



Science and
Technology
Facilities Council

STFC Cloud 2023 Requirements

Presented by John Good, STFC Cloud Delivery Manager

Background

The STFC Cloud is an Infrastructure-as-a-service platform delivering compute and storage resources to facilities both within STFC and externally, such as the Astro and Particle Physics communities including SKA.

£4.5M of investment this year, up from £3.5 in previous two years. The STFC Cloud has approx. 2000 active users across over 500 project workspaces.

The platform is based on OpenStack and currently running on Rocky 8.

This presentation will focus on the GPU requirements but of course, the STFC Cloud makes use of storage, networking and compute resources at STFC.



All requirements (including CPU only nodes)

- Disks must be able to be spanned and mirrored using hardware raid or other bios functions (RAID10 equivalent)
- Network interface adaptors must be able to act as routers without running anything within the OS
 - They must be able to run a network operating system purely on the NIC
 - They must support vxlan offload, ecmp and evpn
 - They must have out of band management access on a separate port on the same nic

Compute GPUs

- Primarily for computation and machine learning
- Looking for High Powered Compute GPUs
 - Nvidia A100 40GB used as minimum performance requirement
- Currently require CUDA and OpenCL support
- Must support PCI-passthrough with kvm\qemu on Rocky 8 and Rocky 9
- Looking for 4-8 GPUs per node optimised for cost
- 32 logical CPU cores per GPU
- 256GB ram per GPU
- 1 TB of SSD per GPU
- Minimum 25gb networking per GPU
- Option for alternate nodes with Intel or AMD GPUs for testing/development purpose
 - These must conform to the other requirements as close as possible

Visualization GPUs

- Primarily for accelerating visualisation in a VDI like use case
- Looking for Visualization/workstation GPUs
 - Nvidia A4000 16GB used as minimum performance requirement
- Currently require OpenGL, CUDA and OpenCL support
- Must support PCI-passthrough with kvm\qemu on Rocky 8 and Rocky 9
- Looking for 4-8 GPUs per node optimised for cost
- 32 logical CPU cores per GPU
- 128GB ram per GPU
- 1 TB of SSD per GPU
- Minimum 25gb networking per GPU
- Option for alternate nodes with Intel or AMD GPUs for testing/development purpose
 - These must conform to the other requirements as close as possible
- vGPU,MIG or other solutions will be considered as long as the requirements can be met without requiring custom kvm/qemu.
 - Any licencing for this should be included for the lifetime of the nodes



Science and
Technology
Facilities Council

Thank you!

Any technical questions contact alexander.Dibbo@stfc.ac.uk

Facebook: Science and
Technology Facilities Council

Twitter: @STFC_matters

YouTube: Science and
Technology Facilities Council