

Research Computing at Texas A&M

CVMBS Research Conversations
March 2, 2022



High Performance
Research Computing
DIVISION OF RESEARCH

HPRC Services

- **Free of charge** to all faculty, postdocs, research staff, students and external collaborators
- Computing cycles for research - campus, regional and national.
Application is required for access
- User Services
 - Helpdesk: New user start-up assistance and general support
 - Training: Short Courses, Workshops, & [YouTube](#) videos
 - Advanced Support: Software and research consulting
 - Expertise in many science and engineering research domains
- Research collaborations around cyberinfrastructure technologies





<https://hprc.tamu.edu>

Quick Links

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

User Guides

- Terra
- Grace
- Portal
- Galaxy


Cluster Status

Grace

Nodes 605/874 (69%)
 Cores 19386/42128 (46%)
 Jobs 1581R-261Q

Terra

Nodes 262/309 (85%)



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

JAN 14

[Research study by Texas A&M Libraries finds HPRC's work is recognized by the Texas A&M community](#)

DEC 2

[Texas A&M HPRC supported course materials are now available on OakTrust](#)

Events


Dec 4

[Expanding Your Horizons - Coding for Fun!](#)

Dec 3

[HPRC Data Workshop at the Texas A&M Conference on Energy](#)


High Performance Research Computing Clusters



Grace



Terra




ViDaL

Total Nodes (Cores)	925 (44,656)	307 (8,512)	24 (1,120)
General Nodes	48 cores 384GB	28 cores 64GB	40 cores 192 GB
Features	GPUs (A100, RTX 6000, T4) Large Memory Nodes	GPUs (K80, V100) KNL	Compliant Computing GPUs (V100) Large Memory Nodes
Interconnect	HDR100 InfiniBand	Omni-Path	40Gb Ethernet
Global Disk (raw)	8.9 PB	7.4 PB	2 PB

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

HPRC Account Allocations

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 (Terra)	Faculty, Research Scientists, Qualified Staff	300,000	5,000,000	5,000,000	Determined by HPRC-RAC	No	HPRC-RAC
Research (Grace)	Faculty, Research Scientists, Qualified Staff	300,000	10,000,000	10,000,000	Determined by HPRC-RAC	No	HPRC-RAC



Students & Postdoctoral researchers can apply for a Basic allocation.

PIs can apply for a Startup or Research allocation and sub-allocate SUs to their researchers.

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

HPRC Account: PI Eligibility

Only **active faculty** members and **permanent research staff** (subject to HPRC-RAC Chair review and approval) of Texas A&M System Members headquartered in Brazos County can serve as a PI.

Adjunct and Visiting professors can use HPRC resources as part of a sponsoring PI's group

Note that:

- A PI can have more than one allocation
- A researcher (student) can work on more than one project and with more than one PI

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

Special Requests

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

- Dedicated Use
 - Requests for dedicated cluster use require the approval of the Director.
- Special case allocations
 - 20% of common SUs are reserved for special case assignments
 - Example special case assignments
 - working with HPRC staff on new capabilities of general value to research communities
 - new faculty startup
 - operations that go beyond normal research projects
 - Granted by the Director or the VPR.
- Committed Allocations
 - PIs can purchase dedicated HPRC infrastructure via the “condo” mechanism

HPRC Training Short Courses

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

Primers:

Linux
HPRC Clusters
Data Management
SLURM
Jupyter Notebook

Technology Lab:

Using AI Frameworks
in Jupyter Notebook

Short Courses:

Python
Scientific Python
PyTorch
TensorFlow
MATLAB
Scientific ML
Julia
CUDA
Drug Docking
Quantum Chemistry
and more...

Short Courses:

NGS Analysis
NGS Metagenomics
NGS RADSeq/GBS
NGS Assembly
HPRC Galaxy
Linux
R
Perl
Fortran
OpenMP
MPI

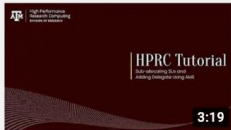







Texas A&M HPRC
539 subscribers

[YouTube training videos](#)

