




Sandia
National
Laboratories

Training for Predictive Science Academic Alliance Program (PSAAP) Users



PRESENTED BY

Heidi Uphoff (hauphof@sandia.gov)

November 19, 2020




Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2020-12526 TR



SNL PSAAP Compute Resources





MANZANO

- Intel Cascade Lake E5-2695 v4; 2.9 GHz
- 1,488 compute nodes with 48 cores each; 71,424 total cores
- 192 GB memory per node
- Intel Omni-Path high speed interconnect / Mellanox ConnectX4
- 4.2 Pflops peak
- In SARAPE, request “SRN Capacity Clusters” (US citizenship required)
- <http://hpc.sandia.gov/HPC%20Production%20Clusters/index.html>

Machine	Total available core-hr/mo	PSAAP allocation	PSAAP total core-hr/mo
MANZANO	52,139,520	1.0%	521,395

HPC Production Clusters



ECLIPSE

- Intel Broadwell E5-2695 v4; 2.1 GHz
- 1,488 compute nodes with 36 cores each; 53,568 total cores
- 128 GB memory per node
- Intel Omni-Path high speed interconnect / Mellanox ConnectX4
- 1.800 Pflops peak
- In SARAPE, request “SRN Capacity Clusters” (US citizenship required)
- <http://hpc.sandia.gov/HPC%20Production%20Clusters/index.html>

Machine	Total available core-hr/mo	PSAAP allocation	PSAAP total core-hr/mo
ECLIPSE	39,104,640	1.0%	391,046



SOLO

- Intel Broadwell E5-2695 v4; 2.1 GHz
- 374 compute nodes with 36 cores each; 13,464 total cores
- 128 GB memory per node
- Intel Omni-Path high speed interconnect / Mellanox ConnectX4
- 460 Tflops peak
- In SARAPE, request “ECN Capacity Clusters”
- <http://hpc.sandia.gov/HPC%20Production%20Clusters/index.html>

Machine	Total available core-hr/mo	PSAAP allocation	PSAAP total core-hr/mo
SOLO	9,828,720	1.0%	98,287



Sandia Mass Storage System (SMSS)

- High Performance Storage System (HPSS) provides high-end near-line storage for HPC systems
- In SARAPE, request “Restricted Sandia Mass Storage System (RSMSS)” for the SRN network and “ECN Capacity Clusters” for the OHPC network.

LYNX

- Sandia’s HPC file transfer agent nodes for the SRN network
- Mounts the same parallel file systems available on the compute clusters
- 10 gigabit and InfiniBand (IB) technology to optimize file transfer performance to/from Sandia SMSS
- High performance data transfer tools HSI, HTAR, and PFTP
- In SARAPE, request “SRN Capacity Clusters”

Advanced Systems Technology Test Beds



- Research systems from several vendors
- Some require NDAs
- Not for production computing
- Systems come and go, lifetime usually a year or less
- Users cannot expect mature hardware/software
- In-depth support typically provided by research teams (e.g. email to <machine-name>-help@sandia.gov)
- Request accounts through SARAPE
- If interested in a testbed not listed in SARAPE, please let us know (hpc-help@sandia.gov)
- http://www.sandia.gov/asc/computational_systems/HAAPS.html

Advanced Systems Technology Test Beds (See our documentation for other available Test Beds)




- Not for production computing cycles - but can be provided to Test Pilot users
- Both hardware and software are intended to be highly dynamic
- Closer to prototypes and technology development drivers
- Multiple nodes available but more important to explore a diverse set of architectural alternatives, than push large scale

	Blake	Mayer	White	Morgan	Stria
CPU	Dual-Socket Intel Xeon Platinum	4 computes B0 ThunderX2 (beta) 43 computes A1 ThunderX2	Dual IBM Power 8, 10 cores	Five dual socket Intel IvyBridge with Xeon Phi co-processor and four Intel 32-core Haswell nodes	2.0 GHz Arm Cavium Thunder-X
Accelerator	None	None	NVIDIA K40 2 per node	Intel Xeon Phi Co-processor (codenamed Knights Corner) 2 per node	None
Nodes	40	47	9	9	266
Interconnect	Intel OmniPath	Mellanox EDR Infiniband with SocketDirect	Mellanox EDR IB	Mellanox Quad Data Rate InfiniBand	Infiniband
Other	Each processor core has dual AVX512 vector processing units that are FMA capable.	n/a	Technology on the path to anticipated CORAL systems	Hetero testbed on restricted network * On SRN network (US citizenship required)	* On SRN network (US citizenship required)



Requesting Access



Resource Accounts - SARAPE



- Accounts are requested by your Center's account authorizer (also called a GPA) through **sarape.sandia.gov**.
- After a SARAPE account request is approved by your Center's account authorizer, it is routed to Sandia for review, collection of additional information, and approval.
- Typically takes 1-2 weeks for US citizens.
- Account processing for non-US citizens requires additional time and paperwork - allow 30-90 days (**plan ahead!**).
- After your accounts are approved, watch your email for instructions on completing the annual online cyber security training to prevent your account from being deactivated.
- Accounts must be renewed annually. You will be notified via email when the renewal date approaches. Work with your Center's account authorizer to renew.



