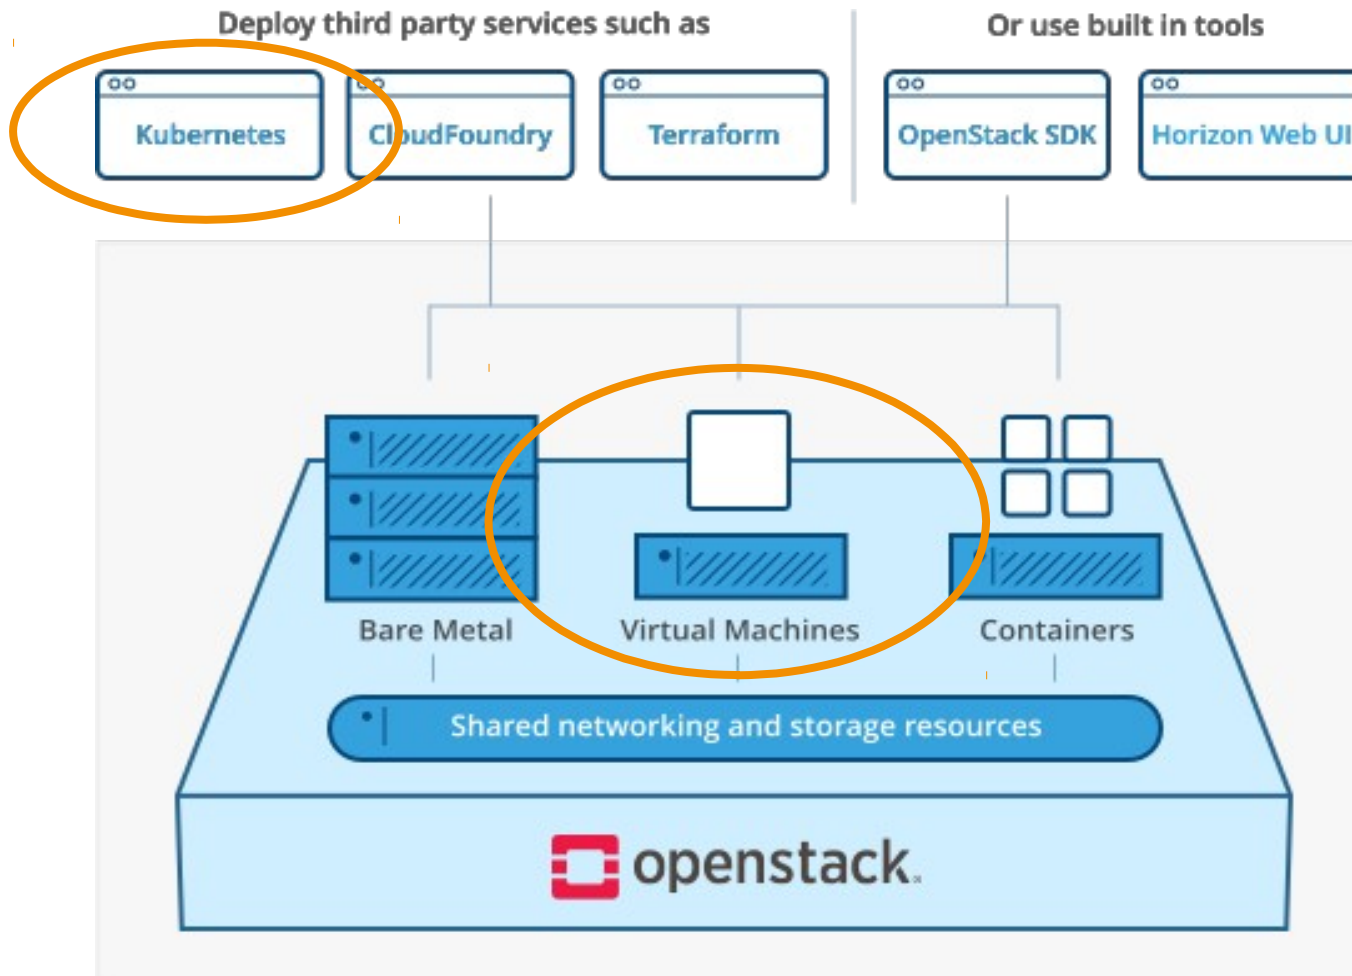


On-demand Scheduling of GPUs for CI/CD with Kubernetes on Openstack

Tim Wetzel, Michael Schuh, Patrick Fuhrmann
ISGC 2022

Primary Field of Application in Cloud Computing

Container Orchestration with Kubernetes



Kubernetes

- for **containerized applications**
