SHARCNET General Interest Webinar Series

# Tips for identifying when job wait times can be reduced by job submission parameter changes

James Desjardins
High Performance Computing Consultant
SHARCNET, Brock University
December 15th, 2021

Compute • Calcul Ontario

SHARCNET™

compute canada | calcul canada

# Overview

- Ranking accounts by “wait time pain”
- Heterogeneity of the general purpose system nodes
- Node partitioning and preferred job parameters
- Viewing a snapshot of a system’s partitions using partition-stats
- Viewing node resources within partitions using cluterstats
- Examining job parameters and their relation to the job’s partition
- Submitting jobs to specific partitions

# Single account job usage (allocation)

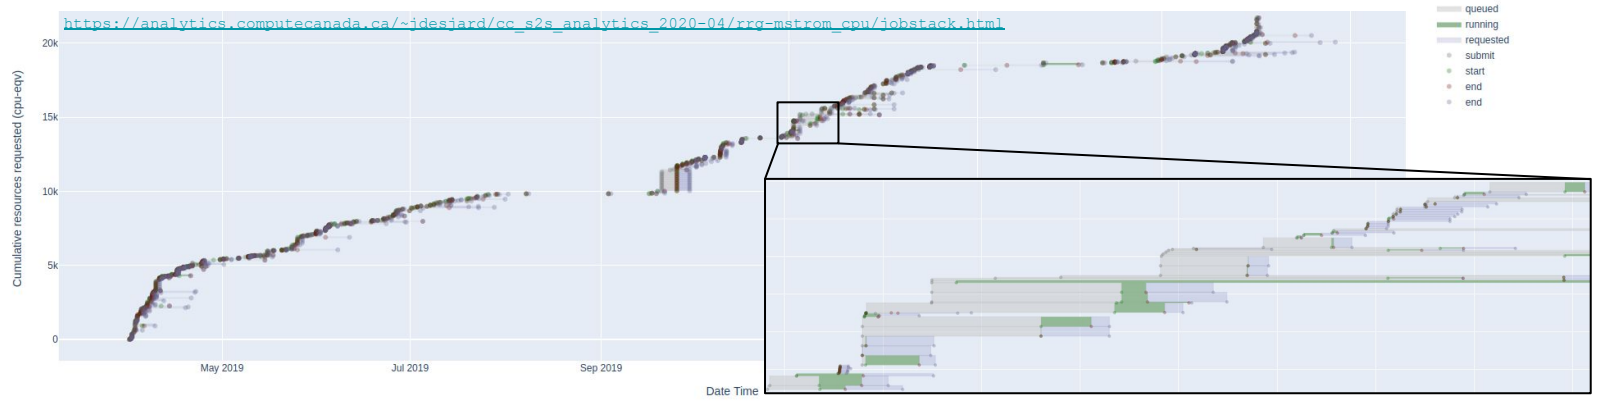


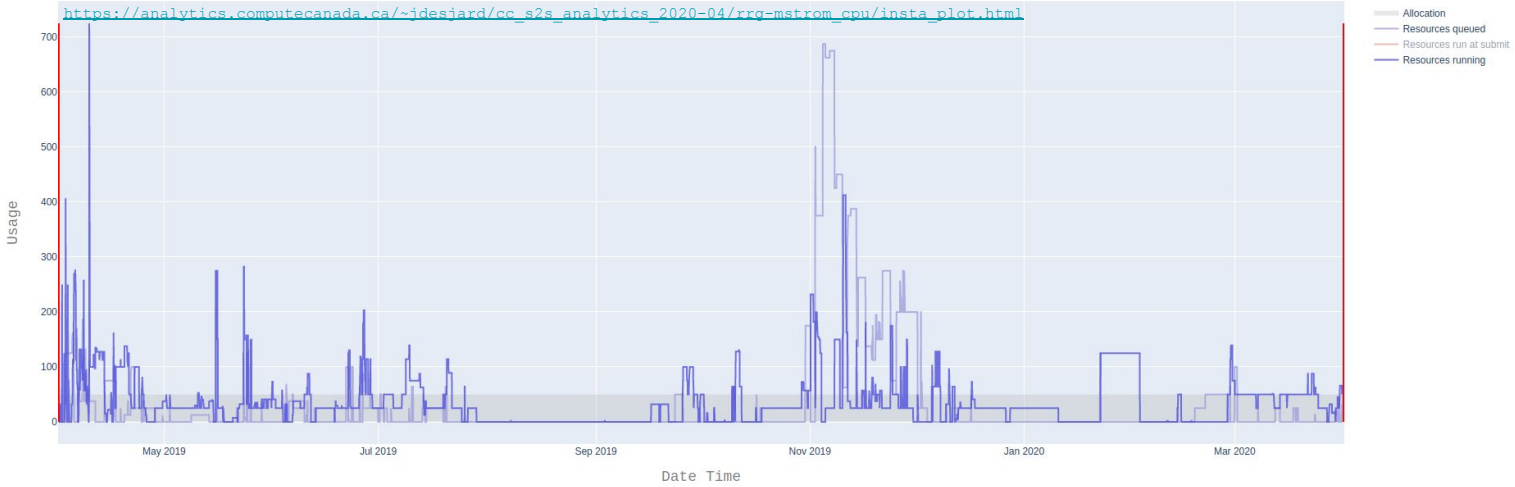

# Single account job usage (allocation) states

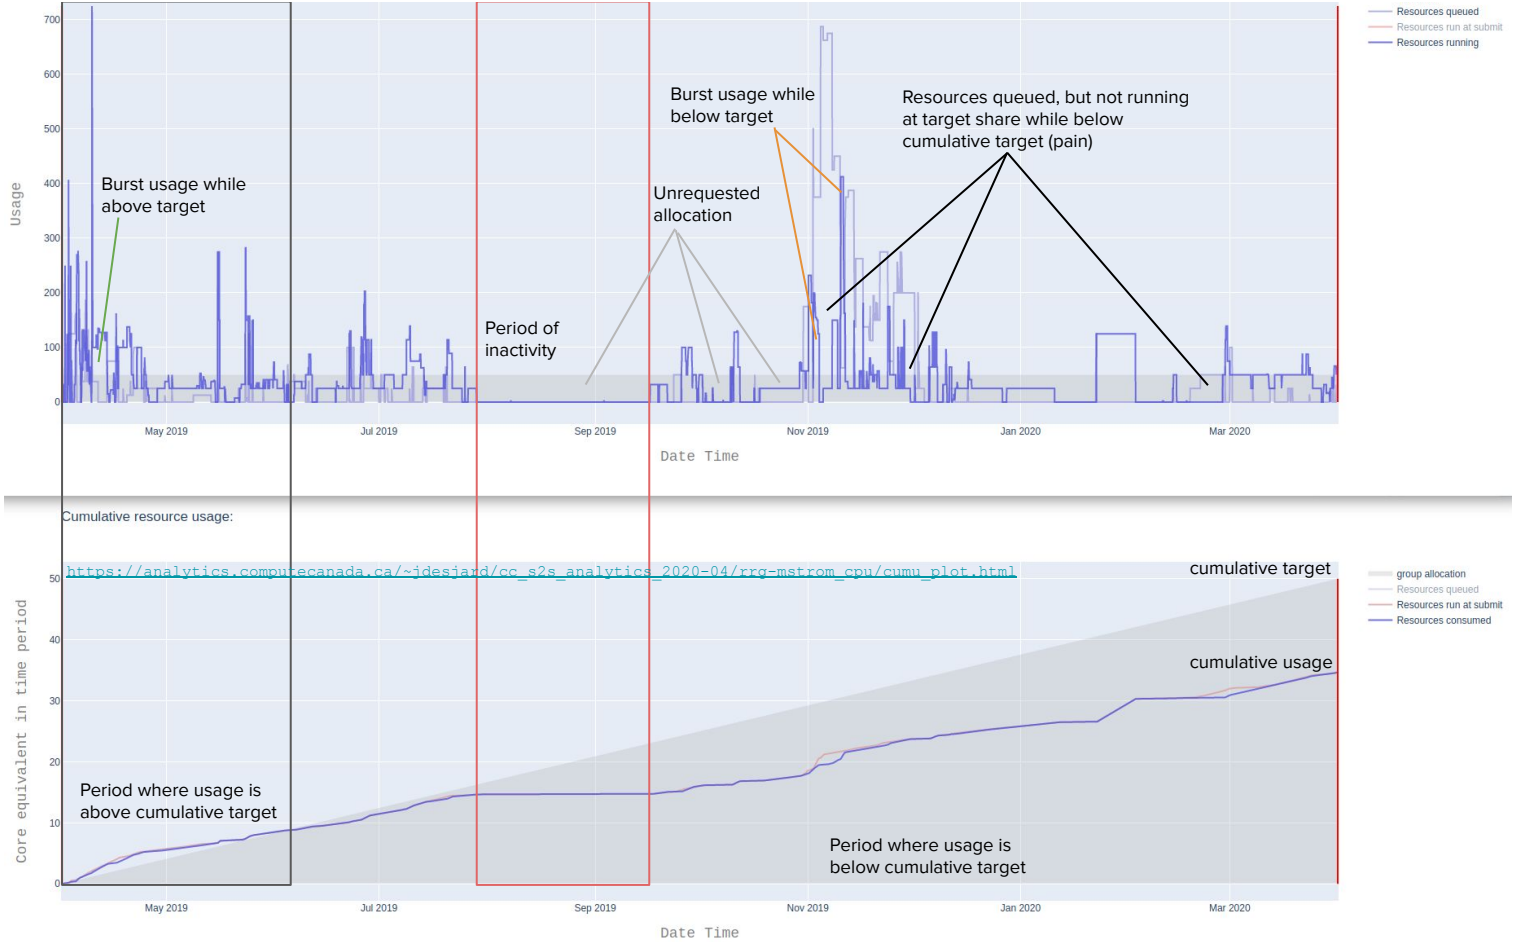

# Single account job usage (allocation) state summary

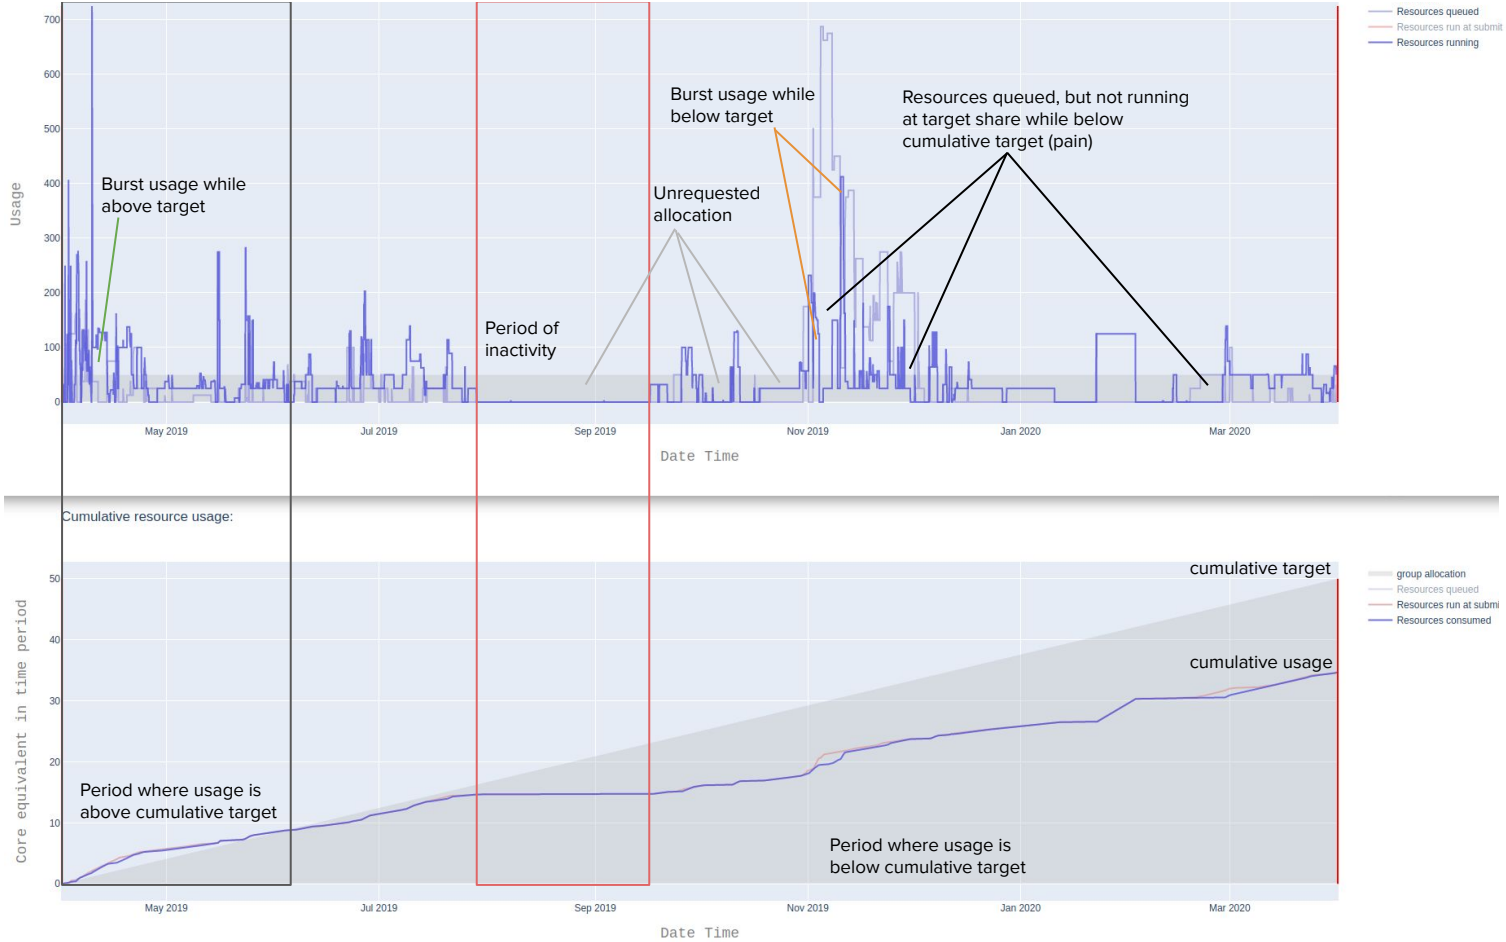


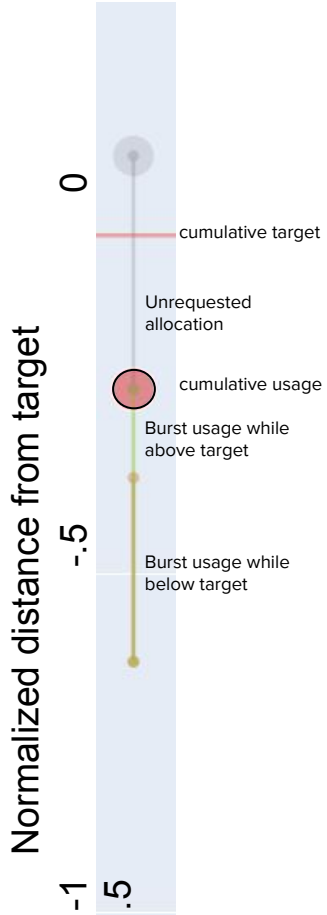

