

Introduction to Odyssey & RC Services

Bob Freeman, PhD
Research Computing Facilitator
XSEDE Campus Champion

robert_freeman@harvard.edu
@DevBioInfoGuy

Plamen Krastev, PhD
Research Computing Associate
(Res. Computational Scientist)

plamenkrastev@fas.harvard.edu

Locally

- Help you to understand the complexities of our services
- Help frame our role in the services landscape
- Submit jobs without impacting 800+ other active users
- Use compute resources efficiently
- Reduce the support burden on the RC/Informatics team

Globally

- “Work smarter, better, faster”
- ... and to Think Differently
- Demonstrate leadership in advanced technologies
- Foster collaborations across institutional boundaries
- Enable you to be successful with your research!



All services are free except additional storage & private compute

- PLEASE MAKE USE of our FREE services
- Did I mention our services are free?
- Most of this information is documented on our website, and reference URLs will be posted on most, if not all, slides
- Please ask questions (I and others here may learn from your inquiries)



- RC Services/People
- Getting/changing your account
- Login & access
- VPN to where?
- Storage/Mounting disk shares
- NX clients
- Other services
- Beyond RC: XSEDE

20 minutes

- HPC & All About Odyssey
- Typical Workflow
 - Login & Access
 - Filesystems & Storage
 - Transferring Files
 - Loading Software
 - Login/Interactive Nodes
 - Choosing Appropriate Resources
 - Submitting/Controlling Jobs
- Best Practices
- Consulting/Training
- Getting Help

60 minutes



Odyssey (cluster)	~60K load-balanced cores (CPUs) and increasing <i>high-throughput / high-performance computing</i>
Storage	14 PB private, public, and scratch storage; and increasing <i>Lab and personal data; scratch/work temporary space</i>
Visualization	65+ GPGPUs and increasing <i>high-end visualization; real-time interactive rendering</i>
Virtual Machines	~600+ KVM images <i>Web portals, license or server-based resources; restricted datasets</i>
Apps/Libraries	>3000, in chemistry, biology, statistics, social sciences, & others <i>Predominantly cluster-based & open source; some desktop/licensed</i>
Other	Application licensing Code Optimization Instrumentation Core facilities support
HUIT	Networking, student computing, admin apps, & desktop support



People of Research Computing



- 14 people (leadership, cloud, sysadmin, client support)
- Strong overlap with Informatics group (10 in leadership, web, science, sci/code)
- Ph.D.s in biology and physics, w/ roles as sysadmins and developers
- Supporting ALL sciences & major HU computing projects (BICEP2, ATLAS, etc.)

Informatics Group



Director: Michele Clamp
mclamp@g.harvard.edu

Senior Web Developers



Reuven Koblick
reuven_koblick@harvard.edu



Post-Docs and Junior Fellows

Tim Sackton (with Edwards lab)
mattison@g.harvard.edu



Bioinformatics

George Marnellos
gmarnellos@fas.harvard.edu



Chris Williams
williams03@g.harvard.edu



RC and Honorary Informatics

Aaron Kitzmiller
akitzmiller@g.harvard.edu



Bob Freeman
robertfreeman@g.harvard.edu

Core-associated:

- Primary sequence analysis i.e. basecalling/multiplexing
- Sequence data QC/trimming
- Alignment to reference genome
- Bioinformatics software installation/support
- Consultation and recommendations for amount and type of sequencing
- Consultation and recommendations for sequence analysis

Non-core activities:

- Transcriptome assembly and differential expression analysis
- Whole genome assembly and genome comparison
- Annotation of assembled transcriptomes
- Database construction
- SNP calling and filtering
- Custom script/algorithm writing
- Phylogenetic analysis

We are happy to meet and help people get up and running with these things.




Getting/Changing Your RC Account

- RC account access to all RC services
- If you did not select the appropriate services up-front (Instrumentation or Odyssey)
 - submit a ticket thru <https://portal.rc.fas.harvard.edu>
- Lock automatically if 5 failed login attempts
 - Unlock automatically after 10 minutes
- Resetting your own password
 - <https://portal.rc.fas.harvard.edu>
- If you are switching labs, please notify us!



Login & Access

- Use your RC credentials for:
 - Logging into Windows machines
 - VPN
 - File Transfer clients
 - Mounting disk shares
 - Terminal sessions to Odyssey
- OpenAuth required for VPN & Odyssey sessions (file transfer & terminal)



Account credentials should not be shared!

