

# Composable Data Infrastructure for Hybrid Cloud

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. Hybrid cloud environments are fast becoming the IT architecture of choice for delivering on this mission.

Ideally, the hybrid cloud model enables organizations to seamlessly deploy different types of workloads across environments. It provides them with the flexibility to utilize on-premises IT resources and shared cloud platforms to achieve operational efficiencies, enabling developers to easily provision resources on-demand, anywhere they are needed.

However, the vast majority of organizations that have embraced hybrid or multi-cloud strategies are not fully reaping the potential benefits.



### **Challenges with Modern Storage**

The primary obstacle for implementing effective hybrid and multi-cloud strategies is the rigidity of current data and storage solutions. When deploying cloud applications, developers need to provision compute, networking, and data resources. While standardization, open source, and the availability of advanced orchestration and automation tools have revolutionized compute and networking, this is hardly the case with storage.

Instead of having common data infrastructure that provides compatibility and simplicity across shared and private clouds and on-premises data centers, organizations with hybrid and multi-cloud environments must use disparate, proprietary storage technologies that are siloed by design. Lacking common feature sets and APIs, managing multi-cloud data infrastructure requires significant DevOps and platform engineering effort and ongoing maintenance.

Consequently, as organizations move further along with their cloud-centric digital transformation journeys they often hit a breaking point. Locked-in to storage systems and cloud services that are based on proprietary architectures, they struggle to create data infrastructure that is optimized to their needs. They must deal with many data services and APIs that are not standardized across platforms, creating a multiplicative effect of complexity that impairs their ability to scale with predictable performance, resiliency, cost, and security.



## **Volumez Composable Data Infrastructure**

Volumez is a universal data layer for shared cloud, private cloud, and on-premises data centers. No matter where a workload needs to run, Volumez provides applications with ultra-high performance and enterprise grade data services with identical features and a common provisioning and management API. By eliminating the need for proprietary storage services in each platform, Volumez abstracts the complexity of multi-cloud data, enabling organizations to run every workload in the right location at the right time, without the complexity and lock-in of siloed data services.

Volumez is based on a composable data architecture, which connects application servers 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 application directly on its server. 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.





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.



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