

Stanford Research Computing

THE SRCC, SHERLOCK &
FARMSHARE

High Performance
Computing Clusters
Physics Department
February 2, 2021





Enabling Research and Discovery

SRCC


Stanford University srcc.stanford.edu

The Stanford Research Computing Center is a joint venture of the Vice Provost and Dean of Research, and the AVP for IT Services.

Our major HPC clusters:

Unless otherwise noted, all environments are available to members of Stanford University and affiliated organizations who have at least a base SUNet (incl. sponsored).

FARM SHARE



SHARED COMPUTE FOR ALL

Free for scholarship and coursework.
14 login nodes with access to AFS.
Five 96 GB (RAM) compute nodes.
Two 767 GB compute nodes.
Ten 128 GB nodes with an nVidia K40.
Shared storage for your group or class.
Full SUNet (or sponsorship) required.
farmshare.stanford.edu

SHERLOCK



COMPUTE FOR RESEARCH

Free for sponsored and departmental research.
130 16-core 64 GB shared nodes.
2 Bigmem (1.5 TB) & 6 GPUs (GTX/K80).
Get access to 700+ nodes by buying compute & GPU nodes for your group!
Storage allocations for each group.
(1 TB long-term + 30 TB fast temp.)
sherlock.stanford.edu

SCG



GENOMICS AT SCALE

2800+ cores and 4+ PB of storage.
350+ software packages ready to use.
Scale out to the cloud with Google Cloud connectivity.
NIH dbGaP compliant.
A charged service of the Genetics Bioinformatics Service Center.
gbsc.stanford.edu

OAK

"CHEAP AND DEEP" STORAGE FOR HPC

\$50/TB a year, billed monthly.
Save 65% by buying 350 TB at once!
Mounted on Sherlock/XStream.
SFTP and Globus support available.
NFSv4 and SMB support in development.
4-year term. Sold in 10 TB increments.
oak-storage.stanford.edu

Depending on the group you're in, you might be able to use one of these clusters:
AHPCRC • BioX3 • GSB • ICME
PHS • SNI • SUGI • XStream
Ask Your PI!

SRCC also provides compute power to:
XSEDE
Open Science Grid

Stanford | Stanford Research Computing Center

Polya Hall
255 Panama Street

srcc.stanford.edu
research-computing-support@stanford.edu

The Stanford Research Computing Center is proud to partner with the following organizations:





Enabling Research and Discovery

SRCC

Stanford University srcc.stanford.edu

THE HIVE

The HANA Immersive Visualization Environment (the HIVE) is our 10- by 24-foot display wall, with 72 million active pixels.

Drive the wall using our Linux, Windows, or macOS systems. Or, bring your own device!

A service of ICME. Free for all Stanford faculty, staff, and students.

Available 8 to 4 weekdays, excluding University holidays. Reservations required.

More information & sign-up: hive.stanford.edu



CONSULTING

We can help you run your code with maximum efficiency. We know MATLAB, Python, R, and more; and can help with parallelization, checkpointing, MPI, and using the GPU. We can help you get the most out of our own clusters; we can refer you to resources at other major academic computing centers.

Free for all Stanford users & affiliates. Access to non-Stanford resources varies.

Reach out to us: research-computing-support@stanford.edu



CLUSTER & LAB SUPPORT

Our experienced system administrators are available to meet your server, cluster, and storage architecture needs. Services we provide include:

- Server and Cluster architecting, provisioning, and administration.
- Network and Storage management. — Workflow and software development.
- Interfacing with other University IT departments.

Charged service available to all Stanford labs & affiliates, including those with High-Risk PHI.

Start a conversation: ruthm@stanford.edu



SINGULARITY

Don't be tied down to just one cluster: Embrace the mobility of compute!

Singularity is containerization—like Docker—but for HPC.

We have experience developing workflows that run in Singularity containers, as well as moving existing workflows into Singularity.

Singularity is open-source, available for anyone to download and use; and is running on all our major clusters.

More information: singularity.lbl.gov



SERVER HOSTING

We can host your servers in the SRCC, our 21k sqft research data center on the SLAC campus.

We provide space, power, and cooling. Multi-10 Gb/sec connections available to campus core, CENIC, Internet2 and the cloud. Dedicated links to the SRN & Internet2's REN available. Or, use dedicated fiber!