- on clusters of **virtual machines**

Openstack

- Used as a **virtualization platform**
- Software defined networking

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

Software stack container orchestration


Software as a Service
SaaS



Container Orchestration
PaaS



Containerization
CI/CD



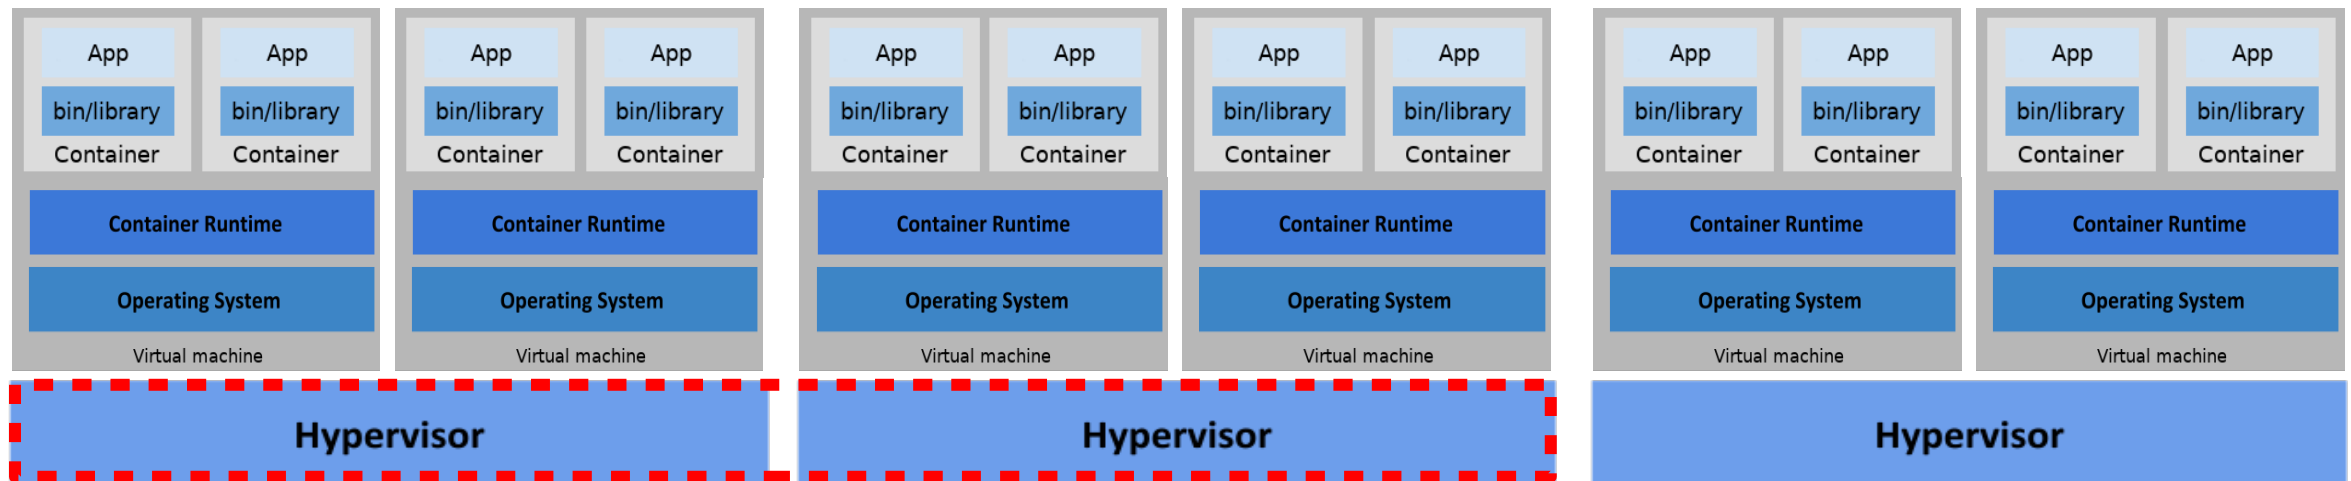
Cloud Computing
IaaS



Architecture

Clusters and Schedulers

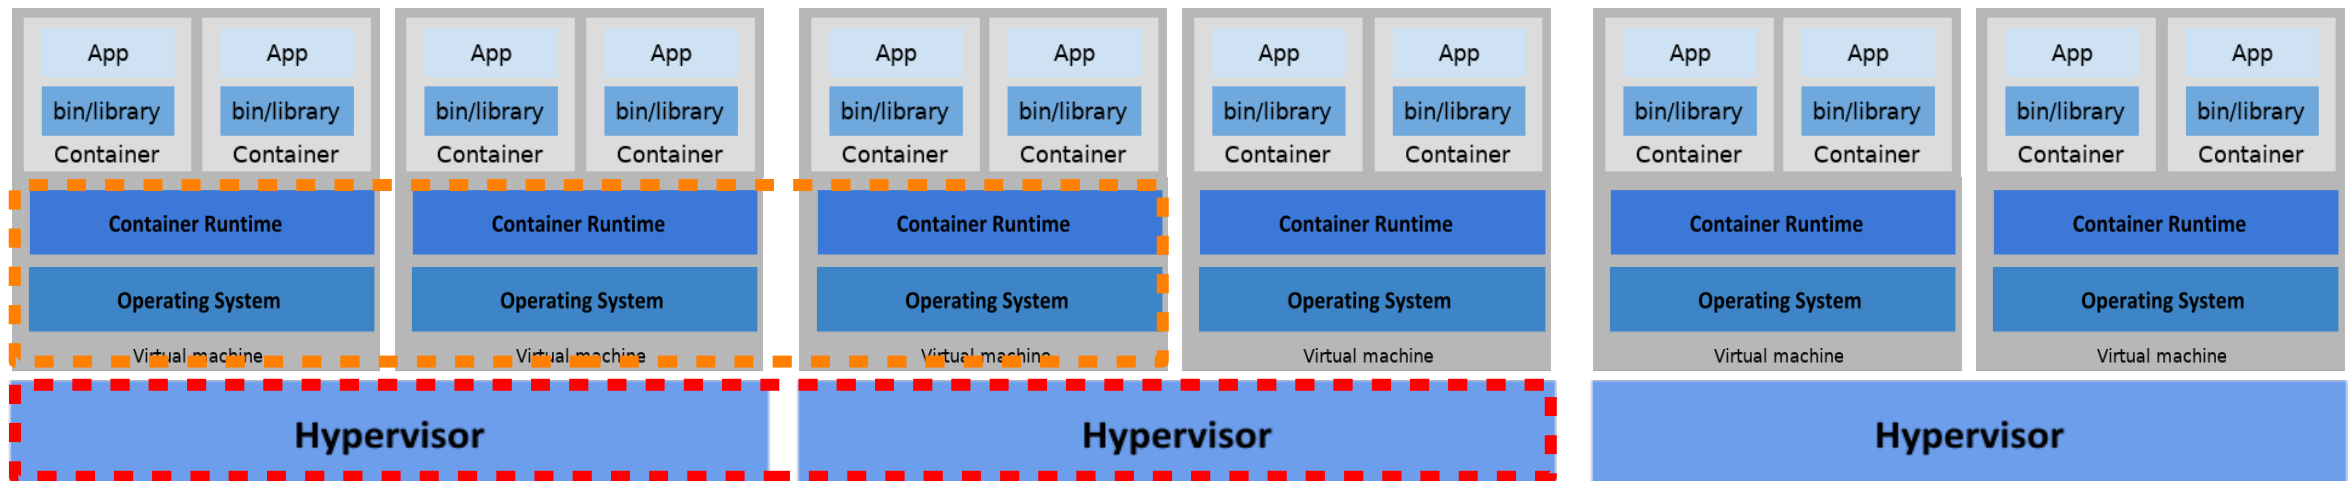
- Openstack Hypervisors organized as **Host Aggregates**
 - Openstack Nova (compute scheduler) places VMs on Hypervisors (hosts)



