

Navigating NERSC



New User Training
September 07, 2023

Charles Lively III, PhD
Science Engagement Engineer/HPC Consultant/Software
User Engagement Group

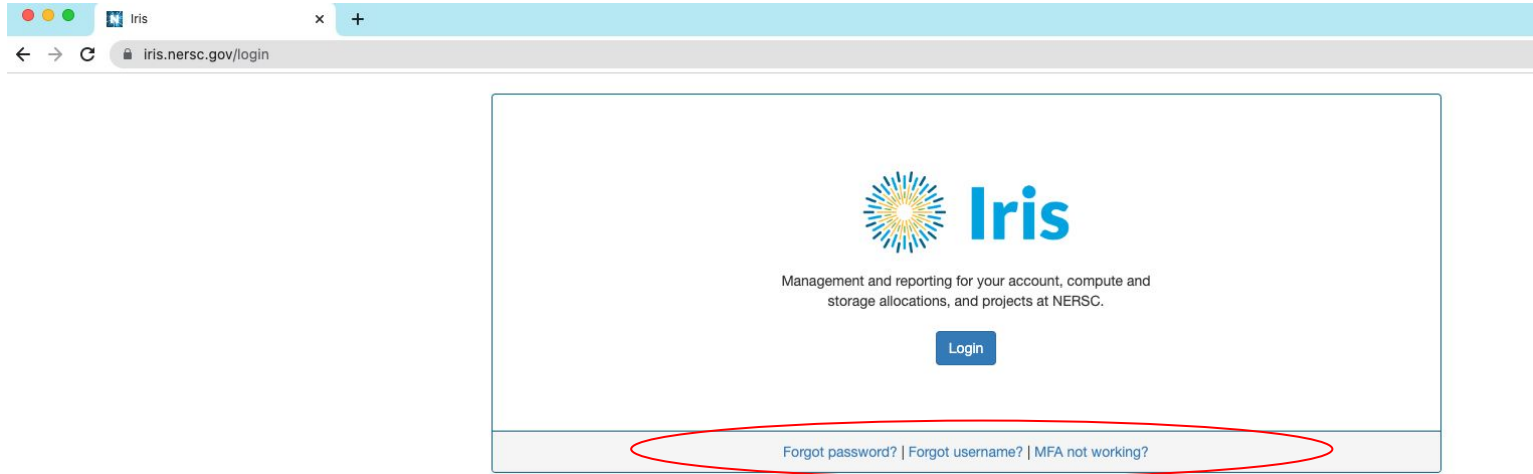
Agenda

- Navigating Iris (<https://iris.nersc.gov>)
- Submitting a User Ticket (<https://help.nersc.gov>)
- MyNERSC (<https://my.nersc.gov>)
- Connecting to Perlmutter
 - a. Connecting with SSH
 - b. <https://jupyter.nersc.gov> notebooks and terminals in your browser
 - c. NoMachine (<https://docs.nersc.gov/connect/nx/>) for GUI apps
- Navigating NERSC Home Page
- Navigating NERSC Documentation

Navigating Iris

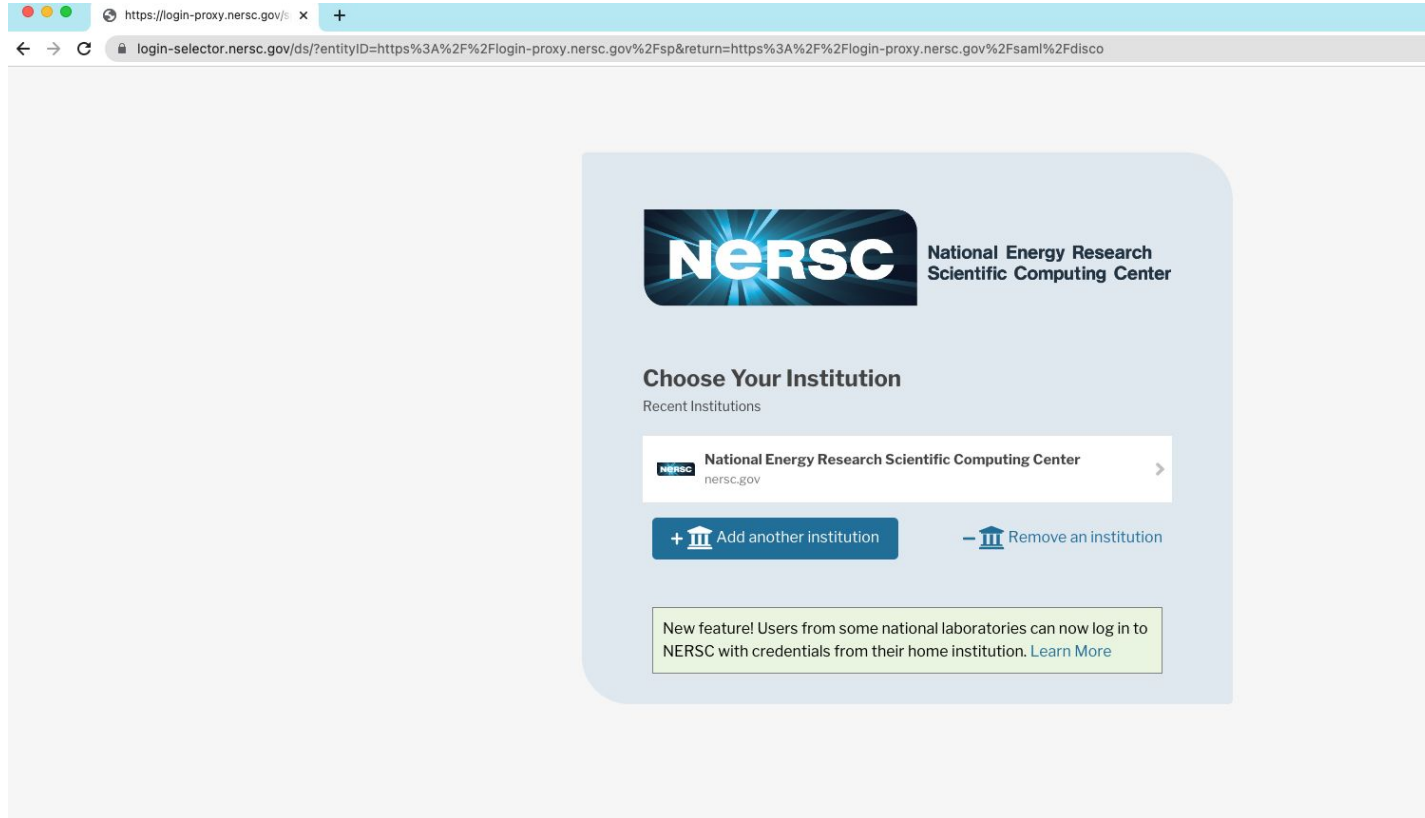


Iris (<https://iris.nersc.gov>) for Your Account



 [Help](#) | [ERCAP](#) | [Docs](#) | [MyNERSC](#)

Iris (<https://iris.nersc.gov>) for Your Account



The screenshot shows a web browser window with the URL `https://login-proxy.nersc.gov/ds/?entityID=https%3A%2F%2Flogin-proxy.nersc.gov%2Fsp&return=https%3A%2F%2Flogin-proxy.nersc.gov%2Fsaml%2Fdisco`. The main content area features the NERSC logo and the text "National Energy Research Scientific Computing Center". Below this is the heading "Choose Your Institution" and the sub-heading "Recent Institutions". A single institution is listed: "National Energy Research Scientific Computing Center" with the domain "nersc.gov". There are two buttons: "+ Add another institution" and "- Remove an institution". A note at the bottom states: "New feature! Users from some national laboratories can now log in to NERSC with credentials from their home institution. [Learn More](#)".

Iris (<https://iris.nersc.gov>) for Your Account

NERSC Login

<https://shib.nersc.gov/idp/profile/SAML2/Redirect/SSO?execution=e1s2>



NERSC National Energy Research Scientific Computing Center

Username

Password

Log In

[Forgot your password?](#)

  U.S. DEPARTMENT OF ENERGY
Office of Science

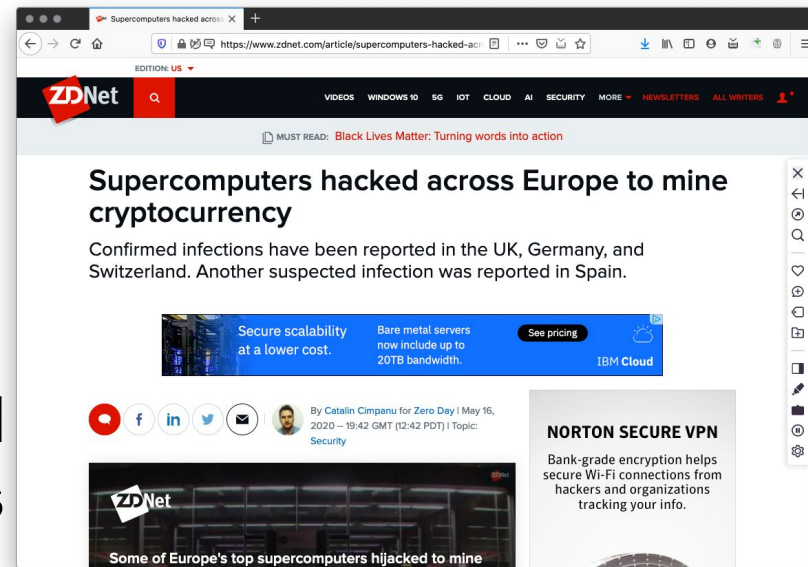
[Contact us](#)
[Privacy & Security Notice](#)



Multi-Factor Authentication (MFA)

Tip: you will use this a LOT

- Protects NERSC users from attacks like this →
- **Log into NERSC resources with your NERSC password plus a one-time code that is provided by an app**




Setting Up MFA in Iris

- First install Google Authenticator on your smartphone (and/or Authy on your computer)

<https://play.google.com/store/apps/details?id=com.google.android.apps.authenticator2&hl=en>
<https://itunes.apple.com/us/app/google-authenticator/id388497605?mt=8>

