

# HIGH PERFORMANCE RESEARCH COMPUTING

Introduction to HPRC, Duo,  
VPN and Clusters



TEXAS A&M UNIVERSITY

Division of Research



# Upcoming HPRC Primers

<https://hprc.tamu.edu/training/>

## Running Jupyter Notebook on HPRC Open on Demand (OOD)

Wednesday, May 27 @ 10am & 2pm

- How to launch the Portal
- How to launch the Notebook in the Portal
- How to use a Jupyter Notebook

## Introduction to Linux using Mobaxterm

Thursday, May 28 @ 2pm

- Basic commands
- Process and filesystem concepts
- Shell scripts

## Introduction to Linux using the HPRC Portal

Friday, May 29 @ 2pm

- How to launch the Portal
- Navigating through the Portal
- Creating and submitting batch jobs from the Portal

## Using the Ada Cluster

Tuesday, June 2 @ 2pm

- Basic system information
- Compiling and running programs
- Creating and submitting LSF batch jobs

## Using the LSF Job Scheduler on Ada Cluster

Wednesday, June 3 @ 2pm

- Building job files
- Job submission on Terra
- Queues, wait times and job management

## Using the Terra Cluster

Thursday, June 4 @ 2pm

- Basic system information
- Compiling and running programs
- Creating and submitting SLURM batch jobs

## Using the SLURM Job Scheduler on Terra Cluster

Friday, June 5 @ 2pm

- Building job files
- Job submission on Terra
- Queues, wait times and job management

## Managing Data on the Clusters

Tuesday, June 9 @ 2pm

- Transferring data to/from Ada/Terra
- Data lifecycle on Ada/Terra
- Data policies on Ada/Terra





<https://hprc.tamu.edu>



### Quick Links


- New User Information
- Accounts
  - Apply for Accounts
  - Manage Accounts
- User Consulting
- Training
- Documentation
- Software
- FAQ

### User Guides

- Ada
- Terra
- Curie
- Portal
- Galaxy

### Cluster Status

Terra	
Nodes	253/317 (80%)
Cores	5990/9428 (64%)



Effect of methylation on local mechanics and hydration structure of DNA by *Xiaojing Teng and Wonmuk Hwang*, Department of Biomedical Engineering, Texas A&M.

### News

**OCT 11** [Texas A&M To Lead Multi-Institution Network Supporting Cyber Expertise](#)

### Events

**NOV 8** [Short Course: Introduction to Code Parallelization using MPI](#)






# Terra: A Lenovo x86 Cluster

- A 8,512-core hybrid system with 48 NVIDIA K80 GPUs
- 304 28-core compute nodes equipped with the INTEL 14-core 2.4GHz Broadwell processor
- 48 nodes have 1 K80 GPU with 128GB memory each
- Interconnect fabric is Intel OmniPath Architecture (OPA)



# Ada: An IBM/Lenovo x86 Cluster

- 17,340-core hybrid system
- 845 20-core nodes equipped with the INTEL 10-core 2.5GHz IvyBridge processor.
- 15 nodes are 1TB and 2TB memory, 4-processor SMPs configured with the Intel 10-core 2.7GHz Westmere processors
- 30 nodes have 2 NVIDIA K20 GPUs
- 4 nodes have dual V100 GPUs
- 9 have 2 Phi coprocessors.



# Curie: 50-node POWER7+ cluster




For Special HPC Codes & Applications

- 16-core ( two 8-core Power7+ processors) 4.2GHz; 537 GFLOPS; 32 nm fabrication
- 256 GB memory; 68GB/s per processor socket; L3 80MB/processor
- 10GbE port (link to Wehner Core 10G/40G switch)
- 4 x 600GB 10K rpm SAS (local disk)
- REDHAT ENTERPRISE LINUX FOR PWR ; GPFS Client; IBM Compilers, ESSL, LSF (batch)
- Target areas: applications & codes requiring fast memory and fast cpus

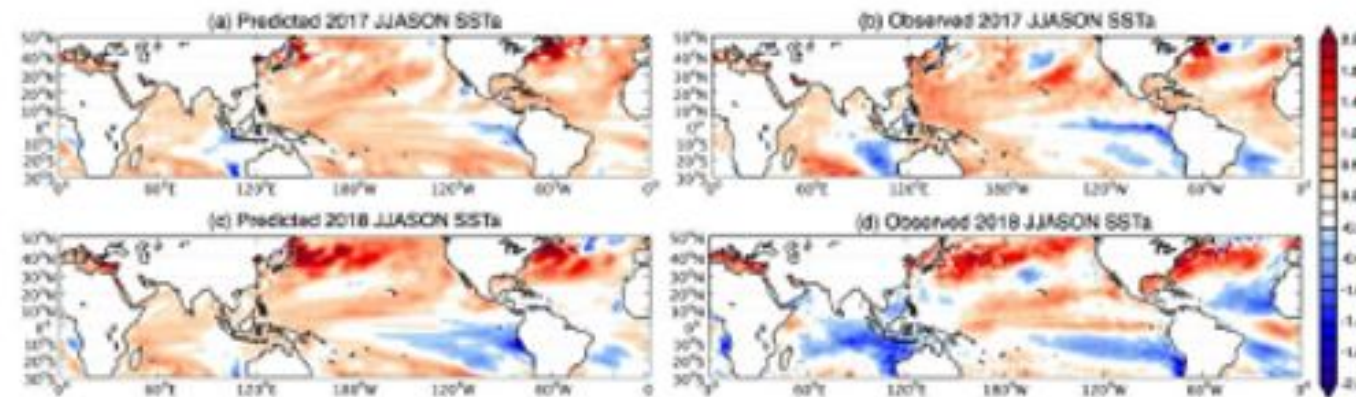




<https://hprc.tamu.edu/events/conferences/sc19/>

research@TEXAS  
A&M



SC19  
Denver, CO | hpc  
is now.





