

HIGH PERFORMANCE COMPUTING (HPC)

AGENDA

The Problem

Introducing HPC

Workflow Demo

Potential Application



THE PROBLEM

SUPPOSE YOU WANT TO TRAIN YOUR AI MODEL

A typical AI job will take a day or two, or even in weeks(1)

- You might want to test with different optimizer
- You might want to test with different learning rate
- You might want to test with different training dataset
- More often, you need to re-train the model if the model failed to converge


CPU TYPE OF WORKFLOW

Suppose

- You have some jobs which take 1 minute to run, but you want to test over 500,000 combinations
- If you run on single CPU, it will take 500,000 minutes to run → which is close to 1 year
- Even if you re-write the code using 64 CPU core:

it will take 500,000 minutes / 64 to run → which still take 5 days (fully optimized)

It's not trivial to re-write the legacy code in the multi-processing
64 CPU core is expensive
Lack of job scheduler/control

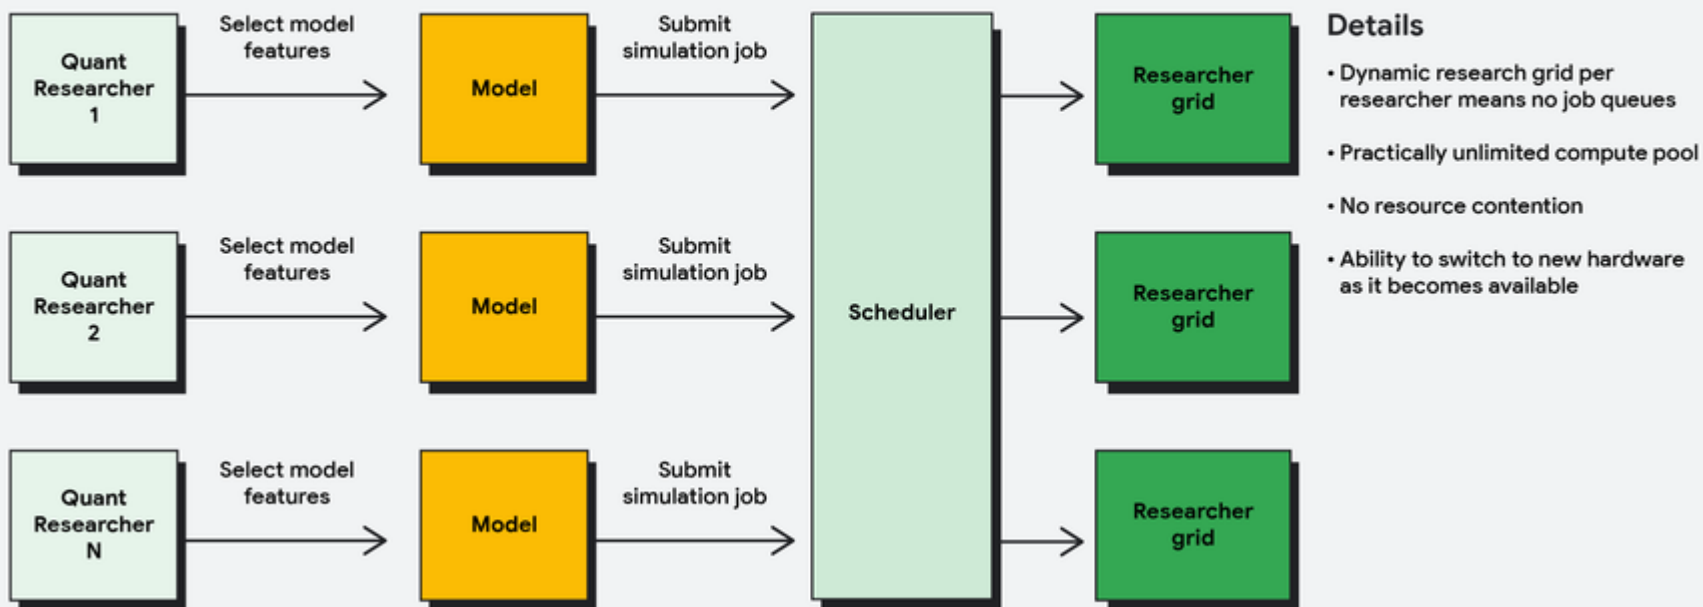


**INTRODUCING
HPC**

How Citadel Securities is reimagining quantitative research on the cloud

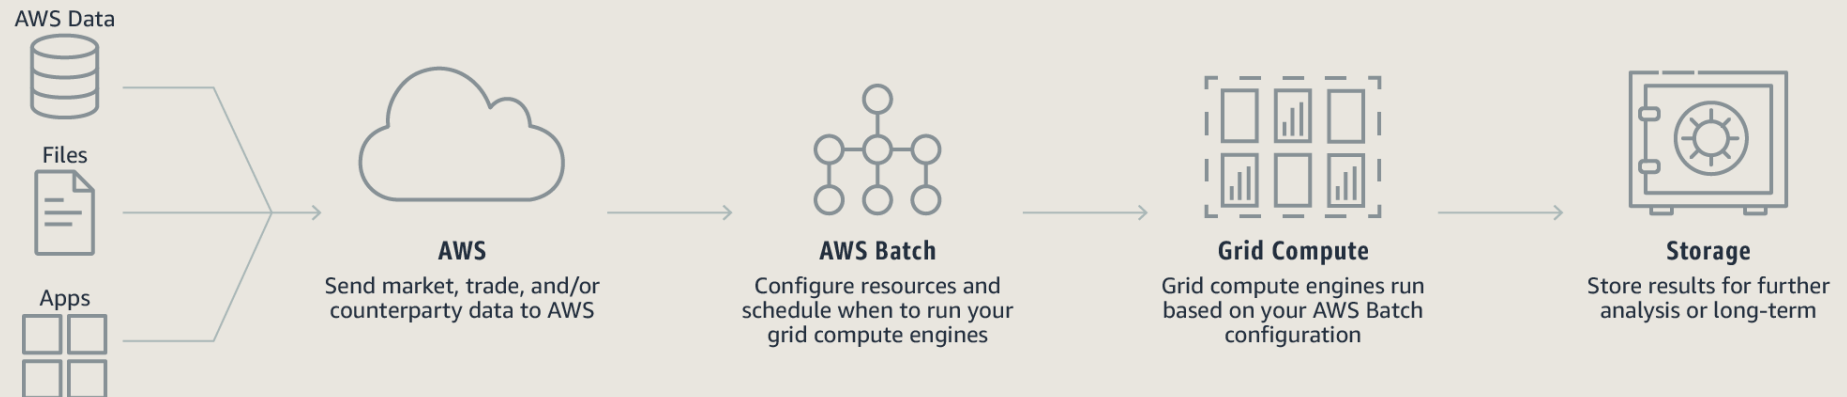
April 10, 2024

Dynamic compute allocation (Google Cloud)