Available to campus researchers & affiliates. Hardware age and capacity limitations apply.

More information: srcc.stanford.edu/srcc-hosting



Stanford | Stanford Research Computing Center

Polya Hall
255 Panama Street

srcc.stanford.edu
research-computing-support@stanford.edu

SRCF: Our Data Center

- 21,000 sq. ft. high density secure computing facility.
- 10 and 100GB networked
- ~4,000 servers
- Located on the SLAC campus, leading edge energy conservation features
- 415V power direct to the server cabinets
- 100% outside air with all outside air relieved from the space via natural convection (zero fan energy for outside air exhaust).



What is High Performance Computing (HPC)

“HPC generally refers to the practice of aggregating computing power in a way that delivers much higher performance than one could get from a typical desktop computer or workstation in order to solve large problems.”

– Inside HPC



- **When Will I Need It?**

Almost every field of *research where simulations, large computation or data is needed*: Astrophysics, Social Sciences, Biology, Chemistry, Economics. Most common software run on Sherlock: R, Matlab, and Python

For computing needs above and beyond what your laptop/desktop can handle, in terms of CPUs, RAM, storage and I/O

High Performance Computing Use Cases

- Analyzing Gravitational Lenses 10 million times faster on Sherlock with neural networks on our GPUs, Advancing the automatic prediction of image quality in MRI from unseen sites, Quantum Chemistry
- SRCC acknowledged publications- <https://srcc.stanford.edu/srcc-enabled-publications>



Pros and Cons of High Performance Computing

Pros-

- Your code/calculations are run on servers that are always on, networked and accessible from anywhere by anyone in your PI group (including a PI's off campus collaborators with basic SUNetIDs)
- High performance parallel file systems (Lustre)- fast i/o, 20TB of group Scratch storage on Sherlock
- Much more compute power, hundreds of CPUs, large memory servers up to 3TB of RAM
- Data sharing among research groups is easy, Globus for large data transfers
- Data in home directories are backed up (snapshots) and replicated
- A job scheduler handles problems with hardware, hardware/nodes can fail but jobs do not, launch jobs and log off

Cons-

- Need to learn how to use a job scheduler and Linux
- Jobs go through a scheduler since the system is shared by thousands of users; so you need to wait
- Sometimes you need to request and wait for software installs, you will not have the same permissions to change/modify the system as you do on your own laptop/desktop, however often users can install software themselves on Sherlock/Farmshare

Personal Computers

Mac Book Pro Laptop

- 2 cores (1 CPU)
- 16 GB RAM
- 512GB Solid State Disk



vs.

High Performance Computers

Typical Sherlock Node

- 24-128 CPUs in two sockets Intel 2.4GHz Xeon Skylake and AMD Rome CPUs, up to 256 CPUs can be run at once on Sherlock, 8,192 CPUs for owners.
- 192GB RAM
- 20TB scratch storage, 1TB group home, 100GB local Solid State Disk
- Infiniband connection 200GB/s between nodes and storage (Scratch, Oak and Home)
- GPU nodes (NVIDIA Kepler K80, K40, Volta V100)
- Big memory nodes (512GB, 1.5 and 3TB RAM)
- New 128 CPU AMD nodes



Sherlock

vs.

Farmshare

<http://www.sherlock.stanford.edu>

- Available for use by all Stanford faculty and their research teams
- Work done must be directly related to the faculty member's research.
- A set of servers are free and open to all at Stanford doing sponsored research. Faculty can supplement shared nodes by purchasing additional servers, and become an owner
 - The owning group will have exclusive use of purchased nodes and access to the over 1,100 owner compute nodes when idle.
 - › **1,472 compute nodes**
 - › **30,872 CPU cores**
 - › **1,195 GPUs**
 - › **1,590 TFlops**
 - › **5,400 users**
 - › **833 research groups**
 - › **160 Owners**

<http://farmshare.stanford.edu>

- Primarily intended for use in coursework and unsponsored research
- Anyone with a Stanford ID (SUNet ID) can log in
- Shared servers
- Helps those doing research have a place to test codes and learn about technical solutions to assist in reaching their research goals, prior to scaling up to Sherlock
 - › **20 128 GB RAM compute nodes.**
 - › **4 767 GB RAM compute nodes.**
 - › **10 128 GB RAM nodes with an Nvidia Tesla K40.**
 - › **Shared storage for your group or class.**

Sherlock is approved for Medium Risk. Neither system is approved for use with high-risk data, PHI, or PII.

