

# Composable Data Infrastructure for Kubernetes

As organizations build systems to convert information into knowledge, they require data infrastructure that helps them move fast, adapting quickly to the changing needs of their customers.

Kubernetes is the industry standard for deploying workloads in the cloud, offering exceptional agility and scalability that enables development teams to innovate and deliver at scale.

The cutting edge of Kubernetes engineering brings stateful applications on-platform. Stateful Kubernetes brings new challenges to platform engineering teams. It requires a persistent data layer that is fast, resilient, and scalable.

Volumez is composable data infrastructure for Kubernetes that delivers breakthrough storage performance and enterprise-grade data services, with a developer-friendly declarative interface.

# Declarative configuration streamlines operations and eliminates complexity

One of Kubernetes' main advantages is its ability to specify the required state of a container using declarative configuration. This allows developers to easily define the amount of CPU, memory (RAM), and other compute resources required. Kubernetes then automatically ensures that each container runs with the resources required to match the declared configuration.

The concept of creating declarative configurations is well established in the cloud - it's commonly used for Infrastructure as Code (IaC), streamlining and accelerating cloud-native application deployment, maintenance, and more.

The declarative nature of Kubernetes makes it easy to manage complex compute and networking configurations for containerized, microservices-based applications at scale.

## However, Kubernetes was not originally designed for stateful workloads

Containers are inherently stateless. Designed for mobility, they contain only the runtime environment that's required to support a specific application, without persistently storing information about their state. This means that when a container is terminated, the data inside it - including the application state and configuration files - is also removed.

On the other hand, most mission-critical enterprise applications, such as cloud databases, AI/ML pipelines, and HPC environments, are stateful. Stateful applications need data to be retained beyond the container's lifecycle. In other words, stateful applications require access to persistent storage in order to operate effectively.

Driven by a rising demand to transform monolithic enterprise applications into cloud-native microservices, objects like StatefulSets enable Kubernetes to run stateful workloads, providing the ability to dynamically attach containers to storage. However, provisioning and releasing storage resources in a declarative manner - in the same way Kubernetes is used for managing compute resources for containers - remains a largely unsolved challenge.

# The Need for Declarative Data Infrastructure

Microservices are typically owned and managed separately by different teams of developers.

With Kubernetes and other cloud-native solutions, each team can use declarative configurations to specify the desired state of their microservice resources.

For cloud-native architectures where DevOps and platform engineers are responsible for infrastructure functions that were previously handled by a traditional IT department, these declarative configurations are key for supporting faster release cycles.

Standardization and open source have turned computing into a low-friction space - using declarative configurations, developers can now write their code once and run it on any IaaS/PaaS environment. However, this is hardly the case with storage. Instead of having a standardized data plane that provides compatibility and simplicity, enterprises must use disparate, proprietary data and storage technologies that are siloed by design.

Unlike the compute layer, the lack of standardization makes it extremely difficult to configure storage in a declarative manner. This multiplies the complexity, impairing enterprises' ability to scale with predictable performance, resiliency, and cost through automation. Without a common control plane that provides a unified management framework for storage, DevOps and platform engineers need to integrate with proprietary storage systems on each cloud platform, creating a matrix of features, performance profiles, and resilience schemes that create vendor lock-in and development overhead. This stands in direct contrast to the write once run anywhere declarative nature of Kubernetes.

The result? When attempting to bring stateful workloads onto Kubernetes, storage creates drag and inhibits agility. As data accumulates, the “data gravity” effect intensifies, further restricting the organization’s ability to respond to changing operational conditions or customer demands.

For example, when I/O performance is degraded (as a result of contention caused by workload surges or “noisy” neighbors), organizations must have the tooling to detect and then manually react with configuration changes.

This means that in order to ensure consistent high performance for mission-critical stateful workloads, organizations have to spend precious time and resources tuning and troubleshooting performance and resolving issues - all of which can drive costs out of proportion.

As cloud-native architectures continue to proliferate, the responsibility for provisioning, consuming, scaling, ensuring resiliency, and observing data is shifting to new users. That’s why change is required **across the stack!**

## Data infrastructure needs to be:



Quickly and easily consumed



Configured and managed in a simple way



Work seamlessly with modern stacks



In a nutshell, storage must become declarative

# Volumez Composable Data Infrastructure

Volumez makes data on Kubernetes behave like compute and networking on Kubernetes, so that teams can quickly execute and deliver new applications and new features to their customers.

Installing the Volumez CSI driver via helm chart takes about 30 seconds. DevOps then specifies the storage capacity, IOPS, latency, and zonal/regional resilience needed, just like with compute and networking resources, and Volumez takes care of the rest.

Volumez is based on a composable data architecture, which connects worker nodes directly to media over the network with NVMe-over-TCP. Volumez eliminates storage controllers from the data path, instead composing a Linux-based storage stack for each stateful container directly on the worker node. The result is exceptionally low latency, high IOPS, and enterprise-grade data services, including snapshots, thin provisioning, encryption-in-flight, and multi-zone resilience.

Since Volumez uses Linux as the data plane, it can compose any Linux-based resource running anywhere, on any hardware, while abstracting the complexity of data path orchestration. Users simply specify storage needs in terms of performance, security and resiliency in an easy declarative way, e.g. a YAML, API command, or web GUI, and Volumez composes a data path that's guaranteed to meet those requirements.

By tackling the key challenges and pain points in cloud storage, Volumez offers a first-of-its-kind data solution that delivers a range of benefits, enabling customers to realize the cloud promise of portability, agility, and efficiency.



## Guaranteed High Performance

Modern applications have strict I/O performance requirements. When deployed in the cloud, these applications often experience performance degradation due to unexpected workload surges, “noisy” neighbors and network issues.

Volumez gives end-to-end-control of storage performance to DevOps, delivering up to 2 million IOPS per volume with sub-millisecond latency and a strict, guaranteed performance SLA.

Volumez automatically profiles each media and reserves IOPS, bandwidth, and network throughput for guaranteed end-to-end application performance.



## Dynamic, Multi-Zone Resilience

Many prominent cloud storage solutions either do not support multi-zone resilience or require static allocation of high-cost volumes for resilience, hampering agility and contributing to cloud opex waste.

Volumez solves this problem with automatic data path resilience that is simple to configure at runtime. Developers simply specify level of zonal resilience their application requires.



## Security and Data Protection

Volumez encrypts data in flight and at rest using policies selected by DevOps during app deployment.

Volumez offers enterprise-grade data services including fast snapshots and restores, rollback, restore to alternate servers, and thin provisioning.

The Volumez data path - from the mount point to the media - resides entirely within the customer’s virtual private cloud and is built end-to-end on Linux. With no proprietary drivers, Volumez open data path provides the highest level of security by giving users full control of their data without even the possibility of vendor lock-in.



## Predictable Costs

Planning for storage growth in the cloud is a major challenge due to the complexity of billing calculation. The common outcome is over-provisioning of cloud storage and storage waste.

By automatically profiling and selecting the most efficient media configuration for every application, Volumez eliminates cloud storage waste and increases application CPU utilization by removing storage bottlenecks.

## About Volumez

Volumez is revolutionizing modern data infrastructure. The pervasive adoption of large-scale data analytics, artificial intelligence, and machine learning systems across industries has created an unprecedented challenge. Businesses need a way to convert knowledge into intelligence quickly, easily and at scale. Volumez has the solution. The company’s innovative controller-less architecture composes direct Linux data paths between media and applications, solving latency and scalability issues and unlocking consistently high performance and high resiliency.

Learn more at <https://volumez.com/> or contact us directly at [sales@volumez.com](mailto:sales@volumez.com)