

UT Southwestern
Medical Center

BioHPC

Workstations & Thin Clients

- **Why use a BioHPC computer?**
- **Comparison between workstation and thin client**
- **Network requirements**
- **Basics of the Linux Graphical User Interface**
- **Running Code Locally**
- **Submitting Jobs to the Cluster**
- **Run Graphical Jobs on the Cluster**
- **Use Windows on a BioHPC system**

- **They run the same environment as the cluster**
 - **Red Hat Enterprise 7 (RHEL 7)**
 - **Features such as bash, sbatch, and modules are preconfigured**

- **It's easier to work with large data sets**
 - **Drag-and-Drop files in the Graphic user interface (GUI)**
 - **You don't have to use an FTP client, Samba shares, or web**



Lenovo P520
6 Core Processor
64GB RAM
1TB NVME SSD
Nvidia GPU 5GB
\$2400



P340 Tiny

6 Core Processor

32GB RAM

512GB NVME SSD

Nvidia GPU 2GB

\$1000

Comparison of Functions




- **10GB Network Card**
 - 1GB file transfers in ~ 2 seconds
- **High end GPU**
 - Image processing, machine learning
- **384GB of space for Windows Virtual Machines**



- **1GB Network Connection**
 - 1GB file transfers in ~ 15 seconds
- **100GB of space for Windows Virtual Machines**

1GB File Transfer Speed Comparison



Monitors



Standard 24" Monitor
Lenovo P24q-10
\$200



Premium 27" IPS Display
ASUS PB278Q
\$400+

Network Requirements

All BioHPC systems require a 10GB network connection.

Price varies for port installation or activation.

Telecom requests can be placed in PeopleSoft. Allow 1 to 2 weeks.

Ask your Administrative Assistants for assistance.

IP reservation requests can be placed here:

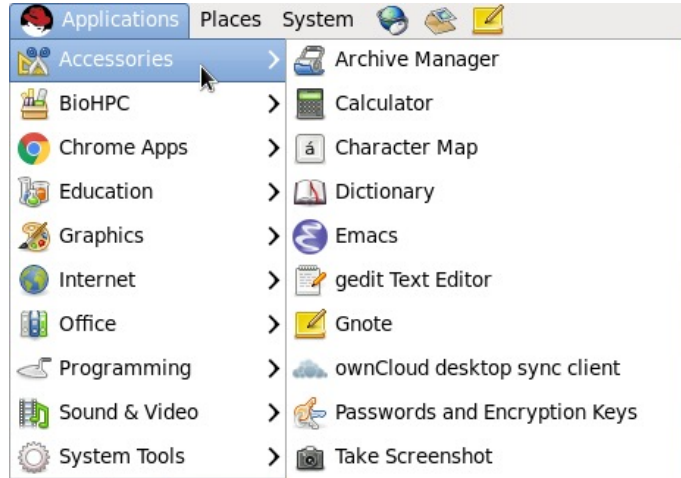
<https://irrequests.swmed.org/SitePages/IP/IPRequests.aspx>

Make sure you put in the notes that you need a 10GB IP address!

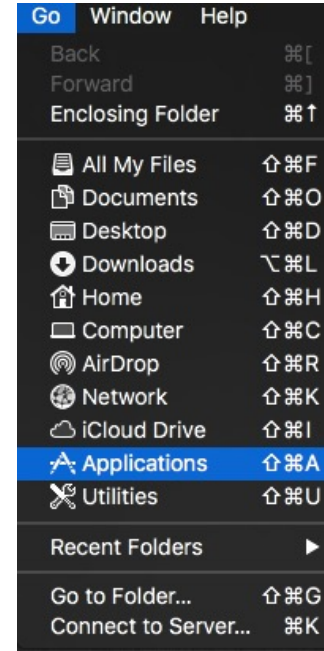
REQUEST DETAILS	
* Description/Purpose of the server or device	This request is for a 10GB IP address!

