

Container Orchestration and GitLab CI/CD with DESY Openstack as Cloud Provider

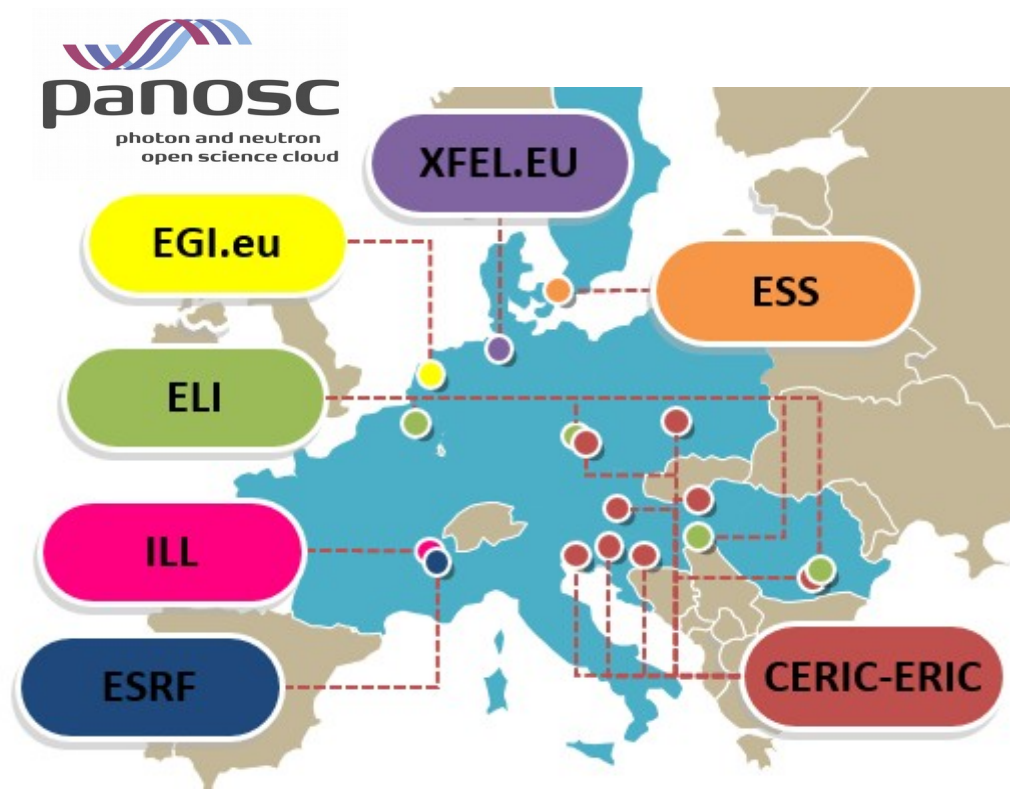
HGF AG Openstack

Michael Schuh, Johannes Reppin
Sep 23 2020



PaNOSC - Photon And Neutron Open Science Cloud

EGI Federated Cloud, European Open Science Cloud



www.eosc-portal.eu

www.panosc.eu

www.egi.eu/federation/egi-federated-cloud

PaNOSC has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 823852.

EGI Federated Cloud

- Virtual Appliance (VA) for Virtual Organizations (VO)

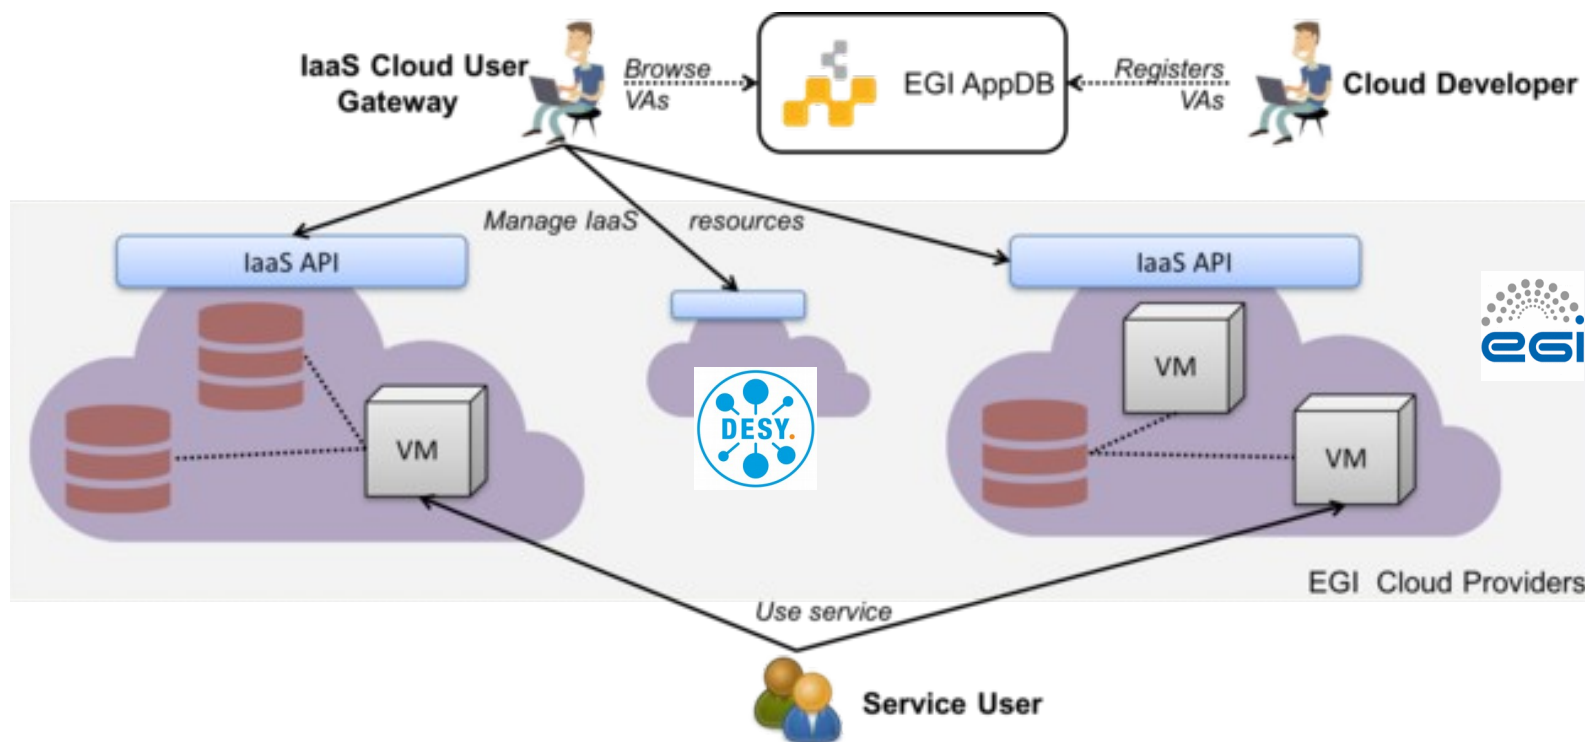


EOSC

- FAIR data, effective Open Science



DESY in the EGI Federated Cloud



wiki.egi.eu/wiki/Federated_Cloud_user_support
source: slideshare.net/EGI_Foundation/egi-federated-cloud-may-2019

DESY provides resources to the EGI Federated Cloud

- 16 servers
- 320 cores, 6 TB RAM
- 1.3 PB block storage
- 200 public IPs
- Self service access to ports 22 (SSH) and (80,443 http/s)

Synchronised services

- Accounting data
- Service discovery
- Virtual machine images
- Authentication
- Authorization
- DNS *.fedcloud.eu

What is Cloud Native?

