


Silk on Microsoft Azure



Silk Makes Applications in Azure Perform Predictably

The Azure Reality

Mission-critical applications don't struggle on Azure because the cloud is slow. They struggle because performance becomes unpredictable as environments scale and applications share data.

As transactional systems, analytics, and AI services compete for the same Azure data layer, latency fluctuates, throughput degrades, and SLAs become harder to maintain. To compensate, teams overprovision compute, storage, and database licenses — turning performance protection into permanent cost and operational burden.

Silk on Azure: Cloud Acceleration by Design

Silk is the cloud acceleration layer that makes Azure behave predictably for mission-critical applications.

Silk is an adaptive, software-defined SAN that sits beneath applications and above Azure infrastructure. Instead of tying performance to VM size, disk tier, or static configuration, Silk governs performance dynamically at the data layer.

By decoupling performance from capacity, Silk enables consistent latency, stable throughput, and reliable SLAs as concurrency and scale increase — without refactoring applications or redesigning architectures.

How Silk Aligns with Azure Architecture


Silk integrates directly with Azure-native constructs, operating transparently within existing environments:

- Azure Virtual Machines and VM Scale Sets
- Azure Managed Disks (Premium SSD v2, Ultra Disk, etc.)
- Azure Virtual Networks and Private Connectivity
- Azure Availability Zones and regional architectures

Silk complements Azure infrastructure rather than replacing it — correcting performance variability where it originates: the shared data layer.

What Changes with Silk on Azure

- 01** 30% lower cloud costs without the tradeoffs
- 02** Predictable application performance under concurrent demand
- 03** Throughput that scales without oversizing compute or disks
- 04** SLAs met by design, not by “just-in-case” capacity
- 05** Live production data safely shared across applications



Availability and Recovery Built In

Silk enables rapid recovery from infrastructure failures, data loss, or ransomware events without disrupting production applications. Recovery is measured in minutes, not hours, and becomes a routine operational capability rather than a high-risk event.

Why Silk for Microsoft Azure

■ Predictable application performance at scale


■ Safe access to live production data for analytics, Dev/Test, and AI

■ Structural cloud cost control by eliminating performance-driven overprovisioning

■ Fast, reliable recovery without operational disruption

Performance Validation

Unlimited application performance under real-world load



Why Silk for Microsoft Azure



10x Faster Performance and Sub-Millisecond Latency

Run your databases on Azure at the same speed – or faster – than on-prem



Always-On Availability

With always on availability and greater resiliency, never worry about unplanned downtime



30% Improved Cost Efficiency

Optimize your database footprint and compute while reducing database licensing costs

About Silk

Silk is an adaptive, software-defined SAN that serves as the cloud acceleration layer for mission-critical workloads, making cloud performance predictable at scale. It delivers real-time data access, consistent application performance, and instant recovery - without performance-driven overprovisioning or application refactoring.

Trusted by enterprises across finance, healthcare, SaaS, and more, Silk accelerates databases, analytics, and AI workloads with built-in resiliency. By enabling independent scaling of performance and capacity, Silk helps organizations unlock the full value of their data.

Learn more at <https://silk.us/>.