



INTRO TO HIGH SECURITY HPC (RIO)

Camden Duy

Computational Scientist

E: cmd7ag@[virginia.edu](mailto:cmd7ag@virginia.edu)

Paul Orndorff

Computational Scientist

E: egg3xa@virginia.edu

High Performance Computing (HPC) at UVA

HPC computing systems tailored to data sensitivity

- Rivanna/Afton (Standard Security)
 - To process and store public data, internal use data, or sensitive data.
 - Highly Sensitive Data is not permissible
 - More information on the SSZ HPC can be found in our [Intro to HPC slides](#).
- Rio (High Security)
 - To process and store Highly Sensitive Data such as HIPAA, FERPA, dbGAP, etc.
 - **Not** CUI certified yet, use Ivy VM

For more information please see [RC Systems Data Sensitivity](#)

Highly Sensitive Data

Data that require restrictions on access under the law or that may be protected from release in accordance with applicable law or regulation.

Examples:

- Personally identifiable information (PII) is any information that can be used to identify a person. Examples include social security number, passport number, driver's license number, military identification number, or biometric records
- Health information that reveals an individual's health condition and/or medical history, including information defined by the Health Insurance Portability and Accountability Act (HIPAA)

Ivy and Rio (Overview)

Ivy

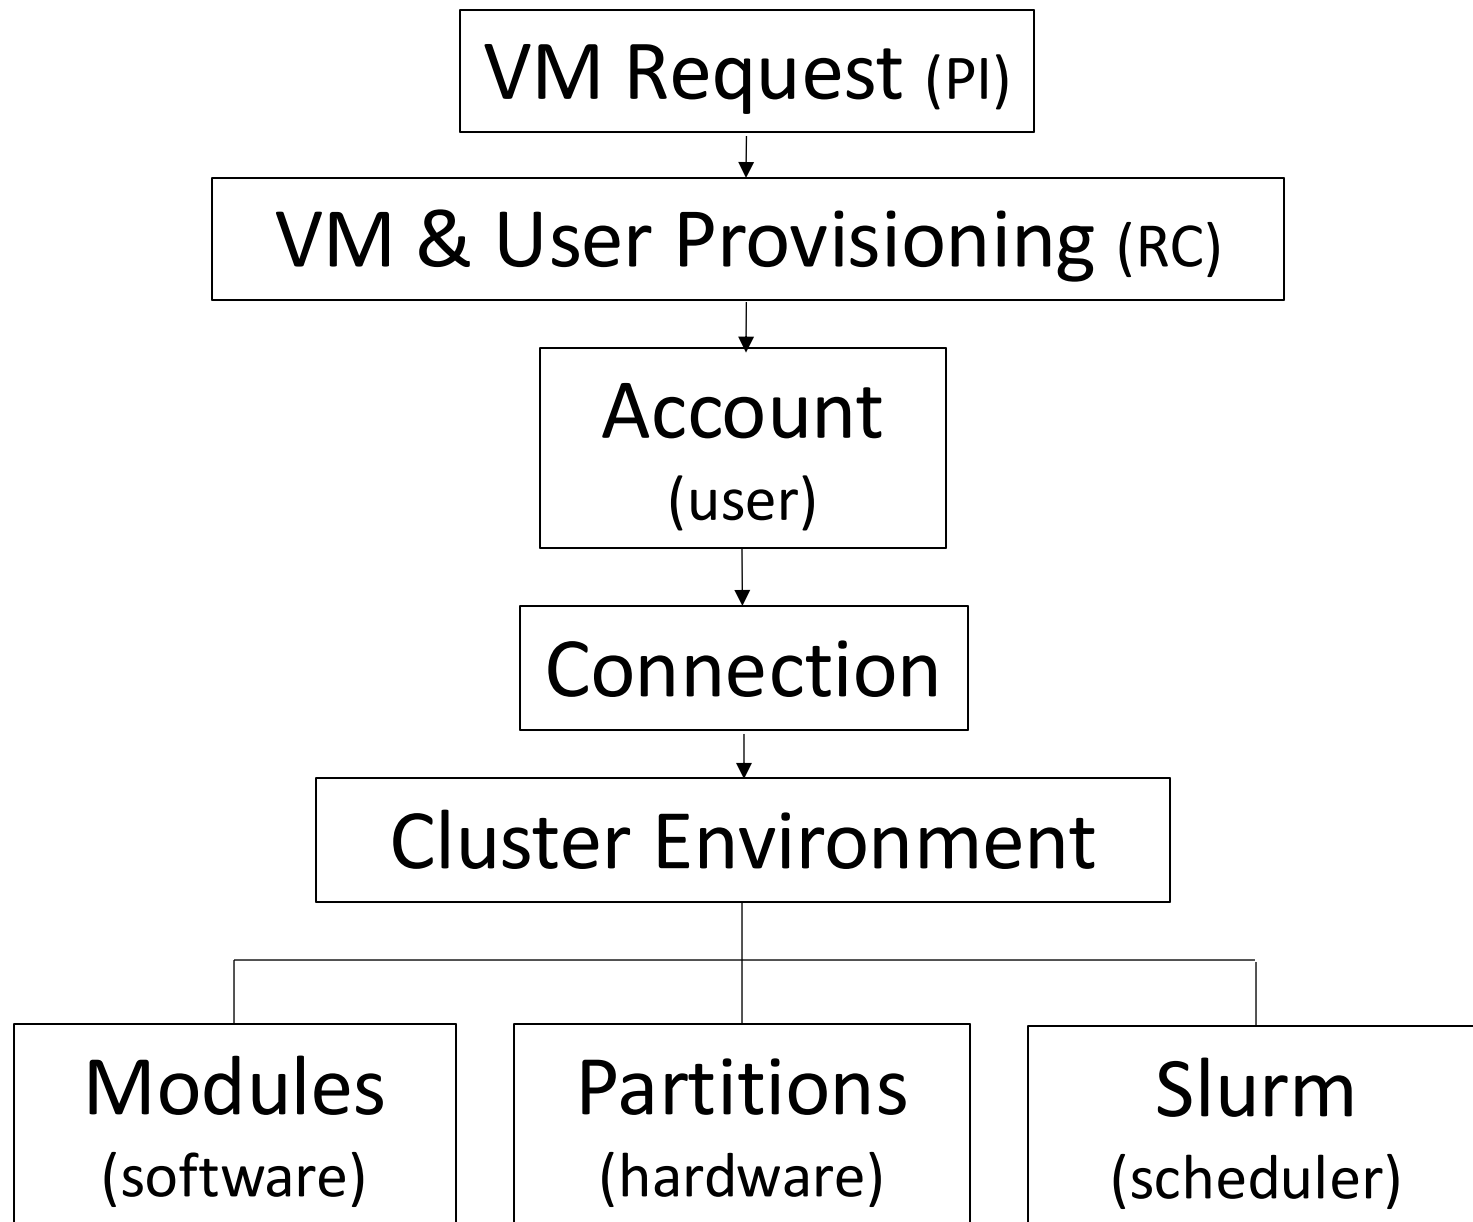
- Secure computing environment, offering Virtual Machines (VM) with both Linux and Windows operating systems.
- Ivy VM is designed for interactive and small-scale analysis of highly sensitive data including HIPAA, FERPA, and CUI.
- Access to Ivy platform is project based. Each project gets a VM.
- Project-specific storage volumes are mounted to the VM and available on the HPC system enabling a seamless transition between tasks on VM and HPC environment.
- Ivy Linux VMs can serve as a frontend for accessing the Rio HPC system (available upon request).

