



The Information Science & Technology Center

ISTeC.colostate.edu

Summit High Performance Computing System

User's Guide

Patrick J. Burns

Version 2.0

Updated 04/26/2017

Table of Contents

- 1. Purpose and Audience for this Document3**
- 2. History3**
- 3. System Architecture3**
- 4. Website6**
- 5. Accounts.....6**
- 6. Remote Login7**
- 7. File Transfer.....7**
- 8. Directories and File Storage9**
- 9. Modules..... 10**
- 10.Compilers..... 11**
- 11.Batch Queues 12**
- 12.Running Jobs 13**
- 13.Condo Model 14**
- 14.Support 16**

1. Purpose and Audience for this Document

This document is intended for users of the Summit High Performance Computer system, a joint activity of Colorado State University (CSU) and the University of Colorado Boulder (CU). Summit is housed at, and operated by, CU IT staff. The emphasis in this manual is for new CSU users. Some of the information presented herein is particular to CSU's environment and some is pertinent to CU's environment, as Summit resides on the CU network and behind the CU IT security infrastructure.

All user and application support for CSU users will be performed by CSU staff, who are resident at CSU. CU staff provide user support and application support for their users, as well as all system administration support.

2. History

The Summit High-Performance Computing (HPC) system was established under NSF MRI Award #1532235 to both CSU and CU. The \$3.55 million system was awarded under the auspices of the Rocky Mountain Advanced Computing Consortium (RMACC; www.rmacc.org), an activity in which ISTE C (istec.colostate.edu) at CSU participates. ISTE C coordinated the proposal for CSU. The system went into full production on February 2, 2017.

Summit allocations are partitioned in proportion to the amount invested by each institution, with 10% reserved for the RMACC, as follows:

RMACC – 10% of the total
CSU - 25% of the remaining 90% = 22.5% of the total
CU – 75% of the remaining 90% = 67.5% of the total

Allocations are approximately on a “use it or lose it” basis monthly, and are somewhat oversubscribed to allow users who need additional resources within a month to use resources available from other users with unused resources in that month.

The initial large Summit system has since been augmented with nodes under the “condo” model, where users buy into add additional nodes to the system that they own. Additionally, the Summit system is to be augmented in Spring 2017 with a small Knights Landing system, using Intel's Many Integrated Cores (MIC) technology, exact timing to be determined.

3. System Architecture

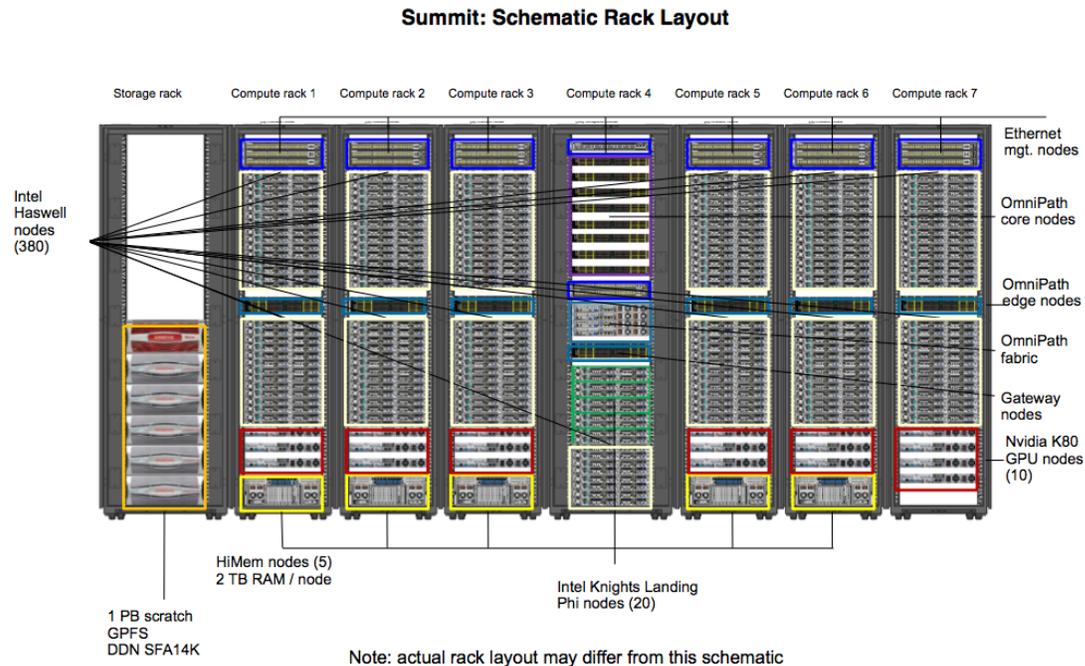
The Summit system architecture is designed as shown below:

