

FLEXIBLE SLURM CONFIGURATION FOR LARGE-SCALE HPC

Steven Robson

HPC Systems Team



THE UNIVERSITY
of EDINBURGH



HPE Clusters at EPCC



ARCHER2: HPE Cray EX

- 5,860 compute nodes
- Dual AMD 64c “Rome” CPUs per node
- Slingshot 10 interconnect
- ARCHER2’s predecessor systems operated with PBS Pro



Cirrus: HPE SGI 8600

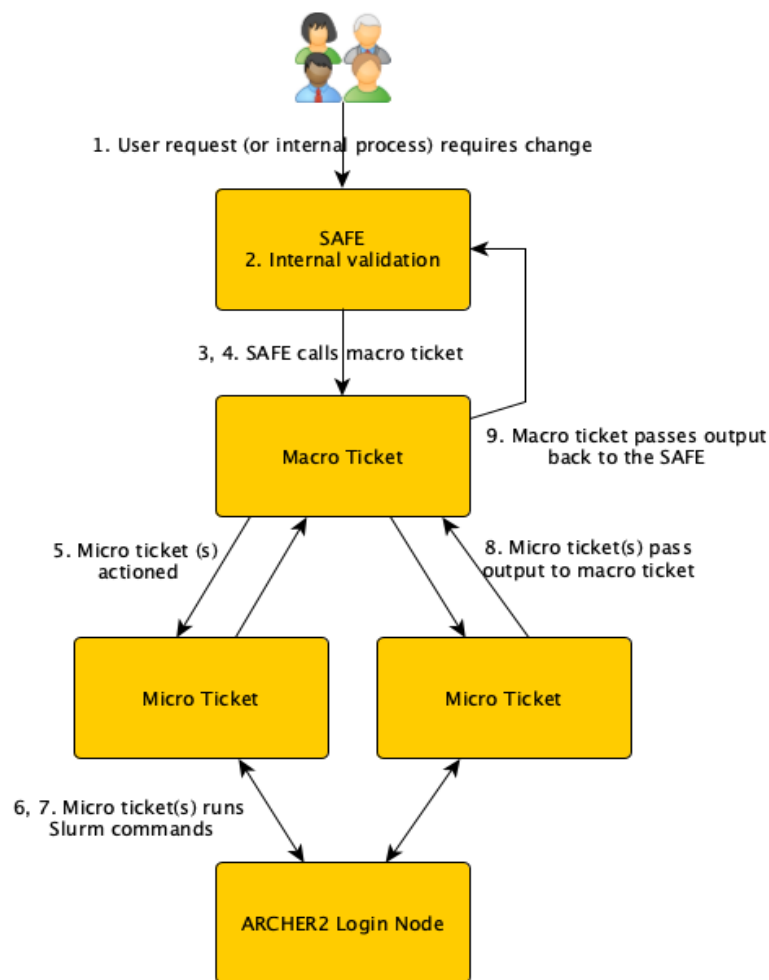
- 368 CPU nodes w/ dual Xeon Broadwell 18-core CPUs
- 38 GPU nodes w/ 4xNvidia V100
- Cirrus operated with PBS Pro previous to it’s 2020 rebuild

Background: The SAFE

- Service management web application
- Developed and operated by EPCC.
- User/project creation and management is handled by SAFE.
- Job accounting is also stored to the SAFE.
- Rundeck behind the scenes

Login account details				
Username	kleach			
Machines	Name	Model	Login	
	archer2-tds		login-tds.archer2.ac.uk	
	archer2	HPE Cray EX Supercomputer	login.archer2.ac.uk	
Status	Active			
Creation Date	03-Jul-2020			
Machine T&C	Fri Jul 08 10:11:56 BST 2022			
Last known login	2023-03-27 11:25			
Last run	2020-11-10 13:27			
Projects	z02 - OSG			
User disk info	Volume	Usage	Files	
	home (a2fs-home1)	33 GiB		
	work (a2fs-work1)	17 GiB	20 Files	
z02 resources	Resource Pool			Remaining Budget
	Archer2			99.2 CUs
	Volume	Usage	Quota	Files
	home (a2fs-home1)	46 GiB	100 GiB	
	home (a2fs-home3)	1 GiB		
	general (rdfaas_general)	668 GiB	2,002 GiB	
	epsrc (rdfaas_epsrc)	668 GiB	2,002 GiB	
work (a2fs-work1)	24 GiB		167,457 Files	
Packages	<ul style="list-style-type: none"> • archer2-assist-access 			