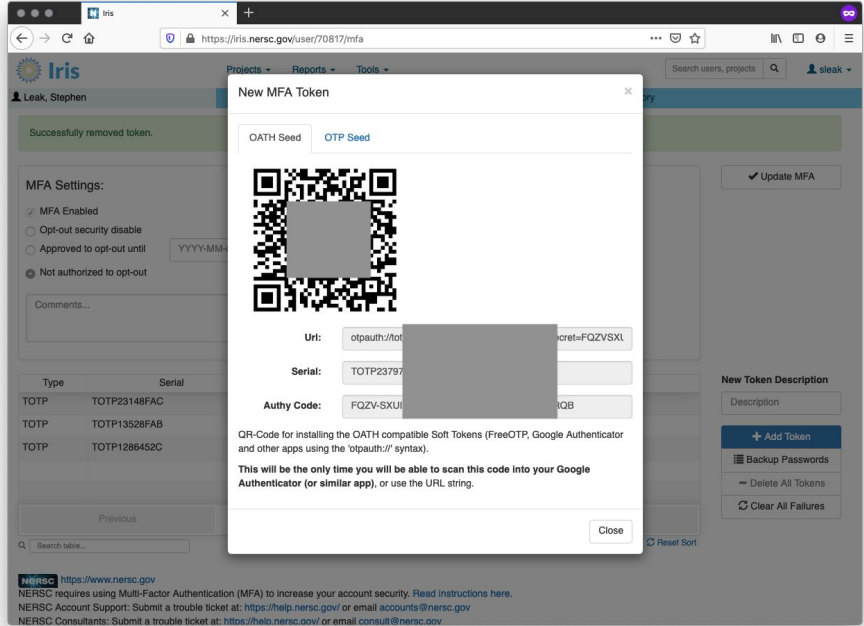
<https://authy.com>

Search "MFA" at
<https://docs.nersc.gov>




Setting Up MFA in Iris

- Click the "MFA" tab
- Click the "Add Token" button
- Scan the QR code with the Authenticator app (or, paste the Authy code into Authy)




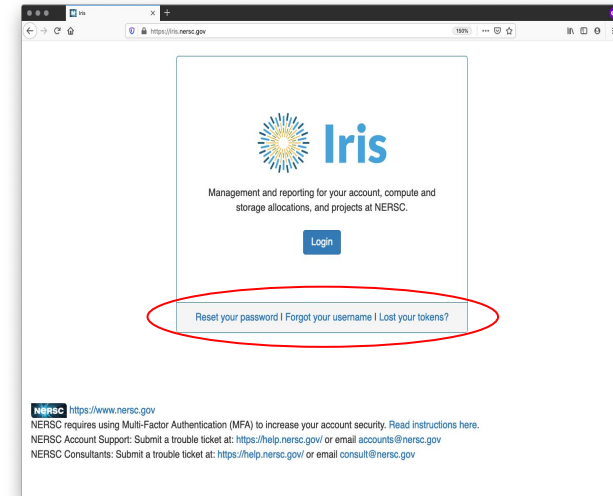
Logging in with MFA

After single-sign-on page you'll be asked for your one-time password (6 digits from app)



Troubleshooting Account Access

- I can't login to Iris
 - New account? It may not be approved yet (can take a few days)
 - Forgot password? Lost MFA tokens? Use the links on the Iris login page
- I can login to Iris, but not Cori or Perlmutter
 - Are you in a project? Check "Roles" tab



Navigating Iris - Menu Bar

The screenshot shows the Iris web application interface. At the top is a navigation menu bar with the following items: CPU, GPU, Jobs, Storage, Roles, Groups, MFA, Profile, and History. A hamburger menu icon is on the left, and a search icon is on the right. Below the menu bar, the main content area is divided into sections. The 'User Organization' section contains a table of user details. The 'Job Details' section contains an ORCID ID field.

Annotations with arrows point to the following menu items and content:

- CPU Account Membership (points to CPU)
- GPU Account Membership (points to GPU)
- Storage Details (points to Storage)
- NERSC Account Membership (points to Roles)
- MFA Token (points to MFA)
- Profile Information (points to Profile)
- Audit Log (points to History)
- Unix Group Membership (points to Groups)

User Organization	
Organization	Lawrence Berkeley National Laboratory - NERSC
Role	PROSTAFF
Org Type	NERSCOPS
Address	1 Cyclotron Road
City	Berkeley
State	CA
Province	
Postal Code	94720
Country	United States of America (US)
Website	www.nersc.gov
ROR ID	https://ror.org/05v8mvq14

All Organizations

Job Details	
ORCID Id	
ORCID	<input type="text" value="0009-0008-4053-2139"/>

All NERSC users are required to have a valid ORCID id. You can easily obtain an ORCID ID by using this [link](#)
If you are unable to create and connect an ORCID id, you can set the value to "n/a".

Navigating Iris - Finding Account Details

Iris Role: consultant

Project Roles:

Project	Description	Role	Accounts	Groupers	Created	Updated
m4232	The Weather Research an...	pi_proxy	m4232_g (DOE) m4232 (DOE)	✓	2023-05-10 10:40	2023-05-10 10:40
m4388	SHI HPC Projects	pi_proxy	m4388 (DOE) m4388_g (DOE)	✓	2023-05-24 04:48	2023-05-24 04:48
nstaff	NERSC Staff Accounts	user	nstaff (DOE) nstaff_g (DOE)	✓	2023-03-24 04:44	2023-03-24 04:44

Iris Search users

m4232 CPU GPU Jobs Storage Roles Groups **Details** History

Project Overview

Click image to edit

Project funding

Allocation Pool: DOE Allocation Pool
Allocation Type: DOE Mission Science
Office: Biological and Environmental Research
Program: Earth and Environmental Systems Sciences Division (EESSD)
Science Category: Earth Systems : Atmosphere
Slurm Category:

ERCAP project details

Organization: Pacific Northwest National Laboratory (PNNL), US
DOE Sensitive Identifiers:
Compute requested in ERCAP: 2.5 K hours
GPU requested in ERCAP: 0.0 hours
HPSS requested in ERCAP: 1.0 TB
CFS storage available: 20.0 TB
CFS files available: 20.0 M
CFS max projectdirs: 10
Funded by DOE Office of Science? Y
Request # ERCAP0025429

Project owners

pi: [Sakaguchi, Koichi](#)
Koichi.Sakaguchi@pnnl.gov

pi proxy: [Leak, Stephen](#)
sleak@lbl.gov

pi proxy: [He, Yun \(Helen\)](#)
yhe@lbl.gov

pi proxy: [Lively, Charles](#)
charleslively@lbl.gov

pi proxy: [Gupta, Lipi](#)
lipigupta@lbl.gov

Navigating Iris - Changing User Shell

Server Logins

Ldap Tree	Home Directory	Login Shell	Username	GID	Group	Actions
alvarez	/global/homes/s/siddiq90	/bin/bash	siddiq90	92503	siddiq90	Edit Delete Search
cori	/global/homes/s/siddiq90	/bin/zsh	siddiq90	92503	siddiq90	Edit Delete Search
datatran	/global/homes/s/siddiq90	/bin/zsh	siddiq90	92503	siddiq90	Edit Delete Search
gerty	/global/homes/s/siddiq90	/bin/zsh	siddiq90	92503	siddiq90	Edit Delete Search
hpss	/home/s/siddiq90	/bin/bash	siddiq90	92503	siddiq90	Edit Delete Search
muller	/global/homes/s/siddiq90	/bin/zsh	siddiq90	92503	siddiq90	Edit Delete Search
nim-ldap	/home/siddiq90	/bin/bash	siddiq90	92503	siddiq90	Edit Search
nx	/global/homes/s/siddiq90	/bin/zsh	siddiq90	92503	siddiq90	Edit Delete Search
perimutter	/global/homes/s/siddiq90	/bin/bash	siddiq90	92503	siddiq90	Edit Delete Search
server	/home/siddiq90	/bin/bash	siddiq90	92503	siddiq90	Edit Delete Search

Previous Page 1 of 2 10 rows Next

Search table... [.csv](#) [Options](#)

Add a new Server Login

Server

Home Directory

Login Shell

Username

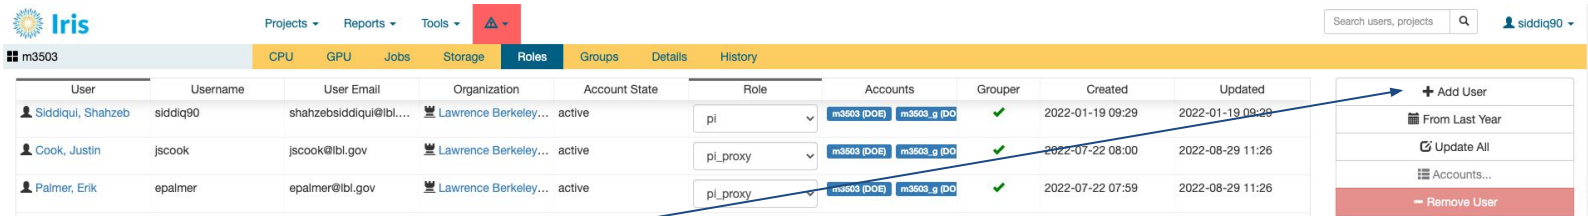
GID

Group Name

Save Changes

Cancel

Navigating Iris - Adding User to Account



Projects ▾ Reports ▾ Tools ▾ ⚠ ▾

Search users, projects 🔍 siddiq90 ▾

m3503 CPU GPU Jobs Storage **Roles** Groups Details History

User	Username	User Email	Organization	Account State	Role	Accounts	Group	Created	Updated
Siddiqui, Shahzeb	siddiq90	shahzebsiddiqui@lbl...	Lawrence Berkeley...	active	pi	m3503 (DOE) m3503_g (DO)	✓	2022-01-19 09:29	2022-01-19 09:29
Cook, Justin	jscook	jscook@lbl.gov	Lawrence Berkeley...	active	pi_proxy	m3503 (DOE) m3503_g (DO)	✓	2022-07-22 08:00	2022-08-29 11:26
Palmer, Erik	epalmer	epalmer@lbl.gov	Lawrence Berkeley...	active	pi_proxy	m3503 (DOE) m3503_g (DO)	✓	2022-07-22 07:59	2022-08-29 11:26

- + Add User
- 📅 From Last Year
- 🔄 Update All
- 📄 Accounts...
- ➖ Remove User

Add a user to Project **m3503** ✕

Use this form to add an **active** NERSC user to this project. To add a **new** or **deactivated** user, please [invite them](#) instead.

Select a user
Please select a valid user.

Role

You can either grant this user a number of node hours they cannot exceed, or a percentage of the project's compute allocation. The latter takes precedence in case they're both specified.

