



KubeCon




CloudNativeCon

Europe 2026

#KubeCon #CloudNativeCon

The Future of Kubernetes Scalability

Challenges of the GigaWatt Computing Power of the AI Era





KubeCon




CloudNativeCon

Europe 2026



Artur Rodrigues

Member of Technical Staff, Anthropic




Maciek Różacki


Product Manager, Google Cloud

Disclaimer: Views shared in this talk are our own

Node count is not enough



Scale is three axes, not one



Training inverts the workload shape

Microservices
what K8s was tuned for

Stateless, loosely coupled

Horizontally scalable

Rolling updates, blue-green

Placement-agnostic

~5K-node benchmark target

AI training
what we actually run

Gang-scheduled, all-or-nothing


Pinned to TPU slice / rack

Weeks-long, restart-averse

Fabric- and topology-aware

Outside the tuned envelope

The Gigawatt stakes



"Anthropic invests \$50 billion in American AI infrastructure"

"Google's Data Center Buildout Could Top \$1 Trillion"

"Oracle and OpenAI Are Rocking AI World with \$500-Billion Stargate Startup"

Control plane can't keep up → accelerators idle → \$M / day

AI labs hit these ceilings first — not only.

Three axes of heterogeneity, one control plane

Multi-cloud

different storage → different failure modes

GKE

Spanner-backed, dozens of API replicas

EKS

offloaded consensus, partitioned key-space

Multi-architecture

each family: own interconnect, own placement rules

TPU

GPU

Trainium

Multi-workload

each with a different mutation-rate profile

Training

Inference

RL

CI bursts

Research

All on Kubernetes

chosen for the ecosystem, not the workload fit

Neither could do the other's work

Anthropic

Custom scheduler

Custom workload CRD

CoreDNS plugin

Read offload (3 gens)

workload-specific, above the API

Google

Spanner-backed storage


Upstream fixes

130K-node demonstration


Joint incident diagnosis

below the API, upstream


Hypothesis: the staffing bar dropped



Faster chips do not reduce control plane load




Workload mix drives the load.




Four futures for cluster scale

Megacluster




One scheduler sees all.
Blast radius is all.

Multi-cluster




Failures stay contained.
Router is its own system.

Bespoke & customise




Swap, extend, diverge.
Feed it back upstream.

Self-managed



Hardware you own.
Ops burden. Tractable?

Bespoke first, upstream later



Extension points where the workload is new. Receptive upstream when the pattern holds.

Areas that would immediately help at scale

Controllers

Work-queue partitioning
reconcile throughput grows with replicas, not capped at one leader

API Server

Reconnection backoff
server paces recovery, not a thundering herd


APF high-churn profile
heartbeats don't lose to bulk LIST

Workloads

WAS companions
workload-level preemption (KEP-5710)

JobSet next steps
fuse admission + scheduling

Bet on modularity, not on one architecture



When the mix shifts, the platform bends — no rewrite.



KubeCon



CloudNativeCon

Europe 2026

Thank you



Putting something behind it

Thank you for
everything you ship.
Claude Max is on us.

Open-source maintainers and contributors keep the ecosystem running. The Claude for Open Source Program is our way of saying thank you for all your hard work, with 6 months of free Claude Max 20x. Apply now.

ANNOUNCEMENT

Linux Foundation Announces \$12.5 Million in Grant Funding from Leading Organizations to Advance Open Source Security



Anthropic, Amazon Web Services (AWS), GitHub, Google, Google DeepMind, Microsoft, and OpenAI Join Forces with the Foundation to Invest in Sustainable Security Solutions for the Open Source Ecosystem

SAN FRANCISCO – March 17, 2026 – [The Linux Foundation](#), the nonprofit organization enabling mass innovation through open source, today announced \$12.5 million in total grants from [Anthropic](#), [AWS](#), [GitHub](#), [Google](#), [Google DeepMind](#), [Microsoft](#), and [OpenAI](#) to strengthen the security of the open source software ecosystem. The funding will be managed by [Alpha-Omega](#) and the [Open Source Security Foundation \(OpenSSF\)](#), trusted security initiatives within the Linux Foundation, to develop long-term, sustainable security solutions that support open source communities worldwide.