The Cloud Native Compute Foundation

- A sub-foundation of **The Linux Foundation** since 2015
- Kubernetes donated as funding project by Google

<https://github.com/cncf/toc/blob/master/DEFINITION.md>

Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds. Containers, service meshes, microservices, immutable infrastructure, and declarative APIs exemplify this approach.

These techniques enable loosely coupled systems that are resilient, manageable, and observable. Combined with robust automation, they allow engineers to make high-impact changes frequently and predictably with minimal toil.

The Cloud Native Computing Foundation seeks to drive adoption of this paradigm by fostering and sustaining an ecosystem of open source, vendor-neutral projects. We democratize state-of-the-art patterns to make these innovations accessible for everyone.

The Cloud Native Landscape



[The Cloud Native Trail Map](#)

is CNCF's recommended path through

[The cloud native landscape](#)

Software stack container orchestration

Infrastructure as a Service
Cloud Computing
Infrastructure as code



Software stack container orchestration

Container as a Service

Cloud Native CI/CD

Docker registry



GitLab



Infrastructure as a Service

Cloud Computing

Infrastructure as code



openstack®

Software stack container orchestration

Kubernetes as a Service

Container Orchestration
Kubernetes Package Manager



Container as a Service

Cloud Native CI/CD
Docker registry



GitLab



Infrastructure as a Service

Cloud Computing
Infrastructure as code



openstack®

Software stack container orchestration

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
Docker registry

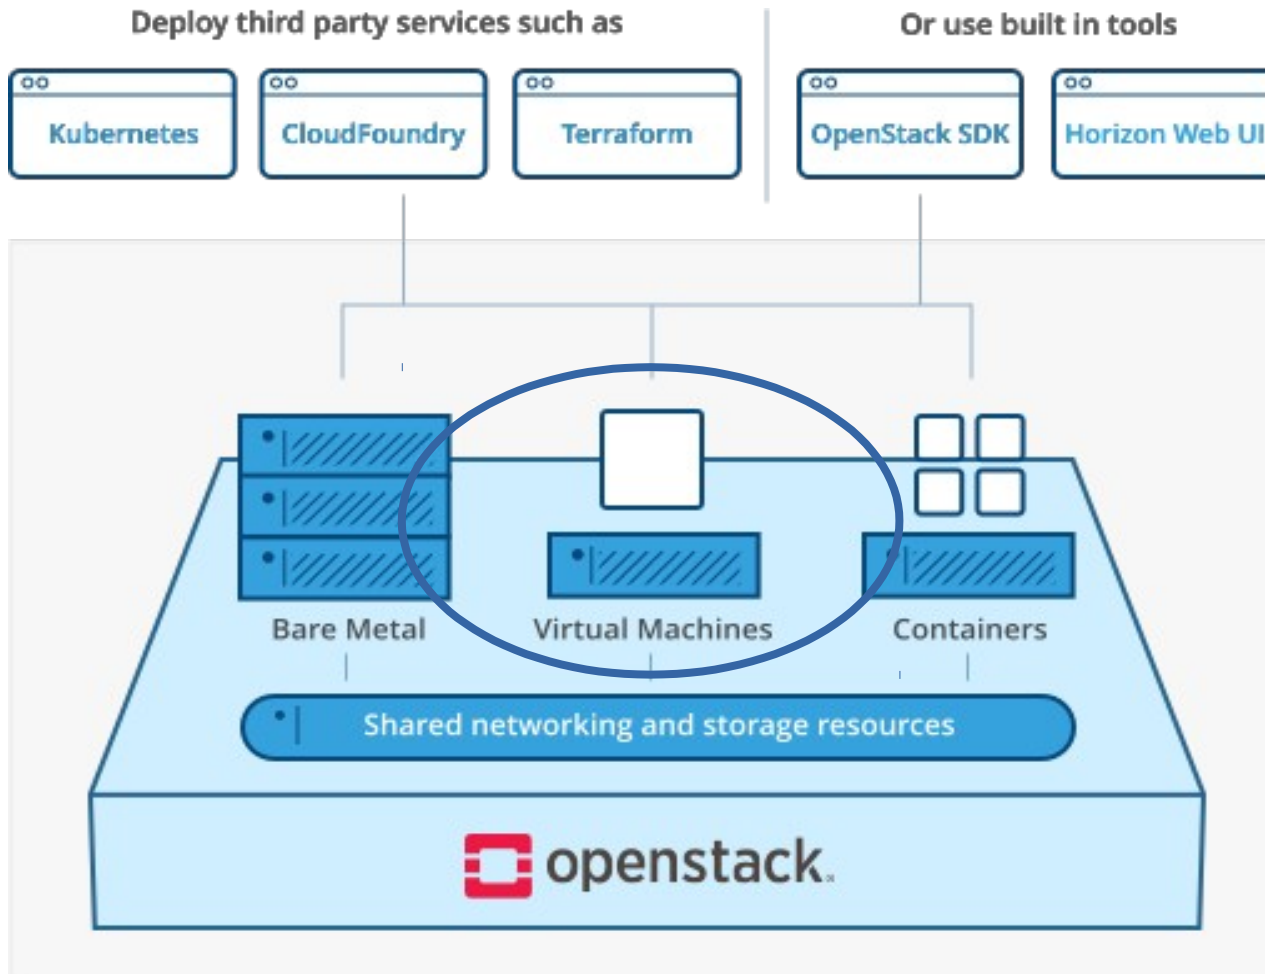


Infrastructure as a Service

Cloud Computing
Infrastructure as code



Primary Field of Application in Cloud Computing: Container Orchestration with Kubernetes

