

Exploring job wait times on Alliance compute clusters: a holistic view

SHARCNET General Interest Webinar
James Desjardins
October 4th, 2023

Overview

Scheduling and account usage states

Assessing unexpected wait times

The show_jobs notebook repository

Population and priority summaries

partition-stats and clusterstats

Measuring job usage and queue load


Jobs request a reservation of a set of resources

The jobs' requested resources are reserved during the run time and sit pending during the queue time

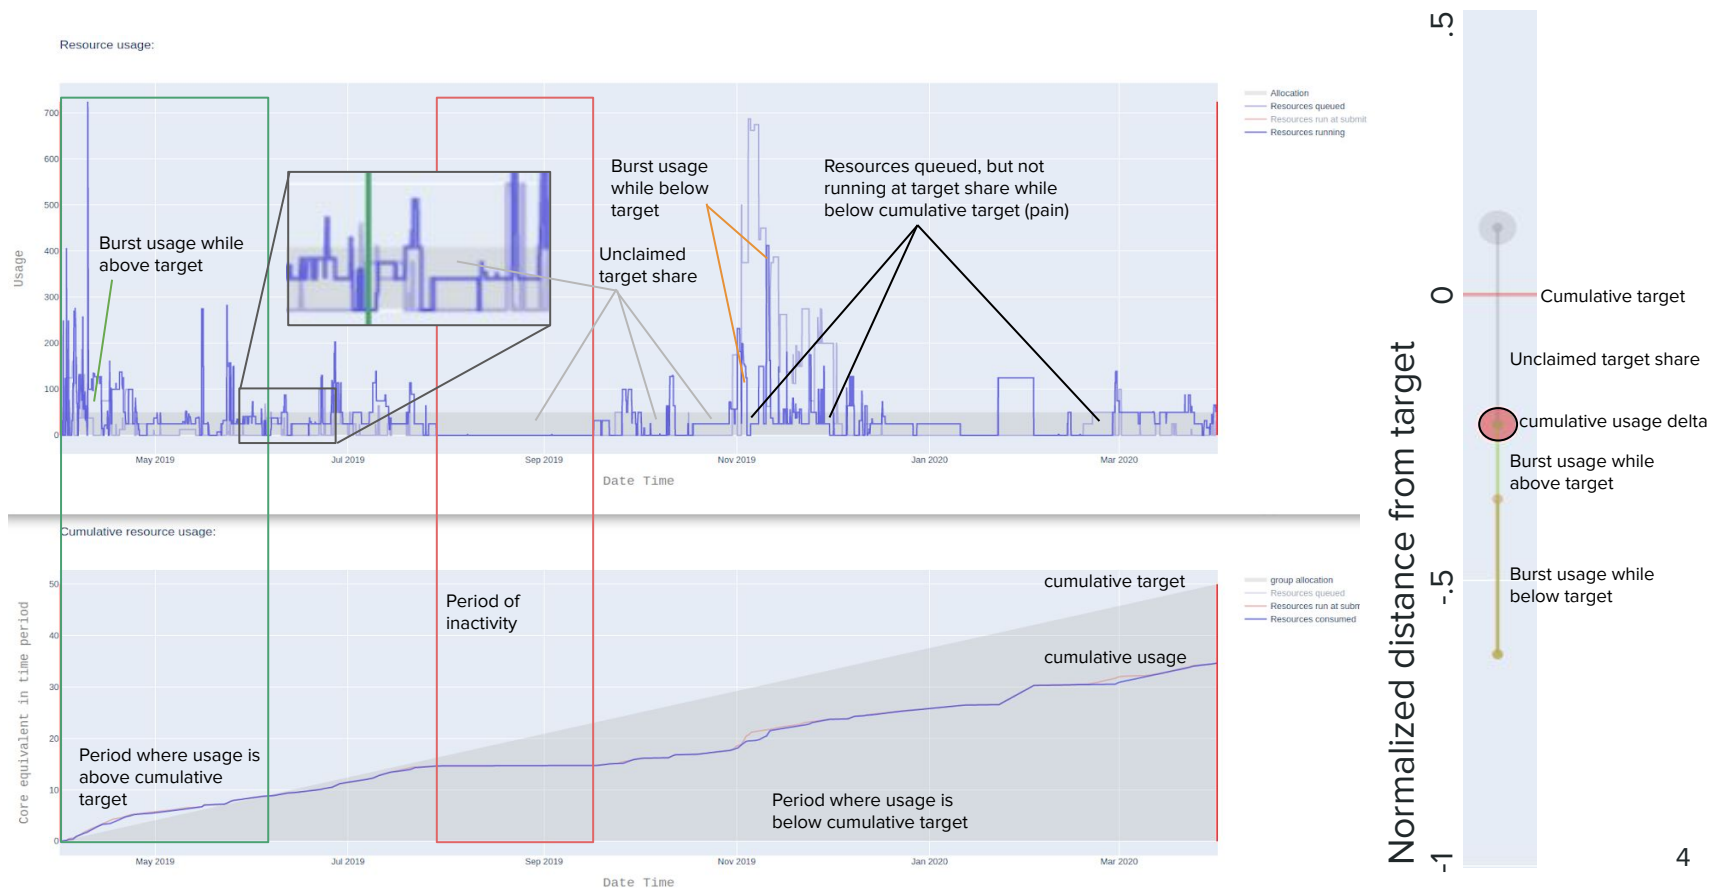


From the job records we can measure each account's resource reservation and queue load

Regardless of the account's queue load, the scheduler tries to maintain usage in line with the account's target share over time.



Account states in relation to job usage and queue load




Assessing unexpected wait times

Unexpected wait times can occur for several identifiable reasons

Some of these reasons can be related to system states which interact with job properties


Resource usage:




```
[ ]: fig_cumu=vcv.cumu_plot(targ, queued, running, fig_out=account_name+'_'+cumu_plot.html', user_run=user_run,
fig_cumu.show())

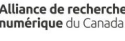
[16]: sj.job_scatter(jobs_df, x_var='submit', y_var='eligible_wait_hours', c_var='partition', s_var='ncpus',
var_labels={'eligible_wait_hours': "wait time (hours)", \
"mem": "Requested memory (MB)", \
"reqcpus": "Requested CPUs", \
"partition": "partition name"},
title='job shape scatter: start time by wait time (colour = partition, size = NCPUs)')
```

job shape scatter: start time by wait time (colour = partition, size = NCPUs)






