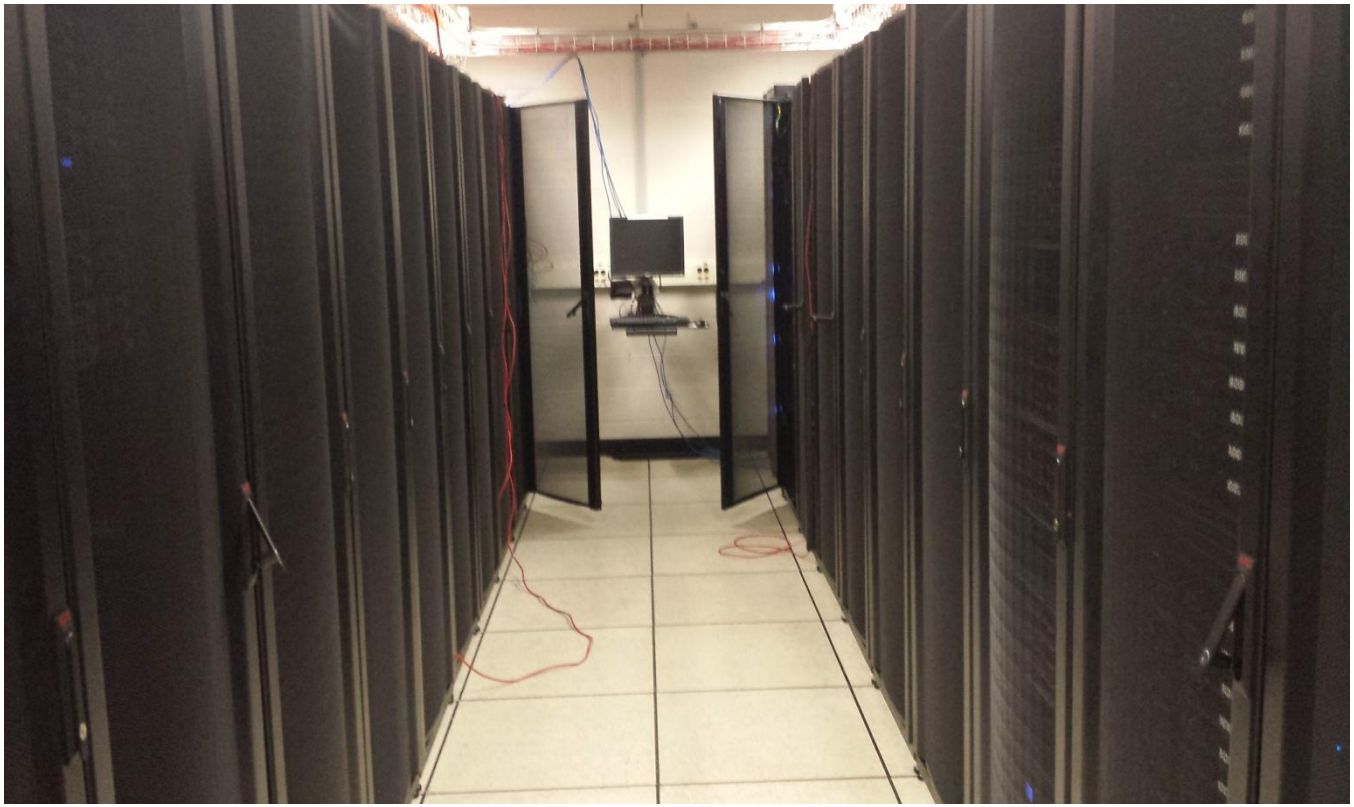

MSU Tier 3 Usage and Troubleshooting

James Koll

Overview

- Dedicated computing for MSU ATLAS members
- Flexible user environment
- ~500 job slots of various configurations
- ~150 TB disk space