Stanford University

Oak Storage for HPC


Oak:

A fast I/O storage for HPC, \$50/TB a year, billed monthly. The Oak storage system is mounted on Sherlock. SFTP and Globus support are available. Oak is a capacity-oriented HPC storage system designed for long term storage that doesn't rely on a single vendor implementation.

More info-

<https://oak-storage.stanford.edu/>

Sherlock's user growth since 2016



5,267 users from 860 research groups, **168 owner partitions**

6 PB scratch, 24 PB long term Oak storage

Sherlock facts

as of January 2021



Users

- **5,267** user accounts
- **860** PI groups
from all Stanford's seven Schools, SLAC,
Stanford Institutes, etc.
- **168** owner groups

Interfaces

- **16** login nodes
- **4** data transfer nodes (DTNs)

Computing

- **2.78** PFLOPs
FP64
- **1,697** compute nodes
19 server models (from 3 different
manufacturers)
- **41,504** CPU cores
4 CPU generations (16 CPU models)
- **608** GPUs
6 GPU generations (11 GPU models)
- **44** racks
56 PDUs (470.13 kW of average power)

Storage

- **6.1 PB** \$SCRATCH
parallel, distributed filesystem, delivering
over **200 GB/s** of I/O bandwidth
- **24.8 PB** \$OAK
long term research data storage

Networking

- **120** Infiniband switches
across 3 Infiniband fabrics (EDR, FDR, HDR)
- **5,675** Infiniband cables
spanning about **29.67 km**
- **56** Ethernet switches

Scheduler

- **137** Slurm partitions
- **114,994,814** CPU.hours
used in the last 6 months, that's over
13,127 years of computing!

Sherlock hns partition

Open to all users from the School of Humanities and Sciences (Physics, Statistics, Sociology, Communication, Economics Departments, etc.)