- 380 “Haswell” Nodes: 303.2 Peak TF
 - (376) Poweredge C6320 Haswell nodes, 2x e5-2680v3 (9,024 cores total)
 - 4 in 2U Design
 - 128GB - 8x16GB
 - 1x200GB Intel S3610 SSD

- 1x OPA HFI in a x16 pcie slot – idrac express, dedicated port per node
 - Transceiver SFP+ to BT
 - HPC BIOS
 - 2x1,400W psu per chassis
 - 5 yrs hardware warranty NBD
- 10 “GPU” Nodes: 15.3 CPU Peak TF + 71.1 GPU Peak TF
 - (19) Poweredge C4130 (config D) GPU nodes, 2x e5-2680v3
 - 1U Chassis
 - 2x NVIDIA K80
 - 128GB RAM
 - 1x200GB SSD
 - 1x OPA HFI in a x16 pcie slot
 - Idrac express, dedicated port
 - HPC BIOS
 - 1x1,600W PSU, Non Redundant
 - 5 yrs hardware warranty NBD
- 5 “Hi-Mem” Nodes: 6.5 Peak TF
 - (5) Poweredge R930 nodes, 4x E5-4830v3
 - 2TB RAM
 - 64x32GB DIMMs
 - H730 Storage Controller
 - JBOD Mode
 - 12TB raw local storage
 - Idrac express, dedicated port
 - 5 yrs hardware warranty NBD
 - Redundant Hot Swap PSU
- 2 Storage Gateway Nodes
 - (2) Poweredge R430
 - 1xE5-2623v3 3GHz 4C
 - 32GB RAM 4x8GB
 - 2x500GB SATA RAID1
 - Redundant PSU
 - Dual port 10Gb Ethernet
 - Idrac Enterprise
 - 5 yrs ProSupport Mission Critical 4 hour response
- 2 Omnipath Architecture (OPA) Interconnect Fabric Management Nodes
 - (2) Poweredge R430 nodes
 - 1xE5-2623v3 3GHz 4C
 - 16GB RAM 4x4GB
 - 1x500GB SATA
 - Redundant PSU
 - Idrac express
 - 5 yrs ProSupport NBD

- Omni-Path 100 GB/sec. OmniPath Interconnect
 - (22) H-series Switches Dell Networking H1048 Omni-Path 1U 48 Port Switches, 14 leaf, 8 spine
 - Fabric easily expands to 24 leaf switches
 - Non-Redundant PSU
- 1 Petabyte DDN SFA14K scratch storage
- 40Gb Uplink to the Science Ethernet Network (fiber connected to CSU)
- 20 Knights Landing “KnL” Nodes: 53.4 Peak TF (pending, not installed yet)
 - (20) KnL Nodes, 2.6GHz
 - 4 in 2U design
 - 96GB RAM
 - 1x200GB SSD
 - Integrated dual port OPA HFI
 - 5 years hardware warranty NBD
 - Redundant PSU

A schematic rack layout of the system is shown below:



4. Website

The Summit website for CSU users is www.hpc.colostate.edu. The website is updated periodically with new information as it becomes available.