Digital Research
Alliance of Canada



Alliance de recherche
numérique du Canada

Home · Systems · Websites · Français



Graham is experiencing a limited outage
Some systems are available with limited access.

✓ Operational

⚠ Available with conditions

🔴 Outage

📅 Scheduled events

Current Incident

Status

Current Incidents

⚠ Reduced Power Capacity For Friday July 22 - Capacité De Puissance Réduite

Issue history

Show 10 entries

Search

Start Date	Status	Title
5 days ago	Open	Reduced Power Capacity For Friday July 22 - Capacité De Puissance Réduite
2022-07-25	Resolved	Project Filesystem problem - Problème avec le système de fichiers (project)
2022-07-21	Resolved	Reduced Power Capacity - Capacité De Puissance Réduite
2022-07-19	Resolved	Reduced Power Capacity - Capacité De Puissance Réduite
2022-07-11	Resolved	Filesystem problem - Problème de système de fichiers
2022-07-06	Resolved	Partial outage
2022-06-17	Resolved	Filesystem problem - Problème de système de fichiers
2022-06-06	Resolved	Slow Transfer Problem
2022-06-06	Resolved	Filesystem problem - Problème de système de fichiers
2022-05-30	Resolved	Scheduler problem - Problème d'ordonnanceur

Showing 1 to 10 of 85 entries

Previous 1 2 3 4 5 ... 9 Next

Maintenance history

Show 10 entries

Search

Start Date	Status	Title
2022-06-23	Resolved	Planned Outage - Arrêt planifié
2022-05-25	Resolved	Planned Outage - Arrêt planifié
2022-03-28	Resolved	Planned Outage - Arrêt planifié
2021-07-17	Resolved	Planned Outage - Arrêt planifié
2021-07-15	Resolved	Planned Outage - Arrêt planifié
2021-07-13	Resolved	Planned Outage - Arrêt planifié
2021-07-14	Resolved	Planned Outage - Arrêt planifié