- HOME
- VIDEOS
- PLAYLISTS
- COMMUNITY
- CHANNELS
- ABOUT

Uploads ▶ PLAY ALL

 3:19	 14:46	 9:29	 12:01	 16:00	 23:53
HPRC AMS: Sub-allocating SUs and adding a delegate... 6 views · 7 hours ago CC	BRICCS Talk: Forough Ghahramani, EdD - CI-based... 9 views · 13 days ago	BRICCS Talk: Prof. Alex Feltus - Bringing CI to... 11 views · 2 weeks ago	BRICCS Talk: Prof. Izzat Alsmadi - TAMUSA Researc... 12 views · 2 weeks ago	BRICCS Talk: Dr. Richard Knepper - Cyberinfrastructu... 7 views · 2 weeks ago	BRICCS Talk: Professor Rajiv Malkan - Data Analytics in t... 8 views · 2 weeks ago

Recommended Playlists

 12	 4	 1	 6	 15
Getting Started with HPRC Texas A&M HPRC VIEW FULL PLAYLIST	HPRC OpenOnDemand Portal Texas A&M HPRC VIEW FULL PLAYLIST	Project Account Management (AMS) Texas A&M HPRC Updated today	HPRC Primers Texas A&M HPRC VIEW FULL PLAYLIST	HPRC Short Courses Texas A&M HPRC VIEW FULL PLAYLIST



Advanced Support Program

HPRC scientists can contribute expertise on computational research projects.

- Software development for **research workflows**
- Developing **GUIs and apps** for research projects
- **Porting** applications to HPRC clusters
- **Code development**, optimizing and analysis
- Leveraging mathematical libraries
- **Workflow automation** in scientific processes

Please send us an e-mail at help@hprc.tamu.edu.

NSF award #1925764, CC* Team: SWEETER -- SouthWest Expertise in Expanding, Training, Education and Research and NSF award # 2112356, ACSS: ACES - Accelerating Computing for Emerging Sciences.

Bring Your Own Genome

- Work one-on-one with Dr. Wesley Brashear!
- Wednesdays 3 PM to 4:30 PM
- In-person or via zoom
- It's for FREE!
- Register at <https://u.tamu.edu/byog>



Documentation

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

Log in

Search TAMU HPRC



HPRC Home Page
Wiki Home Page
Policies
New User Info
Contact Us

User Guides
Grace
Terra
OOD Portal
Galaxy

Helpful Pages

AMS Documentation
Batch Transiation
Software
File Transfer
Two Factor
Systems
Events
FAQ

Tools

What links here
Related changes
Special pages
Printable version
Permanent link
Page information

Welcome to the TAMU HPRC Wiki

• [Grace Guide](#)
• [Terra Guide](#)

• [Software](#)
• [Usage Policies](#)

• [Contact Us](#)

Announcements

- **Grace Cluster Status:** Cluster deployed, with service unit accounting.

Getting an Account

- **Understanding HPRC:** For a brief overview of what services HPRC offers, see [this video](#) in our getting started series on YouTube.
- **New to HPRC's resources?** [This page](#) explains the HPRC resources available to the TAMU community. Also see the [Policies Page](#) to better understand the rules and etiquette of cluster usage..
- **Accessing the clusters:** 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 [Terra](#), and [Grace](#) 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.*

Using the Clusters

- **QuickStart Guides:** For just the "need-to-know" information on getting started with our clusters, visit our QuickStart pages. Topics discussed include cluster access, file management, the batch system, setting up a software environment using modules, creating your own job files, and project account management. [Grace Quickstart Guide](#), [Terra Quickstart Guide](#).
- **Batch Jobs:** As a shared resource between many users, each cluster must employ a batch system to schedule a time for each user's job to run. Without such a system, one user could use a disproportionate amount of resources, and cause other users' work to stall. Grace's and Terra's batch system is called SLURM. Information relevant to each system can be found below.

Grace / SLURM Batch Pages

[Complete Grace Batch Page](#)
[Job Submission \(sbatch\)](#)
[Grace Queue Structure](#)

Terra / SLURM Batch Pages

[Complete Terra Batch Page](#)
[Job Submission \(sbatch\)](#)
[Terra Queue Structure](#)

Creating your own batch jobs: the [tamubatch Page](#) provides information on how to use tamubatch to create and submit jobs easily.

- **Troubleshooting:** While we cannot predict all bugs and errors, some issues on our clusters are common enough to catalog. See the [Common Problems and Quick Solutions Page](#) for a small collection of the most prevalent issues. For further assistance, users can contact help@hprc.tamu.edu to open a support ticket.