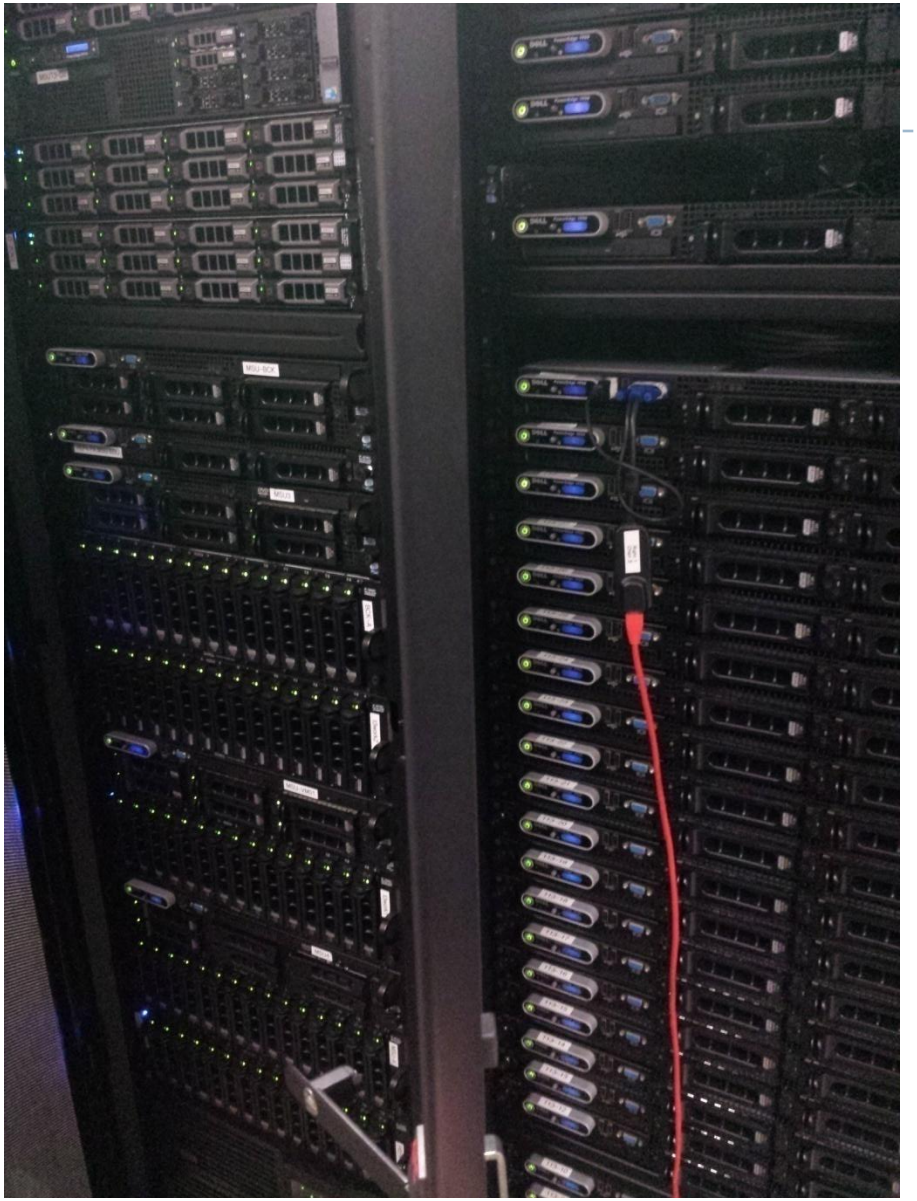

Condor commands

- `condor_q` : What jobs have been submitted from this machine?
- `condor_q –global` : What jobs have been submitted to the cluster?
- `condor_q –better-analyze [jobid]` : What is the status of this job and why?
- `condor_status` How many job slots are available?
- Read the user section of the manual:
  - http://research.cs.wisc.edu/htcondor/manual/v8.0/2_Users_Manual.html

Condor Commands

- condor_q –better-analyze

```
Your job defines the following attributes:

    FileSystemDomain = "msulocal"
    DiskUsage = 1
    RequestDisk = 1
    RequestMemory = 20000

The Requirements expression for your job reduces to these conditions:

         Slots
Step    Matched  Condition
-----  --------  ---------
[0]         258  TARGET.Arch == "X86_64"
[1]         258  TARGET.OpSys == "LINUX"
[3]         258  TARGET.Disk >= RequestDisk
[5]           0  TARGET.Memory >= RequestMemory
[7]         258  TARGET.HasFileTransfer

Suggestions:

    Condition                         Machines Matched    Suggestion
    ---------                         ----------------    ----------
1   ( TARGET.Memory >= 20000 )        0                   MODIFY TO 4096
2   ( TARGET.Arch == "X86_64" )       258
3   ( TARGET.OpSys == "LINUX" )       258
4   ( TARGET.Disk >= 1 )              258
5   ( ( TARGET.HasFileTransfer ) || ( TARGET.FileSystemDomain == "msulocal" ) )
                                      258
```

That's why my job isn't running – I asked for too much memory.