89 Compute Nodes (servers)

2,368 CPUs

2 Large memory nodes .5 TB RAM

4 128 CPU nodes with 1TB RAM

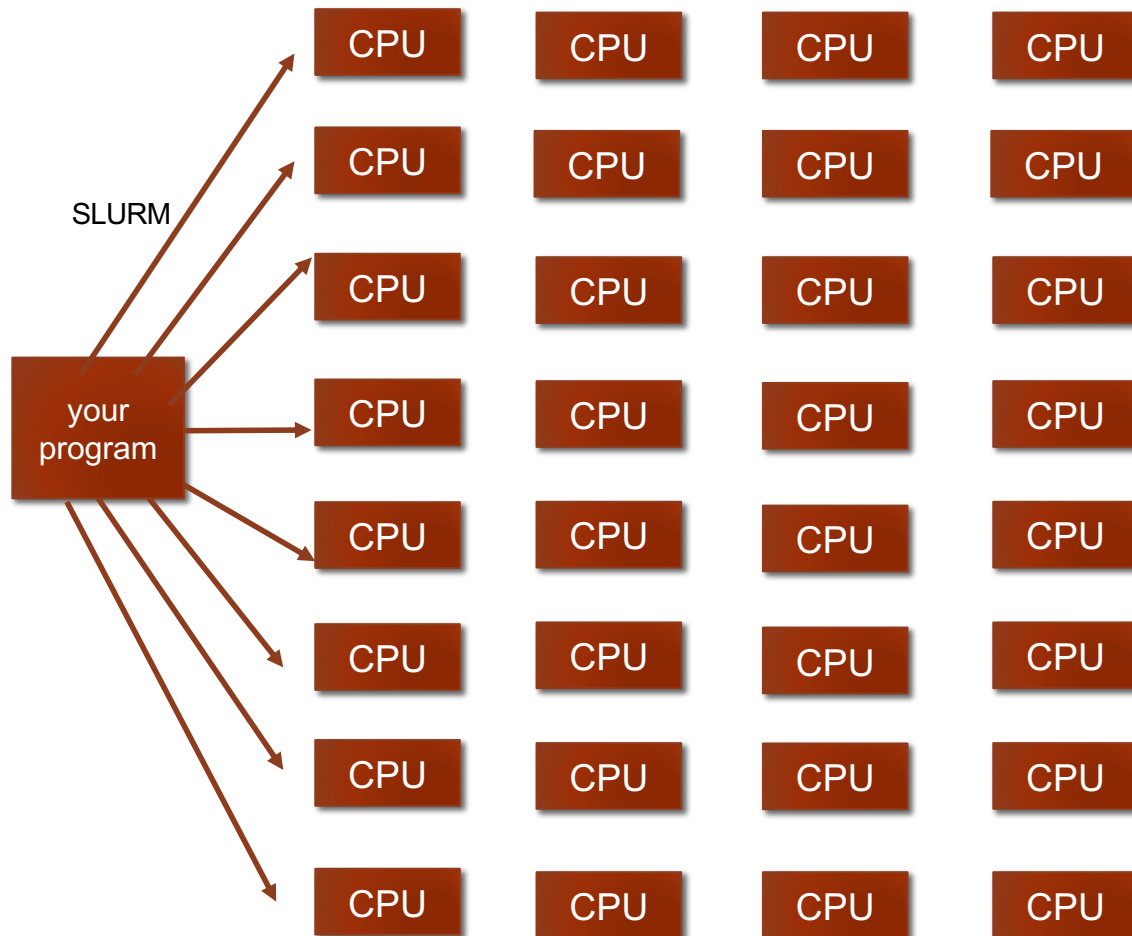
1 GPU node

To use add

```
#SBATCH -p hns
```