HPRC's YouTube Channel

Prefer visual learners? HPRC has launched its official YouTube channel where you can find short versions of our help and training videos, as well as our latest software and news! [Subscribe here!](#)



HPRC NGS Documentation

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

Bioinformatics Tool Categories

- Sequence QC
- Data Normalization, Clustering & Collapsing
- PacBio Tools
- Oxford Nanopore Tools
- Genome Assembly
- Metagenomics
- RNA-seq →
- ChIP-seq
- Sequence Variants (SNPs and indels)
- CNV
- Methylation
- Sequence Alignments
 - : : :
- Common Tools
- Data Visualization
- Statistics
- File Format Tools
- Licenses
- Aspera (SRA, 1000genomes, BioMart)
- Conda/Bioconda
- Biocontainers
- FAQ

RNA-seq

[Back to Bioinformatics Main Menu](#)

Contents [hide]

- 1 RNA-seq
 - 1.1 TUTORIALS
 - 1.2 Transcriptome Assembly
 - 1.2.1 Trinity →
 - 1.2.1.1 Sample Trinity Paired End Assembly job Scripts
 - 1.2.1.1.1 Terra: 54 GB compute nodes
 - 1.2.1.1.2 using \$TMPDIR
 - 1.2.1.2 Example Trinity Tutorials
 - 1.2.2 Scripture
 - 1.2.3 StringTie
 - 1.2.4 Trans-ABYSS
 - 1.2.5 SOAPdenovo-Trans
 - 1.3 Transcriptome Assembly Evaluation
 - 1.3.1 DETONATE
 - 1.3.2 BUSCO
 - 1.3.2.1 version 5.0.x
 - 1.3.2.2 version 4.0.x
 - 1.3.2.3 version 3.0.2b
 - 1.4 Transcriptome Assembly Annotation
 - 1.4.1 Trinotate
 - 1.4.2 TransDecoder
 - 1.4.3 miRNA
 - 1.5 Differential Expression
 - 1.5.1 Bowtie & Bowtie2
 - 1.5.2 Cufflinks

Trinity

GCATemplates available: [grace](#)

```
module spider Trinity
```

[Trinity homepage](#)

[Trinity tutorial videos](#)

[Trinity memory](#) requirements summary

Trinity performs *de novo* or reference guided genome assemblies of transcript sequences from Illumina RNA-Seq data.

R Bioconductor libraries, such as `qvalue`, required by some Trinity analysis scripts are found in this `R_tamu` module

```
module load R_tamu/3.6.0-foss-2018b-recommended-mt
```

Trinity will create 100,000+ temporary files during an analysis run. A quota of 250,000 files (with about 225,000 available) should be enough for one Trinity run at a time. If you want to run multiple Trinity jobs at once, you can use `run` the analysis in `$TMPDIR` and copy the results files to your `$SCRATCH` working directory.

Trinity does have checkpoints. If your job times out and you run the same job script again, it should pick up where it left off unless you are using `$TMPDIR`.

If you are using Galaxy then checkpoints will not work.

NOTE: Trinity 2.4.0+ utilizes `seqtk` for read normalization which requires read headers to be in a specific format.

3,300+ Software Modules!

SOFTWARE MODULES ON THE GRACE CLUSTER

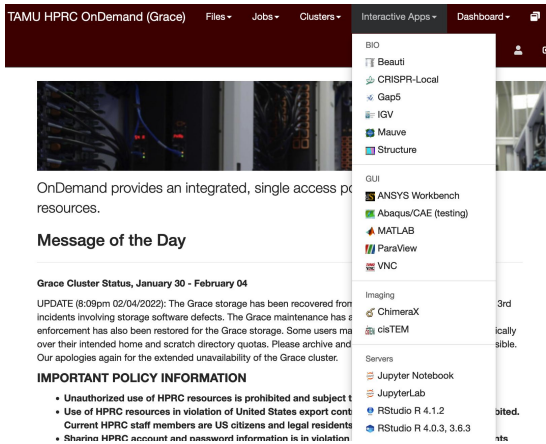
Last Updated: Jul 9 09:33:55 CDT

The available software for the [Grace 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.

Name	Description
AAF	'AAF constructs phylogenies directly from unassembled genome sequence data, bypassing both genome assembly and alignment.'
AGFusion	'AGFusion is a python package for annotating gene fusions from the human or mouse genomes.'
AMOS	'The AMOS consortium is committed to the development of open-source whole genome assembly software'
antiSMASH	'antiSMASH allows the rapid genome-wide identification, annotation and analysis of secondary metabolite biosynthesis gene clusters in bacterial and fungal genomes.'
Bowtie	'Bowtie is an ultrafast, memory-efficient short read aligner. It aligns short DNA sequences (reads) to the human genome.'
Bowtie2	'Bowtie 2 is an ultrafast and memory-efficient tool for aligning sequencing reads to long reference sequences. It is particularly good at aligning reads of about 50 up to 100s or 1,000s of characters, and particularly good at aligning to relatively long (e.g. mammalian) genomes. Bowtie 2 indexes the genome with an FM Index to keep its memory footprint small: for the human genome, its memory footprint is typically around 3.2 GB. Bowtie 2 supports gapped, local, and paired-end alignment modes.'