Getting Help



HPC OneStop – User Support



Mondays – Wednesdays: 8:30 am to 4:30 pm Mountain Time

Thursdays: 8:30 am to 3:00 pm Mountain Time

Fridays: 8:30 am to 4:30 pm Mountain Time



Call **(505) 844-9328** and leave a voicemail



E-Mail hpc-help@sandia.gov or <machine>-help@sandia.gov (e.g. mayer-help@sandia.gov)

PSAAP User Documentation



Predictive Science Academic Alliance Program

<https://psaap.llnl.gov/>

PSAAP 3 Centers

<https://asc.llnl.gov/alliances>

Sandia User Documentation



Public Website

- <http://hpc.sandia.gov>

SRN Website

- <https://onestop.sandia.gov/>
- Authenticate with Yubikey
- Information on specific platforms is under the “Hardware” tab

ECN SharePoint

- https://collaborate.sandia.gov/sites/ecn_hpc_capacity_clusters/solo/SitePages/Home.aspx
- Authenticate with ECN password
- SOLO documentation



HPC Production Clusters - Environment

Manzano, Eclipse, and Solo



Logging In – SOLO (OHPC Network)



- Email solo-help@sandia.gov and provide your home/school/office public IP address. You can see that external address by going to www.myipaddress.com. This will not be the internal 192.168.X.X or 10.x.x.x IP address assigned to your local machine through your router or modem. Once we add your IP, you can ssh to solo directly.
- Connect to Solo cluster via SSH: `ssh <sandia_username>@solo.sandia.gov` or `<sandia_username>@solo-login[1 or 2].sandia.gov`
- Your ECN password is used to login to Solo and to access the ECN collaborative website.
- A Workload Characterization (WC) ID is required to run jobs; request WC ID from your Sandia contact, and use it as your "account" for batch and interactive job submissions. View your current wcid by running the **mywcid** command on the platform.

Logging In – MANZANO and ECLIPSE (SRN Network)



- Connect to Sandia Secure Restricted Network (SRN) via gateway: `ssh <sandia-username>@srngate.sandia.gov`
- Read the “WARNING NOTICE TO USERS” message; then authenticate with your Yubikey
- From Welcome menu, select “kinit” to Acquire Kerberos credentials and enter SRN Kerberos password, then select “ssh” to establish an SSH connection
- Enter desired system name, e.g. “eclipse” or “manzano”
- A Workload Characterization (WC) ID is required to run jobs; request WC ID from your Sandia contact, and use it as your "account" for batch and interactive job submissions. View your current wcid by running the **mywcid** command on the platform.

Changing Your Default Shell



If you have SRN access, go to https://prod-ng.sandia.gov/cgi-bin/asc_ldap/shell-cgi (Yubikey Authentication)

If you don't have SRN access, please email your request to hpc-help@sandia.gov

To see which shell you are using: `echo "$SHELL"`

Available Shells

- sh
- csh
- bash
- zsh
- tcsh
- ksh

Login Node Policies



- Please do not run jobs or applications directly on the login nodes.
 - Login nodes are a shared resource.
 - Compilations may be run on login nodes (or on compute nodes).
 - Use less than half of the processors on the login node.
- Limit of 40 simultaneous logins per user per cluster login node.
 - "Too many logins" error message will be given followed by the session being terminated.
 - This will not affect any other previous established connection.
 - This policy does not apply to compute nodes.

Operating System and Software Stack



- HPC systems run a specialized operating software stack distribution called TOSS (Tri-Lab Operating System Software).
 - TOSS provides a common look and feel for clusters within the Tri-labs (LANL, LLNL and SNL).
 - Based on Red Hat Enterprise Linux 7.
 - TOSS uses Slurm as the cluster's resource manager and scheduler.
- HPC offers compilers and other tools to enable users to run Sandia and other DOE lab codes, and some commercial codes.
 - Core mission is to provide infrastructure to support the development and use of DOE simulation codes.
 - We do allow other codes to run as long as they are customer supported and don't require us to change our core environment.