Showing 1 to 7 of 7 entries

Previous 1 Next

cpubase_bycore_b1

cpubase_interac

cpubackfill, cpubase_interac

cpubackfill


cpubase_bycore_b1, cpubackfill

cpubase_bycore_b4

Assessing unexpected wait times


Jobs in queue are not always eligible to run

job shape scatter: start time by wait time (colour = partition, size = NCPUs)



```
[ ]: sj.job_scatter(jobs_df, x_var='submit', y_var='mem_delta', c_var='partition', s_var='ncpus',  
                  var_labels={"mem_delta": "memory delta", \  
                              "mem": "Requested memory (MB)", \  
                              "reqcpus": "Requested CPUs", \  
                              "partition": "partition name"},  
                  title='job shape scatter: start time by wait time (colour = partition, size = NCPUs)')  
  
[18]: sj.job_scatter(jobs_df, x_var='submit', y_var='eligible_delta_hours', c_var='partition', s_var='reqcpus',  
                    var_labels={"eligible_delta_hours": "eligible delta (hours)", \  
                                "mem": "Requested memory (MB)", \  
                                "reqcpus": "Requested CPUs", \  
                                "partition": "partition name"},  
                    title='job shape scatter: start time by wait time (colour = partition, size = NCPUs)')
```

job shape scatter: start time by wait time (colour = partition, size = NCPUs)

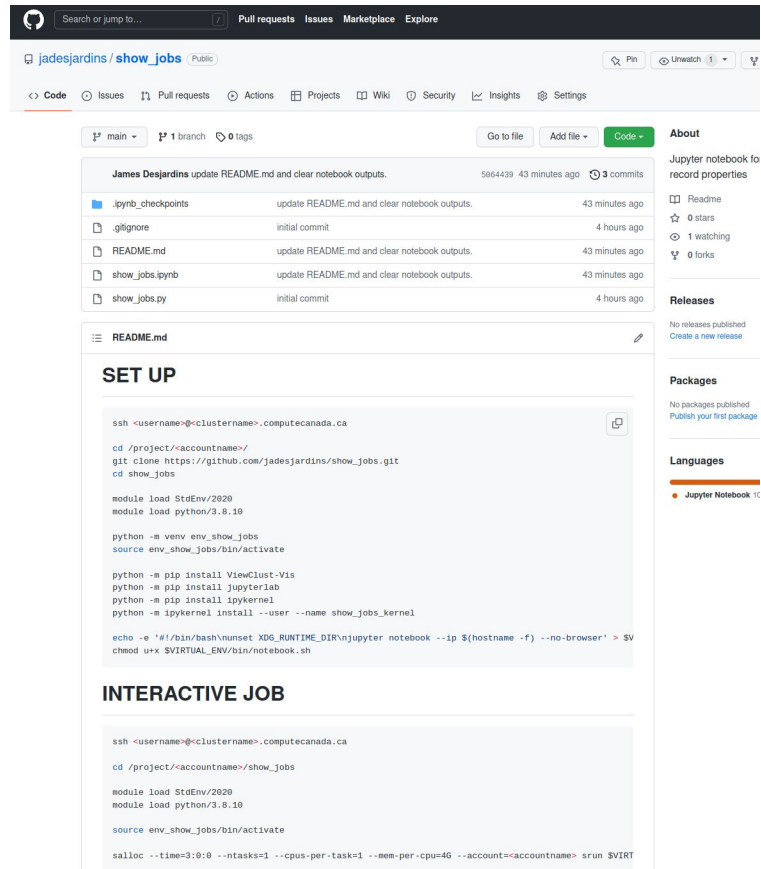


The show_jobs notebook repository

https://github.com/jadesjardins/show_jobs

The show_use notebook repo contains tool for querying job records from the system's Slurm scheduler and provides some interactive visualizations

The repository's README.md provides instructions for cloning the repo, installing the requirements, setting up the virtual environment and kernel, starting the interactive notebook job, tunneling to the compute node, navigating to the notebook from a browser.