Architecture

Clusters and Schedulers

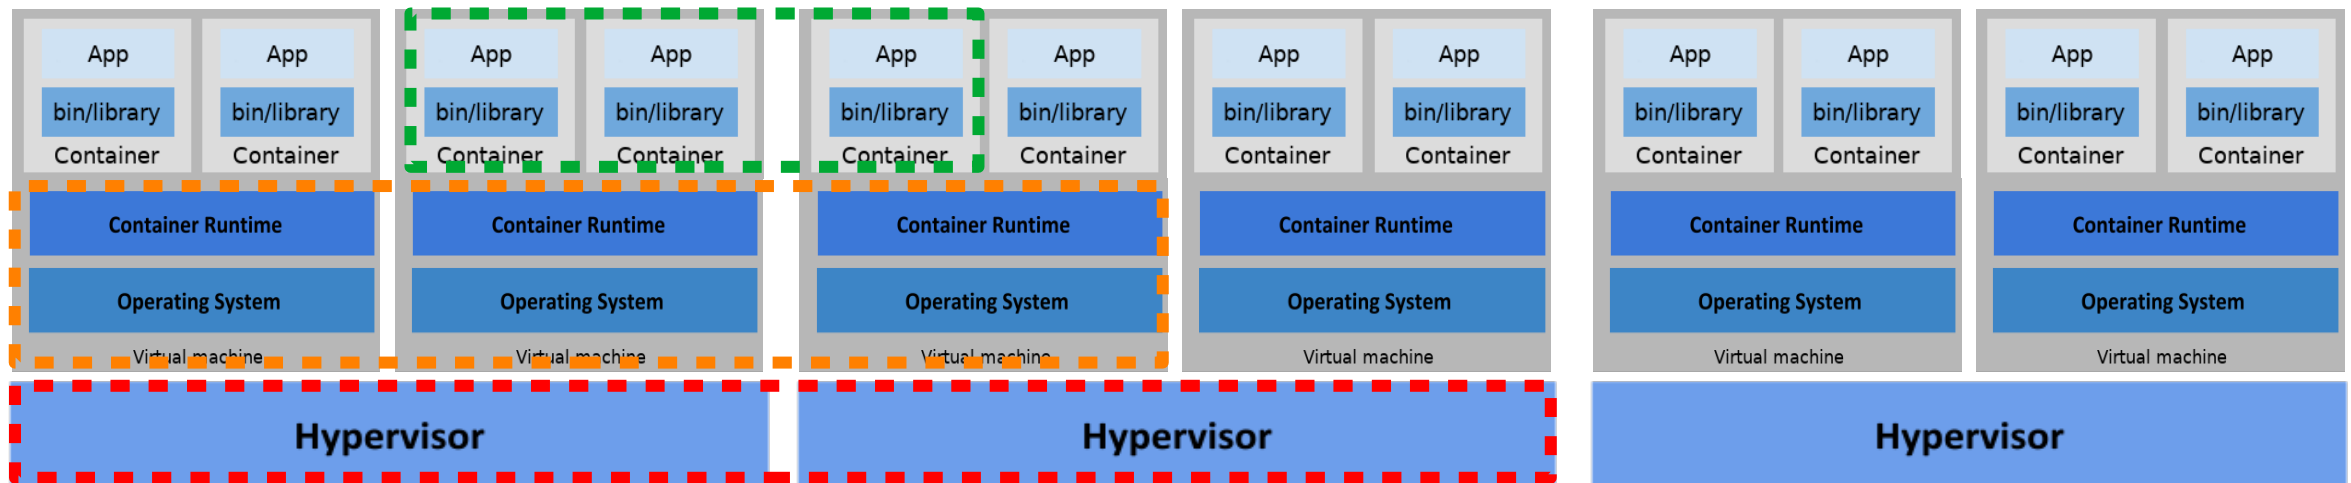

- VMs configured as Kubernetes nodes form a **Kubernetes Cluster**
 - Dynamically created by Rancher Kubernetes management platform
- Openstack Hypervisors organized as **Host Aggregates**
 - Openstack Nova (compute scheduler) places VMs on Hypervisors



Architecture

Clusters and Schedulers

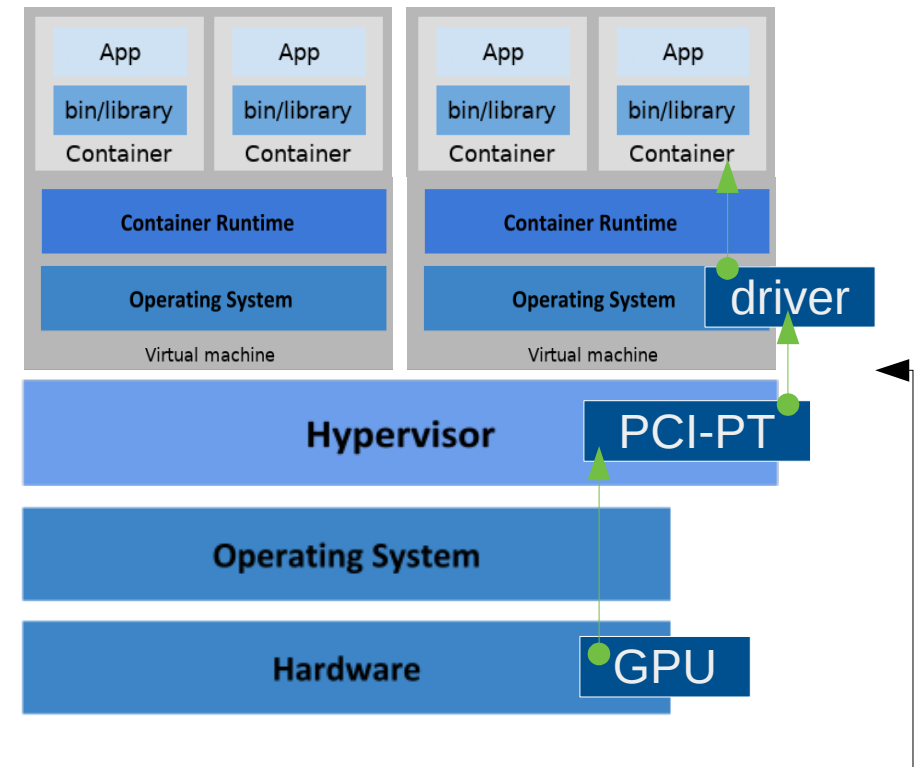

- Kubernetes compute manager places containers on Kubernetes Nodes
 - Application Deployments form **Container Clusters**
- VMs configured as Kubernetes nodes form a **Kubernetes Cluster**
 - Dynamically created by Rancher Kubernetes Management platform
- Openstack Hypervisors organized as **Host Aggregates**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



Architecture

GPU access from kubernetes pods via virtual machine

- Hypervisor prepared for PCI-passthrough of GPU
- VM image incl. GPU driver prepared in CI/CD pipeline
- VM has access to GPU through PCI-PT
- GPU resource nvidia.com/gpu on kubernetes node
- Can be used by all kubernetes pods, jobs, ...