CPU allocation

Allocated Hours

% of Project's Hours

GPU allocation

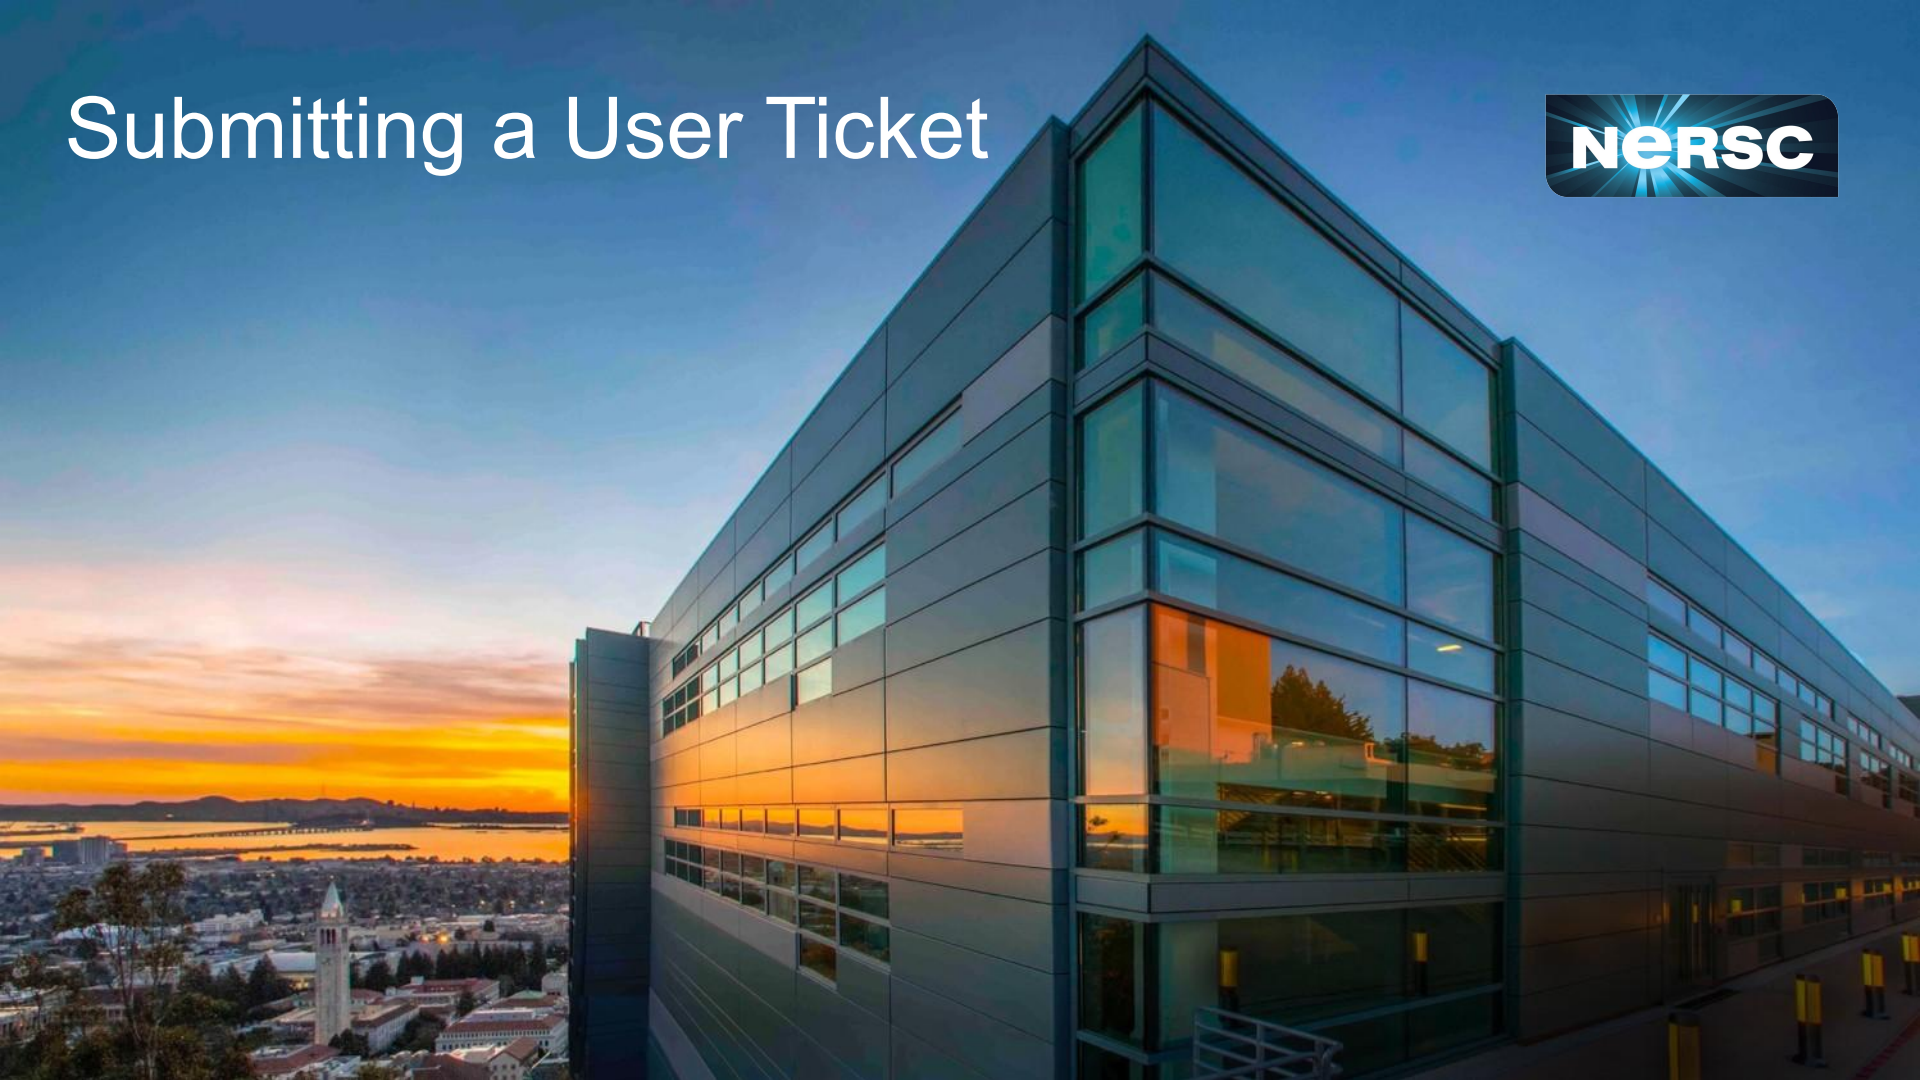
Allocated Hours

% of Project's Hours

Please specify what is the max percent of the project's HPSS allocation that can come from this user.

% of HPSS Storage

Submitting a User Ticket



NERSC Help Portal: <https://help.nersc.gov/>


← → ↻ nersc.servicenowservices.com/sp/ 🔍 🏠 ☆ ⚙️ 🗖


Perlmutter Job Submissions
Project allocation on Perlmutter CPU-only node jobs use (#SBATCH -A mXXXX); for the Perlmutter GPU node jobs you must use the _g_ accounts (#SBATCH -A mXXXX_g).


NERSC Classic View NERSC Homepage

NERSC Help Portal

Search Incidents and Requests 🔍

 **Documentation**
Technical documentation for users, including examples

 **Open Ticket**
Contact NERSC support to report a problem

 **Open Request**
Quota increases, reservations, databases, etc.

Service Announcements

No upcoming maintenances in the next two weeks

Useful Links

- Password Reset
- Book Consulting appointment
- NERSC Status Page
- NERSC Users Slack
- ERCAP
- IRIS

My Recent Incidents

Sije account do not deactivate
INC0189999 • Closed • 3d ago

My Watchlist Open Incidents

BerkeleyGW epsilon run not normal
INC0177149 • 10mo ago • User Updated • Tang, Hong (tang2017)

disk quota
INC0195580 • 9h ago • Awaiting User Info • Prota, James (jprota)

Quantum chemistry programs on Perlmutter
INC0190578 • 8d ago • Awaiting User Info • Greenman, Loren (loreng)

My Projects' Open Incidents

Compute Reservation Request
INC0190690 • 11d ago • Active-Expectations Set • Ross, Hannah (hross)

Unable to login to NERSC systems
INC0190549 • 13d ago • Awaiting User Info • Bhalachandra, Sridutt (sriduttb)


Degraded Performance on Perlmutter nid001532
INC0190502 • 13d ago • New • Bhalachandra, Sridutt (sriduttb)

Compute Reservation Request
INC0185551 • 4mo ago • Active • Blaschke, Johannes (blaschke)

DGEMM performance issues with Perlmutter Nodes (nid003069 and nid003177)
INC0179995 • 7mo ago • Awaiting Vendor • Bhalachandra, Sridutt (sriduttb)

HPSS Allocations
INC0180078 • 8mo ago • User Updated • Nugent, Peter (nugent)

Request Forms



- CNAMEs
 - Request a nersc.gov CNAME
- Collaboration Account
 - Request a collaboration account
- Compute Reservations
 - Reserve time on compute nodes in advance
- Cori GPU Node Access
 - Request access to Cori GPU nodes
- Databases
 - Request a NERSC hosted science database
- IP Addresses
 - Request a NERSC IP address
- Large Memory Nodes Access
 - Request access to Large Memory Nodes
- NERSC VAST Filesystem Account
 - Request a NERSC VAST filesystem account
- Realtime Queue Access
 - Request realtime queue access
- Report a ServiceNow Issue
 - Report an issue on the ServiceNow platform
- Storage Quota
 - Request an increase in storage space or inodes
- Superfacility API or SSH Proxy Extended Access Request
 - Request read/write/execute access to the Superfacility API or an SSH Proxy certificate with an extended lifetime
- Training Accounts
 - Request NERSC Training Accounts
- VASP License Confirmation Request
 - Request to gain access to the VASP binaries provided by NERSC
- Workflow Node Access
 - Request access to workflow nodes



Home > Storage Quota

Search Catalog

Storage Quota
Request an increase in storage space or inodes

Please fill in the information below to request an HPSS Storage Allocation or disk quota change.

*Indicates required

*Which filesystem?

None

New total storage allocation or space quota (TB)

New inode (number of files and directories) quota (millions)

Impact on your work

Low

*Is your project continuing during next AIT?