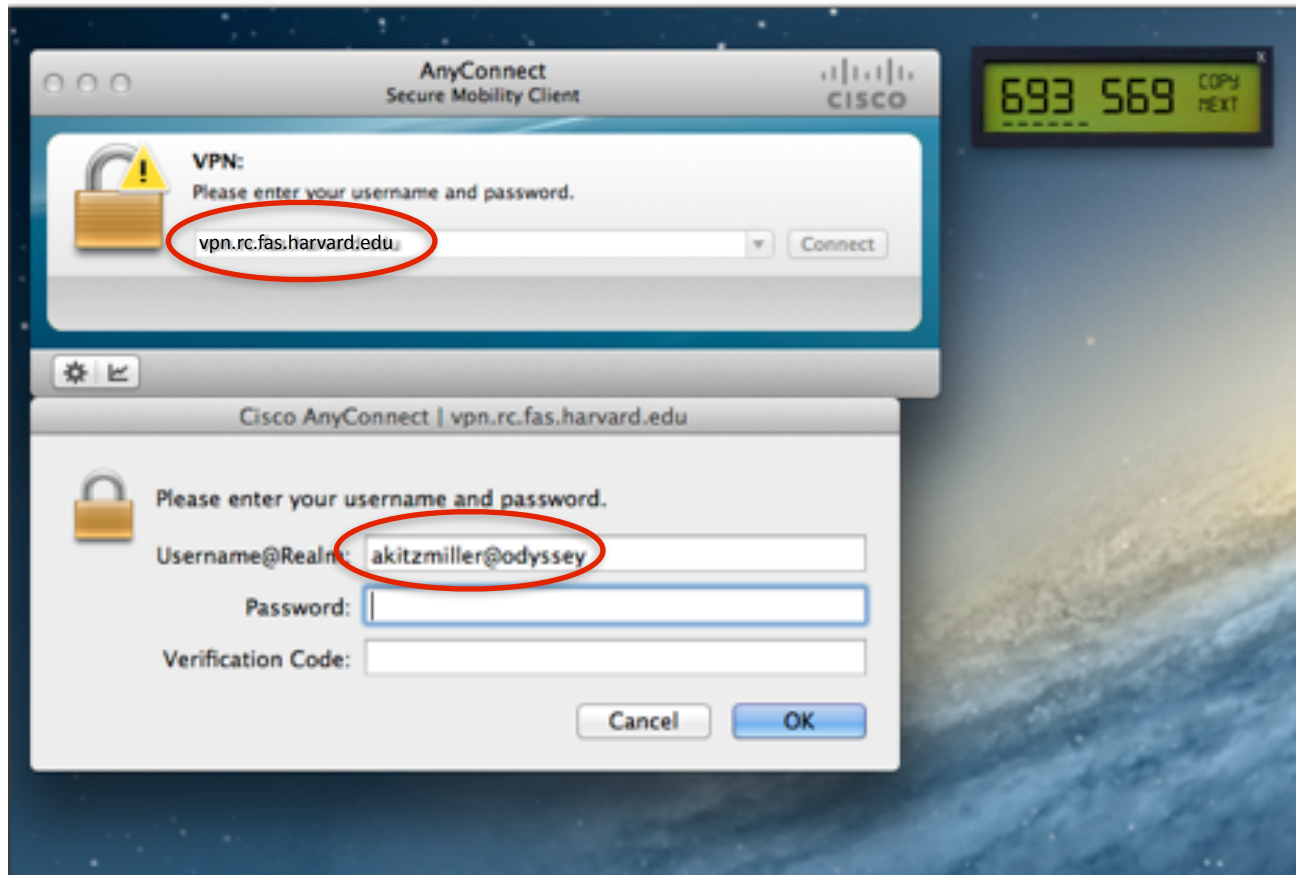
Using RC services in an explicit acceptance of the University Security Policy at

<http://security.harvard.edu/book/information-security-policy>

<https://rc.fas.harvard.edu/resources/access-and-login/>

VPN Is Sometimes Required

- Mounting disk shares using WiFi or from off-campus
- X11/GUI sessions to our GUI login server
- Sessions must be to vpn.rc (not vpn.fas) and within RC domain (@odyssey)*



* exception is for HRCI domains!

<https://rc.fas.harvard.edu/resources/vpn-setup/>

- Lab space fairly common and will vary by lab/department
- Typically this space is backed up. **We make it very clear if not!**
- If Odyssey account:
 - Personal space, 40 GB
 - New labs, 1 TB free space
- Purchased storage available
- Details at <https://rc.fas.harvard.edu/resources/storage/>