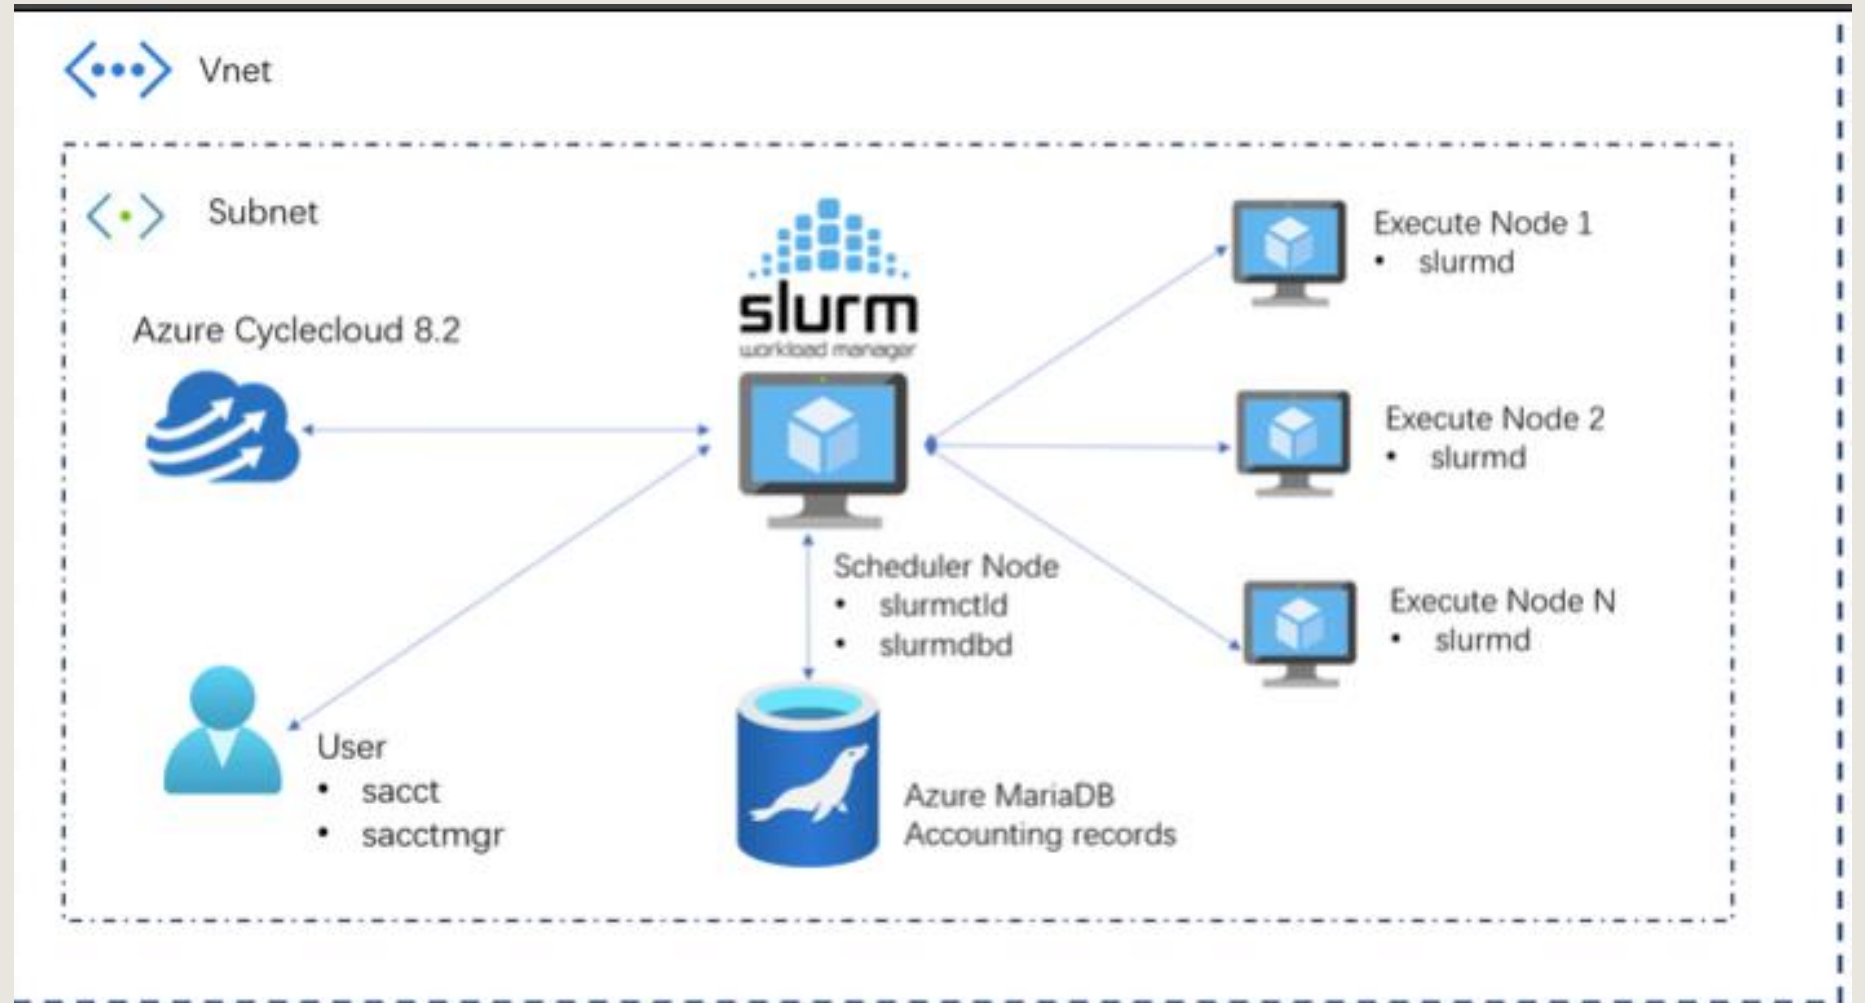


INTRODUCING HPC AND PARALLELCLUSTER(AWS)/CYCLECLOUD(AZURE)