GUI Interface

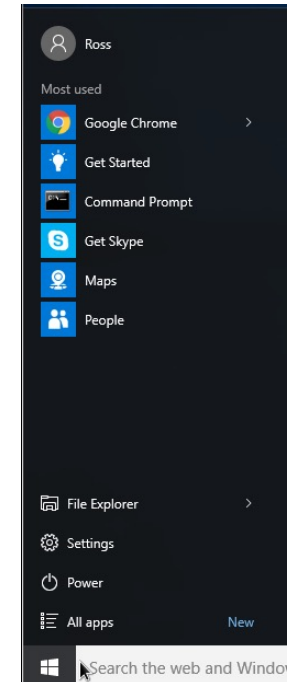
- GUI desktop similar to Windows.



Red Hat



OSX



Windows 10

- Your roaming profile maintains your files and settings on all BioHPC workstations.

Basic Settings

- You must set the proxy configuration in your profile to access the Internet outside of the UTSW firewall.

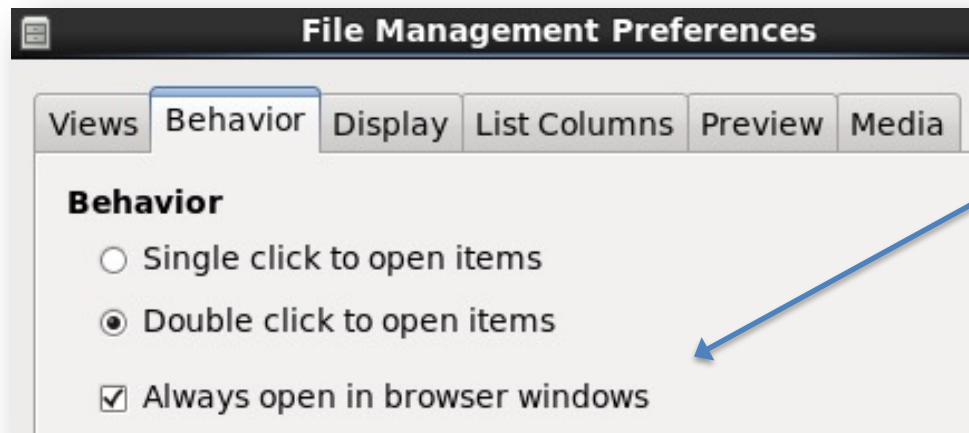
System>Preferences>Proxy Setting

Automatic proxy configuration

Autoconfiguration URL:

- Prevent folders from opening in separate windows

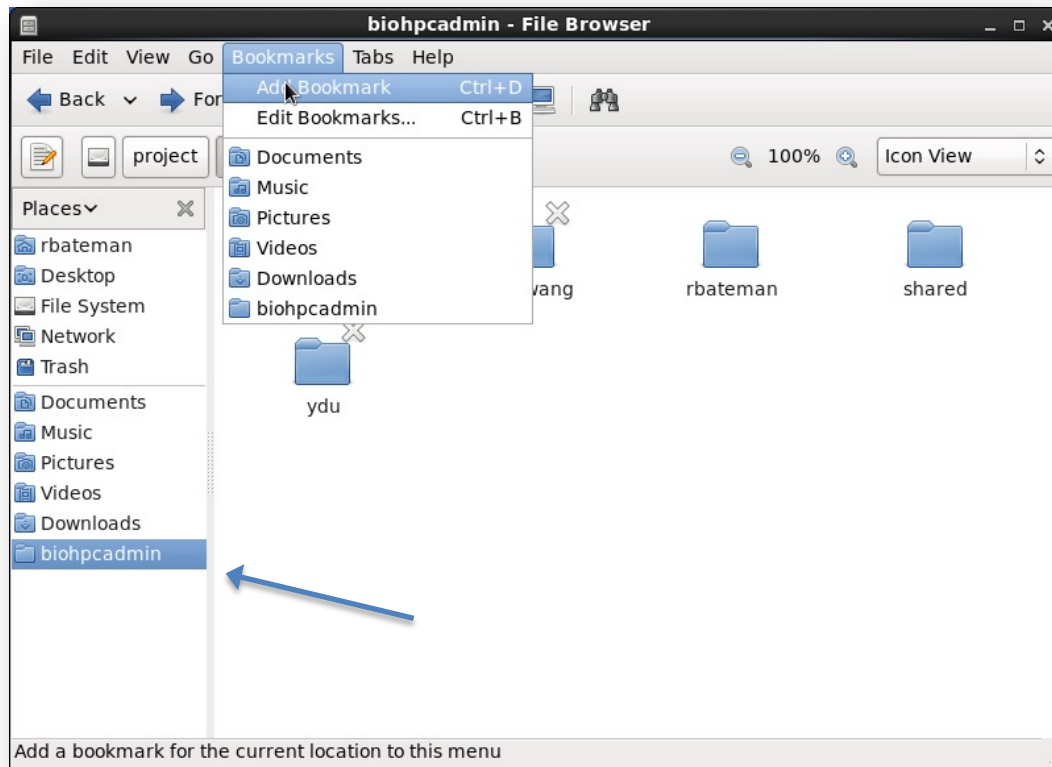
Edit > Preferences > Click on the “Behavior” tab and check “Always open in browser windows”