# Cluster accounts population summaries

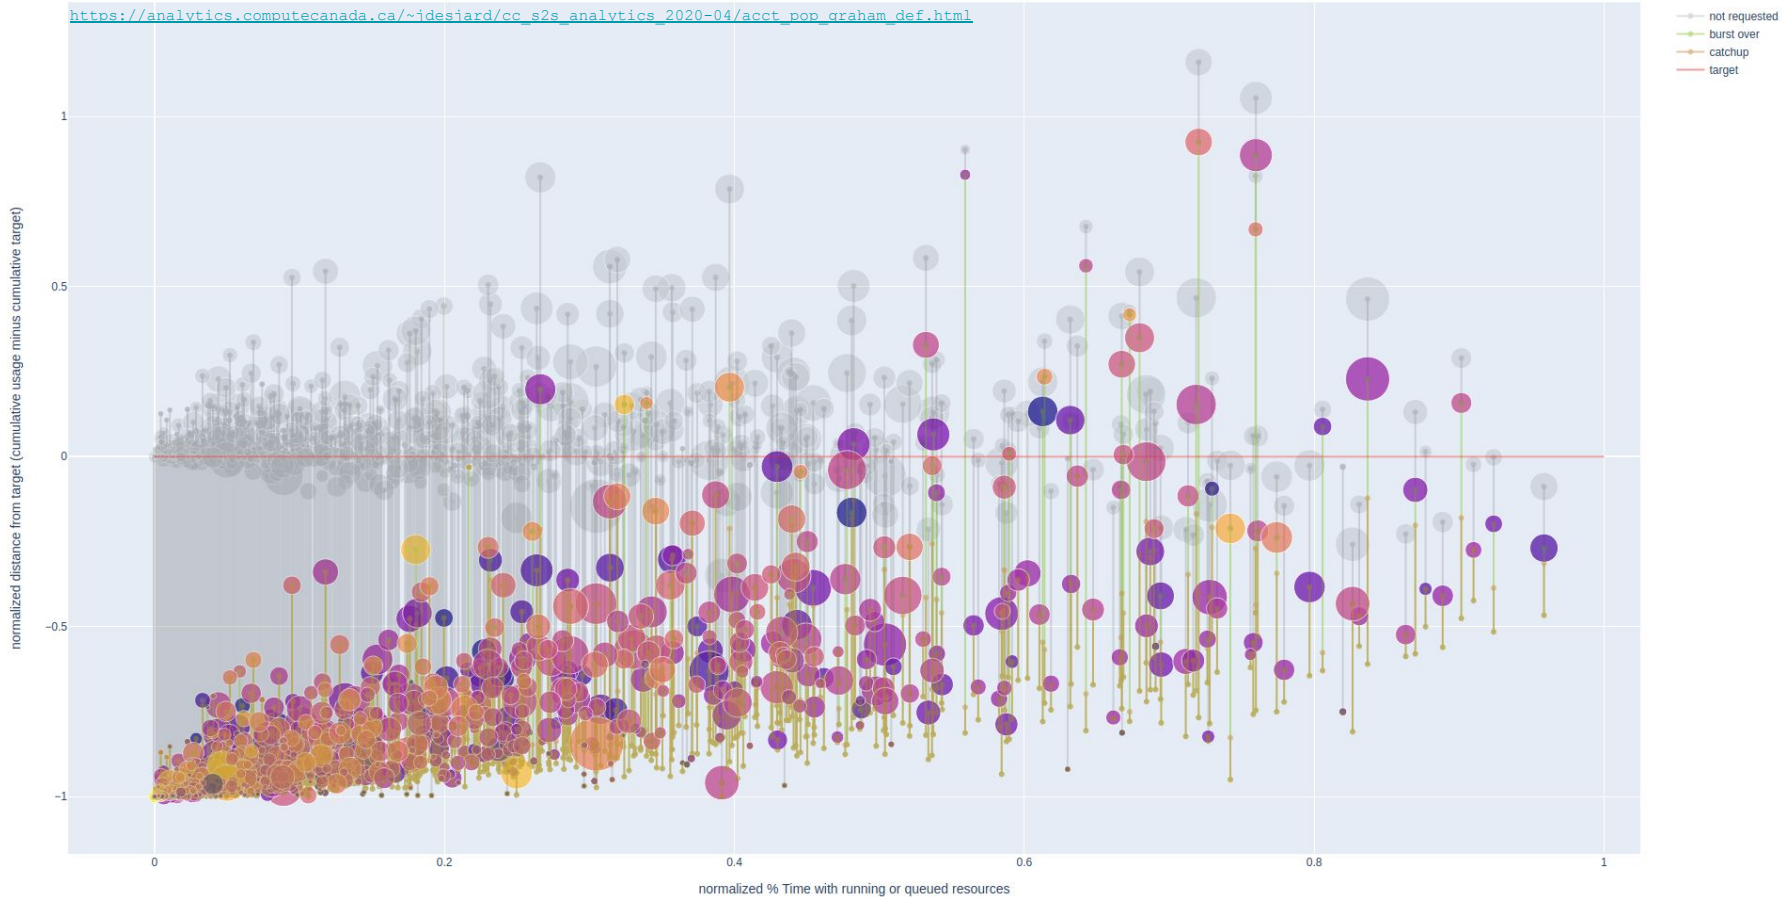

# Cluster node types and their quantities

## Node characteristics [edit]

A total of 41,548 cores and 520 GPU devices, spread across 1,185 nodes of different types; note that Turbo Boost is activated for the ensemble of Graham nodes.

| nodes | cores | available memory | CPU | storage | GPU |
|---|---|---|---|---|---|
| 903 | 32 | 125G or 128000M | 2 x Intel E5-2683 v4 Broadwell @ 2.1GHz | 960GB SATA SSD | - |
| 24 | 32 | 502G or 514500M | 2 x Intel E5-2683 v4 Broadwell @ 2.1GHz | 960GB SATA SSD | - |