- Common compilers and software packages are available through modules
- View modules currently loaded:
 - `module list`
- See details about a module:
 - `module show intel/16.0`
- View all available modules
 - `module av` OR `module avail`
- Load a module
 - `module load canopy/2.1.9`
- Unload a module
 - `module unload canopy/2.1.9`
- Unload all loaded modules
 - `module purge`

* Don't forget to load needed modules in your job submission scripts

Installing Software



- Can install in /home/<username>, /projects/<your_project>, & /projects/netpub
- If you need to download files from the internet, set your proxies first:
 - `export http_proxy="wwwproxy.sandia.gov:80"`
 - `export https_proxy="wwwproxy.sandia.gov:80"`
- How to install python packages in your user directory:
 - `pip install --user <package_name>`

* Follow Sandia's policies for what is safe and appropriate to download and run on Sandia networks and consult your Sandia Technical Contacts.

File Systems Overview



- SOLO (OHPC Network)
 - /home: 16 TB NFS
 - /projects: 23 TB NFS
 - /qscratch: 3.7 PB Lustre
 - Export Controlled Information (ECI) and Unclassified Controlled Information (UCI) are NOT allowed on Solo.
- MANZANO and ECLIPSE (SRN Network)
 - /home: 45 TB NFS
 - /projects: 49 TB NFS
 - /nscratch: 11 PB Lustre
 - /pscratch: 7.8 PB Lustre
 - /gpfs1: 4.6 PB GPFS scratch directory
- You automatically receive a directory under the home and scratch file systems. (e.g. /home/<snl_username>) To access or create a /projects directory for collaboration, work with your Sandia technical contacts or email hpc-help@sandia.gov.

File Systems Policies and Best Practices



- Executables, source code, and small input files should be stored, edited/built, and managed from /home and /projects
 - Backed up - /home/.snapshot or /projects/.snapshot
 - /home has a 100 GB hard limit
- **ALL job output should be written to one of the scratch directories**
 - Built for large, contiguous parallel I/O
 - /gpfs1 is a scratch directory
 - Intended for short-term storage only and not backed up
- Make Data Management a part of your workflow!
 - How long will you need to keep your data? Will it be deleted or moved to long-term storage (SMSS)?
- Users who are responsible for large amounts of data may be blocked from running jobs on the HPC systems
 - When scratch file systems near their limits, top users are blocked



- For the SRN network use LYNX
 - Sandia's HPC file transfer system
 - Consists of 5 login nodes
 - Recommended for all data management operations
 - rsync, scp, sftp, ftp, mv, cp, tar or any of the SMSS utilities.
 - Transfer data to/from the SRN Sandia Mass Storage System (RSMSS)
 - Data transfer examples: <https://onestop.sandia.gov/content-page/item/780> (Yubikey Authentication)
 - Sandia Mass Storage System (RSMSS) documentation: <https://onestop.sandia.gov/content-page/item/783> (Yubikey Authentication)
- For the OHPC network use SOLO
 - Transfer data to/from the OHPC Sandia Mass Storage System (OSMSS)
 - Sandia Mass Storage System (OSMSS) documentation: https://collaborate.sandia.gov/sites/ecn_hpc_capacity_clusters/solo/SitePages/OHPC%20SMS%20File%20Transfer.aspx (ECN password)



HPC Production Clusters - Job Submission & Priority Policies

Manzano, Eclipse, and Solo

Options for Different Jobs



- Partitions = collection of nodes
 - Batch – this is the default
 - Short – for small and short jobs
- qos = collection of policies (limits, priority, etc)
 - Use as an additional job submit parameter
 - Long: use this in the Batch partition for jobs >48 hours ≤ 96 hours



BATCH

- The default partition, most jobs will run in batch
- Jobs are limited to about 50% the size of the cluster
- 48 hour run time limit
- Up to 96 hours with `--qos=long`, additional node limits

SHORT

- For smaller jobs up to 4 hours long
- Submit jobs with the `-p short` or `--partition=short`
 - Or use `-p short,batch` or `--partition=short,batch` to maximize options
 - This will allow the job to run in either the short or batch partition, wherever it can be scheduled first

WCIDs and Fairshare Definitions



- **WCID** - An identifier used to submit jobs. Each Center has a specific WCID. Users will add this to job submission scripts as the “account.”
- **Fairshare** - Method of scheduling jobs by comparing allocation to utilization history. PSAAP is allocated 1% on Solo, Manzano, and Eclipse with the “univ” category. Your category allocation, “univ,” is the most important factor in determining job priority and overrides user utilization. Fairshare is evaluated based on a 28 day lookback period.
- **Priority** - A numeric value assigned to every job by Slurm.
- **Backfill** - Method used by scheduling software to allow jobs to fill in "holes" in the cluster.