NETID - username you use to login to various systems and services  
(howdy.tamu.edu)

For help with your netid: <https://it.tamu.edu/help/>

**ATM** | **TEXAS A&M**  
UNIVERSITY.

Facebook Twitter Instagram Snapchat Pinterest YouTube LinkedIn

# Howdy

Connecting you to Texas A&M

🔌 Log in with your NetID to get Started

- Log In Information
- Information About Applicant Access
- Information About Parent/Guardian Access
- Information About Former Student Access
- Certified Electronic Diploma [CeDiploma]

**Statement of Responsibility (Acceptable Use)**  
It is a violation of university regulations, state laws, and federal laws if you gain or help others gain unauthorized access to the Howdy portal. If unauthorized access or use is gained, you will be subject to university disciplinary action and criminal prosecution, to the fullest extent of the law. By logging into Howdy, you are responsible for strictly





VPN (Virtual Private Network) [connect.tamu.edu](https://connect.tamu.edu)  
Required to access the university's network from off campus.  
For help with VPN: <https://it.tamu.edu/help/>



## Texas A&M SSL VPN Service

DIVISION OF INFORMATION TECHNOLOGY

VPN users are required to use Duo NetID Two-Factor Authentication. If you use the Duo phone call feature, additional steps are required.  
Once VPN is downloaded, type **connect.tamu.edu** in the empty box. Click Connect.

## Log In

Group:

NetID:

Password:

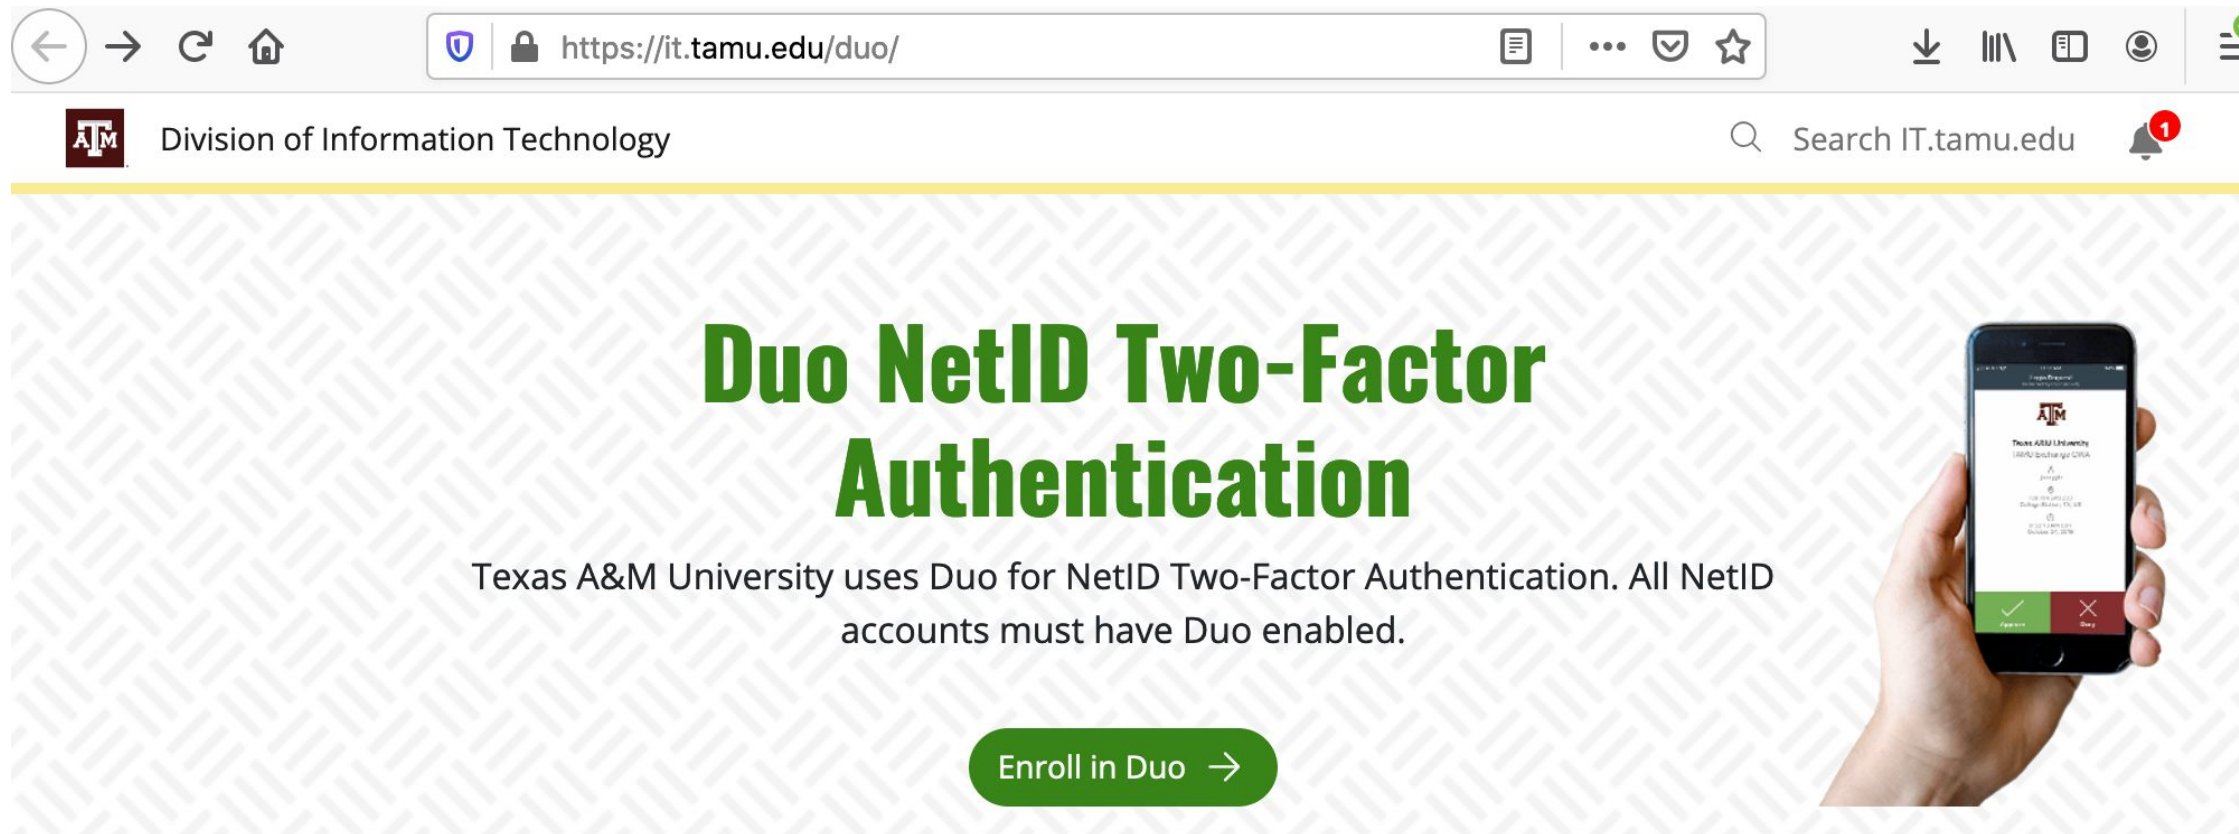
Log In



# Two-Factor Authentication (DUO)

For help with enabling DUO: <https://it.tamu.edu/help/>

- Duo NetID two-factor authentication to enhance security ([it.tamu.edu/duo/](https://it.tamu.edu/duo/))
  - All web login (howdy, portal.hprc.tamu.edu, Globus) through CAS
  - VPN to TAMU campus ([connect.tamu.edu](https://connect.tamu.edu))
  - SSH/SFTP to HPRC clusters ([hprc.tamu.edu/wiki/Two\\_Factor](https://hprc.tamu.edu/wiki/Two_Factor))



The image shows a browser window displaying the Duo NetID Two-Factor Authentication page. The browser's address bar shows the URL <https://it.tamu.edu/duo/>. The page header includes the TAMU logo and the text "Division of Information Technology" on the left, and a search bar with "Search IT.tamu.edu" and a notification icon on the right. The main content area features a large green heading "Duo NetID Two-Factor Authentication" and a sub-heading "Texas A&M University uses Duo for NetID Two-Factor Authentication. All NetID accounts must have Duo enabled." Below this is a green button labeled "Enroll in Duo" with a right-pointing arrow. On the right side of the page, there is an image of a hand holding a smartphone displaying the Duo mobile app interface, which includes the TAMU logo and a "Log In" button.





<https://hprc.tamu.edu>



### Quick Links

- New User Information
- Accounts
  - Apply for Accounts
  - Manage Accounts
- User Consulting
- Training
- Documentation
- Software
- FAQ

### User Guides

- Ada
- Terra
- Curie
- Portal
- Galaxy

### Cluster Status

Terra	
Nodes	253/317 (80%)
Cores	5990/9428 (64%)



### News

**OCT 11** [Texas A&M To Lead Multi-Institution Network Supporting Cyber Expertise](#)

### Events

**NOV 8** [Short Course: Introduction to Code Parallelization using MPI](#)





Texas A&M HPRC

19 subscribers

# YouTube training videos

SUBSCRIBED



HOME

VIDEOS

PLAYLISTS

CHANNELS


DISCUSSION

ABOUT



Uploads PLAY ALL

SORT BY




3:26

Using Open On Demand: Intro

5 views • 3 days ago

CC



3:26

Using Open On Demand: Jobs

6 views • 3 days ago

CC



1:16:17

HPRC Primers: Ada

31 views • 1 week ago

CC



3:03

HPRC Intro: #1 Applying for Accounts

61 views • 1 month ago

CC



5:38

HPRC Intro: #7 Submitting a Job File on Ada/Curie

39 views • 1 month ago

CC



4:28

HPRC intro: #0 What is HPRC?

63 views • 1 month ago

CC




5:13

HPRC Intro: #4 File Management on the Clusters

28 views • 1 month ago

CC



5:32

HPRC Intro #3: Accessing Clusters from Windows...

