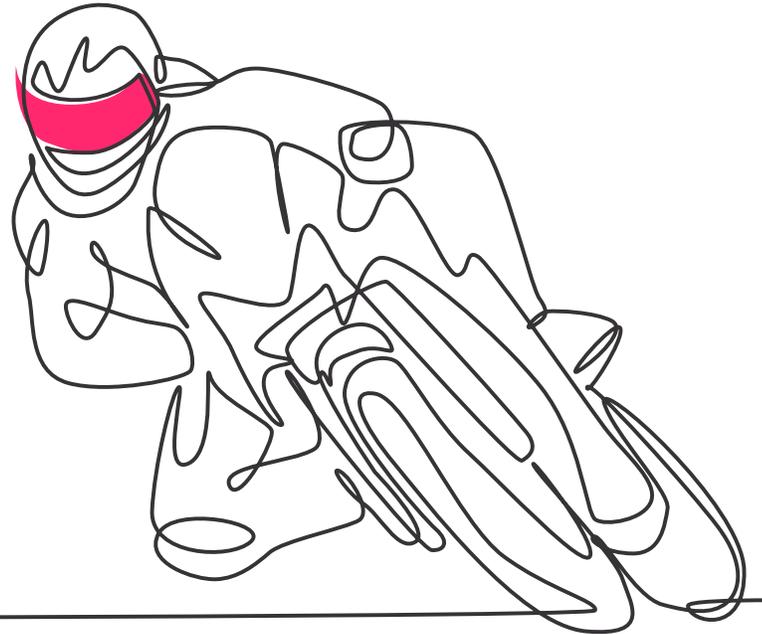


# Silk for Microsoft SQL Server

## Powering SQL Server Workloads on Azure



**As companies continue to move more and more workloads and applications to the cloud, it's becoming apparent that large, complex mission-critical databases are next. Microsoft SQL Server is one of the most widely used databases handling hundreds of gigabytes or terabytes for mission-critical applications. Because of the sheer complexity and demanding IO needed, it is typically very difficult to migrate SQL workloads to the cloud without significant thought and preparation.**

Typically, customers would try to migrate a SQL database towards an Azure PaaS service, but for demanding workloads with gigabytes of IO throughput per second or terabytes of data to support, a more straightforward (and when needed, cloud agnostic solution) is required. But just because the process of moving these workloads to the cloud seem for many too difficult to even consider, Silk makes it simple, fast, and cost-efficient.

This solution brief looks at how the Silk Platform helps you quickly and easily move databases from on-prem to any public cloud, while giving you ultra-fast performance that you can't achieve through native cloud alone and the peace of mind that you can migrate to the cloud while staying within budget and dexterously handle future growth.

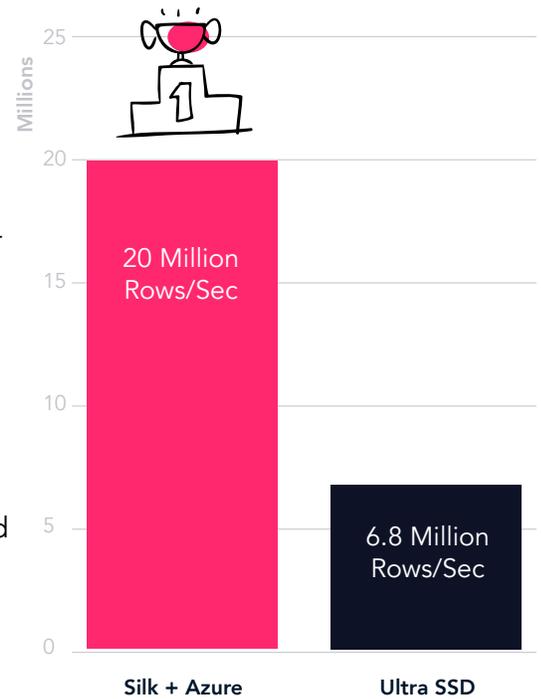
## Supercharging Database Performance

Many SQL Server workloads demand high performance. Data-intensive workloads in SQL that require more than 2+ GB/per second throughput or 80k IOPS are not going to achieve that level of performance on the public cloud. They will either need to be refactored for the cloud vendor of choice or left on-prem. This applies for all cloud vendors, including Azure, where performance for such workloads is throttled by the VM and its various managed disks offerings, limiting possibilities

But not with Silk! The Silk Platform decouples performance from public cloud resources, so high performance is guaranteed. There is no limit on throughput, IOPS, or disk size, so you can easily scale up without the need for more or larger VMs to get higher throughput. Plus a small VM can become supersized on its disks. Silk offers the same as – or better than – on-prem performance with sub-millisecond latency and the ability to dynamically scale your performance up and down automatically or manually as your workloads change. It's even simpler than operating your current SAN.