WCIDs and Fairshare Details



- Job submission requires a valid WCID – this is your Slurm account
 - Users must ensure that they are using the correct WCID
 - WCIDs are used for reporting to Sandia management and DOE HQ
- **mywcid** command
 - What WCIDs you have on a cluster
- Use the **sshare** command to look at the status of your Fairshare category (univ). If your category is over its share, then other jobs may jump ahead of you.



HPC Production Clusters – Basic Job Submission

Manzano, Eclipse, and Solo



Questions for Basic Job Submission



- Batch or interactive?
- What partition? What quality of service? (QOS)
- How long will my job take?
- What happens to the input/output data?
 - Unless specified by other means, job output is placed in the directory where you launched the job. It is a best practice to CD into the scratch directory you wish to use for job output before submitting your batch job.
 - If you specify the scratch file systems used in your submission script, your job will not be started if those file systems are down. If you do not specify, Slurm will assume your job depends on all file systems.
- How many nodes and cores per task do I need?
 - Sandia allocates resources **by compute node** so that it is not possible to share compute nodes with other users' jobs.
- What modules do I need to load?

Sample scripts are located on the platforms here: [/home/samples](#)

Interactive Jobs



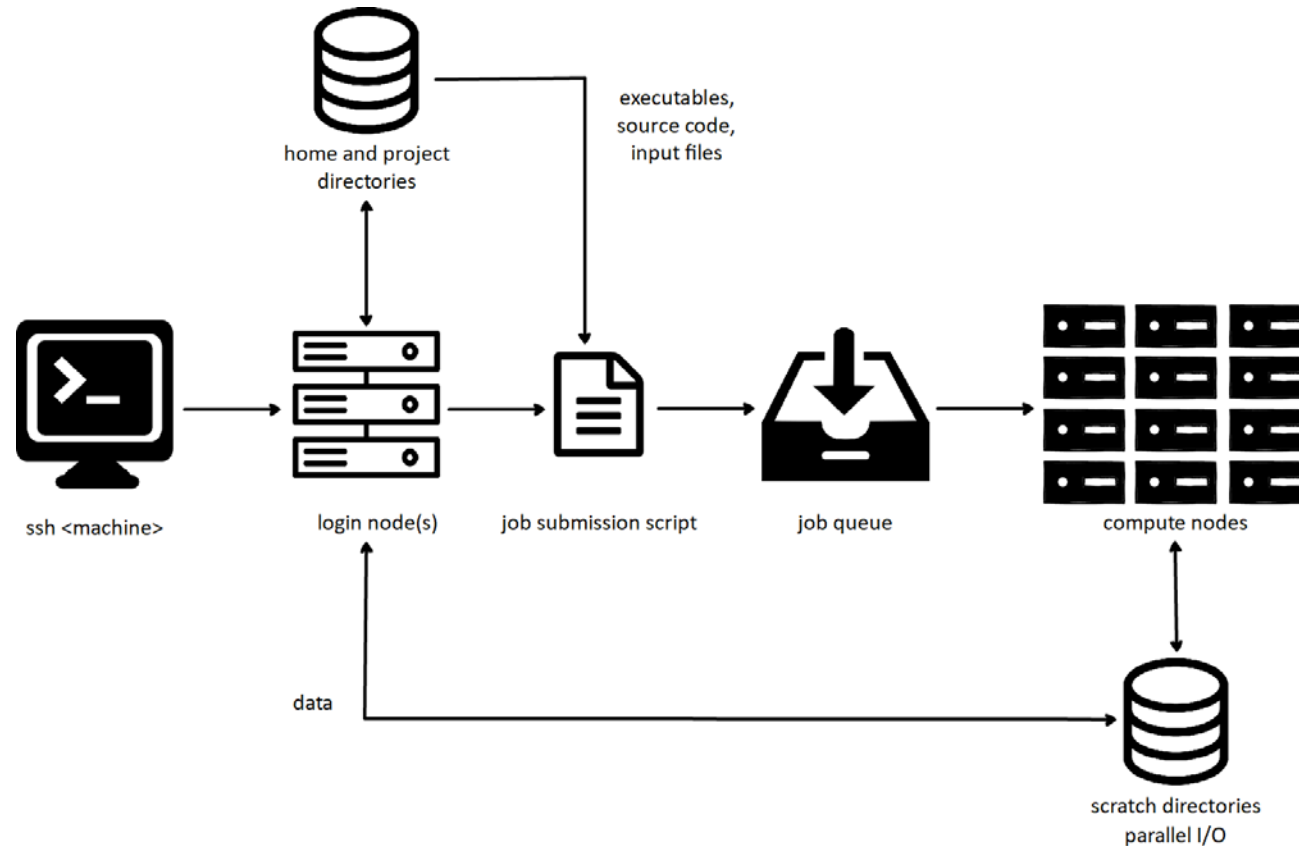
- Request interactive jobs from the command line using `salloc`
 - `salloc -N<number of nodes> --time=<hour:min:sec> --account=your_WC_ID --partition=short,batch`
 - `salloc -N1 --time=00:05:00 --account=fyXXXXXX --partition=short,batch`
- Once the session is allocated, run code using **mpiexec** (OpenMPI) or **srun** (MVAPICH, python mpi4py, non-MPI code) from the login node OR SSH to the compute node
 - `mpiexec --bind-to core --npernode 36 --n 72 ./mpi_hello`
 - `srun -n4 hostname`
- Failure to use `mpiexec` or `srun` on the login node will result in running your job directly on the login node.