Job constraints

- Memory constraints
  - Add to submit file:
    - Request_memory = 4000

Memory*	# of slots (approximate)
1GB	150
4GB	350

*How we define these slot types is malleable to user needs.

- 30 of the slots are available for quick jobs
  - To use them, just add +MSU_QUEUE = "short" to the submit file prior to the queue statement.
  - Jobs marked as "short" will still run on other slots if no short slots are available.

One line submit

- /msu/data/t3work1/scripts/runcommand.sh [Command you want to run]
  - Runs as "short" queue
  - Brings over environmental variables
  - Prints location of further information, such as log files
  - Prints output text to screen until job is finished
  - Ctrl-C before job is finished will remove job from condor

```
koll@green ~ $ /msu/data/t3work1/scripts/runcommand.sh echo Hello, world
Your job has been submitted. Details of your job can be found at
/msu/data/t3work7/tmp/kollBNDOfaxyw/info.txt

The job output will be printed below:
---------------
Hello, world
---------------
Reminder: Your job details are at
/msu/data/t3work7/tmp/kollBNDOfaxyw/info.txt
Job 133759 has finished running.
```

Disk concurrency limits

- Disk I/O is tricky.
- There are voluntary parameters you can set to tell condor how many units of disk I/O your job uses.
  - Instructions to add are on the wiki (link on last slide).
  - It is good practice to include concurrency limits even if you aren't submitting very many jobs.
- Each disk has 10,000 units of disk I/O available.
  - "Unit" is difficult to define.
    - Random vs sequential operations
    - Reads vs writes

Estimating concurrency limits

- Use this calculator to create a conservative estimate of how many units your job uses.
  - http://hep.pa.msu.edu/concurrencycalc.html
- Probably okay to use this number by default if your jobs do not use a lot of I/O.
- If they do use a lot of I/O, you will want to optimize it.
  - You can try decreasing the concurrency limit and see when you run into CPU wait problems .
  - How to detect CPU wait is discussed in a few slides.

Improving disk performance

- There are tricks you can use to improve disk performance.
  - Slim and skim datasets.
  - Remember sequential reads/writes are faster than random ones.
    - Stagger job submission times and have your jobs copy their input files to local scratch.
    - Produce output files locally, then copy output to remote work directory.
  - Split input files between many work disks.
  - Use faster disks (t3fast vs t3work).

“Green is slow”

- Flexible resources => Lots of user freedom => We can run into problems
- If green is slow, there are a number of things you as a user can do to see what may be wrong.
- There are three things you can check:
  - CPU
  - Memory
  - Disk I/O** (sort of)
- Once you identify the problem, you can identify the responsible user and contact them directly.

Troubleshooting CPU load

- Run `top`, then press shift-p to sort by CPU usage:

