

INTEGRATING HPC, CLOUD, AND CONTAINERS FOR DATA-INTENSIVE COMPUTING PLATFORMS

Daniel Murphy-Olson (d@anl.gov)¹, Ryan Aydelott¹

¹ Computing, Environment and Life Sciences Directorate, Argonne National Laboratory

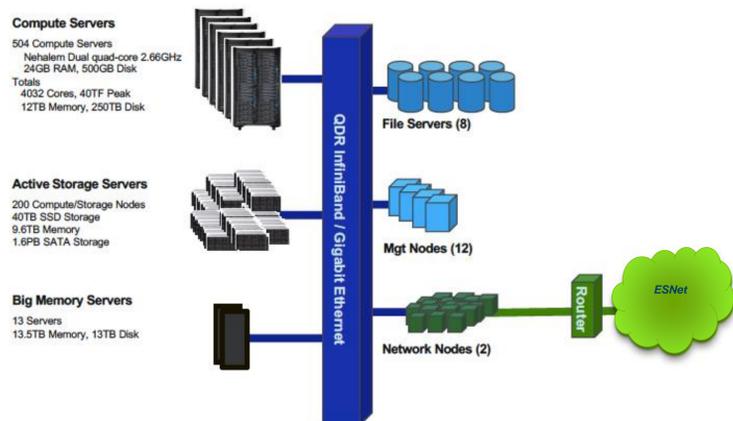
ABSTRACT

- Argonne provides a broad portfolio of computing resources to researchers. Since 2011 we have been providing a cloud computing resource to researchers, primarily using Openstack. Over the last year we've been working to better support containers in the context of HPC. Several of our operating environments now leverage a combination of the three technologies which provides infrastructure tailored to the needs of the specific workload. This poster summarizes some of our experiences integrating HPC, Cloud, and Container environments.

OVERVIEW

- Magellan is a 717 node OpenStack cluster with heterogeneous hardware resources. The system has two interconnects, gigabit ethernet and QDR infiniband. Gigabit ethernet is used for instance communication, and QDR infiniband is used for storage connectivity.

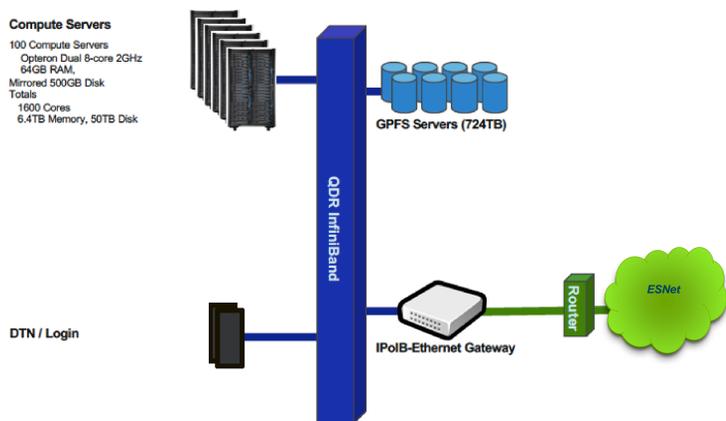
MAGELLAN HARDWARE ARCHITECTURE



Argonne

- Jupiter is a traditional HPC cluster with GPFS storage. The system has two interconnects, gigabit ethernet and QDR infiniband. QDR infiniband is used for node to node communication as well as storage connectivity.

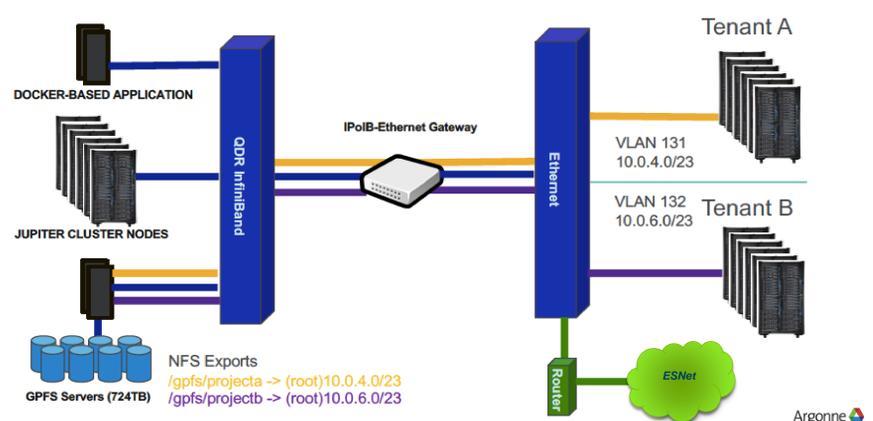
JUPITER HARDWARE ARCHITECTURE



Argonne

INTEGRATION OF MAGELLAN AND JUPITER

- Infiniband to Ethernet gateway (Mellanox SX6036G) is the primary link between the two systems
- NFS is exported from the GPFS server to tenant networks
- Docker-based application, HPC jobs, and cloud images all have access to the same GPFS data store
- Tenant users have full access to the files exported from the GPFS server
- VLANs enforce separation between tenants and various exports from the GPFS server
- Projects commonly use group access controls to files, so this level of access control is sufficient to maintain data security



Argonne

CONCLUSIONS

- It is possible to architect integrated environments that provide both cloud and HPC-like resources
- Simple IP-based access controls along with tenant networks can allow for VM access to HPC data stores
- With this model system users can utilize their own images and directly access HPC data stores without needing to transfer data between systems

NEXT STEPS

- Improve performance between OpenStack instances and GPFS NFS exports
- Link scheduling between application portals on Jupiter and Magellan OpenStack provisioning
- Evaluate provisioning elements of Jupiter with OpenStack
- Experiment with additional HPC data stores (Lustre, GlusterFS)