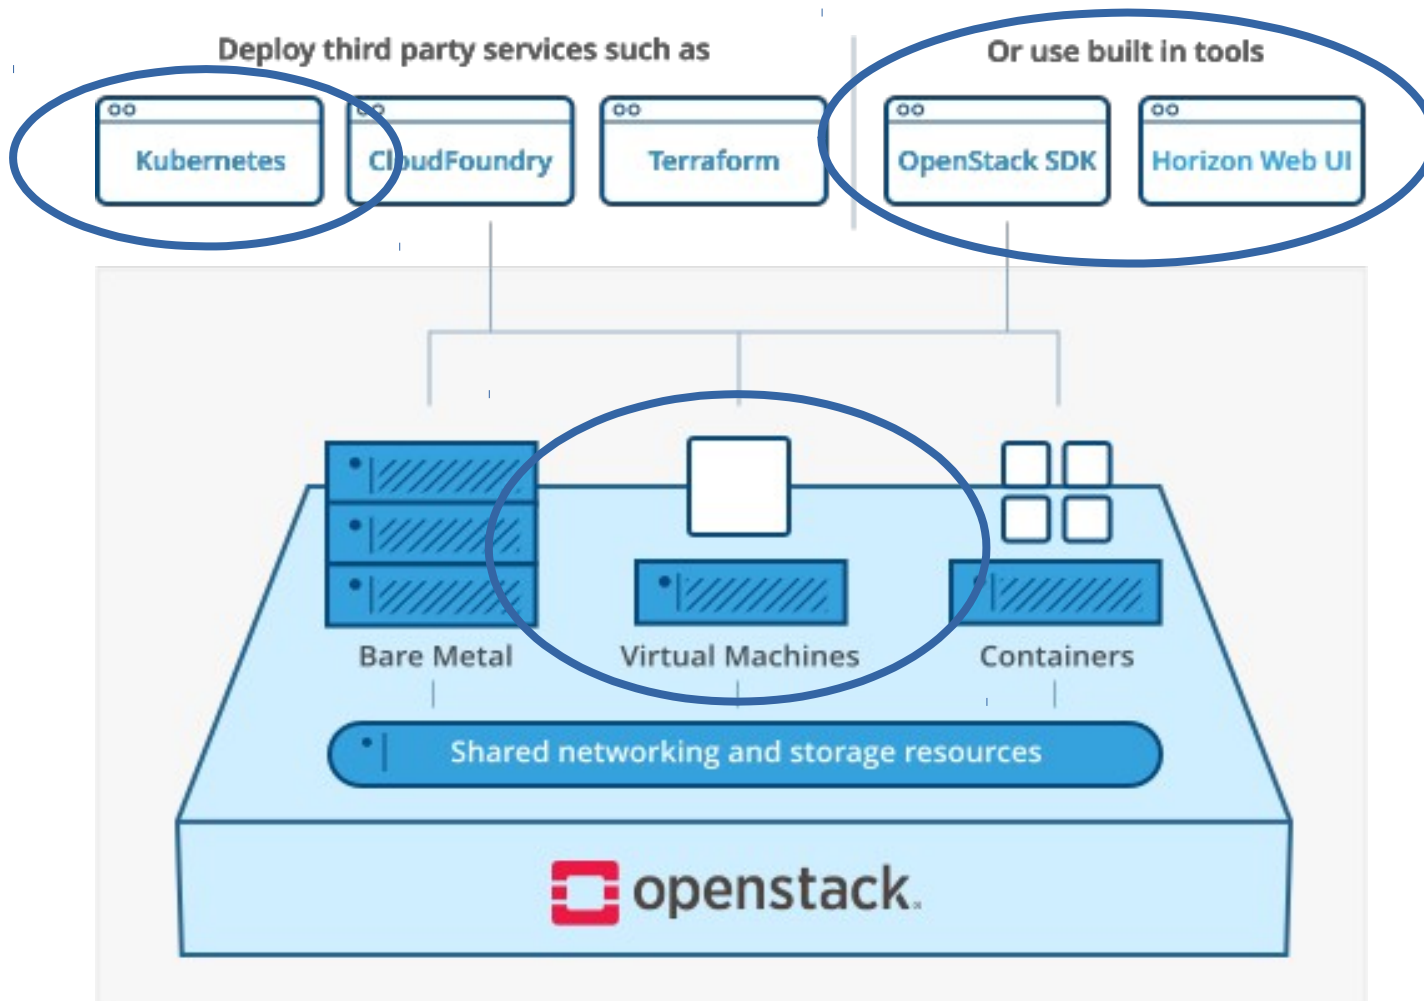


Openstack

- Used as a **virtualization platform**
- Not using modules for management of containers or bare metal servers

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

Primary Field of Application in Cloud Computing: Container Orchestration with Kubernetes



Kubernetes (K8s)

- for **containerized applications**, automating deployment and scaling
- Using it on clusters of **virtual machines**
- Not using it on bare metal servers

Openstack

- Used as a **virtualization platform**
- Not using modules for management of containers or bare metal servers

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

Software stack container orchestration

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
Docker registry



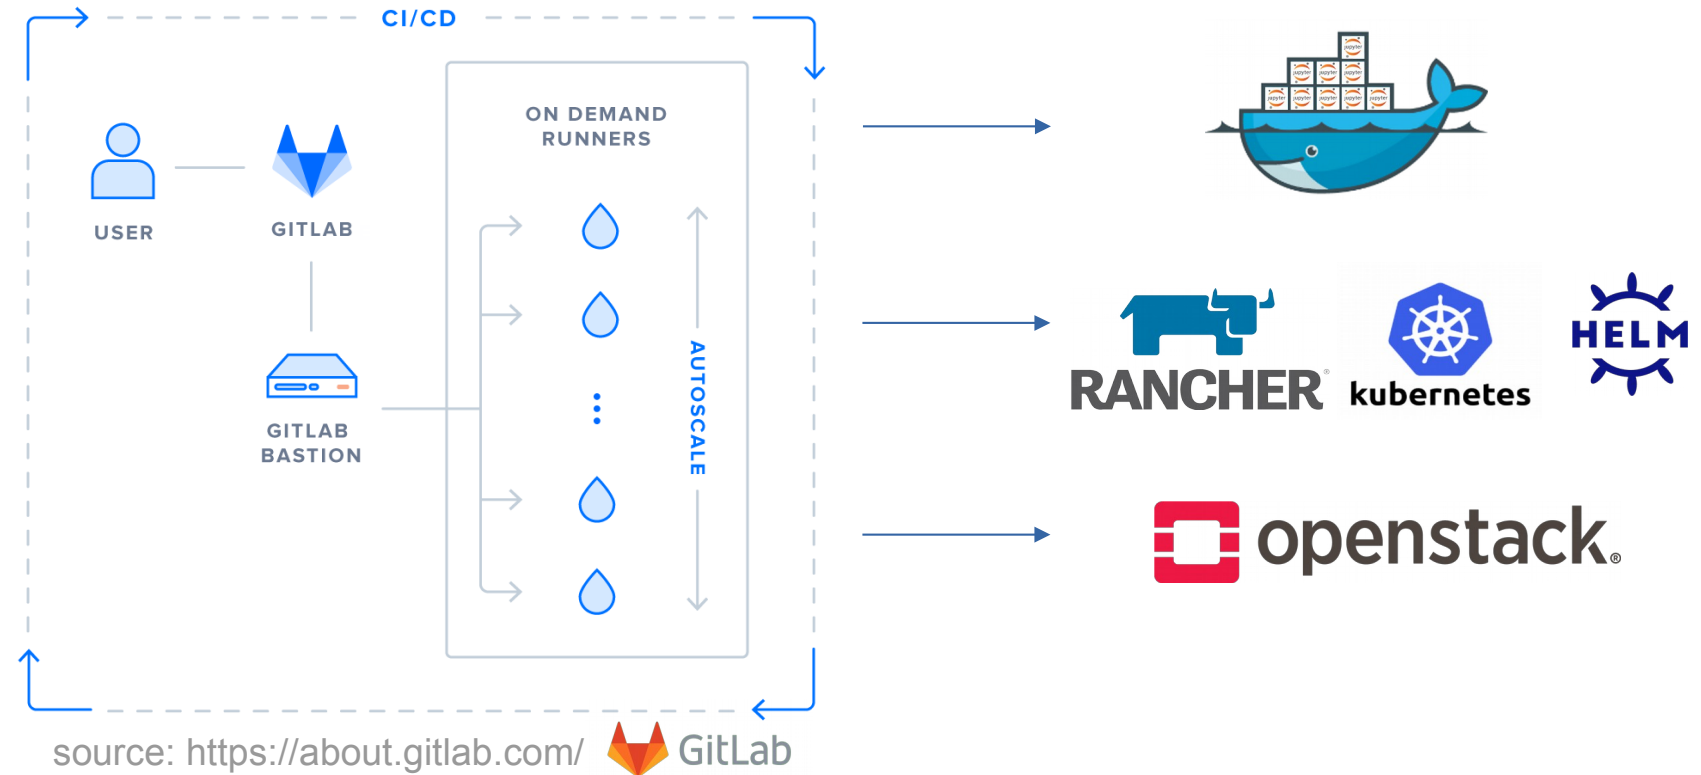
Infrastructure as a Service

Cloud Computing
Infrastructure as code



GitLab CI/CD for Container and Cloud Applications

- Web UI
- Version control
- Auto-scaling CI/CD
- Container Registry
- Secret Management
 - per project, group
 - per CI/CD job



