

CONFLUENT PLATFORM

# Confluent for Kubernetes

Confluent for Kubernetes allows you to build your own private cloud Apache Kafka® service using a complete, declarative API to deploy and operate Confluent Platform. It leverages Kubernetes to enhance the platform’s elasticity, ease of operations, and resiliency for enterprises operating at any scale.



**Achieve faster time to value with a cloud-native experience**



**Reduce costs with a complete, ready-to-use platform**



**Deploy to any private cloud solution**

## Why Confluent for Kubernetes?

Most companies are on a journey to the cloud, with legacy systems either being replaced or slowly migrated. For these companies, [Confluent Cloud](#) offers a fully managed, cloud-native service for building data in motion applications in cloud environments. However, some companies need to keep certain workloads on-premises due to policies and regulations. Other companies have started their transition, but cannot make a sudden shift. These businesses need a platform for data in motion that brings the same benefits and operational experience of the cloud to their on-premises environments.

For enterprises that want to build their own private cloud Kafka service for data in motion workloads that remain on-premises, [Confluent for Kubernetes](#) provides the same cloud-native attributes of Confluent Cloud as a complete, self-managed platform for your private infrastructure. Cloud-native attributes like elasticity, ease of operations, resiliency, and more are key to running modern data platforms and are now available wherever an organization’s applications and data reside, making it easier and faster to harness the full value of data in motion.

## Build your own private cloud Kafka service



**Leverage a complete, declarative API for deploying and operating Confluent in production.**



**Manage topics and RBAC policies through an infrastructure-as-code approach.**



**Integrate with the cloud-native ecosystem for security, reliability, and DevOps automation.**

# Solution

## Achieve faster time-to-value with a cloud-native experience

---

### Elastic scalability

Quickly adjust to changing business conditions and manage costs, scaling up or down using a single command. With Confluent for Kubernetes, you can operate Confluent Platform at any scale, from dev teams just getting started to central platform teams running an enterprise-wide shared service, to meet any data in motion workload.

### DevOps automation

Configure, deploy, and manage platform components and resources using an infrastructure-as-code approach and enterprise-recommended configurations. This accelerates time-to-value and reduces the complexity of managing the platform, freeing up your DevOps team to focus on value-added tasks.

## Reduce costs with a complete, ready-to-use platform

---

### Robust security

Automatically deploy security features with proper configurations and integrate with the cloud-native security ecosystem, such as Hashicorp Vault, Cert Manager, and Kubernetes Secrets.

### Streamlined recovery

Automatically recover after a failure in your brokers or underlying infrastructure with automated fault tolerance. With Confluent for Kubernetes, you get automated rack awareness to spread partition replicas across different racks, improving the availability and limiting the risk of data loss.

## Automate the deployment of Confluent Platform as a cloud-native system on Kubernetes



Confluent for Kubernetes. Cloud-native data in motion platform for your private infrastructure built by the original creators of Apache Kafka. For more information, please visit [confluent.io](https://confluent.io). To contact us, visit [confluent.io/contact](https://confluent.io/contact). For detailed product specifications, please refer to our [documentation](#).

## Complete solution

Leverage a complete, ready-to-use platform to implement data in motion use cases end-to-end. Confluent for Kubernetes makes it easy to programmatically deploy clusters with security and disaster recovery features, connectors for other data systems, and SQL-based stream processing.

## Deploy to any private cloud solution

---

### Broad ecosystem of Kubernetes distributions

Confluent for Kubernetes provides you with a cloud-native, consistent operational experience for market-leading Kubernetes distributions:

- Red Hat OpenShift
- VMWare Tanzu Kubernetes Grid (TKG)
- Rancher
- Google GKE on Anthos
- Amazon EKS on AWS Outposts
- Azure AKS on Azure Stack
- Plus any Kubernetes distribution or managed service meeting the Cloud Native Computing Foundation's (CNCF) conformance standards.

### Proven, battle-tested architecture

Implement standardized, validated architecture designed by the world's foremost Kafka experts, allowing you to operate at scale with confidence. Confluent for Kubernetes operationalizes years of experience running Kafka on Kubernetes at massive scale as part of our managed service, Confluent Cloud.