26 views • 1 month ago

CC



2:51

HPRC Intro: #5 Managing Allocations

27 views • 1 month ago

CC



1:18

HPRC Intro: #2 Cluster Access Using SSH

28 views • 1 month ago

CC



Texas A&M University

High Performance Research Computing

<https://hprc.tamu.edu>



# HPRC Account Allocations

<https://hprc.tamu.edu/policies/allocations.html>

Allocation Type	Who can apply?	Minimum SUs per Allocation per Machine	Maximum SUs per Allocation per Machine	Maximum Total SUs per Machine	Maximum Number of Allocations per Machine	Allowed to spend more than allocation?	Reviewed and approved by
Basic	Faculty, Post-Docs*, Research Associates, Research Scientists, Qualified Staff, Students*, Visiting Scholars/Students*	5,000	5,000	5,000	1	No	HPRC Staff
Startup	Faculty, Research Associates, Research Scientists, Qualified Staff	5,000	200,000	400,000	2	No	HPRC Director
Research (Ada)	Faculty, Research Scientists, Qualified Staff	300,000	8,000,000	8,000,000	Determined by <u>HPRC-RAC</u>	No	<u>HPRC-RAC</u>
Research (Terra)	Faculty, Research Scientists, Qualified Staff	300,000	5,000,000	5,000,000	Determined by <u>HPRC-RAC</u>	No	<u>HPRC-RAC</u>