GitLab CI/CD for Container and Cloud Applications

gitops-demo > apps > My ASP .Net App1 > Pipelines > #88314435

passed Pipeline #88314435 triggered 1 month ago by Brad Downey

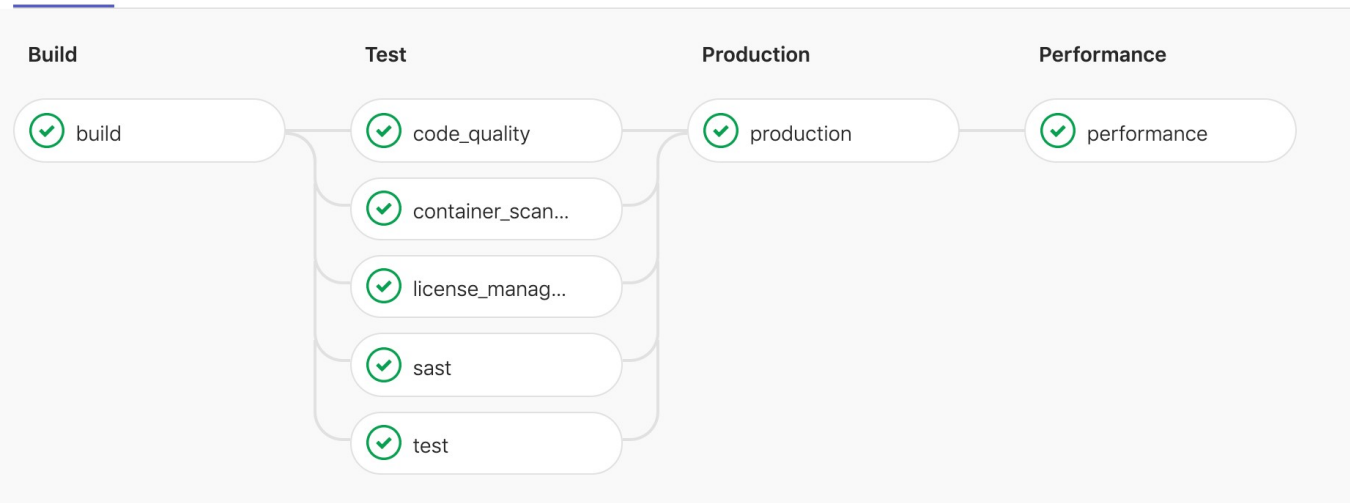
Remove staging

8 jobs for master in 14 minutes and 21 seconds (queued for 9 minutes and 59 seconds)

926f65b8

No related merge requests found.

Pipeline Jobs 8 Security Licenses 12



Git repositories as “single source of truth” for all infrastructure and application deployments

GitOps demo video and article by GitLab on “How To Deploy applications using GitLab CI, Helm and Kubernetes”:

<https://about.gitlab.com/blog/2019/11/18/gitops-prt-3/>

Build Docker Images with Gitlab CI/CD – Group Registry

K k8s

Group overview

Issues 2

Merge Requests 0

Kubernetes

Packages & Registries

Package Registry

Container Registry

Members

Settings

k8s > Container Registry

Container Registry

16 Image repositories

With the GitLab Container Registry, every project can have its own space to store images. [More information](#)

Image Repositories

Filter by name

k8s/jupyterhub-nfs 20 Tags

k8s/custom-pyspark 13 Tags

k8s/custom-pyspark/base36 1 Tag

k8s/helm-ci-image 3 Tags

DESY. | Container Orchestration and GitLab CI/CD with DESY Openstack as Cloud Provider | Michael Schuh, Johannes Reppin | Sep 23 2020

Page 15

Software stack container orchestration

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

Docker registry



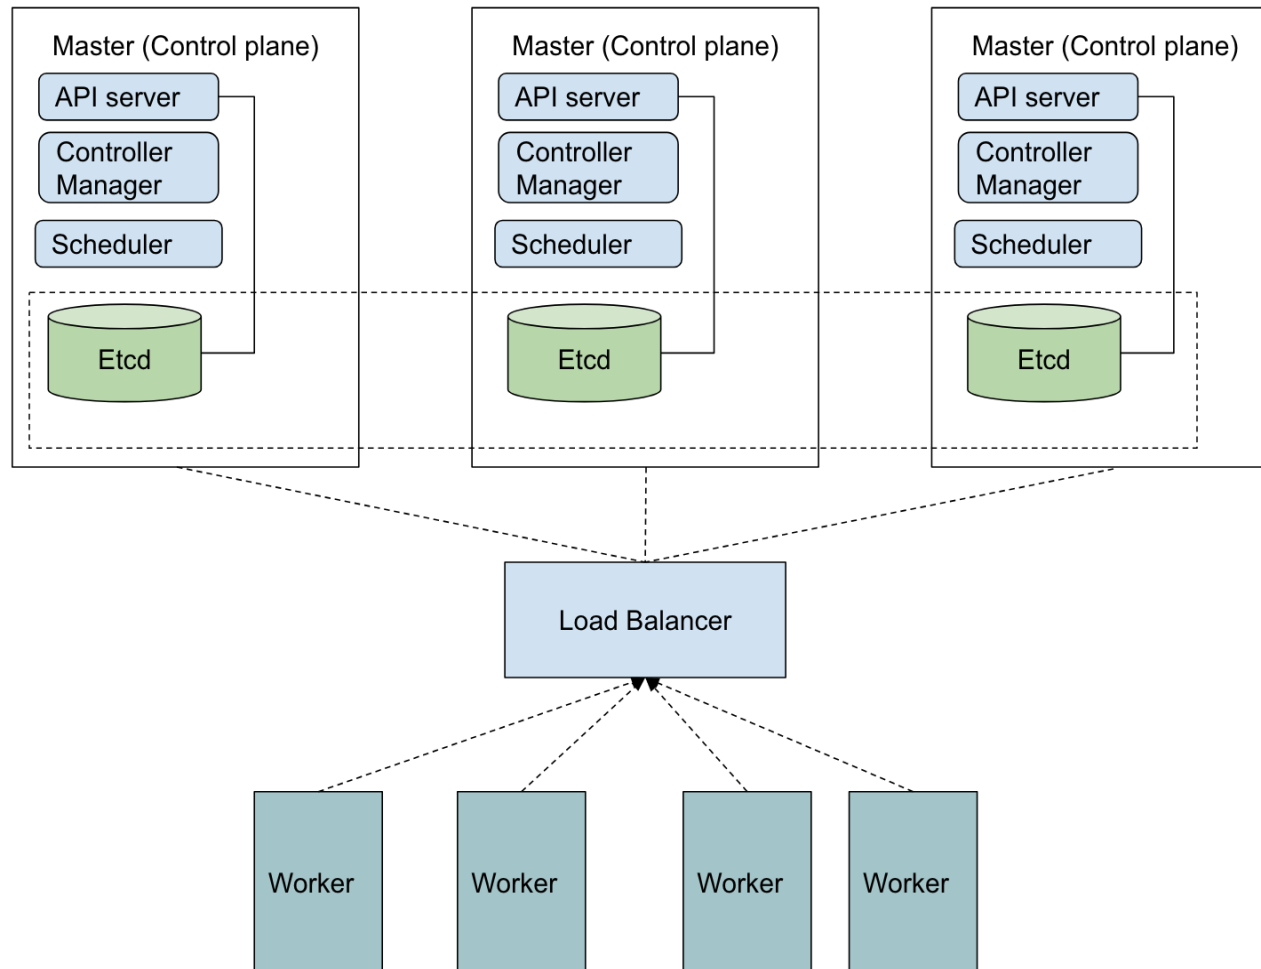
Infrastructure as a Service

Cloud Computing

Infrastructure as code



Rancher Installs Kubernetes Components



- Aims to be simple, fast, work anywhere
- Install VMs for node pools
 - worker node type 1
 - worker node pool 2
 - master nodes
 - ...
- Etcd demanding in disk IO rate
 - Etcd on master nodes assumed to be stable on network block device for < 100 nodes
 - Larger clusters require SSD or High Performance Block device

