Silk

How to Move Mission-Critical Databases to the Cloud

Selecting a Cloud Database for Mission-Critical Workloads

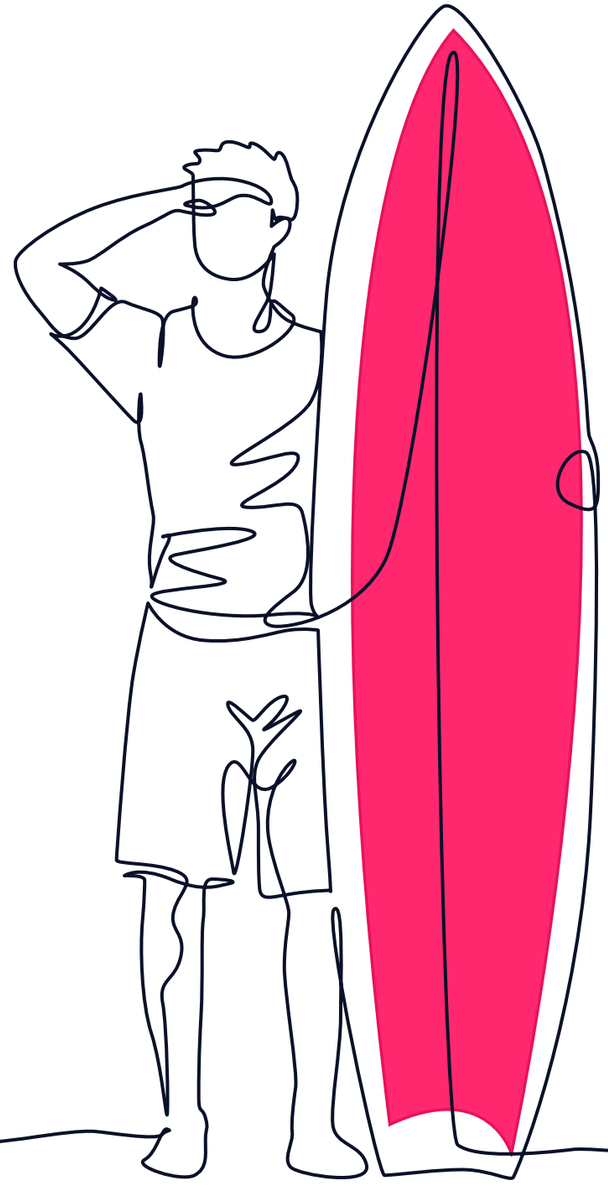

Introduction

Most established enterprises have huge applications and a huge amount of data to move to the cloud. You might be looking at a large relational database, such as Oracle Database or Microsoft SQL Server, that handles data from a transaction-processing application with numerous integrations to other enterprise applications.

Moving this type of database from on-premises to the cloud is hard work given performance, availability, and cost concerns.

Evaluating database options

Architects, developers, and administrators need to understand the different database options available for the cloud. Should you move to the cloud version of your current database? How about a database native to a public cloud provider, such as Aurora on AWS, Cloud Spanner on Google Cloud, or Azure SQL? There are also third-party cloud databases to consider.