Batch Jobs



Submit a job submission shell script using the **sbatch** command. The job will be placed in the queue and run on the compute nodes based on resource availability and job priority.

```
> sbatch ./example.sh
```



Example Batch Job Submission Script – example.sh



```
#!/bin/bash
## Do not put any commands or blank lines before the #SBATCH lines
#
#SBATCH --nodes=2                # Number of nodes; all cores per node are allocated to the job
#SBATCH --job-name=mpi-tutorial_job1  # Name of job
#SBATCH --time=00:05:00          # Wall clock time (HH:MM:SS) once the job exceeds this time, the job will be terminated
#SBATCH --account=PUT_YOUR_WCID_HERE
#SBATCH --partition=short,batch  # partition/queue name: short or batch
#SBATCH --qos=normal            # Quality of Service: long, large, or normal
#SBATCH --licenses=gpbs1        # If you don't specify a file system, all the file systems will be added to your job by default
#SBATCH --mail-type=ALL         # Receive emails when job starts and completes
#
# Add shell commands (load modules, create directories, compile code, etc.)
#
nodes=$SLURM_JOB_NUM_NODES      # Number of nodes - the number of nodes you have requested
                                # (for a list of SLURM environment variables see "man sbatch")
cores=8                         # Number MPI processes to run on each node (a.k.a. PPN)
                                # Eclipse has 36 cores per node, using less cores per node
                                # increases the memory available to each core
# mpiexec --bind-to core --npernode PROC_PER_NODE --n TOTAL_NUMBER_OF_MPI_PROCESSES /executable [--args...]
mpiexec --bind-to core --npernode $cores --n $((($cores*$nodes)) /home/<your_username>/mpi_hello
```



- Putting commands before the #SBATCH lines
 - sbatch: error: Batch job submission failed: Job violates accounting/QOS policy (job submit limit, user's size and/or time limits)
- Syntax
 - Look at the error output for the line number and compare with examples
- Policy Violation
 - Examples: Using a wrong WCID or submitting a longer job than allowed
- Forgetting to load modules in your job script
- Incorrect file paths, be careful when using a relative path instead of a full path to a file or folder
- Incorrect permissions on files and folders



HPC Production Clusters – Basic Slurm Commands

Manzano, Eclipse, and Solo

Monitor Jobs



View your current jobs in the queue:

```
> squeue -u <username>
```

List the job details for one of your jobs:

```
> scontrol show job <job_number>
```

* run **man <command>** OR **<command> --help** to learn usage options

Monitor Queue



List all jobs in the queue:

```
> squeue
```

List pending and running jobs in priority order:

```
> squeue -o "%.10i %.9P %v %.8j %.8u %a %.8T %.10M %.10l %.5D %.10Q %R"
```

* run **man <command>** OR **<command> --help** to learn usage options

Cancel Jobs



Cancel a job, pending or running:

```
> scancel <job_id>
```

Cancel all of your running and pending jobs:

```
> scancel -u <userid>
```

* run **man <command>** OR **<command> --help** to learn usage options

Cluster Policy / Resource Information



Display all queues available:

```
> sinfo
```

Display details about the queues (a.k.a. partitions):

```
> scontrol show partition
```

View status of your wcid's fairshare category (univ):

```
> sshare
```

```
> sshare | grep univ
```

* run **man <command>** OR **<command> --help** to learn usage options

Questions?

Email hpc-help@sandia.gov