The screenshot displays the GitHub interface for the repository `jadesjardins/show_jobs`. The repository is public and has 3 commits. The commit history shows updates to the README.md and clear notebook outputs, as well as initial commits for .ipynb_checkpoints, .gitignore, README.md, show_jobs.ipynb, and show_jobs.py.

The README.md file is open, showing the "SET UP" section with the following instructions:

```
ssh <username>@<clustername>.computeCanada.ca

cd /project/<accountname>/
git clone https://github.com/jadesjardins/show_jobs.git
cd show_jobs

module load StdEnv/2020
module load python/3.8.10

python -m venv env_show_jobs
source env_show_jobs/bin/activate

python -m pip install ViewClust-Vis
python -m pip install jupyterlab
python -m pip install ipynbkernel
python -m ipynbkernel install --user --name show_jobs_kernel

echo -e '#!/bin/bash\nunset XDG_RUNTIME_DIR\njupyter notebook --ip $(hostname -f) --no-browser' > $V
chmod u+x $VIRTUAL_ENV/bin/notebook.sh
```

The "INTERACTIVE JOB" section shows the following instructions:

```
ssh <username>@<clustername>.computeCanada.ca

cd /project/<accountname>/show_jobs

module load StdEnv/2020
module load python/3.8.10

source env_show_jobs/bin/activate

salloc --time=3:0:0 --ntasks=1 --cpus-per-task=1 --mem-per-cpu=4G --account=<accountname> srun $VIRT
```

Population state of accounts on a cluster

All of the jobs on a cluster are competing for resources based on fair share.

The state of all accounts on a cluster can impact job wait times.




Ranking the state of allocation accounts on a cluster

- system: cedar
- node type: cpu
- query period: 2022-04-04 to 2022-04-18

Ranked account plot

Each coloured circle represents an account. The Y position of the coloured circle represents the usage achieved during the query period normalized to the accounts target share (allocation). The length of the marks whisker corresponds to the accounts quantity of unclaimed resources. Unclaimed resources are calculated as the area between the accounts target share (allocation) and the sum of its running and pending job resources on each day while the accounts cumulative usage is below its cumulative target. Unclaimed resources therefore represent the portion of the accounts achieved usage deviation from target that is due to the account not submitting jobs to achieve their allocated resource usage. The colour of the account circles represents the usage achieved (synonymous with Y axis) at a range of 0 to 1 (all accounts that have surpassed their allocation are yellow). Accounts are sorted from left to right based on the rank of their usage achieved plus unclaimed resources (the top of each accounts whisker).



