

Silk on Amazon Web Services (AWS)



Consistent Performance for Mission-Critical Applications on AWS

Silk gives your most important applications the performance consistency they expect on AWS, without refactoring, oversizing, or overpaying.

The AWS Performance Challenge

Mission-critical applications don't struggle on AWS because of speed; they struggle because performance becomes difficult and expensive as environments scale and multiply.

AWS provides powerful building blocks for performance, including provisioned IOPS volumes and high-performance instance types. However, as transactional systems, analytics pipelines, and AI services converge on shared storage, latency variability and throughput contention emerge. Teams typically respond by oversizing or duplicating EC2 instances, upgrading to premium storage tiers, or adding excess capacity as insurance.

Silk on AWS: A High-Performance Cloud Acceleration Layer

Silk operates as an adaptive acceleration layer between your database and AWS infrastructure. Rather than tying performance to instance size, storage tier, or static configuration, Silk dynamically manages throughput and latency at the data layer, helping maintain consistent application performance as demand fluctuates.

How Silk Aligns with AWS Architecture

Silk integrates natively with AWS components, operating transparently within your existing environment:



Runs on your current EC2 and Auto Scaling footprint, adding performance stability without operational disruption



Deployed from the Amazon Marketplace



Works seamlessly inside your existing AWS environment — no refactoring, no architecture changes required



Built for AWS HA patterns, deploy in multi-AZ design architecture for availability and consistent performance across zones

Silk complements AWS infrastructure rather than replacing it, addressing performance variability where it originates: the shared data layer.

What Changes with Silk on AWS

Helps maintain consistent application performance under all workload conditions	Enables throughput scaling independent of EC2 instance sizing	Supports meeting performance SLAs without overprovisioning capacity or compute
Allows controlled access to live production data across applications via application-consistent space-efficient clones	Reduces Day 2 operational noise by stabilizing performance and eliminating constant reactive tuning	Running Existing Applications on AWS Without Refactoring

Silk helps organizations run existing mission-critical applications on AWS with matching or exceeding on-prem performance from day one. Applications can be optimized before, during, or after migration to AWS, eliminating post-migration surprises and long-term overprovisioning driven by performance uncertainty.

Availability and Recovery Built-In

Silk supports rapid recovery from infrastructure failures, data corruption, or ransomware events with minimal disruption to production applications. Silk can insulate applications automatically from infrastructure issues, automatically repairing itself transparently from workloads.

For issues caused by external sources (e.g. logical corruption or ransomware attacks), recovery can happen in minutes rather than hours, making disaster recovery a more routine operational capability rather than a high-risk event.

Architecture Overview

Silk deploys using complementary architectural components that support performance governance and resilience across Availability Zones and regions. Silk Live and Silk Echo work together to deliver consistent performance and recovery capabilities in production and replicated environments.

Why Silk for AWS

- Accelerated and more consistent application performance at scale
- Cost control by reducing performance-driven overprovisioning
- Pooled resources between multiple databases and applications
- Controlled access to live production data for analytics, Dev/Test, and AI workloads
- Faster, more reliable recovery with reduced operational disruption

Organizations running Silk on AWS report significant improvements in application response times and infrastructure efficiency. Silk customers have achieved sub-millisecond latency consistency for database workloads while reducing their EC2 and EBS spending by 30-40% compared to traditional overprovisioning approaches.

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/>.