- CloudFormation + Cluster Management == ParallelCluster
- Simple to install, easy to manage
- Everything you need to get a cluster up and running in a minute
 - Master node with scheduler
 - Compute nodes that grow and shrink on demand
 - Shared NFS storage
 - /shared
 - /home



INTRODUCING HPC AND PARALLELCLUSTER(AWS)/CYCLECLOUD(AZURE)



AWS vs. Azure vs. Google Instance Types

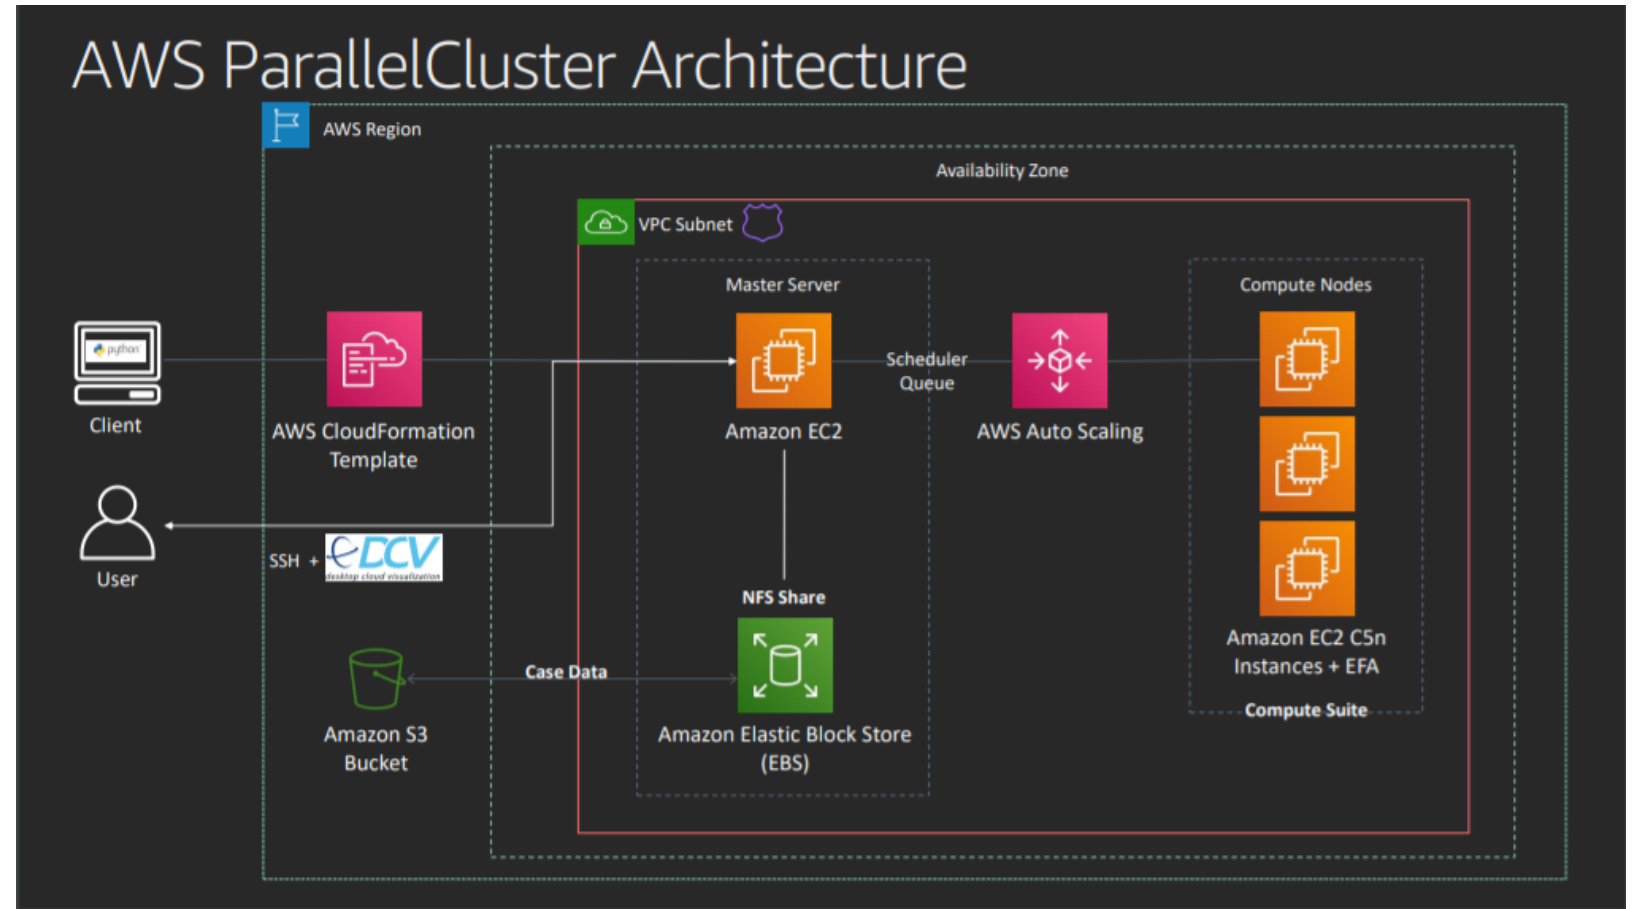
| Resource Type (us-east, Linux) | AWS Instance | AWS Memory | AWS Storage | Azure Instance | Azure Memory | Azure Storage | Google Instance | Google Memory | Google Storage |
|--------------------------------|-----------------|------------|-------------|----------------|--------------|---------------|----------------------|---------------|----------------|
| Standard 2 vCPU w SSD | m3.large | 8 | 32 | D2 v2 | 7 | 100 | n1-standard-2 | 7.5 | 375 |
| Highmem 2 vCPU w SSD | r3.large | 15 | 32 | D11 v2 | 14 | 100 | n1-highmem-2 | 13 | 375 |
| Highcpu 2 vCPU w SSD | c3.large | 3.75 | 32 | F2 | 4 | 32 | n1-highcpu-2 | 1.8 | 375 |
| Standard 2 vCPU no SSD | m4.large | 8 | 0 | D2 v2 | 7 | 100 | n1-standard-2 | 7.5 | 0 |
| Highmem 2 vCPU no SSD | r4.large | 15.25 | 0 | D11 v2 | 14 | 100 | n1-highmem-2 | 13 | 0 |
| Highcpu 2 vCPU no SSD | c4.large | 3.75 | 0 | F2 | 4 | 32 | n1-highcpu-2 | 1.8 | 0 |

