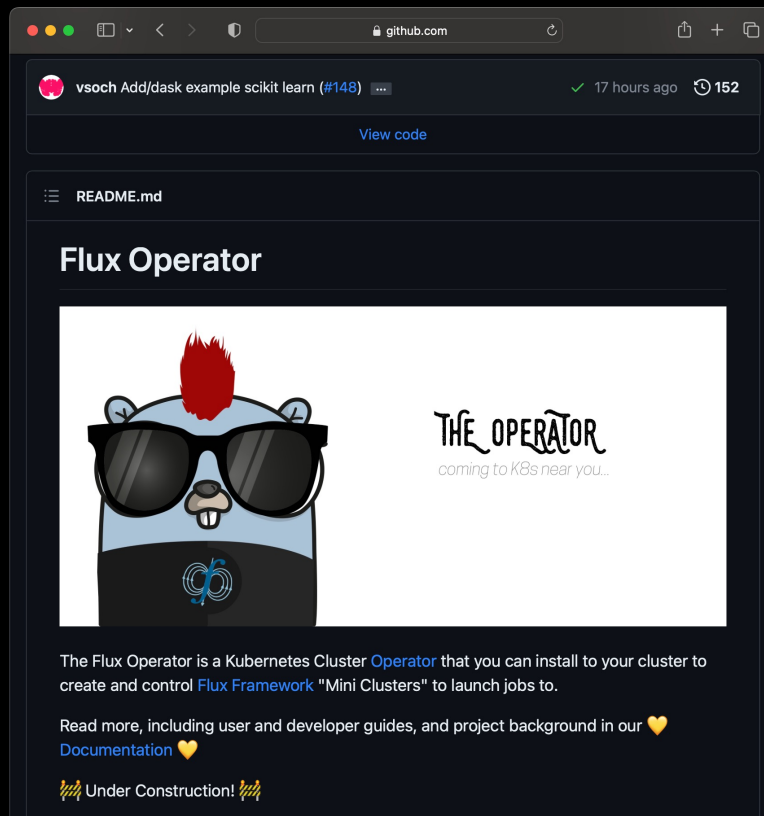
Fluence enables HPC-grade pod scheduling in Kubernetes via the [Kubernetes Scheduling Framework](#). Fluence uses the directed-graph based [Fluxion scheduler](#) to map pods or [podgroups](#) to nodes. Fluence supports all the Fluxion scheduling algorithms (e.g., [hi](#), [low](#), [hi_node](#), etc.). Note that Fluence does not currently support use in conjunction with the kube-scheduler. Pods must all be scheduled by Fluence.

🚧 Under Construction! 🚧

Getting started

For instructions on how to start Fluence on a K8s cluster, see [examples](#). Documentation and instructions for reproducing our CANOPIE2022 paper (citation below) can be found in the [canopie22-artifacts](#) branch.

For background on the Flux framework and the Fluxion scheduler, you can take a look at our award-winning R&D100 submission: https://ipo.llnl.gov/sites/default/files/2022-02/Flux_RD100_Final.pdf



The screenshot shows the GitHub repository page for 'vsoch Add/dask example scikit learn (#148)'. The main content is the README.md file, which features a cartoon character with a red mohawk, sunglasses, and a blue shirt. The character is wearing a blue 'F' logo. The text in the README describes the Flux Operator and its capabilities. On the right side, there are sections for 'Packages' (1 package: flux-k8s-nio-data), 'Contributors' (4 contributors: cmisale, milroy, ArangoGutierrez, dongahn), and 'Languages' (Go 81.6%, Dockerfile 15.3%, Makefile 3.1%).