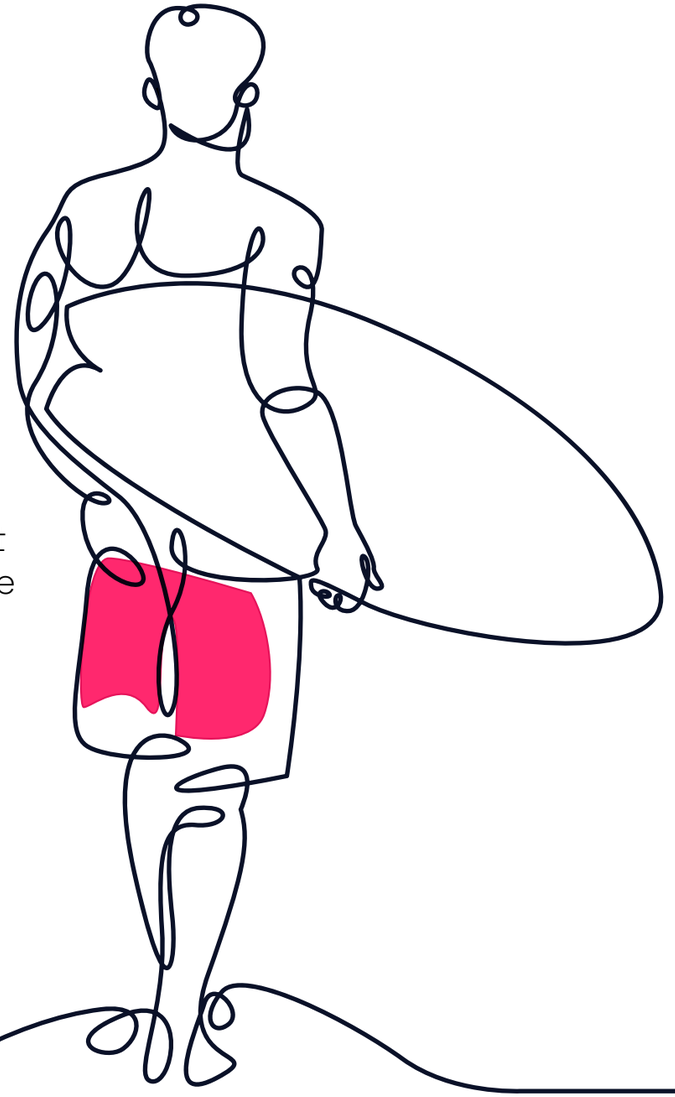

75% of all databases will be deployed or migrated to a cloud platform by 2022.

– Gartner

We will explore four options for moving a mission-critical database to the cloud, including:

1. Lift and shift your current database and host it on the cloud
2. Keep your database on-premises and either replicate data on the cloud or access data from the cloud
3. Migrate your data to a new cloud database and host it on the cloud
4. Migrate your data to a Database-as-a-Service (DBaaS) offering from a cloud provider

Gartner: The Future of the DBMS is Cloud, June 2019

Why database migration is hard work

After moving the “easy stuff” to the cloud with some lessons learned in hand, enterprises are starting to migrate their mission-critical databases. Database migration is expensive and risky, taking months or even years to complete. The process can be complex due to the database structure and applications using the data. The hurdles include data mapping, workload optimization, and dealing with the proprietary artifacts of legacy systems.

83% of data migrations fail or exceed budgets and schedules.
– Gartner

When transferring databases, homogenous, or same-to-same, database migrations are easier to perform than migrating from one source database to a different target database. For example, it is easier to migrate an on-prem PostgreSQL database to Google Cloud’s PostgreSQL, than it is to migrate an Oracle Database to Google Cloud’s PostgreSQL.

Relational databases (or SQL databases) are still widely used on the cloud, but there are other types of cloud databases available, including NoSQL, graph, key-value, document, time-series, wide-column, and in-memory databases. NoSQL databases scale more easily on the cloud than SQL databases, but are not the right solution for different types of applications.

If you decide to migrate from SQL to NoSQL, you will need to rearchitect your data model and application code because data modeling principles vary between the two approaches. Still, some teams prefer NoSQL because of the performance and scaling benefits on the cloud.

Ultimately, database migration can take months or even years to complete when workloads, applications, and analytics tools are factored into the process.

Gartner: Risks and Challenges in Data Migrations and Conversions, March 2012

Determining the way forward

It is certainly more challenging to migrate a mission-critical database to the cloud than it is to start greenfield with a new cloud database or DBaaS offering for new software development. That said, database migration to the cloud is not impossible, and it is a worthy endeavor to future-proof your organization.

Looking back at our database migration options:

If you keep your current database and host it on cloud IaaS or migrate to a new cloud database and host it on cloud IaaS, this is a self-managed model. The database will run on infrastructure made available by the cloud provider and database administration is under your control. You will have full control over your database and the way it is configured. You will need to ensure that administrators have the skills to maintain the underlying infrastructure. There is less risk of vendor lock-in, and you can move the database elsewhere if need be.

If you use a cloud provider's database and take advantage of DBaaS, this is a managed service model. Users do not need to install, configure, and administer the database themselves. Administrative tasks like backup, patching, SLA availability, and scaling up and down is taken care of by the cloud provider. However, you risk vendor lock-in as your mission-critical applications will be more tightly coupled to the cloud provider.

If you keep your database on-prem and either replicate data to the cloud or access data from the cloud, this is considered a hybrid approach. A hybrid approach provides flexibility and allows your IT team to switch over to the cloud as the primary system while running the on-prem database as a backup. There are a limited number of use cases that will work this way.