Note: '\*' - requires a PI



# Examples of **SUs** charged based on Job Cores, Time and Memory Requested

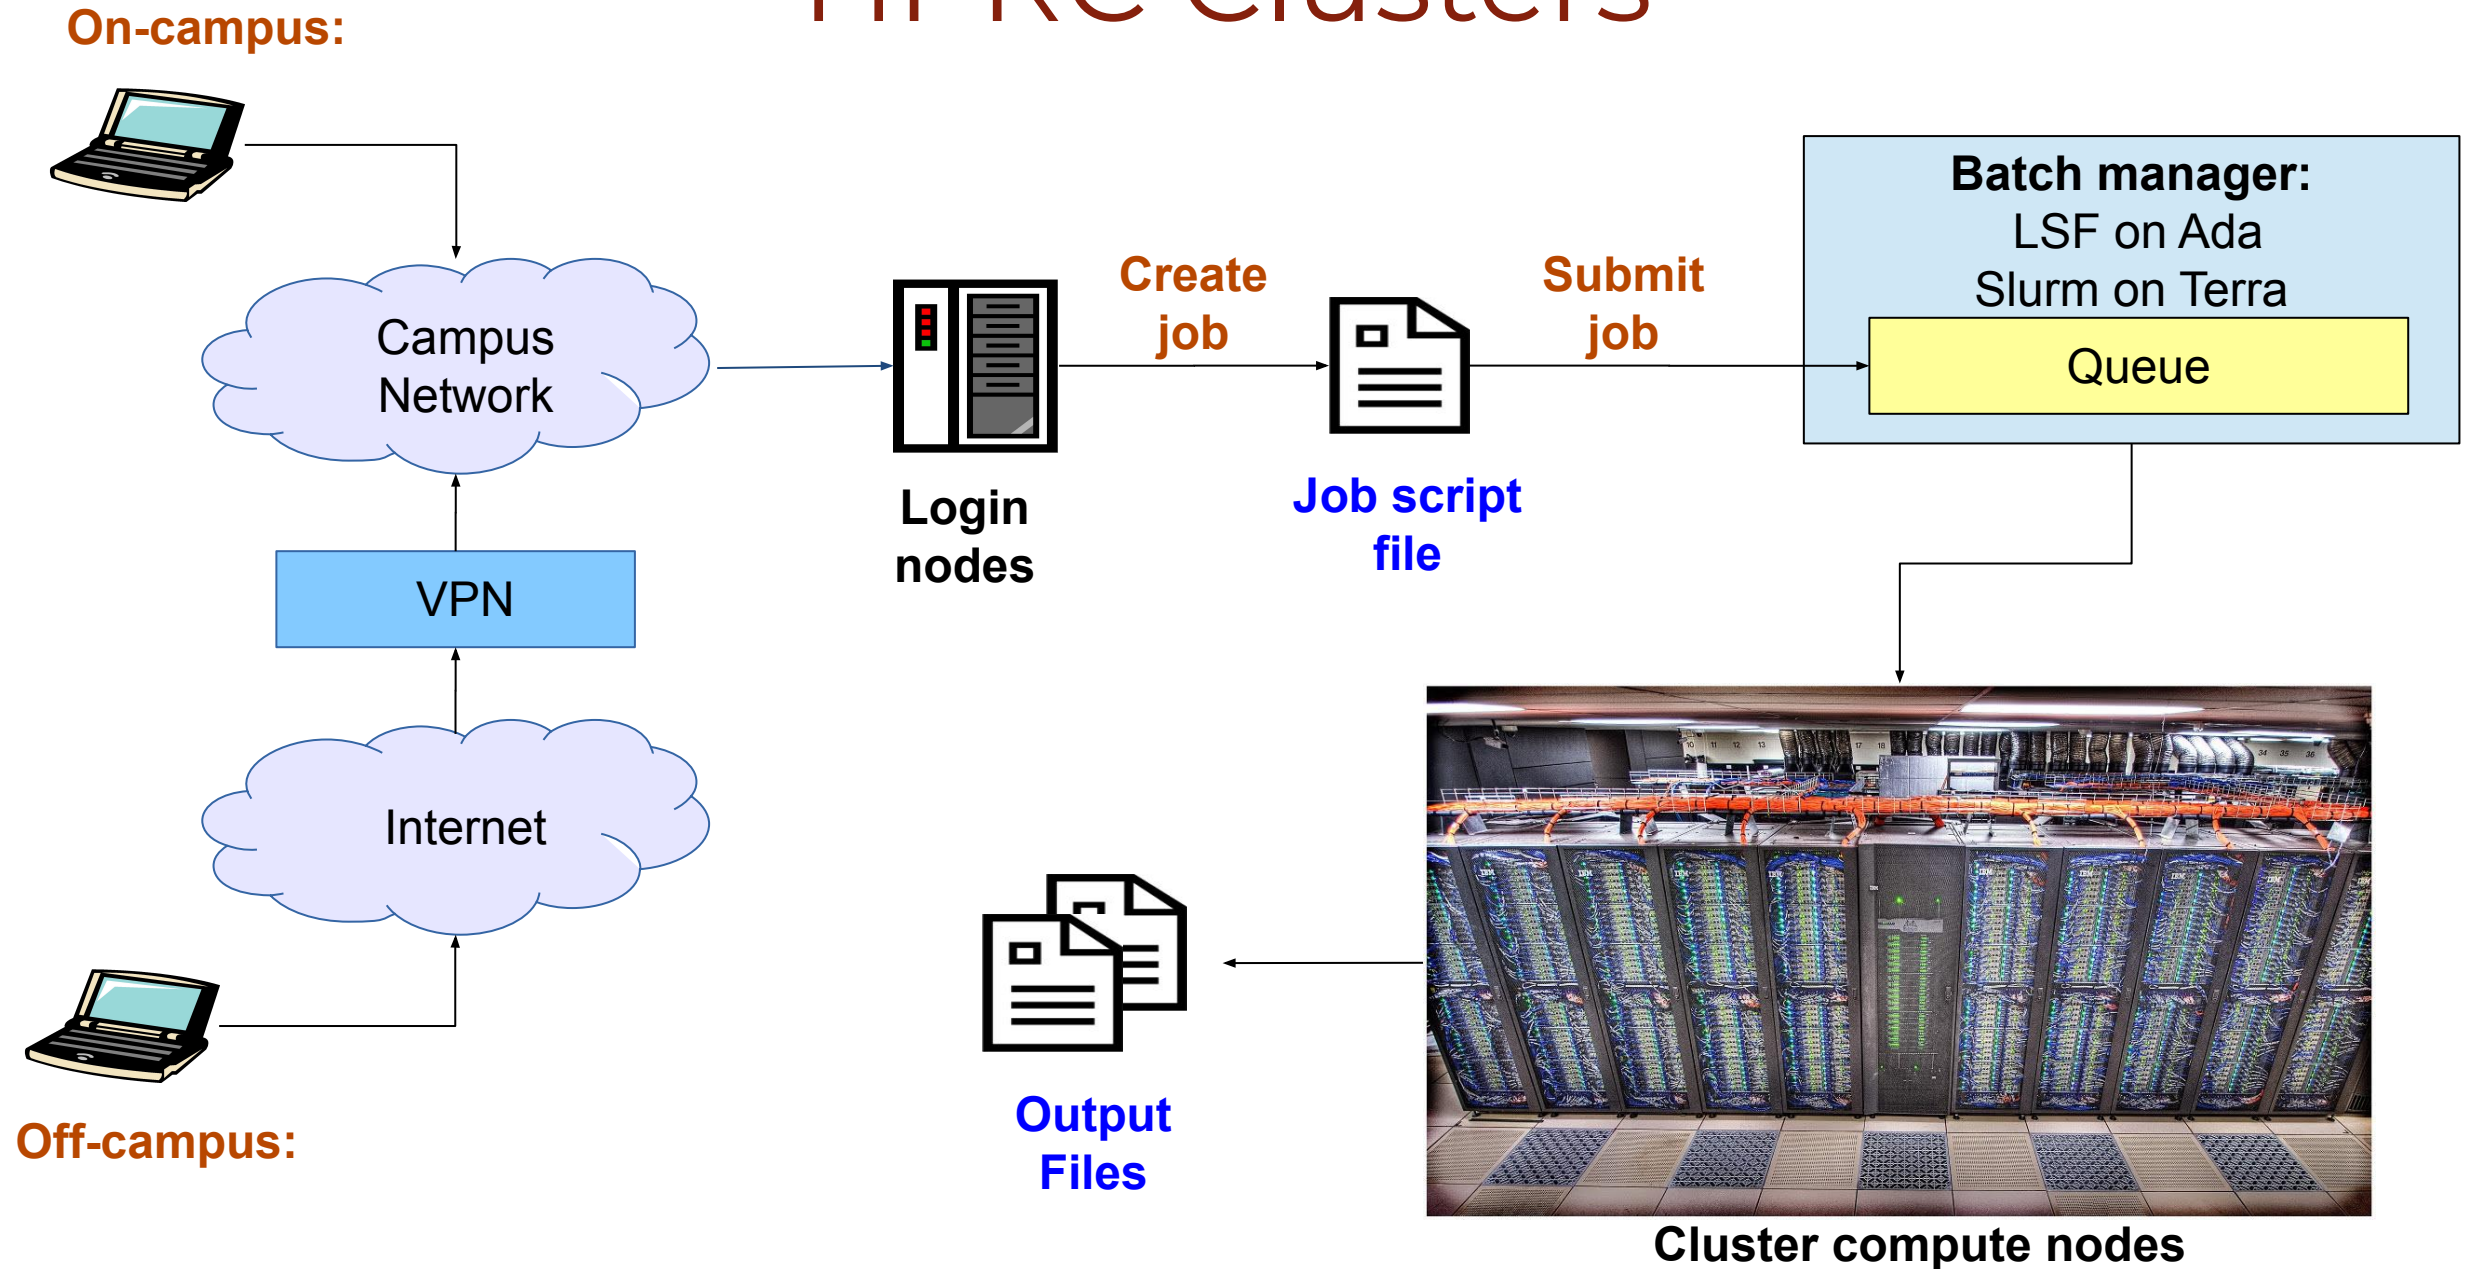
A Service Unit (**SU**) is equivalent to **one core** or **2 GB memory** usage for **one hour**.