Running Bioinformatics Toolkits

HPRC Portal



TAMU HPRC OnDemand (Grace) Files Jobs Clusters Interactive Apps Dashboard

BIO
Beauti
CRISPR-Local
Gap5
IGV
Mauve
Structure

GUI
ANSYS Workbench
Abaqus/CAE (testing)
MATLAB
ParaView
VNC

Imaging
ChimeraX
cisTEM

Servers
Jupyter Notebook
JupyterLab
RStudio R 4.1.2
RStudio R 4.0.3, 3.6.3

OnDemand provides an integrated, single access point for all HPRC resources.

Message of the Day

Grace Cluster Status, January 30 - February 04

UPDATE (8:09pm 02/04/2022): The Grace storage has been recovered from incidents involving storage software defects. The Grace maintenance has been completed and enforcement has also been restored for the Grace storage. Some users may have been affected as their intended home and scratch directory quotas. Please archive and backup your data. Our apologies again for the extended unavailability of the Grace cluster.

IMPORTANT POLICY INFORMATION

- Unauthorized use of HPRC resources is prohibited and subject to the University of Texas System's Information Security Policy.
- Use of HPRC resources in violation of United States export control regulations is prohibited. Current HPRC staff members are US citizens and legal residents.
- Sharing HPRC account and password information is in violation of the University of Texas System's Information Security Policy.

- Web-based
- Software tools usable as a GUI or linux command line
- Best for GUI apps:
 - RStudio
 - IGV

Linux command line

BIOINFORMATICS GCATemplates (grace)

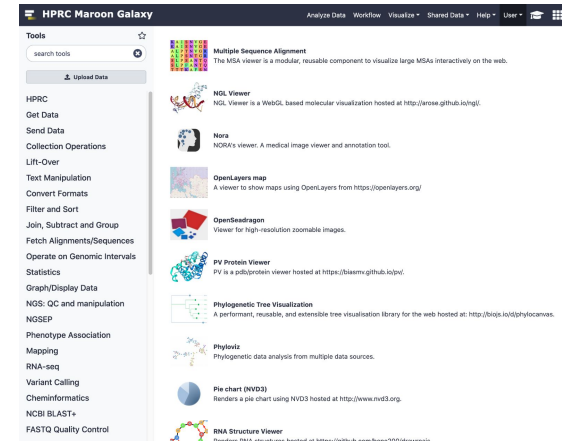
```
CATEGORY
1. FASTA files
2. FASTQ files (QC, trim, SRA)
3. Genome assembly
4. Metagenomics
5. PacBio tools
6. Phylogenetics
7. Population genetics
8. Protein tools
9. RNA-seq
10. SNPs & indels
11. Sequence alignments
12. Simulate data
```

```
s search
q quit
```

```
Select 11
```

- Linux and slurm
- Bioinformatics template scripts
- Need SSH client on your computer

HPRC Maroon Galaxy



HPRC Maroon Galaxy Analyze Data Workflow Visualize Shared Data Help User

Tools
search tools

Multiple Sequence Alignment
The MSA viewer is a modular, reusable component to visualize large MSAs interactively on the web.

HPRC
Get Data
Send Data
Collection Operations
Lift-Over
Text Manipulation
Convert Formats
Filter and Sort
Join, Subtract and Group
Fetch Alignments/Sequences
Operate on Genomic Intervals
Statistics
Graph/Display Data
NGS: QC and manipulation
NGSEP
Phenotype Association
Mapping
RNA-seq
Variant Calling
Cheminformatics
NCBI BLAST+
FASTQ Quality Control

NGL Viewer
NGL Viewer is a WebGL based molecular visualization hosted at <http://arose.github.io/nvgl/>.

Nora
NORA's viewer. A medical image viewer and annotation tool.

OpenLayers map
A viewer to show maps using OpenLayers from <https://openlayers.org/>

OpenSeadragon
Viewer for high-resolution zoomable images.

PV Protein Viewer
PV is a pdb/protein viewer hosted at <https://bluem.github.io/pv/>.

