

HPC Cluster Onboarding

Research Computing
Medical College of Wisconsin

Updated August 2022

Today's Topics

- Accounts & Access
- Resources & Architecture
- Storage & Data Transfer
- Submit and Monitor Jobs
- Practical Session – Open OnDemand
- Getting Help

Research Computing

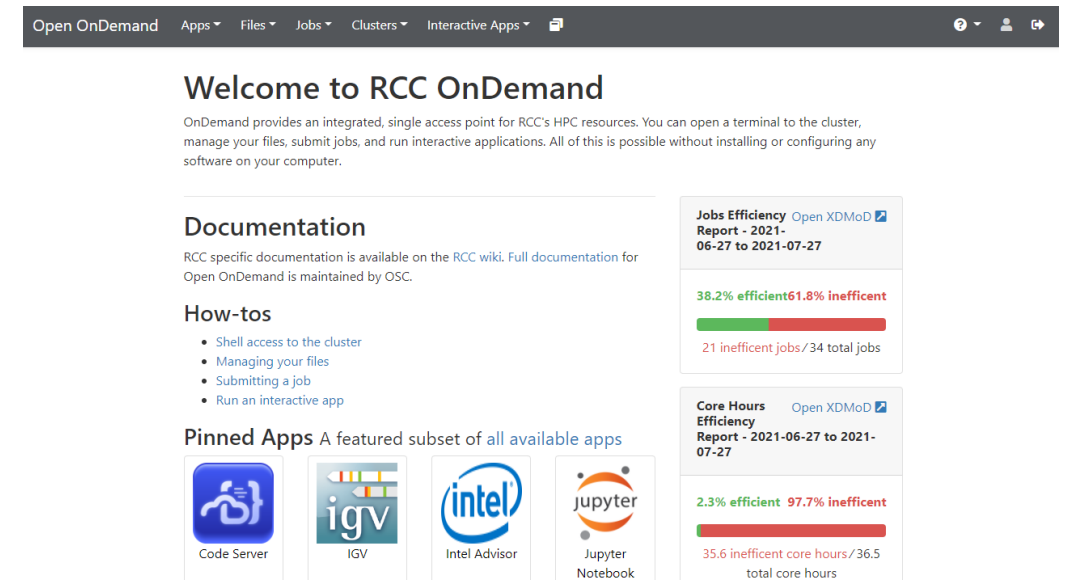
- What is Research Computing?
 - Research Computing is a division of MCW-IS that provides campus-wide access to high performance computing (HPC) resources that are designed for computational biomedical research.
- What is HPC?
 - HPC is the use of high-speed server, storage, and network resources to perform complex calculations. In practice, HPC allows a user to solve problems that are too large for the average desktop or laptop computer. This may include problems that require thousands of cores, or thousands of problems that each require one core.
- How much does it cost?
 - All MCW faculty are eligible for an allocation of free storage and unlimited computing time for their lab. Additional storage is available for fee.

Request an RCC account

- RCC user accounts are available to all MCW researchers
- Visit <https://forms.office.com/r/98QNm6cAyt>
- Most requests are handled same day
- *Please note - to obtain RCC access, you must be a PI or sponsored by a PI and have an active MCW account. A PI may sponsor students, postdoctoral fellows, staff, or colleagues with whom they are collaborating on research. PI sponsors are required to have an RCC account.*

Access HPC Cluster

- Open OnDemand – web-based portal
 - Access cluster CLI
 - Manage files
 - Manage jobs
 - All of these are possible without leaving your browser!
- SSH client
 - Secure method of connecting to an RCC server
 - Command-line interface
 - Windows clients:
 - Secure Shell Client – [download](#)
 - Putty – <http://www.putty.org/>
 - Mac clients:
 - Built-in Terminal App
 - Iterm2 – <https://www.iterm2.com/>
 - Linux
 - Built-in terminal
- Info available on <https://docs.rcc.mcw.edu/user-guide/access/login/>