Number of Cores	GB of memory per core	Total Memory (GB)	Hours	SUs charged
1	2	2	1	1
1	3	3	1	2
1	56	56	1	28
28	2	56	1	28


[hprc.tamu.edu/wiki/HPRC:AMS:Service\\_Unit](https://hprc.tamu.edu/wiki/HPRC:AMS:Service_Unit)



# HPRC Clusters



# HPC Diagram



# Documentation

## <https://hprc.tamu.edu/wiki>



[HPRC Home Page](#)  
[Wiki Home Page](#)  
[Policies](#)  
[New User Info](#)  
[Contact Us](#)

#### User Guides

[Ada](#)  
[Terra](#)  
[Curie](#)  
[OOD Portal](#)  
[Galaxy](#)

#### Helpful Pages

[AMS Documentation](#)  
[Batch Translation](#)  
[Software](#)  
[File Transfer](#)  
[Two Factor](#)  
[Systems](#)  
[Events](#)  
[FAQ](#)

#### Tools

[What links here](#)  
[Related changes](#)  
[Special pages](#)  
[Printable version](#)

## High Performance Research Computing

*A Resource for Research and Discovery*



## Welcome to the TAMU HPRC Wiki

- [Ada Guide](#)
- [Terra Guide](#)
- [Curie Guide](#)
- [Software](#)
- [Usage Policies](#)
- [Contact Us](#)

### Getting Started: Understanding HPC

New to High Performance Computing (HPC)? This HPC Introduction Page explains the "why" and "how" of high performance computing. Also see the [Policies Page](#) to better understand the rules and etiquette of cluster usage.

### Getting Started: Your First Batch Job

#### Getting an Account

All computer systems managed by the HPRC are available for use to TAMU faculty, staff, and students who require large-scale computing capabilities. The HPRC hosts the [Ada](#), [Curie](#) and [Terra](#) clusters at TAMU. To apply for or renew an HPRC account, please visit the [Account Applications](#) page. For information on how to obtain an allocation to run jobs on one of our clusters, please visit the [Allocations Policy](#) page. *All accounts expire and must be renewed in September of each year.*

#### Quick Start Guides

The Quick Start Guides for the Ada and Terra clusters are available to help new users

### Upcoming Change: Two Factor Authentication for SSH

#### Logging in With Two-Factor Authentication

**Starting November 4th, 2019**, two-factor authentication will be required to login to any cluster. This is most commonly done with Duo, which can be set up by following [these instructions](#).

When enabled, the login prompt will add the following information following successful password entry:

```
Duo two-factor login for netid
```

Enter a passcode or select one of the following options:

- Duo Push to XXX-XXX-1234
- Phone call to XXX-XXX-1234
- SMS passcodes to XXX-XXX-1234 (next code starts with: 9)





# Available Software Modules

<https://hprc.tamu.edu/software/ada>  
<https://hprc.tamu.edu/software/terra>

## SOFTWARE MODULES ON THE ADA CLUSTER

Last Updated: Mon Oct 14 00:00:02 CDT

The available software for the Ada cluster is listed in the table. Click on any software package name to get more information such as the available versions, additional documentation if available, etc.

Show 10 entries Search: samtools

Name	Description
<a href="#">FastalIndex</a>	FastA index (.fai) handler compatible with samtools faidx
<a href="#">grabix</a>	grabix leverages the fantastic BGZF library in samtools to provide random access into text files that have been compressed with bgzip. grabix creates it's own index (.gbi) of the bgzipped file. Once indexed, one can extract arbitrary lines from the file with the grab command. Or choose random lines with the, well, random command.
<a href="#">Pysam</a>	Pysam is a python module for reading and manipulating Samfiles. It's a lightweight wrapper of the samtools C-API. Pysam also includes an interface for tabix. URL: <a href="https://github.com/pysam-developers/pysam">https://github.com/pysam-developers/pysam</a>
<a href="#">Sambamba</a>	Sambamba is a high performance modern robust and fast tool (and library), written in the D programming language, for working with SAM and BAM files. Current functionality is an important subset of samtools functionality, including view, index, sort, markup, and depth
<a href="#">SAMtools</a>	SAM Tools provide various utilities for manipulating alignments in the SAM format, including sorting, merging, indexing and generating alignments in a per-position format.