Source: <https://rancher.com/learning-paths/introduction-to-kubernetes-architecture/>

Rancher Server – Node Templates

Node Templates

[Manage Cloud Credentials](#) [Add Template](#)

Delete 1 Item

<input type="checkbox"/>	State	Name	Owner	Provider	Location	Size	
Owner:							
<input checked="" type="checkbox"/>	Active	k8s-centos-master		OpenStack	RegionOne	N/A	
<input type="checkbox"/>	Active	k8s-centos-node		OpenStack	RegionOne	N/A	
<input type="checkbox"/>	Active	rancherOS-k8s-network		OpenStack	RegionOne	N/A	
<input type="checkbox"/>	Active	rancherOS-master		OpenStack	RegionOne	N/A	
<input type="checkbox"/>	Active	rancherOS-node		OpenStack	RegionOne	N/A	
<input type="checkbox"/>	Active	ubuntu-k8s-master		OpenStack	RegionOne	N/A	
<input type="checkbox"/>	Active	ubuntu-k8s-node		OpenStack	RegionOne	N/A	
Owner: Michael Schuh							
<input type="checkbox"/>	Active	k8s-master-ubuntu-18-1908	Michael Schuh	OpenStack	RegionOne	N/A	
<input type="checkbox"/>	Active	ubuntu-18-1908-docker	Michael Schuh	OpenStack	RegionOne	N/A	

- Clusters on any
 - **cloud provider**
 - bare metal server
 - virtualization platform
- Node templates for **Openstack provider**
 - VMS as k8s nodes
 - Configure Docker

Rancher Server – Node Templates

Node Templates

Delete 1 Item

☐ State

Name

Owner:		
<input checked="" type="checkbox"/>	Active	k8s-centos-master
<input type="checkbox"/>	Active	k8s-centos-node
<input type="checkbox"/>	Active	rancherOS-k8s-network
<input type="checkbox"/>	Active	rancherOS-master
<input type="checkbox"/>	Active	rancherOS-node
<input type="checkbox"/>	Active	ubuntu-k8s-master
<input type="checkbox"/>	Active	ubuntu-k8s-node
Owner: Michael Schuh		
<input type="checkbox"/>	Active	k8s-master-ubuntu-18-1908
<input type="checkbox"/>	Active	ubuntu-18-1908-docker

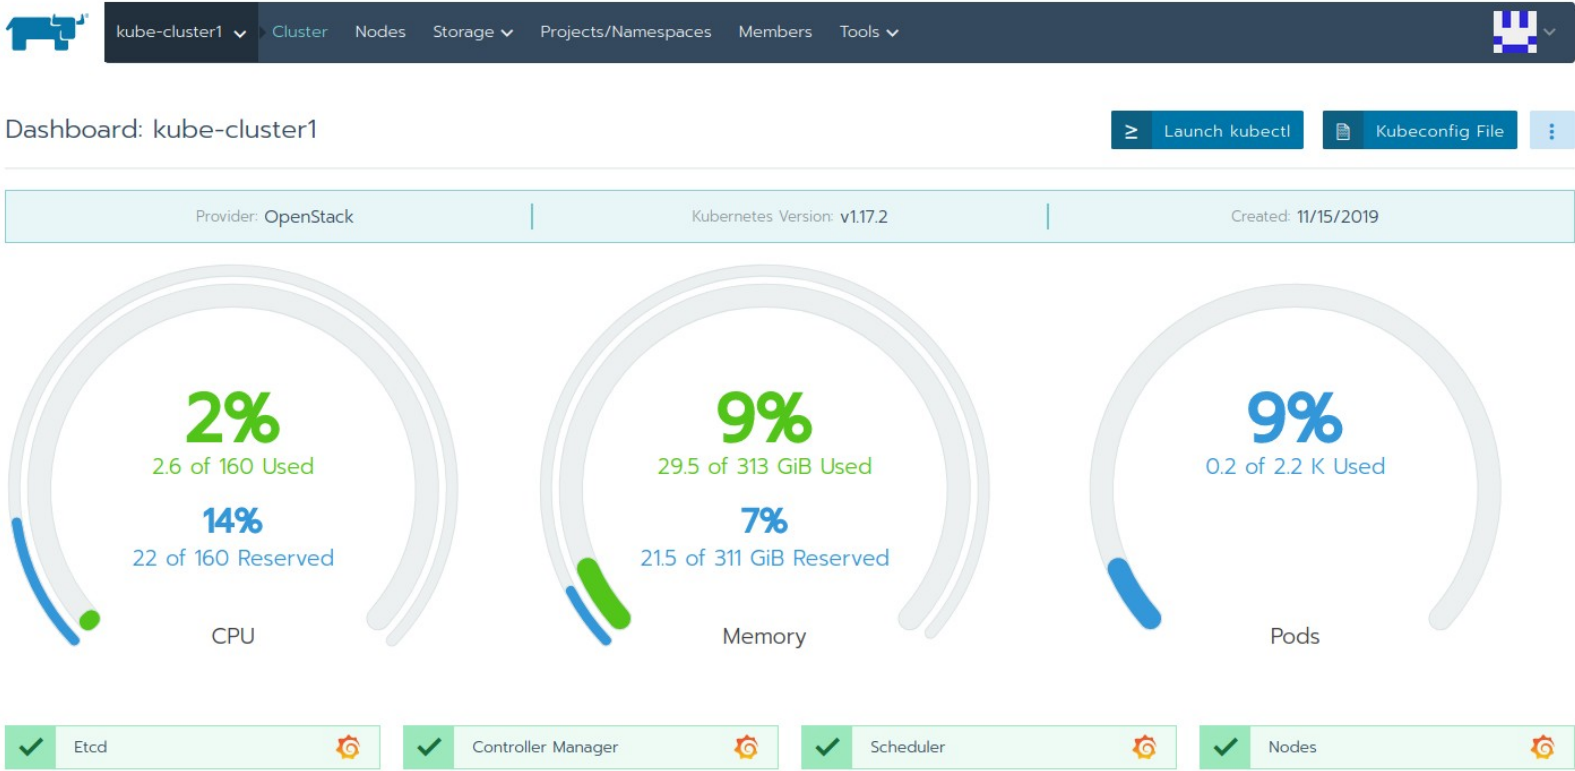
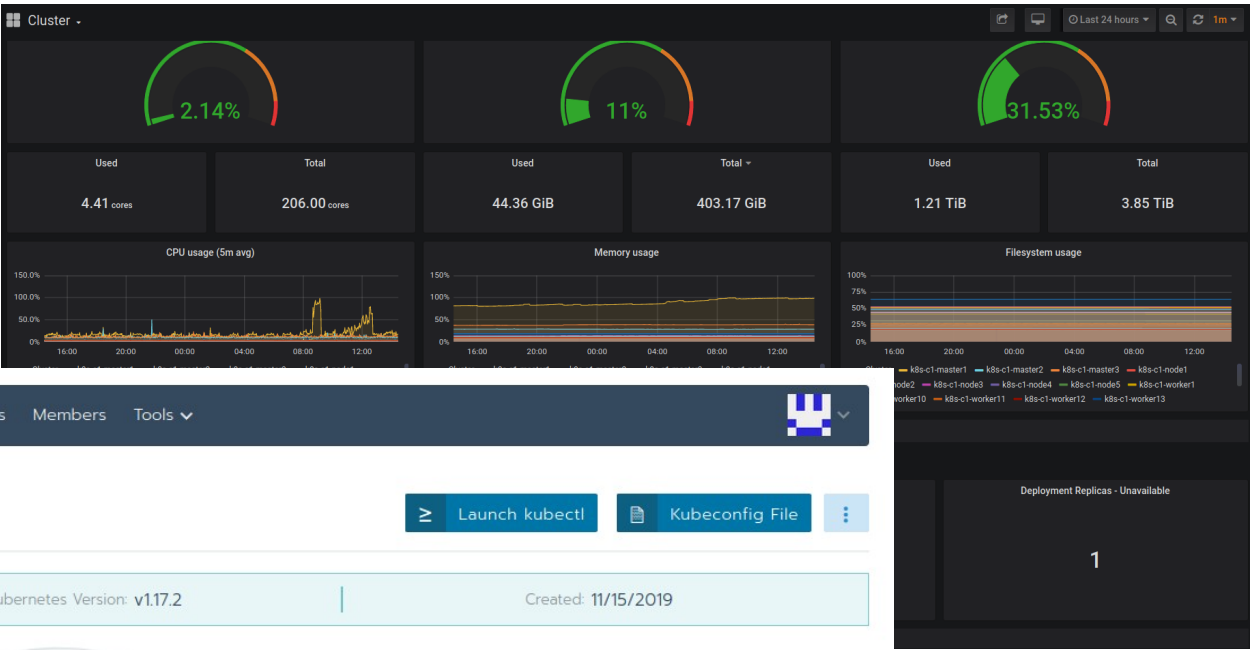
```
1 {
2   "name": "apitemplate-test3",
3   "driver": "openstack",
4   "engineRegistryMirror": [
5     "https://eosc-pan-dhub.desy.de:5000"
6   ],
7   "engineStorageDriver": "overlay2",
8   "openstackConfig":
9   {
10    "activeTimeout": "200",
11    "authUrl": "https://keystone-tank.desy.de:5000/v3/",
12    "availabilityZone": "nova",
13    "configDrive": false,
14    "applicationCredentialId": "APPLICATION_ID",
15    "applicationCredentialSecret": "APPLICATION_SECRET",
16    "domainId": "3d1fb9e6b4744ac9937c8727163ad560",
17    "endpointType": "publicURL",
18    "flavorName": "m1.large",
19    "imageName": "ubuntu-20-focal",
20    "insecure": false,
21    "ipVersion": "4",
22    "netId": "eaab545b-b1e0-49a7-be18-1a5501ad1758",
23    "novaNetwork": false,
24    "region": "RegionOne",
25    "secGroups": "ssh,web,kubernetes",
26    "sshPort": "22",
27    "sshUser": "ubuntu",
28    "userDataFile": null
29  }
30 }
```

