



Datasheet

Ondat for Kafka

Ondat for Kafka

Developed by LinkedIn to handle the immense throughput of their platform, **Kafka** has become the streaming platform of choice for developers who require low latency, high-scale, and persistent event streaming for their applications.

Designed to be highly resilient, Kafka can still benefit enormously from being hosted within Kubernetes. Administrators can utilize the Kubernetes common control plane to create a robust deployment and management tool. Kafka can also take advantage of Kubernetes pod and node scaling features to enhance and complement its own powerful horizontal scaling features.

However, block storage offered by cloud providers does not deliver on these additional advantages, with a lack of both predictable performance and adequate scheduling features.



Ondat Solution

Ondat allows Kafka brokers to take full advantage of Kubernetes. Administrators gain performant, predictable, and secure storage regardless of the cloud provider.

Ondat unlocks the ability to host high-scale Kafka clusters on Kubernetes, without compromise.

Key Features



Predictable disk latency

Business-critical applications need predictable storage performance. And as Kafka is used at scale in more time-sensitive application environments, deterministic performance for both throughput and latency becomes essential. Delivering this to instances running on Kubernetes is a real challenge. Kafka administrators are typically left making difficult choices between reliable storage performance and the resilience and efficiency that Kubernetes delivers. Ondat delivers both.

Hosting demanding Kafka environments on Kubernetes can be highly problematic unless administrators use their platform's local storage. Time-sensitive applications, such as payment gateways or high volume e-commerce, have little or no tolerance against the inherent fluctuations and erratic performance delivered by cloud or network storage. Even where providers do offer guaranteed performance for cloud storage, it comes with a high price tag and still fails to compete on performance against local storage mediums such as NVMe.

At the same time, where your platform's local Persistent Volumes can deliver reliable, predictable performance, they tie applications to a single node, sacrificing the flexibility and power that Kubernetes container scheduling delivers.

Ondat Persistent Volumes are perfectly suited to the architecture of Kafka instances running on Kubernetes.

They provide high-performance local storage with deterministic latency, making the data instantly available to any node in the cluster. Ondat allows operators to pick nodes with high-performance local storage to back Kubernetes Persistent Volumes. This provides the added surety that the data is replicated at the block level to other nodes, adding performance and additional resiliency to Kafka brokers.



Faster broker recovery

Resiliency is a key tenet of Kafka. Kafka brokers are designed to be distributed and replicate data across clusters. In the event of failure, this ensures that a replacement node can be restored using the distributed data.

But when using cloud storage within Kubernetes, this resiliency can act against Kafka. The time needed to reschedule Kubernetes Persistent Volumes is significantly increased. These delays can have severe implications in high-scale Kafka installations, bleeding-off application performance on computationally expensive replication events as volumes are rescheduled. **Ondat removes these limitations and allows Kafka to make the best use of Kubernetes scheduling without performance bottlenecks.**

When a Kafka broker is rescheduled on an Ondat enabled Kubernetes cluster, any pauses waiting for Persistent Volumes are negligible. By ensuring that persistent broker data is constantly replicated at the block level, Kafka nodes can be rescheduled with minimal broker replication required to bring the node into service. **Ondat delivers fewer delays, less performance impact, and increased resiliency.**



Powerful encryption at rest

Despite powerful encryption in transit, Kafka does not have encryption at rest. This lack of encryption can be a critical issue where Kafka is used in secure environments.

Cryptography can be added at the application layer to ensure data is encrypted at all points of its lifecycle. However, this is potentially complex and relies on application teams to add it to every application. And where encryption is in place, defense-in-depth strategies recommend additional layers of encryption in highly secure environments. By default, cloud storage may not enable encryption or even offer it as a feature, leaving a potential gap in the data security lifecycle.

Ondat allows administrators and application developers to enforce disk encryption by default within Kubernetes using annotations. Furthermore, these annotations can be set on the Persistent Volume StorageClass, ensuring that encryption is always enabled, regardless of the underlying storage mechanism. Ondat uses powerful AES encryption to ensure Kafka data is secured both at rest and in transit and further enhances existing security measures.

All prices are calculated directly from the AWS cost calculator and correct at the time of publishing.

Free Your Data

Kubernetes has transformed the way modern applications are deployed. Using Kubernetes, teams can fully utilize techniques such as GitOps and CI/CD to deploy software in a truly agile manner. And while Kubernetes brings significant advances in scheduled workloads by default, Kubernetes persistent volumes are not ready for high-scale production workloads. Inconsistent features, expensive cloud storage, and poor performance can challenge operators of Kubernetes to run mission-critical persistent applications at scale.

Ondat brings peace of mind by offering production-ready Kube-Native persistent volumes. Ondat allows you to run your most mission-critical persistent applications within any Kubernetes cluster, either in the cloud or on-premises, without compromising reliability or performance. And by offering operators a single cohesive persistence layer for Kubernetes, their business data can finally move freely between platforms.

Key features at a glance



Performant

- Deterministic performance
- In-memory caching
- Data locality



Reliable

- Highly resilient
- Real-time data replication
- Block-level checksums



Manageable

- Kube-Native
- Configurable cluster architecture
- Native API



Secure

- Secure by default
- Encrypted at rest
- Kubernetes integrated



Deployable

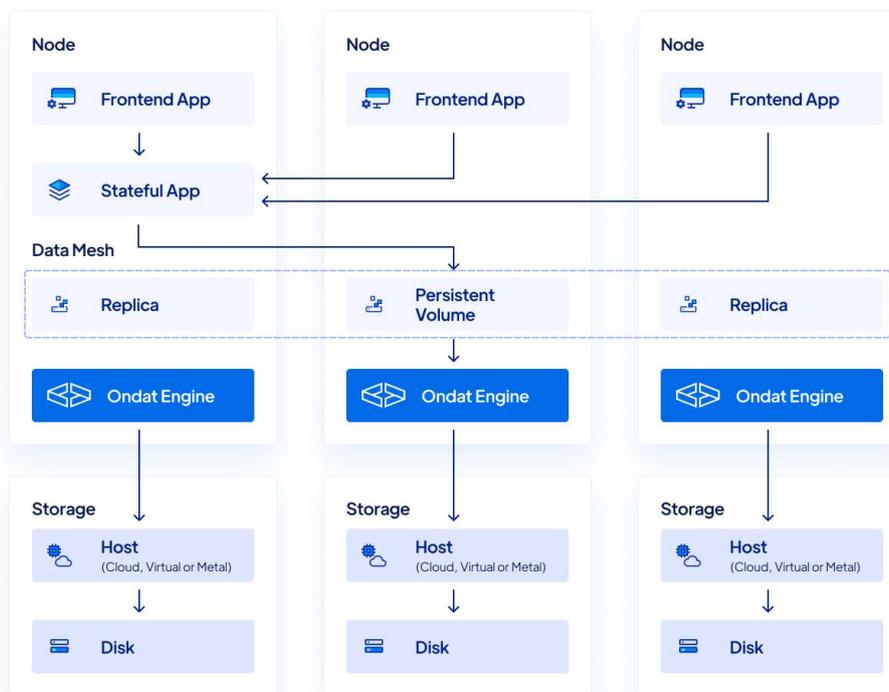
- Runs on any infrastructure
- Easy to install



Observable

- Fully observable
- Granular logging

How Ondat Works Kubernetes cluster



accenture

DHL

T-Systems

LLOYDS BANK

IMT

CIVO

INTERCONNECT

www.ondat.io

ondat™