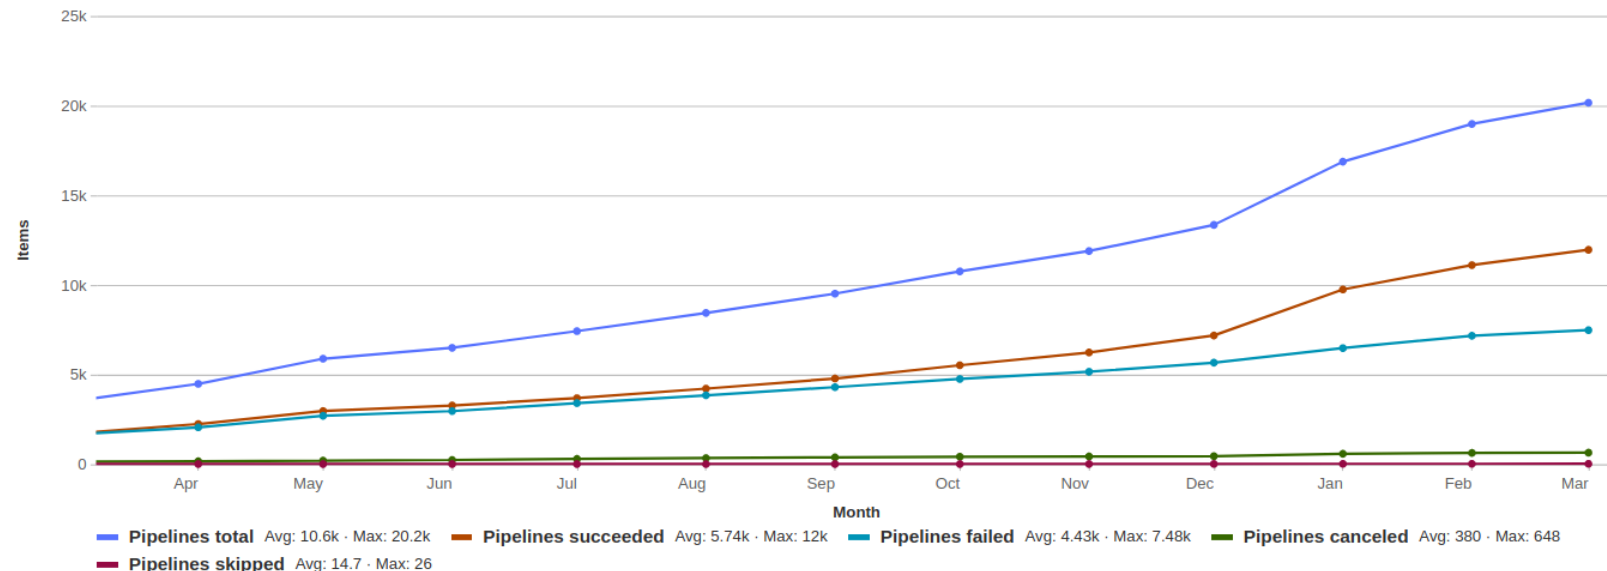
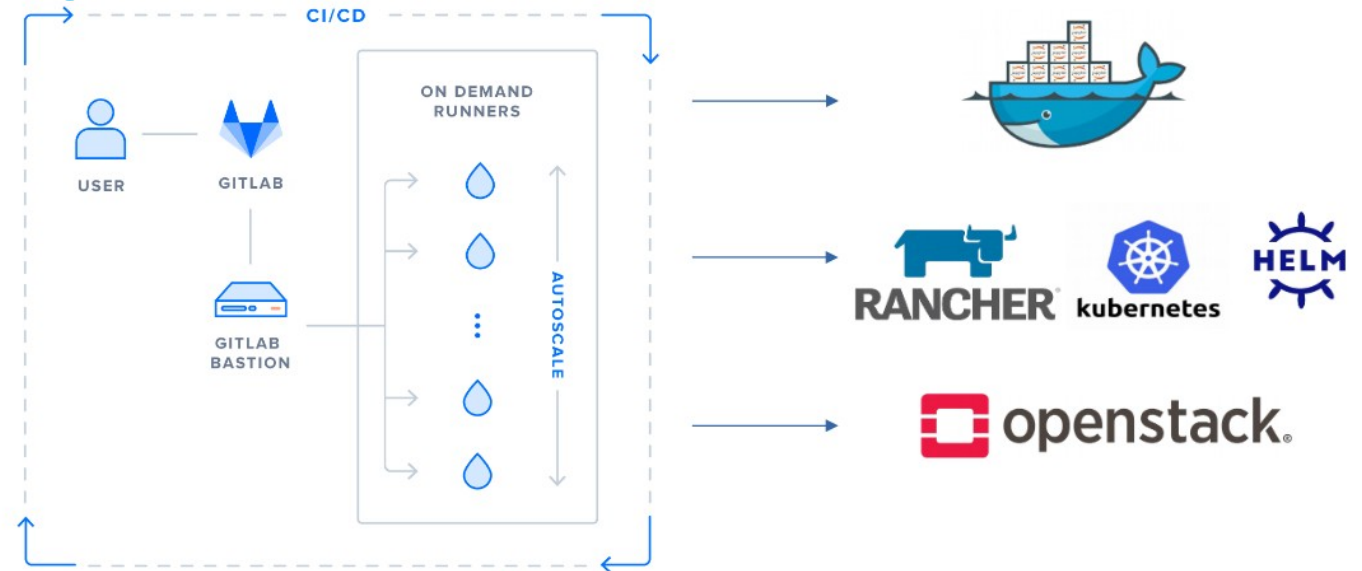