None

Quota change expiration date

Reason for this request

Submit

Required Information

Which filesystem? Is your project continuing during next AIT? Quota change expiration date Reason for this request

How to file a Good Ticket

- NERSC receive thousands of user support tickets every year and we strive to resolve tickets in timely manner.
- In order for us to troubleshoot your user request, we need **as much information** in ticket to best understand the problem and find a solution.

How to File a Good Ticket 📄

NERSC Consultants handle thousands of support requests per year. In order to ensure efficient timely resolution of issues include **as much of the following as possible** when making a request

- error messages
- jobids
- location of relevant files
 - input/output
 - job scripts
 - source code
 - executables
- output of `module list`
- any steps you have tried
- steps to reproduce

Please copy and paste any text directly into the ticket and only include screenshots as attachments when the graphical output is the subject of the support request.

<https://docs.nersc.gov/getting-started/#how-to-file-a-good-ticket>

MyNERSC



<https://my.nersc.gov>

If you only remember **one** URL, <https://my.nersc.gov> will get you everywhere NERSC

The screenshot shows the My NERSC dashboard interface. On the left is a navigation sidebar with links to Dashboard, Jobs, Center Status, File Browser, Service Tickets, Data Dashboard, PI Toolbox, Jupyter Hub, NERSC Homepage, Documentation Portal, and Accounts Portal. The main content area is titled 'Dashboard' and includes sections for 'My Personal Disk Usage', 'My Active Jobs', and 'My Completed Jobs'. On the right, there is a 'System Status' section with the following data:

System Category	System Name	Status
Compute Systems:	Cort	Up
	Perimutter	Degraded
Global Filesystems:	Community File System (CFS)	Degraded
	DNA	Up
	Data Transfer Nodes	Up
	Global Common	Up
	Global Homes	Up
Mass Storage Systems:	HPSS Archive (User)	Up
	HPSS Regent (Backup)	Up

At the bottom of the System Status section, it indicates 'Service Status: All Services Available'.

<https://my.nersc.gov>

The screenshot shows the my.nersc.gov dashboard. On the left is a navigation sidebar with links to Dashboard, Jobs, Center Status, File Browser, Service Tickets, Data Dashboard, PI Toolbox, Jupyter Hub, NERSC Homepage, Documentation Portal, and Accounts Portal. The main content area is titled 'Dashboard' and contains three sections: 'My Personal Disk Usage', 'My Active Jobs', and 'My Completed Jobs'. The 'My Personal Disk Usage' section shows two bar charts: 'HOME' (Used 27 GB of 40 GB) and 'CSCRATCH' (Used 121 GB of 20,971 GB). The 'My Active Jobs' section shows 'No Active Jobs'. The 'My Completed Jobs' section contains a table with columns for Job ID, Host, Completion Time, Wall Hours, and CPU Hours. The 'System Status' section on the right lists the status of various systems: Compute Systems (Cori: Up, Perlmutter: Degraded), Global Filesystems (Community File System (CFS): Degraded, DNA: Up, Data Transfer Nodes: Up, Global Common: Up, Global Homes: Up), Mass Storage Systems (HPSS Archive (User): Up, HPSS Regent (Backup): Up), Service Status (All Services Available), and Planned Outages.

Job ID	Host	Completion Time	Wall Hours	CPU Hours
62890437	Cori	09/19/22 13:58	0.000	0.00
62890424	Cori	09/19/22 13:58	0.001	0.00
62890423	Cori	09/19/22 13:58	0.001	0.00
62890422	Cori	09/19/22 13:58	0.001	0.00
62889438	Cori	09/19/22 13:36	0.000	0.00
62889430	Cori	09/19/22 13:36	0.001	0.00
62889429	Cori	09/19/22 13:36	0.001	0.00
62889427	Cori	09/19/22 13:36	0.001	0.00
62889426	Cori	09/19/22 13:36	0.001	0.00

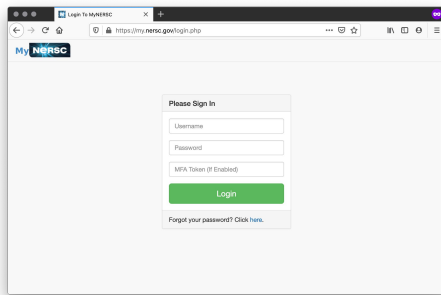
my disk quota

Perlmutter Status

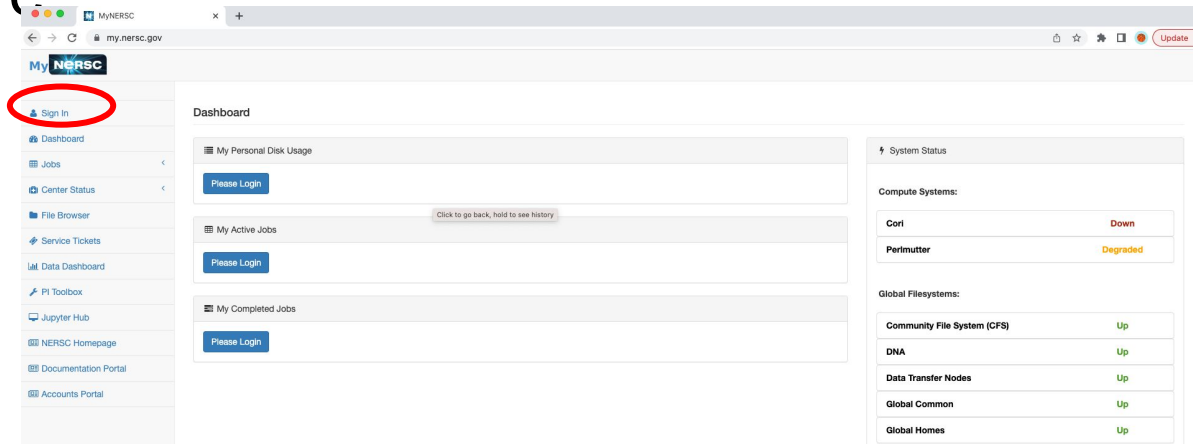
my jobs

<https://my.nersc.gov>

Most things require
login (also MFA)



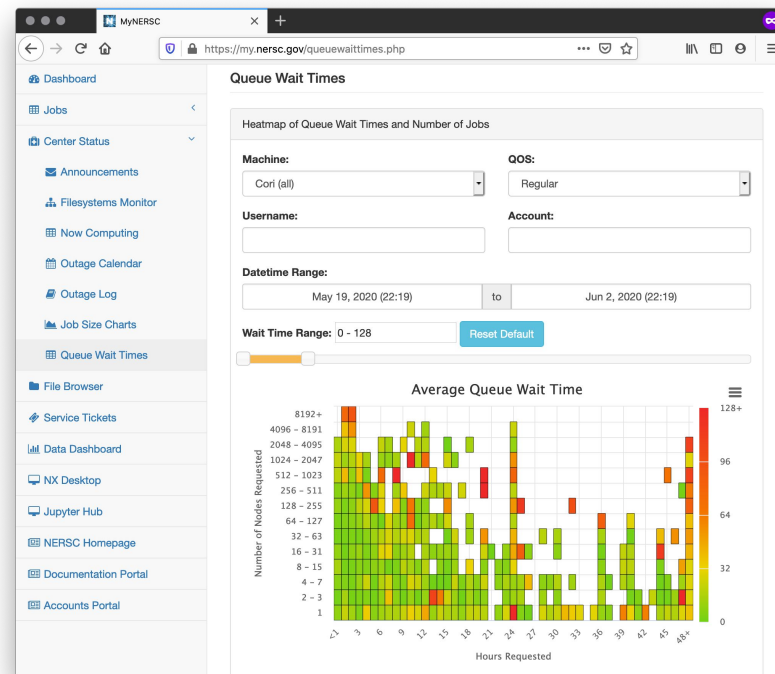
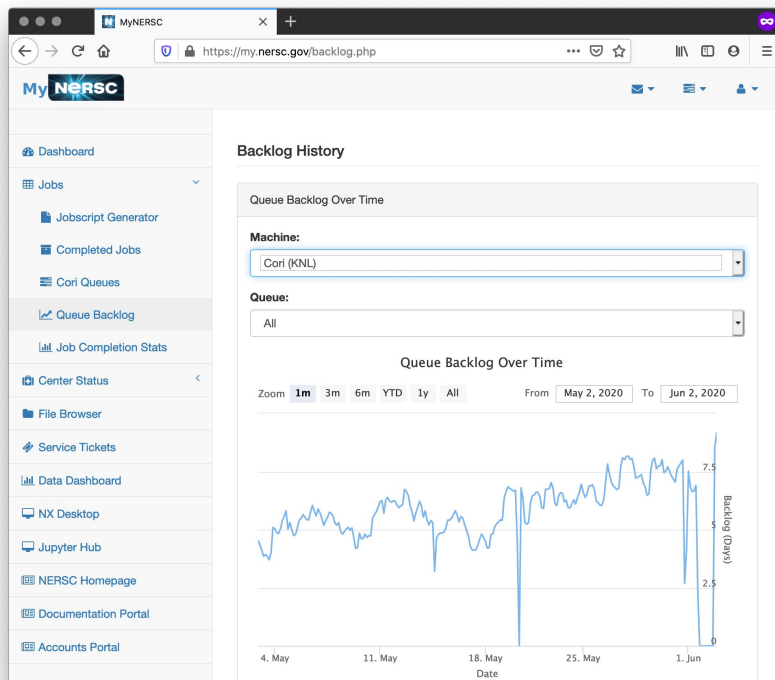
A screenshot of the MyNERSC login page. The page has a white background with a 'Please Sign In' section. It contains three input fields: 'Username', 'Password', and 'MFA Token (if Enabled)'. Below these fields is a green 'Login' button. At the bottom of the login section, there is a link that says 'Forgot your password? Click here.'



A screenshot of the MyNERSC dashboard. The browser address bar shows 'my.nersc.gov'. The 'Sign In' link in the left sidebar is circled in red. The dashboard content includes a 'Dashboard' section with three 'Please Login' buttons under the headings 'My Personal Disk Usage', 'My Active Jobs', and 'My Completed Jobs'. To the right, there is a 'System Status' section with a table of system health:

System Status	
Compute Systems:	
Cori	Down
Perlmutter	Degraded
Global Filesystems:	
Community File System (CFS)	Up
DNA	Up
Data Transfer Nodes	Up
Global Common	Up
Global Homes	Up

<https://my.nersc.gov>



Jobscript Generator

The screenshot shows a web browser window with the URL `my.nersc.gov/script_generator.php`. The page title is "My NERSC" and the main heading is "Jobscript Generator".