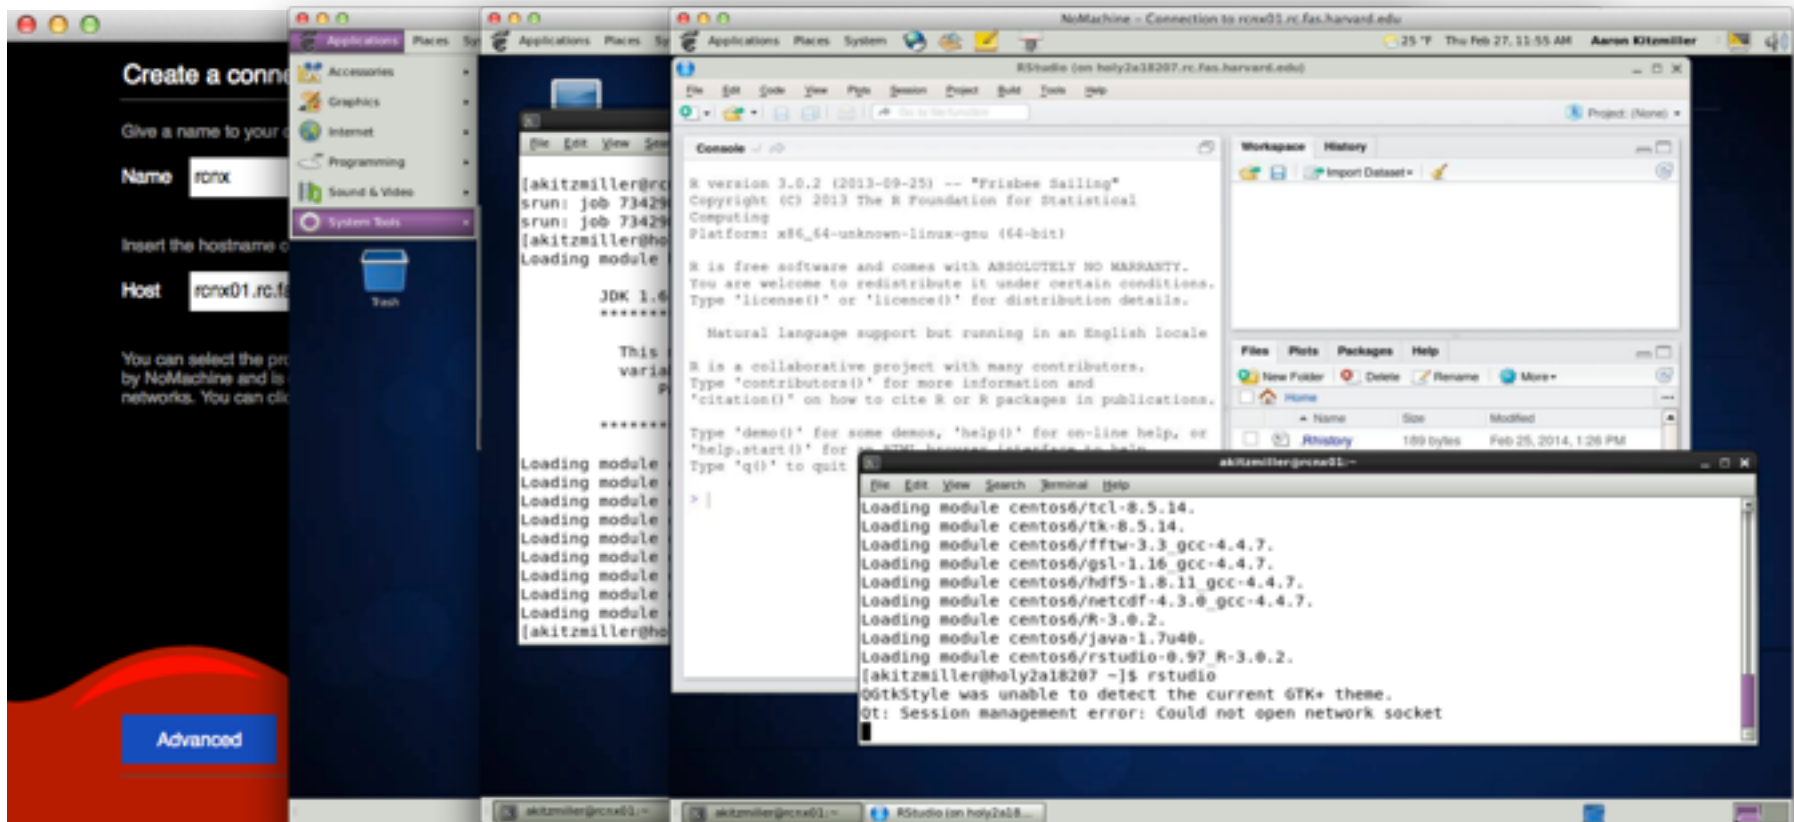
Mounting Disk Shares

- Lab & home folders can be accessed as mounted volumes on Mac, Windows, and Linux OS
- VPN **required** if using WiFi or if off-campus
- `\\rcstore.rc.fas.harvard.edu\homes\home01\bfreeman`
- `smb://rcstore.rc.fas.harvard.edu/homes/home01/bfreeman`
- `\\rcnfs11.rc.fas.harvard.edu\newton_lab\`
- `smb://rcnfs11.rc.fas.harvard.edu/newton_lab`
- More info can be found in “Access and Login” doc:
<https://rc.fas.harvard.edu/resources/access-and-login/>



GUI/X11 Clients

- Some apps require GUI/X11 interface: MATLAB, RStudio, CLCBio, etc..
- X11 performance is sluggish. Use NoMachineX instead!
- This is a login node. Grab an interactive session to compute!



http://wp.me/P42YvN-2Gj/#Consider_an_NX_remote_desktop_for_graphical_applications_like_Matlab_and_RStudio

- Virtual machines:
 - Project- or portal-based web sites
 - License servers
 - Server-based resources
- Special compute setups
 - HRCI
 - Restricted access for datasets, classes

Please contact us!

Beyond RC: XSEDE

- Extreme Science & Engineering Discovery Environment
- Umbrella organization for collection of integrated advanced digital resources and services
- 5-year, \$120m project funded by NSF & successor to TeraGrid
- Supports 16 supercomputers and high-end visualization and data analysis resources across the country
- Campus champion exists to help you apply for and transition to these off-campus resources
- <https://www.xsede.org/>

<https://rc.fas.harvard.edu/partnerships/xsede/>



- RC Services/People
- Getting/changing your account
- Login & access
- VPN to where?
- Storage/Mounting disk shares
- NX clients
- Other services
- Beyond RC: XSEDE

- HPC & All About Odyssey
- Typical Workflow
 - Login & Access
 - Filesystems & Storage
 - Transferring Files
 - Loading Software
 - Login/Interactive Nodes
 - Choosing Appropriate Resources
 - Submitting/Controlling Jobs
- Best Practices
- Consulting/Training
- Getting Help





From **Profiles in Delight: Paul Edmon** (ITC RC Assoc):

What's the biggest misconception about RC or HPC in general?

That if you just put your code on the supercomputer it will run faster. As it turns out the processors we use on the cluster are not much better than what you have in your desktop. At times your desktop may be faster. **What makes HPC work is that we have a vast number of these processors all networked together with a high speed interconnect.** Not even sending it to the cloud will get you that.

In order to get the most out of your code and leverage any HPC resource (whether it be ours, the cloud, or [XSEDE](#)) **you need to optimize your code and workflow. This takes time and effort.** You need to learn about the hardware you are running on, the code you are running, the science you need to get done, and marry all that together to make sure you get things done as quickly and accurately as possible. **Supercomputing isn't a blackbox, and the more you understand the better you can engineer your workflow and code to take advantage of the great resources we have available. We at RC are here to help people achieve that.**

<https://rc.fas.harvard.edu/profiles-delight-paul-edmon/>



High Performance Computing

Old Way: Individual groups maintain their own resources

- Takes time away from research
- Much reinventing of the wheel
- Does not scale well

New Way: RC and HUIT provide resources

- Physical infrastructure: power, cooling, and space
- System administration, algorithm/programming support, & technical advice
- Consolidation of knowledge and expertise
- Economies of scale



Value to you??

- Groups can purchase dedicated hardware or use the general resources
- Leads to better and more efficient use of resources:
 - Easier entry into the world of HPC
 - Groups who own resources get priority on the resources they own
 - Others get to make use of the spare cycles when the resources are not in use

Who can benefit from HPC?

- If your analysis takes longer than 10 min on your desktop/laptop
- If you're doing tens/hundreds of tasks
- Analyses using a GUI that could be done at the command line

As HU pays for infrastructure costs whether machines are busy or idle, fully utilized machines are a very good thing!



What is Odyssey?