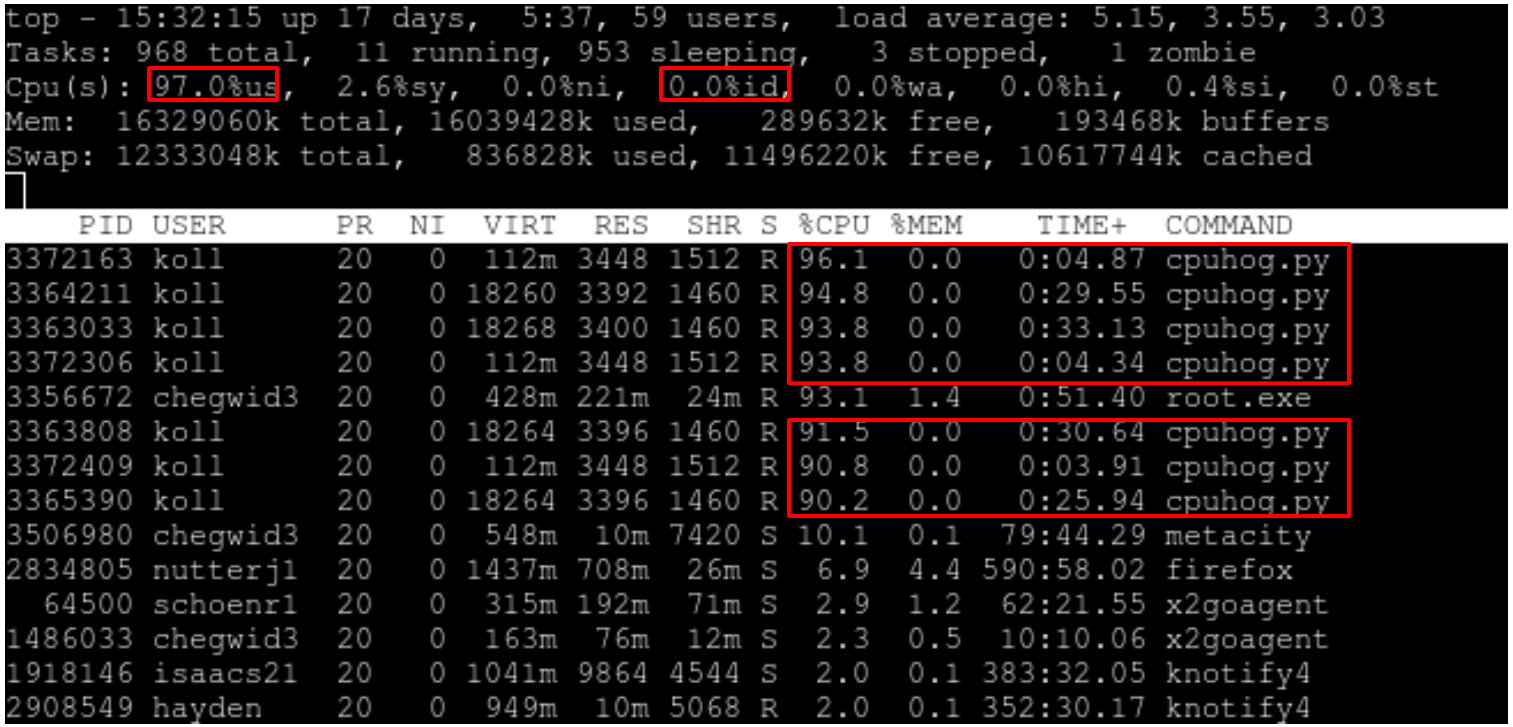

- koll is the bad user. Contact him and tell him to submit his CPU-intensive jobs to condor.

Troubleshooting memory usage

- Is there a memory problem? Run `free`:

```
-bash-4.1$ free
             total       used       free     shared    buffers     cached
Mem:      16329060   16104680     224380          0       7996    1276524
-/+ buffers/cache:   14820160    1508900
Swap:     12333048     997068   11335980
-bash-4.1$
```

Not critically low, but lower than it should be

- Run `top`, then press m to sort by memory usage:

```
top - 15:35:44 up 17 days,  5:40, 57 users,  load average: 4.31, 4.00, 3.32
Tasks: 957 total,   4 running, 949 sleeping,   3 stopped,   1 zombie
Cpu(s): 32.7%us, 12.4%sy,  0.0%ni, 53.6%id,  0.2%wa,  0.0%hi,  1.0%si,  0.0%st

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
3415847 root      20   0 9503m 9.3g  416 R 99.6 59.6   [illegible]
3433108 chegwid3  20   0  426m 219m  24m R 99.6  1.4   0:09.18 root.exe
3506980 chegwid3  20   0  548m 9.9m 6876 S 47.3  0.1  81:08.59 metacity
3957827 zhongdew  20   0  617m  39m 6880 S 33.2  0.2   7:05.34 nautilus
1486033 chegwid3  20   0  162m  75m  11m S  9.5  0.5  10:30.76 x2goagent
3435589 root      20   0  148m 9320 2264 S  3.3  0.1   0:00.10 x2golistsession
2117612 chegwid3  20   0 1074m  30m  16m S  2.6  0.2   1:35.12 nautilus
1918146 isaacs21  20   0 1041m 9532 4236 S  2.3  0.1 383:35.97 knotify4
  67240 schoenr1  20   0  140m 2592 1920 S  2.0  0.0   7:21.74 gvfsd-trash
1486777 chegwid3  20   0  482m 8568 6196 S  2.0  0.1   1:07.30 gnome-settings-
1487187 chegwid3  20   0  336m  14m 8760 S  2.0  0.1   1:18.24 gnome-panel
2824281 nutterj1  20   0 47324 2452 1852 S  2.0  0.0   6:06.98 gvfsd-trash
3308485 gep       20   0 47328 2596 1848 S  2.0  0.0   5:28.93 gvfsd-trash
3958025 zhongdew  20   0 47336 2564 1900 S  2.0  0.0   1:31.35 gvfsd-trash
 806042 truepatr  20   0 47324 2648 1908 S  1.6  0.0   1:23.37 gvfsd-trash
1487931 chegwid3  20   0  140m 3028 2200 S  1.6  0.0   1:10.83 gvfsd-trash
```

Whoa, way too high! Tell this user to check for memory leaks and/or submit to condor!