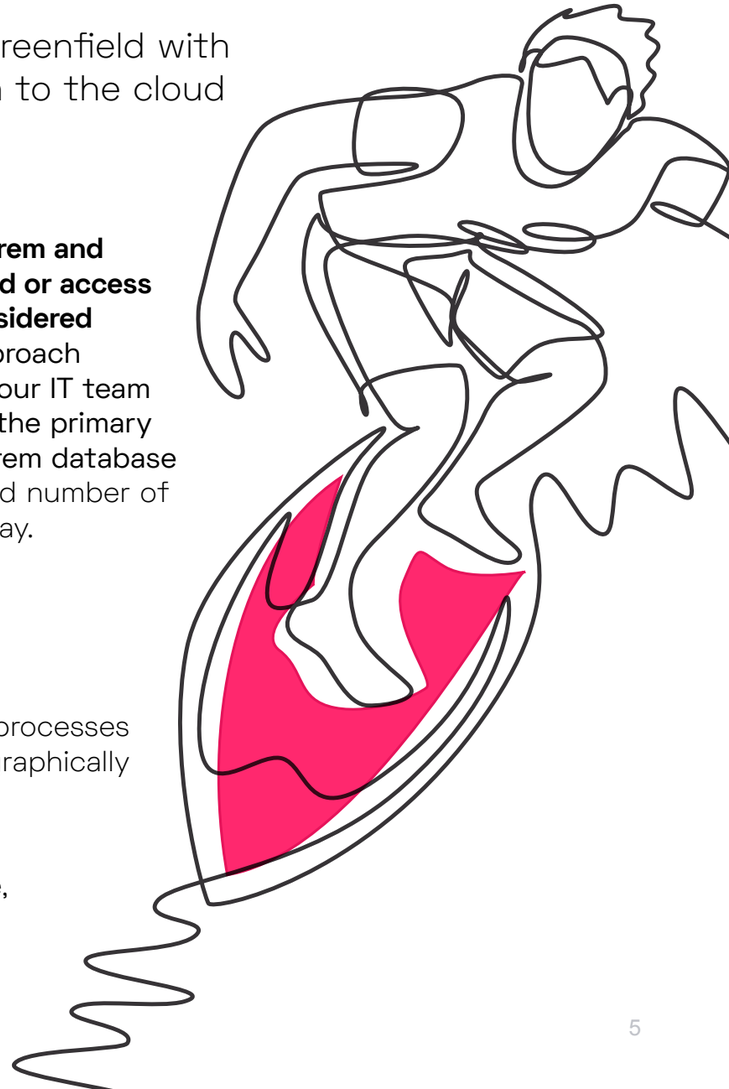

There's been a proliferation of cloud databases over the last decade that have been designed to handle more sophisticated processes than traditional on-prem databases. Users consider cloud databases to be easier to build, deploy, secure, and distribute geographically for policy compliance and redundancy. Cloud databases also offer new capabilities for developers, such as publish/subscribe, asynchronous messaging, data streaming, NoSQL, key value stores, and graph database functions.

Note that sometimes it may not be feasible to migrate mission-critical data to a different database given complexity, timeline, and cost. Demanding performance requirements may also limit your database options on cloud native infrastructure. Users must be aware of capacity and performance constraints for each database instance.

Optimizing database performance on the cloud

The emphasis on database performance is one reason why many enterprises have delayed moving mission-critical databases to the cloud. The cloud is not always a guarantee for better performance. In many cases, cloud database performance may be worse than on-prem, and you simply cannot risk performance gaps.

To optimize database performance, you must select the right data layer. Whether you are hosting a database on cloud IaaS or using a managed cloud database service, the cloud providers offer various database infrastructure options. You can adjust the throughput, latency, and IOPS of virtual machine (VM) disks by choosing the right VM size and storage tiers. Generally, larger and more expensive VMs offer more IOPS and higher throughput for database workloads.