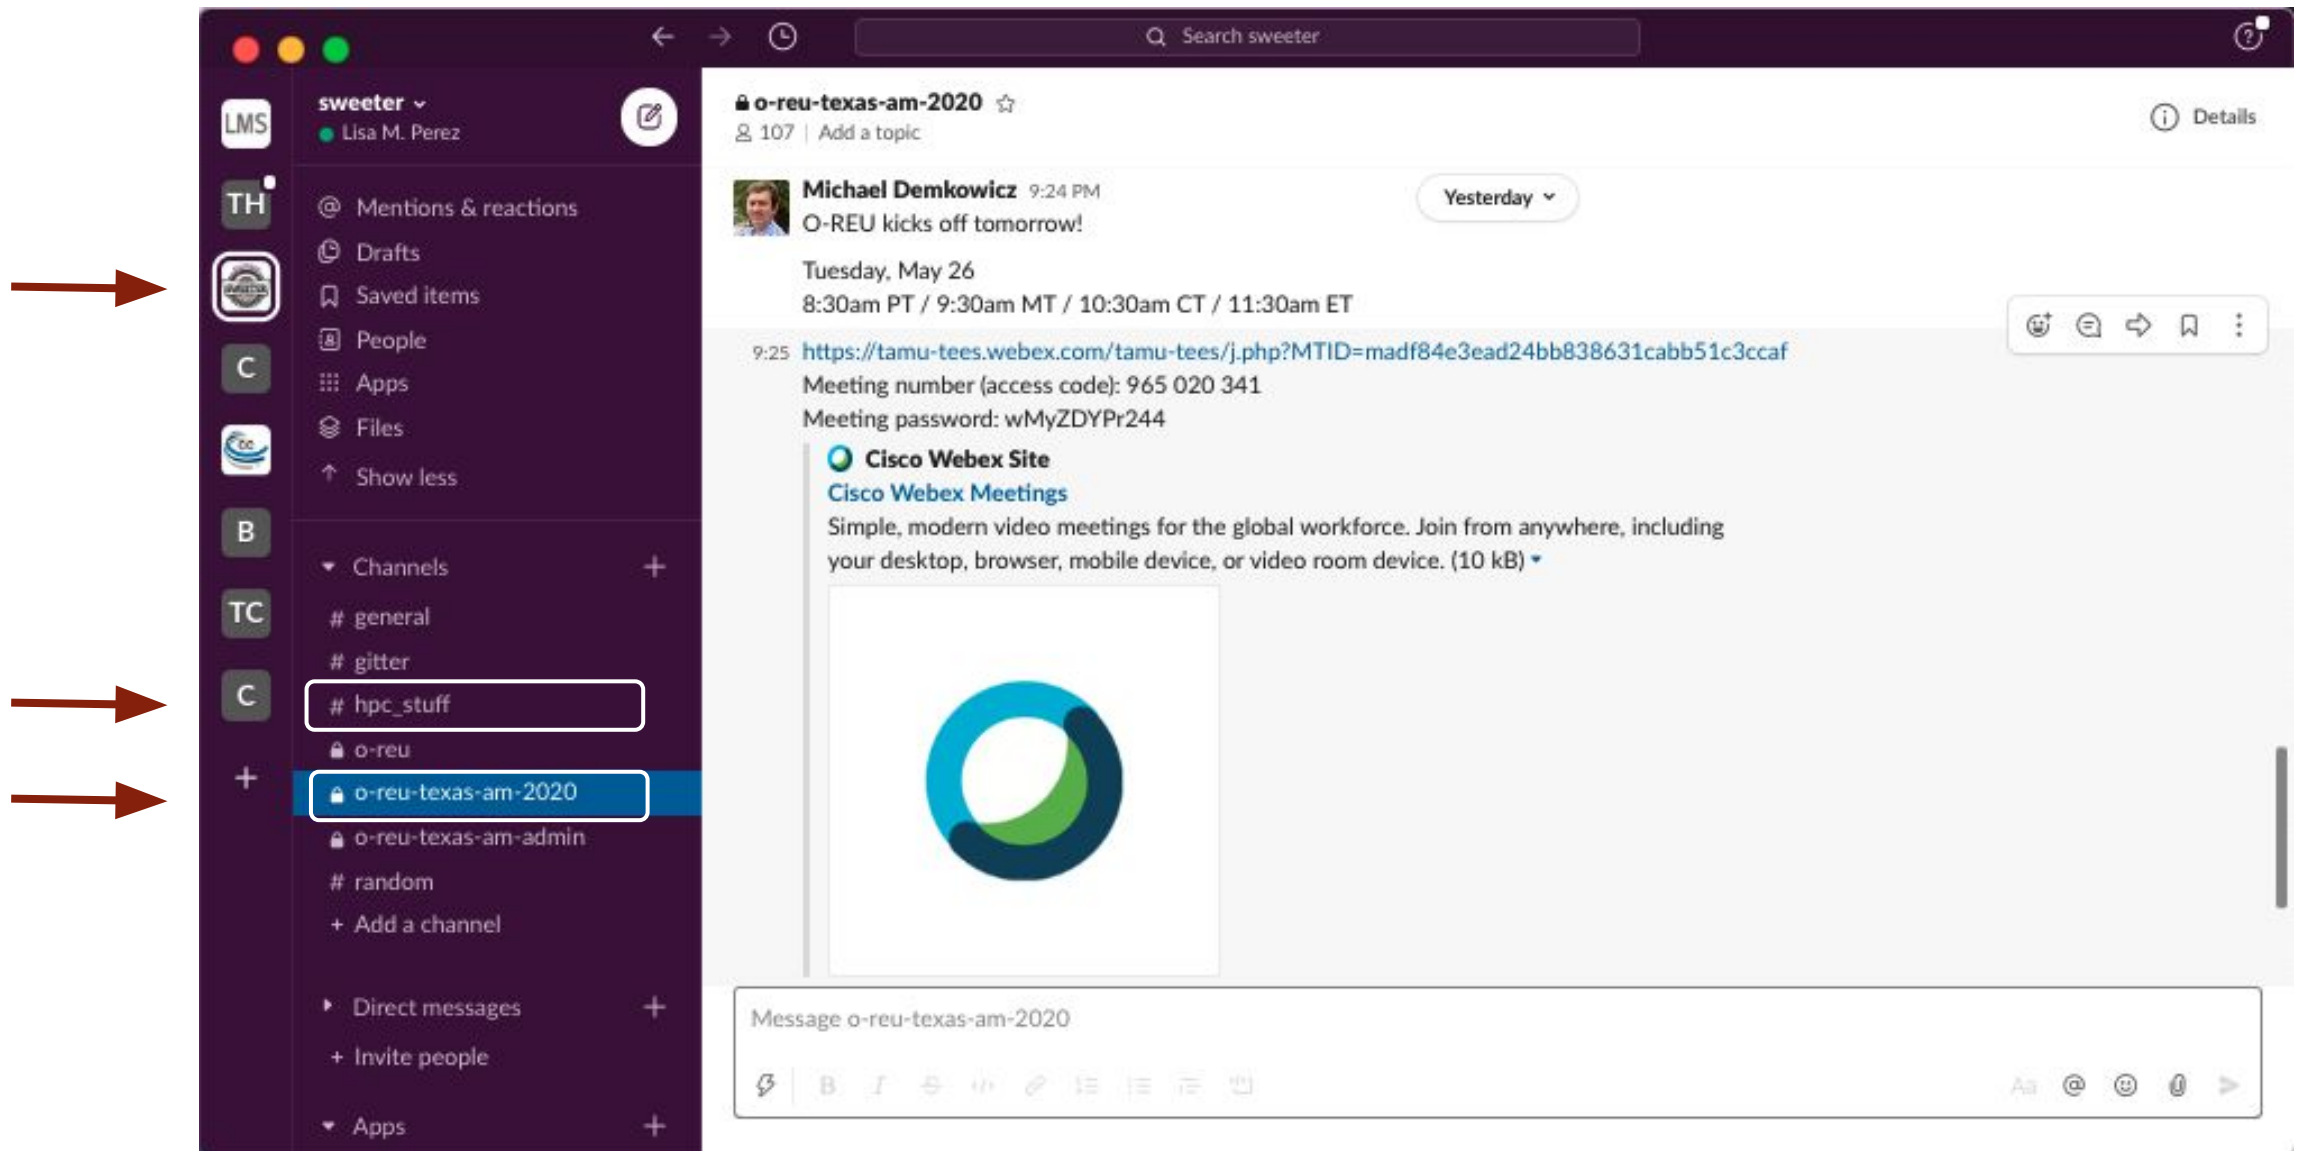
Showing 1 to 6 of 6 entries (filtered from 1,378 total entries) Previous 1 Next