Background: Integrating Slurm and the SAFE



- Communication between the **SAFE** and **Rundeck** is handled by **webhook notifications** over HTTP on a private network
- Actions conducted include:
 - Creation and removal of users
 - Creation and removal of budgets/accounts
 - Locking and unlocking of budgets/accounts
 - Creation of reservations

Automating Slurm Budget Management

- SAFE budget = Slurm Account
 - Rundeck pushes Slurm accounting data to SAFE every night
 - SAFE charges jobs that did not end with `NODE_FAIL`
 - SAFE reconciles compute time vs budget available
- Via Rundeck:
 - SAFE locks an account
 - `sacctmgr -i modify account where cluster=archer2 Account=budget01 set maxtresmins=cpu=0`
 - SAFE unlocks an account
 - `sacctmgr -i modify account where cluster=archer2 Account=budget01 set maxtresmins=cpu=-1`



CPU Frequency Control on ARCHER2

Improve energy efficiency

- 600kW reduction (~20%)

No change to slurm.conf

Environment variable


```
export
```

```
SLURM_CPU_FREQ_REQ=2000000
```

Per application selection using higher frequency:

- GROMACS
- LAMPS
- NAMD

ARCHER2 Self Reported Power



1st November 2022 – 1st March 2023

GPU Frequency Control

- Requirement:
 - Set a default GPU frequency
 - Allow users to override this
- Slurm not compiled with Nvidia support
- Code in Slurm prolog to parse option and configure GPU:

```
/usr/bin/nvidia-smi -ac $gpuFreqMemory,$gpuFreq
```
- Epilog resets frequency

Streamlining Configuration (Cray EX)

- HPE Cray EX uses Jinja2 template engine as part of Slurm configuration
 - Sets slurm controller addresses
 - Builds node list from Hardware State Manager

- Add additional config depending on node memory

```
{{ ',HighMem Weight=1000' if node.RealMemory > 262144 }}
```

```
{{ ',StandardMem Weight=500' if node.RealMemory < 262145 }}
```



Streamlining Configuration (Cray EX)

- Single template for slurm.conf applies to:
 - 5860 node ARCHER2 main system
 - 8 node TDS (Test and Development System)
- Different power monitoring configuration:

```
{% if cluster_name == 'archer2' %}  
AcctGatherEnergyType=acct_gather_energy/pm_counters  
{% elif cluster_name == 'tds' %}  
AcctGatherEnergyType=acct_gather_energy/ipmi  
{% endif %}
```