Rio

- For large scale analysis of HIPAA, FERPA, and dbGaP data

Terminology

- Node
 - Nodes are a type of computer called a server
 - They generally have more power than a typical computer.
 - They may have specialty hardware like graphical processing units
 - Two types of Nodes:
 - **Login node** - a server used for logging in and submitting jobs.
 - **Compute node** - a server that carries out the computational work.
- Core
 - An individual processor on a computer
- Memory
 - The random-access memory on a node
- Storage
 - Disk storage visible from a node



Requesting Access to HSZ HPC (Rio)

- Principal Investigators (PIs) can request access by submitting an Ivy Linux VM request through the [services app](#)
 - Research staff may apply for an exception.
 - Students must be sponsored by a faculty or research staff.
 - In addition to being a login node to Rio, a Linux VM can be used for preliminary computations as well
 - VMs comes in different sizes. Request a VM sized appropriately for your workflow. For larger groups or projects with computationally intensive tasks, choose a larger VM. We recommend Small and above.
 - VMs exist in a private, secure network and cannot reach outside resources on the Internet. Most inbound and outbound data transfer is managed through the Data Transfer Node (DTN).
 - More information on requesting a VM can be found [here](#) on our website.

VM Provisioning

After access is requested, the Research Computing (RC) team will handle the provisioning of the Virtual Machine (VM).

The following steps are carried out by the RC team

- RC will create a Grouper group for Rio access
- RC will request access to the High Security VPN (HSVPN) on your group's behalf and provision 1 TB of High Security Research Standard storage for your project. This storage volume will also be available when using Rio HPC for your project.
- The PI and researchers will be notified upon the completion of the VM provisioning, and the associated IP address and group name will be provided to them.

User Provisioning

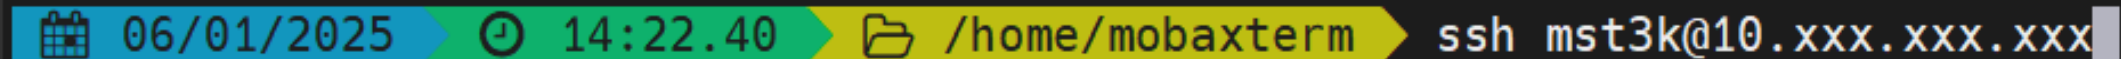
- Any users of the VM must be added to the project in the [services app](#) by the PI
- Once a project is approved, a PI and her/his researchers must sign a RUDA (one for every researcher on each project).
 - Each researcher must also complete the [High Security Awareness Training \(HSAT\)](#)
 - This must be completed to permit each researcher access to the HSVPN filter
- More information on security training and HSVPN access can be found [here](#) on our website

Connecting to the System

- Before connecting to your VM, you must install the following on your personal machine/device.
 - Install the Cisco AnyConnect Secure Mobility Client.
 - Install Opswat.
 - Install Duo MFA on personal smartphone.
 - See [here](#) for details.
- There are two ways to connect to your Linux VM:
 - Accessing the VM from a browser window
 - Using SSH (Secure Shell)
 - Use a terminal window with Mac/Linux
 - Windows users will need to install an SSH client application. We recommend [MobaXTerm](#).
 - In either case, you need the IP address of your VM along with your Netbadge credentials to connect to the VM.

Accessing from SSH

- Start the HSVPN
- Open your SSH client and type: `ssh mst3k@10.xxx.xxx.xxx`, where mst3k is replaced with your user ID and the x's are replaced with your VM's IP address (given in the Services app).

A screenshot of a terminal window with a dark background. The top bar shows a calendar icon, the date '06/01/2025', a clock icon, the time '14:22.40', a folder icon, and the path '/home/mobaxterm'. The main area shows the command 'ssh mst3k@10.xxx.xxx.xxx' with a cursor at the end.

```
06/01/2025 14:22.40 /home/mobaxterm ssh mst3k@10.xxx.xxx.xxx
```

- When prompted for a password, use your EServices password.

This login method provides only command line access to the VM.
No graphical user interfaces (GUIs) are usable.

SSH Clients

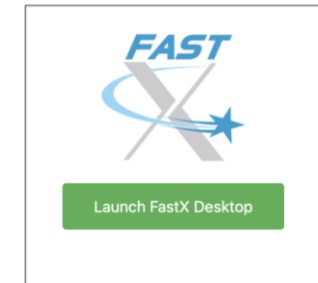


- SSH clients mostly use instructions typed at the command prompt.
- You will need to become familiar with Linux/Unix commands.
 - See learning.rc.virginia.edu/tutorials/unix-tutorial/ to learn more
- Some useful Unix/Linux commands:

```
$ ls  
$ pwd  
$ cd folder_name  
$ cp file_1 file_2  
$ rm file_1  
$ cd ..
```

Accessing from a Browser

- Start the HSVPN
- Open a web browser and enter the IP address for your VM (e.g., <https://10.xxx.xxx.xxx>)
 - If you get a warning message, you may need to click on Advanced Settings and/or a Connect Anyway option, depending on your web browser.
- Upon login, you will see a selection of different graphical applications that you can use.