Phylogenetic Tree Visualization
A performant, reusable, and extensible tree visualization library for the web hosted at <http://biop-lobbio.github.com/>.


Phylovis
Phylogenetic data analysis from multiple data sources.

Pie chart (NVD3)
Renders a pie chart using NVD3 hosted at <http://www.nvd3.org>.

RNA Structure Viewer
RNA Structure Viewer hosted at <https://github.com/robertcraig/structure-viewer>

- Web-based
- Many software tools are available
- First [apply](#) for an HPRC account then request a Galaxy account

HPRC Portal

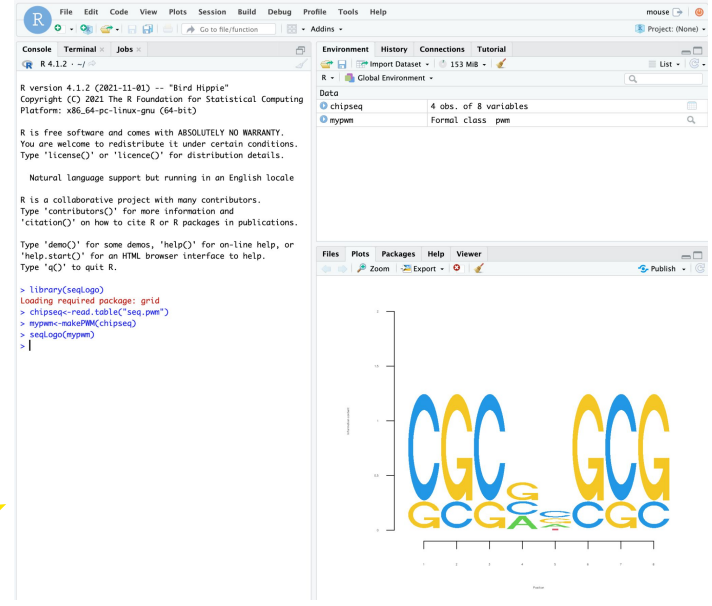
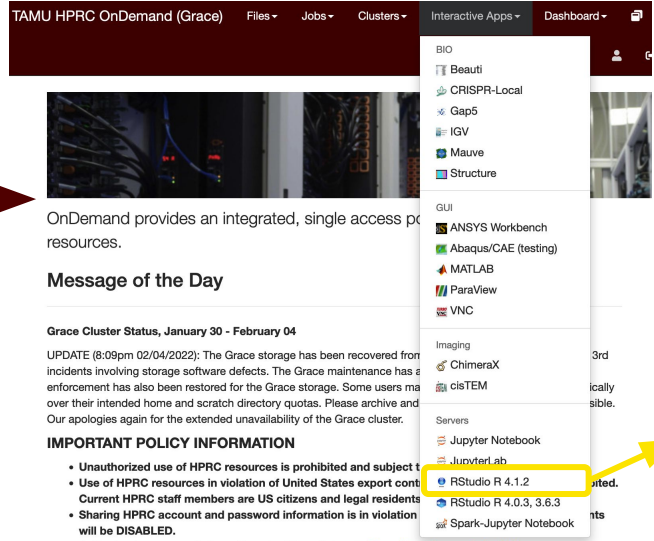


<https://portal.hprc.tamu.edu>

Interactive Apps:
launch R and
more right in your
browser.

Open OnDemand
(OOD) Portal is an
advanced web-based
graphical interface for
HPC users.


[HPRC Portal
YouTube tutorials](#)





CLUSTER STATISTICS

Node Utilization



Core Utilization



Jobs

Running	1333
Pending	533

SUMMARY

Accounts

Account ↑↓	Default ↑↓	Allocation ↑↓	Used ↑↓	Balance ↑↓
132853912556	Set Default	200000	150000	50000
132853918900	default	5000	-2028.54	2971.46

Disk Quotas

Disk	Disk Usage	Limit	File Usage	Limit
/home	5.49 GB (54.89 %)	10 GB	7454 (74.54 %)	10000
/scratch	159.66 GB (15.59 %)	1 TB	226069 (90.43 %)	250000

Request Quota Increase

Your Jobs












Job ID ↑↓	Name ↑↓	State ↑↓	Partition ↑↓	
3558860	...dio_r_tamu	RUNNING	short	Log Error Kill