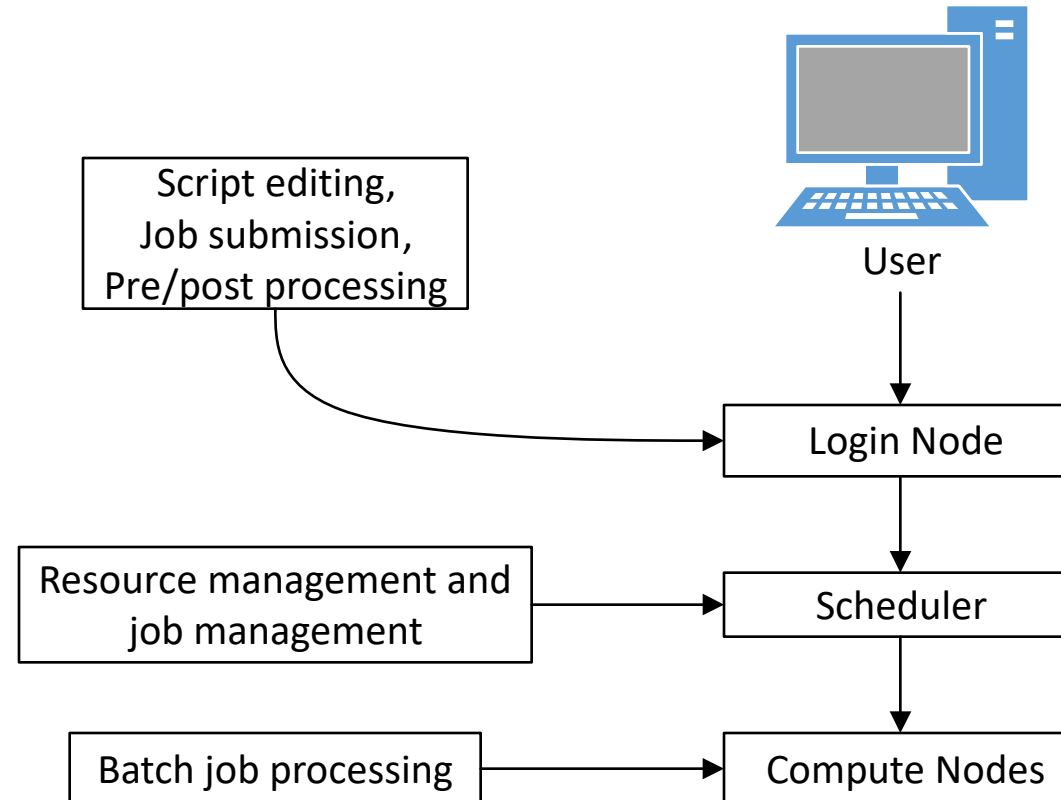
The screenshot shows the Open OnDemand web portal. The top navigation bar includes 'Open OnDemand', 'Apps', 'Files', 'Jobs', 'Clusters', 'Interactive Apps', and user icons. The main content area is titled 'Welcome to RCC OnDemand' and provides a brief overview of the service. Below this, there are sections for 'Documentation', 'How-tos', and 'Pinned Apps'. The 'Pinned Apps' section features icons for Code Server, IGV, Intel Advisor, and Jupyter Notebook. On the right side, there are two efficiency reports: 'Jobs Efficiency Report - 2021-06-27 to 2021-07-27' showing 38.2% efficient and 61.8% inefficient jobs, and 'Core Hours Efficiency Report - 2021-06-27 to 2021-07-27' showing 2.3% efficient and 97.7% inefficient core hours.

Access HPC Cluster

- Remote Access
 - Off-campus access to RCC resources is dependent on MCW-IS standard remote access methods. See <https://docs.rcc.mcw.edu/user-guide/access/remote-access/> for details.
- Open OnDemand
 - Open OnDemand is a web browser-based interface to RCC computing resources. You can manage files, submit and monitor jobs, and run pre-configured interactive apps such as Jupyter and RStudio. All of this is possible without logging in via a traditional SSH terminal.
 - **This is the recommended login method for most users.** For more info, see <https://docs.rcc.mcw.edu/user-guide/access/ondemand/>
- SSH Connection
 - For advanced use case, SSH may be used to login. Please see <https://docs.rcc.mcw.edu/user-guide/access/ssh/> for information about SSH clients and logging in.