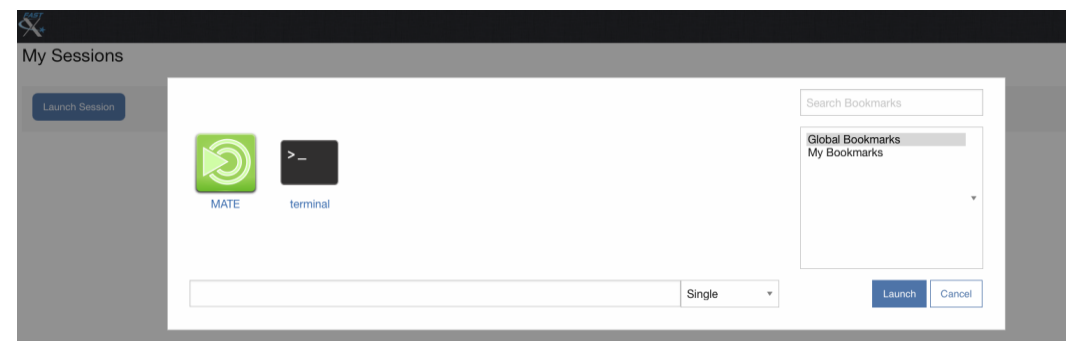
- Jupyterhub
- RStudio
- FastX Desktop



- These applications run on the VM itself and do not use HPC resources.

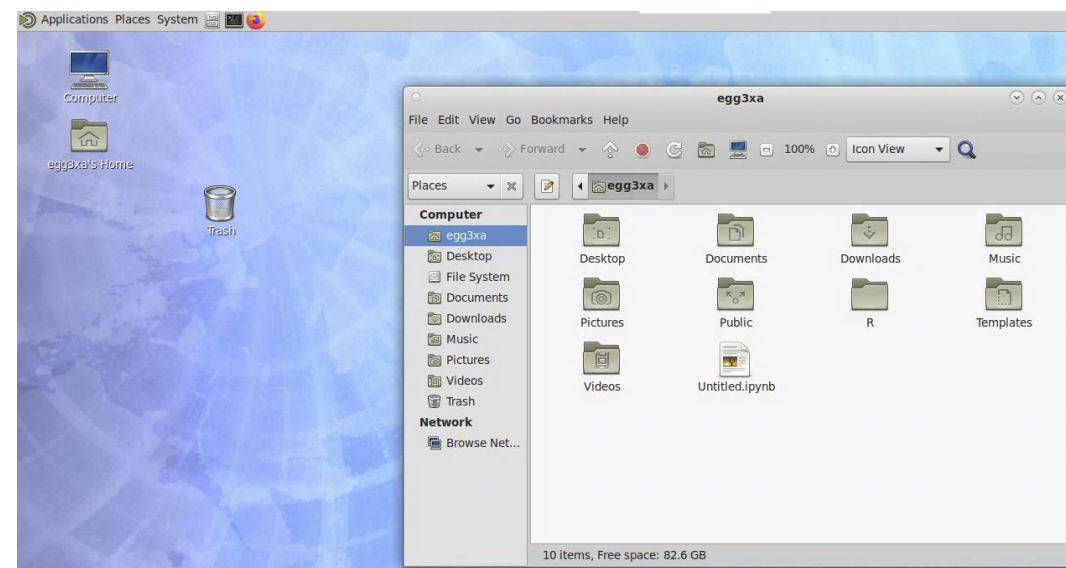
Working with Files Using FastX

- FastX Desktop is a great tool for managing files and launching applications with a graphical user interface.
- Once signed-in, you can select either a "MATE" (GUI Linux Desktop) or "Terminal" session
- MATE sessions provide a terminal (black box in menu bar) and a file browser (white box) in menu bar. The CAJA file browser can help with management and navigation of the filesystem.
- More information on using FastX can be found on [our website](#) and in our [Intro to HPC slides](#).