HPRC Portal Toolbox







- Cluster statistics
- Allocations
- Disk quotas
- Jobs

HPRC Portal - Interactive Apps





BIO

-  Beauti
-  DIYABC
-  FigTree
-  IGV
-  JBrowse
-  Krait
-  Mauve
-  Structure
-  Tracer
-  CRISPR-Local
-  Gap5











GUI

-  ANSYS Workbench
-  Abaqus/CAE
-  LS-PREPOST
-  MATLAB
-  ParaView
-  VNC

Servers

-  Jupyter Notebook
-  JupyterLab
-  RStudio
-  Spark-Jupyter Notebook

Imaging

-  AFNI
-  Chimera
-  Coot
-  Diffusion Toolkit & TrackVis
-  FSL
-  Fiji
-  ICY
-  ImageJ
-  Vaa3D
-  cisTEM

- Tools ☆
- search tools ✕
- Upload Data
- HPRC
 - Get Data
 - Send Data
 - Collection Operations
 - Lift-Over
 - Text Manipulation
 - Convert Formats
 - Filter and Sort
 - Join, Subtract and Group
 - Fetch Alignments/Sequences
 - Operate on Genomic Intervals
 - Statistics
 - Graph/Display Data
 - NGS: QC and manipulation
 - NGSEP
 - Phenotype Association
 - Mapping
 - RNA-seq
 - Variant Calling
 - Cheminformatics
 - NCBI BLAST+
 - FASTQ Quality Control



Multiple Sequence Alignment

The MSA viewer is a modular, reusable component to visualize large MSAs interactively on the web.



NGL Viewer

NGL Viewer is a WebGL based molecular visualization hosted at <http://arose.github.io/ngl/>.



Nora

NORA's viewer. A medical image viewer and annotation tool.



OpenLayers map

A viewer to show maps using OpenLayers from <https://openlayers.org/>



OpenSeadragon

Viewer for high-resolution zoomable images.



PV Protein Viewer

PV is a pdb/protein viewer hosted at <https://biasmv.github.io/pv/>.



Phylogenetic Tree Visualization

A performant, reusable, and extensible tree visualisation library for the web hosted at: <http://biojs.io/d/phylocanvas>.



Phyloviz

Phylogenetic data analysis from multiple data sources.



Pie chart (NVD3)

Renders a pie chart using NVD3 hosted at <http://www.nvd3.org>.



RNA Structure Viewer

Render RNA structure hosted at <https://github.com/ncsa/2016/structure>

<https://hprc.tamu.edu/wiki/SW:Galaxy>

HPRC Genomic Computational Analysis Templates

- GCATemplates resource provides numerous template job scripts
- Template scripts connect to software modules
- Templates are configured with small sample datasets that you can run for testing

```
[mouse@grace4 ~]$ gcatemplates █
```

```
BIOINFORMATICS GCATemplates (grace)
```

CATEGORY

1. FASTA files
2. FASTQ files (QC, trim, SRA)
3. Genome assembly
4. Metagenomics
5. PacBio tools
6. Phylogenetics
7. Population genetics
8. Protein tools
9. RNA-seq
10. SNPs & indels
11. Sequence alignments
12. Simulate data

```
s search
```

```
q quit
```

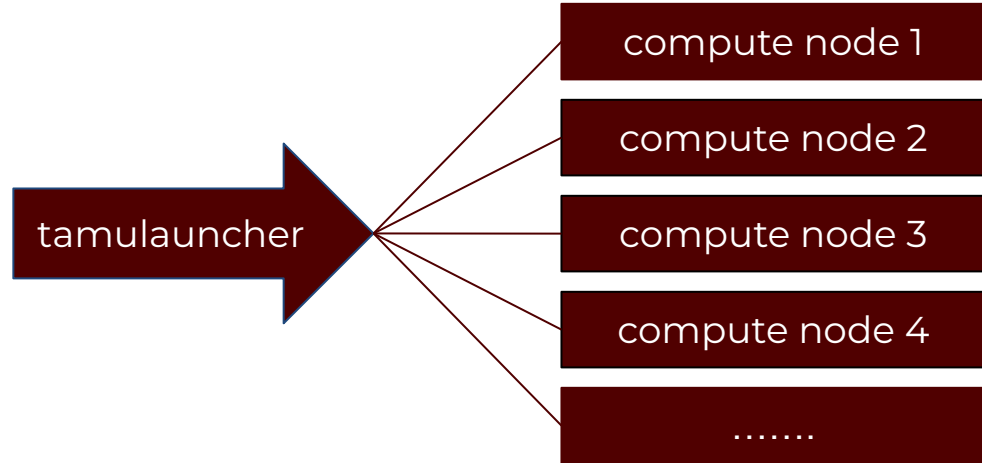
```
Select:11 █
```

HPRC TAMU Launcher

[tamulauncher](#) provides a convenient way to run a large number of serial or multithreaded commands without the need to submit individual jobs or a Job array.

```
commands.txt:
```

```
./blastn ...  
./blastn ...  
./blastn ...  
./blastn ...  
./blastn ...  
: : :
```