HPC Cluster Resources

- Login Nodes
 - 4 login nodes in round-robin DNS
 - 24 cores/node
 - 128GB memory/node
 - 25Gbe network to storage
- Compute Nodes
 - 60 standard nodes
 - 48 cores/node
 - 7.5 GB memory/core
 - 480GB SSD
 - 2 large memory nodes
 - 48 cores/node
 - 32 GB memory/core
 - 480GB SSD
 - 6 GPU nodes
 - 48 cores/node
 - 7.5GB memory/core
 - 480GB SSD
 - 4 V100 NVIDIA GPUs
- Storage
 - 215TB NVMe - /scratch
 - 1.775PB disk - /group (i.e. RGS)



HPC Cluster Architecture



User

SSH, Open OnDemand

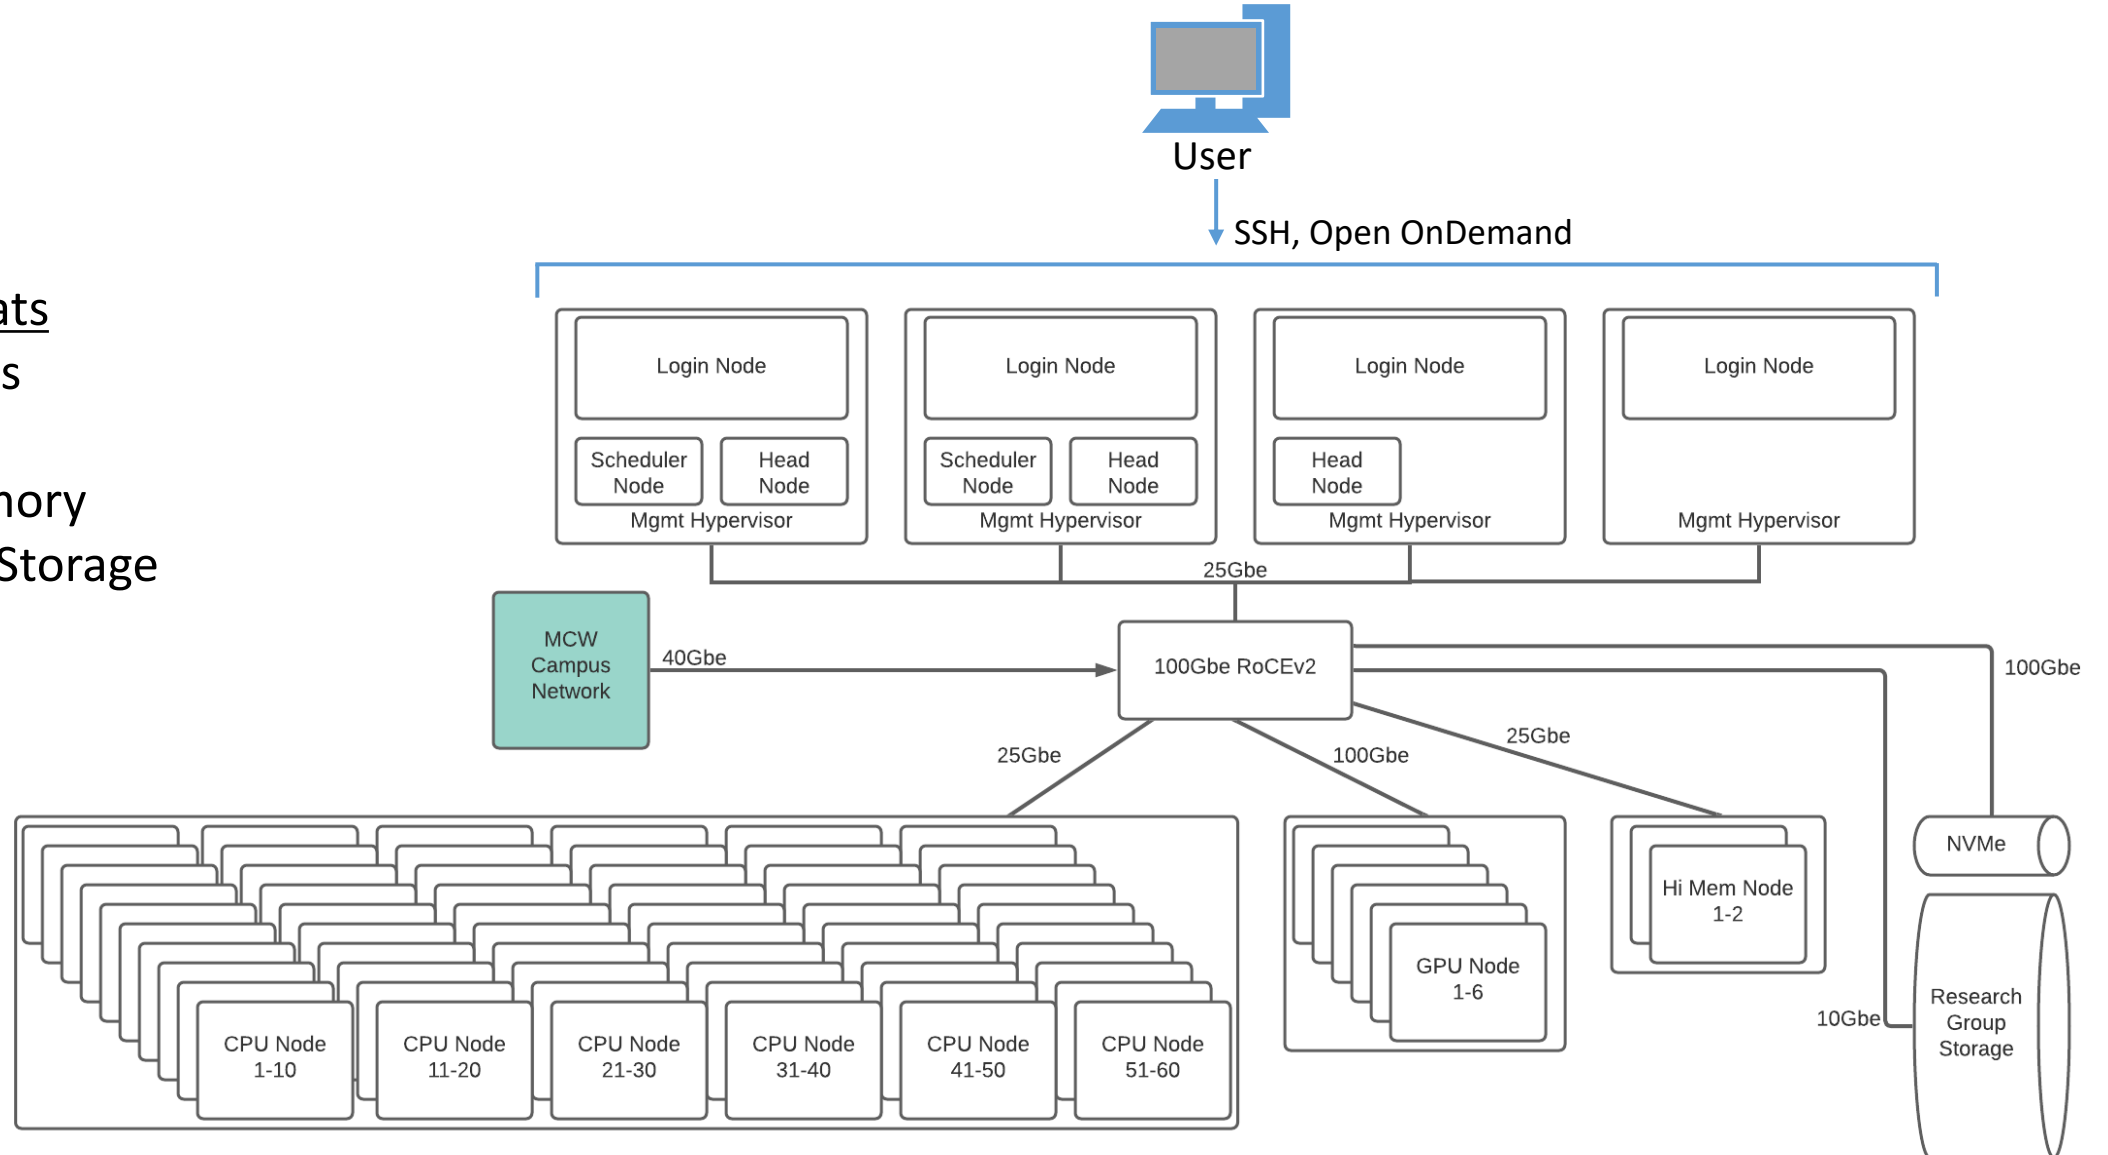
Cluster Stats

3264 Cores

24 GPUs

28TB Memory

32TB SSD Storage



Login vs. Compute Nodes

- Login nodes – not for computationally intensive work
 - Editing scripts
 - Submitting jobs
 - Checking the status of jobs
 - Troubleshooting jobs
 - Interactive tasks
- Compute nodes
 - Running batch jobs from the queueing system
 - Supports interactive and batch workload
 - Standard, GPU, and Hi-mem nodes available

