



KubeCon



CloudNativeCon

Europe 2026

#KubeCon #CloudNativeCon

Building the Next Generation of Kubernetes Workloads Together

Janet Kuo, Google
Maciej Szulik, Defense Unicorns



Focuses on the Application User Experience

Owns the Core Workload APIs

Develops Application-Level Tools and Standards

K8s-native standards for AI

About us



KubeCon



CloudNativeCon

Europe 2026



Janet Kuo, Google

[@janetkuo](#)



Maciej Szulik, Defense Unicorns

[@soltys](#)



Kenneth Owens, Capital One

[@kow3ns](#)

Beta features



KubeCon



CloudNativeCon

Europe 2026

maxUnavailable for StatefulSets - KEP-961

(disabled by default)

Enables concurrent updates

`.spec.updateStrategy.rollingUpdate.maxUnavailable`

Improves update speed and flexibility

Maintains controller availability

*Effective for the `OrderedReady`
`podManagementPolicy`*



*number
or
% value*



Shout out to [@helayoty](#)



Terminating Pods in Deployments - [KEP-3973](#)

New `.status.terminatingReplicas`

in **Deployment** and **ReplicaSet**

Improve visibility into controller actions & resource usage

Part of Deployment Pod Replacement Policy [KEP-5882](#)



Shout out to [@atiratree](#)





Mutable Pod Resources for Suspended Jobs - [KEP-5440](#)


Allows updating container resource specifications

CPU, memory, GPU, and extended resources

Enable external controllers optimize resource usage

Backoff Limit Per Index For Indexed Jobs - [KEP-3850](#)

Per-index restart backoff




```
.spec.backoffLimitPerIndex  
.spec.maxFailedIndexes
```

Improved resilience for Indexed Jobs

Job success/completion policy - [KEP-3998](#)

Flexible success definition

`.spec.successPolicy`



```
1. succeeded indexes  
2. succeeded count
```

Enables advanced workloads

Pod Replacement Policy for Jobs - KEP-3939

Control over pod replacement `.spec.podReplacementPolicy`

- **TerminatingOrFailed**

replace immediately = faster execution

more resources consumed

- **Failed**


Wait for full terminating = slower execution

stable resource consumption

Job API managed-by mechanism - KEP-4368

Introduces `.spec.managedBy`

Enables external workload management



*Kueue uses it to
orchestrate multi-cluster
job dispatching -
MultiKueue*



Shout out to @mimowo and @kannon92



Add Recreate Update Strategy to StatefulSet - [KEP-3541](#)

Expand `.spec.updateStrategy.type` with Recreate

Remove all pods and recreate based on
`.spec.podManagementPolicy`



Shout out to [@helayoty](#) and [@galal-hussein](#)



Consider terminating Pods in Deployments - KEP-3973

Introduces new `.spec.podReplacementPolicy`

- **TerminationStarted**

replace immediately = faster rollout

- **TerminationComplete**

wait for full termination = control resource consumption



Shout out to @atiratree



Agent Sandbox





KubeCon



CloudNativeCon

Europe 2026



Agent Sandbox enables easy management of isolated, stateful, singleton workloads, ideal for use cases like AI agent runtimes.

<https://github.com/kubernetes-sigs/agent-sandbox>

MCP Lifecycle Operator



KubeCon



CloudNativeCon

Europe 2026



MCP Lifecycle Operator is a Kubernetes operator that provides a declarative API to deploy, manage, and safely roll out MCP Servers, handling their full lifecycle with production-grade automation and ecosystem integrations.

(It's just a month old!)

<https://github.com/kubernetes-sigs/mcp-lifecycle-operator>

Biweekly meeting

Monday at 06:00pm CET / 12:00pm EDT / 09:00am PDT

Slack

[#sig-apps](#)

Mailing List

sig-apps@kubernetes.io

