VOLUMEZ

Composable Data Infrastructure for High-Performance Databases

Driven by the rising demand for systems that convert information into knowledge, organizations increasingly rely on high performance database technologies that facilitate the ingestion, processing and analysis of massive streams of data in real-time.

Modern in-memory databases provide online transactional and analytical processing capabilities (OLTP and OLAP, respectively) in a single platform to support a variety of applications that require linear scalability while maintaining sub-millisecond latency for even the most complex queries. These include, for example, real-time analytics, reccomendation engines, sensor data processing, enterprise applications such as ERP and CRM, and many others.

As organizations move forward with their digital transformation strategies, they are hosting mission-critical databases in the cloud to benefit from improved agility, scalability, and more. At the same time, they expect tier-zero performance without having to refactor the application stack.

VOLUMEZ

The Challenge: Ensuring Workload Performance in the Cloud

Many customers who migrate mission-critical databases from on-premises to shared cloud experience a regression in application performance coupled with explosive increases in TCO. The same platform characteristics that enable cloud agility – shared infrastructure and economies of scale – work against applications with stringent I/O requirements.

High performance databases are exceptionally sensitive to write latency, with anomalies even at the P95 tail degrading transaction rates. On premises data paths make extensive use of technologies such as Fibre-Channel, NVMe-over-fabrics, and dedicated enterprise-class storage arrays to deliver on these requirements.

In shared cloud platforms, network and storage resources are, as the name says, shared resources. Service levels are best effort and applications typically experience significant variance in network and storage latencies on a day-to-day basis.

When dealing with mission-critical databases, the inability to achieve consistent high performance may negatively impact the user experience and, consequently, business performance. To address this challenge, organizations are often forced to rearchitect their applications for the cloud, or even undertake massively complex migrations to greenfield cloud-native database platforms.

Alternatively, organizations can choose to over-provision resources up front to mitigate performance issues and handle unexpected variations in latency. However, given the fast-growing amounts of data consumed by modern data platforms, this approach may quickly spiral costs out of control.

VOLUMEZ

Volumez Composable Data Infrastructure

Given these shortcomings, many enterprises choose to remain on-premises with their high performance databases, giving up the benefits of shared cloud.

This is where Volumez comes in.

Volumez is based on a composable data architecture, which connects database 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 database instance 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.

Volumez provides guaranteed performance, with up to 2 million IOPS and 400 microsecond latency per volume on shared cloud platforms. It achieves this by profiling each media and composing slices of capacity and performance across media and network channels into a resilient data path that guarantees IOPS, bandwidth, and latency.

Volumez enables organizations to leverage the agility of shared cloud without giving up the performance of on-premises databases. By providing databases in the cloud with guaranteed low latency, high bandwidth write performance, Volumez enables businesses to dramatically improve their operational capacity and exceed customer needs and expectations.

VOLUMEZ

Case Study: Boosting SAP S/4HANA Performance in the Cloud

SAP HANA is a columnar in-memory database for analyzing large amounts of real-time data. Combining OLTP and OLAP capabilities, SAP HANA enables customers to collect, analyze and act on business-critical data and execute complex transactions faster.

One SAP HANA customer, a large pharmaceutical distributor in the U.S., decided to migrate its SAP S/4HANA environment to shared cloud. The main motivation was to improve their supply chain management capabilities with specific focus on just-in-time inventory management with real-time visibility for their clients. The move to the cloud was aimed at achieving increased scalability while meeting strict performance and disaster recovery objectives. However, achieving the required service levels turned out to be a significant challenge. Hence, the client was looking for solutions that would provide the capabilities required to continue running SAP S/4HANA on AWS and unlock the full potential of the cloud.

Specifically, the client required a solution that would support the ability to collect and analyze large amounts of real-time data to coordinate stock, shipping and manufacturing. Other objectives included increasing customer satisfaction (CSAT) metrics by offering automated forecasting and achieving 10X scalability to support future growth. In terms of performance and resiliency objectives, the customer required a 10X increase in transaction rate and a 60 minute RTO/RPO – with no performance impact from backups.

After deploying Volumez, the customer surpassed these goals and achieved consistent high performance regardless of "noisy neighbors" or failures in the underlying infrastructure. By optimizing storage resource utilization, Volumez eliminated the need to over-provision resources, enabling new operational efficiencies and eliminating cloud opex waste.

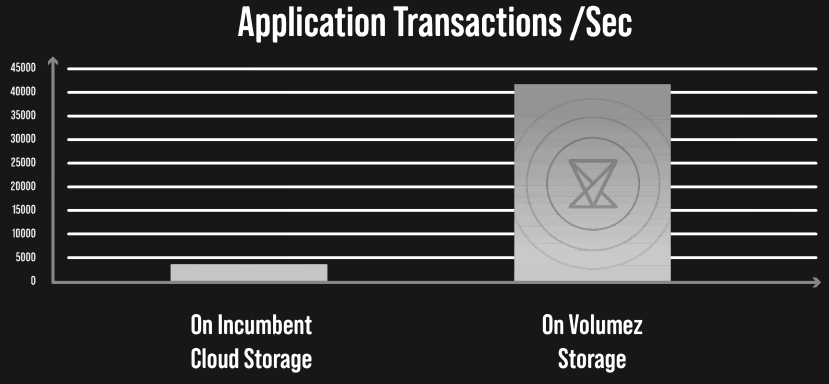

Result

42,000 »»
transactions/sec (11X increase)

IO latency reduced from
3.5ms to 400 uSec »»

Guaranteed storage performance including
4X daily backups »»

Linear storage scalability to more than
1 million IOPS (15X) »»

VOLUMEZ

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.

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.

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.

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