- Clusters on any
 - **cloud provider**
 - bare metal server
 - virtualization platform
- Node templates for **Openstack provider**
 - VMS as k8s nodes
 - Configure Docker

Rancher Client, Dashboard and Cluster Monitoring

```
schuhn@t1pcx34859:~/Desktop/git/hifis-mediawiki-oidc$ rancher cluster
CURRENT ID STATE NAME PROVIDER NODES CPU RAM PODS
* c-754m5 active gitlab-public-ci Rancher Kubernetes Engine 6 3.58/40 1.71/77.83 GB 37/550
c-dsrzl active kube-cluster1 Rancher Kubernetes Engine 28 52.94/200 50.26/389.16 GB 249/2750
c-knighz active guest-k8s Rancher Kubernetes Engine 9 6.58/48 8.82/93.40 GB 81/660
c-w8vl2 active panosc Rancher Kubernetes Engine 3 0.98/24 0.17/46.62 GB 18/330

schuhn@t1pcx34859:~/Desktop/git/hifis-mediawiki-oidc$ rancher nodes
ID NAME STATE POOL DESCRIPTION
c-dsrzl:m-2dfq2 k8s-c1-worker13 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-2tkhq k8s-c1-worker16 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-4x2sw k8s-c1-worker14 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-5qfhn k8s-c1-worker1 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-5lchf k8s-c1-worker8 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-5r9v2 k8s-c1-master3 active k8s-c1-master k8s-c1-master
c-dsrzl:m-64n6t k8s-c1-worker20 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-66g2l k8s-c1-node3 active k8s-c1-node k8s-c1-node
c-dsrzl:m-6qs7g k8s-c1-worker4 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-72gww k8s-c1-worker10 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-8mm4t k8s-c1-worker11 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-8qj2b k8s-c1-worker6 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-9l9xw k8s-c1-node1 active k8s-c1-node k8s-c1-node
c-dsrzl:m-c8shw k8s-c1-worker15 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-cwtwq k8s-c1-worker3 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-dchvw k8s-c1-worker5 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-h4rr7 k8s-c1-master1 active k8s-c1-master k8s-c1-master
c-dsrzl:m-l79sw k8s-c1-node5 active k8s-c1-node k8s-c1-node
c-dsrzl:m-lgm6m k8s-c1-master2 active k8s-c1-master k8s-c1-master
c-dsrzl:m-nlrz6 k8s-c1-node2 active k8s-c1-node k8s-c1-node
c-dsrzl:m-p4c6s k8s-c1-node4 active k8s-c1-node k8s-c1-node
c-dsrzl:m-q2n64 k8s-c1-worker2 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-sbncj k8s-c1-worker18 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-t7jn7 k8s-c1-worker7 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-tshhn k8s-c1-worker9 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-xt8cr k8s-c1-worker19 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-z5z5g k8s-c1-worker12 active k8s-c1-worker k8s-c1-worker
c-dsrzl:m-z8kp5 k8s-c1-worker17 active k8s-c1-worker k8s-c1-worker
```



Integrate Kubernetes with Gitlab

0os-ci

Project overview

Repository

Issues0

Merge Requests0

CI / CD

Operations

Metrics

Environments

Error Tracking

Serverless

Logs

Kubernetes●

Packages

Analytics

Wiki

openstack > os-ci > Kubernetes

Add Kubernetes cluster

Kubernetes clusters can be used to deploy applications and to provide Review Apps for this project

Clusters are utilized by selecting the nearest ancestor with a matching environment scope. For example, project clusters will override group clusters. [More information](#)

Kubernetes cluster	Environment scope
gitlab-public-ci	* <div>Instance cluster</div>

Integrate Kubernetes with Gitlab

Details Applications Advanced Settings

GitLab Integration

Environment scope

*

* is the default environment scope for this cluster. This means that all jobs, regardless of their environment, will use this cluster. [More information](#)

Base domain

Specifying a domain will allow you to use Auto Review Apps and Auto Deploy stages for [Auto DevOps](#). The domain should have a wildcard DNS configured matching the domain. [More information](#).