As of Dec 2, 2016

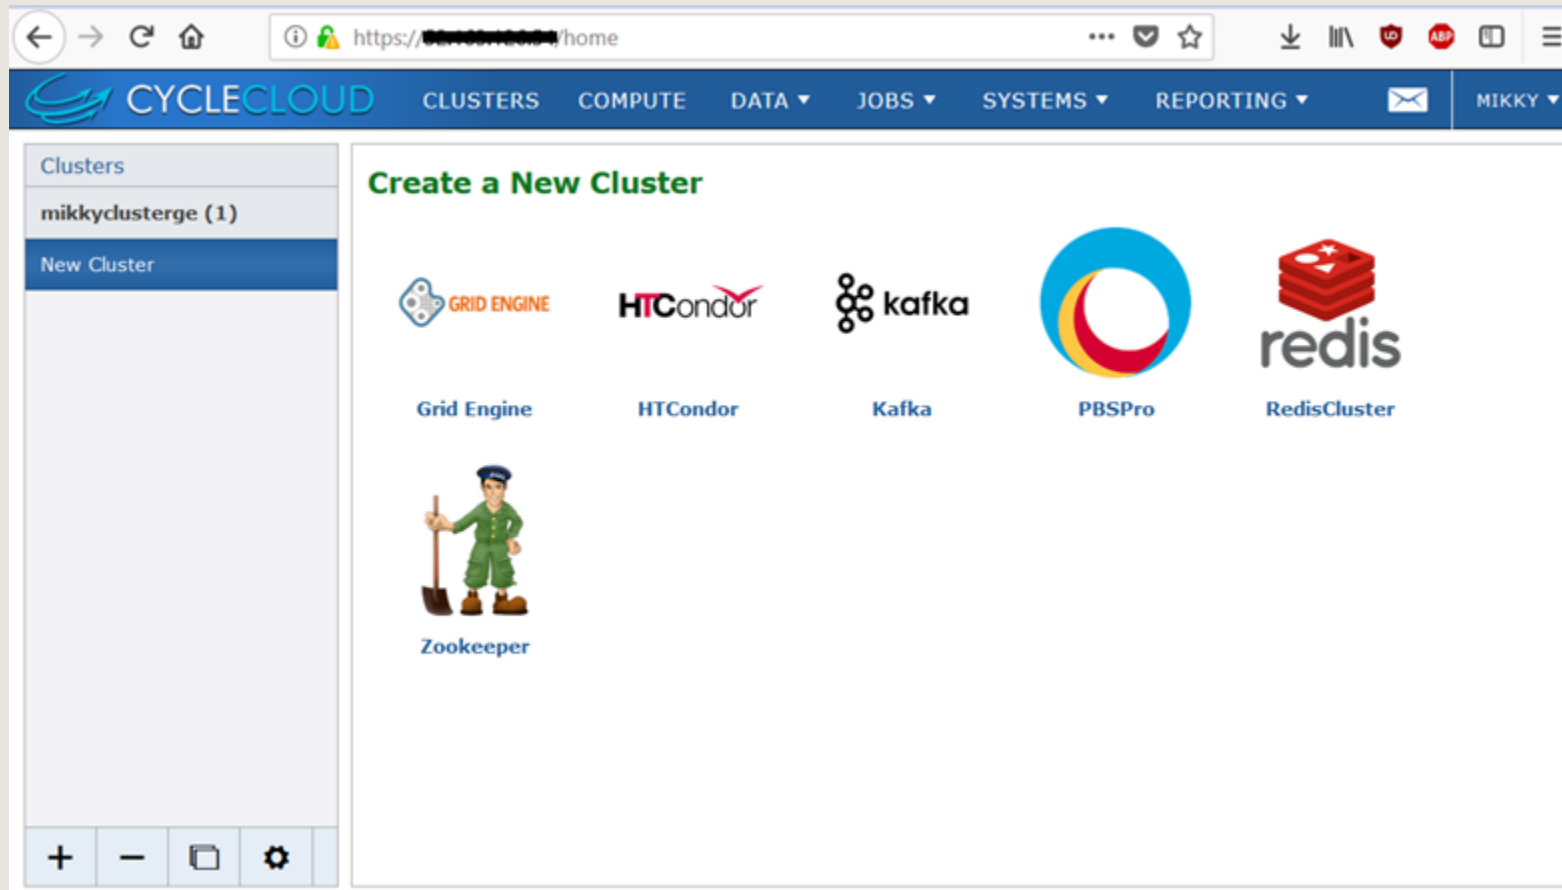
Source: RightScale

INSTANCE TYPE CHOICE

ParallelCluster/CycleCloud Architecture



HOW TO SETUP CYCLECLOUD



Clusters

mikkyclusterge (1)

mikkyclusterge

Start State **Started** at (was up 11m 24s)
 Edit Nodes **1 ready**
 Share Size **1 instance, 4 cores**
 Refresh Usage **0.7 core-hours (\$0)** in the last 24 hours
 Alerts **\$0** of **\$100** monthly budget | Manage
 Volumes **0** volumes, **0 B**

Nodes Arrays

View: Template ▾

Actions ▾

Search

| Template ▲ | Nodes | Cores | Status | Last Message |
|------------|-------|-------|--------|--------------|
|------------|-------|-------|--------|--------------|

| | | | | |
|--------|---|---|--|------|
| master | 1 | 4 | | |
|--------|---|---|--|------|



WORKFLOW DEMO

Workflow demo

```
python analysis.py --source "xxx" --learning_rate "0.02" --model "yyy" --config "myconfig.json" --save_result "True"
```