- Our premier resource is the Odyssey cluster
- A collection of various types of hardware:
 - *A **couple large sets** of tightly interconnected machines with identical hardware*
 - *Several high-powered individual machines and other special purpose machines*
 - *Large amount of **shared storage***
 - *Miscellaneous supporting cast of other servers*
- And some software:
 - **User and group management** and separation of resources
 - **SLURM** (Simple Linux/Unix Resource Manager)
 - Linux OS (CentOS 6)

It's a **shared** system! You are not alone!

<https://rc.fas.harvard.edu/odyssey/>



What is Odyssey?

The typical hardware unit is called a **node**

- Same stuff that's in a desktop/laptop: Motherboard, CPU(s), Memory, Hard drive(s)
- But more powerful and/or more of them compared to a typical desktop
- Nodes are individual hosts with names like rlogin03 or holy2a18208
 - names indicate function and/or location

The basic computational unit in a cluster is a **CPU core**


- Each core runs one process: a job
- Most compute nodes have 4 CPUs and 16 cores/CPU -> 64 cores total
- Cores/node share other resources of the node: memory, network bandwidth, etc.
- Thus, most nodes run 64 batch job processes*

A **typical compute node** is configured:

- 4 CPUs (64 cores)
- 256 GB RAM
- 2 network cards: Infiniband & xGb connections
- Small, local hard disk/SSD for boot and local /scratch



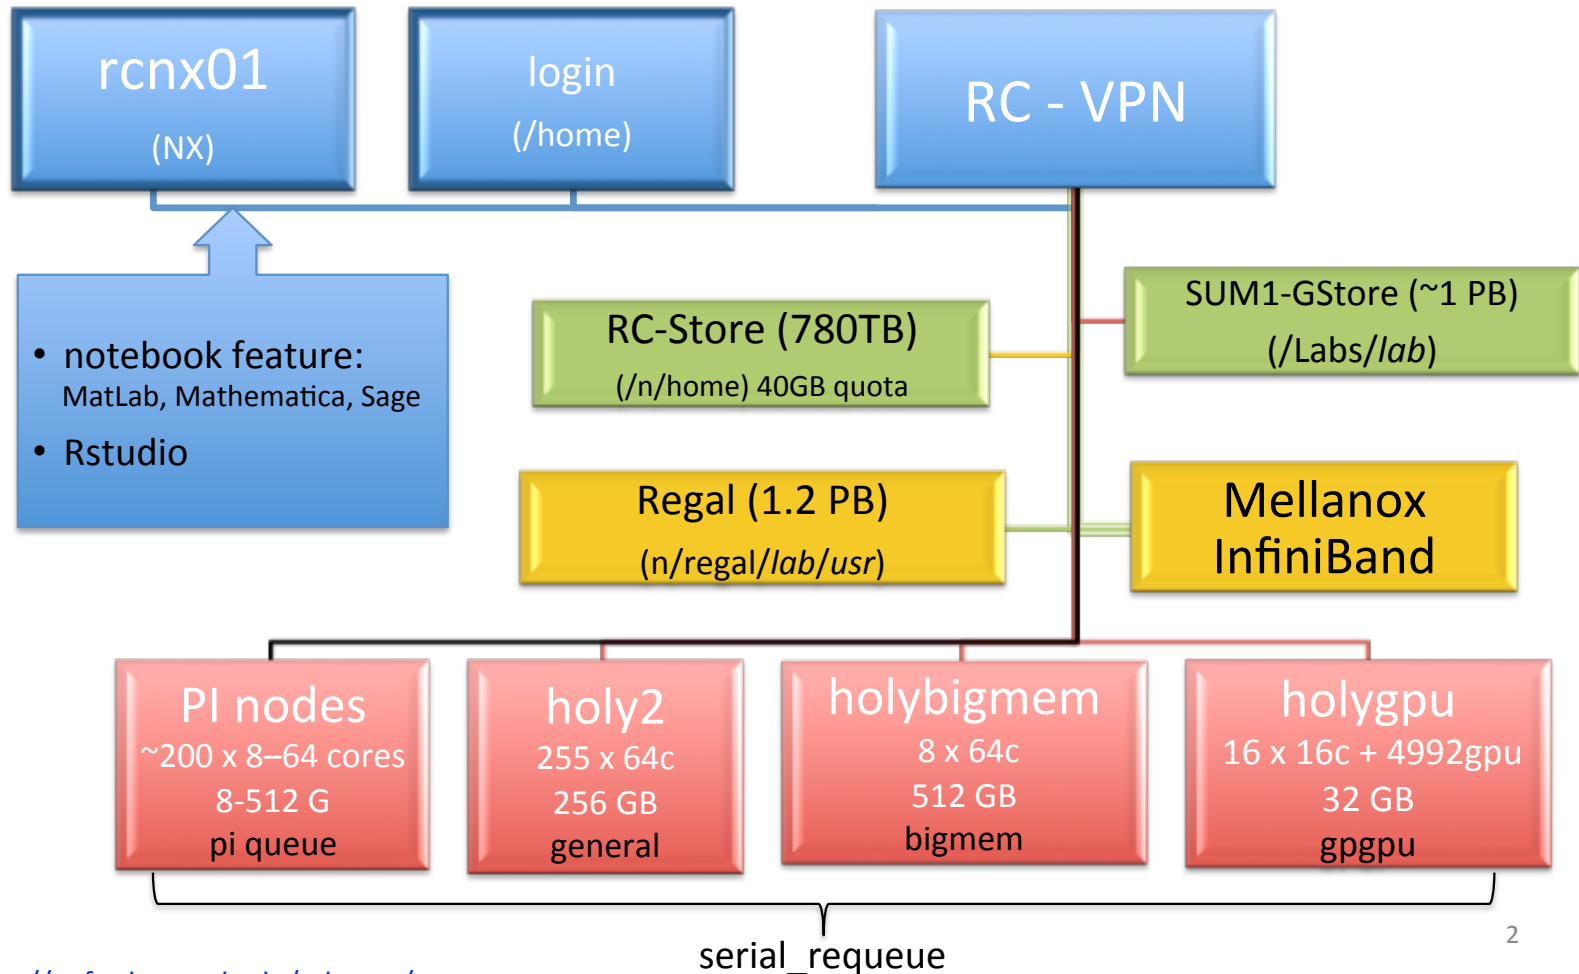
What is Odyssey?