5. Accounts

You can request an account on Summit by completing the Summit Account Request Form provided in Appendix A. Submit the applications via email to: hpc-support@colostate.edu. Also, instructions for obtaining an account are available at: <http://hpc.colostate.edu/accounts>. Under normal circumstances, you may expect your request to be processed within 2 working days. However, as we are so thinly staffed (we have only one person processing these forms), there may be delays due to vacation, sick leave, travel etc. Even in these circumstances, we will do our best to respond expeditiously. If you have an urgent need, and you need to escalate your request, you may do so by forwarding the request to the Vice President for IT, Pat Burns, Patrick.Burns@colostate.edu.

However, please note that there are additional steps you need to perform to obtain access, and these are described in the next section.

Here are instructions to apply for an account on Summit. These instructions must be executed in sequence:

1. Get a CSU eID (<http://eid.colostate.edu>) OR get a CSU Associate's eID. If you are affiliated with CSU, you should already have an eID. If not, go to the website and apply for one. If you are unaffiliated with CSU, you will need an Associate's eID. Fill in the [Associate's eID form](#) and send to your home department.
2. Complete the account application form (Appendix A) and email it to hpc-support@colostate.edu
3. Follow [DUO Two-factor Authentication Instructions](#)
4. Request [CU-Boulder account](#)

Recall that the eID (<http://eid.colostate.edu/About.aspx>) consists of the eName and the ePassword. As an example, one of the author's has an eName of "pburns". The eName is used to obtain accounts and allocations on Summit. The ePassword is the 15-character password used to access CSU central IT systems. Hereinafter, wherever "eName" appears, use your eName that is part of your eID, and whenever "ePassword" appears, use your ePassword that is part of your eID. If you have questions about this, see <http://eid.colostate.edu/About.aspx>, or feel free to contact the ACNS help desk, help@colostate.edu or (970) 491-7276.

On this site, enter the following information:

Organization: Colorado State University

Username: eName

Password: ePassword,DUO_key

Role: choose anything
Preferred login shell: leave it as “bash”
Check “Summit supercomputing cluster”
Submit Request

- Receive account confirmation email from CU-Boulder

At this point your account should be active and you should be able to [login](#) to Summit.

NOTE: The DUO_key mentioned above cycles every 15 seconds. Therefore, in your cell phone DUO app, once you tap the key icon and generate a 6-digit DUO_key, you must enter or use that key within 15 seconds, otherwise it will expire and you’ll have to generate another key.

6. Remote Login

There are several methods for remote login to Summit.

Method 1 (recommended)

```
ssh eName@colostate.edu@login.rc.colorado.edu  
csu_password,push
```

where “push” is a literal string (case insensitive)(see [DUO Two-factor Authentication instructions](#)) Don’t forget the comma “,”.

Method 2

```
ssh eName@colostate.edu@login.rc.colorado.edu  
ePassword,DUO_key
```

where “DUO_key” is the two-factor key obtained from your cell phone DUO app (see [DUO Two-factor Authentication Instructions](#)). Don’t forget the comma “,”.

NOTE: The DUO_key mentioned above cycles every 15 seconds. Therefore, in your cell phone DUO app, once you tap the key icon and generate a 6-digit DUO_key, you must enter or use that key within 15 seconds, otherwise it will expire and you’ll have to generate another key.

7. File Transfer

There are several methods for transferring files to/from Summit.

Method 1

sftp eName@colostate.edu@login.rc.colorado.edu
ePassword,push

where “push” is a literal string (case insensitive)(see [DUO Two-factor Authentication instructions](#)). Don’t forget the comma “,”.

Method 2

sftp eName@colostate.edu@login.rc.colorado.edu
ePassword,DUO_key

where “push” is a literal string (case insensitive)(see [DUO Two-factor Authentication instructions](#)). Don’t forget the comma “,”.

Method 3

Install [Filezilla](#) on your workstation.