| 56 | 32 | 250G or 256500M | 2 x Intel E5-2683 v4 Broadwell @ 2.1GHz | 960GB SATA SSD | - |
| 3 | 64 | 3022G or 3095000M | 4 x Intel E7-4850 v4 Broadwell @ 2.1GHz | 960GB SATA SSD | - |
| 160 | 32 | 124G or 127518M | 2 x Intel E5-2683 v4 Broadwell @ 2.1GHz | 1.6TB NVMe SSD | 2 x NVIDIA P100 Pascal (12GB HBM2 memory) |
| 7 | 28 | 178G or 183105M | 2 x Intel Xeon Gold 5120 Skylake @ 2.2GHz | 4.0TB NVMe SSD | 8 x NVIDIA V100 Volta (16GB HBM2 memory). Note that one node is only populated with 6 GPUs. |
| 2 | 40 | 377G or 386048M | 2 x Intel Xeon Gold 6248 Cascade Lake @ 2.5GHz | 5.0TB NVMe SSD | 8 x NVIDIA V100 Volta (32GB HBM2 memory),NVLINK |
| 6 | 16 | 187G or 191840M | 2 x Intel Xeon Silver 4110 Skylake @ 2.10GHz | 11.0TB SATA SSD | 4 x NVIDIA T4 Turing (16GB GDDR6 memory) |
| 30 | 44 | 187G or 191840M | 2 x Intel Xeon Gold 6238 Cascade Lake @ 2.10GHz | 5.8TB NVMe SSD | 4 x NVIDIA T4 Turing (16GB GDDR6 memory) |
| 136 | 44 | 187G or 191840M | 2 x Intel Xeon Gold 6238 Cascade Lake @ 2.10GHz | 879GB SATA SSD | - |