to your submission scripts

Parallel Processing



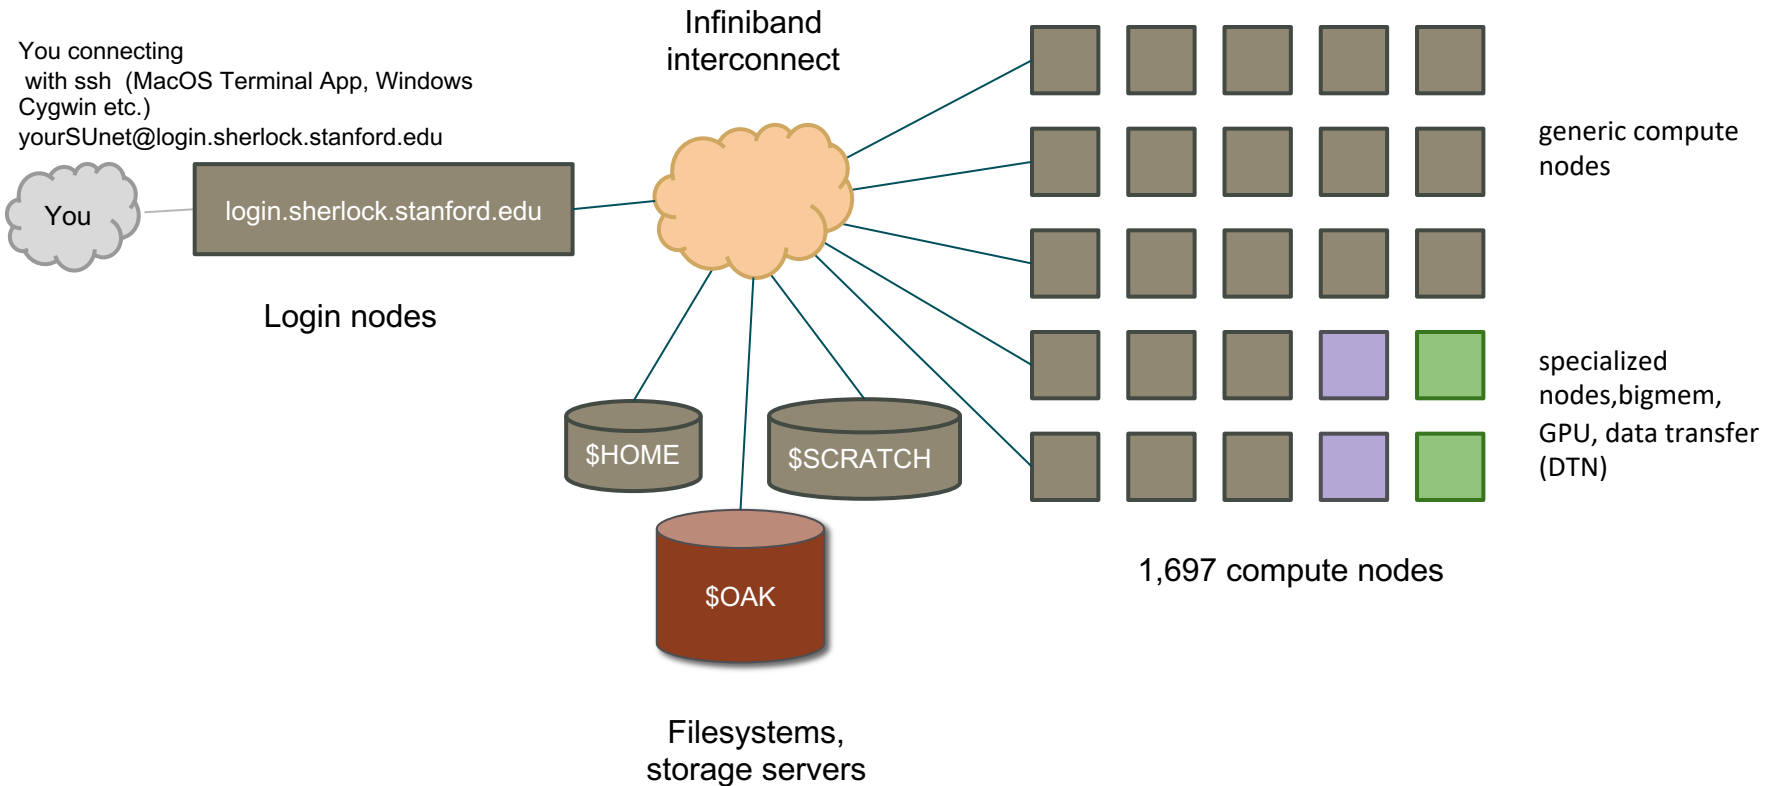
On a cluster multiple tasks can be submitted via a job scheduler to many CPUs and servers at once