Job Information

This tool generates a batch script template which also realizes specific process and thread binding configurations.

Machine
Select the machine on which you want to submit your job.

Application Name
Specify your application including the full path.

Job Name
Specify a name for your job.

Email Address
Specify your email address to get notified when the job enters a certain state.

Quality of Service
Select the QoS you request for your job.

Wallclock Time
Specify the duration of the job. The max walltime for the regular QoS is 48 hours.
 hours minutes seconds

Number of Nodes

Output
Your script will be displayed here.

Navigation Menu:

- Sign In
- Dashboard
- Jobs
 - Jobscript Generator
 - Completed Jobs
 - Cori Queues
 - Queue Backlog
- Center Status
- File Browser
- Service Tickets
- Data Dashboard
- PI Toolbox
- Jupyter Hub
- NERSC Homepage
- Documentation Portal
- Accounts Portal

Connecting to NERSC Systems



Connecting with SSH

"The traditional method"

- For those comfortable working in a terminal, ssh from your local terminal is the most flexible and powerful working environment

You will need a terminal program!

- Mac: terminal (built-in) or "iTerm2" (<https://www.iterm2.com/>)
- Windows: PuTTY (<https://www.putty.org/>), MobaXterm (<https://mobaxterm.mobatek.net/>) or XWin32 or Git BASH
- Linux: Your own favorite :)
- Chromebook: crosh (developer mode) or Crostini (Linux-in-a-container) or SSH App

Connecting to NERSC systems

Connect to NERSC Computational Systems

Please make sure you have configured [Multi-Factor Authentication \(MFA\)](#) prior to login.

To access Perlmutter via `ssh` you can do the following:

```
ssh <user>@perlmutter-p1.nersc.gov
```

or

```
ssh <user>@saul-p1.nersc.gov
```

Similarly, you can access Cori with

```
ssh <user>@cori.nersc.gov
```

Connecting with SSH

```
ssh -l siddiq90 -Y perlmutter-p1.nersc.gov
* ~/ ssh -l siddiq90 -Y perlmutter-p1.nersc.gov
The authenticity of host 'perlmutter-p1.nersc.gov (128.55.126.9)' can't be established.
RSA key fingerprint is SHA256:Db9s2Fa4J3qx7An5oIMgUqUAdK7UWJGTPGoIKD44+Gs.
Are you sure you want to continue connecting (yes/no/[fingerprint])?
```

This means your laptop doesn't recognize the computer. The first time you log in, this is expected. But if your laptop **should** recognize Perlmutter, it's a red flag

The screenshot shows the NERSC website page titled "Connecting to NERSC". The main content area is titled "Key fingerprints" and contains the following information:

NERSC may occasionally update the host keys on the major systems. Check here to confirm the current fingerprints.

Perlmutter

```
4096 SHA256:Db9s2Fa4J3qx7An5oIMgUqUAdK7UWJGTPGoIKD44+Gs perlmutter-p1.nersc.gov [R]
```

Cori

```
4096 SHA256:35yiNfemgwzHCHFrPGWwJBCCqERqLt0VsrR36s1DaPc cori.nersc.gov (RSA)
256 SHA256:Y0ycBUgqcXq5Zi045oG8JKNo9sek07n0C1Xo0MpQZtc cori.nersc.gov (ECDSA)
256 SHA256:/bLLKa0JDbE1rot71v1f+CQC3tFC+e9cKCCObtS+o cori.nersc.gov (ED25519)
```

Connecting with SSH

When you ssh in, you'll see a prompt like:

Password + OTP:

Enter your (iris) password, then the 6 digits from Authenticator, with no spaces etc between eg **Pa\$\$w0rd!123456**

Nothing will appear at prompt as you type! (this is normal)
If you only get "Password: (no "+ OTP)", your account may not be ready yet

```
~/ ssh -l siddiq90 -Y perlmutter-p1.nerisc.gov
Warning: Permanently added the RSA host key for IP address '128.55.126.12' to the list of known hosts.
*****
NOTICE TO USERS
*****
Lawrence Berkeley National Laboratory operates this computer system under
contract to the U.S. Department of Energy. This computer system is the
property of the United States Government and is for authorized use only.
Users (authorized or unauthorized) have no explicit or implicit
expectation of privacy.

Any or all uses of this system and all files on this system may be
intercepted, monitored, recorded, copied, audited, inspected, and disclosed
to authorized site, Department of Energy, and law enforcement personnel,
as well as authorized officials of other agencies, both domestic and foreign.
By using this system, the user consents to such interception, monitoring,
recording, copying, auditing, inspection, and disclosure at the discretion
of authorized site or Department of Energy personnel.

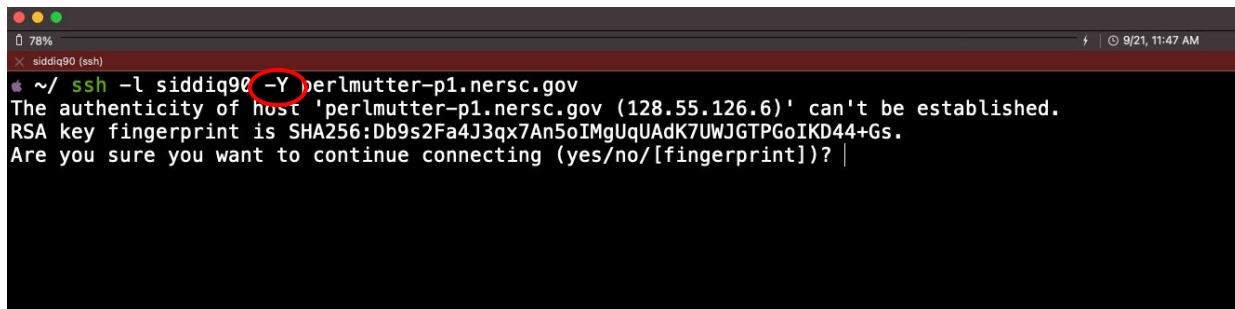
Unauthorized or improper use of this system may result in administrative
disciplinary action and civil and criminal penalties. By continuing to use
this system you indicate your awareness of and consent to these terms and
conditions of use. LOG OFF IMMEDIATELY if you do not agree to the conditions
stated in this warning.

*****

Login connection to host x3116c0s17b0n0:
Password + OTP: ?
```

SSH Options

Wait, what was that "-Y" ?



```
~/ ssh -l siddiq90 -Y perlmutter-p1.nersc.gov
The authenticity of host 'perlmutter-p1.nersc.gov (128.55.126.6)' can't be established.
RSA key fingerprint is SHA256:Db9s2Fa4J3qx7An5oIMgUqUAdK7UWJGTPGoIKD44+Gs.
Are you sure you want to continue connecting (yes/no/[fingerprint])?
```

"ssh -Y" (or "ssh -X")

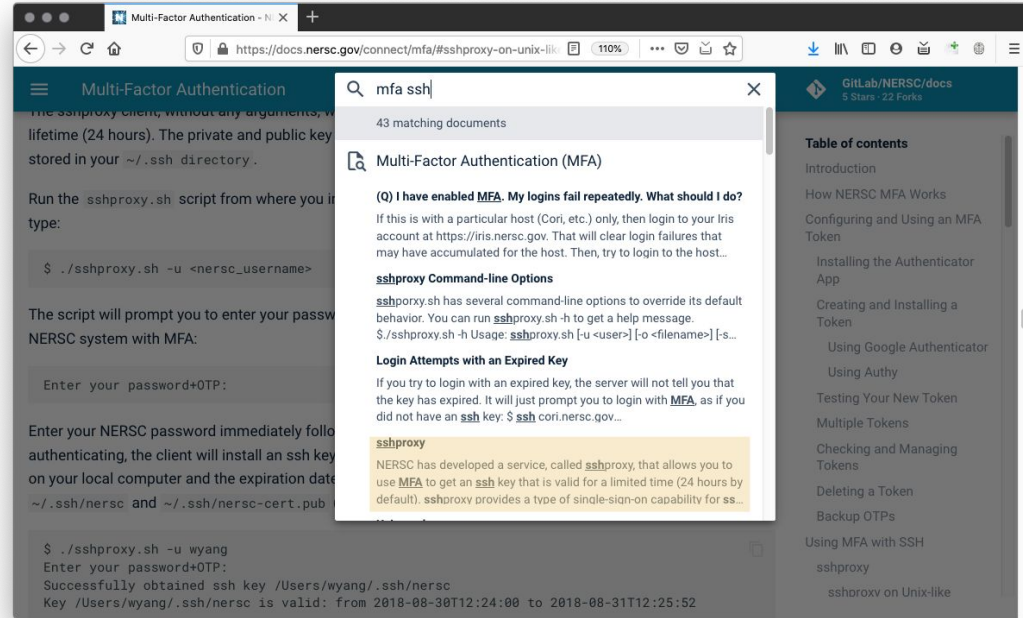
allow X (ie, GUI) programs to display on your local monitor.

- You need an X-server (<https://www.xquartz.org/> for Mac or <http://x.cygwin.com/> for Windows)
- Can be very slow - alternatives coming up!

sshproxy

- Tired of repeatedly typing password + OTP?


- **sshproxy.sh** creates a short-term (24 hours) certificate
- Run **sshproxy.sh** once, then you can ssh to NERSC systems for the next 24 hours before being asked for password+OTP again



- Search "MFA SSH" at <https://docs.nersc.gov>

Jupyter

You can access NERSC systems from any web browser, via <https://jupyter.nersc.gov>




The following NERSC resources that Jupyter depends on appear to be in maintenance or having issues. This may impact Jupyter. See the [NERSC MOTO](#) for further information. Platform status: degraded.

	Shared CPU Node	Shared GPU Node	Exclusive CPU Node	Exclusive GPU Node	Exclusive Large Memory Node	Configurable GPU	Configurable DDX
Alvarez	start						
Muller	start	start	start			start	
Gerty	start						
Perlmutter	start	start	start			start	start
Cori	start	start					

Resources Use a node shared with other users' notebooks but outside the local domain. Use your own node within a job allocation using defaults. Use multiple compute nodes with specialized settings.

Use Cases Visualization and analysis that do not need memory intensive and can run on just a few cores. Visualization, analysis, machine learning that is compute or memory intensive but can be done on a single node. Multi-node analysis jobs, jobs in iterations, custom project changing, and more.