What is Odyssey?



Odyssey 2.0 Topology



<https://rc.fas.harvard.edu/odyssey/>




What is Odyssey?

Compute nodes/disk are located in 3 data centers:

- 60 Oxford St
Personal home folders
Legacy equipment
- 1 Summer Street
Lab disk shares
Compute nodes < 2012 (20K+ cores)
- Holyoke, MA
Compute nodes > 2012 (33K+ cores)
'regal', scratch filesystem

Topology may effect the efficiency of work!




Typical Workflow

1. Login in to Odyssey - you land on a login (head) node
2. Copy/place some files
3. Get interactive session
4. Brief test run of a program or script at the command line
test to ensure it runs properly
5. Brief test run in batch: create batch file & submit to SLURM
continue doing other work while waiting for the results
6. Scale up as necessary (10s, 100s, 1000s)

Login & Access

- Use your RC credentials for:
 - Terminal on Mac/Linux
 - Putty on PC
- OpenAuth required for VPN & Odyssey sessions (file transfer & terminal)
- login.rc.fas.harvard.edu*
- Good Unix chops are required!

*unless in HRCI compute domain



<https://rc.fas.harvard.edu/resources/access-and-login/>

Getting Good Unix Skills

We **expect** you to have the following Unix skills before using Odyssey:

- Know how to traverse directories
- Know how to create & remove files & directories
- Know how to copy/move files & directories
- Know how to edit a file using a simple text editor like nano, vi, or emacs
- Read and write files with commands like head, tail, cat, and less
- Understand and modify unix file/directory permissions
- Understand I/O redirection
- Have some basic understanding of \$PATH and what login files are for

If you cannot perform all these functions, please review the following materials:

- RC's John Brunelle's ComputeFest2014 Unix Intro
https://software.rc.fas.harvard.edu/training/workshop_intro_unix
- Unix screencast from our ACI-REF collaborators at UClemson
<http://citi.clemson.edu/training/> (links in center of page)
- HUID account holders have access to all of Lynda.com's training:
<http://ohrdb.harvard.edu/cwdcourses/description.php?recordID=324>
- Software Carpentry's 'helping lab people compute smartly' Unix Intro:
<http://software-carpentry.org/lessons.html>



The virtual filesystem:

- Note that there are no designations above the root directory, like C:\ and D:\ drives, http:// protocols, etc.
- Physically distinct storage locations are just different directories in the filesystem
- The filesystem is one big virtual filesystem containing these individual filesystems

The distinctions between individual filesystems matter:

- local storage vs. network storage
- performance vs. reliability

memory is **not** storage



Filesystems & Storage

- Personal (home) folders (\$HOME)
 - 40 GB space, avail on all cluster nodes, can be mounted on desktops
 - backed up, no retention
 - OK for small I/O
- Lab folders
 - 1 TB free at start (contact for expansion costs); avail on all cluster nodes, can be mounted on desktops
 - most backed up (NOTED if not); no retention
 - OK for small I/O
- Local scratch (/scratch)
 - 250 GB space for *one node*, avail on all cluster nodes, *cannot* be mounted on desktops
 - no backup, deleted routinely
 - great for high IO, but must be shared with all processes on that node
- Network scratch (/n/regal)
 - 1.2 PB space, avail on all cluster nodes, *cannot* be mounted on desktops
 - no backup, 90 day retention
 - great for high IO

Please note that /n/holyscratch will be decommissioned soon

<https://rc.fas.harvard.edu/resources/odyssey-storage/>



Backups

If you accidentally delete a file from your *home* and *most lab* directories, you can often recover it from the checkpoint directory

- This is a directory named `.snapshot` that's not listed by `ls -a`!

```
[cfest350@rclogin03 ~]$ cd .snapshot
[cfest350@rclogin03 .snapshot]$ ls
rc_homes_daily_2014-01-07-_12-00    rc_homes_hourly_2014-01-14-_00-00
rc_homes_daily_2014-01-08-_12-00    rc_homes_hourly_2014-01-14-_01-00
rc_homes_daily_2014-01-09-_12-00    rc_homes_hourly_2014-01-14-_02-00
rc_homes_daily_2014-01-10-_12-00    rc_homes_hourly_2014-01-14-_03-00
rc_homes_daily_2014-01-11-_12-00    rc_homes_hourly_2014-01-14-_04-00
rc_homes_daily_2014-01-12-_12-00    rc_homes_hourly_2014-01-14-_05-00
rc_homes_daily_2014-01-13-_12-00    rc_homes_monthly_2013-11-01-_00-00
rc_homes_hourly_2014-01-13-_18-00    rc_homes_monthly_2013-12-01-_00-00
rc_homes_hourly_2014-01-13-_19-00    rc_homes_monthly_2014-01-01-_00-00
rc_homes_hourly_2014-01-13-_20-00    rc_homes_weekly_2013-12-22-_12-00
rc_homes_hourly_2014-01-13-_21-00    rc_homes_weekly_2013-12-29-_12-00
rc_homes_hourly_2014-01-13-_22-00    rc_homes_weekly_2014-01-05-_12-00
rc_homes_hourly_2014-01-13-_23-00    rc_homes_weekly_2014-01-12-_12-00
...
```

<https://rc.fas.harvard.edu/resources/faq/accidentally-deleted-data-get-back/>

Transferring files to/from Odyssey

All platforms

- FileZilla
bit.ly/Zd8fNt

Windows

- SecureFX
- WinSCP
<http://winscp.net/eng/download.php>
- PSCP (pscp.exe, from the makers of PuTTY)
<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

Mac/Linux/Unix

- Run scp or rsync from a terminal (example)
- (deprecated) Cyberduck
<http://cyberduck.io/>