Job Submission Scripting

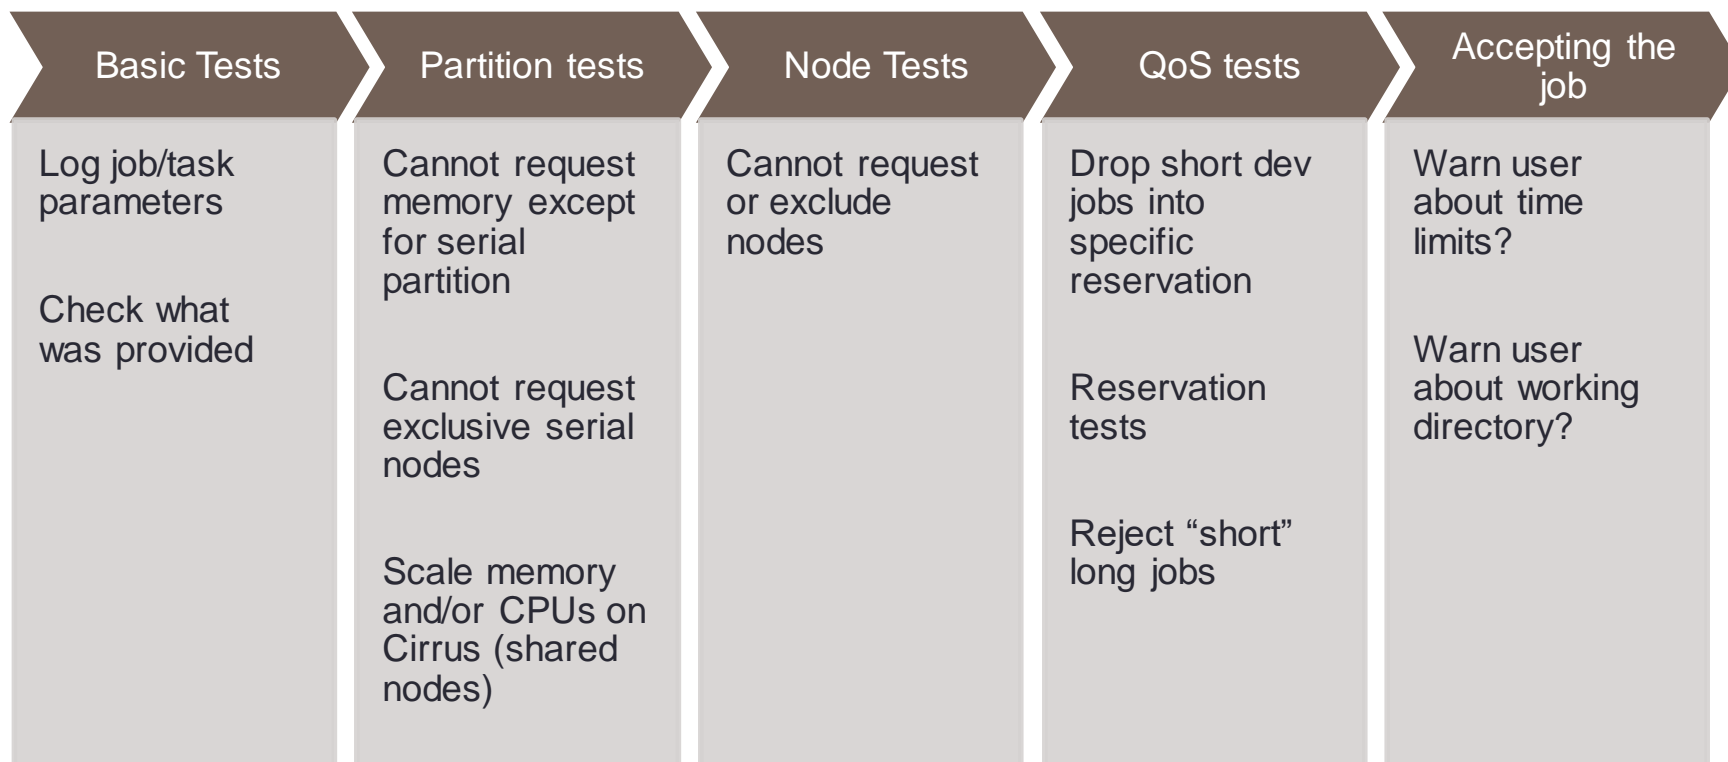
- Slurm plugin
- Lua script on slurm control host

- Avoids ineffective jobs running
 - Warn users so they can resubmit
- Avoids ineffective scheduling
 - Users cannot specify/exclude specific nodes

- Enforce system rules that can't be expressed in other ways

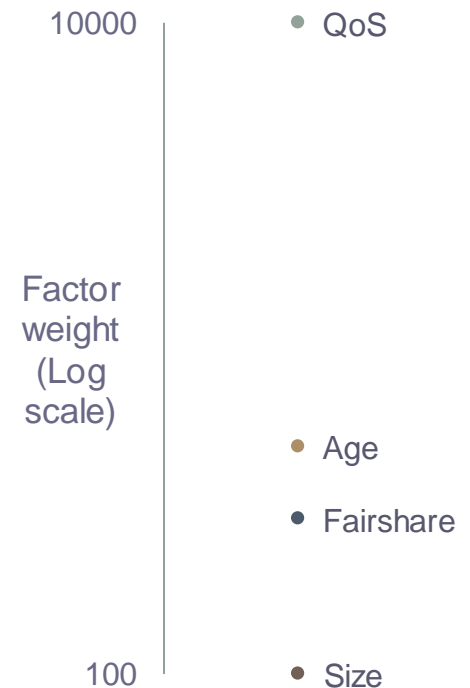


Job Submission Scripting



Priority Management

- **Goal: Reasonably fair scheduling for all users across a diverse user community**
- Slurm's multifactor plugin
- Define factors with weights in config file
- High priority QoS jobs always front of queue
- Weight older jobs (up to 14 days)
- Weight larger jobs (ARCHER2)
- Fairshare
 - Short half-life (48 hours)
 - Longer half lives conflict with time-bound allocations of compute hours



Flexible Slurm Configuration for Large Scale HPC

- Effective scheduling across varying system types/user bases to improve user experience and reduce support staff load:
 - User and budget management
 - Frequency control
 - Streamlining configuration
 - Job submission scripting
 - Priority Management



Thank you. Questions?

s.robson@epcc.ed.ac.uk



Config Repo

Partners



Engineering and
Physical Sciences
Research Council

Natural
Environment
Research Council



THE UNIVERSITY
of EDINBURGH



**Hewlett Packard
Enterprise**



THE UNIVERSITY *of* EDINBURGH