Basic Settings

- Create Directory Shortcuts

Create bookmarks to quickly access commonly used directories - /home2, /project, /work




▪ Disk Usage Analyzer

Applications > System Tools > Disk Usage Analyzer

Shows the status of your storage quotas.

Remember to empty the trash if you need space.



▪ System Monitor

Applications > System Tools > System Monitor

The Resources tab shows the load on the CPU and memory. This is useful for monitoring system resources when testing code on a workstation.

View storage usage on the command line

Home Directory

```
$ quota -s
```

Project Directory

```
$ lfs quota -h /project
```

View storage when connected to the cluster via SSH

```
$ biohpc_quota
```

```
rwbateman — rbateman@Nucleus005:~ — ssh rbateman@nucleus.biohpc.swmed.edu — 80x12
Current BioHPC Storage Quotas for rbateman (group: biohpc_admin):
```

FILE SYSTEM	SPACE USAGE			NUMBER OF FILES		
	USED	SOFT	HARD	USED	SOFT	HARD
home2	94303M	100G	120G	25471	0	0
project	3.966T	0k	0k	1703340	0	0
work	NoT resultsT		forT			

```
Note: project quota is per-lab, others per-user.
[rbateman@Nucleus005 ~]$
```

Need More Space?

Ask your department chair to place a request with BioHPC and include the PI's name, directory paths and amount of space needed.

Modules are used to deploy multiple versions of different software tools.

See what modules are available

\$ module avail

Load a module

\$ module load <name>

Loaded modules need to be added each session.

View loaded modules

\$ module list

Remove a module

\$ module rm <name>

See the different versions of Matlab

\$ module avail matlab

```
[rbateman@Nucleus005 ~]$ module avail matlab  
  
----- /cm/shared/modulefiles -----  
matlab/2008b matlab/2013b matlab/2014b matlab/2015b matlab/2016b  
matlab/2013a matlab/2014a matlab/2015a matlab/2016a
```

Load the Matlab Module

\$ module load matlab/2013a

Run Matlab

\$ matlab

- An admin account is required to install RPM packages.
 - Email biohpc-help@utsouthwestern.edu to request an installation
- You can install binaries in your home2 directory.
 - For example: `./configure --prefix=/home2/username/folder`

Web Job Submission on the portal website can help you create scrips:

<https://portal.biohpc.swmed.edu/sbatch/#/script>

- This provides a GUI interface for writing scrips.
- The same script can be run locally and on the cluster.

bash vs Sbatch

bash	Sbatch
<ul style="list-style-type: none">• Runs <u>Locally</u>• Good for testing• You don't have to wait in the SLURM queue	<ul style="list-style-type: none">• Runs on the <u>Cluster</u>• More powerful resources• Multiple Nodes

Run a Job Locally

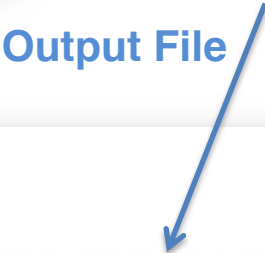
Workstations and thin clients come with Bash preconfigured to test scripts before submitting them to the cluster.