<https://rc.fas.harvard.edu/resources/odyssey-quickstart-guide/#Transfer any files you may need>



Load/Installing software

- ~3000 applications/libraries in chemistry, biology, statistics, social sciences, and more
- Software is loaded incrementally using modules

Rolling out a new module system lmod

- New system is opt-in for old accounts, but will be default soon (or now!)
- <https://rc.fas.harvard.edu/resources/documentation/software-on-odyssey/>

```
source new-modules.sh           # for opt-in folks
module spider fastqc            # find software
module load fastqc              # most recent version
module load fastqc/1.0.0-fasrc01 # specific version
module load legacy              # get to old modules
module load centos6/bowtie-1.2.1
```

Old system: modules at <https://rc.fas.harvard.edu/resources/module-list/>

```
module avail 2>&1 | grep -i 'fastqc' # find software
module load centos6/fastqc-0.10.0
```

Docs for installing software yourself at <https://rc.fas.harvard.edu/resources/documentation/software-on-odyssey/installing-software-yourself/>

Please don't use sudo -- only RC staff have this privilege

Load/Installing software

- You can include module loads in your `.bashrc` login file
- *Caveat Emptor!* Too many, and library clashes may ensue!
 - crashes
 - programs won't run
 - aberrant behaviour

```
Loading module hpc/java-1.7.0_13.
Loading module hpc/perl-5.16.0.
Loading module hpc/intel-mkl-11.0.0.079.
Loading module centos6/tcl-8.5.14.
Loading module centos6/tk-8.5.14.
Loading module centos6/fftw-3.3_gcc-4.4.7.
Loading module centos6/gsl-1.16_gcc-4.4.7.
Loading module centos6/hdf5-1.8.11_gcc-4.4.7.
Loading module centos6/netcdf-4.3.0_gcc-4.4.7.
Loading module centos6/R-3.1.1.
Loading module centos6/python-2.7.3.
Loading module centos6/biopython-1.62b_python-2.7.3.
Loading module centos6/numpy-1.7.1_python-2.7.3.
Loading module centos6/gtk+-2.24.4.
Loading module centos6/cairo-1.12.14.
Loading module centos6/py2cairo-1.10.0_python-2.7.3.
Loading module centos6/matplotlib-1.3.1_python-2.7.3_gtk.
Loading module centos6/bx-python-3-14-2014_python-2.7.3.
Loading module hpc/glib-2.20.4.
Loading module hpc/pixman-0.17.4.
Loading module hpc/cairo-1.8.8.
Loading module hpc/pango-1.24.5.
Loading module hpc/atk-1.20.0.
Loading module hpc/gtk+-2.16.5.
Loading module hpc/python-2.7.3_glib.
Loading module hpc/networkx-1.7_python-2.7.3.
Loading module centos6/python-2.7.3.
Loading module centos6/hdf5-1.8.11_gcc-4.4.7.
Loading module centos6/numpy-1.7.1_python-2.7.3.
Loading module centos6/scipy-0.12.0_python-2.7.3.
Loading module centos6/gtk+-2.24.4.
Loading module centos6/cairo-1.12.14.
Loading module centos6/py2cairo-1.10.0_python-2.7.3.
Loading module centos6/matplotlib-1.3.1_python-2.7.3_gtk.
Loading module centos6/cython-0.19.1_python-2.7.3.
Loading module centos6/python-2.7_modules.
Loading module centos6/pandas-0.11.0_python-2.7.3.
Loading module bio/primer3-2.2.2-beta.
```



Login vs Interactive Nodes

- Terminal sessions to login.rc put you on 1 of 12 head nodes
- Appropriate for light computational work
- Switch to a compute node -- an interactive session -- for other work
- `srun --pty --x11=first --mem 1000 -p interact -t 0-6:00 /bin/bash`
- or submit batch jobs to SLURM

RCNX01 (X11 GUI) is a login node! Get an interactive session to do compute!

https://rc.fas.harvard.edu/resources/odyssey-quickstart-guide/#_or_an_interactive_job



Choosing Appropriate Resources

- Time: determined by your runs
best to shoot for minimum of 10 minutes / job
- Memory: won't know until after a trial run
use `sacct` command to get post-run job info
- Partition

<i>Name</i>	<i>Length</i>	<i>Size (cores)</i>	<i>Memory/node</i>
general	7 days	~14K	256 GB
interact	3 days	512 (8 nodes)	256 GB
unrestricted	no limit	512	256 GB
serial_queue	7 days*	30K+	varies (512 GB max)
bigmem	7 days	512	512 GB
(private)	no limit	varies	256 GB typical

https://rc.fas.harvard.edu/resources/odyssey-quickstart-guide/#_or_an_interactive_job

Submitting Jobs

create file `fastqc.slurm` in your favorite **text** editor*

Required
Recommended
Start
Required

```
#!/bin/bash
```

```
#SBATCH -p serial_requeue           # Partition to submit to (comma separated)
```

```
#SBATCH -n 1                        # Number of cores
```

```
#SBATCH -N 1                        # Ensure that all cores are on one machine
```

```
#SBATCH -t 0-1:00                  # Runtime in D-HH:MM (or use minutes)
```

```
#SBATCH --mem 100                   # Memory in MB (see also --mem-per-cpu)
```

```
#SBATCH -o hostname.out            # File to which standard out will be written
```

```
#SBATCH -e hostname.err            # File to which standard err will be written
```

```
#SBATCH --mail-type=ALL             # Type of email notification- BEGIN,END,FAIL,ALL
```

```
#SBATCH --mail-user=rmf@123.com     # Email to which notifications will be sent
```

```
source new-modules.sh; module load fastqc
```

```
fastqc --casava -o fastqc_reports A01_R1.pair.fastq.gz
```

* TextWranger/BBEdit on Mac
GEdit/NotePad+ on PC
nano, vi, emacs on Linux



Submitting jobs

```
$ sbatch fastq.slurm
```

```
Submitted batch job 29484165
```