Save changes

Provider details

Collapse

See and edit the details for your Kubernetes cluster

Kubernetes cluster name

gitlab-public-ci

API URL

https://131.169.234.119:6443

CA Certificate

```
-----BEGIN CERTIFICATE-----
MIICWjCCAAqgAwIBAgIBADANBgkqhkiG9w0BAQsFADASMRAwDgYDVQQDEwdrdWJl
```

Software stack container orchestration

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
Docker registry



Infrastructure as a Service

Cloud Computing
Infrastructure as code



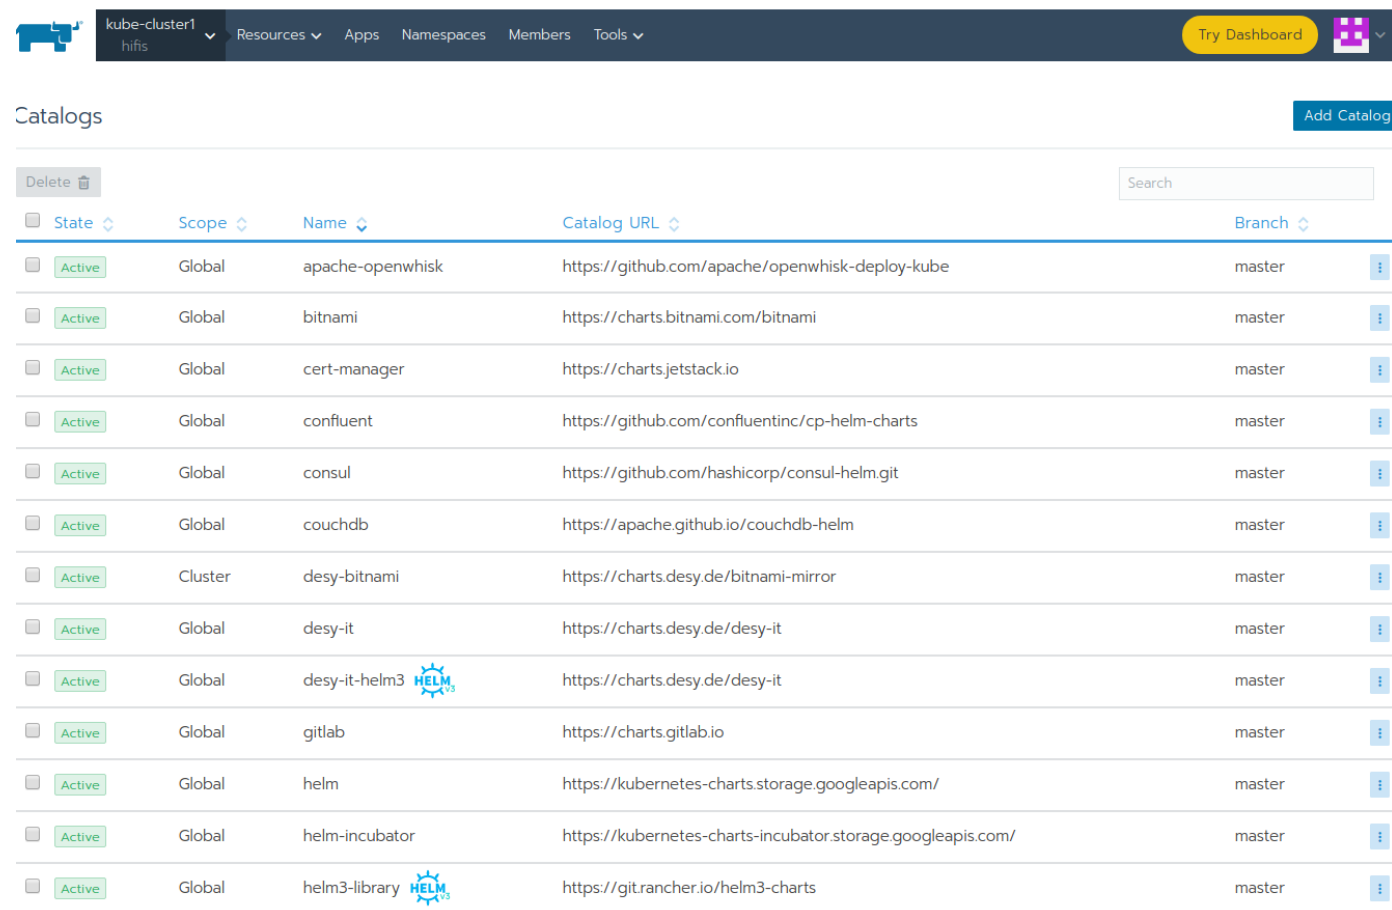
Managing Apps with Helm

Chart Repository





charts.desy.de

- Repository for Helm Chart Tarballs
 - Push new version to GitLab Repo
 - Easy rollback, if update fails
 - Hosted on Kubernetes
 - Simple REST API
- Add chartmuseum as app catalog to Rancher



The screenshot shows the Rancher UI interface. At the top, there's a navigation bar with a kube-cluster1 dropdown, a hifis user indicator, and links for Resources, Apps, Namespaces, Members, and Tools. A 'Try Dashboard' button and a user profile icon are on the right. Below the navigation bar, the 'Catalogs' section is active, showing a list of installed catalogs. A 'Delete' button and a search input are at the top of the list. The list contains 14 entries, each with a checkbox, a state indicator (all are 'Active'), a scope (Global or Cluster), a name, a catalog URL, a branch (all are 'master'), and a three-dot menu icon. The 'desy-it-helm3' entry is highlighted with a blue background and a Helm logo icon.

<input type="checkbox"/>	State	Scope	Name	Catalog URL	Branch	
<input type="checkbox"/>	Active	Global	apache-openwhisk	https://github.com/apache/openwhisk-deploy-kube	master	
<input type="checkbox"/>	Active	Global	bitnami	https://charts.bitnami.com/bitnami	master	
<input type="checkbox"/>	Active	Global	cert-manager	https://charts.jetstack.io	master	
<input type="checkbox"/>	Active	Global	confluent	https://github.com/confluentinc/cp-helm-charts	master	
<input type="checkbox"/>	Active	Global	consul	https://github.com/hashicorp/consul-helm.git	master	
<input type="checkbox"/>	Active	Global	couchdb	https://apache.github.io/couchdb-helm	master	
<input type="checkbox"/>	Active	Cluster	desy-bitnami	https://charts.desy.de/bitnami-mirror	master	
<input type="checkbox"/>	Active	Global	desy-it	https://charts.desy.de/desy-it	master	
<input type="checkbox"/>	Active	Global	desy-it-helm3 	https://charts.desy.de/desy-it	master	
<input type="checkbox"/>	Active	Global	gitlab	https://charts.gitlab.io	master	
<input type="checkbox"/>	Active	Global	helm	https://kubernetes-charts.storage.googleapis.com/	master	
<input type="checkbox"/>	Active	Global	helm-incubator	https://kubernetes-charts-incubator.storage.googleapis.com/	master	
<input type="checkbox"/>	Active	Global	helm3-library 	https://git.rancher.io/helm3-charts	master	

Additional software components

“bare” Kubernetes is not enough

Nginx Ingress Controller

- Direct traffic to pods

MetalLB Loadbalancer

- Level2 Loadbalancer for Kubernetes

Cinder Storage Class

- Automatically Provision Volumes in Ceph
- Currently: one *pod* can attach a Volume at a time

Cert Manager

- Provides Let's Encrypt Certificates
- Installed into Kubernetes Cluster
- Watches the Kubernetes API for *Ingress* Objects