- **Run a script on a workstation with bash**

```
[rbateman@biohpcws028 demo01_serial]$ bash test_single.sh
```

- **Check the Output Folder**

```
[rbateman@biohpcws028 demo01_serial]$ cd output/  
[rbateman@biohpcws028 output]$ ls -l  
total 56  
-rwxr--r-- 1 rbateman biohpc_admin 45721 Apr  3 11:25 FeatureDots_1-persite.jpg  
-rw-r--r-- 1 rbateman biohpc_admin  3705 Apr  3 10:45 remoteGPU.txt  
-rw-r--r-- 1 rbateman biohpc_admin   901 Apr  3 10:31 remoteGUI.txt
```



Output File

Submit a Job to the Cluster

- Run sbatch


```
[rbateman@biohpcws028 demo01_serial]$ sbatch test_single.sh  
Submitted batch job 32440
```

- Check the queue with the squeue command

```
[rbateman@biohpcws028 demo05_srunSingleNode]$ squeue -u rbateman  
JOBID PARTITION      NAME      USER  ST      TIME  NODES NODELIST(REASON)  
32444      super srunSing rbateman  R      0:01      1 Nucleus050
```

- Check the output folder

```
[rbateman@biohpcws028 demo01_serial]$ cd output/  
[rbateman@biohpcws028 output]$ ls  
FeatureDots 1-persite.jpg remoteGPU.txt remoteGUI.txt
```



You can run a remote Graphical User Interface (GUI) session on a single BioHPC node. This is useful if you want to run a graphical program like Matlab and you need more than the 16GB of RAM on your workstation.

- Create a remoteGUI job.

```
$ remoteGUI
```

```
[rbateman@Nucleus005 output]$ remoteGUI  
Submitted batch job 32053
```

- Connect to the GUI session with the TigerVNC client.

```
$ cat remoteGUI.txt
```

```
=====
```

```
To connect via VNC client:  nucleus.biohpc.swmed.edu:30339
```

```
=====
```

Remote GPU

A remote GPU session is just like a GUI session, but it utilizes a node with a high-end Nvidia GPU card.

- Create a GPU job. Requires SSH.

```
$ remoteGPU
```

```
[rbateman@Nucleus005 output]$ remoteGUI  
Submitted batch job 32053
```

