

Peanut Cluster (Slurm)

Infrastructure

Hardware

Our cluster contains nodes with the following specs:

general:

- 16 Cores (2x 8core 3.1GHz Processors), 16 threads
- 64gb RAM
- 2x 500GB SATA 7200RPM in RAID1

fast:

- 24 Cores (2x 24core Intel Xeon Silver 4116 CPU @ 2.10GHz), 48 threads
- 128gb RAM
- OS: 2x 240GB Intel SSD in RAID1
- /local: 2x 960GB Intel SSD RAID0

pascal

- 6 Cores (Intel(R) Core(TM) i7-5930K CPU @ 3.50GHz), 12 threads
- 32gb RAM
- OS: 1x 512gb SSD

quadro

- 16 Cores (Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz), 32 threads
- 128gb RAM
- OS: 2x Samsung SSD 850 PRO 128GB
- /local: ZFS mirror (2x Samsung SSD 850 PRO 1TB)
- 2x Quadro P4000

titan

- 8 Core (Intel(R) Xeon(R) Silver 4110 CPU @ 2.10GHz), 16 threads
- 64gb RAM
- OS: 1x 1TB 7200k spinning disk.
- 4x GeForce GTX 1080 Ti

Storage

There is slow scratch space mounted to /scratch. It is a ZFS pool consisting of 10x 2TB 7200RPM SAS drives connected via a LSI 9211-8i and is made up of 5 mirrored VDEVs, which is similar to a RAID10. The servers uplink is 1G ethernet.

- Files older than 90 days will be deleted automatically.
- Scratch space is shared by all users.

Access

Scratch space is only mounted on nodes associated with the cluster. If you want to be able to transfer files to the scratch space you will want to run an [interactive shell](#). Now you will be able to use standard tools such as scp or rsync to transfer files.

1. You should only do a file transfer via the debug partition: `srun -p debug --pty --mem 500 /bin/bash`
2. Now you can create a directory of your own: `mkdir /scratch/$USER` You should store any files you create in this directory.

Example

Request interactive shell

```
user@csilcomputer:~$ srun --pty --mem 500 /bin/bash
```

Create a directory on the scratch partition if you don't already have one:

```
user@slurm1:~$ mkdir -p /scratch/$USER
```

Change into my scratch directory:

```
user@slurm1:~$ cd /scratch/$USER/
```

Get the files I need:

```
user@slurm1:/scratch/user$ scp user@csilcomputer:~/foo .
foo                               100% 103KB 102.7KB/s   00:00
```

Check that the file now exists:

```
user@slurm1:/scratch/user$ ls -l foo
-rw----- 1 user user 105121 Dec 29 2015 foo
```

I can now exit my interactive shell.

Performance is slow

This is expected. The maximum speed this server will ever be able to achieve is 1Gb/s because of its single 1G ethernet uplink. If this cluster gains in popularity we plan on upgrading the network and storage server.

Partitions / Queues

To find out what partitions we offer, checkout the [sinfo](#) command.

As of December, 2015 we have will have at least 2 partitions in our cluster; 'debug' and 'general'.

Partition Name	Description
debug	The partition your job will be submitted to if none is specified. The purpose of this partition is to make sure your code is running as it should before submitting a long running job to the general queue.
general	All jobs that have been thoroughly tested can be submitted here. This partition will have access to more nodes and will process most of the jobs. If you need to use the --exclusive flag it should be done here.
fast	2019-12-02: 48 threads, 128GB RAM
quadro	2019-12-02: 2x Quadro P4000. *
pascal	2018-05-04: 1x Nvidia GTX1080.
titan	2018-05-04: 4x Nvidia GTX1080Ti. *

* This partition is shared and you MUST use the --gres to specify the resources you wish to use. It is also encouraged to specify cpu and memory.

From:

<https://howto.cs.uchicago.edu/> - **How do I?**

Permanent link:

<https://howto.cs.uchicago.edu/slurm:peanut?rev=1609970812>

Last update: **2021/01/06 16:06**

