Kubernetes at DESY IT

Container Orchestration

Michael Schuh, <u>Tim Wetzel</u>, Johannes Reppin IBERGRID 2022, October 12 2022





HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

Primary Field of Application in Cloud Computing

Container Orchestration with Kubernetes



Source: https://www.openstack.org/software/

DESY. | Kubernetes at DESY IT | Michael Schuh, Tim Wetzel, Johannes Reppin | October 12 2022

Kubernetes

- for containerized applications
- on clusters of virtual machines
- We're not using it on bare metal servers

Openstack

- Used as a virtualization platform
- We're not using it for bare metal servers
- We're not using it for containers

The Cloud Native Landscape

The Cloud Native Trail Map

Container orchestration

1. Containerize applications 2. CI/CD - robust automation 3. Orchestration 4. Monitoring & Analysis



Integration with research infrastructure

5. DNS, Load Balancing 6. Network Operations

- Software-defined networking
- Firewall

argo

Dotaru

Dynamic certificates

7. Scientific data storage

- dCache
- High performance storage
- DBaaS

8. Event streaming platforms

- Data acquisition streams
- Function as a service (FaaS)
- 9. Scale container registry
 - HPC & HTC
 - Docker / Singularity Registry

10.Software Repository

- CVMFS
- NIMS

Source: trailmap.cncf.io



Container as a Service Cloud Native CI/CD Container registry



Kubernetes as a Service Container Orchestration Kubernetes Package Manager

> Container as a Service Cloud Native CI/CD Container registry



Software as a Service Container-based environments App deployments as code

Kubernetes as a Service Container Orchestration Kubernetes Package Manager

> Container as a Service Cloud Native CI/CD Container registry



GitLab Integration

Flux CD and Helm Releases

器 General /	Flux Cluster \$	Stats ☆ ≪			🇤 🛱 🕸 🕐 Last 15 minutes 🗸 🔾 🗸 🖵						
operator_namespace All ~ namespace All ~											
Cluster Reconcilers			Failing Recon	Failing Reconcilers		Kubernetes Manifests Sources		Failing Sources			
•	19		2		2	24		0			
Reconciler ops avg. duration					Source ops avg. duration						
HelmRelease				242 ms	GitRepository				207 ms		
Kustomization				790 ms	HeimRepository				294 ms		
~ Status											
Cluster reconciliation readiness					Source acquisition readiness						
Kind 🖓	Name 🖓	Namespace 🖓		Status 🕈 🐬	Kind 🖓	Name 🖓	Namespace 🖓		Status 🖓		
HelmRelease	dask	dask		Ready	GitRepository	cloud-portal	flux-system		Ready		
HelmRelease	grafana	grafana		Ready	GitRepository	dask	flux-system		Ready		
HelmRelease	int	cloud-portal		Ready	GitRepository	desy-notes	flux-system		Ready		
HelmRelease	int-db	cloud-portal		Ready	GitRepository	flux-system	flux-system		Ready		
HelmRelease	jupyterhub	jupyterhub		Ready	GitRepository	grafana	flux-system		Ready		
HelmRelease	postgres	postgres-test		Ready	GitRepository	jupyterhub	flux-system		Ready		



Kubernetes cluster management

From Openstack Magnum to Rancher



Kubernetes cluster management

From Openstack Magnum to Rancher



Kubernetes cluster management

From Openstack Magnum to Rancher



Architecture

Infrastructure concepts



• A container image is a standalone, executable package of software that includes everything needed to run an application: code, application runtime, system tools, system libraries and settings.

• **Container images become containers at run time** and are processes (or group of processes) running in a cgroup, where the Kernel isolates resource usage (CPU, memory, disk I/O, network)

DESY. | Kubernetes at DESY IT | Michael Schuh, Tim Wetzel, Johannes Reppin | October 12 2022

Architecture

Clusters and Schedulers

- Openstack Hypervisors organized as Host Aggregates
 - Openstack Nova (compute scheduler) places VMs on Hypervisors (hosts)





Architecture

Clusters and Schedulers

- •
- VMs configured as Kubernetes nodes form a Kubernetes Cluster
 - Dynamically created by Rancher Kubernetes management platform
- Openstack Hypervisors organized as Host Aggregates
 - Openstack Nova (compute scheduler) places VMs on Hypervisors





Architecture

Clusters and Schedulers

- Kubernetes compute manager places containers on Kubernetes Nodes
 - Application Deployments form Container Clusters
- VMs configured as Kubernetes nodes form a Kubernetes Cluster
 - Dynamically created by Rancher Kubernetes Management platform
- Openstack Hypervisors organized as Host Aggregates
 - Openstack Nova (compute scheduler) places VMs on Hypervisors







Gitlab Runners for gitlab.desy.de

Johannes Reppin, Michael Schuh – DESY IT / HIFIS, PaNOSC, ExPaNDS



Helm Deployment with Flux CD

Johannes Reppin – DESY IT / HIFIS

- Merge group memberships from Helmholtz AAI in Keycloak
- Run Jupyter Servers, set UID/GIDs
- Integrate with dcache-demo.desy.de



HIFIS Cloud Portal - cloud.helmholtz.de

Thomas Beermann - DESY IT / HIFIS



DESY. | Kubernetes at DESY IT | Michael Schuh, Tim Wetzel, Johannes Reppin | October 12 2022

HIFIS Cloud Portal - cloud.helmholtz.de

Thomas Beermann - DESY IT / HIFIS

- Deploy and manage a staging and production instance
- Automatically deploy full testing instances for each Merge Request.
- Containers help to deploy in a controlled and isolated environment
- Helm packages all parts, makes deployment much easier, e.g. re-using charts for database
- Flux CD keeps all configuration in a central Gitlab repository, keeps track of complex configurations
- Gitlab CI runs full workflow
 - compiling / testing the code
 - packaging / hosting the containers
 - deploying everything to Kubernetes to get review apps for Merge Requests

Vision:

- Better monitoring support, e.g., automatic export of application logs to Elasticsearch
- Better integration with Gitlab, i.e., manage and monitor deployed applications directly from Gitlab, especially for MR deployments.

Collaborative document editing - notes.desy.de

Johannes Reppin – DESY IT / HIFIS

Helm Deployment with Flux CD

- HedgeDoc instance
- Collaborative note taking for Helmholtz
- Storage and backups on S3



Thank you

DESY. | Kubernetes at DESY IT | Michael Schuh, Tim Wetzel, Johannes Reppin | October 12 2022

Contact	Michael Schuh	Tim Wetzel	Johannes Reppin
DECV	IT-RIC	IT-RIC	IT-INFA
DEST.	michael.schuh@desy.de	tim.wetzel@desy.de	johannes.reppin@desy.de

Deutsches Elektronen Synchrotron www.desy.de