Data Transfer Options

- Use HPRC's [fast transfer nodes](#) for high-speed transfers
- For **small files** of less than 2GB
 - Use [HPRC's web-portal](#)
 - [HPRC Galaxy](#) for bio-researchers
 - [MobaXterm](#) from your computer
 - [BaseMount](#) copy files from your Illumina BaseSpace account to Grace
- For **large files** sized hundreds of GB to greater than 1TB
 - [Globus Connect](#) website
 - ftp transfer for [HPRC Galaxy](#)

For more details and options please visit https://hprc.tamu.edu/wiki/HPRC:File_Transfers


Deep Learning in Genomics

- **Deep Learning**



- **Deep Learning Libraries for Genomics**

DragoNN	DragoNN provides a toolkit to learn how to model and interpret regulatory sequence data using deep learning.
Kipoi	Kipoi is an API and a repository of ready-to-use trained models for genomics.



Zou, J., Huss, M., Abid, A. *et al.* A primer on deep learning in genomics. *Nat Genet* 51, 12–18 (2019) doi:10.1038/s41588-018-0295-5

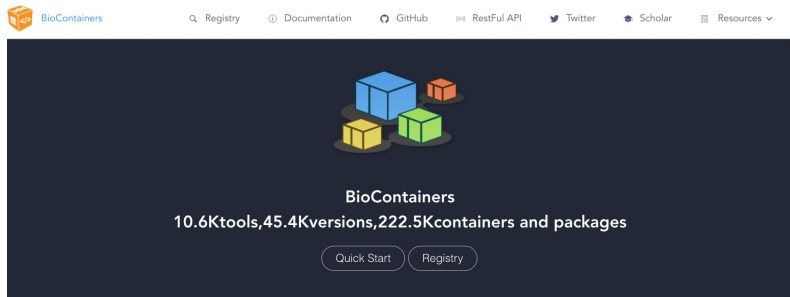




Funded Grants Since 2017

- NSF: ACES (\$6M + \$1M/year for 5 years)
- NSF MRI: FASTER (\$3.09M + \$1.32M TAMU match)
- Texas Virtual Data Library (ViDaL, \$1.4M)
- NSF CyberSEES: Coastal Resilience Collaboratory (Subaward \$100K)
- NSF Frontera (Campus Bridging, Subaward \$500K)
- NSF CC*: Improve Campus Network at PVAMU (Subaward \$65K)
- NSF CC*: SWEETER (\$1.43M)
- NSF CC*: BRICCs (\$250K)
- NSA GenCyber Summer Camps (\$80K, 2019, 2020)
- Texas Workforce Commission (TWC, \$63K, 2019, 2021)
- NSF CyberTraining: CiSE-ProS (\$532K)
- NSF CyberTraining: CMS³ (\$499K)
- THECB (\$301K) & THECB with UTRGV (Subaward \$100K)
- Texas Workforce Commission (\$40.6K, 2022)

Containers and Cloud



BioContainers Flow



OPEN

OnDemand

TAMU HPRC OnDemand (Terra)



Home / My Interactive Sessions / Jupyter Notebook

Jupyter Notebook

This app will launch a [Jupyter Notebook](#) server on the Terra cluster.

Notice: This form has changed. Please pay attention to what options you select and what the defaults are.

Type of environment

Containers (Singularity)

Select the type of environment in which Jupyter is installed. [Help me choose](#)

Path to singularity image file

/scratch/data/Singularity/images/tensorflow_2.4.1-gpu-jupyter.sif

Enter the path to a singularity image file containing the Jupyter app. Recommended that this live under your \$SCRATCH directory.

⋮ (etc)

Node type

GPU

Choose "GPU" if the notebook needs to run on an Nvidia GPU node.

NSF MRI FASTER

Fostering Accelerated Scientific Transformations, Education, and Research

- **composable** software-hardware approach
- 184-Intel Ice Lake nodes (11,520-core) with InfiniBand. (64-core, 256GB memory, and 3.84TB NVMe disk per node)
- **NVIDIA GPUs:** 200x T4, 40x A100, 10x A10, 4x A30, and 8x A40 GPUs
- Each node can compose up to 20 GPUs.