OR

```
sbatch -p serial_requeue -n 1 -N 1 -t 0-1:00 \  
--mem 100 -o hostname.out -e hostname.err \  
--mail-type=ALL --mail-user=rmf@123.com \  
--wrap="rsync -av ~/test/some_set_of_files \  
       /n/regal/my_lab/bfreeman/high_IO_destination/"
```

Submitting jobs

"\" + return allows you to
continue on the next line.
Imperative that there are no
characters after the backslash!

```
sbatch -p serial_requeue -n 1 -N 1 -t 0-1:00 \  
--mem 100 -o hostname.out -e hostname.err \  
--mail-type=ALL --mail-user=rmf@123.com \  
--wrap="rsync -av ~/test/some_set_of_files \  
/n/regal/my_lab/bfreeman/high_IO_destination/"
```

two spaces here ensures
parameter separation and
readability



Controlling Jobs & Getting Job Info

`scancel` may become your best friend

<code>scancel JOBID</code>	# specific job
<code>scancel -u bfreeman</code>	# ALL my jobs
<code>scancel -u bfreeman -J many_blast_jobs</code>	# named jobs
<code>scancel -u bfreeman -p bigmem</code>	# ALL in partition

`squeue` gives info on currently running jobs

<code>squeue -u bfreeman</code>	# jobs for bfreeman
<code>squeue -u bfreeman --states=R wc -l</code>	# # of Running jobs

`sacct` gives current and historical information

<code>sacct -u bfreeman</code>	# jobs for bfreeman
<code>sacct -u bfreeman -p bigmem --starttime=9/1/14</code>	# same+bigmem partition
<code>sacct -j JOBID --format=JobID,JobName,MaxRSS,Elapsed</code>	# RAM usage!!

Check out Common SLURM Commands

<https://rc.fas.harvard.edu/resources/documentation/convenient-slurm-commands/>

SLURM, LSF, SGE, PBS/Torque rosetta stone

<http://slurm.schedmd.com/rosetta.pdf>



Advanced Topics on Odyssey

- FairShare
- Job dependencies
- Job arrays
- Embarrassingly parallel job
- MPI & OpenMP
- Parallel IO
- Parallel R
- Parallel MATLAB
- Parallel Python

Check out documentation at

<https://rc.fas.harvard.edu/resources/documentation/>

Example scripts at

https://github.com/fasrc/slurm_utils

Common Pitfalls

- PEND for very long time
Asking for very large resource requests (cores & memory) or very low fairshare score
- Quick run and FAIL...Not including `-t` parameter
no `-t` means shortest possible in all partitions == 10 min
- Asking for multiple cores but forgetting to specify one node
`-n 4 -N 1` is very different from `-n 4`
- Not specifying enough cores
`prog1 | prog2 | prog3 > outfile` should run with 3 cores!
- Causing massive disk I/O on home folders/lab disk shares
your work & others on the same filesystem slows to a crawl
simple commands like `ls` take forever
- Hundreds/thousands of jobs access one common file
your work & others on the same filesystem slows to a crawl
make copies of file and have jobs access one of the group
- Don't pack more than 5K files in one directory
I/O for your jobs will slow to a crawl
- Bundle your work into ~10 min jobs
Kinder for us, kinder for you, kinder for Odyssey
- Please understand your software -- look at the options!
(who knows what could happen??)
- Trying to `sudo` when installing software
Please don't -- we admin the boxes for you.



Common Problems

- FAILS and others usually require RC admin intervention
- node fails
- memory errors
- disk mounts fail or stuck
- something is stuck, but not sure what
- somebody over uses or over requests (disk or CPU)
usually manifests in unusually slow complete times