Ranked account pain table

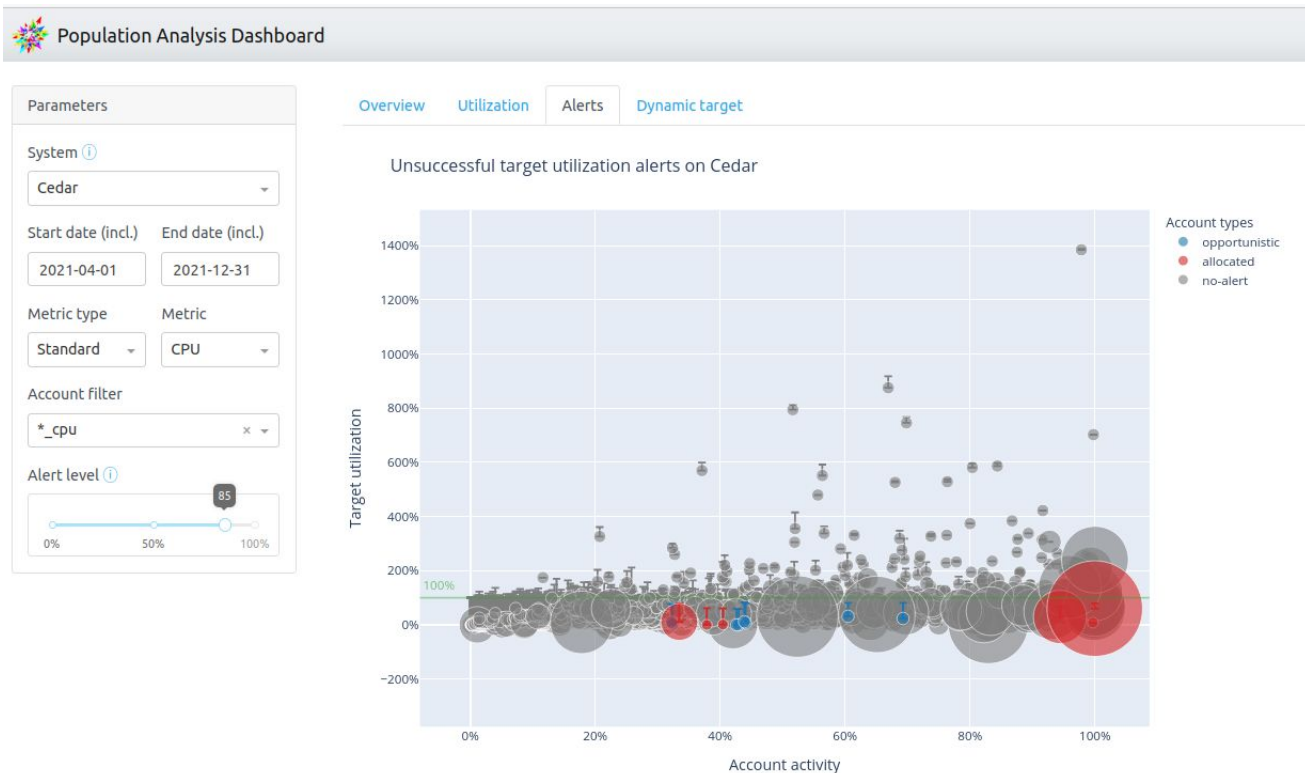
When the normalized usage achieved plus unclaimed resources [top of whiskers in the figure above] does not reach a value of 1 it means that the account had job resources pending in the while the account cumulative usage was below the cumulative target but the usage did not happen for some reason. The table below provides a ranked list of accounts whose normalized usage achieved plus unclaimed resources was below a value of .9. Rows in the table are sorted ascending by the *normalized use plus unclaimed* (sort column indicated by). Assessing the issue of each case typically requires a view of the accounts usage time series. To facilitate viewing the account usage time series the right column usage dash link provides a direct link to the dash.computecanada.ca DANT portal view of the rows account over the reports query period. Edit panels of the usage portal can be used to examine longer time intervals as well the cumulative usage status.

Top pain account table

account	pain rank	target (allocation)	normalized usage achieved	*normalized use plus unclaimed	usage dash link
mp-dlabog-ad	1.0	1979.0	0.15185618304365145	0.15185618304365145	usage dash link
mp-hansi	2.0	456.0	0.19784878863826236	0.19784878863826236	usage dash link
mp-pomes-ad	3.0	561.0	0.06305025040319158	0.27732596468890587	usage dash link
mp-holt	4.0	1636.0	0.39194821597351587	0.39194821597351587	usage dash link
mp-jorniker	5.0	281.0	0.4171217246168606	0.4840323868060074	usage dash link
mp-yearman	6.0	40.0	0.16071549882638335	0.5832439660773705	usage dash link

Alerting of pain states in staff portal.


A staff portal provides an interface for querying the state of a cluster at a given time and alerts of accounts in an unexpected wait state.