Launch Filezilla, then choose “File -> Site Manager...”

Choose “New Site”, enter a name (i.e. Summit) in the My Sites list, and then enter the following information:

Host: login.rc.colorado.edu
Protocol: choose “SFTP – SSH File Transfer Protocol”
Logon Type: Normal
User: eName@colostate.edu
Password: ePassword,push

where “push” is a literal string (case insensitive)(see [DUO Two-factor Authentication instructions](#)). Don’t forget the comma “,”. Then click OK. This step saves the profile information for Summit.

To use Filezilla, choose “File -> Site Manager...”. In the My Sites list choose the profile name for Summit (i.e. Summit), and click Connect. Immediately, the DUO application on your cell phone should prompt you to “Approve” the request. You will then be connected to Summit in sftp mode. A file list in your Summit home directory should display in the “Remote Site” area of Filezilla. At this point you can “drag-and-drop” files to/from Summit.

Method 4

For Windows OS only.

Download and install [WinSCP](#) on your workstation.

Configure WinSCP as follows:

Host: login.rc.colorado.edu
Protocol: SFTP
Username: eName@colostate.edu
Password: ePassword,DUO_key

where “eName” is your regular [CSU eID eName](#), “ePassword” is your regular [CSU eID ePassword](#) and “DUO_key” is the two-factor key obtained from your cell phone DUO app (see [DUO Two-factor Authentication instructions](#)). Don’t forget the comma “,”.

8. Directories and File Storage

Summit’s directory structure is shown below. Again, please note that “eName” is your CSU eID eName (<http://eid.colostate.edu>). After logging into Summit, you are placed in your home directory. There is very little space in your home directory. Instead, most of your file management activities should take place in your project directory. There is a large global scratch directory that can be used by running jobs. Also, each compute node has fast local SSD file storage.

Note: CSU has an Agreement with Azure, where users (for a cost) can subscribe to file storage/backup space. Also, CSU is developing a central, shared file storage system that research groups can subscribe to (for a modest cost, this infrastructure is subsidized centrally).

Home directory

- [/home/eName@colostate.edu](#)
- 2 GB
- Permanent
- Daily incremental backups

Project directory

- [/projects/eName@colostate.edu](#)
- 250 GB
- Permanent
- Daily incremental backups

Scratch directory – global

- [/scratch/summit/eName@colostate.edu](#)
- 1 PB (petabyte)
- Files purged after 90 days
- **NO** backups

Scratch directory – local

- /scratch/local
- 200 GB
- SSD on each node
- Files purged when job is finished
- **NO** backups

9. Modules

Summit includes the [Lmod](#) environment module system to simplify shell configuration and software application management. Some common module commands are shown below.

Command syntax	Purpose
module list	Show modules that are currently loaded
module avail	Show all modules that are available to be loaded
module spider	Show all modules that are available to be loaded (more extensive list)
module load module_name	Load module module_name
module unload module_name	Unload module module_name
module swap module_name_1 module_name_2	Unload module_name_1 and load module_name_2
module show module_name	Show information about module_name
module help module_name	Show information about module_name
module help	Describe module commands and parameters

In the Lmod system, most software is not accessible by default. Instead, it has to be loaded into your Linux shell environment using various module commands above. This system allows users and system administrators to manage multiple versions of software concurrently and to easily switch between versions.

“Loading a module” sets or modifies a user’s environment variables to enable access to the software package provided by that module. For instance, the \$PATH variable might be updated so that appropriate executables for that package can be used.

The Lmod environment module system is hierarchical. There are five layers to support programs built with compiler and library consistency requirements. Modules can only be loaded once their dependencies have been satisfied. This prevents accidental loading of modules that are inconsistent with each other. For example, in order to load an MPI-dependent program, it’s first necessary to load a compiler (i.e. Intel) and then an MPI implementation (i.e. IMPI) consistent with that compiler.