Gitlab Runners for gitlab.desy.de

Autoscaling k8s deployment

Helm Deployment on k8s

- Gitlab CI/CD pipelines / jobs
- Autoscaling service
- Shared runner pilot for gitlab.desy.de
- Used for
 - Container builds
 - Deployments
 - Tests



Gitlab Runners for gitlab.desy.de

Autoscaling k8s deployment

GPU runners

- Test instance with one Nvidia V100 (32GB)
- Tests done
 - Simple CI/CD script calling `nvidia-smi`
 - Compiling & executing CUDA code
 - Tensorflow benchmark (batch size variation for Resnet50 with FP16)
- Resource scheduling config for GPUs is still in development for gitlab-runners ¹
 - Workaround to avoid scheduling clashes on kubernetes: set runner's *concurrency* to 1
- N.b.: Runners not to be used for production scale applications

Tensorflow benchmarks used from <https://github.com/tensorflow/benchmarks> for multiple batch sizes.

Benchmarks on bare metal GPU instances from <https://lambdalabs.com/gpu-benchmarks> show a throughput of **1276 img/s** for the same GPU model with a batch size of 384 images.

Batch size	Throughput (img/s)
64	1027.9
128	1145.9
256	1261.6
384	1299.5
512	1326.3

¹ https://gitlab.com/gitlab-org/gitlab-runner/-/merge_requests/1526

Pilot use case

CI/CD for Helmholtz Imaging



The Helmholtz Imaging team at DESY provides expert knowledge and support to scientists working with imaging data and related AI/ML tasks within the Helmholtz Association

This involves developing pipelines for imaging analysis and AI/ML training in a collaborative fashion

Until now, the code development has to happen on IT systems with GPUs present, which are not tightly integrated into our versioning and collaboration platform gitlab.desy.de

By providing GPU-enabled Gitlab-runners to them, they can construct CI/CD workflows that automatically ensure that they only run executable code on the actual production systems

- less time spent with manual testing
- less production resources consumed during testing
- overall increase in efficiency expected

Benefits also expected for CUDA-based code

- Compile-time errors can be detected, runtime-errors not
- Regression testing with different CUDA versions possible

Summary

- Kubernetes cluster on Openstack VMs
- Pilot for Gitlab-runners with GPU in gitlab.desy.de
- GPU resource request in runners not available yet
- Use cases for Helmholtz Imaging and scientists developing with CUDA

Thank you

Contact

DESY. Deutsches
Elektronen-Synchrotron

www.desy.de

Tim Wetzel, Michael Schuh, Patrick Fuhrmann
Research and Innovation in Scientific Computing
tim.wetzel@desy.de