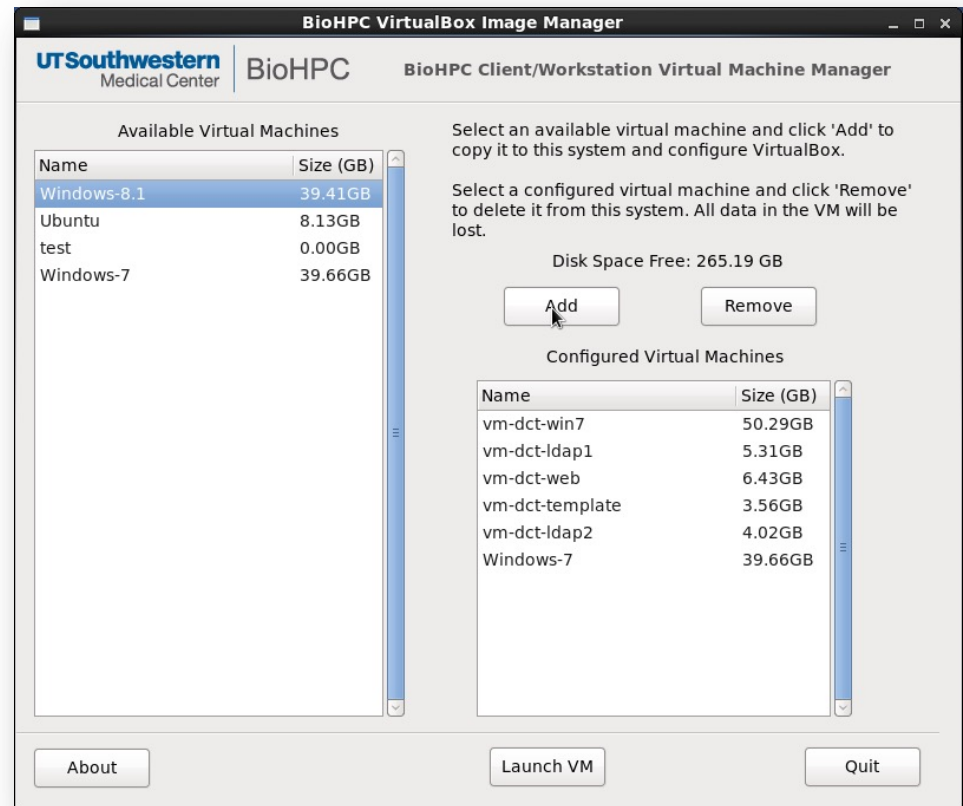
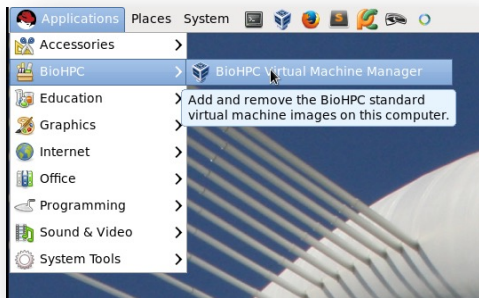
- Connect to the GPU session with the TigerVNC client.

```
$ cat remoteGPU.txt
```

```
=====  
To connect via VNC client:  nucleus.biohpc.swmed.edu:30339  
=====
```

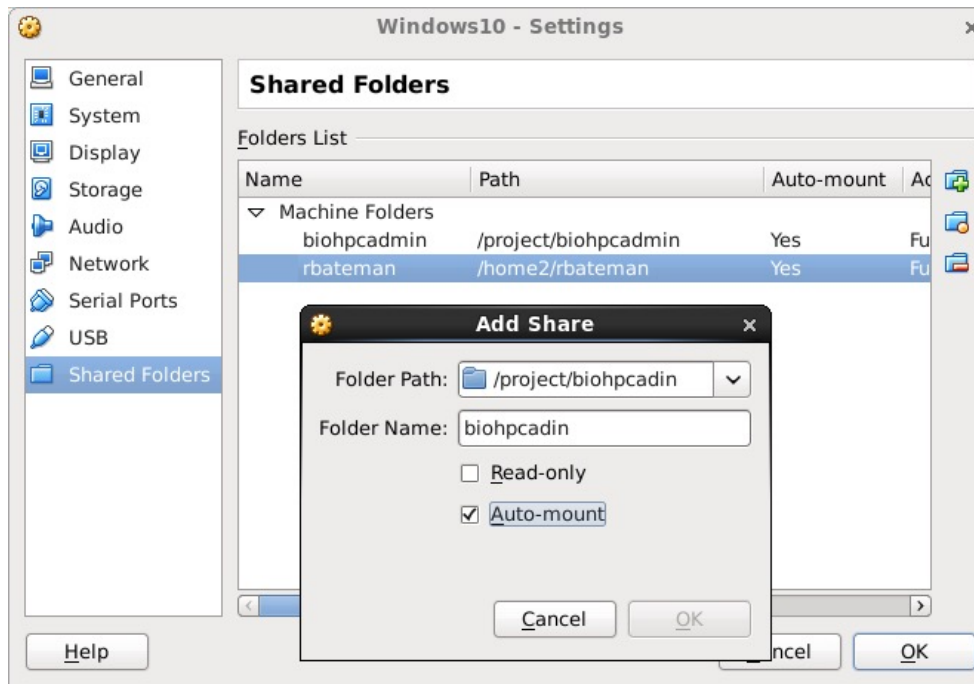
Virtual Machines

The BioHPC VirtualBox Image Manager let's you copy VMs to your local machine. Workstations have a **384GB** limit and Thin Clients have a **100GB** limit.