The five Lmod layers are:

- Independent programs

- Compilers
- Compiler dependent programs
- MPI implementations
- MPI dependent programs

10. Compilers

Summit supports a variety of compilers, interpreters and languages as shown below. There are often several versions of each item in the list. To see which versions are available on Summit enter:

module avail

OR

module spider

If you need to compile your own software applications, [here](#) is additional information about how to do it on Summit.

Name	Description
icc	Intel C compiler
icpc	Intel C++ compiler
ifort	Intel Fortran compiler
gcc	GNU C compiler
g++	GNU C++ compiler
gfortran	GNU Fortran compiler
pgcc	PGI C compiler
pgCC	PGI C++ compiler
pgfortran	PGI Fortran compiler
module load intel module load impi mpicc	Intel MPI C compiler

Name	Description
module load intel module load impi mpicxx	Intel MPI C++ compiler
module load intel module load impi mpif90	Intel MPI Fortran compiler
module load openmpi mpicc	OpenMPI C compiler
module load openmpi mpicxx	OpenMPI C++ compiler
module load openmpi mpif90	OpenMPI Fortran compiler
nvcc	Nvidia CUDA compiler
Python	Python interpreter
Perl	Perl interpreter

11.Batch Queues

[SLURM](#) is the batch queueing system used on Summit. The configuration of the batch queue system is shown below.

Partitions

Running jobs are assigned to one of the following partitions. The four basic partitions on Summit are CPU nodes, GPU nodes, KnL-F nodes and high-memory nodes.

Partition name	Compute node type	Max Time	QoS
shas	Haswell CPU nodes (380 nodes)	24 hr	N, D, C
sgpu	Nvidia K80 GPU nodes (10 nodes)	24 hr	N, D, C
sknl	Intel Knights-Landing Phi nodes (20 nodes)	24 hr	N, D, C
smem	High memory nodes (5 nodes)	168 hr (7 D)	N, D, L, C

QoS (Quality of Service)

Optionally, running jobs can be assigned to a specific level of Quality of Service, which modifies job resources, such as runtimes, priorities and number of jobs per user.

QoS	Description	Limits
Normal (N)	Default	Partition max; normal priority
Debug (D)	Quick turnaround for testing	1 hr; 1 job per user; 32 nodes max; priority boost
Long (L)	For jobs with long runtimes	168 hr (7 D); normal priority
Condo (C)	For users who purchased compute nodes (“condo model”)	168 hr (7 D); normal priority

12. Running Jobs

The following instructions illustrate how to run simple jobs on Summit.

To run jobs in the Slurm batch queues, create a batch script in text file format that includes Slurm parameters and job commands, and then submit the batch script to Slurm.

For example, suppose we have a batch script file in plain text format - filename.txt - that includes the following lines:

```
#!/bin/bash
#SBATCH -J hello           #job name
#SBATCH -p shas           #partition name
#SBATCH --qos debug       #quality of service (note double-dash)
#SBATCH -t 00:01:00      #wall time
#SBATCH -N 1              #number of nodes
gcc -o hello hello.c     #compile
./hello                  #run executable
```

The first line should always be “#!/bin/bash”. “#SBATCH” lines specify Slurm parameters. You should always choose an appropriate partition name for your job; see section “Batch Jobs” for partition names. Optionally, choose an appropriate QoS. Optionally, set a maximum wall clock time for your job to run. Choose the number of nodes to assign to the

job. Choose an appropriate compiler for your application (see section “Compilers”). In this example, source code is available in “hello.c”. Finally, run the executable file.

Submit the batch script to the batch queues:

```
sbatch filename.txt
```

Slurm will display the message:

```
Submitted batch job jobID
```

where jobID is the Job ID number.

Check job progress with:

```
squeue -j jobID
```

where jobID is the Job ID number.

When the job is complete, a new text file will appear in the same directory as filename.txt and will have the syntax:

```
slurm-jobID.out
```

where jobID is the Job ID number.