**SAMtools**  
Home Page: <http://www.htslib.org/>  
Description: SAM Tools provide various utilities for manipulating alignments in the SAM format, including sorting, merging, indexing and generating alignments in a per-position format.  
Notes:  
Versions:  
SAMtools/0.1.9-foss-2018b  
SAMtools/0.1.9-intel-2015B  
SAMtools/0.1.18-intel-2015B  
SAMtools/0.1.18-intel-2016a  
SAMtools/0.1.19-GCCcore-6.3.0  
SAMtools/0.1.19-ictce-7.1.2  
SAMtools/0.1.19-intel-2015B  
SAMtools/0.1.20-foss-2018b  
SAMtools/0.1.20-intel-2018b  
SAMtools/1.0-ictce-6.3.5  
SAMtools/1.1-intel-2015B  
SAMtools/1.2-goolf-1.7.20-HTSlib-1.2.1  
SAMtools/1.2-intel-2015B-HTSlib-1.2.1-r2  
SAMtools/1.2-intel-2015B-HTSlib-1.2.1  
SAMtools/1.3-foss-2016a  
SAMtools/1.3-GCCcore-6.3.0  
SAMtools/1.3-intel-2015B-HTSlib-1.3  
SAMtools/1.3-intel-2015B  
SAMtools/1.3-intel-2016a  
SAMtools/1.3.1-foss-2016a  
SAMtools/1.3.1-GCCcore-6.3.0  
SAMtools/1.3.1-intel-2016a  
SAMtools/1.3.1-intel-2016b  
SAMtools/1.3.1-intel-2018b  
SAMtools/1.6-GCC-6.4.0-2.28  
SAMtools/1.6-GCCcore-6.3.0  
SAMtools/1.6-iccifort-2017.4.196-GCC-6.4.0-2.28  
SAMtools/1.7-GCCcore-6.3.0  
SAMtools/1.8-GCCcore-6.3.0  
SAMtools/1.8-GCCcore-6.4.0  
SAMtools/1.9-foss-2018b  
SAMtools/1.9-GCC-7.3.0-2.30  
SAMtools/1.9-GCC-8.2.0-2.31.1  
SAMtools/1.9-GCCcore-6.4.0  
SAMtools/1.9-iccifort-2019.1.144-GCC-8.2.0-2.31.1  
SAMtools/1.9-intel-2018b



# Slack - Collaboration Software

<https://slack.com/resources/slack-101>





**HIGH PERFORMANCE  
RESEARCH COMPUTING**  
TEXAS A&M UNIVERSITY

Help with HPRC:

[hprc.tamu.edu](http://hprc.tamu.edu)

[help@hprc.tamu.edu](mailto:help@hprc.tamu.edu)

Help with TAMU IT (netid, VPN, enabling DUO):

[it.tamu.edu/help/](http://it.tamu.edu/help/)

[helpdesk@tamu.edu](mailto:helpdesk@tamu.edu)