Troubleshooting Disk I/O

- Easy to cause by accident.
- Check for wait I/O greater than a few percent.

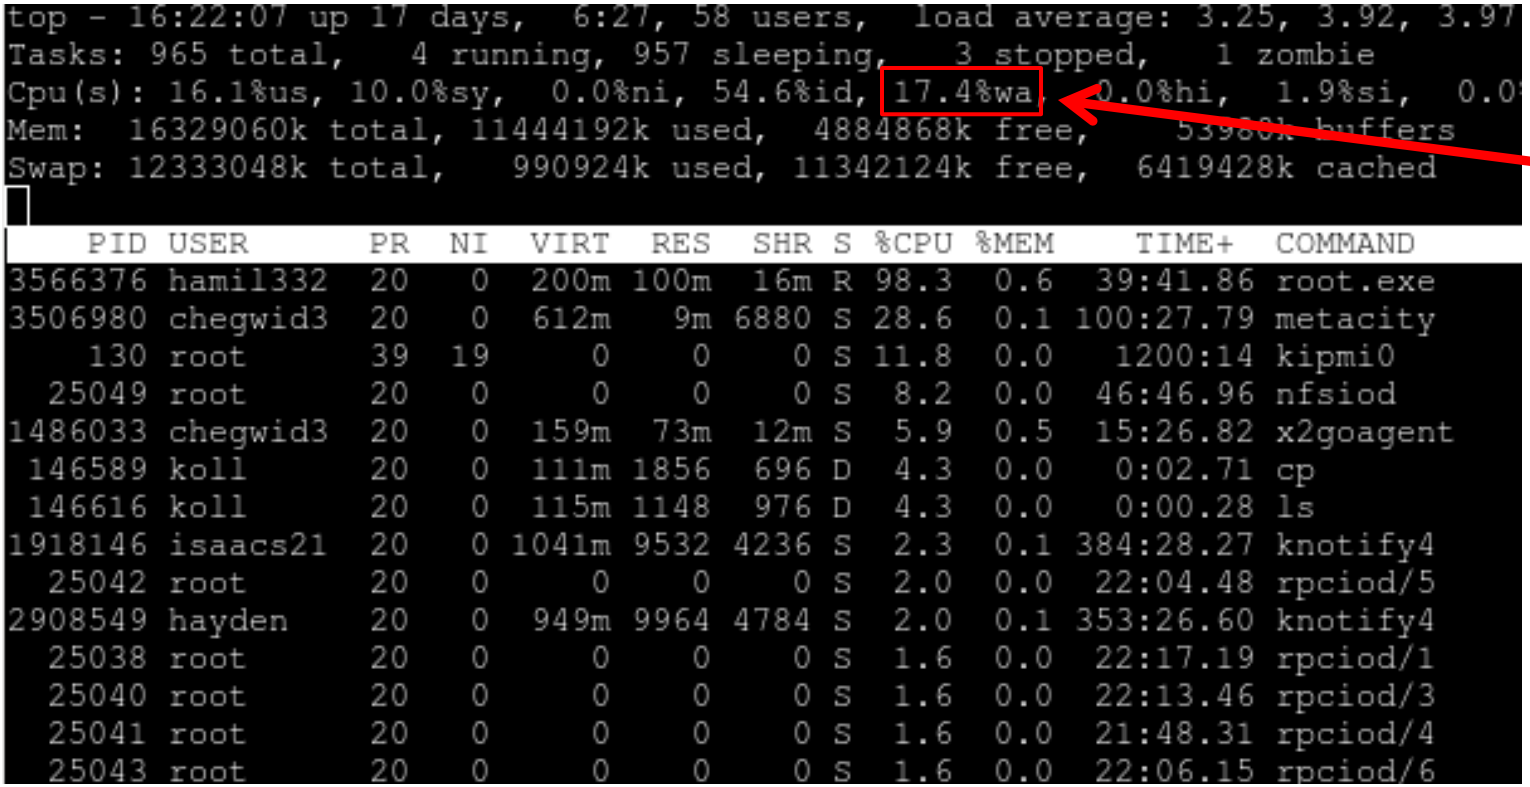

Too high. Look for a user with a large number of jobs running on the condor queue.

- If you can't identify the source, contact me.

Resources

- https://www.aglt2.org/wiki/bin/view/AGLT2/MSUTier3
  - If you think of something that should be in the wiki, please add it.
  - Contact me if you need an account.
- Please feel free to use my contact details, located on the wiki