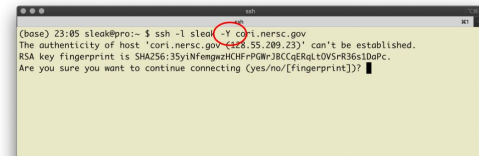


Running GUI Apps

GUI apps eg Matlab, DDT
(debugging), Nsight (performance)
can be painfully slow over a network

Why is this, and how can we fix it?

SSH Options



Wait, what was that "-Y" ?

"ssh -Y" (or "ssh -X")

allow X (ie, GUI) programs on Cori to display on your local monitor.

- You need an X-server (<https://www.xquartz.org/> for Mac or <http://x.cygwin.com/> for Windows)
- **Can be very slow** - alternatives coming up!



NoMachine

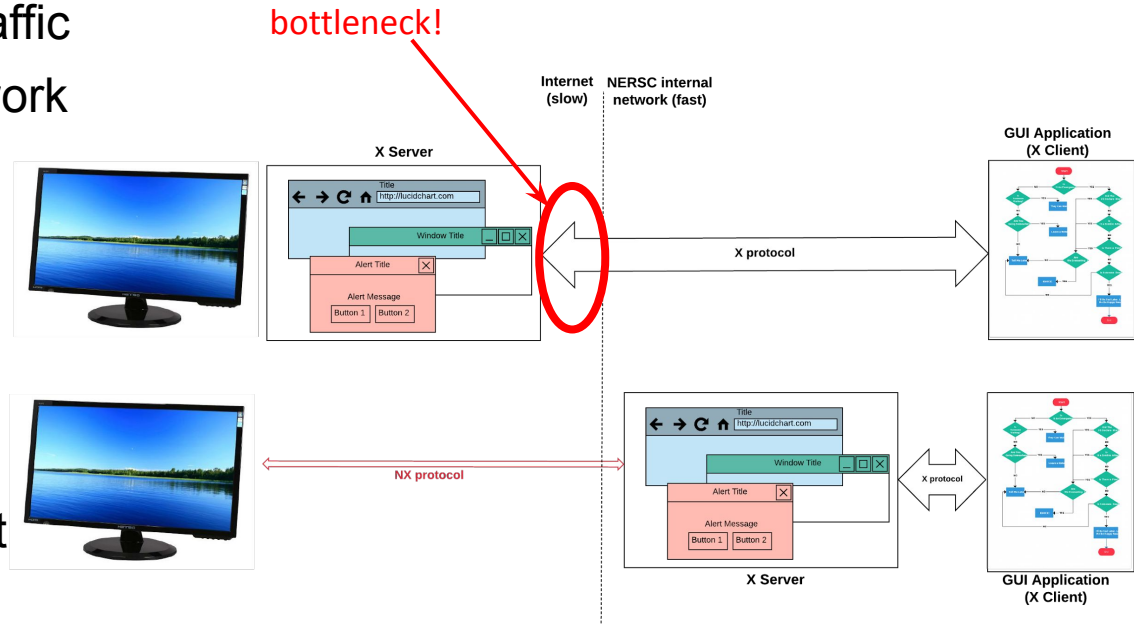


NoMachine: Accelerated X

X protocol makes a lot of traffic

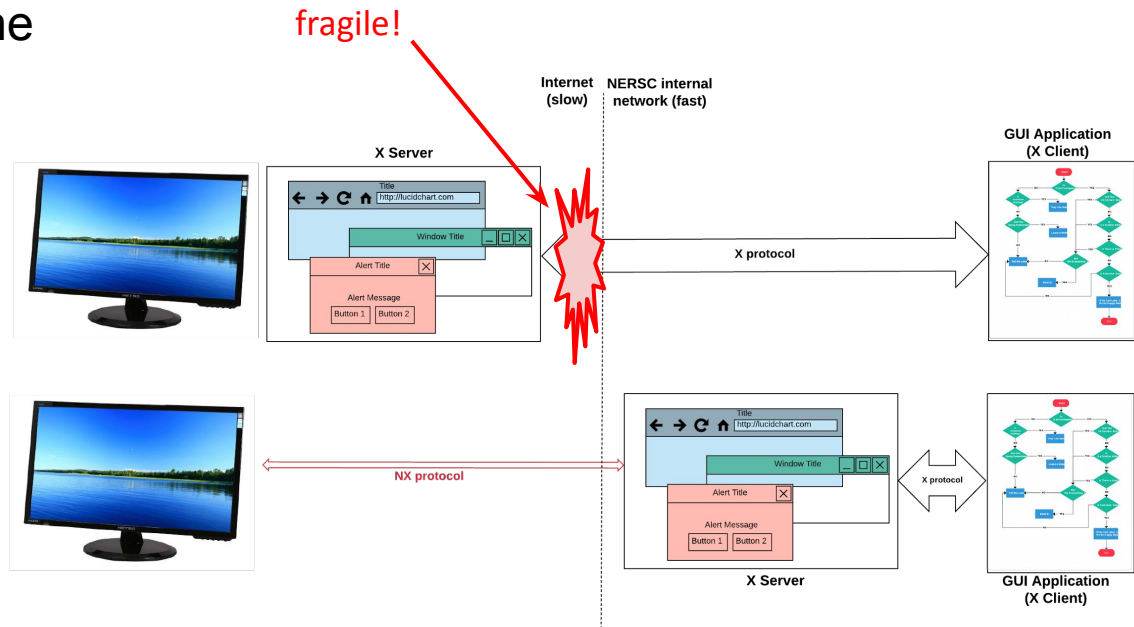
- OK over the (fast) network internal to NERSC
- Not OK over the (slow) internet

NoMachine runs **inside** NERSC, and sends less data over the (slow) internet



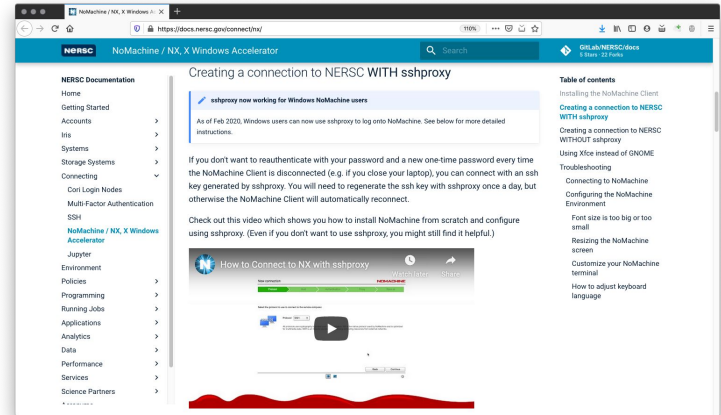
NoMachine: Accelerated X

NoMachine also removes the weakest link, so broken connections don't kill your application




How to Set It Up

- <https://docs.nersc.gov/connect/nx/> has detailed instructions
 - Download the client (<https://www.nomachine.com/download-enterprise#NoMachine-Enterprise-Client>) (Make sure to get the **client**, not the server or workstation)
 - Setup a connection (can optionally use the key you generated with `sshproxy.sh`)



Navigating NERSC Home Page



Navigating www.nersc.gov (NERSC Training)

NERSC Powering Scientific Discovery Since 1974

HOME ABOUT SCIENCE SYSTEMS FOR USERS NEWS R&D EVENTS LIVE STATUS

FOR USERS

- Getting Started
- Accounts & Allocations
- Getting Help
- Live Status
- Documentation
- NERSC Code of Conduct
- Policies
- My NERSC
- Training & Tutorials
- Training Events

NERSC TRAINING

See also the NERSC Events Calendar

Filter by Year

NERSC/NVIDIA AI for Scientific Computing Bootcamp

October 18, 2023

NERSC, in collaboration with the OpenACC organization and NVIDIA, is hosting a virtual, three-day AI for Scientific Computing Bootcamp October 18 - 20, 2023. [Read More](#)

SpinUp Workshop: October 2023

October 18, 2023

Spin is a container-based platform at NERSC designed for you to deploy your own science gateways, workflow managers, databases, API endpoints, and other network services to support your scientific projects. Services in Spin are built with Docker containers and can easily access NERSC systems and storage. Users must apply for and complete the SpinUp instructional workshop before using Spin. Completion of the SpinUp workshop is required for access! October 2023 Location: All... [Read More](#)

GPU Profiling (Performance Profile: Omniperf): Part 5 of HIP Training Series

October 16, 2023

AMD presents a multi-part HIP training series intended to help new and existing GPU programmers understand the main concepts of the HIP programming model. [Read More](#)

OLCF AI Training Series: AI for Science at Scale - Part 2

October 19, 2023

HOME ABOUT SCIENCE SYSTEMS FOR USERS NEWS R&D EVENTS LIVE STATUS STAFF ONLY

FOR USERS

- Getting Started
- Accounts & Allocations
- Getting Help
- Live Status
- Documentation
- NERSC Code of Conduct
- Policies
- My NERSC
- Training & Tutorials

NEW USER TRAINING: SEPTEMBER 7-8, 2023

SEPTEMBER 7, 2023

NERSC is hosting a virtual training event for new users on two half-days on Thursday and Friday, September 7-8, 2023.

The goal is to introduce new users to NERSC

- Computational systems,
- Accounts and allocations,
- Programming environment,
- Running jobs, tools, and best practices, and
- The NERSC data ecosystem.

The training will be focused on *Perlmutter* and using NERSC resources effectively.

Registration

Please use [this form](#) to register and please see the Tentative Agenda for this two half-day virtual event below.

Need Help?

Help Portal help.nersc.gov (requires login)

Accounts Portal iris.nersc.gov Interface (requires login)

Allocations Portal ercap.nersc.gov (requires login)


Code of Conduct

TABLE OF CONTENTS

- Registration
- Day 1 Agenda
- Day 2 Agenda
- Remote Connection Info
- Presentation Materials (Provided After Training)

Back to Top

Navigating www.nersc.gov (NERSC Events)



Powering Scientific Discovery Since 1974

Charles Lively (CharlesLively) | Logout
My NERSC | A-Z Index | Share | Follow

search... >>

HOME ABOUT SCIENCE SYSTEMS FOR USERS NEWS R & D **EVENTS** LIVE STATUS STAFF ONLY

Home » Events » NERSC Events Calendar

NERSC EVENTS CALENDAR

NERSC Training & Events Calendar
Today ◀ ▶ **September 2023** 🖨️ Print

Sun	Mon	Tue	Wed	Thu	Fri	Sat
27	28	29	30	31	Sep 1	2
	10am Training: Po		9am SpinUp: Hact	1pm SpinUp: Hact	10am Spin Office I	
3	4	5	6	7	8	9
		1pm FUN Office H		9am New User Tr	9am New User Tr	10am Spin Office I
10	11	12	13	14	15	16
	9am Perlmutter G				10am Spin Office I	
17	18	19	20	21	22	23
	9am Perlmutter G 10am Training: AN			11am NUG Monthl 11am NUG Monthl 11am NUG Monthl	10am Spin Office I	
24	25	26	27	28	29	30
	9am Perlmutter G	NERSC User Group (NUG) Annual Meeting			10am Spin Office I	

Events shown in time zone: Pacific Time - Los Angeles 📅 Google Calendar

NERSC Training & Events
Today ◀ ▶ **Thursday, September 7** 🖨️ Print

Thursday, September 7
9:00am New User Training, Day 1

Friday, September 8
9:00am New User Training, Day 2
10:00am Spin Office Hours

Monday, September 11

Navigating www.nersc.gov (MOTD)

The screenshot shows the NERSC website interface. At the top, there is a navigation bar with links for HOME, ABOUT, SCIENCE, SYSTEMS, FOR USERS, NEWS, R & D, EVENTS, and LIVE STATUS. The LIVE STATUS menu is expanded, showing options like Now Computing, Queue Look, MOTD, Scheduled Outages, Login Node Status, My NERSC, and Now Computing Highlights. The main content area is titled 'MOTD' and 'Compute Systems:'. It features a table with columns for System, Status, Jobs Running, Queued, Cores in Use, Load, and Description/Notes. Below this, there are sections for 'Global Filesystems:' and 'Mass Storage Systems:', each with their own tables showing system names, statuses, and descriptions.

Compute Systems:

System	Status	Jobs Running	Queued	Cores in Use	Load	Description/Notes
Cori:	Down					09/21/22 7:00-20:00 PDT Scheduled Maintenance
Perlmutter:	Degraded					09/15/22 21:32 PDT - NA. Update by Sep 21 12:00 PDT. System Degraded. Users may see I/O diminished performance and errors accessing files on the scratch file system.

Global Filesystems:

System	Status	Description/Notes
Community File System (CFS):	Up	
DNA:	Up	
Data Transfer Nodes:	Up	
Global Common:	Up	
Global Homes:	Up	

Mass Storage Systems:

System	Status	Description/Notes
HPSS Archive (User):	Up	
HPSS Regent (Backup):	Degraded	09/21/22 9:00-13:00 PDT Scheduled Maintenance. System will remain available during planned maintenance.

Navigating www.nersc.gov (Scheduled System Outages)

NERSC Powering Scientific Discovery Since 1974

Charles Lively (CharlesLively) | Logout
My NERSC | A-Z Index | Share | Follow

search...

HOME ABOUT SCIENCE SYSTEMS FOR USERS NEWS R & D **EVENTS** LIVE STATUS STAFF ONLY

Home » Events » Scheduled System Outages

NERSC SCHEDULED SYSTEM OUTAGES

NERSC Outages

Today Thursday, September 7

Print Week Month Agenda

Wednesday, September 13

- 6:00am Perimutter Scheduled Maintenance
- 9:00am HPSS Regent (Backup) Scheduled Maintenance
- 9:30am Globus Scheduled Maintenance

Tuesday, September 19

- 9:30am Globus Scheduled Maintenance

Wednesday, September 20

- 6:00am Perimutter Scheduled Maintenance

Showing events until 11/15. [Look for more](#)

Events shown in time zone: Pacific Time - Los Angeles

GoogleCalendar

Last edited: 2018-05-02 15:38:06

BERKELEY LAB U.S. DEPARTMENT OF ENERGY Office of Science

- Contact us
- Privacy & Security
- Computing Sciences Area



Navigating www.nerisc.gov (NERSC User Slack)

My NERSC | A-Z Index | Share | Follow

search...

HOME ABOUT SCIENCE SYSTEMS **FOR USERS** NEWS R & D EVENTS LIVE STATUS

FOR USERS

- Getting Help
- Live Status
- Getting Started
- Accounts & Allocations
- Documentation
- Policies
- My NERSC
- Job Logs & Statistics
- Training & Tutorials
- NERSC Users Group
- Monthly NUG Webinars
- Annual Meetings
- SIG Experimental Facility Users
- NUGEX
- NUGEX Nominations
- Charter
- NERSC Users SLACK
- Task Forces/Working Groups

Home » For Users » NERSC Users Group

NERSC USERS GROUP (NUG)

The NERSC Users' Group, NUG, welcomes participation from all NERSC users. NUG provides advice and feedback to NERSC on the current state and future delivery of NERSC resources and services. NUG promotes the effective use of the high performance computing facilities at NERSC by sharing information about experiences in using the facility, suggesting new research and technology directions in scientific computing, and voicing user concerns.

NUG members converse with NERSC and DOE through monthly teleconferences, NUG email lists, and yearly face-to-face meetings.

NERSC holds annual face-to-face meetings.

Monthly NUG Webinars »

This page lists the NERSC User Group (NUG) monthly user telecons and webinars.
[Read More »](#)

NUG Annual Meetings »

NUG holds annual meetings. The annual meetings usually consist of one "business day" and one to three days of High Performance Computing training. [Read More »](#)

SIG Experimental Facility Users »

NUG is sponsoring a Special Interest Group (SIG) within the NERSC Users Group for Experimental Facility Users. [Read More »](#)

NUG Executive Committee (NUGEX) »

NUGEX is the voice of the user community to NERSC and DOE. While all NUG events are open to all NERSC users, NUGEX members regularly participate in the monthly teleconferences and the annual face-to-face meeting. NUGEX is consulted on many NERSC policy issues, e.g., batch configurations, disk quotas, services and training offerings. Members of NUGEX also participate in their office's NERSC Requirements Reviews of High Performance Computing and Storage. There are three representatives from each office and three members-at-large. [Read More »](#)

NUGEX Positions - Now accepting nominations »

We are currently seeking volunteers and nominations for NUGEX - if you would like to participate, or to nominate a potential NUGEX member, please fill in and submit the form at <https://forms.fcl3d/hispd6wlv/77> The Executive Committee (NUGEX) of the NERSC User's Group (NUG) is a group of NERSC users who oversee NUG activities for the benefit of NERSC's user community of over 8,000 researchers across all scientific domains of the DOE Office of Science. NUGEX will meet regularly (up to 1... [Read More »](#)

Need Help?
Help Portal
help.nerisc.gov

Accounts Portal
IRIS Accounts Interface

Allocations Portal
ecsp.nerisc.gov

Charles Lively (CharlesLively) | Logout
My NERSC | A-Z Index | Share | Follow

search...

HOME ABOUT SCIENCE SYSTEMS **FOR USERS** NEWS R & D EVENTS LIVE STATUS STAFF ONLY

Home » For Users » NERSC Users Group » Monthly NUG Webinars

MONTHLY NUG WEBINARS

NUG holds monthly teleconferences with NERSC, usually on the second Thursday of the month, from 11 a.m. to 12 p.m. Pacific Time. All NERSC users, regardless of experience or sophistication, are welcome and encouraged to attend. Connection details are sent monthly via email to all NERSC users.

NUG teleconferences are also listed on the [Events Calendar](#).

NUG Meeting August 24, 2023 »

August 24, 2023
Date: Thursday, August 24, 2023 Time: 11:00 PST The Monthly NUG Meeting is a regular opportunity for our users to show off what they've done, for NERSC to get feedback from users, and for users to exchange ideas. Zoom: <https://lbnl.zoom.us/j/285479463> (full connection details below). We'll also use the NERSC Users Slack #webinars channel for discussion before, during and after the meeting. Add meeting series to calendar/Agenda Announcements and Calls for... [Read More »](#)

NUG Meeting July 20, 2023 »

July 20, 2023
Date: Thursday, July 20, 2023 Time: 11:00 PST The Monthly NUG Meeting is a regular opportunity for our users to show off what they've done, for NERSC to get feedback from users, and for users to exchange ideas. Zoom: <https://lbnl.zoom.us/j/285479463> (full connection details below). We'll also use the NERSC Users Slack #webinars channel for discussion before, during and after the meeting. Add meeting series to calendar/Agenda Announcements and Calls for Participation: Upcoming... [Read More »](#)

NUG Meeting June 15, 2023 »

June 15, 2023
Date: Thursday, June 15, 2023 Time: 11:00 PST The Monthly NUG Meeting is a regular

FOR USERS

- Getting Started
- Accounts & Allocations
- Getting Help
- Live Status
- Documentation
- NERSC Code of Conduct
- Policies
- My NERSC
- Training & Tutorials
- NERSC Users Group
- Monthly NUG Webinars

August 24, 2023 - Best Practices for Reading and Writing Data on Perlmutter

July 20, 2023 - How to Submit a Good Ticket at NERSC

June 15, 2023 - Jupyter at NERSC

May 25, 2023 - SLURM Tips and Tricks at NERSC

April 25, 2023 - Julia at NERSC

March 16, 2023 - Science Highlights

Feb 16, 2023 - Cori Retirement

Jan 19, 2023 - User Community Engagement

Dec 15, 2022 - Allocation Year Transition

Nov 17, 2022 - Migrating from Cori to



Navigating NERSC User Documentation



NERSC Documentation - Main Page



NERSC Documentation
Home
Getting Started
Tutorials
Accounts >
Iris >
Systems >
Storage Systems >
Connecting >
Environment >
Policies >
Development >
Developer Tools >
Running Jobs >
Applications >
Analytics >
Machine Learning >
Performance >
Services >
Science Partners >
Acronyms
Contributed Tips and Tricks
Current Known Issues

NERSC Technical Documentation



Table of contents

- Top documentation pages
- Computing Resources
- Other NERSC web pages

[National Energy Research Scientific Computing \(NERSC\)](#) provides High Performance Computing (HPC) and Storage facilities and support for research sponsored by, and of interest to, the U.S. Department of Energy (DOE) Office of Science (SC).

Top documentation pages

- [Getting Started](#) - Information for new and existing users
- [Getting Help](#) - How to get support
- [Job Queue Policy](#) - Charge factors, run limits, submit limits
- [Example Jobs](#) - Curated example job scripts
- [Jobs overview - Slurm](#) commands, job script basics, submitting, updating jobs
- [Jupyter](#) - Interactive [Jupyter Notebooks](#) at NERSC
- [Globus](#) - High-performance data transfers
- [File permissions](#) - Unix file permissions
- [Multi-Factor Authentication \(MFA\)](#)

Computing Resources

- [Perlmutter](#) - A Cray EX system with AMD EPYC CPUs and NVIDIA A100 GPUs
- [Cori](#) - A Cray XC40 system with Intel Haswell and Intel KNL CPUs



NERSC Documentation - System Overview

NERSC Documentation

Home

Getting Started

Tutorials

Accounts

Iris

Systems

Perlmutter

Cori

Data Transfer Nodes

Storage Systems

Connecting

Environment

Policies

Development

Developer Tools

Running Jobs

Applications

Analytics

Machine Learning

Performance

Services

Science Partners

Acronyms

Contributed Tips and Tricks

Current Known Issues

NERSC Systems

NERSC is one of the largest facilities in the world devoted to providing computational resources for scientific computing.

Perlmutter

[Perlmutter](#) is a HPE (Hewlett Packard Enterprise) Cray EX supercomputer, named in honor of Saul Perlmutter, an astrophysicist at Berkeley Lab who shared the 2011 Nobel Prize in Physics for his contributions to research showing that the expansion of the universe is accelerating.

Perlmutter, based on the HPE Cray Shasta platform, is a heterogeneous system comprising both CPU-only and GPU-accelerated nodes, with a performance of 3-4 times Cori when the installation completes.

We are in the process of Perlmutter Phase 2 integration (adding CPU only nodes and upgrading our system network to Slingshot 11). The final system will consist of 1536 GPU accelerated nodes with 1 AMD Milan processor and 4 NVIDIA A100 GPUs, and 3072 CPU-only nodes with 2 AMD Milan processors. The actual number of nodes available will be in flux during the integration and acceptance of the full system.

Cori (retired)

[Cori](#), a Cray XC40 with a peak performance of about 30 petaflops, was retired on May 31, 2023. The system was named in honor of American biochemist [Gerty Cori](#), the first American woman to win a Nobel Prize and the first woman to be awarded the prize in Physiology or Medicine.

Data transfer nodes

The [data transfer nodes](#) are NERSC servers dedicated to performing transfers between NERSC data storage resources such as HPSS and the NERSC Global File System (NGE), and storage



Table of contents

Perlmutter

Cori (retired)

Data transfer nodes

NERSC Documentation - Storage Overview

The screenshot shows the NERSC Documentation website. The main content area is titled "File System overview" and "Storage System Usage and Characteristics". It includes a "Summary" section and a table of file systems.

NERSC Documentation

- Home
- Getting Started
- Tutorials
- Accounts >
- Iris >
- Systems >
- Storage Systems** >

 - Unix File Permissions
 - Quotas and Purging
 - Perlmutter scratch
 - Cori scratch
 - Community
 - Archive (HPSS) >
 - Global Home
 - Global Common
 - Cori Burst Buffer
 - Project

- Backups
- Connecting >
- Environment >
- Policies >
- Development >
- Developer Tools >

File System overview

Storage System Usage and Characteristics

Summary

File systems are configured for different purposes. Each machine has access to at least three different file systems with different levels of performance, data persistence and available capacity, and each file system is designed to be accessed and used either by a user individually or by their project, as reported in the "Access" column.

File System	Snapshots	Backup	Purging	Access
Home	yes	yes	no	user
Common	no	no	no	project
Community	yes	no	no	project
Cori scratch	no	no	yes	user
Perlmutter scratch	no	no	yes	user
HPSS	no	no	no	user

Table of contents

- Storage System Usage and Characteristics
 - Summary
 - Global storage
 - Home
 - Common
 - Community
 - Scratch
 - Archive (HPSS)
 - Local storage
 - Burst Buffer
 - Temporary per-node Shifter file system
 - Local temporary file system
- Data sharing
 - Sharing Data Inside NERSC
 - Sharing Data Within Your Project
 - Sharing Data Outside Your Project
 - Sharing Data Outside of NERSC

NERSC Documentation - Connecting to NERSC

NERSC NERSC Documentation Search GitHub/NERSC/docs

NERSC Documentation

- Home
- Getting Started
- Tutorials >
- Accounts >
- Iris >
- Systems >
- Storage Systems >
- Connecting** >
 - Multi-Factor Authentication
 - Federated Identity
 - NoMachine / NX, X Windows
 - Accelerator
- Environment >
- Policies >
- Development >
- Developer Tools >
- Running Jobs >
- Applications >
- Analytics >
- Machine Learning >
- Performance >
- Services >
- Science Partners >
- Acronyms
- Contributed Tips and Tricks
- Current Known Issues

Connecting to NERSC

Login Nodes

Opening an [SSH connection](#) to NERSC systems results in a connection to a login node. Typically systems will have multiple login nodes which sit behind a load balancer. New connections will be assigned a random node. If an account has recently connected the load balancer will attempt to connect to the same login node as the previous connection.

Connect to NERSC Computational Systems

Please make sure you have configured [Multi-Factor Authentication \(MFA\)](#) prior to login.

To access Perlmutter via `ssh` you can do the following:

```
ssh <user>@perlmutter.nersc.gov
```

or

```
ssh <user>@saul.nersc.gov
```

If you have configured [sshproxy](#) then you can run the following:

```
ssh -i ~/.ssh/nersc <user>@perlmutter.nersc.gov # or 'ssh -i ~/.ssh/nersc <user>
```

This assumes your identity file is in `~/.ssh/nersc`. The sshproxy route will be convenient if you have multiple ssh connections without having to authenticate every time.

X11 Forwarding

X11 forwarding allows one to display remote computer to your local machine, this can be done as follows:

Table of contents

- Login Nodes
- Connect to NERSC Computational Systems
- X11 Forwarding
- SSH
 - Connecting with SSH
 - Password-less logins and transfers
 - SSH certificate authority
- Key fingerprints
 - Perlmutter
 - DTN[01-04]
 - NoMachine/NX
- Host Keys
 - Perlmutter
- Troubleshooting
 - "Access Denied", "Permission Denied" or "Too many authentication failures"
 - Host authenticity
 - Host identification changed
 - SSH connection disconnects periodically
 - Host key errors
 - Other SSH connection failures

NERSC Documentation - Running Jobs

NERSC NERSC Documentation

Search

GitLab/NERSC/docs

NERSC Documentation

- Home
- Getting Started
- Tutorials >
- Accounts >
- Iris >
- Systems >
- Storage Systems >
- Connecting >
- Environment >
- Policies >
- Development >
- Developer Tools >
- Running Jobs** >
- Scheduling
- Queues and Charges
- Example Jobs
- Best Practices
- Troubleshooting Jobs
- Monitoring
- Affinity
- Interactive
- Reservations
- Workflow Tools >
- Checkpoint/Restart >
- Applications >
- Analytics >
- Machine Learning >
- Performance >
- Services >
- Science Partners >

Running Jobs

NERSC uses [Slurm](#) for cluster/resource management and job scheduling. Slurm is responsible for allocating resources to users, providing a framework for starting, executing and monitoring work on allocated resources and scheduling work for future execution.

Additional Resources

- Documentation: <https://slurm.schedmd.com/documentation.html>
- Tutorial: <https://slurm.schedmd.com/tutorials.html>
- Manual: https://slurm.schedmd.com/man_index.html
- FAQ: <https://slurm.schedmd.com/faq.html>

Jobs

A **job** is an allocation of resources such as compute nodes assigned to a user for an amount of time. Jobs can be interactive or batch (e.g., a script) scheduled for later execution.

Tip

NERSC provides an extensive set of [example job scripts](#)

Once a job is assigned a set of nodes, the user is able to initiate parallel work in the form of job steps (sets of tasks) in any configuration within the allocation.

When you login to a NERSC system you land on a [login node](#). Login nodes are for editing, compiling, or preparing jobs. They are not for running jobs. From the login node you can interact with Slurm to submit job scripts or start interactive jobs.

NERSC's environment is configured to support a diverse workload including high-throughput

Table of contents

- Additional Resources
- Jobs
- Submitting jobs
 - sbatch
 - salloc
 - srun
- Options
 - Commonly Used Options
 - Writing a Job Script
 - Defaults
- Debugging issues
- Available memory for applications on compute nodes
- Quota Enforcement
- Queue Wait Times
- Heterogeneous Jobs
- Further reading about jobs

NERSC Documentation - Programming Models

NERSC NERSC Documentation

Search

GitLab/NERSC/docs

- NERSC Documentation
 - Home
 - Getting Started
 - Tutorials >
 - Accounts >
 - Iris >
 - Systems >
 - Storage Systems >
 - Connecting >
 - Environment >
 - Policies >
 - Development
 - Compilers >
 - Build Tools >
 - Programming Models**
 - MPI >
 - OpenMP >
 - OpenACC
 - CUDA
 - UPC
 - UPC++
 - Coarrays
 - SYCL
 - Kokkos
 - HPX
 - Raja
 - C++ parallel algorithms benchmark
 - Languages >
 - Libraries >
 - Containers >
 - Developer Tools >

Programming Models

A wide variety of programming models are used on NERSC systems. The most common is MPI + OpenMP, but many others are supported.

Parallel programming models at NERSC

Since the transition from vector to distributed memory (MPP) supercomputer architectures, the majority of HPC applications deployed on NERSC resources have evolved to use MPI as their sole means of expressing parallelism. As single processor core compute nodes on MPP architectures gave way to multicore processors, applying the same abstraction (processes passing messages) to each available core remained an attractive alternative - no code changes were required, and vendors made an effort to design optimized fast-paths for on-node communication.

However, as on-node parallelism rapidly increases and competition for shared resources per processing element (memory per core, bandwidth per core, etc.) does as well, now is a good time to assess whether applications can benefit from a different abstraction for expressing on-node parallelism. Examples of desirable functionality potentially available through the latter include more efficient utilization of resources (e.g. through threading) or the ability to exploit unique architectural features (e.g. vectorization).

Perlmutter, and beyond: Performance and portability

Perlmutter has a mixture of CPU-only nodes and CPU + GPU nodes. Each CPU + GPU nodes has 4 GPUs per CPU node.

NERSC has made an effort to provide guidance on parallel programming approaches. Chief among these is the combination of MPI for inter-node parallelism and OpenMP for intra-node parallelism (or potentially MPI per NUMA domain with OpenMP within each).

Table of contents

- Parallel programming models at NERSC
- Perlmutter, and beyond: Performance and portability
- Why MPI + OpenMP?
- Combining Programming Models
 - Distributed memory (inter-node) parallelism
 - Shared memory (intra-node) parallelism
 - CPU
 - GPU

Thank You and
Welcome to
NERSC!

