

Running Workflow Systems at OLCF with Slate

Jason Kincl

HPC Systems Engineer, NCCS

kincljc@ornl.gov

ORNL is managed by UT-Battelle LLC for the US Department of Energy



U.S. DEPARTMENT OF
ENERGY

Agenda

- HPC Workflow Requirements
- Kubernetes for HPC Workflows
- Kubernetes at OLCF
- Demo
- Recap

HPC Workflow Requirements

- Need ways for users to manage their own workflow systems/databases/web portals
 - Diverse ecosystem of tools is difficult for operations group to support
- Why not just use SSH keys for remote access?
 - Our moderate security controls require remote actions to be authenticated with RSA two-factor credentials
 - Instead allow running workflow systems inside of our security boundary
- Upon surveying existing workflow systems we came up with the following requirements:
 - Run a persistent service that can stay running across system reboots
 - Talk to batch submission system for current queue information and job submission
 - Interact with files on GPFS/Lustre/NFS

Kubernetes for HPC Workflows

- For Users

- Access to HPC resources from service containers
- Similar user experience when compared to batch scheduling paradigm
- Growing industry standard for designing portable services that can be run anywhere

- Operations

- Kubernetes allows users to **declare their intent** and operations to **implement and realize that intent**
 - We make use of this in order to enforce operational policy such as running containers as a particular UID/GID and enforcing authentication on exposed services
- Kubernetes provides a single scheduling interface for both users and operations to run their workloads on a resource

Kubernetes at OLCF

- Production cluster running Red Hat OpenShift
 - Use resource quota and limit systems in Kubernetes, projects get Kubernetes namespaces to run their services
- Containers run as project-level automation user
 - Batch job submission from inside container
 - HPC parallel filesystem can be mounted inside container
- Can expose HTTP services externally with NCCS 2-factor authentication accessible to project members
- Can expose TCP/UDP services on higher level ports for access to services from compute nodes inside of OLCF
- Container build service available as part of OpenShift platform

Demo

Lessons Learned

- Diverse ecosystem of workflow tools
- Workflows present a number of production problems that have to be solved in new and unique ways
- Slate allows users to run their workflow tool at OLCF with access to OLCF resources



Looking Forward

- Kubernetes is becoming the standard for container orchestration
- Ongoing Slate pilot has been well received by both staff and users
- Get in touch with us!