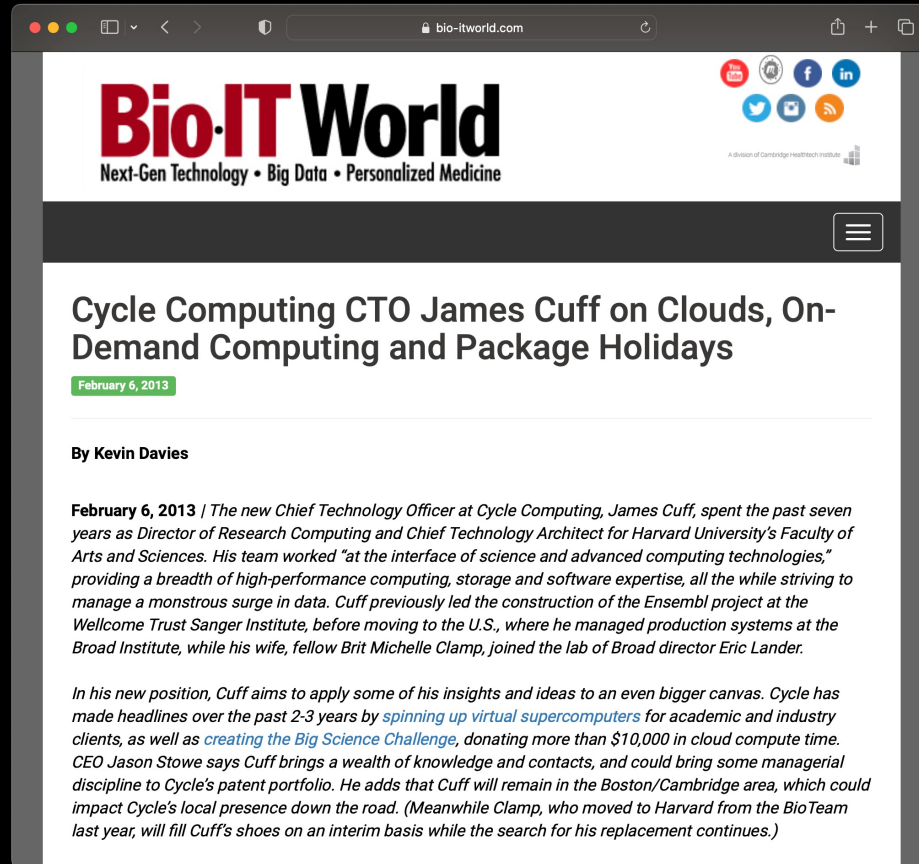
The Flux Operator is a Kubernetes Cluster [Operator](#) that you can install to your cluster to create and control [Flux Framework](#) "Mini Clusters" to launch jobs to.

Read more, including user and developer guides, and project background in our [Documentation](#) ❤️

🚧 Under Construction! 🚧

```
orcd@mit:~# cat /dev/null > ./buzzwords.txt
```

- We've been predicting, and doing all of this stuff for decades
- The Office of Research Computing & Data is all about supporting research **WITH** computing
- The cloud / anti-cloud religion is exhausting & not important, it's just compute, network & storage
- I'll tell you what we're going to do **here @ MIT for everyone @ MIT**



The screenshot shows a web browser window with the address bar displaying "bio-itworld.com". The page header includes the "BioIT World" logo in red and black, with the tagline "Next-Gen Technology • Big Data • Personalized Medicine" below it. To the right of the logo are social media icons for YouTube, Facebook, LinkedIn, Twitter, and SoundCloud. Below the header is a dark navigation bar with a hamburger menu icon. The main content area features an article titled "Cycle Computing CTO James Cuff on Clouds, On-Demand Computing and Package Holidays" with a date tag for "February 6, 2013". The author is listed as "By Kevin Davies". The article text begins with "February 6, 2013 | The new Chief Technology Officer at Cycle Computing, James Cuff, spent the past seven years as Director of Research Computing and Chief Technology Architect for Harvard University's Faculty of Arts and Sciences. His team worked 'at the interface of science and advanced computing technologies,' providing a breadth of high-performance computing, storage and software expertise, all the while striving to manage a monstrous surge in data. Cuff previously led the construction of the Ensembl project at the Wellcome Trust Sanger Institute, before moving to the U.S., where he managed production systems at the Broad Institute, while his wife, fellow Brit Michelle Clamp, joined the lab of Broad director Eric Lander." The article continues with "In his new position, Cuff aims to apply some of his insights and ideas to an even bigger canvas. Cycle has made headlines over the past 2-3 years by spinning up virtual supercomputers for academic and industry clients, as well as creating the Big Science Challenge, donating more than \$10,000 in cloud compute time. CEO Jason Stowe says Cuff brings a wealth of knowledge and contacts, and could bring some managerial discipline to Cycle's patent portfolio. He adds that Cuff will remain in the Boston/Cambridge area, which could impact Cycle's local presence down the road. (Meanwhile Clamp, who moved to Harvard from the Bio Team last year, will fill Cuff's shoes on an interim basis while the search for his replacement continues.)"

```
orcd@mit:~# cat /proc/orcd/mission
```

