KubeCon CloudNativeCon North America 2026

#KubeCon #CloudNativeCon

The Promises We Keep at Scale

A Contract-Driven Guide to Platform SLOs

Vasudev Bongale, Swati Shridhar Nair

LinkedIn



About us



Vasudev



Swati

💼 Compute Infrastructure · LinkedIn

☸️ 3rd KubeCon	☸️ 1st KubeCon
🤖 AI tools	🤖 Cloud-native & AI
📷 Digital photography	✈️ Travel & Food

Our platform at scale

The Kubernetes platform behind LinkedIn's applications

Running stateless and stateful workloads across many clusters and regions

2M+
Containers



500K+
Hosts



3,000+
Applications



60K+
Weekly deployments



Our approach



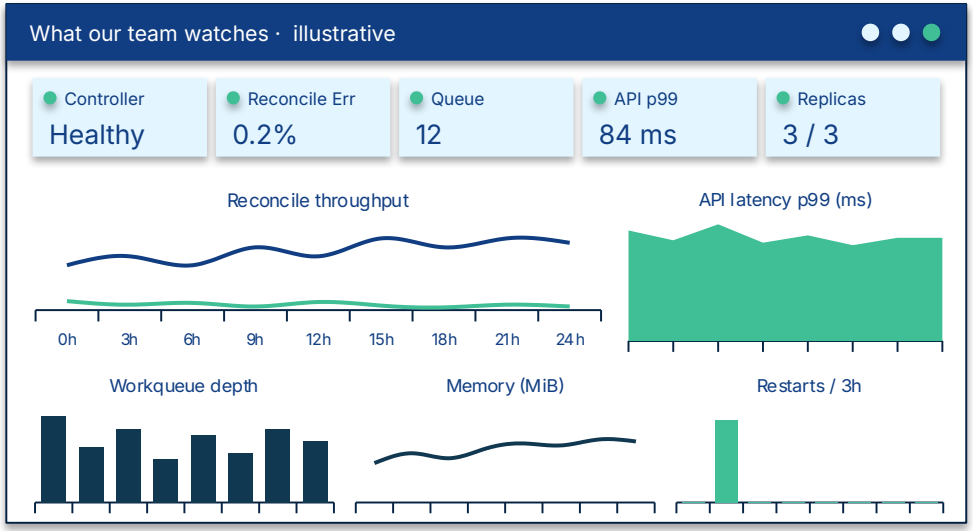
Understand what users care about

Define the promises

Measure what the platform owns

Test and evolve

What users care about



They expect one platform, even though many teams build it

Users experience the platform only through their workloads

Four promises

A promise is what users can expect from the platform



"Your rollout won't fail because of the platform."

Rollout success rate

"The platform won't keep your pods waiting."

Scheduling + preparation

"The platform won't compromise your workload's availability."

Workload readiness

"Your deleted pods are gone on time."

Overdue termination

Same status. Different responsibility.

KubeCon CloudNativeCon North America 2026

```
$ kubectl get lideployments                    # simplified output

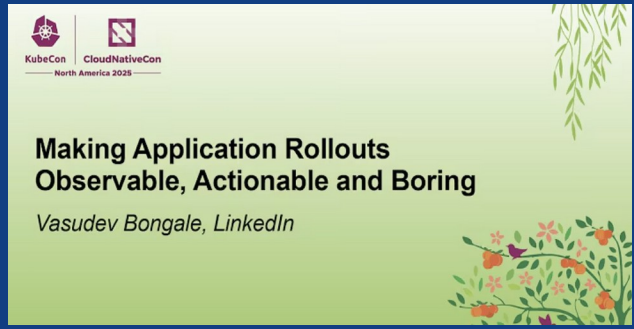
NAME        REPLICAS  READY     CATEGORY  REASON
checkout    1/3       False     App       AppFailedToStart
search      1/3       False     Infra     InitializationFailure
inventory   3/3       Unknown   -         StatusUnknown
feed        3/3       True      -         -
```

▶ Making Application Rollouts Observable, Actionable and Boring, KubeCon NA 2025

KubeCon CloudNativeCon
North America 2026



▶ Watch the talk



▶ Watch the talk

Thank you!

Q & A

Scan to share your feedback →

The Promises We Keep at Scale: A Contract-Driven Guide to Platform SLOs