# Node assignment to partitions

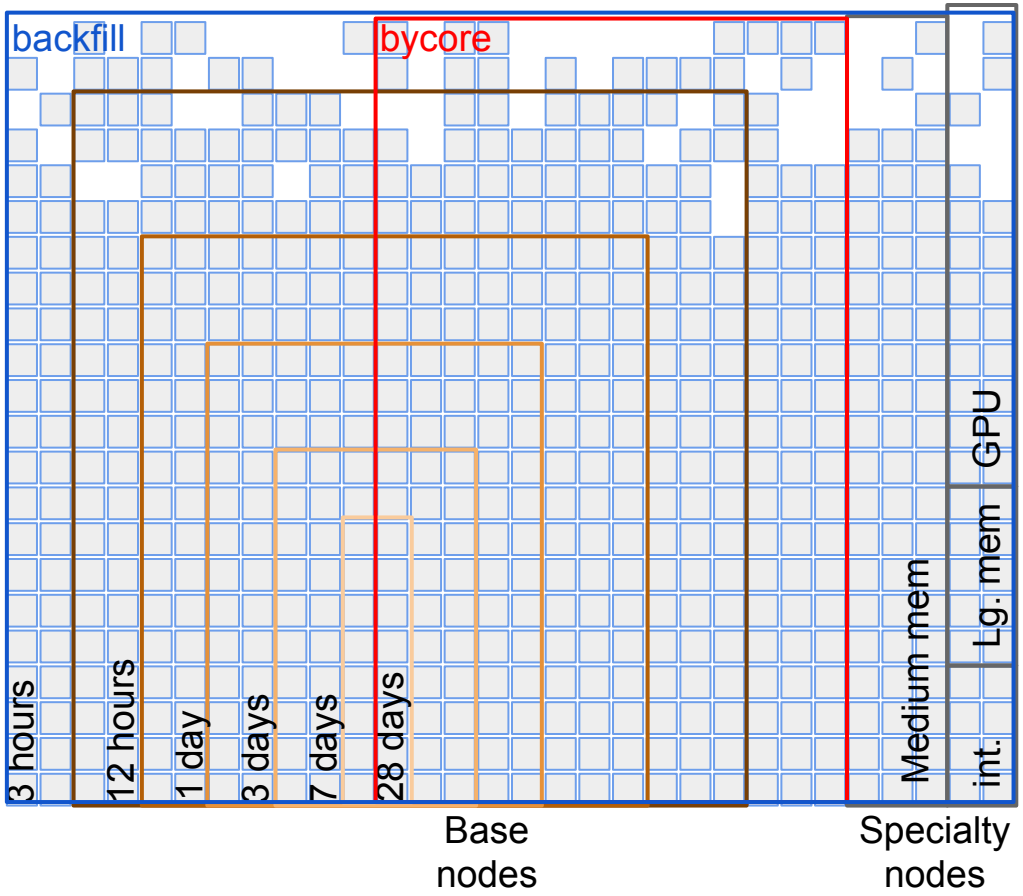


## Partitions

- Restrict jobs of specific shapes to node sets
- Full node jobs can run on most any node (bynode)
- Jobs 24 hours and shorter can run on most any node
- Longer run time jobs have access to fewer nodes
- Partial node jobs (bycore) have access to fewer nodes
- Backfill jobs can run on most any node

# Job triage to node via partition

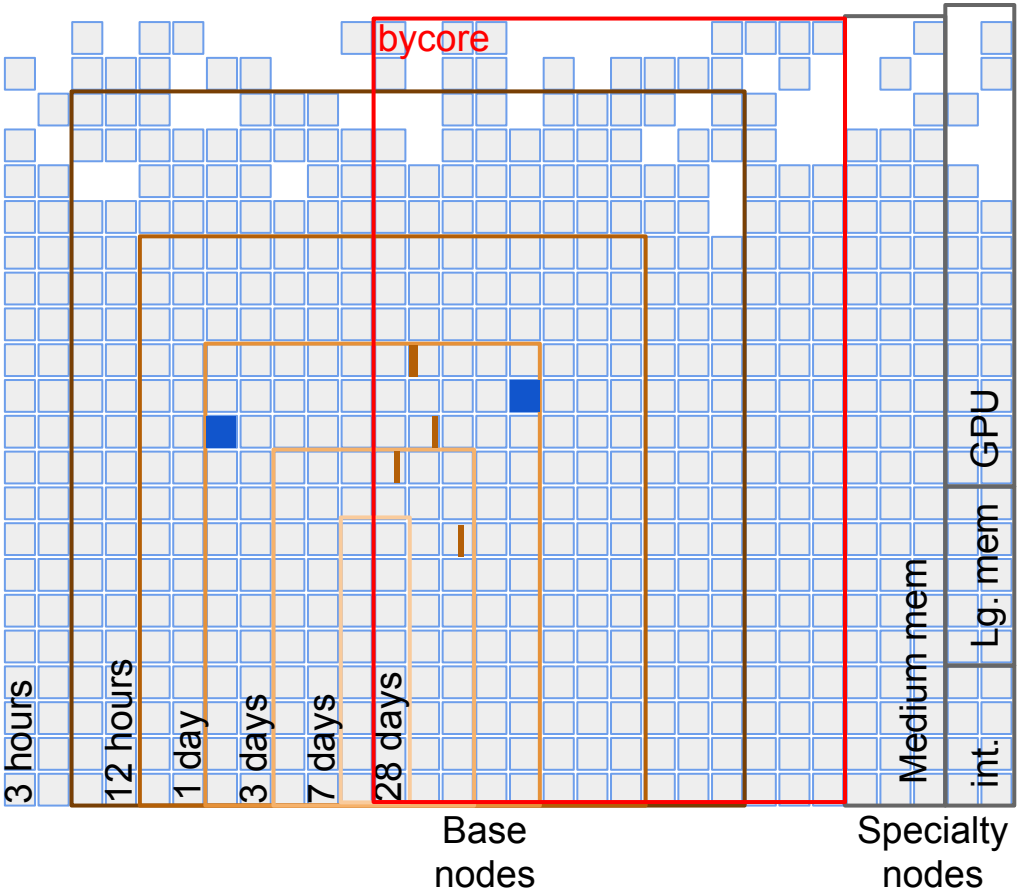


## Partitions

- By node vs by core
  - By node jobs can perform better
  - By core jobs have more opportunity to run
- --time=3-00:00 --nodes=1
- --ntasks-per-node=32
- --time=3-00:00 --ntasks=32