1

Provide support to Principal Investigators (PIs) and other researchers across MIT for advanced research computing through training, documentation, facilitation, and other means.

2

Ensure a basic level of access to advanced computing and data management to PIs across MIT that allows them to complete research projects, develop new ideas, and compete successfully for research support. Includes coordination with IS&T and Libraries.

3

Ensure efficient use of all of MIT's advanced research computing and data management resources through sharing, curation, and renewal.

4

Coordinate with other MIT Offices including IS&T, Libraries, and School/College Deans in providing advanced computing and data management resources to the educational program through the Office of the Vice Chancellor (OVC).

ORCD IS NOT JUST ANOTHER SHADOW IT ORGANIZATION


```
orcd@mit:~# echo "the 2 pillars of ORCD"
```

Platforms
Infrastructure
and Data
Services

HPC &
Scientific
Consulting
Services

Director of PIDS Starts July 31st 2023

orcd@mit:~# cat .plan

1. Funds to support ORCD staff and operations to enable basic access for 11,611 researchers + 4,400 undergraduates
2. Develop funding model for PI-driven initiatives
 - Sharing between PI, DLC, ORCD, School/College, Provost
3. Create and support secure repository for sensitive and shared data
ORCD_ENCLAVE
4. “Get Well Plan” for shared resources

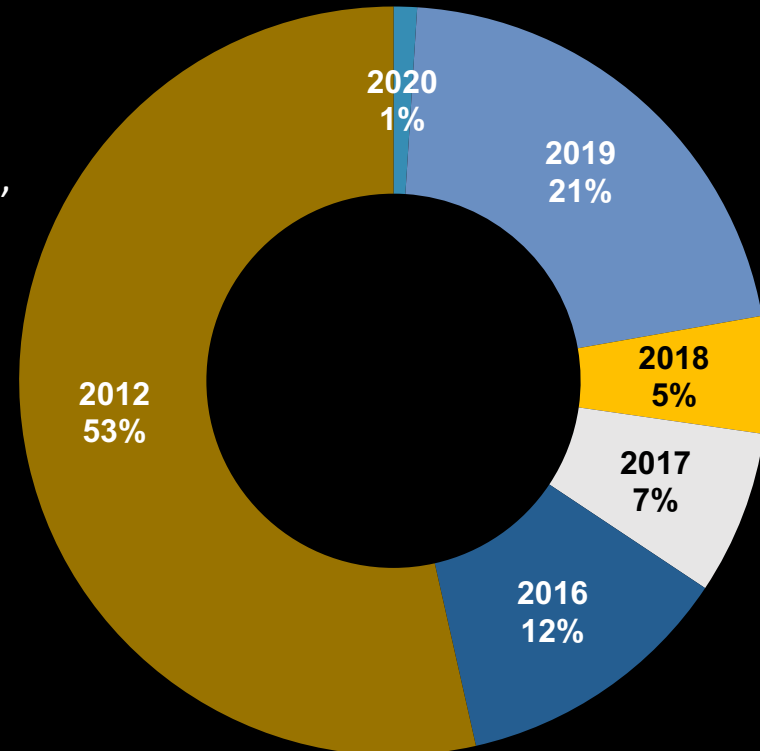


Ensure **basic** access to high performance research computing to **all** 11,611 researchers at MIT, support for PI driven computing initiatives, and consulting help for users and purchasers of computing and data resources.