User Etiquette

- HPC Cluster is a shared resource
- All jobs must be run through the SLURM queueing system
- Please do not start computationally intensive work on the cluster login nodes
- User login is restricted to the cluster head node unless access to a compute node is needed to debug a failed job
- Be respectful of your computational neighbors

Available Storage

- Free to all users:
 - /home/NetID – 50Gb hard limit
 - /group/PI_NetID – 1TB free limit
 - /scratch/u/NetID – User scratch directory
 - 5TB temporary storage for runtime files
 - NVMe storage
 - /scratch/g/PI_NetID – Group scratch directory
 - 5TB temporary storage for runtime files (expanded upon request)
 - NVMe storage
 - /tmp – local scratch storage for active jobs
 - Local disk on compute nodes (do not fill up login node /tmp)
 - 480GB on all compute nodes
- For fee:
 - Research Group Storage - /group/PI_NetID
 - First 1TB is free
 - \$60/TB/year for additional storage

💡 You can easily find your available storage directories and current utilization on the cluster with the `mydisks` command.

```
$ mydisks
====My Lab====
Size Used Avail Use% File
47G 29G 19G 61% /home/user
932G 158G 774G 17% /group/pi
4.6T 0 4.6T 0% /scratch/u/user
4.6T 0 4.6T 0% /scratch/g/pi
```

Disk space will show as less than quota limit. This is due to base-2 vs. base-10 math. Rest assured, the actual limit is being enforced despite **mydisks** output

File Transfer – Command line SCP

File to HPC Cluster:

```
scp local_file NetID@login-hpc.rcc.mcw.edu:/path/to/remote/target-directory
```

Directory to HPC Cluster:

```
scp -r local_directory NetID@login-hpc.rcc.mcw.edu:/path/to/remote/target-directory
```

File from HPC Cluster:

```
scp NetID@login-hpc.rcc.mcw.edu:/path/to/remote_file /path/to/local/target-directory
```

Directory from HPC Cluster:

```
scp -r NetID@login-hpc.rcc.mcw.edu:/path/to/remote_directory /path/to/local/target-directory
```

File Transfer – Command line RSYNC

File to HPC Cluster:

```
rsync -avz local_file user_id@login-hpc.rcc.mcw.edu:/path/to/target-directory
```

Directory to HPC Cluster:

```
rsync -avz local_directory NetID@login-hpc.rcc.mcw.edu:/path/to/target-directory
```

File from HPC Cluster:

```
rsync -avz NetID@login.rcc.mcw.edu:/path/to/remote_file /path/to/local/target-directory
```

Directory from HPC Cluster:

```
rsync -avz NetID@login-hpc.rcc.mcw.edu:/path/to/remote_directory /path/to/local/target-directory
```

HPC Software Environment

- Most packages installed as modules
- Modules dynamically load software package environments
- Commands:
 - *module avail* – list all installed software modules
 - *module avail gcc* – list installed versions of gcc compiler
 - *module load gcc (ml gcc)* – load default module for gcc compiler
 - *module load gcc/9.3.0 (ml gcc/9.3.0)*– load specific version module of gcc compiler
 - *module list (ml)* – display your currently loaded modules
 - *module unload gcc/9.3.0* – unload module
 - *module help gcc/9.3.0* – display help information
- How do I find my software?
 - *module avail*
- Request software install – email help-rcc@mcw.edu

SLURM Workload Manager

- New queuing system – SLURM (previously used Torque/Moab)
 - Reports available resources to the scheduler
 - Manages job submission and runtime
- Commands
 - *sbatch* – submit a job
 - *queue* – show status of jobs
 - *queue -a* – list all jobs with detail
 - *queue -u \$USER* – list all jobs of user
 - *scancel* – kill a job
- <https://docs.rcc.mcw.edu/user-guide/jobs/running-jobs/>