# View a snapshot of partition state on a cluster

[remote $]

```
partition-stats
```

[output]

```
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   |   54:20  |  59:174 | 238:1855|  56:1209|  79:244 |  18:1943|
Large Mem |    0:0   |   0:1   |   1:0   |   0:12  |   0:21  |   0:22  |
GPU       |    0:0   |   0:0   |   2:644 |   3:215 |   1:1   |   0:9   |
----------|-------------------------------------------------------------
      Number of Running Jobs by partition Type (by node:by core)
----------|-------------------------------------------------------------
Regular   |    3:71  |  33:880 |  18:760 |  37:1215|  48:472 |  27:544 |
Large Mem |    0:2   |   0:54  |   5:6   |   1:36  |   1:9   |   0:7   |
GPU       |    0:9   |   1:29  |  19:66  |  17:82  |   0:39  |   0:17  |
----------|-------------------------------------------------------------
       Number of Idle nodes by partition Type (by node:by core)
----------|-------------------------------------------------------------
Regular   |    2:0   |   2:0   |   2:0   |   1:0   |   1:0   |   0:0   |
Large Mem |    0:0   |   0:0   |   0:0   |   0:0   |   0:0   |   0:0   |
GPU       |   13:4   |  11:2   |   6:2   |   2:0   |   0:0   |   0:0   |
----------|-------------------------------------------------------------
       Total Number of nodes by partition Type (by node:by core)
----------|-------------------------------------------------------------
Regular   | 1088:646 |1088:646 |1058:626 | 763:391 | 381:180 | 106:58  |
Large Mem |   27:19  |  27:19  |  24:16  |  20:4   |   5:4   |   3:2   |
GPU       |  200:122 | 194:116 | 182:110 | 139:81  |  42:39  |  29:27  |
----------|-------------------------------------------------------------
```

# View a snapshot of partition state with node information

[remote $]

```
clusterstats
```

[output]

```
[✔] Loading node information (success, loaded cached version that is 18 min old)
[✔] Loading job information (success, loaded cached version that is 18 min old)
[✔] Loading share information (success, loaded cached version that is 15 min old)
Information on? (Use arrow keys, press Enter to select)
▸ User
  Group
  Cluster
  Quit
```

[output]

```
[✔] Loading node information (success, loaded cached version that is 18 min old)
[✔] Loading job information (success, loaded cached version that is 18 min old)
[✔] Loading share information (success, loaded cached version that is 15 min old)
Information on? Cluster
Please select on which part of the cluster would you like more information? CPU, (base) less than 12 GB of RAM per Core
Information on ? Jobs/Partitions/Nodes that allow partial node jobs, ie request by core.
Please select the information you would like to display? Cores with memory

         This table shows all available resources in the partition.
         A resource that is available to run 0-24 hour jobs
         will show up in the (0-3),(3-12) and (12-24) columns.

┌──────────────────────────┬─────────────┬──────────┬──────────┬──────────┬──────────┬──────────┐
│ cpubase_bycore           │ interactive │ 0-3 hr   │ 3-12 hr  │ 12-24 hr │ 1-3 day  │ 3-7 day  │ 7-28 day │
├──────────────────────────┼─────────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┤
│ Total (Cores with memory)│ 204         │ 20644    │ 20644    │ 20004    │ 11952    │ 5532     │ 1916     │
│   cpu=32, Mem=128000     │ 160         │ 14528    │ 14528    │ 13888    │ 8256     │ 4416     │ 1600     │
│   cpu=32, Mem=256500     │ 0           │ 1760     │ 1760     │ 1760     │ 1760     │ 192      │ 96       │
│   cpu=44, Mem=191840     │ 44          │ 4356     │ 4356     │ 4356     │ 1936     │ 924      │ 220      │
│ Idle (Cores with memory) │ 157         │ 756      │ 756      │ 740      │ 411      │ 106      │ 3        │
│   cpu=32, Mem=128000     │ 147         │ 460      │ 460      │ 444      │ 227      │ 80       │ 3        │
│   cpu=32, Mem=256500     │ 0           │ 81       │ 81       │ 81       │ 81       │ 0        │ 0        │
│   cpu=44, Mem=191840     │ 10          │ 215      │ 215      │ 215      │ 103      │ 26       │ 0        │
└──────────────────────────┴─────────────┴──────────┴──────────┴──────────┴──────────┴──────────┴──────────┘

Please select on which part of the cluster would you like more information? (Use arrow keys, press Enter to select)
▸ CPU, (base) less than 12 GB of RAM per Core
  CPU, (highmem or large) more than 12 GB of RAM per Core
  GPU
  Back
  Quit
```

# View job submission parameters

[remote $]

```
sacct -aX -A def-jdesjard_cpu -S 2021-10-21 -o jobid,ncpus,nnodes,reqmem,timelimit,partition%36
```

