ondat™

Datasheet

Ondat for PostgreSQL

PostgreSQL is a powerful Open Source RDBMS chosen by companies worldwide for its enterprise features, performance, and reliability.

Increasingly, operators are choosing to move their PostgreSQL workloads into Kubernetes to take advantage of the power, tooling, and mindshare of the Kubernetes platform. In addition, it allows organizations to offer self-service managed databases to their development teams without locking themselves into a specific cloud-managed database.

But while Kubernetes offers excellent scheduling capabilities for transient workloads, the implementation of persistent volumes varies massively across different providers, with each bringing challenges. In contrast, on-premises deployments of Kubernetes only offer local persistent volumes or HostPath volumes. Both are unsuitable for running production loads at scale.

Ondat allows operators to run PostgreSQL workloads safely, performantly, and at scale. With Ondat, you can be sure that PostgreSQL can make full use of fast local storage while still enjoying the flexibility and power that Kubernetes offers.

Key Features

Fast, predictable query times

Business-critical applications demand high-performance databases with predictable query times. Cloud or network-based storage can lack the performance and consistency for demanding PostgreSQL use cases, leading to wildly varying response times for dependent apps. This need for a high-performance storage fabric can relegate PostgreSQL to use only on dedicated hosts, missing the operational power and flexibility of Kubernetes.

Ondat allows DBAs to use their existing high-performance storage, be instead local NVMe or SAN while allowing PostgreSQL to run seamlessly within a Kubernetes cluster. Furthermore, using its innovative clustering and data replication, Ondat enables you to choose the most suitable storage medium for your workload, without compromising performance.

Increased data safety

There is little that is more important than the data applications process and store. And while PostgreSQL has robust mechanisms to ensure data is consistent, it can be seriously hampered by poor storage technology. Unfortunately, when choosing between speed and safety, some products default too far on the side of speed. These defaults can lead to data loss or inconsistency.

Ondat combines speed and safety. Using powerful in-memory and write persisted caching, Ondat offers high performance with guaranteed read and writes. Within the Ondat cluster, data is 'striped' across nodes, allowing highly performant local disks to be used without any single point of failure. And by using block-level hashing, Ondat ensures that data stored within Kubernetes Persistent Volumes is guaranteed to be consistent. Ondat gives PostgreSQL DBAs the tooling to ensure data is resilient, performant, and trusted.

Higher resiliency

PostgreSQL is a robust and reliable database, and Kubernetes can bolster resiliency even further. Kubernetes ensures that the PostgreSQL engine is self-healing in the face of hardware failures or other disasters, allowing PostgreSQL instances to be quickly rescheduled to rejoin a database cluster. However, cloud storage can massively reduce this advantage by adding arbitrary limits to where and how pods are scheduled and causing lengthy delays while the underlying storage is attached.

With Ondat, there are no delays and no uncertainty. Data is reliably striped across the Ondat cluster, ensuring it is local to the nodes and instantly ready for use. It also removes restrictions around node operations and availability zones, allowing DBAs to fully utilize the resiliency of their infrastructure. By eliminating cumbersome and often unreliable cloud storage operations, Ondat delivers on the reliability and resiliency promise of Kubernetes for PostgreSQL.

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

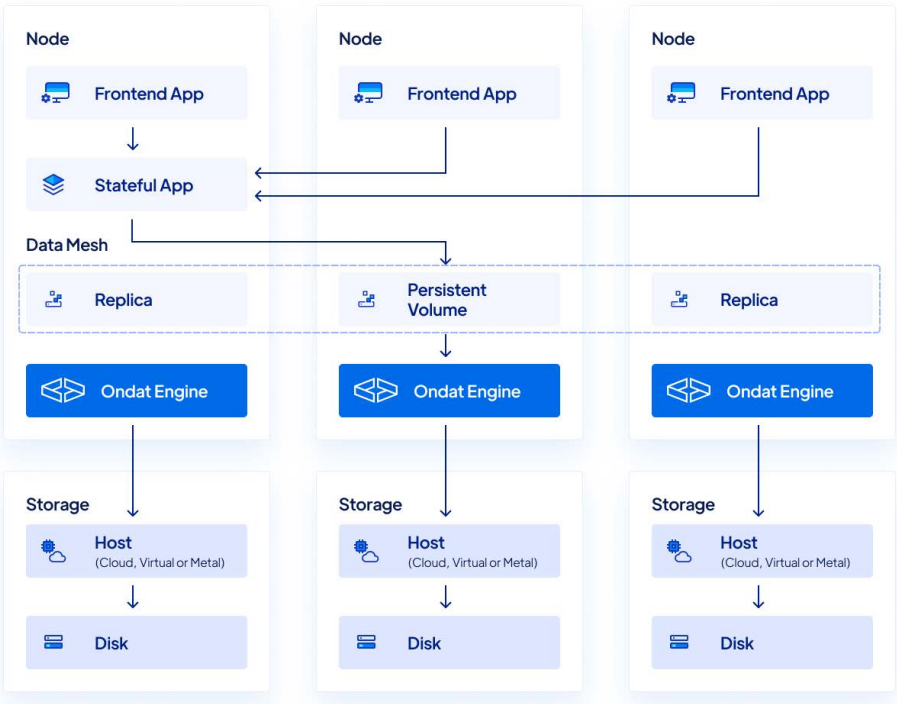

accenture DHL T-Systems LLOYDS BANK IMT CIVO INTERCONNECT

www.ondat.io

ondat™