




KubeCon



CloudNativeCon

North America 2019





KubeCon



CloudNativeCon

North America 2019

CNCF

Research User Group

<https://github.com/cncf/research-user-group>

Bob Killen - co-chair
Klaus Ma - tech lead
Steve Quenette - co-chair

Poll: <https://pollev.com/bobkillen881>



Which industry or group do you represent?

Academia

Government

Non Profit

Private Sector

Other



KubeCon



CloudNativeCon

North America 2019

Why?

Poll: <https://pollev.com/bobkillen881>





KubeCon



CloudNativeCon

North America 2019

Research needs are changing.

Poll: <https://pollev.com/bobkillen881>



Why?



KubeCon



CloudNativeCon

North America 2019

- Increased use of containers...*everywhere*.
- Increasingly complex workflows.
- Adoption of data-streaming and in-flight processing.
- Greater use of interactive Science Gateways.
- Dependence on other more persistent services.

Poll: <https://pollev.com/bobkillen881>

Why form a user group?



KubeCon



CloudNativeCon

North America 2019

Most research oriented workloads are different from typical Enterprise workloads.

- Job/task focused (high rate of churn)
- Resource intensive
- Require more verbose scheduling (MPI)
- Multitenant environment
- Support for large or multiple clusters

Poll: <https://pollev.com/bobkillen881>

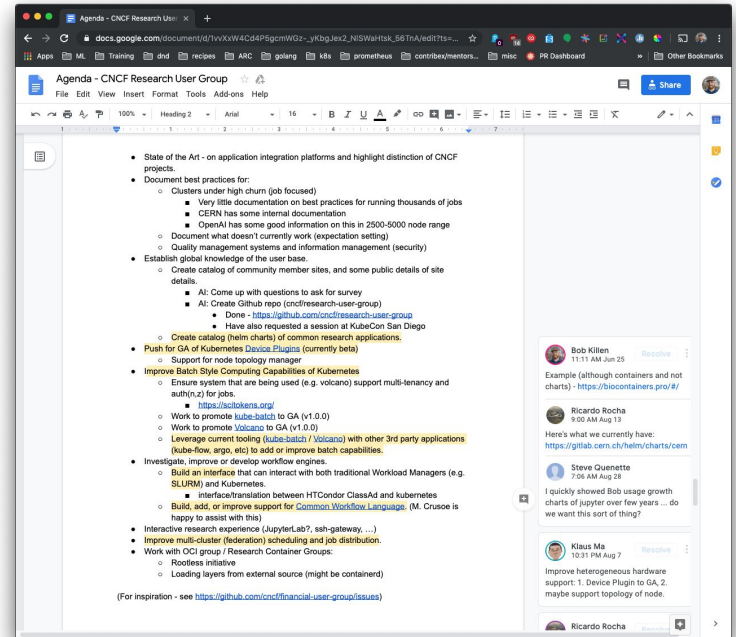


The CNCF Research User Group's purpose is to function as a focal point for the discussion and advancement of Research Computing using "Cloud Native" technologies. This includes enumerating current practices, identifying gaps, and directing effort to improve the Research Cloud Computing ecosystem.

Poll: <https://pollev.com/bobkillen881>

Common themes

- Lack of knowledge of “*what’s out there*”
- No best practices for large shared environments
- Base batch capabilities incomplete
- Multi-cluster/Federation job support lacking
- Multi-tenancy is problematic



Poll: <https://pollev.com/bobkillen881>

Current initiatives



KubeCon



CloudNativeCon

North America 2019

Research Institution Survey

Who is using Kubernetes for research?

What type of workloads are they running?

How have they deployed them?

Index of resources and useful links

“Awesome list” of research focused links

Best practices for running research clusters

Get “*current state*” of landscape

Discussions with various project maintainers

Where should effort be directed?

Poll: <https://pollev.com/bobkillen881>

Which industry or group do you represent?

Academia

Government

Non Profit

Private Sector

Other

Are you using Kubernetes to support research workloads?

Not currently.

Not currently, but
looking into it.

Yes, in development
or testing purposes.

Yes, in production.

What tools or frameworks are you using or would like to use?

What is the biggest pain point you've encountered with using Kubernetes as a research platform?

Where should we focus our efforts?

Top

Get Involved



KubeCon



CloudNativeCon

North America 2019

Mailing List: cncf-research-user-group@lists.cncf.io

GitHub Repo: <https://github.com/cncf/research-user-group>

Meetings:

- Agenda: <https://bit.ly/2WrXgy9>
- Zoom: <https://zoom.us/my/cncfenduser>
- [Second Wednesday of the Month @ 9:00 UTC / 5 AM ET / 2 AM PT](#)
- [Fourth Wednesday of the Month @ 15:00 UTC / 11 AM ET / 8 AM PT](#)



Related Upcoming Sessions



KubeCon



CloudNativeCon

North America 2019

Wednesday November 20th (Today)

- **2:25pm - 3:00pm** - [Intro: Scheduling SIG - Wei Huang, IBM & RaviSantosh Gudimetla, Red Hat](#)
- **3:20pm - 3:55pm** - [Kubeflow: Multi-Tenant, Self-Serve, Accelerated Platform for Practitioners - Kam Kasravi, Intel & Kunming Qu, Google](#)
- **3:20pm - 3:55pm** - [To Infinite Scale and Beyond: Operating Kubernetes Past the Steady State - Austin Lamon, Spotify & Jago Macleod, Google](#)
- **3:20pm - 3:55pm** - [Mitigating Noisy Neighbors: Advanced Container Resource Management - Alexander Kanevskiy, Intel](#)
- **4:25pm - 5:00pm** - [Batch Capability of Kubernetes Intro - Klaus Ma, Huawei](#)
- **5:20pm - 5:55pm** - [Deep Dive: Kubernetes Working Group for Multi-tenancy - Sanjeev Rampal, Cisco](#)

Related Upcoming Sessions



KubeCon



CloudNativeCon

North America 2019

Thursday November 21st (Tomorrow)

- **10:55am - 11:30am** - [Improving Performance of Deep Learning Workloads With Volcano - Ti Zhou, Baidu Inc & Da Ma, Huawei](#)
- **2:25pm - 3:00pm** - [Networking Optimizations for Multi-Node Deep Learning on Kubernetes - Rajat Chopra, NVIDIA & Erez Cohen, Mellanox](#)
- **2:25pm - 3:55pm** - [Tutorial: From Notebook to Kubeflow Pipelines: An End-to-End Data Science Workflow - Michelle Casbon, Google, Stefano Fioravanzo, Fondazione Bruno Kessler, & Ilias Katsakioris, Arrikto](#)
- **3:20pm - 3:55pm** - [Building a Medical AI with Kubernetes and Kubeflow - Jeremie Vallee, Babylon Health](#)
- **4:25pm - 5:00pm** - [GPU as a Service Over K8s: Drive Productivity and Increase Utilization - Yaron Haviv, Iguazio](#)
- **4:25pm - 5:00pm** - [RDMA Enabled Kubernetes for High Performance Computing - Jacob Anders, CSIRO & Feng Pan, Red Hat](#)
- **5:20pm - 5:55pm** - [Supercharge Kubeflow Performance on GPU Clusters - Meenakshi Kaushik & Neelima Mukiri, Cisco](#)

Related Sessions from Contributor Summit



KubeCon



CloudNativeCon

North America 2019

San Diego Kubernetes Contributor Summit:

- [Multi-tenancy in Kubernetes: Let's Talk - Tasha Drew](#)
- [How to Bring Batch into Kubernetes - Klaus Ma](#)
- [Present and Future of Hardware Topology Awareness in Kubelet - Connor Doyle](#)

Related (Past) Sessions



KubeCon



CloudNativeCon

North America 2019

- [Enabling Kubeflow with Enterprise-Grade Auth for On-Prem Deployments - Yannis Zarkadas, Arrikto & Krishna Durai, Cisco](#)
- [Managing Helm Deployments with Gitops at CERN - Ricardo Rocha, CERN](#)
- [Introducing KFServing: Serverless Model Serving on Kubernetes - Ellis Bigelow, Google & Dan Sun, Bloomberg](#)
- [Managing Apache Flink on Kubernetes - FlinkK8sOperator - Anand Swaminathan, Lyft](#)
- [Towards Continuous Computer Vision Model Improvement with Kubeflow - Derek Hao Hu & Yanjia Li, Snap Inc.](#)
- [Measuring and Optimizing Kubeflow Clusters at Lyft - Konstantin Gizdarski, Lyft & Richard Liu, Google](#)
- [Scaling Kubernetes to Thousands of Nodes Across Multiple Clusters, Calmly - Ben Hughes, Airbnb](#)
- [KubeFlow's Serverless Component: 10x Faster, a 1/10 of the Effort - Orit Nissan-Messing, Iguazio](#)
- [Advanced Model Inferencing Leveraging KNative, Istio and Kubeflow Serving - Animesh Singh, IBM & Clive Cox, Seldon](#)
- [Building and Managing a Centralized Kubeflow Platform at Spotify - Keshi Dai & Ryan Clough, Spotify](#)