Pass multiple arguments to your code at once

Also, multiple instances of your code can run simultaneously across the cluster

No need to wait for 2 or 3 cores and limited RAM on a laptop or desktop to be available

Sherlock System Overview

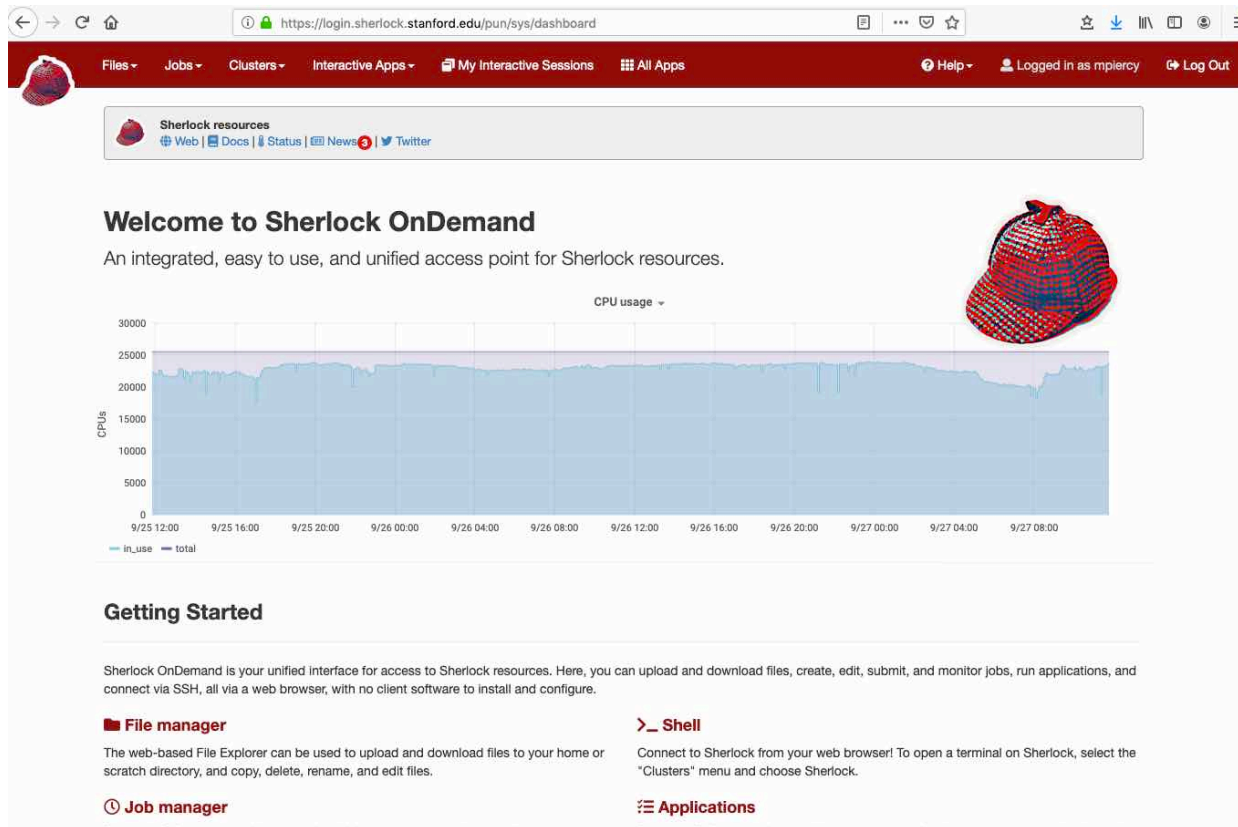


Web browser access to Sherlock

Access Sherlock with your browser - <https://login.sherlock.stanford.edu/>

- open a shell session
- move, copy files from your laptop to Sherlock
- edit files

OnDemand [Documentation](#)



The screenshot shows the Sherlock OnDemand web interface. At the top, there is a navigation bar with a Sherlock logo and menu items: Files, Jobs, Clusters, Interactive Apps, My Interactive Sessions, and All Apps. On the right side of the navigation bar, there are links for Help, Logged in as mpiercy, and Log Out. Below the navigation bar, there is a section titled "Sherlock resources" with links for Web, Docs, Status, News, and Twitter. The main content area features a "Welcome to Sherlock OnDemand" message, followed by a description: "An integrated, easy to use, and unified access point for Sherlock resources." To the right of the text is a 3D wireframe model of a Sherlock hat. Below the text is a "CPU usage" graph showing CPU usage over time from 9/25 12:00 to 9/27 08:00. The graph has a y-axis labeled "CPUs" ranging from 0 to 30000. The legend indicates "in_use" (blue) and "total" (purple). Below the graph is a "Getting Started" section with three sub-sections: "File manager", "Job manager", and "Shell".

File manager

The web-based File Explorer can be used to upload and download files to your home or scratch directory, and copy, delete, rename, and edit files.

Job manager

Check out OnDemand's powerful job management features for your applications.

Shell

Connect to Sherlock from your web browser! To open a terminal on Sherlock, select the "Clusters" menu and choose Sherlock.

Applications

Check out OnDemand's powerful features for running applications on Sherlock.

Scheduling Jobs

Why Do We Need to Schedule a Job?

Some jobs take hours or days to run. So, the compute resources are managed and workloads are balanced using a Scheduler.

How Easy Is It to Schedule a Job?

Basic concept - tell the scheduler:

1. What resources you need- CPUs, RAM, time, partition
2. What it should do- load modules, run your code

Common software pre-installed on Sherlock

We take care of a lot of installations-

Matlab

R

Python

Stata-mp

K-nitro

Gurobi

Sherlock provides 617 software packages, in 7 categories, covering 70 fields of science

A [complete listing](#) of all modules on Sherlock

Search for modules you need with **module spider**

All stored as modules-

\$module avail

--- math -- numerical libraries, statistics, deep-learning, computer science ---

R/3.4.0		py-keras/2.3.1_py36	(g)
R/3.5.1	(D)	py-numpy/1.14.3_py27	(D)
R/3.6.1		py-numpy/1.14.3_py36	
armadillo/8.200.1		py-numpy/1.17.2_py36	
arpack/3.5.0		py-numpy/1.18.1_py36	
caffe2/0.8.1	(g)	py-onnx/1.0.1_py27	
cgal/4.10		py-pytorch/0.2.0_py27	(g)
cuda/5.1	(g)	py-pytorch/0.2.0_py36	(g)
cuda/6.0	(g)	py-pytorch/0.3.0_py27	(g,D)
cuda/7.0.1	(g)	py-pytorch/0.3.0_py36	(g)
cuda/7.0.4	(g)	py-pytorch/1.0.0_py27	(g)
cuda/7.0.5	(g)	py-pytorch/1.0.0_py36	(g)