SLURM Job Script Format

test-job.slurm

```
#!/bin/bash
#SBATCH --job-name=test-job
#SBATCH --ntasks=1
#SBATCH --mem-per-cpu=1gb
#SBATCH --time=00:01:00
#SBATCH --account=PI_NetID
#SBATCH --partition=partition
#SBATCH --output=%x-%j.out
#SBATCH --mail-type=ALL
#SBATCH --mail-user=NetID@mcw.edu

echo "Starting at $(date)"
echo "Job name: ${SLURM_JOB_NAME}, Job ID: ${SLURM_JOB_ID}"
echo "I have ${SLURM_CPUS_ON_NODE} CPUs on compute node $(hostname -s)"
```


Submit a Job

- Most jobs that run on HPC are batch jobs. A batch job is submitted with the *sbatch* command and requires a job script. This is the best method for production job as it allows you to submit many jobs and let SLURM do the work. With a batch job, there is no requirement that you sit and watch the command-line. You can submit the job and come back later.
- Submit the job
 - *sbatch hpc-run.slurm*
- Check the job status
 - *squeue -u NetID*

Monitoring Jobs

- A little time spent here dramatically increases success rate
- Tools
 - *queue*
 - *sacct -j JobID*
 - *scontrol show job JobID (only works with running job)*
 - Output/error files
 - SSH to compute node and run “top” command
- Make sure your job is doing what you intended

Troubleshooting Jobs

- What is the status of the job? – *queue -j Job_ID*
 - Job status is CG – job has cleared, check output for errors
 - Job status is PD – job is waiting to run
- Does the job violate HPC queue limits?
 - Job status is PD(*Some_Reason*) – Job is potentially blocked – job will not run due to some violation of policy
 - *scontrol show job JOBID* – look for reason
 - Contact RCC
- Check XDMoD for job information (<https://xdmod.rcc.mcw.edu>)

Practical Session

<https://ondemand.rcc.mcw.edu/>

OPEN

 **nDemand**

Log in with your MCW username and password.

Username

Password

Remember me

Log In

Getting Help

- RCC wiki – <https://docs.rcc.mcw.edu/>
 - Help documentation specific to our systems
 - Subjects include accessing resources, running jobs, troubleshooting jobs, and software package
 - FAQ – <https://docs.rcc.mcw.edu/faq/>
- Man pages
 - Manual pages are included within the Linux operating system for many commands
 - *man ls*
- Forums
 - Stack Overflow
 - Stack Exchange
 - Many others
- Send a help request to help-rcc@mcw.edu

Getting Help

- Sending a help request to RCC
 - Contact help-rcc@mcw.edu
 - Include the following when applicable:
 - Your MCW NetID (username)
 - Machine name where the problem occurred (usually at the top of output file)
 - Job number of problem job (if applicable)
 - Name of problem software package (if applicable)
 - Brief explanation of what happened
 - Steps to reproduce the issue (if applicable)
 - Any fixes you've tried (if applicable)
 - RCC may request further information
 - Please be timely in your response

Recommended

- Explore the RCC wiki – <https://docs.rcc.mcw.edu/>
- Work on being comfortable in the command line
- <https://www.codecademy.com/learn/learn-the-command-line>
- Man pages
 - Linux includes manual pages for most common commands
 - Information about flags and functions of each command is presented in a uniform way
- Linux Forums – Google is your friend

Supporting RCC

- Publications and Acknowledgement

- For projects that have received support from the RCC, we appreciate an acknowledgement in your publication
- Suggested:

"This research was completed in part with computational resources and technical support provided by the Research Computing Center at the Medical College of Wisconsin."

- We're excited when our work helps result in publications. Please send us an email to let us know!

Getting Started

- RCC wiki – <https://docs.rcc.mcw.edu/>
- New Users – HPC Quick Start
 - Request an account
 - <https://docs.rcc.mcw.edu/user-guide/quickstart/>
 - <https://docs.rcc.mcw.edu/user-guide/access/remote-access/>
 - <https://docs.rcc.mcw.edu/user-guide/jobs/running-jobs/>
- Experienced Users – SLURM Guide
 - <https://docs.rcc.mcw.edu/user-guide/jobs/running-jobs/>
- Send a help request to help-rcc@mcw.edu