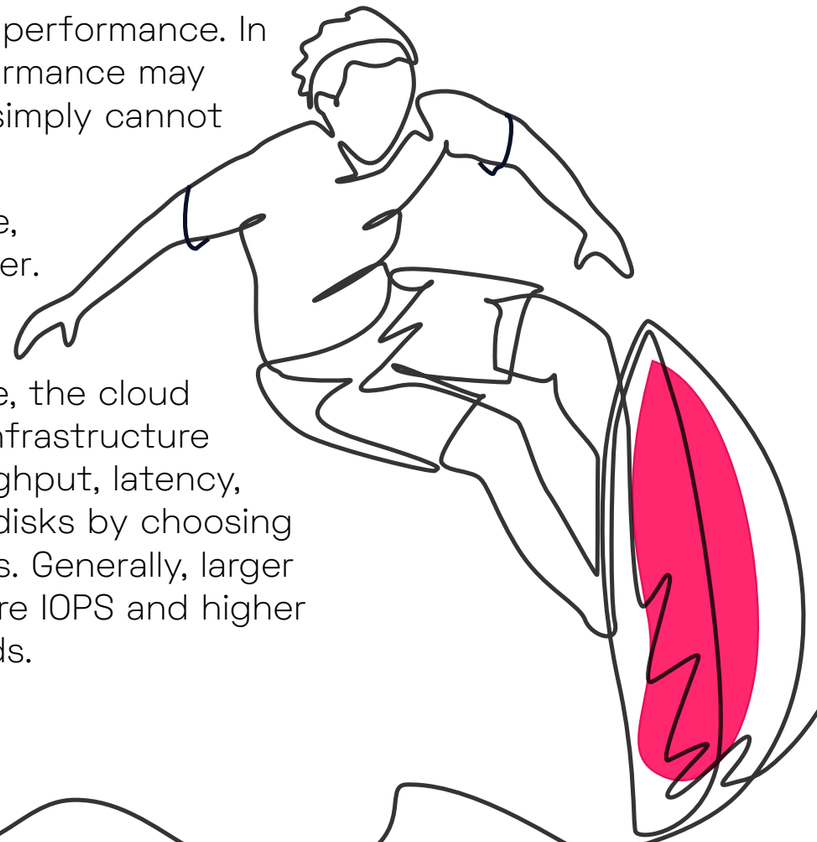

Here are the current database performance maximums from Amazon, Microsoft, and Google cloud native IaaS:

1. **AWS io2:** 256k IOPS and 4GB/s throughput
2. **Azure Ultra SSD:** 160k IOPS and 2GB/s throughput
3. **GCP SSD PD:** 100k IOPS, 1.2GB/s throughput

There are other options, such as managed SQL instances, with different performance thresholds.

While these numbers are fine for many applications, they may not be enough for mission-critical transactional or analytic database applications.

A word of caution: Those who experience performance issues may be quick to increase cloud resources – adding more storage to deal with I/O performance, more cores/CPUs to deal with processor-bound workloads or increasing available memory – which can lead to waste and higher cloud bills.

Using a data platform on top of cloud native infrastructure can help you scale database performance without overprovisioning cloud resources, so you can achieve mission-critical performance levels without a budget-busting cloud bill.

Silk supercharges database performance

If concerns about performance, availability, and cost are keeping your mission-critical workloads from moving to the cloud, then look no further. The Silk Platform is a data layer that sits invisibly between your databases and cloud infrastructure to supercharge database performance. Silk decouples performance from cloud capacity – with virtually no limit on IOPS or throughput – to deliver the fast performance you're used to on-prem, but on any cloud. There is no need to refactor or redesign around performance limitations, so you can take advantage of full data mobility while avoiding vendor lock-in.

Get 10x performance

Silk scales database performance without needing to provision larger VMs in the cloud or add more cloud resources. Your databases receive 10x the performance of cloud native alone to power mission-critical workloads on the public cloud with ease.

With full data mobility

Many enterprises will use a mix of cloud and on-premises databases for the foreseeable future. With Silk, you can seamlessly move data across clouds without refactoring or re-engineering your data for a specific cloud platform to enable true hybrid cloud and multi-cloud capabilities.

Enabling lift and shift

Lift and shift mission-critical workloads to the cloud without worrying about vendor lock-in. Get the right performance at the right cost, and flexibly move data when and where as needed. Refactor or redesign on your own timetable while taking advantage of the benefits of cloud computing.

At the best price point

Silk's Tier 1 data services including thin provisioning, deduplication, compression, and zero-footprint clones reduce your cloud footprint and save 30% on cloud costs. Database licensing costs will be reduced as you dynamically scale performance without needing more VMs.

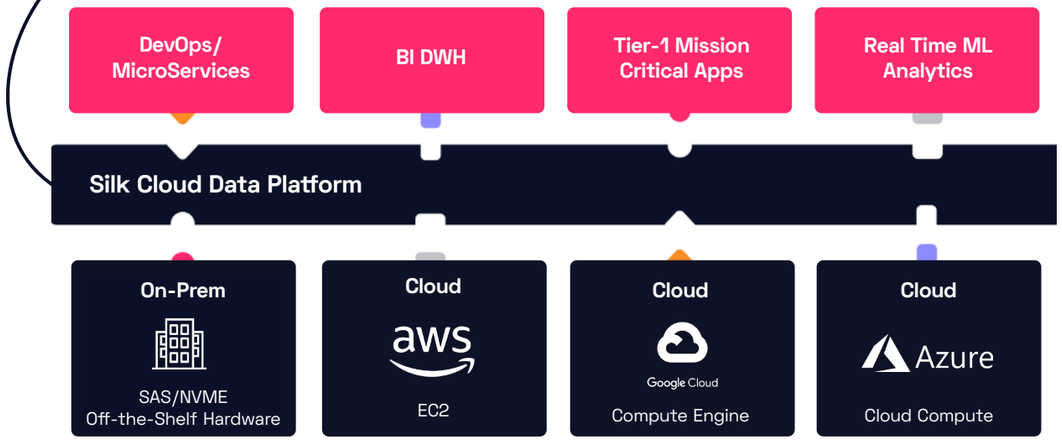

Ready to migrate your mission-critical databases to the cloud with supercharged performance?
Learn more at www.silk.us.

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.