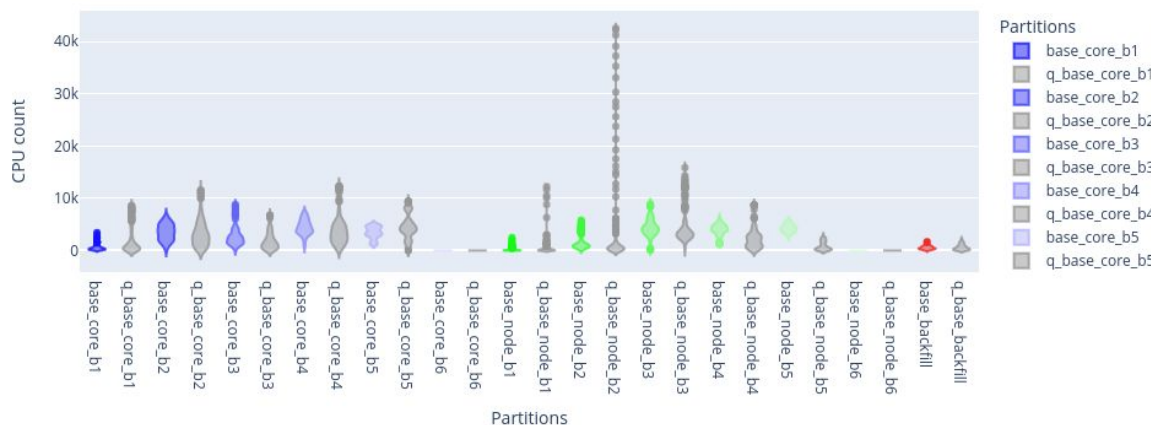
Assessing a cluster's state of job priorities

The start time priority of job records on a cluster can be used to assess the types of load and the fairness with which heterogeneous workloads are deployed to the compute nodes.

Distributions of job start time priority by partition




Distributions of CPUs allocated to running jobs (colour) and pending in the queue (gray) over time by partition



Assessing a cluster's state of job priorities

Exploring the populations of job start time priorities over different periods of time allows staff to identify potential fairness issues.



partition-stats snapshot of the partition and job state

Users can get a high-level view of the current state of a cluster's partition and job state using "partition-stats"

```
[jdesjard@gra-login1 hw demo]$ partition-stats
```

Node type	Max walltime					
	3 hr	12 hr	24 hr	72 hr	168 hr	672 hr
Number of Queued Jobs by partition Type (by node:by core)						
Regular	0:6	16:297	4:594	62:69	45:627	0:0
Large Mem	0:1	126:397	1:333	5:55	2:135	0:0
GPU	0:0	0:69	0:129	12:222	31:4	0:0
Number of Running Jobs by partition Type (by node:by core)						
Regular	3:11	11:29	63:738	52:551	63:1082	0:0
Large Mem	0:9	5:123	1:28	7:9	1:21	0:0
GPU	0:2	0:37	2:29	50:78	9:5	0:0
Number of Idle nodes by partition Type (by node:by core)						
Regular	1:0	1:0	1:0	1:0	0:0	0:0
Large Mem	0:0	0:0	0:0	0:0	0:0	0:0
GPU	37:32	25:20	20:20	18:12	12:10	0:0
Total Number of nodes by partition Type (by node:by core)						
Regular	1074:632	1074:632	1044:612	753:381	371:170	0:0
Large Mem	27:19	27:19	24:16	20:4	5:4	0:0
GPU	214:136	194:116	182:110	139:81	42:39	0:0