NGINX



MetalLB



Managing Apps with Helm

Example Deployment of Jupyterhub



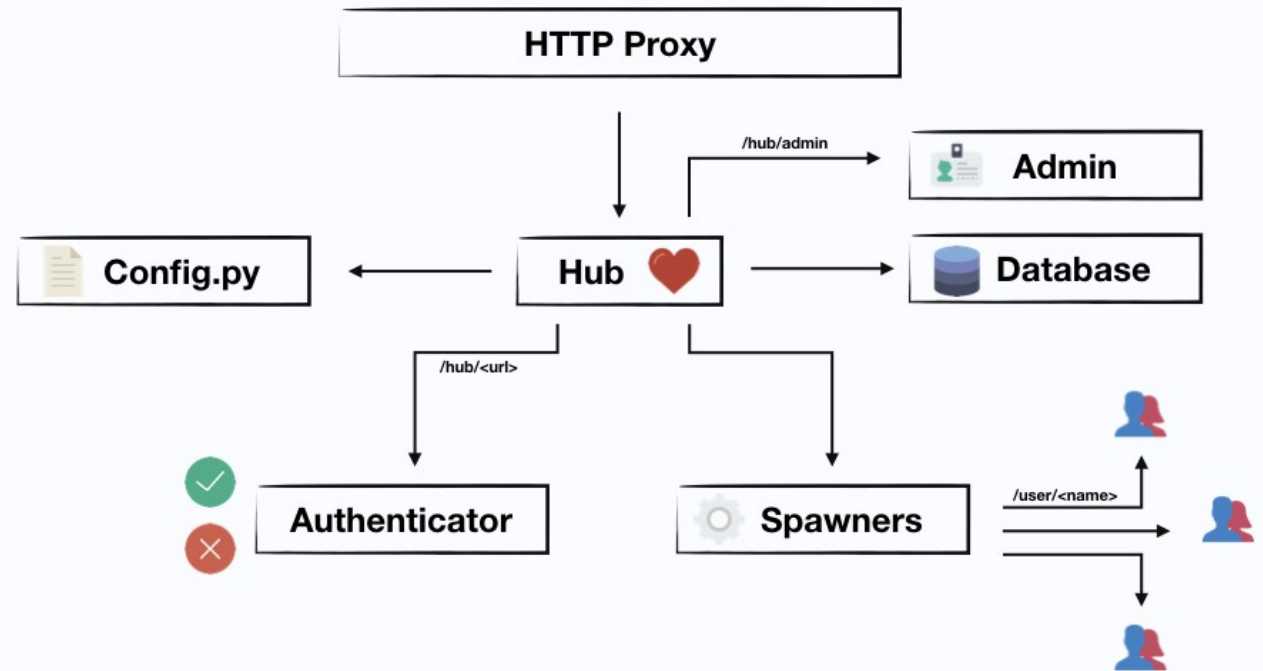
eosc-pan-jhub.desy.de

- Hub, proxy, user-scheduler, image-puller
- User pods spawned on demand
- Persistent data volumes
- LE Certificate

Deployment


- Add DNS alias to Loadbalancer IP
- Add helm chart repository
- Get customizable Values and edit
- Install to k8s (helm install)

JupyterHub




All icons where obtain on Flaticon (<https://www.flaticon.com/packs/essential-collection>)

Rancher Server – Deploy Helm Charts as Applications



kube-cluster1
Data-Analysis


Resources Apps Namespaces Members Tools



Apps

Manage CatalogsLaunch

Search




dask

Up to date (4.16)Active

443/https, 30310/tcp, 30730/tcp

3




jupyterhub

Up to date (0.9.0-beta.4.n000.h72e54ee)Active

443/https

24




project-monitoring

Up to date (0.0.7)Active

/index.html, /index.html

2



spark-pure

Up to date (1.2.6)Active

443/https

3

Jupyter Notebooks

Extensible environments for interactive data analysis

Plot Slow Data


With `run` you can plot slow data by generating a pandas dataframe of the same. This is implemented in `karabo_data` by the `get_dataframe` function. Here we get the *beam position* and the *photon flux* from the XGM data of the run.

```
1 # get data as pandas dataframe
2 df = run.get_dataframe(fields=[("XGM/**", "**.i[xy]Pos"), ("XGM/**", "**.photonFlux")])
```

Everything else is basic plotting commands with the created dataframe.

```
1 ##-----
2 ## from here, basic plotting in python
3 fig, ax = subplots(1, 2, figsize=(9,4)) # make a figure with two subplots
4 ax2 = ax.flat[0].twinx() # twin the x axis to plot on two different y scales
5
6 ax.flat[0].set_xlabel('train index',)
7
8 # plot beam x on left y axis
9 color = 'orangered'
10 df.reset_index().plot(y=df.columns[0], ax=ax.flat[0], color=color, legend=False)
11 ax.flat[0].set_ylabel('beam x', color=color)
12 ax.flat[0].tick_params(axis='y', labelcolor=color)
13
14 # plot beam y on right y axis
15 color = 'navy'
16 df.reset_index().plot(y=df.columns[1], ax=ax2, color=color, legend=False)
17 ax2.set_ylabel('beam y', color=color)
18 ax2.tick_params(axis='y', labelcolor=color)
19
20 # plot photon flux on second graph
21 df.reset_index().plot(y=df.columns[2], ax=ax.flat[1], label='photon flux', legend=False)
22 ax.flat[1].set_xlabel('train index')
23 ax.flat[1].set_ylabel('photon flux')
24
25 # make the plots look better
26 for a in [ax, ax2]:
27     niceplot(a, grid=0)
28 tight_layout()
```


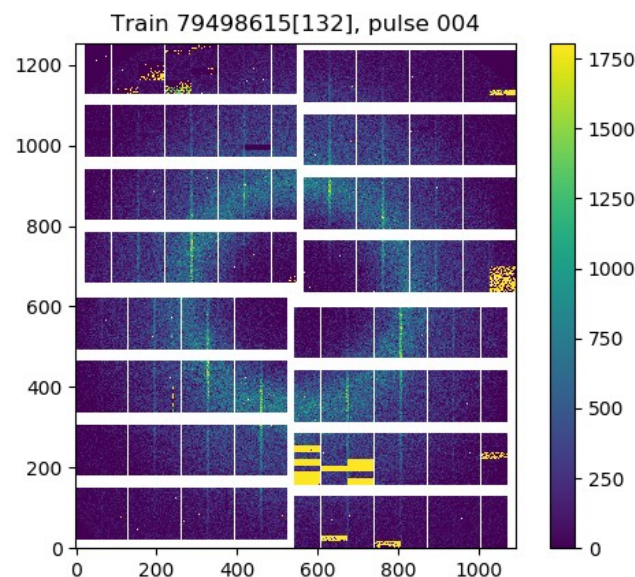
<IPython.core.display.Javascript object>



```
[6]: demo = InteractiveGeom(geom, run)
```

```
[7]: demo.interactive()
```

Figure 1



European XFEL Python data tools

Karabo Data

x-ray gas monitor data,
Beam position
and photon flux,
(Mario Reiser, Eu-XFEL)

Karabo Interactive
x-ray femto second
crystallography data,
interactive visualisation
(Robert Rosca, Eu-XFEL)

karabo-data.readthedocs.io

Software stack container orchestration

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
Docker registry



Infrastructure as a Service

Cloud Computing
Infrastructure as code



The Cloud Native Landscape – The Road Ahead



Integrate elastic environments

- Network Operations
 - DNS
 - LBaaS
 - X509 Certificates
- Scientific Data
 - dCache
 - High performance storage
- Data Acquisition streams
 - Event streaming platforms
- Scaling container registry to
 - HPC clusters
 - HTC clusters
- Software Repository

Contact

DESY Deutsches
Elektronen-Synchrotron

www.desy.de

Michael Schuh
Research and Innovation in Computing
michael.schuh@desy.de
+49 040 8998 2316

Johannes Reppin
Information Fabrics
johannes.reppin@desy.de
+49 040 8998 2096