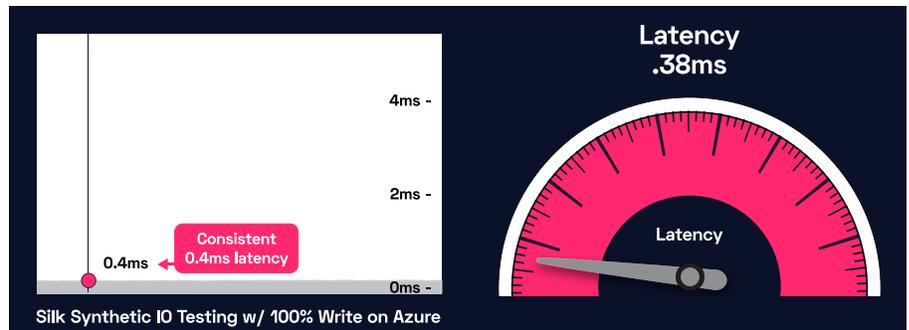
### Silk's proven performance makes it possible to:

- Bulk load more than 20 million rows in 1 second,
- Load more than 8TB of data into a SQL database in less than an hour,
- Perform a full table scan and retrieve a full 75+ GB data table into memory within half a minute. All while easily reading sustained 4-5 GB/sec without the need of partitioning tricks. That's performance from the very first second!



## Consistently High Performance Regardless of Block Size

If your SQL workloads combine transactional and analytic workflows, this can cause the database to generate an array of IO sizes. The changing nature and demand of these workloads can create inconsistent SLAs and inconsistent latency in the cloud. Large and inconsistent block sizes can dramatically hinder the performance of your entire database.



*Low latency and high throughput to smoothly support any of your SQL workloads.*

Silk's patent-protected technology allows for consistent low latency and ultra-high performance that is completely independent of Read/Write block size or pattern. So no matter what your workloads consist of, you will get the same, consistent experience while servicing your SQL workloads.

## Enabling Cloud Adoption with Lift and Shift

Refactoring your databases and applications to fit PaaS offerings maybe just not be worth it. It simply may be too time-consuming, too risky, too costly, or simply impossible due to latency requirements. With Silk powering SQL Server on Azure, we offer an enterprise cloud-ready alternative. Supporting lift and shift to the cloud, avoiding refactoring and allowing innovation over time.

Silk makes it possible to lift and shift your large-scale SQL Server database and achieve the same consistent experience as you did on-prem along with much faster performance than cloud native and refactoring alone. You can easily move your data across all clouds or even back to on-prem if needed without having to refactor the databases. And if your ultimate goal is to refactor SQL, Silk is a great way to get up and running in the meantime.

## Rich Data Services for Cost-Efficiency

Your costs can easily increase when migrating demanding workloads to the public cloud. The culprits are the cost of database licenses needed, and “hidden” costs related to data inflation and/or to support full restores of database clones within Dev/Test/Acceptance environments.

1. **Database Licensing** – Customers are often compelled to provision much larger VMs than are necessary from a compute perspective, just to get the data throughput performance they require from the attached resources. This has a significant effect on the license requirements.

With Silk, customers get the exact power and performance they need without paying for additional CPU, memory, and database licensing costs. Silk’s architecture overcomes the performance limitations and throttles, allowing for the use of fewer VMs with less vCPUs.

2. **Data Reduction** – Data reduction services – such as deduplication, data compression, and thin-provisioning – are features known from on-prem enterprise SAN solutions and are available to reduce your cloud capacity footprint. On average, data reduction services reduce your footprint by a 3:1 ratio. The native cloud offerings do not offer these data reduction capabilities. Silk does without extra license costs – making it possible to significantly reduce how many cloud resources you have to buy.

3. **Zero-Footprint Clones** – Database clones are frequently used in the CI/CD process for Dev/Test/Acceptance environments for ad hoc demanding analytics on a snapshot of a production. Snapshots are taken within a second without impacting the source (production) system. Typically moving to the cloud means you have to rethink these processes. Many customers are surprised to find that in order to create writeable clones, they will have to create full clones – which wastes cloud capacity and is time-consuming.

With Silk, you can continue your development and testing in the cloud in a similar way you are used to on-prem. Clones can be mounted for read/write purposes, serving to create additional working environments, at a very low cost to capacity. These read/write clones deliver the same performance as the production copies without any impact on the actual production database.

With Silk’s patented snapshot architecture, zero-footprint clones are created instantly, with no performance impact and without taking up any additional capacity. The time it takes to create a clone isn’t dependent on the number or size of the data being cloned. A clone is created within microseconds. These clones can be made available in different zones and regions, supporting multi-cloud data replication.

**Start getting a better cloud experience for your SQL Server today. Learn more at [www.silk.us](http://www.silk.us).**

**Silk is the database supercharger** – the smart platform that delivers game-changing database performance without changing a thing about your underlying apps or database infrastructure, whether you’re running real-time transactional workloads or analytical workloads – so your entire stack runs 10x faster. And with always-on availability across regions, zones, and clouds, your databases keep going strong no matter what the cloud throws at you. Industry leaders like Priceline, Cisco, and Telefonica rely on Silk for unlimited cloud flexibility, unbreakable data resiliency, and the greatest database performance of their lives.