We have been fully funded.

Modernizing \$20M (replacement value) of over 10 years of legacy computing infrastructure and processes

1. Develop new organization based on **operational excellence and obsessive customer support** working with IS&T, Libraries, Facilities, Sustainability, Audit
 - Decommission legacy systems and infrastructure
 - Build a **reliable fault-tolerant foundation** and sharing platform (login, home, auth, access controls, storage, standards-based networks and on-call FTE)
 - **Secure, harden** all research computing endpoints & data at rest
 - Provide **frictionless** access through 2fA, Globus & Open on Demand
2. Build the ORCD_ENCLAVE: A **secure, sustainable** repository for all **sensitive and shared data**
3. Integrate remaining viable systems into a new shared research computing platform (“Engaging TNG”)



`orcd@mit:~# setenforce 1`

The new ORCD_ENCLAVE is critical

Over the next 12 months will build MVP

Multi petabyte scale shared, secure storage you can compute against

Don't be clever – be reliable, be predictable, be available!

SAFEGUARDING THE U.S. RESEARCH ENTERPRISE:

Transparency, Integrity, and Reciprocity

A FBI ACADEMIC RESEARCH SECURITY CONFERENCE
hosted by Harvard University

NOVEMBER 1, 2022 | 9:00 AM – 2:00 PM ET

Spangler Center, Harvard University
117 Western Ave, Allston, MA 02163



```
orcd@mit:~# mghpcc --show --spare -orcd
```

MGHPCC Capacity report

Date: Apr 24 2023

Racks:	64
RU:	3000
Rack:	50KW
Date:	Oct 2023
MAX:	3200KW
Cooling:	Liquid
Invest:	\$3,000,000
Tech Debt:	NONE



```
orcd@mit:~# python ./install.py
```

Over \$3M of shared community hardware *installed*,
and managed by ORCD since October 2022!