Preinstalled physics software on Sherlock

----- physics -- astronomy, astrophysics, HEP, geoscience -----

atat/3.36	meep/1.4.3	py-astropy/4.0.1_py36
cdo/1.9.7.1	mpb/1.5	py-lenstools/1.0_py36
gdal/2.2.1	mpb/1.6.2	py-opendrift/1.0.3_py27
geos/3.6.2	oommf/1.2b4	py-pyproj/1.9.5.1_py27
heasoft/6.22.1	openmc/0.10.0	py-pyproj/1.9.5.1_py36
heasoft/6.26.1	opensees/2.5.0	su2/7.0.3
meep/1.3	proj/4.9.3	udunits/2.2.26

Detailed listing -

<https://www.sherlock.stanford.edu/docs/software/list/#physics>

How Can I Obtain Access to HPC?

Sherlock

Your PI or Faculty Sponsor needs to request an account - <http://www.sherlock.stanford.edu/>

Becoming a Sherlock owner

Faculty can purchase compute nodes on Sherlock for a 4-year fixed term. Members of their team will have exclusive use and have access to idle nodes in addition to multi-terabyte scale storage.

Node orders- <https://srcc.stanford.edu/private/sherlock-orders>

Farmshare

Everyone with a SUNet ID can access:

Log into rice.stanford.edu using a SSH client. If you have a Mac, open the Terminal application and type: `ssh [sunset]@rice.stanford.edu`

Nero

For high risk data computing environment

Both Cloud and on premises option available

<http://med.stanford.edu/nero.html>

Stanford Data Risk Classifications

<https://uit.stanford.edu/guide/riskclassifications>

Low Risk

Data and systems are classified as Low Risk if they are not considered to be Moderate or High Risk, and:

1. The data is intended for public disclosure, or
2. The loss of confidentiality, integrity, or availability of the data or system would have no adverse impact on our mission, safety, finances, or reputation.

Moderate Risk

Data and systems are classified as Moderate Risk if they are not considered to be High Risk, and:

1. The data is not generally available to the public, or
2. The loss of confidentiality, integrity, or availability of the data or system could have a mildly adverse impact on our mission, safety, finances, or reputation.

High Risk

Data and systems are classified as High Risk if:

1. Protection of the data is required by law/regulation,
2. Stanford is required to self-report to the government and/or provide notice to the individual if the data is inappropriately accessed, or
3. The loss of confidentiality, integrity, or availability of the data or system could have a significant adverse impact on our mission, safety, finances, or reputation.

Want to Display Your Work? Welcome to HIVE



10-foot-tall by 24-foot-wide display with 13440x5400 resolution and 72 million total active pixels. Located in Huang B050

Use multiple displays simultaneously to investigate various aspects of data collection, simulation, and visualization and to zoom in to see detail at previously unheard-of levels. Free to use for teaching, research by any groups at Stanford.

<https://icme.stanford.edu/resources/hive> Stanford University

To Learn More

Documentation

Sherlock: <http://www.sherlock.stanford.edu/>

Farmshare: <http://farmshare.stanford.edu>

Nero: <http://med.stanford.edu/nero.html>

Office Hours

On Zoom

Tuesdays 10:00 - 11:00 am and Thursdays 3:00pm - 4:00pm

Zoom Link-

<https://stanford.zoom.us/j/95962823750?pwd=cFM2U2ZRQ243Zkx0Ry83akdtWU9zUT09>

<https://www.sherlock.stanford.edu/docs/overview/introduction/ - office-hours>

Sherlock On-boarding Slides-

https://srcc.stanford.edu/sites/g/files/sbiybj10261/f/sherlock_onboarding_6.pdf

Contact

Questions/Answers: srcc-support@stanford.edu

SRCC group: <http://srcc.stanford.edu>

Mark Piercy: mpiercy@stanford.edu