[output]

```
        JobID      NCPUS   NNodes     ReqMem  Timelimit                            Partition
------------ ---------- -------- ---------- ---------- ------------------------------------
53538234_0            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538234_1            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538234_2            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538234_3            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538234_4            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538235_0            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538235_1            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538235_2            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538235_3            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538235_4            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538236_0            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538236_1            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538236_2            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538236_3            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538236_4            1        1      256Mc   00:02:00                  cpubase_bycore_b1
53538392_0            1        1      256Mc   00:02:00                        cpubackfill
53538392_1            1        1      256Mc   00:02:00                        cpubackfill
53538392_2            1        1      256Mc   00:02:00                        cpubackfill
```

# View job usage metrics: case study

```
[jdesjard@cedar1 ~]$ sacct -aX -A <account_name> -S 2021-11-01 -o jobid%18,submit,start,timelimit%20,elapsed,ncpus,nnodes,reqmem,partition%18,state
             JobID             Submit              Start            Timelimit    Elapsed      NCPUS   NNodes     ReqMem          Partition      State
------------------ ------------------- ------------------- -------------------- ---------- ---------- -------- ---------- ------------------ ----------
          19388564 2021-11-16T21:06:16 2021-11-16T21:36:06          4-00:00:00   07:37:36         32        1     7000Mc  cpubase_bynode_b5  COMPLETED
          19388598 2021-11-16T21:07:07 2021-11-17T05:13:47          4-00:00:00   07:35:25         32        1     7000Mc  cpubase_bynode_b5  COMPLETED
...
          20018464 2021-11-23T23:09:10 2021-11-24T00:22:49          4-00:00:00   04:50:37         30        1     5000Mc  cpubase_bycore_b5  COMPLETED
          20018478 2021-11-23T23:09:44 2021-11-24T00:22:49          4-00:00:00   05:48:49         30        1     5000Mc  cpubase_bycore_b5  COMPLETED
...
          20038407 2021-11-24T07:27:59 2021-11-24T07:45:55          4-00:00:00   06:06:33         30        1     4000Mc  cpubase_bycore_b5  COMPLETED
          20038444 2021-11-24T07:29:18 2021-11-24T07:45:56          4-00:00:00   00:00:08         30        1     4000Mc  cpubase_bycore_b5     FAILED
...
          20060696 2021-11-24T11:53:27 2021-12-03T16:26:05          5-00:00:00   06:50:25         30        1     7000Mc  cpubase_bycore_b5  COMPLETED
          20060713 2021-11-24T11:54:19 2021-12-03T23:26:19          5-00:00:00   06:40:40         30        1     7000Mc  cpubase_bycore_b5  COMPLETED
          20060730 2021-11-24T11:55:27 2021-12-04T02:18:20          5-00:00:00   06:39:54         30        1     7000Mc  cpubase_bycore_b5  COMPLETED
          20060760 2021-11-24T11:56:10 2021-12-04T06:11:14          5-00:00:00   06:51:48         30        1     7000Mc  cpubase_bycore_b5  COMPLETED
...
          20061633 2021-11-24T12:08:19 2021-12-02T00:48:03          4-00:00:00   05:57:14         30        1     4000Mc  cpubase_bycore_b5  COMPLETED
          20061682 2021-11-24T12:08:47 2021-12-02T02:19:13          4-00:00:00   00:00:03         30        1     4000Mc  cpubase_bycore_b5     FAILED
```

```
[jdesjard@cedar1 ~]$ sacct -a -A <account_name> -S 2021-11-01 -o jobid%18,submit,start,timelimit%20,elapsed,ncpus,avecpu,nnodes,reqmem,maxrss,partition%18,state
             JobID             Submit              Start            Timelimit    Elapsed      NCPUS     AveCPU   NNodes     ReqMem     MaxRSS          Partition      State
------------------ ------------------- ------------------- -------------------- ---------- ---------- ---------- -------- ---------- ---------- ------------------ ----------
...
          20060730 2021-11-24T11:55:27 2021-12-04T02:18:20          5-00:00:00   06:39:54         30                    1     7000Mc             cpubase_bycore_b5  COMPLETED
    20060730.batch 2021-12-04T02:18:20 2021-12-04T02:18:20                       06:39:54         30 6-12:28:25        1     7000Mc 143561672K                    COMPLETED
   20060730.extern 2021-12-04T02:18:20 2021-12-04T02:18:20                       06:39:56         30   00:00:00        1     7000Mc          0                    COMPLETED
          20060760 2021-11-24T11:56:10 2021-12-04T06:11:14          5-00:00:00   06:51:48         30                    1     7000Mc             cpubase_bycore_b5  COMPLETED
    20060760.batch 2021-12-04T06:11:14 2021-12-04T06:11:14                       06:51:48         30 6-21:10:49        1     7000Mc 148288388K                    COMPLETED
   20060760.extern 2021-12-04T06:11:14 2021-12-04T06:11:14                       06:51:50         30   00:00:00        1     7000Mc          0                    COMPLETED
```

# In summary:

- In a saturated system wait times are required to distribute usage according to a fair share principle.
- Some wait times however are extended by the scarcity of the requested resource combination
- Understanding the node types and quantities can help in designing job submissions for maximal resource access
- Understanding the distribution on nodes across partitions can further help increase access to nodes on the general purpose systems