13 departments and 19 PIs

160 computers

13,768 CPU cores

130 Ampere and Hopper class GPUs

Operating savings

~400 tons of CO2 emissions avoided

~0.13MWh energy cost reduction

orcd@mit:~# nvidia-smi

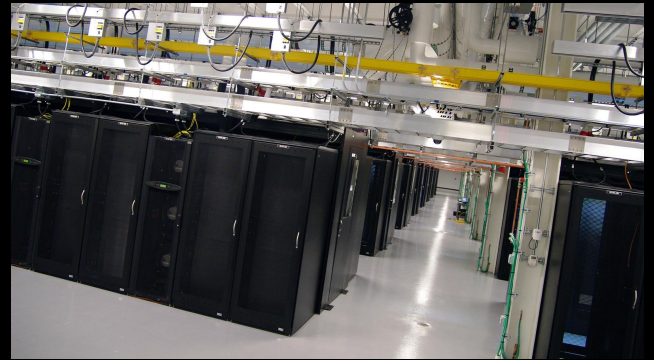
Mon Jun 19 07:55:42 2023

NVIDIA-SMI 530.30.02		Driver Version: 530.30.02		CUDA Version: 12.1			
GPU	Name	Persistence-M	Bus-Id	Disp.A	Volatile	Uncorr. ECC	ECC
Fan	Temp	Pwr:Usage/Cap		Memory-Usage	GPU-Util	Compute M.	MIG M.
Perf							
0	NVIDIA H100 80GB HBM3	On	00000000:19:00.0	Off		0	
N/A	53C P0	551W / 700W	77753MiB / 81559MiB		97%	Default Disabled	
1	NVIDIA H100 80GB HBM3	On	00000000:3B:00.0	Off		0	
N/A	52C P0	685W / 700W	77653MiB / 81559MiB		97%	Default Disabled	
2	NVIDIA H100 80GB HBM3	On	00000000:4C:00.0	Off		0	
N/A	51C P0	549W / 700W	64345MiB / 81559MiB		97%	Default Disabled	
3	NVIDIA H100 80GB HBM3	On	00000000:5D:00.0	Off		0	
N/A	57C P0	726W / 700W	77573MiB / 81559MiB		97%	Default Disabled	
4	NVIDIA H100 80GB HBM3	On	00000000:9B:00.0	Off		0	
N/A	56C P0	639W / 700W	77945MiB / 81559MiB		97%	Default Disabled	
5	NVIDIA H100 80GB HBM3	On	00000000:BB:00.0	Off		0	
N/A	55C P0	683W / 700W	70665MiB / 81559MiB		97%	Default Disabled	
6	NVIDIA H100 80GB HBM3	On	00000000:CB:00.0	Off		0	
N/A	53C P0	600W / 700W	64345MiB / 81559MiB		97%	Default Disabled	
7	NVIDIA H100 80GB HBM3	On	00000000:DB:00.0	Off		0	
N/A	52C P0	554W / 700W	77389MiB / 81559MiB		97%	Default Disabled	



MGHPCC: Massachusetts Green High Performance Computing Center

Power comes from a hydroelectric plant
Electricity costs 1/2 of what it does in Cambridge
Building operates 2x as efficiently as Cambridge
LEED Platinum building and operations since 2012



BOSTON
UNIVERSITY



HARVARD
UNIVERSITY

Northeastern

Search by Name:

MGHPCC [Export ' XML] [Click to show statistics](#)

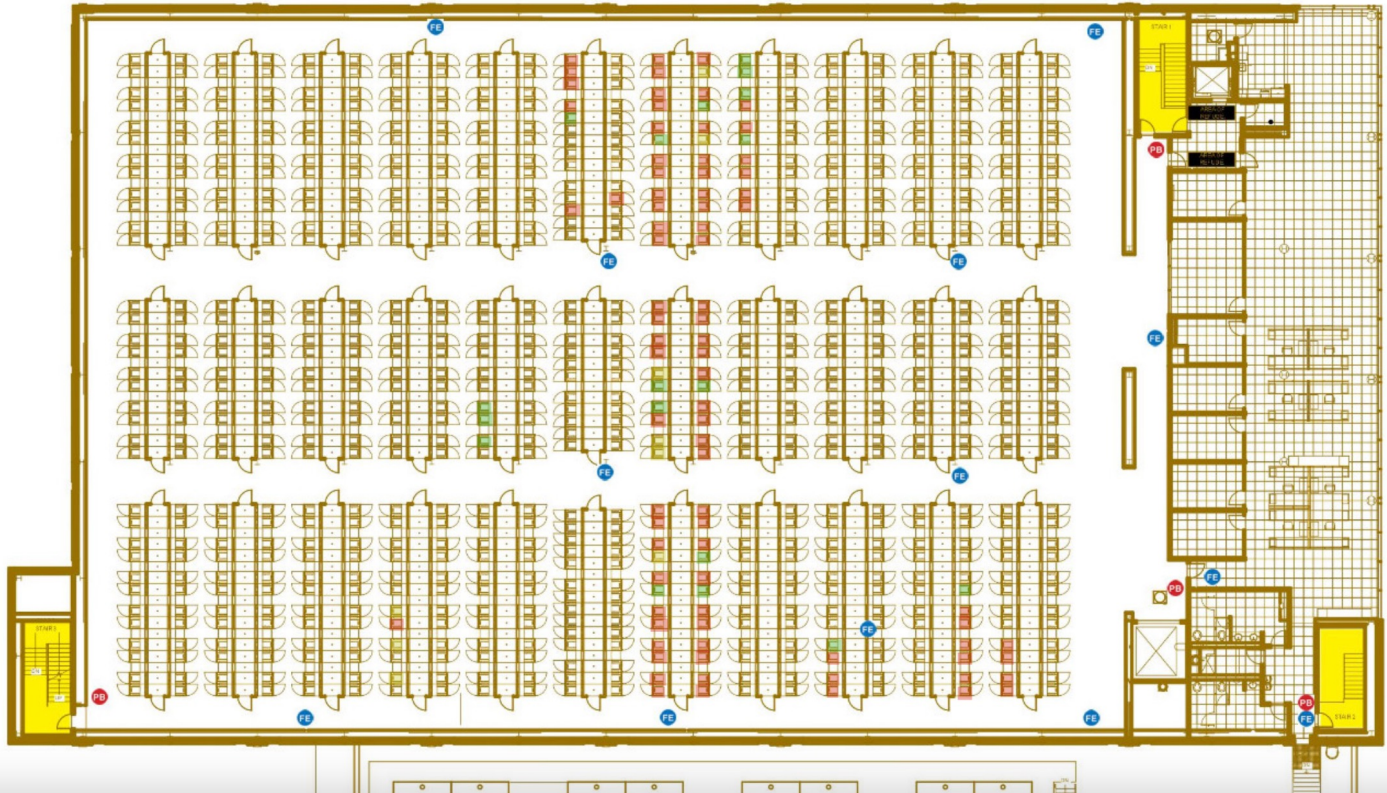
[Advanced] [Custom Search]