clusterstats snapshot of the partition and job state

Users can also dig down into details about the priority of status of their jobs within current partition and queue state of the cluster using “clusterstats”

```
[jdesjard@gra-login1 hw_demo]$ clusterstats
[✓] Loading node information (success, loaded cached version that is 0 min old)
[✓] Loading job information (success, loaded cached version that is 0 min old)
[✓] Loading share information (success, loaded cached version that is 2 min old)
Information about? User
Information about ? Jobs
Information about Job ? 11734779 (pending)
Information on ? Basic
Job:11734779 state: pending partition: cpularge_bycore_b5 priority: 1816197
    This job is ranked 50 of 1260 in terms of priority
Information on ? Report
Job 11734779:
    This pending job belongs to user jdesjard, accounting group def-jdesjard_cpu in partition cpularge_bycore_b5
    Nodes that can possibly run the job:
        Total: 4 Busy: 3 Down: 1 Idle: 0
        Node Type (cpu=64, Mem=3095000): Total 2 Down 1 Idle 0
        Node Type (cpu=32, Mem=514500): Total 2 Down 0 Idle 0
        This job is ranked 50 of 1260 in terms of priority on these nodes

Information on ? Long Report
Job 11734779:
    This pending job belongs to user jdesjard, accounting group def-jdesjard_cpu in partition cpularge_bycore_b5
    Nodes that can possibly run the job:
        Total: 4 Busy: 3 Down: 1 Idle: 0
        Node Type (cpu=64, Mem=3095000): Total 2 Down 1 Idle 0
        Node Type (cpu=32, Mem=514500): Total 2 Down 0 Idle 0
    This job is ranked 50 of 1260 in terms of priority on these nodes
    Jobs on these nodes are in Partitions:
        cpularge_interac,cpularge_bynode_b1,cpularge_bynode_b2,cpularge_bynode_b3
        cpularge_bynode_b4,cpularge_bynode_b5,cpularge_bynode_b6,cpularge_bycore_b1
        cpularge_bycore_b2,cpularge_bycore_b3,cpularge_bycore_b4,cpularge_bycore_b5
        cpularge_bycore_b6

    Some nodes are reserved for jobs that are no more than <168> hours long:
        Total: 4 Busy: 3 Down: 1 Idle: 0
        Node Type (cpu=64, Mem=3095000): Total 2 Down 1 Idle 0
        Node Type (cpu=32, Mem=514500): Total 2 Down 0 Idle 0
    This job is ranked 50 of 1260 in terms of priority on these nodes
    Jobs on these nodes are in Partitions:
        cpularge_interac,cpularge_bynode_b1,cpularge_bynode_b2,cpularge_bynode_b3
        cpularge_bynode_b4,cpularge_bynode_b5,cpularge_bycore_b1,cpularge_bycore_b2
        cpularge_bycore_b3,cpularge_bycore_b4,cpularge_bycore_b5
```


clusterstats snapshot of the partition and job state

Given all of the dynamic states of the job population of a cluster, a job's rank within the queue of jobs competing for the same resources is one of the best predictors of wait time ... but wait time is still impossible to predict.

```
[jdesjard@gra-login1 hw_demo]$ clusterstats
[✓] Loading node information (success, loaded cached version that is 0 min old)
[✓] Loading job information (success, loaded cached version that is 0 min old)
[✓] Loading share information (success, loaded cached version that is 3 min old)
Information about? User
Information about ? Jobs
Information about Job ? 11734877 (pending)
Information on ? Basic
Job:11734877 state: pending partition: cpularge_bycore_b5 priority: 2268095
  This job is ranked 3 of 1262 in terms of priority
Information on ? Back
Information about Job ? 11734779 (pending)
Information on ? Basic
Job:11734779 state: pending partition: cpularge_bycore_b5 priority: 2268099
  This job is ranked 2 of 1262 in terms of priority
Information on ? Quit
[jdesjard@gra-login1 hw_demo]$ sacct -aX -A def-jdesjard_cpu -S 2023-10-04 -o jobid,submit,start,partition%32
```

JobID	Submit	Start	Partition
11720155	2023-10-04T06:39:01	2023-10-04T06:39:02	cpubase_bycore_b1
11720173	2023-10-04T06:44:59	2023-10-04T06:49:39	cpubase_bycore_b5
11720208	2023-10-04T06:56:13	2023-10-04T07:11:39	cpularge_bycore_b5
11734727	2023-10-04T10:28:57	2023-10-04T10:31:36	cpubase_bycore_b5
11734760	2023-10-04T10:34:02	2023-10-04T10:35:01	cpubase_bycore_b1
11734779	2023-10-04T10:35:44	Unknown	cpularge_bycore_b5
11734877	2023-10-04T10:50:41	Unknown	cpularge_bycore_b5

```
[jdesjard@gra-login1 hw_demo]$
```

Thank you for your attention!

Questions?