This project is supported by NSF award #[2019129](#)

NSF ACES - Accelerating Computing for Emerging Sciences (Spring 2022)



Component	Quantity	Description
Allocatable resources		Total cores: 11,520, 512 GB memory/node
<u>Graphcore IPU</u>	16	16 IPUs direct-attached to a server
Future <u>Intel Agilex FPGA</u> replacement	20	Agilex FPGA with a broad hierarchy of memory
<u>NextSilicon</u> coprocessor	20	Reconfigurable accelerator
<u>NEC Vector Engine</u>	24	Vector computing card and HBM2 memory
Intel Ponte Vecchio GPU	100	Intel GPU for HPC, DL Training, AI Inference
<u>Liquid Intel Optane PCIe SSDs</u>	6	3 TB PCIe SSD cards addressable as memory

This project is supported by NSF award #[2112356](#)



NSF VIRTUAL CAMPUS CYBER INFRASTRUCTURE PI WORKSHOP SEPTEMBER 14 – 15 & 21-22, 2021

Quad Chart for: *SWEETER: South West Expertise in Expanding Training, Education, and Research*

Challenge Project Seeks to Address:

- Multi-disciplinary research will be strengthened by offering opportunities to researchers to collaborate
- There is a need for computing research support at institutions at all levels of learning
- Research projects need more than enablement to succeed

Deliverables:

- Research exchange supports 20+ fields of science
- A boots-on-the-ground approach using existing CI resources is adopted
- All institutions are providers and consumers of research CI
- Site ambassadors support campuses
- Educational resources developed
- Engage community colleges
- Rotating annual conference and annual site activities



Scientific Impact:

- Holistic vision for researcher success envisioned
- Uses CI as the means for researcher engagement and collaborations
- Enablement is reimagined
- Regional MSI, emergent MSIs non-profits and industry learn together
- Several CI projects implemented

Team:

- Texas A&M, UT Austin, New Mexico State, West Texas A&M, UT San Antonio, Texas A&M San Antonio, University of Arizona, Prairie View A&M, UT Rio Grande Valley, University of Alabama (B), LEARN, and the National Center for Genome Research

Project Updates:

- Stay tuned for updates at hprc.tamu.edu/sweeter/
- Lots of faculty and student programs
- Need more funds to support programs at other regional MSIs!



NSF VIRTUAL CAMPUS CI INFRASTRUCTURE PI WORKSHOP SEPTEMBER 14 – 15 & 21-22, 2021


Quad Chart for: *BRICCs: Building Research Innovation at Community Colleges*

Challenge Project Seeks to Address:

- Expand research cyberinfrastructure adoption at smaller institutions and community colleges
- Develop a communication mechanism to identify and ameliorate local issues
- Offer (local) campus CI expertise for researchers

Deliverables:

- Developed a CI adoption plan for two-year institutions
- Targeted technical, policy and learning resources are offered on our website: <https://hprc.tamu.edu/briccs/>
- Engage with CIO, faculty and administration at smaller institutions
- Rotating annual workshop and site activities
- Assist institutions and groups working on CC* proposals



Scientific Impact:

- Building support mechanisms for curricular and research involving CI
- Engaged with the science team in CC* SWEETER Cyberteam
- Engaged community – volunteers are CIOs and senior faculty members
- Assisting in faculty-engagement at smaller schools

Upcoming Community Workshop:

- Join us in mid-October for our BRICCs workshop virtually or at South Plains College in Levelland, TX.
- Workshop collocated with Texas Association of Community Colleges CIO group

Project Updates:

- Stay tuned for updates on twitter and the HPRC YouTube channel
- Resources available for faculty and student CI programs

Partnering on Outreach

Leverage our programs to strengthen your broader impacts

- Teach a short course on computing to the TAMU community
- Participate as an instructor in the Summer Computing Academy program camps for middle and high school students
- Join the the NSF SWEETER CyberTeam to explore computing-driven research and educational partnerships with universities in Texas, New Mexico and Arizona
- Participate in our NSF BRICCs community to support research computing at smaller institutions and community colleges
- Make your computing products available on the NSF ACES, NSF FASTER, and NSF Frontera machines





High Performance
Research Computing
DIVISION OF RESEARCH

<https://hprc.tamu.edu>

HPRC Helpdesk:

help@hprc.tamu.edu

Phone: 979-845-0219