FastX Dashboard

MATE Desktop with CAJA file browser open (white box on menu bar)



How to edit a file in Linux VM

Editors you can use:

- When in a Terminal Window, use vi/emacs/pluma.
 - vi/emacs uses all keyboard shortcuts and may take a while to learn.
 - pluma is a very basic editor. To launch, type: `pluma myfile.txt`
- In Desktop, you can use Pluma Text Editor.

Storing Data

- VMs do not come with any significant disk storage of their own.
- A PI specifies the storage space they would like to have when requesting access to Ivy.
 - Default of 1TB of Research Standard Storage
 - Mounted under */data/ivy-hip-name* where “ivy-hip-name” is replaced by the name of your ivy project’s Grouper group name
 - Storage can be resized upon request. Submit your request [here](#)
- JupyterLab, RStudio, Rio, and FastX can all save data into the shared storage at */data/ivy-hip-name*.

Transferring Data

- Globus is the only permitted data-transfer protocol for highly sensitive data.
- ***UVA IVY-DTN*** is the the official collection for moving files into High-Security Research Standard Storage.
 - To transfer data to High-Security Research Standard Storage, please see the special instructions [here](#) on our website.
 - Ensure that you are **NOT** connected to the HSVPN. Data transfer will not work if you are connected to the HSVPN.
- Additional instructions and information on using Globus can be found [here](#) on our learning site.

Preinstalled Software

- Every VM comes with a base installation of software by default. See [here](#) for a list of preinstalled software and how to request new software installations.
 - Additional software packages are pre-approved and available for installation upon request.
 - Software not already approved will need to undergo a security evaluation to determine if it is suitable for a High Security environment.
 - Python and R packages are available to users through the normal pip, conda, and CRAN library installation methods.
- Access to installed software is done using the **lmod** Module System.

Modules

- Any application software that you want to use will need to be loaded with the **module load** command.
- For example:
 - `module load apptainer/1.3.1`
- You will need to load the module any time that you create a new shell
 - Every time that you log out and back in
 - Every time that you run a batch job on a compute node

Modules Basics

Some basic commands to use lmod:

- `module avail` – lists every module available to load
- `module key <keyword>` - searches for modules in a category specified
 - Example: `module key bio`
- `module spider <module-name>` - prints information about the software including different offered versions
- `module load <module-name>` - load the desired software
 - Use `module load <module-name/ver-number>` to load a specific version
- `module unload <module-name>` – unloads the desired software
- `module list` – prints all currently loaded modules
- `module purge` – unloads all modules

Partitions (Queues)

- Rio has several partitions (or queues) for job submissions.
 - You will need to specify a partition when you submit a job.
 - To see the partitions that are available to you, type **qlist** at the command-line prompt.

Queue (partition)	Total Cores	Free Cores	Jobs Running	Jobs Pending	Time Limit
neo	960	836	0	0	7-00:00:00
neo-gpu	104	104	0	0	7-00:00:00
standard	1520	1200	0	0	7-00:00:00
gpu	96	96	0	0	3-00:00:00

- The neo partitions are exclusive to a specific group and are not for general users of Rio.

Slurm

- Slurm manages the hardware resources on the cluster (e.g. compute nodes/cpu cores, compute memory, etc.).
- Slurm allows you to request resources within the cluster to run your code.
 - It is used for submitting jobs to compute nodes from an access point (Ivy VM).
 - Frontends are intended for editing, compiling, and very short test runs.
 - Production jobs go to the compute nodes through the resources manager.
- Slurm documentation:
 - <https://www.rc.virginia.edu/userinfo/hpc/slurm/>
 - <https://slurm.schedmd.com/>