To run the same job in HPC

```
qsub python analysis.py --source "xxx" --learning_rate "0.02" --model "yyy" --config "myconfig.json" --save_result "True"
```

Every 2.0s: qstat

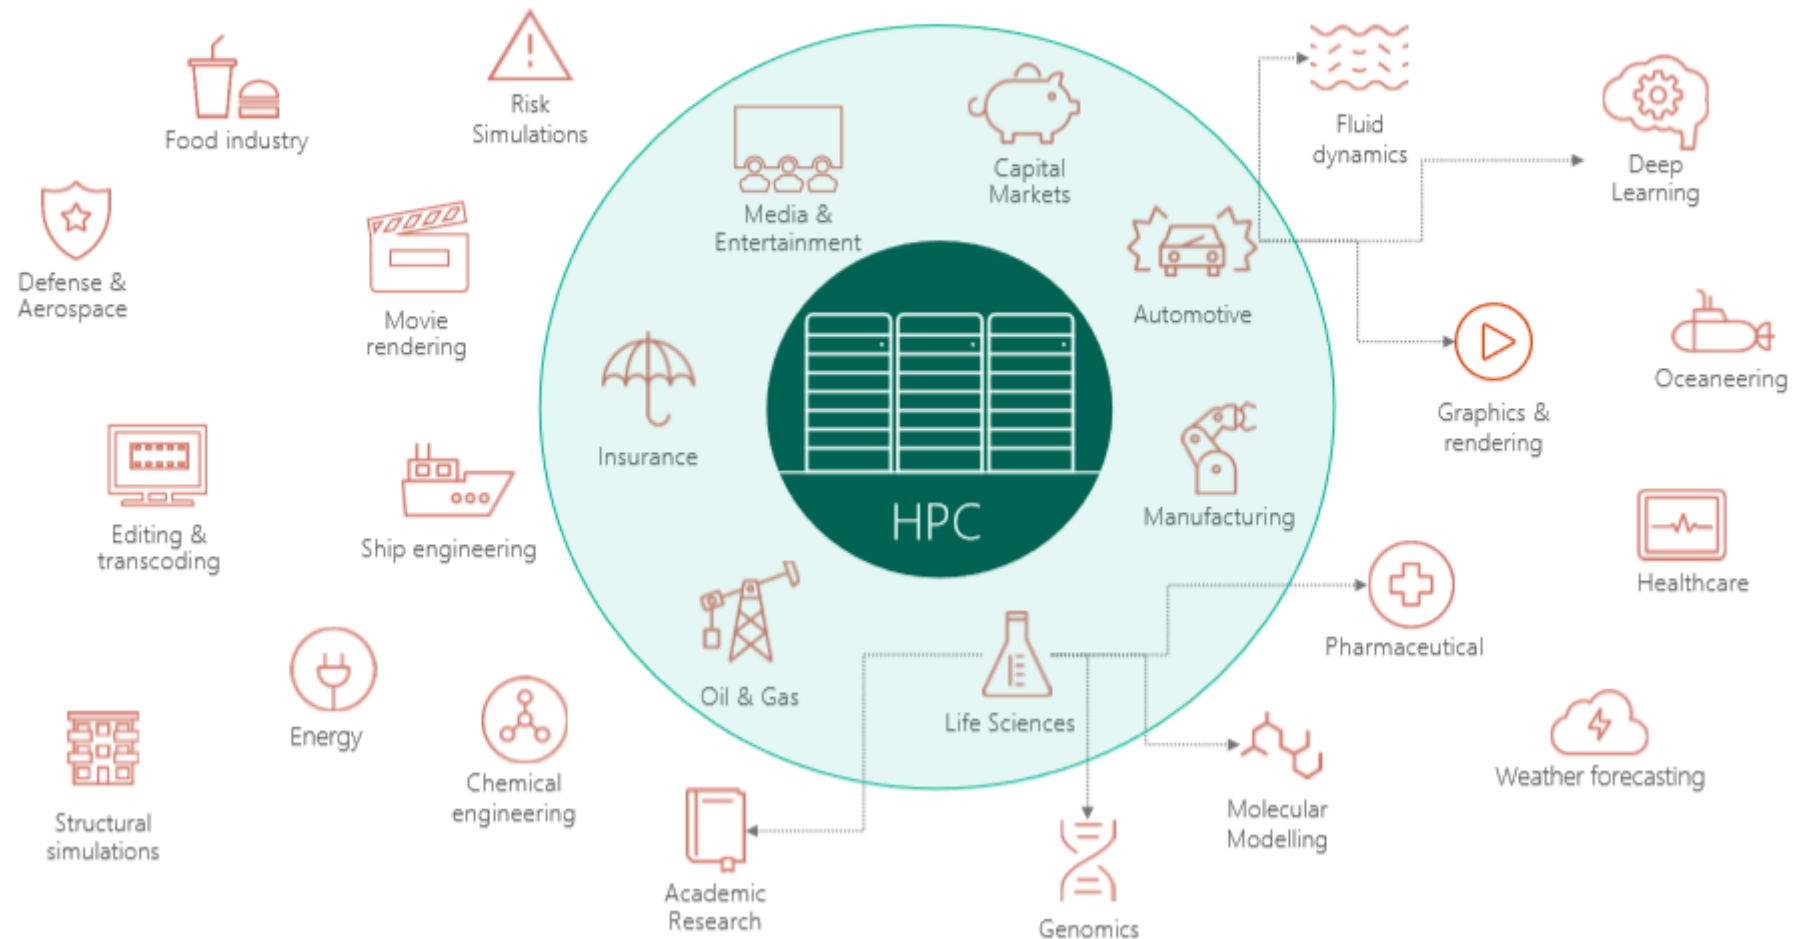
| job-ID | prior | name | user | state | submit/start at | queue | slots | ja-task-ID |
|---------|---------|-----------|------|-------|---------------------|-------------------|-------|------------|
| 6060662 | 0.56000 | single.sh | blau | r | 05/03/2024 00:51:49 | all.q@ip-0A03010A | 1 | |
| 6060663 | 0.55500 | single.sh | blau | r | 05/03/2024 00:51:49 | all.q@ip-0A03010A | 1 | |
| 6060664 | 0.55333 | single.sh | blau | r | 05/03/2024 00:51:49 | all.q@ip-0A03010A | 1 | |
| 6060665 | 0.55250 | single.sh | blau | r | 05/03/2024 00:51:49 | all.q@ip-0A03010A | 1 | |
| 6060666 | 0.55200 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060667 | 0.55167 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060668 | 0.55143 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060669 | 0.55125 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060670 | 0.55111 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060671 | 0.55100 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060672 | 0.55091 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060673 | 0.55083 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060674 | 0.55077 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060675 | 0.55071 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060676 | 0.55067 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060677 | 0.55063 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060678 | 0.55059 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060679 | 0.55056 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060680 | 0.55053 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060681 | 0.55050 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060682 | 0.55048 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060683 | 0.55045 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060684 | 0.55043 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060685 | 0.55042 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060686 | 0.55040 | single.sh | blau | qw | 05/03/2024 00:51:37 | | 1 | |
| 6060687 | 0.55038 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060688 | 0.55037 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060689 | 0.55036 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060690 | 0.55034 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060691 | 0.55033 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060692 | 0.55032 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060693 | 0.55031 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060694 | 0.55030 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060695 | 0.55029 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060696 | 0.55029 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060697 | 0.55028 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060698 | 0.55027 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060699 | 0.55026 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060700 | 0.55026 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060701 | 0.55025 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060702 | 0.55024 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060703 | 0.55024 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060704 | 0.55023 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060705 | 0.55023 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060706 | 0.55022 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060707 | 0.55022 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060708 | 0.55021 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060709 | 0.55021 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |
| 6060710 | 0.55020 | single.sh | blau | qw | 05/03/2024 00:51:38 | | 1 | |

Advantages

- You can choose the smallest available machine to run
- All the computed nodes shared the same EBS drive, so all the results can be saved in same place (I normally choose the cheapest magnetic drive)
- All the log file is saved, you can use simple unix command to check the results (imagine you use GUI and need to click each log file manually)
- All the compute nodes will have same config as the master node, no need to worry all about the virtual environment/docker/setup etc

POTENTIAL APPLICATION

Where is Big Compute used?





THANK YOU