Composite View of Cabinets

Overview

- Reports
 - Rack Request Form
 - User Administration >
 - Issue Escalation >
 - Project Catalog
 - Template Management >
 - Infrastructure Management >
 - Bulk Importer >
 - Material Management >
 - Power Management >
 - Path Connections >
 - Edit Configuration

- Home
 - Massachusetts
 - MGHPCC
 - R1-PC
 - R2-PC
 - R3-PC
 - R4-PA
 - R5-PA
 - R5-PB
 - R5-PC
 - R6-PA
 - R7-PB
 - R8-PC
 - R9-PA
 - Storage Room
 - NESE
 - General Storage Room



orcd@mit:~\$ play hua_rong_dao.wav

張 飛	曹 操		趙 雲
馬 超	關 羽		黃 忠
	卒	卒	
卒			卒

In the year 208 Cao Cao led 220,000 troops of the Wei army to fight against an army of 50,000 Shu troops in a mountainous area near Chibi in today's Hubei province. Because of some strategic errors, Cao Cao's troops were badly defeated in the Battle of Chibi, and he fled with only a handful of his soldiers.

The opening at the bottom of the board is Huarong Pass. Initially the blocks are arranged as shown here, with Cao Cao's block trapped by the other nine. The player's job is to slide blocks horizontally and vertically so that Cao Cao can eventually escape through the pass.

(Hint: The game can be solved in 81 moves of 25,955)

So, in summary...

- Provide a “One MIT” research computing and data infrastructure
- Enable MIT to be competitive for MRI & other \$10-\$100M+ computationally intensive next generation awards
- Provide a **safe, professional and secure environment** for research scientists and their data
- Develop **highly skilled technical systems support staff** and research facilitators & engineers to support faculty and researchers
- Provide significant **economies of scale** and be a forcing function for research computing **best practices**, a **safe place** for investments, and USG data requirements
- Direct **sustainable core funding to make every dollar count** towards research -- more compute per KW

Be a first-class citizen of science to support MIT's research mission



orcd@mit:~# finger orcd

```
Login: orcd                               Name: Office of Research Computing
Directory: /mit/vpr/orcd                  Shell: /bin/tcsh
Last login Thur Jan  5 04:56 2023 (EDT) on pts/0 from 192.168.99.5
New mail received Fri Jan  6 05:25 2023 (EST)
      Unread since Fri Jan  6 05:24 2023 (EST)
```

Plan:

Build a great, sustainable, reliable and professional “One MIT” research computing experience for faculty, staff and students

```
URL:                                     https://orcd.mit.edu
JAMES CUFF:                              https://mit.edu/~jcuff
EMAIL:                                    jcuff@mit.edu; fisherp@mit.edu
MAILING LIST:                             orcd-admin@mit.edu
```

```
orcd@mit:~# cat ~/acknowledgements.txt
```

Chris Hill

Lauren Milechin

Paul Hsi

Wendy Zhang & Yichen Gao

Heather Williams

Christina Andujar

Mary Markel Murphy

The VPR web team & all @ TechSquare

Prof. Peter H. Fisher

Mark Silis

Jeremy Gregory

Joe Higgins

Greg Moffatt

Glen Shor

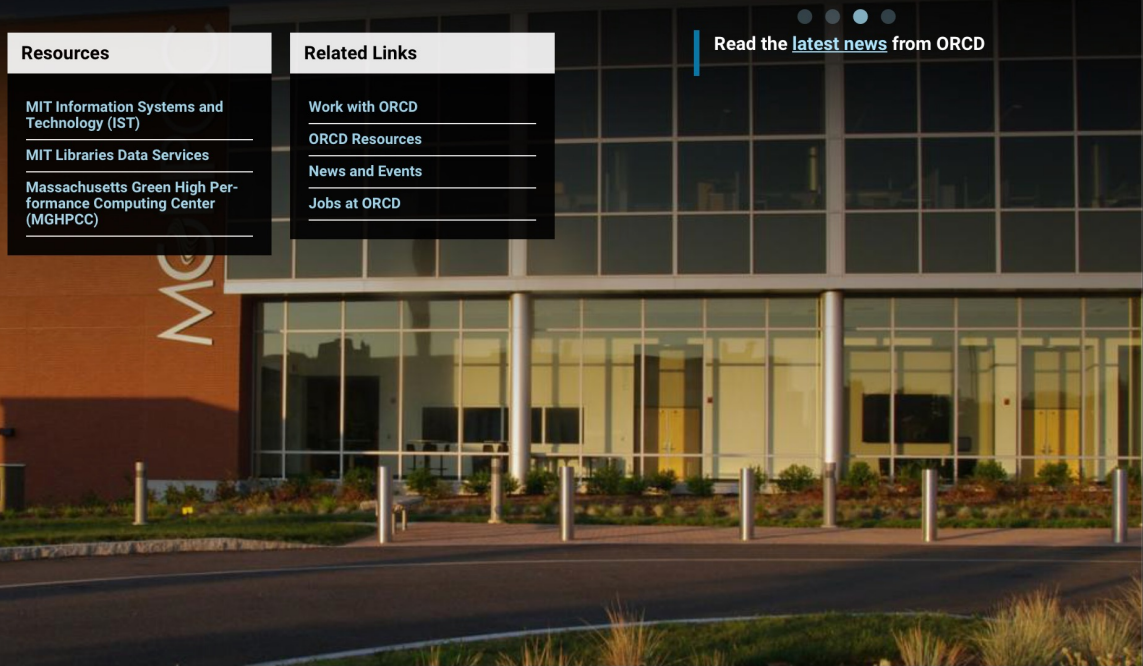
Prof. Maria Zuber

Prof. Cynthia Barnhart

```
orcd@mit:~$ echo 'echo anyone?' \  
> questions.sh
```

```
orcd@mit:~$ while true; \  
do sh ./questions.sh; done
```


MIT Office of Research Computing and Data (ORCD)



Resources

- [MIT Information Systems and Technology \(IST\)](#)
- [MIT Libraries Data Services](#)
- [Massachusetts Green High Performance Computing Center \(MGHPCC\)](#)

Related Links

- [Work with ORCD](#)
- [ORCD Resources](#)
- [News and Events](#)
- [Jobs at ORCD](#)

Read the [latest news](#) from ORCD