



Accelerate Dev/Test Velocity With Silk

Solution Brief

Executive Summary

Development and test teams only move as fast as their data. Silk's software-defined cloud storage platform and Silk Echo remove long waits for environment builds and eliminate dependence on stale, error-prone test data by making instant, production-accurate copies available on demand — without inflating storage. The result is faster iterations, fewer defects, and materially lower cloud costs.

The Challenge: Data Speed vs. Cloud Cost

Traditional copy workflows force teams to choose between current data and budget. Full copies consume capacity, drag down performance, and delay releases. The increased time investment in manually managing production-quality data means product and development teams often settle for mocked or stale data. This approach creates rework and risk later in the cycle. Silk's approach removes these trade-offs by virtualizing data copies at the storage layer and automating them at the database layer.

The Silk Platform: Under the Hood

Silk sits between your databases and cloud infrastructure and provides a shared, software-defined data layer optimized for performance, resiliency, and cost efficiency. Core architecture highlights include:

- **Two-tier virtualization with shared datastore:** Compute (performance) and capacity scale independently across a multi-tenant datastore, improving utilization and eliminating overprovisioning.
- **Silk DataPods:** DataPods virtualize resources using c.nodes, d.nodes, and m.nodes. Orchestration is handled by Silk Flex, and analytics/automation by Silk Clarity.
- **Enterprise data services:** Inline deduplication, compression, thin provisioning, zero-elimination, snapshots, and replication reduce footprint and boost performance.
- **Performance and resiliency:** The Silk platform consistently delivers sub millisecond latency and scales to 26 GB/s or 2.1 million IOPS per DataPod, with active-active design and self-healing to weather cloud turbulence.
- **API-first operations:** Teams automate through a REST API, CLI, PowerShell cmdlets, and a Terraform provider. Datadog integrations support observability.

Silk Snapshots, Clones, and Views

Snapshots: Instantaneous, pointer-based, zero-footprint point-in-time copies. Created with no performance impact, capacity grows only with changes. Per-volume-group limits apply. Can be instantiated into clones and views.



Thin clones: Read/write copies based on snapshots. Created instantly with no performance hit and initially no extra capacity. Support second-level cascading clones. Total volumes drive practical limits within a DataPod. Can be mounted read/write.



Views: Real-time, zero-footprint projections of production data that offload read demand (for example, AI/BI/RAG) to non-production without impacting production. Can be mounted read/write.



Silk Echo: How Database-level Copies Work

Silk Echo elevates copy data management to the database layer and is operated through a simple host/DB-level UI and full API. It uses a lightweight database host agent with instantaneous, zero-footprint snapshots and clones to deliver full-function database copies while leaving production untouched — at effectively no initial storage cost. Silk Echo supports both crash- and application-consistent operations and can clone to one or many destination hosts in a single run. Tasks are trackable through REST endpoints, enabling clean integration with CI/CD.

- **Performance of copies:** Clones and views use a metadata read cache, so reads from copies nearly mirror base-volume performance.
- **Automation examples:** Teams commonly schedule Silk Echo operations in pipelines to refresh lower environments on every build.

Why This Matters for Developers and Testers

- **Provisioning time drops from days to minutes:** Silk customers use Echo to create instant copies of production databases for testing and feature validation, turning days of manual work into minutes.
- **Accurate, writable test data:** Thin clones are read-write, so automated and manual tests mimic realistic scenarios without touching production.
- **No performance penalty on production:** Snapshots and clones are pointer-based and created instantly; they consume capacity only as changes are written.
- **Flexible placement:** Snapshots can be mounted across zones and regions — and even across cloud providers — supporting distributed teams and DR exercises.

Silk Supercharges Your FinOps Practice

Storage efficiency: Zero-footprint snapshots and clones avoid paying twice for data. Only changed blocks consume capacity. Silk customers have seen up to a 50 percent reduction in cloud data spend with these techniques. Track copy counts, retention, and differential growth.

Right-sizing compute and licenses: Pooled performance and separation of performance from capacity cut the need to overprovision large virtual machines (VMs) or database cores.

Time is money: Quantify the cost of developer/tester idle time waiting for copies, plus the cost of a DBA manually creating these copies. Silk Echo's minutes-not-days provisioning – and the ability of team members to create their own zero-footprint copies – reduces that waste and shortens release cycles. Bake copy automation into CI/CD to keep environments current.

Governance and chargeback: Use API telemetry to attribute storage deltas, snapshot counts, and refresh frequency per team. Set auto-expiration policies for clones/views and ensure masked datasets where appropriate.

Technical Details at a Glance


Placement: Runs on public clouds. Datapods can be built on VMs for peak performance or managed disks for peak durability.




Protection: Active-active architecture with self-healing. Native encryption at rest, SAML SSO, and CHAP for host authentication.



Performance: Up to 26 GB/s or 2.1 million IOPS per DataPod. Latency consistently under 1MS.



Automation: Full REST API, CLI, Terraform provider, and PowerShell cmdlets. CI/CD-friendly.



Silk Echo operations: Crash-consistent or application-consistent snapshots. clone one-to-many destinations. Monitor tasks via API.



Reference Implementation

The following example can help establish your implementation approach for using Silk for your development and testing use.

1. **Deploy Silk Flex** from your cloud marketplace and stand up a Silk DataPod sized for performance and capacity, then add hosts.
2. **Install the Silk Echo host agent** on source and destination database hosts; register hosts with Flex.
3. **Enable app-consistent snapshots** where required (for example, VSS for SQL Server).
4. **Automate** with the REST API, Terraform, and your CI/CD system for scheduled refreshes; pair with data-masking if needed.
5. **Govern** access securely with role-based controls (SSO/SAML) and CHAP; ensure data remains protected at rest through native encryption.

Expected Outcomes

Using Silk for your development and testing needs can help you achieve and surpass enterprise KPIs, including:

- **Velocity:** Self-service, one-click data copies is easily aligned to sprint cadence.
- **Quality:** Fresh, production-accurate data earlier in the cycle reduces rework and defects.
- **Cost:** Zero-footprint snapshots and clones shrink storage, and pooled performance reduces VM and license overprovisioning. Silk customers report material reductions in data spend.
- **Bottom line:** If your teams are waiting on data or paying for duplicate capacity, Silk and Silk Echo give you a faster, cheaper, and safer way to keep dev/test environments current — without touching production SLAs.

Getting Started

Ready to implement? Here are your first steps to using Silk for your dev/test needs.

-  **Identify target databases and SLAs** (RTO/RPO, refresh cadence, data-masking requirements).
-  **Deploy Silk Flex and a Silk DataPod** sized to initial performance and capacity needs; connect hosts.
-  **Install Silk Echo and register hosts**, then script snapshot/clone refreshes for dev/test pipelines.
-  **Pilot with one product team**, measure provisioning time, defect trends, and storage deltas; expand broadly.

About Silk

Silk fuels AI innovation by enabling real-time access to production data in the cloud. Seamlessly integrating high-performance cloud storage into AI workflows, Silk empowers organizations to enhance innovation while maintaining performance, reliability, and control over trusted enterprise data. Silk's agile data delivery eliminates the need to copy production data — enhancing flexibility and enabling that data to be leveraged by Generative AI.

Backed by over 20 technology patents, Silk helps customers unlock the full potential of the public cloud with speed and ease. Learn more at www.silk.us.