

HPC Cloud workshop 2024

Welcome

HPC-Cloud Overview

- call for windows, defender deleting executables. Code signing certificate.
- nvidia license agreement per VM launch, so each project agrees. If we have a image with pre installed.
- management of old images, keep them to best extent.
- interacting with openstack from Gitlab CI
- provide pre built images with module like software (nvidia drivers)
- not much call for hyperthreading, meaning many core VMs.
- generalz GPU node need only GPU and often not lots of RAM. But there is a probable call for GPU with high mem. The HPC systems, raven etc, are obvious targets for such jobs.
- Multi attache volumes. would be good to document whet they really are.
- double check the default SG rules to ensure cross project traffic is not allowed by default.
- port security and forwarding issue.
- ssh access from external via FIP. Currently blocked due to security policy. can we consider relaxing.
- Auto DNS for private networks. Re-look at designate for DNS as a service.
- How large can a Manila NFS/CEPHFS be. Often seen as a disadvantage to use nexus-posix for large shared storage. would like to be in full control.
- possible table to show pros/cons for nexus-posix vs Manila.
- Posit as a service? Many users ask for Posit rather than jupyterhub.

Openstack Infrastructure

- clarification about what is a core institute service (email etc).
- Could possibly be run at GWDG?
- how to manage autoscaling of clusters and nodes.

Globus

- clarification: Globus Compute is different from moving data to Nexus-POSIX/Raven. It is, e.g., running slurm jobs after transfer finishes
- Only one side of a transfer has to be licenced

- Globus & Archive: we are careful about this level of automation on the archive

Project talks

Full Stack Bioinformatics

- their own Terraform based k8s recipe
- Terraform SLURM recipe can scale cluster - John wants to have a look
- ftp vs s3 tempurls

Operating NOMAD: A FAIR data service on the MPCDF HPC Cloud with Kubernetes

- NOMAD Oasis is isolated resources, can hide data here
-

Workstations, SLURM cluster, Analysis server, User Storage, Network Integration - Use cases and challenges of the FHI

- aim to move into cloud-local-fhi
- considering how to use manilla rather than nexus for home and software. generally inode number is an issue.
- Dynamic scaling of slurm cluster
- workstations going very well. stable wit no real issues. VNC working well enough that no complaints have been had.
- minio server for smaller user data. restic backup of laptops + any random data

Projects II

DCOR

Bioinformatics

- Coordinate AAI with GWDG? Discuss what makes sense

Storage

Proxy for S3 access that provides similar functionality to datashae :-)

Projects III

Sci Server

- k8s based data and analysis service. Jupyter notebooks etc.
- currently containers for users actually are launched on VMs (not in k8s). This may change. Considering what is the best solution VMs vs Containers.

- mounting and sharing mounts of data between users (running code in containers) is probably going to be the core challenge.
-