Submit at ticket


<http://portal.rc.fas.harvard.edu>

Troubleshooting Odyssey jobs slides


<https://rc.fas.harvard.edu/education/training/training-materials/>



Accounts by Domain




Monthly Job Counts by Domain



Job efficiency



Grid CPU Allocation vs Utilization last week



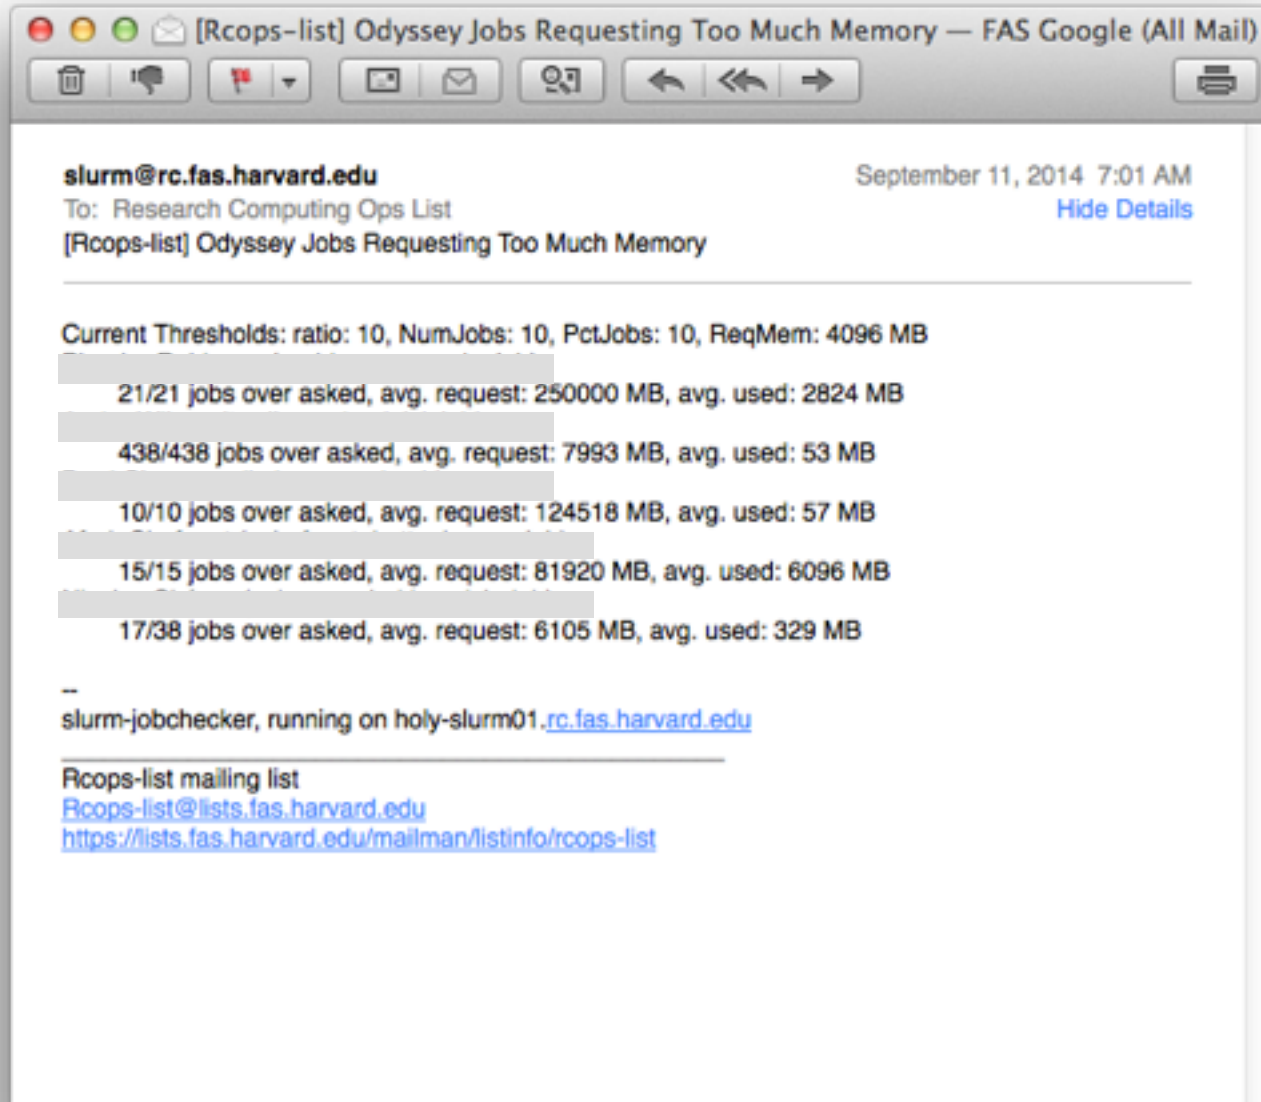
completed jobs

CPU wasters

top CPU-wasteful COMPLETED jobs in last 1 day, 0:00:00

user	job	CPU days wasted	CPU efficiency	cores allocated	
	18991789	610	5%	512	arun -n 512 \$exe -c auto-rockstar.cfg
	18955779	128	9%	64	Trinity.pl --seqType fq --JM 20G --SS_lib_type RF --lef
	19029202	58	0%	100	R CMD BATCH county.pd.parallel.R county.pd.parallel.Rout
	19013733	34	2%	50	R < test_model.r --vanilla;
	19003806	33	66%	64	sbatch \$0
	19003852	25	66%	64	sbatch \$0
	19018259	24	20%	40	mpirun -np 40 fdtd-engine-ompi-lcl \$1
	19012643	21	66%	64	sbatch \$0
	19019916	20	20%	40	mpirun -np 40 fdtd-engine-ompi-lcl \$1
	19011109	17	20%	40	mpirun -np 40 fdtd-engine-ompi-lcl \$1

Job efficiency



- Consulting

Given appropriate expertise, some significant engagements to help with code and analysis optimizations are possible

- Adv CyberInfrastructure – Research & Education Facilitation

Special 2-yr NSF grant for cluster computing education/facilitation

Partner with you to enable your research: use compute resources, work more efficiently, and be more competitive

Collaborations: Clemson U., U. Southern California, U. Utah, U. Wisconsin (Madison), U. Hawaii

Bob Freeman, PhD & Aaron Kitzmiller, PhD



Workflow Use Cases

- Transition to HPC from GUI
 - video processing workflow, all GUI driven
 - uses ImageJ and MATLAB, with 6 streams and ~10 steps
 - in progress, but down to 30 minutes from 12 hrs
- Video processing workflow
 - complex cluster pipeline for tracking animal movements
 - custom code developed by post-doc who left; no version control
 - in progress, adapting to current cluster, version control, & multicore processing
- Recoding of geophysical simulations
 - visco-elastic models for simulating earthquake movements
 - transitioned from MATLAB to fortran, & parallelized code via MPI
 - reduced compute time to 10 hrs from weeks



- Training opportunities: New & Evolving!
 - Office Hours: every Wed 12 - 3 pm @ RC conference room
 - Intro to Odyssey & RC Services
 - Using HPC Resources Efficiently series
 - Highlighted topics in the first ½ hour of office hours
 - Guest lectures in courses
 - On demand
 - Parallel MATLAB (11/21, spring '15)
 - Software Carpentry (<http://software-carpentry.org/>) (11/8-9; spring '15)



Accessing RC Resources

RC Website & Documentation (only authoritative source!)

<https://rc.fas.harvard.edu/>

OdyBot: for quick fix problems (free prizes for quiz!)

<https://odybot.org/>

Getting help

- Submit a ticket via web form: <https://portal.rc.fas.harvard.edu/rchelp@fas.harvard.edu>
- or email: rchelp@fas.harvard.edu

Best way to help us? Give us...

Description of problem:

Steps to Reproduce (1., 2., 3...)

Actual results

Expected results

Additional info (login/batch? partition? JobIDs?)

Research Computing

Please talk to your peers, and
We wish you success in your research!

<http://rc.fas.harvard.edu>

<https://portal.rc.fas.harvard.edu>

rchelp@fas.harvard.edu

@fasrc

Harvard Informatics

@harvardifx

robert_freeman@harvard.edu

@DevBioInfoGuy



Questions?