Slurm Script Example

- A Slurm script is a bash script with Slurm directives (#SBATCH) and command-line instructions for running your program.

```
#!/bin/bash
#SBATCH --nodes=1                #total number of nodes for the job
#SBATCH --ntasks-per-node=1     #how many copies of code to run
#SBATCH --time=1-12:00:00       #amount of time for the whole job
#SBATCH --partition=standard     #the queue/partition to run on
#SBATCH --account=myGroupName   #the account/allocation to use

module purge
module load goolf R             #load modules that my job needs
Rscript myProg.R                #command-line execution of my job
```

Submitting a Slurm Job

- To submit the Slurm command file to the queue, use the **sbatch** command at the command line prompt.
 - /home only exists on VM. Compute nodes do not have access to /home.
 - Job submission needs to be done from /standard/storage/
 - Recommended to include the complete file paths in the script
- For example, if the script on the previous slide is in a file named job_script.slurm, we can submit it as follows:

```
[jus2yw@ivy-tst-rc-1 ivy-tst-rc]$ sbatch job_script.slurm  
Submitted batch job 18316
```


Submitting an Interactive Slurm Job

- To submit an interactive SLURM job to the queue, use the **ijob** command at the command line prompt.
- For example, if you want to run an interactive application on a compute node in the standard queue using one cpu, we can submit it as follows:

```
-bash-4.1$ ijob -c 1 -p standard -A MyGroupName -t 06:00:00
salloc: Pending job allocation 21640112
salloc: job 21640112 queued and waiting for resources
salloc: job 21640112 has been allocated resources
salloc: Granted job allocation 21640112
srun: Step created for job 21640112
udc-aw34-21c0-teh1m$
```

Checking Job Status

- To display the status of only your **active** jobs, type: **squeue -u <your_user_id>**
- The squeue command will show pending jobs and running jobs, but not failed, canceled or completed job.

```
-bash-4.2$ squeue -u mst3k
```

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST (REASON)
18316	standard	job_sci	mst3k	R	1:45	1	udc-aw38-34-1

- A job's status is indicated by:
 - PD – pending, R – running, CG – exiting

Checking Job Status

- To display the status of all jobs, type: `sacct -S <start_date>`

```
[jus2yw@ivy-tst-rc-1 ivy-tst-rc]$ sacct -S 2025-01-01
```

JobID	JobName	Partition	Account	AllocCPUS	State	ExitCode
-----	-----	-----	-----	-----	-----	-----
3104009	RAxML_NoC+	standard	hpc_build	20	COMPLETED	0:0
3104009.bat+	batch		hpc_build	20	COMPLETED	0:0
3104009.0	raxmlHPC-+		hpc_build	20	COMPLETED	0:0
3108537	sys/dashb+	gpu	hpc_build	1	CANCELLED+	0:0
3108537.bat+	batch		hpc_build	1	CANCELLED	0:15
3108562	sys/dashb+	gpu	hpc_build	1	TIMEOUT	0:0
3108562.bat+	batch		hpc_build	1	CANCELLED	0:15
3109392	sys/dashb+	gpu	hpc_build	1	TIMEOUT	0:0
3109392.bat+	batch		hpc_build	1	CANCELLED	0:15
3112064	srun	gpu	hpc_build	1	FAILED	1:0
3112064.0	bash		hpc_build	1	FAILED	1:0

- The `sacct` command lists all jobs (pending, running, completed, canceled, failed, etc.) since the specified date.

Canceling a Job

- To delete a job from the queue, use the **scancel** command with the job ID number at the command line prompt:

```
-bash-4.2$ scancel 18316
```

- To cancel all your jobs, run this command:

```
-bash-4.2$ scancel -u $USER
```

- We use scontrol to print information of a running job
 - scontrol show job <jobid>

More Commands

- `sacct` will return accounting information about your job. See [Slurm's documentation](#) for a full list of options.
 - Use the option `-j <jobid>` to inspect a particular job
- `seff` will return information about the utilization (called the “efficiency”) of core and memory.

NEED HELP?

Research Computing Zoom Office Hours

Tuesdays: 3 pm – 5 pm
Thursdays: 10 am – noon

To connect to the Zoom sessions,
go to
<https://www.rc.virginia.edu/support/#office-hours> and click on the
“Join us via Zoom” button

Or, contact us through the forms at:

<https://www.rc.virginia.edu/support/>