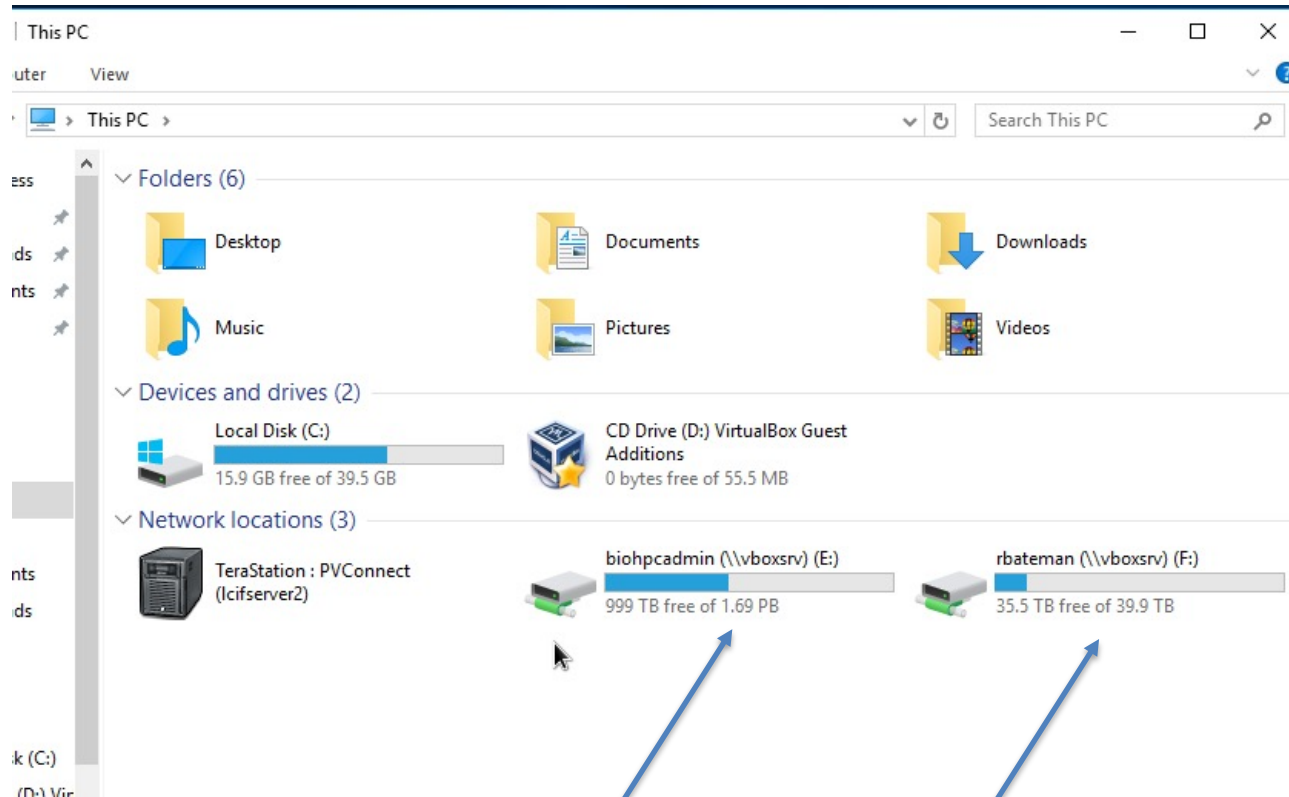




Mount BioHPC Storage



Access BioHPC Storage in Windows

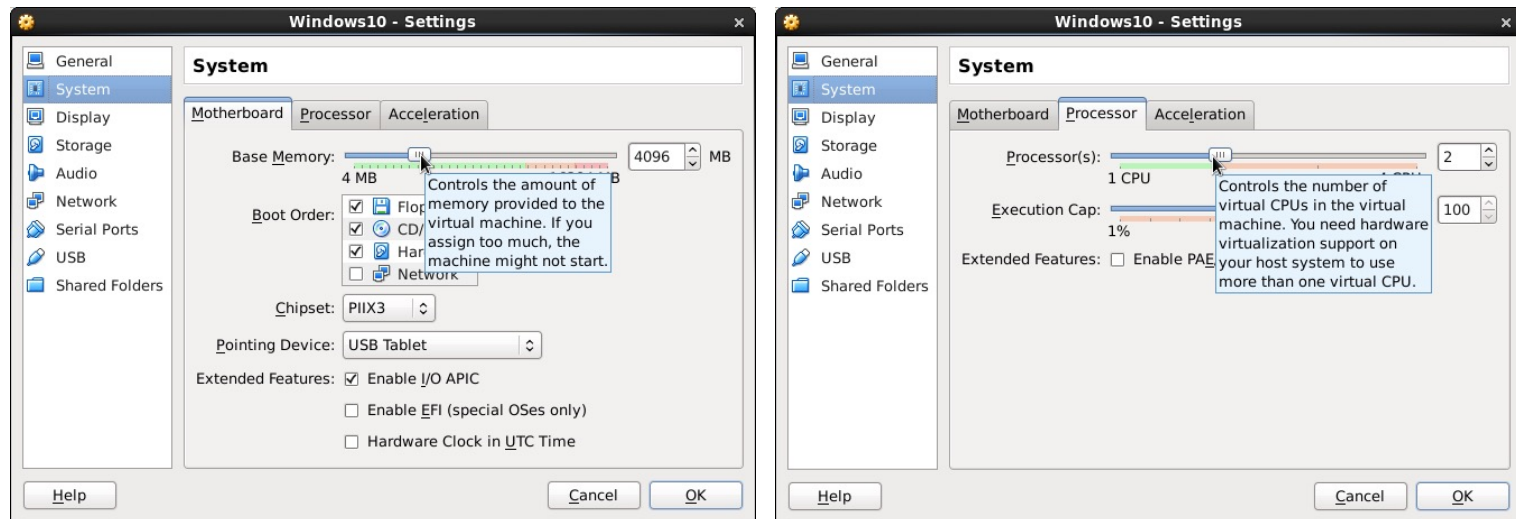


Project

Home

Virtual Box Settings

Adjust the amount of RAM and number of cores the VM can utilize.



4GB of RAM and 2 cores works well for most applications.

Increase the Size of the Drive

\$ VBoxManage modifyhd <absolute path to file> --resize <size in MB>

Example

I have a 40GB drive that I want to be 45GB the command would be.

\$ VBoxManage modifyhd <absolute path including the name and extension> --
resize 46080

```
[rbateman@biohpcws083 Win8.1]$ VBoxManage modifyhd /shared/rbateman/virtualbox_images/Win8.1/Win8dyn.vdi --resize 46080  
0%...10%...20%...30%...40%...50%...60%...70%...80%...90%...100%
```

*You **can not** shrink a guest drive with VirtualBox due to the inherent danger of loosing data or making the guest non-bootable.*

Downtime

Downtime for system upgrades will be announced two weeks in advance.

Major downtimes of a day or more are scheduled biannually

You will not be able to login to workstations or thin clients.

Incompatible Applications

OpenMPI, RDMA, GPU Direct will only work on the cluster

Connect to Nucleus with SSH to submit these jobs.

Power

Leave BioHPC systems powered on at all times, unless instructed otherwise.

For More Information

- <https://portal.biohpc.swmed.edu>
- **Email the ticket system: biohpc-help@utsouthwestern.edu**