

Activities around Kubernetes in the CNCF Research User Group

Ricardo Rocha
CERN IT-CM
CNCF RUG , CNCF TOC

CNCF Research User Group

<https://github.com/cncf/research-user-group> : CNCF Slack: #ug-research

Meetings : [Agenda](#) (recordings linked in the agenda, live youtube soon)

Zoom meetup every 1st and 3rd Wednesday 8am PT / 5pm CET

Goals

Advance research computing using Cloud Native technologies

Bring people from research institutions together

Learn how others have solved common problems


Build a research-focussed cloud-native community

CNCF Research User Group

Participation from our community has been increasing recently

Including some presentations (CERN, EGI, UNL)

Representation from multiple domains




Recent Survey

Full results in an upcoming kubecon session...

Which industry or group do you represent?


12 responses



- Academia
- Government
- Non Profit
- Private Sector
- I represent multiple groups: Academia / Non-Profit / Private Sector


What areas of research do you primarily support?

12 responses



What is your current usage of Kubernetes?


12 responses



- Already using in production
- Still in evaluation phase
- Only for testing purposes
- Learning it for grant work to support researchers

What types of workloads do you run in your clusters?

12 responses



Topics Covered

Multi-cluster with Admiralty

POSIX Integrations (how to access hpc storage, slurm jobs etc)

OIDC (Lab / Institute SSO) integration with Kubernetes

Notebooks on Kubernetes (JupyterHub, Dask, GPU Integration, ...)

Image distribution / replicated registries / Pull-Through Caching

Lazy Loading of Container Images

Batch on Kubernetes (G-Research Armada)

Rootless Kubernetes

Cloud Bursting

GPU and GPU Sharing

Topic Covered

POSIX Integrations (how to access hpc storage, slurm jobs, etc)

Jeremy Rogers, ORNL/OLCF

Covering UID based access to shared filesystems from non root containers

Open Policy Agent (OPA) used to mutate workloads (admission controller)

Annotations with filesystem access information

sshexec-operator to manage user ssh keys for remote submission

kube-namespace-manager to keep user metadata (uid, ...)

Considering open sourcing the tools above

Topic Covered

POSIX Integrations (

Jeremy Rogers, ORN

Covering UID bas

Open Policy Agen

Annotations

sshexec-operator

kube-namespace

Considering ope

The screenshot shows a terminal window titled "deployment.yaml" with tabs for "TERMINAL", "PROBLEMS", "52", and "OUTPUT". The terminal output is as follows:

```
TERMINAL PROBLEMS 52 OUTPUT
1: tmux
"ccs.ornl.gov/supplementalGroups": "1099,2001,2324,24665,26694,27493,27537,28490",
"modified-by.stackrox.io/namespace-label-patcher": "true",
"openshift.io/node-selector": "region=primary",
"openshift.io/requester": "rogersjn3",
"openshift.io/sa.scc.mcs": "s0:c251,c603",
"openshift.io/sa.scc.supplemental-groups": "2324/1",
"openshift.io/sa.scc.uid-range": "15771/1"
}
→ ./rug-demo.sh oc get ns stf002-rug-demo -ojsonpath='{.metadata.labels}' | jq
{
  "ccs.ornl.gov/project": "STF002",
  "ccs.ornl.gov/projectID": "3865",
  "ccs.ornl.gov/projectType": "user",
  "namespace.metadata.stackrox.io/name": "stf002-rug-demo"
}
→ ./rug-demo.sh cat shared-file-system-pod.yaml
---
apiVersion: v1
kind: Pod
metadata:
  name: shared-file-system-pod
  namespace: stf002-rug-demo
  annotations:
    ccs.ornl.gov/fs: olcf
spec:
  containers:
    - name: user-image
      image: registry.apps.granite.ccs.ornl.gov/platform-public-images/user-image:latest
      stdin: true
      tty: true
      command:
        - /bin/bash
→ ./rug-demo.sh oc apply -f shared-file-system-pod.yaml
pod/shared-file-system-pod created
→ ./rug-demo.sh oc get pods -n stf002-rug-demo --watch
NAME          READY   STATUS    RESTARTS   AGE
shared-file-system-pod  1/1     Running  0           6s
AC→ ./rug-demo.sh oc get pod -n stf002-rug-demo shared-file-system-pod -oyaml | less
→ ./rug-demo.sh oc exec -it -n stf002-rug-demo shared-file-system-pod -- /bin/sh
```

A video call inset in the top right corner shows a man with glasses and a beard, identified as "Jeremy". At the bottom of the terminal window, there is a video player interface with a play button, a progress bar showing "00:14:16 / 00:52:19", a volume icon, a speed control icon, and a full-screen icon.

Topic Covered

Lazy Loading of Container Images - containerd remote snapshotter

Kohei Kotunaga, NTT

On average 6% of an image is actually read, but can be 76% of startup time

Much worse if images are very large - we're guilty as charged!

Support added in containerd, linuxkit, ko (knative), ...

Several group members testing at scale, including CERN

[Speeding up Analysis with Remote Container Images](#)

Topic Covered

Lazy Loading of Container Images - containerd remote snapshotter

Kohei Kotunaga, NTT

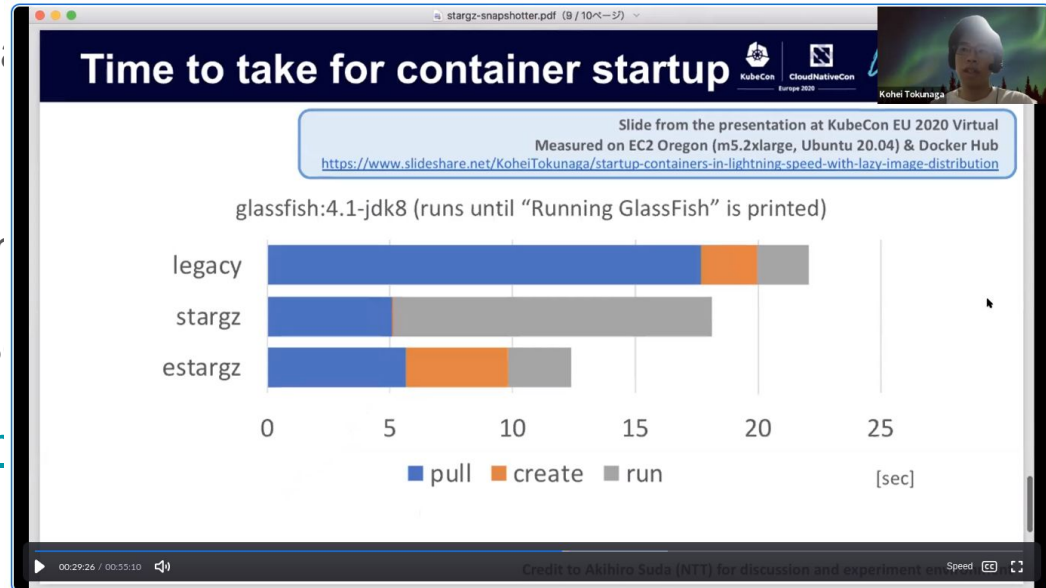
On average 6% of an image is

Much worse if images are very

Support added in containerd, li

Several group members testing at s

[Speeding up Analysis with Rem](#)



Topic Covered

Notebooks on Kubernetes (JupyterHub, Dask, GPU Integration, ...)

Enol Fernández: JupyterHub on EGI

Notebooks as a Service for EOSC, deployed across K8S EGI Cloud sites

Integration with EGI check-in (OIDC), NFS for persistent storage

Brian Bockelman, UNL

Notebooks and Dask Scheduling from Jupyter environments

Dask as a Service, submission in cluster or offloading to HTCondor

Discussion on challenges for auth/authz and credential conversion

Work on common recipe / recommendations: <https://github.com/cncf/research-user-group/issues/33>

Topic Covered

Notebooks on Kubernetes (JupyterHub, Dask, G

Enol Fernández: JupyterHub on EGI

Notebooks as a Service for EOSC, deployed as

Integration with EGI check-in (OIDC), NFS for persistent storage

Brian Bockelman, UNL

Notebooks and Dask Scheduling from Jupyter environ

Dask as a Service, submission in cluster or offloadin

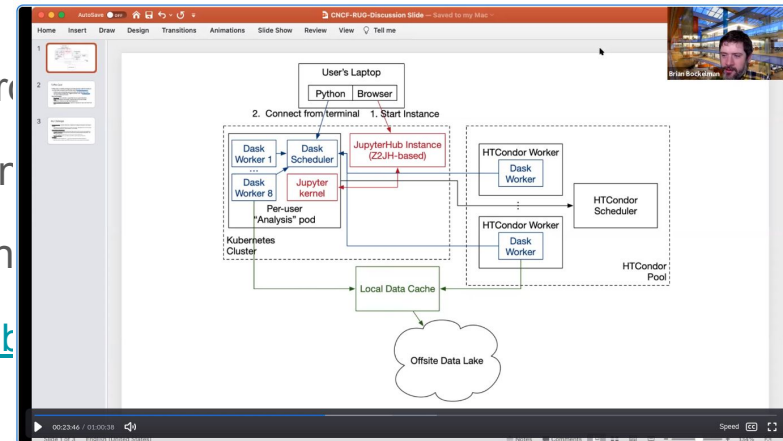
Discussion on challenges for auth/authz and creden

Work on common recipe / recommendations: <https://github.com>

The slide is titled "EGI Notebooks" and features the EGI logo. It describes a hosted JupyterHub solution for EOSC. The text includes:

- JupyterHub hosted in the EGI Cloud
- Offers Jupyter notebooks 'as Service': login with Check-in and start using
- Persistent storage
- Runs on k8s deployed on EGI Cloud providers
- Started mid-2017 with manual creation of all yml files
- Discovered zero-to-jupyterhub early 2018 and never looked back :)

On the right side of the slide, there is a screenshot of the "Notebooks" web interface, which shows a login button and some introductory text. The video player interface at the bottom shows a timestamp of 00:08:10 / 01:00:38.



Upcoming Session at Kubecon Europe 2021

Maintainers Track

Cloud-Native Computing For Research Users

Jamie Poole (G-Research)

Ricardo Rocha (CERN IT)

<https://sched.co/iE82>



Virtual Event
May 4 - May 7

[Learn More](#) and [Register to Attend](#)

The Sched app allows you to build your schedule but is not a substitute for your event registration. You **must be registered** for KubeCon + CloudNativeCon Europe 2021 Virtual to participate in the sessions. If you have not registered but would like to join us, please go to [the event registration page](#) to purchase a registration.

Please note: This schedule is automatically displayed in Central European Summer Time (UTC +2). To see the schedule in your preferred timezone, please select from the drop-down menu to the right, above "Filter by Date." The schedule is **subject to change**.



[Schedule](#) ▾ [Speakers](#)

Thursday, May 6 • 12:25 - 13:00

[Back To Schedule](#)

Cloud-Native Computing For Research Users - Jamie Poole, G-Research & Ricardo Rocha, CERN

Other Activities

Considering dedicated Rootless WG within the CNCF SIG-Runtime

Lots of progress for rootless support across the whole stack

Help Akihiro and others disseminating and validating all the different pieces

Reach out if you want to help

Some ongoing discussion on OpenTelemetry usage

Anyone has experiences to report? **Tracing**, Metrics, Logging?

Questions?