If you need to cancel a running job enter:

```
scancel jobID
```

where jobID is the Job ID number.

13. Condo Model

The Summit system supports a Condominium Computing Model (“Condo Model”) for researchers who choose this method of participation. In the Condo Model, costs are split between researchers and Central IT. Researchers purchase their own compute nodes and Central IT provides the hosting environment and support services for those nodes.

Researchers purchase

- CPU compute nodes
- GPU accelerators (if applicable)
- KnL-F accelerators (if applicable; available Q2 2017)
- Memory

- Disk storage

Central IT provides

- Data center facility
- Shared service nodes (i.e. login nodes)
- Shared OmniPath interconnect switches and cables
- Ethernet management switches and cables
- Shared scratch storage
- Server racks
- Power
- Cooling
- Security
- Purchase, order & install equipment
- Install OS
- System administration
- Assist with software application installation

Through academic discounts and volume purchase agreements, Central IT can negotiate on behalf of researchers to get reduced prices for compute node resources.

There are currently three types of compute nodes available to researchers

- Haswell CPU nodes
- GPU nodes
- HiMem nodes

The CPU nodes include Intel Haswell CPU's. The GPU nodes include 2X Nvidia Tesla K80 GPU accelerators. The HiMem nodes include 2 TB RAM. (Note: Intel Knights-Landing Phi (KnL-F) nodes will become available in Q3 2017).

Condo jobs have the following privileges

- request longer run times (up to 168 hrs. (7 D))
- get queue priority boost (equal to 1 D boost)
- access all nodes

To properly activate Condo shares, Condo users should send the following info to hpc-support@colostate.edu

- full name
- csu_eName
- condo group ID (see table below)

Note: If you are unsure about which condo group ID to use, submit the info anyway. We'll determine the right condo group ID for you.

You will receive an email note when your condo group ID assignment is complete. You'll then be able to submit jobs using your Condo allocation.

The table below shows the condo group ID that has been assigned to each principal investigator and their department affiliation.

PI	Dept.	Condo group ID
Michael Antolin	Biology	bio
Wolfgang Bangerth	Mathematics	mat
Asa Ben-Hur	Computer Science	hal
Stephen Guzik	Mechanical Engineering	cfid
Tony Rappe	Chemistry	akr
Chris Weinberger	Mechanical Engineering	crw
Ander Wilson	Statistics	fhw

Condo job submission

To submit jobs using your Condo allocation, include the following lines in your Slurm batch job file

```
#SBATCH --qos condo  
#SBATCH -A csu-smmit-xxx
```

where “xxx” is your condo group ID from the table above. Note the double-dash for the “qos” parameter. When you submit a Slurm batch job file with these parameters, the job will run with the additional privileges described above.

14.Support

Technical Support

Submit Summit support requests to:
hpc-support@colostate.edu

System Status

To receive system updates and other announcements for Summit, send a message to:
rc-announce-join@colorado.edu

Appendix A

Account Application Form for Summit HPC System

Complete and return via email to hpc-support@colostate.edu

1. Your name, First Last:
2. Your CSU eName (see <http://eid.colostate.edu>):
3. Your status:
 - CSU faculty
 - CSU administrative professional
 - CSU graduate student
 - CSU undergraduate student
 - CSU state classified staff, or temporary staff
 - CSU Affiliate
 - Other: _____
4. Your home department
 - a. Department name
 - b. Department number:
5. Describe your area(s) of use for Summit, keywords are fine:
6. List of funded research projects that this account or project request will support:

Funding Agency Name and Grant Number	Title
--------------------------------------	-------

7. Is this a request for a Project with multiple participants?
 - No – done, please submit completed request to: hpc-support@colostate.edu
 - Yes, please complete questions 8 and 9
8. List of project participants eNames (attach and submit an Account Application form for each):

9. Allocation request
 - a. Number of Service Units requested: _____
 - b. Justification for this Number of Service Units: