

Summary of Kubecon 2022

Pre-GDB on Kubernetes, June 7th 2022

<https://indico.cern.ch/event/1096043/>

Ricardo Rocha, CERN

Co-Located Events

28 additional events

Day 0 and -1 of the conference (Monday and Tuesday)

Cloud Native SecurityCon (2 days)

Kubernetes AI Day

Cloud Native WASM Day

Kubernetes Data Workshop Hosted by Portworx

GitOpsCon Europe

KnativeCon Europe

Kubernetes Batch + HPC Day

...



**KUBERNETES
BATCH + HPC DAY
EUROPE**

17 MAY

VALENCIA, SPAIN

#K8SBATCH + #K8SHPC

Tuesday, May 17

07:30 CEST

Registration + Badge Pick-up

13:00 CEST

Opening + Welcome - Abdullah Gharaibeh & Ricardo Rocha, Kubernetes Batch + HPC Day
Program Committee Members

13:15 CEST

Keynote: High Performance Computing on Google Kubernetes Engine- Maciek Różacki, Google
Cloud

13:25 CEST

Kueue: A Kubernetes-native Job Queueing - Abdullah Gharaibeh, Google

13:55 CEST

Resource Orchestration of HPC on Kubernetes: Where We Are Now and the Journey Ahead! -
Swati Sehgal & Francesco Romani, Red Hat

14:25 CEST

Volcano – Cloud Native Batch System for AI, BigData and HPC - William(LeiBo) Wang, Huawei
Cloud Computing Co., Ltd

14:55 CEST

Get More Computing Power by Helping the OS Scheduler - Antti Kervinen, Intel & Alexander
Kanevskiy, Intel

15:25 CEST

How to Handle Fair Scheduling in a Private Academic K8s infrastructure - Lukas Hejtmanek, Masaryk University & Dalibor Klusacek, CESNET

15:35 CEST

Fast Data on-Ramp with Apache Pulsar on K8 - Timothy Spann, StreamNative

15:50 CEST

Apache YuniKorn A Kubernetes Scheduler Plugin for Batch Workloads - Wilfred Spiegelburg, Cloudera & Craig Condit, Cloudera

16:20 CEST

Efficient Deep Learning Training with Ludwig AutoML, Ray, and Nodeless Kubernetes - Anne Marie Holler, Elotl & Travis Addair, Predibase

16:45 CEST

Closing - Aldo Culquicondor, Kubernetes Batch + HPC Day Program Committee Member

17:00 CEST

CNCF-hosted Co-located Events Happy Hour

Co-Located Event: Batch+HPC Day

Available on Youtube:

<https://www.youtube.com/c/cloudnativefdn/videos>

Also available on the meeting play platform from the conference

Good discussion during the event

Also clarifications between Kubernetes WG Batch and the CNCF BSI

More on this later...

Main Conference



KubeCon



CloudNativeCon

Europe 2022

INITIAL DEMOGRAPHIC DATA



18,550

ATTENDEES

65% 1ST-TIME ATTENDEES

158

MEDIA +
ANALYSTS

12

CNCF-HOSTED
CO-LOCATED
EVENTS

13

SPONSOR-HOSTED
CO-LOCATED
EVENTS



45%

MEN



7%

WOMEN



<1%

NON-BINARY/
OTHER GENDERS

47% PREFER NOT TO ANSWER

ATTENDEES PER REGION





KubeCon



CloudNative

Europe 2022

INITIAL
DEMO
DATA

12 550

TOP 3 COUNTRIES REPRESENTED



USA, GERMANY, INDIA

TOP THREE JOB FUNCTIONS: DEVELOPER, DEVOPS/SRE/SYSADMIN, ARCHITECT

1187

CFP SUBMISSIONS

76

MAINTAINER TRACK SESSIONS

17

KEYNOTES

158

BREAKOUTS

CFP ACCEPTANCE RATE: 12%

END USER TALKS: 33%

GENDER MINORITY SPEAKERS: 18%

% OF GENDER MINORITY KEYNOTE SPEAKERS: 48%

IN-PERSON DATA

7,084

ATTENDEES

ATTENDEES PER REGION



66% 1ST-TIME ATTENDEES

111 MEDIA + ANALYSTS



49%

MEN



8%

WOMEN



<1%

NON-BINARY/
OTHER GENDERS

42% PREFERRED NOT TO ANSWER

TOP 3 COUNTRIES REPRESENTED



**USA, GERMANY,
UNITED KINGDOM**

TOP THREE JOB FUNCTIONS:

**DEVOPS/SRE/SYSADMIN,
DEVELOPER, ARCHITECT**

7,084

IN-PE

VIRTUAL DATA

11,466

ATTENDEES

ATTEND

ATTENDEES PER REGION



65% 1ST-TIME ATTENDEES

47 MEDIA + ANALYSTS



43%

MEN



6%

WOMEN



<1%

NON-BINARY/
OTHER GENDERS

50% PREFERRED NOT TO ANSWER

TOP 3 COUNTRIES REPRESENTED



USA, INDIA, GERMANY

TOP THREE JOB FUNCTIONS:
DEVOPS/SRE/SYSADMIN,
DEVELOPER, ARCHITECT

TOP 3 C



U
UN

TRACK	ATTENDANCE	PERCENTAGE
Maintainer Track	6,427	13.40%
101 Track	5,655	11.79%
App + Dev	4,805	10.02%
Security + Identity + Policy	4,703	9.81%
Customizing + Extending	4,229	8.82%
Observability	3,944	8.22%
Networking	3,103	6.47%
CI/CD	2,724	5.68%
Operations	2,724	5.68%
Service Mesh	2,206	4.60%
Community	1,595	3.33%
Business Value	1,393	2.90%
Student	1293	2.70%
Machine Learning + Data	1,041	2.17%
Performance	832	1.73%
Serverless	396	0.83%
Storage	363	0.76%
Research + Academia	300	0.63%
Runtimes	222	0.46%
TOTAL	47,955	

Main Conference

130+ CFP Sessions

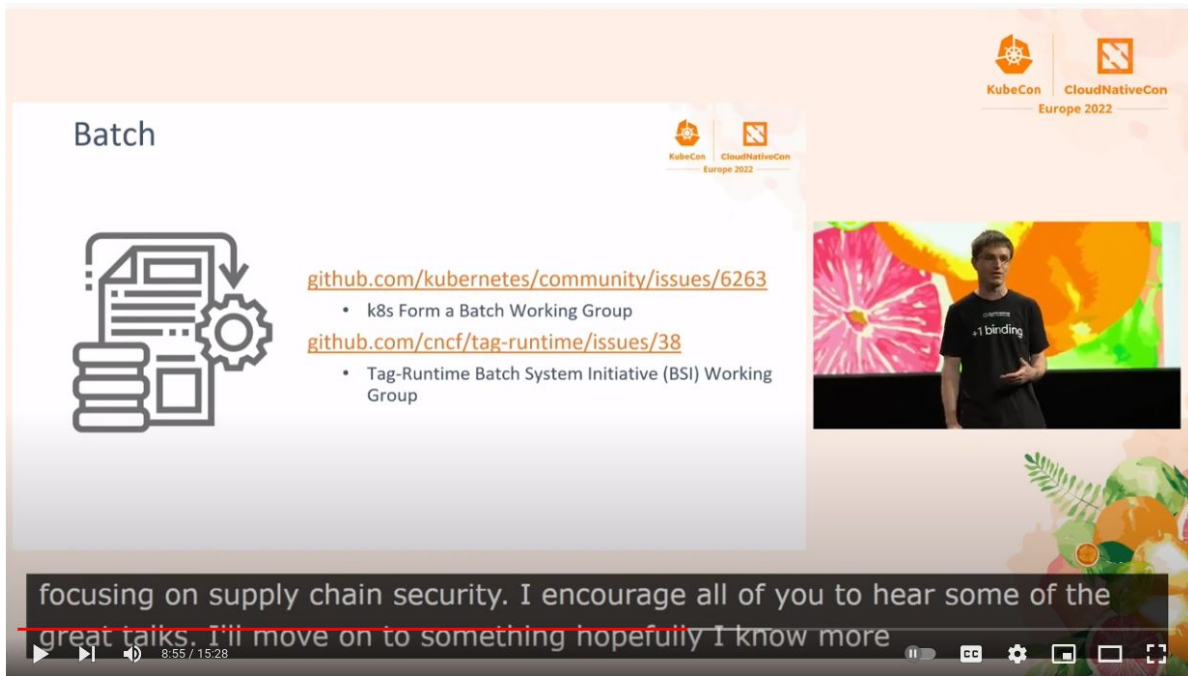
70+ Maintainer Track Sessions

Filter By Type

- 101 Track
- Application + Development
- Birds of a Feather
- Breaks
- Business Value
- CI/CD
- Co-Located Events
- Community
- Customizing + Extending Kubernetes
- Diversity + Equity + Inclusion
- Experiences
- Keynote Sessions
- Machine Learning + Data
- Maintainer Track
- Networking
- Observability
- Operations
- Performance
- Project Meeting
- Project Office Hours
- Registration
- Research + Academia
- Runtimes
- Security + Identity + Policy
- Serverless
- Service Mesh
- Solutions Showcase
- Storage
- Student
- Wellness Activities

<https://events.linuxfoundation.org/kubecon-cloudnativecon-europe/program/schedule/>

Keynotes: Pillars of Cloud Native Growth



Batch

github.com/kubernetes/community/issues/6263

- k8s Form a Batch Working Group

github.com/cncf/tag-runtime/issues/38

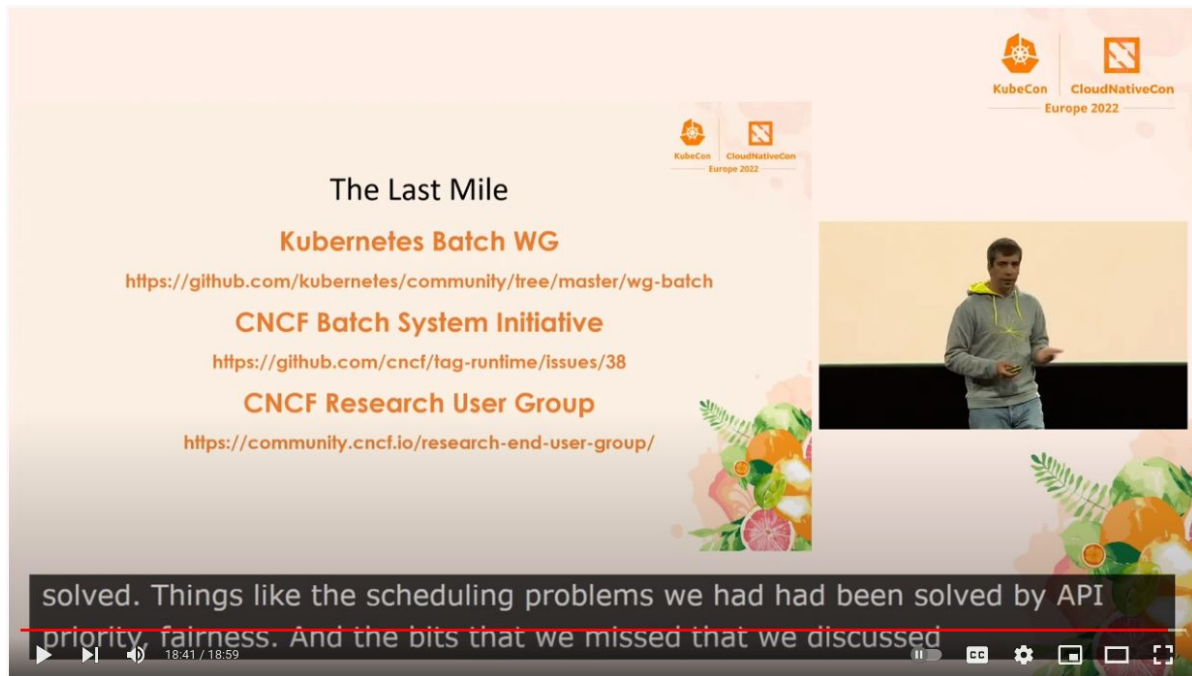
- Tag-Runtime Batch System Initiative (BSI) Working Group

focusing on supply chain security. I encourage all of you to hear some of the great talks. I'll move on to something hopefully I know more

8:55 / 15:28

<https://www.youtube.com/watch?v=YWZsXdAXFO8>

Keynotes: Building Bridges: Cloud Native and HPC



The Last Mile

Kubernetes Batch WG
<https://github.com/kubernetes/community/tree/master/wg-batch>

CNCF Batch System Initiative
<https://github.com/cncf/tag-runtime/issues/38>

CNCF Research User Group
<https://community.cncf.io/research-end-user-group/>

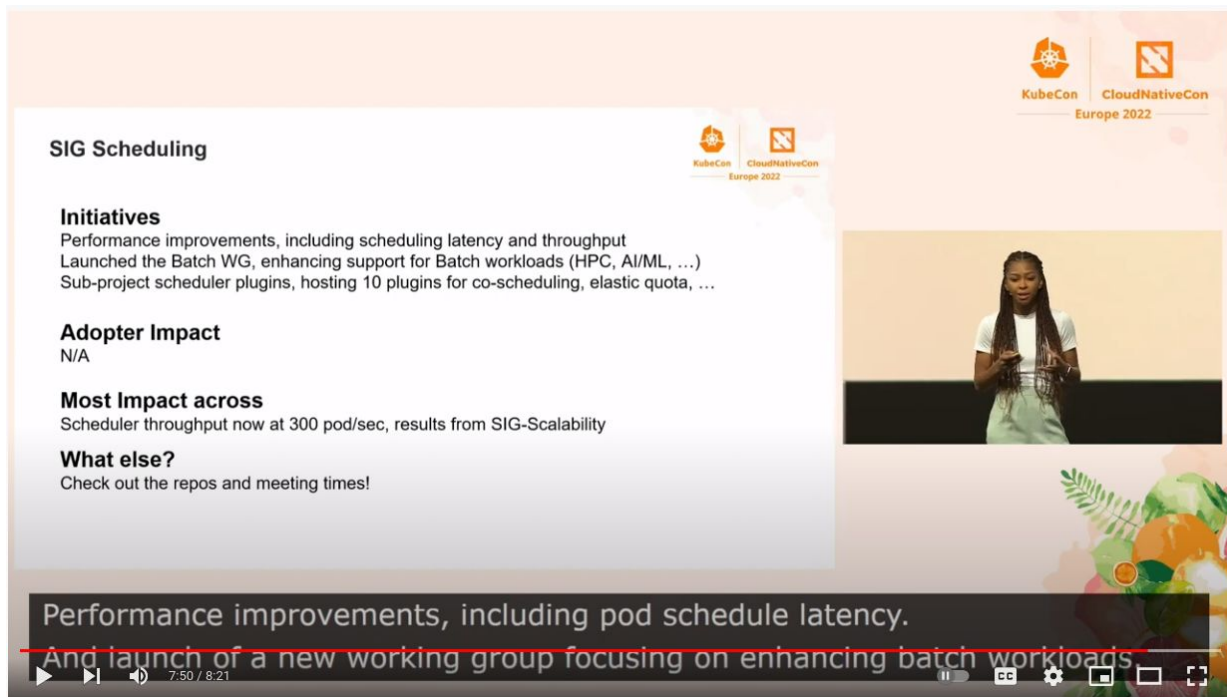
solved. Things like the scheduling problems we had had been solved by API
priority, fairness. And the bits that we missed that we discussed

18:41 / 18:59

The image shows a video player interface. The main content is a presentation slide with a light orange background and a decorative floral pattern at the bottom. The slide contains the title 'The Last Mile' and three bullet points, each with a title and a URL. A small inset video shows a man in a grey hoodie speaking. The video player controls at the bottom show the video is at 18:41 of 18:59. The video title is partially visible at the bottom of the slide.

https://www.youtube.com/watch?v=M95AQi1wA_s

Keynotes: Kubernetes Project Updates



SIG Scheduling

Initiatives
Performance improvements, including scheduling latency and throughput
Launched the Batch WG, enhancing support for Batch workloads (HPC, AI/ML, ...)
Sub-project scheduler plugins, hosting 10 plugins for co-scheduling, elastic quota, ...

Adopter Impact
N/A

Most Impact across
Scheduler throughput now at 300 pod/sec, results from SIG-Scalability

What else?
Check out the repos and meeting times!

Performance improvements, including pod schedule latency.
~~And launch of a new working group focusing on enhancing batch workloads~~

7:50 / 8:21

<https://www.youtube.com/watch?v=Ar4CgGHhpls>


Sessions: Kubernetes WG Batch

Very nice attendance, developers and end users

Aldo went through the ongoing work and plans

Sessions: Kubernetes WG Batch

Motivation



- k8s was originally built for serving applications and there is increasing support for stateful applications.
- Batch workloads can be run on k8s, but there are feature gaps:
 - Advanced completion and failure modes in Job API
 - Support for specialized devices and pinning
 - All-or-nothing pod scheduling
 - Job/workload queueing
- The status-quo was to support those features through CRDs
- There is fragmentation in the ecosystem:
 - Forked pod schedulers
 - Forked Job APIs
 - New CRIs

Sessions: Kubernetes WG Batch



Motivation

- k8s was originally built for service support for stateful applications
- Batch workloads can be run on
 - Advanced completion and failure
 - Support for specialized devices
 - All-or-nothing pod scheduling
 - Job/workload queueing
- The status-quo was to support
- There is fragmentation in the ecosystem
 - Forked pod schedulers
 - Forked Job APIs
 - New CRIs

WG Batch charter

- **Mission:** Discuss and enhance the support for Batch workloads in core Kubernetes. The goal is to unify the way users deploy batch workloads to improve portability and to simplify supportability for Kubernetes providers.
- **Stakeholders**
 - SIG Apps
 - SIG Autoscaling
 - SIG Node
 - SIG Scheduling

git.k8s.io/community/wg-batch


Sessions: Kubernetes WG Batch

Motivati

- k8s was or support for
- Batch work
 - Advanc
 - Support
 - All-or-n
 - Job/wo
- The status
- There is fra
 - Forked
 - Forked
 - New CF

WG Batch workstreams

- Advance the Job API to support a wider range of workloads (static partitioning, MPI, ML, AI).
- Job management, queueing, provisioning, scheduling and autoscaling.
- Runtime and scheduling support for specialized hardware (accelerators, NUMA, RDMA, etc.)



Batch
fy the way
ility and to

community/wg-batch

Session: Kubernetes as Substrate of ATLAS Compute

Kubernetes for Batch processing in the ATLAS experiment

Scaling public cloud deployments up to 100k cores

Key points

Managed, scalable solutions out of the box

Better elasticity for *exotic resources*, improve analysis offering

Next Steps

Zero-to-grid-site with GitOps and Helm next

Similar deployments in HPC environments

<https://www.youtube.com/watch?v=6dA43w08wLI>

Session: Kubernetes as Substrate of ATLAS Compute

Kubernetes for Batch processes

Scaling public cloud deployment

Key points

Managed, scalable solution

Better elasticity for exotic workloads

Next Steps

Zero-to-grid-site with Google Cloud

Similar deployments in other environments

The video player shows a presentation slide with the following content:

- Elastic cloud scale out**
- Running vCPUs on regional GKE cluster in eu-west-1**
- Processing type** (simul, gile, evgen, deriv)
- 1% failure on Spot VMs!**
- Complexity metrics:**
 - $O(10^8)$ events
 - $O(10^5)$ vCPUs
 - $O(10^4)$ Pods
 - $O(10^3)$ Nodes
 - $O(1)$ day
 - $O(1)$ Harvester instance
 - <1 Engineer
- Globally running vCPUs on 30 April 2022**
- Processing type** (Vega, GDSILE, MERZ, CERN-T0, CERN-P1)
- Image credits: ICT Network News**

The slide also features logos for KubeCon and CloudNativeCon Europe 2022, a photo of two people at a presentation, and a decorative graphic of fruits in the bottom right corner. The video player interface at the bottom shows a progress bar at 18:08 / 33:52 and various control icons.

<https://www.youtube.com/watch?v=6dA43w08wLI>

Integration of heterogeneous architectures

- Cloud queues backed by resources not commonly available on prem
- Straightforward to integrate different architectures, e.g. ARM, GPU
- Multi-arch Docker images doing the heavy lifting


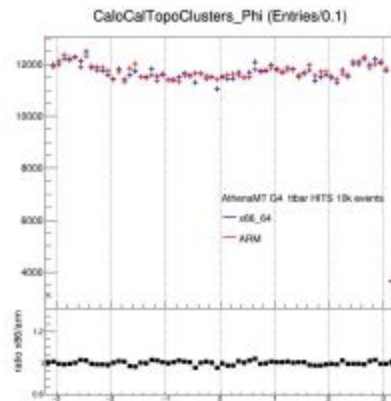
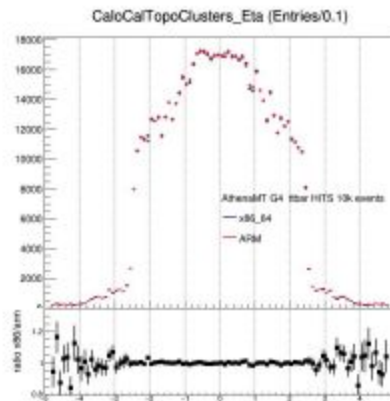


Image credits: Communications of the ACM, February 2019

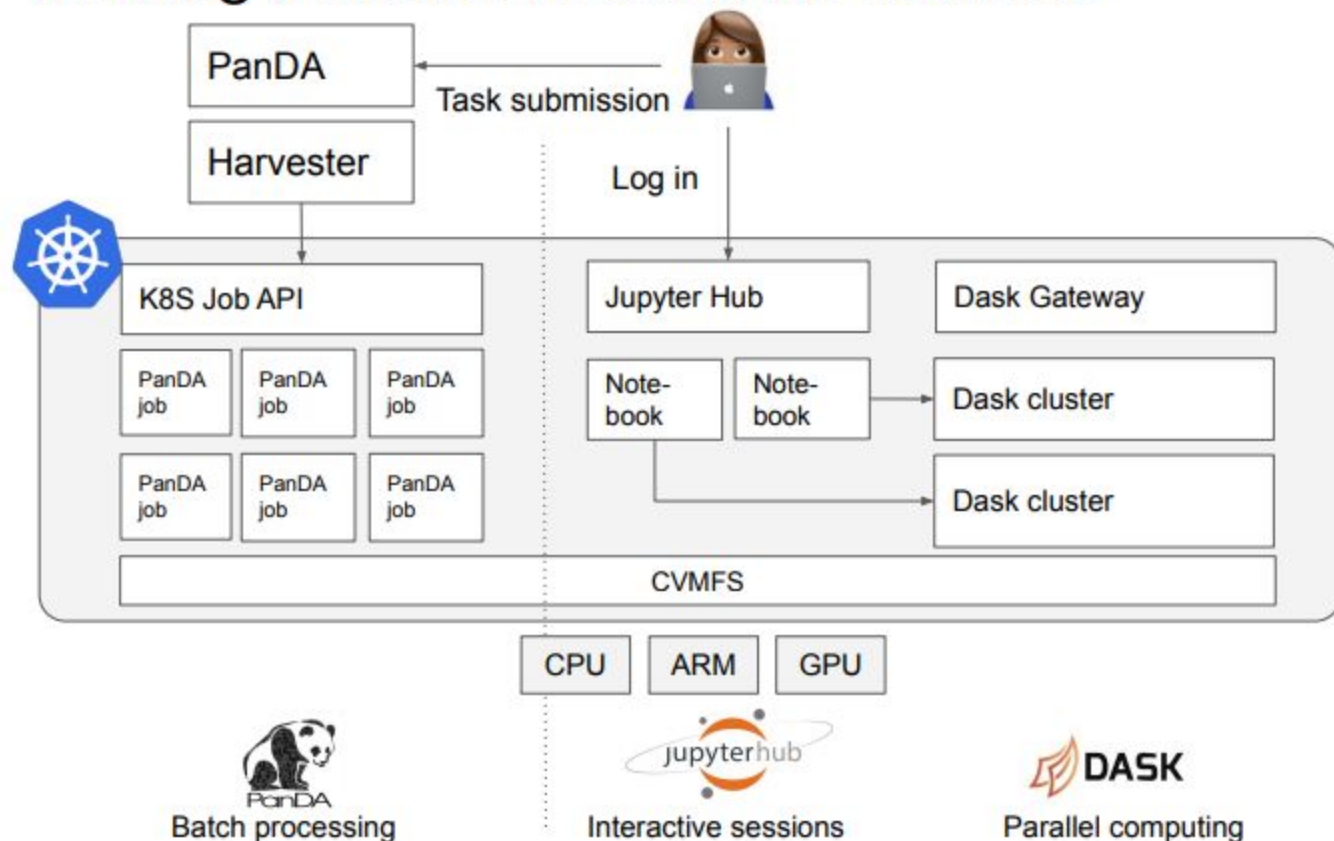


Image credits: NVIDIA GTC May 2020 Keynote



First ATLAS simulation task on ARM processors. Currently under physics validation. Generated on Amazon EKS cluster backed by Graviton 2 nodes.

Adding interactive analysis facilities



Other Sessions

[Improving GPU Utilization using Kubernetes - Maulin Patel & Pradeep Venkatachalam, Google](#)

[How to Migrate 700 Kubernetes Clusters to Cluster API with Zero Downtime - Tobias Giese & Sean Schneeweiss, Mercedes-Benz Tech Innovation](#)

[Logs Told Us It Was DNS, It Felt Like DNS, It Had To Be DNS, It Wasn't DNS - Laurent Bernaille & Elijah Andrews, Datadog](#)

[Seeing is Believing: Debugging with Ephemeral Containers - Aaron Alpar, Kasten](#)

[It's All for the Users. More Durable, Secure, and Pluggable. KubeVirt v0.53 - Alice Frosi, Red Hat](#)

[Autoscaling Kubernetes Deployments: A \(Mostly\) Practical Guide - Natalie Serrino, New Relic \(Pixie team\)](#)

[Reproducing Production Issues in your CI Pipeline Using eBPF - Matthew LeRay, Speedscale & Omid Azizi, New Relic](#)



KubeCon



CloudNativeCon

North America 2022

OCTOBER 24 – 28

DETROIT, MICHIGAN

<https://events.linuxfoundation.org/kubecon-cloudnativecon-north-america/>